STATE DEPENDENT PERCEPTUAL LEARNING

Craig Scott

Thesis submitted to the University of Nottingham for the degree of Doctor of Philosophy

March 2021

Abstract

Visual perception is vitally important for most everyday tasks. Whether reading a book, crossing a busy road or picking up on subtle changes in the facial expression of a friend, we rely on our brains to accurately and rapidly interpret the patterns of light captured by our eyes. Whilst seemingly effortless, these tasks rely upon complex processing, which is refined by practice and experience – via a process termed perceptual learning. Due to the prevalence of perceptual learning it is considered to be a widespread property of visual processing.

The majority of research on perceptual learning has focussed on a simple question: exactly *what* is learnt? A core proxy that most research has adopted to answer this question has focussed on training and transfer; whereby participants first perform two tasks, then train on one of those tasks for several days, before again performing the initial two tasks. This protocol has shown that in some cases improvements in performance are highly specific, whereas in others, improvements in performance generalise. This thesis adopts a fundamentally different approach to understanding perceptual learning by probing the internal state of the visual system – we refer to this as 'state-dependent perceptual learning'.

One factor that could be important for state-dependent perceptual learning is adaptation, which is another form of perceptual processing in which the response properties of the visual system change in response to regularities in sensory input. Typically, adaptation is studied by asking participants to passively view a given stimulus for a period of time (e.g. several seconds or minutes) and then measure changes in their perception to subsequently presented changes. Conventionally adaptation and perceptual learning have been investigated in isolation, however only recently research has shown that altering the adapted state that a participant learns in can determine whether training on a given visual task is beneficial, or detrimental to performance. In Chapter 3 of this thesis, I explore this interaction between adaptation and perceptual learning further by characterising its properties and limits. I found that training in an adapted state alters perceptual sensitivity (a core hallmark of perceptual learning) but does not alter perceptual bias (a core hallmark of adaptation). In Chapter 4 of this thesis, I explored whether state-dependent perceptual

learning generalises to other visual dimensions if adapted to across different timescales, however found learning to be specific to the trained visual dimension.

Another important factor for state-dependent perceptual learning could be self-initiated action. Our visual and motor systems are highly integrated and in most everyday tasks their functioning is integral for optimal performance. Some research has shown that visual processing is altered following self-initiated control over a stimulus and that physiological change in areas that prepare motor actions may underlie perceptual performance. In Chapter 5, I first explore whether the internal adapted-state is altered when the visual stimulus is self-initiated, however found that perceptual bias to the moving stimulus was consistent with conditions requiring no self-initiated control. In Chapter 6, I did however find that detection performance was modulated by the congruence between self-initiated action and the visual stimulus. I hypothesised, in line with other research, that visual-motor expectations could conceivably drive this finding and investigated whether repetitive training in an environment where detectable motion is incongruent to these expectations, could drive a consistent change to visual expectations. This further exploration indicated that visual motor expectations are flexible and state-dependent perceptual learning *may* alter these types of visual expectations.

Throughout this thesis state-dependent perceptual learning is explored and the outcomes are discussed in light of theories of perceptual learning in humans, and as a potential practical application for both clinical rehabilitation and expertise attainment.

Acknowledgements

I would really like to thank first and foremost, my supervisor Neil Roach who I have had the pleasure of working with and learning from throughout my MSc and PhD. Thank you for always having your door open and for your patience with my many questions over the years. I have gained a lot from working with you, your guidance has been invaluable to my PhD and also my personal development.

I would also really like to thank my other supervisor Ben Webb for being a great support over the years. Thank you for your guidance, especially in our team meetings where we were often scratching our heads at what my experimental results could mean for several hours. I wish you every success in whatever the future holds for your new business.

I of course cannot forget to thank the Economic and Social Research Council (ESRC). Thankyou first of all for supporting this project, for laying on many networking events over the years and pushing me to learn about topics outside of my own research area. Whilst the thought of doing extra modules alongside of my PhD was at first slightly daunting, in reflection it was a really useful exercise in applying my knowledge and gaining perspective on the type of skills you learn as a PhD student.

Another big thank you to my wonderful friends, colleagues and family. To my office mates and the Visual Neuroscience Group, thanks for being there over the last three year, what a ride – thanks for your support and for the cake and coffee mornings. To all my friends outside of work, thank you so much for checking in to see how it was all going and for being a great set of ears for when I needed to chat. To my family thank you for being an amazing support, telling me to keep going and not to give up.

Finally, I would like to thank my partner Camilla Babbage. I would not have met you if we both had not decided to continue on in academia and start PhDs with the ESRC. I may dare to say that it has therefore been the best decision I have ever made. That introductory event with the ESRC does not seem that long ago now and I cannot believe we have made it through together. Thank you for everything!

Contents

Chapter 1: Introduction

If we think of all the actions that we perform on a day-to-day basis, whether that is pouring our morning cup of coffee or riding our bicycle along a cycle path, the necessity for our visual system to accurately represent sensory information and feed into decision-making processes, so that we can decide when to stop pouring our coffee or navigate away from something stationary in our environment, becomes quite obvious.

An integral part to the successful completion of many of these everyday tasks are the decisions we make based on visual information. A visual decision is defined as a process where patterns of neural activation are converted into a categorical proposition about the location, presence, or identity of those task-relevant neural patterns (Gold & Stocker 2017). They represent a best guess at what the decision ought to be given the current evidence (Helmhotz, 1924). The seemingly effortless nature of these decisions is in part due to *neural plasticity*, which is the capacity for our neural systems to be modified as a result of sensory input across different timescales.

One obvious instance of neural plasticity that influences visual decisions relates to practice and experience. If we take an example of a senior radiologist inspecting an X-ray chart, their ability to detect diseased tissue will be informed by their formal training and experience completing the same task on countless other X-ray images over several years. The repetitive nature of this will have enabled their visual system to hone its processing and become highly sensitive to detecting the diseased tissue necessary for the diagnostics task. Their decision will be confident, well-informed and highly accurate, which is the opposite to how non-specialist trainee radiologists will likely perform given the same task. Another factor, that may not be so obvious, is an internal adapted state in response to the immediate sensory environment that also informs visual decisions. To continue with our example, before inspecting the X-ray chart, the radiologist will ensure the room is dark. In doing so their visual system adapts to this light change, and after several minutes will be ready to locate any diseased tissue on the computer screen or illuminator. If the trainee radiologist suddenly opens the door and fills the room with light, the senior radiologists adapted state will be altered and they will find the diagnostics task impossible. This is because their visual systems adapted state will be suddenly altered due to light adaptation, which will cause a drop in their ability to perform the diagnostics task. Therefore, their ability to perform this task to such a high standard is in some way dependent on the adapted state.

An additional internal state factor that influences visual decisions is motor actions. Consider the following situation: you borrow your friends’ computer but find the mouse cursor moves faster than you expect. This results in inaccuracies that make it hard to navigate the cursor. For instance, perhaps you find that you go to look for the cursor, but it is often in an unexpected location on the screen. Luckily, after several minutes your brain will adjust, and you will likely be making smaller adjusted hand movements to compensate for the cursors speed and this will allow you to better predict where the mouse will be and track it more consistently. However, when you return to your own computer, you might temporarily make the opposite error (i.e. underestimating where the cursor will be) as a result of your experience on your friends computer. In this example it is clear that our visual and motor systems play an important role in shaping our perception through self-initiated action.

In combination, the visual decisions we make (whether explicitly or implicitly) are shaped by the internal factors just described: *perceptual learning*, the timescale over which our visual system learns through experience and practice, *adaptation*, the timescale over which our visual system adjusts to our immediate sensory environment, and *self-initiated action*, the process whereby our visual perceptions are shaped through how we interact with the world. To date, the vast majority of research has focussed on these factors independently. Although taking this approach has led to many advances in our knowledge, it is important to understand however, how these processes may share a functional relationship.

Is there any evidence that the internal adapted state of the visual system is an important factor during the perceptual learning process? Do they share a functional relationship? Can our internal visuo-motor system alter adapted state processing? And can visuo-motor expectations be altered through perceptual learning? In this chapter, I will explore the background research for each of these questions in turn and build a case for the rest of the work in this thesis. In section 1.1, I will consider how adaptation and perceptual learning may exist within the same functional framework which I will refer to as *state-dependent* *perceptual learning*. I end this section by describing how the rest of this thesis will explore state-dependent perceptual learning. In the section 1.2 that follows, I will consider how our visuo-motor system may shape our visual perception and bring together some common themes to inspire some justified aims for how this thesis will investigate how the visuo-motor system may shape adapted state processing and perceptual learning.

1.1 A Functional Relationship between Perceptual Learning and Adaptation: State-Dependent Perceptual Learning

1.1.1 Definition and Introduction to Perceptual Learning

A much-studied behavioural manifestation of neural plasticity is *perceptual learning*, where repetitive practice at a challenging perceptual task can lead to substantial and long-lasting improvements in task performance (for reviews see: Sagi, 2011; McGraw et al., 2008). Perceptual learning has been widely documented in adulthood and has been found on the vast majority of visual tasks (Fine & Jacobs, 2002; Fahle & Poggio, 2002) and across different sensory modalities. For instance, perceptual learning has been demonstrated for taste (Balleine, Espinet & Gonzalez, 2005; Bende and Nordin; 1997), smell (Wilson & Stevenson, 2003) and auditory tasks (Polley, Steinberg & Merzenich, 2006). In the visual domain, perceptual learning has been shown to improve performance in a number of tasks including orientation discrimination (Schoups et al., 2001; Shiu & Pashler, 1992), spatial frequency discrimination (Fiorentini & Berardi, 1980; Astle, Webb & McGraw, 2010), contrast detection (Sowden, Rose & Davies, 2002; Polat, 2009), motion discrimination (Ball & Sekuler, 1987), vernier acuity (Poggio, Fahle & Edelman, 1992; Fahle & Morgan, 1996) and depth perception (Fendick & Westheimer, 1983; Ramachandran & Braddick, 1973). In addition, studies have revealed that feedback is not a requirement for learning to occur on these tasks (Herzhog & Fahle, 1997; Fahle and Edelman, 1993) and these learnt improvements are long-lasting as they have been shown to still hold for several years after the initial training period (Ball and Sekuler, 1987).

When measured psychophysically, the perceptual consequences induced by perceptual learning result in a reduction (improvement) in either detection or discrimination thresholds as the training period progresses. A typical training period lasts over several days. Detection thresholds refer to the lowest intensity at which the presence of a stimulus can be detected. For example, the lowest motion coherence percentage of a moving dot pattern at which one can detect any coherent motion. Whereas, discrimination thresholds refer to the smallest perceivable difference in a stimulus dimension, also referred to as the just noticeable difference (JND). For example, if presented with two patterns that differ in either orientation/contrast/spatial frequency or colour, what is the smallest difference along one of these attributes that you can reliably identify. Perceptual learning has also been investigated using reaction times (Harwerth & Levi, 1977; Polat, 2009), however this approach will not be explored in this thesis.

As highlighted earlier, perceptual learning-induced changes to either detection or discrimination performance have been shown for a wide range of different visual attributes that are thought to be processed in different parts of the visual cortex. A major hallmark of these trained visual improvements is that they are often strongly coupled to the visual attributes used during the trained task. For instance, training to discriminate orientations around vertical (0°) leads to performance improvement for that trained orientation (Schoups et al., 2001), but not for an orthogonal orientation or a different spatial frequency dimension (Fiorentini & Berardi, 1980; Fahle & Morgan, 1996). Other features that learning is specific toward includes retinal position (Ball & Sekuler, 1987; Shiu & Pashler, 1992, Crist et al., 1997; Ahissar & Hochstein, 1997; Fahle, 2004; Poggio et al., 1992), size and binocular disparity (O'Toole & Kersten, 1992) and even the temporal order of the visual stimulus (Zhang et al., 2008) used during training.

The high specificity of perceptual learning has led some to theorise that physiologically, perceptual learning could be due to modifications at relatively early stages of visual processing (e.g. primary visual cortex – V1) where neurons "encode" simple stimulus attributes (for review see Florentini & Berardi, 1997; Fahle, 2004; Weiss, Edelman & Fahle, 1993; Teich & Qian, 2003). In models, the perceptual consequences are explained by amplifying the signal from neurons in the early stages of visual processing (encoding neurons) most useful for the task as a consequence of stimulus repetition (Weiss et al., 1993). In these cases, the high-level parts of visual cortex continuously interpret the encoding level area and benefits from the amount of noise reduction at the encoding stage through amplification (Fahle, 2004).

Support for this theory comes from two sources. The first is the plethora of psychophysical training tasks that utilise low-level visual characteristics that neurons in V1 show a response toward (contrast: Geisler & Hamilton, 1992, spatial frequency: Foster, Gaska, Nagler, & Pollen, 1985, orientation: Hubel & Weisel, 1962, temporal frequency: Hawken, Shapley & Grosof, 1996). If learning is specific and does not generalise, then it seems plausible that neural plasticity underlying these changes might occur in these early-stage cortical areas only. The second source of support is various electrophysiological and functional magnetic resonance imagining (fMRI) studies that have shown changes in V1 linked to improved behavioural changes (Crist, Li & Gilbert, 2001; Schoups, Vogels, Qian & Orban, 2001; Furmanski, Schluppeck, Engel, 2004; Schwartz, Maquet, Frith, 2002).

However, the specificity of perceptual learning when measured behaviourally, does not necessarily imply that perceptual learning improvements are confined to these early stages of visual processing (Mollon & Danilova, 1996). It is entirely possible that the neural plasticity that underlies these perceptual improvements may occur in cortical areas of the visual system outside of V1, that “read-out” (decode) from these early-stage encoding areas to form decisions. This idea overcomes the argument that encoding neurons are not *aware* of the wider task requirements, which is a problem for theories that confine perceptual learning to early processing stages (Seriès, Stocker & Simoncelli, 2009). For example, if perceptual learning is due to a sharpening of V1 neurons that are tuned to that specific visual attribute, then this could be deleterious to the performance of other tasks that draw on the same visual areas (Gilbert, Sigman & Crist, 2001) – assuming neurons in V1 are not aware of the task demands. Therefore, it seems intuitive that perceptual learning may occur outside of these early cortical areas.

A growing body of work supports the idea that learnt changes might occur at the level at which visual information is decoded to form decisions and behaviours (Law & Gold, 2008;

Dosher & Lu, 1998; 1999; Petrov, Dosher & Lu, 2005; Vania, Sundareswaran & Harris, 1995). One model proposed by Dosher & Lu (1999 – see figure 1.1) explains the neural plasticity that underlies perceptual learning improvement as a process where higher-level decoding areas add weight to the most relevant encoding neurons for the perceptual task. In this way, the reweighting of encoding neurons ensures there is a reduction of internal noise (encoding neurons not relevant for the task) and increased efficiency in the decoding areas that read-out from these early stages of visual processing. This model has a central mechanism that analyses the inputs from lower-level stages of processing. The inputs from the low-level stage are tuned to different properties of the input (i.e. orientation) image which in combination compute a noisy representation the image. The differing amounts of external noise are excluded along with internal noise suppression through channel reweighting – where the most optimal neurons for the task are favoured.

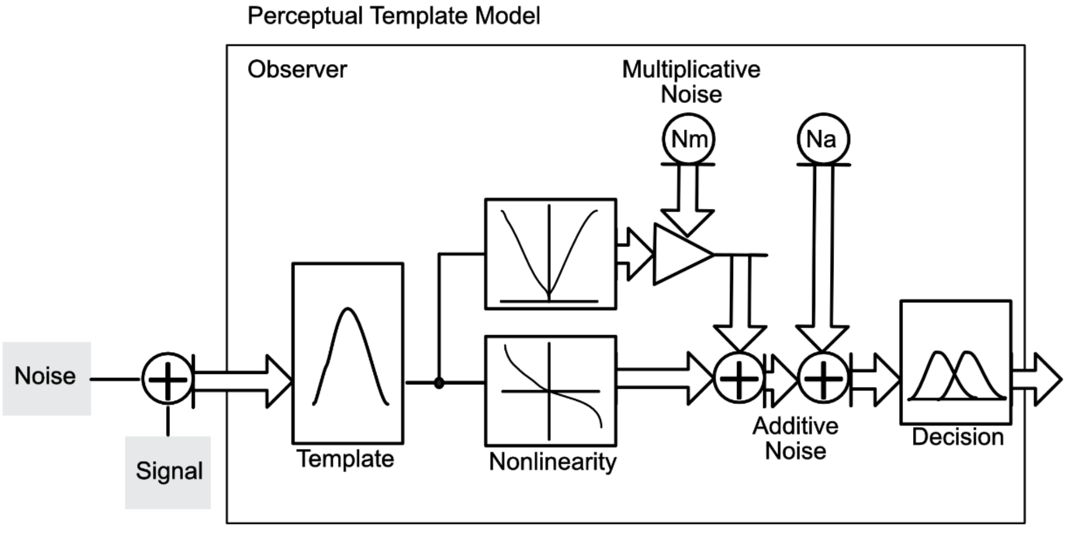

Figure 1.1. Perceptual Template Model of Perceptual Learning. The signal and noise is first processed, in a perceptual template relevant to the task. Nonlinearity transducer functions, along with multiplicative noise (Nm) and additive noise (Na) are then added to the signal before the final decision process (adapted from Dosher & Lu, 1999).

This idea is supported by electrophysiological research that measured neuronal populations in lateral intraparietal area (LIP) and middle-temporal area (MT) in two monkeys (Macaca mulatta), who were tasked with determining the direction of visual motion (Law & Gold, 2008). Area LIP is part of the parietal cortex and has been implicated in many different visuo-motor and cognitive functions (for review see: Snyder, Batista & Andersen, 2000; Colby & Goldberg, 1999), as well as being an indicator of saccadic choice (Shadlen & Newsome, 2001). Thus, LIP can be considered a higher-level area that decodes visual information to form decisions. Area MT on the other hand, has been implicated as an area that encodes visual information regarding location and direction of moving stimuli (Britten, Newsome, Shadlen, Celebrini & Movshon, 1996). The results showed that neurons in LIP were insensitive to visual motion at the start of training but developed responses as performance improved, and this correlated with task performance. Whereas neurons in MT responded consistently to that motion stimuli throughout the experiment.

This electrophysiological evidence certainly adds support to the reweighting model (Dosher & Lu, 1998) and this idea that perceptual improvements may correspond to increasingly selective readout from encoding level areas that represent visual information. Whilst this idea makes a good case for perceptual learning it is challenged however by some experiments that have shown learning generalisation across task (Xiao, Zhang, Wang, Klein, Levi & Yu, 2008; Zhang, Xiao, Klein, Levi & Yu, 2010). For instance, in these examples where learning generalises across retinal locations, despite no task training and thus neuronal reweighting to occur, it raises the possibility that there is some other neural mechanism. This will be covered in more depth in section 1.2.2 of this chapter.

One alternative theory that has gained popularity, known as the Reverse Hierarchy Theory, proposes that perceptual learning maps onto different receptive fields within the visual hierarchy and its generalisation or specificity is dictated by which cortical area is modified through learning (Ahissar & Hochstein, 1997, 2004). This theory proposes that perceptual learning is a top-down guided process that progresses backwards through the higher and early input stages of the visual hierarchy, until a better signal to noise ratio is reached. The authors propose that depending on the level of task difficulty, learning can generalise when training conditions are easy as they match the spatial generalisation of higher visual areas with the largest receptive fields. However, as task difficulty increases, learning becomes more specific as the learning process works along the visual pathway to match the finer spatial retinotopy exhibited by early-level areas with the smallest receptive fields. This model can explain why under some circumstances learning can generalise, which explains some of the contradictory data within the field. However, this model has been challenged by evidence that suggests that it is not necessarily the difficulty of the training task that dictates learning generalisation, rather it is the difficulty in the task to which learning is transferred to that is important (Jeter, Dosher, Petrov & Lu, 2009).

Overall, it is a challenge for the perceptual learning field to distinguish between these different interpretations. Perhaps an alternative view is that perceptual learning can be the combination of both, depending on the way learning is tested within the experimental paradigm. Petrov, Dosher & Lu (2005) proposed a set of learning predictions based on the stimuli and type of training and transfer design (see Figure 1.2). In some cases, depending on the training and transfer task, the two tasks may rely on distinctly separate units (such as different retinal locations) and distinct connections from those units to decision units (orientation to spatial frequency task – see Figure 1.2a). In other cases, the sensory units (low-level areas) are shared between tasks (i.e. have the same retinal location) but the weighted connections to decision units are separate (orientation to spatial frequency task – see Figure 1.2c). In each of these examples Petrov and colleagues give a prediction of learning and transfer. For instance, if there are two tasks that rely on the same stimulus input and there is a change to the low-level representation of the stimulus through perceptual learning then we might expect learning to transfer if the same stimulus is used (see Figure 1.2d). Alternatively, we might expect no transfer as the two stimuli rely on different set of connections to decision units (see Figure 1.2c). Ultimately this framework is valuable when thinking about how perceptual learning can be achieved through either reweighting, representation modification or both. It also highlights that learning might not just be represented by a brief period of training at the transfer task but could be seen by multiple switches between tasks.

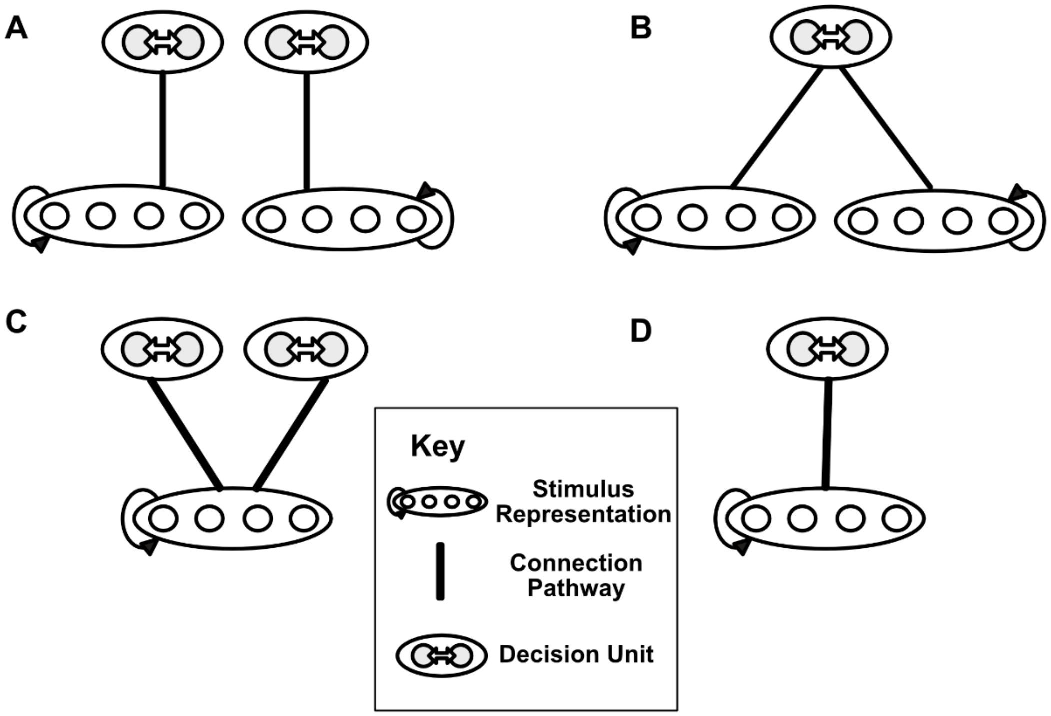

Figure 1.2. Four Key Task Situations Probing Perceptual Learning. Each type of task situation (A, B, C, D) has a stimulus representation phase, a decision unit and a path that connects them. In A training and transfer conditions rely on separate sensory units, connections and decision units. In B, training and transfer conditions rely on the same decision unit but separate stimulus inputs. In A and B, learning would be expected to be specific, but it is not possible to determine which mechanism is responsible. In C, training and transfer conditions share the same stimulus representation but different decision units. Here we expect learning to transfer if it involves change at the sensory representation stage and specificity if the change occurs at the decision unit. In D, both the stimulus representation and decision unit are shared between conditions. Under this condition, learning could transfer under both hypotheses (adapted from Petrov, Dosher & Lu, 2005).

In summary, it is a matter of considerable debate whether learning reflects neural plasticity at the encoding or decoding level of the visual system. In addition, the specificity of learning is a contentious issue and one that has predominately been looked at through the lens of training and transfer tasks. Whilst this has advanced our understanding considerably on perceptual learning and neural plasticity within visual processing, a fundamentally different approach, where learning is considered in light of other forms of neural plasticity, could shed further light on the underlying mechanisms and perceptual consequences of perceptual learning.

1.1.2 Definition and Introduction to Adaptation

Adaptation is a well studied behavioural manifestation of neural plasticity, in which the response properties of the visual system change in response to regularities of the immediate sensory environment (for reviews see: Webster, 2011; Wark, Lundstrom & Fairhall, 2007). Adaptation is considered to be a ubiquitous property of visual neurons and has been shown to occur at multiple stages of the visual hierarchy (Helson, 1964). For instance, like perceptual learning, adaptation aftereffects have been measured for a range of low-level stimulus attributes, including motion (Hol & Treue, 2001), orientation (Georgeson, 1973; Ware & Mitchell, 1974), colour (Jones & Tulunay-Keesey, 1975; Webster & Mollon, 1997), spatial frequency (Blakemore & Sutton, 1969), and even high-level attributes such as facial processing (including expression, gender, ethnicity and identity – Leopold, O'Toole, Vetter & Blanz, 2001; Webster & MacLeod, 2011). The study of visual adaptation has been a part of vision science for a long time and is often commonly referred to as the 'psychophysicists electrode' with the idea that 'if it adapts, it's there' (Webster, 2015). Thus, visual adaptation is a useful tool that has been used to increase our understanding of the visual system.

As mentioned in the previous section, neurons in V1 are sensitive to low-level stimulus attributes. A classic example is the tight tuning curves neurons display in V1 for orientation (Hubel & Wiesel, 1962), which is characterised as a bell-shaped curve with the largest response at the preferred orientation and sharp drops in responsivity for orientations further from the neurons preferred direction. Physiologically, the effects of adaptation are typically characterised as a reduction in this peak response of neurons tuned to that stimulus feature (Blakemore & Sutton, 1969; Clifford, Wenderoth, & Spehar, 2000; Coltheart, 1971). Through prolonged presentation of an oriented stimulus, it has been reported that the firing rate of neurons tuned to the adapting orientation reduce (Blakemore & Campbell, 1969; Carandini, Movshon & Ferster, 1998, Hammond, Pomfrett & Ahnmed, 1989). After adaptation, this manifests in a reduction in the amplitude of orientation tuning curves of neurons whose preferred orientation was near the adapted orientation (Teich & Qian, 2003 – see Figure 1.3). Thus, when a subsequent oriented stimulus is presented after the adapting stimulus, there is brief temporary change in

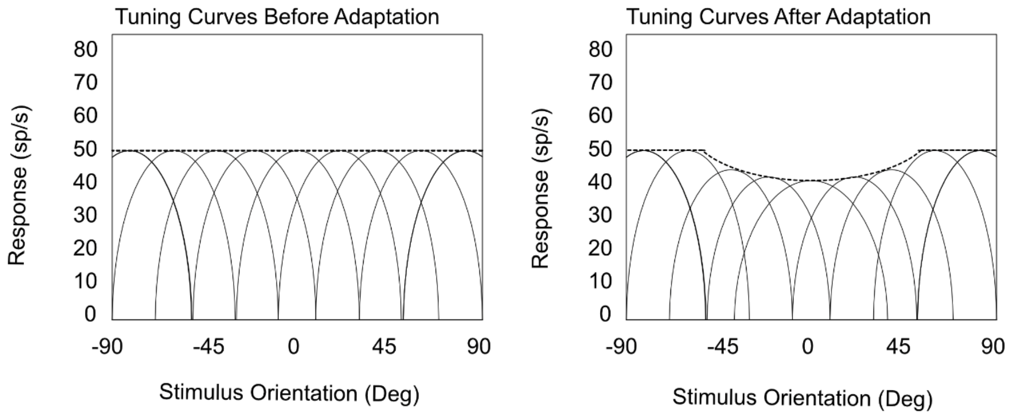

Figure 1.3. Orientation Tuning Curves in V1 Before and After Adaptation. Simulated orientation tuning curves before (left)and after (right) adaptation to a vertically orientated stimulus (0°). Adaptation to oriented stimulus results in a reduction in firing rates to the neuron tuned to that orientation and to neurons with tuning curves near to the adapting orientation. Adapted from Teich (2003).

sensitivity as the visual system processes the subsequent stimulus in light of the lingering aftereffect of the past adapting stimulus (Webster, 2011).

Whilst reductions in neuronal firing occur for many stimulus features when measured physiologically, other complex changes have also been recorded in neurons following adaptation. For instance, Dragoi, Sharma and Sur (2000) measured changes in neuron’s orientation tuning curves following prolonged adaptation to a particular orientation using pattern adaptation. Following 2 minutes of exposure to a tilted stimulus 15-20° away from the neurons preferred orientation, their results showed a response on the flank of the tuning curve being depressed, whilst responses on the opposite flank showed the opposite and were enhanced. Thus, instead of reducing the firing rate, the neurons were reorganised and shifted their preferred orientation away from the adapting stimulus. The magnitude of this shift depended on the relationship between the adapting orientation and the neurons preferred orientation, which was maximal for differences 15-20° away. For instance, following adaptation to a stimulus 60-90° away had little effect on the neurons tuning which shows specificity for selective adaptation. The authors suggest that the reorganisation of the neurons tuning curve may improve detection performance for stimuli orientated near the neurons preferred orientation (Dragoi, Sharma & Sur, 2000).

The perceptual consequences of adaptation are usually measured as a perceptual bias or a change in perceptual sensitivity. Perceptual biases are characterised as a change in the appearance of subsequently presented stimuli and are often referred to as aftereffects. Two well studied aftereffects are known as the Tilt-Aftereffect (TAE) and the Motion Aftereffect (MAE). These aftereffects can be measured by recording the duration of the adaptation aftereffect (typically in seconds) or by having participants estimate the appearance of stimuli when subsequently viewed. For instance, a well-studied example is the MAE (for review see: Mather, Pavan, Campana & Casco, 2008), where following exposure to motion, subsequent stationary stimuli can appear to move in the opposite direction for a period of time (discussed further is section 1.2.1) or can also induce a form of repulsion bias to the motion direction of subsequently presented visual stimuli away from its true direction (for review see: Rauber & Treue, 1999).

Another type of measurement for perceptual bias is the point-of-subjective equality (PSE). The point of subjective equality is a value that a given participant, is equally likely to judge as equal to some standard stimulus. For example, if participants are performing a discrimination task between two stimuli, the PSE correspond to a stimulus value where the participant is equally likely to choose either of them.

Perceptual sensitivity has also been shown to be altered following adaptation. In section 1.1.2 of this chapter, I defined discrimination and detection for perceptual learning and introduced the JND as a performance threshold. These definitions still hold under adaptation paradigms, however the way they are affected differs. For instance, typically any change in discrimination and detection thresholds are measured after prolonged exposure (usually for several seconds) to a stimulus before each stimulus presentation (trial). In addition, thresholds are altered given the exact relationship between adaptor and test stimuli (probe). For example, passively exposing participants to either upwards (0°) or bi-directional motion, inclined equally either ±20°, ±40°, ±60°, ±80° and ±100° from upwards can be used to nullify the direction aftereffect so that changes to perceptual sensitivity can be measured independent of any bias (Levinson & Sekuler, 1976; Regan & Beverley, 1985). Interestingly, the relationship between adaptor and a slightly offset from vertical oriented (±2°) test probe, reveals there is little deviation in orientation discrimination when the adaptor and test probe are highly similar, i.e. when the adaptor is

0°. In some cases, this can cause an advantage to orientation discrimination (Kristjansson, 2011). However, as the adaptor orientation deviates further from the test probe, the discrimination thresholds reveal a tuning function which peaks (i.e. performance is most inhibited) around ±40°, ±60° (McGovern, Roach & Webb 2014; Price & Prescott, 2012).

The relationship between the adaptor and test probe has been important for understanding of how visual information is encoded but also read-out to form decisions. For instance, it has been demonstrated that the visual system may rely on different subsets of encoding neurons depending on the type of task (Hol & Treue, 2001 – see Figure 1.4). Hol & Treue (2001) demonstrated this for orientation discrimination and detection and modelled this using a population of direction selective neurons. The authors revealed that for orientation discrimination, adapting 50° away from the test orientation was most effective at raising detection thresholds (inhibiting performance). This is evidence that when the visual system is performing a discrimination task it may rely on the neurons that have the greatest change in response between the two presented stimuli it is comparing. Thus, adapting away from vertical (50°) causes reduced responsiveness in neurons tuned slightly away from each stimulus direction (off-axis neurons) that are most useful for the task. For orientation detection on the other hand, adapting the test direction is most effective at lowering detection thresholds (facilitating performance). This reveals that the visual system depends on the signal to noise ratio across the population of direction tuned neurons for orientation detection. The noise caused by stimulus noise leads to an even activation across the neural population, whilst the to-be-detected-motion embedded in this noise results in an increase in neurons tuned to that direction. A similar pattern of impaired detection and discrimination is seen across the visual domains including motion (Phinney, Bowd & Patterson, 1997), spatial frequency (Regan & Beverley, 1983), colour (Krauskopf & Karl, 1992) and contrast (Blakemore & Campbell, 1969; Pantle & Sekuler, 1968).

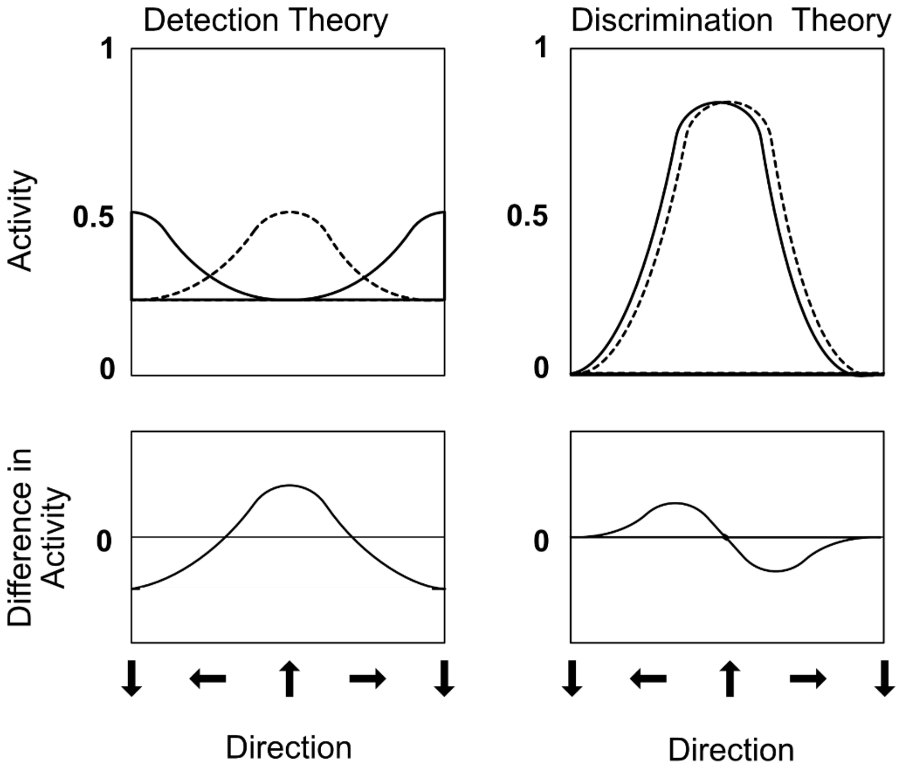

Figure 1.4. Neuronal Population Activity for Detection and Discrimination tasks. For detection (left – upper panel) the population activity (y-axis) elicited by stimuli moving up (solid line) or down (dashed line). The x-axis represents the preferred direction of the sampled neuron. In a detection task the activities of the signals (upwards or downwards) are embedded in the activity of the noise (straight black line). The difference in activity (left - lower panel) is largest for those neurons preferring the direction of the signal. For discrimination (right – upper panel) the population activity elicited by a stimuli moving slightly leftward (solid line) or rightward (solid line), means that the problem for the visual system is to distinguish whether an even activation is caused by one or the other. The largest response differences (right – lower panel) occur for neurons preferring a direction at the flanks of the curve, i.e. neurons approximately 45° from the stimulus direction.

Despite its utility as an investigative tool, why adaptation occurs at all in the first place remains a curiosity. As Webster’s (2015) review highlights, its prevalence across many stages of visual cortex points to a common purpose. Nevertheless, much like the different functions that each stage of the visual hierarchy performs, the functional purpose of adaptation may also change depending on the visual attribute or task.

One role that has been proposed is that adaptation functions to maximise efficient coding and optimise information transmission (Wainwright, 1999; Clifford, Webster, Stanley, Stocker, Kohn, Sharpee & Schwartz, 2007; Wark, Lundstrom & Fairhall, 2007; Stocker &

Simoncelli, 2006, Rieke & Rudd, 2009). Neurons in the visual system are faced with encoding a large range of different stimulus attributes whilst only having a limited dynamic range to do so (i.e. a preferred orientation). To do this, neurons must function optimally by increasing the transmission of new information (Attneave, 1954) and maximise the bandwidth available for novel information transmission (Srinivasan, Laughlin & Dubs, 1982). In addition, it is important that neurons are not responding continuously so that there is balanced sensitivity across the entire neural population. This could be achieved through a process of adaptive decorrelation (Barlow & Foldiak, 1989). Clifford and colleagues (2000) argued that adaptation was a result of 'centring', where the dynamic range of responses is shifted in on a particular range and 'scaling', where the visual system adapts to the variability in the stimulus to decorrelate their responses for efficient information transmission (Clifford, Wenderoth & Spehar, 2000).

Together these approaches are highly successful at predicting the properties of adaptation at early levels of the visual system. Contrast adaptation is a good example of this. For instance, much like retinal ganglion cells, neurons in V1 have a dynamic range for contrast and cannot respond across the entire range of contrasts. Adaptation for different contrasts has been shown to shift neurons contrast response function and centre on the adapting stimulus contrast (Durant, Clifford, Crowder, Price & Ibbotson, 2007). Without this process, there would be response saturation and information loss in neurons. Which would render simple processes for everyday vision (i.e. discriminating small changes in contrast) hard. Where these approaches fall short however is at higher levels of the visual system where its properties become less clear. For instance, the characteristics of faces do not vary over the same range as contrast levels do, therefore it is difficult to demonstrate and understand why we would need to continuously adjust to faces.

Another theory that has been proposed for adaptation is that it functions to maintain perceptual stability through error correction (Andrews, 1964). In this theory, both the observer and the stimulus have variations that the visual system through adaptation needs to discount. By discounting it is possible to remove extraneous information (Webster, 2011). Support for this theory comes from perceptual aftereffects that do not change perceptual sensitivity. For instance, altering the lighting of the environment (Foster, 2011) or alterations to the observer themselves, such as the yellowing of the eyes lens (Werner &

Schefrin, 1993), that leads to bias in the neural code and is error corrected by adaptation to provide stable perception over the environment. This theory could also help explain why adaptation occurs in higher level areas– for example, facial recognition might benefit from tying percept's more directly to the individuals unique attributes, by adapting out the average configural properties of one's social group (Webster & MacLeod, 2011).

A final theory for adaptation suggests that it is integral for building predictions about the world (Chopin & Mamassian, 2012). Under this theory the past is important for the visual system to estimate the future. Thus, a predictive adaptive code efficiently saves metabolic resource by signalling only when there are errors – i.e. percept's that do not match the remote past (Barlow, 1990; Gardner, Sun, Waggoner, Ueno, Tanaka & Cheng 2005). Support for this theory comes from the fact that adaptation tends to null out stimuli we are exposed to overtime, thus it would be efficient for the visual system to do this and build a current prediction of the world that determines what captures our attention.

In summary, adaptation is a well-established phenomenon, and its prevalence throughout the hierarchy of the visual system suggests that it is a fundamental process in determining how and what we see. However, as we have explored in the latter half of this section, there are debates that still remain regarding its functionality and even why we should adapt at all. This latter point is something that is not contested in the perceptual learning literature, as the goal of learning is obvious. However, the physiological consequences of adaptation have been more widely researched and without it theories of learning would not be as well established. In combination, both forms of neural plasticity operate over the same neural architecture and more than likely do not operate as distinct mechanisms when processing visual information in the real-world. Thus, it would be informative for both fields to understand how they may operate together.

1.1.3 Where the hallmarks of Perceptual Learning and Adaptation become blurred

As set out in Sections 1.1.1 and 1.1.2, the fields of perceptual learning and adaptation have been studied extensively and as consequence our understanding of these types of neural plasticity have advanced separately. An obvious distinction between them is that learning is traditionally distinguished from adaptation because it reflects changes in performance rather than appearance and has a longer time course (Teich & Qian, 2003). However, these perceptual distinctions are not absolute (Webster, 2011). For example, like adaptation, learning paradigms can also bias the appearance of visual stimuli (Haijiang, Saunders, Stone & Backus, 2006). In addition, like learning, adaptation can also facilitate some discriminations (Kristjansson, 2011; Barlow, 1972) and facilitate learning generalisation through repetitive passive exposure (Gutnisky, Hansen, Iliescu & Dragoi, 2009). Thus, there may be some functional relationship between adaptation and perceptual learning that could further our understanding of them. This will now be discussed further in the following sections.

The Effects of Long-Term Adaptation

Adaptation is a useful tool for probing different parts of the visual system and its resulting aftereffects are *usually* short and non-permanent. What evidence is there to suggest that the perceptual aftereffects of adaptation may however be modulated through long-term experience? An interesting early example comes from De Valois (1977) who investigated the characteristics of spatial frequency adaptation. Her procedure specifically looked at adapting to gratings at certain spatial frequencies and the aftereffects on participants ability to detect subsequent gratings at other spatial frequencies. In part of her report, she details how some subjects *"were run in either this experiment or closely related tasks almost daily over a period of 1.5yr"* and therefore she was able to investigate whether there were any long-term changes to the adaptation aftereffects overtime.

Figure 1.5 shows the effects of repetitive adaptation to a grating at 1.19 cycles per degree (cpd) for one year, on spatial frequency detection. A noticeable difference between these timepoints is the decrease in the bandwidth of the adaptation aftereffect, and supposedly (although it is not graphically represented) the increase in the amplitude of the aftereffect, which increased *"often"* throughout the year. Whilst this only shows data for one out of five participants tested and represents data for two discrete periods, these results suggest that the aftereffects of adaptation for neighbouring spatial frequencies was becoming weaker, perhaps as a consequence of increasing training and experience of the task in an adapted-state overtime. A potential reason for this result could be that, perceptual learning, was altering the adaptation aftereffects by adding more weight to those cells involved in the task, therefore restricting the sensitivity range of those cells affected by adaptation. Whatever the potential explanation is behind this change, as De Valois (1977) explains adaptation experiments "*such as this one … cannot be taken as simple, direct reflections of a simple, unvarying, underlying physiological organisation*".

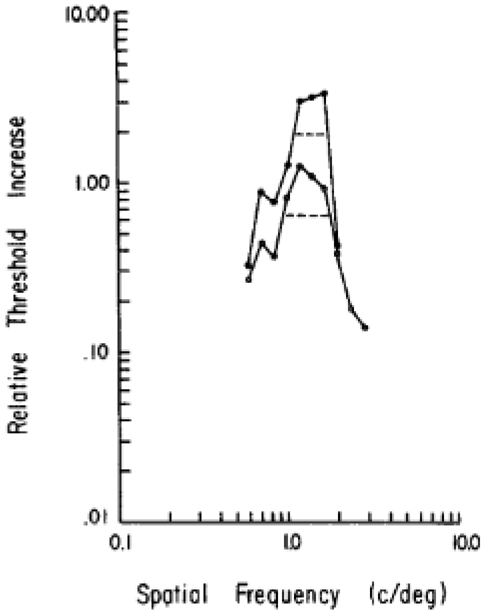

Figure 1.5. Relative threshold increase following adaptation at 1.19 c/deg taken at two different timepoints. The black circles (top line) represent data from three sessions conducted 1 year prior to the black squares (bottom line). The data is taken for one subject who run almost daily sessions of the same task within these two timepoints. Taken from Figure 7 in De Valois (1977).

Since this research, many other laboratories have investigated adaptation under a range of other different scenarios. One avenue of research has investigated adaptation aftereffects for prolonged timescales of adaptation – spanning several hours and days with the use of portable head-mounted displays (Kwon, Legge, Fang, Cheong & He, 2009; Zhang, Bao, Kwon, He & Engel, 2009; Haak, Fast, Bao, Lee & Engel, 2014). This is ecologically interesting since in many real-life settings environmental changes stretch over many hours or multiple days. For instance, after four hours of adaptation to a natural setting with reduced contrast, contrast discrimination improved and the time course of this change revealed different coding strategies as adaptation progressed (Kwon et al., 2009). In similar experiment, depriving participants of a certain orientation in their environment for several hours, increased orientation-specific sensitivity following deprivation (Zhang et al., 2009). These increases in sensitivity can be explained as adaptation leading to increases in neuronal gain overtime, rather than reductions. The limits of this have been reported by Haak and colleagues (2014) who deprived subjects of vertical contrast adaptation over four days. The authors results revealed that after two days the aftereffects of long-term adaptation asymptotes, after which it continued at a slower rate.

The shared finding of increased perceptual sensitivity following prolonged adaptation (Kwon et al., 2009; Zhang et al., 2009; Haak et al., 2009) has also been shown in a small number of experiments investigating orientation adaptation over some shorter timescales (Kristjansson, 2011, Barlow, 1972). For instance, Kristjansson (2011) has shown that adaptation to tilt can improve discrimination around the orientation of the adapting stimulus, and this improves steadily with increased adaptation time (30 seconds). In combination, this improvement in discrimination, which is a commonly associated hallmark of perceptual learning, raises the intriguing question as to whether the adapted state during learning is functionally useful for perceptual learning.

One informative investigation related to state-dependent perceptual learning sought to understand whether experience with adaptative states can change how quickly one adapts (Yehezkel, Sagi, Sterkin, Belkin & Polat, 2010). In this example, they asked subjects to wear prism glasses that induced orientation blur and measured the adaptation aftereffect (in this case percent perceptual distortion) following 2- and 4-hours of adaptation. This was repeated for each participant up to four times on different days within a 10-day period. The average result across sessions showed that adapting for just 2 hours led to minimal reduction in perceptual distortion caused by the lenses and led to strong aftereffects when the lens was taken off. However, similar to those researchers using head mounted display units (Kwon et al., 2009; Zhang et al., 2009), 4 hours of adaptation led to significant reductions in both the distortions and the aftereffect. Strikingly, the aftereffect was completely eliminated when the glasses were taken off after 4 hours. The authors also tested the relationship across the repeated sessions and revealed that between session 1 and 2 the degree of adaptation began where it left off from the previous session and continued to reduce (Yehezkel et al., 2010). This highlights that the visual system may have processed the input under the two different adapted states differently (glasses on vs. off) and was perhaps able to consolidate this perceptual adapted state as a form of learning that enabled them to processes these adapted states faster.

Whilst these results are promising, there are two reasons for caution before making strong conclusions regarding state-dependent perceptual learning. The first is that Yehezkel and colleagues base their learning effect on 9 subjects, 4 of whom completed up to four days of training and the remaining 5 participants who only completed the first session and were included in their learning estimate. It is possible therefore that inclusion of these 5 subjects increased the degree of adaptation between session 1 and 2, meaning that the average degree of adaptation would be much greater. If all 9 participants had completed all 4 days of the experiment, then we would potentially see a different learning effect entirely.

Another point of contention is that Yehezkel and colleagues did not show whether the adaptation aftereffect from taking the glasses off changed across days. This would be important to understand in light of their claim that the adapted state was being consolidated as a form of learning between sessions. Furthermore, it would be interesting to understand whether any learnt change in one adapted state altered anything about the processing of another. For example, if learning relied on the internal adapted state then would we expect to see any interaction between the glasses on vs. off states given that one of the hallmarks perceptual learning is its specificity? Alternatively, if learning in one adapted state altered the other adapted state by causing some detriment to performance then would this be an indication of some of the limits of state-dependent perceptual learning?

Despite these contentions, it is encouraging that Yehezkel and colleagues work reveals that the adapted states can be consolidated between sessions. This has also been shown by researchers investigating the effects of multiple daily sessions on repetitive exposure to an adapting stimulus (Dong, Gao, Lv & Bao, 2016; Dong & Bao, 2019). Dong and colleagues (2016; 2019) have shown that perceptual bias induced by motion adaptation can be altered through multiple daily adaptation sessions. In both experiments, subjects repeatedly adapted to a random dot kinematotogram (RDK) in each daily session and reported the direction of a subsequent test RDK that either contained motion that moved in the adapting direction or the opposite (180°) direction. The motion coherence of the test RDK (noise vs signal) was controlled using the nulling percentage method. The authors also measured the duration of the MAE at the end of each session (Mather, Verstraten & Anstic, 1998). In both experiments, the MAE when measured using the nulling method, significantly reduced over repetitive training. The MAE duration also reduced over training however this did not reach statistical significance (Dong et al., 2016; 2019). This reduction also transferred across both the retinotopic and spatiotopic locations as long as the trained adapting direction was consistent (Dong et al., 2019).

Dong and colleagues concluded that the Stimulus-Model Comparator theory (Sokolov, 1960) was a likely candidate that explains their result. Under this model, the repetitive exposure of a stimulus creates an internal model within the nervous system, of the expected stimulus. Each time this stimulus is presented, the internal model of what is expected inhibits the response. This model would explain why at the start of training, perceptual bias following adaptation was at its strongest. However, with more repetitions of the adapting stimulus, the internal model led to increased inhibition on the response and thus, less adaptation as the training period continued.

There are however exceptions to Dong and colleague's account. For instance, other research involving the MAE has shown that passively adapting participants to a RDK stimulus for several daily sessions whilst performing a direction discrimination task, alters the discrimination performance for motion detection in the trained adapted state, but the strength of the perceptual bias remains the same (McGovern, Roach & Webb 2012). In another informative investigation related to state-dependent perceptual learning, McGovern and colleagues (2012) trained participants for 9 days on a motion discrimination task whilst keeping them in a consistent adapted state (adaptation to two fixed directions of motion - ±20° vertical). Before and after training, participants performed the same discrimination task and measured the duration of the MAE under a range of different adapted states (0°, ±10°, ±20°, ±30°, ±40°, ±50°). The authors results revealed that participants task performance improved dramatically for the trained adapted state, meaning that they were getting better at discriminating motion whilst under an adapted state. In addition, this gain in sensitivity resulted in a deterioration to performance for the same task but in untrained adapted states. Interestingly, the adapted states that before training were a benefit to performance (0°), was now a detriment to performance. The MAE duration however, in either the trained or untrained adapted states was not altered throughout training and even when measured after each daily session remained consistent. This suggests the opposite to Dong and colleagues (2016; 2019) work, as the adapted state bias remained constant whilst perceptual sensitivity improved over the course of training.

This finding is also supported by Petrov and van Horn (2012), who measured whether training on a direction discrimination task altered the strength of the MAE in the trained direction relative to a control direction. In their measurements of perceptual bias, they measured the duration for both static (sMAE) and dynamic (dMAE) MAE, which are two types of aftereffects resulting from motion that appear to arise from different levels of the motion processing pathway (Nishida and Ashida, 2000; Nishida and Sato, 1995). The authors results revealed that training on a direction-discrimination task caused large performance improvements on the task and no significant changes to the MAE duration when measured post training.

In combination, these results provide a mixed picture as to whether adapted-state bias is modulated. Before discussing this further, there are a number of important methodological differences between these experiments that are of note. The first is that participants in Petrov and van Horn (2012) were not trained in an adapted state. Therefore, it is possible that they could have measured some form of decrease to the MAE if they had tested subjects for their discrimination performance in a consistent adapted state. A second difference is that they measured perceptual bias after the training period, whereas Dong and colleagues (2016; 2019) and in a follow-up experiment in McGovern and colleagues work (2012), measured it after each daily session. Potentially measuring bias regularly over a training period, opposed to after training, is important for decreased adaptation aftereffects. However, why reductions would be found in one situation and not the other is not immediately clear.

A final difference between the studies is how perceptual bias was quantified. The MAE duration, which was used by McGovern, Roach and Webb (2012) and Petrov and van Horn (2012), is a commonly used method for measuring perceptual bias having been used for many decades (for review see: Mather, Pavan, Campana & Casco, 2008). However, measurement of MAE requires a subjective judgement on whether a stimulus is moving or not, which may be susceptible to participant motivation, bias or attention. The nulling method however may not be as susceptible to such factors as it relies on the participant making a binary judgement on a single stimulus. In addition, these two methods may activate distinct neural populations. For instance, the MAE duration may rely on low-speed units (Verstraten, van der Smagt, Freidricksen & van der Grind, 1999; Verstraten, van der Smagt, & van der Grind, 1998), whereas the nulling method may be engaging different sensory units involved in motion processing (see discussion: Castet, Keeble & Verstraten, 2002).

The research laid out in this section highlights the need to explore further the relationship between perceptual learning and adaptation within a common framework – state-dependent perceptual learning. If mechanisms that drive neural plasticity based on long-term experience (i.e. perceptual learning) alter the neural plasticity of short-term neural plasticity (i.e. adaptation), then how is this learning manifested? In Yehezkel and colleagues (2010) and Dong and colleagues (2016; 2019), learning in their paradigms is manifested as reductions in perceptual bias. On the other hand, in McGovern and colleagues (2012) work learning is manifested as an improvement in perceptual sensitivity. Both forms can be rightly viewed as form of learning, as ultimately visual functioning is improving over time. What is needed is a paradigm where both sensitivity and perceptual bias are measured at the same time each day whilst in a consistent adapted state. This would help understand the relationship between bias and sensitivity during adapted state learning – a question that the experiments in this thesis aim to address. It would also help this literature to measure both MAE bias under the nulling method and duration, to better understand these forms of MAE bias under repetitive training.

Where Passive Stimulus Exposure enables Perceptual Learning

As described in section 1.1.2, perceptual learning is a process that requires practice and experience to induce long-term perceptual sensitivity improvements. Many experiments have shown that we become more sensitive to the visual attributes we are trained on (i.e. what is task-relevant). However, what evidence is there to suggest that learning may be induced through adaptation-like exposure, when there is no repetitive training on a task requiring attention (i.e. when it is task-irrelevant)?

An interesting example of where perceptual improvements may be induced through repetitive adaptation like-exposure comes from Watanabe, Nanez and Sasaki (2001). In this example, participants were trained on a letter identification task, whilst being passively exposed to random-dot motion that on some trials throughout training contained 5% coherent motion. The motion direction was consistent throughout training. Before and after training, participants were tested on the same motion stimulus, however this time were tasked with detecting coherent motion which could appear in any direction. This motion was either 0%, 5% or 10% coherent within the dot stimulus. Watanabe and colleagues (2001) results found that after training coherence thresholds at the exposed direction and +-45° from it significantly improved when it was 10% coherent. Motion detection for 5% and 0% did not show any significant differences between 5% and 0% coherent motion for any direction.

In follow-on experiments, the same group using a similar design aimed to understand how passive this type of exposure-based perceptual learning was (Seitz & Watanabe, 2003; Tsushima, Seitz & Watanabe, 2008). In the first of these experiments they matched the 5% coherent motion in the random dot motion with signal letters in the letter identification task (Seitz & Watanabe, 2003). Their results mirrored their earlier report that detection thresholds improved through exposure, however they showed that detection performance improvement was positively correlated with the presentation of the target letter. In the second experiment, they showed that the beneficial effects of exposure disappeared when the stimuli was attended to during training (Tsushima et al., 2008), thus highlighting that this type of learning was dictated by passive exposure.

One issue with these accounts however is how they quantify the exposure-based performance improvements. In the studies mentioned above (Watanabe et al., 2001; Seitz & Watanabe, 2003; Tsushima et al., 2008) they use percent correct as their measure of performance. Many other studies investigating perceptual learning have used discrimination or detection thresholds which are determined using various psychophysical methods such as a staircase or method of constant stimuli (see Methods section for more details). The advantage to these methods is that the performance level under these procedures can be fixed (70% or 79% correct) in advance, based on how many correct or incorrect responses there are from each individual participant for an increase or decrease in stimulus intensity. In addition, it provides a measure of physical performance where one can generate the absolute difference along that physical attribute that the participant is able to reliably detect or discriminate. Using percent correct however does not yield this advantage and any changes in performance can be masked by ceiling or floor effects in participants performance, which make it difficult to understand how the performance level varies.

Many other investigations have now examined this type of exposure-based learning in what has become referred to as ‘training-plus-exposure’ or ‘double-training’ paradigms (Xiao, Zhang, Wang, Klein, Levi & Yu, 2008; Zhang et al., 2012; Wang, Zhang, Klein, Levi & Yu 2012 Harris, Gilksberg & Sagi, 2012; Szpiro & Carrasco, 2015; Szpiro, Wright & Carrasco, 2014). In one of these studies, researchers conducted daily exposure sessions (10 days) followed by orientation discrimination tests using the JND (Gutnisky et al., 2009). During the exposure sessions, participants fixated on a dot on the screen whilst oriented Gabors in three separate spatial locations (attended, unattended and control) rapidly flashed on the screen. At the attended and unattended locations, the Gabors alternated between orthogonal orientations and 10% of the time were presented at a lower contrast to the rest. Participants were tasked with detecting the low-contrast Gabors at the attended location and were asked to ignore the unattended location. At the control location, random orientations were exposed. During the orientation discrimination tests at the three locations, there was a gradual improvement in discrimination performance (JND) across all locations over the daily sessions. However, performance at the attended location was significantly more improved across session than at the unattended and control location.

In another investigation, Tanaka and colleagues (2007) explored the relationship between perceptual learning and prism adaptation. In their experiment, participants either adapted to prism reversing goggles for up to six consecutive days (adaptation period) or performed the exact same task over 6 days without prism goggles. In the group that wore the prism goggles, they wore them continuously throughout the day except for the sleep period (blindfolded) and each daily experimental session where their detection performance for a Gabor was measured in two locations in space. Before the adaptation period, participants performed a pre-adaptation learning session where learning was facilitated on the trained side of space. After the adaptation period, performance was tested on both sides 3 months later. The authors results showed that during prism adaptation learning was persistent on the trained side of space and it transferred to the untrained side of space during the adaptation period. In the group that performed the task without prism adaptation, learning persisted on the trained side but not the untrained side of space (Tanaka, Miyauchi, Misaki and Tahiro, 2007).

Together, these experiments seem to suggest that sensory adaptation dictated by passive stimulus exposure, enhances perceptual learning as it reduces the constraints of learning specificity (Gutnisky et al., 2009; Tanaka et al., 2007). However, it is difficult to discern the exact relationship between learning and adaptation under these paradigms. For instance, it is entirely conceivable in both cases that the performance improvements were a consequence of repetitive testing – i.e. perceptual learning through training as opposed to adaptation. As an example, in Gutnisky and colleagues, the perceptual improvement which their results showed across session and stimulus location, could have been a consequence of learning that happens between each daily orientation discrimination measurement.

Another potential explanation is that the performance improvements for orientation discrimination were facilitated by performing contrast judgements for the same stimulus at different orientations close to the orientation discrimination tasks discrimination boundary. This may appear to be a peculiar explanation for Gutnisky and colleagues result, however stimulus-roving (for review see: Tartaglia, Aberg & Herzog, 2009), which is a field within perceptual learning where different stimulus types and tasks alternate between trials within a block, has in some instances been shown to facilitate perceptual performance between the tasks (Kuai, Zhang, Klein, Levi & Yu, 2005). Importantly, the experiments discussed so far (Gutnisky et al., 2009; Tanaka et al., 2007, Watanabe et al., 2001; Seitz & Watanabe, 2003; Tsushima et al., 2008) could better try to understand this adaptation period if it truly benefits perceptual performance. Thus, this literature would benefit from research that can capture the perceptual state during adaptation and attempt to understand how this may facilitate learning performance.

If perceptual learning can be facilitated by prolonged exposure to some stimulus dimension, then understanding the adapted-state and how it may be facilitating performance would help further our understanding of adapted state perceptual learning and the functional relationship they may share. What is needed is an experiment where two task dimensions share the same multidimensional space, thus creating a paradigm where a participant can be cued to attend to one task dimension (i.e. training) whilst being passively exposed to the other (i.e. adaptation). If the participant is cued to attend to one attribute first before the other, then we might expect that this exposure would build up overtime and facilitate their learning if the task is switched. However, if adaptation exposure does not facilitate learning then we might expect learning to be constant through time or only improve through training.

1.2 Adaptation, Perceptual Learning and the Visuo-Motor System

1.2.1 Definition and Introduction to the Visuo-Motor Systems

Vision is an active process that is caused by our visuo-motor (self-initiated) actions. This makes sense since we are subjected to a constant stream of sensory information, a sizeable portion of which is caused by our own action. However, as Goodale (1998) comments in an early review of this literature, there is a sharp division between the way research investigates the visual and motor systems, with most research focussing on these two systems separately. However, one interesting and emerging debate, that is attempting to bridge this gap, asks how embodied visual processing and visual decision-making is to motor action (for review see: Lepora & Pezzulo, 2015). This debate is concerned with whether decision making is a process of accumulating evidence before passing a decision on to the motor systems, or whether the motor systems are a part of the visual decision-making process? In this section, I review some evidence considering how this information is processed in the brain.

The classic model of information processing in the brain is based on a hierarchical organisation, where feedforward connections from early level areas (i.e. the dorsolateral geniculate nucleus and V1) pass information up through increasingly higher level cortical areas (Van Essen, 1979). Other influential theories have expanded upon this by proposing that there are two pathways within this hierarchy that process sensory information. These are the ‘what’ pathway (dorsal), that processes information about an object’s features and the ‘where’ pathway (ventral), that processes information about spatial location (Ungerleider and Miskin, 1982). Whilst this is still the cornerstone of our understanding of visual processing, some argue that we understand only a fraction of how information is processed in early level visual areas (Olshausen & Field, 2005) and passed on through the visual hierarchy (Carandini, Demb, Mante, Tolhurst, Dan, Olshausen, Gallant & Rust, 2005). In addition, some research has shown that these separate streams may not be purely visual pathways (Snyder, Barista & Andersen, 1997). For instance, Snyder, Batista and Andersen (1997) recorded from cells in the posterior parietal cortex (part of the dorsal stream) that is known to have visual receptive fields for visual information in the periphery. Their results indicated that many of the neurons recorded maintained their activity long after the visual stimulus had been removed and remained active depending on the type of motor action the animal intended to make.

Other evidence comes from Law and Gold (2008), whose research was discussed in section 1.1.2 of this chapter. They measured the response properties of neurons in MT and LIP and showed neural response change in LIP was significantly associated with visual perceptual improvement. Whilst the exact mechanism of LIP is not known, it is interesting that performance improvements were associated with an area known to produce saccadic movements – i.e. the motor movements needed to perform the task. This evidence is perhaps indicative of visual perceptual learning being more embodied with the motor systems involved in executing the motor behaviour for the task.

There are three alternative modelling approaches that explain how visuo-motor decisions are made and how these systems may interact. Serial models propose that visuo-motor systems are independent and are part of a serial process (Newell & Simon, 1972). Thus, according to these types of models a decision regarding a visual stimulus is reached first and then the action (e.g. pressing a button to report your choice or walking towards an object) is executed after. An alternative model also views vision and motor systems as distinct processes; however, theorises both as being deployed in parallel. In this account known as the continuous flow model (Coles, Gratton, Bashore, Eriksen & Donchin 1985), the ongoing decision process is continuously transferred to the action component, which allows action to start before completion of the visual decision. In comparison to serial models, action is not just accounted for as something that is instantaneous but as a process that also requires time to execute and change should sensory information needed to reach a decision be altered.

Another alternative model however considers action as an embodied part of the decision-making process and not as separate components. In this model referred to as the Embodied Choice model (Lepora & Pezzulo, 2015), the brain prepares several action plans in parallel whilst visual information is processed and has a continuous feedback flow between the action stage and decision process. The ‘correct’ action is eventually chosen and executed as the movement with the least amount of cost. The basis for this model is in the ecological nature of action and decision making in the real-world, whereby often visual information may change and we will have to change our actions in light of this. Rather than the visual and motor systems being viewed as distinct systems, the motor system and intended actions are a part of visual processing and together form dynamic decision-making.

Ultimately, how the visual and motor systems interact to form subsequent action is still a matter of debate. In addition, considering the sheer number of studies that have studied visual and motor systems in isolation, bridging the gap between them will require methodologies and theories to come together and form novel ways of understanding this topic. One interesting similarity between the visual system and motor system is that both can adapt and learn (for review see: Bastian, 2008). For example, within the motor literature, it is theorised that motor adaptation is a form of motor learning, where the repetitive adaptation to a novel motor calibration (i.e. the locomotor pattern of riding a bike for the first time), leads to an error-driven learning process of this new motor calibration that can be stored and immediately used in the appropriate context (Martin, Keating, Goodkin, Bastian & Thach, 1996). Another form of motor learning can be viewed as an additive process where many different component movements are combined, simultaneously or in a series, to generate and learn a completely new movement (for review see: Doyon & Benali, 2005). Whilst interesting similarities exist between the neural plasticity involved in motor learning and adaptation and the visual systems equivalent (as detailed in Sections 1.1.1 and 1.1.2), this thesis will focus on how our actions may influence our visual processing.

Some evidence suggests that the nervous system constantly builds expectations about the sensory consequences of our own actions, and subsequently alters perceptual processing. A prominent example is the cancellation of visual motion during saccadic (Bridgeman, Van der Heijden & Velichkovsky, 1994) and pursuit eye movements (Haarmeier, Bunjes, Lindner, Berret, & Thier, 2001), as we constantly move our eyes but perceive the world as stable. Similar sensory consequences from our own actions have also been shown in other sensory modalities including auditory (Aliu, Houde & Nagarajan, 2009) and tactile (Blakemore, Wolpert & Frith, 1998; Hesse, Nishitani, Fink, Jousmaki & Hari, 2010; Tsakiris & Haggard, 2003) with the general finding that stimuli are perceived as less intense (sensory attenuation) when caused by self-generated action compared to when they are viewed passively.

One hypothesis for why self-initiated action is perceived as less intense comes from predictive forward models (for review see: Waszak, Cardoso-Leite & Hughes, 2012). Under these models, sensorimotor systems predict the future state of the system which help to interpret sensory information and thus shape perception. For instance, an upcoming visual perception, will be based on an additive or subtractive (cancellation) combination of the current visual information, efferent (outgoing) motor signals and expectations based on our knowledge about the likely sensory consequences of an action in a particular environment based on past experience (Kording & Wolpert, 2006). An explanation for why visual stimuli from self-generated action are perceived as less intense, is that when the action is prepared it triggers a forward model which activates perceptual areas in the cortex. This additive or subtractive activation makes the incoming sensory signal less intense, compared to situations with no prior expectations or an incorrect expectation. One proposed functional role is that it may improve the perceptual performance in the detection of other events (Bays & Wolpert, 2007).

These predictive coding frameworks (Kording & Wolpert, 2006; Rao & Ballard, 1999; Clark, 2013), suggest that stimulus representation in the early visual cortical areas is modulated by top-down prediction signals conveyed by higher-level brain regions involved in motor processing. Feed-forward information from early-level areas carries signals to the high-level areas in the form of a prediction error, which is the difference between what is predicted and the actual sensory input. Any error, or mismatch between the prediction and sensory input is then used to update the model. According to Khan & Hofer (2018), the crux of these models is the relationship between top-down predictions and early level sensory representation is that they may be flexible. For example, depending on the certainty of prior expectations (predictions), which may be weak if the task is novel, or the reliability of early stage sensory inputs as well as other task demands.

An interesting test of these predictive coding frameworks comes from two studies (Lally, Frendo & Diedrichsen, 2011; Christensen, Ilg & Giese, 2011) which compared the ability to detect visual motion following self-generated actions of the hand and arm. In Christensen and colleagues work, participants were required to detect motion that was linked to their arm movement. This motion was embedded in a field of moving distractors. Whereas, in Lally and colleagues work, participants self-generated arm movement either followed the visual distractors in a dot motion stimulus or did not match the arm movement. The general hypothesis in both studies was that the visuo-motor system would use the internal expectation created by their motor action to facilitate visual detection performance in a task dependent manner. For instance, in Christensen and colleagues (2011) the to-be-detected motion was linked to the movement of the participants arm, and this signal was embedded in noise. Whereas, the to-be-detected motion in Lally and colleagues (2011) was not related to the movements of the participant, instead the participants movements controlled the noise within the stimulus. In combination, the results suggest that the visuo-motor predictions can attenuate or facilitate the perception of self-generated stimuli depending on the type of task.

Other research has shown that these visuo-motor expectations can also alter the appearance of subsequent stimuli after prolonged practice (Bompas & Regan, 2006a; Bompas & Regan, 2006b; Richters & Eskew 2009). In research by Richters and Eskew, participants performed leftward eye movements toward a red stimulus and rightward eye movements toward a green stimulus repetitively over a period of training. Before and after training, each participant's colour judgements were assessed by asking them to compare the colour of two stimuli after either leftward or rightward saccades. Their results revealed that after training each participants PSE for colour judgement shifted depending on whether they made a right or leftward saccade. This bias in perceptual judgement was also a truly visuomotor effect as associating the same colours with an auditory tone did not produce the same effect (Richters & Eskew 2009). Other research from the same lab, has shown a similar effect when associating coloured stimulus with a leftward or rightward movement of a joystick, demonstrating that the visuo-motor nature of this effect is not confined just to eye movements (Richters, 2008).

Overall, the literature surrounding the visual-motor systems and how self-initiated action shapes our visual processing is vast. In addition, it has been researched through many diverse techniques. Whilst many questions about the influence of self-initiated action on visual processing remain open to debate, there are several interesting avenues of research that this thesis will aim to address. The next section of this report will explore each of these further.

1.2.2 Vision as an Active Process: Visual-Motor Systems, Adaptation and Perceptual Learning

As detailed in section 1.2.1, our visual perception can be strongly shaped by our self-initiated actions. A common finding is sensory suppression, where when the stimulus is self-initiated it appears as less intense. According to predictive forward models (Waszak, Cardoso-Leite & Hughes, 2012) this may be due to expectations that the visuo-motor system has regarding our self-initiated actions and future state, which may function to improve the perceptual performance in the detection of other events (Bays & Wolpert, 2007). This raises some intriguing questions for how the visuo-motor system as an internal state factor may affect visual processing, in particular adaptation and perceptual learning.

A natural question that deserves further investigation is whether an adapted state can be altered through self-initiated action. Often in adaptation paradigms it is the passive exposure to a stimulus that leads to short-lived perceptual bias. The perceptual bias resulting from motion is a classic example of adaptation in the laboratory (see section 1.1.2 for further detail), in addition motion is a perceptual feature that we are subjected to in our everyday lives. However, in the real-world there are many instances where we are exposed to motion but without the resulting perceptual aftereffects that are found in the lab. There are several differences between how we might experience motion in the laboratory compared to the real world. The first may be at the stimulus level and that often assumptions regarding adaptation are based on unnaturalistic stimuli (Felsen & Dan, 2005). The second, which will be explored in this thesis, may be how we experience motion in the real-world as it often involves us physically moving through space under our own control and direction.

Some early experiments investigated the relationship between visual motion and movement, by examining whether the congruency with participant altered the resulting perceptual bias to motion (Pavard & Berthoz, 1977; Harris, Morgan & Still, 1981). In a study by Harris and colleagues (1981), participants stood on a trolley whilst viewing a dot-motion stimulus and were pushed by the experimenter in either a forward or backwards direction. The stimulus either contracted or expanded during these movements before a test stimulus appeared. This created conditions where the physical movement was either congruent (i.e. moving in a forward direction with pattern expanding) or incongruent (i.e. moving in a forward direction with pattern contracting) to the viewable motion. The results showed that congruency between viewed and visual motion caused perception to be less biased when moving forward. When the opposite congruence was presented (travelling backwards with contracting stimuli) the adapted sate was still biased. Whilst Harris and colleagues (1981) concluded that this was a consequence of vestibular input (movements that involve the head moving through space), the fact this would only work for one direction casts doubt on whether this is the only factor that would alter the perceptual bias.

Other researchers have followed on this research by investigating whether self-initiated vestibular input would alter adaptation (Bai, Bao, Zhang & Jiang, 2019). In this experiment, participants wore a head-mounted display whilst seated on a rotating chair that they could rotate with their legs. Their experimental protocol tested congruency between visual motion and head rotation, and participants either initiated the movement or were spun on the chair by the experimenter. The authors results indicated that the MAE lasted a significantly less time when the subjects rotated than when they were stationary. However, this effect remained regardless of congruency and whether it was self-initiated or not. Therefore, it is difficult to understand the exact contribution that self-initiated action might have had on the adapted state and the resulting adapted state aftereffect.

Interestingly, although no experiment has directly tested the effects of self-initiated action on adaptation, a few experiments have investigated how manual control of a stimulus alters visual processing (Ichikawa & Masakura, 2006; 2013; Dewey & Carr 2013; Lally et al., 2011; Christensen et al., 2011). Aside from the examples highlighted in section 1.1.3 of this chapter, Ichikawa and Masakura (2006) investigated the flash lag effect under several conditions that differed in degree of observers control of the stimulus. The authors found that when participants controlled the stimulus using a computer mouse, the flash-lag effect was significantly reduced in comparison to passively viewing the stimulus. In another experiment, Dewey and Carr (2013) showed that self-initiated motion that was congruent with the participants action before stimulus onset, was perceived as slower in comparison to passive motion and motion that moved in an incongruent direction.

A natural progression from this research is to ask whether self-initiated control over a stimulus for a prolonged period of time would alter adapted state processing and the resulting aftereffect. This could help understand whether in the absence of vestibular input, self-initiated action (visual-motor system) has an influence over visual processing when viewing motion. A possibility would be to use the hand to control a stimulus for a period of time (adaptation period) before making a judgement on subsequent stimuli. This would add to literature that has tested this with vestibular input, predominately in the absence of self-initiated action. It would also be interesting for our understanding of how self-initiated actions may influence adaptation and the bias to visual perception.

Another question that deserves further research relates to how flexible our visual motor systems expectations are over shaping our visual perception. As discussed in section 1.2.1 of this chapter, predictive forward models assume that an upcoming visual perception will be based on a combination of the current visual information, efferent motor signals and expectations about the likely sensory consequences of our actions in the environment (Kording & Wolpert, 2004). A possible functional role for this mechanism in the brain would be to improve the perceptual performance in the detection of other events. Two experiments already discussed have tested this hypothesis (Lally et al., 2011; Christensen et al., 2011) and in combination, their results reveal that self-initiated action may involve expectations that can be used to facilitate visual performance depending on the task. It would be interesting therefore, to ask whether these expectations that shape visual perception can be altered through training paradigms. A possibility for how to test this that this thesis will test, would be to perceptually train participants in an environment where their perceptual expectations become altered overtime though training.

1.3 Thesis Plan

One of the main aims of this thesis is to explore the relationship between perceptual learning and adaptation. For example, some experiments have shown that adaptation has been shown to improve perceptual sensitivity after longer exposure timescales which may enable perceptual learning despite no task training to the exposed visual attribute. These findings have obvious implications for the field of perceptual learning and raises the intriguing possibility that the adapted state is important for perceptual improvement. In this thesis I will explore the effects of training in a consistent adapted state by encapsulating bias and sensitivity during the training period. Rather than only measure perceptual bias (Dong et al., 2016; 2019), performance will be collected both on visual sensitivity (learning) and bias (adaptation) during state-dependent learning to understand more about the internal state over the course of learning.

This will also enable an investigation into whether adapted state exposure to a visual attribute during training improves perceptual learning sensitivity. To achieve this, experimental groups will complete training in a multidimensional environment where two visual dimensions are coextensive. Learning and exposure will be controlled using a cueing paradigm on the dimensions will be set over different timescales and orders of task training. This should help understand more about the potential benefits to adaptation that have been shown in the literature where learning generalisation and improvement is facilitated through passive exposure that is akin to adaptation.

The final area that will be explored relates to adaptation and perceptual learning when the stimuli is self-initiated through our actions. In our daily life we are exposed to many perceptual attributes, such as motion, which in many instances is caused by our self-initiated actions. Research has shown that vestibular input may be an important factor for attenuating adaptation (Harris et al., 1981; Bai et al, 2018), however no research has explored whether adaptation is modulated by self-initiated motor control. Some experiments have shown that visual processing is often attenuated (less intense) for stimuli that is under motor control compared to stimuli viewed under passive conditions (Ichikawa & Masakura, 2004; 2006; Ichikawa et al., 2010; Konkle et al, 2009; Dewey & Carr, 2013). In addition, other researchers have shown that the brain may use the signals from self-initiated actions to shape our visual perception flexibly depending on the task (Lally et al., 2011; Christensen et al., 2011). Further exploration of this research areas could add to this debate regarding motion adaptation under vestibular input and whether adapted state sensitivity and bias are altered under our self-driven actions.

In summary the experiments in this thesis addressed the following questions:

- Does training in a consistent adapted state change adaption bias or perceptual sensitivity?
- Can exposure to a task dimension, in the absence of training, improve perceptual performance?
- Can adapted state processing be altered when it is self-generated?
- Can perceptual learning alter expectations to self-generated motion signals?

Chapter 2: General Methods

The following methods section applies to all experiments reported in this thesis. Any differences from the methodology described here will be outlined in detail in the methods section for each experiment.

2.1 Apparatus

The data for experiments described in Chapter 3 and Chapter 4 were obtained out on a computer controlled cathode-ray-tube (CRT) monitor (20-inch NEC Multisync FP2141sb) with a display resolution of 1152x870 pixels. The refresh rate was 75Hz and the mean luminance was set at 85 candela/square meter (cd/m^2). The stimuli were generated on an Apple Macintosh Mac-Mini running the Psychtoolbox Software (Kleiner et al., 2007) on MATLAB 2015b. The observers head was stabilised throughout the experiment using a chin-rest 71.6cm away from the monitor. The viewable area of the monitor subtended 35° of visual angle and the angle subtended by 1 pixel was 1.62 arcminute (arcmin) or 0.027° of visual angle.

In Chapter 5, the experimental data were collected by projecting visual stimuli onto a large semi-circular screen (radius = 2.5m, height = 2m). The screen wrapped 180° around the participant horizontally and subtended 38.6° vertically. The stimuli were projected using 3 overlapping projectors. The separate feeds from the projectors were corrected for geometric distortions resulting from the projection onto a circular surface and luminance distortions resulting from the overlapping displays of two projector images. These corrections were achieved using Immersaview's Sol7 software (https://www.immersaview.com/). The refresh rate were 75 Hz and the mean luminance was set at 48 cd/m^2. The stimuli were generated on Dell computer running Psychopy (Peirce, 2007) a custom computer software written in Python. The observers head was secured using a chin-rest 2.5 m away from the centre of the screen. A wireless mouse and keyboard were used to record the responses.

In Chapter 6, the data were collected on an LCD monitor (Cambridge Research Systems Display++) with a display resolution of 1920 x 1080 pixels. The refresh rate was 120Hz and the mean luminance was set at 56 cd/m^2. The stimuli were generated on a Lenovo Thinkstation P520c computer running Psychopy (Peirce, 2007). The observers head was stabilised using a chin-rest 31.6 cm away from the monitor. The viewable area of the monitor was full field (covered the entirety of the observers vision) and at this distance 1 pixel subtended 4 arcmin or 0.067° visual angle.

2.2 Monitor Calibration

2.2.1 Gamma Correction

The input voltage from the computer that specifies a luminance to display on the output screen of a CRT monitor, is non-linear. Due to this non-linearity, it is important for the purpose of vision research that these monitors are gamma corrected to ensure that the presented stimuli can be displayed with a precise luminance values and stimuli can be presented in linear increments. Thus, the relationship between the input voltage from the computer and the CRT monitor screen luminance was characterised and linearised.

To achieve this, a program presented a rectangular test patch in the centre of the screen (size varied for the test patch depending on set-up). The luminance of the test patch was measured using a photo-spectrometer (LS-150, Konica Minolta). The test patch varied across 8 different luminance levels, starting at the lowest luminance level and finishing at the highest luminance level. These were evenly spaced between gun output values of 0 and 255. From these uncorrected luminance values (see the dotted lines in figure 2.1 for this procedure in chapter 3), it was possible to generate a set of gamma functions.

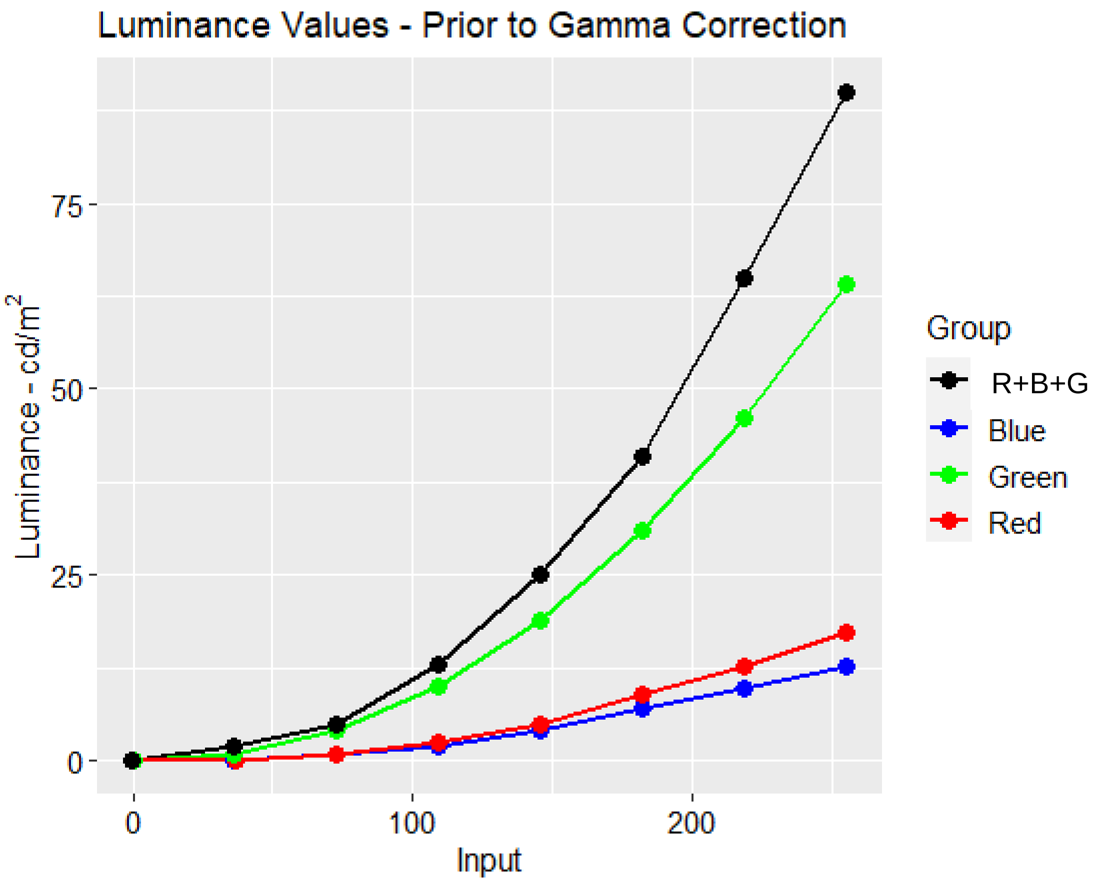

Figure 2.1 - Luminance Values Prior to Gamma Correction for the 20-inch NEC Multisync FP2141sb monitor.

The following equation was fit to the 8 luminance values in an uncorrected look-up-table (LUT), to generate the corrected gamma functions:

$$Lum = a + bI^{y}$$

where Lum is the Luminance, and parameter a is a constant. Parameter b is represented by the luminance range, I is the gun intensity requested as a decimal ranging from 0-1 and γ represented the gamma value respectively. Parameters a, b and γ were free to vary during fitting.

Having determined the gamma value γ, this was applied to create an inverse LUT using the following equation:

$$LUT_{ii} = \left(\frac{ii}{255}\right)^{\frac{1}{y}}$$

The indices ii, in the LUT ranged from 0 to 255. The resulting physical luminance is linearly related to the index in the LUT by the equation:

$$Lum = a + b\left(\frac{ii}{255}\right)$$

The luminance values were again remeasured to confirm gamma correction for each of the experiments. Figure 2.2 shows these remeasured values luminance values with gamma correction. This correction was then applied to the luminance values in all experiments in this thesis which was controlled automatically at the start of each experiment.

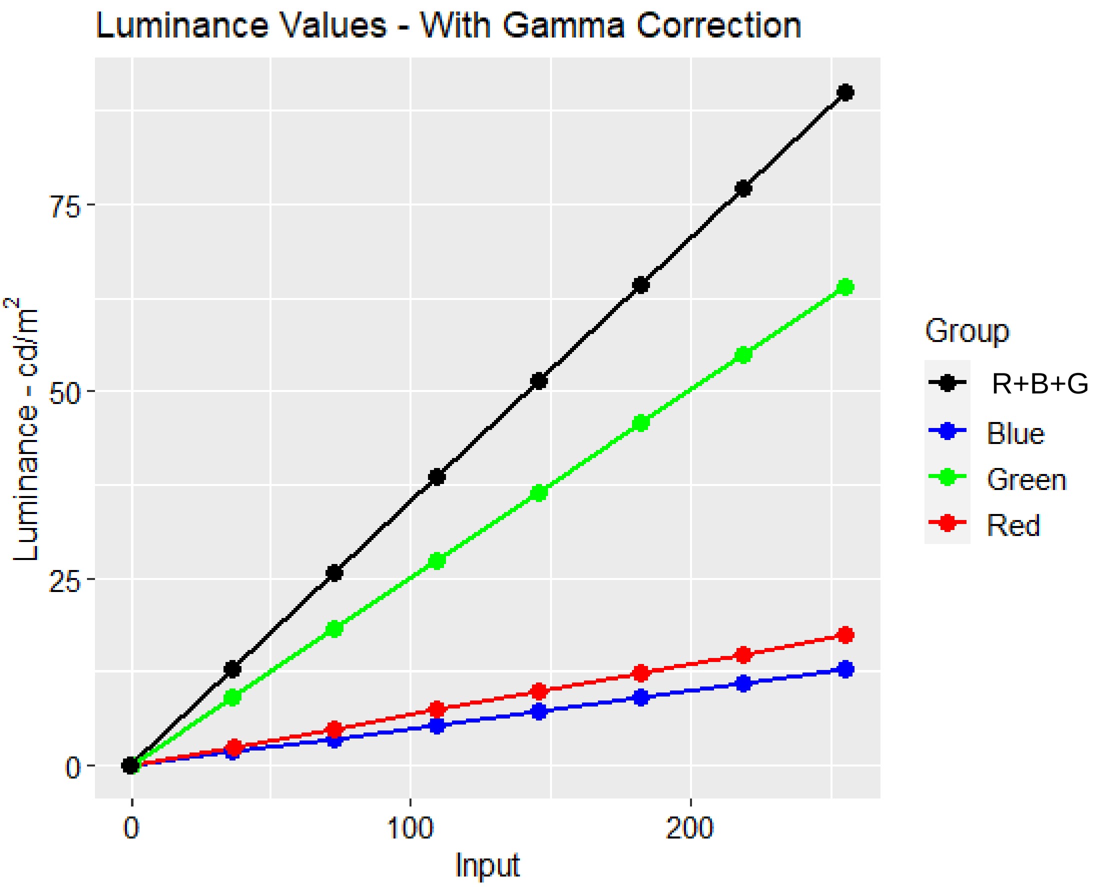

Figure 2.2 – Gamma Corrected Luminance Values for the 20-inch NEC Multisync FP2141sb monitor.

2.2.2 Spatial Calibration

To accurately display stimuli to the correct specifications, it is necessary to determine the pixel sizes of the display monitors. To confirm the spatial calibration of a monitor, one can first predict the size of one pixel by dividing the internal dimensions of the monitor (in millimetres) by the resolution of the monitor in pixels. These were first acquired for the each of the monitor set ups outlined in the table 1 below:

Screen Monitor Brand	Experimental Chapter Used	Internal dimensions/Pixel Resolution (Predicted)
20-inch NEC Multisync FP2141sb	Chapter 3 and 4	Height: 305/870 = 0.350mm Width: 406/1152 = 0.352mm
Cambridge Research Systems Display++ LCD Monitor	Chapter 6	Height: 37/1080 = 0.342mm Width: 67/1920 = 0.354mm

Table 1 – Screen Monitor and Spatial Calibration Specifications.

To confirm these measurements a program was created for all experiments using the computer programming software they were run in which displayed a test patch on the monitor. These were then physically measured using a ruler with a 1mm scale. This was repeated numerous times with different sized test patches. An example of these measurements are plotted in Figure 2.3 and were taken on the screen monitor used in Chapter 3 and 4.

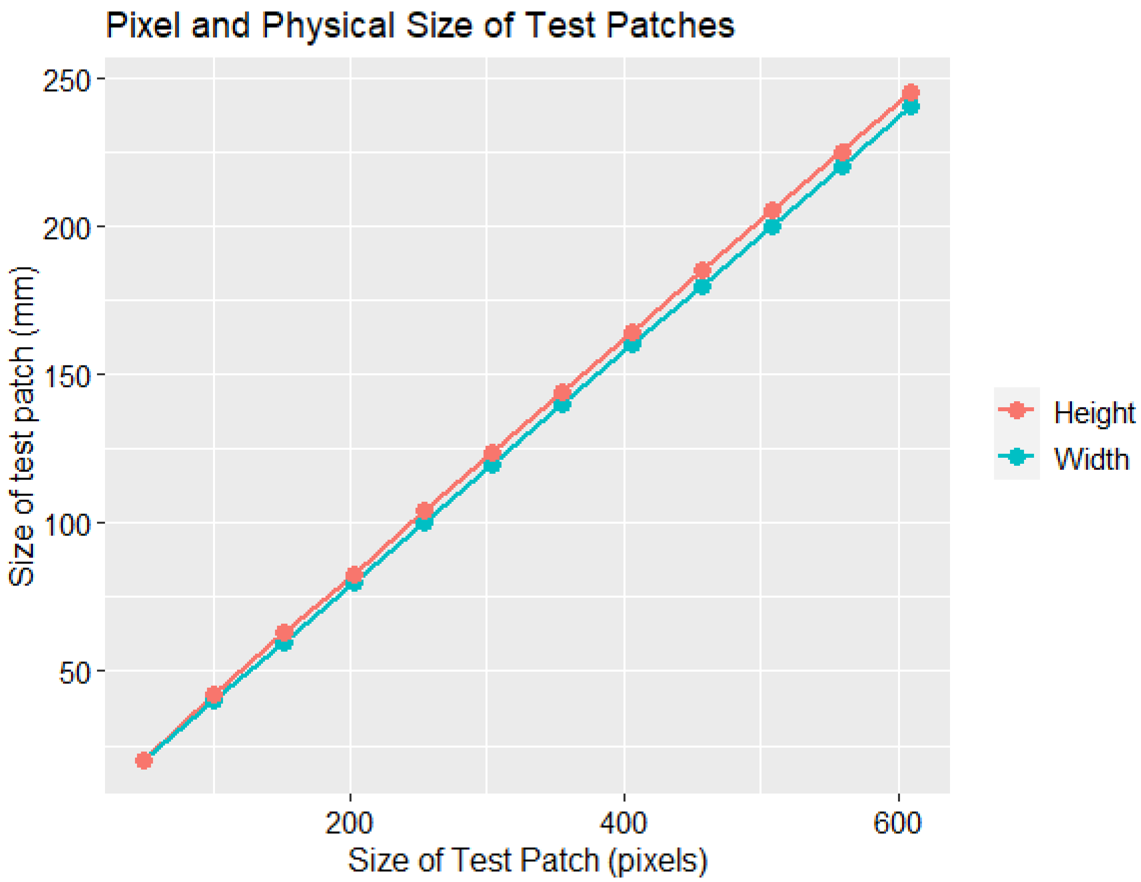

Figure 2.3 – The Pixel and Physical Height and Width Measurements of test patches displayed on the 20-inch NEC Multisync FP2141sb screen monitor.

2.3 Stimuli

Two types of stimuli were used in this this thesis; Gabor patches and random-dot kinematograms.

2.3.1 Gabor Patches

Gabor patches were used in experimental Chapters 3, 4 and 5 of this thesis to measure orientation, spatial frequency and motion direction discrimination. They can be described as a vertical sinusoidal carrier modulated on a uniform grey background (85 cd/m^2) and windowed by a two-dimensional Gaussian envelope. This can be defined mathematically in the following equation:

$$L(x, y) = Lm\left\{1 + Cp\ \cos[2\pi x f c + \theta c] \times exp\left[-\frac{1}{2}\left(\frac{xy}{\sigma xy}\right)^2\right]\right\}$$

where Lm is the mean luminance of the display, fc is the spatial frequency (in cycles per pixel), Cp is the peak contrast of the Gabor and σx and σy is the standard deviations of the Gaussian envelope. Example code for this relationship can be found in Appendix 1. An example of the type of Gabor patch used is shown in Figure 2.4.

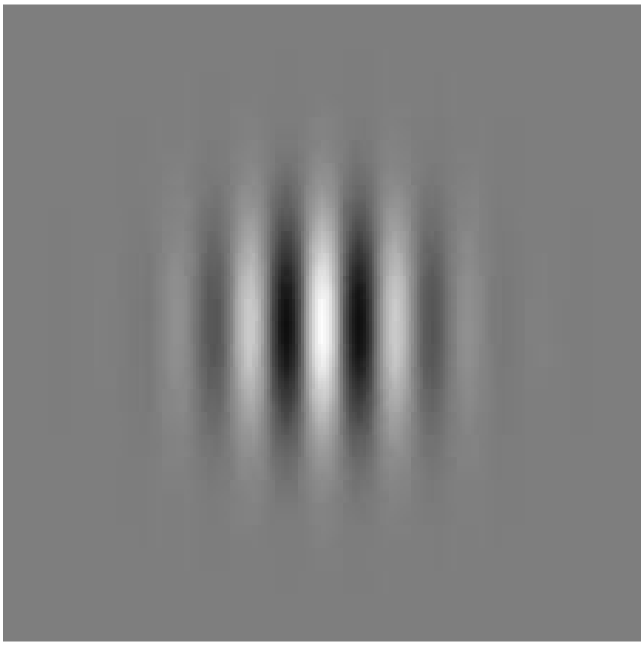

Figure 2.4 – Example of a Gabor Patch used in Chapter 3, 4 and 5 of this thesis.

The Gabor patches used in the experiments here varied according to the type of task. For example, by changing fc and altering the spatial frequency, experimenters can test spatial frequency discrimination of the visual system. There are many dimensions along which a Gabor can be used to test aspects of the visual system, details of which will be explained in each chapter.

2.3.2 Random-dot Kinematograms

Random-dot kinematograms (RDKs) were used in Chapter 6 of this thesis to measure motion coherence detection. They can be described as a field of moving dots, presented in sequence, either moving coherently in a uniform direction or in ambiguous directions, at a given speed. The number of dots, the coherence and direction of motion and their speed at which they are displaced are all dimensions that comprise a task. Example code of this can be found in Appendix 2. An example of a single frame from a RDK can be found in Figure 2.5.

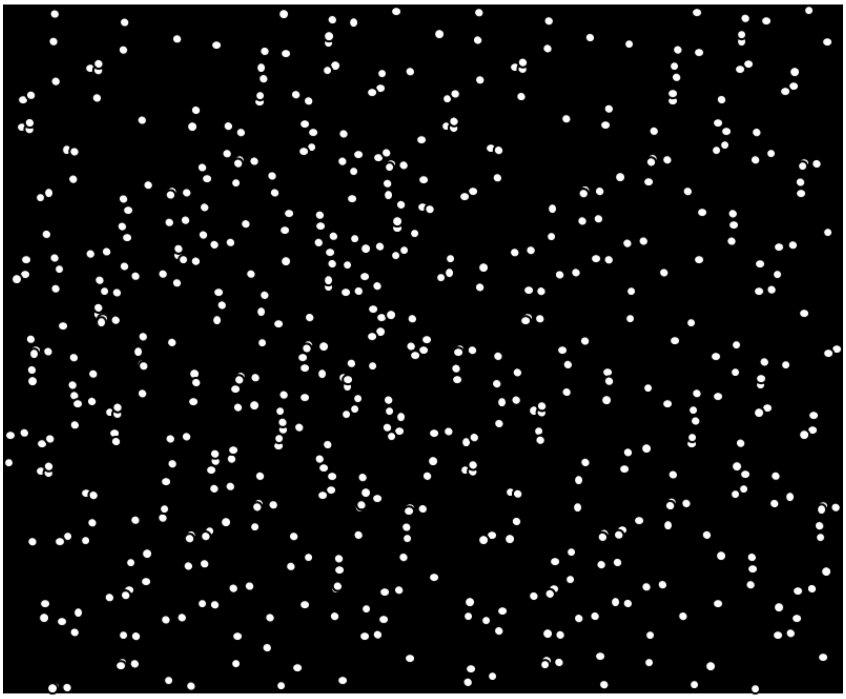

Figure 2.5 – Example of a Random-Dot Kinematogram used in Chapter 6 of this thesis.

2.4 Psychophysical Methods

Psychophysical methods provide a means of measuring the relationship between physical stimuli (e.g. luminance, frequency) and sensations/perceptions (e.g. brightness, pitch). By using psychophysical methods in experiments, it is possible to determine the rules that these sensory systems follow in generating the perception that is experienced and acted upon (Fechner, 1860). There are many fields within vision research that benefit from employing psychophysical designs. The first, Perceptual Learning (PL) uses psychophysical methods to understand how our sensory systems improve on tasks over prolonged periods of time. The second, Adaptation uses psychophysical methods to understand how our sensory systems change to the regularities in sensory information over much shorter periods of time, this is more commonly referred to as adaptation aftereffects. Both fields employ similar psychophysical design principles and can also be used to investigate both in conjunction. These methods and designs will now be described in more detail below.

2.4.1 Psychophysical Decision-Making

Psychophysical methods are underpinned by decisions and the experimental tasks used in this thesis use a forced-choice decision paradigm. Under this protocol, the observer is forced to make a choice between two possible options after each trial. In Chapters 3, 4 and 5 of this thesis the forced choice was between two-alternatives after a trial that contained the single presentation of a stimulus. For example, in Chapter 3, a Gabor was presented that was orientated either to the left or right of a decision boundary orientation and their choice was forced between whether they perceived this to be orientated more to the left or right. This was the case even if they did not feel confident in making their decision on the presented stimulus. In chapter 6 of this thesis, a forced-choice decision paradigm was used that is referred to more commonly referred to as a two-interval forced choice (2IFC) paradigm. In this type of design each trial consists of two intervals presented sequentially in the same location. One interval contains the signal and the other does not and the participant is forced to choose in which interval contained the stimulus signal.

2.4.2 Method of Constant Stimuli

Under the method of constant stimuli (MCS), a fixed set of stimuli are specified that are evenly spaced either side of a discrimination boundary. For example, a common decision boundary in a psychophysical orientation task may be 90° and the fixed set of stimuli would be orientations either side of 90° (e.g. 75°,80°,85°,95°,100°,105°). Each fixed set of stimuli are presented an equal number of times and in a random order. The frequency with which each stimulus elicits a stimulus response is counted and used to plot and fit a psychometric function (see section 2.6.2). As each stimulus is presented an equal number of times this can be a strong way of deriving performance measurements from the psychometric function. Mainly, the just noticeable difference (JND), a measure of how accurate the participant is in their responses, and the point of subjective equality (PSE), a measure of how biased they are in their responses.

A potential disadvantage of the MCS however is in the range of the stimuli used to test the observer. In fields such as PL and adaptation, where accurate measuring of the change in perceptual performance overtime is important, repetitively testing the same set of stimuli between day and session may lead to redundancy in some trials and a missed opportunity to study more important stimulus levels. For example, in a typical PL experiment utilising the MCS, asobserver's performance improves with practice, recycling the same stimulus levels over training day would be inefficient as performance accuracy inevitably improves to a point where stimulus levels are no longer challenging but *too easy*. One approach to overcome this inefficiency is to tailor the range of stimulus intensities between day or session based on past performance and is described in more detail in the next section.

2.4.3 Adaptive Method of Constant Stimuli

To accurately gauge how the participant is learning across session, performance thresholds in Chapter 3 were obtained using an adaptive MCS. This method maintains a fixed set of stimulus values that are presented an equal number of times within a training block, however the stimulus intensities between each block of trials are determined based on the performance level of the previous block (see schematic in Figure 2.6). Specifically, this was 2 JND widths either side of the decision boundary. The advantage of this method is that trials that become *too easy* through practice are eliminated and a more accurate estimate of PL abilities around the threshold boundary are obtained.

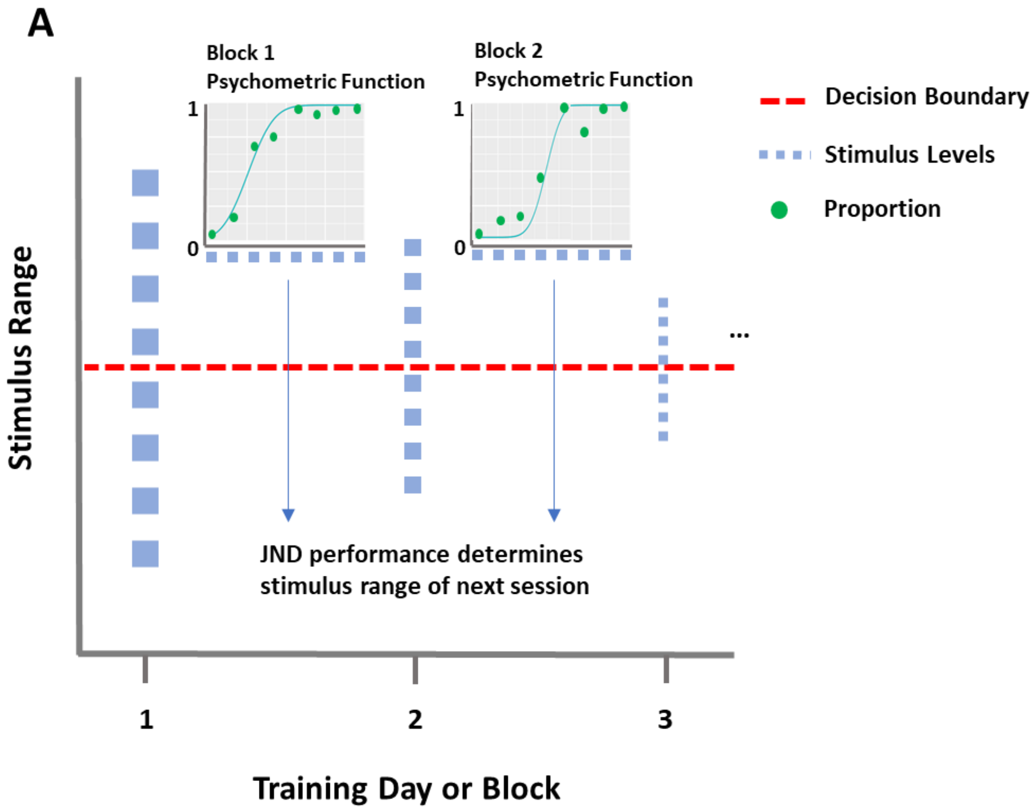

Figure 2.6 – Adaptive Method of Constant Stimuli Schematic. This type of paradigm shares many features as conventional method of constant stimuli such as the fixed set of stimuli within a block, however the stimulus values for each block is determined by the performance in the previous block.

2.4.4 Staircase Methods

Another psychophysical method used in the experimental Chapters 4, 5 and 6 of this these was the staircase method. In this method the stimulus range and current stimulus intensity for a given trial is adjusted according to how the observer is performing. Thus, the present trial is based on an observer's response history. Usually, the next trial typically becomes easier following an incorrect response from the observer or more difficult if the observer makes a string of correct responses. Under the staircase method, the experimenter has control over where the stimulus intensities start, how large the stimulus intensities can vary based on correct or incorrect answers and when the staircase should end. For a 1 - up, 1 - down descending staircase, if a correct response is made, then the stimulus intensity is reduced after every correct answer and increased after every incorrect response. A 1 - up, 3 - down staircase follows a similar rule however requires 3 correct responses before the stimulus value is decreased and only one error for the stimulus intensity to increase. The advantage of this method is that the observer's perceptual ability is captured efficiently as trials will converge on a fixed performance level. For a 1 – up, n-down staircase this equates to a percent correct of $\frac{100}{2^{1/n}}$.

Combinations of these staircase criteria are typically utilised and have been in experimental Chapter 6 of this thesis (see Figure 2.7 for schematic). For instance, at the start of a run a 1 – up, 1 - down staircase can be used until the observer made the first mistake. This ensures relatively rapid descent toward the threshold level making the procedure more efficient. After the first mistake the staircase can then change to a 1 – up, 3 - down staircase. This would converge at a performance level of around 79.4% (Levitt, 1971).

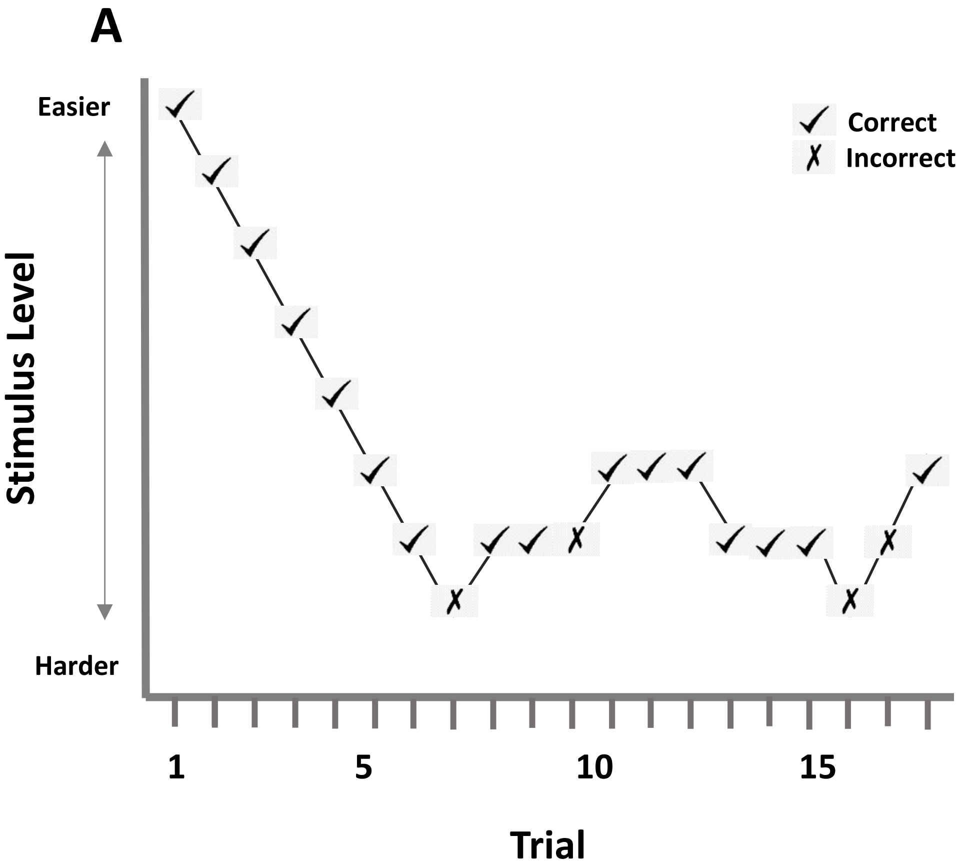

Figure 2.7 – Staircase Method Schematic. In this method observer's performance dictates the range of stimulus values over time. Here the staircase begins with a 1 – up , 1 – down staircase until the first mistake is made and then a 1 – up, 3 - down staircase is used.

In experimental Chapters 4 and 5, perceptual performance was gauged using two 1 - up, 1 - down interleaving staircases, one ascending and one descending, that can converge (see Figure 2.8). The difference between these two interleaving staircases is that responses are not classed as correct or incorrect. Rather they are traveling in opposite directions. All staircases terminated after a set amount of reversals. A reversal is where the staircase stimulus intensity reverses direction (e.g. decreases after increasing).

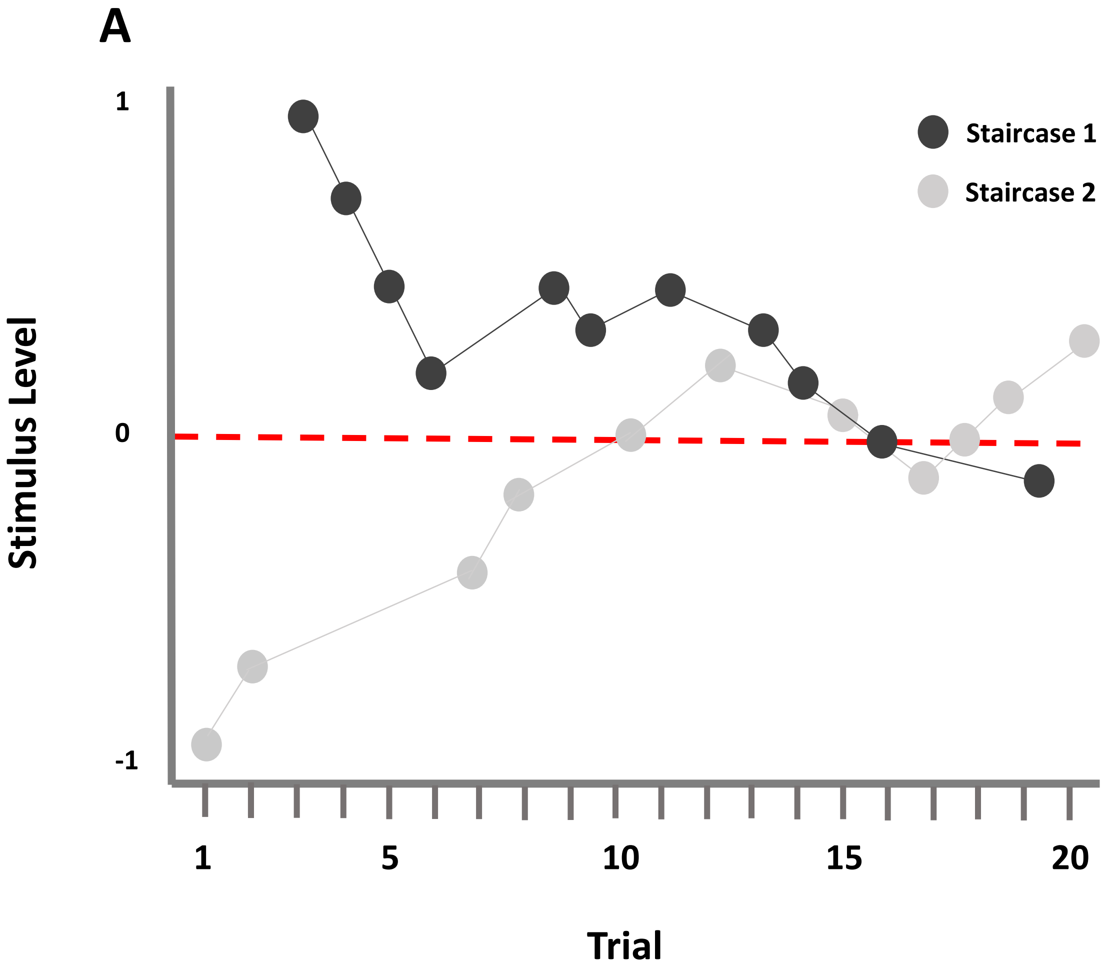

Figure 2.8 – Interleaving Staircases. In this method, two staircases are randomly interleaved between trials and depending on the staircase the observer's performance would vary the stimulus intensity. Each staircase begins with a 1 – up , 1 – down staircase until the first mistake is made and then a 1 – up, 3 - down staircase is used.

2.4.5 Adaptation Aftereffect Methods

The perceptual consequences of Adaptation can be measured using MCS and staircase methods. The main difference is that these methods are used to estimate perceptual bias, rather than sensitivity. A common form of adaptation bias that has been extensively studied are known as aftereffects, which is typically studied by exposing a subject to a visual attribute for a period of time and then measuring the aftereffect when the stimulus is removed. One well studied aftereffect that is used in Chapter 4 of this thesis, is the motion aftereffect (MAE) which measures the effects of motion adaptation. There are different methods used to test MAE which are outlined below.

2.4.5.1 Motion Aftereffect duration

A commonly used method for measuring the MAE involves presenting a stimulus for a prolonged period and then presenting a neutral stimulus in the same retinal location and asking the observer to report when the aftereffect ends. This method was used in Chapter 4 of this thesis and has shown a reliable direct relationship between the length of the initial stimulus and the length of the adaptation-aftereffect (Greenlee, Georgeson, Magnussen & Harris, 1991). However, the problem with this technique is that it is susceptible to confounding factors such as participant attention (Chaudhuri, 1990). An alternative approach is to use a nulling method approach which is described in more detail below.

2.4.5.2 Motion Aftereffect Nulling Procedure

The nulling procedure is a procedure to estimate the percentage of signal needed to perceptually null the MAE. After a prolonged period of exposure to a moving stimulus in one direction, a test stimulus is presented for a short amount of time that contains either both directions of motion superimposed onto each other or only a percentage of motion moving in the nulling direction against incoherent motion. For instance, if the adapting direction was leftward, the nulling direction would be leftward for the test stimulus. The nulling percentage procedure was used in Chapters 4 and 5 of this thesis and consisted of two superimposed Gabors travelling in opposing directions with their contrast manipulated. After each test trial, the log contrast ratio of the two superimposed gratings would either increase or decrease depending on the participants response. This can be expressed in the following equations:

$$R = \frac{c_l}{c_r}$$

Where R is the ratio contrast, manipulated in the staircase in using the logarithm base 10 $[\log_{10} R]$, c_l and c_r are the contrast for the Gabor moving to the left and the Gabor moving to the right. The contrast for each Gabor was manipulated as follows:

$$c_r = \frac{k}{(R+1)}$$

$$c_l = k - c_r$$

Where k is the total contrast which is held constant. When $c_l = c_r$ this made a counter-phasing Gabor.

2.5 Testing Procedures

2.5.1 Perceptual Learning Training Protocols

The perceptual learning experiments detailed throughout this thesis all use a similar testing protocol that can be thought of as a three-stage process: pre-training, training and post-training. Prior to the pre-training phrase, each participant underwent a very quick practice (maximum of 10 trials) of the task after an explanation of the experimental task. This was to ensure they understood the task and minimise the amount of learning. Both the pre-training and post-training phases were identical and consisted of training on a task over a single day. The training phase focussed on training on the same task over several days. Where possible the training was completed on as close to consecutive days as possible. By having the pre- and post-training phases, it is possible to determine whether there was any improvements in perceptual performance on a particular task dimension and understand if there is any learning transfer.

2.5.2 Long-term Adaptation Training Protocols

Long-term adaptation experiments detailed in this thesis follow a similar protocol regime to that detailed in the section 2.5.1 above. The major difference is in manipulating the adapted state of the observer overtime and quantifying bias in observer's response. To manipulate the adapted state of the observer usually an adaptation stimulus appears on the screen for a prolonged period before each test stimulus. In the interest of keeping training blocks shorter for the participant, the adaptation periods follow a top-up manner. In Chapter 4 of this thesis, the adaptation period lasts for 30 seconds before the first trial and 5 seconds between trials.

2.6 Data Analysis

2.6.1 Psychophysical Thresholds

In the methods described above there are several ways that thresholds can be derived. One method for deriving a threshold based on the described methods is by taking the geometric mean of a certain amount of reversals in a staircase design (see section 2.4.4). The geometric mean is the nth root of the product of the n terms. The geometric mean is appropriate when logarithmic steps are used in staircases and the last n reversals are used because they are likely to be near the threshold level. In experimental chapter 4 of this thesis, the geometric mean of the last 6 reversals of the two interleaving staircases was taken as the accuracy threshold for the session which can be expressed in the following equation:

$$\left(\prod_{i=1}^{6} s1_i \ , s2_i\right)^{\frac{1}{12}} = \sqrt[\frac{1}{12}]{s1_1 s1_2 s1_3 s1_4 s1_5 s1_6 s2_1 s2_2 s3_3 s4_4 s5_5 s6_6}$$

Where $s1$ and $s2$ are the staircases 1 and 2, and $[s1_{1-6} s2_{1-6}]$ are the last 6 reversals of each staircase.

2.6.2 Quantifying Perceptual Learning effects

All experiments adhered to the following conventions regarding quantifying perceptual learning effects. For each stimulus intensity presented during either an adaptive staircase or method of constant stimuli procedure, the outcome of a particular response was counted and then divided by the total amount of presentations at that stimulus intensity level. Psychometric functions were then fit using maximum likelihood method of estimation of the form:

$$\mathrm{P}(R) = g\,\frac{1}{\left(1+\exp\left(\frac{\alpha - x}{\beta}\right)\right)}$$

Where $\mathrm{P}(R)$ is the proportion of making a particular response computed as a function of the particular stimulus intensity level, x. The fitted parameters α and β, correspond respectively to the PSE (50% point) and the slope of the function also known as JND. From this fit, the JND was used and represented the degree of accuracy that the observer could distinguish either side of the discrimination boundary.

2.6.3 Permutation Testing

Permutation testing is a method that is used to test the null hypothesis of no difference between groups or conditions. This method resamples the observed data (without replacement) under the assumptions that the null hypothesis is true. For example, in a design that has a treatment and a control group, the resampling would randomly swap the observed values in the treatment group for some of the observed values in the control group. This changes the means of the treatment ($\bar{x}_1$) and the control group ($\bar{x}_2$), making a new observed mean difference ($\bar{x}_1 - \bar{x}_2$). Repeating this procedure thousands of times and plotting each permuted mean difference builds up a permutation distribution (see Figure 2.9). By locating the mean difference from the observed data (black line – Figure 2.9) you can then show how likely or unlikely the observed mean is when the null hypothesis is true. The permutation method when testing the null hypothesis of no difference within groups is slightly different in that the procedure requires you to pair each participant within group observations and flip them. For instance, the *pre* value for one participant may become the *post* value.

If the null hypothesis is true then the mean of the observed data sets will be very similar to the centre of the permutation distribution. If the null hypothesis is not true the mean of the observed data will be close to either of the two tails of the distribution. To estimate a one-tailed P-value from this distribution you obtain the proportion of resampled values that give a result at least as great as your observed mean difference.

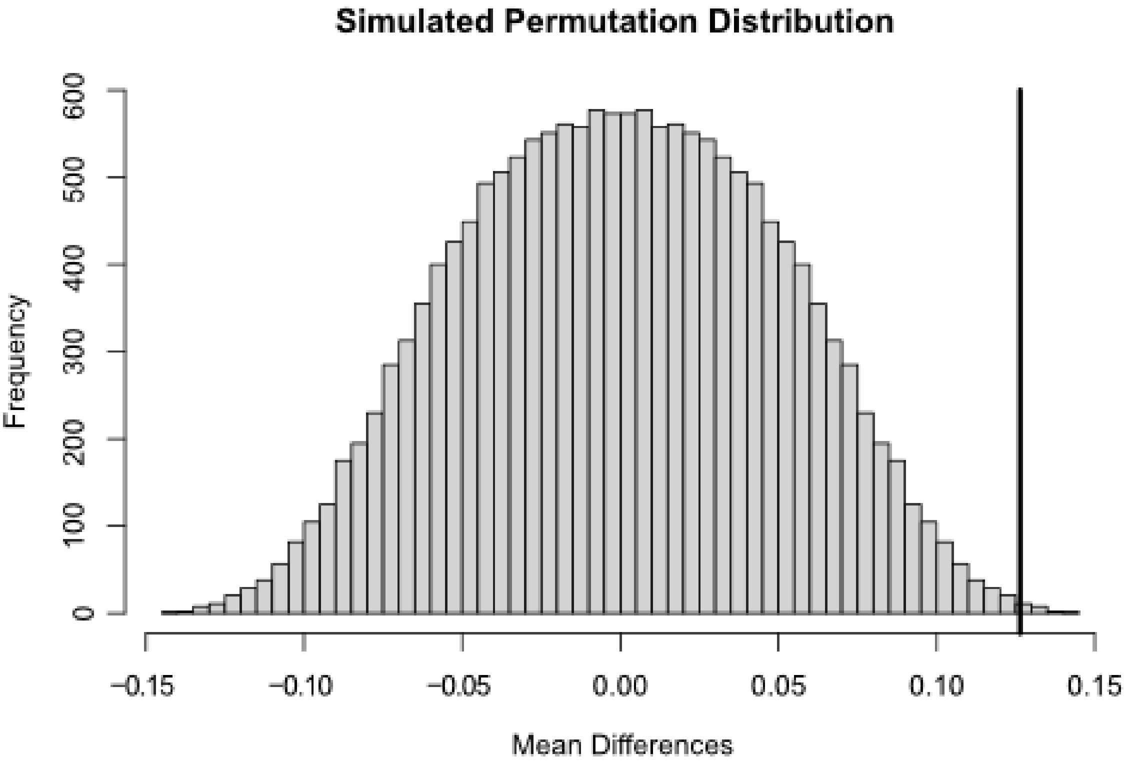

Figure 2.9 – Permutation Distribution for Simulated Data. Shows the distribution of the statistic $\bar{x}_1 - \bar{x}_2$ between a treatment group and a control group. The solid black line marks the actual observed mean difference (0.126). Its location shows that the value is unlikely to occur when the null hypothesis is true. The P-value for this would be derived from the amount of mean differences greater than the observed mean divided by the total amount of permutations.

2.6.4 Resampling Methods: Bootstrapping with replacement

The bootstrapping method is a modern resampling method that draws thousands of resamples of your data by essentially swapping the data points from your sample, collecting a central tendency or similar statistic from this distribution and then plotting the resampled statistics in a histogram form (Hesterberg et al., 2005). The bootstrap distribution approximates the same shape and size as a sampling distribution (original data) and this can be used to check normality of the sampling distribution. It is also used to calculate the standard deviation of the bootstrap distribution which is a powerful way of estimating the variation in a statistic based on the original data (Hesterberg et al., 2005, Wichmann & Hill, 2001b).

In the following experiments, we used the bootstrapping method to determine the standard error of the obtained JND values. A total of 10000 bootstrap resamples were generated from the data, a Weibull function was fitted to the bootstrapped resamples and a JND value was calculated. From these fits, the standard deviation of the bootstrap distribution and its 95% confidence interval were determined for each observer.

Chapter 3: Learning and adaptation: no habituation of the dynamic motion aftereffect measured with grating stimuli

3.1 Introduction

The visual environment is complex and full of variety, yet the human brain manages to make sense of it. Two well-studied mechanisms, known as adaptation and perceptual learning, are involved in this process. Both operate within the same areas of visual cortex (Teich & Qian, 2003) and are driven by our experience of the past. Adaptation is driven by the recent past on a short-term millisecond to second timescale (for review: Webster 2011). Perceptual learning is driven by long-term experiences on the hour to week timescale (for a review see: Gold & Stocker, 2017).

Most research has treated these mechanisms as isolated processes with distinct characteristics. Unlike perceptual learning, which is characterised by a permanent performance improvement, a core characteristic of adaptation is its ability to induce temporary biases in perception. However, only recently have researchers begun to show that these core characteristics for each mechanism might not be as independent as previously thought. For instance, like learning, adaptation has been shown to change over much longer timescales (Dong et al., 2016; 2019; Bao and Engel, 2011; McDermott, Malkoc, Mulligan & Webster 2010) and even facilitate perceptual performance (Clifford et al., 2001; Oruc & Barton; 2011; Wissig, Patterson & Kohn, 2013; Webster et al., 2015). Similarly, some perceptual learning experiments have shown that performance can get worse over training (Mednick et al., 2003), can show changes in visual appearance through training (Haijiang et al., 2006) and that simple exposure, in the absence of dedicated training, can facilitate perceptual performance improvements (Watanabe et al., 2001) analogous to adaptation. This raises the intriguing possibility that the two mechanisms share a functional relationship.

Recently experimenters have investigated the effects of repetitive adaptation over several daily training sessions (Dong et al., 2016; 2019). In their design, they explored the motion aftereffect (MAE) and quantified bias using the nulling method and by measuring its duration. The nulling method was used to obtain the percentage (coherence) of signal against the noise dots needed to perceptually null the MAE, a technique useful for measuring adaptation bias introduced by Hiris and Blake (1992). Over 8 training days, participants completed blocks of the nulling task at a specified retinal location then immediately after each block reported the duration of the MAE at the same location. Before and after this training period, participants completed the same method but at two retinal locations.

Dong and colleagues (2016) showed a significant reduction in the nulling percentage over the course of training and a decreasing trend in the duration of the MAE. In their follow up experiment (Dong et al., 2019), the authors demonstrated that this reduction could transfer across retinotopic and spatiotopic locations as long the adaptation direction remained the same (i.e. a consistent adapted state). The was interpreted as evidence that the strength of the adaptation was getting weaker as a result of repetitive training. The authors suggested that this was the outcome of both attention and habituation-like learning mechanisms and consistent with the Stimulus-Model Comparator theory (Sokolov, 1963, for review see Thompson, 2009). Under this descriptive model, the repeated experience of a stimulus creates an internal model within the nervous system, of the expected stimulus. Each time the stimulus is presented, it is compared to the internal model and if it matches responding will be inhibited. The core feature of this model that lends itself to Dong and colleagues results, is that at first the stimulus model is weak as it has not had much experience of the stimulus. This would explain why adaptation bias at the start of training is at its strongest. However, with more repetitions, the stimulus model represents the stimulus more precisely, leading to increased inhibition on the response and thus, less adaptation as training progresses.

The theory of habituation and its account of the reduction in adaptation is intriguing, however an alternative explanation is that participants were becoming more accurate at performing the task in a specific trained adapted state. Under this view, rather than the strength of the adaptor getting weaker through repetition, it is the perceptual accuracy of the visual system to accurately detect the coherence of dots within noise that is increasing through training. Under this explanation, the visual system is training whilst in an adapted state. Performance is at its worst at first because the visual system is not accurate at performing the task in that adapted state. However, through training it becomes more perceptually accurate in that trained adapted state. This explanation would also explain the lack of effect to untrained adapted states (i.e. different adapting motion directions) and unadapted thresholds after training.

Perhaps the most valuable insight into state-dependent perceptual learning and a result that is consistent with this idea, comes from McGovern and colleagues (2012). In this experiment they trained participants on an direction discrimination task whilst in a consistent adapted state. Before and after training, participants completed the same task but in different adapted states, these were for directions of motion offset symmetrically either side of vertical (0°, ±10°, ±20°, ±30°, ±40°, ±50°). The authors also measured perceptual bias in each of these adapted states after training, by measuring the duration of the adaptation aftereffect. The results for discrimination performance before training revealed that accuracy was most inhibited at ±20° and slightly facilitated, relative to unadapted discrimination performance, at 0°. This is in line with other experiments that have investigated orientation discrimination following adaptation (Kristjansson, 2011; Hol & Treue, 2001, McGovern et al., 2014; Price & Prescott, 2012). However, over the course of 9 days of repetitive training sessions at ±20°, the initial detriment to performance changed into a gain. In addition, when the authors re-examined the tuning function at the end of training, discrimination performance for other adapted directions reversed. What was a detriment to performance before training became a benefit and what was a benefit to performance before training became a detriment. Despite these changes in perceptual sensitivity however, perceptual bias in the trained and untrained adapted states remained the same as there was no change in the duration of the adaptation aftereffect.

In combination, both experiments highlight that the effects of adaptation are clearly not fixed and in fact can change. What is unclear however is the relationship between perceptual sensitivity and bias under these repetitive adaptation training protocols. For instance, in direct contrast with one another, Dong and colleagues show that perceptual bias reduces whereas, McGovern and colleagues show that it remains constant whilst it is adapted perceptual sensitivity that improves overtime. A similar finding of consistency in adapted state bias following training has also been shown by other researchers (Petrov &

Van Horn, 2012). Scrutiny of these experiments reveals several obvious differences between their methodologies that may explain the opposing outcomes, the first is in the measuring of perceptual bias (i.e. nulling method or aftereffect duration using static or dynamic test probe). Another important difference is in how adaptation bias is captured (i.e. after every daily training session or at the end of training). A final difference is how perceptual sensitivity is quantified overtime (i.e. change in nulling percentage or JND performance thresholds), which is related to the type of methodology used.

What is needed in order to bring these findings together is an experiment that can characterise this relationship between bias and sensitivity at the same time over training. One promising technique that could achieve this is the nulling method, which can be viewed as a robust method for investigating adaptation aftereffects (for review see: Castet, Keeble, & Verstraten, 2002). More importantly, it conceivably can be used to extract the PSE (a measure of bias), also referred to as the nulling point under this method, and the JND (measure of sensitivity) for a task dimension, which is something that Dong and colleagues (2016; 2019) were not able to do with their stimulus design. Therefore, it would be interesting to use the nulling method and a comparable design using periodic grating stimuli to extract both bias and sensitivity over training day. This would allow for better understanding of whether repetitive training in an adapted state is altering bias, sensitivity or both.

In this chapter, I examine the effects of multiple daily adaptation sessions on perceptual bias and sensitivity. The MAE has been investigated in the past using the relative contrast between two opposing gratings (Ledgeway, 1993). Furthermore, it should provide an interesting comparison between motion mechanisms activated by the relative contrast between two opposing gratings and the coherence detection of dot-motion stimuli used by Dong and colleagues (2016; 2019).

3.2 Method

3.2.1 Participants

Fourteen participants (Mean Age = 22.6, Std = 2.34) with normal or corrected to normal vision took part in this research. They were seated in a dark room and viewed the computer monitor at a distance of 75.6cm. Their heads were secured using a chin rest to minimize head movement. All participants gave informed consent and were paid an inconvenience allowance for their time.

3.2.2 Stimuli

The stimuli were computer generated grating patches that were displayed on a mid-grey background (65 cd/m^2) either side of a central fixation point. The fixation point measured 0.25° of visual angle. The adaptor and test grating had the same features (90°, 1 cpd) and were each displayed within a circular aperture window that measured 5° in diameter and drifted at a constant speed of 5Hz. Both adaptor and test probe were either presented 5° to the left or right of the fixation point depending on the condition. The adaptor grating was always displayed at 100% contrast and drifted in a coherent direction, either to the left or right and always drifting away from the centre. The test grating consisted of two superimposed gratings, drifting in opposing directions (5Hz) at different contrasts that together summated to 100% contrast. When the contrast of each of the two opposing gratings was 50%, this produced a counter phasing grating.

3.2.3 Procedure

The experiment lasted a total of 11 days and were completed on as close to consecutive days as possible. On the pre- and post-training sessions (day 1 and 11), participants completed 6 blocks of the nulling procedure whilst adapting (3 blocks for each side) and 6 blocks of the nulling procedure whilst unadapted (3 blocks for each side). The adapting and unadapted blocks were completed separately to avoid cross contamination with any adaptation blocks. Both types of blocks were randomised between each side. In total the pre- and post-sessions lasted approximately 1 hour 45 minutes, with adapting blocks lasting an average of 12 minutes. During the training days (day 2 to 9), participants were assigned a fixed adapting direction and location for the duration of the 9-day training period. Each training day session lasted approximately 45 minutes and participants completed 3 blocks of the nulling procedure.

The nulling method employed in this design used a 2-AFC forced choice design where on each trial participants reported whether the test grating drifted more to the right or to the left (see schematic in figure 3.1). In the adapting blocks, participants were passively exposed to a coherently moving grating before being presented with the test grating (0.5 sec). The adaptation protocol between each trial followed a top-up manner where participants adapted for 30 seconds before the first trial and then for 5 seconds in between trials as a top-up. An inter-stimulus-interval (ISI) of 0.5 seconds separated the adaptor and test grating. In the unadapted conditions, participants performed this without the adaptor in between trials.

In the adapted and unadapted conditions, the contrast ratio of the test grating was controlled using two-interleaved staircases that dictated the relative contrast of the two opposing gratings. Each staircase was controlled by a 1-up 1-down staircase method that was controlled in logarithmic units. It initially started either -0.95 or 0.95 log units and reduced by a step size of 0.25 log units until the first reversal in the staircase. It then decreased by 0.1 log units and then finally by 0.025 log units after the third reversal in the staircase. Each block ended after 30 trials in each staircase had been completed.

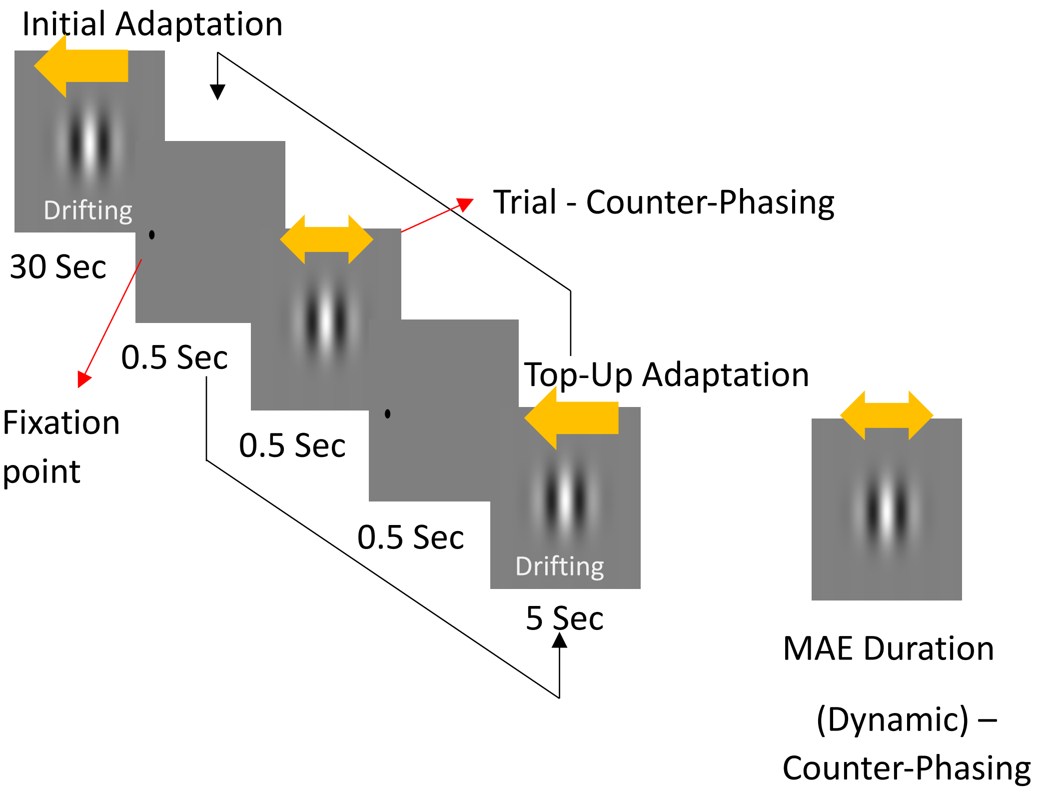

Figure 3.1 – *Nulling Method Trial Sequence*. After the initial adaptation period of 30 seconds (sec), the following adaptation and top-up sequence followed where participants were presented a trial for 0.5 sec followed by ISI of 0.5 sec and a top-up adaptation period of 5 sec. This process repeated until the two interleaving-staircases ended after a limit on amount of trials. At the end of each block, the dynamic MAE duration was measured.

After both staircases ended, an on-screen message appeared and advised that the task was now to respond to the duration of the MAE. After a 10-second countdown, the adapting grating re-appeared on screen for 30 seconds in the same retinal location and adapting direction. After 0.5 seconds, a counter phasing grating appeared and remained on screen until the participant indicated, by button press, when the MAE had ended. This type of MAE measurement is known as the dynamic MAE. In total 3 measurements were averaged for each session. Participants were given a 5-minute rest at the end of each block.

3.3 Results

3.3.1 Training effects on MAE Duration

The average MAE durations across each testing day are plotted in Figure 3.2. The data reveal that the MAE across participants remained relatively flat across training day. The average MAE duration did decrease slightly for both sides across pre and post measures. For the trained side, the MAE reduced by 2.08 seconds between the start and end of the experiment (trained pre: mean = 8.8 secs, std = 4.51 secs; trained post: mean = 6.72 secs, std = 3.39 secs). For the untrained side, the reduction in MAE was 1 second exactly between the start and end of the experiment (untrained pre: mean = 8.68 secs, std = 6.82 secs; untrained post: mean = 7.68 secs, std = 3.3). To understand whether these reductions were significant, a 2 (Session: Pre vs Post) x 2 (Side: Trained vs Untrained) repeated

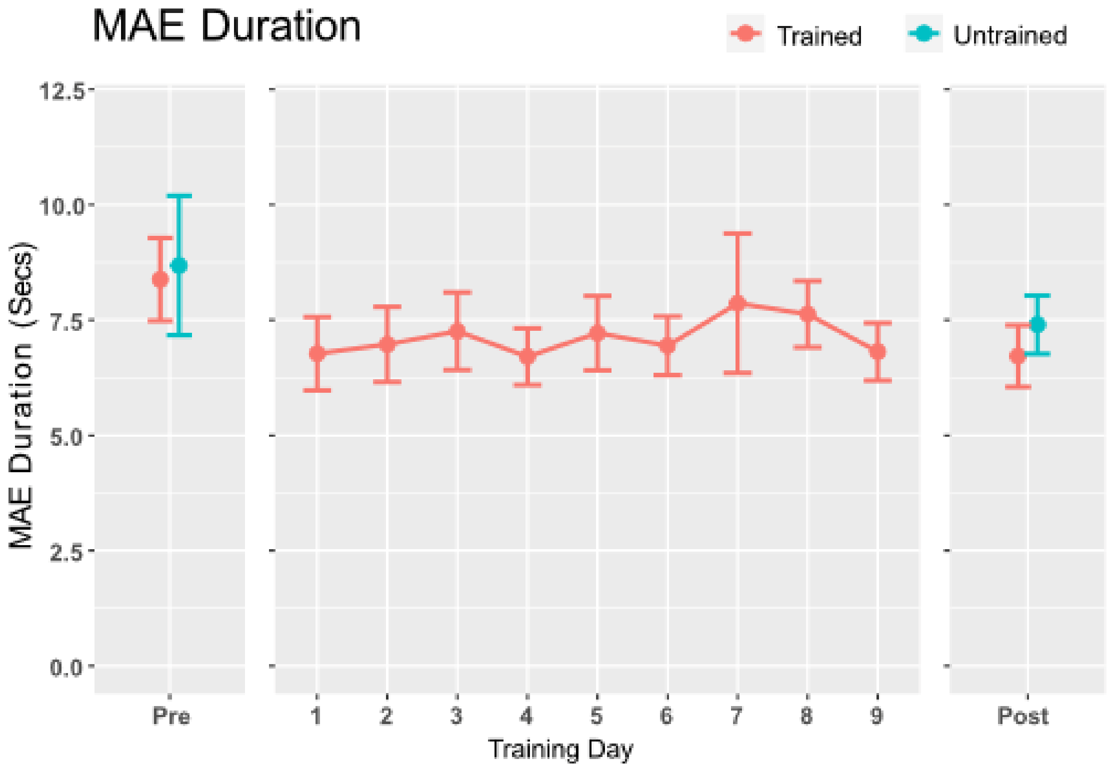

Figure 3.2 – *Motion Aftereffect (MAE) Duration across Session.* The results revealed a slight decrease in MAE between Pre and Post sessions however this was not significant. Error bars represented 1 standard error of the mean.

measures ANOVA was run on the MAE data and revealed that this decrease was not significant. There was a non-significant within effect of session on MAE duration (f(1,51) = 1.283, p = .263, ηp^2 = 0.02), indicating that 9 days of training had no effect on the MAE duration. There was also a non-significant effect of training between sides (f(1,51) = 0.095, p = .76, ηp^2 = 0.002). The ANOVA also revealed a non-significant interaction between session and side (f(1,51) = 0.158, p = .693, ηp^2 = 0.004).

The results reported here for the MAE duration are strikingly different to the MAE duration reported by Dong et al., (2016) shown in Figure 3.3. The most notable difference is in the MAE durations reported in the two experiments which is much higher in Dong and colleagues report. For instance, their average pre-training estimates for both trained and untrained sides were estimated to be 22.6 and 23.1 seconds in duration, more than double the aftereffect durations we report in our pre-training measurement. Perhaps even more

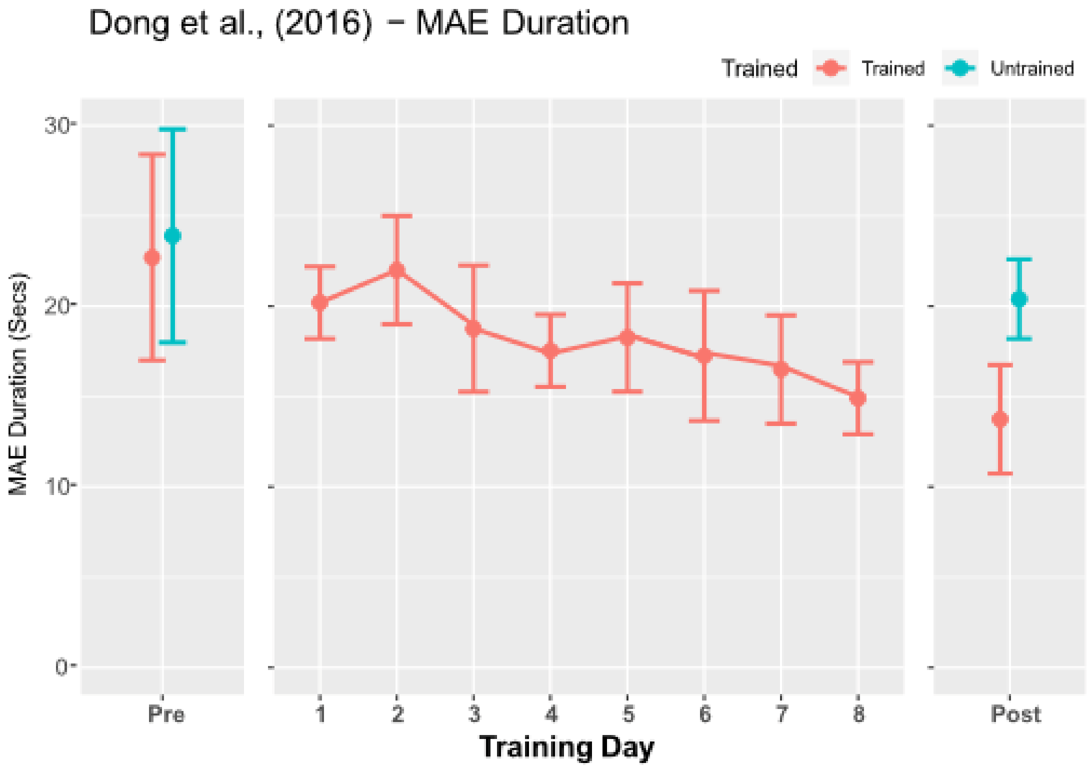

Figure 3.3 – *Motion Aftereffect (MAE) Duration in Dong et al., (2016).* The duration of the MAE declined over session and revealed a gradual decrease in MAE between Pre and Post sessions. However, this did not reach significance. The duration of the MAE was taken using a static RDK stimulus. Error bars represented 1 standard error of the mean.

strikingly, their trained side post training MAE duration (14.2 seconds), where they showed an overall reduction, was still higher than our pre-training estimates.

One obvious methodological difference that could have caused this difference is that Dong and colleagues measured the MAE using a static field of dots, whereas we measured the dynamic MAE using a counter phasing grating. These two methods have been shown to activate different parts of the visual system that process motion (for review see Mather et al., 2008). Thus, it could be that the cortical areas that are adapted and tested for either static or dynamic motion have different decay times for the aftereffect to fade. However, our duration estimates are actually in line with the duration results that Petrov and Van Horn (2012) reported, who measured the effects of training on the duration of the MAE for both static and dynamic probes following 20 seconds of adaptation to motion. In their experimental results (not represented), the average pre- and post-training MAE durations were between 4 and 5.5 seconds.

4.3.2 Training effects on Bias

The average PSE (referred to as the nulling point from herein) across each training day is reported in Figure 3.4. The PSE values for each day were obtained by fitting a psychometric function, using a maximum likelihood estimation procedure, to the raw data for each participant. The PSE threshold was the stimulus value along this fitted curve that yielded 50% leftward response. These PSE values for each participant were then recoded such that the values indicate net motion in the same direction as the adaptor. These values were then averaged across all participants to create the average nulling point displayed in Figure 3.4. The data reveal a very slight increase in the average nulling point for both adapting sides between pre- and post-sessions, however the nulling point across training day suggest that generally this was relatively flat and unvarying across session. For the trained side, the average nulling point increased by 0.038 log units (trained pre: mean = 0.361, std = 0.209; trained post: mean = 0.399, std = 0.188). For the untrained side the average nulling point increased by 0.018 log units (untrained pre: mean = 0.391, std = 0.204; untrained post: mean = 0.409, std = 0.169). To understand whether this slight increase was significant a 2 (Session: Pre vs Post) x 2 (Side: Trained vs Untrained) repeated measures ANOVA was run on the nulling point data and revealed that these differences were not significant. There

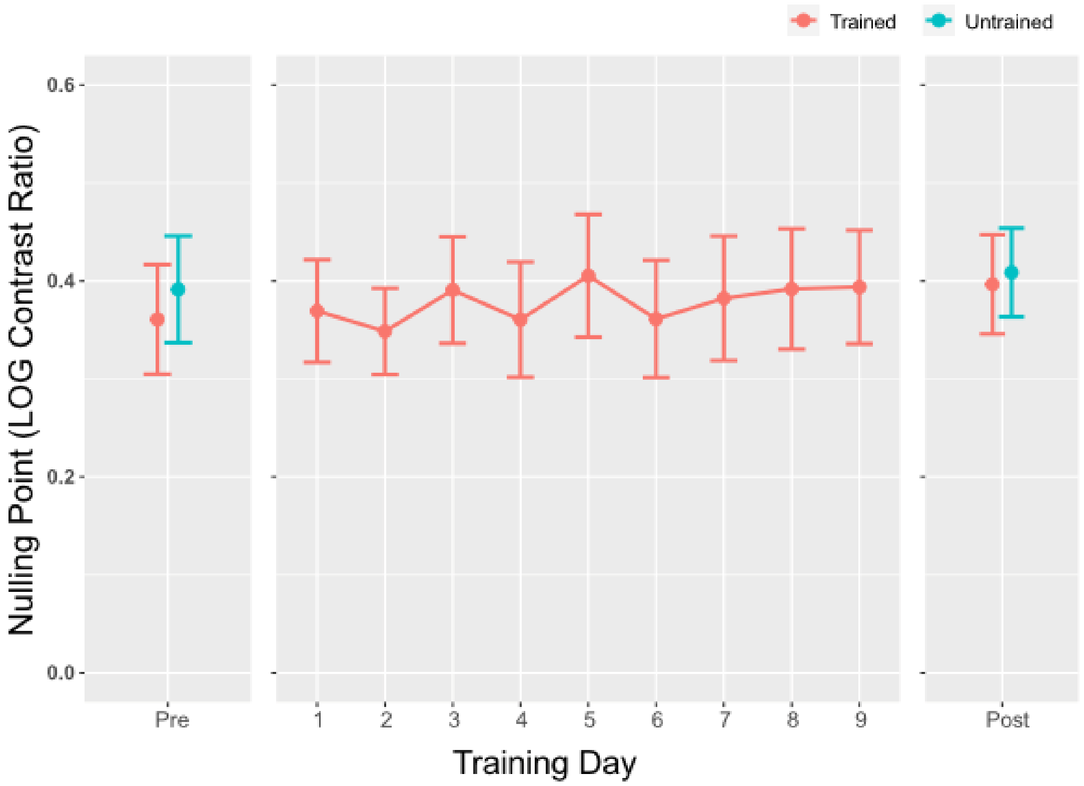

Figure 3.4 – *Average Nulling Point in Log Contrast Ratio units across Session.* The data reveals a slight increase in the nulling point for both sides between pre and post sessions. However, this was not significant. Generally, the nulling point remained flat across training session. Error bars represented 1 standard error of the mean.

was a non-significant effect of session on the nulling point estimates (f(1,51) = 0.263, p = .610, ηp^2 = 0.005). The analysis also revealed a non-significant effect of training between sides (f(1,51) = 0.171, p = .681, ηp^2 = 0.003). The ANOVA also revealed a non-significant interaction between session and side (f(1,51) = 0.04, p = .859, ηp^2 = 0.0006).

Figure 3.5 shows the nulling percent coherence that Dong and colleagues (2016; 2019) reported. It is important to note that this is not directly comparable to our nulling point, because their nulling percentage is a coherence level of signal dots required to null the adaptation aftereffect. However, what is clear is that in both of their experiments the nulling percentage is clearly reducing overtime for the trained side, with untrained nulling percentages not changing overtime. For example, in both experiments their nulling percentage decreased by around 20 percentage points in their 2016 experiment, and 15

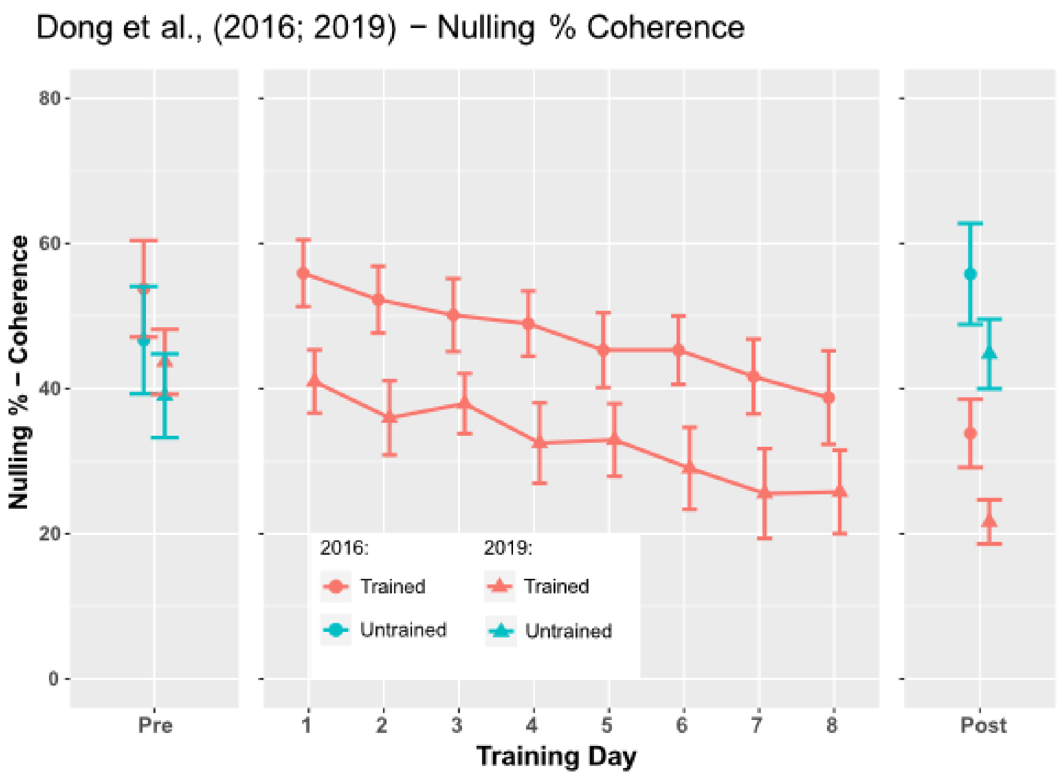

Figure 3.5 – *Nulling Percentage for Dong et al., (2016; 2019).* The results for this experiment show a systematic significant decrease in the nulling percent coherence overtime. In 2016, Dong and colleagues recruited 14 participants and in 2019 recruited 10. Error bars in both graphs represented 1 standard error of the mean.

percentage points in their 2019 study, whereas, their untrained nulling percentage points showed slight increases of about 5 percentage points in both.

4.3.3 Training effects on Sensitivity

The average JND thresholds across session is represented in Figure 3.6. The JND values for each day were obtained by fitting a psychometric function, using a maximum likelihood estimation procedure, to the raw data for each participant. The JND threshold was estimated from the slope of this psychometric function. These JND values were then averaged across all participants to create the average nulling point displayed in Figure 3.6. The data reveal a decrease in the average JND for the trained side across pre- and post-measurements (trained pre: mean = 0.308, std = 0.157; trained post: mean = 0.181, std = 0.08) indicating that participants were becoming more sensitive to discriminating the relative contrast of two opposing gratings. However, there was a slight increase in JND threshold for the untrained side across pre- and post-measurements (untrained pre: mean

= 0.295, std = 0.148; untrained post: mean = 0.317, std = 0.147). The group average JND thresholds across training day also show a clear downward trend as JND performance continued to decrease between days. To test whether these differences were significant, a 2 (Session: Pre vs Post) x 2 (Side: Trained vs Untrained) repeated measures ANOVA was run on the pre and post JND threshold data. The analysis revealed a non-significant main effect of session (f(1,51) = 2.076, p = .156, ηp^2= 0.038) and a non-significant main effect of side (f(1,51) = 2.851, p = .097, ηp^2 = 0.052). The ANOVA did however reveal a significant interaction between session and side (f(1,51) = 4.118, p = .048, ηp^2 = 0.073). To explore this interaction further, paired samples t-tests were conducted on the pre- and post-measurements for the trained and untrained JND thresholds. There was a significant difference between pre and post JND thresholds on the trained side (t(13) = -3.35, p = 0.005) indicating performance was significantly better after training. The analysis revealed a non-significant difference for the untrained side (t(13) = 0.66, p = 0.51) indicating that there was no transfer to the untrained side.

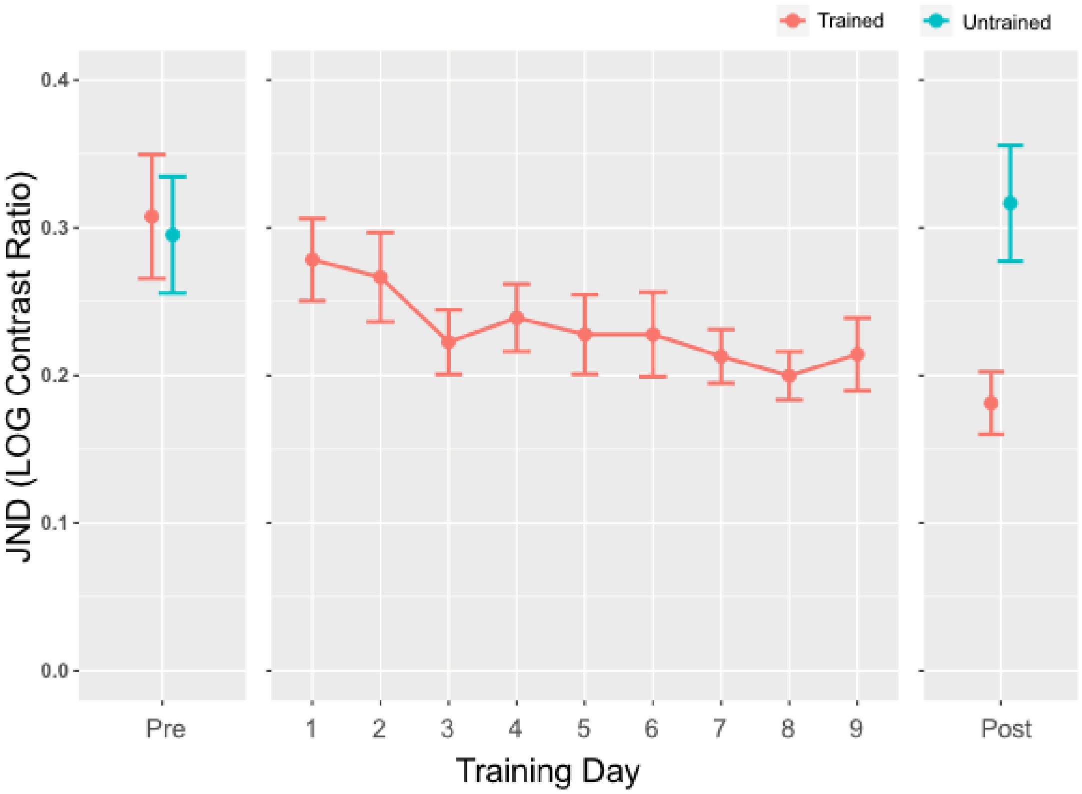

Figure 3.6 – *Just-Noticeable Difference (JND) in Log Contrast Ratio units across Session.* The results revealed a decrease in JND between Pre and Post sessions which was also seen across training day. The difference between pre and post session was significantly different on the trained side. Error bars represented 1 standard error of the mean.

An example of the fitted psychometric functions over session from one participant in this experiment can be found in Figure 3.7. As the functions show, the JND did not change for the untrained side as the slope of the functions remained the relatively similar (Figure 3.7a). However, there is a noticeable difference between the fitted functions on the trained side as the slope of the post session was steeper indicating increased accuracy (Figure 3.7a). This is supported by the progressive steepness of the fitted functions between training day 1 through to day 9 (Figure 3.7b – left to right column).

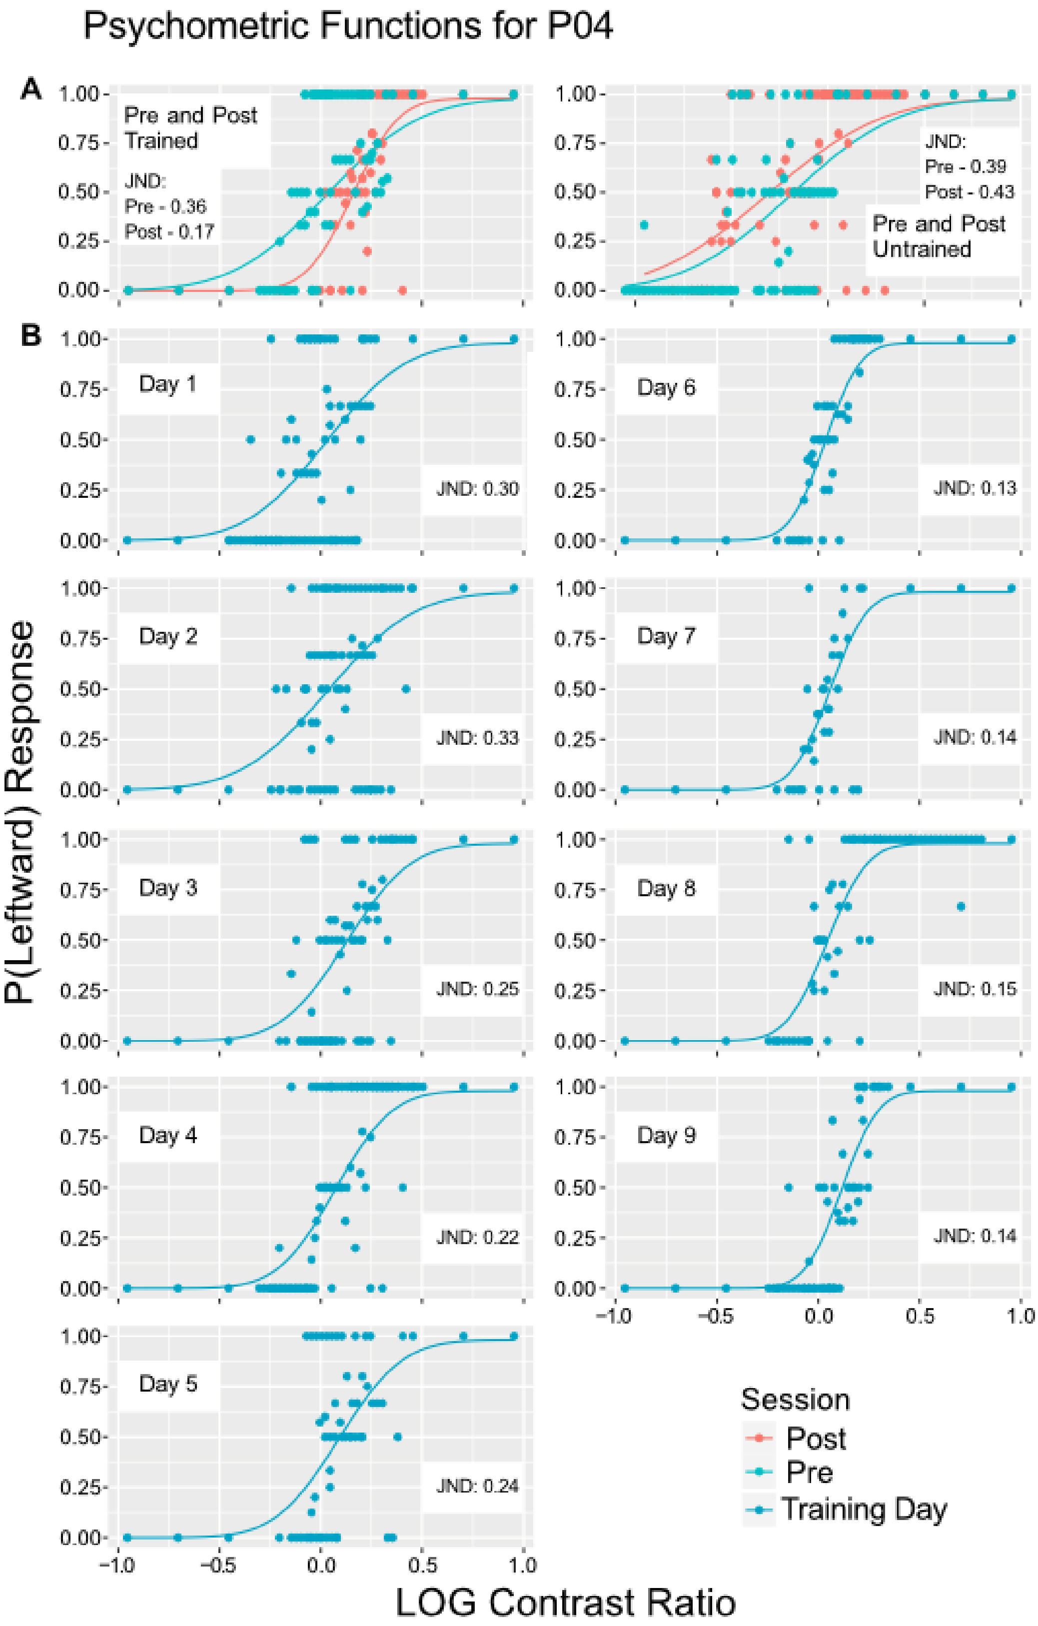

Figure 3.7 – *Fitted Psychometric Functions for Participant 4 and associated JND values.* Each psychometric function was fitted using maximum likelihood method of estimation. **Figure 3.7a** – psychometric functions for Trained and Untrained pre- and post-measurement estimates. The JND for the untrained side increased slightly, however the JND for trained side decreased. **Figure 3.7b** – Training day psychometric functions for day 1-9. The JND progressively decreased across training session.

4.3.4 Training effects on Unadapted Thresholds

The average unadapted nulling points for the trained and untrained sides are represented in boxplots in Figure 3.8a. The data reveal that the unadapted nulling points were all biased toward responding more leftward. For the trained side, the average pre nulling point was 0.006 log units more than the post nulling point (trained pre nulling point: mean = -0.008, std = 0.1; trained post nulling point: mean = -0.002, std = 0.2). For the untrained side, the average pre nulling point was 0.062 log units more than post nulling point (untrained pre nulling point: mean = -0.021, std = 0.16; untrained post nulling point: mean = -0.083, std = 0.19). To understand whether these differences were significantly different, a 2 (Session: Pre vs Post) x 2 (Side: Trained vs Untrained) repeated measures ANOVA was run on the unadapted nulling point data. The analysis revealed a non-significant effect of session (f(1,48) = 0.342, p = .561, ηp^2 = 0.007) and a non-significant effect of training between sides (f(1,48) = 0.982, p = .372, ηp^2 = 0.02). The ANOVA also revealed a non-significant interaction between session and side (f(1,48) = 0.502, p = .482, ηp^2 = 0.01)

To ensure that the these baseline biases did not obscure the pattern of adaptation results, the adapted nulling points for each participant were baseline corrected and these nulling point values are represented in Figure 3.8b. In comparison to the uncorrected nulling point values represented in Figure 3.4, the overall shape is consistent, there is a rise in the nulling point overtime on the trained side and the untrained side shows a very slight increase. The only difference between the corrected and uncorrected nulling points is a slight lowering overall of the trained side values. For instance, on the trained side, the nulling point after baseline correction decreased by 0.03 log units. However, on the untrained side the nulling point after baseline correction decreased by 0.008 log units. Just to ensure that our initial analysis on the nulling point was accurate, a 2 (Session: Pre vs Post) x 2 (Side: Trained vs Untrained) repeated measures ANOVA was run on the baseline corrected nulling point data and the same non-significant results were obtained. There was a non-significant effect of session on the nulling point estimates (f(1,51) = 0.129, p = .721, ηp^2 = 0.002) indicating that there was no effect of training. The analysis also revealed a non-significant effect of training between sides (f(1,51) = 0.936, p = .338, ηp^2 = 0.018). The ANOVA also revealed a non-significant interaction between session and side (f(1,51) = 0.06, p = .794, ηp^2 = 0.001).

The average unadapted JND thresholds for the trained and untrained sides are represented in boxplot in figure 3.9. The data reveal that unadapted JND thresholds decreased on the trained side by 0.092 log units between pre- and post-training (trained JND pre: mean = 0.343, std = 0.124; trained JND post: mean = 0.258, std = 0.168). The data for the unadapted JND thresholds also show a slight decrease by 0.023 log units pre- and post-training (untrained JND pre: mean = 0.357, std = 0.13; untrained JND post: mean = 0.334, std = 0.09). To understand whether these differences were significant a 2 (Session: Pre vs Post) x 2 (Side: Trained vs Untrained) repeated measures ANOVA was run on the data. The analysis revealed a non-significant effect of session (f(1,48) = 2.291, p = .137, ηp^2 = 0.046) and a non-significant effect of training between sides (f(1,48) = 1.565, p = .217, ηp^2 = 0.032). The ANOVA also revealed a non-significant interaction between session and side (f(1,48) = 0.749, p = .391, ηp^2 = 0.015).

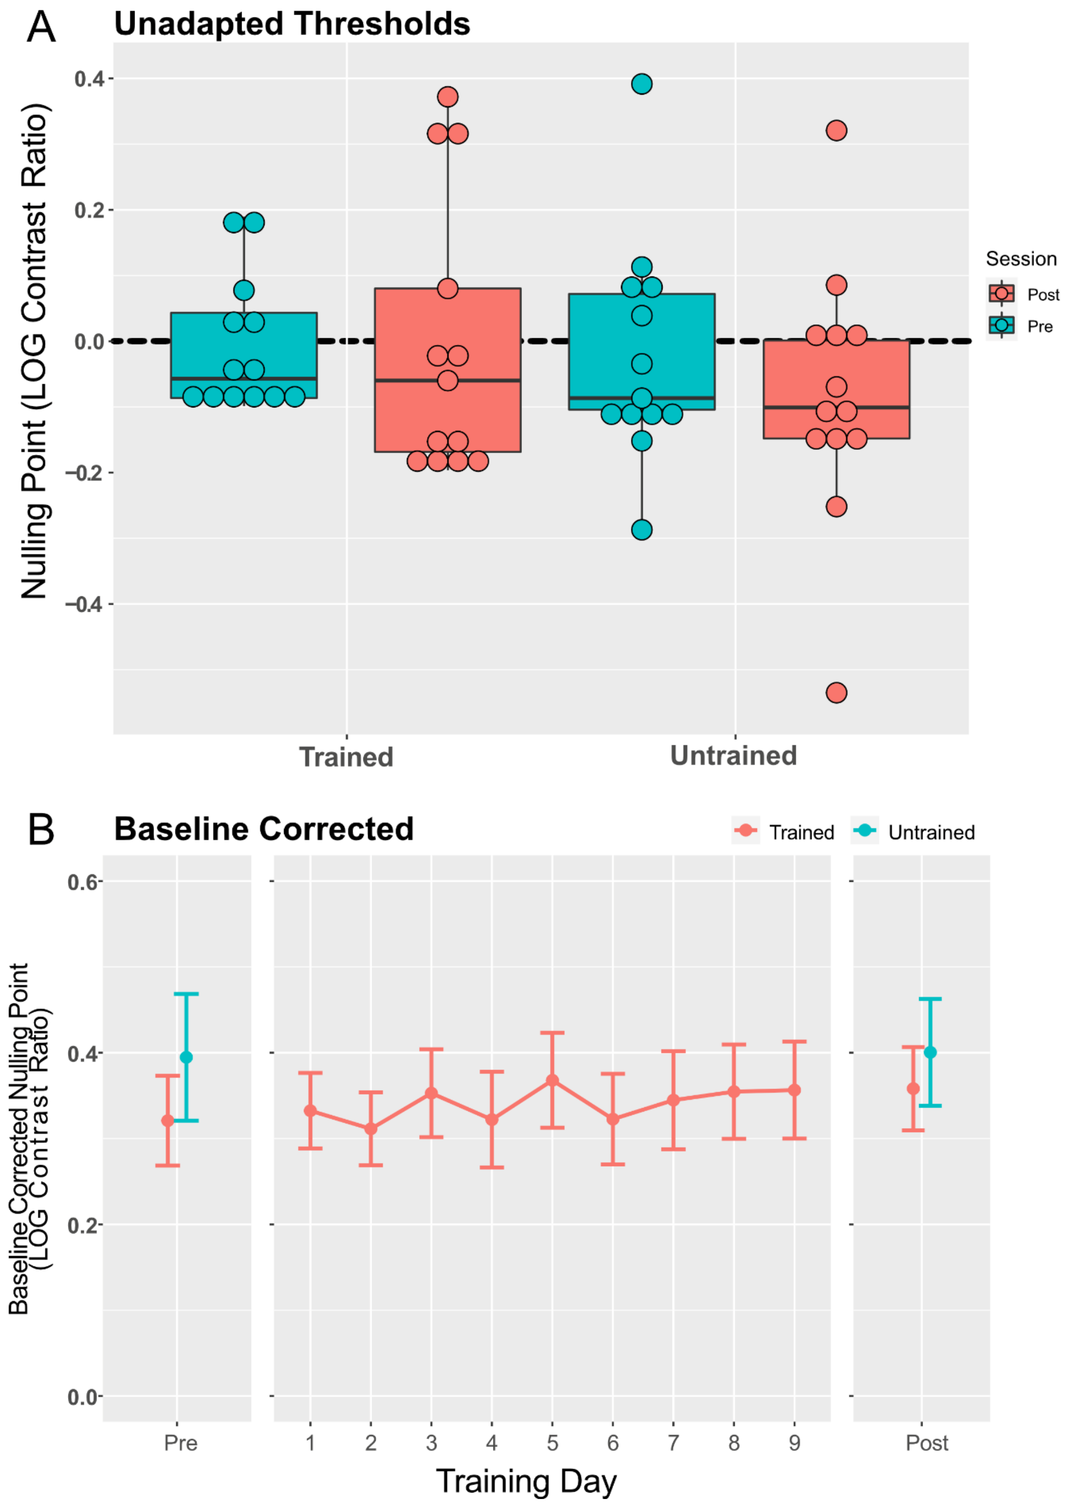

Figure 3.8 – ***A)*** *Unadapted Nulling Point Threshold Boxplots for the Trained and Untrained side. Participants in the task were more biased in their unadapted response.* ***B)*** *Baseline Corrected Nulling Point Training Thresholds.* To account for the slight bias in unadapted thresholds, the training nulling point thresholds were baseline corrected. There was a slight overall reduction but analysis agreed with our uncorrected thresholds. Error bars represented 1 standard error of the mean.

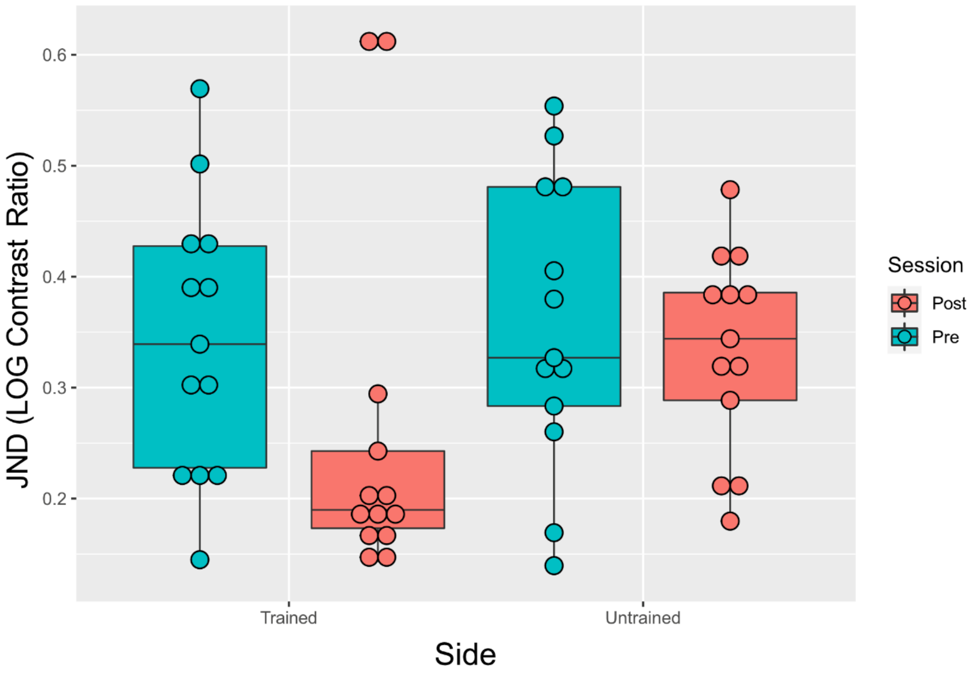

Figure 3.9 – *Unadapted JND Threshold Boxplots for the Trained and Untrained side.*

3.4 Discussion

In this chapter, I explored the effects of multiple daily adaptation sessions (training) on the adapted state. Specifically, I was interested in whether biases induced by adaptation, when measured using the nulling method, would habituate overtime for the trained adapted state. I was also interested in whether other measures of adaptation bias, the duration of the aftereffect, would reduce for the trained adapted state. Both forms of bias have been shown to habituate overtime by others in the literature (Dong et al., 2016; 2019). Finally, I was interested in whether adapted state sensitivity for the task would be modified for the trained side as a result of training. Our intention with this experimental data was to better characterise this relationship between adapted state bias and sensitivity over time to elucidate some conflicting reports on how adaptation aftereffects change as a consequence of training (McGovern et al., 2012; Dong et al., 2016; 2019; Petrov & Van Horn, 2012).

Using a comparable design to Dong and colleagues (2016; 2019), my results showed no signs of adaptation habituation across training day when measuring the MAE using the relative contrast between two opposing gratings. In agreement with this lack of change to adapted state bias, my results also showed that the duration of the MAE remained consistent throughout training for both the trained and untrained sides. I did however find that discrimination thresholds significantly reduced as a result of training for the trained adapted state. This indicated that participants were getting more accurate at discriminating changes in the contrast ratio of the opposing gratings and that this improvement was specific to the trained adapted state. In combination, these results suggest that the habituation of visual adaptation is perhaps not a general characteristic of all motion mechanisms. Instead, this chapter suggests that training in an adapted state primarily results in improved performance of the trained task.

This dissociation between learning induced changes in adapted state sensitivity rather than bias is the opposite to Dong and colleagues (2016; 2019) research. Dong and colleagues (2016; 2019) explained their results using the Stimulus-Model Comparator theory of habituation (Sokolov, 1963), whereby with repeated experience of a stimulus the sensory system creates a model of what is expected and through repetitive exposure, the sensory system builds a more accurate representation of the expected stimulus, leading to gradually increased inhibition on the response. As training progresses, the stimulus model gets more developed and consequently exerts increasing descending inhibition to reduce the neural response of the adaptor. Therefore, given their results this theory appears to make sense because the coherence of the dot motion signal needed to perceptually null the MAE became less as a consequence of adapting to a weaker motion stimulus.

Whilst we cannot directly compare our results, as Dong and colleagues data represents percent motion coherence, our results cannot be explained as a reduction in the effective strength of the adapter over training because our nulling point measure of bias remained consistent throughout training. If habituation was reflected in our results we would expect this nulling point to have reduced, and like Dong and colleagues, we may have also expected the MAE duration to have reduced over time. Therefore, as my results do not fit this theory, what could conceivably explain my results?

The MAE has been studied extensively (for recent review: Mather et al., 2008) using a variety of different methods and stimuli. The nulling method has been used with random dot-motion stimuli (Raymond & Braddick, 1996) and sinusoidal gratings (Ledgeway, 1994). Both show quite clearly that when measured psychophysically, both sets of stimuli can cause significant bias in the perception of motion. Physiologically, the processing of visual motion relies on direction selective neurons that respond preferentially to motion in defined directions. These direction selective neurons have been found throughout the visual motion pathway, including the retina, retinal ganglion cells, V1, middle temporal area (MT) and medial superior temporal area (Ölveczky, Baccus & Meister, 2003; Barlow, Hill & Levick, 1964; Adleson & Movshon, 1982; Movshon, Adleson, Gizzi & Newsome, 1985; Movshon & Newsome, 1996; Pack, Berezovskii & Born, 2001, Orban, 2008), which have all shown to have a diverse response range involved in different aspects of motion processing. In addition, areas along the visual pathway that respond to motion have been found to project onto each other, for instance, layers in V1 have been found to project onto MT (Gur & Snodderly, 2007).

The processing of visual motion is thought to start in V1 neurons, that represent the first stage of computing the oriented elements in visual scenes that are selective to orientation and direction (Hubel & Wiesel, 1962). The activity of these V1 neurons represent single components, however they would not reveal the overall motion of a pattern, as the components of a pattern can move in different directions (Adleson & Movshon, 1982). The visual system overcomes this by combining motion signals from multiple V1 cells representing different components, to compute pattern motion represented by the activity in direction selective neurons in area MT (Adleson & Movshon, 1982; Movshon & Newsome, 1996). How the visual system does this is an important issue, and one that is outside of the scope of this experiment, however it may be an important issue for the two opposing outcomes in ours and Dong and colleagues (2016; 2019) results. For instance, as Kumano and Uka (2013) comment, it is not fully understood how pattern motion selectivity using gratings, where there are clear directions of motion, is related to pattern motion selectivity using random dots, that contain multiple orientation and spatiotemporal frequency components on a single-neuron basis. Whilst further physiological research is needed to understand this issue, it could be the case that the use of gratings or random-dot stimuli, are processed more by a certain area of the visual motion pathway which are more susceptible to habituation like mechanisms. For instance, it could be the case that the use of sinusoidal gratings relies more on the processing from specific populations of component cells in V1 which could be less sensitive to habituation-like mechanisms than random dot motion stimuli, where the processing may be processed more as a pattern in area MT.

Another possibility that could help understand the differences between results, is in the perceptual decision that the participants were making in each experiment. As already highlighted for dot-motion stimuli, the coherence of the nulling direction is presented with dots travelling in random directions. Thus, the problem that the visual system must decode, is to detect the signal direction against the equally represented noise dot directions. For sinusoidal gratings on the other hand, the relative contrast of the nulling direction is presented with an opposing grating travelling in the opposite direction. Thus, the problem that the visual system must decode, is discriminating between the contrast of the two opposing directions.

Perhaps rather than viewing one experiment as reflecting a change in perceptual bias and the other sensitivity, these findings show a form of learnt improvement in the decoding stage of visual processing and thus reflect different forms of the same process – state-dependent perceptual learning. For instance, possibly through the repetitive nature of performing the dot motion coherence task whilst in a consistent adapted state, the visual system gets better at detecting the nulling direction amongst the noise in the trained adapted state. The same explanation could be applied to the relative contrast of the gratings that we used in our study. This could be plausible since in neither experiment did the altered adapted state transfer to untrained adapted states or unadapted performance.

Overall, our results agree the most with those presented by McGovern and colleagues (2012) who showed that repetitive adaptation improved discrimination accuracy in that particular trained adapted state, whilst adapted state bias remained consistent throughout training. However, the research presented here adds to this work because in this study I measured both bias and sensitivity in conjunction with one another, opposed to measuring them one after the other. This is important because the adapted state could have changed between these measurements.

Models of perceptual learning theorise that increased perceptual sensitivity on a task can arise either through some change to the early levels of the visual cortex (Schoups et al., 2001; Yang & Maunsell, 2004) or through higher-level decision related areas that read-out to form decisions (Dosher & Lu, 1998; Law & Gold, 2008). In addition, models of adaptation whilst diverse (Kohn & Movshon, 2003; Dragoi et al., 2000), are more commonly described as a bias in visual functioning due to a reduction in neuronal gain (Seriès et al., 2009, Clifford et al., 2000). Both of these established fields have accomplished their findings in the absence of one another. Therefore, a common framework that can account for adapted state perceptual learning, or put another way, the effects of learning and adaptation within the same theoretical framework, is hard to theorise on and should be the focus of further work.

Whatever form this common framework will take, it is unlikely to be the process of purely altering the response properties of early levels of visual cortex. If we consider perceptual learning, if training altered the tuning of adapted neurons most relevant to the task at early stages of visual processing by boosting their gains, this could explain discrimination accuracy in the adapted state, however cannot explain why this change is specific to the trained adapted state. For instance, conceivably it would have uniformly improved adapting discrimination across untrained and unadapted conditions (McGovern et al., 2012). This common framework is also unlikely to be the process of purely reweighting connections (Dosher & Lu, 1998; Law & Gold, 2008). For example, according to Series et al., (2009) adaptation-induced costs to discrimination accuracy cannot be overcome by optimising the readout. This is because thresholds are limited by a lower bound determined by Fisher information during adaptation (Seriès et al., 2009), meaning it is unlikely that my results and those of McGovern et al (2012) to be the consequence of placing greater weight on neuronal populations most relevant to discriminating the relative contrast of opposing gratings during adaptation.

If we consider adaptation, which is typically modelled as a reduction in neuronal gain, if through training these reductions switched to an increase in neuronal gain, this could explain the increase in discrimination accuracy. Long-term adaptation has been studied in the past and has shown, like this study, that the aftereffects are not fixed and evolve overtime (Kwon et al., 2009; Zhang et al., 2009; Haak et al., 2014). This could be a possible explanation however, would we also not expect some change in the MAE duration or a reverse of the aftereffect (McGovern et al., 2012)? This latter point is particularly interesting because it relates to how aware visual neurons are to the nature of the task demands (Mollon & Danilova, 1996). In addition, how can reconcile that the adapted state itself, may have a memory component that allows it to be consolidated and accessed between training sessions.

In conclusion, this work adds to the growing body of literature that suggests that the neural mechanisms of adaptation are not fixed. In this research, I believe I was able to accurately measure both perceptual bias and sensitivity at the same time, which revealed that adapted state bias remains consistent throughout training, whereas adapted state sensitivity is altered for the trained adapted state. Whilst this is in opposition to the habituation account proposed by Dong and colleagues (2016; 2019), my results are more in line with previous results shown by McGovern and colleagues (2012). This first experimental chapter is an important step for my next experimental chapter that will investigate the interaction more between adaptation exposure and perceptual learning between distinct task dimensions.

Chapter 4: Multidimensional Perceptual Learning across Distinct Timescales

4.1 Introduction

Human beings display a remarkable capacity to improve their performance on tasks through practice and experience. A good example of this is a trainee radiologist, who must undergo several months of training before perceptually distinguishing diseased from non-diseased tissue on X-ray images with the same accuracy as their more senior counterparts. In this example, one of the key elements that underlies this improvement is a change to the trainee's visual perceptual accuracy. Their visual ability to detect and discriminate across the multiple visual dimensions that make up every unique X-ray image and distinguish between what is task-relevant and task-irrelevant for accurate diagnosis (i.e., the physical contrast of diseased tissue against non-diseased tissue or bone) becomes greatly enhanced. This ability to take in more visual information from the environment, is a phenomenon referred to as perceptual learning (Gibson, 1969, *for review see* Sagi, 2011) and is an example of neural plasticity within the visual cortex.

One aspect of perceptual learning is an enduring change to the important features that are relevant for the task. If we consider natural sensory inputs that we process and operate within each day, they rarely consist of isolated unidimensions that are stable across time (Barlow, 1997). Rather the visual system contends and learns within a multidimensional environment where task-relevant dimensions are coextensive across space and time with task-irrelevant dimensions (Felsen & Dan, 2005).

In sharp contrast with the real-world, many lab-based perceptual leaning experiments adopt a training and testing protocol where this multidimensional spatial and temporal context is absent. A common feature to many lab-based paradigms is to train participants for several days with simple stimuli that vary along a single dimension (see Figure 4.1A) and briefly test participants on a task-irrelevant dimension when that task is switched after training (see Figure 4.1B). Such paradigms are used to understand more about visual

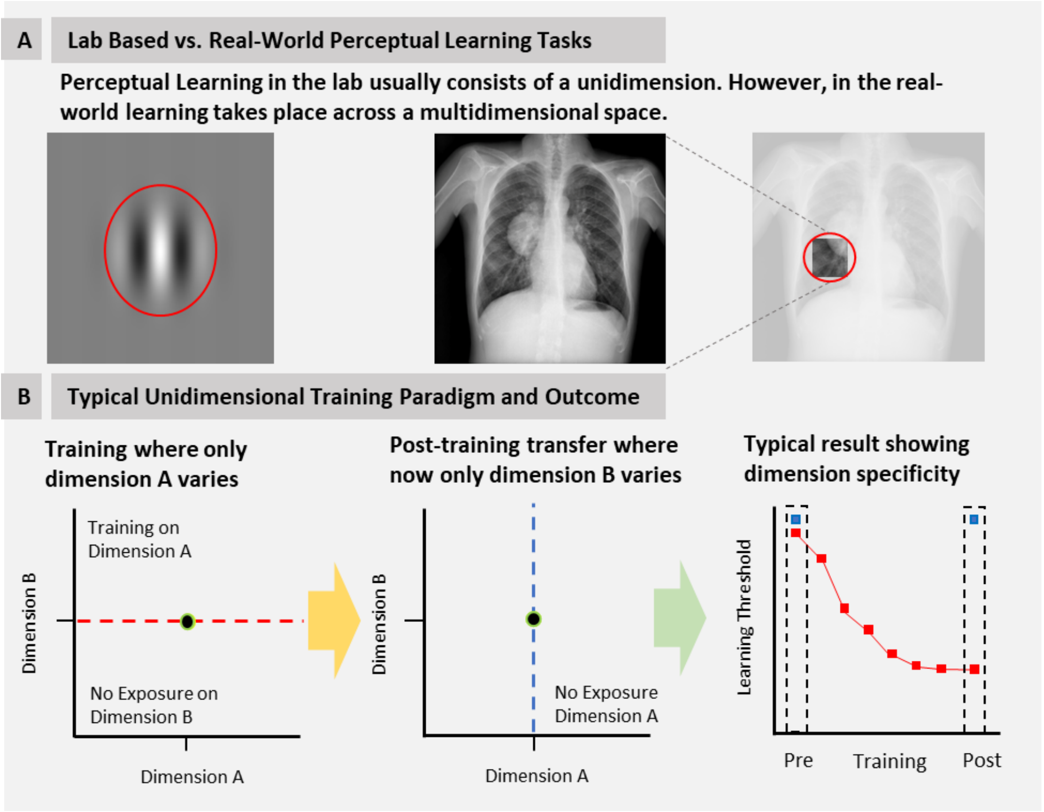

Figure 4.1 – ***A*** Lab-Based vs. Real World Perceptual Learning Tasks. The signal between lab-based and real-world perceptual tasks is very different. For lab-based paradigms, the signal that is trained on is unidimensional there are two dimensions but only one will fluctuate throughout training. In the real-world, the signal is multidimensional and is coextensive with other visual dimensions that fluctuates throughout training. ***B*** Typical Unidimensional Training Paradigm and Outcome. Under the typical training paradigm observers training on one dimension and then switch to another task dimension post-training period. The black dots indicate the decision boundary that is shared between tasks. In both cases, only one dimension fluctuates at a time anda prominent finding is no transfer between dimensions with no exposure to the other task dimension

processing with a view to make perceptual learning more functionally optimal. Overwhelmingly most experiments adopting this paradigm show that learning is specific and non-transferable to other task-irrelevant dimensions outside of the trained task-relevant dimension (Ahissar & Hochstein, 1997; Crist et al., 1997; Karni & Sagi, 1991; Yu; Klein & Levi, 2004; Fiorentini & Berardi, 1980, Fahle, 2004; Fahle and Morgan, 1996; Shiu and Pashler, 1992; Vogels & Orban, 1985). However, with such stark differences between real-world and lab-based learning environments, it is vital for our understanding of perceptual learning to use more naturalistic settings.

Recent research has adopted an approach for testing learning that utilises a different training structure and relationship between tested visual dimensions (Wang et al., 2012; Xiao et al., 2008; Zhang et al, 2010; Szpiro et al., 2016; Watanabe et al ., 2001; Seitz et al,. 2009; Huang, Lu & Dosher, 2012; Petrov et al., 2005, Gutnisky et al., 2009). Whilst there are many differences between these experimental approaches, a shared feature is that these studies presented additional stimuli that were passively exposed to the participant - similar to how adaptation is induced. Adaptation, like perceptual learning, is also a type of neural plasticity that operates on the visual system, however unlike perceptual learning it is characterised as a bias in visual processing following passive sensory exposure (for review see: Webster 2015). These paradigms then briefly test or train on the exposed task-irrelevant dimension after training and exposure ends to measure sensitivity improvement on the exposed task. Research adopting this approach, which have been referred to as "training-plus exposure" or "double-training" paradigms, have shown performance improvements following exposure (Watanabe et al, 2001; Wang et al., 2012; Xiao et al., 2008; Seitz & Watanabe, 2009; Tsushima et al., 2008; Zhang et al, 2010).

While the effectiveness of these procedures remains contentious (Cong, Wang, Yu & Zhang, 2016; Liang, Zhou, Fahle & Liu, 2015a; 2015b; Zhang & Yu, 2016), it has been suggested that passive exposure in this way may lead to a release from adaptation (Harris, GIlksberg and Sagi, 2012). Recently, prolonged adaptation has been shown to promote its own learning-like process over long-timescales that improve visual processing by reducing adapted state bias (Dong et al., 2016; 2019). However, as we showed the in last experimental chapter, repetitive training in an adapted state leads to sensitivity improvements rather than alterations to perceptual bias.

A missing element in the research on perceptual improvements following adaptation exposure is an account of this adapted state that could help explain learnt improvement. One promising line of enquiry may be to interleave training and exposure on tasks to understand the progression of performance improvement following adaptation exposure. Two recent experiments have used such designs to understand more about learning and its transfer across visual dimensions (Szpiro et al., 2014; Huang et al., 2012). Szpiro and colleagues (2014) interleaved training on a task-relevant dimension (orientation discrimination task) with practice exposure on a task-irrelevant dimension (spatial frequency discrimination task). Importantly, both task dimensions shared the same retinal locations and discrimination boundaries (orientation - 30°, spatial frequency - 4 cpd). Whilst exposure to the spatial frequency dimension was not passive, the objective was to understand whether this exposure to a spatial frequency task would facilitate learning on the task relevant orientation dimension. The authors results indicated that learning was facilitated for orientation when training was switching numerous times with the spatial frequency task. This was in comparison to the group that just trained on orientation. This result suggests that the perceptual performance on one dimension was improving performance on another when the two alternated.

In direct contradiction to this, Huang et al., (2012) trained subjects on two task-dimensions that were coextensive with each other throughout training (i.e. occurred at the same time and spatial location). In contrast to Szpiro and colleagues investigation, because dimensions were coextensive in space and time throughout successive switches between task dimensions, training and passive exposure could be accurately assessed against each other. If exposure to the task-irrelevant signal improved performance on that task when it was switched, then one might expect a more accurate initial performance threshold (Figure 4.2). However, despite numerous switches between the task-dimensions across 7 training days, the authors findings suggested that learning was independent (Huang et al., 2012), and exposure to the other dimension did not impact on perceptual performance.

Together these empirical results are at odds. In the first instance, learning between two task dimensions shows some dependency and facilitates learning (Szpiro et al., 2016), in the other, learning is specific for two task dimensions that are coextensive and share identical stimulus inputs (Huang et al., 2012). An interesting test of these two opposing outcomes requires a training and switching paradigm between task dimensions that share the same multidimensional space, and where attention to the task-relevant dimension can be cued before each trial to understand the effects of learning and exposure to a task-irrelevant dimension over time (see Figure 2). It would particularly be interesting to add to Szpiro and colleagues (2016) work and report physical performance results for both task dimensions rather than percent performance for the task-relevant dimension.

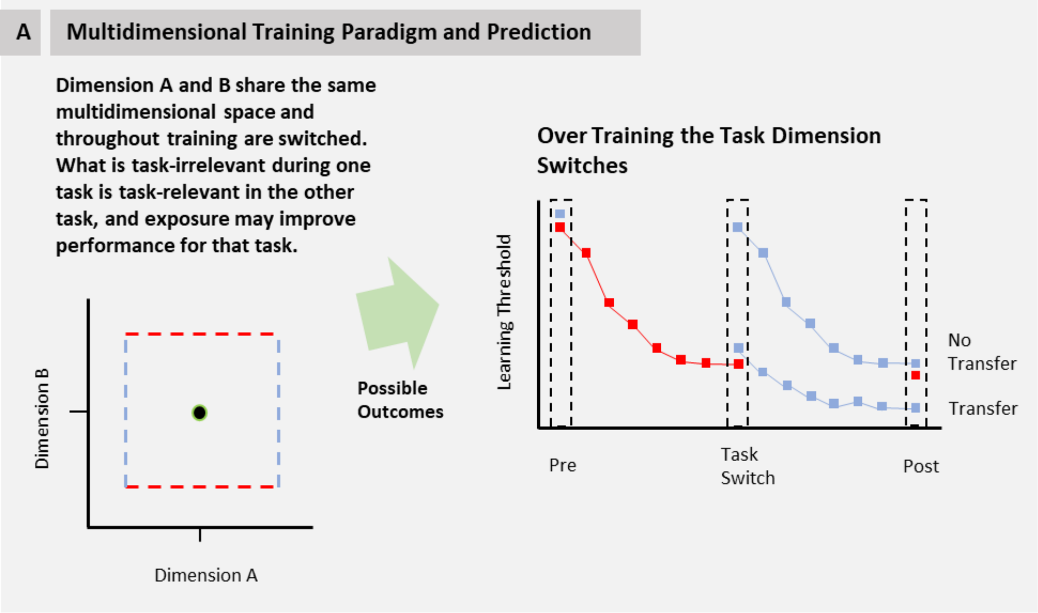

Figure 4.2 – ***A*** Multidimensional Training Paradigm and Prediction. In this type of training, participants are cued to attend to a task-dimension that is coextensive with another task dimension. Over training as the task switches between dimensions, the learning threshold between switches could indicate whether there is a dependency between task dimensions. One possible outcome is that learning will transfer at switch and continue on a similar trend through training. Another possible outcome is that learning will be specific and exposure will cause no performance improvements.

In summary, the aim of this experimental chapter is to investigate the perceptual costs or benefits of learning two perceptual tasks that are coextensive within a multidimensional stimulus space and where learning is switched between them at different timescales. As I discovered in the last chapter, adapted state learning causes sensitivity improvements whilst bias remains consistent throughout. Could I test this further by training subjects on two task dimensions that share the same multidimensional space, and where state-dependent perceptual learning includes exposure to a dimension? In the first part of this experimental chapter, I find the number of trials required to cause significant learning on a task-relevant dimension under a typical unidimensional environment and measure the degree of transfer between them (Experiment 1). In the next experiment, I test whether exposing subjects to a task-irrelevant dimension whilst learning on a task-relevant dimension facilitates learning acquisition; and whether switching after a set of trials causes any learning facilitation on that task dimension (Experiment 2). Finally, I ask whether learning is facilitated when the task-dimensions are switched on a much finer timescale, between trials (Experiment 3).

4.2 Method

4.2.1 Participants

Fifty adults (Mean Age = 24.35, range 19-31 years) with normal or corrected-to-normal vision participated in the study. All participants were untrained at orientation and spatial frequency discrimination and were unaware of the purpose of the experiment. There were ten participants in each of the five experimental groups. Written informed consent was obtained for each participant. All participants were paid an inconvenience allowance for their time.

4.2.2 Apparatus and Stimuli

Participants sat approximately 75.6cm away from the screen with their head secured using a chin rest. The experiment consisted of a standard and test stimulus which were both computer generated Gabor patch stimuli. Both were displayed at 50% contrast and subtended 3° of visual angle. Both Gabor's were presented approximately 7° of visual angle away from the central fixation cross. The test Gabor was directly below the fixation and the standard was displayed to the left of the fixation point. The test Gabor varied along the orientation and spatial frequency dimensions depending on experimental group (more details in each experiment section). The standard Gabor to the left of the fixation point had a fixed orientation of 45° and spatial frequency of 2 cpd. These were the discrimination boundaries for the orientation and spatial frequency tasks and served as a comparison Gabor throughout each task. A noise stimulus was also used that consisted of small black and white pixels. The noise stimulus subtended 4° of visual angle was presented after each presentation of the test Gabor, 7° of visual angle beneath the central fixation cross.

4.2.3 Task and Procedure

All groups shared the same basic task structure during each block of trials (see Figure 4.3). A fixation cross appeared for 750-ms and primed the observer to the task depending on its colour (yellow fixation for orientation judgement; blue fixation for spatial frequency judgement). This was followed by 150-ms test stimulus and then a 400-ms noise mask. After the mask, a standard Gabor was presented to the left of the fixation cross and remained on the screen until the participant made a response using the computer keyboard in front of them. Participants then received feedback and the next trial began exactly 500ms after feedback. All participants were asked to maintain focus on the fixation cross throughout the trial sequence and respond as accurately as possible.

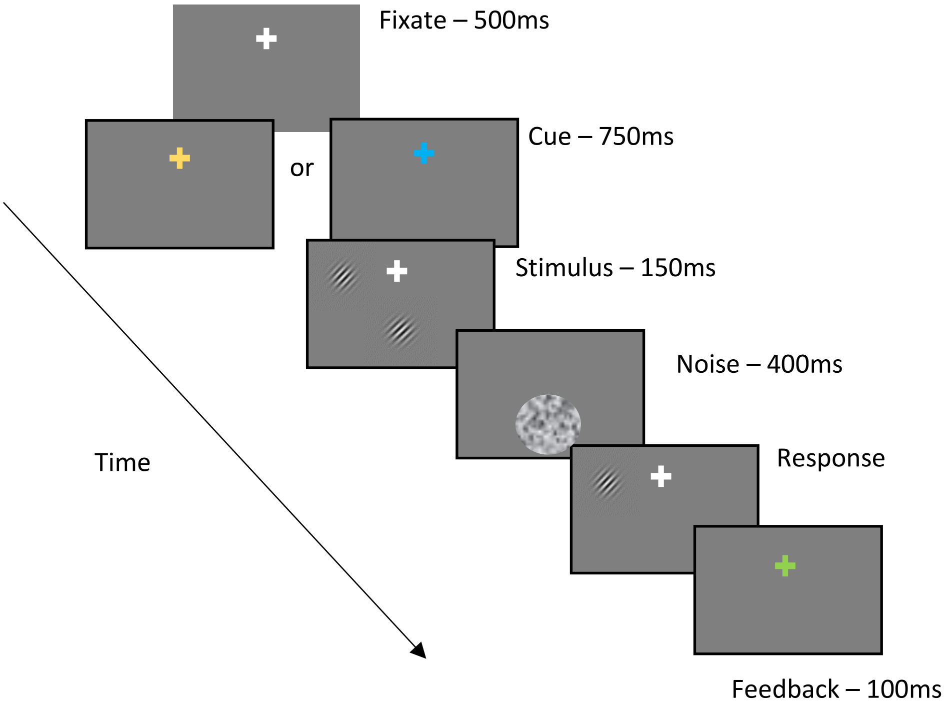

Figure 4.3 – *Trial sequence*. After the initial fixation, the participant was cued to the task dimension. Yellow indicated an orientation trial. Blue indicated a spatial frequency trial.

All participants were given limited practice before beginning the experiment. The practice consisted of first showing the participant the standard stimulus and 5 alternatives either side of the orientation and spatial frequency discrimination boundaries. Following this, each observer was given 15 practice trials of both tasks using the exact same trial structure as in the real task but at a slowed pace (e.g. 1 second stimulus display followed by 1 second noise). This ensured observers understood the nature of what was being asked of them, but without producing significant task related learning.

After the practice period, all participants completed a pre-training session (1 hour), followed by a training period (3 days) and a post-training session (1 hour). The pre- and post-training sessions were identical, and participants completed 3 separate blocks (~180 trials total) of the orientation and spatial frequency tasks in a randomised order. During the training period, depending on the experimental group, participants completed 640 trials of one task (2 days) or both tasks in a switching paradigm (3 days – 640 trials each task).

The initial thresholds obtained for each observer in the pre-training session were used to determine the range of stimulus intensities in the first block of trials in the training period. The stimulus range was within three just-noticeable difference's (JND) from the discrimination boundary for that dimension. In the multidimensional groups, the test stimuli varied along both the orientation and spatial frequency dimensions within this range. However, in the unidimensional groups only the trained dimension would vary with learning whilst the task-irrelevant dimension remained constant at the discrimination boundary. The stimulus intensities in the rest of the training blocks, were determined using the same procedure but based on the JND performance obtained in the previous block.

4.3 Experiment 1

4.3.1 Rationale

Experiment 1 aimed to first establish our chosen paradigm and secondly establish an amount of training trials in a unidimensional environment that leads to learning on the trained dimension. These visual dimensions were the same as Szpiro and colleagues (2016) testing dimensions and these learning functions for orientation and spatial frequency have been well studied in the literature. According to many perceptual learning accounts, training in a unidimensional environment will lead to performance improvements for that trained task dimension and will not transfer to task dimensions outside of training (Ahissar & Hochstein, 1997; Crist et al., 1997; Karni & Sagi, 1991; Yu; Klien & Levi, 2004; Fiorentini & Berardi, 1980, Fahle, 1997; Fahle and Morgan, 1996; Shiu and Pashler, 1992; Vogels & Orban, 1985). Therefore, we do not expect the task-irrelevant dimension to show task improvement after training.

4.3.2 Method

Twenty participants took part in Experiment 1. They were randomly separated into either the Orientation Learning group (O) or Spatial Frequency Learning group (SF). There were 10 participants in each group.

Participants first completed the pre-training session where they completed both the orientation and spatial frequency task. During the training sessions, which lasted 2 days, participants trained on their assigned task-relevant dimension. In the O group, participants trained only on the orientation dimension and completed 640 trials (256 trials on the day 1 and 384 on day 2). In Spatial Frequency Learning group (SF), the training was identical but for the Spatial Frequency dimension. In the post-training session, participants completed both the Orientation and Spatial Frequency tasks again in a randomized order.

4.3.3 Results

The first objective of Experiment 1 was to characterise the learning function for each task dimension under a conventional unidimensional training paradigm. Figure 4.4 shows the average log normalised learning function as a function of trial for the O and SF groups. These were calculated by first fitting a psychometric function, using a maximum likelihood estimation procedure, to the raw data for each participant to obtain the JND. The JND threshold was estimated from the slope of this psychometric tuning function. These JND values for each participant were then log normalised and then averaged across all participants. Converting them to a logarithmic score makes it possible to compare spatial frequency and orientation thresholds on the same axis. The data show noticeable learning effects in both groups relative to pre-training thresholds (black dashed line). For the O group (Figure 4.4 top tier), perceptual performance improved across trials and there was

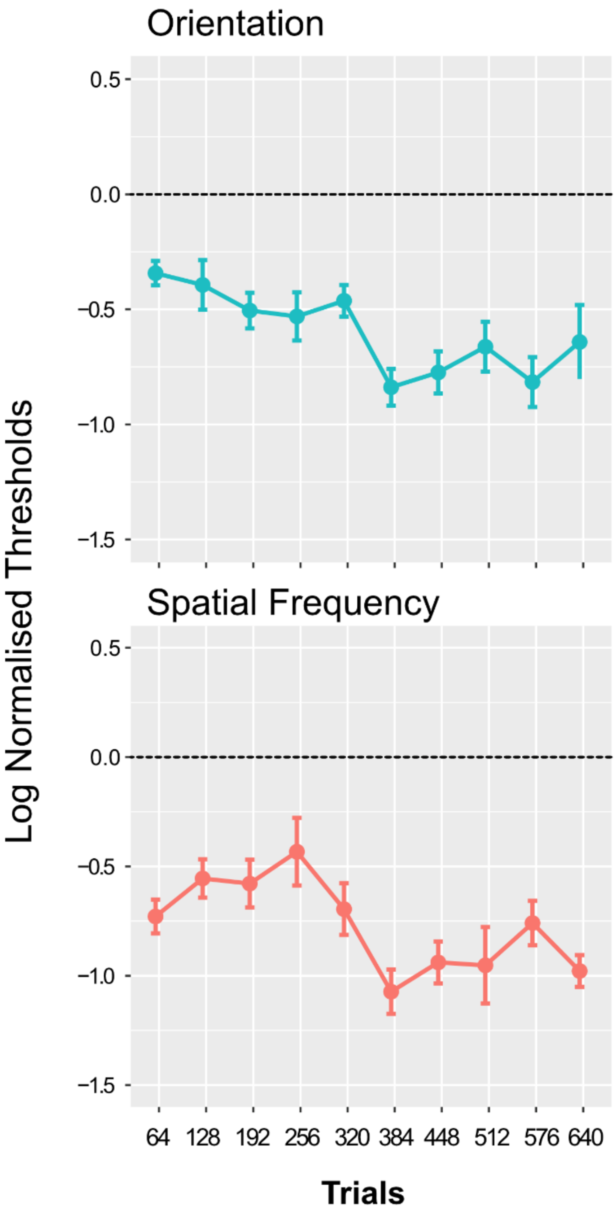

Figure 4.4 – Group-averaged log normalised learning curves for (top tier) Orientation Learning Group and, (bottom tier) Spatial Frequency Learning Group as a function of training trial. Both groups demonstrate learning on both tasks. Error bars indicate SEM across observers.

minimal slippage from day 1 (trial 256: -0.3) to day 2 (trial 320: -0.46). For the SF group (Figure 4.4 bottom tier), learning on day 1 increased toward pre-training threshold levels (trial 64: -0.72, trial 256: -0.43). However, this learning was consolidated continued to reach a similar level as the O group by the end of day 2.

The second objective of Experiment 1 was to examine the learning effects within and between dimensions post-training. To achieve this, the post-training performance was compared to the pre-training performance session for both the task-relevant and task irrelevant dimensions (Figure 4.5). In both groups, observers learnt on the task-relevant dimension they were exposed to during training. In the O group, the average post training discrimination threshold (mean = 3.16°) was significantly lower than pre training discrimination threshold (mean = 6.23° - t(9) = 2.21, $p < 0.05$). In the SF group, the average post training discrimination threshold (mean = 0.21) was significantly lower than pre training discrimination threshold (M = 0.09 - t(9) = 5.31, $p < 0.05$).

However, there was a different pattern of results for the between dimension comparisons for each group. In the O group, participants did not learn on the spatial frequency task-dimension indicating no transfer of learning between dimensions (mean pre: 0.16; mean post: 0.13; t(9) = 0.62, p = 0.54). In contrast, participants in the SF group did learn on the orientation task-dimension indicating transfer of learning to the task-irrelevant dimension (mean pre = 5.33°; mean post = 2.79°; t(9) = 2.64, $p < 0.05$). This suggests that training on an orientation dimension does not improve perceptual performance for the spatial frequency dimension.However, training on the spatial frequency dimension does improve perceptual performance for the orientation dimension.

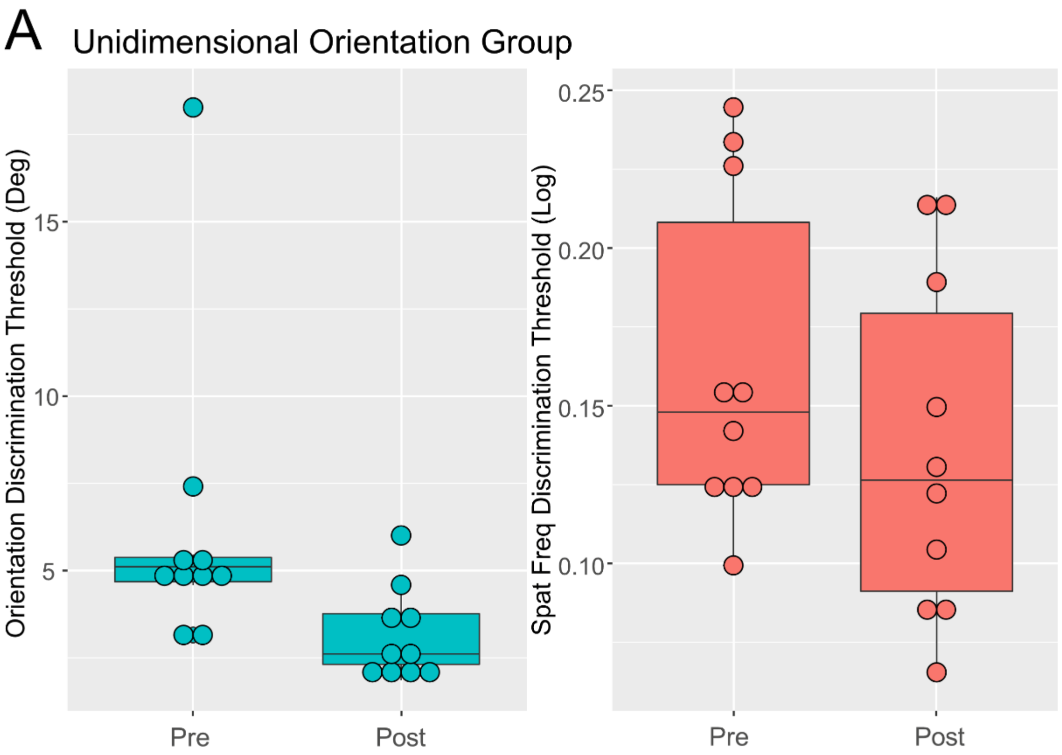
A Unidimensional Orientation Group

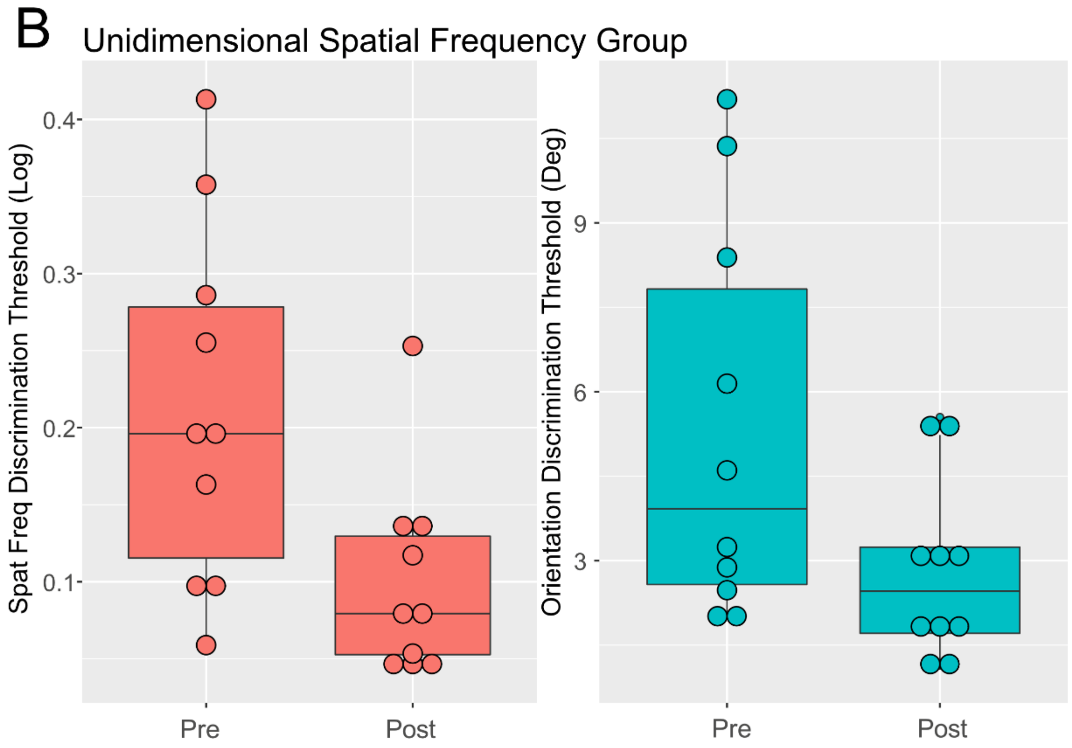
B Unidimensional Spatial Frequency Group

Figure 4.5 – Group-averaged Pre and Post performance thresholds for (A) Orientation Learning Group and, (B) Spatial Frequency Learning Group. Both groups demonstrate learning on their trained task. Learning in Orientation Learning group did not transfer to the spatial frequency dimension (A, right hand graph). Learning in the Spatial Frequency Learning group did transfer to the orientation dimension (B, right hand graph).

4.3.4 Discussion

The results from Experiment 1, reveal that both O and SF groups showed significant learning for the trained, task-relevant dimension. However, there was a difference between the groups in task-irrelevant learning, despite neither group being exposed to the other task dimension. In O group, learning did not transfer to the task-irrelevant spatial frequency dimension. However, in the SF group performance for the orientation dimension did improve post training despite not being exposed to the orientation dimension during training.

This is contrary to what has been observed in previous experiments where specificity to the trained spatial frequency dimension has been demonstrated (Fiorentini & Berardi, 1980, Szpiro et al., 2016). Of the previous studies that have looked at the interaction of PL across task dimensions after unidimensional training, all have reported a lack of transfer between task-dimensions (Fahle, 1997; Shiu and Pashler, 1992; Petrov and Van Horn, 2012). Therefore, what could explain the transfer of visual ability for the orientation dimension through training on the spatial frequency dimension in this experiment?

Despite some similarity at the level of visual input, the two tasks did share the same procedural structure. For instance, they shared the same base timing structure, including a 750-ms task cue, 150-ms stimulus presentation followed by a 500-ms noise mask. In addition, they also shared the same decision boundary that the stimulus levels in each block were drawn around. The sharing of procedural factors such as these across multiple tasks, has been found to be important for facilitating learning acquisition on subsequent tasks (Kattner et al., 2017). In this research, Kattner and colleagues (2017) used a range of different tasks (Gabor orientation categorisation, dot bisection tasks etc.) and either trained subjects on all five of the task or just one before performing a dot motion categorisation task to compare the rate of learning between the groups. Importantly, all tasks shared the same procedure, and they completed the same amount of trials during training. The group that performed all five tasks sequentially, had a faster rate of learning each time they switched to a new task and learnt more than those who only trained on one task, when the authors compared learning for both groups in a task after training. Possibly, training on the spatial frequency dimension facilitates learning on the orientation dimension provided the two task dimensions share the same procedural structure. Why the reverse of this was not true is a curiosity and perhaps discounts this line of thought.

Another factor that could explain between dimension transfer could be to do with the relationship between the training and transfer tasks (Jeter, Dosher, Petrov & Lu. 2009; McGovern, Webb & Peirce, 2012). For example, Jeter and colleagues (2009) showed that learning transfer across retinal locations was dictated by the relationship between the difficulty of the tasks. For instance, more transfer and less specificity was observed when switching to a low-precision task, while less transfer and more specificity was observed when switching to a higher-precision task which could be the case when training unidimensionally on spatial frequency and testing on orientation.

Another result consistent with Jeter and colleagues, showed that the type of trained task dimension determined the degree of transfer to other task dimensions (McGovern et al., 2012). McGovern and colleagues (2012) showed that the relative complexity of the trained task determined the level of transfer to other tasks. Training on a highly complex task, transferred more to a moderately complex task opposed to a less complex task. Potentially the spatial frequency dimension can be considered a higher precision task (more complex) and an orientation task a lower precision task (less complex). Therefore, given this possible dimension complexity between orientation and spatial frequency dimensions, it could be that learning transfers more to lesser complex orientation dimension after perceptual learning on the spatial frequency dimension and not the reverse. However, further research should explore this further by restricting the range of stimulus values (to control the precision of training and transfer) similar to Jeter and colleagues (2009).

Given that Experiment 1 has now established my training protocol, it will now be interesting to use the same paradigm for Experiment 2 however under a multidimensional environment. It will be interesting to compare the learning and transfer to the multidimensional group at the same timepoints, and whether state-dependent learning is still possible with the additional exposure to the task-irrelevant dimension. In addition, it would also be interesting whether a similar trend will also be found when participants train on the spatial frequency dimension first before task-relevancy switches to the orientation dimension.

4.4 Experiment 2

4.4.1 Rationale

In Experiment 1 I established an amount of training trials sufficient to result learning on a task dimension within a unidimensional environment and established the degree of transfer to each dimension after training without exposure on the other. The aims of Experiment 2 are twofold. The first aim is to understand whether participants can still learn on a task whilst passively exposed to variation along a coextensive task dimension that is irrelevant for the task. The second aim is to understand whether exposure to a task-irrelevant dimension facilitates transfer to that dimension when the task relevancy switches and this will be tested in two different task orders.

Passive exposure to a dimension that is temporally coextensive with the task relevant signal improves performance on that dimension (task-irrelevant) despite receiving no training (Watanabe et al., 2001; Wang et al., 2012; Xiao et al., 2008; Seitz et al., 2009; Tsushima et al., 2008; Huang et al., 2012; Zhang et al, 2010). In addition, exposure to the spatial frequency dimension that is spatially coextensive with the orientation dimension and shares the same decision boundary, facilitates learning on the orientation dimension (Szpiro et al., 2016). Therefore, I expect that when both task dimensions are spatially and temporally coextensive, that adaptation to the task-irrelevant dimension will improve learning on that dimension only when it becomes task-relevant. In addition, I also expect learning may be facilitated and transfer in the group who learns on spatial frequency first before learning on orientation. This is based on Experiment 1 where learning transferred when participants trained on the spatial frequency dimension when it was task-relevant.

4.4.2 Method

Twenty new participants took part in Experiment 2. They were randomly separated into either of two groups (10 participants each group) which were different depending on the order of dimension training. In the Orientation | Spatial Frequency group (O_SF) participants trained on the Orientation dimension first before the Spatial Frequency dimension. In the Spatial Frequency | Orientation group (SF_O) participants trained on the Spatial Frequency dimension first before the Orientation dimension.

The experimental procedure was similar in structure to Experiment 1. Participants completed the same pre- and post-training sessions. However, the training period differed in terms of the multidimensional environment and total amount of trials. In both the O_SF and SF_O group,s participants trained in a multidimensional environment where both the orientation and spatial frequency dimensions varied on each trial according to pre-training thresholds. The stimulus levels that were set for each dimension were always 3 JND widths either side of the discrimination boundary, which was determined by the psychometric function fitted to the previous block's performance. For the first training block this was Pre-training performance. Over the first 2 days of training, participants were trained on a task-relevant dimension whilst being exposed to variation along the task-irrelevant dimension (640 trials in total, 256 on day 1 and 384 on day 2). During the second day, the task-relevancy switched to the task-irrelevant dimension where participants then completed training on that task over 2 days (640 trials) whilst being exposed to the previous task-relevant dimension. Thus, the task relevancy switched in that what was task-irrelevant became task-relevant and vice-versa. In total participants completed 1280 trials during training but did the same amount of trials that participants completed in Experiment 1 on each dimension.

4.4.3 Results

The first objective of Experiment 2 was to understand whether participants learnt in a multidimensional environment. To achieve this the learning functions across trial were log normalised and averaged across participants following the procedure detailed in section 4.3.3 of this chapter. Figure 4.6 illustrates the average log normalised learning as a function of trials for the O_SF and SF_O groups (solid lines) against the average log normalised learning for the unidimensional groups in Experiment 1 (softer lines). The data for the O_SF

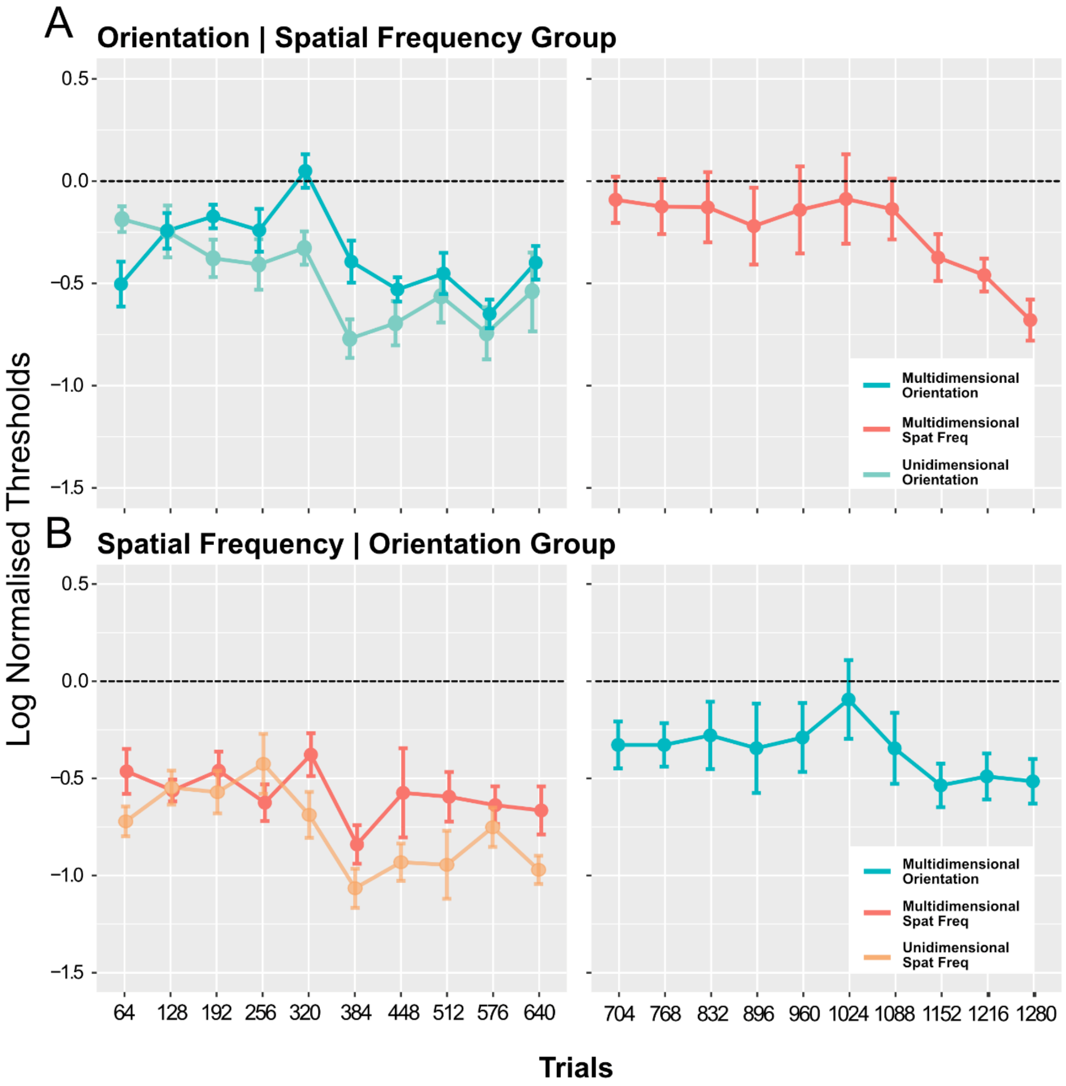

Figure 4.6 – A) Group-averaged log normalised learning curves for Orientation | Spatial Frequency Learning Group and B) Spatial Frequency | Orientation Learning Group as a function of training trial. Error bars indicate SEM.

group (Figure 4.6a) shows performance deteriorating within the first day of training on the orientation dimension (solid blue line). The average log normalised thresholds for observers in the first block (trial 64: -0.54) was lower than the average log normalised threshold at the end of the first day (trial 256: -0.24) suggesting that performance was getting worse. In addition, there were no signs of consolidation as the first block back on day 2 of training returned to pre-training levels (pre training raw score: 4.95°; trial 320 raw score: 5.13°). However, on the second day of training participant thresholds did improve over training. Interestingly, comparing the performance at the end of training (trial 640: -0.39) with performance at the start of training (trial 64: -0.50) shows that whilst performance was better than pre-training performance (log normalised thresholds are less than zero), there was a slight increase in performance over training but this was not significant (t(9) = -0.39, p = 0.70).

Comparing the orientation learning trajectories for the O_SF (Figure 4.6a – blue solid line) and O (Figure 4.6a – light blue line) groups reveal a clear difference in the trajectory that learning took. For instance, unidimensional orientation learning showed no obvious signs of performance deterioration across training and no slippage in performance at the start of day 2 (trial 320). In addition, participants were more accurate at the end of training following unidimensional orientation learning compared to a multidimensional environment (trial 640 unidimensional: -0.86; trial 640 multidimensional: -0.39) but this was not significant (t(9) = 0.66, p = 0.52).

The data for the SF_O group (figure 4.6b) shows performance improving on the first day of training and the log normalised threshold for observers in the first block (trial 64: -0.46) and at the end of the first day (trial 256: -0.62) were below pre-training performance levels indicating learning. In addition, there were signs of consolidation as the first block on day 2 of training was lower than the pre-training performance (pre training raw score: 0.25; trial 320 raw score: 0.13). Interestingly, comparing the performance at the end of training (trial 640: -0.66) with performance at the start of training (trial 64: -0.46) shows that performance did improve through training but this was not significantly different (t(9) = 0.52, p = 0.60).

Comparing the spatial frequency learning trajectories for the SF_O (Figure 4.6b – red solid line) and SF (Figure 4.6b – light orange line) groups also reveals differences in the trajectory that learning took. For the first day of training, learning for the SF_O group showed a slight deterioration on the first day and minimal slippage in performance at the start of day 2. Thus, in comparison to unidimensional SF group learning at least on the first day gradually improved opposed to a deterioration. In addition, participants performance reached the same point at the end of day 2 (trial 640) in both groups (trial 640 SF Group: -0.67; trial 640 SF_O Group: -0.64).

Another objective of Experiment 2 was to understand whether training in a state where participants are also exposed to the task-irrelevant dimension would facilitate performance when the task was switched, and when that dimension became task-relevant. The spatial frequency data for the O_SF group (Figure 4.6a – red solid line) suggests that this exposure caused a slight reduction in thresholds when training was switched to that task (raw threshold trial 704: 0.12) in comparison to pre-training thresholds (raw pre training thresholds: 0.18), however this was not significant ($t(9) = 1.7$, $p = 0.12$). In addition, the learning when the task switched deteriorated within day (trial 704: -0.32; trial 1024: -0.09), until the final day of training where thresholds performance improved to a similar level as the orientation dimension that preceded it.

A similar trend was also observed for the orientation data for the SF_O group (Figure 4.6b – blue solid line). Exposure caused a very slight reduction in threshold when training was switched to that task (raw threshold trial 704: 6.29°) in comparison to pre-training thresholds (raw pre threshold: 6.84°), however this was not significant ($t(9) = 0.15$, $p = 0.88$). Learning for that day plateaued for the rest of day 2 training (trial 1024: 7.17°), until the final day of training where thresholds performance improved to a similar level as the orientation dimension that preceded it (trial 640: -0.66, trial 1280: -0.65).

The final objective of Experiment 2 was to understand perceptual learning after training and exposure on both dimensions. To do this the Pre- and Post-training thresholds were compared. Figure 4.7 illustrates the pre and post training scores for the orientation (blue box) and spatial frequency (red box) dimensions for O_SF and SF_O groups. In the O_SF group, observers learnt on the orientation ($t(9) = 5.29$, $p < 0.05$) and the spatial frequency dimension ($t(9) = 2.3$, $p < 0.05$). Similarly, in the SF_O group, observers also learnt on the orientation ($t(9) = 4.84$, $p < 0.05$) and the spatial frequency dimension ($t(9) = 2.18$, $p < 0.05$).

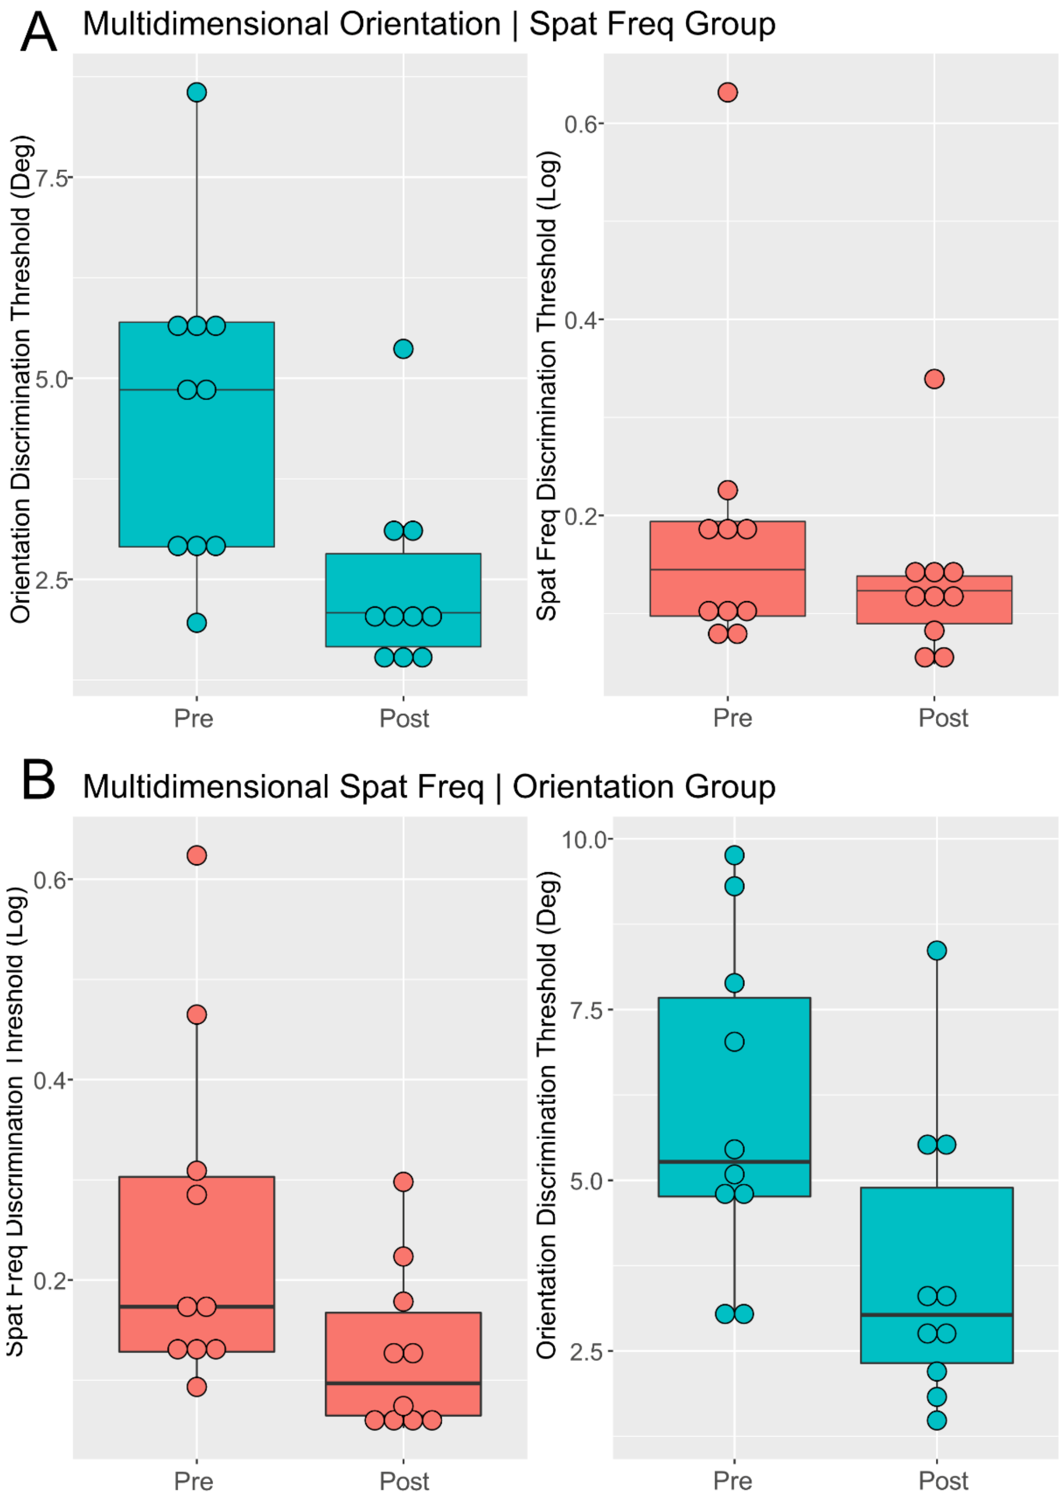

Figure 4.7 – Group-averaged pre- and post performance thresholds for (A) Orientation | Spatial Frequency Learning Group and, (B) Spatial Frequency | Orientation Learning Group. Observers in both groups learnt on both dimensions.

3.4.4 Discussion

The results from Experiment 2 reveal that exposure to variation along a task-irrelevant stimulus dimension neither hinders learning on the primary task or facilitates transfer. This was seen in our paradigm when the task relevancy was switched between tasks after adaptive exposure. There are some subtle differences between learning trajectories at the start of training for both groups which appears to depend on the order of training. When orientation was the first trained dimension (O_SF group) performance deteriorated within the first day of training and was not consolidated when they returned to complete training on the orientation task the next day. In the SF_O group however, there was no detriment to learning as there was no deterioration in performance on the first day and no slippage between training days.

One consistency in both the O_SF and SF_O groups was that adaptation did not facilitate any learning performance when the exposed task-irrelevant dimension became task-relevant. In both groups, performance was flat after the switch and did not improve until the final day of training. This is in disagreement with research suggesting that prolonged exposure to a task-irrelevant dimension facilitates perceptual performance (Wang et al., 2012; Xiao et al., 2008; Seitz & Watanabe, 2009; Tsushima et al., 2008; Zhang et al, 2010, Harris, Gilksberg and Sagi, 2012; Watanabe et al., 2001).

An interesting comparison to our experiments is that of Watanabe and colleagues (2001). The authors exposed observers to a random-dot stimulus attribute whilst participants performed a letter-identification task. Importantly, a sub-threshold percentage of the random dots drifted at a certain orientation throughout training and this repetitive exposure improved post training coherence thresholds at that orientation. There are many differences between Watanabe and colleague's experimental paradigm and the one used in Experiment 2, including stimulus used and task structure, however another obvious difference is the coextensive nature of the stimulus dimensions in space. In Watanabe and colleagues (2001) paradigm the exposed dimension was temporally present at the same time as the task-relevant signal, but was in another location in space. Perhaps together, this suggests that separating out and presenting the trained and adapted dimensions in distinct retinal locations in space, rather than being truly coextensive and multidimensional, is an important feature for improving performance after exposure.

Another important difference is that in our study, the adapted dimension was 3 JND widths either side of the discrimination boundary, which was based on each individuals pre-training performance. However, in Watanabe and colleagues (2001) and in other similar designs (Wang et al., 2012; Xiao et al., 2008; Seitz & Watanabe, 2009; Tsushima et al., 2008; Zhang et al, 2010), the adapted dimension is “sub-threshold” – based on another groups performance. Perhaps adapting participants to a dimension that is 1 JND width around the discrimination boundary for each individual participant, is important for adaptive learning to be effective. Further research, using a similar design to this one should follow this up and by adapting participants to stimulus intervals closer to their discrimination performance to understand this further.

An additional factor that could be important is the frequency of switching between task dimensions. By design, this experiment kept the amount of training and exposure exact across task dimension training and the task-relevancy switched for participants once training had ended for that dimension. However, in Szpiro and colleagues (2016) experiment, exposure consisted of training on the task-irrelevant dimension at frequent switching intervals within day. For instance, participants completed 800 trials within a training day split into four 200 trials blocks of either an orientation task, spatial frequency task or where both were successively interleaved with one another. It is a possibility that sequentially interleaving training like this on a much more frequent within day basis, may cause significant learning on the task dimensions.

A stimulus-roving paradigm (Berliner & Durlach, 1973) is an experimental procedure that has been investigated within the perceptual learning field where two or more stimulus tasks or variants are randomly interleaved with each other between trials. In some cases, alternating the task in such a way does not affect learning for each of the tasks. For instance, when the stimulus position is randomly interleaved (Censor, Karni & Sagi, 2006; Otto, Herzog, Fahle & Zhaoping, 2006; Karni & Sagi, 1991; Sireteanu & Rettenbach, 1995). However, in other examples stimulus-roving can prevent perceptual learning. For example, when four or more stimulus variants are randomly interleaved during training, no improvement of performance occurs (Kuai, Zhang, Klein, Levi & Yu, 2005; Yu, Klein & Levi, 2004).

It would be interesting to follow on the current experimental findings of Experiment 2 by trying to understand more about state-dependent perceptual learning over much finer timescales. In the current experiment I have shown that adaptation does not facilitate performance once learning is complete after 2 days, however interleaving the learning states on a much finer timescale could facilitate performance. For example, in Szpiro et al., (2016) both interleaved (roved) tasks within the same day. In other roving examples, tasks are interleaved between trials. For instance, performing a bisection task at different directions between trials (Otto et al., 2006). Thus, it would be interesting to understand whether adapted state perceptual learning between task dimensions is facilitated on much finer timescales. This will be explored in the next experiment.

4.5 Experiment 3

4.5.1 Rationale

In Experiment 2, the results revealed that learning can still be achieved in a multidimensional environment, however, exposure to variation along a task-irrelevant dimension does not benefit learning. In addition, the trajectory that learning took when the task-relevancy switched after they had been exposed to that dimension, suggesting that learning may have even been inhibited as learning on that day deteriorated or plateaued until the final day of training.

Following on from these findings and the work of Szpiro and colleagues (2016), in this experiment we consider whether the frequency of switching between trained adapted states will alter the path that learning might take on both dimensions. It could be the case that learning will be enhanced if the task-relevancy and adapted learning states are switched on a much more frequent basis. For example, participants in Szpiro's work switched task four times within training day, compared to just once in the entire training paradigm in Experiment 2. This experiment aims to test these questions further to understand whether switching the task-relevancy between tasks on a much finer time scale alters anything about learning on both task dimensions.

4.5.2 Method

Ten new participants took part in Experiment 3. The experimental set-up was very similar to Experiment 2 (see section 4.4.2); however, participants were now cued to one task dimension on one trial then the other on the next trial. The order of task-relevancy switching always followed an alternating manner. In order to compare the amount of training trials for each dimension against Experiments 1 and 2, the total amount of training for each dimension matched the amount of unidimensional training in Experiment 1 and the total amount of training trials in Experiment 2.

4.5.3 Results

The first objective of Experiment 3 was to understand whether participants still learn in a multidimensional environment when participants are cued between both task dimensions on a trial to trial basis. Figure 4.8a and c illustrates the average log normalised orientation (solid blue line) and spatial frequency (solid red line) learning data as a function of trial for the TS group against the average log normalised unidimensional O and SF (Figure 4.8a and c – lighter blue and red lines) learning functions. The data reveals performance improvements followed a similar trajectory to the unidimensional groups overtime. For orientation, both unidimensional and trial switch multidimensional learning followed almost the exact same path, until the end when trial switch learning performance surpassed unidimensional performance (unidimensional trial 640: -0.64, multidimensional trial switch group trial 640: -0.98). A similar pattern is seen between these groups for spatial frequency however, unidimensional learning outperformed trial switch performance by the end (unidimensional trial 640: -0.97, multidimensional trial switch group trial 640: -0.74).Figure 4.8b and d illustrate the average log normalised orientation (solid blue lines) and spatial frequency (solid red lines) learning data as a function of trial for the TS group against the average log normalised unidimensional O_SF and SF_O (Figure 4.8b and d – lighter blue and red lines) learning functions. The lower panels Figure4.8b and d show the TS learning functions against the learning functions when the task-relevancy switched in Experiment 2.

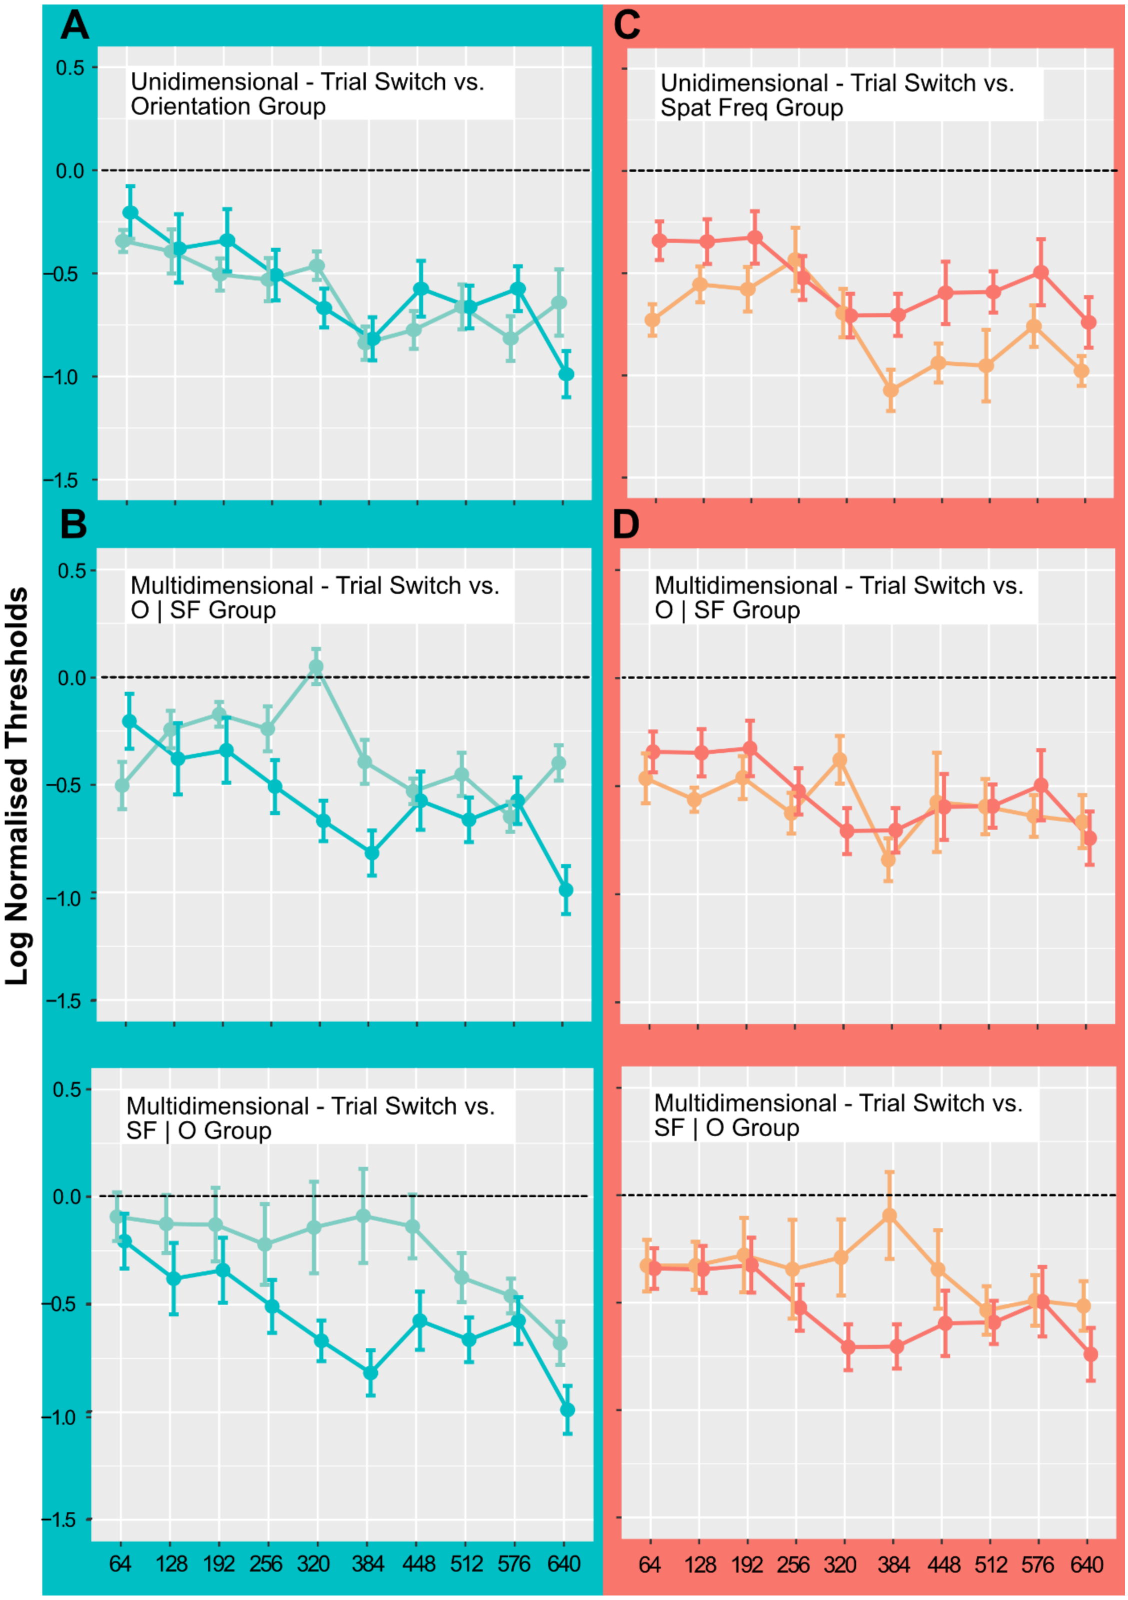

Figure 4.8 – Group-averaged log normalised learning curves for Orientation (left column) and Spatial Frequency (right column) dimensions for the Trial Switch Group against the unidimensional (A, C) and O_SF and SF_O groups (B,D). Error bars indicate 1 SEM.

For orientation (Figure 4.8b), the TS group in comparison to the O_SF group (top-tier Figure 4.8b), learning showed no deterioration or slippage between days. The same comparison can be made with the TS group with when the task relevancy switched to orientation in the SF_O group (bottom tier Figure 4.8b). For spatial frequency (Figure 4.8d), the TS group in comparison to the SF_O group (top-tier Figure 4.8d), the learning trajectories closely followed each other. A similar comparison can be made with the TS group with when the task relevancy switched to spatial frequency in the O_SF group (bottom tier Figure 4.8d). Here we find that apart from it plateauing and rising in the O_SF group, learning in both groups reached the same point at the end of training. In combination, this suggests that multidimensional learning for orientation may have been facilitated in more successive switches, whereas multidimensional spatial frequency learning may not have benefited from successive switching.

The second objective of the Experiment 3 was to understand whether there was any effect of training in a trial switch manner on pre- and post-thresholds (Figure 4.9). The data

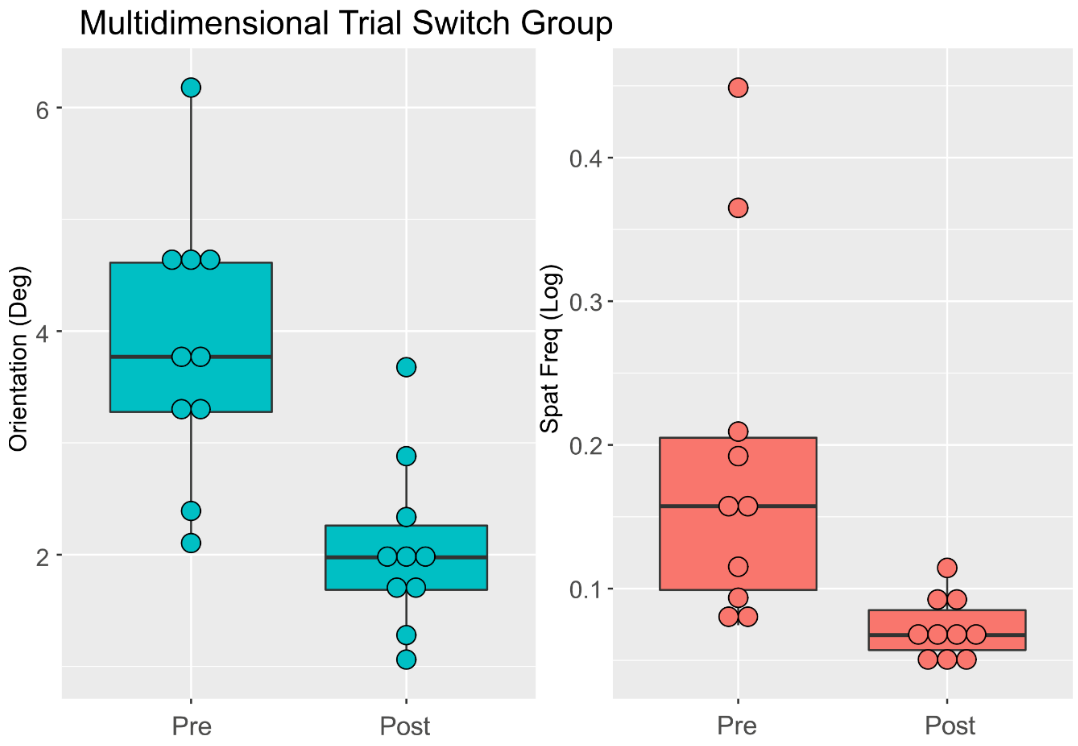

Figure 4.9 – Group-averaged pre- and post-performance thresholds for Trial Switch Group.

revealed that participants learnt on both the orientation ($t(9) = 3.61$, $p < 0.05$) and spatial frequency task-dimensions ($t(9) = 4.12$, $p < 0.05$) after training, indicating that learning was achieved on both dimensions with successive switching between them on trials.

3.5.4 Discussion

The results for Experiment 3 reveal that state dependent perceptual learning is still achieved even after interleaving training on two task dimensions between trials. In Experiment 2, in the O_SF and SF_O groups we showed that adaptation to another dimension does not facilitate learning when it becomes the trained dimension, and in this experiment wondered whether having more successive switching would facilitate learning on either dimension. The results reveal that switching successively did not affect the time course of learning on either dimension, like it appeared to do when the task relevancy switched in the O_SF and SF_O groups. This would suggest that whatever was driving the learning to plateau in Experiment 2, was diminished by successive switching in this experiment. In addition, in comparison to Experiment 1 where participants completed the training in a unidimensional environment, learning followed a very similar trajectory overtime for both groups learning trajectories.

It is important to note that under the current paradigm it is difficult to observe the exact effects of adaptation exposure on learning performance. Apart from comparing the trajectories that learning took overtime, there are no timepoints along these learning trajectories for either dimension that would reveal anything useful about adaptive exposure. However, in comparison to the other groups, the fact that both dimensions followed the same path is encouraging and suggests that state-dependent learning between the dimensions may have had some dependency with one another. In addition, switching between trials reduces the effects of adaptation exposure that we witnessed when the task switched in the O_SF and SF_O groups.

3.6 Between Group Analysis

To test whether there were any differences between the Pre and Post thresholds of the experimental groups, a 5 (group: O, SF, O_SF, SF_O, TS) x 2 (Pre vs Post) mixed ANOVA was conducted on the raw thresholds for each task dimension separately. For orientation, the analysis revealed a non-significant main effect of group (f(4,45) = 1.38, p = 0.25), a significant within effect between pre and post-performance (f(1,45) = 42.5, p < 0.001) and a non-significant interaction (f(4,45) = 0.22, p = 0.9). For spatial frequency, the analysis revealed a non-significant main effect of group (f(4,45) = 1.12, p = 0.35), a significant within effect between pre and post-performance (f(1,45) = 32.26, p < 0.001) and a non-significant interaction (f(4,45) = 1.98, p = 0.11). Overall, this indicates that there was no difference between each group in pre-training performance and no group outperformed the other as a result of training.

Finally, to understand the effect of training in a multidimensional environment, a 2 x (Task: Orientation or Spatial Frequency), 2 x (Dimensionality: Unidimensional or Multidimensional) x 10 (Training Block) three-factor between-within ANOVA was conducted on the normalised data. This analysis violated Mauchly Test of Sphericity for the within-subject effect block X^2(44) = 126.72, p = <. 001), therefore, Greenhouse-Geisser correction will be applied to the degrees of freedom to assess a critical F-value. The results revealed a significant within subject effect for Training block (f(6.49,493.42) = 3.932, p = 0.001) and this did not significantly interact with Task (f(6.49,493.42) = 1.225, p=0.290) or Dimensionality (f(6.49,493.42) = 0.679, p = 0.67). There was also no significant interaction between Block, Task and Dimensionality (f(6.49,493.42) = 0.730, p = 0.636). Between subject comparisons revealed a significant main-effect for Dimensionality (f(1,76) = 4.548, p = 0.048). The collapsed means across dimensionality revealed that those in the multidimensional groups had higher thresholds across block compared to those in the unidimensional groups. The analysis revealed a non-significant main effect of Task (f(1,76) = 0.572 , p =0.452). The analysis also showed a non-significant interaction between Task and Dimensionality (f(1,76) = 0.025 , p =0.875).

4.7 Discussion

In this chapter, I explored the limits of state-dependent perceptual learning through training and exposure on a task-relevant and task-irrelevant dimension when they are coextensive with each other. To achieve this, 5 different experimental groups either completed training in a unidimensional environment (Experiment 1) or in a multidimensional environment where task-relevancy switched once training on the task had completed (Experiment 2) or switched between trials (Experiment 3). One of the main purposes of training in a multidimensional environment was to understand whether adaptation exposure to another dimension would facilitate learning when it became task relevant.

Although the experiments provide a somewhat mixed pattern of results, overall they demonstrate that perceptual learning can still occur even in multidimensional environments where the addition of another dimension can be considered noise. In somewhat contradiction to previous experiments that have shown perceptual improvement following passive exposure to a visual dimension, the results here suggest that adaption does not facilitate perceptual performance. In addition it appears that the overall magnitude of learning is unaffected by the multidimensional nature of stimuli and although different task-relevancy switching can change the shape of this learning, participants generally reach the same end goal regardless of manipulation.

Contrary to previous results that have shown cross-dimension transfer (Szpiro et al, 2014; McGovern et al., 2012; Watanabe et al., 2001, Wang et al., 2012; Xiao et al., 2008; Seitz & Watanabe, 2009; Tsushima et al., 2008; Zhang et al, 2010), we found no benefit for state-dependent perceptual learning when participants were adapted to a task-irrelevant dimension. Instead the research agrees with the results that show learning is specific and independent to the task trained (Crist et al., 1997; Fahle, 1997; Shiu and Pashler, 1992; Petrov and Van Horn, 2012), even if the task dimensions share identical stimulus inputs (Huang et al., 2012). One of the differences between the research reported here and most of the research investigating cross-task transfer, is the separation of task dimensions across retinal space. Therefore, it is possible that separating the dimensions in space, a feature that is unique to lab-based paradigms, leads to better transfer across dimensions and tasks.

A good real-life example of this comes from an experiment that used lab-based perceptual paradigms to improve the performance of baseball players on the field (Deveau & Seitz 2014).

However, a surprising result in the current experiments is the asymmetric transfer of visual ability when the task-relevancy switched in the O_SF and SF_O groups. In comparison to Huang et al., (2012), learning in experiment did not show any deterioration like the results reported here when task relevancy switched following exposure. Some research has reported within-session deterioration on perceptual tasks (Mednick, Arman & Boynton, 2005; Aberg, Tartaglia & Herzog, 2009) and show that within-session deterioration may be linked to the amount of trials a participant is subjected to within day. Thus, perhaps a likely candidate for this deterioration in the O_SF and SF_O groups may be attention related mechanisms. However, it is hard to ascertain whether this is truly the cause or whether adaptation to that visual dimension before switching did inhibit learning until the following day.

One core assumption in this experiment was that passively exposing participants to a dimension would act like adaptation, which would allow us to test whether this passive based exposure would facilitate perceptual performance. An alternative explanation, that could help explain the asymmetric transfer, could be that participants were not adapting to the task-irrelevant dimension during the first task. Rather, as the task-irrelevant dimension was coextensive with the task-relevant dimension during training, part of the learning process for the initial task-relevant dimension was ignoring and filtering out the task-irrelevant dimension. A study that investigated the importance of attention in perceptual learning, showed that attentional suppression of task-irrelevant stimuli becomes more efficient over practice (Vidnyanszky & Sohn, 2005). In this example, they had a dot-motion stimulus that included a signal that was task relevant and another separate coextensive signal that was task-irrelevant. Importantly, the task-irrelevant signal was adding a level of noise to the task-relevant signal and became task-relevant later on in the experiment to assess any change in processing. They showed that through practice that requires a specific subset of the visual input to be attended resulted in more efficient suppression of the processing of task-irrelevant visual stimuli. Therefore, the results for the O_SF and SF_O groups, could reflect a similar process which could explain the deterioration and plateauing of learning when task-relevancy switched. The lack of this in the TS group, where relevancy switches between trials, perhaps is an indication that attentional suppression of the task-irrelevant dimension was not given a chance to alter learning trajectories of either dimension.

One of the challenging aspects to the design of these experiments was matching the two tasks in terms of difficulty. Care was taken during the piloting phase to ensure that the two task dimensions were relatively matched when either discriminating orientation or spatial frequency so that state-dependent adaptation and perceptual learning can be accurately compared between them. However, one caveat to this is that when ultimately the final design was tested on a much larger cohort, some of these new participants would have found certain tasks easier to learn over time than the other and initial performance for each task would have been different. Interobserver variance in perceptual learning paradigms has been investigated before and the results have indicated that even in homogeneous participant groups (age, intelligence and motivation matched), baseline performance and the extent of learning vary widely (Fahle & Henke-Fahle, 1996). The authors of this study, who recruited 70 participants to complete a bisection discrimination task, also indicate that initial threshold variation also dictated how much and how quickly that participant would learn, with those having the worst initial performance levels learning the most. What this might suggest for this experiment is that relatively homogenous initial baseline performance is hard to achieve for a task (let alone for two tasks), and that potentially the state-dependent adaptation and learning effects are just a product of this variation.

To conclude, this chapter found that alternating between task dimensions that share the same multidimensional environment during training neither facilitates or inhibits perceptual learning across task. Instead the results suggest that participants can learn despite these manipulations as the brain has an impressive capacity to achieve learning. Surprisingly, there are differences in the learning trajectories between groups where the switching between task-dimensions differed and this may be due to the attentional mechanisms. Overall, these findings shed light on some of perceptual learning in a multidimensional environment that is similar to perceptual environments in the real-world.

Chapter 5: No differences in the perceptual consequences of adapting to externally-generated or self-generated visual motion

5.1 Introduction

Our experience of the environment is manifested through our various senses. As our complex environment changes, our sensory systems effortlessly keep our experience continual as well as stable, and this process is an integral part to many of our different everyday tasks. One such task is our ability to navigate through our surroundings – which is a process that does not exclusively rely on any one of our senses. For example, when riding a bicycle, our vestibular system keeps us balanced on the seat whilst our visual attention keeps us from colliding into objects in our surroundings by processing what is in front of us.

The natural viewing condition just described creates a type of visual motion, such as the pavement moving underneath us or objects passing us by in our periphery, that is commonly referred to as radial optic flow (Gibson, 1950). When this motion is recreated under laboratory settings, it induces a motion aftereffect (MAE) whereby after exposure to coherent motion in the visual field, stationary stimuli appear to move in the opposite direction (for example: Patterson et al., 2009). This resulting aftereffect has been used to elucidate several properties of the various stages within visual cortex that process motion (for review see Mather et al., 2008). As valuable as the laboratory examples are, there may be a difference in how we process motion in the real-world compared to the laboratory setting. For instance, if we think about our example involving the bicycle, if we suddenly stop at a traffic light, our visual world does not suddenly appear to move in the opposite direction, rather it remains stable.

Whilst this in no way casts doubt on the existence of the MAE, it does however raise interesting questions for adapted state bias and factors that influence visual processing that warrant further research. One factor that has been explored as an explanation for this difference between real-world and laboratory viewing settings, has been vestibular input (change in position or movement of the head). This was first tested by Harris and colleagues (1981), who examined whether congruency between vestibular and visual motion input altered the resulting MAE. In this research, participants stood on a trolley and were pushed either forward or backwards whilst at the same time either being shown radial dot-motion that expanded (as if moving toward) or contracted (as if moving away). Immediately after being pushed in a direction and viewing the motion, participants reported the magnitude of the MAE using the method of cancellation.

The author's results revealed a reduced MAE after viewing expanding radial motion during forward motion, in comparison to much larger adapted state bias following incongruent conditions (i.e. backward motion during expanding motion) or when the participant was stationary and the motion was externally generated. However, this same congruency effect was not true for the reverse despite the congruency between visual and vestibular inputs, i.e. viewing contracting motion during backwards motion. Harris and colleagues (1981) concluded that the visual and vestibular systems calibrate with each other and as result reduce adapted state bias. This would explain the lack of aftereffect under real-world settings. In addition, it is more natural that we move forward in space, which they argue explains why the MAE was unaltered despite congruency when moving backwards through space.

Perhaps another factor that could change adapted state bias in natural settings, is whether the motion is self-generated. This is plausible since in many real-world examples we are in control of the viewed motion. Going back to our bicycle example, our motor system is an integral part to generating and controlling the radial motion, as we pedal with our feet, use our arms to turn the handlebars and squeeze the breaks with our fingers to control our speed and thus the viewable motion. Whereas, in the laboratory, the motion and therefore adaptation is externally generated as participants typically make their response after passively viewing the motion.

A recent experiment explored whether self-generated head-movements altered adapted state bias for the MAE (Bai et al., 2018). In their paradigm participants moved their head whilst a head-mounted display showed a grating that drifted either in the same, opposite or perpendicular direction to the head's rotation. After their head movement, participants reported the duration of the resulting MAE. Importantly, unlike in Harris and colleagues (1981) research, the participants head-movement generated the visual motion that was displayed and thus were active in controlling it. The authors results revealed that the MAE lasted significantly shorter when participants generated the motion in comparison to when their heads were still. In addition, this effect was not specific to a certain direction of motion or directed movement of the head. Instead, it appeared that the combination of self-generated movement of the head and motion caused a reduction in motion aftereffect bias across all directions of adapted motion (Bai et al., 2018).

Bai and colleagues (2018) results are compelling and perhaps they can be interpreted that factors relating to self-generated motion are also responsible for reducing adapted state processing and bias. It is important however to note that in their experimental design, it is difficult to disentangle the separate effects of the vestibular factor and the self-generated motion factor on the resulting adapted state bias. This is because by moving their heads, participants were generating the viewed motion which was also activating their vestibular system. Thus, it remains to be determined how the adapted bias is altered through self-generated motion without vestibular input.

To date, a few studies have explored how manual control of a stimulus in a task influences various other visual aftereffects involving motion (Ichikawa & Masakura, 2004; 2006; Ichikawa et al., 2010; Konkle et al, 2009; Dewey & Carr, 2013), as well as other studies that have explored visual perception after self-generated motion more generally, such as the attenuation of self-generated stimuli (Rhode et al., 2014; Lally et al., 2011; Christensen et al., 2011). To begin with the former, Ichikawa & Masakura (2006) set-up an experiment where participants manipulated and controlled the state of a moving stimulus across a computer monitor. Participants moved a computer mouse either in an upward or downward motion which matched the movement of the stimulus on screen. At a point along the moving stimulus's trajectory a secondary stimulus flashed on screen briefly (13ms). This paradigm has typically been conducted when the motion is externally generated, and is known as the flash-lag effect (Nijhawan, 1994) as results typically show a bias to the secondary flashing stimulus position, where it is perceived in a lagged position relative to its true position next to the moving stimulus. This effect however, on the true position of the moving stimulus was significantly reduced when the participant generated the movement of the stimulus with their hand in comparison to the passive viewing condition and when the moving stimulus went in directions opposite to the direction the hand moved in. This result implies that there may be some unique function between motion and temporal visual motion processing under self-generated conditions.

Another interesting group of experiments have shown that perceptual discrimination and detection performance may be more accurate when the viewable stimulus is self-generated in comparison to when the stimulus might appear to be externally-generally (Rhode et al., 2014; Lally et al., 2011; Christensen et al., 2011). In Christensen and colleagues (2011) they measured the effects of motion detection when it either matched the participants' movement of the arm, or was temporally delayed, rotated 90° or inverted. Importantly, these manipulations detracted from the motion appearing to be self-generated which interestingly, appeared to significantly impair when it was spatially or temporally incongruent. This suggests that self-generated stimuli are processed differently in the form of attenuation.

In combination, these findings highlight that self-generated visual stimuli in our environment can alter visual processing. Although, somewhat amazingly, no study I could find has tested whether the MAE is altered when it self-generated. The work by, Harris et al., (1981) and Bai et al., (2018) have shown that vestibular input can alter the aftereffects whilst adapting to motion, however it remains unclear what role self-generated adaptation has on the MAE. Other studies have shown that visual processing is altered, a general finding is that it appears to be more accurate (Ichikawa & Masakura, 2004; 2006; Ichikawa et al., 2010; Konkle et al, 2009; Dewey & Carr, 2013; Rhode et al., 2014; Lally et al., 2011; Christensen et al., 2011). The purpose of this investigation is to add to this debate by understanding whether adaptation to motion is altered when the motion is self-generated by the participant who is stationary, thus removing vestibular input.

5.2 Method

5.2.1 Participants

A total of 5 right-handed participants (Mean Age = 24.4, STD = 1.85) gave informed consent and took part in the experiment. All had normal or corrected to normal vision and were paid an inconvenience allowance for their time.

5.2.2 Apparatus and Stimuli

Participants sat approximately 2 meters away from a semi-circular screen (radius = 2.5m, height = 2.5m) that wrapped 180° around the participant horizontally (see Figure 5.1). Each participant had their head secured using a chin rest and in front of them sat a trackball computer mouse (Kensington, Orbit Wireless Trackball M01380-M). The computer mouse's trackball was used to record participants swipe movements, it could rotate along two axes and could also record participant responses. Three projectors were situated behind and above the participant which projected onto the semi-circular screen in front of them. More details regarding the projector set-up can be found in the Methods chapter (Chapter 2 – section 2.1 of this thesis).

The presented stimuli were controlled using PsychoPy (Peirce, 2007) a custom computer software written in Python. The stimuli consisted of an adapting and test grating (both were 90°, 1.25 cpd) that were each displayed within a circular aperture that measured 200cm in diameter. The adapting grating drifted at a maximum temporal frequency of 4.8 deg/s (6 Hz) and was controlled via the swiping motions participants made using the trackball of the mouse in the experimental conditions (details see 5.2.3). The adapting grating was static until the participant started moving the trackball. After 1 full swipe of the trackball the temporal frequency of the grating increased by 0.24 deg/s (0.3Hz) in that direction of the swipe. This was a linear increase in speed over 60 frames, so that there was no dramatic jump in the speed of the grating. The temporal frequency of the adapting grating could also slow down due to inactivity (2 seconds) or swiping in the other direction and decreased in speed at the same rate as it increased.

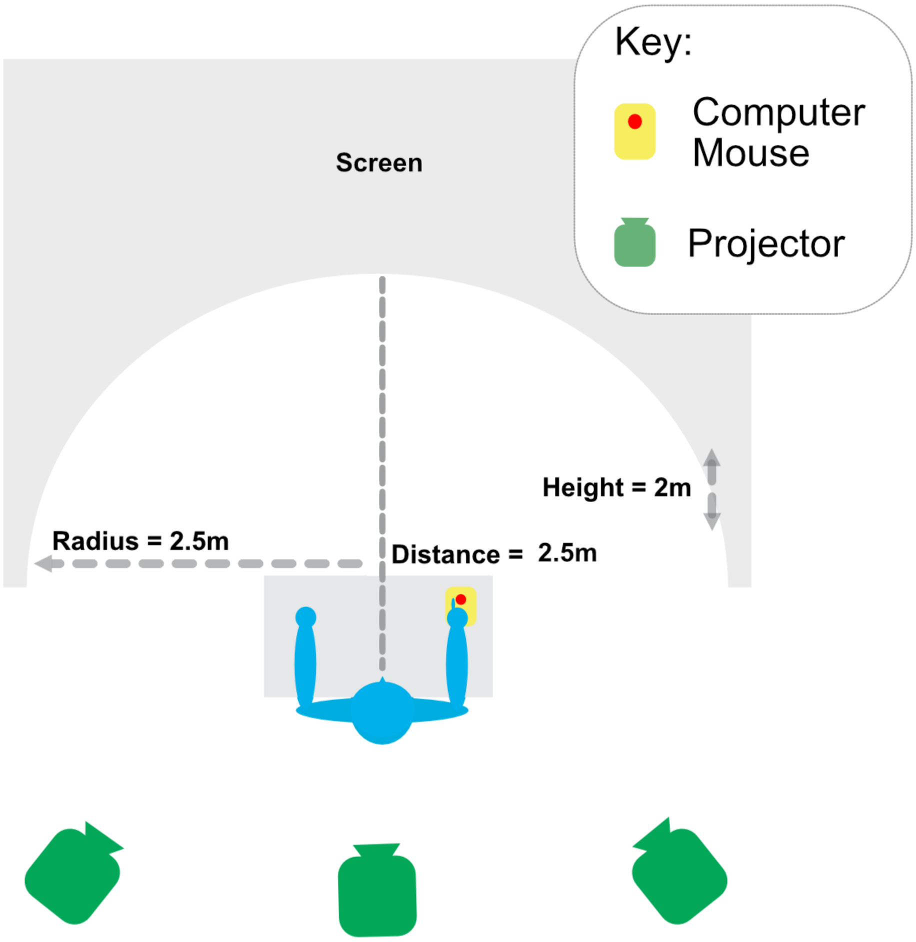

Figure 5.1 – *Experimental Set-up*. Participants sat approximately 2.5 meters (m) from a semi-circular screen (2.5m radius, 2m height) with their head secured with a chin rest. Participants used the mouse to move the grating on the screen during the adapting period. They did this by moving the trackball on the computer mouse. Three projectors projected the stimulus on the screen.

The test grating composed of two superimposed gratings presented on top of each other, drifting in opposite directions (6Hz). The contrast-ratio of each opposing grating was controlled under the nulling method. A summary of the nulling method is given in the next section but has been described in more detail in Chapter 3 (section 3.2.4) of this thesis.

5.2.3 Procedure

The experiment lasted a total of 3 hours for each participant and was completed over two days. Each session lasted approximately 1.5 hours. On the first day participants were given practice at the experiment before moving onto the main experimental blocks. Participants were first asked which hand they predominately used to operate a computer mouse. Depending on their answer, the computer mouse was situated on that side and the experimental grating stimulus was also displayed on that side, i.e. right-hand dominant led to the mouse and the stimulus being projected onto the right side of space.

Before the main task, participants completed a brief exercise where they were shown a static grating on the screen and were instructed to make the grating drift either in the left or right direction by using the trackball in the computer mouse in front of them. The grating drifted at a maximum temporal frequency of 4.8 deg/s in the direction they swiped in and participants were encouraged to slow the grating down and speed it up during this exercise. The purpose of this exercise was to ensure they gained familiarity with moving the grating for the purpose of adaptation in the main experiment.

The next part of the practice exercise consisted of participants completing a brief 15 trial block of the nulling method task. During this block participants were passively adapted to a coherently moving grating before each trial. This was the same as the procedure described in Chapter 3 (section 3.2.4) of this thesis. In brief, after the adapting grating appeared on screen (30 secs before first trial, 5 secs there-after between trials) a test grating appeared (0.5 secs) on the screen that was comprised of two opposing gratings travelling in opposite directions. The contrast-ratio of the two opposing test gratings was controlled via two 1-up 1-down interleaving staircases to perceptually null the MAE.

Once participants had completed the practice exercise and were comfortable with what was required of them, they were then briefed about the task. There were three experimental conditions that participants completed. There were two adapting conditions (see Figure 5.2) and an unadapting condition. In the *active* (Self-Generated) adapting condition participants generated the motion adaptation before each trial. In the *passive*

(Externally Generated) adaptation condition participants were played back a cycle of adaptation. In the *passive* condition, the cycle of adaptation mirrored the exact pace of adaptation during the *active* condition. This was achieved by cycling through the swipe trajectories that the participants themselves had made during the *active* conditions.

In the *unadapted* (no adaptation) condition, participants completed the nulling procedure with no adaptation. All conditions were counterbalanced appropriately. If participants had a *passive* block first, the swipe pace of the adaptation was taken from a bank of swiping trajectories taken from other participants. During piloting of this experiment, it was established that generally participants swiped their finger across the trackball in a relatively consistent manner. This was probably due to the instructions that they should swipe their finger across the trackball of the mouse using their fingers to keep the grating drifting consistently.

Each adapting condition block under the nulling method lasted approximately 15 minutes and was comprised of 60 trials. The unadapted condition lasted approximately 3 minutes and was comprised of the same 60 trials. A 5-minute break was given in-between blocks to avoid any cross contamination between blocks.

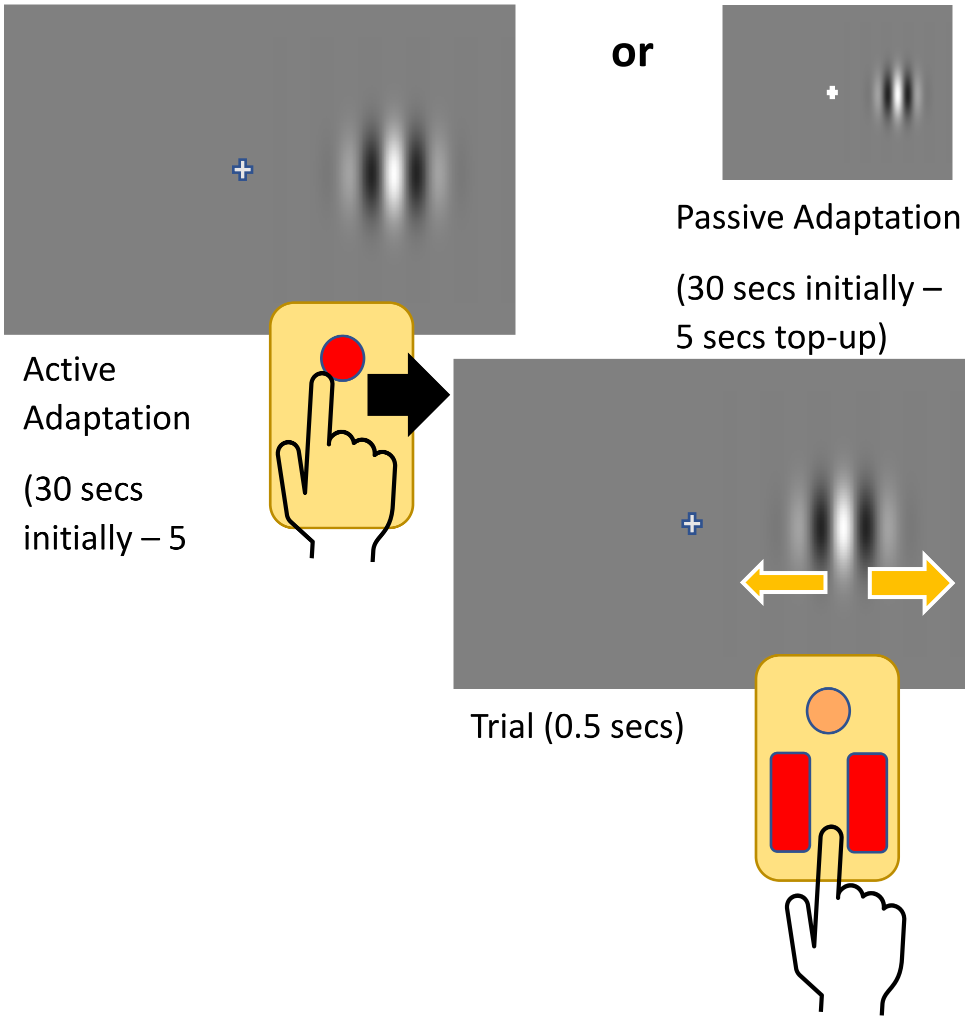

Figure 5.2 – *Active and Passive conditions*. During the active condition participants swiped their finger(s) across the trackball of the mouse to move the grating on the screen during the adaptation period. In the passive condition they were played back an adaptation run. After each trial, participants responded whether to the test grating was drifting more to the left or to the right.

5.3 Results

The psychometric functions for each participant, fitted using maximum likelihood estimation, can be found in Figure 5.3, with their associated nulling point (PSE) values underneath each plot. Each participant's psychometric function shows the proportion a rightward response as a function of the log contrast ratio of the leftward and rightward test grating components. To begin, the unadapted (green line and dots) psychometric function for each participant were roughly centred on the log contrast ratio axis, whereas the

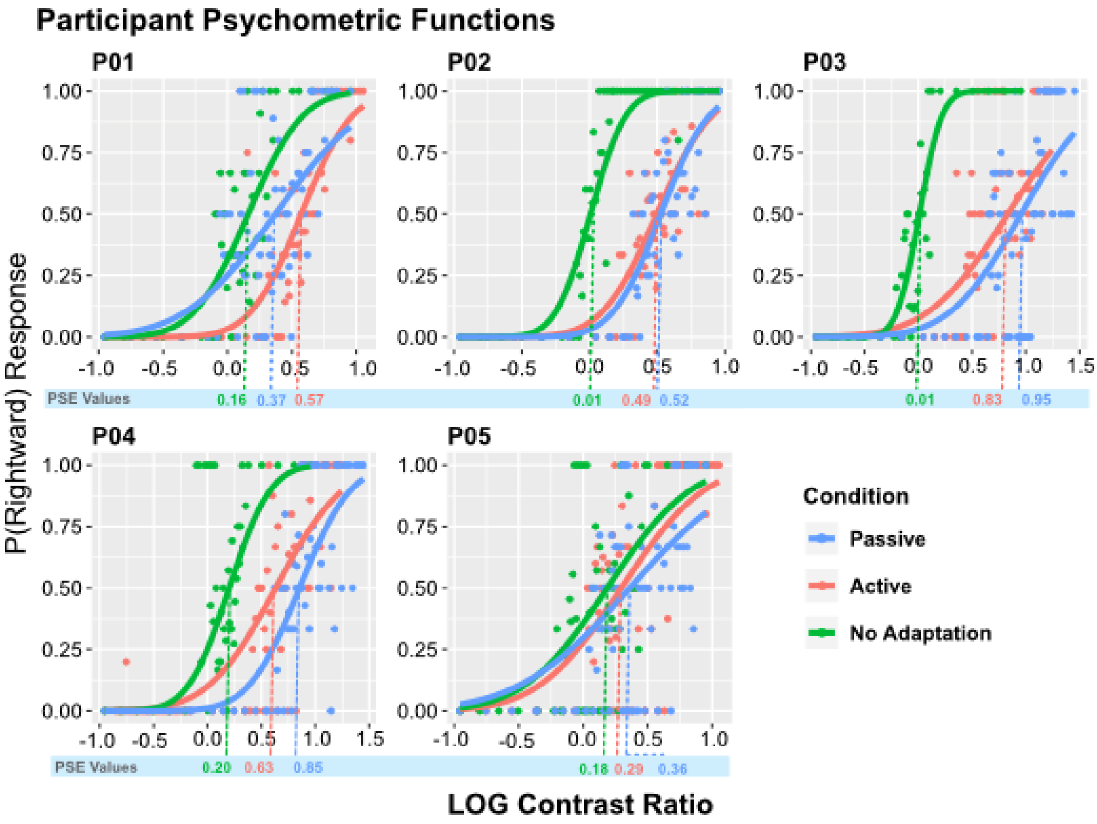

Figure 5.3 – Participant Psychometric Functions for the Passive, Active Adaptation and No Adaptation Conditions. Each line and associated dots represent the fitted psychometric function and proportion rightward response for each log contrast ratio level. The blue line and associated blue dots represent the passive (externally-generated) adaptation condition. The red line and associated red dots represent the active (self-generated) adaptation condition. The green line and associated green dots represent the no adaptation condition. Each graph represents the psychometric functions for the conditions for each participant. Nulling point (PSE) values are displayed underneath each participant graph.

adapted psychometric functions, passive (blue line and dots) and active (red line and dots) were shifted rightwards along the log contrast ratio axis. The unadapted psychometric functions were also roughly central along this axis. Participants 1, 4 and 5 showed a slight bias in their response for net motion in the rightwards direction, which is reflected in their nulling point values.

The psychometric functions for each participant in the adapting conditions were clearly biased in their responses in comparison to unadapted condition, which is indicative of a MAE under these adapting conditions.

Further scrutiny of the psychometric functions and corresponding nulling points reveals that for participants 1, 3, 4 and 5 there were differences between the psychometric function and nulling points between the adapting conditions. For instance, for participants 3, 4 and 5 their active psychometric functions were shifted slightly more centre and to the left of the their passive adaptation psychometric functions, indicating that adapted state bias was reduced in self-generated adaptation conditions. This similar pattern is also represented by the same shifts in their nulling points. However, for participant 1 this difference between adapting condition psychometric functions was the reverse, as the passive psychometric function was shifted more to the centre and leftward of the active adapting psychometric function. This was also reflected in their nulling point between the adapting conditions.

The psychometric functions and nulling points for participant 2 on the other hand, showed small differences between the adapting conditions. Their psychometric functions were roughly in the same position along the log contrast ratio axis and had a similar shape to them. The nulling point for the active adapting conditions was slightly less by about 0.03 log contrast points than the passive adapting condition, however generally in comparison to the other participants this was a small difference.

Figure 5.4a shows boxplots for the adapting nulling points for both the active and passive conditions and Figure 5.4b shows a boxplot for the nulling point for the unadapted condition. As the nulling point data shows, the active and passive adaptation conditions were more biased than the unadapted conditions. In comparison, the average for each condition across participants revealed an average nulling point for the unadapting condition of 0.11 (std = 0.09), and an average nulling point for the adapting conditions of 0.56 (std = 0.19) for the active condition and 0.61 (std = 0.27) for the passive adapting condition.

To understand whether these differences were significant, a repeated measures one-way ANOVA with condition (Passive, Active, No Adaptation) as a factor was run on the nulling point data. The analysis revealed a significant effect of condition on the nulling point thresholds (f(8) = 11.952, p < 0.01, ηp^2 = 0.6). Post-hoc paired t-tests revealed that there was a statistically significant difference between the no adaptation and active adaptation condition (t(4) = 3.99, p < 0.01) and a statistically significant difference between the no adaptation and passive adaptation condition (t(4) = 3.49, p = 0.02). This highlights a significant MAE was produced in each adapting condition. However, the difference between the active and passive adapting conditions was not significant (t(4) = -0.69, p = 0.52) indicating that there was no difference the MAEs induced by whether self- or externally-generated motion.

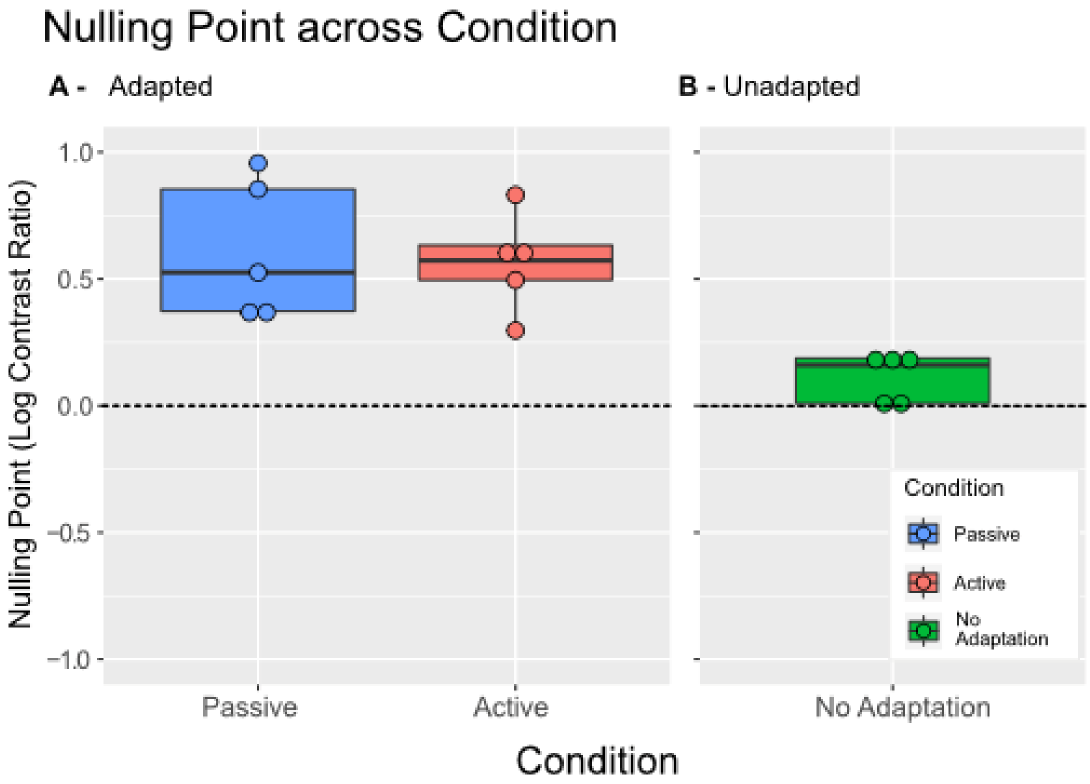

Figure 5.4 – Boxplot Nulling Point Thresholds across Passive, Active Adaptation and No Adaptation Conditions. **5.4a** – Nulling Point thresholds for the Passive and Active adaption conditions. **5.4b** – Nulling point thresholds for the unadapted condition.

Figure 5.5 shows boxplots of JNDs for the active and passive adapting conditions and for the unadapted condition. The results indicate that participants were slightly more accurate at discriminating between the contrast of the superimposed gratings in the self-generated adaptation condition (Mean = 0.44, STD = 0.12), compared to the passive adaptation condition (Mean = 0.48, STD = 0.15). In the unadapted condition, participants were much more accurate at performing the task (Mean = 0.30, STD = 0.13) in comparison to the adapting conditions. To understand whether these differences were significant, a repeated measures ANOVA with Condition (Passive, Active, No Adaptation) as the factor was run on the data and revealed a non-significant effect of condition on JND thresholds (F(8) = 3.742, p = 0.07, n2 = 0.27). This indicates that there was no difference in adapted state discrimination under the adapting conditions and unadapted conditions.

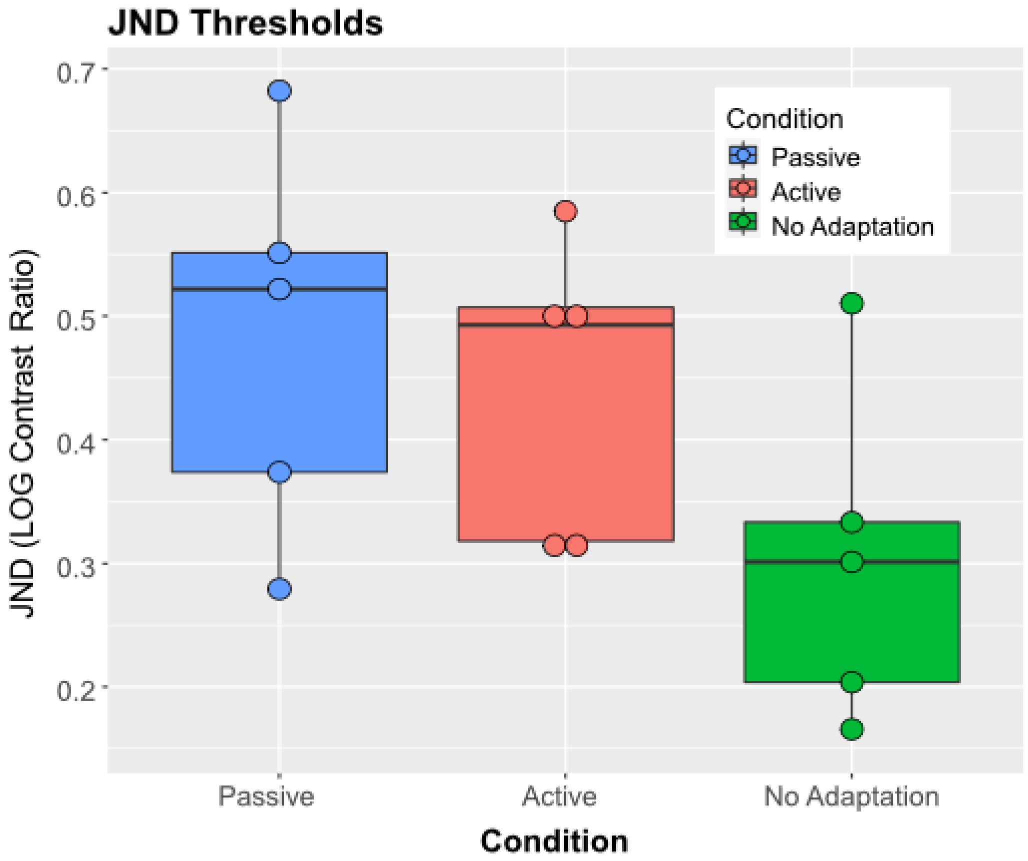

Figure 5.5 – Boxplot JND Thresholds across Passive, Active and No Adaptation Conditions.

5.4 Discussion

This experimental chapter investigated how adaptation to self- or externally generated motion altered perceptual bias and discrimination. This was in part motivated by the absence of aftereffects found in many real-world viewing settings, especially following motion, and also by studies that have shown perceptual sensitivity following self-generated stimuli is more accurate when there is a congruency between the visual and motor inputs (Konkle et al, 2009; Ichikawa & Masakura, 2004; 2006; Ichikawa et al., 2010; Dewey & Carr, 2013). I also wanted to directly compare our results with previous studies that have investigated adaptation bias in combination with vestibular input (Harris et al., 1981; Bai et al., 2018). The experimental results from these studies indicated that vestibular input reduced adaptation to motion.

In this experiment, there were three experimental conditions where participants generated the motion during the adaptation period, were passive to the adaptation as it was externally generated, and finally performed the task without adaptation. The results of this experiment showed statistically significant biases in perceived motion in both the adaptation conditions in comparison to the no adaptation condition. However, there was no statistically significant difference between self-generated or externally generated motion conditions. This indicates that self-generated adaptation does not reduce the effects of motion adaptation.

There are several reasons why this experiment might not have shown statistically significant differences between the self-generated and externally-generated adaptation conditions, and why the results indicate a significant MAE overall in the adapting conditions. To begin, part of the motivation for this experiment was to understand how self-generated adaptation affected the MAE in the absence of vestibular input. Vestibular input has been assessed by both Harris et al., (1981) and Bai et al., (2018) and was either congruent or incongruent with the motion participants adapted too. In these experiments it was determined that vestibular input is a important factor in suppressing the adapted state bias and aftereffects to motion. Perhaps in combination with these studies, the results in this experimental chapter highlight that self-generated motion adaptation is not in itself a key factor to reducing adapted state bias or visual motion processing.

Another possibility is that self-generated control over adapting stimuli is important for altering other aspects of visual processing. For instance, in the studies by Ichikawa & Masakura (2004; 2006) and Ichikawa et al., (2010), they showed that participants awareness and motor control over a moving stimulus on the screen, was sufficient and necessary to cause a reduction to the flash-lag effect. This effect is used to understand temporal and spatial aspects of visual motion perception as the stimulus moves across the screen in a short amount of time. As the effect reduced when under self-generated control, in comparison to when it was externally generated or incongruent with the participants hand movement, this suggested that visual performance was becoming more accurate as a consequence. This finding is also in agreement with other studies that have shown improvements to visual motion processing under self-generated motor control, such as motion speed perception (Dewey & Carr, 2013) and motion detection (Rhode et al., 2014; Lally et al., 2011; Christensen et al., 2011). Therefore, in light of this experiment it seems that self-generated adaptation and the resulting aftereffect to motion remains stable at least in our laboratory environment.

Perhaps another possibility is the distinction between voluntary and restricted self-generated control over a stimulus and the resulting visual processing in light of the other experiments that have shown differences between self-generated and externally generated visual processing, a shared feature to these examples is that the stimuli are under voluntary control (i.e. they choose when to initiate the trial) and they have less restrictions on the movements they can make (i.e. can choose to change the type of movement they can make). Whereas, in this experiment the self-generated control over the stimulus was much more restrictive. For instance, participants had to control the adapting stimulus as soon as it appeared on the screen and were explicitly told to keep the stimulus moving in a certain direction at the maximum speed. It is possible that these restrictions on self-generated adaptation, meant that adapted state bias and visual processing built up as if they were externally generated.

It is unclear whether voluntary self-generated action is important for reducing bias, however this could potentially be an interesting future direction for this research investigating the MAE. For instance, further research could investigate whether allowing greater participant freedom in how they choose to move the stimulus (i.e. no constraint on stimulus direction) would alter adaptation to self-generated motion. Another future direction, which will be the focus of the next chapter of this thesis, would be to explore other aspects of visual perceptual processing under self-generated motor control. As highlighted by other research visual processing is more accurate when under self-generated control (Ichikawa & Masakura, 2004; 2006; Ichikawa et al., 2010; Konkle et al, 2009; Dewey & Carr, 2013; Rhode et al., 2014; Lally et al., 2011; Christensen et al., 2011). According to predictive forward models, the visual-motor interaction maps an upcoming visual perception based on current sense data, efferent motor signals and knowledge of past experiences of the world (Körding & Wolpert, 2004). Thus, in the next chapter I will investigate whether our expectations of visual motion under self-generated control can influence visual processing and whether this performance can be altered.

To summarize and conclude, the results of this experimental chapter highlight that adaptation to motion is not altered when the motion is self-generated motion in comparison to when it is externally generated. This highlights that reductions in adapted state motion aftereffects, that have previously seen in the literature, are a consequence of the matching vestibular inputs that also accompany many instances of how we perceive self-generated motion in our everyday life. The focus of the next chapter is on how self-generated control over a visual stimulus reduces other aspects of visual processing.

Chapter 6: Expectations to self-generated motion signals can be altered through Perceptual Learning

6.1 Introduction

Co-ordinating the brain's visual and motor processing systems is key to interacting successfully with the world, e.g. performing visual guided movements. For instance, if our visual attention is suddenly alerted to something in the environment, this visual processing is vital for informing our motor system as we begin to interact with whatever has caught our attention. However, this interaction can also work the other way. For example, once our motor system interacts with an object, our actions alter the pattern of sensory input.

Numerous theories on this relationship have proposed a tight link between visual and motor systems, suggesting that they share a bidirectional relationship with every executed action (Prinz, 1997; Hommel et al., 2001; Miall and Wolpert, 1996; Wolpert et al., 2003; Earlhagen et a., 2006; Kilner et al., 2007). A shared feature between these theories is that the motor system creates an internal simulation (or an efference copy) of the action and its expected consequences. The purpose of these signals is to anticipate and cancel the sensory effect of motor movement should sensory information change and permit expectations about the actual associated sensory signals. These expectations have been theorised as playing a role in influencing the way we perceive the world.

The interaction between visual and motor processing has been investigated by few researchers, who in a bid to understand the influence on perception, have manipulated the relationship between temporal and spatial congruency of observed and executed action (Christensen et al., 2011; Lally et al., 2011; Miall et al., 1993; Heron et al., 2009; Rohde & Ernst, 2013; Rohde et al., 2014). In one example Christensen et al., (2011) created a novel experimental paradigm that combined real-time full body tracking with online generation of motion stimuli. This made the viewed motion stimuli dependent on the actual executed action of the observer. Over a series of experiments, they systematically varied the temporal (i.e. delaying the viewed motion in time) and spatial congruency (i.e. rotating the viewed motion at different angles) between executed action and observed motion. On each trial, participants were required to detect this motion amongst an array of moving distractors. They revealed that performance was facilitated when the to be detected motion was congruent with the self-produced movement of the arm, but impaired when there were substantial degrees of incongruency.

In a similar experiment, Lally et al., (2011) employed a task where participants used their hand to move a robotic arm in a figure-of-eight movement. This movement controlled a field of dots on a computer screen that either matched or mismatched (moved with a delay and with a slight random trajectory) in real time with the participants movement. The authors also ran a passive condition, where participants did not move their hand, and a condition where participants moved their hand but were not controlling the stimulus. In each condition, there were noise dots and target dots. Depending on the condition, the noise dots either matched or mismatched the participants movements, and on each trial participants detected and estimated the coherent motion of the target dots. The target dots moved in a given direction at a random interval during the 10 second trial. The results showed that visual motion detection was attenuated for motion that was congruent to the actions of the participant. Thus, the noise dots in this case were being attenuated when they were being self-generated. In combination with Christensen et al., (2011), where they found that self-generated motion was facilitated, this suggests that depending on the task the visual motor systems can either attenuate or facilitate the perception of self-generated visual motion. It also suggests that visual motor expectations rely on an efference copy of the self-generated action which is necessary for either perceptual predictability or cancellation.

In parallel to the research investigating visual-motor control and perception, a growing body of work has investigated visual expectations (also referred to as priors) and how they influence our perception (for review: Seriés & Seitz, 2013). One type of visual expectation, known as structural expectations, is characterised as an expectation that generalises across environmental circumstances. These are based on implicit learning of the statistics of the natural environment and reflect long-term learning over a lifetime. The characteristics of these structural expectation are similar to the visual motor expectations that have been observed (Christensen et al., 2011; Lally et al., 2011). For instance, one could argue that the self-generated action and the resulting perceptual predictability or cancellation, is based on long-term learning in the natural environment that self-generated stimuli behave in a certain way (i.e. congruent with the intended path of the self-action). Therefore, it would be valuable to explore these similarities to understand more about visual motor expectations and how they shape perception.

One interesting avenue in the visual expectation literature has shown that new visual expectations can be learnt over a training period (Chalk et al., 2010) and even long-term structural expectations can be recalibrated and influence our perception (Adams et al., 2010, Sotiropoulos et al., 2011). One experiment by Adams et al., (2010) tested participants using the 'light-from-above' prior, an expectation that light in our environment comes from above, which the visual system uses to recognise objects or determine which way is up. Throughout the experiment participants made convex-concave judgements on bump-dimple stimuli at different orientations. In the pre-training session, participants viewed the stimuli as convex, confirming the 'light-from-above' prior in their lab setting. However, over the course of the training phase a new shape was added in the trials that conflicted with this initial prior by shifting the light source by a 30° angle. At the end of training, this prior was remeasured in the same way as the start, and their light direction prior had shifted significantly in the direction of the 30° shift used during training.

In another example, Sotiropoulos et al., (2011) tested participants slow-speed prior, an expectation where moving objects are judged to be slower than they truly are. This expectation is thought to develop over the course of a lifetime where static or slowly moving objects are more common than fast objects (Weiss et al., 2002; Stocker & Simoncelli, 2006). In this experiment, participants were presented with a field of parallel lines that either moved in a normal direction relative to the rotation of the line (perpendicular – 50% of trials) or in an abnormal direction relative to the rotation of the line (oblique – 50% of trials). Participants were also separated into two groups, one group that trained on this task at a slow speed (4°/s) and another at a fast speed (8°/s). Importantly, owing to the nature of the slow-speed prior, participants in both groups initially viewed the direction of motion of the stimulus as more perpendicular than it truly was. However, over the course of training in the fast speed group, they updated their expectations toward the fast speed, and this updated expectation changed the direction bias of the stimulus toward the oblique direction. For the slow speed group, the initial prior remained as the illusion of direction continued to be perceived as more perpendicular.

These experiments highlight how prior expectations about the fundamental statistics of the world can be applied to novel experimental settings that the participant has no experience of. In addition, it also reveals that these expectations can be updated through repetitive training regimes. Using a similar training paradigm, could be particularly interesting to explore with what has been shown for visual-motor expectations. For example, visual motor expectations have been shown to facilitate or attenuate perceptual performance depending the spatial congruency and the type of task. However, an interesting question would be whether these expectations could be altered through training. For instance, it would be interesting to understand whether training in an environment where this expectation is not useful for performance, could be altered at the end of training.

Given the current information, one of the main aims of this chapter is to understand more about the perceptual consequences to self-generated stimuli. Research investigating visual-motor interaction suggests that our visual perception may be facilitated or impaired depending on the type of task. This may be based on an efference copy that the motor system creates that informs the visual system through expectations. Thus, in Experiment 1, I aim to first set-up an experiment that characterises the profile of sensitivity to self-generated motion following varying directions of that motion. Following on from this, I will then investigate whether this mechanism, that may be based on visual-motor expectations, can be altered through training (perceptual learning) in Experiment 2.

6.2 Experiment 1

6.2.1 Rationale

It is theorised that the visual motor system continually builds an internal simulation of the sensory consequences of our self-generated actions (Prinz, 1997; Hommel et al., 2001, Wolpert and Miall, 1996; Wolpert et al., 2003; Earlhagen et a., 2006; Kilner et al., 2007). A few studies have investigated how these predictions are used by the brain to influence our perception (Christensen et al., 2011; Lally et al., 2011; Miall et al., 1993; Heron et al., 2009; Rohde & Ernst, 2013; Rohde et al., 2014). The results from two key experiments that have manipulated the spatial and temporal congruencies between visual motion and self-generated action (Christensen et al., 2011; Lally et al., 2011), suggest that the brain may be able to either cancel or facilitate visual motion perception in a task-dependent manner.

In the first experiment of this chapter, I aim to establish an experimental paradigm that can adequately test this visual-motor interaction to understand more about how the brain makes visual-motor expectations and how they shape visual perception. To do this I use a 2-interval forced choice (2IFC) motion detection paradigm that utilises wide field Random-Dot Kinematograms (RDKs). In two conditions, participants either actively controlled the motion they were detecting with a swipe of their finger in real-time (*Active Condition*) or were passive to motion (*Passive Condition*). In both conditions, the to-be-detected motion could move in any direction, however importantly in the active condition this either matched the participants movement or varied in a direction offset from the participants executed action trajectory.

6.2.2 Methods

6.2.2.1 Participants

Seventeen participants (10 Male, Median = 27, age range, 20-51) gave informed consent and completed the experiment. All had normal or corrected to normal vision and were paid an inconvenience allowance for their time (£10 per hour).

6.2.2.2 Apparatus

Participants sat in front of touch sensitive a Cambridge Research Systems Display++ LCD Monitor approximately 33cm away, with their head secured on a chin rest. The monitor had a 120Hz frame rate (8.3ms inter-frame interval) and a pixel resolution 1920x1080. The experiment was controlled and run using custom software written in PsychoPy (Peirce et al., 2007) running Ubuntu operating system.

6.2.2.3 Stimuli

The stimuli presented were computer generated RDKs that filled the entire computer screen. The RDKs were comprised of 5000 white dots, each measuring 4 pixels in diameter and displayed on a black background.

Throughout the block, the dots were updated every 10 frames (83ms) and had a new set of random locations each time. Between the trial intervals, which were initiated by the swipe of the participants finger (more detail in section 6.2.2.4), all dots moved in random directions. During one of the trial intervals, a subset of the dots was randomly selected to have their positions updated in real time in line with the exact speed as the participants swipe trajectory. The direction of the dots was either in the same direction as the participants swipe trajectory (0°) or travelled in directions either side of this swipe trajectory (± 45°, ± 90°, ± 135°, 180°). The selection of dots was random and from a uniform distribution (each dot had an equal probability of being selected). The proportion of signal dots (that moved with the participants swipe movement) selected was dependent on motion coherence for that trial and these dots were updated on each frame (every 8.3ms). This was the closest to real-time that the dots could move with the participants touch input. The random selection of noise dots was also updated on every frame.

6.2.2.4 Design and Procedure

All participants were given limited practice before beginning the main experiment. The practice consisted of first showing the participants the dots on the screen and allowing the participant to get used to the swiping structure of the experiment. It was explained at the beginning of the experiment that throughout the experiment participants should fixate on the screen in front of them and not follow their finger across the screen. Explicit instructions were also given to participants that they could swipe their finger however they choose, however should avoid simply tapping their finger on the screen without swiping as this would activate the trial without moving any of the dots.

In total the experiment lasted approximately 5 hours and participants completed 8 blocks of the active and passive condition (Figure 6.1a). Both conditions used a 2IFC design. In the active condition, each interval was initiated by the participant swiping their finger across the touch screen monitor in front of them. In one of those intervals, the RDK rotated in a

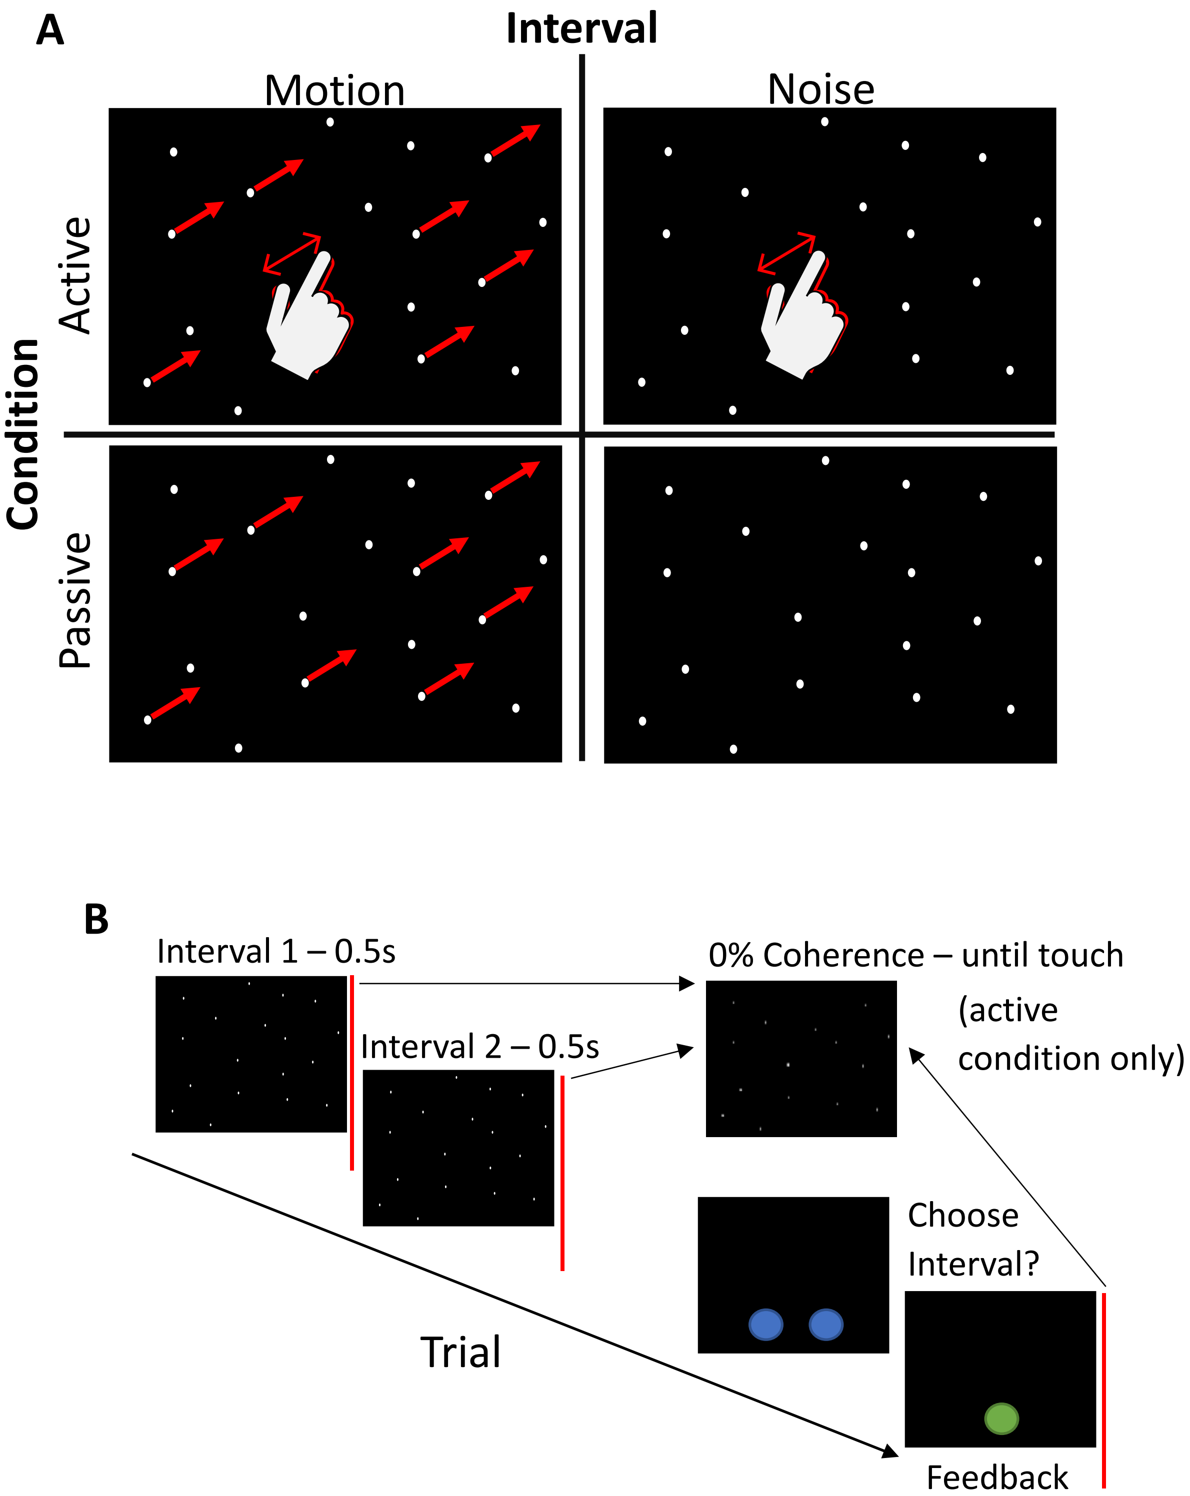

Figure 6.1 – A) Active and Passive Condition Structure. During the Active condition participants controlled the dot motion by swiping across the screen. During the Passive condition, participants were played back their swiping trajectories. The decision on each trial was to identify which interval contained coherent motion. **6.1 B)** Trial Structure. Each interval lasted approximately 0.5s and participants choose which interval contained the motion. In the active condition, there was 0% motion coherence intervals between each trial interval which lasted until they began the trial or after making their response and receiving feedback. Feedback was given before the next trial.

direction at a determined coherence level controlled by the staircase. During the non-motion interval, the dot positions were randomly selected from a uniform distribution and had a flickering appearance. Participants generally completed the experiment over a period of 5 days and each condition was sequentially interleaved after one another.

During the active condition, participants controlled the dot motion by swiping their finger across the screen. During the passive condition, participants were played back their own motion patterns made during the active condition. Blocks lasted approximately 15 minutes. Each block was comprised of four staircases that randomly interleaved between trials. Each staircase controlled the coherence for each of the directions and randomly switched between its inverse (i.e. ± 90° staircase switched between 90° and -90°). The coherence of each staircase started at 40% (5000 dots * 40% = 2000 dots) and was controlled using a 1 up – 1 down staircase until the first mistake was made, then it was controlled by a 1 up – 3 down staircase to achieve an 80% correct discrimination threshold which were derived from a fitted psychometric function using maximum likelihood estimation. Each staircase terminated after 6 reversals.

Each trial was comprised of a two-interval-forced-choice (2IFC) paradigm (Figure 6.1b), participants were instructed to choose which interval contained coherent motion. Participants indicated on which interval they thought the motion occurred in and responded using a virtual button choice on the screen. This appeared after the second interval. Participants then received on screen visual feedback. Each interval lasted 500ms and subsequent response were separated by a 500ms inter-stimulus-interval (ISI). There was also a 500ms inter-trial-interval once the participants gave their response.

6.2.3 Results

Figure 6.2 shows the group motion coherence thresholds for the active (red line) and passive (light blue line) conditions. The average thresholds for the active condition systematically changed depending on thedirection offset. For instance, relative to when the motion trajectory was congruent (0°), the motion detection thresholds were elevated the most when the trajectory direction was at ± 90° and ± 135°. The best detection performance was for the congruent motion trajectory at 0°. This perhaps suggests that motion coherence is facilitated when the motion trajectory matches the participants swiping trajectory and inhibited the further it deviates from the swipe trajectory. In addition, participants also showed a reduced threshold following a full 180° direction, suggesting a sensitivity to the axis of motion.

The passive condition (blue line) by comparison, showed no systematic variation that was consistent with the passive play back of each direction offset. The threshold at 0° is elevated in this condition but this may be due to participant variation to detect motion coherence during passive viewing. Interestingly, the average passive level threshold across direction was 2.04% (Std Error = 0.13) which was slightly higher compared to the group average active direction detection threshold for 0° which was 1.67% (Std Error = 0.15). Therefore, this may suggest

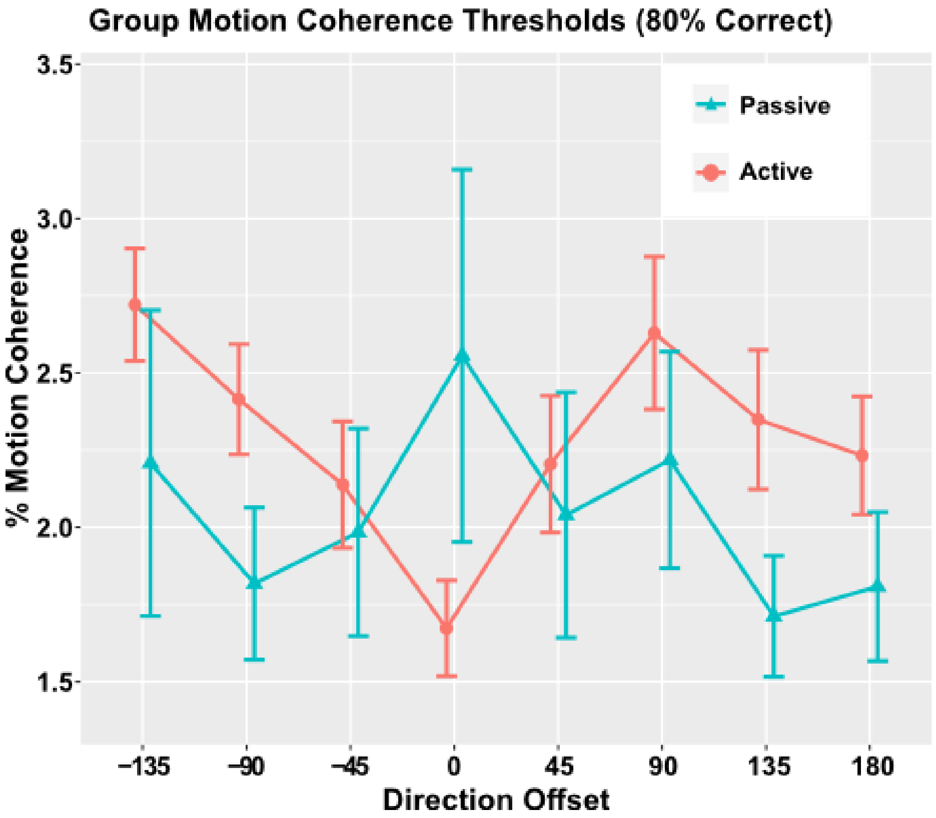

Figure 6.2 – Group Motion Coherence Thresholds per Direction Offset for the Active and Passive Conditions. The Active condition shows a clear tuning function with performance improvements at 0° and deteriorations at ± 90°. The Passive condition shows no real variation across direction. Error bars show 1 standard error.

there is some enhancement for congruent self-generated motion, however an alternative view is that 0° is representative of baseline or ceiling performance.

To understand this pattern of results further, analysis was run on the threshold data for the effects of Condition (Active vs. Passive) and Direction Offset (0°, 45°, 90°, 135°, - 45°, - 90°, - 135°, 180°). A two-way repeated measures ANOVA with Condition and Direction Offset as factors revealed a non-significant effect of Condition (F(1,16) = 0.861, p = 0.367, ηp^2 = 0.011) and a non-significant effect of Direction Offset (F(2.73,43.61) = 1.833, p = 0.160, ηp^2 = 0.018). The analysis did however reveal a significant interaction between Condition and Direction Offset (F(2.74,43.92) = 4.317, p = 0.011, ηp^2 = 0.034).

A one-way ANOVA (DIrection Offset: (0°, 45°, 90°, 135°, - 45°, - 90°, - 135°, 180°) was run on each condition to understand this interaction further. The result of this analysis revealed a significant simple effect of direction for the active condition F(7,128) = 2.55, p = 0.017, ηp^2 = 0.122, and a non-significant simple effect of direction for the passive condition F(7,128) = 0.53, p = .809, ηp^2 = 0.028. Post hoc analysis for the active condition revealed a significant difference between 0° and the 90° (p < 0.01) and the -90° (p <0.001) directions. There was also a significant difference between 0° and the 135° (p < 0.001) and the -135° (p <0.01) directions.

We next investigated participant’s swipe trajectories during the active condition. As a reminder, participants in this condition could swipe their finger in any direction they liked, and as soon as they began their swipe movement the motion trajectory lasted approximately 0.5 seconds. On inspection of the swipe trajectory data, participants could be categorized into two different groups based on the pattern of their swiping data. Figure 6.3 shows the x-y coordinate swipe trajectory data for the first group which contains the majority of participants (N = 14). This group is characterised by their swipe trajectories following a straight line during each interval. Whilst some of this group chose the same swipe trajectory on each interval (e.g. P07, P10, P14), others chose to swipe using the full 360° and varied their swiping direction between intervals (e.g. P02, P12). Another interesting trait to this group of participants is that the majority stuck to the same swipe trajectory throughout the entire active condition and did not change their swiping tactics between blocks. This however was not the case for P04 who chose a different tactic in the latter half of the condition.

The other category of participants is displayed in Figure 6.4 which shows the x-y coordinate swipe trajectory data for a small minority of participants (3 participants). This group is categorised by swipe trajectories that followed less of a straight-line trajectory. One of the participants (P09) in Figure 6.4 stuck to their chosen swipe trajectory throughout the active condition, similar to the majority of participants in Figure 2. The remaining two however, adjusted their swipe trajectories between each block in the Active condition. For instance, P05 alternated between quick circular trajectories with their finger to straight line swipes, and P13 made ‘S’ shaped trajectory patterns as well as the occasional straight-line and ‘V’ shaped swipe between intervals within each block.

Given the nature of the experimental design, having a variety of swipe trajectories between participants might have meant that some participants had a different pattern of results due to the swipe strategy they employed. For instance, this may be the case for P07, P10 and P14 who chose the same straight-line swipe tactics for each interval throughout the entire active condition, in comparison to P02 or P12 who used cardinal directions between intervals, or P09 who chose circular swiping motion. The reason there could be a difference could be due to participants learning to detect the same motion if they chose to swipe in the same way throughout each block.

Figure 6.5 plots the active (blue lines) and passive (red lines) thresholds for all 17 participants. Most participants showed a characteristic tuning function in the active condition. A prominent feature for the majority of participants in the active condition is the lowest detection threshold being centred around 0°, when the motion was congruent. Although in some participants ± 45° was lower or at the same level as the 0° direction. Another prominent feature is the elevation at orthogonal directions (± 90°, ± 135°). This varied throughout participants but generally seemed to centre around either ± 90° or ± 135°. For example, in 9 out of the 17 participants recruited in this experiment, the highest threshold was for motion trajectory directions at ± 90°.

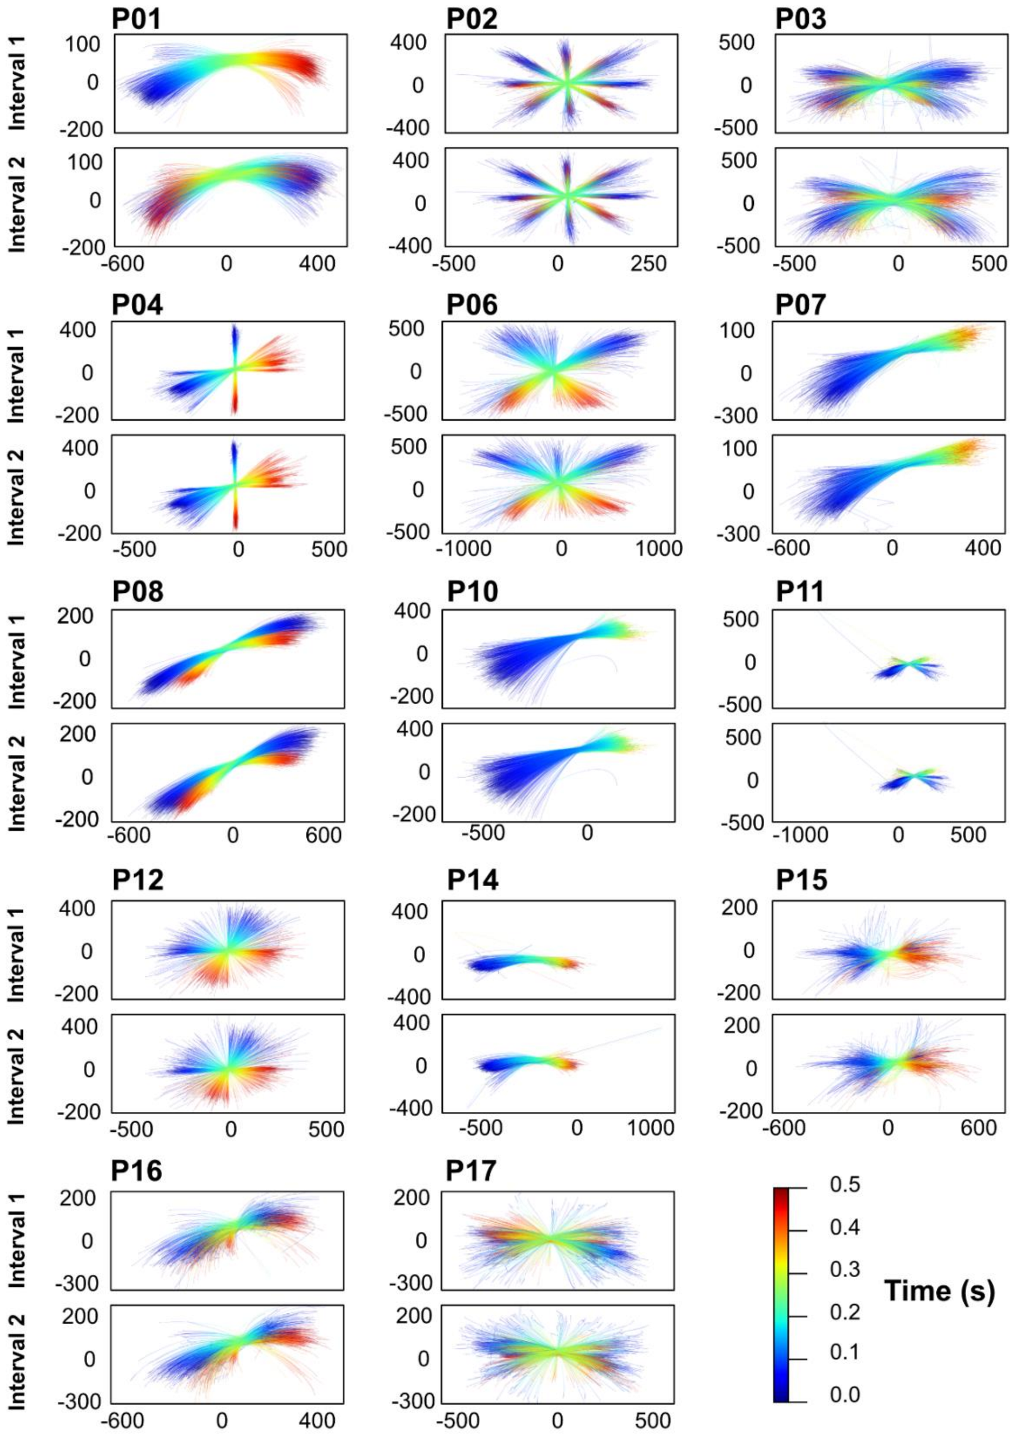

Figure 6.3 - Swipe Trajectory Plots for 14 Participants who on each Interval Swiped in Straight Lines. The X and Y axis are in pixel units of the screen. Generally, participants in this group stuck to the same strategy for each interval and throughout the Active condition.

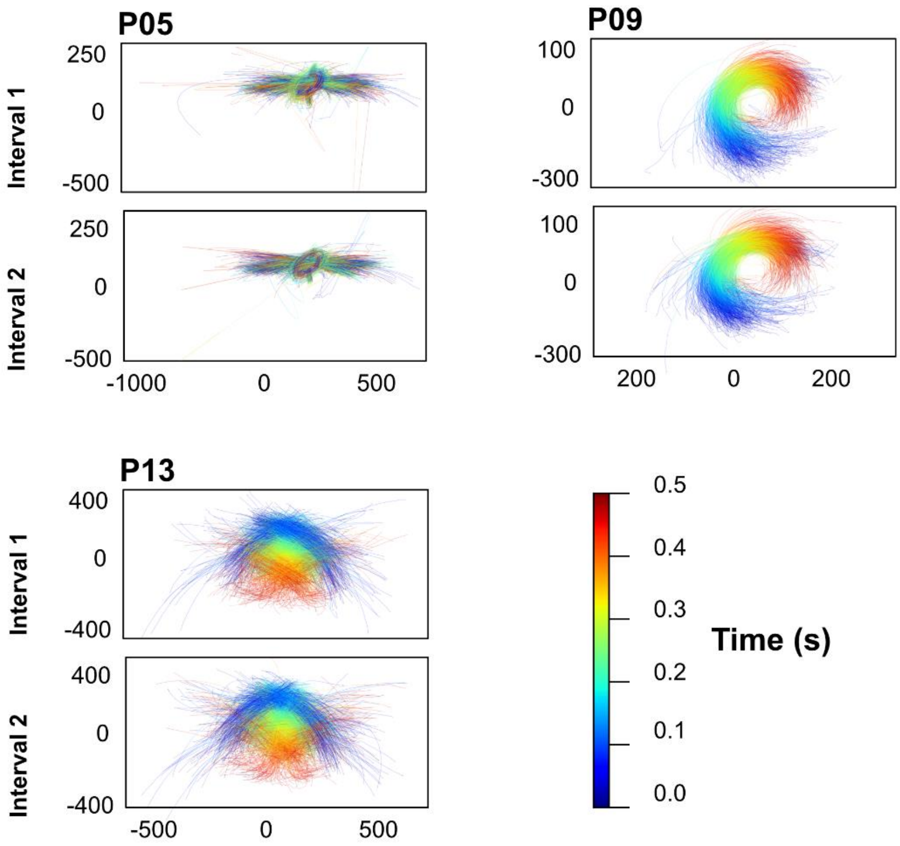

Figure 6.4 - Swipe Trajectory Plots for 3 Participants who on each Interval Swiped in Odd Shaped Trajectories. The X and Y axis are in pixel units of the screen. Generally, participants in this group chose more circular, oval or 'S' and 'V' shaped trajectories throughout this block. In addition, apart from P03 they changed this between block and trials.

Interestingly, comparing the swipe trajectories with the individual threshold data does not point to any advantage to maintaining the same swipe trajectory throughout experiment. For instance, the threshold plots for P07, P10 and P14 who swiped in a straight line throughout the Active condition, are not different to the thresholds seen for P05, P09, P13 who swiped in circular or 'S' or 'V' shaped motion. Therefore, it appears that different swipe strategies do not lead to strikingly different thresholds in either the active or passive conditions.

The passive thresholds in Figure 6.5 for each participant do not show systematic variation per motion trajectory direction and thresholds vary between participants. For some participants, their passive thresholds were lower than the active condition thresholds or were comparable to the active direction performance threshold for 0°. For two participants

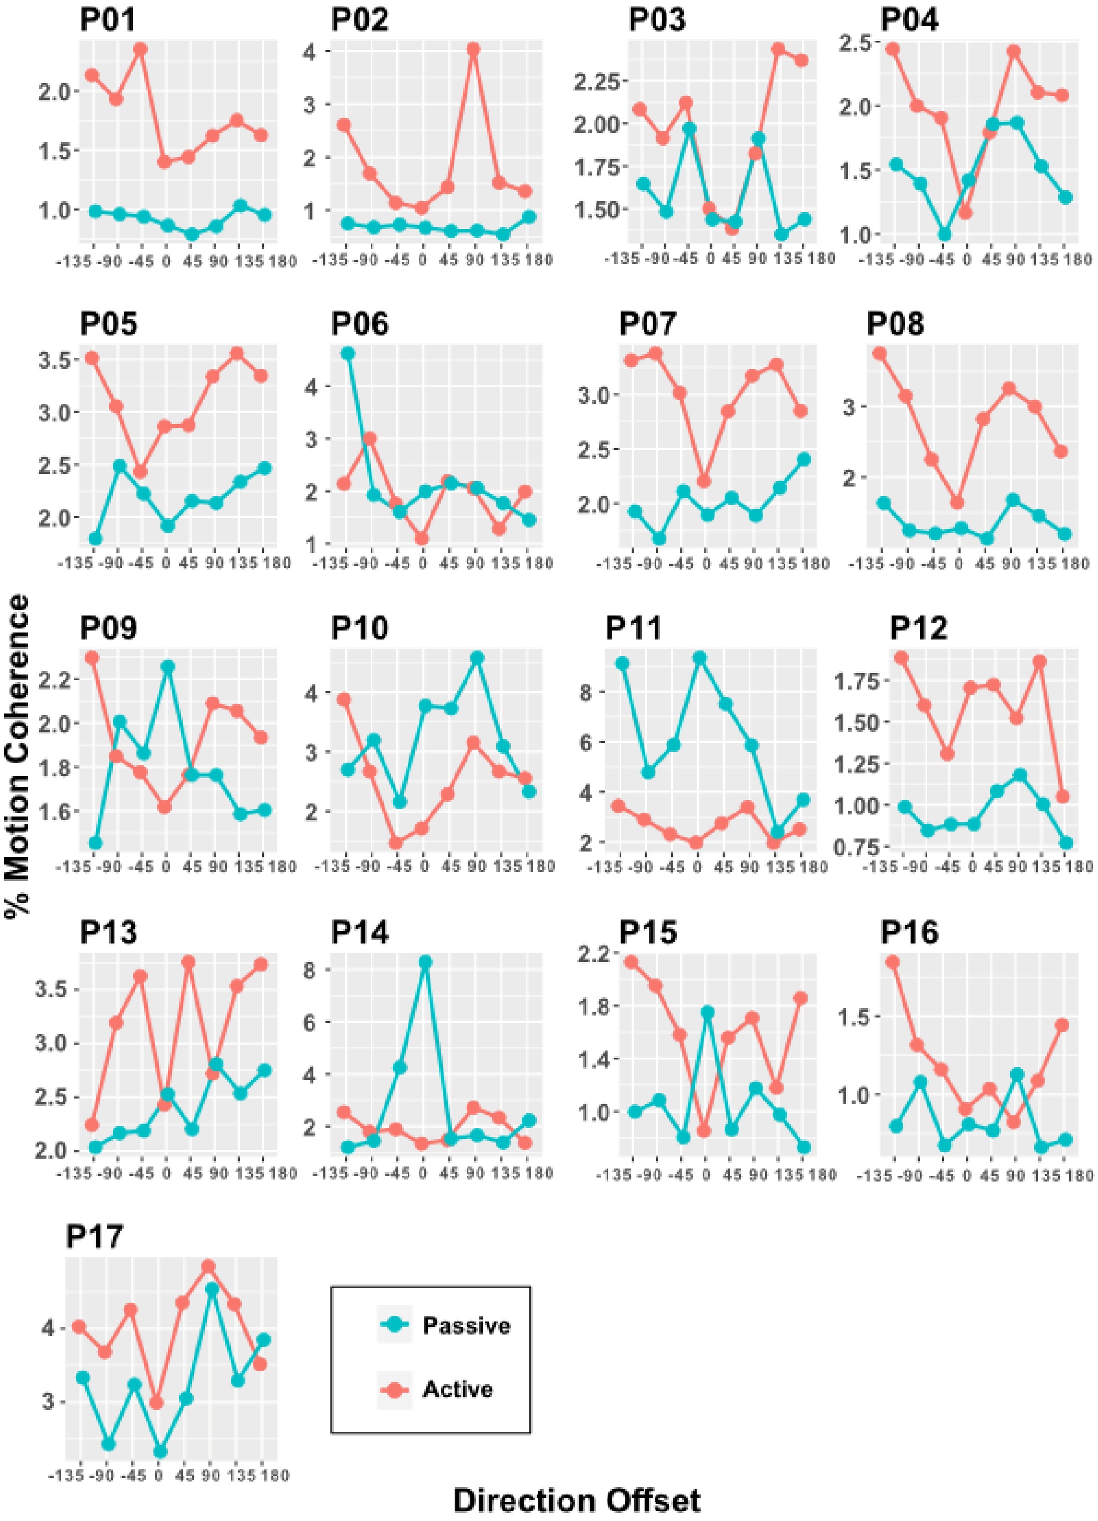

Figure 6.5 – Active and Passive Motion Coherence Thresholds for Participants 1 – 17. Red lines indicate the Active condition and the blue lines represent the Passive condition.

(P11 and P14) however, their passive thresholds for the direction at 0° was elevated to around 8% coherence threshold level. Further inspection of the psychometric function fits revealed that for P11 these fits were in keeping with their other psychometric function fits for different directions. In addition, the fits met our inclusion criteria. Thus, whilst their thresholds were elevated relative to other participants, it seems this was in keeping the rest of their performance in the passive condition. For P14 however, the psychometric function fit for 0° did not meet our criteria as the function fit started at 0.75 and did not reach 1.0 for portion correct. Further investigation revealed that this was due to one staircase for 0° in a block. Removal of this staircase reduced P14 threshold for 0° to 1.4%. This was included in the group motion coherence thresholds and the repeated measurers ANOVA re-run on the data. This revealed the same pattern of non-significant main effects and significant interaction between the factors Condition and Direction Offset.

6.2.4 Discussion

The interaction between our visual and motor processing systems plays an important role in our everyday functioning. Theories postulate that the motor system makes an efference copy with every action which may provide other sensory systems an expectation for the likely consequences of this action, thus shaping our perception of the world (Prinz, 1997; Hommel et al., 2001; Wolpert and Miall, 1996; Wolpert et al., 2003; Earlhagen et a., 2006; Kilner et al., 2007). Research has shown that this expectation may be used flexibly within the visual system depending on the type of task. For instance, some have shown that the perceptual system facilitates perceptual sensitivity to visual motion when it is congruent with our motor behaviour (Christensen et al., 2011) or attenuates self-generated signals to increase sensitivity to unpredictable visual motion (Lally et al., 2011).

In this experiment we used two conditions where participants were either active or passive to the visual motion in each trial. During the active conditions, the viewable motion was either congruent with the participants swipe action or travelled in a direction offset from the participants swipe action. The results showed that relative to the passive group, the active group showed a systematic and statistically significant sensitivity to motion direction. The results in the active condition, suggest response suppression to self-generated visual motion signals when they are incongruent with motor signals in terms of the axis of motion. When the motion was congruent (0°) the motion threshold was comparable to the average passive threshold. Performance was significantly elevated, relative to 0°, when the motion moved at an orthogonal either ± 90° or ± 135° to the swipe trajectory. However, performance did decrease when the direction of motion was 180°, suggesting that the relative axis of finger and stimulus motion is the critical factor.

The results of this experiment suggest that in the presence of angular deviations away from motion that matches the swipe trajectory of the participants (0°), motion coherence detection is suppressed. This may be interpreted as support for the idea that the brain creates a representation of our self-generated actions which is shared with our visual system to shape our perception (Prinz, 1997; Hommel et al., 2001; Wolpert and Miall, 1996; Wolpert et al., 2003; Earlhagen et a., 2006; Kilner et al., 2007). Thus, our results may be due to the predicted perceptual information being anti-cyclic with the incoming sensory information, thus impairing the detection of the visual motion that does not match. Put another way, when we interact with the world, we expect that when we interact with something it will travel in the same direction as we intend it. However, when this moves in another direction our ability to detect it is inhibited because this is different to our expectations.

Research has shown similar results between visual motor congruency, at different temporal delays as well as spatial delays (Christensen et al., 2011). Whist we would expect similar results for temporal incongruencies as we have observed for spatial incongruencies, a direction that future research should take could be to explore whether temporal and spatial incongruencies show a similar tuning function to the axis of motion under our current paradigm. This could be useful in understanding more about this mechanism and the detection thresholds following self-generated visual motion.

Other research on the other hand, has shown the opposite to Christensen and colleagues (2011) that congruent motion under self-generated control can be suppressed (Lally et al., 2011). These two experiments are not in disagreement, rather they suggest that the visual system can be flexible in how it uses the expectations from self-generated motor actions. In light of our experimental findings, it would be interesting for further research to understand whether a similar sensitivity to the axis of motion can be found in a task where suppression of this signal is useful for the task.

One interesting method, which will be explored in the next experiment of this chapter (section 6.3), could be to train participants at the peaks of this tuning function. Literature on visual expectations have shown that long-term structural expectations, that may appear to be hard-wired, can be updated and altered through long-term training (Adams et al., 2010; Sotiropoulos et al., 2011). If this mechanism is flexible in how it chooses to use the self-generated signals (Christensen et al., 2011; Lally et al., 2011), then could training conceivably alter how these self-generated motion signals are used to detect motion? For instance, we might expect that training at ± 90° (the peak) could reverse this suppression and perhaps even reverse the lack of suppression at 0°.

The results outlined in this experiment also raise other interesting questions that future research should explore. For instance, an alternative view for our experimental findings is that the tuning function for self-generated motion can be explained instead by the motion-streak mechanism. Temporal integration in the visual systems causes the motion from fast-moving objects to streak – these are often perceived as static, orientated traces (for review see: Burr, 2000). Psychophysics experiments investigating motion streaking has shown that estimates of velocity are affected and often disappear for fast motions, whereas the direction becomes more reliable (Geisler, 1999; Burr & Ross, 2002 Francis & Kim, 2001). In addition, this can improve motion thresholds (Edwards & Crane, 2007).

In a recent experiment investigating motion streaking, the authors showed that adaptation to slow speeds (1.63 deg/s) adapted orthogonal (90°) motion directions, whereas fast motion (13.02 deg/s) adapted parallel motion orientations (0° - Apthrop, Schwarzkopf, Kaul, Bahrami, Alais & Rees 2013). Respectively these speeds were well above and below the speed of dot motion critical to the generation of motion streaks. The results revealed that slow and fast motion are processed by distinct neural mechanisms. In relation to the tuning function we present in this experiment, perhaps the distribution for motion coherence thresholds is the mechanism for fast streaking motion. An interesting test of this hypothesis could be to adapt participants to vertical or orthogonal directions of motion at fast or slow speeds before they perform their swipe action. This however would come at the cost to constraining the participants swipe movements between two points on the screen.

It should be noted that, based on the swipe trajectory data we collected, the different swiping trajectories that participants employ, did not appear to alter the pattern of active thresholds. Therefore, putting constraints on the type of self-generated movements, at least in this paradigm, should not alter the results I have found. However, it would be useful for further research to extract a meaningful metric regarding the self-generated swipe to understand how this applies to performance. In this experiment I tried to isolate the direction of the swipe for each participant. However, I could not provide a meaningful enough metric for the direction of the swipe that could be applied to all of our participants. This was in part due to the variety of different swipe trajectories employed. Thus, a further direction that research should address is understanding whether any one swiping strategy or swipe direction was optimal for any one direction or type of performance in the active conditions.

Another interesting question for studies could be how motion detection is influenced by temporal mismatches when that delay precedes the actual movement. This could be possible if the movements were pre-recorded and participants were making more consistent swipe trajectories. If the visual-motor system has and utilises expectations regarding self-generated motion, would we expect a similar pattern in the tuning function if these visual signals preceded the participants swiping movements. In addition, future research could also explore whether response inhibition to self-generated motion when the motion signal is incongruent, also biases perception to the direction of motion. This could be possible by having the participant estimate the direction of motion after each trial and comparing the angular error (in degrees) for each direction.

6.3 Experiment 2

6.3.1 Rationale

In Experiment 1, I created a paradigm that characterised the profile of sensitivity to self-generated motion following varying direction offsets of that motion. The results showed that detection thresholds were sensitive to the axis of motion as they were elevated when the self-generated visual motion moved at orthogonal directions off the participants motion trajectory, up until 180° where thresholds reduced again. These results could be due to a mismatch between the expectations the brain was making about the sensory visual consequences for each participant's motor trajectory swipe. However, owing to the sensitivity to the axis of motion, an alternative theory for our results could be that detection thresholds were facilitated by a motion-streak system.

This next experiment is designed to understand more about this mechanism. Past research into visual sensitivity following self-generated motion has shown that this mechanism is flexible and can either facilitate or suppress self-generated motion (Christensen et al., 2011; Lally et al., 2011). In addition, studies have shown that visual expectations play a major role in the way we perceive the world and are built up overtime based on past and present input of the world. However, some research has shown that visual expectations are not necessarily fixed and can be altered through training in a paradigm (Adams et al., 2010, Sotiropoulos et al., 2011). In this experiment, I ask whether the threshold modulation we found in Experiment 1 in the active condition can be altered through training at the peaks of this tuning function (± 90°). If expectations can be updated through training, then we might anticipate a reversal in the way these visual motor expectations shape our perceptions. For instance, what was a detriment to motion detection performance pre-training could be a benefit to performance post training. Similarly, we might also expect that motion detection performance at 0° post-training may be elevated following training.

6.3.2 Methods

6.3.2.1 Participants

Thirteen new participants (6 Male, Median = 26, age range, 19-31) gave informed consent to participate in the experiment. Due to the Coronavirus Pandemic, 4 participants did not complete the full experiment. As a result, only 9 participants (2 Male, Median = 23, age range, 21-27) completed the experiment. All had normal or corrected to normal vision and were paid an inconvenience allowance for their time (£10 per hour).

6.3.2.2 Apparatus

The apparatus was the same as Experiment 1, further details can be found in section 6.2.2.2.

6.3.2.3 Stimuli

The stimuli were the same as Experiment 1, further details can be found in section 6.2.2.3.

6.3.2.4 Design and Procedure

The design and procedure were similar to Experiment 1 except in this experiment there was no passive condition. Details on the 2IFC procedure and trial structure can be found in 6.2.2.4.

Participants underwent 9 days of training and completed the experiment as close to consecutive days as possible. Training consisted of a pre and post-training session and 7 days of training in-between. Throughout the experiment, participants were active in generating the motion. In the pre- and post-training sessions, each block contained 5 separate staircases which controlled the direction of the motion intervals relative to the participant’s swipe. The directions were either 0°, ±45°, ±90°, ±135° or 180° and were randomly interleaved between trials. In the training sessions, each block contained one staircase that controlled directions offset at ±90°. Participants were unaware of this during training. This was the angle that caused the greatest threshold increase in most participants in Experiment 1. During each session participants completed 6 blocks.

6.3.3 Results

Can training in an environment where visual motion travels at the orthogonal to each swipe action alter visual sensory predictions or detection performance for this motion? To examine this question Figure 6.6 shows the averaged JND thresholds across all participants who completed the experiment (n=9). The JND values for each day were obtained by fitting a psychometric function, using a maximum likelihood estimation procedure, to the raw data for each participant. The JND threshold was calculated as the stimulus value yielding 75% correct performance. Both 90° (grey line) and -90° (solid black line) showed a gradual decline in thresholds over the course of the experiment, suggesting that performance for detecting self-generated motion at the orthogonal was improving. Overall, between pre- and post-thresholds, 90° decreased by 41% and -90° decreased by 39%. The shape of the learning function overtime is characteristic of perceptual learning and it is clear that participants motion detection performance was getting better with training for both ±90° as they followed the same trajectory.

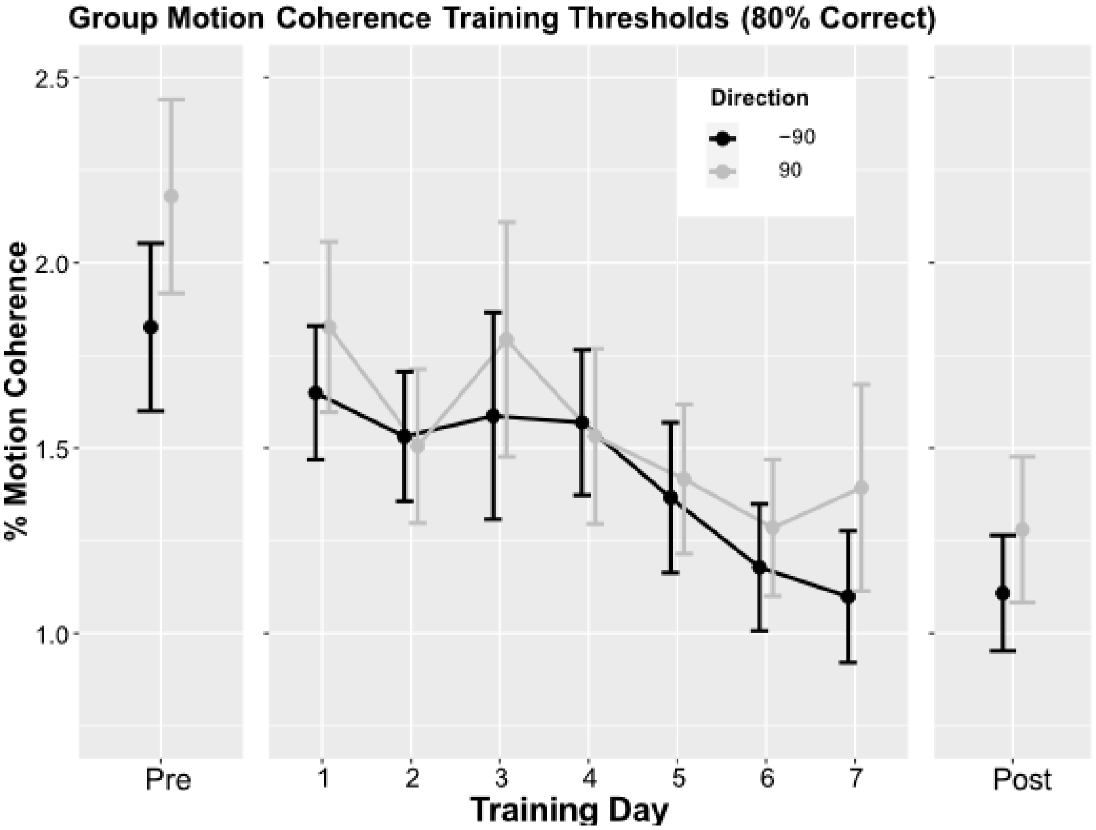

Figure 6.6 – Group Motion Coherence Performance for ±90° across Training Day. Error bars represent 1 standard error. N = 9.

Whilst it is clear they were getting better at performing the task at ±90°, can training to detect motion at these orthogonal directions alter anything about the pre-training tuning function when it is remeasured? Figure 6.7a shows the group motion coherence thresholds pre- and post-training thresholds for those participants who completed the experiment (n = 9). Figure 6.7b shows the pre-training motion coherence performance with all of the participants who completed the pre-training session (n = 14). Firstly, as expected the pre thresholds in both Figure 6.7a and 6.7b are characteristic of the tuning function that was observed in the active condition of Experiment 1 (see Figure 6.2). The lowest threshold for detection performance was for motion that was congruent with the swiping trajectory of the participant. However, motion coherence detection got progressively worse the further away this motion gets from the true trajectory, which peaked around ±90° and ±135° and started to reduce at 180° showing sensitivity to the axis of motion. The only difference between Figure 6.7a and 6.7b is pre-training thresholds are much more even in terms of motion coherence threshold on both sides in the group with more participants (Figure 6.7b).

Inspection of the post-training tuning function blue line – (Figure 6.7a) reveals a non-uniform reduction in thresholds. The post training tuning function showed no systematic variation and instead performance was relatively consistent. To understand this further Figure 6.8 shows the percentage change between the pre- and post-training thresholds for each direction. The greatest percent change was at the trained direction (±90° - highlighted in red). However, observing the percent change for other directions, this change was not unique to ±90°, instead training caused performance improvements at the other directions that deviated (±45°, ±135° and 180°). Only a small performance improvement was recorded at 0° (5%).

Figure 6.9 shows a scatter plot for pre- and post-training thresholds for each directionn offset. For 0° the majority of participants thresholds did not change as a result of training. However, for all the other directions the scatter plots reveal improvements in post-training motion coherence thresholds. This is shown by the majority of dots located in the grey box. In combination with Figures 6.7 and 6.8, this suggests that training at ±90° improved motion detection performance by minimised the elevation of thresholds to motion trajectories that

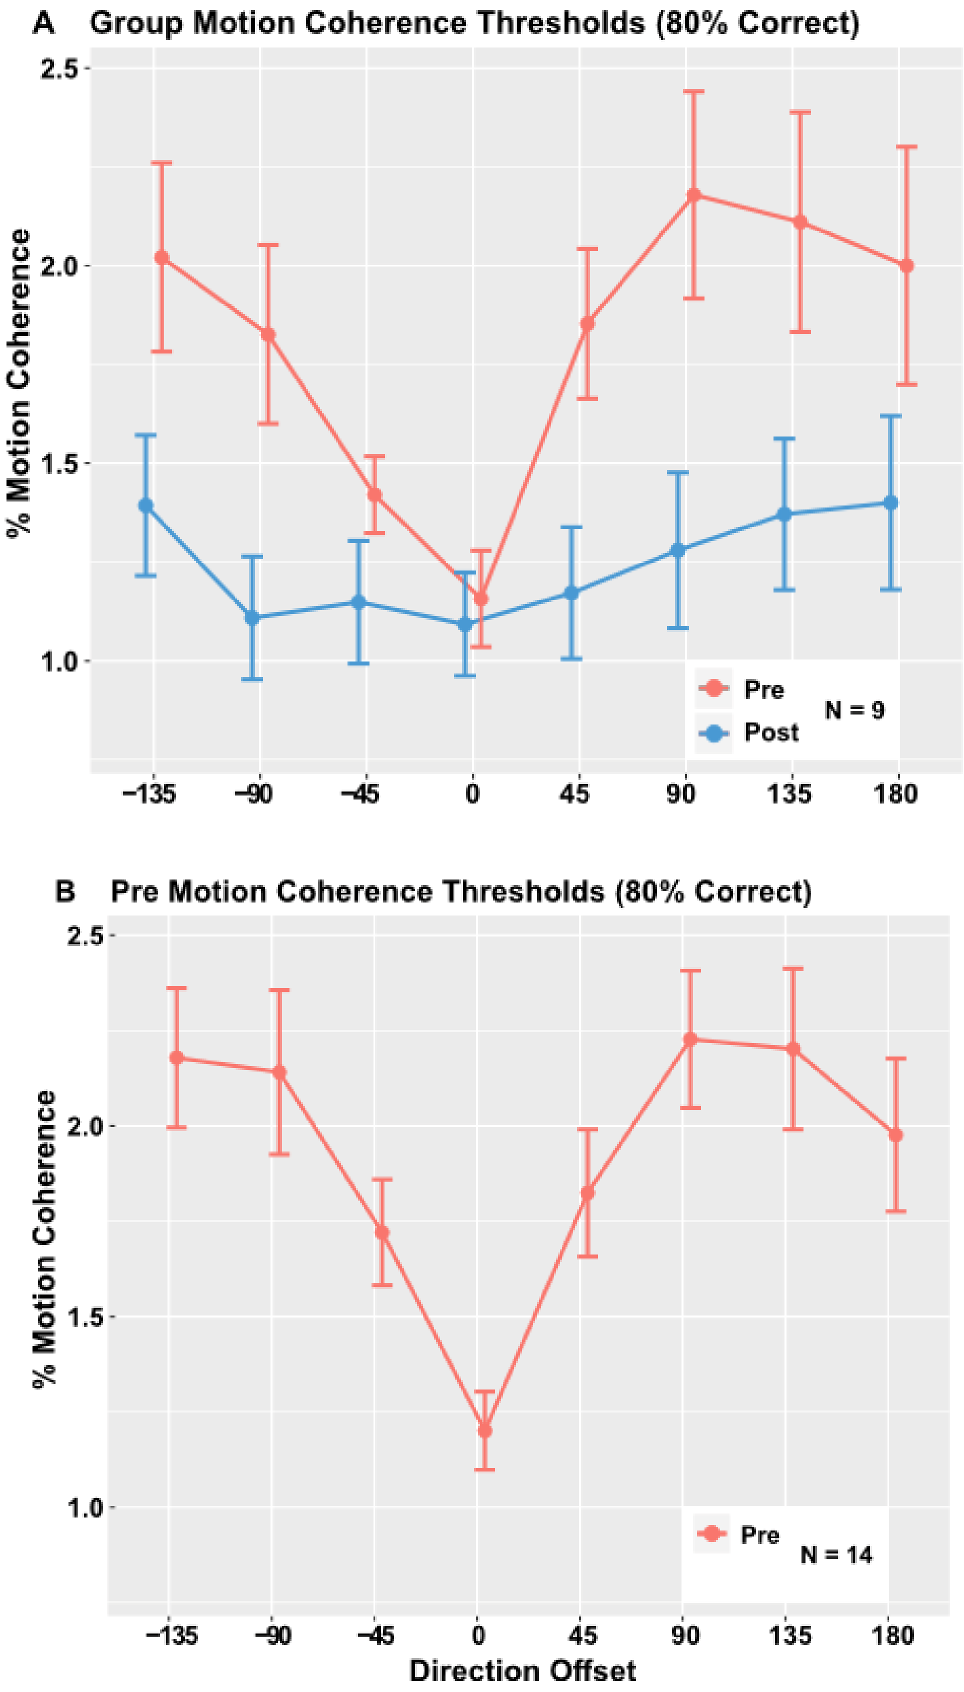

Figure 6.7 – A) Pre and Post Training Motion Coherence Thresholds. Blue lines indicate the Pre condition and the red line represents the Post training condition. **B)** Pre Training Motion Coherence Thresholds for participants who completed Pre Training. Error bars represent 1 standard error.

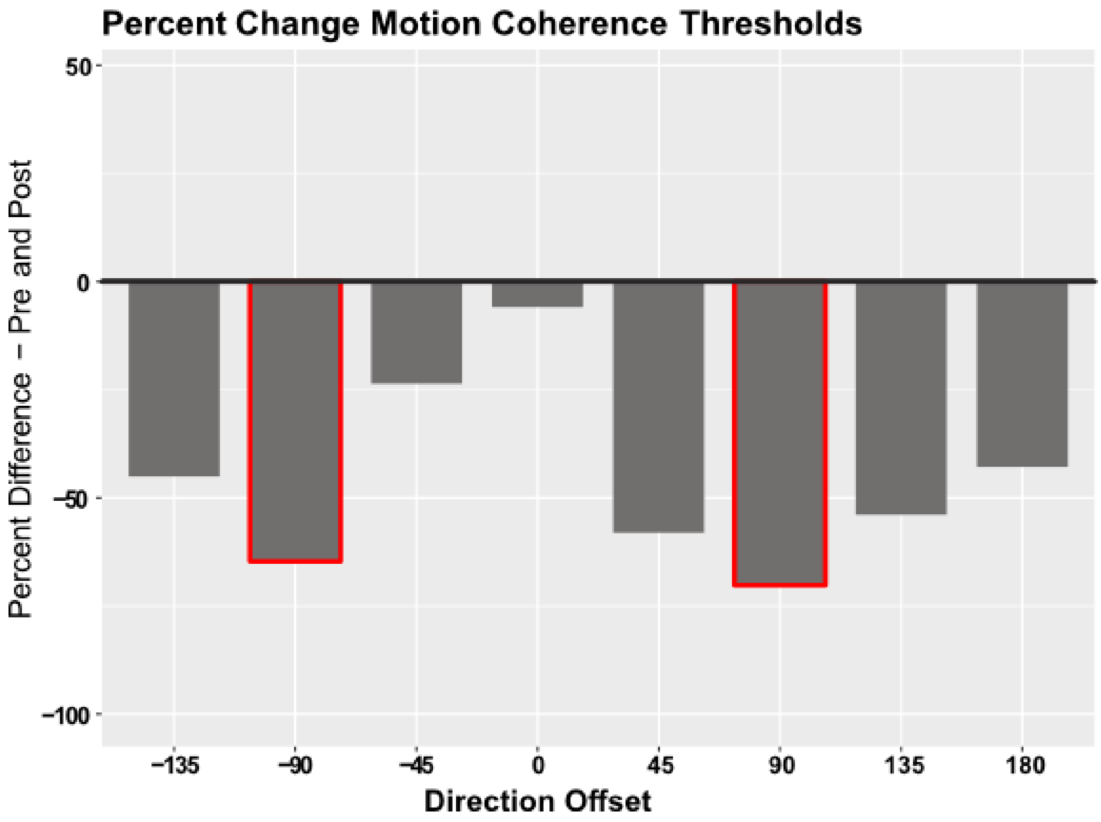

Figure 6.8 – Percent Change in Motion Coherence Thresholds Pre and Post Training per Direction Offset. The smallest change is recorded at 0°, the biggest changes were recorded at the trained direction ±90° (grey solid bars with red outline).

deviated away from 0°. This is suggestive of an overall effect of training that improved coherence thresholds across all rotations, albeit to different degrees.

To understand this pattern of results further an analysis was run on the threshold data for the effects of Training (Pre vs. Post) and Direction Offset (0°, 45°, 90°, 135°, - 45°, - 90°, - 135°, 180°). A two-way repeated measures ANOVA with Training and Direction Offset as factors revealed a significant main effect of Training (F(1, 8) = 27.01, $p < 0.001$, $\eta p^2 = 0.202$) which indicates that there was significant reduction in thresholds post training and a significant main effect of Direction Offset (F(2.96,23.72) = 6.50, $p < 0.001$, $\eta p^2 = 0.125$). There was also a significant interaction between Training and Direction Offset (F(7,56) = 3.56, $p < 0.001$, $\eta p^2 = 0.047$).

To examine this interaction further a one-way ANOVA (Direction Offset: 0°, 45°, 90°, 135°, - 45°, - 90°, - 135°, 180°) was run on the pre- and post-training thresholds. The analysis for the pre-training session revealed a significant simple effect of Direction Offset (F(7,64) = 2.504, p = 0.02, ηp^2 = 0.215) and the analysis for the post-session revealed a non-significant simple effect of direction (F(7,64) = 0.55, p = .793, ηp^2 = 0.057). Post hoc analysis on the pre-training session thresholds revealed a significant difference between 0° and the 135° (p = 0.008) and the -135° (p = 0.003) direction. There was also a significant difference between 0° and the 90° (p = 0.002) and the -90° (p = 0.03) directions. Significant differences were also recorded between 0° and 180° (p = 0.01), 0° and 45° (p = 0.03) and 90° and 45° (p = 0.02) directions.

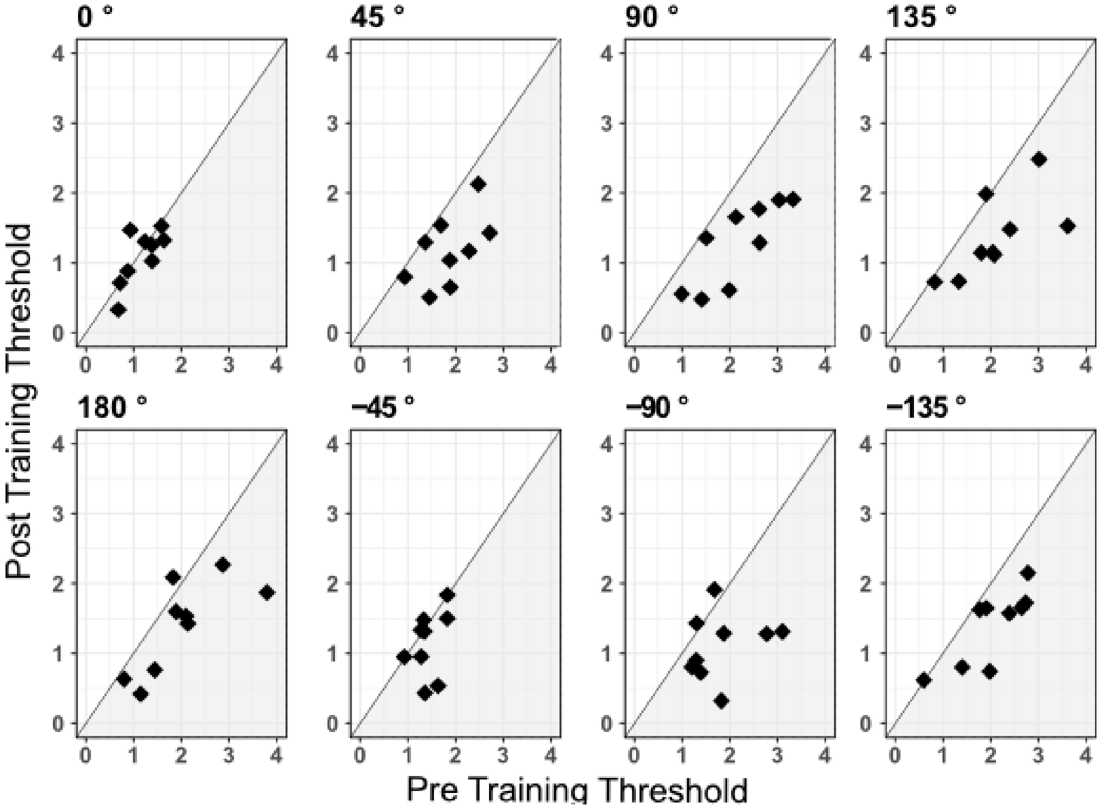

Figure 6.9 – Scatter Plots for Pre- and Post-Training thresholds..

Like participants in Experiment 1, participants in Experiment 2 could swipe their finger in any direction they liked within and between trials. On inspection of the swipe trajectory data, there were no participants that chose to swipe their finger in a circular, ‘S’ or ‘V’ shaped motion. Instead, the majority chose a relatively straight line or semi arc like trajectory on each interval (all swipe trajectories displayed in Appendix). Interestingly, one noticeable difference between participants was in swiping strategy between training days. For instance, the majority of participants maintained the same swipe trajectory over the course of training and there were little noticeable differences between training days. Figure 6.10 shows the x-y coordinate swipe trajectory data for one participant who is an example of a participant who chose a consistent swipe trajectory over the course of training. For this

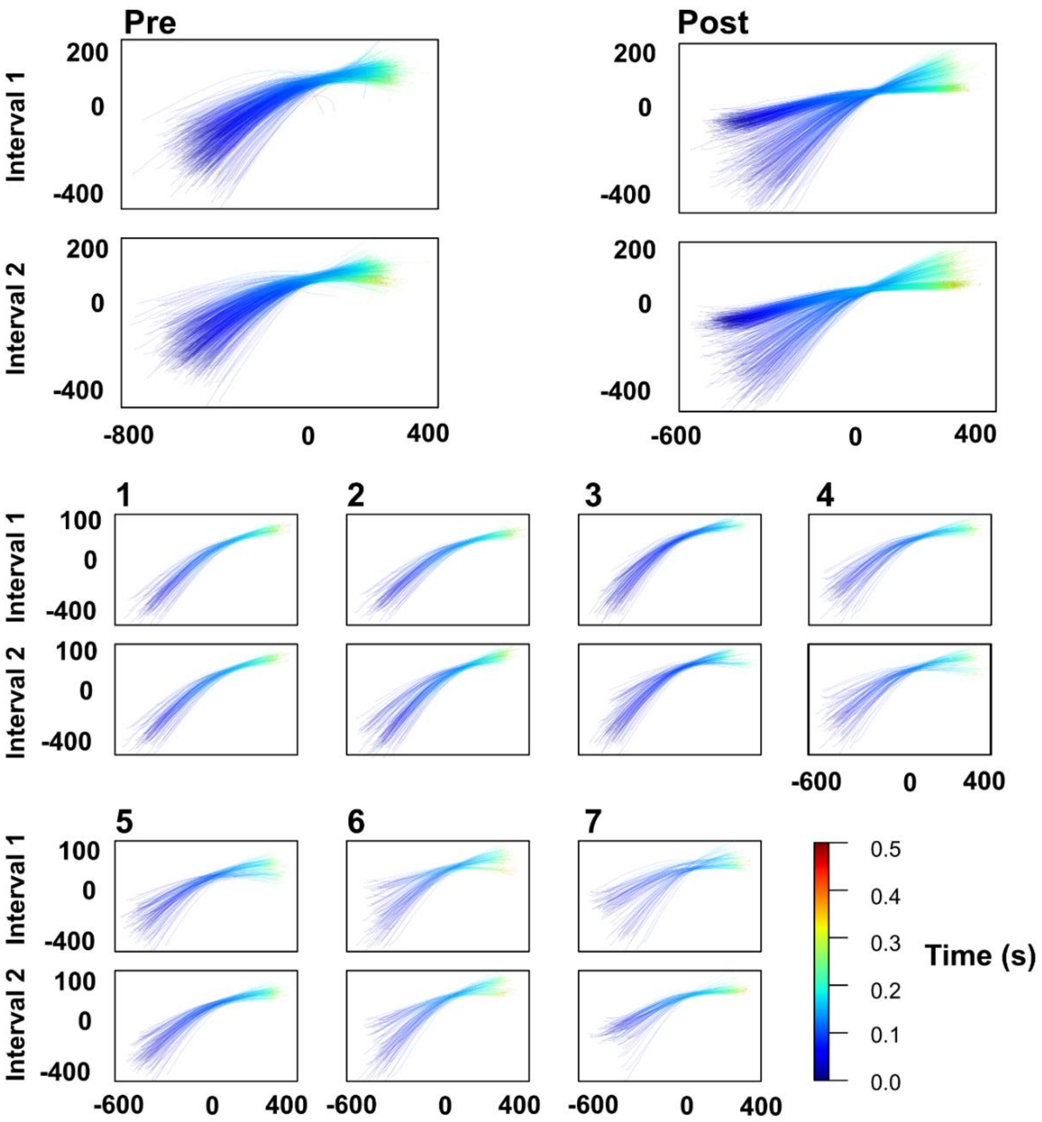

Figure 6.10 – Swipe Trajectory Plots throughout Training for Participant 4. The X and Y axis are in pixel units of the screen. Generally, this participant maintained the same swipe trajectory throughout the entire experiment and did not change their trajectory between session, trial or interval.

participant, there were only slight differences between training day, such as the angle and curvature of their swipe, however generally they maintained a similar trajectory through to the post-training session. Another interesting feature to this participant concerns the duration of their swipe which ends 350ms into the interval.

The remaining participants in this group, changed their swiping tactics over the course of training. Figure 6.11 shows the x-y coordinate swipe trajectory data for one participant who is an example of a participant where there is a noticeable change in swipe trajectories over the course of training. For instance, in their pre-training session they alternated their swiping trajectories between trials and intervals along two different axes. However, over the course of training changed their swipe trajectory in favour of just one of those directions, which they stuck with until the end of the experiment. Comparison between their pre- and post-training swipe trajectories reveals the extent of their change in tactics.

Finally, Figure 6.12 shows the pre- (red line) and post-training (blue line) thresholds for participants 1-13 in separate plots. It is important to note that participants 10-13 did not complete the post-training session. In comparison to their swipe plots (Figure 6.10 and 6.11), the pre- and post-training thresholds for participants 4 and 9 are not strikingly different. The only slight difference is that perhaps participant 9 maintained a bit more of the tuning function shape post-training. This lack of difference may indicate that swiping tactics throughout training were of no advantage to motion coherence performance. Overall, the majority of participants showed a characteristic tuning function pre-training, and the effects of training appear to flatten or reduce this tuning function post-training.

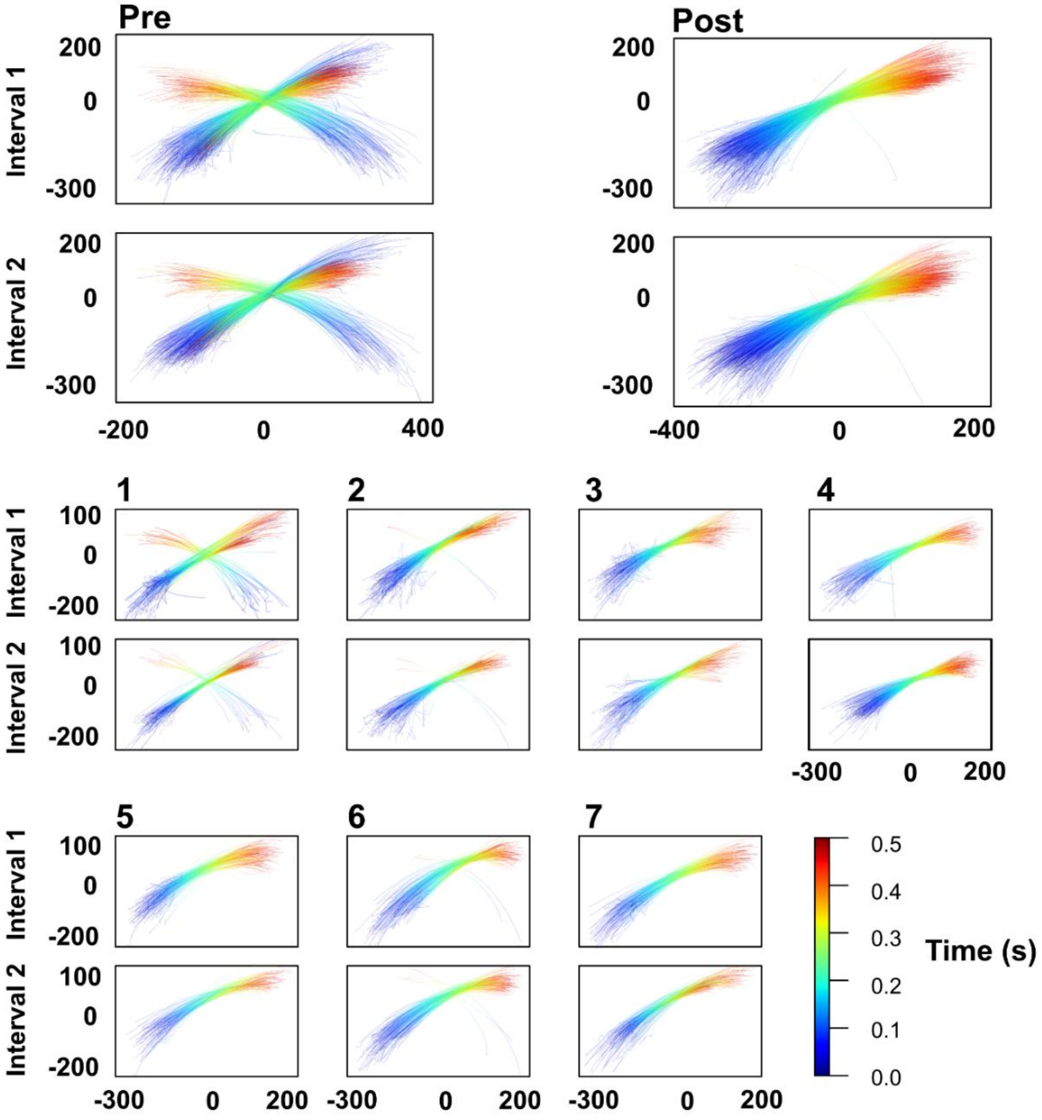

Figure 6.11 – Swipe Trajectory Plots throughout Training for Participant 9. The X and Y axis are in pixel units of the screen. This participant changed their swipe trajectory over the course of training where on day 3 they did not make any swipes along one of the directions they had used in the Pre training session.

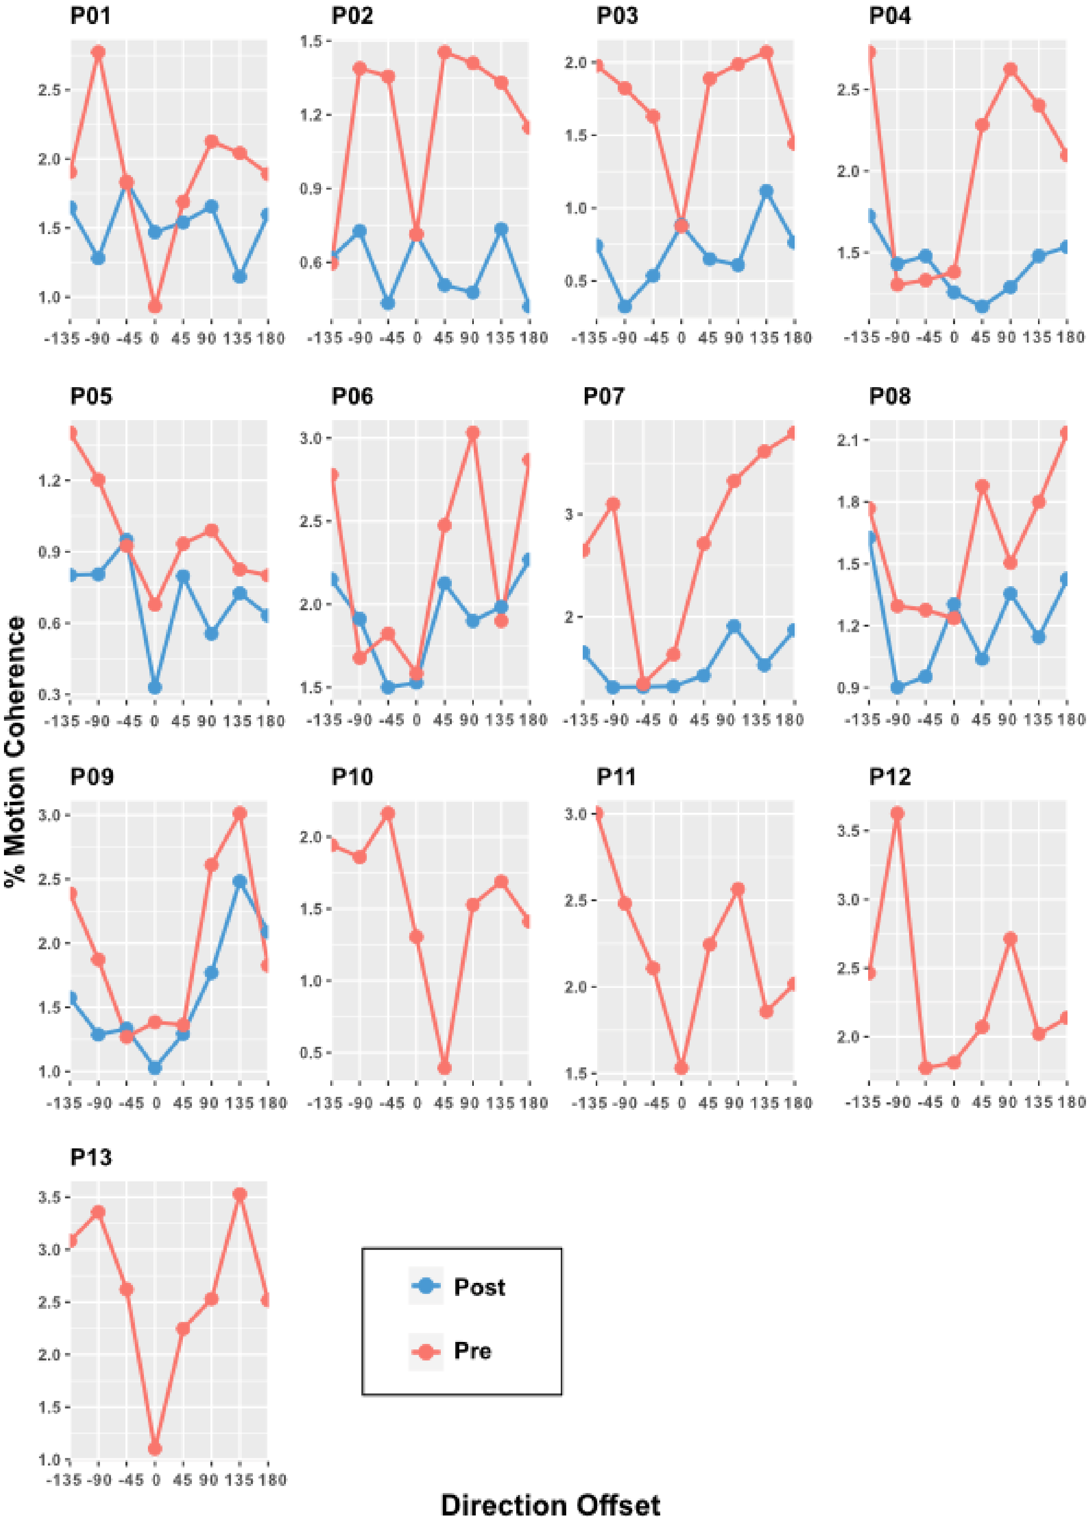

Figure 6.12 – Pre- and Post-Training Motion Coherence Thresholds for Participants 1 – 13. Participants 10-13 did not complete the experiment, so data differs for them. Red lines indicate the pre-training session and the blue lines represent the post-training session.

6.3.4 Discussion

In this experiment, I sought to understand more about the mechanism that was driving my findings in Experiment 1. One possible explanation for my findings was visual-motor expectations, which has been shown to facilitate or suppress visual motion that was either congruent or incongruent to the motor action. A similar finding has been reported previously (Christensen et al., 2011) and this mechanism has been shown to be flexible, as researchers have also shown that it can attenuate the self-generated signal (Lally et al., 2011). Based on this, as well as evidence that structural visual expectations can be altered through training (Adams et al., 2010; Sotiropoulos et al., 2011), we wondered whether long-term training at the peak of this tuning function, would alter the tuning function after training.

The results showed that average pre-training motion detection performance for each direction, was highly similar to the active tuning function we showed in Experiment 1. This was characterised by impaired ability to detect visual motion when the dots moved in a direction ± 90° and ± 135° relative to finger movement, but not when the motion axis was congruent (0°) or 180°. Over the course of training where participants repetitively trained at ±90° for seven days and motion coherence thresholds for this direction improved across day. When the motion detection performance for each direction was remeasured in the post-training session, the tuning function was significantly altered. The average post-training thresholds showed that the greatest percentage change in thresholds was at the trained orientations (± 90°). However, there were also large reduction for the other directions (± 45°, ± 90°, ± 135° and 180°) that were elevated, relative to 0°, pre-training. This is suggestive of an overall PL effect for the detection of self-generated motion.

Before discussing the interpretation of these findings, see footnote for detail on how the Covid-19 pandemic impacted upon the data collection for this experiment which should be taken into consideration in the interpretation of these findings.[1] That said, as already

---

[1] One caveat to this experiment is that we did not reach our intended number of participants due to the Covid-19 pandemic and therefore our interpretations should be treated in light of this. However, it is encouraging that the new group of participants tested

explored in Experiment 1, the brain may have expectations about the way objects in our environment should behave when our motor system interacts with them, i.e. move in the direction that we have intended them to travel in. Thus, one interpretation of the results in Experiment 2 is that these visual-motor expectations are not always fixed and can be changed following training in a certain environment. This agrees in some way with Adams et al., (2010) and Sotiropoulos et al., (2011) who both showed that long-standing visual expectations are not only applied to novel experimental contexts but are constantly updated through training, which modifies perception within that specific trained context.

However, an alternative view is that through training, motion detection performance reached ceiling level or a similar level to 0° which improved performance more broadly. For instance, training improvement was not specific at the trained direction (± 90°), rather detection performance was improved across all directions post-training. Perhaps participants got so used to repetitively swiping their finger on each interval, and this action became associated with simply activating each interval, instead of an action that could be used to detect motion.

There are two interesting avenues that could be used to understand this issue further. Firstly, a new group of participants could be trained at 0° to understand whether this would alter anything about post-training thresholds. If the same pattern of results is found post-training this could be suggestive of a general improvement to motion detection performance more generally as a consequence of training. Alternatively, if training at 0° maintains the same tuning function post-training this may be suggestive of something interesting about training at ± 90° in this experiment. A second interesting avenue that could be explored is the participants swipe trajectories and the influence this plays in motion detection overtime. As participants in Experiment 2 could swipe their finger in any direction or pattern, this led to us not being able to find a useful direction metric that could

---

in Experiment 2 showed a strikingly similar tuning function as those tested in the active condition of Experiment 1, which would suggest that the Post tuning function we collected on those who did manage to complete 7 days of training, may not be too dissimilar to what could have been collected with our desired number of participants.

be applied to all participants. However, future investigations could control participant swipe trajectories more by controlling how the participants swipe and which direction to swipe for each interval (i.e. between two dots on the screen). It could be interesting to compare detection performance of participants who swiped in a consistent direction to say participants who were told to change the direction of their swipe in each block. Using this with a direction metric would be useful to determine how that relates to improvements to a particular direction of motion.

6.4 General Discussion

This chapter has explored how the visual system responds to different directions of self-generated motion using a touch screen paradigm where participants used their finger to control the motion in real-time. Experiment 1 showed that there was a clear tuning function for visual motion under self-generated control. The results showed that motion coherence was sensitive to the axis of visual motion as performance was most sensitive when it was congruent to the swipe trajectory the participant made with their finger, inhibited for orthogonal directions and reduced again for 180°. Further research is needed to explore whether this result is related to visual motor expectations or the motion streaking mechanism. Experiment 2 of this chapter next explored whether these visual-motor expectations can be updated through training. In this paradigm participants trained at the peaks of this tuning function found in Experiment 1 for seven days. After training, this tuning function was remeasured. The results showed that participants coherence thresholds were significantly altered and motion detection performance was improved across all of the different directions. This result could be interpreted as an alteration to visual expectations or an overall perceptual learning effect for motion coherence detection through training. More work is needed to understand the mechanism behind this result.

The interaction between visual and motor processing is an important area for the field to investigate, as visual perception can be argued to be dependent on internal state factors such as actions and expectations (Khan & Hofer, 2018). Whilst it has been acknowledged for some time that it may not be useful to categorise between the sensory and motor regions of the brain when determining their relationship (Goodale 1998), it could still be argued that there is a sharp division in the way the field thinks about this issue today.

Nevertheless, I believe this chapter demonstrates a paradigm whereby these processes can be studied in combination and it would be particularly valuable for further research to build a metric for the swipe trajectories so that our understanding of this mechanism can be furthered.

A limitation however to my design, that relates to both experiments in this chapter, could be eye movements. A prerequisite in my experiments was that subjects-maintained fixation on the screen in front of them throughout each interval. However, I did not use eye tracking or a fixation point in the experiment to control for this. Part of the reasoning that eye-tracking was not used was due to the physical distance the touch monitor was to participants heads. In addition, the reason that a fixation point was not used was in case participants used it in some way as an indicator to detect coherent motion. Therefore, a potential caveat is that because I did not control for participants eye movements, they could have followed their finger across the screen when they activated each interval. Whilst we do not expect that the latter is true, owing to the size of the monitor and that all participants rested on their elbow to swipe in the corner of the screen, I lack objective supporting data for these potential issues. Thus, the use of either eye-tracking or fixation should be accounted for in further studies, so that the role of fixation can be further explored in this current paradigm.

That said, an important avenue that further research should continue to explore is the distinction between the mechanisms we know about in the visual and visual-motor fields. For example, in Experiment 1 the view that visual-motor expectations shape our perception, can certainly fit the data I provide. However, the sensitivity the visual system displayed for the axis of motion direction could also equally fit the mechanism for motion streaking. However, it remains to be determined how this mechanism relates to self-generated motor actions. An important experiment that could determine which one explains my data better, could confine participants swipe movements to a predetermined path, where the swipe action would activate viewable motion that would streak. It would also be important to follow this up by adapting parallel or orthogonal populations to understand which are relied upon for the task.

A similar endeavour could also be addressed for Experiment 2. For instance, whilst we made the assumption that visual-motor expectations may be updated through training, what we showed was similar to what you might find in the perceptual learning literature. In both Adams et al., (2010) and Sotiropoulos et al., (2011), they sought to understand the learning of expectations and specifically how they shape our perception in terms of bias. The perceptual learning field in its entirety is broad, and many investigations have not been explicitly designed to understand expectations. However, a growing view in the field is that top-down processes are necessary to produce PL (Zhang et al., 2010; Dosher & Lu, 1998; Li et al., 2004; Gold et al., 1999). Expectations and actions therefore may be a key part to this learning process however it remains to be determined what the exact mechanism looks like.

One might wonder whether the altering of expectations and perceptual learning improvements, could be explained by a similar mechanism, such as either top-down areas strengthening signals to the relevant neurons for the task (Gold et al., 1999), noise reduction (Dosher and Lu, 2000) or changes to decision rules or read-out strategies (Zhang et al., 2010). In light of these perceptual learning mechanisms, it is interesting in Experiment 2 that detection performance was generalised across direction post-training and was not specific to one particular direction. In comparison to other perceptual learning based experiments that have psychophysically demonstrated that motion detection is usually highly specific to features such as orientation, even with training-plus exposure paradigms that have shown to promote learning generalisation through exposure (Watanabe et al., 2001). Our result is suggestive more of a rule-based learning mechanism (Zhang et al., 2010) which dictates that learning can transfer provided the rules of the task have been learnt and the new features (e.g. other directions) to which the new rules must be applied have been adequately exposed. However, this still does not explain this generalisation, as the different directions were not necessarily exposed throughout training as the majority of participants chose the same swiping strategy. Further work could address this to understand how perceptual learning maps onto these sensitivity changes to the axis of self-generated motion post-training.

To conclude, the present study has introduced a touch-screen approach to investigating the integration between visual and motor processing systems. My results reveal that the visual system was sensitive to the axis of motion following self-generated motion and that this sensitivity can be flattened following repetitive training at the peak of the tuning function. Further research is however needed to help distinguish some of the potential mechanisms that could account for these results.

Chapter 7: General Discussion

7.1 Summary of Findings

The aim of this thesis was to explore internal factors relating to the adapted state of the visual system during perceptual learning. This thesis did this by addressing three main areas:

1) What is the relationship between bias and sensitivity whilst learning in an adapted state?
2) Can perceptual learning still occur on a task-relevant dimension when also passively exposed to variation along a coextensive task-irrelevant dimension? How might the relationship unfold between these two dimensions if the task relevancy switches over different timescales? Is passive exposure to a visual dimension enough to induce learnt perceptual improvements?
3) What are the characteristics of adaptation and perceptual sensitivity when the stimulus is self-generated compared to when it is externally-generated? Can perceptual sensitivity to self-generated stimuli be updated through perceptual learning?

The experiments in Chapters 3-6 of this thesis addressed these areas. An outline of each experimental chapter and their findings is provided below.

Chapter 3 of this thesis explored changes in perceptual bias and sensitivity during perceptual learning in an adapted state. The experiment utilised the nulling method and aftereffect duration for the MAE using the contrast between two gratings moving in opposite directions. The results indicated that perceptual bias induced by adaptation was consistent throughout the training protocol as both the nulling point and MAE duration followed the same consistent pattern and showed no systematic change in bias over the course of training. However, adapted state sensitivity did change. This change was characterised as a significant improvement in discrimination performance over the course of training which was specific to the trained adapted state. These results suggest that adapted state bias remains consistent, whilst it is adapted state sensitivity that improves during training. Overall, this raises questions about previous work demonstrating an interaction between perceptual learning and adaptation (Dong et al., 2016; 2019; McGovern et al., 2012). However, to the best of our knowledge, our results cannot be explained under the habituation learning account of visual adaptation, as our measure of perceptual bias remained consistent. Instead, our work agrees with previous research conducted in this laboratory (McGovern et al., 2012), that shows that training whilst in an adapted state improves perceptual sensitivity in that specific adapted state and does not alter perceptual bias (McGovern et al., 2012). It is unclear exactly what this mechanism might look like that can reflect adaptation and perceptual learning within the same framework. As such, one important goal for future work will be to further explore the nature of this state-dependent perceptual learning, both behaviourally and through neuroimaging.

In Chapter 4 of this thesis, I explored whether adaptation to a stimulus dimension that varied over a period of training on another coextensive stimulus dimension, would alter perceptual learning and how it would transfer. Previous research has indicated that adaptation (which is predominantly induced through passive exposure) to a stimulus dimension can cause performance improvements for that stimulus, despite no task training. Therefore, I investigated whether a similar finding could be found if the stimuli presented during training contained variation along a task-irrelevant dimension. This potential interaction between dimensions was explored over two different timescales of training and switching (e.g. one after the other or between trials). It was also compared to learning and transfer without variation along another task-irrelevant dimension. Overall, this chapter was comprised of five separate experimental groups. The results indicated that perceptual learning was not facilitated for the task-irrelevant dimension through adaptation exposure. Instead perceptual performance was facilitated by training on what was task-relevant. One caveat to this experiment may be inter-observer variability across our 50 participants, in their pre-training estimates for two task dimensions. Further research is certainly needed to understand more about adaptation and perceptual learning under more naturalistic settings.

In Chapter 5 of this thesis, I explored perceptual sensitivity and bias during adaptation to self-generated stimuli. Motivated by previous work that has shown that visual processing can be attenuated following self-generated control over stimuli, I tested whether bias and sensitivity whilst adapting may be altered if participants were responsible for generating and controlling it. This experiment was motivated by a lack of research into motion adaptation for self-generated stimuli and reports that self-generated control over stimuli in our environment may attenuate visual processing. There were three experimental conditions. Two of these were adaptation conditions, where adaptation to motion was either self-generated or externally generated. The final condition was a no adaptation condition. The results revealed no difference in adapted state bias or sensitivity between the self-generated and externally generated adaptation conditions. Instead, they revealed an overall effect of adaptation relative to the no adapted condition. This suggests that perceptual bias and sensitivity is not altered when adaptation is self-generated and instead, points to past research that has highlighted vestibular input as a factor that can modulate perceptual bias during adaptation.

While there were no significant changes between self-generated and externally generated adapted states, in the first part of Chapter 6 I showed that self-generated control over a motion stimulus improved detection performance when the viewable motion was congruent with the participants movement. The results indicated that participants were most sensitive when the motion matched their movements and most inhibited at detecting the motion when the motion was offset at 90°. This is in broad agreement with previous work that has shown sensory attenuation with self-generated stimuli. However, in this experiment I show a tuning function for this mechanism which shows a sensitivity to the axis of motion, rather than its direction. This could potentially be the mechanism for motion streaking although further research is needed to understand this. In the second part of Chapter 6, I wondered whether this tuning function for self-generated control is related to visual-motor expectations. Visual expectations are internal state factors that have been shown to shape perception. In addition, recent work has shown that these expectations can be reversed following repetitive training. To investigate this, participants completed several days of training where the viewable motion was continually offset by 90° with each self-generated action. Before and after this training period, participants completed the task in the same way that participants in the first part of this chapter had done. Results revealed that performance improved through training at 90°. However, when the tuning function was remeasured after training, they were not significantly different from each other and were at the same threshold level as the congruent (0°), suggesting an overall effect of training.

These brief summaries outline the main findings of this thesis. The proceeding sections will now consider these experimental findings in turn by addressing these findings in a wider context. In addition, issues and future research directions will be explored. This discussion begins with state-dependent perceptual learning (section 7.2) before moving onto self-generated adapted states (section 7.3).

7.2 The importance of State-Dependent Perceptual Learning

In Chapters 3 and 4 of this thesis, I attempted to understand more about the internal state that participants learn to perform a task in, which I characterised as state-dependent perceptual learning. Although framing perceptual learning and adaptation within the same framework makes intuitive sense, given that in the real-world our visual system more than likely achieves both at the same time, probing their existence within the same experimental design is complicated. Not only is it difficult to capture accurate thresholds for bias and sensitivity for a participant across several training days but given the relative scarcity of prior research on state-dependent perceptual learning, the conditions under which the internal state might be functionally useful for perceptual learning (if they do indeed interact) is poorly understood.

Undeniably, there are also a plethora of other internal state factors that may be important for perceptual learning. In particular, variables such as learning incentives, motivation, confidence, experiment duration as well as other unknown internal variables that may be highly relevant to the internal state that participants learn in. For obvious reasons related to resource and time, it is not feasible to attempt to test and understand all of these parameters. Thus, whilst this thesis focussed on quantifying adaptation and perceptual learning using psychophysical methodology, further research should continue to explore this field and those related to these alternative parameters, as there is very little prior research on the internal state during perceptual learning. That said, I believe that the work presented in this thesis provides a good starting account for state-dependent perceptual learning. However, inevitably this research also taps into several considerations that further research should consider. These will now be explored in the proceeding sections.

Theoretical Considerations for the Nature of State-Dependent Perceptual Learning

As already discussed in Chapter 3, one obvious theoretical consideration for further research should be to attempt to understand the underlying physiology involved in state-dependent perceptual learning. Unlike previous research that has shown perceptual bias decreases over repetitive training whilst in an adapted state (Dong et al., 2016; 2019; McGovern et al., 2012, Yehezkel et al., 2010), my research shows that perceptual bias does not change, instead perceptual sensitivity does. Thus, my findings are in obvious disagreement with this work as well as the Stimulus Model Comparator Theory of habituation (Sokolov, 1963) that can explain how perceptual bias decreases over repetitive adaptation. Further physiological investigation into this matter may help us understand the potential mechanism underlying both. For example, an experiment designed to measure the underlying physiological change for high- and low-level visual areas following repetitive adaptation to a stimulus could help determine whether the behavioural change I examined, is associated with any underlying physiological change in these areas.

As established in the Introduction to this thesis (Chapter 1), a popular theory of perceptual learning is that learning is the result of an improved extraction of task-relevant information via reweighting those neurons that represent the stimulus (Lu & Dosher, 2004). This is a process that takes place between low-level visual cortex responsible for representing the stimulus and high-level visual cortex responsible for reading-out from the low-level areas to reach a decision. Models of adaptation on the other hand, view this relationship between low- and high-level visual areas differently. For instance, adaptation mechanisms are typically modelled as the result of an alteration (mainly reductions in firing) to the underlying low-level sensory representation areas and biases a result of a faulty, but unchanged, read-out from higher-level cortical areas. Given what we understand about these mechanisms independently, it would be interesting to understand how they may operate physiologically together to try which would also help elucidate my findings.

fMRI experiments may also help identify a mapping between brain anatomy and function during interactions between perceptual learning and adaptation, and which brain areas are changing their activity levels as a result. To date many fMRI investigations into perceptual learning have shown increased BOLD activity to signals in V1 that correlates to the learning of a particular task (Furmanski et al,. 2004; Yotsumoto, Watanabe and Sasaki, 2008; Walker, Stickgold, Jolesz & Yoo, 2005). On the contrary, fMRI adaptation has shown that prolonged exposure to a stimulus leads to a selective reduction in BOLD signal in low-level sensory representation areas (Grill-Spector & Malach, 2001). Thus, an fMRI equivalent study to the protocol we employed in Chapter 3, where participants performed the task for the trained and untrained locations pre- and post-training whilst in the scanner, could help reveal functional and anatomical differences between adapted states following training. In addition, it would also be interesting to see whether the BOLD signal would be suppressed or increased in V1 for the trained side following such training.

The findings from Chapter 3 and 4 also have theoretical implications for perceptual learning as a potential therapeutic or skill-based intervention that further research should take into consideration. As my results in Chapter 3 show that there is some form of interaction between perceptual learning and adaptation, it could be the case that for some visual skill-based tasks, more desirable learning outcomes (i.e. faster learning, greater accuracy) may be found if the adapted state that the person was trained in was brought into consideration. For example, many real-world contexts have tightly linked statistical regularities in their visual input (i.e.. the high prevalence of near vertical contrast energy because of the quantity of trees found in the forest). As such, when the visual system performs and learns on a task within that natural environment it may use elements of that adapted state in the learning of that task.

In Chapter 4, I explored exposure to fixed and varying stimuli and whether this exposure could have a beneficial impact on perceptual learning. Whilst not directly comparable to the results in Chapter 3, I was specifically interested in whether passive exposure to a visual stimulus dimension that varied, at the same time as training on another, would increase perceptual sensitivity for the exposed visual dimension despite no task training. This finding has been shown by other researchers in the field (Wang et al., 2012; Xiao et al., 2008; Zhang et al, 2010; Szpiro et al., 2016; Watanabe et al., 2001; Seitz et al., 2009; Huang, Lu & Dosher, 2012; Petrov et al., 2005, Gutnisky et al., 2009), which is important for perceptual learning as a potential therapeutic intervention. Whilst my results in Chapter 4 did not agree with previous research, further research should continue to explore exposure to fixed and varying stimuli during perceptual learning, as the combination of task-relevant and task-irrelevant dimensions creates a more naturalistic approach to understanding visual processing during perceptual learning and adaptation.

As already outlined in Chapter 4, perceptual tasks outside of the laboratory contain dimensions that are coextensive with each other and what is task-relevant dictates which feature is exposed, and which is learnt. One consideration that future research could explore is the temporal and spatial relationship between the learnt on and exposed task dimensions. The separation of task dimensions in space may be an important component for exposure-based performance improvements, as it is one key distinction between the research that has shown performance improvements following exposure (Wang et al., 2012; Xiao et al., 2008; Zhang et al, 2010; Szpiro et al., 2016; Watanabe et al., 2001; Seitz et al., 2009; Gutnisky et al., 2009), as well as my and other researchers results (Huang, Lu & Dosher, 2012; Petrov et al., 2005).

Another way in which future work could address this could be to follow-on from Deveau and colleagues (2014) work, who showed how interleaving lab-based perceptual learning paradigms with real-world tasks, that share certain characteristics, may improve performance on a real-world task. However, rather than focus exclusively on performance on the task, researchers could emulate the adapted state of the real-world task as much as possible during training. For instance, perhaps this research could use filtered video methodologies, where a recording of the real-world task are replayed and particular surroundings in the video recording altered. Given this, it would be interesting if this could achieve similar learning outcomes if they are just exposed to this video before performing a task. This could be practical for learning outcomes and would be especially interesting for those real-world tasks that cannot be practiced without adequate support (i.e. trainee jet pilots).

Another interesting theoretical consideration is the extent to which this type of exposure can be considered adaptation in our and other experimental paradigms. The assumption we made in Chapter 4, was that the passive exposure of a stimulus would be encoded in the visual system like adaptation would. Arguably, this is also the assumption that other researchers make, who view the passive exposure to the stimulus as some form of "release from adaptation" which explains their performance improvement (Gutnisky et al., 2009). As such however, many of these exposure-based learning accounts have no measure of whether the visual system is adapting to the visual attribute over the course of training. This would be important to understand and is perhaps a question better suited to the next section on methodological considerations.

Methodological Considerations for State-Dependent Perceptual Learning

In Chapters 3 and 4 of this thesis, I used novel methods to attempt to understand more about the dynamics of state-dependent perceptual learning. Long-term training protocols on a challenging task, lasting for several days and hours, is the traditional method for investigating perceptual learning in isolation. However, probing the adapted state of the visual system whilst participants learn on that task has seldom been studied. In addition, only a handful of experiments have measured the dynamics of the adapted state over timescales lasting several days with a view to understanding how it changes overtime. Thus, whilst I believe that the experimental methods used in Chapters 3 and 4 of this thesis were valid, even though they are somewhat novel and undertested, I also believe that the experimental protocols used inevitably taps into some methodological issues that future research should consider.

To begin, I adopted a nulling method in Chapter 3, that allowed my research to measure changes in both bias and sensitivity, the key hallmarks for adaptation and perceptual learning, at the same time. This was important because in the handful of experiments that have measured bias and sensitivity within the same experimental protocol, they are captured using distinct methods at separate times where the adapted state could change between them.

In addition, to understand more about state-dependent perceptual learning between states, the method used in Chapter 4 aimed to understand more about the dynamics of exposure-based adaptation and its potential for increasing perceptual learning benefits. Previous reports that have shown performance improvements following passive exposure, typically measure performance on the exposed task briefly before and after the training period or after every daily session. In addition, task dimensions are not coextensive, which makes it difficult to understand exactly how perceptual performance is improving at that location if the participant does not directly attend to its features. Therefore, to overcome these issues, the method we adopted utilised a multidimensional environment that could be accurately comprised of two task dimensions that were coextensive. This way, when participants learnt on one task they were passively exposed (adapting) to the other task dimension. Moreover, rather than measuring task performance briefly at a given time point, interleaving the learning and exposure on the two tasks was thought to be a better way of capturing any influence on performance following exposure.

In light of our methodological choices, a key methodological issue for future consideration is how to better capture the adapted state during learning. Adaptation is a rapid process that is modulated by brief as well as long periods of passive exposure to a stimulus, thus quantifying the dynamics of adaptation overtime and the resulting bias is complicated. For instance, the adapted state I measured in Chapter 3, was inevitably modulated by the several seconds of the adaptor as well as the short periods of the test probe used for measurement. Instead of measuring adaptation over this type of timescale, I quantified it over a much more prolonged period - across several trials within a block. This single estimate for bias is a rather coarse measurement in comparison to other methods that can measure bias at much higher temporal resolutions.

A good alternative to explore may be matching tasks (for introduction see: Pelli & Farell, 1995; for examples with adaptation see: Wexler, Duyck & Mamassian, 2015; Wexler; 2018), that are methods where participants adapt and then reproduce the displayed stimulus by providing estimates for their perception of the stimulus. As such, these methods have a much finer temporal resolution when it comes to measuring the dynamics of the adapted state and can capture the shape of adaptation build up and decay. An alternative psychophysical paradigm that could address this could be to utilise both the matching task and nulling method within a block of adaptation. This could be possible by interleaving them within each block and would be useful for understanding each time point in adaptation build up or decay, within and across sessions. Furthermore, comparing this against some of the nulling point data could help us understand more about the adaptation dynamics during state-dependent learning and how adaptation dynamics may be improving perceptual performance.

7.1.3 Self-Generated Adaptation and Visual-Motor Expectations

In Chapters 5 and 6 of this thesis, I focused on visual processing following self-generated control over stimuli. My research was motivated in part by studies that have shown that visual motion processing can be altered when stimuli are self-generated. This influenced my first question in Chapter 5 which investigated whether adaptation would habituate if the adaptation probe was self-generated opposed to externally-generated. To follow on from this I next focused on visual sensitivity for self-generated stimuli and investigated whether the pattern of sensitivity for self-generated motion detection, could be modulated if participants train in an altered state of expectations.

The key finding from Chapter 5 was that self-generated motion adaptation was not quantitatively different from externally generated adaptation involving moving gratings. The key findings of Chapter 6 however, suggest that sensitivity to self-generated motion is inhibited the further the axis of motion is offset relative to the swiping action of the observer. Although the exact mechanism behind this however should be investigated further and attempt to understand whether this is a motion-streaking mechanism (Geisler, 1999). In addition, in the second part of this chapter the shape of this tuning function could be altered through prolonged training but my results could have reflected an overall perceptual learning effect. Further research could resolve this by measuring visual detection performance to the axis of motion, pre- and post- a period of no training. In combination, these experiments highlight some other issues for further research to consider for adapted state processing and internal factors related to visual-motor expectations and perceptual learning.

Chapters 5 and 6 also raise several methodological considerations in probing aspects of visual processing when stimuli are self-generated. In both chapters I used mechanical devices, such the track-ball mouse in Chapter 5 and the touch screen monitor in Chapter 6. As apparatus, these mechanical devices are useful for studying aspects to multisensory integration and visual processing, however, are restricted by the type of self-motion they can measure. For instance, in both chapters the type of self-motion was limited to very specific movements made by the participants fingers and hand, which is limited when compared to how we interact with objects under real-world natural settings. Thus, mechanical complexity and lack of portability may restrict the generality of the methods we utilised and therefore should be a consideration for methodological choices in future research.

One factor that should inform methodological choice is ensuring that the sense of self-agency over the mechanical device and manipulated stimuli is created and maintained. It is possible that the sense of self-agency participants felt in Chapter 5 for instance, quickly disappeared due to the repetitive nature of the swiping motion over the trackball within consistently timed intervals, and as a consequence may have held back any mechanisms that could have altered adapted state visual processing. Thus, although participants actions were directly generating the motion, they were just reactions to the unnatural constraints of the experimental protocol and as such that sense of recognising oneself as the agent of a behaviour may have disappeared. Interestingly this was not the case in Chapter 6, thus future research could investigate whether the results in Chapter 5 would be different if it was under the conditions used in Chapter 6.

Another consideration for future research is the length of the adaptation period when it is self-generated. In other research that has investigated motor control and visual processing, the self-generated stimuli are often constrained to much shorter time windows spanning milliseconds (Ichikawa & Masakura, 2004; 2006; Ichikawa et al., 2010; Konkle et al, 2009; Dewey & Carr, 2013). Perhaps the sense of self-agency is easier to set up in these experiments than in Chapter 5 of this thesis, where the motion lasts much longer and where the maximum speed of the stimulus was reached in a relatively quick time. Although further research should explore this issue further.

An alternative method that future research could also utilise is the use of virtual reality (VR) devices. The use of VR devices has been used previously by other research laboratories to investigate prolonged adaptation (Kwon et al., 2009; Zhang et al., 2009; Haak et al.,2014) and response suppression of visual processing following self-motion (Bai et al., 2019). They are also becoming more easily accessible in cost and in calibration, as MATLAB and the Psychophysics Toolbox (Brainard, 1997) have programming code that can be easily customisable with their use (see Bai et al., 2019 for details). The main advantage of such devices for investigating self-generated motion is its capability to satisfy various kinds of research goals in the field owing to their adaptability and efficient, lightweight size. For instance, sensors on various parts of the body can be used in conjunction with altered real-time visual input being displayed on the head-mounted display. Altering where these sensors go on the body or the types of movements they record can easily be adapted and calibrated so that any restrictiveness and thus lack of generalisability that may be present in the use of mechanical devices can be overcome. The only factor that researchers should consider if choosing to use VR in this way, is whether to control for eye movement because saccades have been shown to have suppressive properties on visual motion processing (Bridgeman et al., 1994). Therefore, a key issue that future research should consider if attempting to understand whether self-generated action suppresses or alters visual processing, should be whether to use some method of keeping good central fixation to rule that out as a factor.

Another issue for future research, which was discussed in Chapter 6, is choosing a method that would allow for a meaningful metric to be derived from participant's motor movements. In Chapter 6, I did not constrain participant swiping movements. This meant that participants were able to freely swipe in any direction following any path. Whilst this approach carried advantages and was interesting, it came at the cost of not being able to derive a meaningful swiping metric that could be applied across all participants. This metric would have been interesting in particular for understanding more about this mechanism that due to its sensitivity to the axis of motion direction, may be the same mechanism that has been shown by those investigating motion streaking (Geisler, 1999; Apthorp et al., 2013). Thus, one direction that future research should take is to understand how participants swipe trajectories and swipe speed related to their detection performance for each motion direction offset. As such, a possible experimental design for future work could be to constrain the swiping movements more by having the participants start and end their swiping trajectory between two points on the touch screen. This could be between two random locations which would mean that the direction and speed could be controlled for. This analysis could help understand further whether estimates of motion orientation become more reliable as speed increases or certain swiping strategies may aid perceptual performance. In addition, this metric could also help us compare participant overall swipe trajectory within and across session, to understand whether certain swipe trajectory strategies facilitated motion detection.

7.4 Concluding Remarks

In this thesis, I have explored internal state factors of the visual system during perceptual learning. I began by exploring the adapted state during perceptual learning where we identified that adapted state-bias remains consistent, whereas adapted state sensitivity improves for the trained adapted state. Subsequently I examined whether exposure to a task dimension could facilitate perceptual learning. Specifically, I was interested in testing whether a build-up of adaptation exposure to a visual dimension would be beneficial to perceptual learning. However, I showed no such benefit between adapted states and learning improvement was only achieved through training.

In the following Chapters 5 and 6, I next tested how internal factors relating to adapted state processing. I was specifically interested in testing the hypothesis that self-generated control over visual stimuli leads to attenuation in visual processing. In Chapter 5, I showed very little support for this hypothesis, with no quantitative difference between self-generated or externally generated adaptation. In Chapter 6 however, I did find support for this hypothesis and showed that motion detection performance was altered when it was self-generated and congruent with the participants action. The exact mechanism of this effect remains to be determined by future research. However, as I showed that this mechanism can be altered through prolonged training this highlights that perceptual learning may be a factor.

The empirical findings of Chapters 3-6 highlight the needs for a greater understanding of the state-dependent perceptual learning and factors relating to visual processing during perceptual learning. Exploring these factors further and developing this framework will help guide future research on better functional outcomes of perceptual learning as a real-world practical intervention.

References

Aberg, K. C., Tartaglia, E. M., & Herzog, M. H. (2009). Perceptual learning with Chevrons requires a minimal number of trials, transfers to untrained directions, but does not require sleep. Vision Research, 49(16), 2087-2094.

Adelson, E. H., & Movshon, J. A. (1982). Phenomenal coherence of moving visual patterns. Nature, 300(5892), 523-525.

Ahissar, M., & Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. Nature, 387(6631), 401-406.

Ahissar, M., & Hochstein, S. (2004). The reverse hierarchy theory of visual perceptual learning. Trends in Cognitive Sciences, 8(10), 457-464.

Aliu, S. O., Houde, J. F., & Nagarajan, S. S. (2009). Motor-induced suppression of the auditory cortex. Journal of Cognitive Neuroscience, 21(4), 791-802.

Andrews, D. (1964). Error-correcting perceptual mechanisms. Quarterly Journal of Experimental Psychology, 16(2), 104-115.

Apthorp, D., Schwarzkopf, D. S., Kaul, C., Bahrami, B., Alais, D., & Rees, G. (2013). Direct evidence for encoding of motion streaks in human visual cortex. Proceedings of the Royal Society B: Biological Sciences, 280(1752), 20122339.

Attneave, F. (1954). Some informational aspects of visual perception. Psychological Review, 61(3), 183.

Bai, J., Bao, M., Zhang, T., & Jiang, Y. (2019). A virtual reality approach identifies flexible inhibition of motion aftereffects induced by head rotation. Behavior Research Methods, 51(1), 96-107.

Ball, K., & Sekuler, R. (1987). Direction-specific improvement in motion discrimination. Vision Res, 27(6), 953-965.

Barlow, H. (1990). A theory about the functional role and synaptic mechanism of visual after-effects. Vision: Coding and efficiency, 363375.

Barlow, H., & Foldiak, P. (1989). The computing neuron. Adaptation and decorrelation in the cortex, 54-72.

Barlow, H., Hill, R., & Levick, W. (1964). Retinal ganglion cells responding selectively to direction and speed of image motion in the rabbit. The Journal of Physiology, 173(3), 377.

Barlow, H. B. (1997). The knowledge used in vision and where it comes from. Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences, 352(1358), 1141-1147.

Bastian, A. J. (2008). Understanding sensorimotor adaptation and learning for rehabilitation. Current Opinion in Neurology, 21(6), 628.

Bays, P. M., & Wolpert, D. M. (2007). Computational principles of sensorimotor control that minimize uncertainty and variability. The Journal of Physiology, 578(2), 387-396.

Berliner, J., & Durlach, N. (1973). Intensity perception. IV. Resolution in roving-level discrimination. The Journal of the Acoustical Society of America, 53(5), 1270-1287.

Blakemore, C., & Campbell, F. W. (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.

Blakemore, C., & Sutton, P. (1969). Size adaptation: A new aftereffect. Science, 166(3902), 245-247.

Blakemore, S.-J., Wolpert, D. M., & Frith, C. D. (1998). Central cancellation of self-produced tickle sensation. Nature Neuroscience, 1(7), 635-640.

Bompas, A., & O'Regan, J. K. (2006a). Evidence for a role of action in colour perception. Perception, 35(1), 65-78.

Bompas, A., & O'Regan, J. K. (2006b). More evidence for sensorimotor adaptation in color perception. Journal of Vision, 6(2), 5-5.

Bridgeman, B., Van der Heijden, A., & Velichkovsky, B. M. (1994). A theory of visual stability across saccadic eye movements. Behavioral and Brain Sciences, 17(2), 247-257.

Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., & Movshon, J. A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. Visual Neuroscience, 13(1), 87-100.

Burr, D. C., & Ross, J. (2002). Direct evidence that “speedlines” influence motion mechanisms. Journal of Neuroscience, 22(19), 8661-8664.

Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., . . . Rust, N. C. (2005). Do we know what the early visual system does? Journal of Neuroscience, 25(46), 10577-10597.

Carandini, M., Movshon, J. A., & Ferster, D. (1998). Pattern adaptation and cross-orientation interactions in the primary visual cortex. Neuropharmacology, 37(4-5), 501-511.

Castet, E., Keeble, D. R., & Verstraten, F. A. (2002). Nulling the motion aftereffect with dynamic random-dot stimuli: limitations and implications. Journal of Vision, 2(4), 3-3.

Censor, N., Karni, A., & Sagi, D. (2006). A link between perceptual learning, adaptation and sleep. Vision Research, 46(23), 4071-4074.

Chaudhuri, A. (1990). Modulation of the motion aftereffect by selective attention. Nature, 344(6261), 60-62.

Chopin, A., & Mamassian, P. (2012). Predictive properties of visual adaptation. Current Biology, 22(7), 622-626.

Christensen, A., Ilg, W., & Giese, M. A. (2011). Spatiotemporal tuning of the facilitation of biological motion perception by concurrent motor execution. Journal of Neuroscience, 31(9), 3493-3499.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behavioral and Brain Sciences, 36(3), 181-204.

Clifford, C. W., Webster, M. A., Stanley, G. B., Stocker, A. A., Kohn, A., Sharpee, T. O., & Schwartz, O. (2007). Visual adaptation: Neural, Psychological and Computational Aspects. Vision research, 47(25), 3125-3131.

Clifford, C. W., Wenderoth, P., & Spehar, B. (2000). A functional angle on some after-effects in cortical vision. Proceedings of the Royal Society of London. Series B: Biological Sciences, 267(1454), 1705-1710.

Colby, C. L., & Goldberg, M. E. (1999). Space and attention in parietal cortex. Annual Review of Neuroscience, 22(1), 319-349.

Coles, M. G., Gratton, G., Bashore, T. R., Eriksen, C. W., & Donchin, E. (1985). A psychophysiological investigation of the continuous flow model of human information processing. Journal of Experimental Psychology: Human Perception and Performance, 11(5), 529.

Coltheart, M. (1971). Visual feature-analyzers and aftereffects of tilt and curvature. Psychological Review, 78(2), 114.

Cong, L.-J., Wang, R.-J., Yu, C., & Zhang, J.-Y. (2016). Perceptual learning of basic visual features remains task specific with Training-Plus-Exposure (TPE) training. Journal of Vision, 16(3), 13-13.

Crist, R. E., Kapadia, M. K., Westheimer, G., & Gilbert, C. D. (1997). Perceptual learning of spatial localization: specificity for orientation, position, and context. J Neurophysiol, 78(6), 2889-2894.

Crist, R. E., Li, W., & Gilbert, C. D. (2001). Learning to see: experience and attention in primary visual cortex. Nature Neuroscience, 4(5), 519-525.

De Valois, K. K. (1977). Spatial frequency adaptation can enhance contrast sensitivity. Vision Research, 17(9), 1057-1065.

Deveau, J., & Seitz, A. R. (2014). Applying perceptual learning to achieve practical changes in vision. Frontiers in Psychology, 5, 1166.

Dewey, J. A., & Carr, T. H. (2013). Predictable and self-initiated visual motion is judged to be slower than computer generated motion. Consciousness and Cognition, 22(3), 987-995.

Dong, X., & Bao, M. (2019). Direction selective habituation of motion adaptation. Journal of Vision, 19(4), 6-6.

Dong, X., Gao, Y., Lv, L., & Bao, M. (2016). Habituation of visual adaptation. Scientific reports, 6, 19152.

Dosher, B. A., & Lu, Z.-L. (1998). Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. Proceedings of the National Academy of Sciences, 95(23), 13988-13993.

Dosher, B. A., & Lu, Z.-L. (1999). Mechanisms of perceptual learning. Vision Research, 39(19), 3197-3221.

Dragoi, V., Sharma, J., & Sur, M. (2000). Adaptation-induced plasticity of orientation tuning in adult visual cortex. Neuron, 28(1), 287-298.

Durant, S., Clifford, C. W., Crowder, N. A., Price, N. S., & Ibbotson, M. R. (2007). Characterizing contrast adaptation in a population of cat primary visual cortical neurons using Fisher information. JOSA A, 24(6), 1529-1537.

Edwards, M., & Crane, M. F. (2007). Motion streaks improve motion detection. Vision Research, 47(6), 828-833.

Fahle, M. (2004). Perceptual learning: a case for early selection. J Vis, 4(10), 879-890.

Fahle, M., & Henke-Fahle, S. (1996). Interobserver variance in perceptual performance and learning. Investigative Ophthalmology & Visual Science, 37(5), 869-877.

Fahle, M., & Morgan, M. (1996). No transfer of perceptual learning between similar stimuli in the same retinal position. Curr Biol, 6(3), 292-297.

Felsen, G., & Dan, Y. (2005). A natural approach to studying vision. Nature Neuroscience, 8(12), 1643-1646.

Fiorentini, A., & Berardi, N. (1980). Perceptual learning specific for orientation and spatial frequency. Nature, 287(5777), 43-44.

Fiorentini, A., & Berardi, N. (1997). Visual perceptual learning: a sign of neural plasticity at early stages of visual processing. Arch Ital Biol, 135(2), 157-167.

Foster, D. H. (2011). Color constancy. Vision Research, 51(7), 674-700.

Foster, K., Gaska, J. P., Nagler, M., & Pollen, D. (1985). Spatial and temporal frequency selectivity of neurones in visual cortical areas V1 and V2 of the macaque monkey. The Journal of Physiology, 365(1), 331-363.

Francis, G., & Kim, H. (2001). Perceived motion in orientational afterimages: direction and speed. Vision Research, 41(2), 161-172.

Furmanski, C. S., Schluppeck, D., & Engel, S. A. (2004). Learning strengthens the response of primary visual cortex to simple patterns. Current Biology, 14(7), 573-578.

Gardner, J. L., Sun, P., Waggoner, R. A., Ueno, K., Tanaka, K., & Cheng, K. (2005). Contrast adaptation and representation in human early visual cortex. Neuron, 47(4), 607-620.

Geisler, W. S. (1999). Motion streaks provide a spatial code for motion direction. Nature, 400(6739), 65-69.

Geisler, W. S., & Albrecht, D. G. (1992). Cortical neurons: isolation of contrast gain control. Vision Research, 32(8), 1409-1410.

Georgeson, M. A. (1973). Spatial frequency selectivity of a visual tilt illusion. Nature, 245(5419), 43-45.

Gibson, E. J. (1969). Principles of perceptual learning and development.

Gilbert, C. D., Sigman, M., & Crist, R. E. (2001). The neural basis of perceptual learning. Neuron, 31(5), 681-697.

Gold, J. I., & Stocker, A. A. (2017). Visual decision-making in an uncertain and dynamic world. Annual Review of Vision Science, 3, 227-250.

Goodale, M. A. (1998). Visuomotor control: Where does vision end and action begin? Current Biology, 8(14), R489-R491.

Greenlee, M. W., Georgeson, M. A., Magnussen, S., & Harris, J. P. (1991). The time course of adaptation to spatial contrast. Vision Research, 31(2), 223-236.

Gur, M., & Snodderly, D. M. (2007). Direction selectivity in V1 of alert monkeys: evidence for parallel pathways for motion processing. The Journal of Physiology, 585(2), 383-400.

Gutnisky, D. A., Hansen, B. J., Iliescu, B. F., & Dragoi, V. (2009). Attention alters visual plasticity during exposure-based learning. Current Biology, 19(7), 555-560.

Haak, K. V., Fast, E., Bao, M., Lee, M., & Engel, S. A. (2014). Four days of visual contrast deprivation reveals limits of neuronal adaptation. Current Biology, 24(21), 2575-2579.

Haarmeier, T., Bunjes, F., Lindner, A., Berret, E., & Thier, P. (2001). Optimizing visual motion perception during eye movements. Neuron, 32(3), 527-535.

Haijiang, Q., Saunders, J. A., Stone, R. W., & Backus, B. T. (2006). Demonstration of cue recruitment: Change in visual appearance by means of Pavlovian conditioning. Proceedings of the National Academy of Sciences, 103(2), 483-488.

Hammond, P., Pomfrett, C., & Ahmed, B. (1989). Neural motion after-effects in the cat's striate cortex: orientation selectivity. Vision Research, 29(12), 1671-1683.

Harris, H., Gliksberg, M., & Sagi, D. (2012). Generalized perceptual learning in the absence of sensory adaptation. Current Biology, 22(19), 1813-1817.

Harris, L., Morgan, M., & Still, A. (1981). Moving and the motion after-effect. Nature, 293(5828), 139-141.

Harwerth, R. S., & Levi, D. M. (1978). Reaction time as a measure of suprathreshold grating detection. Vision Res, 18(11), 1579-1586.

Hawken, M., Shapley, R. M., & Grosof, D. (1996). Temporal-frequency selectivity in monkey visual cortex. Vis Neurosci, 13(3), 477-492.

Helson, H. (1964). Adaptation-level theory: an experimental and systematic approach to behavior.

Hesse, M. D., Nishitani, N., Fink, G. R., Jousmäki, V., & Hari, R. (2010). Attenuation of somatosensory responses to self-produced tactile stimulation. Cerebral Cortex, 20(2), 425-432.

Hiris, E., & Blake, R. (1995). Discrimination of coherent motion when local motion varies in speed and direction. Journal of Experimental Psychology: Human Perception and Performance, 21(2), 308.

Hol, K., & Treue, S. (2001). Different populations of neurons contribute to the detection and discrimination of visual motion. Vision Research, 41(6), 685-689.

Huang, C.-B., Lu, Z.-L., & Dosher, B. A. (2012). Co-learning analysis of two perceptual learning tasks with identical input stimuli supports the reweighting hypothesis. Vision Research, 61, 25-32.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. The Journal of Physiology, 160(1), 106.

Ichikawa, M., & Masakura, Y. (2006). Manual control of the visual stimulus reduces the flash-lag effect. Vision Research, 46(14), 2192-2203.

Ichikawa, M., & Masakura, Y. (2013). Effects of consciousness and consistency in manual control of visual stimulus on reduction of the flash-lag effect for luminance change. Frontiers in Psychology, 4, 120.

Jeter, P. E., Dosher, B. A., Liu, S.-H., & Lu, Z.-L. (2010). Specificity of perceptual learning increases with increased training. Vision Research, 50(19), 1928-1940.

Jeter, P. E., Dosher, B. A., Petrov, A., & Lu, Z.-L. (2009). Task precision at transfer determines specificity of perceptual learning. Journal of Vision, 9(3), 1-1.

Jones, R. M., & Tulunay-Keesey, U. (1975). Local retinal adaptation and spatial frequency channels. Vision Research, 15(11), 1239-1244.

Karni, A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. Proceedings of the National Academy of Sciences, 88(11), 4966-4970.

Kattner, F., Cochrane, A., Cox, C. R., Gorman, T. E., & Green, C. S. (2017). Perceptual learning generalization from sequential perceptual training as a change in learning rate. Current Biology, 27(6), 840-846.

Khan, A. G., & Hofer, S. B. (2018). Contextual signals in visual cortex. Current Opinion in Neurobiology, 52, 131-138.

Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. Trends in Cognitive Sciences, 10(7), 319-326.

Krauskopf, J., & Karl, G. (1992). Color discrimination and adaptation. Vision Research, 32(11), 2165-2175.

Kristjánsson, Á. (2011). The functional benefits of tilt adaptation. Seeing and perceiving, 24(1), 37-51.

Kuai, S.-G., Zhang, J.-Y., Klein, S. A., Levi, D. M., & Yu, C. (2005). The essential role of stimulus temporal patterning in enabling perceptual learning. Nature Neuroscience, 8(11), 1497-1499.

Kwon, M., Legge, G., Fang, F., Cheong, A., & He, S. (2009). Identifying the mechanism of adaptation to prolonged contrast reduction. Journal of Vision, 9(8), 976-976.

Lally, N., Frendo, B., & Diedrichsen, J. (2011). Sensory cancellation of self-movement facilitates visual motion detection. Journal of Vision, 11(14), 5-5.

Law, C.-T., & Gold, J. I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nature Neuroscience, 11(4), 505-513.

Ledgeway, T. (1994). Adaptation to second-order motion results in a motion aftereffect for directionally-ambiguous test stimuli. Vision Research, 34(21), 2879-2889.

Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. Nature Neuroscience, 4(1), 89-94.

Lepora, N. F., & Pezzulo, G. (2015). Embodied choice: how action influences perceptual decision making. PLoS Comput Biol, 11(4), e1004110.

Levi, D. M., & Polat, U. (1996). Neural plasticity in adults with amblyopia. Proceedings of the National Academy of Sciences, 93(13), 6830-6834.

Levinson, E., & Sekuler, R. (1976). Adaptation alters perceived direction of motion. Vision Research.

Li, R. W., Klein, S. A., & Levi, D. M. (2008). Prolonged perceptual learning of positional acuity in adult amblyopia: perceptual template retuning dynamics. Journal of Neuroscience, 28(52), 14223-14229.

Liang, J., Zhou, Y., Fahle, M., & Liu, Z. (2015a). Limited transfer of long-term motion perceptual learning with double training. Journal of Vision, 15(10), 1-1.

Liang, J., Zhou, Y., Fahle, M., & Liu, Z. (2015b). Specificity of motion discrimination learning even with double training and staircase. Journal of Vision, 15(10), 3-3.

Martin, T., Keating, J., Goodkin, H., Bastian, A., & Thach, W. (1996). Throwing while looking through prisms: II. Specificity and storage of multiple gaze—throw calibrations. Brain, 119(4), 1199-1211.

Mather, G., Pavan, A., Campana, G., & Casco, C. (2008). The motion aftereffect reloaded. Trends in Cognitive Sciences, 12(12), 481-487.

Mather, G. E., Verstraten, F. E., & Anstis, S. E. (1998). The motion aftereffect: A modern perspective: The MIT Press.

McDermott, K. C., Malkoc, G., Mulligan, J. B., & Webster, M. A. (2010). Adaptation and visual salience. Journal of Vision, 10(13), 17-17.

McGovern, D. P., Roach, N. W., & Webb, B. S. (2012). Perceptual learning reconfigures the effects of visual adaptation. Journal of Neuroscience, 32(39), 13621-13629.

McGovern, D. P., Roach, N. W., & Webb, B. S. (2014). Characterizing the effects of multidirectional motion adaptation. Journal of Vision, 14(13), 2-2.

McGovern, D. P., Webb, B. S., & Peirce, J. W. (2012). Transfer of perceptual learning between different visual tasks. Journal of Vision, 12(11), 4-4.

Mednick, S. C., Arman, A. C., & Boynton, G. M. (2005). The time course and specificity of perceptual deterioration. Proceedings of the National Academy of Sciences, 102(10), 3881-3885.

Mitchell, D. E., & Ware, C. (1974). Interocular transfer of a visual after-effect in normal and stereoblind humans. The Journal of Physiology, 236(3), 707-721.

Mollon, J. D., & Danilova, M. V. (1996). Three remarks on perceptual learning. Spatial Vision, 10(1), 51-58.

Movshon, J., Adelson, E., Gizzi, M., Newsome, W., Chagas, C., Gattass, R., & Gross, C. (1985). Pattern Recognition Mechanisms.

Movshon, J. A., & Newsome, W. T. (1996). Visual response properties of striate cortical neurons projecting to area MT in macaque monkeys. Journal of Neuroscience, 16(23), 7733-7741.

Newell, A., & Simon, H. A. (1972). Human problem solving (Vol. 104): Prentice-Hall Englewood Cliffs, NJ.

Nishida, S. y., & Ashida, H. (2000). A hierarchical structure of motion system revealed by interocular transfer of flicker motion aftereffects. Vision Research, 40(3), 265-278.

Nishida, S. y., & Sato, T. (1995). Motion aftereffect with flickering test patterns reveals higher stages of motion processing. Vision Research, 35(4), 477-490.

O'Toole, A. J., & Kersten, D. J. (1992). Learning to see random-dot stereograms. Perception, 21(2), 227-243.

Olshausen, B. A., & Field, D. J. (2005). How close are we to understanding V1? Neural Computation, 17(8), 1665-1699.

Ölveczky, B. P., Baccus, S. A., & Meister, M. (2003). Segregation of object and background motion in the retina. Nature, 423(6938), 401-408.

Orban, G. A. (2008). Higher order visual processing in macaque extrastriate cortex. Physiological Reviews.

Oruç, I., & Barton, J. J. (2011). Adaptation improves discrimination of face identity. Proceedings of the Royal Society B: Biological Sciences, 278(1718), 2591-2597.

Otto, T. U., Herzog, M. H., Fahle, M., & Zhaoping, L. (2006). Perceptual learning with spatial uncertainties. Vision Research, 46(19), 3223-3233.

Pack, C. C., Berezovskii, V. K., & Born, R. T. (2001). Dynamic properties of neurons in cortical area MT in alert and anaesthetized macaque monkeys. Nature, 414(6866), 905-908.

Pantle, A., & Sekuler, R. (1968). Size-detecting mechanisms in human vision. Science, 162(3858), 1146-1148.

Pavard, B., & Berthoz, A. (1977). Linear acceleration modifies the perceived velocity of a moving visual scene. Perception, 6(5), 529-540.

Peirce, J. W. (2007). PsychoPy—psychophysics software in Python. Journal of Neuroscience methods, 162(1-2), 8-13.

Pelli, D. G., & Farell, B. (1995). Psychophysical methods. Handbook of optics, 1, 29.21-29.13.

Petrov, A. A., Dosher, B. A., & Lu, Z.-L. (2005). The dynamics of perceptual learning: an incremental reweighting model. Psychological Review, 112(4), 715.

Petrov, A. A., & Van Horn, N. M. (2012). Motion aftereffect duration is not changed by perceptual learning: Evidence against the representation modification hypothesis. Vision Research, 61, 4-14.

Phinney, R. E., Bowd, C., & Patterson, R. (1997). Direction-selective coding of stereoscopic (cyclopean) motion. Vision Research, 37(7), 865-869.

Poggio, T., Fahle, M., & Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. Science, 256(5059), 1018-1021.

Polat, U. (2009). Making perceptual learning practical to improve visual functions. Vision Res, 49(21), 2566-2573.

Price, N. S., & Prescott, D. L. (2012). Adaptation to direction statistics modulates perceptual discrimination. Journal of Vision, 12(6), 32-32.

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

Rauber, H.-J., & Treue, S. (1999). Revisiting motion repulsion: evidence for a general phenomenon? Vision Research, 39(19), 3187-3196.

RAYMOND, J., & BRADDICK, O. (1996). Responses to Opposed Directions of Motion:: Continuum or Independent Mechanisms? Vision Research, 36(13), 1931-1937.

Regan, D., & Beverley, K. (1983). Spatial-frequency discrimination and detection: comparison of postadaptation thresholds. JOSA, 73(12), 1684-1690.

Regan, D., & Beverley, K. (1985). Postadaptation orientation discrimination. JOSA A, 2(2), 147-155.

Richters, D. P., & Eskew, R. T. (2009). Quantifying the effect of natural and arbitrary sensorimotor contingencies on chromatic judgments. Journal of Vision, 9(4), 27-27.

Rieke, F., & Rudd, M. E. (2009). The challenges natural images pose for visual adaptation. Neuron, 64(5), 605-616.

Roach, N. W., & McGraw, P. V. (2009). Dynamics of spatial distortions reveal multiple time scales of motion adaptation. Journal of Neurophysiology, 102(6), 3619-3626.

Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. Nature, 412(6846), 549-553.

Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. Nature, 412(6846), 549-553.

Schwartz, S., Maquet, P., & Frith, C. (2002). Neural correlates of perceptual learning: a functional MRI study of visual texture discrimination. Proceedings of the National Academy of Sciences, 99(26), 17137-17142.

Seitz, A. R., & Watanabe, T. (2003). Is subliminal learning really passive? Nature, 422(6927), 36-36.

Seitz, A. R., & Watanabe, T. (2009). The phenomenon of task-irrelevant perceptual learning. Vision Research, 49(21), 2604-2610.

Seriès, P., & Seitz, A. (2013). Learning what to expect (in visual perception). Frontiers in Human Neuroscience, 7, 668.

Seriès, P., Stocker, A. A., & Simoncelli, E. P. (2009). Is the homunculus “aware” of sensory adaptation? Neural Computation, 21(12), 3271-3304.

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. Journal of Neurophysiology, 86(4), 1916-1936.

Shiu, L. P., & Pashler, H. (1992). Improvement in line orientation discrimination is retinally local but dependent on cognitive set. Percept Psychophys, 52(5), 582-588.

Sireteanu, R., & Rettenbach, R. (1995). Perceptual learning in visual search: Fast, enduring, but non-specific. Vision Research, 35(14), 2037-2043.

Snyder, L. H., Batista, A., & Andersen, R. A. (1997). Coding of intention in the posterior parietal cortex. Nature, 386(6621), 167-170.

Snyder, L. H., Batista, A. P., & Andersen, R. A. (2000). Intention-related activity in the posterior parietal cortex: a review. Vision Research, 40(10-12), 1433-1441.

Sokolov, E. N. (1960). The neural model of the stimulus and the orienting reflex. Voprosy Psychologii, 4, 61-72.

Srinivasan, M. V., Laughlin, S. B., & Dubs, A. (1982). Predictive coding: a fresh view of inhibition in the retina. Proceedings of the Royal Society of London. Series B. Biological Sciences, 216(1205), 427-459.

Stocker, A. A., & Simoncelli, E. P. (2006). Sensory adaptation within a Bayesian framework for perception. Paper presented at the Advances in neural information processing systems.

Szpiro, S. F., & Carrasco, M. (2015). Exogenous attention enables perceptual learning. Psychological Science, 26(12), 1854-1862.

Szpiro, S. F., Wright, B. A., & Carrasco, M. (2014). Learning one task by interleaving practice with another task. Vision Research, 101, 118-124.

Tanaka, Y., Miyauchi, S., Misaki, M., & Tashiro, T. (2007). Mirror symmetrical transfer of perceptual learning by prism adaptation. Vision Research, 47(10), 1350-1361.

Tartaglia, E. M., Aberg, K. C., & Herzog, M. H. (2009). Perceptual learning and roving: Stimulus types and overlapping neural populations. Vision Research, 49(11), 1420-1427.

Teich, A. F., & Qian, N. (2003). Learning and adaptation in a recurrent model of V1 orientation selectivity. Journal of Neurophysiology, 89(4), 2086-2100.

Thompson, R. F. (2009). Habituation: a history. Neurobiology of Learning and Memory, 92(2), 127.

Tsakiris, M., & Haggard, P. (2003). Awareness of somatic events associated with a voluntary action. Experimental Brain Research, 149(4), 439-446.

Tsushima, Y., Seitz, A. R., & Watanabe, T. (2008). Task-irrelevant learning occurs only when the irrelevant feature is weak. Current Biology, 18(12), R516-R517.

Ungerleider, L., & Miskin, M. (1982). Two cortical visual systems In DJ Engle, MA Goodale, & RJ Mansfield (Eds.) Analysis of Visual Behavior, 549-586. In: Cambridge, MA: MIT Press.

Vaina, L. M., Sundareswaran, V., & Harris, J. G. (1995). Learning to ignore: psychophysics and computational modeling of fast learning of direction in noisy motion stimuli. Brain Res Cogn Brain Res, 2(3), 155-163.

Van Essen, D. C. (1979). Visual areas of the mammalian cerebral cortex. Annual Review of Neuroscience, 2(1), 227-261.

Verstraten, F. A., van der Smagt, M. J., Fredericksen, R. E., & van de Grind, W. A. (1999). Integration after adaptation to transparent motion: static and dynamic test patterns result in different aftereffect directions. Vision Research, 39(4), 803-810.

Verstraten, F. A., Van Der Smagt, M. J., & Van De Grind, W. A. (1998). Aftereffect of high-speed motion. Perception, 27(9), 1055-1066.

Vidnyánszky, Z., & Sohn, W. (2005). Learning to suppress task-irrelevant visual stimuli with attention. Vision Research, 45(6), 677-685.

Vogels, R., & Orban, G. A. (1985). The effect of practice on the oblique effect in line orientation judgments. Vision Research, 25(11), 1679-1687.

Wainwright, M. J. (1999). Visual adaptation as optimal information transmission. Vision Research, 39(23), 3960-3974.

Wang, R., Zhang, J.-Y., Klein, S. A., Levi, D. M., & Yu, C. (2012). Task relevancy and demand modulate double-training enabled transfer of perceptual learning. Vision Research, 61, 33-38.

Wark, B., Lundstrom, B. N., & Fairhall, A. (2007). Sensory adaptation. Current Opinion in Neurobiology, 17(4), 423-429.

Waszak, F., Cardoso-Leite, P., & Hughes, G. (2012). Action effect anticipation: neurophysiological basis and functional consequences. Neuroscience & Biobehavioral Reviews, 36(2), 943-959.

Watanabe, T., Náñez, J. E., & Sasaki, Y. (2001). Perceptual learning without perception. Nature, 413(6858), 844-848.

Webster, M. A. (2011). Adaptation and visual coding. Journal of Vision, 11(5), 3-3.

Webster, M. A. (2015). Visual adaptation. Annual Review of Vision Science, 1, 547-567.

Webster, M. A., & MacLeod, D. I. (2011). Visual adaptation and face perception. Philosophical Transactions of the Royal Society B: Biological Sciences, 366(1571), 1702-1725.

Webster, M. A., & Mollon, J. D. (1997). Adaptation and the color statistics of natural images. Vision research, 37(23), 3283-3298.

Weiss, Y., Edelman, S., & Fahle, M. (1993). Models of perceptual learning in vernier hyperacuity. Neural Computation, 5(5), 695-718.

Werner, J. S., & Schefrin, B. E. (1993). Loci of achromatic points throughout the life span. JOSA A, 10(7), 1509-1516.

Wexler, M. (2018). Multidimensional internal dynamics underlying the perception of motion. Journal of Vision, 18(5), 7-7.

Wexler, M., Duyck, M., & Mamassian, P. (2015). Persistent states in vision break universality and time invariance. Proceedings of the National Academy of Sciences, 112(48), 14990-14995.

Wissig, S. C., Patterson, C. A., & Kohn, A. (2013). Adaptation improves performance on a visual search task. Journal of Vision, 13(2), 6-6.

Xiao, L.-Q., Zhang, J.-Y., Wang, R., Klein, S. A., Levi, D. M., & Yu, C. (2008). Complete transfer of perceptual learning across retinal locations enabled by double training. Current Biology, 18(24), 1922-1926.

Yehezkel, O., Sagi, D., Sterkin, A., Belkin, M., & Polat, U. (2010). Learning to adapt: Dynamics of readaptation to geometrical distortions. Vision Research, 50(16), 1550-1558.

Yu, C., Klein, S. A., & Levi, D. M. (2004). Perceptual learning in contrast discrimination and the (minimal) role of context. Journal of Vision, 4(3), 4-4.

Zhang, J.-Y., & Yu, C. (2016). The transfer of motion direction learning to an opposite direction enabled by double training: A reply to Liang et al.(2015). Journal of Vision, 16(3), 29-29.

Zhang, J. Y., Kuai, S. G., Xiao, L. Q., Klein, S. A., Levi, D. M., & Yu, C. (2008). Stimulus coding rules for perceptual learning. PLoS Biol, 6(8), e197.

Zhang, P., Bao, M., Kwon, M., He, S., & Engel, S. A. (2009). Effects of orientation-specific visual deprivation induced with altered reality. Current Biology, 19(22), 1956-1960.

Zhang, T., Xiao, L.-Q., Klein, S. A., Levi, D. M., & Yu, C. (2010). Decoupling location specificity from perceptual learning of orientation discrimination. Vision Research, 50(4), 368-374.