THE TIME COURSE OF VISUAL ADAPTATION

Sarah Maddison

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

October 2018

Abstract

Visual neurons persistently alter their operating properties in response to changes in the statistics of the visual environment – a process termed adaptation. The prevalence of adaptation has led it to be considered a universal law of visual processing.

Neural adaptation produces measurable perceptual aftereffects. One such aftereffect is the tilt aftereffect; whereby prolonged viewing of a tilted stimulus causes subsequently viewed stimuli to appear tilted away from the adapting stimulus orientation. This aftereffect is used as a proxy for adaptation throughout this thesis. Longer adapting durations result in larger perceptual aftereffects. These aftereffects show a characteristic time course of logarithmic growth and decay, however little is known about whether this time course is fixed or whether it can be altered by previous visual experience.

To measure changes in the time course of adaptation, it must be measured with high temporal resolution. Traditionally, adaptation growth is measured using multiple discrete periods of continuous adaptation, increasing in duration. This method is time consuming and has low temporal resolution. Therefore, we explored the use of a single period of adaptation that was briefly interrupted by the measurement of adaptation magnitude – we refer to this as 'interrupted adaptation'.

Interrupted adaptation was shown to be a valid high temporal resolution measure of the time course of the tilt aftereffect. However, overall adaptation magnitude was suppressed using this measure when compared to the same duration of adaptation measured with tradition continuous adaptation. We explored the effects of passive recovery, attention, and memory on the suppression of the tilt aftereffect, but found suppression was produced by passive recovery alone. The magnitude of the suppression

could not be accounted for using a simple model of adaptation with a fixed rate of growth and decay, suggesting rates of growth and decay may differ.

Once a valid high temporal resolution measure of the time course had been achieved, we explored whether the time course of visual adaptation could be altered by previous visual experience. We hypothesised that the rate of adaptation could be determined by the rate of change recently encountered in the environment. Observers were exposed to temporally high-pass and low-pass filtered videos, followed by the measurement of the time course of the tilt aftereffect. Exposure to predominantly high temporal frequencies resulted in the increased rate and magnitude of tilt aftereffect growth, consistent with the hypothesis that the rate of adaptation growth reflects the rate of recently encountered change. Exposure to predominantly low temporal frequencies had negligible effects on the time course of the tilt aftereffect, likely due to their prevalence in natural statistics.

The original finding of increased rate and magnitude of adaptation following exposure to fast temporal statistics failed to be replicated using high temporal frequency contrast-modulated artificial stimuli. This led to the hypothesis that differences in the average contrast of the filtered videos, rather than differences in temporal statistics, drove the original effect. In line with this hypothesis, increased rate and magnitude of the tilt aftereffect was observed following exposure to a temporally unfiltered video with contrast matched to the original high-pass filtered video. Further exploration revealed that exposure to videos with increased contrast relative to natural scenes leads to a suppression of the tilt aftereffect in line with a response gain control mechanism.

Throughout this thesis the mechanisms underlying the time course of adaptation were explored through modelling. Consistent with previous research, many of the empirical findings are well captured by a two-mechanism leaky integrator model. However, further work is needed to identify the precise nature and number of mechanisms driving adaptation across a wide range of timescales.

Acknowledgements

First and foremost, I would like to thank my supervisor Ben Webb for being not only a fantastic supervisor full of patience, direction, and support, but also for being a great friend. I have learnt a great deal from you as an academic and as a person and couldn't have done it without you. So here it is: '**my** PhD'.

I would also like to say thank you to my secondary supervisor Neil Roach, whose happy demeanour and sense of humour made supervision meetings a genuine pleasure. Your guidance throughout the last three years was fundamental in producing this thesis and your patience when explaining the more complex mathematical components was admirable.

A big thank you to my wonderful friends and family. Your love and encouragement kept me going throughout good times and bad, I'm so lucky to have you all in my life. A specific shout out to the motley crew that is the Visual Neuroscience Group - it's not often that you can switch between Fourier Transforms and Great British Bake Off in the same conversation.

A special thank you goes to Luke Holden for being there to remind me that there is more to life than a PhD. Your love, support, and reality checks made writing this thesis a more pleasant experience.

Last but definitely not least, thank you to Becca. You've been there through it all. Thank you for being my biggest cheerleader, the best shoulder to cry on, and the best friend anyone could ask for. Haven't we learnt a lot.

Table of contents

1. Introduction 1
1.1. VISUAL ADAPTATION 1
1.1.1. Definition and prevalence throughout the visual hierarchy 1
1.1.2. Perceptual Consequences 9
1.1.3. Timescales 14
1.1.4. Modelling adaptation 20
1.1.5. Function 23
1.2. THE TILT AFTEREFFECT 26
1.2.1. Definition 26
1.2.2. Tuning 27
1.2.3. Psychophysical measurement 29
1.2.4. Physiology 33
1.2.5. Time course 34
1.3. VISUAL PROCESSING IN NATURAL SCENES 37
1.3.1. Statistical structure of natural scenes 37
1.3.2. Visual processing and natural scenes 40
1.3.3. A case for using naturalistic stimuli to study adaptation 44
1.3.4. Adaptation to naturalistic stimuli 46
1.4. INTRODUCTION TO THE EXPERIMENTS IN THIS THESIS 49

2. General methods 50
2.1. APPARATUS 50
2.1.1. Monitor luminance calibration 50
2.1.2. Monitor temporal precision 52
2.2. STIMULI 53
2.3. REPRODUCTION TASK 54
2.4. MEASURING THE TIME COURSE OF THE TILT AFTEREFFECT 55
2.5. ANALYSIS 56
2.5.1. Outlier removal 57
2.5.2. Descriptive statistics 57
2.5.3. Inferential statistics – permutation testing 59

3. Biases in the reproduction of orientations 60

3.1. INTRODUCTION 60

3.2. METHODS 64

3.3. RESULTS 65

3.4. MODELLING BIASES IN THE REPRODUCTION OF ORIENTATIONS 69

3.5. INDIVIDUAL DIFFERENCES OF BIASES IN THE REPRODUCTION OF ORIENTATIONS 71

3.6. DISCUSSION 73

4. Interrupted measurement of the tilt aftereffect 77

4.1. INTRODUCTION 77

4.2. EXPERIMENT 1: INTERRUPTED ADAPTATION 79

4.2.1. Methods 79

4.2.2. Results 80

4.3. EXPERIMENT 2: TYPE OF INTERRUPTION 83

4.3.1. Methods 84

4.3.2. Results 87

4.3.3. Modelling the growth of the tilt aftereffect 88

4.4. DISCUSSION 90

5. The time course of the tilt aftereffect following exposure to altered temporal frequency content of naturalistic videos 95

5.1. INTRODUCTION 95

5.2. METHODS 97

5.2.1. Temporal frequency filtered videos 97

5.2.2. Data analysis 98

5.3. RESULTS 100

5.4. MODELLING THE TIME COURSE OF THE TILT AFTEREFFECT 102

5.5. DISCUSSION 106

6. The time course of the tilt aftereffect following exposure to altered temporal frequency content of artificial stimuli 110

6.1. EXPERIMENT 1 110

6.1.1. Introduction 110

6.1.2. Methods 111

6.1.3. Results 115

6.2. EXPERIMENT 2 117

6.2.1. Introduction 117

6.2.2. Methods 119

6.2.3. Results 125

6.3. DISCUSSION 131

7. The time course of the tilt aftereffect following exposure to altered RMS contrast 136

7.1. INTRODUCTION 136

7.2. METHODS 139

7.2.1. Stimuli 139

7.2.2. Data analysis 140

7.3. RESULTS 142

7.4. MODELLING THE TIME COURSE OF THE TILT AFTEREFFECT 144

7.5. DISCUSSION 146

8. General discussion 150

8.1. SUMMARY OF RESULTS 150

8.2. WHAT IS AN APPROPRIATE METHODOLOGY TO MEASURE THE TIME COURSE OF THE TILT AFTEREFFECT WITH FINE TEMPORAL RESOLUTION 152

8.2.1. Future directions 156

8.3. HOW IS THE TIME COURSE OF THE TILT AFTEREFFECT, SPECIFICALLY THE RATE OF ADAPTATION, AFFECTED BY PRIOR EXPOSURE TO ALTERED TEMPORAL ENVIRONMENTS? 157

8.3.1. Future directions 162

8.4. WHAT ARE THE CHARACTERISTICS OF THE MECHANISM(S) THAT GOVERN THE GROWTH AND DECAY OF THE TILT AFTEREFFECT?.....164

8.4.1. Future directions.....166

8.5. CONCLUDING REMARKS 167

References 169

1

Introduction

1.1. VISUAL ADAPTATION

1.1.1. Definition and prevalence throughout the visual hierarchy

Neurons in the visual system are forced to persistently alter the way they encode visual information in response to changes to the statistical structure of the visual environment. This process is termed visual adaptation and is a ubiquitous coding strategy that adjusts how we see the world. The study of visual adaptation has long been a central tenet of vision science. In part, this is due to the use of adaptation as a method to identify visual processing mechanisms. Adaptation is thus commonly referred to as 'the psychophysicist's electrode' with the idea that 'if it adapts, it's there'. This use of adaptation as a tool has led to an increased understanding of visual processing mechanisms. The sheer prevalence of adaptation effects discovered in this way revealed that the processes of adaptation were in fact fundamental to the way visual information is encoded and decoded throughout the visual hierarchy, such that 'if it's there, it adapts'. Adaptation is suggested to be a universal law (Helson, 1964), operating in a functionally similar way across multiple stimulus domains. The effects of adaptation can be seen throughout the visual stream, showing plasticity for the type of visual information that those neural circuits are designed to represent. Here, I discuss the physiological underpinnings of adaptation for a select few stimulus properties that are processed at different levels of the visual hierarchy (see Figure 1.1). How the physiological effects of visual adaptation result in the perceptual consequences measured psychophysically is discussed in Section 1.1.2.

Earliest in the visual hierarchy (see Figure 1.1) exists a form of adaptation without which we would be left blind most of the time: retinal light adaptation (Shapley & Enroth-cugell, 1984). Retinal ganglion neurons are faced with light intensity that varies over 10 orders of magnitude (Frazor & Geisler, 2006). Neurons in the retina maintain responsivity across this entire range using two strategies. Firstly, retinal ganglion cells adapt to the mean light intensity, which is termed 'light adaptation' and mostly involves photoreceptor properties, switching between rods and cones, and switching post-receptoral circuits (Spillman & Werner, 1990). Light adaptation allows the visual system to maintain sensitivity when switching between environments with vastly different mean luminance. Secondly, as we scan a scene, the retina is also faced with different ranges of intensities

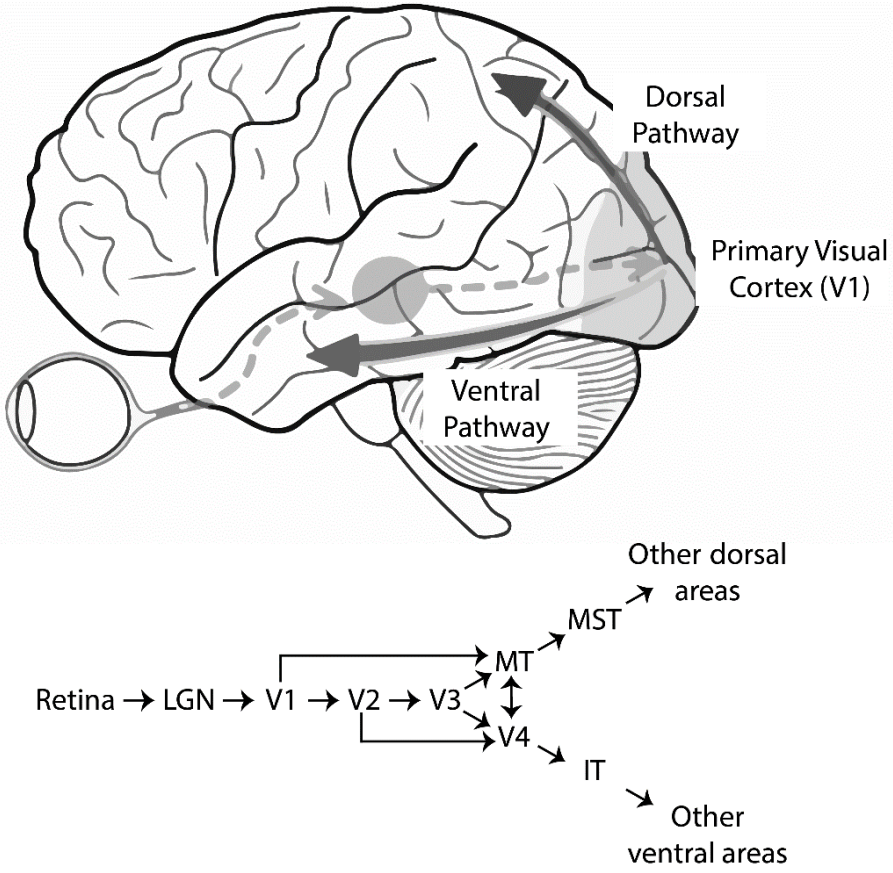


Figure 1.1 Hierarchical organization of the visual system. Top: Dashed grey line shows flow of information from retina to LGN (grey circle) to primary visual cortex (shaded area). From here, processing continues in dorsal and ventral pathways. Bottom: The connections between early visual areas, showing the dorsal and ventral processing streams. For clarity, only key neural structures and connections are shown. Adapted from Kafaligonal (2014).

around the mean (i.e. different contrasts). Retinal ganglion cells therefore adapt to the range of intensities in a process known as 'contrast adaptation' (Bonin, Mante, & Carandini, 2006; Mante, Frazor, Bonin, Geisler, & Carandini, 2005; Shapley & Victor, 1978). Contrast adaptation serves two purposes: when input signals are weak, adaptation serves to increase neuronal sensitivity and improve the signal to noise ratio; and when input signals are strong, adaptation serves to decrease neuronal sensitivity preventing response saturation and information loss (Gaudry & Reinagel, 2007; Kohn, 2007). Switching from a low-contrast environment to a high-contrast environment therefore quickly results in the reduced gain of spiking responses in retinal ganglion cells, preventing response saturation (Smirnakis, Berry, Warland, Bialek, & Meister, 1997). Functionally, this allows discrimination of relatively small changes in contrast regardless of the prevailing light intensity, which is fundamental in everyday vision. This is the earliest locus of contrast adaptation, with further adaptation occurring in downstream areas of the visual system such as V1.

Progressing along the visual stream, there exist a considerable number of adaptation mechanisms in primary visual cortex (V1). Neurons in V1 are tuned for simple stimulus attributes such as contrast, spatial frequency, and

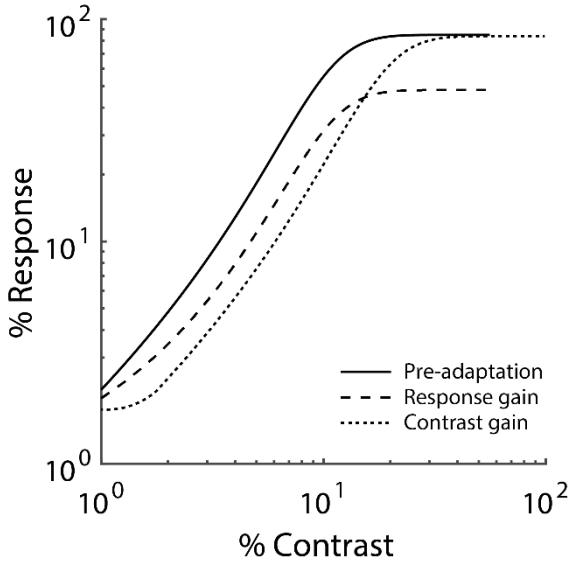


Figure 1.2. Example contrast response function from a V1 neuron in a cat. The response increases linearly as the stimulus contrast increases before the slope rapidly declines. Adaptation to contrast causes both vertical (response gain – dashed line) and horizontal (contrast gain – dotted line) shifts in the function. Adapted from Albrecht & Hamilton (1982).

orientation. Neurons in V1 do not respond equally across all contrasts, rather the response curve is sigmoid-shaped (see Figure 1.2). As the contrast of a pattern increases, a V1 neuron's response increases in a steep linear fashion, covering only a small contrast range, before the slope of the function rapidly declines (Albrecht & Hamilton, 1982). Much like retinal ganglion cells, the dynamic range of V1 neurons is not wide enough to respond across the entire range of contrasts, which without adaptation would lead to response saturation and information loss. Adaptation has been shown to horizontally shift the contrast response function so it is centred on the adapting stimulus contrast (Durant, Clifford, Crowder, Price, & Ibbotson, 2007; Gardner et al., 2005; Ohzawa, Sclar, & Freeman, 1982). This form of contrast gain control is functionally beneficial in allowing neurons to respond optimally to the current environmental contrast. Exposure to a high-contrast environment also results in the overall reduced firing rate of contrast sensitive neurons (Crowder, Hietanen, Price, Clifford, & Ibbotson, 2008; Vautin & Berkley, 1977), known as 'response gain control', which is advantageous in preventing response saturation.

V1 neurons are also tightly tuned to spatial frequency, showing a distinct peak in responsivity to a specific spatial frequency and a sharp drop off in responsivity for higher and lower spatial frequencies (Foster, Gaska, Nagler, & Pollen, 1985). This is due to the elongated receptive field structure of simple cells in V1 that have multiple on and off regions that are excited and inhibited by strips of light (i.e. light bars of a grating stimulus) respectively. A cell will be maximally stimulated by a grating with light and dark stripes that match the on-off regions of its receptive field (Movshon, Thompson, & Tolhurst, 1978). Therefore, a larger number of on-off regions will result in a cell tuned to high spatial frequencies (DeValois & DeValois, 1988). Physiological evidence suggests the existence of a number of spatial frequency selective channels that pool the responses from V1 neurons tuned to similar spatial frequencies (Campbell, Cooper, & Enroth-Cugell, 1969; Enroth-Cugell & Robson, 1966). Selective adaptation to spatial frequency was first shown by Movshon & Lennie (1979) who showed that

adaptation to a grating with specific spatial frequency resulted in a reduced response to gratings of the adapting spatial frequency, but not to gratings of differing spatial frequency.

V1 neurons also show selectivity for temporal frequency. Foster, Gaska, Nagler, & Pollen (1985) identified three types of tuning curve in macaque monkey area V1 and V2. The first were considered band-pass, whereby the neuron showed a distinct peak for a specific temporal frequency and attenuated responses for temporal frequencies above or below the preferred temporal frequency. The second showed a peak for a specific temporal frequency, but a shallower drop-off for lower temporal frequencies. And the final type were flat, low-pass tuning curves, responding maximally for low temporal frequencies with a sharp drop-off in responsivity to higher temporal frequencies. The majority of V1 cells were shown to exhibit the latter two types of temporal frequency tuning curves, suggesting low-pass tuning specificity. Again, selective adaptation to temporal frequency results in the reduced responsivity of neurons tuned to the adapting temporal frequency, minimising response saturation (Saul & Cynader, 1989).

The experiments in this thesis focus on the effects of adaptation to orientation, thus a more detailed explanation of the physiology of orientation tuning and adaptation is given here. Neurons in V1 show tight tuning to stimulus orientation, again due to their elongated receptive fields (Hubel & Wiesel, 1962). Orientation selectivity is described by tuning curves that show neuronal response rate as a function of stimulus angle, these curves are bell-shaped with the largest response at the preferred orientation and responsivity sharply dropping off for orientations either side (see Figure 1.3 (a)).

Psychophysical reports of altered perception of tilt following prolonged viewing of a tilted stimulus (Gibson & Radner, 1937; Greenlee & Magnussen, 1987; Magnussen & Johnsen, 1986) suggested that the orientation tuning of

V1 neurons could be altered by the orientation of recently encountered stimuli. This was initially reported as a reduced firing rate of neurons similarly tuned to the adapting orientation (Blakemore & Campbell, 1969; Carandini, Anthony Movshon, & Ferster, 1998; Hammond, Pomfrett, & Ahmed, 1989; Müller, Metha, Krauskopf, & Lennie, 1999; Nelson, 1991), however more complex changes in individual neurons' tuning curves following adaptation have since been revealed.

Dragoi, Sharma, & Sur (2000b) measured changes in orientation tuning curves following prolonged adaptation to a particular stimulus orientation using pattern adaptation (Movshon & Lennie, 1979). Their results showed an interesting dynamic plasticity whereby 2 minutes of exposure to a tilted stimulus ~15-20deg away from the neuron's preferred orientation resulted in responses on the flank of the tuning curve nearest the adapting orientation being depressed, whilst responses on the opposite flank were enhanced (see Figure 1.3b). This re-organisation shifted the preferred orientation of the neuron away from the adapting stimulus, with the magnitude of this shift dependent on the difference between the adapting orientation and the neuron's preferred orientation and being maximal for differences ~15-20deg. Conversely, adaptation to a stimulus with orientation 60-90deg away from the neuron's preferred orientation had little effect on the neuron's tuning, showing the orientation specificity of selective adaptation. The authors suggest that the reorganization of the neuron's tuning curve following prolonged adaptation may result in improved detection of stimuli oriented near the neuron's new preferred orientation and may help to improve discriminability in the face of prolonged exposure to the statistics of natural images (Coppola, Purves, Mccoy, & Purves, 1998; Whitaker & McGraw, 2000).

Subsequent studies have also looked at how adaptation effects the response of orientation-tuned V1 neurons at the population level rather than single cells. Adaptation to an oriented stimulus results in a reduction in firing rates of neurons with tuning curves near the adapting orientation, a shift of close-by tuning curve peaks away from the adapting orientation (in support of

(Dragoi, Sharma, & Sur, 2000)), and an increase in firing rate of neurons with tuning curves far from the adapting orientation (Figure 1.3 (c)) (Valentin Dragoi, Sharma, & Sur, 2000a; Teich et al., 2003). The perceptual consequences of the alteration of orientation tuning curves are of direct relevance to this thesis and will be discussed in detail in section 1.2.4.

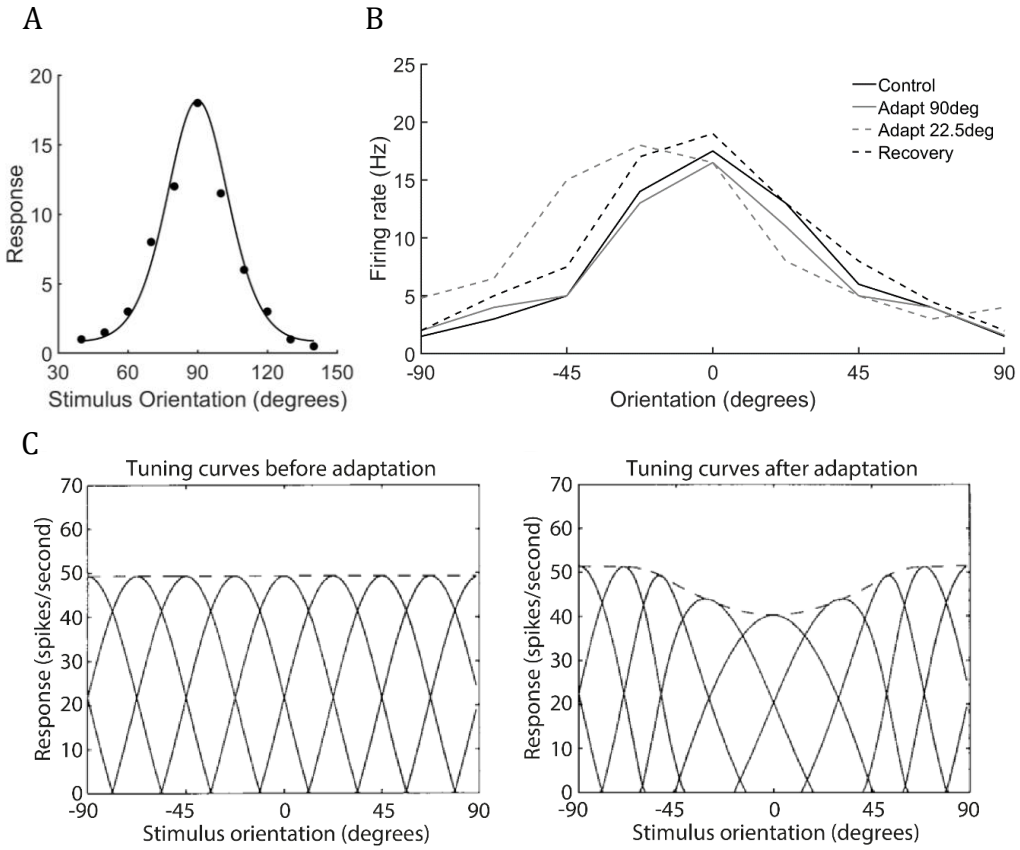


Figure 1.3 (a) a measured orientation tuning curve from a cat V1 neuron with vertical preferred orientation. Adapted from Swindale (1998) (b) Cat V1 neuron's orientation tuning curves pre-adaptation (solid black), following adaptation to a 90deg oriented stimulus (solid grey), following adaptation to a 22.5deg oriented stimulus (dashed grey), and following adaptation recovery (dashed black). Adaptation to a 90deg oriented stimulus had little effect on the neuron tuning curve. Adaptation to a 22.5deg oriented stimulus resulted in increased responses on the left flank of the tuning curve (closest to the adapting orientation) and decreased responses on the opposite flank. Following adaptation offset the tuning curve begins to return to pre-adaptation state. Adapted from Dragoi, Sharma & Sur (2000). (c) simulated orientation tuning curves pre (left) and post (right) adaptation to a vertically oriented stimulus. Adapted from Teich (2003).

Primary visual cortex also contains large populations of colour-selective neurons. Three types of neural pathways have been discovered that represent colour, a red-green pathway, a blue-yellow pathway, and a luminance pathway. However, physiological experiments show mixed results in finding neurons with response properties reflecting these pathways (Lennie, Krauskopf, & Sclar, 1990; Livingstone & Hubel, 1984; Thorell, Devalois, & Albrecht, 1984; Vautin & Dow, 1985). fMRI evidence suggests that colour signals relevant for perception are encoded by large populations of neurons (Engel, Zhang, & Wandell, 1997) and selective adaptation through the prolonged presentation of a coloured stimulus is known to reduce the firing rate of neurons tuned to the adapting colour, sparing those tuned to other colours (Engel & Furmanski, 2001). The experiments in this thesis use only achromatic stimuli, but for a more thorough explanation of colour adaptation see Fairchild & Lennie (1992) and Rinner & Gegenfurtner (2000).

More recent research has studied adaptation at higher levels of the visual hierarchy. Adaptation to contrast is shown in area V2 in a similar manner to that in V1 (Crowder, 2005). Adaptation to colour, in terms of maintaining a stable percept of colour irrespective of lighting conditions, is observed in V4 (Walsh, 1995). And adaptation to direction of motion is found in area MT, with prolonged viewing of motion stimulus causing reduced responsiveness in neurons tuned to that direction of motion (Huk, Ress, & Heeger, 2001; Tolias, Smirnakis, Augath, Trinath, & Logothetis, 2001). However, this is suggested to be due to inheritance of adaptation-induced contrast gain changes at earlier levels (Kohn & Movshon, 2003). Adaptation is also seen to more complex, abstract stimulus attributes at even higher levels of visual cortex. Thus far, adaptation has been shown to faces (including identity, gender, expression, and ethnicity - Leopold, O'Toole, Vetter, & Blanz, 2001; Webster & MacLeod, 2011), biological motion (Jordan, Fallah, & Stoner, 2006; Troje & Westhoff, 2006), surface texture (Motoyoshi, Nishida, Sharan, & Adelson, 2007), illusory contours (Paradiso,

Shimojo, & Nakayama, 1989), and temporal synchrony (Roseboom, Linares, & Nishida, 2015) to name but a few.

The prevalence of visual adaptation throughout the hierarchy suggests that it is a fundamental process in determining how and what we see. A greater understanding of the perceptual consequences, timescales, underlying mechanisms, and functional benefits of adaptation will shed light on visual processing as a whole.

1.1.2. Perceptual Consequences

The physiological changes induced by adaptation result in two main perceptual consequences – altering detection and discrimination thresholds and inducing biases in the perception of subsequent stimuli.

Detection thresholds refer to the lowest intensity at which the presence of a stimulus can be detected. For example, the lowest contrast of a pattern before you cannot detect its presence. Discrimination thresholds refer to the smallest perceivable difference in a stimulus dimension, sometimes known as the just noticeable difference (JND). For example, if presented with two patterns, what is the smallest difference in their orientation/contrast/spatial frequency/colour you can perceive before they look identical.

Physiologically, detection and discrimination performance are suggested to be driven by different subsets of neurons (Hol & Treue, 2001; Regan & Beverley, 1983, 1985). For example, detection of stimulus direction depends on the signal to noise ratio across the population of direction tuned neurons. Noise leads to an even activation of all neurons, whilst the presented stimulus results in an increase in activation of neurons tuned to that direction. A reduction in the amplitude of neurons tuned to the presented orientation (through adaptation) will therefore result in the

greatest increase in detection thresholds. In comparison, discrimination depends on the greatest response change between the two presented stimuli, which is shown between neurons tuned slightly away from each stimulus direction. Selectively reducing the responsiveness of off-axis neurons will therefore result in reduced discrimination performance. Hol & Treue (2001) revealed that adapting along the test direction was most effective in increasing detection thresholds and had little effect on discrimination performance, whilst adapting ~50deg away from the test orientation was most effective in reducing discrimination performance. These effects are shown to be paralleled for the effects of orientation adaptation on orientation detection and discrimination performance (Clifford, 2002; Regan & Beverley, 1985).

Adaptation-induced changes in detection and discrimination performance are seen throughout the visual hierarchy (Phinney, Bowd, & Patterson, 1997). For example, prolonged adaptation to a high contrast stimulus increases the contrast needed for detection of stimuli with similar contrast to the adapter (Blakemore & Campbell, 1969; Pantle & Sekuler, 1968), and increases the difference in contrast needed to discriminate stimuli with similar, but not identical, orientation and spatial frequency (Greenlee & Heitger, 1988). The same pattern of impaired detection and discrimination is seen across stimulus domains such as orientation (Regan & Beverley, 1985), spatial frequency (Regan & Beverley, 1983), and colour discrimination (Krauskopf & Karl, 1992).

The second consequence of adaptation is perceptual bias for subsequently presented stimuli. Adaptation-induced perceptual biases are usually repulsive in nature, with the stimulus appearing repelled away from the adapting dimension (see Figure 1.4). Repulsive biases of this nature are seen for simple properties such as orientation (Gibson & Radner, 1937), spatial frequency (Colin Blakemore & Sutton, 1969), colour (Webster & Mollon, 1997), and motion (Anstis, Verstraten, & Mather, 1998), as well as higher level properties such as face identity (Leopold, Rhodes, Müller, &

Jeffery, 2005) and gender (Webster, Kaping, Mizukami, & Duhamel, 2004), gaits implied by biological motion (Troje, Sadr, Geyer, & Nakayama, 2006), the openness of natural scenes (Greene & Oliva, 2010), and many more (see Webster, 2015).

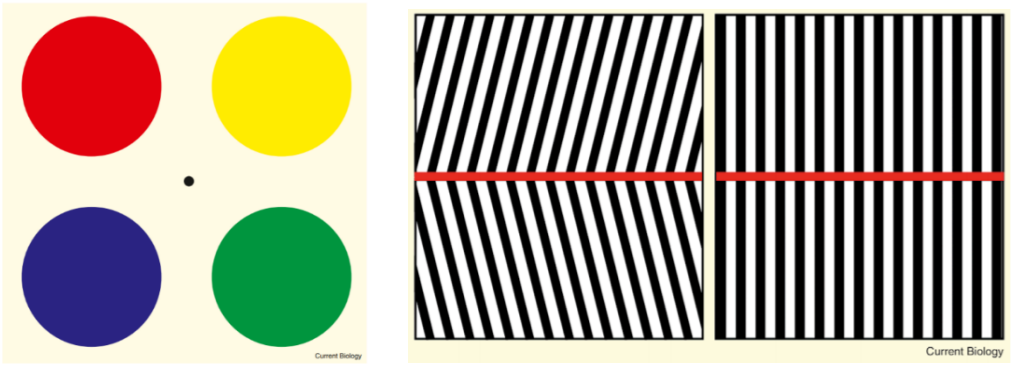


Figure 1.4 Adaptation induced perceptual biases. Left: Colour aftereffects. Staring at the dot for 30 seconds then looking at a plain surface will result in the complementary colours being seen. The red ang green dots with appear to have switched places, as will the blue and yellow. Right: The tilt aftereffect. Moving ones eyes along the red line of the left pattern for 30 seconds then looking at the right pattern will result in the lines of the right pattern appearing titled in the opposite direction to the left pattern. Both taken from Thompson & Burr (2009).

The physiological basis for repulsive perceptual biases is the reduction in responsiveness of neurons tuned to the adapting stimulus property, creating an asymmetrical bias in the population response (Blakemore & Sutton, 1969; Clifford, Wenderoth, & Spehar, 2000; Coltheart, 1971). The physiological underpinnings of the perceptual effects of contrast, spatial frequency, temporal frequency, and orientation adaptation are discussed below. It is important to note that there are more complex physiological consequences of adaptation on neuronal tuning curves beyond simple amplitude suppression that affect the perception of subsequent stimuli. These will be discussed in detail for orientation adaptation in Section 1.2.4 as it is specifically relevant to the experiments in this thesis.

Prolonged viewing of a high contrast stimulus will reduce the perceived contrast of subsequently presented stimuli (Blakemore, Muncey, & Ridley, 1973; Georgeson, 1985). This effect is specific to the spatial frequency of the adapting stimulus. At a single contrast, stimuli with mid-range spatial

frequencies are most perceptible (4-6cpd), with very low and very high spatial frequencies being harder to detect – forming the contrast sensitivity function (Figure 1.5) (Campbell et al., 1969; Campbell & Robson, 1968). Physiologically, selective adaptation to a high contrast pattern with a specific spatial frequency results in a reduction in responsivity of neurons tuned to the adapting spatial frequency, and a change in shape of the CSF (Figure 1.5(a)). Perceptually, this results in reduced perceived contrast of subsequent stimuli and increased contrast detection thresholds for stimuli similar to the adapting stimulus (Sachs, Nachmias, & Robson, 1971). Adaptation at a specific spatial frequency will only affect the perception of stimuli with similar spatial frequency to the adapter, the narrow specificity of the adaptation effect led to the discovery that there exist distinct channels for the processing of different spatial frequencies (Blakemore & Campbell, 1969).

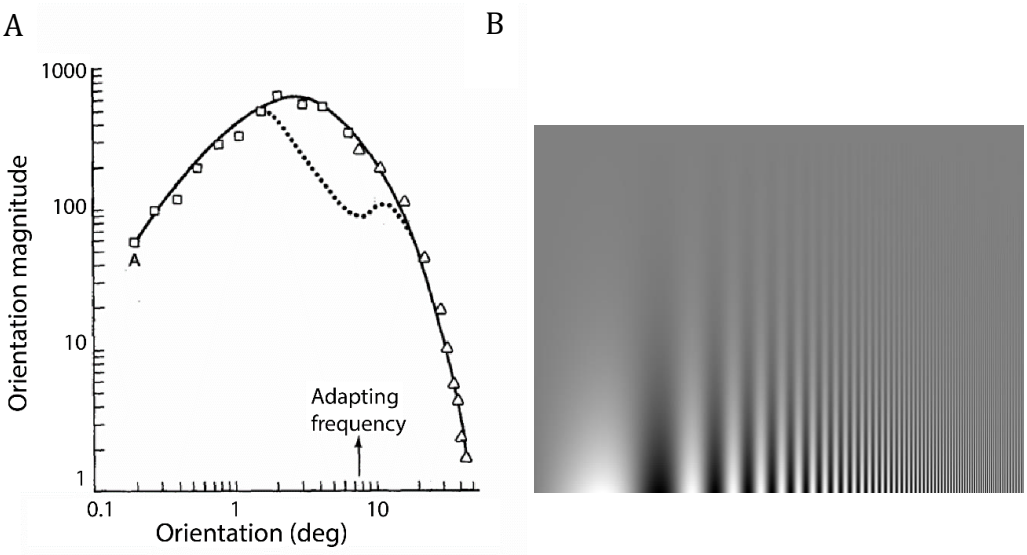


Figure 1.5 (a) The contrast sensitivity function pre (solid line) and post (dotted line) adaptation. Adapted from Campbell & Robson (1968). (b) image that decreases in contrast moving bottom to top and increases in spatial frequency moving left to right. The observer should notice the visible space has the same shape as the pre-adaptation CSF shown in (a).

Similarly, selective adaptation to a high temporal frequency stimulus results in subsequently presented stimuli appearing lower temporal frequency and elevated temporal contrast detection thresholds for similar stimuli (Smith, 1971). This perceptual effect is the result of the reduced

responsivity of neurons tuned to high-temporal frequencies, creating a bias in the population code. The responses of neurons tuned to temporal frequency are pooled over a wider range than spatial frequency, therefore adapting to a specific temporal frequency will affect the perception of a wider range of test stimulus temporal frequencies. Figure 1.6 (right) shows the magnitude of adaptation (increase in temporal contrast threshold) at different test frequencies after adaptation to 15c/s. Although the largest adaptation effect is for stimuli with similar temporal frequencies as the adapter, this tuning is quite broad in comparison to spatial frequency tuning.

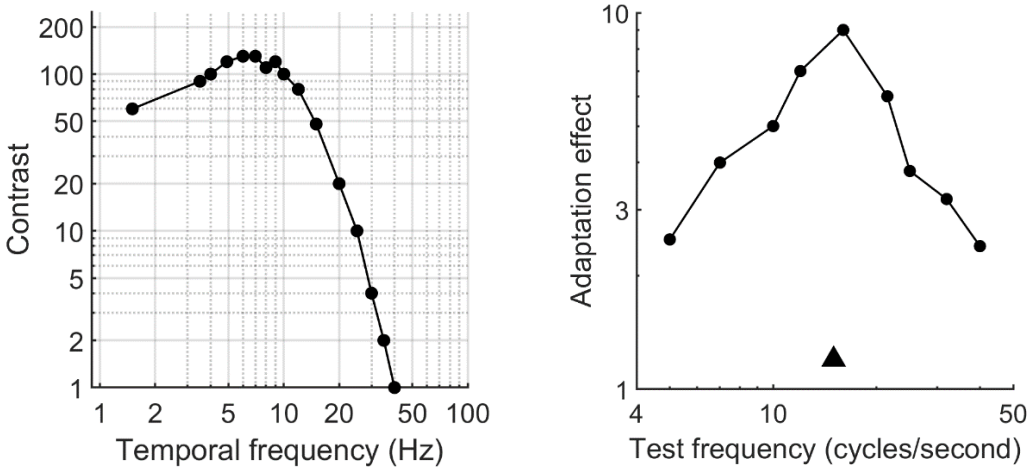


Figure 1.6. Left: The human temporal contrast sensitivity function. Adapted from Wandell (1995). Right: The magnitude of adaptation (increase in temporal contrast threshold) at different test frequencies after adaptation to 15c/s. Although the largest adaptation effect is for stimuli with similar temporal frequencies as the adapter, this tuning is quite broad. Adapted from Smith (1971).

As well as repulsive biases, adaptation sometimes results in attractive shifts in perception, with stimuli appearing more like the adapting stimulus. These effects are said to be related to active prediction of the world and were first shown in the priming literature. For example, viewing a moving pattern that is alternating in direction causes a subsequent ambiguous test pattern to appear to be moving in the direction that would complete the alternation (Maloney, Martello, Sahm, & Spillmann, 2005).

These positive aftereffects have been recently conceptualised as 'serial dependence'. This was initially shown for orientation, whereby an oriented test stimulus appeared biased towards the orientation of a previously presented stimulus - the effect is maximal if the previous orientation was close to that of the current (Fischer & Whitney, 2014). Serial dependence effects have now been shown for stimulus domains such as face identity (Liberman, Fischer, & Whitney, 2014), attractiveness (Xia et al., 2016), and emotion (Liberman & Whitney, 2015) as well as biological motion (Chaney, Liberman, & Whitney, 2016) and object identity (Liberman, Zhang, & Whitney, 2016). These effects are proposed to be the result of the visual system exploiting the temporal redundancies in natural scenes, with most properties remaining stable over time with the best prediction of the future being the current state of the environment (Burr & Cicchini, 2014).

An important distinction between the attractive effects of serial dependence and negative effects of adaptation is the timescales over which they occur. Serial dependence only occurs at sub-second timescales, and is likely a different form of temporal contextual effect driven by different mechanisms to adaptation (Burr & Cicchini, 2014). The focus of this thesis is on adaptation that occurs at longer, and a larger range, of timescales

1.1.3. Timescales

Changes in our visual environment occur over timescales ranging from milliseconds to millennia. Rapid changes occur with fixational eye movements (Martinez-Conde, Macknik, & Hubel, 2004); medium term changes result from moving between environments e.g. inside to outside; and long term changes occur over the course of the development of the visual system (Banks, 1988). However, traditional studies of adaptation focused on timescales from seconds to minutes.

Over these timescales, the magnitude of adaptation aftereffects increases as a logarithmic function of time, and decays with a similar profile following the offset of the adapting stimulus (Greenlee, Georgeson, Magnussen, & Harris, 1991). This classical time course of adaptation growth and decay is consistent across stimulus domains, suggesting common adapting mechanisms (Leopold et al., 2005). The precise dynamics of adaptation decay often depend on the stimulus input following adaptation offset. In some cases, adaptation aftereffects do not passively decay, but are 'stored' until the presentation of new sensory information (Machulla, Di Luca, Froehlich, & Ernst, 2012; Spigel, 1960), which in rare cases leads to permanent adaptation if the appropriate de-adapting stimulus is not presented (Vul, Krizay, & MacLeod, 2008). However, this is not the case for some stimulus domains, such as adaptation to orientation which shows no storage (Magnussen & Johnsen, 1986).

More recent work has expanded the study of adaptation over numerous timescales. Much like the prevalence of adaptation across stimulus domains, it appears adaptation operates over a similarly broad range of timescales. Adaptation occurring at very short (<1 second) timescales highlights the role of adaptation not as a mere consequence of overstimulation, but as a continually active process altering the way we see. Sub-second adaptation has been demonstrated for motion (Roach & McGraw, 2009), contrast (Pavan, Marotti, & Campana, 2012), orientation (Sekuler & Littlejohn, 1974), contour orientation (Suzuki, 2001), and spatial and temporal multisensory timing (Roach, Heron, Whitaker, & McGraw, 2010; Van der Burg, Alais, & Cass, 2013; Wozny & Shams, 2011).

Adaptation has also been studied at durations spanning hours or days. For example, prolonged exposure to a coloured bias induces aftereffects that exceed the duration of traditional short-term chromatic adaptation (Belmore & Shevell, 2011; Eisner & Enoch, 1982; Neitz, Carroll, Yamauchi, Neitz, & Williams, 2002), and the removal of cataracts results in months of

re-adaptation before normal colour vision is restored (Delahunt, Webster, Ma, & Werner, 2004).

The use of portable head-mounted display systems has allowed the exploration of adaptation at never-before studied durations. Four hours of exposure to reduced contrast revealed improvements in contrast discrimination and a switch in coding from contrast gain to response gain changes (Kwon, Legge, Fang, Cheong, & He, 2009). This finding revealed different adapting strategies, with short-timescale adaptation being driven by the current mean contrast and long-timescale adaptation driven by the overall range of contrasts (Webster, 2015). At similar timescales, increases in orientation-specific sensitivity were seen following deprivation to contrast at a specific orientation (Zhang, Bao, Kwon, He, & Engel, 2009). These increases in sensitivity can be explained as adaptation to a 'loss' leading to increases, rather than reductions, in neuronal gain.

However, the limits of neuronal adaptation are reported to have been reached following four days of vertical contrast deprivation (Haak, Fast, Bao, Lee, & Engel, 2014). Over the course of the first day adaptation showed power law growth as expected. However, adaptation appeared to asymptote at day 2, after which it continued to grow but at a slower rate.

The change in the rate of adaptation in Haak et al.'s (2014) study highlights the need to understand what sets the rate of adaptation's time course – a question that the experiments in this thesis aim to address. To adapt optimally in our ever-changing environment, adapting mechanisms should track the rate of stimulus change and reflect this in the dynamics of adaptation. Faster changes in stimulus statistics should be met with faster adaptation, whilst slower changes are met with more gradual adaptation. This characteristic has been shown in motor adaptation with faster adaptation to transient changes, such as muscle fatigue, and slower adaptation to prolonged changes such as development or injury (Kording,

Tenenbaum, & Shadmehr, 2007; Shadmehr, Smith, & Krakauer, 2010; Smith, Ghazizadeh, & Shadmehr, 2006).

In addition to adapting over different timescales, the visual system must distinguish actual changes in stimulus input, from noise in neuronal responses. Recent physiological evidence suggests that retinal ganglion cells tackle this problem by accumulating evidence for a change in stimulus statistics (Wark, Fairhall, & Rieke, 2009). These cells adapt faster when there is more evidence for a change, thus avoiding expending energy adapting to noise. Similar findings are seen in guinea pig auditory midbrain neurons with increased rates of adaptation when an adapting stimulus is re-encountered (Robinson, Harper, & McAlpine, 2016). This suggests that the system uses recent visual experience to find an efficient code more rapidly.

Therefore, as shown in other domains, it is feasible that the rate of visual adaptation might show similar flexibility, altering the dynamics of adaptation in response to the frequency and type of changes in recent visual history. The ability of adapting mechanisms to themselves adapt is termed meta-adaptation and will form a central tenet of this thesis.

Due to the wide range of timescales over which contrast adaptation operates, there is ongoing debate over the number and characteristics of its underlying mechanisms. One theory is that a single neural mechanism operates across multiple timescales, increasing confidence in its estimates of the current environment as adaptation duration increases (see Figure 1.7 (A)). This increased confidence results in an increase in the evidence required to convince the system that the environment has changed back to its original state and thus produces longer lasting adaptation aftereffects (Grzywacz & de Juan, 2003; Wark et al., 2009). An alternative theory is that multiple mechanisms might drive adaptation, each tuned to distinct timescales (Kording et al., 2007)(see Figure 1.7 (B)). As adaptation duration

increases, mechanisms tuned to longer timescales exert more control over adaptation, and their dynamics result in longer lasting aftereffects.

Pitting long and short adapting durations against each other confirmed the existence of at least two adapting mechanisms operating over distinct timescales (Bao, Engel, Fast, & Mesik, 2013). The aftereffects induced from four hours of adaptation to contrast deprivation were temporarily nulled by 15 minutes of adaptation to natural viewing. This effect was short-lived as with continued testing the effects of the initial adaptation period returned – a phenomenon termed 'spontaneous recovery'. This pattern of results suggests that the initial adaptation activated both a short and long-timescale mechanism (see Figure 1.7 (C)). The short-timescale mechanism has fast dynamics, with adaptation growing and decaying quickly. In contrast, the long-timescale mechanism has slower dynamics, with adaptation growing more slowly and being stored for longer before decaying.

The initial adaptation activated both the short-timescale and the long-timescale mechanism, whilst the de-adaptation re-activated the short-term mechanism in the opposing direction. The short-timescale de-adaptation temporarily cancelled out the effects of the initial adaptation. Over time, the effects of the short-timescale mechanism decayed, revealing that the effects of the initial adaptation had been stored in the long-timescale mechanism which had yet to decay (Bao et al., 2013). This 'spontaneous recovery' is the signature of two-mechanisms and has been shown in motor learning (Smith et al., 2006), saccadic motor memory (Ethier, Zee, & Shadmehr, 2008), and adaptation to motion (Mesik, Bao, & Engel, 2013), and tilt (Magnussen & Johnsen, 1986).

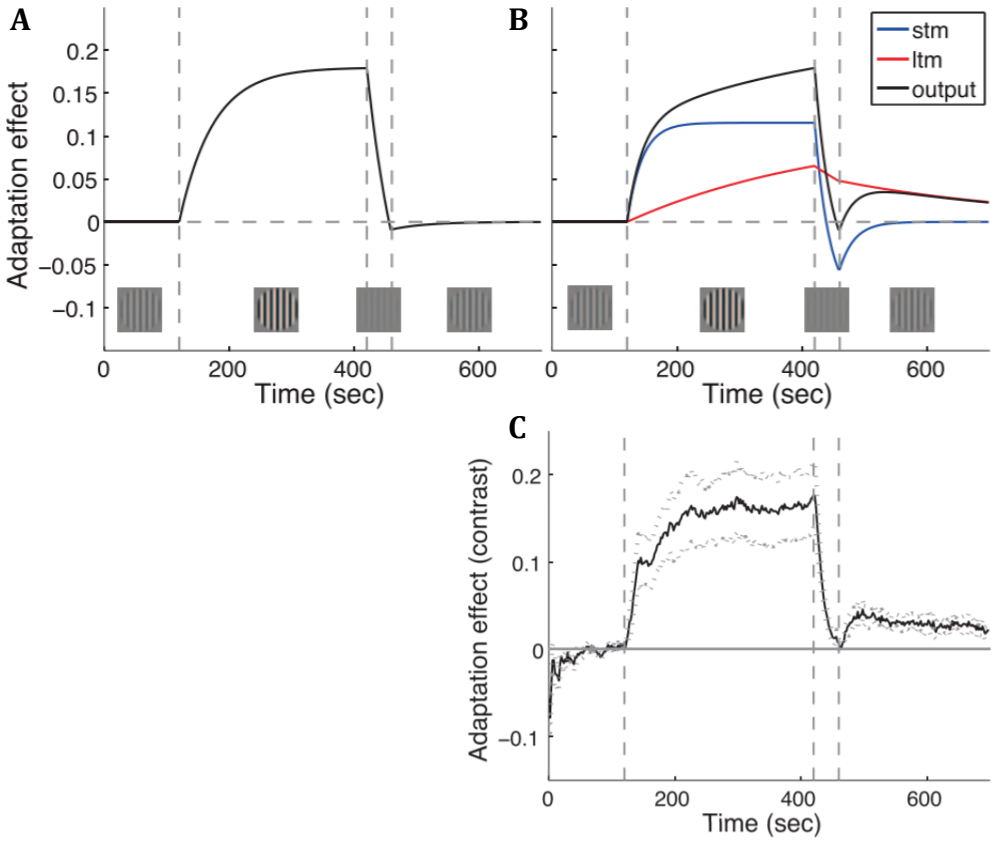


Figure 1.7 Top panel: The time course of adaptation as predicted by two mechanistic theories. (A) A single mechanism drives adaptation (black line). During adaptation to high contrast it grows in strength but is entirely cancelled out by de-adaptation to low contrast resulting in no adaptation. (B) Adaptation (black line) is the summed output of a short-timescale mechanism (blue line) and a long timescale mechanism (red line). Adaptation to high contrast activates both mechanisms, the short timescale mechanism adapts faster than the long timescale mechanism. De-adaptation to low contrast causes the short timescale mechanism to quickly shift in the opposite direction but has little effect on the slower long-timescale mechanism. Initially, this nulls adaptation as the mechanisms signal in opposing directions. Over time, the short-timescale mechanism rapidly decays, leaving the remaining output of the long timescale mechanism, resulting in 'spontaneous recovery'. Bottom panel: (C) Bao et al. (2013) found results consistent with a two-mechanism model of adaptation.

The perceptual consequences of adaptation are therefore produced by the summation of the responses from independent mechanisms with distinct temporal tuning. The nature, and number, of these adapting mechanisms remain unknown, as do the factors that determine the timescales to which these mechanisms are tuned. Having several mechanisms that are tuned to different timescales of adaptation allows the visual system to flexibly adapt to both transient and prolonged changes in environmental statistics. It is

possible, that the tuning of these mechanisms has been fixed through evolution in our natural environment, reflective of the timescales over which changes typically occur. To optimally adapt to the large number of timescales over which changes occur, this would involve many mechanisms with relatively narrow tuning. Alternatively, there could be fewer mechanisms, which possess the ability to alter their temporal tuning to match the prevailing temporal regularities in the environment. In this way, the system could track the temporal regularities in the environment and alter the rate of adaptation to reflect the rate of changes in the stimulus statistics (meta-adaptation). Disambiguating these possibilities psychophysically is difficult, as both schemes would produce the same measurable response. Logically, we must first identify whether the rate of visual adaptation is flexibly altered by the rate of change in the environment, before exploring the mechanisms that might drive this process. The experiments in this thesis take the first step in identifying whether the rate of adaptation is flexibly altered in response to changes in our temporal environment. The findings of these experiments will guide future research on the number and nature of visual adapting mechanisms.

1.1.4. Modelling adaptation

Exploring the number and characteristics of adapting mechanisms can be achieved through modelling. A parsimonious model of adaptation is the leaky integrator, which can be well described by a 'leaky bucket' analogy where the bucket is the adapting mechanism. The bucket has a hole in it, meaning adaptation decays if the rate of water flowing into the bucket does not counteract the rate of water flowing out of the bucket.

The time course of adaptation can therefore be considered in three phases (see Figure 1.8). Pre-adaptation, the bucket is empty and adaptation rests at zero. During adaptation, water flows into the bucket, counteracting the leak it begins to fill. The fuller the bucket gets, the faster water flows out of

the hole, creating adaptation's characteristic power law growth. Eventually the bucket is full, signalling the adapting mechanism has reached saturation. Following adaptation offset, the water leaks out of the bucket, decaying back to baseline.

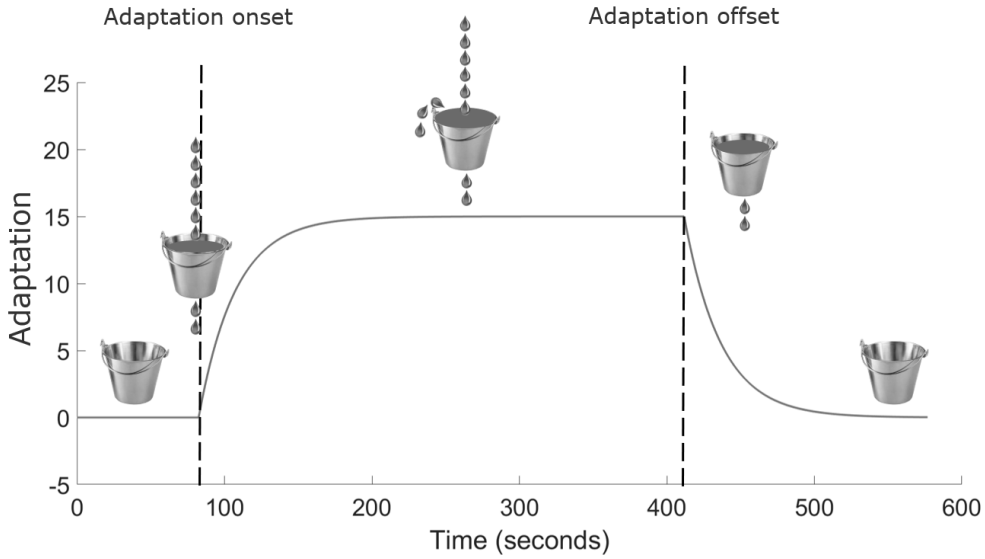


Figure 1.8 the time course of adaptation as predicted by a leaky integrator model. The buckets represent the adapting mechanism, the water represents adaptation itself. Prior to adaptation onset, the bucket is empty as we are at baseline. The onset of adaptation is represented by the flow of water into the bucket, the bucket begins to fill but some water leaks out of the bucket, setting the rate at which it fills. Saturation of adaptation occurs when the bucket is full. Following adaptation offset the water leaks from the bucket representing the decay of adaptation. Once all the water has leaked out, we once again reach a baseline state of no adaptation.

The time course of adaptation can be altered by two properties of the bucket. The size of the bucket sets the magnitude of adaptation and is referred to as the 'gain' of the mechanism. The size of the hole in the bucket determines the rate at which adaptation grows and is known as the 'time constant' of the mechanism. Adaptation's time course has been successfully modelled by the linear summation of multiple adapting mechanisms with different gains and time constants (Bao et al., 2013) reflecting the finding that adaptation is driven by multiple adapting mechanisms.

The mathematics behind the leaky bucket analogy is a leaky integrator adapting mechanism. Generically, models of adaptation consist of an

element that integrates the response to the adapting stimulus and feeds back the integrated signal to suppress the response (Drew & Abbott, 2006; see Figure 1.9).

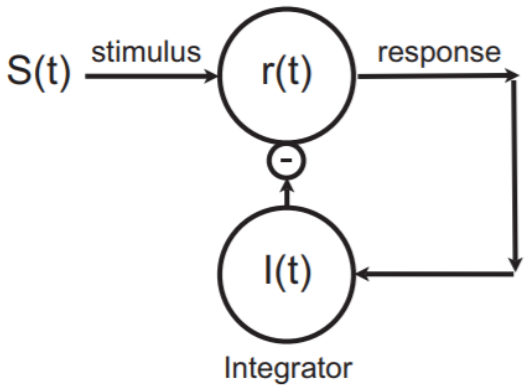


Figure 1.9 General model of adaptation. The stimulus S(t) creates a response r(t) that feeds back into an integrator I(t) that suppresses the response so that r(t) = S(t) – I(t). Different types of integrators lead to different types of adaptation. Adapted from Drew & Abbott (2006).

The most common type of integration for modelling neural adaptation is leaky integration. The suppressive effect of the response is accumulated with exponential forgetting at a fixed time constant. Adaptation therefore grows and decays at the rate of the time constant of the integrator. The equation for the leaky integrator is shown below, where t is time, and g and c are the gain and time constant of the leaky integrator respectively.

$$y = g\frac{e^{\frac{-t}{c}}}{\int e^{\frac{-t}{c}}dt} \qquad \text{Equation 1.1}$$

To implement the model, the box-car of the stimulus time course is convolved with the leaky integrator. The magnitude and rate of adaptation can be altered by changing the gain and time constant of the integrator respectively. Multiple leaky integrator mechanisms with different parameters can also be summed, allowing the modelling of adaptation driven by multiple adapting mechanisms. The leaky integrator is therefore

a simple but powerful model of adaptation and will be frequently used to model the empirical findings of the experiments in this thesis.

1.1.5. Function

The prevalence of adaptation throughout the visual hierarchy suggests an important functional role in determining how we see. However, much like the different nature of processing at each stage of the visual hierarchy, the functional purpose of adaptation may also change depending on stimulus domain or task. Currently, there are three main hypotheses on the function of adaptation: efficient coding, error correction, and predictive coding.

Information theoretic approaches suggest the function of adaptation is to maximise efficient coding and optimise information transmission (Clifford et al., 2007; Rieke & Rudd, 2009; Stocker & Simoncelli, 2006; Wainwright, 1999). Visual neurons are faced with encoding a large range of stimulus levels with only a limited dynamic range of responses. Adaptation allows this problem to be overcome by adjusting neurons' dynamic range to match the prevailing stimulus statistics, preventing response saturation and balancing sensitivity across neurons. In single neurons, reducing the transmission of redundant information (Attneave 1954) maximises the effective bandwidth available for novel information transmission (Srinivasan et al, 1982). In populations of neurons, responsiveness is balanced through adaptive decorrelation, ensuring neurons are not responding to the same stimuli (Barlow & Foldiak 1989; Atick et al 1993). Adaptation has thus been modelled as the combined effects of 'centring' and 'scaling' whereby the dynamic range of responses is shifted to be centred on the prevailing stimulus distribution and the responses of populations of neurons are scaled to decorrelate their responses and result in efficient information transmission (Clifford, Wenderoth, & Spehar, 2000). This model accounts for both attractive and repulsive aftereffects in tilt, motion, and chromatic adaptation.

This approach can be readily applied to light adaptation in the retina. Characterised by a horizontal shift in the luminance-response relationship of retinal ganglion cells, retinal adaptation allows us to discriminate relatively small changes in luminance despite variations in light intensity that occur over ten orders of magnitude. Adaptation in the retina is therefore vital to our ability to encode the sensory environment but is one of the only forms of visual adaptation to demonstrate such a clear functional benefit. Moving to higher sensory areas, it is possible that adaptation serves to calibrate these mechanisms, adjusting their sensitivity to the range of stimuli they're likely to receive (Webster, 2015). However, most higher-level stimuli do not vary across as large a range as light level, so it is unclear why they would need continual readjustment.

Another caveat of information theoretic approaches is that adaptation induced changes in visual coding, such as centring dynamic ranges on the prevailing stimulus statistics, would predict improvement in perception of these stimuli. Psychophysical data to support this prediction is sparse. For example, improvements in contrast discrimination have been observed but they are sensitive to the adaptation protocol (specific contrast levels) and discrimination task (monocular or binocular) (Abbonizio, Langley, & Clifford, 2002), and improvements are seen for stimuli with contrasts distinct from the adapting stimulus (Greenlee & Heitger, 1988). Changes in performance have been investigated for repulsive (Muller et al 2009) and attractive (Krekelberg et al 2006) shifts in preference. But, discrimination thresholds for stimuli similar to the adapter are either unaffected or only slightly enhanced by adaptation (Dragoi et al 2002; Hol & Treue 2001; Krekelberg 2006; Regan & Beverly 1985). Although, improvements have been seen in discriminating temporal luminance variations with increasing adapting durations (Kristjánsson, 2012). And it is suggested that the variation in stimulus input may be more important than the absolute level of stimulation. Hence, adaptation may appear harmful in the fact it alters

perception of absolute level, but it is functionally beneficial in improving discriminability around the mean level at that time.

In contrast to information theoretic approaches, error correction theories suggest that the role of adaptation is to maintain perceptual stability rather than alter perceptual sensitivity (Andrews, 1967). This theory gains support from the demonstration of large perceptual aftereffects when there is little change in sensitivity. Such accounts suggest that alterations in the observer, such as the yellowing of the lens of the eye (Werner & Schefrin, 1993), or the environment, such as lighting (Foster, 2011), lead to a bias in the neural code which is then compensated for by adaptation.

Another proposed function of adaptation is its involvement in making predictions about the world (Chopin & Mamassian, 2012; Srinivasan, Laughlin, & Dubs, 1982). Adaptation is suggested to 'null' the stimuli we are exposed to, forming perceptual norms and a null response to the current state of the world. This function improves the efficiency of information transmission by only signalling errors or unexpected stimuli (Barlow, 1990; Boehnke et al., 2011; Gardner et al., 2005; Ranganath & Rainer, 2003), which is supported by reports of enhanced salience and visual search of novel stimuli (McDermott, Malkoc, Mulligan, & Webster, 2010; Wissig, Patterson, & Kohn, 2013). This idea suggests that salient visual information can be interpreted as aftereffects from our current adapted state (Barlow, 1990; Webster, Werner, & Field, 2005).

Some of these functional accounts are incompatible. For example, different mechanisms are needed to optimize identification over simple detection (Hillis & Brainard, 2007), and there can be a trade off when trying to achieve both constancy and efficiency (Webster & Mollon, 1995). These functions all seem important, there is evidence in support of each one, and a definitive test to distinguish them remains elusive. However, Webster (2014) suggests that trying to distinguish them is likely a futile task, as all of the

functional benefits can be produced by a single adjustment – normalizing neural responses around the mean stimulus level.

Recent physiological evidence suggests that adaptation serves to maintain homeostasis in neuronal populations (Benucci, Saleem, & Carandini, 2013). Recording populations of V1 orientation selective neurons during the presentation of a heavily biased stimulus distribution revealed adaptation equalized the time averaged response across the neuronal population. That is, despite one orientation being presented three times more than all others, neurons selective for this orientation did not respond more than the rest of the population when averaged over time. In addition to equalization, adaptation served to decorrelate the responses of neurons, ensuring that a biased stimulus ensemble didn't result in high correlation of neurons with similar tuning around the adapting orientation. Further to this, Benucci et al. (2013) note that this adaptation occurred in under 2 seconds, with the visual system only needing to observe the same orientation in 18/52 presentations to adapt to a change in statistics. These findings support previous proposals of centring and scaling (Clifford et al., 2000), response independence through decorrelation (Barlow & Földiák, 1989), and response normalization discussed earlier

1.2. THE TILT AFTEREFFECT

1.2.1. Definition

Prolonged viewing of a tilted stimulus causes subsequently viewed stimuli to appear tilted away from the adapting stimulus orientation (Gibson & Radner, 1937). For example, viewing of a stimulus tilted 20deg clockwise from vertical will result in a subsequently presented vertical stimulus appearing tilted counter-clockwise from vertical (see Figure 1.10).

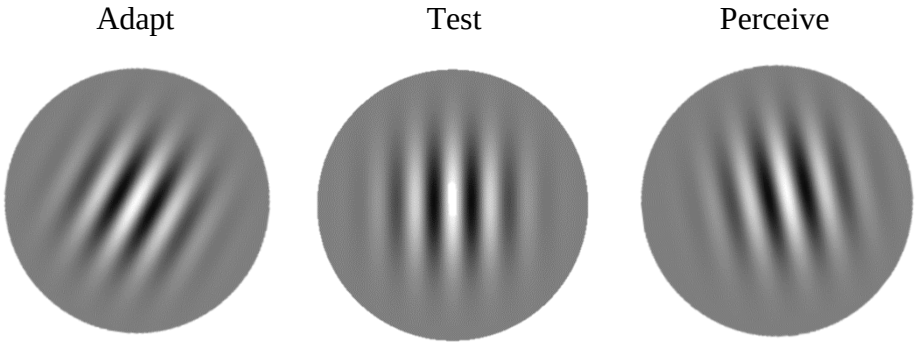


Figure 1.10 Graphical representation of the tilt aftereffect (TAE). Prolonged viewing of the tilted adapt stimulus causes the vertical test stimulus to appear tilted away from the adapting orientation.

1.2.2. Tuning

The magnitude and direction of the TAE is dependent on the angular difference between the adapting and test stimulus orientations (see Figure 1.11). Adapting stimuli oriented 0-50deg clockwise from vertical will cause a vertical test stimulus to appear tilted counter-clockwise, with the strongest effect for stimuli oriented 10-20deg. For larger angular differences, there is a smaller attractive effect, with test stimuli appearing tilted towards the adapter. This effect peaks at 75-80deg (Wenderoth & Johnstone, 1987) and is robust for large adapting and test stimuli (Gibson & Radner, 1937; Köhler & Wallach, 1944; Morant & Harris, 1965). This angular relationship exists regardless of the test stimulus orientation, such that the adapter will result in the largest TAE on presentation of a stimulus oriented 10-20deg away.

The magnitude of the tilt aftereffect is also sensitive to changes in adapter and test contrast and spatial frequency. Low contrast stimuli produce smaller TAEs at short test durations, and larger TAEs at long test durations, than high contrast stimuli. A similar interaction is found between contrast and spatial frequency, with low spatial frequency stimuli producing larger TAEs at short test durations, and smaller TAEs at long test durations, than

high spatial frequency stimuli. These effects are suggested to result from differential activation of transient and sustained channels (Harris & Calvert, 1989).

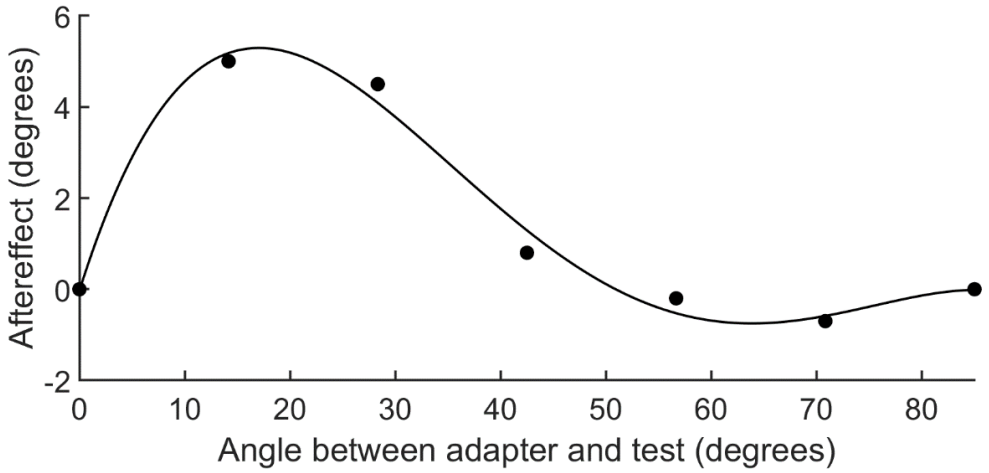


Figure 1.11 The angular tuning of the tilt aftereffect (TAE). The magnitude of the TAE depends on the angular difference between adapter and test stimuli. Repulsive aftereffects (positive on y axis) are seen for angular differences between 0-50deg, peaking at 10-20deg. Smaller attractive aftereffects (negative on y axis) are seen for angular differences greater than 50deg and peak at 75-80deg. Adapted from Clifford, Wenderoth & Spehar (2000).

The duration of adapter and test stimuli can also affect the magnitude of the TAE. Shorter test durations produce larger TAEs (Wolfe, 1984), and TAE magnitude increases as a logarithmic function of adapting duration (Magnussen & Johnsen, 1986). The dynamics of TAE growth and decay will be discussed further in Section 1.2.5.

The TAE can also be affected by adapter and test stimulus position and observer attention. The TAE is retinotopic in nature, meaning the largest TAEs are produced when adapting and test stimuli fall on the same retinal location (Knapen, Rolfs, Wexler, & Cavanagh, 2011; Mathôt & Theeuwes, 2013). In addition to spatial location the magnitude of the TAE is also affected by top-down attention (Festman & Ahissar, 2004), with TAE's produced by an attended adapting stimulus being 39% larger than TAEs produced by an unattended stimulus (Spivey & Spirn, 2000).

Thus far the focus has been on the properties of simplistic adapting and test stimuli such as gratings and Gabors, but tilt aftereffects can also be produced by more complex stimuli. For example the TAE has been demonstrated to oriented information in natural images (Dekel & Sagi, 2015a), translational glass patterns (Pavan, Hocketstaller, Contillo, & Greenlee, 2016), and even 'invisible' patterns so high in spatial frequency that they were indistinguishable from a uniform grey field (He & MacLeod, 2001). The reproducibility and depth of study of the tilt aftereffect makes it a suitable proxy for adaptation, hence its use in the experiments in this thesis.

1.2.3. Psychophysical measurement

Traditional measurement of the tilt aftereffect (TAE) involves the presentation of an oriented adapting stimulus followed by measuring subsequent bias in orientation perception. There are several ways in which biases in orientation perception can be measured. The first two experimental chapters of this thesis explore appropriate methodologies for the measurement of the time course of the tilt aftereffect, some advantages and disadvantages of popular methodologies are discussed below.

A common way to measure adaptation induced biases in orientation perception is the method of adjustment. The observer is briefly presented with an oriented test stimulus, following which they adjust the orientation of a response stimulus until they believe it perceptually matches that of the test stimulus. The response stimulus is presented in a different retinal location to the adapting and test stimuli. Due to the retinal specificity of the tilt aftereffect (Knapen et al., 2011; Mathôt & Theeuwes, 2013), this ensures the response stimulus is unaffected by the adapting stimulus. Bias can then be calculated as the angular difference between the test stimulus and the observer's adjustment response (see Figure 1.12). An advantage of this method is that an estimate of the tilt aftereffect can be gathered from the

presentation of a single stimulus. However, this method is very transparent and can lead to observer bias. If an observer is aware of the experimental hypothesis, they could consciously or subconsciously alter their responses to match the perceived goals of the experiment.

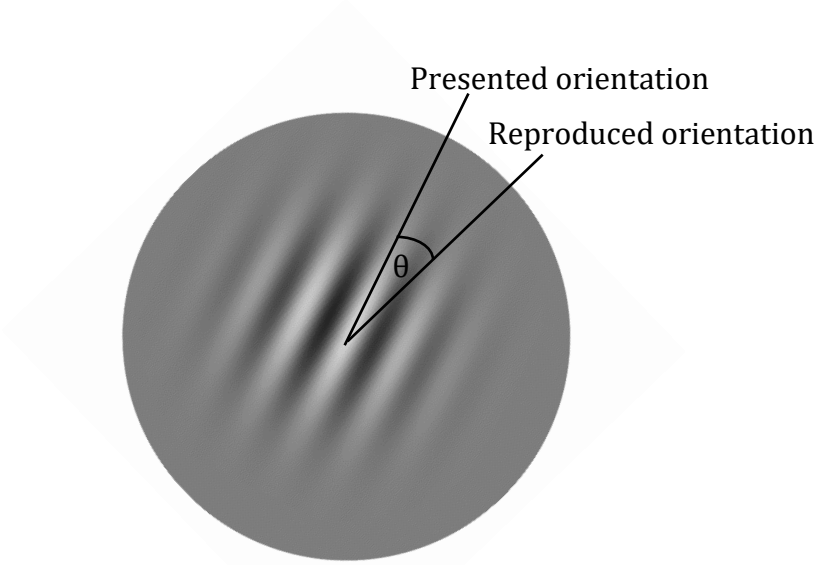


Figure 1.12. Adaptation-induced perceptual bias (θ) calculated as the angular difference between the orientation of the presented test stimulus and the observer's reproduced orientation.

A second way to measure perceptual bias is the method of constant stimuli. Observers are presented with stimuli with a range of orientations around 10-20deg from the adapting stimulus and must decide whether they are clockwise or counter-clockwise from a reference point. For example, if the adapting stimulus is 15deg counter-clockwise, the largest tilt aftereffect will be observed in response to near-vertical stimuli (Gibson & Radner, 1937). Therefore, the reference point for the observer is vertical, and they are presented with near-vertical orientations which they must judge as being clockwise or counter-clockwise from vertical.

When using the method of constant stimuli, the judgement the observer makes on the test stimulus can be around an implicit reference such as vertical or horizontal or can be marked by an explicit reference near the test

stimulus such as a reference line, bar, or grating. The observer judges whether the test stimulus is more clockwise or counter-clockwise than the reference point.

The selection of the test stimulus orientations depends on the predicted magnitude and direction of the effect. In this example, a -15deg adapting stimulus will likely shift perceived vertical to being around +2-5deg. Therefore, the distribution of test stimulus orientations should include stimuli to capture this shift and should consist of at least 5 different orientations to properly sample perceived orientation in this space. Each test stimulus orientation is presented multiple times to produce a reliable estimate of bias at each orientation.

An estimate of the tilt aftereffect is produced by plotting the test stimulus orientations against the proportion of times observers judged them as being tilted clockwise from the reference point. Pre-adaptation, when the stimulus is counter-clockwise from vertical, the proportion of times observers judged this stimulus to be clockwise will be very low. When the stimulus is clockwise from vertical, the proportion of times the observer judged the stimulus to be clockwise will be very high. When the stimulus was vertical, the proportion of times the observer judged the stimulus to be clockwise will be roughly 50%. This produces a s-shaped pattern of data that can be well fit by several functions such as the logistic and Weibull.

Bias in orientation perception is calculated from the point at which observers perceive presented stimuli to be indistinguishable from vertical which corresponds to when observers' proportion of clockwise responses is at 50% and is known as the point of subjective equality (PSE, dashed lines in Figure 1.13). This point of subjective equality shifts following adaptation so that perceived verticality is biased away from the adapting orientation (shift from black to grey dashed line in Figure 1.13). The magnitude of this shift, is the magnitude of the TAE.

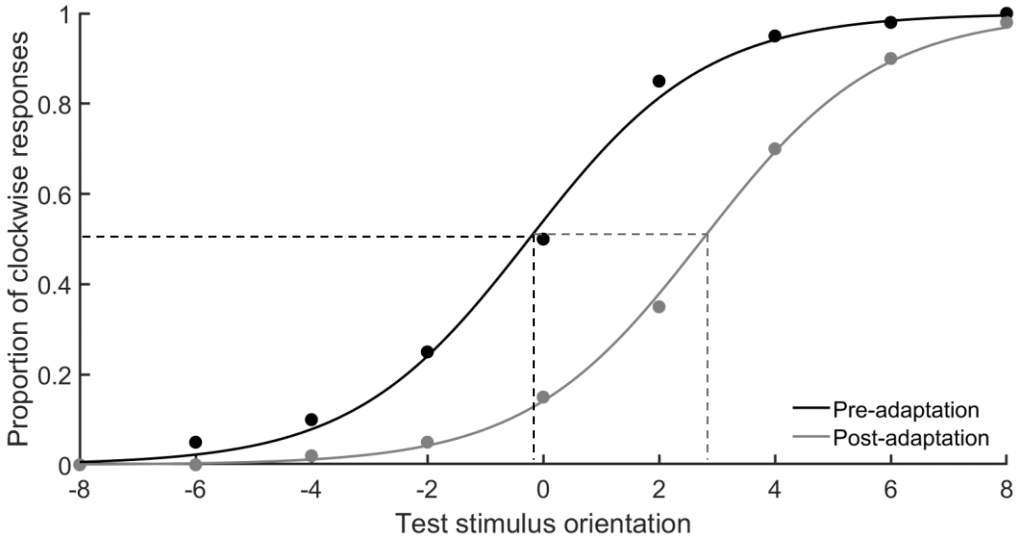


Figure 1.13. Simulated psychometric functions pre-adaptation (solid black line) and post-adaptation to a counter-clockwise tilted adapter (solid grey line). Dashed lines indicate the orientation at which the observer perceived the presented stimulus to be indistinguishable from vertical (the point of subjective equality, PSE). Note the positive shift in the PSE following adaptation to a counter-clockwise stimulus.

An advantage of a two-alternative forced-choice response is that it is slightly less transparent, with the observer being less likely to influence the results due to any demand characteristics. Although there does remain the possibility that observers could response 'clockwise' on proportionally more stimuli to enhance the TAE.

Additionally, the method of constant stimuli is disadvantageous in requiring multiple trials to gain an estimate of the TAE. This therefore requires maintaining the observer in an adapted state throughout the test period using an adapt-test-top-up design. In this design, a prolonged period of initial adaptation is followed by the presentation of a test stimulus and gathering of a response, after which there is a shorter period of top-up adaptation followed by the presentation of the next test stimulus. The top-up period aims to counteract the decay of adaptation caused by interrupting to gather a response, maintaining the observer's adapted state throughout

testing. Guidelines for the ratio of initial adaptation to top-up adaptation do not exist, a caveat which will be explored further in Chapter 4.

1.2.4. Physiology

The physiological consequences of adaptation to orientation were discussed in detail in Section 1.1.1 but will be summarised here. Primary visual cortex contains populations of neurons selectively tuned to orientation (Hubel & Wiesel, 1962). Each neuron's selectivity can be represented as a tuning curve with a bell-shape (See Figure 1.3) - neurons respond preferentially to the orientation they are selective for, with orientations on either side provoking little response.

Following adaptation, there are two changes to the response properties of these neurons. Adaptation results in suppressed responses of neurons tuned to the adapting orientation, and the preferred orientation of neurons shifts away from the adapting orientation (Dragoi, Sharma, & Sur, 2000; Dragoi, Rivadulla, & Sur, 2001). In terms of individual tuning functions, this results in a depression of the bell-curve flank near the adapting orientation, an increase in response on the opposite flank, and a shift in the peak of the curve away from the adapting orientation (Dragoi et al., 2001).

The observer's perceived orientation depends on the balance of activity across a population of orientation tuned neurons (Clifford et al., 2000; Gilbert & Wiesel, 1990; Pouget, Dayan, & Zemel, 2000; Vogels, 1990; Wainwright, 1999). Repulsive biases have been shown to be the result of the reduced responsiveness of neurons tuned to the adapting orientation, creating an asymmetrical bias in the population response (Dragoi et al., 2000b; Seriès, Stocker, & Simoncelli, 2009). This was supported by a population coding model, which showed a reduction in the firing rate of neurons tuned near to the adapting orientation resulted in a biased population response (Jin, Dragoi, Sur, & Seung, 2005). Somewhat

surprisingly, the repulsive shifts in the preferred orientation of tuning curves away from the adapting orientation predict that oriented stimuli would appear tilted towards the adapting stimulus, and therefore reduce the magnitude of the TAE (Gilbert & Wiesel, 1990; Teich et al., 2003; Yao & Dan, 2001). In summary, the reduced amplitude of tuning curves near the adapting orientation produces the perceptual TAE, whilst repulsive shifts in neurons' preferred orientation might act to reduce the perceptual error induced by neuronal response suppression (Jin et al., 2005).

1.2.5. Time course

The magnitude of the tilt aftereffect increases with adapting duration until reaching saturation after approximately 1 hour (Greenlee & Magnussen, 1987) (see Figure 1.14). The rate of growth increases as a logarithmic function of time and is thus often plotted on logarithmic time axes. The decay of adaptation following adapter offset follows a similar trajectory (see Figure 1.14).

To date, measurement of the growth of the TAE is rare, with most studies measuring the final magnitude of the aftereffect and its subsequent decay. Measuring the growth of adaptation can be achieved in two ways. The traditional way, shown in Figure 1.14, is to use discrete continuous periods of adaptation increasing in duration. The magnitude of the TAE is measured following each duration of adaptation to sample the growth of the TAE over time. This method is advantageous in its simplicity and is useful when measuring TAE growth over short durations. However, measuring longer periods of adaptation in this way quickly becomes time consuming, especially if high temporal sampling is needed, as the number of discrete periods of adaptation increases.

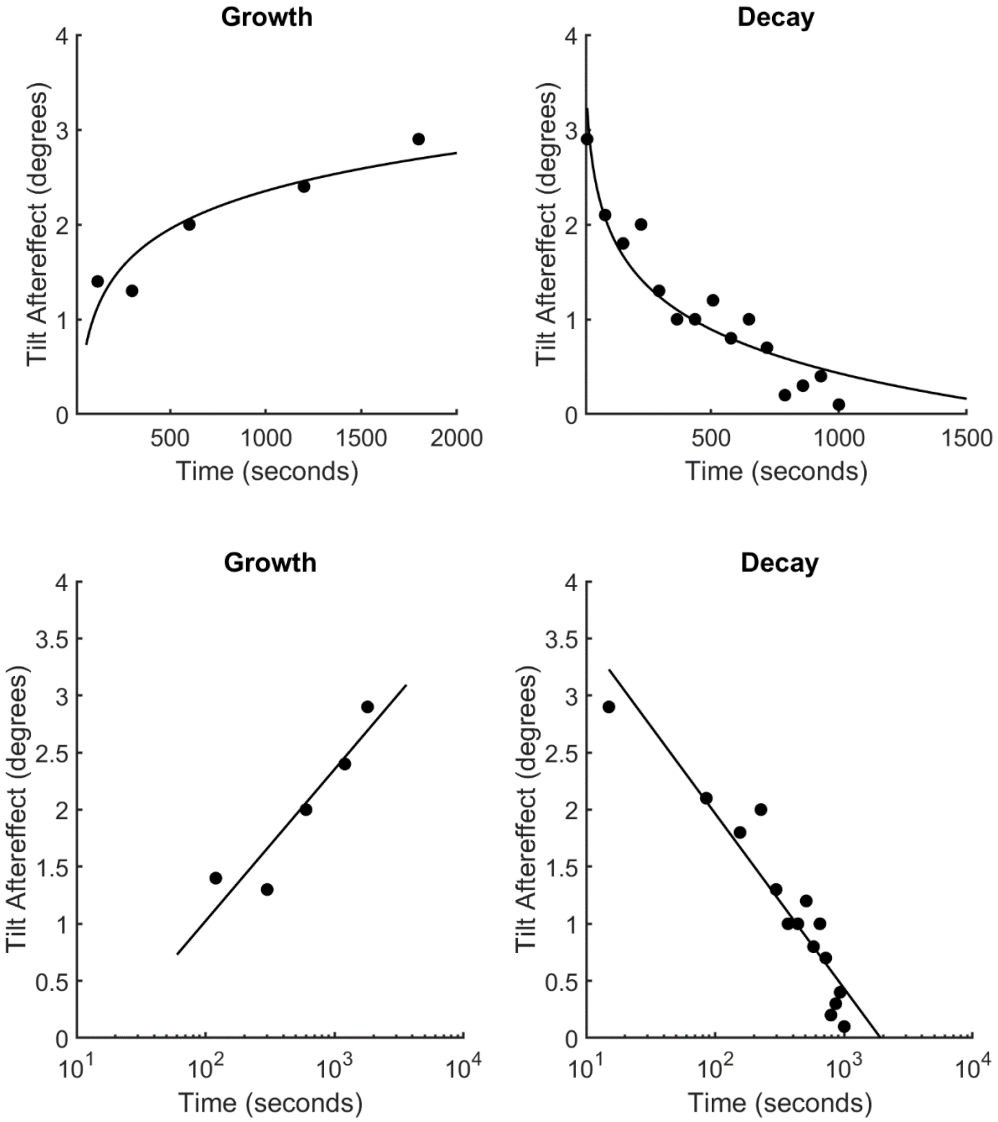


Figure 1.14. The growth (left) and decay (right) of adaptation to a 10deg grating and tested with a vertical test grating plotted with a linear (top) and logarithmic (bottom) x axis. Adaptation grows and decays as a logarithmic function of time and is well fit by a logarithmic function (solid lines). Adapted from Magnussen & Johnsen (1986).

A less common method to measure the growth of adaptation is through interrupted adaptation. A period of adaptation is briefly interrupted by the presentation of a test stimulus and gathering of a response at a chosen frequency. Each interruption therefore samples the magnitude of the TAE at that point in time, and the frequency of interruption determines the temporal resolution at which the time course of adaptation is sampled. This method allows efficient measurement of longer durations of adaptation, but

the effects of interrupting adaptation on the TAE time course remain unexplored.

As mentioned earlier, some aftereffects display 'storage', whereby presentation of a uniform field between adaptation offset and measurement of the aftereffect does not result in decreased aftereffect magnitude (Machulla et al., 2012; Spigel, 1960). In other words, adaptation's effects are stored until presentation of an appropriate de-adapting stimulus, which in some cases can even lead to permanent adaptation (Machulla et al., 2012; Spigel, 1960). Interestingly, the TAE does not show storage (Magnussen & Johnsen, 1986). Therefore, increasing the duration between adaptation offset and presentation of the test stimulus will result in a smaller magnitude TAE as adaptation passively decays during a period of no stimulation. This suggests that interrupting adaptation in the manner described above could alter the time course of the TAE due to frequent periods of passive decay. This hypothesis will be formally tested in Chapter 4.

In general, TAE growth can be summarised by two properties - its magnitude and its rate. The former is known to be altered by the duration of adaptation, the angular difference between adapt and test patterns, the contrast and spatial frequency of adapt and test patterns, attention, and retinal location. The latter, however, remains relatively unexplored. Mei, Dong, & Bao (2017) compared adaptation to compound stimuli typically processed in mid-level visual areas to component stimuli in low-level visual areas and revealed slower rates of adaptation growth and decay in mid-level areas. This is suggested to be the product of adapting mechanisms in mid-level areas giving more weight to their past states than those in early visual areas (known as a longer temporal receptive window - Chaudhuri, Knoblauch, Gariel, Kennedy, & Wang, 2015; Honey et al., 2012). This finding suggests that the rate of adaptation may differ between stimuli of different complexity but does not elucidate what determines the rate of adaptation within a single stimulus property, and whether this rate it fixed or flexible.

Intuitively, the rate of adaptation should reflect the rate of change in the current environment, with fast changes being met with rapid adaptation and slow changes with gradual adaptation. The effect of prior temporal context on the time course of adaptation is termed 'meta-adaptation', whereby adapting mechanisms are themselves plastic, altering the dynamics of adaptation in response to changes in the temporal statistics of the environment. Whether visual adapting mechanisms exhibit meta-adaptation will be formally tested in Chapters 5, 6, and 7.

Adaptation effects are temporal by nature and are considered a hallmark of a system optimized to the temporal structure of the natural environment. (Barlow & Földiák, 1989; Dayan, Sahani, & Deback, 2002; Fairhall, Lewen, Bialek, & de Ruyter Van Steveninck, 2001; Lochmann, Ernst, & Deneve, 2012; Wainwright, Schwartz, & Simoncelli, 2001). As demonstrated by the lack of research on the rate of adaptation, there are many unanswered questions on how adaptation functions under natural viewing and specifically in temporally complex environments. It is therefore useful to study adaptation processes in contexts informed by the statistical properties of natural scenes.

1.3. VISUAL PROCESSING IN NATURAL SCENES

1.3.1. Statistical structure of natural scenes

Natural scenes show significant statistical regularities in both the spatial and temporal domains. Spatially, neighbouring regions of a natural image are correlated in intensity. This is intuitive when observing the similarities of neighbouring pixels in an image. The standard measurement of spatial correlation is the autocorrelation function, which gives the correlation of intensity at two separate locations as a function of their relative positions. The magnitude of correlation decreases with increasing distance between positions (Reinagel & Zador, 1999), showing that neighbouring positions

have more similar intensity statistics than far positions (see Figure 1.15 (left)).

Natural images can be decorrelated by transforming to the frequency (Fourier) domain, where the two dimensional power spectrum can be reduced to a one dimensional function of spatial frequency (see Simoncelli & Olshausen (2001) for further explanation). The spectral power of natural images is seen to fall with spatial frequency (f), with power law behaviour ($1/f^p$), where values of p are around 2 (Ruderman & Bialek, 1994; Tolhurst, Tadmor, & Chao, 1992) (see Figure 1.15 (right)).

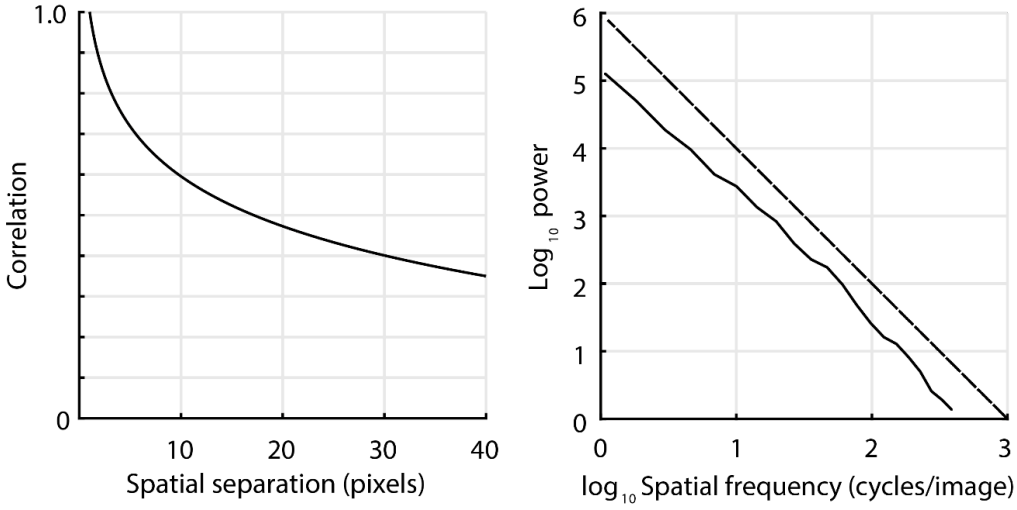


Figure 1.15 Left: Autocorrelation function. Neighbouring regions of a natural scene show more correlation in intensity values than regions far apart. Right: Power spectrum of a natural image (solid line) compared with $1/f^2$ (dashed line). Adapted from Simoncelli & Olshausen (2001).

Natural images also demonstrate statistical regularities in oriented contours, with a prevalence of vertical and horizontal contours in indoor, outdoor, and entirely natural scenes. The anisotropic distribution of orientations in natural scenes has been shown in both the Fourier spectra of natural images (Switkes, Mayer, & Sloan, 1978) (see Figure 1.16B), and by using a Sobel direction filter to obtain the distribution of orientations in natural images (Coppola et al., 1998) (see Figure 1.16 A).

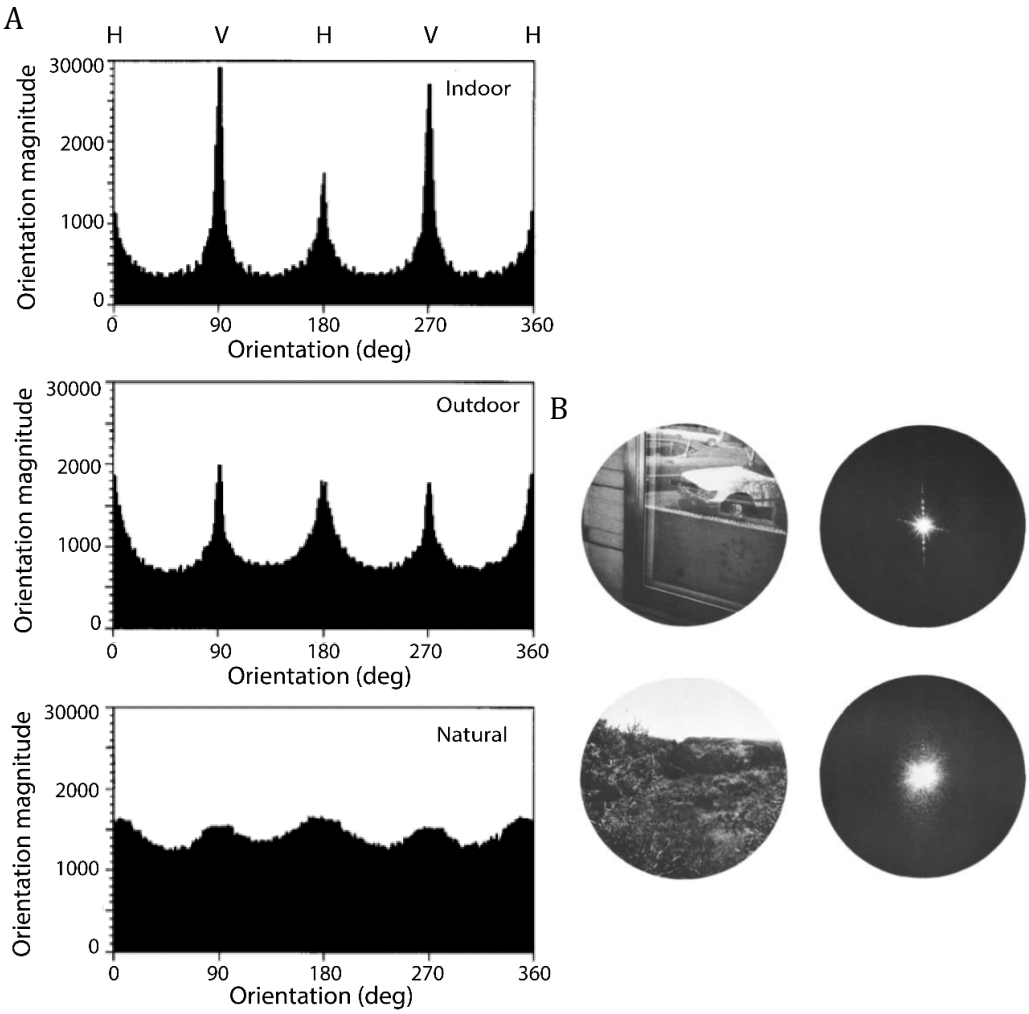


Figure 1.16 (A) Distribution of oriented contours in indoor, outdoor, and entirely natural scenes. (B) Power spectra of outdoor carpentered (top) and pastoral (bottom) images. Taken from Coppola et al (2001) and Switkes et al (1978) respectively.

Images falling on the retina must also be considered in terms of their temporal context, which is influenced both by self-movement and the movement of objects in the world. As before, the temporal statistics of natural movies are revealed by Fourier analysis, in this case with time as the third dimension (see Figure 1.17). Analysing spatiotemporal statistics reveals interesting interactions between the two domains. The spatial frequency spectrum exhibits $1/f$ structure, but the slope of the function shallows with increasing temporal frequency. The temporal frequency spectrum also exhibits $1/f$ structure, with more energy at low temporal frequencies, but again the slope of the function shallows with increasing

spatial frequency (Dong & Atick, 1995). It is likely that the consistent spatial and temporal regularities in naturalistic visual input are exploited by the visual system.

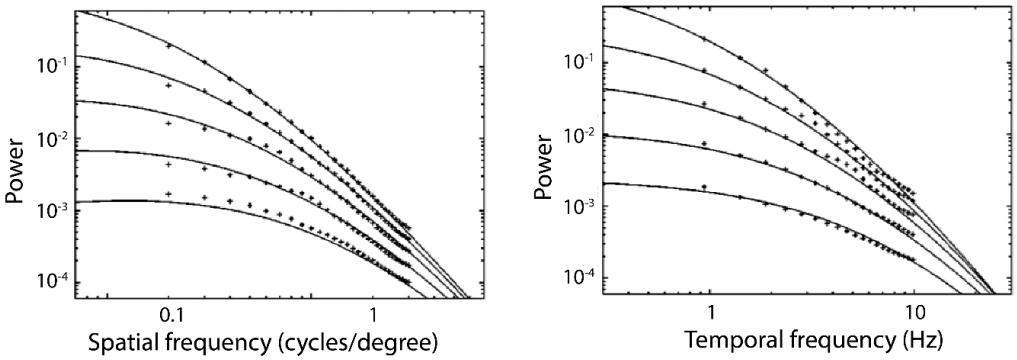


Figure 1.17 Spatiotemporal power spectrum of natural movies. Left: Joint spatiotemporal power spectrum shown as a function of spatial-frequency for different temporal frequencies (1.4, 2.3, 3.8, 6, and 10 Hz, from top to bottom). Right: Same data, replotted as a function of temporal frequency for different spatial frequencies (0.3, 0.5, 0.8, 1.3, and 2.1 cy/deg., from top to bottom). Solid lines indicate model fits according to a power-law distribution of object velocities (from Dong & Atick (1995)).

1.3.2. Visual processing and natural scenes

Early studies of visual processing held an important assumption: the way the brain processes visual stimuli can be understood by the responses to simplistic artificial stimuli. However, the architecture, processing, and efficiency of the visual system is no accident, having been shaped by evolution and refined by experience-dependent plasticity during development. The structure of the natural environment is therefore inherently related to the function of the visual system.

One way in which the visual system is optimized to process naturalistic stimuli is the reduction of redundancy in neuronal responses, which can occur for neurons that respond to nearby regions of space with similar visual properties. Removing the correlation between neuronal responses requires exploitation of the spatial statistical regularities of natural scenes. The inhibitory surrounds of photoreceptor fields were the first evidence for

effective spatial decorrelation in early visual processing (Srinivasan et al., 1982), and in combination with the biphasic impulse response functions of retinal and LGN neurons (Dong & Atick, 1995), leads to a significant reduction in coding redundancy for naturalistic images (Barlow, 1961). Since then, it has been revealed that the visual system 'whitens' (equivalent to decorrelation) the power spectrum of natural images to reduce redundancy of information transmission. Spatial whitening has been shown to be consistent with single cell physiology and the psychophysically measured contrast response function (Atick & Redlich, 1992) and is an example of how the visual system has evolved to efficiently process naturalistic visual input.

In addition to the decorrelation of neuronal responses, the visual system must optimally use a neuron's limited dynamic range to encode a vast array of incoming stimuli. A plausible solution to this problem is that neurons match their range to the range of signals to be encoded in the natural environment. Dynamic range matching ensures that the neuron uses its entire range and is not over- or under-responding. Histogram equalization is the proposed mechanism, adjusting the shape of the neuron's response function so that all response levels occur evenly with natural stimulation (Geisler, 2008). Physiological evidence has shown that monopolar neurons in the eye of the blowfly (Laughlin, 1981), and contrast responses of neurons in the LGN and V1 show histogram equalization to prevent under and over- responding to naturalistic stimuli (Brady & Field, 2000; Clatworthy, Chirimuuta, Lauritzen, & Tolhurst, 2003; Tadmor & Tolhurst, 2000). These efficient coding processes were first characterized with artificial stimuli, but their functional relevance has been revealed through the study of naturalistic stimuli (Felsen & Dan, 2005).

With specific relevance to the topic of this thesis, the anisotropic distribution of orientations in natural scenes in also thought to have affected processing of oriented information in the visual system. The anisotropic distribution of stimuli is mirrored in the distribution of

orientation selectivity in V1 neuronal populations – with significantly more neurons tuned to vertical and horizontal orientations (Li, Peterson, & Freeman, 2003). This, in turn, is thought to underlie the 'oblique effect' – a perceptual phenomenon whereby observers exhibit better discrimination performance around the cardinal orientations (Appelle, 1972). Matching the tuning characteristics of a population of neurons to the distribution of natural signals is also shown for binocular disparities that occur during navigation in a natural environment in monkeys (Liu, Bovik, & Cormack, 2008). The visual system has therefore evolved to process naturalistic stimuli, exploiting their statistical regularities through both the structural architecture of the visual system and visual processing mechanisms.

These three examples show that the structure and function of the visual system is inherently linked to the statistical properties of naturalistic input. However, rather than present the visual system with naturalistic input and measure its responses, studies exploring the function of the visual system traditionally adopt parametric sets of artificial stimuli such as spots of light and sinusoidal gratings. There are many benefits of artificial stimuli (Rust & Movshon, 2005), and their use has been paramount in elucidating the neural underpinnings of sensory processing, such as the orientation selectivity of cortical neurons (Hubel & Wiesel, 1962). The ability to readily parameterize these stimuli allows the measurement of changes in the neural response when a parameter, e.g. orientation of a bar, is changed, producing the tuning curves seen earlier. However, the non-linearity of visual processing means the responses to natural stimuli often cannot be predicted by the sum of responses to artificial stimulus components. Therefore, the use of artificial stimuli is limited in predicting how a neuron, or population of neurons, will respond in the face of naturalistic input.

In comparison to simple parametric stimuli, the use of naturalistic stimuli requires fewer assumptions about relevant stimulus parameters. Physiologically, the use of naturalistic stimuli has validated the receptive field properties discovered by artificial stimulation in the LGN (Dan, Atick,

& Reid, 1996) and primary visual cortex (Prenger, Wu, David, & Gallant, 2004; Ringach, Hawken, & Shapley, 2002), as well as addressing the functional significance of a range of response properties previously characterized with artificial stimuli. In addition, certain neural properties have become apparent only under naturalistic stimulation. For example, a model of receptive field properties of V1 neurons based on recorded responses to naturalistic stimuli was more than twice as accurate at predicting neural responses to naturalistic stimuli than a model created using artificial stimuli (David, 2004). The model based on naturalistic stimuli helped to reveal a functionally relevant response component that would have otherwise remained undiscovered using artificial stimuli.

In addition, the use of naturalistic stimuli in psychophysical and neuroimaging research has increased, furthering the understanding of many visual processing domains. Examples include the differential processing of naturalistic images with regards to saliency and attention (Zipser, 2017), modelling how V1 processes differences in naturalistic images (To, 2010) and dynamic movies (To, Gilchrist, & Tolhurst, 2015), and optimised discrimination for images that have naturalistic 2nd order statistics (Parraga, Troscianko, & Tolhurst, 2000).

The study of visual processing in response to naturalistic stimuli therefore suggests the brain is optimally adapted to real-world everyday scenes. In support of this, humans are better at classifying complex, naturalistic stimuli compared to geometric shapes (Li, VanRullen, Koch, & Perona, 2002) and can process complex images very quickly (VanRullen & Thorpe, 2001) and in parallel (Rousselet, Fabre-Thorpe, & Thorpe, 2002). Therefore, the experiments in this thesis aim to explore the temporal aspects of visual processing, specifically adaptation, to naturalistic stimuli.

1.3.3. A case for using naturalistic stimuli to study adaptation

The goal in the study of visual adaptation is to understand how and why the visual system adapts to stimuli encountered in our natural environment. Achieving this goal is suggested to require both knowledge of the characteristics of natural stimuli, which we have explored above, and knowledge of the response properties of sensory neurons under naturalistic stimulation, which remains a promising area of study (Felsen & Dan, 2005).

The study of adaptation can undoubtedly be benefitted through use of naturalistic stimuli. If adaptation serves to improve the processing of naturalistic information, then it seems essential to study adaptation's processes in response to naturalistic stimulation. There are several differences between the stimuli used in traditional lab-based experiments and the stimuli the visual system has evolved to process. These differences may underlie the lack of success in identifying the functional benefits of adaptation throughout the visual hierarchy.

Firstly, adaptation is not limited to low-level spatial features such as contrast and orientation, as demonstrated by the pervasiveness of adaptation across stimulus domains and throughout the visual hierarchy. The use of more spatially complex adapting stimuli has increased, but they are still often presented in isolation. Under natural viewing conditions, the visual system must adapt to a host of different stimuli simultaneously, and adaptation should be studied with such complex inputs.

Secondly, traditional studies still adopt temporally simplistic stimuli that are not representative of the complex spatio-temporal structure of naturalistic visual input. Visual signals in the real world have complex temporal dynamics e.g. dynamically changing images in a single location (Weiss & Adelson, 1998), or an entirely new scene due to eye movements. Given that adaptation is defined as a temporal contextual effect, it is likely

that the temporal dynamics of both previous visual experience and of the adapting stimulus itself may alter the properties of adaptation.

Thirdly, traditional studies distinguished adaptation from other forms of plasticity as operating over short, seconds to minutes, timescales. We now know this is not the case, with adaptation already measured over days (Haak et al., 2014), and weeks (Delahunt et al., 2004), and the acceptance that our 'baseline' state is in fact permanently adapted. Adaptation should continue to be explored over timescales representative of change in the natural world. Of interest would be to identify whether it is possible to permanently alter the system when presented with a change in stimulus statistics that appears permanent. For example, how long would it take to see an alteration in oblique effects when exposed to an environment where inter-cardinal orientations dominate over cardinal orientations?

Finally, in traditional studies of adaptation, the observer plays an entirely passive role, where even eye movements are eliminated due to fixation. In the natural world, the observer is interacting in the environment to which they are constantly adapting. The neuronal and perceptual effects of passive adaptation may be entirely different from those induced in active viewing of a scene. For example, the negative aftereffect induced by adaptation to optic flow is known to be nulled when the observer is moving in the direction that would cause such optic flow patterns (Harris & Still, 1981). Supporting the notion of calibration between the senses is the finding that manual control over the visual stimulus in an audio-visual adaptation paradigm reduces the adaptation aftereffect (Ichikawa & Masakura, 2006). Adaptation should therefore be studied in response to free-viewing of naturalistic time-varying input.

In addition to passively observing the stimuli, in most psychophysical tasks there is no goal in the task the observer is completing. When we interact with our visual environment we are usually completing some goal and studying visual processing during goal-directed tasks may further reveal

any functional benefits. For example, colour constancy is starting to be explored with paradigms closer to naturalistic situations where constancy is paramount to achieving a goal (Radonjić, Cottaris, & Brainard, 2015). It is therefore important to study adaptation during active viewing conditions that are representative of the input the visual system is faced with in the natural environment; the advancement of mobile head mounted displays for altered and virtual reality is likely to help achieve this.

These four main differences in the visual input an observer receives in lab-based experiments versus naturalistic viewing provide abundant avenues for the study of adaptation in response to naturalistic stimulation. Exploring adaptation in response to spatially and temporally complex stimuli across a wide range of timescales and interactive tasks may shed light on the function of adaptation throughout the visual hierarchy. The use of naturalistic stimuli in the study of adaptation has therefore become more prominent and will be adopted in the experiments in this thesis.

1.3.4. Adaptation to naturalistic stimuli

The advancement in understanding the statistical structure of natural scenes has led to the study of adaptation with more naturalistic stimuli. Until recently, visual adaptation had been studied using simple, artificial stimuli, presented for prolonged durations of time. This approach is still used today and has undoubtedly led to a greater understanding of the mechanisms of adaptation but lacks the ability to extrapolate to how the visual system adapts to stimuli encountered in the natural environment.

In addition, exploring visual adaptation to naturalistic image statistics may help to further understand the functional benefits of adaptation that have been evasive when using artificial stimuli. A promising assumption is that visual adaptation should reflect the statistics to which the visual system is tuned, namely the statistics of natural scenes over space and time (Geisler,

2008; Rieke & Rudd, 2009b; Schwartz, Hsu, & Dayan, 2007; Simoncelli & Olshausen, 2001; Wark, Lundstrom, & Fairhall, 2007). The implementation of inherently complex naturalistic images for adapting stimuli is often avoided due to difficulty in parameterising stimuli to create controlled experiments (Solomon & Kohn, 2014). However, there have been several recent experiments using static natural images in place of simple artificial stimuli.

The tilt aftereffect has been successfully demonstrated with static natural images, selected for their dominant orientation (Dekel & Sagi, 2015b), and with images of carpentered scenes such as buildings (Hussain Ismail, Solomon, Hansard, & Mareschal, 2016), albeit with a smaller magnitude than that shown to artificial stimuli. The latter experiment also found that the TAE for natural images occurs irrespective of stimulus size and spatial frequency, suggesting adaptation beyond low-level mechanisms. Interestingly, this study also revealed negative aftereffects for large angular differences between adapter and test stimuli, in direct opposition to the attractive effects seen with artificial stimuli (Gibson & Radner, 1937). The authors suggest the existence of mechanisms selective for the orientations of images with higher order structure, but do not propose a neural locus. The use of naturalistic stimuli has therefore already highlighted subtle differences in orientation aftereffects for natural and artificial stimuli that may be important in further exploring the functional benefits of adaptation.

In addition, the ubiquity of adaptation was recently highlighted in a medical context, when observers were revealed to adapt to radiological images of mammograms (Kompaniez, Abbey, Boone, & Webster, 2013). Images were altered along the dimension of tissue density from 'dense' to 'fatty'. Adaptation to mammograms showing dense tissue caused subsequent neutral mammograms to appear fattier. Given that radiologists must classify and interpret these images by visual inspection, which can sometimes lead to diagnosis, it is important to take into consideration the effects of adaptation on subsequent decision making. The study of

adaptation to naturalistic images may therefore reveal situations in which adaptation can be both beneficial and detrimental to everyday tasks, and important contexts in which the consequences of adaptation must be carefully considered.

The use of static images is still far removed from the temporally complex input the visual system receives during natural viewing. Recent advances in altered reality technology, however, have allowed the exploration of adaptation to naturalistic time-varying images at never-before studied timescales. Four day's exposure to a visual environment with reduced energy at vertical orientations revealed an attractive tilt aftereffect that peaked in magnitude during the first day of adaptation, but continued to grow with decreased rate for the remaining days (Haak et al., 2014). Studying adaptation over these timescales revealed interesting dynamics of adaptation growth that would not have been possible with artificial stimuli. Again, properties of adaptation that have thus far gone unknown may be discovered when the visual system is exposed to the naturalistic input it has evolved to process.

Haak et al.'s (2014) novel findings on the rate of adaptation growth at long timescales suggest that the use of naturalistic time-varying stimuli may be useful in further exploring what sets the rate of adaptation's time course. One hypothesis, outlined earlier, was that the rate of adaptation growth might be matched to the rate of recently encountered environmental change. Naturalistic time-varying stimuli would allow the presentation of a variety of both typical and atypical temporal environments, for extended durations, following which changes in the rate of adaptation might be observed. This methodology is implemented in Chapter 4 to explore one of the fundamental aims of this thesis.

Overall, the study of adaptation in response to naturalistic stimulation may serve to both corroborate findings from artificial stimuli and elucidate

novel properties and functions of adaptation that have thus far gone undiscovered.

1.4. INTRODUCTION TO THE EXPERIMENTS IN THIS THESIS

The main aim of this thesis is to explore what sets the rate of visual adaptation, using selective adaptation to orientation and measuring the tilt aftereffect (TAE) as a proxy for visual adaptation in general. To identify changes in the rate of adaptation, the time course of the TAE will need to be sampled with fine temporal resolution. The first two chapters of this thesis therefore focus on the methodology of measuring the time course of the TAE. Following the successful implementation of an appropriate methodology to measure the time course of the TAE, subsequent chapters aim to explore what sets the rate of adaptation's time course, using both naturalistic and artificial stimuli. From the empirical findings of these chapters, hypotheses about the number and nature of adapting mechanisms will be explored.

The main questions addressed by this thesis are:

i) How can the time course of the tilt aftereffect be measured at a single trial temporal resolution?
ii) Is the time course of the tilt aftereffect fixed or can it be flexibly altered by the rate of change in the environment?
iii) What are the characteristics of the mechanism(s) that govern the growth and decay of the tilt aftereffect?

2

General Methods

An overview of the methodologies common to the experiments in this thesis is provided in the following chapter. Additional methods sections detail any specific methodological manipulations used in individual experiments.

2.1 APPARATUS

A 22" (50.8cm) CRT monitor (LaCie Electron22blueIV, 1152x870 pixels) with a 75Hz refresh rate (1 frame = 13.33ms) was used for visual stimulus presentation in all experiments. Observers collected data in a darkened room with their head stabilised in a chin rest at a distance of 76.1 cm from the monitor (1 pixel = 0.026 degrees of visual angle). Experimental tasks were programmed in Python using PsychoPy software (Peirce, 2007, 2009). The auditory counterpart to presented video stimuli was played through Sennheiser 25 HD light headphones.

2.1.1 Monitor luminance calibration

The majority of CRT monitors have a non-linear relationship between input voltage and output intensity, which must be corrected for (Rodieck, 1983). The voltage input signal to the monitor must be 'gamma corrected' to ensure the output of the monitor is linearly related to the voltage input.

To calculate the exact gamma value (g) of the monitor, a photometer (Konica Minolta, LS100) was used to measure the luminance of a 500x500 pixel test patch. The output value of all three guns (red, green, and blue) was changed to 21 linearly spaced values between 0 and 255. At each, the

luminance of the patch was measured with an uncorrected look-up table (LUT) loaded. From these luminance values it was possible to generate a gamma function for the monitor (Figure 2.1, red circles). Equation 2.1 was fitted to these data:

$$y = ax^g + b \qquad \text{Equation 2.1}$$

Where y is the luminance, a is a constant, x is the gun intensity requested, g is the gamma, and b is the minimum luminance. The value of x was an integer ranging between 0 and 255. Parameters a, g, and b were free to vary during fitting.

Over repeated calibrations, the luminance gamma value of the CRT monitor ranged between 2.52 and 2.62. The value at the time of measurement was input to the stimulus presentation software, PsychoPy (Peirce, 2007, 2009), which generated a gamma corrected look up table.

To check this, another 21 measurements were taken using the photometer following gamma correction. An example of the linear relationship between the input and output luminance after the gamma correction is shown in Figure 2.1 (black).

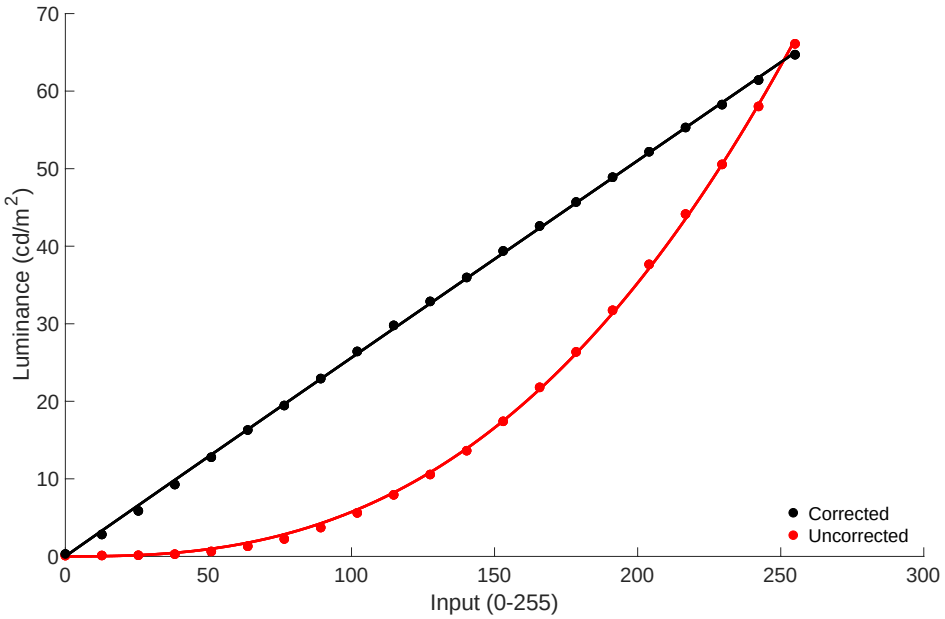


Figure 2.1 Example gamma functions for the LaCie Electron 22 Blue CRT monitor both before (red) and after (black) gamma correction. The r^2 value for each function fitted to the data was 0.99 or greater.

2.1.2 Monitor temporal precision

To present visual stimuli in a temporally precise manner, they were displayed for a number of video frames. As the monitor had a frame rate of 75Hz, each frame equates to 13.3ms.

'Dropped frames' occur when a stimulus cannot be updated within one screen refresh, they result in the stimulus being presented for an extra frame, and thus imprecise stimulus duration. Recording the duration of each frame during the presentation of a temporally varying stimulus allows the identification of any dropped frames. Figure 2.2 shows the duration of frames during the presentation of a 15Hz sinusoidally modulating Gaussian noise patch on the monitor used in all experiments. The orientation and contrast of this patch were updated on each frame.

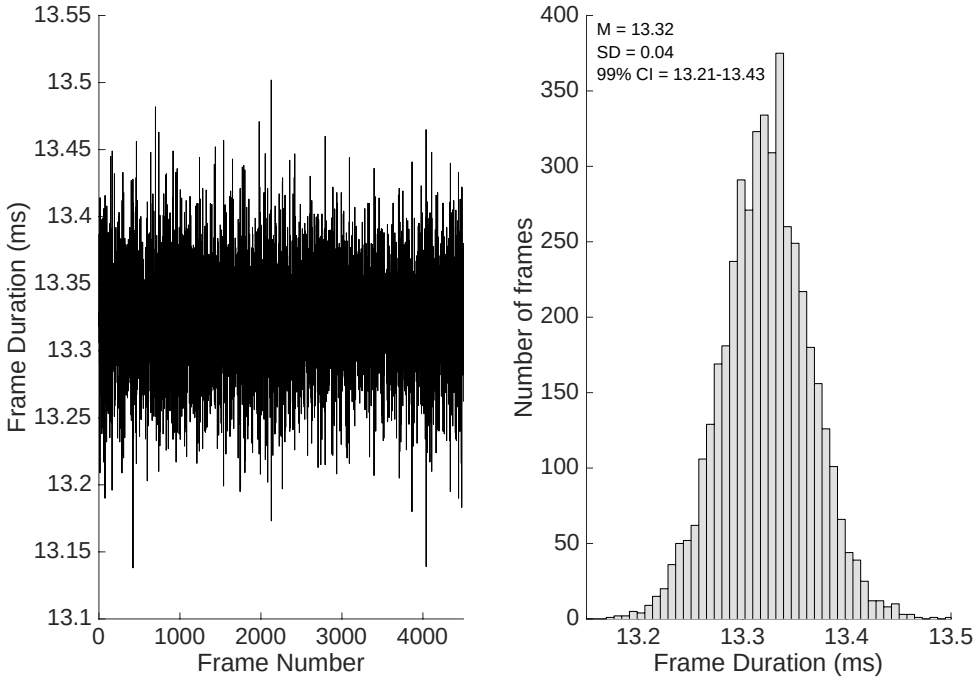


Figure 2.2. The duration of frames during the 1-minute presentation of an updating stimulus. Right: Histogram of recorded frame durations.

The mean frame duration was 13.32ms (SD = 0.04ms) with a 99% confidence interval of 13.21-13.43ms. There were no dropped frames during stimulus presentation, thus a full screen retrace was completed within each frame.

Frame durations were monitored in this manner during the development of all experimental programs to allow for the rectification of any programs producing dropped frames prior to testing.

2.2 STIMULI

Adapting and test stimuli were oriented Gabor patches (sinusoidal gratings windowed by a Gaussian envelope, see Figure 2.3, left) with a spatial frequency of 0.77 cycles per degree, 0.8 Michelson contrast, and a size of 6 times the Gaussian envelope standard deviation (sd = 58.3pix /

1.52dva). Adapting stimuli were vertically oriented and counter phased at a rate of 3Hz. Test stimuli were static and oriented ±18-22 degrees from vertical in one-degree increments. The mask stimulus, used to prevent retinal afterimages of the test stimulus, was a patch of normalised Gaussian noise of the same size and location as the test stimulus (see Figure 2.3, right). The response stimulus was a 0.039dva-wide white line that originated at the central fixation point and terminated at the current position of the mouse cursor.

Adapting stimuli were presented either side of fixation at 9.1dva eccentricity. The test stimulus was randomly presented to either the left or right of fixation, at the same eccentricity, on an equal number of trials per side. In contrast to the adapting and test stimuli, the response stimulus was presented at fixation; due to the spatial specificity of the TAE (Knapen, Rolfs, Wexler, & Cavanagh, 2011; Mathôt & Theeuwes, 2013), this prevented the response stimulus being contaminated by adaptation.

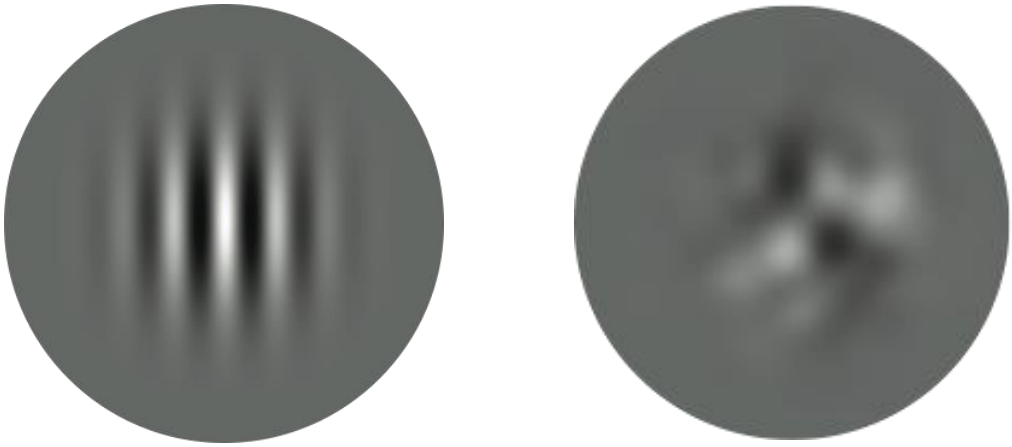


Figure 2.3. Left: Gabor stimulus (sinusoidal grating windowed by a Gaussian envelope). Right: Mask stimulus (normalised white Gaussian noise)

2.3 REPRODUCTION TASK

Biases in the perception of oriented test stimuli were measured using the psychophysical method of reproduction. The method of reproduction was chosen to allow the measurement of the growth and decay of the TAE on a trial-by-trial timescale. The most robust, bias free estimate of the TAE would be gained using a combination of the method of constant stimuli for stimulus presentation and a two-alternative forced choice response. This method, however, requires collation of responses across multiple test trials, making it unsuitable for the measurement of perceptual bias on a trial by trial timescale. To gain an estimate of the TAE on a single trial, it was therefore necessary to use the method of reproduction.

Following the presentation of a test stimulus, observers were instructed to accurately reproduce the orientation of the stimulus by drawing a line on the screen with the mouse. The line could be extended from fixation anywhere above the horizontal meridian (creating angles of -90 to 90 degrees); the length of the line was limited only by the size of the display. Observers pressed the spacebar when they were confident that the orientation of their drawn line accurately reproduced the orientation of the previously presented test stimulus.

2.4 MEASURING THE TIME COURSE OF THE TILT AFTEREFFECT

Throughout this thesis 'time course' will be used to refer to the growth and decay of the tilt aftereffect (TAE); this may sometimes be accompanied by a period of pre-adaptation to show baseline orientation perception. To measure the time course of the TAE, there were two types of trial: adaptation and no-adaptation. In the pre-adaptation phase, a number of no-adaptation trials could be presented to measure baseline orientation perception. Following this, a number of adaptation trials allowed measurement of the growth of the TAE. Subsequently, the decay

of the TAE was measured by a number of no-adaptation trials. Figure 2.4 represents the trial structure used throughout this thesis.

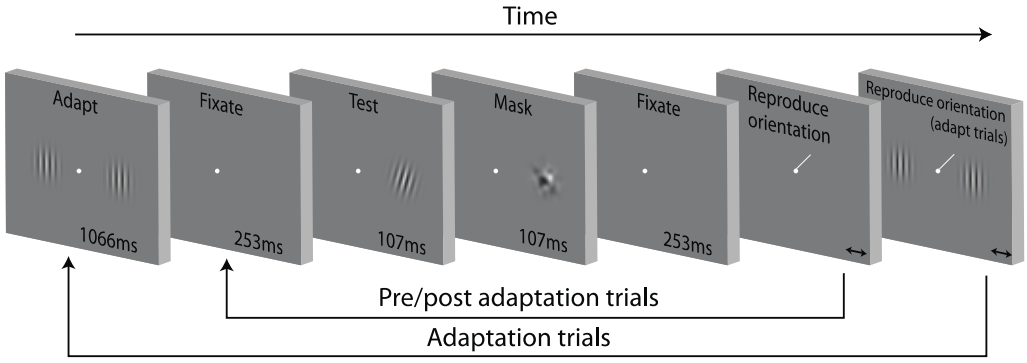


Figure 2.4. Stimulus timings for adapt and no-adapt trials on the orientation reproduction task used to measure the time course of the tilt aftereffect (TAE). Subjects completed a number of pre/post adaptation trials either side of a number of adaptation trials. During adaptation trials the adapting stimulus was also present during the response period (last two tiles of figure represent the response period during no-adaptation and adaptation trials respectively).

In all experiments, observers completed 5 runs of each experimental condition. To prevent practice effects, each experimental session consisted of 1 of each experimental condition in a randomised order.

2.5 ANALYSIS

Adaptation-induced perceptual bias (θ) was calculated as the angular difference between the presented orientation of the test stimulus and the observer's reproduced orientation (see Figure 2.5). Observers completed

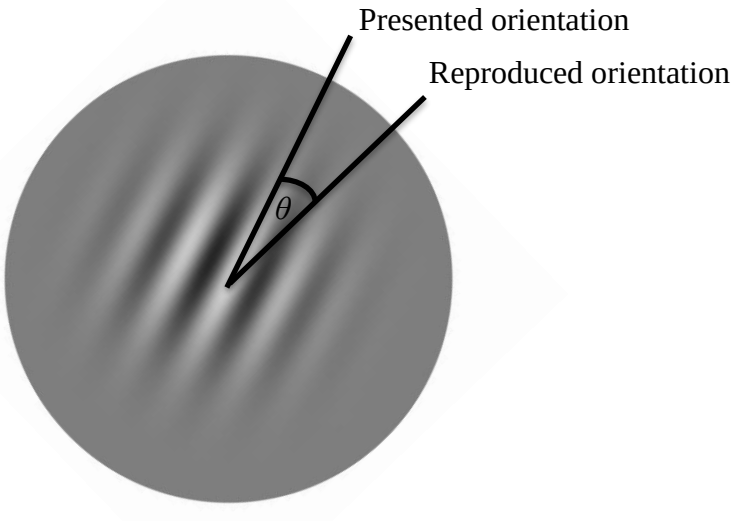


Figure 2.5. Adaptation-induced perceptual bias (θ) calculated as the angular difference between the orientation of the presented test stimulus and the observer's reproduced orientation.

5 runs of each experimental condition. As is traditional in most psychophysical experiments, large quantities of data were collected for a small number of observers. Data was pooled across observers to gain a robust measurement of the TAE at each time point. This has implications for statistical generalization, as small samples may be considered non-representative of the general population. However, it is assumed that the low-level visual processes explored in this thesis are common and relatively stable across all human observers, allowing for the generalisation of statistical results to the wider population.

2.5.1 Outlier removal

Observers inevitably showed lapses in attention and occasionally made responses that were unrelated to the orientation of the test stimulus, introducing extraneous variability into the data set. These outliers were removed prior to calculating descriptive statistics. As recommended by Leys, Ley, Klein, Bernard, & Licata (2013), perceptual bias values at each time point that were more 2.5 median absolute deviation (MAD) units away from the distribution median were removed. A criterion of 2.5 is a moderately conservative approach according to Miller (1991).

To calculate the MAD, the median perceptual bias at each time point is subtracted from each value in the distribution. The median of the resulting distribution is calculated and multiplied by 1.4826 (a constant used assuming normally distributed data) to give the MAD (see Equation 2.2).

$$MAD = b\, M_i\left(\left|x_i - M_j(x_j)\right|\right) \qquad \text{Equation 2.2}$$

Where b is a constant, x_i is each value in the original distribution, and M_j is the median of the original distribution.

2.5.2 Descriptive statistics

The mean perceptual bias at each point in the time course was calculated, collapsing across observers and runs. 95% confidence intervals were calculated using non-parametric bootstrapping. The bootstrapping approach readily allows computation of confidence intervals for subsequently calculated composite metrics, such as baseline corrected biases and parameters of any functions fitted to the data.

The original distribution of values at each time point was resampled with replacement to create a new distribution of equal size from which a new, resampled, mean was calculated. This procedure was repeated 1000 times to produce a distribution of resampled means at each time point (see Figure 2.6), from which the 2.5th and 97.5th percentiles were taken to give a 95% confidence interval.

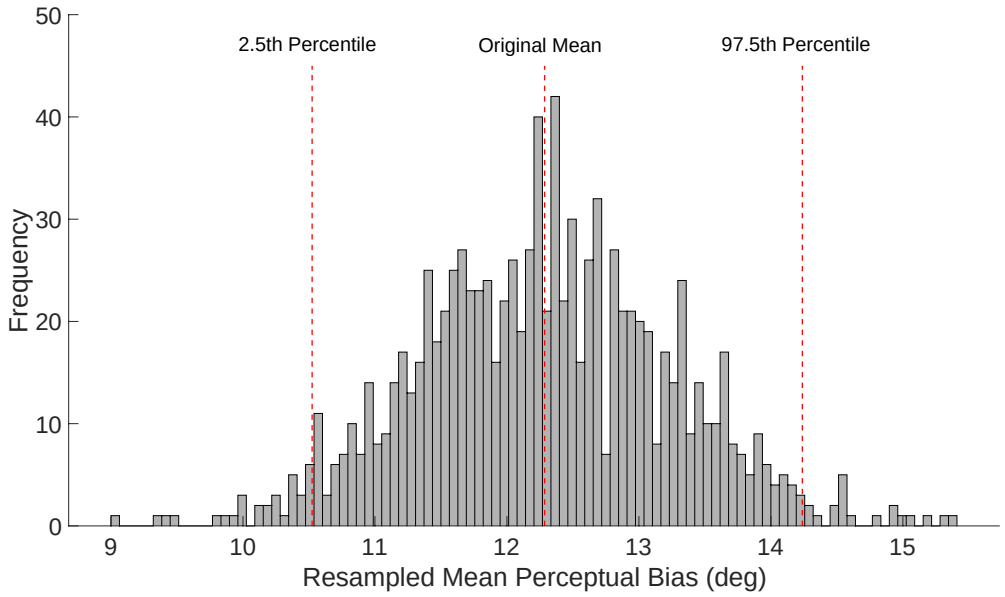


Figure 2.6. Histogram of 1000 resampled mean bias values from a single time point in a representative tilt aftereffect time course. Red dashed lines mark the original mean and the 2.5th and 97.5th percentiles used as 95% confidence intervals.

As will be discussed in Chapter 3 there are inherent biases in the reproduction of oriented stimuli. These biases are seen in baseline

orientation reproduction trials, with observer's reproduced orientations being consistently tilted ~5-10 degrees further away from vertical than the presented stimulus. These biases are consistent within observers and across conditions. For graphical simplicity, the time courses presented in this thesis have been 'baseline corrected' so that pre-adaptation bias lies on zero. This was achieved by subtracting the mean bias during pre-adaptation time points from all time points in the time course. This has the effect of shifting the entire time course downwards, so that pre-adaptation time points lie on zero.

Where differences in the time courses between the conditions are difficult to identify in the raw data, descriptive models have been fit to aid interpretation. The growth of adaptation is well characterised by a power law function (Equation 2.3) (Drew & Abbott, 2006) which was fit to the data by minimising the sum of squared residuals. Goodness of fit was assessed by the root mean square statistic.

$$y(x) = (ax^b) + c \qquad \text{Equation 2.3}$$

2.5.3 Inferential statistics – permutation testing

Non-parametric permutation testing was conducted to identify time points where there was a significant difference in mean perceptual bias between two conditions. Permutation tests followed a similar resampling procedure to the bootstrapping, with the exception that the relevant condition labels were also randomly shuffled. The resulting permutation distributions thus reflect the sampling distribution of the bias difference obtained under the null hypothesis that conditions do not differ. Two-tailed *P* values were derived non-parametrically, by calculating the proportion of permuted differences with an absolute value greater than or equal to the observed value. A criterion of $p<.05$ was used to identify statistical significance.

3

Biases in the reproduction of orientations

3.1. INTRODUCTION

In the following experiments in this thesis, estimates of the tilt aftereffect (TAE) were obtained on a trial-by-trial timescale before, during, and after periods of visual adaptation. Prior to carrying out these studies, a pilot experiment was completed by 3 observers to ensure the time course of the TAE could be measured in this way. Experimental trial structure was identical to that outlined in General Methods (2.4).

Each experimental run consisted of 25 trials (See Figure 2.4) measuring observers' baseline perception of orientation. The presented test stimuli differed by 1 angular degree, which is well above the usual human orientation discrimination threshold of 0.3-0.5deg for cardinal orientations (Burbeck & Regan, 1983), but just below the discrimination threshold for stimuli oriented at 20deg from vertical which is ~1.2deg (Orban, Vandenbussche, & Vogels, 1984). Therefore, it was important to identify whether observers could discriminate the test stimuli and if so, whether they were also capable of accurately reproducing the orientation of the test stimuli. Of specific interest was whether observers could differentially respond to each test stimulus. Identifying this was important in assessing the validity of using trial-by-trial orientation reproduction as an estimate of the TAE in future experiments.

The difference between the orientation of the presented test stimuli and the observers' mean reproduced orientation revealed the accuracy of their responses. When mean reproduced orientation is plotted against test stimulus orientation, data points that lie on the unity line represent

veridical performance. Figure 3.1 shows observers' mean reproduced orientation as a function of presented orientation. The angle of observers' reproduced orientation increased as a function of the angle of the test stimulus orientation. The slope parameters of fitted simple linear regression functions are significantly different to zero (counter clockwise: t(3)=15.17, *p*=.00062; clockwise: t(3)=8.42, *p*=.0035), confirming that observers were making different responses depending on the orientation of the test stimulus.

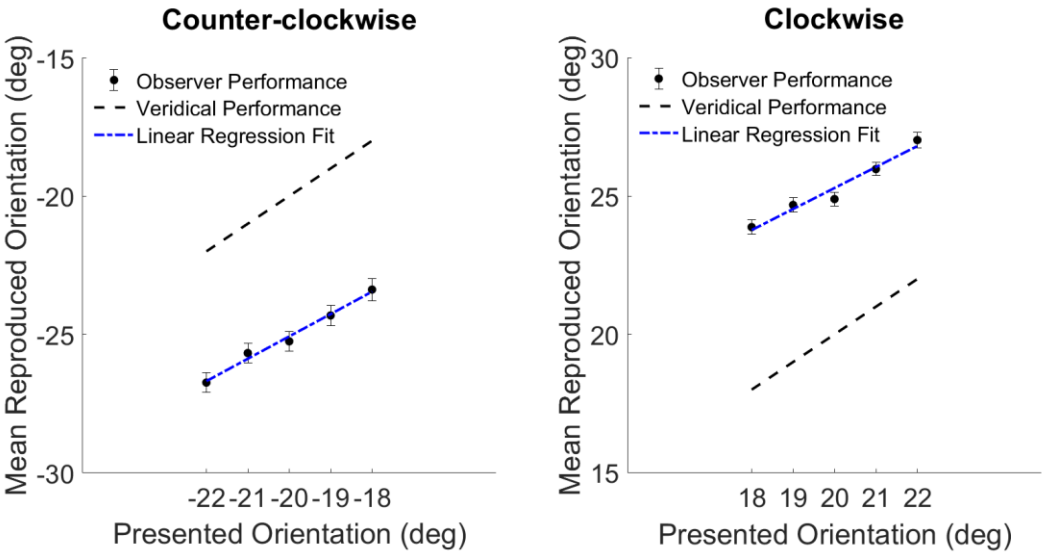


Figure 3.1. Mean reproduced orientation as a function of presented orientation for clockwise (right) and counter-clockwise (left) test stimulus orientations. Fitted linear regression lines: counter-clockwise: y=0.8x–9, R2 = 0.98, P slope= .00062; clockwise: y = 0.8x+10, R2 = 0.95, P slope = .004.

It is also evident that whilst their responses varied systematically with stimulus orientation, they were not accurate. That is, the regression line was shifted downwards/upwards from the unity line for the counter-clockwise and clockwise stimuli respectively. This indicates that observers perceived counter-clockwise and clockwise titled stimuli as being tilted further away from vertical than they were. These shifts occur in the absence of any selective adaptation, and therefore represent an inherent bias in either the perception or reproduction of stimuli oriented ±18-22 degrees

from vertical. This chapter will explore and characterize biases in the reproduction of oriented stimuli in an unadapted state.

Biases in the perception of oriented stimuli in an unadapted state are well documented, with both anisotropies in observers' ability to discriminate between oriented stimuli, as well as biases in their perception of orientation. Observers exhibit better discrimination performance around the cardinal orientations, known commonly as the 'oblique effect' (Appelle, 1972). Oblique effects are thought to result from the anisotropic distribution of oriented contours in the real world, as well as neural anisotropies in orientation processing. Cardinal orientations dominate over oblique orientations in natural scenes (Coppola, Purves, Mccoy, & Purves, 1998; Dakin, Mareschal, & Bex, 2005; Dragoi, Turcu, & Sur, 2001; Hansen & Essock, 2004)), and physiological research (Li, 2003; Wang, Ding, & Yunokuchi, 2003; Xu, Collins, Khaytin, Kaas, & Casagrande, 2006) and functional imaging (Yacoub, Harel, & Ugurbil, 2008) has shown there are more neurons tuned to cardinal than oblique orientations in primary visual cortex, and that these may have sharper tuning curves (Li, 2003; Wang et al., 2003).

These bottom-up anisotropies also appear to affect perceptual bias as well as discrimination performance. Early behavioural studies report that a line tilted slightly away from the horizontal or vertical cardinal axis is perceived as more titled than it actually is (Jastrow, 1892; Smith, 1962). In fact, Wei & Stocker (2016) suggest a simple mathematical relationship between discrimination threshold and perceptual bias. They suggest that perceptual bias $b(\theta)$ (as a function of the stimulus variable) is proportional to the slope of the discrimination threshold $D(\theta)$ squared, shown in Equation 3.1:

$$b(\theta) \propto (D(\theta)^2)' \qquad \text{Equation 3.1}$$

This simple relationship predicts that perceptual bias will be at zero at cardinal orientations (0, 180, 90, -90 degrees) and intercardinal

orientations (45, -45, 135, -135 degrees), and non-zero at all other orientations. This law accurately describes the biases seen in our (and others' - de Gardelle, Kouider, & Sackur, 2010) orientation reproduction data.

The biases seen in the pilot data (Figure 3.1) are also nearly identical to those found in the motion literature on reference repulsion, whereby the perception of the direction of motion of moving dot stimuli is biased away from cardinal directions (Rauber & Treue, 1998). Importantly, the reference from which perception is repelled does not have to be physically present, with biases occurring due to intrinsic reference points such as the cardinals (Rauber & Treue, 1998). It could be that the biases in pre-adaptation orientation reproduction are due to a similar form of reference repulsion. To fully characterise the biases seen in the pilot data, orientation reproduction was measured in the same way as before, but across the entirety of orientation space.

When the test stimulus distribution is drawn from a subset of orientation space, there also exists the possibility of central tendency effects. When asked to reproduce a property of a visual stimulus, with that property being drawn from a fixed distribution, observer's responses are biased towards the mean of the distribution where there forms an 'indifference point' of zero bias (Hollingworth, 1910). Either side of this point, observers underestimate larger values and over-estimate smaller values. This central tendency bias is shown in a range of stimulus domains such as size (Hollingworth, 1910), colour (Olkkonen & Allred, 2014), line length (Ashourian & Loewenstein, 2011) and interval duration (Jazayeri & Shadlen, 2010), but has not currently been reported in orientation. Central tendency biases have been explained as the integration of acquired knowledge about the stimulus distribution with noisy sensory estimates, resulting in maximum reproduction precision (Hollingworth, 1910; Jazayeri & Shadlen, 2010). To investigate the contribution of central

tendency effects towards bias in reproduction orientation, test distributions were drawn that were a subset of orientation space.

Careful selection of the test stimulus distributions allowed for the exploration of the two types of effects discussed above. For example, a 90-degree wide uniform distribution of test stimuli centred on zero (vertical) will produce a pattern of biases where there is zero bias at vertical, and stimuli either side of vertical are biased away from vertical. When this distribution is subsequently shifted 10 degrees clockwise from vertical, reference repulsion and central tendency biases would result in a different pattern of biases. Reference repulsion effects would produce the same pattern of biases as in the original distribution, with a shifted selection of orientations. In contrast central tendency effects would shift the 'indifference point' of zero bias to be the new centre of the distribution, 10 degrees clockwise of vertical. Test stimulus selection can therefore be used to explore the contribution of reference repulsion and central tendency effects on biases in unadapted orientation reproduction.

3.2. METHODS

To measure biases in the reproduction of oriented stimuli across the whole circular variable, the reproduction of oriented Gabor stimuli was measured at orientations drawn from three uniform distributions. Distributions were 180 degrees wide, in 5-degree steps, centred on vertical (0deg), 90 degrees clockwise of vertical (90deg), and 90 degrees counter clockwise of vertical (-90deg). The distributions therefore overlapped, allowing the measurement of the reproduction of the same stimulus orientation when the centre of the distribution was a different cardinal orientation (see Figure 3.3).

To explore the effects of central tendency, a further three narrower distributions of oriented stimuli were used. Distributions were 90-degrees

wide, in steps of 5 degrees, centred on 10deg, 20deg, and 30deg clockwise from vertical.

Trial structure was identical to that of the pre-adaptation trials outlined in the General Methods. An oriented Gabor was presented to the left or right of fixation, after which the observer (n=1) reproduced the orientation by drawing a line on the display with the mouse. Each orientation was presented 5 times in a randomised order across the duration of a run. 4 runs were completed for each distribution of orientations.

3.3. RESULTS

The test stimuli orientations used in the pilot experiment are contained in the 180deg distribution centred on vertical. Figure 3.2 (left) shows the mean reproduced orientation for each presented orientation. The observer's mean reproduced orientation is only veridical for 0deg, -45deg and +45deg oriented test stimuli.

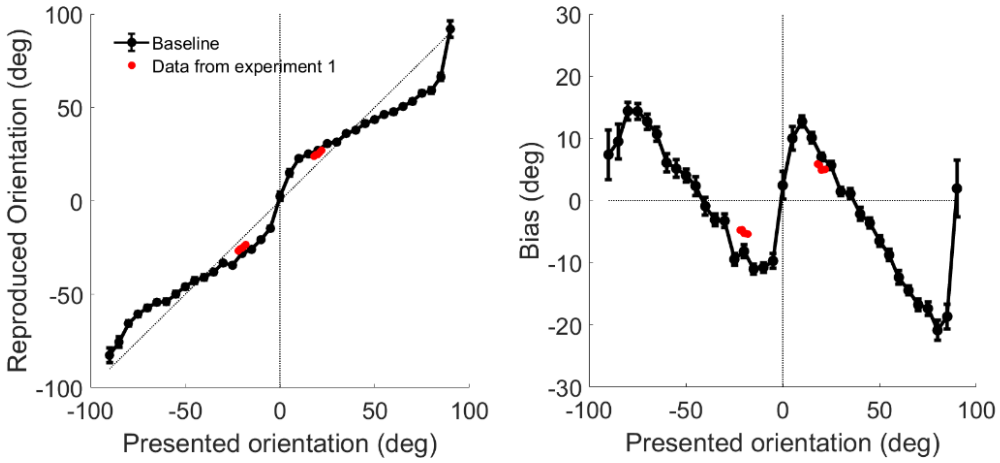


Figure 3.2 Mean reproduced orientation (left) and bias in reproduced orientation (right) plotted as a function of presented orientation. Red data points are the biases shown in the reproduction of test stimulus orientations in experiment one. Error bars represent 1 standard deviation of the mean, data is from the same subject in all graphs.

Orientations between 0deg and 45deg are reproduced as being oriented further clockwise than they are. In comparison, orientations between 45deg and 90deg are reproduced as being more counter-clockwise than they are. This pattern of repulsion from the nearest cardinal is seen across all orientations, with peak repulsion occurring approximately 20deg away from each cardinal.

The pattern of bias as a function of orientation can be seen in Figure 3.2 (right), with positive deviation from zero bias representing more clockwise stimulus reproductions and negative deviations representing more counter-clockwise stimulus reproductions. This clearly shows a repulsion in reproduced orientation away from all four cardinal orientations, with the largest bias for stimuli ~20deg from the cardinal. There is zero bias at inter-cardinal orientations (±45deg, ±135deg) which are equidistant from two cardinal orientations that would produce bias in different directions. Overlaying the original data reveals that the bias shown in the reproduction of ±18-22deg stimuli in the pilot experiment (red) is in alignment with the biases in reproduction of the 180deg wide distribution.

Repulsive biases are present at all cardinal orientations, with nearly identical magnitude and tuning to those centred on vertical (see Figure 3.3). Where distributions overlap, regardless of which cardinal is the centre of the distribution, the magnitude of biases around each cardinal are relatively stable (see overlapping of blue and black, and red and black lines in Figure 3.3).

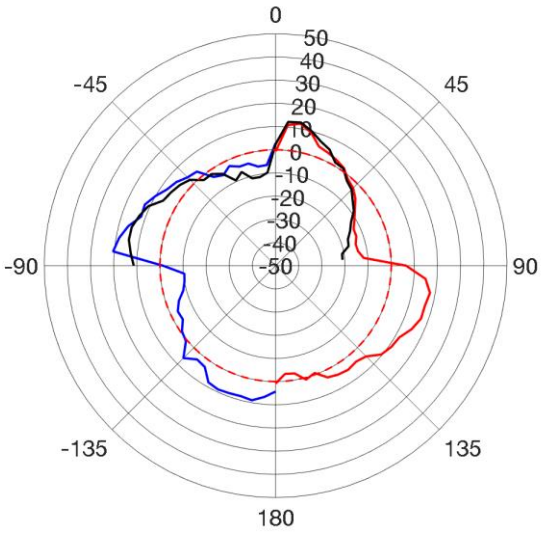


Figure 3.3 Polar plot of bias in orientation reproduction as a function of test stimulus orientation for test stimulus ranges -90:90 (black), 0:-180 (blue), and 0:+180 (red).

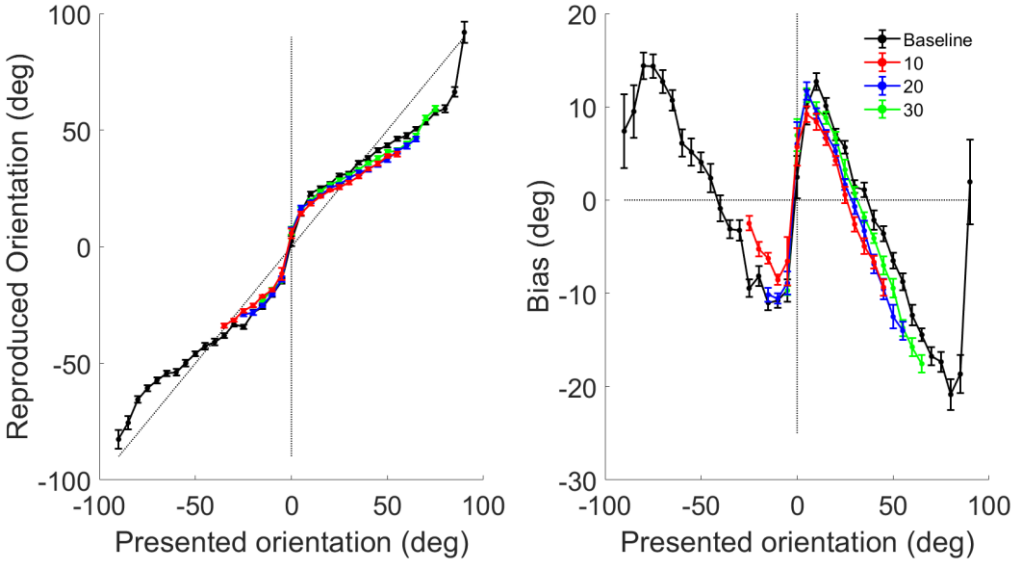


Figure 3.4 Reproduced orientation (left) and bias in reproduction (right) as a function of presented test stimulus orientation for the original 180-degree distribution and three 90-degree distributions centred on 10deg, 20deg, and 30deg.

The data from the three 90-degree wide distributions centred on 10, 20 and 30 degrees clockwise from vertical are plotted in Figure 3.4 alongside the 180-degree wide distribution centred on zero.

In all test stimulus distributions, there is no bias at vertical, regardless of the centre of the distribution. However, in the three narrower distributions, the point at which there is zero bias is shifted from 45 degrees towards the centre of the narrower distribution. A linear equation (Equation 3.2) was fit to the profile of biases observed in each condition across the subset of orientations that were common to all four distributions (+10 to +55 deg).

$$y = mx + c \qquad \text{Equation 3.2}$$

Where *m* represents the gradient of the line, and *c* the intercept. The value of y when x = 0 was taken as the 'indifference point' for each distribution. The indifference points for distributions centred on 10, 20, and 30deg were 26.7, 28.2, and 31.2deg respectively.

3.4. MODELLING BIASES IN THE REPRODUCTION OF ORIENTATIONS

The data suggest two sources of bias. Reference repulsion produces biases in the reproduction of oriented stimuli irrespective of the range of orientations tested. Additionally, central tendency biases are seen when a distribution of orientations is used that spans a subset of orientation space. Biases seen to stimuli in these narrower distributions can therefore be modelled as the summation of reference repulsion effects and central tendency effects.

To characterize the biases observed with the three narrow distributions of test stimulus orientations, a Bayesian observer model was developed. The magnitude of central tendency bias was determined by the weighted sum of the observer's noisy sensory estimates, or likelihood function; and their learned knowledge of the test stimulus distribution, or prior.

Observer's sensory estimates were modelled by a Gaussian distribution with the mean centred on the test stimulus orientation plus the repulsive bias and the standard deviation representing uncertainty in the sensory estimate. The standard deviation was modulated sinusoidally over orientation to account for oblique effects, where variation in sensory estimates is smaller around the cardinal orientations using Equation 3.3:

$$\sigma_s = a + b|sin(2x)| \quad \text{Equation 3.3}$$

Where *a* represents the minimum standard deviation and *b* represents the magnitude of the oblique effect.

Observer's knowledge of the test stimulus distribution was also represented by a Gaussian, centred on the mean of the stimulus distribution. The weighting of the likelihood and prior were calculated using Equations 3.4 and 3.5:

$$w_S = \frac{\frac{1}{{\sigma_S}^2}}{\frac{1}{{\sigma_S}^2}+\frac{1}{{\sigma_p}^2}} \quad \text{Equation 3.4}$$

$$w_p = \frac{\frac{1}{{\sigma_p}^2}}{\frac{1}{{\sigma_S}^2}+\frac{1}{{\sigma_p}^2}} \quad \text{Equation 3.5}$$

The estimated reproduced orientation (θ_r) is a weighted sum of the biased sensory estimate and the prior mean (Equation 3.6).

$$\theta_r = w_s s + w_p \mu \quad \text{Equation 3.6}$$

Whereby s is the biased stimulus orientation from the measured reference repulsion effects, and μ is the mean of the test stimulus distribution.

The Bayesian observer model had three free parameters, the standard deviation of the prior (σ_p), and the two determinants of the standard deviation of the sensory estimate (σ_s), a and b. These parameters were free to vary during the simultaneous fitting of the three narrow distributions. The best fitting model was that that minimised the squared error between the data and the model.

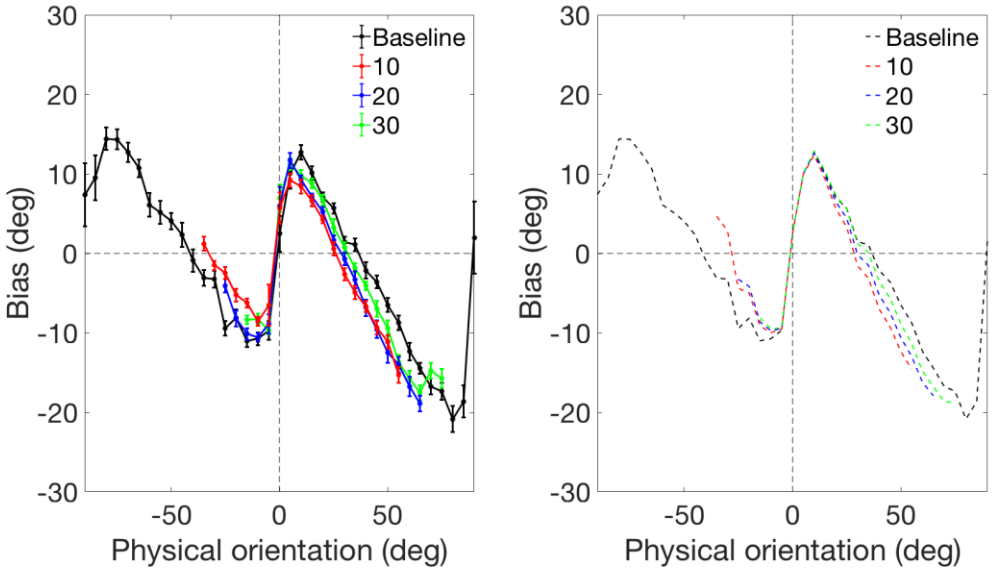


Figure 3.5. Left: bias in reproduction as a function of presented test stimulus orientation for the original 180-degree distribution and three 90-degree distributions centred on 10deg, 20deg, and 30deg. Right: Bayesian observer model prediction

Figure 3.5 shows the pattern of biases in each test stimulus distribution (left) and the model's predicted pattern of biases (right). The model accurately reproduces the shifts in the pattern of biases seen when the centre of the distribution is altered ($r^2 = 0.89$). The best fitting model had a prior standard deviation of 1.45 degrees, and a sinusoidally modulating sensory estimate standard deviation that varied between 0.0008 and 0.68 degrees. Although these model parameters result in predictions that closely match the measured data, the standard deviation of the prior is so narrow that it makes little psychological sense. If true, observers would have a prior that was very narrowly focused on the mean of the stimulus distribution.

Therefore, further modelling of these biases is recommended to identify whether the biases in pre-adaptation orientation reproduction shown in the pilot experiment can be modelled as a combination of reference repulsion and central tendency effects.

3.5. INDIVIDUAL DIFFERENCES OF BIASES IN THE REPRODUCTION OF ORIENTATION

To assess individual differences in these biases, the same data were collected for three additional observers who had completed the pilot experiment. Each observer completed four runs of the 4 test stimulus distributions described above. These observers were naïve to the hypotheses of the experiment.

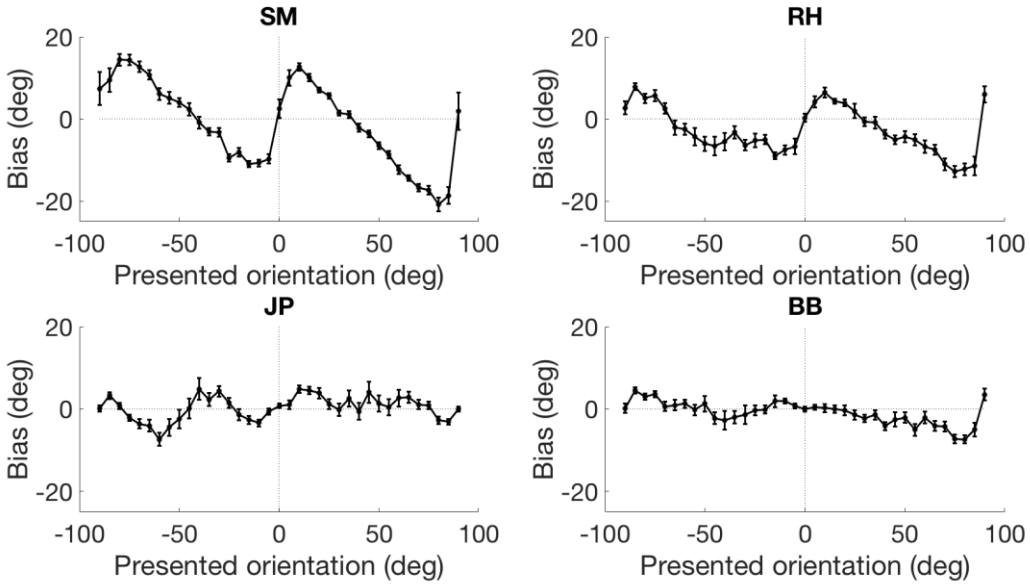


Figure 3.6. Observers' bias in the reproduction of test stimuli oriented -90 to 90 degrees. All four observers completed experiment one.

Figure 3.6 shows the pattern of biases for the -90:90deg range of test stimuli. Observer RH shows a similar pattern of biases to the original observer (SM). Observer BB shows similar biases for stimuli at the extremes of the distribution but shows relatively veridical reproduction of orientations ±20deg from vertical. These three observers' data is consistent with reference repulsion from the cardinal orientations.

Interestingly, observer JP shows a remarkably different pattern of biases. Rather than showing a 90deg period of modulation (biases away from the cardinal orientations), JP shows a 45deg period – a pattern consistent with reference repulsion from both the cardinal and intercardinal (±45deg) orientations. JP is an experienced observer who has completed >1000 trials of direction of motion discrimination around a ±45deg reference point. This previous experience could explain JP's use of intercardinal orientations as reference points, and thus the repulsion from these references seen in their data.

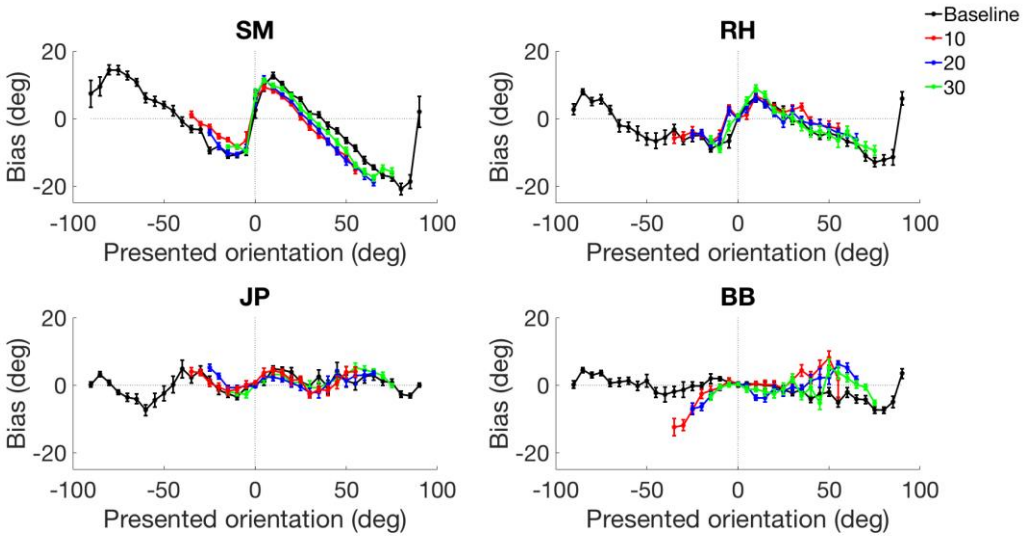


Figure 3.7. Observer's bias in the reproduction of test stimuli drawn from the original -90 to 90deg baseline distribution and three narrower 90deg wide distributions centred on 10deg 20deg and 30deg clockwise from vertical.

Figure 3.7 shows the same observer's data for the three narrower distributions of test stimuli centred on 10deg, 20deg, and 30deg. In contrast to the original observer (SM) the other three observers do not appear to show any shifts in the indifference point towards the centre of each distribution.

3.6. DISCUSSION

A clear pattern of biases is seen in the reproduction of orientations around the clock. Oblique orientations are consistently reproduced as being repelled away from the nearest cardinal, whilst the reproduction of cardinal and inter-cardinal orientations remains unbiased. For narrower distributions, this pattern of biases is altered by the centre of the test stimulus distribution. These central tendency effects manifest as shifts in the 'indifference point' towards the centre of the distribution of test stimulus orientations.

These biases in orientation reproduction are well described by reference repulsion, whereby stimuli close to a reference point, whether it be physically present or absent, are repelled away from the reference (Rauber & Treue, 1998). This is reported most commonly in direction of motion but is directly applicable to orientation. Here we see the reproduction of Gabor stimuli orientated close to implicit cardinal reference points being repelled away from these orientations. However, the specific distribution of test stimulus orientations also appears to affect the pattern of biases shown. When the centre of the distribution was systematically manipulated, the point at which there is zero bias was shifted in the same direction as the manipulation. This result suggests that central tendency affects (Hollingworth, 1910) influence the biases in the reproduction of orientations seen pre-adaptation.

A Bayesian observer model corroborated the assumption from the empirical data that biases seen in orientation reproduction are a driven by a combination of reference repulsion and central tendency effects can explain 89% of the variance in the data when the centre of the test stimulus distribution was manipulated. In all subsequent experiments in this thesis, test stimuli were oriented ±18-22deg. From the findings above, biases seen in the reproduction of these stimuli are likely a combination of both reference repulsion and central tendency effects.

It is important to note that although reference repulsion effects were demonstrated across all observers (with varying magnitude and pattern), central tendency biases were less clear in the three additional observers. Firstly, these observers completed relatively few trials per point (n=20), resulting in noisy data in which may obscure small shifts. Secondly a small number of trials means observers have more uncertainty regarding the stimulus distribution, leading to a larger standard deviation of the prior. Observers likely relied less on their prior knowledge of the distribution of stimulus orientations, and more on their sensory estimates of the stimulus, reducing central tendency effects.

It is important to note that reference repulsion effects occur due to implicit cardinal reference points, whilst central tendency effects occur due to the selected distribution of test stimuli. Neither of these factors are altered during the adaptation phase of time course measurement. It is therefore assumed that the biases seen in pre-adaptation orientation reproduction remain constant throughout the measurement of the rest of the time course. For this reason, in all future experiments, observers' average pre-adaptation bias in orientation reproduction is subtracted from the biases measured throughout the rest of the time course. This has the effect of shifting the entire time course downwards, so that bias in pre-adaptation trials rests at zero, and is referred to in future as 'baseline correction'.

Baseline correction in all further experiments was completed for each observer individually as there appears to be significant variation between observers in both the pattern and magnitude of these biases. This variation is likely due to observers' prior psychophysical experience. Observers such as JAP, practiced in making direction of motion judgements around intercardinal orientations, develop strong internal reference points that drive reference repulsion at orientations in addition to the cardinals.

The effect of prior experience and environmental distribution of orientations would be an interesting avenue for future research. A plausible experiment would be to expose observers, through a virtual reality headset, to an altered distribution of orientations in natural viewing. A virtual environment could be created in which intercardinal, rather than cardinal, orientations dominate the distribution. From the results seen here, biases in subsequent orientation reproduction might show reference repulsion around the intercardinal, as well as cardinal orientations (as seen in observer JAP). This would be due to observers developing implicit reference points at the now-dominant inter-cardinal orientations. The timescales over which 'new' dominant orientations are learned would be interesting to examine, as well as how the relative strength of biases at each set of orientations is altered with increasing exposure to the 'new' cardinals.

4

Interrupted measurement of the tilt aftereffect

4.1. INTRODUCTION

Traditional measurement of the tilt aftereffect (TAE) involves the continuous presentation of an adapting stimulus for a prolonged period, followed (after variable delay) by the measurement of the observer's bias in orientation perception. To measure the time course of the TAE, observer's bias, quantified as changes to the point of subjective equality, is measured following increasing durations of continuous adaptation. Measurement in this manner has revealed that the magnitude of the tilt aftereffect increases logarithmically as a function of adapting duration, and decays logarithmically following adaptation offset (Magnussen & Johnsen, 1986).

The experiments in this thesis aimed to identify how the time course of adaptation to tilt may be altered by previous visual experience. For example, previous visual exposure to different temporal or spatial statistics may lead to changes in the rate and magnitude of the time course of the TAE. To best capture changes in the time course of the TAE, it is fundamental that it can be measured with fine temporal resolution. Measuring the time course of the TAE using the traditional method of continuous periods of adaptation is time consuming and the sampling of the growth and decay of the TAE is temporally coarse; often using only four or five adapting durations (Magnussen & Johnsen, 1986). Subtle changes in the rate and magnitude of the time course of the TAE would be hard to detect using this traditional methodology.

A means of measuring the TAE with greater temporal resolution was needed for the experiments in this thesis. This was achieved by briefly interrupting adaptation with a test stimulus at a fixed rate. Following each interruption, an observer's bias in orientation perception was measured at a different retinal location to the adapt and test stimuli, allowing observers to continue adapting during the response period. This produced a measurement of the time course of the TAE with high temporal fidelity. The temporal scale at which the time course is sampled is determined by the rate of interruption. Collapsing data across multiple runs and observers allowed a robust estimate of the TAE at each interruption. Hereafter, this measurement of the TAE will be referred to as 'interrupted adaptation' and compared to the more traditional method of 'continuous adaptation'.

Measuring the time course of the TAE using interrupted adaptation is not often used (Bao, Engel, Fast, & Mesik, 2013; Mei, Dong, & Bao, 2017). It was therefore of interest to identify whether interrupted adaptation would produce the logarithmic growth and decay of the TAE as previously shown using continuous adaptation, and if not, how the process of interruption altered the time course of the TAE.

The following experiments aimed to explore the effects of interrupting adaptation on the time course of the TAE. The implementation of interrupted adaptation to measure the time course of the TAE required a trial-by-trial estimate of observers' bias in perceived orientation (as outlined in General Methods (2.4)). To do so, observers reproduced the orientation of a presented test stimulus on each trial. Bias was quantified as the angular difference between the orientation of the presented test stimulus and the orientation of the observer's response.

4.2. EXPERIMENT 1: INTERRUPTED ADAPTATION

Experiment one aimed to identify whether interrupting adaptation altered the time course of the TAE in comparison to continuous adaptation, and if so, whether this effect was tuned to the frequency of interruption.

4.2.1. Methods

The experiment measured the time course of the TAE produced by 327.3 seconds (5.45 minutes) of adaptation. This was measured by both traditional continuous adaptation and the interrupted adaptation technique used in the following experiments in this thesis (see General Methods (2.4)).

Traditional continuous adaptation consisted of 5 durations of passive exposure to a fixed orientation: 0.53, 13.2, 63.9, 127.2, and 327.3 seconds. In contrast, in the interrupted adaptation condition, a period of 327.3 seconds of adaptation was interrupted by the measurement of the TAE at one of 4 fixed intervals: 2.53, 4.33, 6.59 and 12.66 seconds. The interruption, which consisted of the presentation of an ISI, test stimulus, mask, and ISI, was 0.76 seconds in duration.

In both conditions, the adaptation period was preceded by 25, and followed by 50, test trials which consisted of the presentation of stimuli with variable orientations and a response period to measure the magnitude of the TAE. Figure 4.1 shows the trial timings in both the continuous and interrupted adaptation conditions.

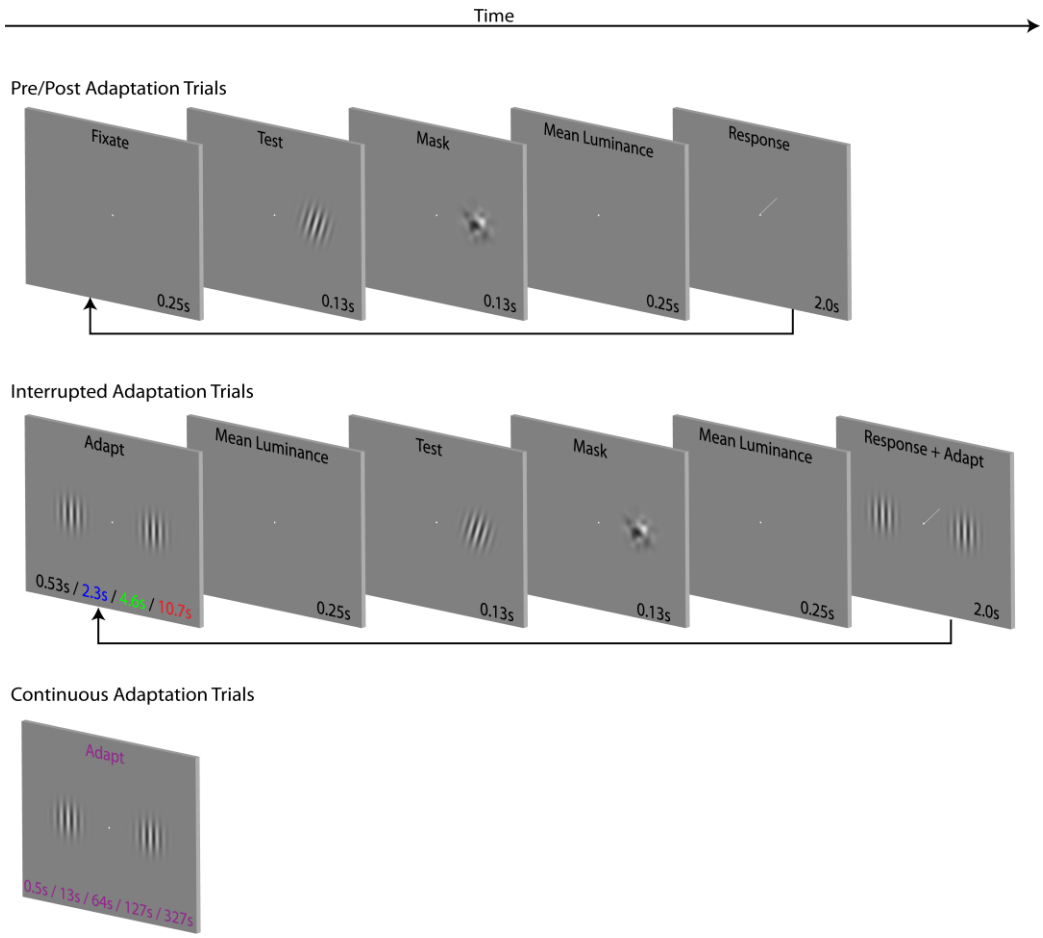


Figure 4.1 Trial timings used in Experiment 1. Top row: stimuli and timings used in the 25 pre-adaptation and 50 post-adaptation trials. Middle row: stimuli and timings used in the interrupted adaptation trials. Bottom row: stimuli used in the continuous periods of adaptation.

4.2.2. Results

The time course of the TAE for each condition was calculated as described in chapter 2.5.2.

Figure 4.2 shows the time course of the TAE produced by both continuous and interrupted adaptation. In all conditions, pre-adaptation onset (60 seconds) the TAE is stable and corrected to zero degrees. Following adaptation onset, the TAE grows as a function of time, and decays after adaptation offset.

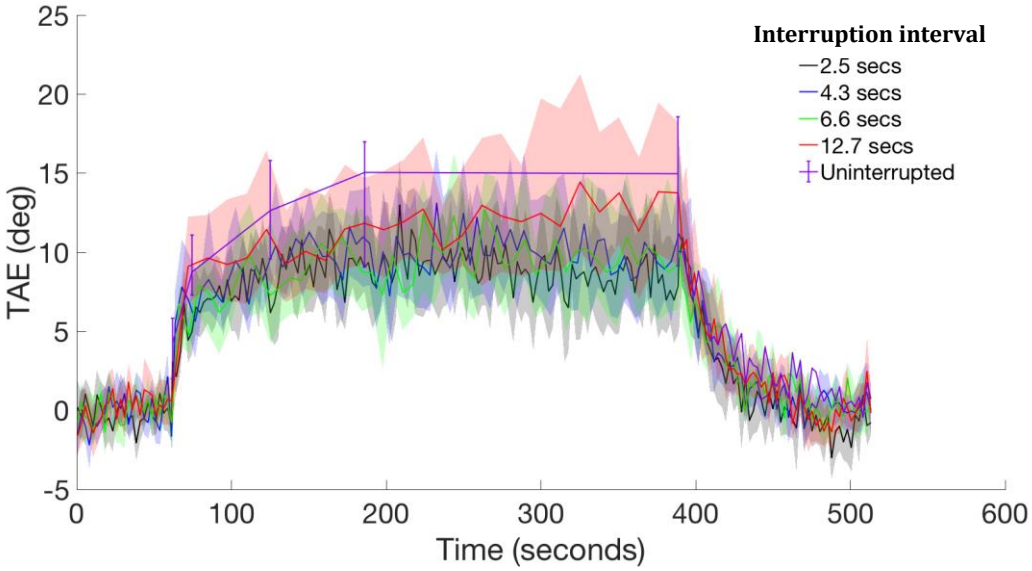


Figure 4.2 Magnitude of the tilt aftereffect (TAE) as a function of time for each adapting condition. Shaded regions represent 95% bootstrapped confidence intervals.

Interrupting adaptation reduces the magnitude of the TAE produced by 327.3 seconds of adaptation. This effect is tuned to the frequency of interruption, with more frequent interruptions resulting in smaller magnitude TAEs. The differences in the magnitude of the TAE between conditions can be seen more clearly in the descriptive power law fits to the time course of each condition (Figure 4.3). The largest final magnitude TAE, at 15.0 degrees, is produced following 327.3 seconds of continuous adaptation. In comparison, the smallest final magnitude TAE, at 7.9 degrees, occurs following the 327.3 seconds of adaptation that was interrupted most frequently (every 2.5 seconds). The magnitude of the TAE in the remaining interrupted conditions lie between these two extremes (4.3 = 11.0deg; 6.6 = 9.2deg; 12.7 = 13.8deg).

The magnitude of the TAE on the last trial of adaptation (327.3 seconds) was compared between conditions using a series of pairwise permutation tests. Significant differences in the final magnitude of the TAE between conditions were revealed through permutation testing. The magnitude of the TAE following continuous adaptation was significantly greater than

the magnitude of the TAE in all interrupted conditions except the least interrupted (*p's* < .05). The magnitude of growth in the least interrupted condition did not differ significantly from the continuous condition (p=.34) or the 4.3 second interrupted condition (p = .080) but did differ significantly from the 2.5 second and 6.7 second interrupted conditions (*p's*<.01). The 6.7 second interrupted condition did not differ significantly from the 2.5 second (p = 0.48) or 4.3 second (p = 0.26) interrupted conditions. The 2.5 second interrupted condition differed significantly from the 4.3 second interrupted condition (p = 0.022). It should be noted however, that there will be inflation of the familywise error rate to due multiple comparisons. In future, this could be corrected for using a Bonferroni correction.

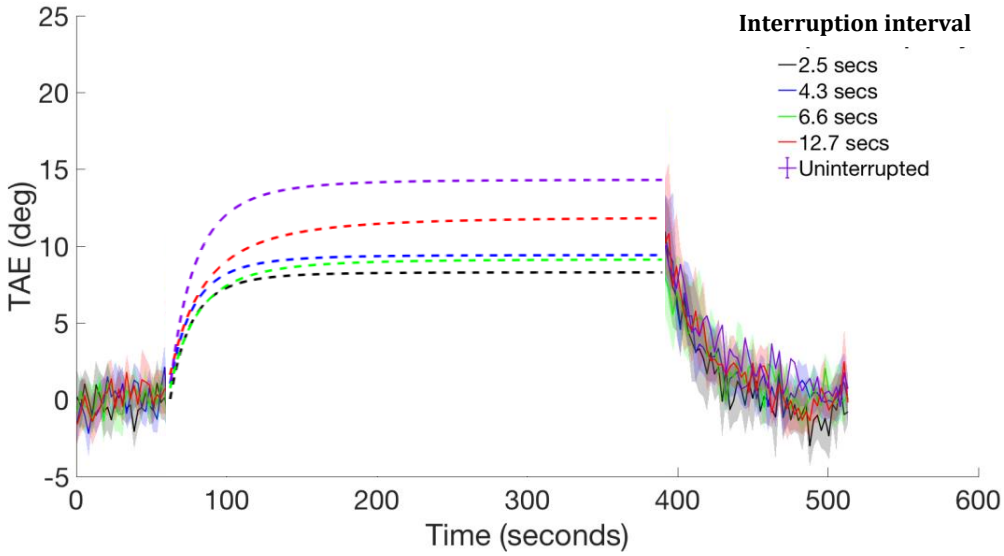


Figure 4.3 Magnitude of the TAE as a function of time for each adapting condition. Mean bias during adaptation trials for all conditions are replaced by descriptive power law fits to those data (see Equation 2.5).

These results suggest that interrupting adaptation reduces the magnitude of TAE growth relative to that seen with traditional continuous adaptation. This suppressive effect was tuned to the frequency of interruption, with the largest amount of suppression seen in the most interrupted condition (every 2.5 seconds), and the smallest amount of suppression seen in the

least-interrupted condition (every 12.7 seconds). Following the offset of adaptation, the magnitude of the TAE decays back to zero. There are no differences in the rate of decay between the conditions.

4.3. EXPERIMENT 2: TYPE OF INTERRUPTION

The results of Experiment 1 suggest that the magnitude of the TAE was suppressed by interrupting adaptation. The interruption period might result in passive recovery, during which previously accumulated adaptation begins to decay. This can be thought of in terms of the leaky bucket analogy covered in Chapter 1, section 1.1.5. During the interruption period, there is no water flowing into the bucket and water continues to leak out of the bucket. Therefore, with no adaptation input, the effects of adaptation begin to decay – a process known as passive recovery. However, there are also events in the interruption period that may be contributing to the suppression of adaptation alongside passive recovery.

Firstly, there was the presentation of a test stimulus with an orientation that differed from the adapting orientation. Adaptation processes occur over a wide range of timescales, from single stimulus presentation (Fischer & Whitney, 2014) to four days of altered stimulus statistics (Haak, Fast, Bao, Lee, & Engel, 2014). Therefore, interrupting a period of adaptation with a change in stimulus statistics, i.e. the test stimulus orientation, could contribute to the decay of the TAE.

Secondly, immediately following the interruption period, the observer is required to accurately reproduce the orientation of the previously presented test stimulus. Therefore, during the interruption period the observer is required to memorise the orientation of the test stimulus. Holding the orientation of the test stimulus in working memory may interfere with adapting mechanisms during the adaptation period that

coincides with the response period (Kang, Hong, Blake, & Woodman, 2011).

Finally, also during the response period, they were told to maintain fixation, ensuring adapting stimuli remained in the correct retinal location. Regardless, their attention will have been diverted to the response stimulus, resulting in less attention towards the adapting stimuli. Spivey & Spirn (2000) have shown that the TAE produced by an attended stimulus is 39% larger than that produced by an unattended stimulus. Therefore, it is possible that reduced attention to the adapters during the response period produces the suppression of the TAE.

The following experiment aims to disentangle the effects of passive recovery, higher level memory processes, and attention during the response period, on the suppression of the TAE during interrupted adaptation.

4.3.1. Methods

Like Experiment one, Experiment two measured the time course of the TAE using a period of continuous or interrupted adaptation. The durations of continuous adaptation were 0.53, 13.2, 63.9 and 127.2 seconds. These are the first four durations used in Experiment one; the longest duration was not used in order to expedite data collection.

Three different interrupted adaptation conditions were compared. In the standard condition, the interruption period consisted of the presentation of an ISI, test stimulus, mask, and ISI, and was immediately followed by the onset of adaptation whilst the observer reproduced the orientation of the test stimulus. In the video playback condition, the interruption period remained the same, but instead of making a response to the test stimulus, a video-recorded response from the observer's previous experimental run

was played. In the mean-luminance interruption condition, the interruption period was replaced by a mean luminance screen and the observer did not make a response during the subsequent onset of adaptation. Details of stimulus presentation timings can be seen in Figure 4.4.

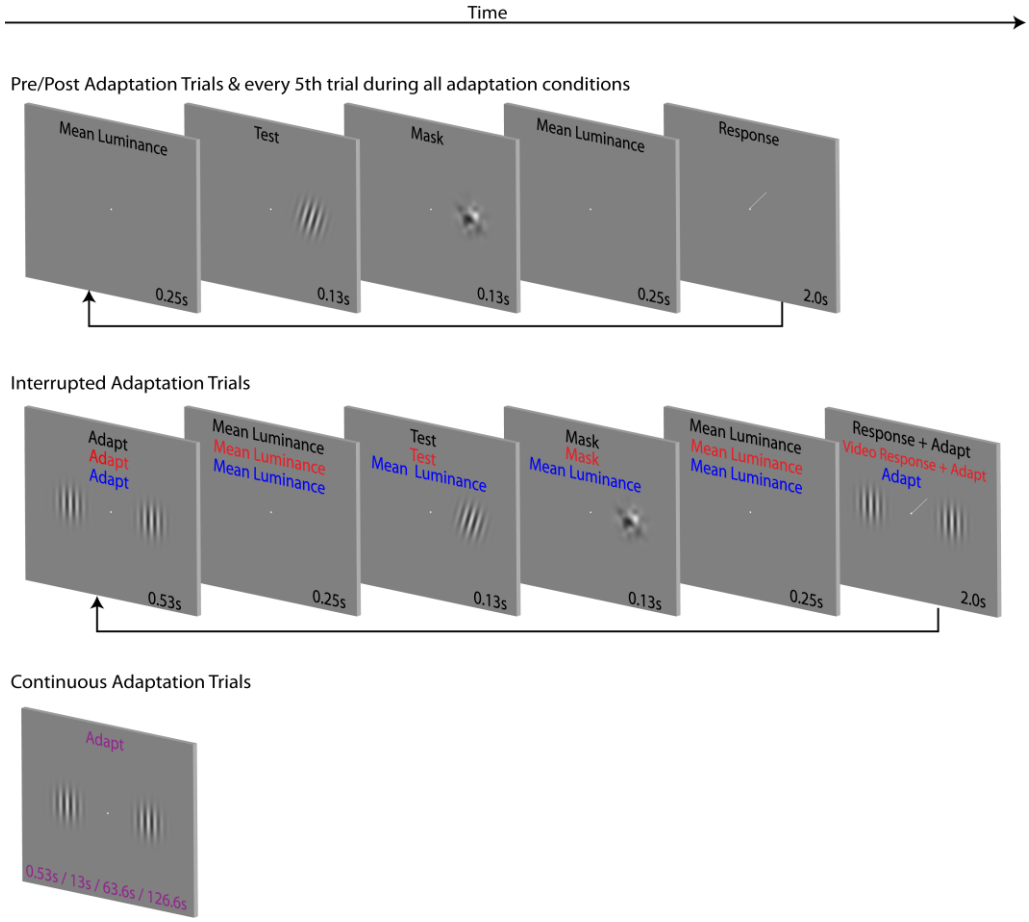


Figure 4.4 Trial timings used in experiment 2. Top row: stimuli and timings used in the 25 pre-adaptation and 50 post-adaptation trials. Middle row: stimuli and timings used in the interrupted adaptation trials (Black = standard interrupted, blue = mean luminance interrupted, red = video playback) . Bottom row: stimuli used in the continuous periods of adaptation.

The orientation and order of stimuli were randomised for each of 5 experimental runs per condition but were identical across conditions. This allowed the presentation of identical stimulus sequences for each of the three interrupted conditions. In each experimental session the observers completed one run of all four conditions. The order of the conditions was counterbalanced for each observer however the standard condition

always preceded the video playback condition to allow the recording of the observer's responses for subsequent playback.

To have no responses on the video playback and mean luminance interrupted conditions would make it impossible to measure the time course of the TAE. Therefore, in these conditions, on every 5th trial the stimulus sequence was identical to that of the standard condition. This allowed the measurement of the time course of the TAE at a reasonable temporal resolution.

The results of Experiment one suggested that the standard interrupted condition would result in a supressed TAE in comparison to the continuous condition. The leaky integrator model (see Chapter 1, section 1.1.5) suggests that all three interrupted conditions will lead to a suppressed TAE due to the time out from adaptation leading to passive recovery. The mean-luminance condition is therefore likely to produce some suppression of the TAE. If the presentation of a test stimulus, the decision-making process involved in making a response, and the reduced attention in the response period also affect the growth of the TAE, then the largest suppression may be seen in the standard interrupted condition.

To disentangle presentation of the test stimulus from having to make a response, the suppression seen in the video—playback condition can be compared to the standard condition. If removing the need to make a response in the video-playback condition reduced the suppression of the TAE, this would suggest that attention during the response period affects the growth of the TAE beyond the presentation of a test stimulus.

If, however, the suppression in the standard and video-playback conditions is equal, then the presentation of the test stimulus alone could be suggested to produce the suppression seen in Experiment one.

Finally, if all three interrupted conditions show the same suppressed time course, then the suppressive effects seen in Experiment one could be attributed simply to time out from adaptation resulting in passive recovery.

4.3.2. Results

As can be seen in Figure 4.5, the results of Experiment one were replicated in Experiment two. The magnitude of the TAE following 127.2 seconds of continuous adaptation was 14 degrees compared to a smaller 8-degree magnitude following the same duration of adaptation interrupted by the measurement of the TAE. The magnitude of the TAE following continuous adaptation was significantly greater than the magnitude of the TAE following all three types of interrupted condition (Permutation test: all comparisons $p < .01$)

In contrast, there were no significant differences in the magnitude of the TAE between the three interrupted conditions (permutation test: all comparisons $p > .05$). The magnitude of the TAE in each condition grows to ~8 degrees.

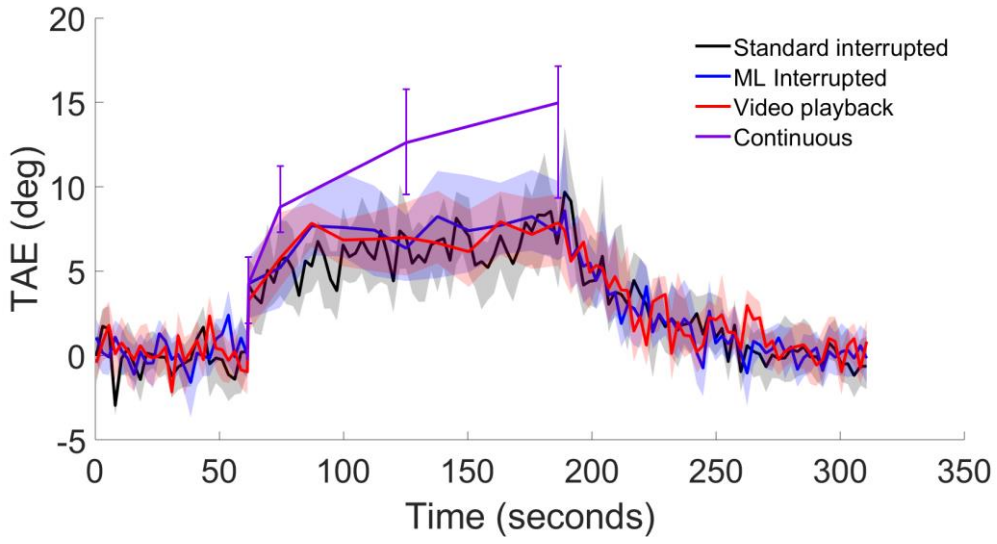


Figure 4.5 Magnitude of the TAE as a function of time for each adapting condition. Shaded regions and error bars represent 95% bootstrapped confidence intervals.

4.3.3. Modelling the growth of the TAE

Adaptation growth and decay is commonly modelled through leaky integration, whereby the magnitude of adaptation at each time point is approximated by the convolution of a leaky integrator with the stimulus input. The leaky integrator has two free parameters, the time constant and the gain, which determine the rate and magnitude of adaptation growth and decay.

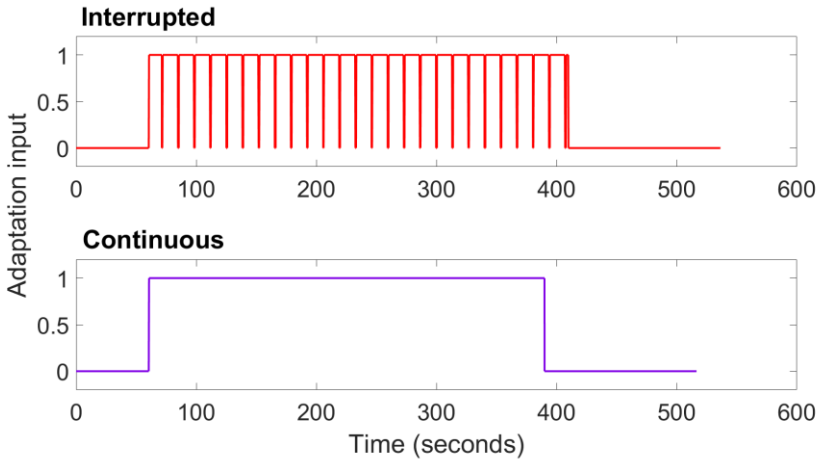


Figure 4.6. Interrupted and continuous adaptation inputs to the leaky integrator model. The y-axis represents the presence of adaptation at each time point, with 0 being absent and 1 being present. The continuous input (purple) shows adaptation is present for the full adapting duration. In contrast, the interrupted input shows the brief periods where adaptation is absent during each of 26 interruptions. These interruptions cause the time course of adaptation input to be longer than that of the continuous input, however the total duration of adaptation is identical.

The time courses from each condition in Experiment 1 were modelled by convolving the time course (box car) from each condition with an exponential leaky integrator. The time course consisted of ones and zeros to represent the presence (1) or absence (0) of adaptation at that time point. Figure 4.6 shows example stimulus inputs for interrupted and continuous conditions. The equation for the leaky integrator is shown below (Equation 4.1), where t is time, and g and c are the gain and time constant of the leaky integrator respectively.

$$y = g\frac{e^{\frac{-t}{c}}}{\int e^{\frac{-t}{c}}dt}$$ Equation 4.1

The model was fit to all five conditions simultaneously to find the values for the time constant and gain that minimised the squared error between the model and the data.

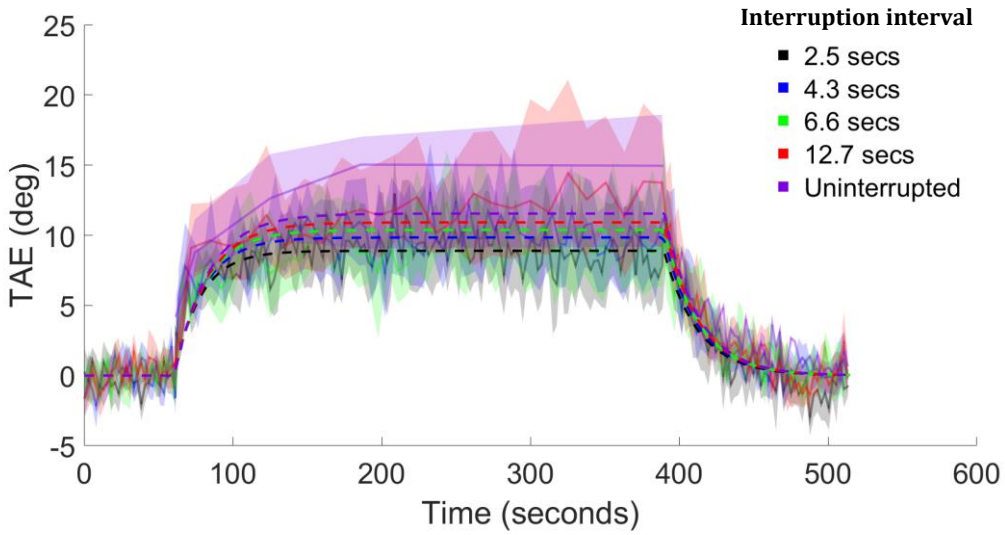


Figure 4.7. The time course of the tilt after-effect for each condition in Experiment one (faint solid lines). Overlaid is the best-fitting model's predicted time course for each condition (dashed lines). The best fitting model had a gain of 11.55 degrees and a time constant of 22.73 seconds. R^2 = 0.635.

Figure 4.7 shows the time course of adaptation for each condition in Experiment 1, with power law descriptive fits for the growth of adaptation (solid lines). Overlaid is the best-fitting model (dashed lines), which had a time constant of 22.73 seconds and a gain of 11.55 degrees. This model explained 63.5% of the variance in the data.

The model successfully captures the pattern of suppression, with larger suppression being seen in conditions with more frequent interruptions. However, the model does not accurately predict the magnitude of suppression caused by increasing the frequency of interruption. For example, the model predicts a 2.66deg difference in final magnitude

between the continuous and 2.53sec interrupted condition, compared to the actual difference of 7.1deg.

This simple model therefore cannot account for the magnitude of suppression caused by interrupting adaptation. This suggests that the process of interrupting adaptation may alter the nature of adaptation growth beyond simple decay during periods of interruption. In the above model, the rate of growth is identical to the rate of decay. However, it is possible that different mechanisms control the rate of growth and decay. If the rate of decay were to be higher than the rate of growth, the magnitude of suppression caused by interruption might be better captured.

4.4. DISCUSSION

Experiment 1 showed that interrupting adaptation to gain an estimate of the TAE reduced the growth of the TAE. This effect was tuned to the frequency of interruption, with more frequent interruption leading to greater reduction in the final magnitude of the TAE. There were no differences seen in the rate of decay of the TAE between conditions. This is expected, as following adaptation offset the rate of interruption was identical for all conditions.

In each condition, the cumulative time spent adapting was identical at 327.3 seconds. The difference between conditions lies in the frequency of interruption, which can be thought of as time spent not adapting. The results therefore suggest that during the interruption periods, where there is no adapting stimulus, previously accumulated adaptation began to decay. With a higher frequency of interruption, there was more opportunity for decay, and thus the TAE reached a smaller magnitude. This is in line with the findings of Magnussen & Johnsen (1986) who show there is no storage of the TAE, suggesting that without adaptation input the effect will begin to decay.

The results of Experiment 1 can be interpreted in terms of a leaky integrator adapting mechanism. Using the 'leaky bucket' analogy, in the continuous condition, the flow of water to the bucket is continuous, meaning adaptation grows until it reaches asymptote. In the interrupted condition however, the flow of water into the bucket switches on and off at regular intervals. This means adaptation grows during the time in which the water is flowing in, but when it stops, water continues to leak out of the bucket, and adaptation partially decays (see Figure 4.8). The more the flow is interrupted, the more water leaks out of the bucket, and the smaller the eventual magnitude of adaptation. This therefore explains why the same cumulative duration of adaptation produces a smaller magnitude TAE during interrupted conditions.

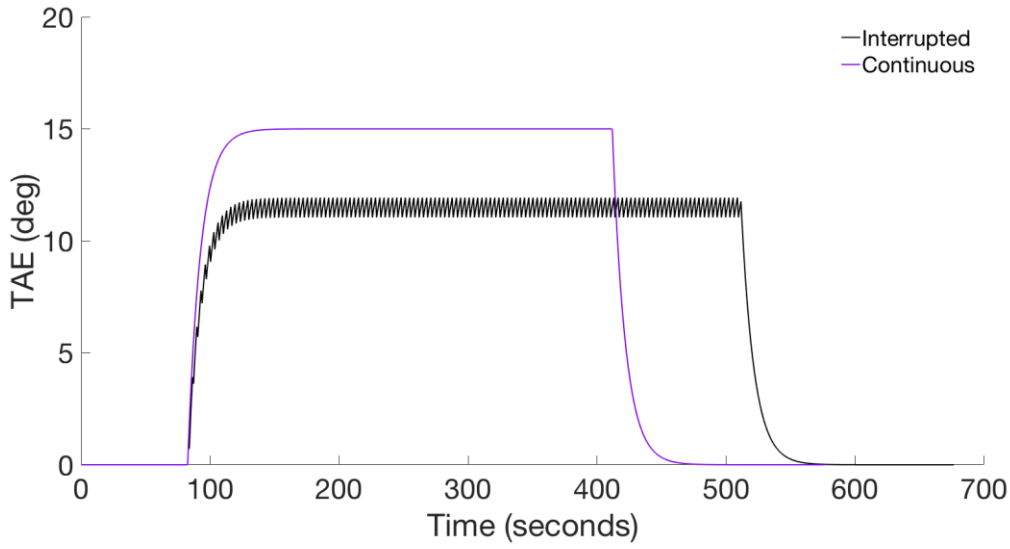


Figure 4.8 visual representation of the leaky bucket analogy. Continuous adaptation leads to smooth growth of the TAE. Interrupted adaptation leads to supressed growth. The 'zig-zag' nature of the growth shows the decay of the TAE during interruptions of adaptation.

The effect of interrupting adaptation to measure its growth has wider implications on traditional adapt-test-top-up experimental designs. Experiments adopting these designs are often not measuring adaptation per se, but are measuring another property whilst the observer is maintained in an adapted state. To do so, a period of initial adaptation is

presented to induce an adapted state in an observer. This is followed by the presentation of a test period during which there is no adaptation, and a subsequent top-up period of adaptation to maintain a fixed level of adaptation.

These adapt-test-top-up experiments assume that the period of top-up adaptation will counteract any adaptation decay during the test phase. However, as the results of Experiment 1 demonstrate, this will depend heavily on the ratio of initial adaptation to top-up adaptation, as well as the length of the test phase where there is no adaptation. Should the period of top-up adaptation be too short, then the adapted state throughout the experimental trials will be smaller in magnitude than that produced by the initial adaptation. This is due to the decay of adaptation during the test phase not being offset by a sufficient period of top-up adaptation. The opposite will be found if the top-up period is too long.

There does not appear to be consensus on the ratio of initial adaptation to top-up adaptation that maintains adaptation at a fixed level, with most researchers using simple heuristics or arbitrary values. Picking just three of the many examples inducing tilt adaptation we see 30sec of initial adaptation followed by 10sec of top-up (Pavan, Hocketstaller, Contillo, & Greenlee, 2016); 20sec of initial adaptation followed by 5sec top-up (Hussain Ismail, Solomon, Hansard, & Mareschal, 2016), and 10sec initial adaptation followed by 5 sec top-up (Zou, He, & Zhang, 2016). This approach of arbitrary initial adapt and top-up durations is unlikely to induce a consistent adapted state in the observers used in these studies, and makes comparisons between studies difficult as the level of adapted state is unknown. However, the data suggest that within a given experiment, an observer's adapted state will be consistent if conditions have the same initial and top-up durations. Therefore, despite the level of an observer's adaptation potentially being larger or smaller than expected, it will be consistent and allow for valid comparisons across conditions.

From the data in Experiment 1, it is possible to identify the amount of adaptation suppression produced by a fixed duration and frequency of interruption. This knowledge would permit the creation of a simple model whereby periods of interruption result in the decay of adaptation and the eventual suppression of the magnitude of the TAE seen in Experiment 1.

However, due to the nature of the interruption in Experiment 1, Experiment 2 aimed to identify whether solely time out from interruption caused adaptation decay, or whether factors such as working memory or attention were contributing. If this was the case, such a simple model of adaptation would be insufficient to explain the suppression of the TAE produced by interrupting adaptation.

The data in Experiment two replicated the findings of Experiment one; interrupting a period of adaptation to measure the TAE supresses the growth of the TAE by ~5 degrees. The type of interruption, whether this be with a test stimulus and a response, a test stimulus and video response, or just with mean luminance, did not affect the amount of suppression caused by the interruption.

The magnitude of the TAE following mean luminance interruption did not differ from the other interrupted conditions, suggesting that solely time out from adaptation produces the suppressive effect seen in Experiments One and Two. This suggests that interrupting adaptation with the presentation of a test stimulus, holding the orientation of this stimulus in working memory, and the attentional cost of making a response do not supress the growth of the TAE to a greater extent than interruption with mean luminance alone.

These findings drove the creation of a simple model in which adaptation decayed during periods of interruption at the same rate as adaptation growth. However, this simple model was unable to fully capture the magnitude of suppression present in the empirical data. This model is

therefore unable to aid prediction of the duration of top-up adaptation needed to counteract adaptation decay during a test-interruption. A more complex model, perhaps where the rates of adaptation growth and decay differ, might be able to better make predictions about appropriate initial and top-up adaptation durations.

5

The time course of the tilt aftereffect following exposure to altered temporal frequency content of naturalistic videos

5.1. INTRODUCTION

Perception adjusts to statistical regularities in natural environments by modifying the balance of excitatory and suppressive brain signals– a process known as visual adaptation (Gibson & Radner, 1937; Levinson & Sekuler, 1976; Solomon & Kohn, 2014). The perceptual consequences of adaptation typically manifest as changes to the appearance of visual input and perceptual sensitivity of the organism. Changes to the appearance of incoming visual input often present as a repulsive perceptual bias away from repetitive visual experiences (Thompson & Burr, 2009). For example, repeated exposure to a tilted line causes a subsequently presented vertical line to be perceived as tilted in the opposite direction (Gibson & Radner, 1937). This ability to adapt to changes in incoming stimulus statistics is thought to enable efficient encoding of a rich and dynamic visual environment (Clifford et al., 2007; Rieke & Rudd, 2009; Stocker & Simoncelli, 2006; Wainwright, 1999).

Visual aftereffects have a characteristic time course, growing and decaying in strength in proportion to the logarithm of adapting stimulus duration (Magnussen & Johnsen, 1986). Yet we do not understand what determines this time course. Changes in our environment can occur at timescales ranging several orders of magnitude and optimal adaptation would require matching the time course of adaptation to these timescales (Wark, Fairhall, & Rieke, 2009). It remains unclear whether, and if so how, the

brain matches the dynamics of adaptation to statistics that can change over a large range of timescales - from years to months to milliseconds.

The dynamics of adaptation may have evolved to match persistent temporal regularities in natural environments, resulting in visual aftereffects with a fixed time course. The temporal power spectrum of light intensity in the natural visual environment approximates 1/f (f: temporal frequency), with proportionally more energy at low frequencies (Dong & Atick, 1995). One possibility therefore, is that the visual system has evolved fixed adapting mechanisms to adapt at rates that are optimally tuned to this 1/f spectrum. In this scheme, the time course of adaptation to a fixed adapting duration would be fixed, unaffected by current environmental temporal statistics.

Alternatively, adaptation mechanisms might have the capacity to alter adaptation's time course, should the temporal structure of the environment change. Adjustable adaptation mechanisms would possess the ability to gradually compensate for slow changes to visual input (e.g. maintaining perceptual stability during ocular development (Banks, 1988), while retaining sufficient sensitivity to rapid, transient changes in the environment (e.g. compensating for discrepant optic flow signals (Anstis, Verstraten, & Mather, 1998). In this scheme, the same duration of adaptation may produce distinct time courses depending on the prevailing temporal statistics, with faster rates of change producing more rapid adaptation. Recent exposure to statistical temporal regularities would therefore set the current dynamics of the time course of adaptation.

Here we ask whether perturbing the temporal statistics of natural environments changes the build-up and decay of adaptation. We tested whether prolonged exposure to video sequences that were filtered to alter the temporal frequency spectra from 1/f, would modulate the time course

of subsequently measured adaption. High-pass filtered videos were dominated by fast changes in contrast, whilst low-pass filtered videos possessed slow changes in contrast, an unfiltered video served as baseline 1/f exposure. Adaptation was quantified using the tilt aftereffect (TAE), whereby prolonged exposure to a tilted stimulus results in a measurable bias in the perceived orientation of subsequently presented stimuli away from the adapting stimulus orientation (Gibson & Radner, 1937).

5.2. METHODS

To expose observers to perturbed temporal frequency statistics they viewed 10-minute video stimuli that had been filtered in the temporal domain. Following this, the time course of the tilt aftereffect was measured as described in General Methods 2.4.

5.2.1. Temporal frequency filtered videos

Video stimuli were 10-minute, black and white segments from a popular television show. Temporal high-pass and low-pass filtering was conducted in Matlab using a 9^{th} order Butterworth filter with a cut off frequency of 0.8Hz. High-pass filtering attenuated temporal frequencies below 0.8Hz, whilst low-pass filtering attenuated frequencies above 0.8Hz. A 0.8Hz cut-off was chosen to ensure a significant deviation from natural scene statistics, whilst retaining a pleasant viewing experience. Increasing the cut-off would produce very fast flashing images in the high-pass filtered video that would be uncomfortable to view, whilst decreasing the cut-off would reduce the deviation from 1/f statistics. Filtered videos were spliced with their unfiltered audio counterparts and corrected for audio-visual asynchrony.

Figure 5.1 shows the temporal (left) and spatial (right) frequency spectra of the video stimuli used in this experiment. The temporal frequency spectra show that the unfiltered video exhibits classic 1/f statistics (black line), with most of the power at low temporal frequencies. In contrast, high-pass filtering (red line) significantly reduced the power at low temporal frequencies that are predominant in natural scenes. All three videos retained the 1/f spatial frequency spectra which is typical of natural scenes, with most power at low spatial frequencies.

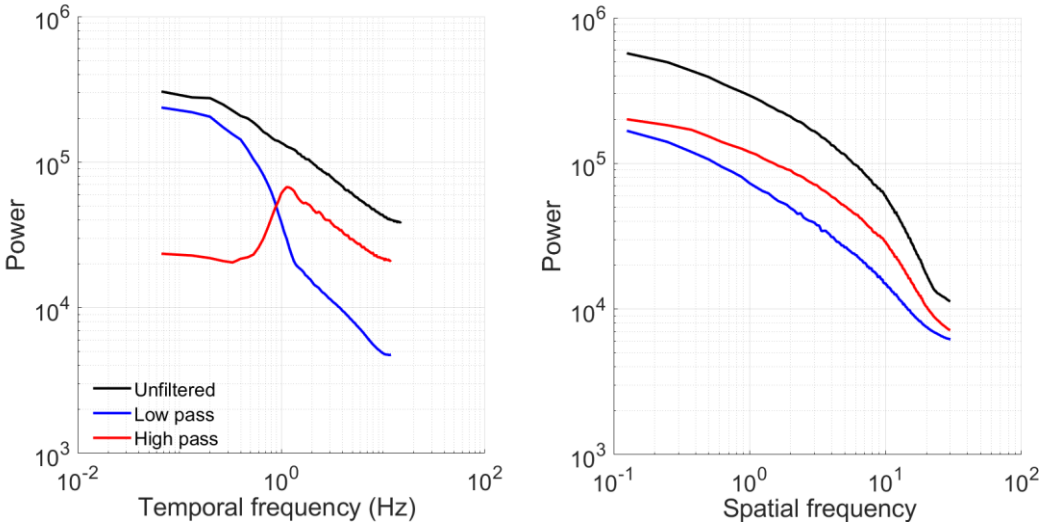


Figure 5.1 Pre-exposure video filtering: Participants watched 10-minute clips of a popular TV show that were either unfiltered (black), high-pass (red) filtered, or low-pass filtered (blue) in the temporal domain (9th order Butterworth filter, 0.8Hz cutoff). Unfiltered videos show a characteristic 1/f temporal frequency spectrum. High-pass filtering attenuated low temporal frequency components predominant in natural scenes, whilst low-pass filtering attenuated high temporal frequency components. The spatial frequency spectra of all videos showed classic 1/f characteristics with more energy at low spatial frequencies.

5.2.2. Data analysis

The magnitude of the TAE on each trial was quantified as the angular difference between the presented and reproduced orientation, averaged

across participants. Outlier removal using the median absolute deviation (MAD), and correction for biases in baseline orientation perception was conducted for each participant's data (see General Methods 2.5.1). Data was pooled across 6 observers for subsequent analysis.

Adapting stimuli were presented during the period in which observers made their response. Due to the variability in response times, the total duration of adaptation on each trial was inconsistent. The mean trial time, across observers during adapting trials was 3.4 seconds, but this ranged from 1.9 seconds to 7.8 seconds. On each trial, there was a fixed period of 0.7 seconds no adaptation meaning the duration of adaptation ranged from 1.2 seconds to 7.1 seconds. Looking at bias as a function of trial number would pollute estimates of the TAE across different adapting durations. Those taking longer to respond would have adapted more on each trial than those with faster response times. To circumvent this potential confound, cumulative adapting time was calculated by subtracting the constant timings of stimulus presentation (0.7sec) away from the total duration of the trial, leaving the duration of adaptation on each trial for each observer. Estimates of the TAE were then binned in 3.5 second time bins to ensure that responses for the same duration of adaptation were aligned. Trials during the growth of adaptation are therefore plotted as cumulative adapting time rather than time from adapter onset.

5.3. RESULTS

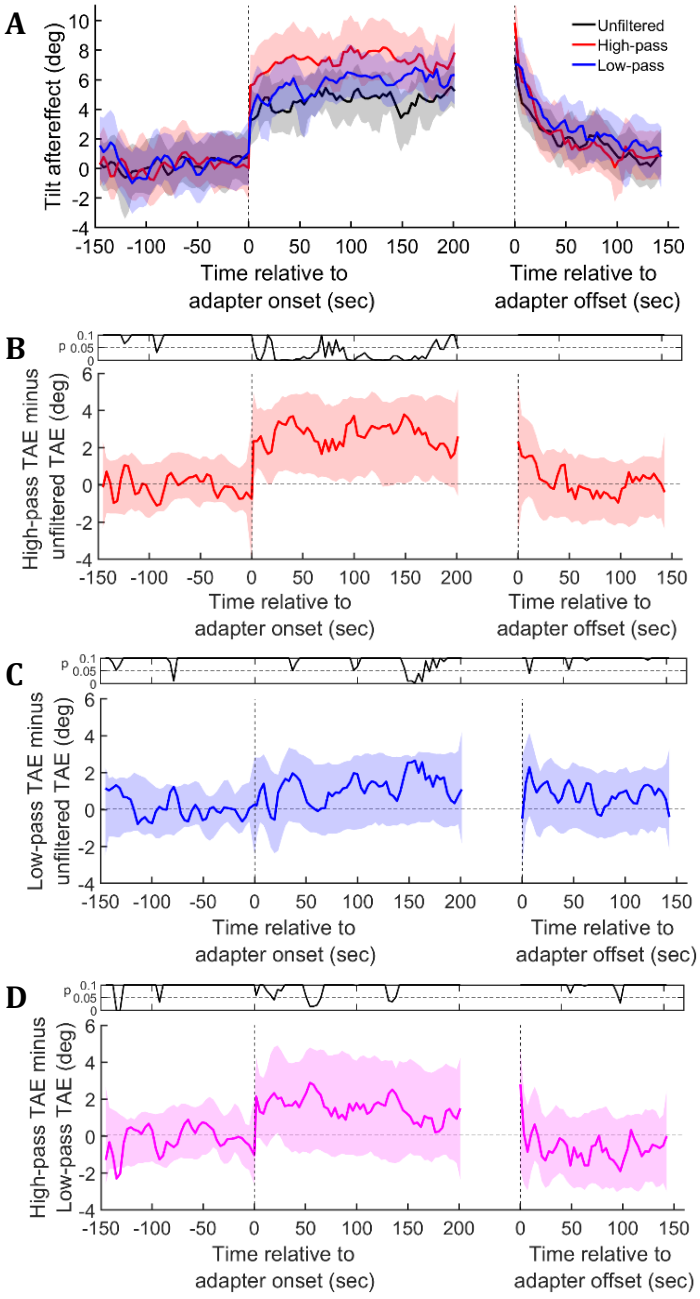


Figure 5.2 (A) The time course of the tilt aftereffect (TAE): the magnitude of the TAE was quantified as the angular difference between the presented and reproduced orientation on each trial, averaged across participants (n=6). The TAE reaches a greater magnitude following prior exposure to a high-pass filtered video (red) than an unfiltered video (black). (B-D) Difference in the magnitude of the TAE between high-pass and unfiltered conditions (red), low-pass and unfiltered conditions (blue) and high-pass and low-pass conditions (magenta). Permutation testing was conducted at each time point (top of each figure). Shaded regions represent 95% bootstrapped confidence intervals.

Figure 5.2A shows that before the onset of adaptation, bias remains stable around zero for all conditions. This is important, because it confirms that prolonged exposure to the video sequences did not directly alter orientation perception.

Following the onset of adaptation, the TAE grows in magnitude with adaptation duration. The TAE grows more rapidly over the first few seconds after adapter onset, and has a larger overall magnitude following exposure to temporally high-pass filtered (red line) than unfiltered videos (black line). Permutation testing reveals significant differences in the average TAE during adaptation growth (0-200 secs after adaptation onset, mean difference =2.69deg, p = .04).

More temporally fine differences between the time courses following high-pass pre-exposure and unfiltered pre-exposure can be seen in Figure 5.2B, where the difference in TAE magnitude between the two conditions at each time point is plotted. Above, permutation derived p values are plotted for the difference between conditions at each time point. Both the initial magnitude of the TAE, and its eventual magnitude at asymptote are significantly greater following exposure to high-pass filtered videos than exposure to an unfiltered video.

In contrast, the growth of the TAE following exposure to the low-pass filtered video does not differ significantly from that following unfiltered video exposure (mean difference =1.15deg, p = .55). The difference in magnitude between these conditions, and the permutation derived p values at each time point can be seen in Figure 5.2C. Both the initial rate of growth, and eventual magnitude of the TAE at asymptote are not significantly different.

As can be seen in Figure 5.2D, there were also no significant differences in either the rate or eventual magnitude of TAE growth following high-pass filtered and low-pass filtered exposure videos (mean difference =1.54deg, p = .40). And no significant differences in TAE decay between any conditions (all permutation tests on average differences p >.05)

5.4. MODELLING THE TIME COURSE OF THE TAE

Recent exposure to an environment with low temporal frequencies attenuated was sufficient to significantly increase both the initial rate and final magnitude of TAE growth. The time course of the TAE can therefore be altered in response to changes in the temporal frequency content of recent visual experience. Here, we aim to model the mechanisms underlying this meta-adaptation effect.

As discussed in Chapter 3, the time course of adaptation can be successfully modelled using a leaky integrator model (Drew & Abbott, 2006). Each adapting mechanism has two properties that determine the time course of adaptation. The rate of adaptation growth is determined by the time constant of the mechanism, whilst its magnitude is determined by the gain of the mechanism. Recent work suggests that there are at least two distinct mechanisms underlying adaption, one operating at short timescales and one at long timescales (Bao, Engel, Fast, & Mesik, 2013). Mechanisms operating at different timescales makes adaptation more efficient, being able to gradually adapt to slow changes in the environment whilst retaining sensitivity to fast alterations.

Within this explanatory framework, there are at least two possible mechanistic explanations for the meta-adaptation effect described above. One possibility is that there is a single adapting mechanism with a time

constant and gain that can be altered to match the prevailing temporal statistics of the environment. To model this mechanism seems unnecessary as any time course could be captured by a model with this flexibility. In addition, this mechanistic explanation is inconsistent with the finding that at least two mechanisms underlie adaptation at different timescales.

The alternative mechanistic explanation is that multiple adapting mechanisms exist with fixed time constants but variable gains, and the time course of adaptation is determined by the relative contribution of each mechanism. These mechanisms could therefore operate in various temporal environments by adjusting their gain to produce the appropriate rate of adaptation. For example, in fast-changing environments the adaptation output would be driven mostly by a short-timescale mechanism, achieved by increasing the gain of this mechanism relative to mechanisms tuned to longer-timescales that would produce slower adaptation.

The empirical data collected cannot determine the number of adapting mechanisms. While previous research has discovered at least two mechanisms, potentially more exist. To produce the most simplistic model, and build on previous findings of two adapting mechanisms, the time course of adaptation was modelled with only two adapting mechanisms, and one variable parameter, as a proof of concept for the mechanistic underpinnings of the effect described above.

The magnitude of adaptation was modelled with a short-timescale mechanism and a long-timescale mechanism. The weighting of each mechanisms' contribution to the adaptation output was altered by their relative gains. Intuitively, the predominance of fast timescale changes in recent history due to high-pass video pre-exposure may result in the increased contribution of the short-timescale mechanism to the overall adaptation output. In this way, the dynamics of adaptation would shift from being matched to 1/f natural statistics to the statistics of the pre-exposure stimulus.

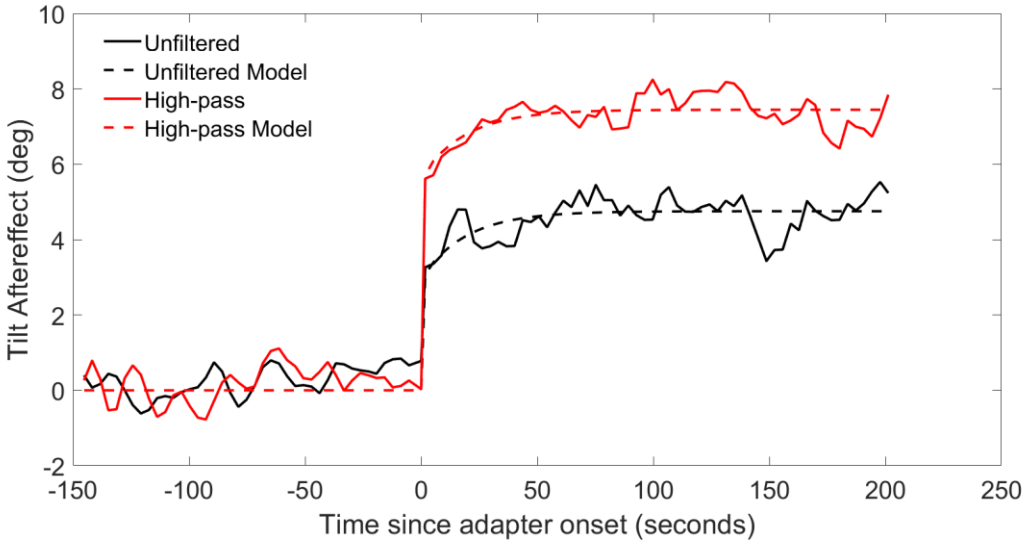


Figure 5.3 Modelling the time course of the TAE. The time course of the TAE following exposure to unfiltered (black) and high-pass filtered (red) videos is depicted by solid lines. Dashed lines show the best-fitting two-mechanism leaky integrator model. The dynamics of adaptation following exposure to a high-pass filtered video are captured by increasing the gain of the short-timescale mechanism whilst all other mechanistic parameters remain constant.

The model therefore consisted of two adapting mechanisms, one short timescale and one long timescale, for each time course. Each of the four mechanisms had a time constant and gain parameter. There was one free parameter to vary during fitting: the gain of the short timescale mechanism. All other parameters remained constant for all mechanisms.

Doubling the gain of the short timescale mechanism from 0.14deg to 0.28deg produced a model that explained 91% of the variance in the data (see Figure 5.3). This is consistent with the idea that exposure to an environment with predominantly fast changes would increase the input from the short-timescale mechanism. Each model had a long timescale mechanism with a gain of 0.1deg and a time constant of 5.5 seconds. The short timescale mechanism had a smaller time constant of 0.26 seconds which had a gain of 0.14deg for the unfiltered exposure condition and 0.28deg for the high-pass exposure condition. Altering the relative contributions of a short and long timescale mechanism could therefore produce time courses of adaptation matched to the rate of change in the environment.

As stated previously, it is also possible that rather than their being only two mechanisms that alter their gain to match the time course of adaptation to recently encountered dynamics, there are multiple mechanisms with fixed time constants that become activated in different temporal environments. The two-mechanism model above explains a large proportion of the data and is preferred to a more complex model. To fully understand the mechanistic underpinnings of flexible adaptation there would need to be subsequent research into the number of adapting mechanisms operating at these timescales.

In conclusion, the empirical data suggest that the dynamics of adaptation are altered by the temporal statistics of previous visual experience. However, neither the empirical data nor modelling can determine whether this flexible adaptation is driven by a single mechanism with flexible temporal tuning and gain or multiple fixed-timescale mechanisms that become active in different temporal environments. Further experimental and modelling studies are needed to elucidate the mechanistic underpinnings of meta-adaptation.

5.5. DISCUSSION

The time course of the TAE is not fixed but can be altered through exposure to temporal frequency content that deviates from naturalistic 1/f statistics. Exposure to a video filtered to selectively attenuate frequencies with proportionally more energy in natural scenes was sufficient to cause the TAE to grow more rapidly and reach a greater magnitude than exposure to a video with naturalistic 1/f temporal structure. The increase in rate and magnitude of the TAE following exposure to a video with predominantly fast changes in contrast suggests that adapting mechanisms can alter the dynamics of adaptation to match the temporal structure of recent visual experience, a process referred to as meta-adaptation or meta-plasticity (Abraham & Bear, 1996; Robinson, Harper, & McAlpine, 2016)

Meta-adaptation would allow for gradual compensation for slow changes in visual input whilst retaining sufficient sensitivity to rapid, transient changes in the environment. Moreover, this could be achieved with a small number of mechanisms that could alter their properties to match the dynamics of adaptation to current statistics. Meta-adaptation is therefore an efficient solution allowing the visual system to adapt in a temporally complex environment.

As discussed earlier, the dynamics of the TAE are captured by a two mechanism model (Bao & Engel, 2012). Each mechanism has a time constant, controlling the rate of adaptation, and a gain controlling the magnitude of adaptation. Adaptation output is suggested to be a simple summation of a short-timescale mechanism with a short time constant, and a long-timescale mechanism with a long time constant. There are two mechanistic possibilities underlying the observed meta-adaptation effect.

Firstly, a single adapting mechanism could exist that has flexible time constant and gain properties. This mechanism would be capable of altering its temporal tuning and overall strength (gain) to produce almost any time course of adaptation. However, to alter adaptation across a wide range of timescales would demand considerable flexibility, and previous research has shown at least two mechanisms of adaptation operating at distinct timescales.

Alternatively, the time course of the TAE could result from the activation of multiple adapting mechanisms tuned to fixed distinct timescales. Alterations in the relative contribution of each mechanism to the adaptation output could adjust the dynamics of adaptation. Exposure to a video with predominantly high temporal frequencies may therefore have led to the reduced contribution of a long-timescale mechanism, and the increased contribution of a distinct mechanism tuned to shorter timescales. This was successfully modelled using a two-mechanism leaky integrator model. Increasing the gain of the short timescale mechanism did result in more rapid and larger adaptation growth, as seen following exposure to the high-pass filtered video. Subtle changes in the gains of both mechanisms could therefore allow distinct time courses of adaptation to a wide variety of recently experienced temporal statistics.

Elucidating the precise neural mechanisms of meta-adaptation cannot be achieved with the current empirical data and is an interesting challenge for future empirical and computational studies. Initial work suggests that cellular and circuit mechanisms can simultaneously store multiple timescales of adaptation, which may themselves be modifiable (Drew & Abbott, 2006; Lundstrom, Higgs, Spain, & Fairhall, 2008; Pozzorini, Naud, Mensi, & Gerstner, 2013; Ulanovsky, Las, Farkas, & Nelken, 2004). For example, recordings from mouse retinal ganglion cells revealed that

timescales of contrast adaptation are coupled to temporal constancies in the environment (Wark et al., 2009). Exploring how visual adaptation is itself adaptable, and how it operates in response to naturalistic stimulation, may lead to a better understanding of its function in visual processing.

Worth noting, is that exposure to videos with high temporal frequencies attenuated did not alter the dynamics of adaptation relative to exposure to the unfiltered video. This finding is likely due to low-pass filtering attenuating power at high temporal frequencies, which are already less predominant in natural scenes. Therefore, it is considered that high-pass filtered videos represented a larger deviation from natural scene statistics than low-pass filtered videos, resulting in the asymmetry of the observed effect. If a small deviation in naturalistic statistics is insufficient to alter the time course of the TAE, this suggests a change in temporal statistics may have to be greater than some threshold to produce alterations in the time course of adaptation. Future work identifying the degree to which the dynamics of adaptation can be altered could be achieved using temporal notch filtering. Exposure to more narrow bands of temporal frequency content, that represent large deviations from natural statistics, may elucidate how finely the time course of adaptation can be altered. It could be that subtle variations in the time course of adaptation are seen to each temporal frequency band, or the time course of adaptation may be more broadly fixed to short and long timescales.

In this experiment, naturalistic video stimuli were used to expose observers to different temporal statistics for two reasons. Firstly, there is strong argument that visual processing should be studied in response to naturalistic stimulus input (Felsen & Dan, 2005). Adaptation is a process that enables the visual system to function efficiently in a complex and dynamic environment. Intuitively, the first step in exploring if adaptation

is sensitive to the temporal structure of natural scenes is to probe its effects in response to naturalistic scenes themselves. Secondly, the timescales over which adaptation operates are broad, and the largest effects will likely be seen when the visual system is exposed to a change in stimulus statistics that is long-lasting. Moreover, adaptation effects are largest when observers maintain attention (Festman & Ahissar, 2004; Spivey & Spirn, 2000). The use of naturalistic stimuli therefore has the added benefit of maintaining observer attention during lengthy pre-exposure procedures.

However, there are also disadvantages to the use of naturalistic stimuli (Rust & Movshon, 2005). Mainly, that they are complex by nature and are therefore not readily parameterized. This results in the inability to test specific hypotheses, as there are many properties of the naturalistic stimuli that may be driving the observed effect. In the current experiment, our naturalistic pre-exposure stimulus is no exception. Although there are clear and measurable differences in the temporal structure of the filtered videos, there are many complexities within this that could drive our observed effect. For example, altering the rate at which contrast changes over time will also have changed the rate at which other spatial properties change over time, and the filtering itself affected the average RMS contrast of each video stimulus. As Rust and Movshon (2005) suggest, to explore the precise mechanisms driving the observed effect, carefully designed simplistic artificial stimuli must be used to test specific hypotheses. Further experiments in this thesis therefore make us of artificial stimuli to unpick the interesting observed effect.

6

The time course of the tilt aftereffect following exposure to altered temporal frequency content of artificial stimuli

6.1. EXPERIMENT 1

6.1.1. Introduction

The results discussed in Chapter 5 revealed that the time course of adaptation can be altered by exposure to naturalistic stimuli with perturbed temporal frequency spectra. Exposure to a high-pass filtered video with predominantly fast changes in luminance, led to a tilt aftereffect (TAE) that grew more rapidly and to a greater magnitude than that following exposure to a video with naturalistic 1/f temporal statistics. These results were compatible with the theory that at least two mechanisms, tuned to short and long timescales respectively, set the time course of adaptation (Bao & Engel, 2012), and were successfully modelled by an increase in the gain of a short timescale mechanism.

The use of artificial stimuli allows the testing of specific hypotheses that are not confounded by the complexities of natural images that could not be controlled for in Chapter 5. For example, complex interactions between spatial and temporal statistics that are not yet fully explored in natural scenes may differ between the videos. In addition, the temporal filtering process altered the root mean square (RMS) contrast of the videos which could affect the rate and magnitude of subsequent adaptation. Therefore, the observed effects in Chapter 5 cannot be attributed solely to our manipulations of temporal frequency content. To do so, the observed effect

must be replicated with stimuli that differ only in temporal frequency content.

This experiment aimed to dissociate the potential causes of the effects observed in chapter 5, using readily controllable artificial pre-exposure stimuli. The pre-exposure stimuli consisted of narrowband isotropic visual noise, at a fixed maximum contrast, that was sinusoidally modulated from zero to maximum at either a slow, or fast, temporal frequency. Pre-exposure stimuli therefore only differed in the rate of contrast change over time. Replication of the results observed in Chapter 5 would suggest that the original observed effect was driven by exposure to videos with temporal frequency spectra deviating from naturalistic statistics.

6.1.2. Methods

6.1.2.1. Pre-exposure noise stimulus creation

Gabor elements (n = 10,000) were combined to create a 1024x1024 pixel square of broadband visual noise. These elements were 90 pixels in diameter, with a spatial frequency of 0.6cpd (see Figure 6.1 – left). The orientation of each element was randomly sampled from a uniform distribution between 0° and 180°. Each element was set at 20% opacity, enabling overlaying elements to blend – the result being a narrowband isotropic visual noise stimulus (See Figure 6.1- right; note this stimulus is equivalent to taking random pixel noise and passing it through a band-pass isotropic filter). The root mean square (RMS) contrast of this stimulus was then fixed at 8% by altering the standard deviation of the pixel intensities to be 0.08 using equation 6.1:

$$\left(\frac{i's}{sd}\right) \times RMS \qquad \text{Equation 6.1}$$

Where $i's$ represents the pixel intensities minus the mean pixel intensity; sd represents the standard deviation of the pixel intensities; and RMS is the desired RMS contrast of the stimulus.

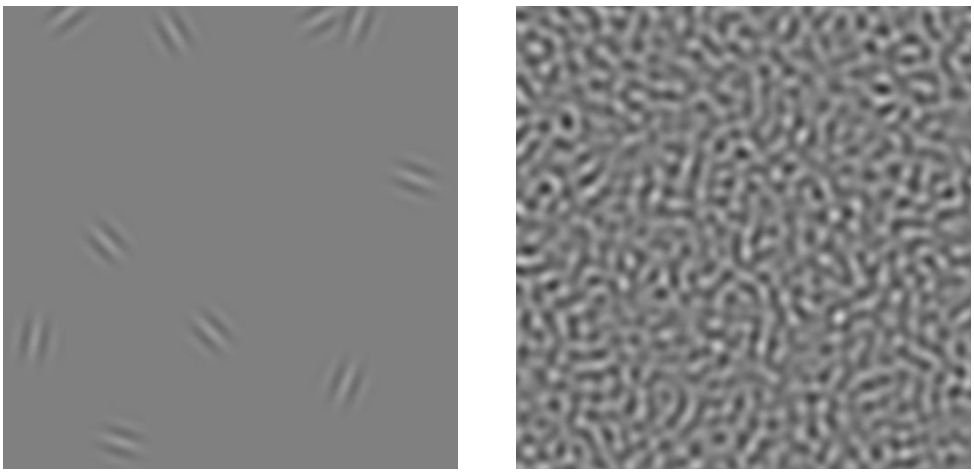


Figure 6.1. Left shows 10 individual sinusoidal grating elements with randomly sampled orientations (0-180°). Right shows 10,000 of these elements overlaid to produce the noise stimulus.

Over the course of the five minutes of pre-exposure, the contrast of this stimulus was modulated sinusoidally, according to Equation 6.2. This stimulus therefore modulated between 0RMS and 0.08RMS as a sinusoidal function of time (See Figure 6.2). The temporal frequency at which contrast was modulated was 1Hz or 15Hz, creating two pre-exposure conditions that are comparable with the low-pass filtered and high-pass filtered videos used in Chapter 5.

$$y(t) = 0.04cos(2\,\pi\,\omega_t) + 0.04 \qquad \text{Equation 6.2}$$

Where $y(t)$ represents the RMS contrast at time t, and TF represents the temporal frequency that was either 1Hz or 15Hz.

For precise timing, the temporal frequency of the modulation was determined by the frame rate of the monitor used for stimulus presentation (75Hz). This was achieved by altering the number of frames over which one cycle of the sinusoid occurred. For example, completing one cycle of the sinusoid over 75 frames results in a contrast modulation of 1Hz as one cycle of the sinusoid is completed per second (Figure 6.2 – blue). In comparison, completing one cycle of the sinusoid in 5 frames results in a 15Hz stimulus, as 15 cycles are completed per second (Figure 6.2 – red). Therefore, the temporal frequency of the stimulus is defined by altering the number of frames that set the period of the sine wave modulation. To further ensure precise timing, the experiment was monitored to ensure no frames were being dropped.

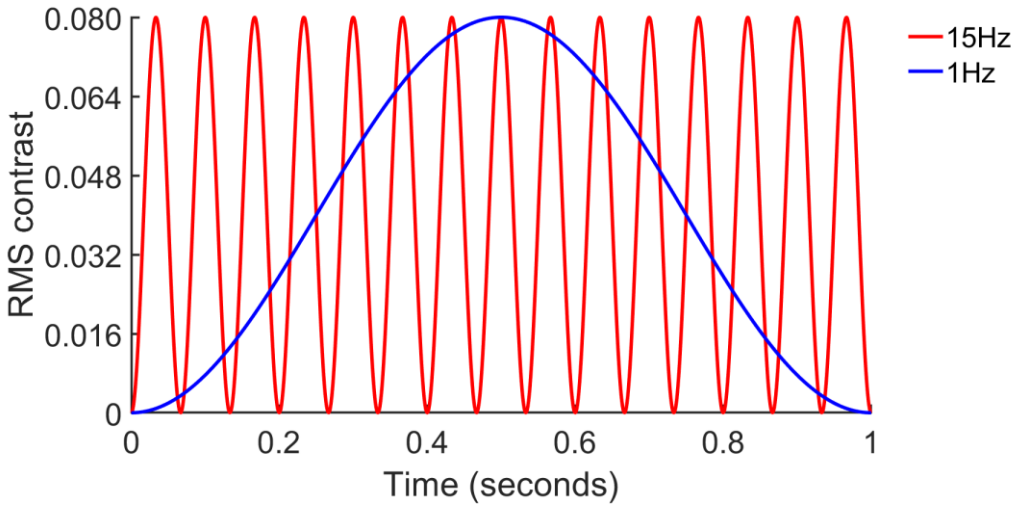


Figure 6.2. Sinusoidal modulation of the RMS contrast of the noise stimulus at 15Hz (red) and 1Hz (blue).

To avoid retinal afterimages the noise stimulus was updated every few seconds. To create the noise stimulus anew on each presentation was computationally demanding and led to dropped frames and imprecise stimulus presentation. To circumvent this, a single noise stimulus was presented throughout the pre-exposure with a circular mask, 22.6 degrees

in diameter. The circular mask allowed the stimulus to be both rotated (to a random orientation between 0-360deg) and flipped along the vertical meridian, effectively creating a new noise stimulus on each presentation. Flips and rotations were conducted at the point at which contrast was zero, so as not to distort the appearance of the stimulus and were therefore also synchronized with the temporal frequency of the stimulus. This created a dynamic visual noise stimulus with contrast modulating at a predefined temporal frequency.

The duration of the pre-exposure stimulus was reduced to 5 minutes (compared to 10 minutes in Chapter 5) because 10 minutes of fixating on a 15Hz flickering stimulus was deemed too long to be comfortable and would likely lead to lapses in observers' attention.

To further encourage observers to maintain attention, they were asked to complete a task during the pre-exposure phase. Small dots (1.82deg radius) were presented at a random location within the circular noise stimulus, for a duration of 800ms. The dots were made up of small subsections of the noise patch and rotated out of synchrony but at the same temporal frequency as the noise patch. This gave the appearance that they were 'floating' on top of the noise patch. Participants were asked to move the mouse and click when the cursor was over the dot as quickly and accurately as they could. Their overall percentage correct score was taken and if above 60% the run was included in the analysis. Runs where their performance was below 60% were excluded and the data recollected.

6.1.2.2. Measurement of the time course of the tilt aftereffect (TAE)

The subsequent measurement of the time course of the TAE was identical to that in Chapter 5 with two exceptions that aimed to increase the speed of

data collection. Firstly, the number of baseline trials was reduced from 50 to 25, as 25 trials is sufficient to see if bias remains stable around zero. Secondly, the number of adapting trials was reduced from 100 to 50 due to differences in the rate of adaptation growth in Chapter 5 being seen mainly during the first 50 trials.

One further alteration was to reduce the initial period of adaptation, presented before the first test stimulus, from 1066ms to 530ms. Adaptation grew rapidly over the course of the first trial in Chapter 5 meaning changes in the rate of initial adaptation were hard to identify. Halving the amount of initial adaptation aimed to reduce the rate at which adaptation grew on the first few trials.

6.1.2.3. Data analysis

As in Chapter 5, cumulative adapting time was calculated by subtracting the constant timings of stimulus presentation (0.7sec) away from the total duration of the trial, leaving the duration of adaptation on each trial for each observer. Estimates of the TAE were then binned in 3.5 second time bins to ensure that responses for the same duration of adaptation were aligned. Trials during the growth of adaptation are therefore plotted as cumulative adapting time rather than time from adapter onset.

6.1.3. Results

The time course of the TAE shows the characteristic growth and decay observed in Chapter 5. As can be seen in Figure 6.3, the magnitude of the TAE in all conditions is stable around zero prior to the onset of adaptation. Following the onset of adaptation, the magnitude of the TAE grows rapidly

over the first few seconds, reaching a magnitude of ~6deg in all conditions. Following adaptation offset, the TAE decays back to zero at a similar rate for all conditions.

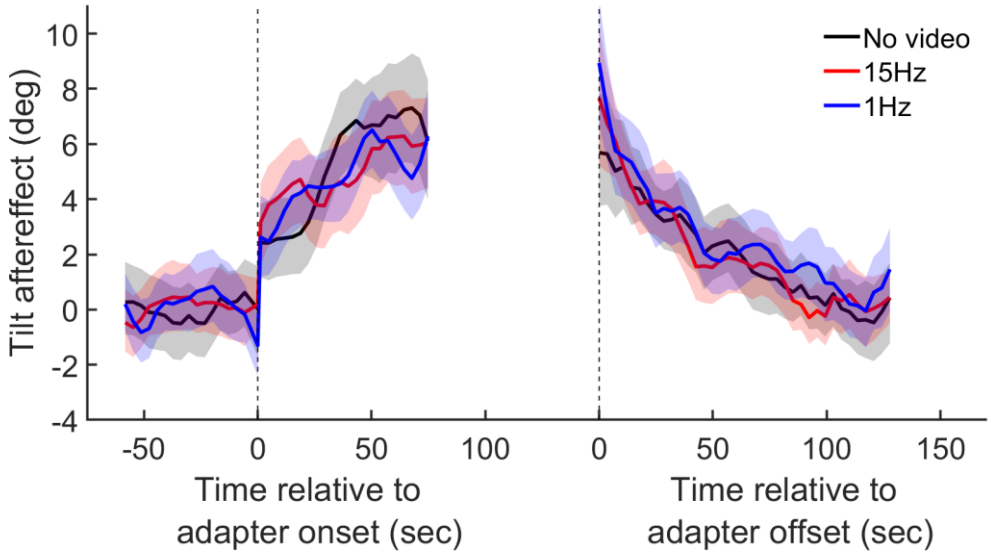


Figure 6.3 The time course of the tilt aftereffect. 5 minutes of pre-exposure to 15Hz (red) or 1Hz (blue) contrast modulated visual noise stimuli did not affect the time course of the TAE compared to no pre-exposure. Shaded regions represent 95% bootstrapped confidence intervals.

Consistent with the results of Chapter 5, permutation testing on the mean TAE during pre-adaptation trials (<0 sec) revealed no significant difference between the no-video and 15Hz conditions ($p = 0.22$), the no-video and 1Hz conditions ($p = 0.32$), or the 15Hz and 1Hz conditions ($p = 0.71$). Additionally, there were no significant differences in the average TAE during decay between the no-video and 15Hz conditions ($p = 0.93$), the no-video and 1Hz conditions ($p = 0.25$), or the 15Hz and 1Hz conditions ($p = 0.23$).

However, in contrast to Chapter 5, there were no significant differences in the mean growth of the TAE (0-75sec) between the no-video and 15Hz conditions ($p = 0.79$), the no-video and 1Hz conditions ($p = 0.64$), or the

15Hz and 1Hz conditions (p = 0.89). Five minutes of exposure to contrast modulations at 1hz or 15hz did not lead to a significant alteration in the subsequent time course of the TAE as compared to a condition in which there was no pre-exposure.

6.2. EXPERIMENT 2

6.2.1. Introduction

In Experiment 1, five minutes of pre-exposure to a noise patch modulating in contrast at either 1Hz or 15Hz did not affect the time course of the TAE relative to no pre-exposure. These results do not corroborate the findings in Chapter 5, where exposure to a video with predominantly high temporal frequency content led to an increase in TAE magnitude compared to exposure to a video with 1/f temporal frequency content.

The lack of replication of the findings in Chapter 5 could be due to differences in experimental design, that although seemed minor, may have affected the time course of the TAE. For example, the reduction of the pre-exposure duration from 10 minutes to 5 minutes was necessary for observer comfort and attention. However, the temporal tuning of observers' adapting mechanisms has likely been shaped by evolution and development in an environment with 1/f temporal statistics (Parraga, Troscianko, & Tolhurst, 2000). It is therefore not unreasonable to question whether 5 minutes of exposure to altered temporal statistics is sufficient to induce those adapting mechanisms to change their properties and alter the dynamics of adaptation. It could be that 10 minutes exposure naturalistic stimuli that have had their temporal spectra perturbed is sufficient to induce change, but 5 minutes exposure to an artificial stimulus is not.

Moreover, changes in the rate of adaptation are difficult to detect due to the rapid growth of adaptation after a single trial. As seen in both Experiment 1 and in Chapter 5, the magnitude of the TAE grows rapidly following the first few seconds of adaptation. Despite the reduction of the initial adapting duration in Experiment 1, the TAE reaches almost a third of its final magnitude in the first 3 seconds following adaptation onset. The dynamics of this growth are difficult to capture when the rate of sampling the time course is limited by the frequency of test interruptions, which in this case were roughly every three seconds. Measuring the initial few seconds of adaptation growth with higher temporal resolution could reveal subtler differences in the rate at which adaptation grows.

Therefore, although the effect observed in Chapter 5 was not replicated with artificial stimuli, it is possible that changes in the experimental design could reveal effects that were not observed in Experiment 1. For example, a longer pre-exposure duration and a temporally fine measurement of the first few seconds of adaptation may be necessary to observe differences in the dynamics of adaptation following exposure to contrast modulated noise stimuli.

Here, significant changes in experimental design were made to measure the effect of the pre-exposure stimuli used in Experiment 1 on the rate and magnitude of TAE growth. Firstly, the duration of adaptation was shorted to 4 seconds, and was sampled at a higher frequency than previous experiments to capture the rapid growth seen in the first few seconds of previous measurement. As there were no differences in the decay rates of the TAE in either previous experiments, it was decided to measure only the growth of the TAE.

Secondly, shortening the adapting duration allowed the use of continuous periods of adaptation followed by measurement of the TAE with a two-

alternative forced-choice paradigm that is more robust to changes in observer bias (Morgan, 2014). This is a more traditional measurement of the growth of the TAE, and is documented to produce logarithmic growth (Magnussen & Johnsen, 1986).

Finally, the pre-exposure period in both Chapter 5 and Experiment 1 occurred prior to the measurement of the time course of the TAE. It is possible that during measurement the effects of the pre-exposure are decaying in strength, especially in Experiment 1 where the pre-exposure was only 5 minutes in duration. Here, the pre-exposure was presented before the onset of TAE measurement, and was also 'topped up' before each adaptation period. The temporal statistics of the environment were therefore altered throughout measurement, setting the current statistics of the environment, rather than as a block beforehand.

6.2.2. Methods

The preceding experiments used the method of constant stimuli coupled with observers responding by reproducing the orientation of the test stimulus. This choice of methodology allowed an estimate of perceptual bias that could be averaged across observers and runs, with trial-by-trial temporal resolution (sampled every ~3 seconds). This methodology revealed the characteristic time course of the tilt aftereffect (TAE), but also highlighted that a large proportion of growth occurs during the first few seconds following adaptation onset.

Measuring only the first four seconds of adaptation growth negated the need for a trial-by-trial estimate of observers' perceptual bias. Instead, it was possible to present discrete continuous periods of adaptation followed by a test stimulus and the gathering of a response. The method of matching

used previously is subject to observer bias. Observers familiar with the process of visual adaptation might subconsciously alter the magnitude of their responses to increase or decrease the size of the TAE, depending on their perceived aims of the experimental task. To try and obtain a bias-free estimate of the TAE, this experiment adopted a spatial two-alternative forced-choice method (Morgan, 2014; Morgan, Melmoth, & Solomon, 2013).

6.2.2.1. Stimuli

The pre-exposure noise stimulus was identical to that used in Experiment 1, modulating sinusoidally in contrast at a rate of 1Hz or 15Hz. The timing of this stimulus was altered so that it was presented 30 seconds before the onset of adaptation and was then present constantly thereafter. This meant the initial 30 seconds acted to set the temporal statistics of the preceding environment, and subsequent presence kept this 'topped up' as the prevailing temporal statistics. Stimuli used to measure the growth of the TAE were presented on top of this pre-exposure stimulus.

The size and eccentricity of adapt, test, and mask stimuli was identical to previous experiments. Unlike previous experiments, where one stimulus was presented to either the left or right of fixation, here pairs of stimuli were presented either side of fixation at equal eccentricity. The orientation of adapting stimuli was 20 degrees clockwise from vertical. The optimal orientation for test stimuli was therefore near-vertical, which gave observers an inherent reference point around which to make their orientation judgements. The two 'test' stimuli consisted of a 'standard' stimulus and a 'test' stimulus. The standard stimulus was oriented either -2 or 2 degrees from vertical on an equal number of interleaved trials. The location of the standard stimulus was randomised on each trial to be either to the left or right of fixation. The test stimulus orientation was the standard

orientation plus one of 9 increments (-6, -4.5, -3, -1.5, 0, 1.5, 3, 4.5, 6). This results in two different test distributions, each centred on one of the standard orientations.

Multiple adapting trials were completed in a single experimental session. To prevent the accumulation of adaptation effects across trials, stimulus pair position varied between three locations. Location 1: 30deg and 210deg clockwise from vertical; location 2: 90deg and 270deg clockwise from vertical; and location 3: 150deg and 270deg clockwise from vertical. The location of stimulus pairs on each trial was randomized, with the limitation that the stimuli could not be in the same location on consecutive trials.

6.2.2.2. Procedure

Pre-exposure conditions were blocked with each of 6 observers completing the no pre-exposure, 1Hz and 15Hz conditions in a randomized order. Adapting durations were blocked, with each run being one of 5 adapting durations (0, 0.5, 1, 2, 4 seconds). The zero second adapting condition was included to identify observer's baseline bias in orientation perception in each of the pre-exposure conditions. The duration selected was randomised on each run to ensure there were not order or practice effects. Each observer completed four runs at each of the 5 adapting durations, for each of the three pre-exposure conditions – totalling 60 runs (~8 hours).

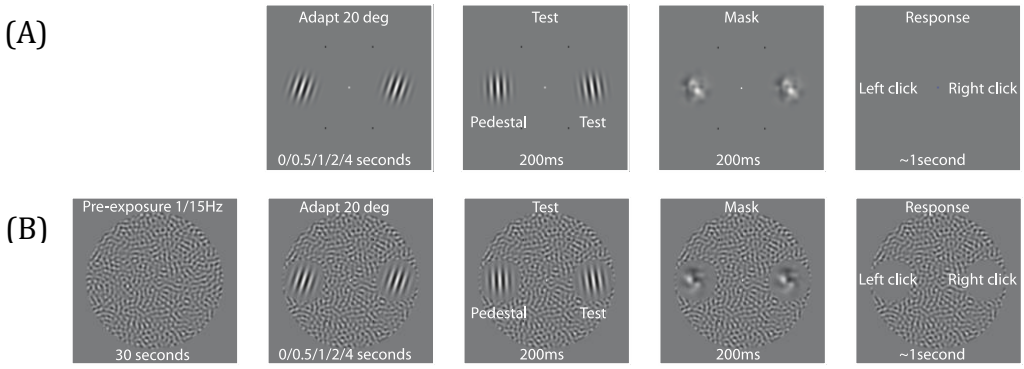


Figure 6.4. (A) Trial structure on no pre-exposure runs. (B) Trial structure on pre-exposure runs. Pre-exposure condition and adapting duration were blocked with observers completing four runs for each adapting duration at each of three pre-exposure conditions.

Each run consisted of six repetitions of each standard-increment pairing, to create a total of 108 trials – the duration of each run varied depending on the adapting duration of that run. On pre-exposure runs, the pre-exposure stimulus was presented 30 seconds before the first adaptation trial, and then constantly for the remainder of the run.

On each trial, observers were presented with 0, 0.5, 1, 2 or 4 seconds of adaptation. There was then a 250ms frame ISI followed by two test Gabors (200ms), and a noise mask (200ms). Following this, observers had to decide which of the two test stimuli (left or right of fixation) was closer to vertical. They responded using the left and right mouse buttons and had 1350ms to respond before the start of the next trial (see Figure 6.4).

6.2.2.3. Data analysis

Data was collapsed across runs and observers. The data from each adapting duration was analysed separately, starting by separating the two standard orientations. Following this, the proportion of times the standard stimulus

was chosen as being closer to vertical was plotted against the orientation of the test stimulus.

An observer model was then fitted to these data, similar to that described by Morgan (2014). In keeping with signal detection theory (Green and Swets, 1966), we modelled the perceived orientation of standard and test stimuli as normal distributions S and T, such that:

$$S \sim N(p + \mu, \sigma^2/2) \text{ and } T \sim N(p + t + \mu, \sigma^2/2)$$

where σ^2 is the variance of the internal noise, p and $p + t$ are the orientations of standard and test stimuli respectively, and μ represents any perceptual bias, induced by the adapting stimuli for example. On any given trial, the model chooses the stimulus with a perceived orientation closest to zero (vertical). Therefore, the probability of a 'standard more vertical' response in our orientation judgement task is given by:

$$\begin{aligned} \Pr(s) &= \Pr(|S| < |T|) \\ &= \Pr\left(\frac{S^2}{T^2} < 1\right). \end{aligned}$$

S^2/T^2 is a random variable having a doubly non-central F-distribution. Its numerator's non-centrality parameter is $2(p + \mu)^2/\sigma^2$ and its denominator's non-centrality parameter is $2(p + t + \mu)^2/\sigma^2$, with both numerator and denominator having 1 degree of freedom (Morgan, 2014). Like Yarrow, Martin, Di Costa, Solomon, & Arnold (2016), we made use of a saddle-point approximation to the doubly non-central F cumulative distribution function (Butler & Paolella, 2002; Paolella, 2007). This Matlab code is available online (Yarrow, 2018, Retrieved from: http://www.hexicon.co.uk/Kielan/).

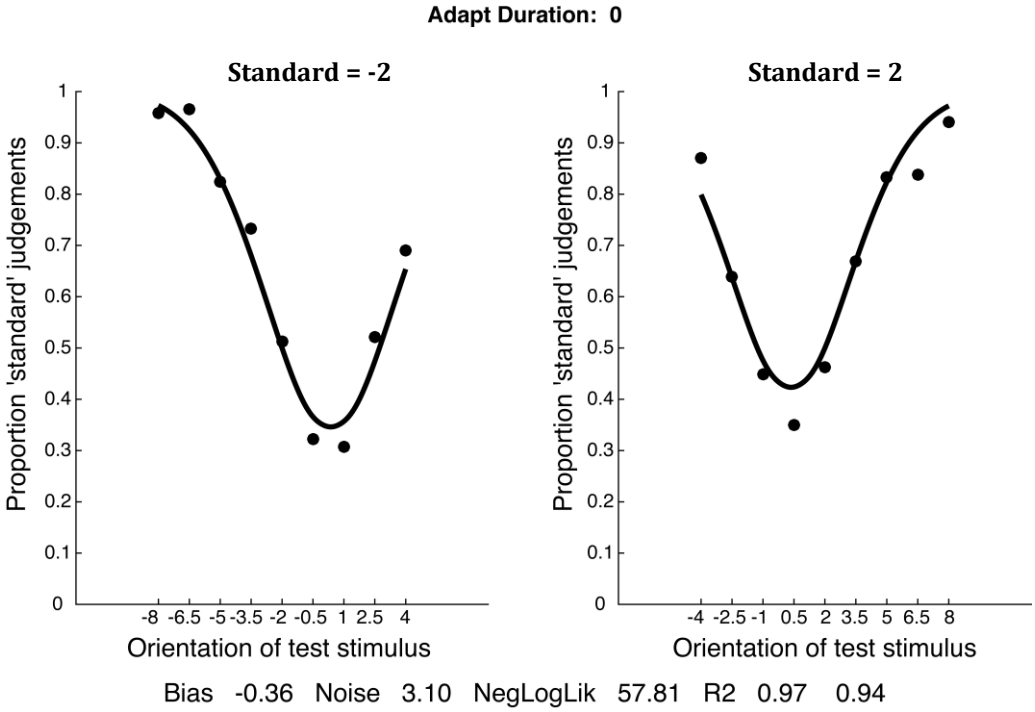


Figure 6.5. Data from each standard orientation are separated (-2 left, +2 right). The proportion of times observers picked the standard stimulus as a function of the orientation of the test stimulus is plotted. An observer model is fitted to both data sets simultaneously. The model has two free parameters: Bias, which sets the minima of the functions and represents perceived vertical, and noise which sets the width of the functions and represents error.

This observer model was fitted to the data for both standard orientations simultaneously, resulting in a single set of parameter estimates across the two standards (Figure 6.5). The model has two free parameters: the noise, which determines the width of the functions (internal error); and the bias, which determines the minimum point of the functions (subjective vertical).

Best fitting parameter values were obtained using a maximum likelihood criterion. That is, we selected the parameters that maximised the probability of observing the experimental data given the predictions of the observer model. In practice, this was achieved by minimising the model's negative log likelihood using the fmincon function in Matlab. On each

iteration of the fitting procedure, the binomial distribution was first used to calculate the probability of each data point in the psychometric function given the corresponding estimate of Pr(s) from the model. To avoid rounding errors for very small values, we took the logarithm of each of these likelihoods, then summed them across all combinations of standard and test orientations (note, this is numerically equivalent to the logarithm of the product of all likelihoods).

The observer model was fit to the data from each of the 5 adapting durations, and the parameters from the best fitting model returned for each. The main parameter of interest is the bias, which was our measure of the TAE at each adapting duration.

95% confidence intervals on the bias parameter were achieved through non-parametric bootstrapping. A binomial distribution was used to simulate new data sets, using the number of trials, and the probability of picking the standard. The same observer model used to fit the original data was then used to fit the simulated data set and gain a new estimate of the bias parameter. This was repeated 100 times to create a distribution of parameter estimates centred on the original estimate. The 2.5th and 97.5th percentiles of the new parameter estimate distribution are were taken as 95% confidence intervals.

6.2.3. Results

The bias parameter returned from the best fitting model was plotted for each adapting duration and a logarithmic function fitted (Figure 6.6, Equation 6.1):

$$y = intercept + slope \times \log_{10}(x) \qquad \text{Equation 6.1}$$

From this, the slope of the function indicates the rate of the adaptation growth, which was compared across pre-exposure conditions.

When there was no adapting stimulus presented (0sec adaptation condition) there was a small, ~0.5deg, bias in orientation perception. This bias in baseline orientation perception was consistent across the no pre-exposure, 1Hz pre-exposure, and 15Hz pre-exposure conditions, suggesting that pre-exposure does not alter the perception of oriented stimuli per se. As in previous experiments, 'baseline correction' was conducted whereby the bias in orientation perception observed in the zero-adaptation condition was subtracted from the bias observed at all other adapting durations for each pre-exposure condition.

In the baseline condition, with no pre-exposure, the TAE grew logarithmically as adapting duration increased from 0.5 to 4 seconds, reaching a final magnitude of 1.3deg (see Figure 6.6). The slope of the logarithmic regression fit was significantly different to zero ($t(2) = 19.69$, $p = .003$), suggesting significant growth of the TAE with increasing adapting duration.

In the low temporal frequency pre-exposure condition, the TAE grew to a similar magnitude as that in the baseline condition following 0.5 seconds of adaptation (~0.9deg). However, adaptation did not increase in magnitude for increasing adapting durations of 1 second, 2 seconds, or 4 seconds. This was confirmed by a logarithmic regression slope that is not significantly different to zero ($t(2) = 0.44$, $p = .71$). The magnitude of the TAE was smaller still in the high temporal frequency condition, with an initial magnitude of 0.5deg following 0.5sec of adaptation and growing to only 0.6deg following 4 seconds of adaptation. This rate of growth was not significantly different to zero($t(2) = 1.48$, $p = .28$).

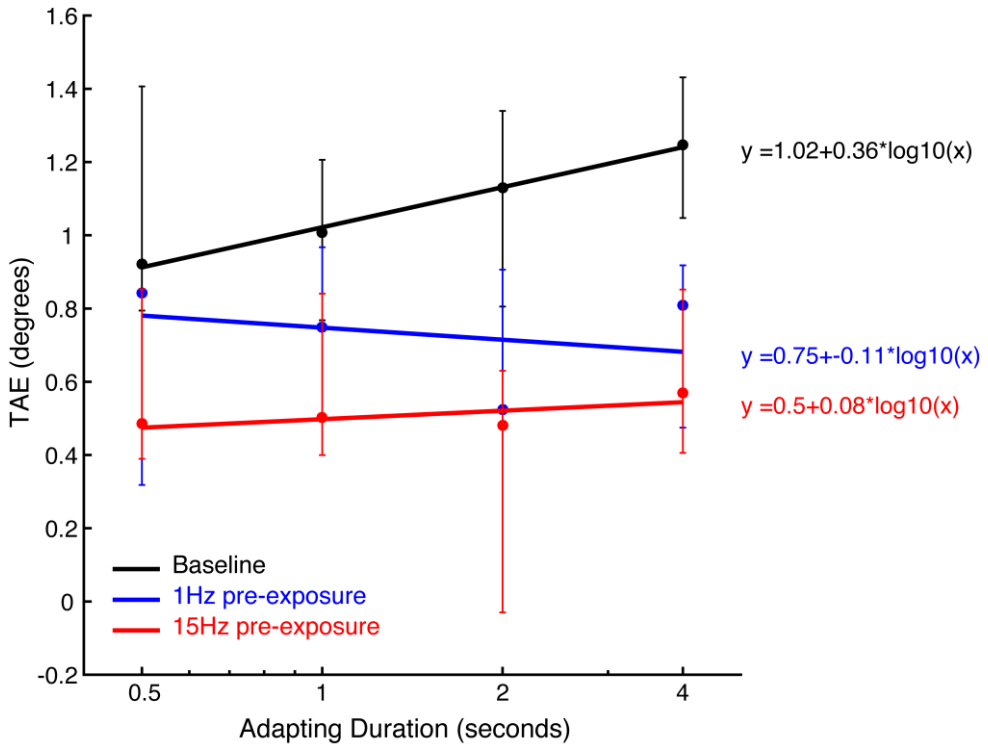


Figure 6.6. The magnitude of the TAE (degrees) is plotted against the duration of continuous adaptation. Error bars represent 95% bootstrapped confidence intervals.

Exposure to contrast modulating noise stimuli therefore suppressed the growth of the TAE relative to no pre-exposure. The magnitude of this suppression was larger when observers were exposed to 15Hz noise stimuli compared to 1Hz noise stimuli. To disambiguate these findings, a single observer (SM) completed a further experiment examining the magnitude of the suppressive effect when altering the contrast and temporal frequency of the pre-exposure noise stimulus. This experiment aimed to identify whether the suppressive effect was tuned to the contrast and temporal frequency of the pre-exposure stimulus. The magnitude of the TAE following 2 seconds of adaptation was measured during pre-exposure to the same noise patch used in the original experiment modulating in contrast at 4 temporal frequencies (1, 3.25,7.5,15Hz). At each of the temporal frequencies, there was a 0.08RMS contrast condition and a 0.04RMS

contrast condition. The adapting duration of two seconds was chosen as it produced a measurable effect in the original experiment and enabled fast data collection.

Figure 6.7 shows the magnitude of the TAE following two seconds of adaptation for each of 8 contrast-temporal frequency pairings. The black dashed line represents the magnitude of the TAE following no-pre-exposure as baseline comparison. The magnitude of adaptation suppression is largest when the pre-exposure stimulus is at the highest contrast (Max = 0.08RMS, time averaged = 0.04RMS), with the magnitude of the TAE at all temporal frequencies being smaller than that when the pre-exposure stimulus was lower in contrast (Max = 0.04RMS, time averaged = 0.02RMS). The suppressive effect is also tuned to temporal frequency at

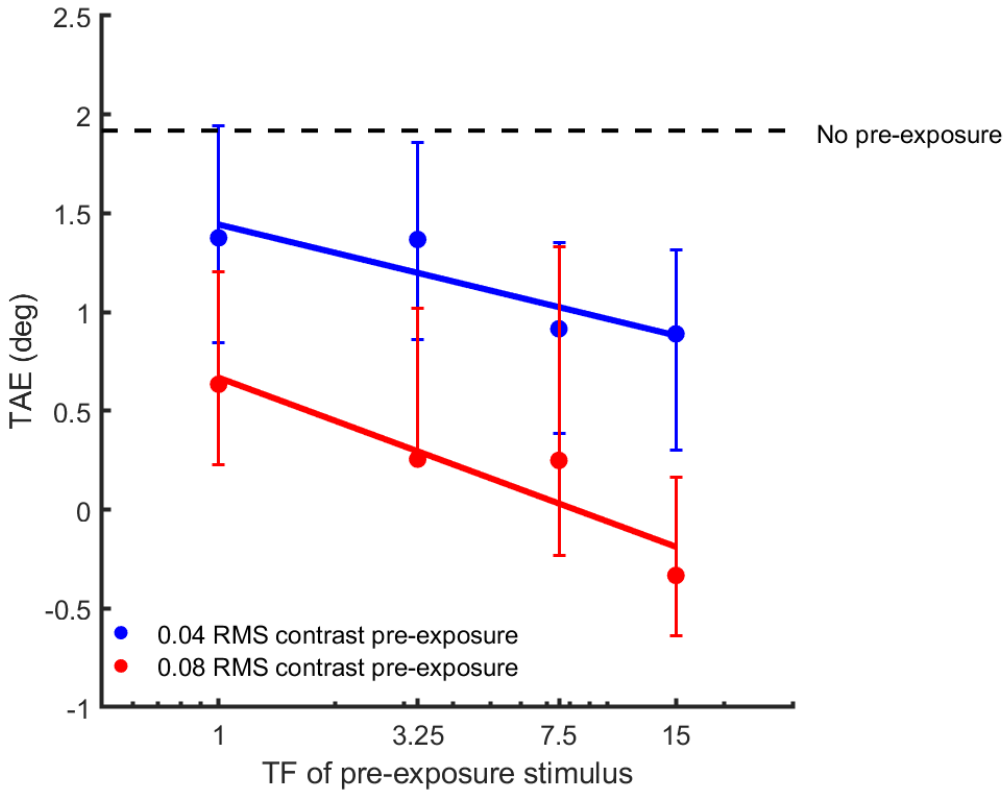


Figure 6.7 The magnitude of the TAE in 8 different pre-exposure conditions produced by a combination of two contrasts and four temporal frequencies. The dashed line is the magnitude of the TAE following 2 seconds of adaptation in the no pre-exposure condition for subject SM.

both contrasts measured, with lower temporal frequency stimuli producing less suppression of adaptation.

These results suggest that the presence of the pre-exposure stimulus reduced the effectiveness of the adapting stimulus. A further experiment aimed to identify if there was an accompanying perceptual change in contrast or temporal frequency of the adapting stimulus when the pre-exposure stimulus was present. Therefore, a single observer (SM) completed a contrast and temporal frequency matching task.

To the right of the vertical meridian, a vertically oriented 'test' adapting Gabor stimulus was presented alongside either 1Hz, 15Hz, or no pre-exposure stimulus depending on the block. The parameters of the stimuli to the right of the vertical meridian were therefore identical to those in the above experiment. To the left of the vertical meridian was the 'matching' adapting stimulus, which was identical to the 'test' adapting Gabor and in the same location as the stimuli presented in the above experiment – the only difference was its contrast on contrast matching trials and its temporal frequency on temporal frequency matching trials.

When contrast was the property to match, the contrast of the matching stimulus began at 0.3, 0.4, 0.5, 0.6 or 0.7 Michelson contrast, randomised on each trial. The observer then adjusted the contrast of the matching stimulus in 0.05 steps until they perceived the contrast of the matching stimulus and the test stimulus to be equated. The temporal frequency of the matching stimulus during contrast matching conditions was set to the same as that in the above experiment (3Hz).

During temporal frequency matching conditions, the temporal frequency of matching stimulus began at 1, 2, 3, 4, or 5Hz, randomized on each trial. The observer could then adjust the temporal frequency of the matching stimulus

until they believed it to perceptually match that of the test adapting stimulus. The contrast of both stimuli was held constant at 1 Michelson contrast during temporal frequency matching conditions, which was the same as the above experiment.

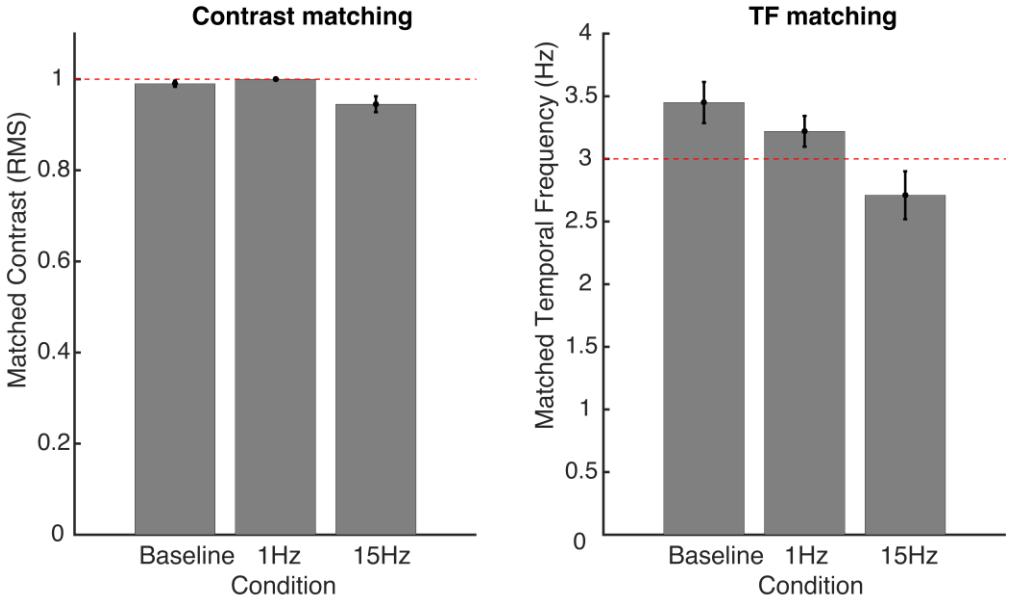


Figure 6.8. Results of the contrast and temporal frequency matching control experiment. Red dashed lines represent the actual contrast and temporal frequency of the 'test' stimulus. The contrast/ temporal frequency values needed to perceptually match the 'matching' stimulus to the 'test' stimulus are plotted for the different pre-exposure conditions. The pre-exposure stimuli make little difference to the perceived contrast of the adapting stimuli. However, the 15Hz stimulus significantly reduces the perceived temporal frequency of the adapting stimulus.

Figure 6.8 shows that neither the 1Hz nor 15Hz pre-exposure stimulus altered the perceived contrast of the adapting stimulus when compared to baseline perception. However, the 15Hz pre-exposure stimulus significantly reduced the perceived temporal frequency of the adapter relative to baseline perception. The temporal frequency of the 3Hz-adapting stimulus appears to be counter-phasing significantly slower (2.7Hz) when the 15Hz stimulus is present than when it is not (3.5Hz) ($t(9) = 2.51$, $p = 0.03$).

When the 15Hz noise was presented, the adapting stimulus was perceived to have a contrast of 0.94 Michelson and be counter-phasing at a rate of 2.71Hz. Although these seem like subtle alterations, it is possible that this led to the suppressed magnitude of the TAE. To test this, the magnitude of the TAE was measured following 2 seconds of adaptation to an adapting stimulus with 0.94 Michelson contrast and counter phasing at 2.71Hz, with no pre-exposure present. The magnitude of the TAE was 1.9 degrees and was therefore not comparable to that seen during the 15Hz pre-exposure (-0.4deg) in the original experiment. It was instead close to that seen during the no-pre-exposure condition (1.8deg). Thus, the suppressive effect of exposure to 1Hz and 15Hz contrast modulating noise stimuli in Experiment 2 was not the result of reduced perceived contrast or temporal frequency of the adapting stimulus.

6.3. DISCUSSION

Chapter 5 showed that 10 minutes of exposure to a naturalistic video with predominantly high temporal frequency content was sufficient to increase the rate and magnitude of the subsequently measured TAE time course. Here, Experiment 1 aimed to replicate this effect with 5 minutes of pre-exposure to artificial stimuli modulating in contrast at high (15Hz) or low (1Hz) temporal frequencies. The results of Experiment 1 did not replicate the observed effect, with no significant differences in the time course of the TAE between pre-exposure and baseline conditions.

One possible explanation for the lack of replication was the reduced pre-exposure stimulus duration. Using a continuous block of pre-exposure before measurement of the TAE time course meant that the duration of the artificial pre-exposure stimulus was limited. In addition, the methodology

used in Chapter 5 and Experiment 1 led to a TAE that grew rapidly during the first few seconds of adaptation.

To address these issues, Experiment 2 aimed to explore the effect of exposure to altered temporal frequency content using the same artificial stimuli but focusing on the first few seconds of adaptation. A shorter adapting duration permitted the measurement of adaptation using continuous periods of adaptation and a robust two-alternative forced-choice response. In addition, the pre-exposure stimulus was present both before and during adaptation ensuring the observer's exposure to altered temporal statistics was consistent throughout the measurement of the TAE.

The results of Experiment 2 showed logarithmic adaptation growth in the no pre-exposure condition that is consistent with previous continuous measurement of TAE growth at longer timescales (Magnussen & Johnsen, 1986). The magnitude of adaptation, however, is smaller than expected from previous experiments in this thesis. Following only 2.5 seconds of uninterrupted adaptation in Chapter 5 and Experiment 1, the magnitude of the TAE had reached ~3deg. In contrast, measured aftereffects here reach only 1.2deg after 4 seconds of uninterrupted adaptation. As the adapting stimuli were identical in these experiments, differences in the magnitude of the TAE in baseline conditions could be a result of the different response methodologies used. Comparing the magnitude of the TAE following equivalent durations of adaptation measured with both matching and two-alternative forced-choice measures would be an interesting avenue for future research.

Experiment 2 also showed that exposure to a 1Hz or 15Hz contrast modulating noise stimulus suppressed the growth of the TAE relative to the no pre-exposure baseline condition, with larger suppression shown for higher contrast and higher temporal frequency pre-exposure stimuli. This

contradicts the findings from Experiment 1, where 5 minutes of exposure to a 1Hz or 15Hz contrast modulating noise stimulus had no measurable effect on the time course of the TAE in comparison to a no pre-exposure condition.

Although this effect is unexpected, and difficult to explain, it does show lawful tuning of the degree of suppression to the contrast and temporal frequency of the pre-exposure stimulus. The main difference in stimulus presentation was the continuous nature of the pre-exposure stimulus throughout adaptation measurement, rather than a fixed block of pre-exposure. This resulted in the pre-exposure stimulus surrounding the adapting stimulus during presentation, which did not occur in Experiment 1.

It was therefore proposed that the presence of the pre-exposure stimulus during viewing of the adapting stimulus might change the perceived contrast or temporal frequency of the adapter. The magnitude of the TAE is larger for high contrast stimuli (Harris & Calvert, 1989) and may also be affected by adapter temporal frequency. If the presence of the pre-exposure stimuli reduced the perceived contrast or temporal frequency of the adapting stimulus, it may have led to suppression of the TAE.

However, a contrast and temporal frequency matching experiment revealed a small reduction in the perceived temporal frequency of the adapting stimulus only during 15 Hz pre-exposure, and no changes in perceived contrast for either 15Hz or 1Hz pre-exposure. When the TAE was measured with the perceived contrast and temporal frequency from the matching experiment, there was no suppression of the aftereffect. The suppressive effect seen in the original experiment is therefore unlikely to be due to changes in the perceived contrast and temporal frequency of the adapting stimulus.

The tuning of the suppressive effect to pre-exposure contrast and temporal frequency still suggests that the pre-exposure stimulus is exerting a direct effect on subsequent adaptation. The lack of suppressive effect in Experiment 1 also suggests that the pre-exposure stimuli only result in adaptation suppression when presented alongside the adapting stimulus rather than preceding it. However, the precise mechanism by which this occurs cannot be elucidated from the current experiments.

The results of neither experiment support the finding from Chapter 5 that pre-exposure to high temporal frequency content results in an increased rate and magnitude of TAE growth compared to exposure to 1/f temporal statistics. There are marked differences in the experimental design of these two experiments, but it appears that pre-exposure to altered temporal frequency in artificial stimuli does not result in matching of the dynamics of adaptation.

However, it is important to note that the baseline conditions in Chapters 5 and 6 were different, complicating comparisons between the two experiments. The baseline condition in Chapter 5 was an unfiltered video that contained 1/f naturalistic statistics, whilst the baseline here was no-video pre-exposure. Although both likely expose the observer to 1/f temporal statistics, there are still differences in their properties. Natural viewing brings with it everything associated with viewing the world from first person perspective, such as optic flow signals (Lee, 1980), motion induced by eye and head movements (Steinman & Collewijn, 1980), and interaction with one's environment (Harris & Still, 1981; Ichikawa & Masakura, 2006), which may affect the dynamics of adaptation in a way that a TV show does not. Future experiments aim to use consistent baseline conditions to allow for easy comparison between experimental conditions.

It is possible that the artificial stimuli used here are unable to effectively simulate environmental changes that would result in meta-adaptation. Little is known about how adaptation is modulated by the magnitude and reliability of a change in environmental statistics. Exposure to a stimulus which is narrowband in terms of spatial and temporal statistics is a very large deviation from the natural environment. Exposure to environments that are very far from natural statistics may evoke different responses to those which are more subtly perturbed. It could be that large deviations are considered noise, or anomalies that will be short-lived, to which adjusting the dynamics of adaptation would be a waste of neural resources. Exploring how the magnitude, duration, and reliability of a change in stimulus statistics affects the time course of adaptation would be an interesting avenue for future research.

Chapters 5 and 6 thus highlight the trade-off between stimulating the visual system with the input it has evolved to process and having readily parameterized stimuli that are easily controlled. Due to the complexity of naturalistic stimuli (Rust & Movshon, 2005), there are marked differences between the pre-exposure stimuli used here and in Chapter 5. It is plausible that the observed effect in Chapter 5 was driven by a property of the naturalistic stimulus that was not present in the artificial stimulus used here. One possibility is that the filtered videos used for pre-exposure in Chapter 5 differed in contrast, whilst the noise stimuli used here were matched in contrast. Further analysis of the differences between the pre-exposure video stimuli and how they compare to the artificial stimuli used here is needed to identify if a property other than temporal frequency statistics drove the effect observed in Chapter 5.

7

The time course of the tilt aftereffect following exposure to altered RMS contrast

7.1. INTRODUCTION

In Chapter 5, 10 minutes of pre-exposure to a temporally high-pass filtered video resulted in the increased rate and magnitude of the tilt aftereffect (TAE) in comparison to exposure to an unfiltered video with naturalistic 1/f temporal statistics. Due to the complex nature of the pre-exposure stimuli, further experiments aimed to identify if the observed effect was driven solely by the temporal frequency perturbations in the videos.

Two experiments in Chapter 6 aimed to replicate the observed effect with artificial stimuli that differed only in the rate at which contrast modulated sinusoidally over time - 1Hz versus 15Hz. Using the same methodology as Chapter 5, Chapter 6 (Experiment 1) showed that 5 minutes of exposure to contrast modulated noise stimuli with fast or slow temporal frequency was insufficient to alter the time course of the TAE. Differences in the baseline pre-exposure stimuli used in each experiment made direct comparisons between conditions difficult and highlighted the need for a representative baseline common to all experiments. Moreover, in both Chapters 5 and 6 the TAE grew rapidly over the first few seconds, making changes in the rate of growth difficult to observe.

To address the methodological limitations of previous chapters, Chapter 6 (Experiment 2) adopted a more traditional methodology, measuring

blocks of continuous adaptation, and used a two-alternative forced-choice measurement of the TAE. To increase exposure to altered temporal statistics, the pre-exposure stimulus was presented both before and during the measurement of adaptation. Altering the pre-exposure in this way led to the suppression of the TAE, which was larger for pre-exposure stimuli of higher contrast and temporal frequency. Whilst the precise mechanism of this effect is not known, exposure to artificial stimuli with altered temporal frequency content does not appear to affect the dynamics of subsequent adaptation in the way hypothesised in Chapter 5. This could be due to the artificial stimuli being too simplistic to reproduce the original effect, or it could be that the original effect was driven by a property other than the temporal frequency content of the videos.

One such property is the average contrast of the video stimuli. The tilt aftereffect is modulated by adapter contrast (Harris & Calvert, 1989), therefore exposure to altered visual contrast could affect the magnitude of subsequent adaptation. The average RMS contrast over a selection of 500 frames was calculated using Equation 7.1 for the pixels in the circular radius of the stimuli locations for each pre-exposure stimulus.

$$\sqrt{\frac{\sum \frac{\sum (I-\bar{I})^2}{WH}}{500}} \qquad \text{Equation 7.1}$$

Where *'I'* is a matrix of intensity values for a single video frame with width '*W*' and height '*H*', and $\bar{I}$ represents the mean of all pixel intensities. The unfiltered video had the highest contrast at 0.13RMS, followed by the low-pass filtered video at 0.09RMS, and the high-pass filtered video at 0.04 RMS. The same analysis was completed on the visual noise stimuli which

had a maximum contrast of 0.08RMS but an average contrast of 0.04RMS over 500 frames.

In Chapter 5, the largest TAE was produced following exposure to the high-pass filtered video, followed by the low-pass filtered video, and the smallest TAE was produced by the unfiltered video. Therefore, the magnitude of the TAE decreases as a function of increasing pre-exposure contrast.

In Chapter 6, the artificial pre-exposure stimuli were matched in contrast, and there was no difference in the magnitude of the TAE following the 1Hz and 15Hz pre-exposure. There was also no difference in the magnitude of the TAE between either pre-exposure condition and the no pre-exposure baseline. This is likely due to the no pre-exposure baseline (natural viewing) having a similar RMS contrast to the noise pre-exposure stimuli, at ~0.01RMS (Bex & Makous, 2002). Moreover, the TAE following no-pre-exposure baseline in Chapter 6 was larger in magnitude than the TAE following an unfiltered video baseline in Chapter 5. This could be due to the higher contrast of the unfiltered video (0.13RMS) compared to natural viewing (0.01RMS). This analysis therefore suggests that the original effect with naturalistic stimuli in Chapter 5 was due to exposure to altered contrast rather than the temporal frequency content of the videos.

To formally test this hypothesis four of the observers that completed the experiment in Chapter 5 completed two new pre-exposure conditions. A contrast-matched pre-exposure condition consisted of a video that was temporally unfiltered but matched in contrast to the high-pass filtered video used in Chapter 5. Should the time course of the TAE following exposure to this stimulus be of similar rate and magnitude to that following exposure to the high-pass stimulus, it would confirm that altered contrast, rather than temporal frequency, drove the original finding. A no

pre-exposure baseline condition was also included to allow comparison between all previous experiments.

7.2. METHODS

Four subjects who had completed the experiment in Chapter 5 acted as observers in this experiment. In the contrast-matched condition they viewed a 10-minute segment of a popular TV show that was temporally unfiltered but matched in average RMS contrast to the high-pass filtered video from Chapter 5. Immediately following, the time course of the TAE was measured on a trial-by-trial basis identical to that in Chapter 5. In the no-exposure condition, subjects completed the measurement of the time course of the TAE without previous exposure to any video stimuli.

7.2.1. Stimuli

To create the contrast-matched videos, the RMS contrast of each frame of an unfiltered, 10-minute, black and white video was matched to the RMS contrast of a frame from the high-pass filtered video. This was achieved by matching the 2D histogram of each frame in the unfiltered video to the histogram of a representative selected frame from the high-pass filtered video.

The histogram of grayscale values was computed for the unfiltered frame and the high-pass reference frame. Following this, the cumulative distribution functions (CDF) of each images' histogram were calculated. For each grey level in the CDF the grey level needed to match the unfiltered value to the high-pass value was calculated. This function was

applied to each pixel in the frame. The CDF of each unfiltered video frame was matched to that of the high-pass filtered frame.

The temporal frequency spectra of all videos are shown in Figure 7.1. The unfiltered and contrast-matched videos show classic 1/f temporal frequency spectra, whilst the high-pass and low-pass filtered videos have low temporal frequencies and high temporal frequencies attenuated respectively. The method for measuring the time course of the TAE was identical to that in Chapter 5.

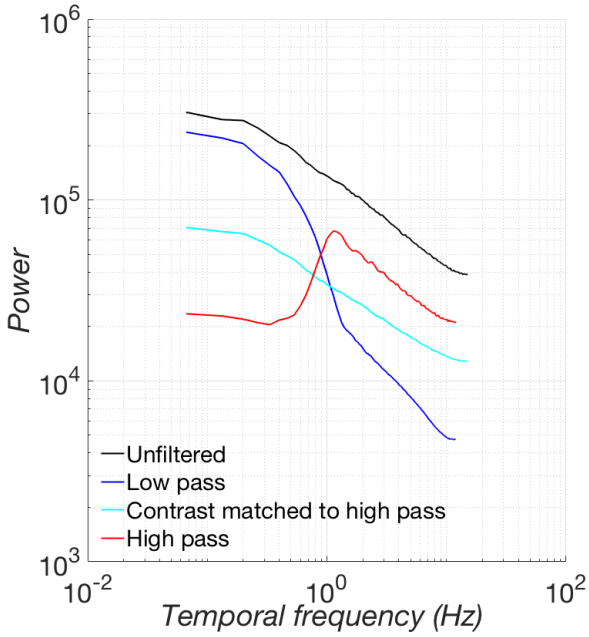


Figure 7.1. Temporal frequency spectra of pre-exposure video stimuli. The unfiltered and contrast-matched videos show characteristic 1/f temporal frequency statistics. The high-pass filtered video had temporal frequencies below 0.8Hz attenuated whilst the low-pass filtered video had temporal frequencies above 0.8Hz attenuated.

7.2.2. Data analysis

The TAE was calculated identically to that in Chapters 5 and 6. Data from the four observers was extracted for the high-pass, low-pass, and unfiltered video pre-exposure conditions from Chapter 5 to allow

comparison between the new conditions and those previously collected. These data were collected 2 years apart, and it was interesting to identify if there was consistency in the time course of the TAE.

Prior to them completing the high-pass, low-pass, and unfiltered video conditions in Chapter 5, three subjects completed no-video pre-exposure runs to pilot the task. Figure 7.2 shows the time course of the TAE following no-video exposure when the data for Chapter 5 were collected (2015) and following no-video exposure in the current experiment (2018). The time course of the TAE changes very little from the first measurement to the second, confirming the suitability of comparing data collected 3 years apart from the same observers. Permutation testing revealed there were no significant differences between the two time courses during pre-adaptation (mean difference= 0.11deg, p= 0.92) adaptation growth (mean difference= 1.10 , p= 0.97), or adaptation decay (mean difference = -1.58, p= 1).

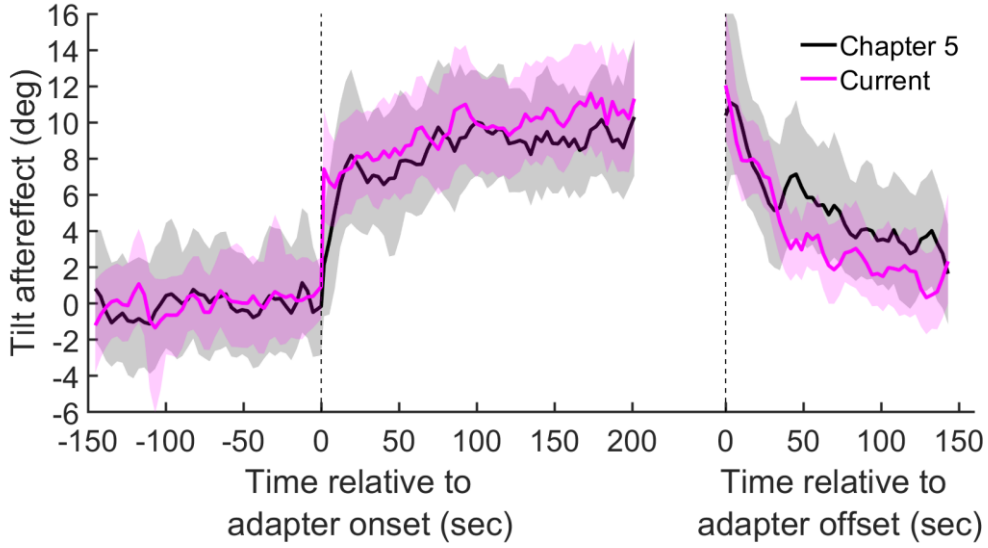


Figure 7.2. The time course of the tilt aftereffect following no-video pre-exposure when the data for Chapter 5 were collected in 2015 (magenta) and the current no-video data from 2018 (black).

7.3. RESULTS

The time course of the TAE following no-video (magenta), contrast-matched video (cyan), high-pass video (red), low-pass video (blue) and unfiltered video (black) pre-exposure is plotted in Figure 7.3. Prior to adaptation onset (time < 0) each time course is corrected for baseline bias and is stable at zero degrees. Following adaptation onset (time = 0) the magnitude of the TAE in all conditions grows as a function of adapting duration, and decays following adaptation offset. Each condition therefore shows the classical time course of the TAE.

The TAE time course following high-pass, low-pass and unfiltered pre-exposure shows the same pattern of results for four observers as was seen in the full data set. The TAE following high-pass pre-exposure built more rapidly and reached a greater magnitude than that following unfiltered pre-exposure. However, permutation testing revealed that the differences in the growth of adaptation in these two conditions was not statistically significant (mean difference 0-200sec after adaptation onset =2.24deg, p = .14). As seen in Chapter 5, exposure to a video with predominantly low temporal frequencies does not alter the time course of the TAE when compared to that following exposure to an unfiltered video (mean difference 0-200sec after adaptation onset =0.71deg, p = .55).

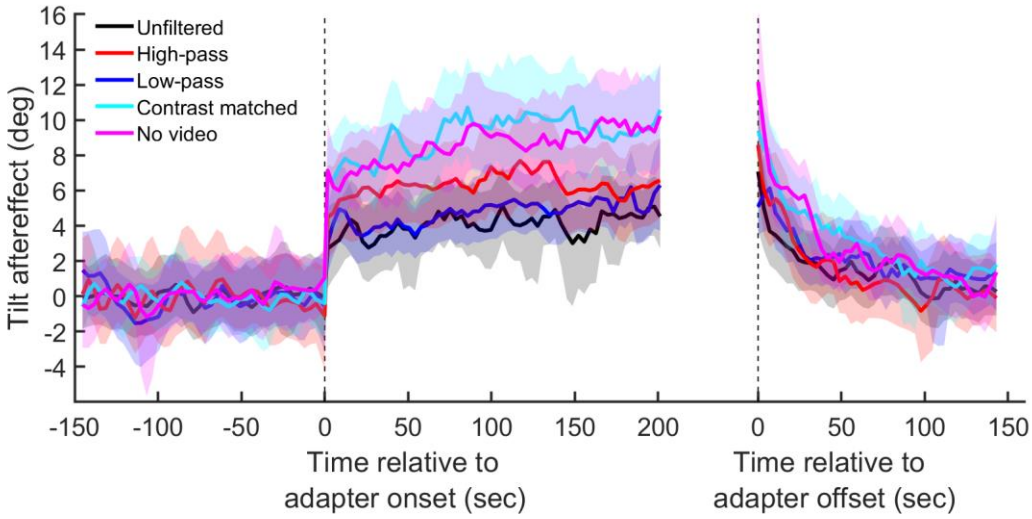


Figure 7.3. The time course of the TAE following Contrast matched (cyan) and no-video (magenta) pre-exposure. High-pass (red), low-pass (blue), and unfiltered (black) video data is reproduced for 4 observers from Chapter 5.

The unfiltered contrast-matched video and the high-pass filtered video differed in temporal frequency statistics, but not in contrast. In line with the hypothesis that contrast drove the observed effect in Chapter 5, the time course of the TAE following exposure to the contrast matched video reached a similar magnitude to that following high-pass pre-exposure. Permutation testing revealed there was no statistically significant difference in TAE growth between these two conditions (mean difference 0-200sec after adaptation onset =2.84deg, $p = .24$).

In further support, the magnitude of the TAE following the unfiltered, contrast matched video was significantly greater than that following an unfiltered video that had higher contrast (mean difference 0-200sec after adaptation onset =5.10deg, $p = .001$), as well as the low-pass filtered video that also had higher contrast (mean difference 0-200sec after adaptation onset =4.39deg, $p = .006$).

The time course of the TAE following no video pre-exposure reached the second largest magnitude (8.57deg), only slightly below that following contrast matched pre-exposure. This is consistent with the findings of Chapter 6, and the proposal that natural scenes have relatively low contrast (Bex & Makous, 2002). Permutation testing revealed the magnitude of TAE growth following no-video pre-exposure was significantly greater than that following the unfiltered video, which had much higher contrast, and was used as a baseline measure in Chapter 5 (mean difference 0-200sec after adaptation onset =4.52deg, $p = .02$). In addition, the TAE following no-video pre-exposure was significantly greater than the low-pass filtered video that also had higher contrast (mean difference 0-200sec after adaptation onset =3.80deg, $p = .01$).

There were no further significant differences in the magnitude of TAE growth in any other condition comparisons. Additionally, there were no significant differences in TAE decay between any conditions.

7.4. MODELLING THE TIME COURSE OF THE TAE

The above results suggest that the time course of the TAE was altered by the contrast of the pre-exposure stimuli. Pre-exposure to increased contrast suppresses the magnitude of the TAE relative to that following no-video pre-exposure; the magnitude of the suppression increases with increasing pre-exposure contrast.

The suppression of the TAE following exposure to an environment with high contrast is consistent with a response gain control mechanism. Repeated pre-exposure to relatively high contrast stimuli, such as an unfiltered video, causes populations of visually responsive neurons to lower their firing rate (Schwartz & Simoncelli, 2001). This reduction in gain results in reduced responsiveness to subsequently presented adapting stimuli, which reduces their perceived contrast, and leads to suppression of the growth of the TAE (Harris & Calvert, 1989).

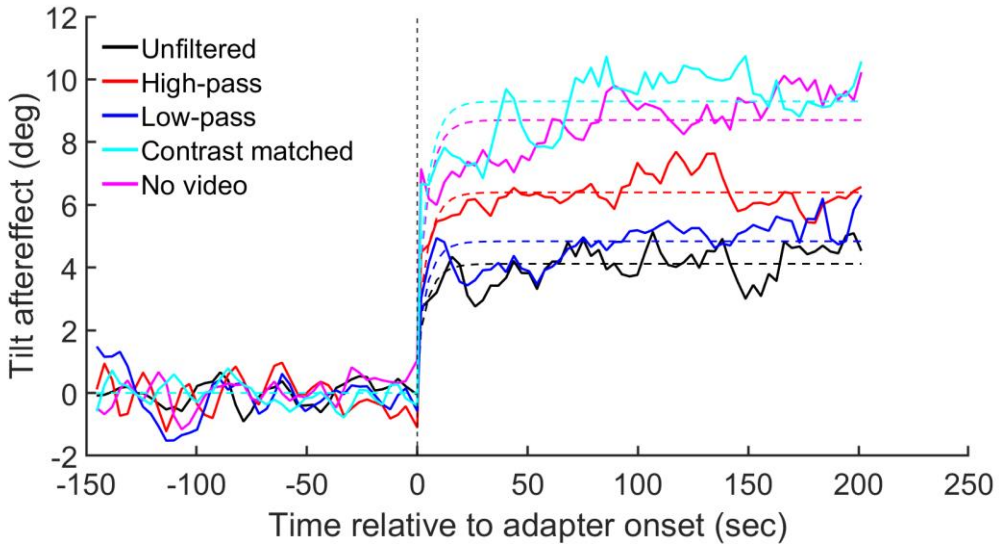


Figure 7.4. The time course of the TAE for each pre-exposure condition (solid lines) overlaid with the model's predicted time course for each condition (dashed lines).

Changes in response gain were modelled using a single-mechanism leaky integrator model with gain as a variable parameter. Each time course was fit simultaneously with a single mechanism that has a fixed time constant of 1.47 seconds and a variable gain parameter. Figure 7.4 shows the time course of the TAE for each pre-exposure condition (solid lines) overlaid with the best fitting model's predicted time courses (dashed lines). This model explained 93% of the variance in the data. Figure 7.4 shows that the gain of each mechanism in the model is linearly tuned to the contrast of the pre-exposure stimulus, with a liner model accounting for 79% of the

variance (see Equation 7.2). The time course of the TAE can therefore be successfully modelled as a function of pre-exposure contrast, with increasing contrast leading to a reduction in adaptation gain in a linear fashion.

$$g = -2x + 0.5 \qquad \text{Equation 7.2}$$

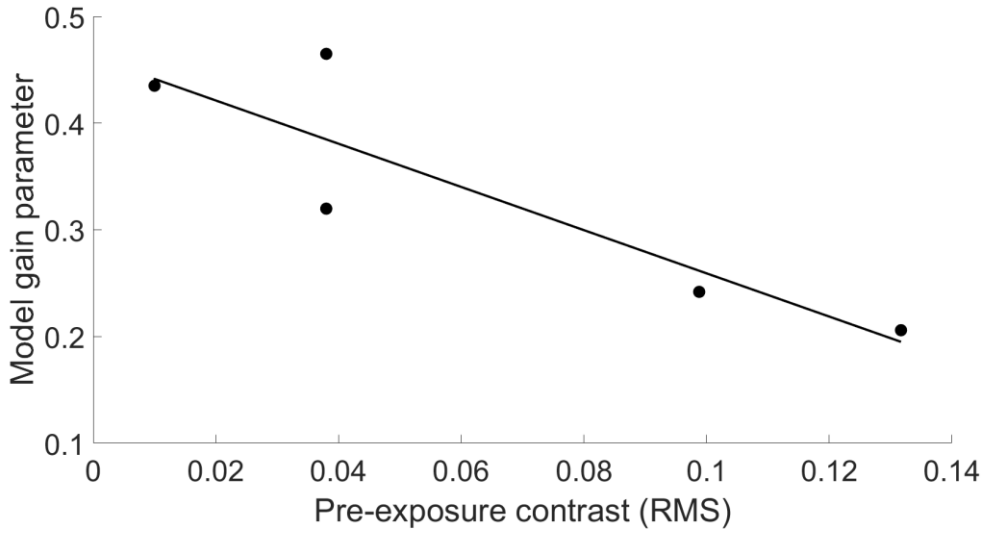


Figure 7.5. The model's gain parameter as a function of pre-exposure contrast. Solid line depicts the best fitting linear equation.

7.5. DISCUSSION

First, it is important to note the lack of any significant differences in TAE decay between any conditions in Chapters 5 and 7, even when there were significant differences in TAE growth. There are two possible explanations for this finding. Firstly, the effects of the pre-exposure stimulus may be passively decaying during the measurement of the TAE time course, and thus might be minimal at the time of TAE decay measurement. Secondly, there is evidence that the rate of growth and decay of the TAE may be

driven by independent mechanisms (Magnussen & Johnsen, 1986), therefore pre-exposure may affect the rate of adaptation growth but not decay. Due to the lack of any significant differences in the rate of decay between conditions, further discussion focuses on differences in TAE growth.

Chapter 5 showed that exposure to a temporally high-pass filtered video was sufficient to increase the rate and magnitude of the subsequently measured tilt aftereffect (TAE) relative to an unfiltered video containing naturalistic 1/f statistics. Subsequent experiments in Chapter 6 failed to replicate this effect with artificial stimuli that were modulated in contrast at fast (15hz) and slow (1Hz) temporal frequencies. The lack of replication with temporally modulated artificial stimuli suggested that the temporal statistics of the original videos did not drive the observed effect in Chapter 5. Rather, this effect was driven by differences in the videos other than temporal frequency statistics, such as altered RMS contrast.

Here, TAE growth following exposure to a high-pass filtered video was not significantly different in magnitude to that following exposure to an unfiltered video matched in average RMS contrast. These two videos differed markedly in their temporal statistics, but this was not sufficient to cause a significant difference in the TAE time course. In Chapter 5, the increase in the rate and magnitude of the TAE following high-pass pre-exposure was therefore likely a result of exposure to a low contrast stimulus rather than exposure to predominantly high temporal frequency content. Furthermore, the lack of significant difference in TAE magnitude between the low-pass and unfiltered conditions can be explained by their similarity in average RMS contrast rather than their similarity in temporal frequency content.

In addition, TAE magnitude following unfiltered and low-pass pre-exposure was significantly smaller than that following both contrast-matched and no-video pre-exposure. This is could be due to the low-pass and unfiltered videos having much higher contrast than the contrast matched and no-video exposure. The observed pattern of results in Chapter 5 is therefore consistent with the idea that exposure to altered RMS contrast affects the subsequent dynamics of adaptation.

Moreover, there was a large difference in TAE magnitude between the two different 'baseline' conditions used in Chapters 5 and 6. TAE magnitude following no-video pre-exposure was significantly larger than that following exposure to an unfiltered video. In terms of temporal frequency statistics, this result would be surprising as both exhibit 1/f statistics representative of natural scenes. However, the RMS contrast of a natural scene is much lower than the unfiltered video used in Chapter 5 (Bex & Makous, 2002), meaning the no-video pre-exposure is a more representative baseline.

Using the no-video baseline as a comparison, exposure to videos with higher RMS contrast than naturalistic viewing appears to suppress the magnitude of the TAE. The TAE following no-video pre-exposure is similar in magnitude to that following the high-pass and contrast matched videos which had low RMS contrast. Exposure to videos with higher RMS contrast, such as the low-pass and unfiltered videos, leads to significantly smaller magnitude TAEs. The higher the RMS contrast of previous exposure, the greater the suppression of the TAE.

The suppression of the TAE following exposure to an environment with high RMS contrast is consistent with a response gain control mechanism. This was modelled using a single mechanism leaky integrator with a fixed time constant but flexible gain for each time course. Repeated exposure to

increased contrast causes a reduction in gain in visually responsive neurons (Schwartz & Simoncelli, 2001). This was well captured by reducing the gain of the leaky integrator mechanism. The degree of gain reduction is roughly linearly related to the contrast of pre-exposure.

Although the effect observed in Chapter 5 can be explained in terms of exposure to videos with different RMS contrast, the way the visual system matches the dynamics of adaptation to the rate of change in the visual environment remains elusive. In addition, the dynamics of adaptation may also be altered by other properties of the environment. For example, recent physiological research in the auditory midbrain has shown that the rate of adaptation is increased when an environment with altered statistics is re-encountered (Robinson, Harper, & McAlpine, 2016). The dynamics of adaptation may therefore depend on the magnitude or reliability of the change in stimulus statistics, with repeated environments inducing more rapid adaptation as the system learns to adapt. What governs the time course of visual adaptation therefore remains an interesting question for future research.

8

General Discussion

8.1. SUMMARY OF RESULTS

This thesis addressed three main questions:

i) How can the time course of the tilt aftereffect be measured at a single trial temporal resolution?
ii) Is the time course of the tilt aftereffect fixed or can it be altered by the rate of change in the environment?
iii) What are the characteristics of the mechanism(s) that govern the growth and decay of the tilt aftereffect?

Experiments in Chapters 3-7 addressed these questions. An outline of the chapters and their main findings is provided below.

Chapter 3 explored the use of the method of reproduction as a single trial estimate of orientation bias. Consistent biases in orientation reproduction during baseline trials were seen for the full circular variable of orientation. Oblique orientations were consistently reproduced as being repelled away from the nearest cardinal, whilst the reproduction of cardinal and inter-cardinal orientations remained unbiased. In addition, the specific distribution of test stimulus orientations affected the pattern of biases shown. When the centre of the distribution was systematically manipulated, the point at which there was zero bias was shifted in the same direction as the manipulation. A Bayesian observer model corroborated the experimental data, revealing that biases seen in orientation reproduction are a combination of reference repulsion and central tendency effects.

Chapter 4 explored whether interrupting adaptation, by taking a measurement of the TAE, would alter its time course relative to continuous adaptation. Interrupting adaptation was found to suppress the magnitude of the TAE relative to the same duration of continuous adaptation. The magnitude of suppression was tuned to the frequency of interruption, with more frequent interruption leading to larger suppression. In addition, Chapter 4 explored the roles of passive recovery, higher level memory processes, and attention, on the suppression of the TAE during interrupted adaptation. Passive recovery alone was shown to explain the suppression of the TAE in interrupted adaptation, however this could not be captured by a simple leaky integrator model.

Chapter 5 explored whether the dynamics of adaptation could be altered by previous exposure to perturbed temporal statistics in naturalistic movie stimuli. Ten minutes of exposure to a video with predominantly high temporal frequency content was sufficient to significantly increase the rate and magnitude of the tilt aftereffect (TAE) relative to that following exposure to the unfiltered video. In contrast, ten minutes of exposure to the video with predominantly low temporal frequencies did not significantly alter the time course of the TAE. This empirical finding was successfully modelled by altering the contribution of two leaky integrator mechanisms tuned to distinct timescales.

Chapter 6 explored whether the observed effect in Chapter 5 could be replicated using artificial stimuli. Using the same methodology as Chapter 5, five minutes of pre-exposure to contrast modulated noise stimuli at 1Hz and 15Hz was insufficient to alter the time course of the TAE relative to a condition where there was no pre-exposure. A further attempt was made to replicate the observed effect in Chapter 5, measuring only the first few seconds of adaptation with a continuously present pre-exposure stimulus. Despite these methodological alterations, neither the 1Hz nor 15Hz pre-exposure condition resulted in changes in the time course of the TAE

reflective of those seen in Chapter 5. The observed effect in Chapter 5 could not be replicated with artificial stimuli.

Chapter 7 formally tested the hypothesis that the observed effect in Chapter 5 was driven by the contrast, rather than the temporal frequency content, of the original video stimuli. The time course of the TAE following a temporally unfiltered video that was matched in contrast to the high pass filtered video used in Chapter 5 led to the increased rate and magnitude of the TAE relative to the previously measured unfiltered condition. This confirmed that the effect observed in Chapter 5 could be explained by the contrast, rather than temporal frequency content of the pre-exposure stimulus. This empirical finding was consistent with a gain control mechanism.

These brief summaries outline the main findings in each chapter. The following sections consider each main experimental question in turn. Each question is addressed by considering the experimental findings of this thesis in the context of previous research. Unanswered questions and promising avenues for future research are considered for each experimental question.

8.2. WHAT IS AN APPROPRIATE METHODOLOGY TO MEASURE THE TIME COURSE OF THE TILT AFTEREFFECT WITH FINE TEMPORAL RESOLUTION?

To explore the dynamics of adaptation, a valid and fine-timescale measurement of its time course was needed. The traditional measurement of the time course of the TAE, using fixed discrete blocks of adaptation increasing in duration (see Magnussen & Johnsen (1986)), was unsuitable to measure changes in the dynamics of adaptation as the rate at which the time course is sampled is too coarse. Interrupted adaptation, using a fixed duration of adaptation with frequent interruptions for bias measurement,

was adopted instead. Bias was measured using the method of adjustment rather than a two-alternative forced choice method as the former allows estimation of bias from a single test stimulus presentation.

Two factors were considered when deciding whether interrupted adaptation with an adjustment response method was an appropriate methodology to measure the time course of the TAE with fine temporal resolution. The first was whether the adjustment response produced a valid representation of orientation perception and was explored following the observation of consistent biases in observers' baseline responses during initial pilot experiments. The second was whether interrupting adaptation altered the classic logarithmic nature of adaptation growth and decay (Magnussen & Johnsen, 1986).

Therefore, Chapter 3 explored the use of the method of adjustment as a single trial estimate of orientation bias. Observers reproduced the orientation of a tilted Gabor stimulus by drawing a line on the monitor with the mouse. Initial pilot experiments had revealed consistent biases in orientation reproduction during baseline trials for stimuli oriented ±18-22 from vertical. To further explore these biases observers' orientation reproduction responses were gathered for the full circular variable of orientation space in 5deg steps. A clear pattern of biases was observed with oblique orientations being consistently reproduced as being repelled away from the nearest cardinal, whilst the reproduction of cardinal and inter-cardinal orientations was relatively veridical. These effects are typical of reference repulsion effects, first demonstrated in the way the direction of a pattern of moving dots is repelled away from a reference line (Rauber & Treue, 1998).

The explanation of these effects was complicated by the finding that the specific distribution of test stimulus orientations affected the pattern of biases shown. When the centre of the test stimulus orientation distribution was systematically varied, the pattern of biases shifted in the same direction

as the manipulation. This result suggested that central tendency effects (Hollingworth, 1910) influence the biases in the reproduction of orientations seen pre-adaptation. A Bayesian observer model corroborated the experimental finding that biases seen in orientation reproduction are a combination of reference repulsion and central tendency effects. Importantly, these biases are due to implicit reference points and the distribution of test stimuli used – neither of which change during measurement of the time course of the TAE. Therefore, these biases are assumed to remain constant throughout measurement, justifying their correction on an observer-by-observer basis in the early stages of data analysis in subsequent experiments in this thesis.

From the data collected, it is impossible to disentangle an observer's response from their perception, making it difficult to identify if the adjustment response used provided a valid measure of orientation perception. It could be that observers' perception of the test stimuli was not biased, and instead the bias is introduced at level of the response rather than at the level of perception. If this is the case, then the adjustment response is not a valid measure of orientation perception. Alternatively, if the bias is introduced at the level of perception, then the adjustment response is a valid measure of perception. There is ongoing debate over the level at which these biases originate, with some suggesting they occur at the level of perception (Dakin, Apthorp, & Alais, 2010; Jazayeri & Movshon, 2007) and others suggesting they occur at the level of response (Zamboni, Ledgeway, McGraw, & Schluppeck, 2016). Regardless of the perceptual or cognitive locus of these biases, they are systematic and consistent for the orientations used in the experiments in this thesis and can therefore be easily corrected for during data analysis. The method of adjustment was therefore considered an appropriate response method to gather a valid estimate of the TAE on a single trial.

Chapter 4 explored whether interrupting adaptation altered the classic logarithmic nature of adaptation growth and decay (Magnussen & Johnsen,

1986). The first experiment in Chapter 4 aimed to identify how the frequency of interruption altered the time course of the TAE. Five minutes of adaptation was interrupted at one of four interruption frequencies (2.5, 4.3, 6.6, and 12.7 seconds) by the presentation of a test stimulus (interruption duration = 0.76s). This revealed that interrupting adaptation suppressed the magnitude of the TAE relative to the same duration of continuous adaptation, with more frequent interruption leading to larger suppression. It was concluded that adaptation was passively decaying during the interruption period, resulting in the suppressed magnitude of the TAE.

The first experiment in Chapter 4 therefore simplified the interruption period as time spent not adapting, during which passive decay could occur. However, the interruption period included the presentation of a test stimulus and the observer making a decision about its orientation and was followed by a response period where the observer had to hold the orientation of the test stimulus in memory and reproduce the orientation using the mouse. These elements were absent during continuous adaptation and could therefore have contributed to the suppression in magnitude of the TAE observed in the previous experiment.

Therefore, the second experiment in Chapter 4 aimed to disentangle the effects of passive recovery (Magnussen & Johnsen, 1986), higher level memory processes (Kang, Hong, Blake, & Woodman, 2011), and attention (Spivey & Spirn, 2000) during the response period, on the suppression of the TAE during interrupted adaptation. Observers completed three conditions: standard interruption was identical to that in the first experiment; video playback interruption removed the need for the observer to make a response; and mean luminance interruption involved no stimulus presentation or response period. The magnitude of suppression was comparable in all three conditions, suggesting that interrupting adaptation with the presentation of a test stimulus and the attentional cost of making

a response do not supress the growth of the TAE to a greater extent than interruption with mean luminance alone.

These experiments suggest that using interrupted adaptation to measure the time course of the TAE does alter the classic logarithmic nature of TAE growth. The final magnitude of adaptation is suppressed and adaptation asymptotes long before the saturation point of 1 hour reported by Greenlee & Magnussen (1986). However, the time course of adaptation measured in this way is consistent if the interruption period is kept constant and allows for the comparison of both adaptation rate and magnitude across conditions at fine temporal resolution. In conclusion, the use of interrupted adaptation and an adjustment response was therefore deemed an appropriate methodology to measure the time course of the TAE with fine temporal resolution.

8.2.1. Future directions

Chapter three identified reference repulsion effects in pre-adaptation orientation reproduction. The anisotropic distribution of orientations in the natural world (Coppola, Purves, Mccoy, & Purves, 1998; Dakin, Mareschal, & Bex, 2005; Dragoi, Turcu, & Sur, 2001; Hansen & Essock, 2004), is suggested to have resulted in parallel neural anisotropies (Li, 2003; Wang, Ding, & Yunokuchi, 2003; Xu, Collins, Khaytin, Kaas, & Casagrande, 2006), and could be the basis for the reference repulsion biases seen in orientation perception measured in Chapter 3.

Intuitively, one might think that these neural anisotropies and perceptual biases are fixed in nature due to being shaped over long timescales of evolution and development in an environment dominated by cardinal orientations. However, looking at the pattern of biases for individual observers in Chapter 3 revealed that one subject, who had hours of experience judging inter-cardinal orientations, showed reference repulsion to intercardinal as well as cardinal orientations. This suggests that these

biases may be flexibly altered in an experience-dependent manner. A plausible experiment would be to expose observers, through a virtual reality headset, to an altered distribution of orientations in natural viewing. A virtual environment could be created in which intercardinal, rather than cardinal, orientations dominate the orientation distribution. The timescales over which these 'new' dominant orientations are learned would be interesting to examine, as well as how the relative strength of biases at each set of orientations is altered with increasing exposure to the 'new' cardinals.

Chapter 4 showed that interrupting adaptation results in suppression of the magnitude of the TAE. This has implications for commonly used adapt-test-top-up designs where an initial period of adaptation is followed by periods of 'top-up' adaptation to prevent adaptation decay during test trials. The results of Chapter 4 suggest that a precise amount of top-up adaptation will be needed to counteract the passive decay during test trials. This is currently not standard in psychophysical studies, with most using arbitrary durations of initial adaptation and top-up adaptation. A more detailed understanding of the mechanisms underlying the suppressive effect is needed to produce a model that could be used to predict the appropriate amount of top-up adaptation given the initial adaptation duration and the duration of the interruptions. This would ensure consistent adapted states throughout individual experiments and allow better comparisons across studies through methodological consistency.

8.3. HOW IS THE TIME COURSE OF THE TILT AFTEREFFECT, SPECIFICALLY THE RATE OF ADAPTATION, AFFECTED BY PRIOR EXPOSURE TO ALTERED TEMPORAL ENVIRONMENTS?

Chapters 5, 6, and 7 aimed to identify whether the dynamics of adaptation could be altered by previous exposure to perturbed temporal statistics. In Chapter 5 observers were exposed to high-pass and low-pass temporally filtered videos followed by the measurement of the time course of the TAE

using interrupted adaptation. Temporally filtered videos represented deviations from naturalistic 1/f statistics. It was hypothesised that the dynamics of subsequent adaptation might match the rate of change in the pre-exposure video, with faster adaptation following high-pass pre-exposure and slower adaptation following low-pass pre-exposure. The time course of the TAE was compared to a baseline condition where observers were exposed to an unfiltered video with naturalistic 1/f statistics.

The results of Chapter 5 showed that ten minutes of exposure to a video with predominantly high temporal frequency content was sufficient to significantly increase the rate and magnitude of the tilt aftereffect (TAE) relative to that following exposure to the unfiltered video. However, ten minutes of exposure to the video with predominantly low temporal frequencies did not significantly alter the time course of the TAE relative to that following the unfiltered video. The asymmetrical findings suggested that exposure to a large deviation from natural temporal statistics was necessary to alter the dynamics of the TAE. High-pass temporal filtering created a video with a large deviation from natural statistics by attenuating the low temporal frequencies that are predominant in natural environments. In comparison, low-pass temporal filtering attenuated the high temporal frequencies that are less abundant in natural temporal statistics, resulting in a video with an arguably smaller deviation from natural temporal statistics. Therefore, the results of Chapter 5 led to the conclusion that prior exposure to a large deviation in the temporal statistics of the environment was sufficient to alter the time course of the TAE. However, the advantages and disadvantages of using naturalistic stimuli were considered to assess the validity of this finding.

The use of naturalistic pre-exposure stimuli in Chapter 5 was advantageous both theoretically and practically. Theoretically, there is a strong argument for studying visual processing in response to stimuli that the visual system has evolved to process (Felsen & Dan, 2005). Any successful model of visual processing must be able to predict neural responses to the natural

environment, therefore any processes identified using artificial stimuli must also hold in response to naturalistic stimuli. Moreover, adaptation specifically is a ubiquitous process that allows the visual system to efficiently encode our complex natural environment, therefore to fully understand adaptation we need to study it in response to complex naturalistic stimuli. Practically, the use of video stimuli led to the maintained attention of the observers throughout the pre-exposure period. Attention has a large effect on visual processing, and in specific on adaptation, therefore it was important to use stimuli that maintained observer attention.

However, the complex nature of the naturalistic stimuli used in Chapter 5 prevents the conclusion that the observed effects were solely due to the manipulation of the videos' temporal statistics. Confounding differences existed between the videos, such as complex spatial and temporal relationships, and the RMS contrast of each video. It was therefore necessary to replicate the observed effect with artificial stimuli that differed only in their temporal frequency statistics (See Rust & Movshon, 2005).

Chapter 6 aimed to replicate the observed effect in Chapter 5 using artificial stimuli. Pre-exposure stimuli consisted of a visual noise patch modulating in contrast at one of two temporal frequencies: 1hz and 15Hz. Chapter 6: Experiment 1 presented observers with these pre-exposure noise stimuli for 5 minutes followed by the measurement of the TAE time course. Neither the 1Hz nor 15Hz noise pre-exposure stimuli altered the time course of the TAE relative to a condition where there was no pre-exposure. This experiment therefore failed to replicate the observed effect in Chapter 5. The lack of replication was hypothesised to be due to the reduction in the pre-exposure duration, with 5 minutes of pre-exposure to artificial stimuli being insufficient to alter the properties of adapting mechanisms that have been shaped through evolution and development. Additionally, any differences in the rate of growth were hard to identify in both Chapter 5 and

Chapter 6: Experiment 1 due to the rapid growth of the TAE in the initial few seconds of adaptation.

Chapter 6: Experiment 2 addressed the methodological limitations of previous experiments. The first few seconds of adaptation were sampled more frequently using continuous periods of adaptation, and a bias-free estimate of the TAE was gathered using a two-alternative forced-choice response (Morgan, 2014). Moreover, the pre-exposure stimulus was presented both prior to and during the measurement of adaptation's time course, ensuring the temporal statistics of the environment were consistent. The results of this experiment showed traditional logarithmic growth of adaptation when there was no pre-exposure, but this growth was suppressed in both pre-exposure stimulus conditions. The suppression of the TAE was tuned to the contrast and temporal frequency of the pre-exposure stimuli, with higher contrast, higher temporal frequency stimuli producing larger suppression. This tuning, and the lack of suppression using the same pre-exposure stimuli in Chapter 6: Experiment 1, suggests that the simultaneous presentation of the pre-exposure and adapting stimuli resulted in the reduced effectiveness of the adapting stimulus. However, a matching experiment exploring whether the pre-exposure stimulus reduced the perceived contrast or temporal frequency of the adapter showed this not to be the case. The precise nature of this effect therefore remains to be characterised and would be an interesting avenue for further study.

Regardless, neither experiment in Chapter 6 using artificial stimuli that differed only in temporal frequency content replicated the observed effect in Chapter 5. From this, it was concluded that the increased rate and magnitude of the TAE following exposure to a temporally high-pass filtered video in Chapter 5 was not driven solely by the temporal frequency content of the video, but by some other property.

One such property was the average contrast of each pre-exposure video stimulus used in Chapter 5. Analysis of the video stimuli revealed differences in their time-averaged root mean square contrast - the highest contrast video was the unfiltered video, followed by the low-pass filtered video and the high-pass filtered video. The increased magnitude of the TAE following high-pass pre-exposure could therefore be explained by exposure to reduced contrast, and the lack of any differences between the low-pass and unfiltered pre-exposure conditions could be due to their relatively similar contrast rather than a similarity in temporal statistics.

Chapter 7 explored whether the observed effect in Chapter 5 was driven by the contrast, rather than the temporal frequency content, of the pre-exposure video stimuli. Observers were exposed to a video stimulus that was temporally unfiltered but matched in contrast to the high-pass filtered video that produced the original effect. If the time course of the TAE following this stimulus showed increased rate and magnitude, then it would be concluded that the original effect was driven by the contrast of the stimulus not the temporal frequency content. Alternatively, if the time course of the TAE was similar to that following unfiltered video pre-exposure, it would be concluded that another aspect, such as a complex spatiotemporal relationship drove the original effect in Chapter 5.

Exposure to a temporally unfiltered video that was matched in contrast to the high pass filtered video used in Chapter 5 (low contrast) led to the increased rate and magnitude of the TAE relative to exposure to a temporally unfiltered video with higher contrast. This confirmed that exposure to altered contrast, rather than temporal frequency content drove the original effect observed in Chapter 5.

In addition, the observers were exposed to a no-video pre-exposure condition to allow for easier comparison between the results of Chapters 5 and 6. The TAE following no video exposure was larger in magnitude than that following exposure to the temporally unfiltered video, despite both

possessing 1/f temporal statistics. This supports the proposal that pre-exposure contrast alters the time course of the TAE as natural stimuli have considerably lower contrast than the unfiltered video stimulus (Bex & Makous, 2002).

Considering these findings, the results of Chapter 5 can be simply interpreted in terms of exposure to altered contrast. The magnitude of the TAE is roughly linearly related to the contrast of previous visual experience, with exposure to stimuli with higher contrast than natural scenes leading to suppression of the TAE.

Worth noting, are the outlying results of Chapter 6. The artificial noise stimuli used in Chapter 6 had a time average contrast of 0.04RMS, the same as the high-pass filtered video used in Chapter 5. The magnitude of the TAE following ~80 seconds of adaptation (~6deg) is comparable between these conditions. However, the magnitude of the TAE following no-video pre-exposure in Chapter 6 (~7deg) is smaller than that following no-video exposure in Chapter 7 (~9deg). It is therefore possible that simple manipulation of contrast is insufficient to explain the observed meta-adaptation effects.

In conclusion, the time course of the TAE is not fixed, but rather is flexibly altered by exposure to environments with different contrast, rather than temporal statistics. Whether, and if so how, adapting mechanisms match the dynamics of adaptation to the rate of change in the environment remains unknown, but possible avenues for future exploration are discussed below.

8.3.1. Future directions

Chapters 5, 6, and 7 explored whether the rate of adaptation was determined by the rate of change in recent environmental statistics. Exposure to different rates of contrast change did not affect the rate of the

tilt aftereffect (TAE), suggesting that the dynamics of adaptation are not matched to the rate of recently encountered change. However, it is possible that the rate of change in properties other than contrast may determine the dynamics of the TAE.

One possibility is that the rate of change in a specific property sets the rate at which the system adapts to that property. For example, the rate of change in orientation statistics, rather than contrast, might alter the rate of TAE growth. It would be interesting to explore whether fast changes in orientation statistics might lead to increased rates of tilt adaptation and vice versa.

Alternatively, the rate of adaptation may not be set by the recently encountered rate of change in the adapting property, but rather whether the adapting stimulus has been previously encountered. Recent physiological research in the auditory system has shown faster rates of adaptation when an adapting stimulus is re-encountered following an interruption with a different stimulus (Robinson, Harper, & McAlpine, 2016). It is not known whether these physiologically measured changes in the rate of adaptation would be measurable psychophysically, but if so it could be a promising avenue of research. The adapting stimulus could be presented, and the time course of adaptation measured, followed by an interruption during which a different stimulus is presented, after which the initial adapting stimulus is again presented, and the time course of adaptation measured. Based on the above findings, we might expect to see a faster rate of adaptation in response to the second presentation of the adapting stimulus.

A final interesting avenue for future research is the use of distributions of adapting stimuli to measure the dynamics of adaptation in response to changes in adapter strength and duration. This could be explored using distributions of adapting stimuli similar to those used by Benucci, Saleem, & Carandini (2013). For example, a natural distribution of orientations

would be roughly uniform, with small peaks at the cardinals. Adapting distributions could be biased by creating a peak at a single orientation, the strength of which could be increased to create larger deviations from natural statistics. It can be assumed that adapting distributions with greater strength (more biased distributions) will induce larger magnitude adaptation - but how does this interact with the duration of adapting stimuli? For example, weak but long duration adapting stimuli may result in greater adaptation than large but brief adapting stimuli as the latter is treated as noise or an anomaly that the system would waste resources adapting to. Exploring different combinations of adapting stimulus strength and duration may elucidate whether and how the system distinguishes between noise and actual changes in the environment, and how this affects the dynamics of adaptation.

8.4. WHAT ARE THE CHARACTERISTICS OF THE MECHANISM(S) THAT GOVERN THE GROWTH AND DECAY OF THE TILT AFTEREFFECT?

Chapter 4 explored the effect of interrupting adaptation on the time course of the tilt aftereffect. Interrupting adaptation was found to suppress the magnitude of the TAE, with more frequent interruption leading to larger suppression. This suppressive effect was unaffected by observers' higher-level memory processes or attention during the interruption period.

The mechanism underlying this suppression was hypothesised to be due to passive recovery of adaptation during interruption periods, with more frequent interruption resulting in more opportunity for passive recovery. This hypothesis was explored using a leaky integrator model. This model had two free parameters, the time constant, controlling the rate of adaptation growth; and the gain, controlling the magnitude of adaptation growth. The output of the model was TAE magnitude, which grew during periods of adaptation, and decayed during periods of interruption.

Conditions in which there were more frequent interruptions therefore resulted in more opportunity for passive recovery.

The pattern of results was well-captured by a model whereby adaptation passively recovered during interruption periods, with greater suppression for higher frequency interruption. However, this simple model could not capture the magnitude of suppression in the measured data. This is hypothesised to be due to differences in the rate of adaptation growth and decay (Magnussen & Johnsen, 1986). If this is the case, adaptation might decay more rapidly than it grows, and a model with a fixed rate for adaptation growth and decay would be unable to capture the dynamics of interrupted adaptation. A model with a faster rate for adaptation decay might therefore capture the pattern seen in the empirical data. To explore this hypothesis, differences in the rate of adaptation growth and decay at very short timescales could be measured and a more flexible model created.

Developing a model that can account for the different rates of adaptation growth and decay would enable better prediction of how the time course of adaptation would be affected by different interruption frequencies and durations. Consequently, this would allow future researchers using adapt-test-top-up designs to better predict the ratio of initial adaptation to top-up adaptation needed to maintain an observer in a consistent adapted state throughout testing.

Chapters 5, 6, and 7 explored whether the time course of adaptation was altered by previous exposure to perturbed temporal statistics. The results of Chapter 5 initially suggested that the rate and magnitude of adaptation were increased following exposure to a temporally high-pass filtered video. This meta-adaptation effect was consistent with two mechanistic possibilities.

Firstly, there could be a single mechanism with flexible temporal tuning and gain properties. This mechanism could alter its time constant and gain to

produce a time course of adaptation matched to the prevailing temporal statistics. This mechanistic explanation could easily account for the increased rate and magnitude of adaptation seen in response to high temporal frequency exposure in Chapter 5. However, this explanation contradicts previous research suggesting at least two mechanisms tuned to distinct timescales underlie adaptation (Bao & Engel, 2012; Bao, Engel, Fast, & Mesik, 2013).

Alternatively, there could be multiple mechanisms, tuned to distinct timescales but with flexible gains. When exposed to an environment with abundant short-timescale changes, the gain of the short-timescale mechanism could be increased, resulting in the faster, larger magnitude adaptation observed in Chapter 5. The empirical data and leaky integrator modelling could not distinguish between these two possible mechanistic explanations of meta-adaptation effects. To do so, future research is needed to explore the number of adapting mechanisms tuned to different timescales.

Chapter 7 revealed that the time course of adaptation can be flexibly altered by previous exposure to altered contrast, rather than temporal frequency statistics. This suggests that the original finding in Chapter 5 is likely not due to the change in gain of a short-timescale mechanism or the activation of different mechanisms following high-pass pre-exposure. Instead, this finding is consistent with a response gain control mechanism whereby repeated exposure to increased contrast causes visually responsive neurons to reduce their responsiveness, which in turn reduces the magnitude of the TAE (Schwartz & Simoncelli, 2001).

This response gain mechanism was successfully modelled as the reduction in gain of a leaky integrator mechanism (Drew & Abbott, 2006) with a fixed time constant. Here, a response gain mechanism would prevent the saturation of responses of contrast-tuned neurons when confronted with an increase in the average environmental contrast.

8.4.1. Future directions

The precise mechanistic underpinnings of adaptation remain to be explored. Of particular interest is the number of adapting mechanisms operating across the wide range of adaptation timescales. As described in detail in section 1.1.3, at least two mechanisms operating at distinct timescales have been observed (Bao & Engel, 2012; Bao et al., 2013). The method of pitting long-timescale and short time-scale mechanisms against each other has thus far involved vastly distinct timescales of adaptation (4 hours vs 15 mins in Bao & Engel (2012) 5 mins vs 40sec in Bao et al. (2013)). Using such distinct timescales of adaptation makes the observation of distinct mechanisms more likely considering we do not know how broadly each mechanism is tuned. The same adapt-deadapt procedure could be used to explore the number of mechanisms driving adaptation and the width of their temporal tuning. Decreasing the difference between the adaptation and de-adaptation durations would eventually yield a scenario in which there was no characteristic spontaneous recovery - indicating both adapting durations are driven by the same mechanism. Sufficient comparisons could therefore indicate both the number of mechanisms within a specific range of timescales and the width of their temporal tuning.

8.5. CONCLUDING REMARKS

The first Chapters in this thesis explored the measurement of the time course of the tilt aftereffect (TAE). Initially, significant biases were identified in observers' reproduction of orientations around the clock, with oblique orientations consistently reproduced as being biased away from cardinal orientations in a manner consistent with reference repulsion effects seen in motion perception. Subsequently, the use of interrupted adaptation was shown to be a valid, temporally fine, measurement of the

time course of the TAE, despite a reduction in overall TAE magnitude due to passive decay during interruption periods.

The final three experimental chapters tested the hypothesis that the dynamics of the TAE could be altered by the rate of contrast change in recent visual experience. There was little support for this hypothesis, with the magnitude of adaptation being driven instead by the time-averaged root mean square contrast of the video stimuli. Exposure to increased contrast relative to the natural environment led to suppression of the TAE, consistent with a response gain mechanism. How the rate of adaptation is determined remains an interesting topic for future research. Possible avenues are the rate at which the adapting property changes over time, the frequency with which the adapting property is encountered, and the interaction between the magnitude of the deviation from natural statistics and the duration of the adapting stimulus.

The empirical findings of Chapters 3-7 highlight the need for a greater understanding of the mechanistic underpinnings of meta-adaptation and adaptation at different timescales. Exploring the number and characteristics of adapting mechanisms will help to guide further research on the functional benefits of visual adaptation.

References

Abbonizio, G., Langley, K., & Clifford, C. W. G. (2002). Contrast adaptation may enhance contrast discrimination. *Spatial Vision*, *16*(1), 45–58. https://doi.org/10.1163/15685680260433904

Abraham, W. C., & Bear, M. F. (1996). Metaplasticity: The plasticity of synaptic plasticity. *Trends in Neurosciences*, *19*(4), 126–130. https://doi.org/10.1016/S0166-2236(96)80018-X

Albrecht, D. G., & Hamilton, D. B. (1982). Striate cortex of monkey and cat: contrast response function. *Journal of Neurophysiology*, *48*(1), 217–237. https://doi.org/0022-3077/82/0000-0000$01.25

Andrews, D. (1967). Perception of contour orientation in the central fovea part 1: short lines. *Vision Research*, *7*, 975–997. https://doi.org/10.1016/0042-6989(67)90015-6

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

Appelle, S. (1972). Perception and discrimination as a function of stimulus orientation: The "oblique effect" in man and animals. *Psychological Bulletin*, *78*(4), 266–278. https://doi.org/10.1037/h0033117

Ashourian, P., & Loewenstein, Y. (2011). Bayesian inference underlies the contraction bias in delayed comparison tasks. *PLoS ONE*, *6*(5). https://doi.org/10.1371/journal.pone.0019551

Atick, J. J., & Redlich, A. N. (1992). What Does the Retina Know about Natural Scenes? *Neural Computation*, *4*(2), 196–210. https://doi.org/10.1162/neco.1992.4.2.196

Banks, M. S. (1988). Visual recalibration and the development of contrast and optical flow perception. In *Perceptual development in infancy.* (pp.

145–196). Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.

Bao, M., & Engel, S. A. (2012a). Distinct mechanism for long-term contrast adaptation. *Proceedings of the National Academy of Sciences*. https://doi.org/10.1073/pnas.1113503109

Bao, M., & Engel, S. A. (2012b). Distinct mechanism for long-term contrast adaptation. *Proceedings of the National Academy of Sciences*. https://doi.org/10.1073/pnas.1113503109

Bao, M., Engel, S., Fast, E., & Mesik, J. (2013). Distinct mechanisms control contrast adaptation over different timescales. *Journal of Vision*. https://doi.org/10.1167/12.9.698

Barlow, H. (1961). Possible principles underlying the transformations of sensory messages. *Sensory Communication*. https://doi.org/10.7551/mitpress/9780262518420.003.0013

Barlow, H. (1990). Conditions for versatile learning, Helmholtz's unconscious inference, and the task of perception. *Vision Research*, *30*(11), 1561–1571. https://doi.org/10.1016/0042-6989(90)90144-A

Barlow, H. B. (1990). A theory about the functional role and synaptic mechanism of visual after-effects. *Vision: Coding and Efficiency*, 363–375. https://doi.org/10.1017/CBO9780511626197.034

Barlow, H. B., & Földiák, P. (1989). Adaptation and decorrelation in the cortex. *The Computing Neuron*, (1988), 54–72.

Belmore, S. C., & Shevell, S. K. (2011). Very-long-term and short-term chromatic adaptation: Are their influences cumulative? *Vision Research*, *51*(3), 362–366. https://doi.org/10.1016/j.visres.2010.11.011

Benucci, A., Saleem, A. B., & Carandini, M. (2013). Adaptation maintains population homeostasis in primary visual cortex. *Nature Neuroscience*, *16*(6), 724–729. https://doi.org/10.1038/nn.3382

Bex, P. J., & Makous, W. (2002). Spatial frequency, phase, and the contrast of natural images. *Journal of the Optical Society of America A-Optics Image Science and Vision*, *19*(6), 1096–1106. https://doi.org/10.1364/JOSAA.19.001096

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. https://doi.org/10.1113/jphysiol.1969.sp008862

Blakemore, C., Muncey, J. P. J., & Ridley, R. M. (1973). Stimulus specificity in the human visual system. *Vision Research*, *13*(10), 1915–1931. https://doi.org/10.1016/0042-6989(73)90063-1

Blakemore, C., & Sutton, P. (1969). Size adaptation: A new aftereffect. *Science*, *166*(3902), 245–247. https://doi.org/10.1126/science.166.3902.245

Boehnke, S. E., Berg, D. J., Marino, R. A., Baldi, P. F., Itti, L., & Munoz, D. P. (2011). Visual adaptation and novelty responses in the superior colliculus. *European Journal of Neuroscience*, *34*(5), 766–779. https://doi.org/10.1111/j.1460-9568.2011.07805.x

Bonin, V., Mante, V., & Carandini, M. (2006). The Statistical Computation Underlying Contrast Gain Control. *Journal of Neuroscience*, *26*(23), 6346–6353. https://doi.org/10.1523/JNEUROSCI.0284-06.2006

Brady, N., & Field, D. J. (2000). Local contrast in natural images: Normalisation and coding efficiency. *Perception*, *29*(9), 1041–1055. https://doi.org/10.1068/p2996

Burbeck, C. a, & Regan, D. (1983). Independence of orientation and size in spatial discriminations. *J Opt Soc Am*, *73*(12), 1691–4. https://doi.org/10.1364/JOSA.73.001691

Burr, D., & Cicchini, G. M. (2014). Vision: Efficient Adaptive Coding. *Current Biology*, *24*(22), R1096–R1098. https://doi.org/10.1016/j.cub.2014.10.002

Butler, R. W., & Paolella, M. S. (2002). Calculating the density and distribution function for the singly and doubly noncentral F. *Statistics and Computing*, *12*(1), 9–16. https://doi.org/10.1023/A:1013160019893

Campbell, F. W., Cooper, G. F., & Enroth-Cugell, C. (1969). The spatial selectivity of the visual cells of the cat. *The Journal of Physiology*, *203*(1), 223–235. https://doi.org/10.1113/jphysiol.1969.sp008861

Campbell, F. W., & Robson, J. G. (1968). Application of fourier analysis to the visibility of gratings. *The Journal of Physiology*, *197*(3), 551–566. https://doi.org/10.1113/jphysiol.1968.sp008574

Carandini, M., Anthony Movshon, J., & Ferster, D. (1998). Pattern adaptation and cross-orientation interactions in the primary visual cortex. In *Neuropharmacology* (Vol. 37, pp. 501–511). https://doi.org/10.1016/S0028-3908(98)00069-0

Chaney, W., Liberman, A., & Whitney, D. (2016). Serial dependence in perception of biological motion. *Journal of Vision*, *16*(12), 274. https://doi.org/10.1167/16.15.28.doi

Chaudhuri, R., Knoblauch, K., Gariel, M. A., Kennedy, H., & Wang, X. J. (2015). A Large-Scale Circuit Mechanism for Hierarchical Dynamical Processing in the Primate Cortex. *Neuron*, *88*(2), 419–431. https://doi.org/10.1016/j.neuron.2015.09.008

Chopin, A., & Mamassian, P. (2012). Predictive Properties of Visual Adaptation. *Current Biology*, *22*(7), 622–626. https://doi.org/10.1016/j.cub.2012.02.021

Clatworthy, P. L., Chirimuuta, M., Lauritzen, J. S., & Tolhurst, D. J. (2003). Coding of the contrasts in natural images by populations of neurons in primary visual cortex (V1). *Vision Research*, *43*(18), 1983–2001. https://doi.org/10.1016/S0042-6989(03)00277-3

Clifford, C. W. G. (2002). Perceptual adaptation: Motion parallels orientation. *Trends in Cognitive Sciences*.

https://doi.org/10.1016/S1364-6613(00)01856-8

Clifford, C. W. G., 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. https://doi.org/10.1016/j.visres.2007.08.023

Clifford, C. W., Wenderoth, P., & Spehar, B. (2000). A functional angle on some after-effects in cortical vision. *Proceedings. Biological Sciences / The Royal Society*, *267*(1454), 1705–1710. https://doi.org/10.1098/rspb.2000.1198

Coltheart, M. (1971). Visual feature-analyzers and after-effects of tilt and curvature. *Psychological Review*, *78*(2), 114–121. https://doi.org/10.1037/h0030639

Coppola, D. M., Purves, H. R., Mccoy, A. N., & Purves, D. (1998). The distribution of oriented contours in the real world. *Neurobiology*, *95*, 4002–4006. https://doi.org/10.1073/pnas.95.7.4002

Crowder, N. A. (2005). Relationship Between Contrast Adaptation and Orientation Tuning in V1 and V2 of Cat Visual Cortex. *Journal of Neurophysiology*, *95*(1), 271–283. https://doi.org/10.1152/jn.00871.2005

Crowder, N. a, Hietanen, M. a, Price, N. S. C., Clifford, C. W. G., & Ibbotson, M. R. (2008). Dynamic contrast change produces rapid gain control in visual cortex. *The Journal of Physiology*, *586*(Pt 17), 4107–4119. https://doi.org/10.1113/jphysiol.2008.156273

Dakin, S. C., Apthorp, D., & Alais, D. (2010). Anisotropies in judging the direction of moving natural scenes. *Journal of Vision*, *10*(11), 5. https://doi.org/10.1167/10.11.5

Dakin, S. C., Mareschal, I., & Bex, P. J. (2005). An oblique effect for local motion: psychophysics and natural movie statistics. *Journal of Vision*, *5*(10), 878–887. https://doi.org/10.1167/5.10.9

Dan, Y., Atick, J. J., & Reid, R. C. (1996). Efficient coding of natural scenes in

the lateral geniculate nucleus: experimental test of a computational theory. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *16*(10), 3351–3362.

David, S. V. (2004). Natural Stimulus Statistics Alter the Receptive Field Structure of V1 Neurons. *Journal of Neuroscience*, *24*(31), 6991–7006. https://doi.org/10.1523/JNEUROSCI.1422-04.2004

Dayan, P., Sahani, M., & Deback, G. (2002). Adaptation and Unsupervised Learning. *Advances in Neural Information Processing Systems 15*, 237–244. Retrieved from http://papers.nips.cc/paper/2310-adaptation-and-unsupervised-learning.pdf%5Cnhttp://discovery.ucl.ac.uk/85050/

de Gardelle, V., Kouider, S., & Sackur, J. (2010). An oblique illusion modulated by visibility: non-monotonic sensory integration in orientation processing. *Journal of Vision*, *10*(10), 6. https://doi.org/10.1167/10.10.6

Dekel, R., & Sagi, D. (2015a). Tilt aftereffect due to adaptation to natural stimuli. *Vision Research*, *117*, 91–99. https://doi.org/10.1016/j.visres.2015.10.014

Dekel, R., & Sagi, D. (2015b). Tilt aftereffect due to adaptation to natural stimuli. *Vision Research*, *117*, 91–99. https://doi.org/10.1016/j.visres.2015.10.014

Delahunt, P. B., Webster, M. A., Ma, L., & Werner, J. S. (2004). Long-term renormalization of chromatic mechanisms following cataract surgery. In *Visual Neuroscience* (Vol. 21, pp. 301–307). https://doi.org/10.1017/S0952523804213025

DeValois, R. L., & DeValois, K. K. (1988). *Spatial Vision*. *Oxford Psychology Series*. https://doi.org/10.1093/acprof:oso/9780195066579.001.0001

Dong, D., & Atick, J. (1995). Statistics of natural time-varying images. *Network: Computation in Neural Systems*, *6*(3), 345–358.

https://doi.org/10.1088/0954-898X/6/3/003

Dong, D. W., & Atick, J. J. (1995). Temporal decorrelation: A theory of lagged and nonlagged responses in the lateral geniculate nucleus. *Network: Computation in Neural Systems*, *6*(2), 159–178. https://doi.org/10.1088/0954-898X_6_2_003

Dragoi, V., Rivadulla, C., & Sur, M. (2001). Foci of orientation plasticity in visual cortex. *Nature*, *411*(6833), 80–86. https://doi.org/10.1038/35075070

Dragoi, V., Sharma, J., & Sur, M. (2000). Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron*, *28*(1), 287–298. https://doi.org/10.1016/S0896-6273(00)00103-3

Dragoi, V., Sharma, J., & Sur, M. (2000a). Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron*, *28*(1), 287–298. https://doi.org/10.1016/S0896-6273(00)00103-3

Dragoi, V., Sharma, J., & Sur, M. (2000b). Adaptation-Induced Plasticity of Orientation Tuning in Adult Visual Cortex These phenomena raise the question of whether orientation selectivity in adult V1 can undergo. *Neuron*, *28*, 287–298. Retrieved from http://ac.els-cdn.com/S0896627300001033/1-s2.0-S0896627300001033-main.pdf?_tid=f81bb412-6bee-11e7-8822-00000aacb361&acdnat=1500406157_f32320fe15df092254fd96599ac38330

Dragoi, V., Turcu, C. M., & Sur, M. (2001). Stability of cortical responses and the statistics of natural scenes. *Neuron*, *32*(6), 1181–1192.

Drew, P. J., & Abbott, L. F. (2006). Models and properties of power-law adaptation in neural systems. *Journal of Neurophysiology*, *96*(2), 826–833. https://doi.org/10.1152/jn.00134.2006

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. *J Opt Soc Am*

A Opt Image Sci Vis, *24*(6), 1529–1537. Retrieved from http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=17491620

Eisner, A., & Enoch, J. M. (1982). Some effects of 1 week's monocular exposure to long-wavelength stimuli. *Perception & Psychophysics*, *31*(2), 169–174. https://doi.org/10.3758/BF03206217

Engel, S. a, & Furmanski, C. S. (2001). Selective adaptation to color contrast in human primary visual cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *21*(11), 3949–3954. https://doi.org/21/11/3949 [pii]

Engel, S., Zhang, X., & Wandell, B. (1997). Colour tuning in human visual cortex measured with functional magnetic resonance imaging. *Nature*, *388*(6637), 68–71. https://doi.org/10.1038/40398

Enroth-Cugell, C., & Robson, J. G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *The Journal of Physiology*, *187*(3), 517–552. https://doi.org/10.1113/jphysiol.1966.sp008107

Ethier, V., Zee, D. S., & Shadmehr, R. (2008). Spontaneous recovery of motor memory during saccade adaptation. *Journal of Neurophysiology*, *99*(5), 2577–2583. https://doi.org/10.1152/jn.00015.2008

Fairchild, M. D., & Lennie, P. (1992). Chromatic adaptation to natural and incandescent illuminants. *Vision Research*, *32*(11), 2077–2085. https://doi.org/10.1016/0042-6989(92)90069-U

Fairhall, a L., Lewen, G. D., Bialek, W., & de Ruyter Van Steveninck, R. R. (2001). Efficiency and ambiguity in an adaptive neural code. *Nature*, *412*(6849), 787–792. https://doi.org/10.1038/35090500

Felsen, G., & Dan, Y. (2005). A natural approach to studying vision. *Nature Neuroscience*, *8*(12), 1643–1646. https://doi.org/10.1038/nn1608

Festman, Y., & Ahissar, M. (2004). Attentional states and the degree of visual adaptation to gratings. *Neural Networks*, *17*(5–6), 849–860. https://doi.org/10.1016/j.neunet.2004.02.006

Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *NATURE NEUROSCIENCE*, *24*(21), 2569–2574. https://doi.org/10.1016/j.cub.2014.09.025

Foster, D. H. (2011). Color constancy. *Vision Research*. https://doi.org/10.1016/j.visres.2010.09.006

Foster, K. H., Gaska, J. P., Nagler, M., & Pollen, D. A. (1985). Spatial and temporal frequency selectivity of neurones in visual cortical areas V1 and V2 of the macaque monkey. *The Journal of Physiology*, *365*, 331–363. Retrieved from http://jp.physoc.org/content/365/1/331.short

Frazor, R. A., & Geisler, W. S. (2006). Local luminance and contrast in natural images. *Vision Research*, *46*(10), 1585–1598. https://doi.org/10.1016/j.visres.2005.06.038

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. https://doi.org/10.1016/j.neuron.2005.07.016

Gaudry, K. S., & Reinagel, P. (2007). Benefits of Contrast Normalization Demonstrated in Neurons and Model Cells. *Journal of Neuroscience*, *27*(30), 8071–8079. https://doi.org/10.1523/JNEUROSCI.1093-07.2007

Geisler, W. S. (2008). Visual Perception and the Statistical Properties of Natural Scenes. *Annual Review of Psychology*, *59*(1), 167–192. https://doi.org/10.1146/annurev.psych.58.110405.085632

Georgeson, M. A. (1985). The effect of spatial adaptation on perceived contrast. *Spatial Vision*, *1*(2), 103–112. https://doi.org/10.1163/156856885X00125

Gibson, B. Y. J. J., & Radner, M. (1937). Adaptation, after-effect and contrast in the perception of tilted lines. 1. Quantitative studes, 186–196.

Gilbert, C. D., & Wiesel, T. N. (1990). The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat.

Vision Research, *30*(11), 1689–1701. https://doi.org/10.1016/0042-6989(90)90153-C

Glasser, D. M., Tsui, J. M. G., Pack, C. C., & Tadin, D. (2011). PNAS Plus: Perceptual and neural consequences of rapid motion adaptation. *Proceedings of the National Academy of Sciences*, *108*(45), E1080–E1088. https://doi.org/10.1073/pnas.1101141108

Greene, M. R., & Oliva, A. (2010). High-level aftereffects to global scene properties. *Journal of Experimental Psychology. Human Perception and Performance*, *36*(6), 1430–1442. https://doi.org/10.1037/a0019058

Greenlee, M., Georgeson, M., Magnussen, S., & Harris, J. (1991). The time course of adaptation to spatial contrast. *Vision Research*, *31*(2), 223–236. https://doi.org/10.1016/0042-6989(91)90113-J

Greenlee, M. W., & Heitger, F. (1988). The functional role of contrast adaptation. *Vision Research*, *28*(7), 791–797. https://doi.org/10.1016/0042-6989(88)90026-0

Greenlee, M. W., & Magnussen, S. (1986). Saturation of the tilt aftereffect.

Greenlee, M. W., & Magnussen, S. (1987). Saturation of the tilt aftereffect. *Vision Research*, *27*(6), 1041–1043. https://doi.org/10.1016/0042-6989(87)90017-4

Grzywacz, N. M., & de Juan, J. (2003). Sensory adaptation as Kalman filtering: theory and illustration with contrast adaptation. *Network*, *14*, 465–482. https://doi.org/10.1088/0954-898X/14/3/305

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. https://doi.org/10.1016/j.cub.2014.09.027

Hammond, P., Pomfrett, C. J. D., & Ahmed, B. (1989). Neural motion after-effects in the cat's striate cortex: Orientation selectivity. *Vision Research*, *29*(12), 1671–1683. https://doi.org/10.1016/0042-6989(89)90149-1

Hansen, B. C., & Essock, E. A. (2004). A horizontal bias in human visual processing of orientation and its correspondence to the structural components of natural scenes. *Journal of Vision*, *4*(12), 5. https://doi.org/10.1167/4.12.5

Harris, J. P., & Calvert, J. E. (1989). Contrast, spatial frequency and test duration effects on the tilt aftereffect: Implications for underlying mechanisms. *Vision Research*, *29*(1), 129–135. https://doi.org/10.1016/0042-6989(89)90179-X

Harris, L. R., M.J., M., & Still, A. . (1981). Moving and the motion after-effect.

He, S., & MacLeod, D. I. (2001). Orientation-selective adaptation and tilt after-effect from invisible patterns. *Nature*, *411*(6836), 473–476. https://doi.org/Doi 10.1038/35078072

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

Hillis, J. M., & Brainard, D. H. (2007). Distinct mechanisms mediate visual detection and identification. *Current Biology*, *17*(19), 1714–1719. https://doi.org/DOI 10.1016/j.cub.2007.09.012

Hol, K., & Treue, S. (2001). Different populations of neurons contribute to the detection and discrimination of visual motion. *Vision Research*, *41*(6), 685–689. https://doi.org/10.1016/S0042-6989(00)00314-X

Hollingworth, H. L. (1910). The Central Tendency of Judgment. *The Journal of Philosophy, Psychology and Scientific Methods*, *7*(17), 461. https://doi.org/10.2307/2012819

Honey, C. J., Thesen, T., Donner, T. H., Silbert, L. J., Carlson, C. E., Devinsky, O., … Hasson, U. (2012). Slow Cortical Dynamics and the Accumulation of Information over Long Timescales. *Neuron*, *76*(2), 423–434. https://doi.org/10.1016/j.neuron.2012.08.011

Hubel, D., & Wiesel, T. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of*

Physiology, *160*(1), 106–154. https://doi.org/10.1523/JNEUROSCI.1991-09.2009

Huk, A. C., Ress, D., & Heeger, D. J. (2001). Neuronal basis of the motion aftereffect reconsidered. *Neuron*, *32*(1), 161–172. https://doi.org/10.1016/S0896-6273(01)00452-4

Hussain Ismail, A. M., Solomon, J. A., Hansard, M., & Mareschal, I. (2016). A tilt after-effect for images of buildings: evidence of selectivity for the orientation of everyday scenes. *Royal Society Open Science*, *3*(11), 160551. https://doi.org/10.1098/rsos.160551

Ichikawa, M., & Masakura, Y. (2006). Manual control of the visual stimulus reduces the flash-lag effect. *Vision Research*, *46*(14), 2192–2203. https://doi.org/10.1016/j.visres.2005.12.021

Jastrow, J. (1893). On the judgement of angles and positions of lines. *American Journal of Psychology*, *5*(2), 214–248.

Jazayeri, M., & Movshon, J. A. (2007). A new perceptual illusion reveals mechanisms of sensory decoding. *Nature*, *446*(7138), 912–915. https://doi.org/10.1038/nature05739

Jazayeri, M., & Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nature Neuroscience*, *13*(8), 1020–1026. https://doi.org/10.1038/nn.2590

Jin, D. Z., Dragoi, V., Sur, M., & Seung, H. S. (2005). Tilt Aftereffect and Adaptation-Induced Changes in Orientation Tuning in Visual Cortex Tilt Aftereffect and Adaptation-Induced Changes in Orientation Tuning in Visual Cortex. *Journal of Neurophysiology*, *94*(6), 4038–4050. https://doi.org/10.1152/jn.00571.2004

Jordan, H., Fallah, M., & Stoner, G. R. (2006). Adaptation of gender derived from biological motion. *Nature Neuroscience*, *9*(6), 738–739. https://doi.org/10.1038/nn1710

Kang, M. S., Hong, S. W., Blake, R., & Woodman, G. F. (2011). Visual working memory contaminates perception. *Psychonomic Bulletin and Review*,

18(5), 860–869. https://doi.org/10.3758/s13423-011-0126-5

Knapen, T., Rolfs, M., Wexler, M., & Cavanagh, P. (2011). The reference frame of the tilt aftereffect. *Journal of Vision*, *10*(1), 8–8. https://doi.org/10.1167/10.1.8

Köhler, W., & Wallach, H. (1944). Figural After-Effects. An Investigation of Visual Processes. *Source: Proceedings of the American Philosophical Society*, *88*(4), 269–357. https://doi.org/10.2307/985451

Kohn, a. (2007). Visual Adaptation: Physiology, Mechanisms, and Functional Benefits. *Journal of Neurophysiology*, *97*(5), 3155–3164. https://doi.org/10.1152/jn.00086.2007

Kohn, A., & Movshon, J. A. (2003). Neuronal adaptation to visual motion in area MT of the macaque. *Neuron*, *39*(4), 681–691. https://doi.org/10.1016/S0896-6273(03)00438-0

Kompaniez, E., Abbey, C. K., Boone, J. M., & Webster, M. A. (2013). Adaptation Aftereffects in the Perception of Radiological Images. *PLoS ONE*, *8*(10). https://doi.org/10.1371/journal.pone.0076175

Kording, K. P., Tenenbaum, J. B., & Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience*, *10*(6), 779–786. https://doi.org/10.1038/nn1901

Krauskopf, J., & Karl, G. (1992). Color discrimination and adaptation. *Vision Research*, *32*(11), 2165–2175. https://doi.org/10.1016/0042-6989(92)90077-V

Kristjánsson, Á. (2012). Dynamic coding of temporal luminance variation. *Journal of the Optical Society of America A*, *29*(6), 1180. https://doi.org/10.1364/JOSAA.29.001180

Kwon, M., Legge, G. E., Fang, F., Cheong, A. M. Y., & He, S. (2009). Adaptive changes in visual cortex following prolonged contrast reduction, *9*(2009), 1–16. https://doi.org/10.1167/9.2.20.Introduction

Laughlin, S. (1981). A simple coding procedure enhances a neuron???s information capacity. *Zeitschrift Fur Naturforschung - Section C Journal of Biosciences*. https://doi.org/10.1515/znc-1981-9-1040

Lee, D. N. (1980). The optic flow field: the foundation of vision. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *290*(1038), 169–179. https://doi.org/10.1098/rstb.1980.0089

Lennie, P., Krauskopf, J., & Sclar, G. (1990). Chromatic mechanisms in striate cortex of macaque. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *10*(2), 649–669. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/2303866

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. https://doi.org/10.1038/82947

Leopold, D. a, Rhodes, G., Müller, K.-M., & Jeffery, L. (2005). The dynamics of visual adaptation to faces. *Proceedings of The Royal Society B*, *272*(1566), 897–904. https://doi.org/10.1098/rspb.2004.3022

Levinson, E., & Sekuler, R. (1976). Adaptation alters perceived direction of motion. *Vision Research*, *16*(7). https://doi.org/10.1016/0042-6989(76)90189-9

Leys, C., Ley, C., Klein, O., Bernard, P., & Licata, L. (2013). Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology*, *49*(4), 764–766. https://doi.org/10.1016/j.jesp.2013.03.013

Li, B. (2003). Oblique Effect: A Neural Basis in the Visual Cortex. *Journal of Neurophysiology*, *90*(1), 204–217. https://doi.org/10.1152/jn.00954.2002

Li, B., Peterson, M. R., & Freeman, R. D. (2003). Oblique Effect : A Neural Basis in the Visual Cortex Oblique Effect : A Neural Basis in the Visual

Cortex. *Journal of Neurophysiology*, *90*(February 2003), 204–217. https://doi.org/10.1152/jn.00954.2002

Li, F. F., VanRullen, R., Koch, C., & Perona, P. (2002). Rapid natural scene categorization in the near absence of attention. *Proceedings of the National Academy of Sciences*, *99*(14), 9596–9601. https://doi.org/10.1073/pnas.092277599

Liberman, A., Fischer, J., & Whitney, D. (2014). Serial dependence in the perception of faces. *Current Biology*, *24*(21), 2569–2574. https://doi.org/10.1016/j.cub.2014.09.025

Liberman, A., & Whitney, D. (2015). The serial dependence of perceived emotional expression. *Journal of Vision*, *15*(12), 929. Retrieved from http://dx.doi.org/10.1167/15.12.929

Liberman, A., Zhang, K., & Whitney, D. (2016). Serial dependence promotes object stability during occlusion. *Journal of Vision*, *16*(15), 1–10. https://doi.org/10.1167/16.15.16

Liu, Y., Bovik, A. C., & Cormack, L. K. (2008). Disparity statistics in natural scenes. *Journal of Vision*, *8*(11), 19.1-14. https://doi.org/10.1167/8.11.19

Livingstone, M. S., & Hubel, D. H. (1984). Anatomy and physiology of a color system in the primate visual cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *4*(1), 309–356. https://doi.org/10.1523/JNEUROSCI.04-01-00309.1984

Lochmann, T., Ernst, U. A., & Deneve, S. (2012). Perceptual Inference Predicts Contextual Modulations of Sensory Responses. *Journal of Neuroscience*, *32*(12), 4179–4195. https://doi.org/10.1523/JNEUROSCI.0817-11.2012

Lundstrom, B. N., Higgs, M. H., Spain, W. J., & Fairhall, A. L. (2008). Fractional differentiation by neocortical pyramidal neurons. *Nature Neuroscience*, *11*(11), 1335–1342. https://doi.org/10.1038/nn.2212

Machulla, T. K., Di Luca, M., Froehlich, E., & Ernst, M. O. (2012).

Multisensory simultaneity recalibration: Storage of the aftereffect in the absence of counterevidence. *Experimental Brain Research*, *217*(1), 89–97. https://doi.org/10.1007/s00221-011-2976-5

Magnussen, S., & Johnsen, T. (1986). Temporal aspects of spatial adaptation. A study of the tilt aftereffect. *Vision Research*, *26*(4), 661–672. https://doi.org/10.1016/0042-6989(86)90014-3

Maloney, L. T., Martello, M. F. D., Sahm, C., & Spillmann, L. (2005). Past trials influence perception of ambiguous motion quartets through pattern completion. *Proceedings of the National Academy of Sciences*, *102*(8), 3164–3169. https://doi.org/10.1073/pnas.0407157102

Mante, V., Frazor, R. A., Bonin, V., Geisler, W. S., & Carandini, M. (2005). Independence of luminance and contrast in natural scenes and in the early visual system. *Nature Neuroscience*, *8*(12), 1690–1697. https://doi.org/10.1038/nn1556

Martinez-Conde, S., Macknik, S. L., & Hubel, D. H. (2004). The role of fixational eye movements in visual perception. *Nat Rev Neurosci*, *5*(3), 229–240. https://doi.org/10.1038/nrn1348

Mathôt, S., & Theeuwes, J. (2013). A reinvestigation of the reference frame of the tilt-adaptation aftereffect. *Scientific Reports*, *3*. https://doi.org/10.1038/srep01152

McDermott, K. C., Malkoc, G., Mulligan, J. B., & Webster, M. a. (2010). Adaptation and visual salience. *Journal of Vision*, *10*(13), 17. https://doi.org/10.1167/10.7.405

Mei, G., Dong, X., & Bao, M. (2017). The timescale of adaptation at early and mid-level stages of visual processing. *Journal of Vision*, *17*(1), 1. https://doi.org/10.1167/17.1.1

Mesik, J., Bao, M., & Engel, S. A. (2013). Spontaneous recovery of motion and face aftereffects. *Vision Research*, *89*, 72–78. https://doi.org/10.1016/j.visres.2013.07.004

Miller, J. (1991). Short Report: Reaction Time Analysis with Outlier

Exclusion: Bias Varies with Sample Size. *The Quarterly Journal of Experimental Psychology Section A*, *43*(4), 907–912. https://doi.org/10.1080/14640749108400962

Morant, R. B., & Harris, J. R. (1965). Two different after-effects of exposure to viual tilts. *The American Journal of Psychology*, *78*, 218–226. https://doi.org/10.2307/1420493

Morgan, M. (2014). A bias-free measure of retinotopic tilt adaptation. *Journal of Vision*, *14*(2014), 1–9. https://doi.org/10.1167/14.1.7.doi

Morgan, M. J., Melmoth, D., & Solomon, J. a. (2013). Linking hypotheses underlying Class A and Class B methods. *Visual Neuroscience*, *30*(5–6), 197–206. https://doi.org/10.1017/S095252381300045X

Motoyoshi, I., Nishida, S., Sharan, L., & Adelson, E. H. (2007). Image statistics and the perception of surface qualities. *Nature*, *447*(7141), 206–209. https://doi.org/10.1038/nature05724

Movshon, J. A., & Lennie, P. (1979). Pattern-selective adaptation in visual cortical neurones. *Nature*. https://doi.org/10.1038/278850a0

Movshon, J. A., Thompson, I. D., & Tolhurst, D. J. (1978). Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat's visual cortex. *The Journal of Physiology*, *283*(1), 101–120. https://doi.org/10.1113/jphysiol.1978.sp012490

Müller, J. R., Metha, a B., Krauskopf, J., & Lennie, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science (New York, N.Y.)*, *285*(5432), 1405–1408. https://doi.org/10.1126/science.285.5432.1405

Neitz, J., Carroll, J., Yamauchi, Y., Neitz, M., & Williams, D. R. (2002). Color perception is mediated by a plastic neural mechanism that is adjustable in adults. *Neuron*, *35*(4), 783–792. https://doi.org/10.1016/S0896-6273(02)00818-8

Nelson, S. B. (1991). Temporal interactions in the cat visual system. I. Orientation-selective suppression in the visual cortex. *The Journal of*

Neuroscience : The Official Journal of the Society for Neuroscience, *11*(2), 344–356.

Ohzawa, I., Sclar, G., & Freeman, R. D. (1982). Contrast gain control in the cat visual cortex. *Nature*, *298*(5871), 266–268. https://doi.org/10.1038/298266a0

Olkkonen, M., & Allred, S. R. (2014). Short-term memory affects color perception in context. *PLoS ONE*, *9*(1). https://doi.org/10.1371/journal.pone.0086488

Orban, G. A., Vandenbussche, E., & Vogels, R. (1984). Human orientation discrimination tested with long stimuli. *Vision Research*, *24*(2), 121–128. https://doi.org/10.1016/0042-6989(84)90097-X

Pantle, A., & Sekuler, R. (1968). Size-Detecting Mechanisms in Human Vision. *Science*, *162*(3858), 1146–1148. https://doi.org/10.1126/science.162.3858.1146-a

Paolella, M. S. (2007). *Intermediate Probability: A Computational Approach*. *Intermediate Probability: A Computational Approach*. https://doi.org/10.1002/9780470035061

Paradiso, M. A., Shimojo, S., & Nakayama, K. (1989). Subjective contours, tilt aftereffects, and visual cortical organization. *Vision Research*, *29*(9), 1205–1213. https://doi.org/10.1016/0042-6989(89)90066-7

Parraga, C. A., Troscianko, T., & Tolhurst, D. J. (2000). The human visual system is optimised for processing the spatial information in natural visual images. *Current Biology*, *10*(1), 35–38. https://doi.org/10.1016/S0960-9822(99)00262-6

Pavan, A., Hocketstaller, J., Contillo, A., & Greenlee, M. W. (2016). Tilt aftereffect following adaptation to translational Glass patterns. *Scientific Reports*, *6*, 23567. https://doi.org/10.1038/srep23567

Pavan, A., Marotti, R. B., & Campana, G. (2012). The temporal course of recovery from brief (sub-second) adaptations to spatial contrast. *Vision Research*, *62*, 116–124.

https://doi.org/10.1016/j.visres.2012.04.001

Peirce, J. W. (2007). PsychoPy-Psychophysics software in Python. *Journal of Neuroscience Methods*, *162*, 8–13. https://doi.org/10.1016/j.jneumeth.2006.11.017

Peirce, J. W. (2009). Generating Stimuli for Neuroscience Using PsychoPy. *Frontiers in Neuroinformatics*, *2*, 10. https://doi.org/10.3389/neuro.11.010.2008

Phinney, R. E., Bowd, C., & Patterson, R. (1997). Direction-selective coding of stereoscopic (cyclopean) motion. *Vision Research*, *37*(7), 865–869. https://doi.org/10.1016/S0042-6989(96)00244-1

Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews. Neuroscience*, *1*(2), 125–132. https://doi.org/10.1038/35039062

Pozzorini, C., Naud, R., Mensi, S., & Gerstner, W. (2013). Temporal whitening by power-law adaptation in neocortical neurons. *Nature Neuroscience*, *16*(7), 942–948. https://doi.org/10.1038/nn.3431

Prenger, R., Wu, M. C. K., David, S. V., & Gallant, J. L. (2004). Nonlinear V1 responses to natural scenes revealed by neural network analysis. *Neural Networks*, *17*(5–6), 663–679. https://doi.org/10.1016/j.neunet.2004.03.008

Radonjić, A., Cottaris, N. P., & Brainard, D. H. (2015). Color constancy in a naturalistic goal-directed task. *Journal of Vision*, *15*(13), 3.1-3.21. https://doi.org/10.1167/15.13.3.doi

Ranganath, C., & Rainer, G. (2003). Neural mechanisms for detecting and remembering novel events. *Nat Rev Neurosci*, *4*(3), 193–202. https://doi.org/10.1038/nrn1052

Rauber, H. J., & Treue, S. (1998). Reference repulsion when judging the direction of visual motion. *Perception*, *27*(4), 393–402. https://doi.org/10.1068/p270393

Regan, D., & Beverley, K. I. (1983). Spatial-frequency discrimination and detection: comparison of postadaptation thresholds. *Journal of the Optical Society of America*, *73*(12), 1684. https://doi.org/10.1364/JOSA.73.001684

Regan, D., & Beverley, K. I. (1985). Postadaptation orientation discrimination. *Journal of the Optical Society of America. A*, *2*(2), 147–155. https://doi.org/10.1364/JOSAA.2.000147

Reinagel, P., & Zador, A. M. (1999). Natural scene statistics at the centre of gaze. *Network: Computation in Neural Systems*, *10*(4), 341–350. https://doi.org/10.1088/0954-898X_10_4_304

Rieke, F., & Rudd, M. E. (2009a). The Challenges Natural Images Pose for Visual Adaptation. *Neuron*, *64*(5), 605–616. https://doi.org/10.1016/j.neuron.2009.11.028

Rieke, F., & Rudd, M. E. (2009b). The Challenges Natural Images Pose for Visual Adaptation. *Neuron*. https://doi.org/10.1016/j.neuron.2009.11.028

Ringach, D. L., Hawken, M. J., & Shapley, R. (2002). Receptive field structure of neurons in monkey primary visual cortex revealed by stimulation with natural image sequences. *Journal of Vision*, *2*(1), 2. https://doi.org/10.1167/2.1.2

Rinner, O., & Gegenfurtner, K. R. (2000). Time course of chromatic adaptation. *Vision Research*, *40*, 1–4. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10837828

Roach, N. W., Heron, J., Whitaker, D., & McGraw, P. V. (2010). Asynchrony adaptation reveals neural population code for audio-visual timing. *Proceedings. Biological Sciences / The Royal Society*, *278*(1710), 1314–1322. https://doi.org/10.1098/rspb.2010.1737

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.

https://doi.org/10.1152/jn.00548.2009

Robinson, B. L., Harper, N. S., & McAlpine, D. (2016). Meta-adaptation in the auditory midbrain under cortical influence. *Nature Communications*, *7*, 13442. https://doi.org/10.1038/ncomms13442

Rodieck, R. W. (1983). Raster-based colour stimulators. In J. D. Mollon & L. T. Sharpe (Eds.), *Colour Vision* (1st ed., pp. 133–144). London: Academic Press.

Roseboom, W., Linares, D., & Nishida, S. (2015). Sensory adaptation for timing perception. *Proceedings of the Royal Society. B, Biological Sciences*, *21*(12), 2833–2841. https://doi.org/10.1098/rspb.2014.2833

Rousselet, G. A., Fabre-Thorpe, M., & Thorpe, S. J. (2002). Parallel processing in high-level categorization of natural images. *Nature Neuroscience*, *5*(7), 629–630. https://doi.org/10.1038/nn866

Ruderman, D. L., & Bialek, W. (1994). Statistics of natural images: Scaling in the woods. *Physical Review Letters*, *73*(6), 814–817. https://doi.org/10.1103/PhysRevLett.73.814

Rust, N. C., & Movshon, J. A. (2005). In praise of artifice. *Nature Neuroscience*, *8*(12), 1647–1650. https://doi.org/10.1038/nn1606

Sachs, M. B., Nachmias, J., & Robson, J. G. (1971). Spatial-frequency channels in human vision. *Journal of the Optical Society of America*, *61*(9), 1176–1186. https://doi.org/10.1364/JOSA.61.001176

Saul, A. B., & Cynader, M. S. (1989). Adaptation in single units in visual cortex: The tuning of aftereffects in the temporal domain. *Visual Neuroscience*, *2*(6), 609–620. https://doi.org/10.1017/S0952523800003539

Schwartz, O., Hsu, A., & Dayan, P. (2007). Space and time in visual context. *Nature Reviews Neuroscience*, *8*(7), 522–535. https://doi.org/10.1038/nrn2155

Schwartz, O., & Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nature Neuroscience*, *4*(8), 819–825. https://doi.org/10.1038/90526

Sekuler, R., & Littlejohn, J. (1974). Tilt aftereffect following very brief exposures. *Vision Research*. https://doi.org/10.1016/0042-6989(74)90133-3

Seriès, P., Stocker, A. A., & Simoncelli, E. P. (2009). Is the homunculus "aware" of sensory adaptation? *Neural Computation*, *21*(12), 3271–3304. https://doi.org/10.1162/neco.2009.09-08-869

Shadmehr, R., Smith, M. a, & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, *33*, 89–108. https://doi.org/10.1146/annurev-neuro-060909-153135

Shapley, R., & Enroth-cugell, C. (1984). Visual adaptation and retinal gain control. *Progress in Retinal Research*, *3*, 263–346.

Shapley, R. M., & Victor, J. D. (1978). The effect of contrast on the transfer properties of cat retinal ganglion cells. *The Journal of Physiology*, *285*(1), 275–298. https://doi.org/10.1113/jphysiol.1978.sp012571

Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation.

Smirnakis, S. M., Berry, M. J., Warland, D. K., Bialek, W., & Meister, M. (1997). Adaptation of retinal processing to image contrast and spatial scale. *Nature*, *386*(6620), 69–73. https://doi.org/10.1038/386069a0

Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biology*, *4*(6), 1035–1043. https://doi.org/10.1063/1.2184639

Smith, R. A. (1971). Studies of temporal frequency adaptation in visual contrast sensitivity. *The Journal of Physiology*, *216*(3), 531–552. https://doi.org/10.1113/jphysiol.1971.sp009539

Smith, S. L. (1962). Angular estimation. *Journal of Applied Psychology*, *46*(4), 240–246. https://doi.org/10.1037/h0044445

Solomon, S. G., & Kohn, A. (2014). Moving sensory adaptation beyond suppressive effects in single neurons. *Current Biology*. https://doi.org/10.1016/j.cub.2014.09.001

Spigel, I. M. (1960). The Effects of Differential Post-Exposure Illumination on the Decay of a Movement After-Effect. *The Journal of Psychology*, *50*(2), 209–210. https://doi.org/10.1080/00223980.1960.9916438

Spillman, L., & Werner, J. S. (1990). Visual Perception. The Neurophysiological Foundations. *Neural Networks*.

Spivey, M. J., & Spirn, M. J. (2000). Selective visual attention modulates the direct tilt aftereffect. *Perception and Psychophysics*, *62*(8), 1525–1533. https://doi.org/10.3758/BF03212153

Srinivasan, M. V., Laughlin, S. B., & Dubs, A. (1982). Predictive coding a fresh view of inhibition in the retina.pdf. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, *216*(1205), 427–459. https://doi.org/10.1098/rspb.1982.0085

Steinman, R. M., & Collewijn, H. (1980). Binocular retinal image motion during active head rotation. *Vision Research*, *20*(5), 415–429. https://doi.org/10.1016/0042-6989(80)90032-2

Stocker, A. a, & Simoncelli, E. P. (2006). Sensory Adaptation within a Bayesian Framework for Perception. *Advances in Neural Information Processing Systems*, *18*, 1291–1298.

Suzuki, S. (2001). Attention-dependent brief adaptation to contour orientation: A high-level aftereffect for convexity? *Vision Research*, *41*(28), 3883–3902. https://doi.org/10.1016/S0042-6989(01)00249-8

Switkes, E., Mayer, M. J., & Sloan, J. A. (1978). Spatial frequency analysis of the visual environment: Anisotropy and the carpentered environment hypothesis. *Vision Research*, *18*(10), 1393–1399.

https://doi.org/10.1016/0042-6989(78)90232-8

Tadmor, Y., & Tolhurst, D. J. (2000). Calculating the contrasts that retinal ganglion cells and LGN neurones encounter in natural scenes. *Vision Research*, *40*(22), 3145–3157. https://doi.org/10.1016/S0042-6989(00)00166-8

Teich, A. F., Qian, N., Abbott, L., Varela, J., Sen, K., Nelson, S., … Dan, Y. (2003). Learning and Adaptation in a Recurrent Model of V1 Orientation Selectivity. *Journal of Neurophysiology*, *89*(4), 220–224. https://doi.org/10.1152/jn.00970.2002

Thompson, P., & Burr, D. (2009). Visual aftereffects. *Current Biology*, *19*(1), 11–14. https://doi.org/10.1016/j.cub.2008.10.014

Thorell, L., Devalois, R., & Albrecht, D. (1984). Spatial-mapping of monkey V1-cells with pure color and luminance stimuli. *Vision Research*, *24*(7), 751–769. Retrieved from http://links.isiglobalnet2.com/gateway/Gateway.cgi?GWVersion=2&SrcAuth=mekentosj&SrcApp=Papers&DestLinkType=FullRecord&DestApp=WOS&KeyUT=A1984TB32300014

To, M. P. S. (2010). Perception of suprathreshold naturalistic changes in colored natural images. *Journal of Vision*, *10*(4), 1–22. https://doi.org/10.1167/10.4.12

To, M. P. S., Gilchrist, I. D., & Tolhurst, D. J. (2015). Perception of differences in naturalistic dynamic scenes, and a V1-based model. *Journal of Vision*, *15*(1), 19–19. https://doi.org/10.1167/15.1.19

Tolhurst, D. J., Tadmor, Y., & Chao, T. (1992). The amplitude spectra of natural images. In *Ophthalmic and Physiological Optics* (Vol. 12, p. 97). https://doi.org/10.1016/0275-5408(92)90078-B

Tolias, A. S., Smirnakis, S. M., Augath, M. A., Trinath, T., & Logothetis, N. K. (2001). Motion processing in the macaque: revisited with functional magnetic resonance imaging. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *21*(21), 8594–8601.

https://doi.org/21/21/8594 [pii]

Troje, N. F., Sadr, J., Geyer, H., & Nakayama, K. (2006). Adaptation aftereffects in the perception of gender from biological motion. *Journal of Vision*, *6*(8), 850–857. https://doi.org/10.1167/6.8.7

Troje, N. F., & Westhoff, C. (2006). The Inversion Effect in Biological Motion Perception: Evidence for a "Life Detector"? *Current Biology*, *16*(8), 821–824. https://doi.org/10.1016/j.cub.2006.03.022

Ulanovsky, N., Las, L., Farkas, D., & Nelken, I. (2004). Multiple Time Scales of Adaptation in Auditory Cortex Neurons. *J. Neurosci.*, *24*(46), 10440–10453. https://doi.org/10.1523/JNEUROSCI.1905-04.2004

Van der Burg, E., Alais, D., & Cass, J. (2013). Rapid recalibration to audiovisual asynchrony. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *33*(37), 14633–14637. https://doi.org/10.1523/JNEUROSCI.1182-13.2013

VanRullen, R., & Thorpe, S. J. (2001). Is it a bird? Is it a plane? Ultra-rapid visual categorisation of natural and artifactual objects. *Perception*, *30*(6), 655–668. https://doi.org/10.1068/p3029

Vautin, R. G., & Berkley, M. A. (1977). Responses of single cells in cat visual cortex to prolonged stimulus movement: neural correlates of visual aftereffects. *Journal of Neurophysiology*, *40*(5), 1051–1065. https://doi.org/10.1152/jn.1977.40.5.1051

Vautin, R. G., & Dow, B. M. (1985). Color cell groups in foveal striate cortex of the behaving macaque. *Journal of Neurophysiology*, *54*(2), 273–292. https://doi.org/10.1152/jn.1985.54.2.273

Vogels, R. (1990). Population coding of stimulus orientation by striate cortical cells. *Biological Cybernetics*, *64*(1), 25–31. https://doi.org/10.1007/BF00203627

Vul, E., Krizay, E., & MacLeod, D. I. a. (2008). The McCollough effect reflects permanent and transient adaptation in early visual cortex. *Journal of Vision*, *8*(12), 4.1-12. https://doi.org/10.1167/8.12.4

Wainwright, M. J. (1999). Visual adaptation as optimal information transmission. *Vision Research*. https://doi.org/10.1016/S0042-6989(99)00101-7

Wainwright, M. J., Schwartz, O., & Simoncelli, E. P. (2001). Natural Image Statistics and Divisive Normalization: Modeling Nonlinearities and Adaptation in Cortical Neurons. *Statistical Theories of the Brain*, 203–222.

Walsh, V. (1995). Colour Vision: Adapting to change. *Current Biology*, *5*(7), 703–705. https://doi.org/10.1016/S0960-9822(95)00138-2

Wang, G., Ding, S., & Yunokuchi, K. (2003). Representation of cardinal contour overlaps less with representation of nearby angles in cat visual cortex. *Journal of Neurophysiology*, *90*(6), 3912–3920. https://doi.org/10.1152/jn.00219.2003

Wark, B., Fairhall, A., & Rieke, F. (2009). Timescales of Inference in Visual Adaptation. *Neuron*, *61*(5), 750–761. https://doi.org/10.1016/j.neuron.2009.01.019

Wark, B., Lundstrom, B. N., & Fairhall, A. (2007). Sensory adaptation. *Current Opinion in Neurobiology*. https://doi.org/10.1016/j.conb.2007.07.001

Webster, M. a. (2014). Probing the functions of contextual modulation by adapting images rather than observers. *Vision Research*, *104*, 68–79. https://doi.org/10.1016/j.visres.2014.09.003

Webster, M. A. (2015). Visual Adaptation. *Annual Review of Vision Science*, *1*(1), 547–567. https://doi.org/10.1146/annurev-vision-082114-035509

Webster, M. A., Kaping, D., Mizukami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, *428*(6982), 557–561. https://doi.org/10.1038/nature02420

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

https://doi.org/10.1016/S0042-6989(97)00125-9

Webster, M. A., & Mollon, V. (1995). Colour constancy influenced by contrast adaptation. *Nature*, *373*(6516), 692–694. https://doi.org/10.1038/373694a0

Webster, M. A., Werner, J. S., & Field, D. J. (2005). Adaptation and the Phenomenology of Perception. In *Fitting the Mind to the World: Adaptation and After-Effects in High-Level Vision*. https://doi.org/10.1093/acprof:oso/9780198529699.003.0010

Webster, M. a, & MacLeod, D. I. a. (2011). Visual adaptation and face perception. *Philosophical Transactions of the Royal Society of London. B*, *366*(1571), 1702–1725. https://doi.org/10.1098/rstb.2010.0360

Wei, X.-X., & Stocker, A. A. (2016). A new law of human perception. *BioRxiv*, 2010. https://doi.org/http://dx.doi.org/10.1101/091918

Weiss, Y., & Adelson, E. H. (1998). Slow and Smooth: a Bayesian theory for the combination of local motion signals in human vision. *Center for Biological and Computational Learning Paper*, *158*(1624), 1–42. Retrieved from http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Slow+and+Smooth:+a+Bayesian+theory+for+the+combination+of+local+motion+signals+in+human+vision#0

Wenderoth, P., & Johnstone, S. (1987). Possible neural substrates for orientation analysis and perception. *Perception*.

Werner, J. S., & Schefrin, B. E. (1993). Loci of achromatic points throughout the life span. *Journal of the Optical Society of America A*, *10*(7), 1509–1516. https://doi.org/10.1364/JOSAA.10.001509

Whitaker, D., & McGraw, P. V. (2000). Long-term visual experience recalibrates human orientation perception. *Nature Neuroscience*, *3*(1), 13--. https://doi.org/10.1038/71088

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

https://doi.org/10.1167/13.2.6

Wolfe, J. M. (1984). Short test flashes produce large tilt aftereffects. *Vision Research*, *24*(12), 1959–1964. https://doi.org/10.1016/0042-6989(84)90030-0

Wozny, D. R., & Shams, L. (2011). Recalibration of Auditory Space following Milliseconds of Cross-Modal Discrepancy. *Journal of Neuroscience*, *31*(12), 4607–4612. https://doi.org/10.1523/JNEUROSCI.6079-10.2011

Xia, Y., Leib, A. Y., Whitney, D., Arnheim, R., Shaffer, L., Solso, R. L., … M. (2016). Serial dependence in the perception of attractiveness. *Journal of Vision*, *16*(15), 28. https://doi.org/10.1167/16.15.28

Xu, X., Collins, C. E., Khaytin, I., Kaas, J. H., & Casagrande, V. a. (2006). Unequal representation of cardinal vs. oblique orientations in the middle temporal visual area. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(46), 17490–17495. https://doi.org/10.1073/pnas.0608502103

Yacoub, E., Harel, N., & Ugurbil, K. (2008). High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences*, *105*(30), 10607–10612. https://doi.org/10.1073/pnas.0804110105

Yao, H., & Dan, Y. (2001). Stimulus timing-dependent plasticity in cortical processing of orientation. *Neuron*, *32*(2), 315–323. https://doi.org/10.1016/S0896-6273(01)00460-3

Yarrow, K., Martin, S. E., Di Costa, S., Solomon, J. A., & Arnold, D. H. (2016). A roving dual-presentation simultaneity-judgment task to estimate the point of subjective simultaneity. *Frontiers in Psychology*, *7*(MAR), 1–19. https://doi.org/10.3389/fpsyg.2016.00416

Zamboni, E., Ledgeway, T., McGraw, P. V., & Schluppeck, D. (2016). Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proceedings of the Royal Society B:*

Biological Sciences, *283*(1833). https://doi.org/10.1098/rspb.2016.0263

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. https://doi.org/10.1016/j.cub.2009.10.018

Zipser, K. (2017). Visualizing fMRI BOLD responses to diverse naturalistic scenes using retinotopic projection. *Journal of Vision*, *17*(6), 18. https://doi.org/10.1167/17.6.18

Zou, J., He, S., & Zhang, P. (2016). Binocular rivalry from invisible patterns. *Proceedings of the National Academy of Sciences*, *113*(30), 8408–8413. https://doi.org/10.1073/pnas.1604816113