SIGNAL TRANSDUCTION AND CONTEXTUAL MODULATION IN THE HUMAN VISUAL SYSTEM

Martin Timothy Wilkinson Scott

PhD

The University of York

Psychology

April 2022

Thesis abstract

The transduction of luminance contrast is nonlinear. As contrast increases, responses first accelerate, and then saturate. This acceleration is thought to be responsible for the "dipper effect": the improvement in stimulus sensitivity at low absolute contrasts. Moreover, a stimulus that cannot provoke a neuronal response can still modulate a neuron's response to its preferred stimuli. Recent evidence suggests that similar computations exist in more complex visual parameter spaces. Furthermore, there is indication that the divisive inhibitory mechanisms of the brain are affected in aging. In this thesis, I present three experiments that explored these possibilities.

My first experiment investigated reports that surround suppression of contrast strengthens with age, implying disruption of excitatory/inhibitory balance. I suspected that previous reports may reflect contributions from overlay masking, a distinct form of suppression. Using stimuli that preclude overlay masking, I found surround suppression to be similar in younger and older observers, suggesting that spatial suppression is stable with age. I also reported a novel finding of untuned suprathreshold suppression in central vision.

My second experiment explored whether the "dipper effect" can be found in the mid-level perception of global form. I also expanded on reports of enhanced sensitivity to concentric and radial form. By manipulating the saliency of the global form percept elicited by Glass patterns, I reported a dipper effect that was similar across the pattern axes I examined. In my third experiment, using electroencephalography, I investigated whether this global form "dipper effect" is predicted by neuronal responses. At ventral electrodes, I found evidence for a transducer that would produce a dipper effect, though the level of global form for which it predicted peak sensitivity differed from psychophysical predictions. I reasoned that this discrepancy is likely due to methodological limitations, but also provided an explanation from attentional modulation and gain control.

Table of contents

Thesis abstract 1
Table of contents 2
Table of figures 5
Acknowledgements 6
Author's declaration 7
Chapter 1. General introduction 8
1.1 An introduction to nonlinear phenomena in the perception of luminance contrast 8
1.1.1 The linear-nonlinear model of early visual perception 8
1.1.2 The behavioural outcome of response nonlinearity 11
1.1.3 Spatial contributions to divisive nonlinearity 14
1.2 Gaps in our understanding of nonlinearity in visual perception 18
1.2.1 What is the effect of age on surround suppression? 19
1.2.2 Is there a "dipper effect" beyond low-level visual processing, and is it reflected in neuronal responses? 21
1.2.3 Is there a suprathreshold bias for the perception of polar form? 25
1.3 Rationale for the current thesis 25
Chapter 2. Measuring the effects of age on contrast suppression 28
2.1 Abstract 28
2.2 Introduction 28
2.3 Methods 32
2.3.1 Participants 32
2.3.2 Stimuli & apparatus 32
2.3.3 Procedure 33
2.3.4 Model fitting and statistical analysis 35
2.4 Results 36
2.4.1 Elevated contrast detection thresholds in older observers 36
2.4.2 Surround suppression does not change with age 37
2.5 Discussion 41
2.5.1 An alternative explanation for age effects in contrast suppression 41
2.5.2 Broadly tuned central supra-threshold surround suppression 44
2.6 Conclusion 46
Chapter 3. Human sensitivity to global form violates Weber's law 47

3.1 Abstract 47
3.2 Introduction 47
3.3 Methods 51
3.3.1 Participants 51
3.3.2 Online design and calibration 51
3.3.3 Visual Stimuli 52
3.3.4 Experimental procedure 56
3.3.5 Model fitting 57
3.4 Results 59
3.4.1 Example psychometric function 59
3.4.2 Normalisation to variable detection thresholds 60
3.4.3 Violation of Weber's law at low global form coherence 62
3.4.4 No evidence for enhanced sensitivity to polar form 66
3.5 Discussion 68
3.5.1 A "dipper effect" for global form coherence? 68
3.5.2 Nonlinear transduction or uncertainty reduction? 69
3.5.3 Equal sensitivity along different global form axes 71
3.6 Conclusion 72
Chapter 4. Evidence of nonlinear transduction in the perception of global form 73
4.1 Abstract 73
4.2 Introduction 73
4.3 Methods 76
4.3.1 Participants 76
4.3.2 Visual Stimuli 77
4.3.3 Experimental procedure 81
4.3.4 Data acquisition and pre-processing 81
4.3.5 Reliable components analysis 83
4.3.6 Signal-to-noise ratio calculation 84
4.3.7 Model fitting 85
4.4 Results 86
4.4.1 Grand-average responses to global and local switching 86
4.4.2 Spatially extensive responses to global form 87
4.4.3 Responses to local resampling 89
4.4.4 Lateralised responses scale with global form coherence 90

4.4.5 An accelerating nonlinearity independent of pattern axis 92
4.4.6 Response dynamic range is higher than behavioural predictions 95
4.5 Discussion 97
4.5.1 Relating global form responses to behavioural thresholds 98
4.5.2 Little evidence for polar form enhancement 100
4.6 Conclusion 101
Chapter 5. General discussion 102
5.1 Summary of experimental work 102
5.2 Future directions 103
5.3 Final conclusions 107
References 108
Appendices 131
Appendix A: Individual contrast-matching psychometric functions 131
Appendix B: ANOVA tables for Chapter 2 with and without outliers 143
Appendix C: Global form discrimination psychometric functions for individual observers 145
Appendix D: Psychometric function slopes for global form thresholds 166
Appendix E: Contrasts of local component SNR across global form coherence levels 167
Appendix F: SSVEP model fits with individual datapoints 170

Table of figures

Figure 1.1: Contrast response functions of macaque striate neurons 9
Figure 1.2: Producing a sigmoidal nonlinearity via divisive inhibition 11
Figure 1.3: Relating the neuronal transducer to behavioural sensitivity 12
Figure 1.4: Stimuli producing distinct forms of contrast suppression/masking 15
Figure 1.5: Spatial tuning of surround suppression and overlay masking 17
Figure 1.6: Stimuli previously used to investigate age and surround suppression of contrast. 20
Figure 1.7: Glass patterns across three alignment axes 23
Figure 2.1: Centre and surround spatial configuration 33
Figure 2.2: Contrast matching procedure 35
Figure 2.3: Detection thresholds compared across age-groups 37
Figure 2.4: Example psychometric function fits 38
Figure 2.5: Points of subjective equivalence split by age-group 39
Figure 2.6: Points of subjective equivalence collapsed across all observers 40
Figure 3.1: Spatial configuration of Glass patterns 52
Figure 3.2: Parameterisation of global form 54
Figure 3.3: Transformation of kappa to coherence 56
Figure 3.4: Stimulus presentation 57
Figure 3.5: Example global form psychometric function fits 60
Figure 3.6: Coherence discrimination and detection thresholds 61
Figure 3.7: Normalised coherence discrimination threshold 62
Figure 3.8: Model fits to normalised thresholds 63
Figure 3.9: Contrast of model residuals 64
Figure 3.10: Hyperbolic-ratio exponent parameter 65
Figure 3.11: Fit of a power law suprathreshold discrimination 66
Figure 3.12: Cross-condition contrast of model fit parameters 67
Figure 3.13: Reaction times as function of c50 69
Figure 4.1: Glass pattern spatial specification 78
Figure 4.2: Presentation protocol for a single pattern condition 80
Figure 4.3: Grand average sensor timeseries and spectrum 87
Figure 4.4: Hierarchy of reliable components 88
Figure 4.5: Topographies of reliable components 89
Figure 4.6: SNR of local RCs across global form coherence levels 90
Figure 4.7: SNR of global RCs across global form coherence levels and pattern axes 91
Figure 4.8: Contrasts of SNR between each RC, across coherence and pattern axes 92
Figure 4.9: Bootstrapped model fits to median SNR 93
Figure 4.10: Comparison of fit parameters between Glass pattern axes 94
Figure 4.11: Comparison of SNR at maximum coherence across pattern axes 95
Figure 4.12: Comparison of SSVEP models with psychophysical predictions 96
Figure 4.13: Contrast of SSVEP and psychophysical fit parameters 97

Acknowledgements

Many individuals have helped me to complete this thesis. Principal amongst them is my supervisor - Doctor Heidi Baseler - who has carefully and effectively guided me through a turbulent time to be a scientist. I would also like to thank Professor Alex Wade, whose door was always open (unlocked, at least) for conversations on other aspects of being an effective scientist. I also thank Professor Antony Morland and the rest of the team at the York Neuroimaging centre for their support and willingness to share their advice and experience. I would also like to thank Dr. Daniel Baker for his advice on experimental design. Several projects fell by the wayside due to the Coronavirus, but my then-collaborators Doctor.. Aneurin Kennerley and Professor Glen Jeffrey should still be thanked for their input, though it has not made it into my thesis. I wish them well in their continued scientific endeavours.

I would also like to thank my thesis advisory panel members: Professor Sven Mattys and Doctor Emma-Hayiou-Thomas for their time and support. Also, Professors Jonathan Smallwood and Elizabeth Jeffries, and Drs. Scott Cairney and Andre Gouws for employing me as a research assistant prior to the commencement of my PhD studies. Special thanks goes to Dr. Mladen Szormaz for initially making the case for my employment. My work with them built a foundation of skill and confidence that I believe allowed me to obtain a scholarship. Thanks goes also to my cherished friends and peers Juliana and Kevin Olivier, Anna Guttensen, Miaomiao Yu, Catia Oliveira, and Laia Bosque Mercader. A special mention also goes to Barbara Molz, my good friend and PhD-problem confidante. They helped me keep issues of academia in perspective: I hope I have done the same thing for them.

Last, but far from least, I thank my family, extended and immediate. This thesis, and my eventual doctoral award, is a monument to their unyielding support. Particularly my late mother, whose passion for the arts, teaching, and learning will always be with me, as will her memory. My Father too: a deep well of level-headed advice and perspective, and finally my sister, who has often entertained stressed-out phone calls.

Author’s declaration

I declare that this thesis is a presentation of original work, and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References.

Chapter 1. General introduction

This thesis is focused on nonlinear visual processing phenomena, that is, operations carried out by neuronal populations that go beyond simple additions or subtractions. In this chapter, I will begin by providing an overview of nonlinear transduction in low level visual perception and the different forms of spatially-dependant contrast suppression. This will be followed by an outline of the gaps in the literature that have motivated my experimental chapters. I will conclude this chapter with more specific rationale for each experiment I have carried out.

1.1 An introduction to nonlinear phenomena in the perception of luminance contrast

1.1.1 The linear-nonlinear model of early visual perception

In the early visual system, we tend to model the initial response of a neuron to a visual stimulus as the sum of its weighted ON/OFF subregion activations (Carandini et al., 2005). These subregions receive feedforward responses from earlier processing nodes, originating from the photoreceptors of the retina. If the area of visual space subtending the ON subregion increases in luminance, the excitatory potential of the neuron is increased, when an increase in luminance subtends the OFF subregion, the inhibitory potential is increased. When the spatial distribution of luminance is sufficient for excitation to escape inhibition, the neuron will produce an action potential. Therefore, the spatial distribution of luminance in the perceived environment (i.e., edges, contours, patterns) is encoded via the unique shape, size, and numerosity of ON/OFF subregions within a neuron's receptive field. The visual systems' ability to encode the degree to which a stimulus matches these spatial filters ultimately allows perceptually informed decision making.

For a given point in space, the granularity of spatial information is not carried by any single neuron, but rather by an entire population of neurons with different and complementary subregion configurations. Retinal ganglion cells, and many neurons within the lateral geniculate nucleus (LGN), have centre-surround receptive fields (Hubel & Wiesel, 1960; Kuffler, 1953; Wiesel & Hubel, 1966), with ON and OFF subregions often modelled as two concentric circles. These encode luminance "blobs", and have a preferred spatial frequency dictated by the size of the central subregion

(which can be ON or OFF weighted). Though these neurons will fire vigorously for spatial luminance modulation at their preferred spatial frequency, they still produce a gradient of responses for suboptimal spatial frequencies (Enroth-Cugell & Robson, 1966; Hawken et al., 1997). This is because stimuli that partially match the spatial filter described by the receptive field can still evoke a response. The same is true of the elliptical receptive fields of V1 simple cells (to which LGN neurons have feedforward connections) though their asymmetric configuration additionally confers orientation selectivity (Hubel & Wiesel, 1968; Ringach et al., 1997, 2002). Therefore, up to V1 (and indeed beyond), the response of a neuron is contingent on the spatial modulation of luminance. However, responses also depend on the amplitude of this modulation: the luminance contrast. Individual neuronal responses of the early visual system tend to monotonically increase with luminance contrast, aside from a minority of neurons that reduce in response after reaching their maximum (supersaturation) (Albrecht & Hamilton, 1982). Furthermore, some forms of spatial tuning (such as orientation & spatial frequency) are often found to be contrast invariant (Sclar & Freeman, 1982; Skottun et al., 1987), but receptive field size may vary with contrast (Sceniak et al., 1999).

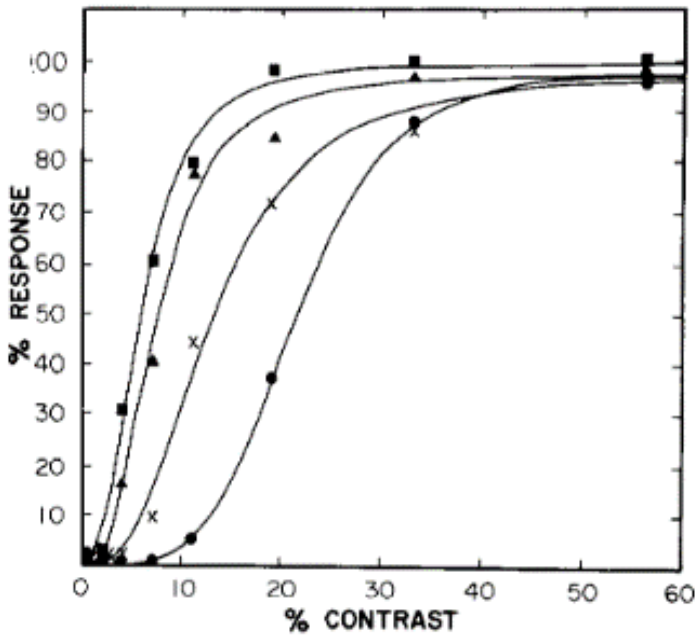

Figure 1.1: Contrast response functions of macaque striate neurons
The responses of four striate neurons (different markers) across a range of luminance contrasts. Curves represent the best-fit of a hyperbolic ratio function. Reprinted from the Journal of Neurophysiology, 48, 217-37, Albrecht & Hamilton (1982), "Striate cortex of monkey and cat: contrast response function", with permission from The American Physiological Society.

If the visual system were a strictly linear system, we might expect neural response to scale proportionally with contrast, such that doubling the contrast doubles the response. Instead, we find a remarkable set of response nonlinearities. The work of Albrecht and Hamilton (1982) demonstrated that many macaque V1 cells have a contrast response function that is sigmoidal when presented on linear axes. The responses they recorded for several V1 neurons are shown in Figure 1.1 - as contrast increases, neuronal responses tend to first accelerate, often reaching half of their maximum response by 10-30% contrast, and then to compress to the point of saturation, sometimes as early as ~30%. This means that a small range of lower contrasts - those within the accelerating regime of the function - occupy a large proportion of the neuron's dynamic range, and thus have high response granularity relative to the later compressive regime. The end-product is a bank of neurons with a discrete contrast range of high sensitivity that can differ considerably between cells (Albrecht & Hamilton, 1982), and cortical layers (Tootell et al., 1988). The neuronal mechanism that produces response acceleration and saturation is still a topic for investigation, but there is evidence to suggest (at least for saturation) that they are not simply the consequence of a static biophysical limitation (Peirce, 2007), but rather an active control process that can adjust the response rate of the cell according to the spatiotemporal parameters of the stimulus.

Whatever the generative cellular process may be, it has long been known that computational models incorporating exponential and divisive nonlinearities (originating from the work of Naka and Rushton (1966)) often provide good fits to recorded responses (Albrecht & Hamilton, 1982; Heeger, 1992a; Somers et al., 1998). It has even been proposed that response expansion and compression are fundamental response transformations that can be found throughout the brain (Carandini & Heeger, 2011) that allow neuronal responses to be much more sensitive over a discrete range of contrasts, relative to a linear response relationship. These models, an example of which is inlaid in Figure 1.2B, typically include an inhibitory (denominator) input that is outpaced by an excitatory (numerator) input across lower contrasts. At low contrasts (c), due to an additive constant ($c50$) that is applied to the denominator, the ratio of excitation to inhibition (i.e., model neuronal response) accelerates freely thanks to an exponentiation term (n). As the input contrast approaches and exceeds the additive constant, the denominator acts to more substantially limit the response of the cell.

Figure 1.2 graphs the transducer model used by Albrecht and Hamilton (1982) alongside its theoretical relationship to excitation and inhibition (Figure 1.2A is essentially the quotient of the lines graphed in Figure 1.2B). At low contrasts, we can see that inhibition remains quite stable, while the excitatory drive ramps-up, meaning the neural response is free to accelerate. As contrast increases, inhibitory drive increases alongside excitation, meaning the response function asymptotes. It should be noted similar response nonlinearities for contrast have been identified at the population level in human observers using non-invasive recording techniques such as fMRI (Boynton et al., 1999) and EEG (Baker et al., 2021). Irrespective of the best fitting model, an interesting aspect of the observed response nonlinearity is the degree to which it is reflected in behavioural thresholds. If behavioural detection and discrimination thresholds for contrast are purely the derivative of neuronal response functions, then we should expect observers to be most sensitive to stimulus changes around the steepest part of the neuronal response function.

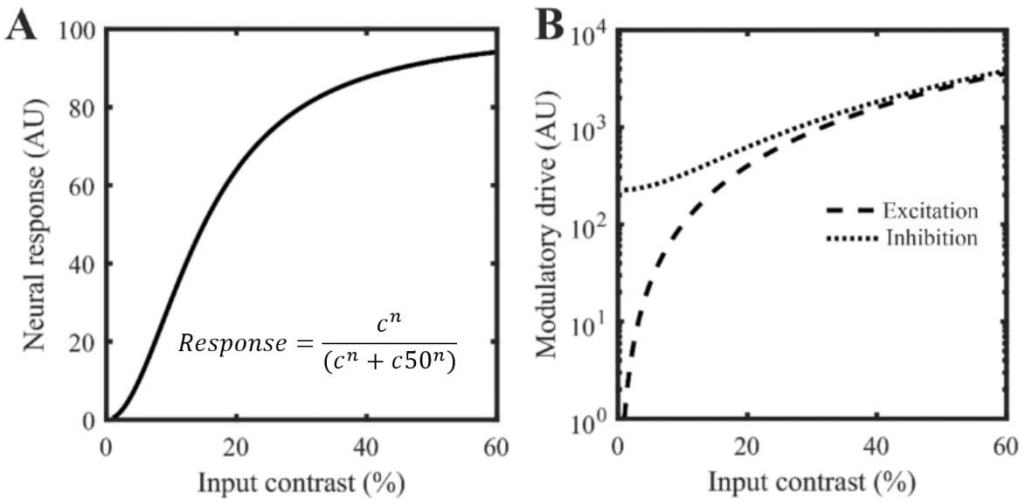

Figure 1.2: Producing a sigmoidal nonlinearity via divisive inhibition
A: The model response of single V1 neuron to a gradient of luminance contrasts. B: The numerator and denominator (excitation and inhibition, respectively) of the hyperbolic ratio function shown in A. The pointwise quotient of B (excitation/inhibition) produces A. This figure was generated using the "H-Ratio" function from Albrecht and Hamilton (1982)*, which is inlaid in panel A. The rMAX and rMIN parameters of this function have been omitted for brevity.*

1.1.2 The behavioural outcome of response nonlinearity

Observers' sensitivity to changes in stimulus intensity can be measured along any continuous parameter space through the use of a pedestal experiment. Here, measurements of just-noticeable differences (JNDs) are used to assess the absolute detection threshold (the JND from zero intensity), and discrimination thresholds (the

JND from a non-zero "pedestal"). By measuring JNDs at multiple pedestal intensities, we can discretely sample an observer's JND threshold vs intensity (TvI) function, which describes sensitivity as a function of the absolute intensity of the stimulus. Assuming behavioural responses are limited by a static source of noise after any compressive or expansive nonlinearity, it is the integral of this function that should approximate the shape of the neuronal transducer. Figure 1.3 illustrates this relationship; panel B is the shape of the behavioural TvI function predicted by the neuronal transducer in panel A (B is the first derivative of A). Notice that where the neuronal response is accelerating, behavioural sensitivity is predicted to be at its best ("facilitated"), with discrimination thresholds smaller even than the threshold for absolute detection (the threshold at a pedestal of zero). Likewise, as the response function saturates discrimination thresholds increase (are "suppressed"), reflecting the need for greater increases in contrast to achieve the same change in neuronal response. However, the question remains: do we find evidence for such a pattern of behavioural sensitivity to contrast as predicted by neuronal responses?

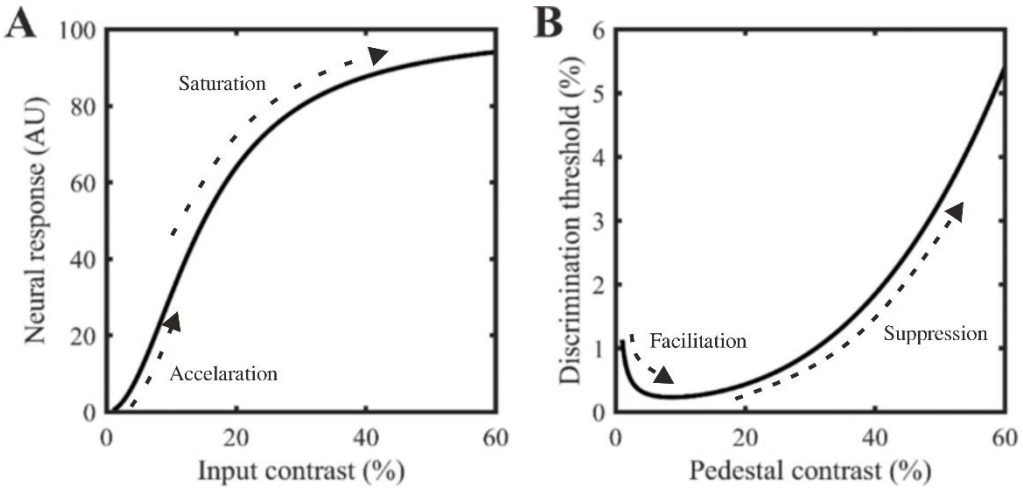

Figure 1.3: Relating the neuronal transducer to behavioural sensitivity
A: The model response of single V1 neuron to a gradient of luminance contrasts (the same as Figure 1.2A). B: The first derivative of the curve in panel A, the rate of change in A dictates the form of B. Inlaid text and arrows describe the distinct regimes of each function.

Evidence for a saturating nonlinearity can be found at the inception of formal psychophysical experimentation. This was implied in the sensation of weight from the work of Ernst Weber, later published and codified by Theodore Fechner as "Weber's law" (Fechner, 1860). Weber's observation was that the amount of added weight required to discern the weight of two held objects was a constant fraction of the weights being compared (a line of unity slope on logarithmic axes). In a derivation

published alongside Weber's law, Fechner attributed this proportionality to a logarithmic scaling of internal sensations (similar to the excitatory curve of Figure 1.2B), which we may now interpret as a form of compressive nonlinearity in neuronal responses. Aside from weight sensation, there are multiple sensory parameter spaces that approximately adhere to Weber's law (with some deviations at high/low intensities), including luminance (Barlow, 1957) and loudness (Parker & Schneider, 1980), and the detection of global form in corruptive noise (Maloney et al., 1987).

Interestingly, for luminance *contrast* however, discrimination thresholds at small (but suprathreshold) pedestals show a distinct non-monotonicity that is consistent with an accelerating nonlinearity. This was initially recognised by Nachmias and Sansbury (1974), who themselves proposed an expansive nonlinearity as an explanation. They found that notice that discrimination thresholds are best at pedestals near the absolute detection threshold (i.e., they are "facilitated" relative to detection), but worsen as contrast increases – consistent with a compressive nonlinearity. This pattern of discrimination thresholds (see Figure 1 of Nachmias and Sansbury (1974)) often referred to as the "dipper" or "pedestal" effect, has been quite convincingly replicated (Bird et al., 2002; Boynton et al., 1999; Boynton & Foley, 1999; G. J. Burton, 1981; Foley & Legge, 1981), and is indeed well described by a model incorporating response exponentiation and compression (Foley, 1994). That behavioural thresholds are well described by a nonlinear contrast transducer is further supported by fMRI work from Boynton and colleagues (1999), who, in human observers, measured the blood-oxygen level dependant (BOLD) signal in V1 and discrimination thresholds as a function of luminance contrast. By fitting a variation of the transducer model used by Albrecht and Hamilton (1982) to both JNDs and V1 BOLD contrast responses, they found the best fit was produced by an accelerating and saturating nonlinearity.

While an observable response nonlinearity in agreement with psychophysical sensitivity heavily favours an explanation from nonlinear transduction, it should be noted that there are competing explanations for the dipper effect that do not rely on a nonlinear transducer. These typically dispute the aforementioned assumption of a late, static, and additive noise source, contending that noise proportional (or related to) the input signal could also produce the facilitation and suppression of the dipper effect.

There are several such explanations: uncertainty reduction could account for facilitation at low contrasts (Pelli, 1985), and multiplicative noise with a linear transducer rather than additive noise with a nonlinear transducer could produce suppression as contrast increases (Kontsevich et al., 2002). These alternative explanations will be elaborated upon when appropriate in the forthcoming chapters, but Solomon (2009) provides a relatively recent review of alternative accounts for the dipper effect (see also Sanborn and Dayan (2011)).

1.1.3 Spatial contributions to divisive nonlinearity

Thus far, I have outlined the nonlinearity of the contrast response function and its behavioural consequence. However, there is a wealth of evidence that nonlinearity can also be observed in the lateral and feedback interactions between neurons with different receptive fields across the visual hierarchy (Angelucci et al., 2017). One of the advantages of the accelerating and saturating nonlinearity is greater response resolution over the steepest section of a neuron's transducer function. However, this also introduces a problem: what if the average contrast of a stimulus is below or beyond a neuron's range of sensitivity? Early observations in feline and primate electrophysiology revealed that the firing rate of a neuron to a stimulus of optimal spatial frequency and orientation can be reduced by the simultaneous presentation of a non-optimal stimulus, even when this stimulus itself cannot provoke a response (Bonds, 1989; DeAngelis et al., 1992), suggestive of some active contribution from other neurons with different receptive field properties. Initially referred to as "nonspecific suppression", this phenomenon cannot be explained by a purely linear computation. Within a receptive field, summing or subtracting responses between an optimal stimulus and a stimulus that cannot provoke any response should yield only the response of the optimal stimulus. The prevalent explanation for this reduction in neuronal activity is an inhibitory input from other neurons with proximal receptive field locations, and that this inhibition is not a just linear subtraction, but rather a division (Carandini et al., 1997; Cavanaugh et al., 2002). This computation is thought to reflect a form of contrast normalisation employed by the visual system, such that neurons' response functions are horizontally shifted to cover the average contrast of the stimulus. Contrast suppression is predicted by an extension of the model used to fit responses to increasing contrast, but the denominator now receives contributions from a "normalisation pool", which contains the summed responses from other

neurons with different or similar spatial feature sensitivity (orientation, spatial frequency, receptive field location, etc.), which can inform the spatial selectivity of suppression (Carandini & Heeger, 2011; Heeger, 1992a).

Response normalisation is thought to be another canonical computation employed by neurons that serves to increase the effective dynamic range for contrast (Carandini & Heeger, 2011), and there is evidence that this computation is applied differently for different stimulus configurations. We tend to make a qualitative distinction between suppression that is triggered by stimuli that are superimposed onto the original stimulus (often called overlay masking) and suppression triggered by a non-overlapping stimulus, usually via an annular surround (referred to as surround suppression). An illustration of this spatial distinction is presented in Figure 1.4, via the stimuli used in Petrov et al., (2005), who's work will be discussed in a later section. In terms of projections to V1, this distinction is analogous to the difference between suppression that occurs within a cortical hyper-column (overlapping receptive fields) and suppression that occurs across hyper-columns (nearby/surrounding receptive fields), though there is evidence of suppressive contributions from pre-cortical mechanisms in both cases.

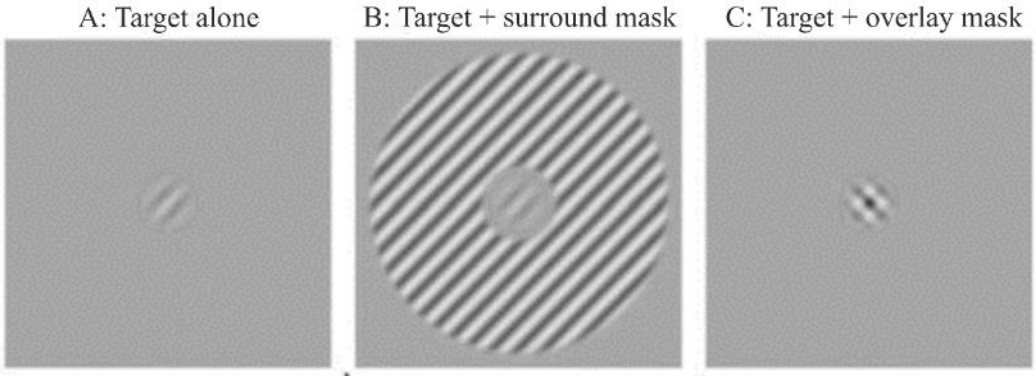

Figure 1.4: Stimuli producing distinct forms of contrast suppression/masking
A: The central target to be suppressed/masked. B: An annular mask designed to produce surround suppression. C: A superimposed mask designed to produce overlay masking. Note, the term "mask" is often used to refer to the suppressive stimulus, particularly in psychophysical experimentation. Figure taken from Petrov et al. (2005), *Figure 1 (with minor formatting alterations). Copyright 2022 Society for Neuroscience.*

It is challenging to discern the cortical and pre-cortical contributions to overlay masking and surround suppression. Both electrophysiology and psychophysical assessment tend to do so by comparing the spatiotemporal selectivity profiles of overlay masking and surround suppression (i.e., the conditions in which they are strongest). These suppression/masking tuning functions can then be related to

previously established receptive field configurations and temporal sensitivities. As previously described, different visual processing regions contain neurons with distinct receptive field properties. While there is some overlap in these properties between processing regions, this can still provide some indication of the generative neuronal population. Psychophysically, Meier and Carandini (2002) demonstrated that overlay masking is still strong when using gratings that are drifting faster than the temporal resolution of most cortical neurons, consistent with similar findings from feline electrophysiology (Freeman et al., 2002; Sengpiel & Vorobyov, 2005). The inverse has been observed for surround suppression, the effect of which can be considerably diminished by the same manipulation (Durand et al., 2007). This finding suggests that surround suppression has a significant basis in cortex, as this manipulation is designed to attenuate responses from cortical neurons (which are selective for slower drift-rates), but preserve responses from pre-cortical neurons which response to a greater range of drift frequencies. Although, since Durand and colleagues did not find surround suppression to be entirely abolished by high drift rates, both cortical and pre-cortical sites may still be implicated in surround suppression. This is compatible with eye-of-origin experiments that find distinct spatiotemporal tuning functions for within-eye and across-eye surround suppression in the responses of macaque V1 (Webb et al., 2005) and human contrast detection thresholds (Cai et al., 2008; Petrov & McKee, 2009; Schallmo & Murray, 2016), suggesting different contributions from suppressive mechanisms occurring before and after binocular fusion. Similar eye-of-origin experiments for overlay masking also suggest the presence of multiple generative components (Baker, Meese, & Summers, 2007; Meese & Baker, 2009). Though, it should be noted that a phenomenon being abolished by dichoptic presentation does not guarantee a pre-cortical locus, as monocularly driven neurons with centre-surround receptive fields can still be found in primary visual cortex (Hubel & Wiesel, 1968).

Few experiments have directly compared the tuning characteristics of overlay masking and surround suppression in human observers; the work of Petrov et al. (2005) is currently authoritative in this regard. Using binocular stimuli, they examined the effect of both overlay masking and surround suppression on contrast detection thresholds, but varied the degree of spatial similarity between the target and suppressive stimulus (the latter is often referred to as the “masker”). They found

surround suppression to be more dependent on the degree of similarity between the target stimulus and the suppressive surround, as well as the eccentricity at which stimuli were presented (Figure 1.5 shows their suppression/masking tuning curves). Surround suppression was effectively absent at the fovea, but also into the perifovea when surrounds were oriented orthogonal to the target stimulus (indicative of sharp orientation tuning). In contrast, overlay masking was relatively stable across all eccentricities, and only broadly orientation tuned, still doubling the detection threshold even with an orthogonal mask. Both forms of suppression/masking showed some dependence on the spatial frequency of the surround, but overlay masking was only vulnerable to severe spatial frequency discontinuities, whereas surround suppression had a clear preference for similarity. The orientation specificity of surround suppression is consistent with a cortical contribution, as elongated receptive fields first arise in V1 (Scholl et al., 2013), and orientation biases are quite weak in pre-cortical neurons (Xu et al., 2002), likely due to slightly elliptical receptive centre-surround receptive fields.

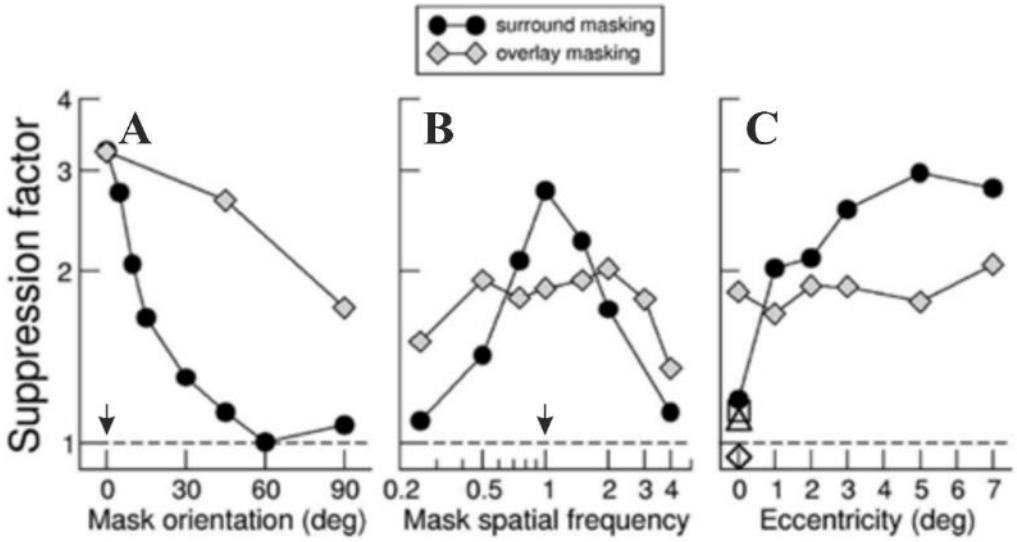

Figure 1.5: Spatial tuning of surround suppression and overlay masking
A: Selectivity for mask orientation. B: Selectivity for mask spatial frequency. C: Selectivity for eccentricity (both central probe and mask were moved together). A suppression factor of 1 is indicative of no suppression, positive factors indicate suppression. Figure taken from Petrov et al., (2005)*, Figure 2, with the addition of arrows on the abscissae of panels A and B to illustrate the parameters of the central target stimulus. Copyright 2022 Society for Neuroscience.*

Petrov and colleagues (2005) also examined the order in which the surround and overlay suppression are applied. They accomplished this by superimposing an orthogonal mask onto a suppressive annulus, finding that the suppressive effect of an annular surround could be severely attenuated under this condition. This implies that

overlay masking is applied before surround suppression as the annulus was suppressed by the orthogonal mask before surround suppression could be enacted. This seriality does not confirm a pre-cortical origin for overlay masking, though it is compatible with it. Overall, Petrov and colleagues' report overlay masking to be broadly spatially tuned, consistent with previous work from feline and macaque models (Bonds, 1989; DeAngelis et al., 1992), and to be present throughout the visual field. Conversely, surround suppression (which occurs later in the processing hierarchy) is orientation and spatial frequency selective, but effectively absent in central vision.

Interestingly, it appears that Petrov and colleagues' finding that surround suppression is absent in the fovea may only be true of detection thresholds (subthreshold stimuli). Several studies have found surround suppression of contrast in central vision when using readily detectable (suprathreshold) target probes (Cai et al., 2008; Cannon & Fullenkamp, 1991; Meese & Hess, 2004; Nurminen et al., 2010; Vanegas et al., 2015; Xing & Heeger, 2000). The subset of experiments that also examined the orientation specificity of central-vision suprathreshold surround suppression have found it to be orientation tuned (Cannon & Fullenkamp, 1991; Xing & Heeger, 2000), with Xing and Heeger reporting it strengthen at greater eccentricities, much like the suppression reported by Petrov and colleagues. That the effect of surround suppression depends on the absolute contrast of the stimulus has also been suggested by primate electrophysiology, where surround suppression has been observed to broaden in its orientation tuning at lower contrasts (Levitt & Lund, 1997; Webb et al., 2005), though these measurements are obtained from neurons with perifoveal receptive fields, so their sensitivity profiles may not generalise to fixation. Overall, it is clear that the terms "surround" and "overlay" suppression/masking only succeed in describing the physical attributes of a stimulus that can produce them, but do not convey the variety of spatiotemporal parameters over which they can be distinguished, including sensitivities to orientation, spatial frequency, temporal frequency, eccentricity, and absolute contrast.

1.2 Gaps in our understanding of nonlinearity in visual perception

This section will more specifically outline the context for the research questions addressed by the experimental chapters of this thesis, which are based in the

content of the broader overview provided in the previous section. I begin with the observations leading to my first experimental chapter, which explores the possibility of a change in surround suppression strength across the human lifespan. My second and third experimental chapters focus on investigating nonlinear transduction beyond the representation of luminance contrast.

1.2.1 What is the effect of age on surround suppression?

One of the most robust effects of age on visual perception is an increase contrast detection thresholds at intermediate and high spatial frequencies (Arundale, 1978; Beard et al., 1994; Owsley, 1983), which can mostly be attributed to changes to the optical media (Owsley, 2011). While there is evidence for some role of post-optical neuronal factors at the retina and beyond (Spear, 1993), there is little evidence that these actually lead to any practical reductions in contrast perception (Owsley, 2011). Nevertheless, there are functional and structural changes to neurons in primary visual cortex with age in both humans (Brewer & Barton, 2012) and in the macaque (Schmolesky et al., 2000; Yu et al., 2006), so investigating age-related changes in cortically dependant visual processing is worthwhile. One such form of processing, as I have described in the previous section, is surround suppression of contrast. There is already evidence that alterations in neuronal inhibitory/excitatory balance accompany aging (Betts et al., 2005), but also Alzheimer's disease (Zhuang et al., 2016), autism (Schallmo et al., 2020) and schizophrenia (Serrano-Pedraza et al., 2014). Although these findings come from experiments focused on the suppression of motion direction perception (a phenomenon identified by Tadin (2003)), they highlight the diagnostic potential of simple psychophysical experiments designed to measure processes of spatial suppression. Indeed, if processes of normalisation (a suspected "canonical" computation) are generally affected throughout the brain in these populations, we should also expect surround suppression of contrast to be compromised.

Betts and colleagues' (2005) findings suggested that surround suppression of motion direction weakens with age, a finding that has provoked a decade of research investigating the consequences of age on surround suppression of contrast. Interestingly, this body of work (which probes suprathreshold surround suppression of contrast at fixation) routinely finds evidence to the contrary of Betts et al. – suppression of contrast may strengthen with age (Karas & McKendrick, 2009, 2011,

2015), such that older observers (>60 years of age) experience a more severe reduction in apparent target contrast. Importantly, this effect does not appear to be related to the increase in absolute detection thresholds that accompanies advancing age (Karas & McKendrick, 2011). It has also been shown to be robust to different stimulus configurations, with initial work using the Chubb illusion (Chubb et al., 1989; Karas & McKendrick, 2009) and more recent work showing the same effect with static and drifting sinusoidal gratings (Karas & McKendrick, 2012, 2015). The augmentation of surround suppression in older observers is predominant at low contrasts (20%) absent for very short stimulus durations of 40ms (Pitchaimuthu, Nguyen, et al., 2017), strongest at intermediary presentation times of 100ms, and still present in some older observers up to 500ms (Karas & McKendrick, 2015). Interestingly, the eye-of-origin work from Pitchaimuthu et al. (2017) found the age-effect to be abolished by dichoptic presentation of the centre and surround, suggesting that the neuronal circuitry affected by age occurs prior to binocular fusion. That the effect occurs pre-fusion and is strongest at low contrasts may suggest that it is based in a mechanism analogous to the low-contrast and monoptic (within-eye) component of surround suppression identified in macaque V1 by Webb et al. (2005).

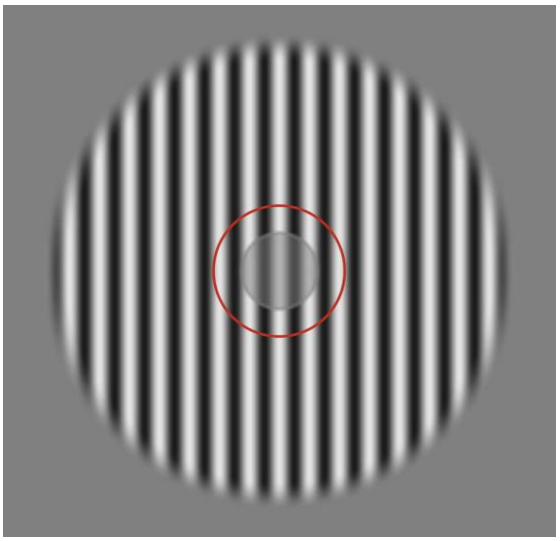

Figure 1.6: Stimuli previously used to investigate age and surround suppression of contrast.
The spatial configuration of the stimulus typically used by McKendrick and colleagues. The red circle has been added to illustrate the extent to which the surround may encroach on mechanisms of overlay masking according to Petrov et al., (2005). *This image has been generated based on the parameters of McKendrick et al. (2015)*

Indeed, the mechanism affected by age has often been interpreted to be surround suppression in previous reports. However, there remains the possibility that

overlay masking is the mechanism affected by age. Petrov and colleagues' (2005) measurement of the spatial tuning characteristics of overlay masking and surround suppression provided an important reference point for understanding how stimulus configuration affects suppressive tuning characteristics. Importantly, they also took steps to increase the likelihood of separately measuring these two types of suppression, though these precautions are scarcely adopted by subsequent experimentation probing suprathreshold surround suppression. When testing surround suppression, Petrov and colleagues principal precaution was the addition of a mean luminance gap between the central grating and surrounding annulus (see Figure 1.4). This gap, a width of 1 grating cycle, is important because it (in theory) reduces the probability of overlay masking "leaking" into surround suppression. When using a surrounding annulus with an inner border that directly (or very closely) abuts the border of the central probe, some receptive fields will be occupied by both the centre and surround, effectively producing overlay masking in these neuronal populations. By introducing a one cycle gap between the centre and surround, contamination from overlay masking can no longer be sourced from neurons with a preference for the spatial frequency of the central stimulus and only a single ON-OFF subregion cycle. Neurons with very large receptive fields, or more than a single subregion cycle (which do occur in V1 – see Chen et al.,(2020)) could still contribute, but suppression from the centre-surround receptive fields of the retina and LGN (which are a strong candidate for overlay masking) would be heavily attenuated. The stimuli used by McKendrick and colleagues regularly employ a surround that closely (or directly) abuts the central probe (c.f. the red circle in Figure 1.6 and Figure 1.4B). This means that there is an open question as to whether their age-related effect persists in stimuli designed to preclude contributions from overlay masking.

1.2.2 Is there a "dipper effect" beyond low-level visual processing, and is it reflected in neuronal responses?

Most of our current understanding of nonlinear transduction, and comparisons to behavioural sensitivity, have been made in the luminance contrast parameter space. The lack of similar explorations for higher-level visual parameter spaces is justifiable: luminance contrast is encoded as early as the retina, and homologues of early visual processing nodes can be found in other mammalian species (Orban et al., 2004; Payne, 1993). However, extensive investigation of low-level luminance contrast encoding has

supported the development of predictive models which can be compared to human behaviour. Moreover, it has led to the proposition of "canonical computations" – mathematical functions that are employed throughout the brain. Exponentiation and saturation of neuronal responses have been proposed to be such computations (Carandini & Heeger, 2011), implying that sigmoidal transducer nonlinearities may be found throughout the visual system, along with its theorised consequence in behaviour – the "dipper effect". This leads to the following question: do we find evidence for nonlinear transduction in more derivative parameter spaces, beyond luminance contrast?

Accepting that the "dipper effect" is a consequence of nonlinear transduction, then evidence for an accelerating and saturating nonlinearity can be found in other domains. For example, in the visual and tactile perception of speed, speed discrimination (that of a spinning wheel) is facilitated at low speeds relative to the speed required to detect movement (Gori et al., 2011; Simpson & Finsten, 1995). A similar finding comes from Huang and Chen (2014) for the visual perception of 2nd order patterns (gratings defined by contrast boundaries rather than luminance boundaries), and Morgan et al., (2008) for the discrimination of orientation variance in a field of small Gabor patches. A dipper effect has even been identified in a parameter space as complex as human facial expression. Gray et al. (2020) collected discrimination thresholds and EEGs for face images drawn from a bespoke facial expression parameter space. They found that the discrimination threshold for a small (but suprathreshold) difference in expressivity is lower than the threshold for detecting an expression in a face, and that discrimination thresholds increase with emotivity, consistent with a saturating nonlinearity. By measuring steady-state visually evoked potentials (SSVEPs), the authors further demonstrated that population neuronal responses are consistent with response saturation as facial emotivity increases, though they were unable to confirm the presence of an accelerating regime due to signal-to-noise limitations. That a parameter space as complex as facial expression shows behavioural thresholds consistent with an accelerating and saturating nonlinearity is an exciting finding. However, evidence from more intermediate processing nodes is required to support the ubiquity of an accelerating and saturating nonlinearity in visual processing. One example of a percept dependent on mid-level pooling processes is the global form elicited by Glass patterns.

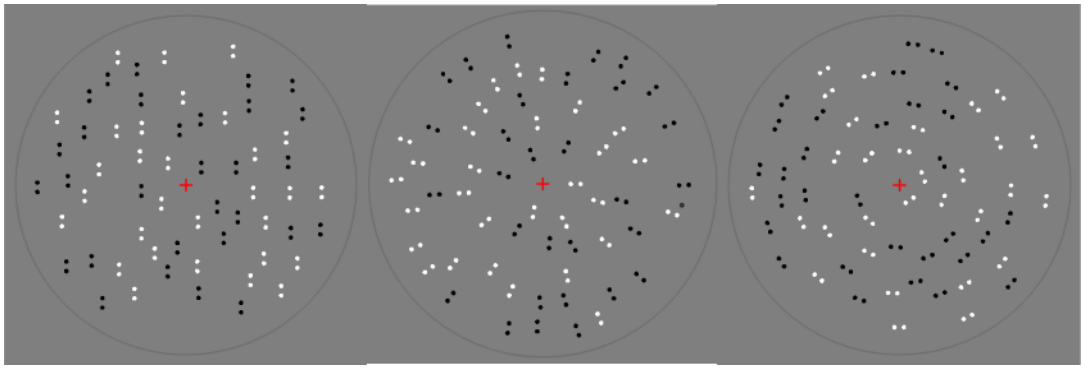

Figure 1.7: Glass patterns across three alignment axes
Three varieties of Glass patterns: translational, radial, and concentric.

A Glass pattern is a type of global form stimulus, initially recognised by Leon Glass (1969). Glass patterns, several forms of which are illustrated in Figure 1.7, can be generated by drawing a field of randomly positioned dots, applying some geometric transformation to their coordinates, and drawing them again. These patterns are remarkable because they are able to provoke a global form percept that emerges from local dot-pair (dipole) orientations, though no individual element of the pattern is innately oriented. Glass patterns have several interesting qualities. First, their perception must rely on at least two integration processes, one that can pool two dots into an oriented dipole, and another that can pool across multiple dipoles to encode information about their wide-field global correlation (H. R. Wilson & Wilkinson, 1998). The percept of a Glass pattern is severely eroded by within-dipole luminance differences (i.e., a dipole consisting of one black and one white dot) but relatively robust to differences across dipoles (Anstis, 1970; Glass & Switkes, 1976; J. A. Wilson et al., 2004). The sensitivity of the global form percept to equiluminance within a dipole is consistent with primate work indicating that dots can be pooled into locally oriented elements ("virtual lines" – (Prazdny, 1986)) by the elongated receptive fields of V1 simple cells, such that two white dots within and parallel to an elongated ON subregion can produce a suprathreshold action potential (Movshon et al., 2010; Smith et al., 2002), while a black and white dot would produce net-zero excitation.

The configuration of the mechanism that integrates across dipoles is still a topic for investigation. For polar form, early (H. R. Wilson & Wilkinson, 1998; H. R. Wilson et al., 1997) and more recent (Lin et al., 2017) models incorporate a global orientation summation stage containing subunits with a preference for orientated elements perpendicular or parallel to its central origin. Wilson and colleagues have

often found that increasing the "signal area" of a pattern (the proportion of the circular aperture containing geometrically correlated dipoles) improves the salience of concentric and radial form, but not the salience of translational form. When interpreted as an absence of spatial summation for translational form, this implies the existence of distinct mid-level circuitry specific to representing polar form. Indeed, there is evidence for neurons with a preference for polar form in extrastriate visual areas such as macaque V4 (Gallant et al., 1996) and for concentric form in the population responses of human lateral occipital cortex (Ostwald et al., 2008), though some global form selectivity can be found as early as V1 (Mannion et al., 2009; Ostwald et al., 2008).

Overall, Glass patterns are an ideal stimulus for probing mid-level local-to-global integration processes. However, most electroencephalography and fMRI experiments (an exception being Mannion et al., (2013)) that investigate Glass patterns use maximally coherent dipole fields, and the majority of psychophysical experiments focus on the detection of Glass patterns in noise. Thus, we have little information on how behavioural sensitivity differs along a gradient of suprathreshold global form intensities. While finding a "dipper effect" is a strong indicator of transducer nonlinearity, it is not a guarantee. As mentioned previously, there are alternative explanations for the dipper effect that do not require nonlinear transduction (Kontsevich et al., 2002; Pelli, 1985; Solomon, 2009), though it is difficult to separate these explanations psychophysically. One approach to establishing the explanatory merit of nonlinear transduction in human observers is to non-invasively measure the transducer (at a population-response level) using fMRI (Boynton et al., 1999) or EEG (Gray et al., 2020), and compare the behavioural sensitivity predicted by the measured transducer with that obtained psychophysically. If both functions predict a similar profile of sensitivity (i.e., the intensity at which thresholds are facilitated), an explanation from nonlinear transduction is strongly supported. The intensity of the global form percept elicited by a Glass pattern can be manipulated in several ways: one can embed Glass pattern dipoles in a field of randomly positioned unpaired dots (Maloney et al., 1987); or manipulate the proportion of geometrically aligned signal dipoles to randomly oriented noise dipoles (H. R. Wilson & Wilkinson, 1998, 2003; H. R. Wilson et al., 1997), or set the orientation of all dipoles according to a Gaussian distribution (Dakin, 1997). These different approaches produce subtly different pattern

fields, and any could be used to assess the discrimination of the global form percept elicited by Glass patterns.

1.2.3 Is there a suprathreshold bias for the perception of polar form?

As previously described, there are a variety of transformations that can be applied to generate Glass patterns. Many experiments focus on horizontal, vertical, radial, and concentric patterns. These transformations are particularly interesting because they are based in two different coordinate systems: horizontal and vertical translations are based on a cartesian reference frame, and only contain a single dipole orientation (when using perfectly aligned patterns) relative to the vertical or horizontal meridian, while radial and concentric patterns are polar forms that can contain any orientation, and are correlated relative to the centre of the pattern field. This distinction is important because encoding polar global form likely requires an extra processing step (as mentioned in the previous section), but also because there is psychophysical evidence to suggest that human observers are better at detecting polar form embedded in fields of randomly oriented dipoles, particularly so for concentric form (Seu & Ferrera, 2001; H. R. Wilson & Wilkinson, 1998; H. R. Wilson et al., 1997). Indeed, there is evidence from magnetoencephalography and EEG showing that maximally aligned radial and concentric forms elicit a higher neuronal response than translational form (Pei et al., 2005; Rampone & Makin, 2020; Swettenham et al., 2010), and that polar form selectivity can be found at extrastriate processing nodes (Gallant et al., 1996; Ostwald et al., 2008). Thus, there is reason to suspect that a unique and more sensitive mechanism may exist for polar form integration, and that increased sensitivity to concentric form does not merely reflect aperture artefacts, as has been previously suggested (Dakin & Bex, 2002).

1.3 Rationale for the current thesis

Above, I have detailed several gaps in the literature pertaining to nonlinear phenomena in visual perception. Each experimental chapter of this thesis (of which there are three) is designed to explore one or more of these gaps in our collective knowledge. In the present section – the final part of my general introduction – the rationale and approach of each experimental chapter is briefly summarised.

First, there are several reports that surround suppression of suprathreshold contrast in central vision increases with age (Karas & McKendrick, 2009, 2011, 2015). Surround suppression of contrast, interpreted as a form of automatic input gain control, acts to keep neuronal responses at the most sensitive region of their response function. If older adults are experiencing a strengthening of surround suppression, an imbalance of excitation and inhibition is implied. It is unclear how this may affect their natural perceptual experiences, but there may be implications for segmentation operations that serve to highlight spatial contrast discontinuities (object boundaries, contours etc.). For this reason, it is important to ensure that the age-related increase in contrast suppression in previous reports truly does reflect surround suppression, and that it is not being contaminated purely by the use of surrounds that directly/closely abut the central stimulus (as is the case in existing reports). Such stimuli may lead to difficulties segmenting the central and surrounding stimulus (Appelbaum et al., 2008), which could be mistaken for "suppression". Although, that the age-effect can still be observed even with out-of-phase surrounds (Karas & McKendrick, 2011) suggests that difficulties related to segmentation are unlikely to be the basis for the effect observed. Alternatively, closely abutting surrounds could lead to contributions from untuned overlay masking (Petrov et al., 2005), a separate (likely pre-cortical) inhibitory mechanism. Unlike surround suppression, overlay masking is generated from within a neuron's receptive field, and strengthening of it may produce perceptual challenges unique to the normalisation of local image contrast. Therefore, the first experimental chapter of my thesis (Chapter 2) is designed to compare surround suppression in younger and older observers, but using stimuli designed to attenuate contributions from edge effects and overlay masking. In this chapter, we have also assessed the orientation tuning of suprathreshold surround suppression at central fixation, as previous investigations with the same aim have also used stimuli vulnerable to contributions from overlay masking.

Second, our understanding of expansive and compressive nonlinearities (and their behavioural consequences in visual perception) is mostly based in low-level visual perception, though there are proposals that they may exist throughout the visual system as "canonical" computations (Carandini & Heeger, 2011). Recently, this proposal has gained some traction, as evidence for a dipper effect and a saturating nonlinearity has been found in the perception of human facial expression (Gray et al.,

2020). However, the extent to which this finding persists at intermediate (mid-level) forms of visual processing is unclear, suggesting that behavioural and electrophysiological explorations of global form perception would be fruitful. Exploring the possibility of a sigmoidal transducer in the perception of global form would add to our understanding of how common such nonlinearities are throughout the visual processing hierarchy. Additionally, we have observed that reports of increased sensitivity to polar global form (radial and concentric Glass patterns) are limited to psychophysically assessed global form detection thresholds (Seu & Ferrera, 2001; H. R. Wilson & Wilkinson, 1998; H. R. Wilson et al., 1997) and neuronal responses to maximally aligned Glass patterns (Ostwald et al., 2008; Pei et al., 2005; Rampone & Makin, 2020; Swettenham et al., 2010). Suprathreshold sensitivity at intermediate global form coherences is scarcely investigated. Thus, it is unknown whether the polar form biases at detection and maximal coherence can also be found along a gradient of suprathreshold global form intensities. Assuming the presence of a sigmoidal nonlinearity: it is possible that the shape of the transducer (or TvI curve) varies as a function pattern axis. As such, the 2nd and 3rd experimental chapters of this thesis focus on measuring the profile of suprathreshold sensitivity for the global form percept elicited by translational, radial, and concentric Glass patterns. In Chapter 3, we present the results of an online pedestal experiment designed to measure Glass pattern global form discrimination thresholds. To predict the neuronal response function that would generate the thresholds we observed, we fit them with a differentiated neuronal transducer function. In Chapter 4, we compared these model predictions with steady-state visually evoked potentials (SSVEPs) along the same global form continuum that we tested psychophysically.

Chapter 2. Measuring the effects of age on contrast suppression

2.1 Abstract

The perceived contrast of a suprathreshold central stimulus can be reduced by the presence of a surrounding stimulus. It has been suggested that this this effect increases with age when measured using foveal stimuli. The underlying mechanism proposed for this age dependence (a general change in the balance of inhibition and excitation in cortex) makes this psychophysical phenomenon potentially interesting as a biomarker of neurological dysfunction. In this chapter, we attempt to repeat these measurements using stimuli that are designed to eliminate potential confounds that were present in the early reports. Principally, we control for untuned 'overlay' masking which is thought to have a pre-cortical origin. We measured contrast matching thresholds in twenty younger (< 30) and seventeen older (>60) observers. Across all observers, we find suppression that has little or no orientation tuning and, importantly, no effect of age. Our findings contradict those from earlier studies and suggest that effects relating to age may arise from an 'overlay masking' mechanism that could originate as early as the retina.

2.2 Introduction

The average human lifespan is increasing in the United Kingdom, and is projected to be above 90 years for new-borns by 2045 (Office for National Statistics, 2022). This trend calls for an increase in our understanding of the unique perceptual challenges faced by the elderly. In visual perception, age-related changes in visual acuity and contrast sensitivity are mostly (though not entirely) driven by changes to the optical media (Owsley, 2011) in healthy subjects. However, there is clear evidence that the human visual cortex is affected by structural and functional senescence (Brewer & Barton, 2012), and investigations of post-retinal visual processing across the lifespan may prove useful in understanding how best to prevent or palliate age-related changes in visual perception. One line of enquiry could be the perceptual normalisation processes that rely on the delicate balance between neuronal excitation and inhibition. In visual perception, normalisation is thought to shift a neuron's dynamic range along the input axis to minimise output saturation, effectively

preserving sensitivity over a wide range of inputs (i.e., a range of luminance contrasts). It has been previously suggested that response normalisation is a "canonical" computation – a fundamental neural operation that is enacted throughout the subdivisions of the brain (Carandini & Heeger, 2011). If these mechanisms de-sensitise or diminish with age, it could produce age-related changes in visual perception, but also other sensory modalities, and perhaps memory and cognition. Therefore, measuring the action of these mechanisms using relatively simple visual experiments can inform our understanding of the general effect of age on the human brain. In visual perception, neuronal populations that contribute to the regulation of neuronal responses are though to exist throughout the low and mid-level visual processing hierarchy (Angelucci et al., 2017), and have been shown to become less effective in aging monkeys (Fu et al., 2010; Schmolesky et al., 2000).

In humans, a plausible example of neuronal inhibition comes from motion direction perception. Tadin et al. (2003) asked observers to judge the direction of a moving stimulus, bur parametrically varied its size. They reported that larger stimuli made this task more difficult, likely due to an inhibitory mechanism that is dependent on the spatial extent of the stimulus. Interestingly, later work from Betts et al. (2005) suggests that the strength of this spatial antagonism may reduce with age, finding that older observers were better than younger observers at judging the direction of large stimuli. Though this interpretation is not universally accepted (Hutchinson et al., 2014), it has encouraged further investigation of the relationship between spatial inhibition and advancing age. Examples of spatial inhibition are not limited to motion perception. When viewing a texture of set luminance contrast surrounded by a texture of greater contrast, the apparent contrast of the central texture is reduced (Cannon & Fullenkamp, 1991; Chubb et al., 1989; Petrov & McKee, 2006; Xing & Heeger, 2001). This effect ('surround suppression') appears to be physiologically distinct from a contrast suppression that is found when two stimuli are overlaid (Petrov et al., 2005): it operates over long ranges (i.e., neighbouring receptive fields) and is tuned for orientation and spatial frequency – suggestive of an origin in the elongated receptive fields found in the primary visual cortex. In comparison, overlay masking appears to be short range and untuned and could arise as early as the centre-surround receptive fields of the retinal ganglion cells. Later work has shown that surround suppression may be the result of at least two serial mechanisms with different tuning properties

(Petrov & McKee, 2009; Webb et al., 2005). When measuring contrast detection thresholds, most labs find that surround suppression is weak or effectively absent at fixation (Meese & Hess, 2004; Petrov et al., 2005; Saarela & Herzog, 2008), and some find it to increases in strength into and beyond the perifovea (Petrov et al., 2005; Snowden & Hammett, 1998).

Over the past decade, work from one lab has investigated whether surround suppression of contrast is subject to the same age-related weakening reported by Betts and colleagues' in perception of motion. Surprisingly, initial work using a version of the Chubb illusion (Chubb et al., 1989) suggested that surround suppression strengthened with age (Karas & McKendrick, 2009). This pattern of results was has also been replicated using drifting luminance gratings (Karas & McKendrick, 2012), and static gratings (Karas & McKendrick, 2011), with the latter work finding the age-effect to be present even when the centre and surround are 180° out of phase. Probing a subset of spatiotemporal tuning characteristics, Karas & McKendrick (2015) reported that older adults experience enhanced suppression primarily a low contrasts (centres of 20% contrast), that surrounds at the same or twice the contrast produce the greatest age-related difference in suppression strength. The same study also found the age difference in surround suppression to be most profound with shorter stimulus presentation times of 150ms, but still observable at durations of 500ms. Though later work has indicated that very short presentation times (~40ms) abolish the age-related difference (Pitchaimuthu, Nguyen, et al., 2017). Overall, McKendrick and colleagues' work has often found surround suppression to strengthen with age, despite previous evidence suggesting a weakening of spatial inhibition in older adults (Betts et al., 2005).

Surround suppression specifically describes the reduction of perceived contrast produced by presenting a target stimulus in the context of a non-overlapping surround, but this is not the only route to contrast masking. As previously described, overlay masking produces a similar effect but is instead achieved by super-imposing a masking stimulus over the target. The aforementioned differences in the spatiotemporal tuning characteristics of overlay and surround suppression (see also section 1.1.3), as well as the absence of long-range pre-cortical receptive fields (Bonin et al., 2005) suggests that they are based on entirely different mechanisms. Importantly, when Petrov and

colleagues (2005) measured the distinct spatial tuning properties of surround and overlay masking, they included stimulus features to reduce cross-contamination between these two suppression types. Primarily, they introduced a gap between the central 'probe' region and the surround (equal to one stimulus wavelength) to reduce the effect of overlay masking when using surround stimuli. This manipulation reduces the occurrence of neuronal receptive fields that are partially subtended both by the centre and surround, as such neurons would effectively be experiencing overlay masking. It also prevents the intrusion of difficulties related to the segmentation of the centre from the surround (Appelbaum et al., 2008). Secondly, though less relevant to McKendrick's suprathreshold experiments, to reduce contamination from spatial uncertainty they ensured that the position of the target probe was constantly cued by the presence of a faint, grey ring. This low contrast cue was too subtle to contribute meaningfully to a contrast gain control mechanism. At detection, the presence of a surround (which provides an unambiguous cue to location) may appear to increase probe contrast, as has been shown using collinear flankers (Petrov et al., 2006). By removing these confounds, the authors effectively isolated and characterised two distinct mechanisms of contrast gain control: an initial spatially untuned overlay masking and a later spatially tuned surround suppression. It is possible that the absence of these controls in stimuli used by other groups (principally the lack of a one cycle gap) may explain some of the paradoxical effects that they observe. McKendrick and colleagues typically present centres and surrounds that directly/closely abut each other, and have no spatial cues to target location.

The present experiment explores the possibility that the increase in surround suppression found by previous reports is due to such stimulus confounds. Using a contrast matching experiment, we psychophysically assessed the strength of surround suppression in younger (< 30) and older (> 60) observers. Our stimuli were similar to those used in previous similar experiments: luminance contrast gratings centred on the fovea, but with the addition of a mean luminance gap spanning 1-cycle (1λ) of the grating's spatial frequency to preclude any possible contributions from overlay masking. If surround suppression alone is the main driver of the age-related changes to contrast suppression in central vision, then we expected to find a significant difference in suppression strength between age-groups despite this manipulation. We additionally assessed the orientation tuning of surround suppression, as the spatial

tuning of suprathreshold surround suppression in central vision has not been investigated with stimuli such as ours, and there is evidence from macaque electrophysiology that the orientation tuning of surround suppression broadens with age (Fu et al., 2010).

2.3 Methods

2.3.1 Participants

Seventeen older observers (mean: 69, range: 60 – 81) and twenty younger observers (mean: 18.8, range: 18 – 22) with self-reported normal or corrected-to-normal visual acuity and no personal history of neurological disease or disorder were recruited. This age-range was selected to mirror that of previous studies that have found the effect of age on surround suppression (Karas & McKendrick, 2011, 2015). All observers were aware of the purpose of the study, but not aware of any of our predictions. Ethical approval for the study was given by the Department of Psychology at the University of York, and all observers were reimbursed for their time and travel costs.

2.3.2 Stimuli & apparatus

The spatial parameters of the stimuli used are illustrated in Figure 2.1. Stimuli frames consisted of centrally presented luminance gratings and annuli with a spatial frequency of 3.33 cycles per degree. Central gratings always subtended 0.6 degrees of visual angle (meaning there were precisely 2 cycles of luminance modulation). Central gratings were surrounded by a thin (2 pixel) grey ring to prevent uncertainty pertaining to the location of the target stimulus, which is known to bias detection thresholds as previously described. When presented, the surrounding annulus subtended 3.2 degrees of visual angle. The outer and inner edges of all stimuli were smoothed by a raised-cosine mask, and the inner and outer edges of the raised cosine plateau were separated by 0.3 degrees of mean luminance (~1λ). Stimuli were displayed on a ViewPixx 3D Lite display (1920x1080, 120Hz) (https://vpixx.com/products/viewpixx-3d/), via an Apple Mac Pro 6.1 running macOS High Sierra (version 10.13.6). The display used 8-bits of grayscale resolution and was gamma corrected using a Minolta LS110 photometer. The mean luminance of the display after correction was approximately 50cd/cm^2. All stimuli were created and presented using the Python programming language (https://www.python.org/) and

PsychoPy3 (https://www.psychopy.org/). Observers were seated 1 metre away from the display and used a chin rest while the experiment was in progress. Despite wearing correction, two observers in the older age-group had difficulty seeing the thin grey ring at a distance of one metre, and were moved closer to the display (~0.68m), but all stimuli were temporarily rescaled to accommodate this.

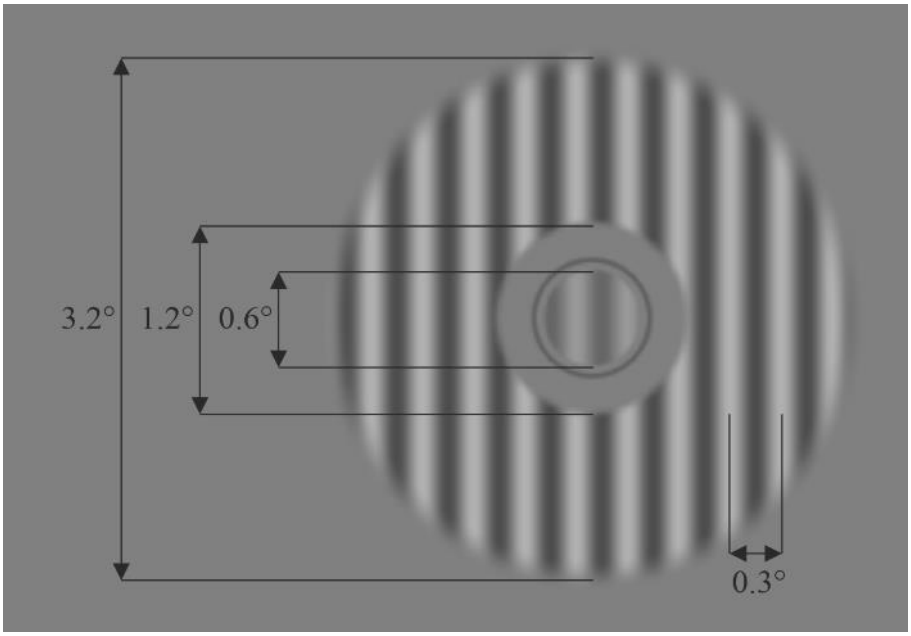

Figure 2.1: Centre and surround spatial configuration
Spatial definition of the stimuli used for our contrast matching experiment. The same central grating and thin ring were used when measuring contrast detection thresholds.

2.3.3 Procedure

First, to ensure all observers would be able to perceive the stimuli used in the contrast matching task, we measured contrast detection thresholds using the method of adjustment. The method of adjustment is not as accurate as a staircase procedure, but is sufficient for a quick and coarse guarantee that all observers should be able to perform the main experiment reliably. To obtain detection thresholds, observers were asked to manipulate the contrast of a centrally presented grating using a trackball mouse. This grating shared the spatial parameters of the central grating used in the contrast matching procedure (shown in Figure 2.1). Rolling the trackball away from them increased the contrast, while rolling it towards them decreased the contrast in equal linear increments. Observers were given as much time as they needed to adjust the contrast of the grating until it was just barely visible, which they indicated by a keypress. While observers were adjusting the contrast, the grating slowly cycled on

and off (0.15s ON : 2.15s OFF), and its orientation was randomised on every cycle. This procedure was repeated six times, with the median of these six thresholds taken as the observers absolute contrast detection threshold.

To assess the strength of surround suppression across different surround configurations, observers took part in a "1-up 1-down" staircase procedure designed to measure the point of subjective equivalence (PSE) between the perceived contrast of an unmasked reference grating and a surround-masked target grating. Observers viewed target and reference gratings sequentially via a two-interval forced-choice (2IFC) procedure, and their task was to indicate the interval (1^{st} or 2^{nd}) that contained the central grating of highest contrast. Many observers in the older age-group would be completely unfamiliar with visual psychophysics, so we wanted to ensure that the experiment was as simple as possible. For this reason, the reference (unmasked) grating was always presented in the first interval (at 20% contrast), and the masked target grating was always presented in the second interval, with a contrast dictated by the staircase procedure over the course of 70 trials. Note, the correct response could be in either interval, as the contrast of the target grating can be greater or less than that of the reference grating. The target either had no mask (a control condition), a collinear mask, or an orthogonal mask, and surrounds were fixed at 40% contrast. The phase and orientation of presented stimuli was randomised such that the reference grating in interval one had a different phase and orientation to the grating and annulus presented in interval two. This was done to ensure observers were only making responses based on the perceived contrast, and to prevent adaptation to the spatial parameters of the stimulus. As shown in Figure 2.2, a single trial consisted of the reference grating for 150ms, followed by an inter-stimulus-interval of 500ms. Observers then saw 500ms of surround, followed by 150ms of target grating and surround, and a final 500ms of surround alone. We included this 500ms onset delay for the target grating to prevent any contribution from temporal onset masking. Observers were then given unlimited time to indicate (via a button-press) the interval that contained the central grating of highest contrast. The first 5 trials of the staircase procedure had a step-size of 2.25% contrast, but subsequent trials were reduced to a step-size of 0.75% to approximate the observer's true PSE more closely. Each of the surround-on conditions were repeated up to 3 times by each observer. If the PSE estimates of the 1^{st} and 2^{nd} run (based on a least-squares Weibull fit) differed by less than the minimum staircase step-size

(0.75%), the final repetition was skipped. This was done to avoid unnecessarily straining attentive observers. The no-surround control condition was always conducted first, and only once. For each observer, this resulted in 5 – 7 completed staircases. To give observers the opportunity to obtain a stable response criterion prior to the main experiment, each practised with an abbreviated version of the experiment containing one surround and one no-surround condition, each with 25 trials of the staircase procedure. Additional practice runs were permitted if necessary.

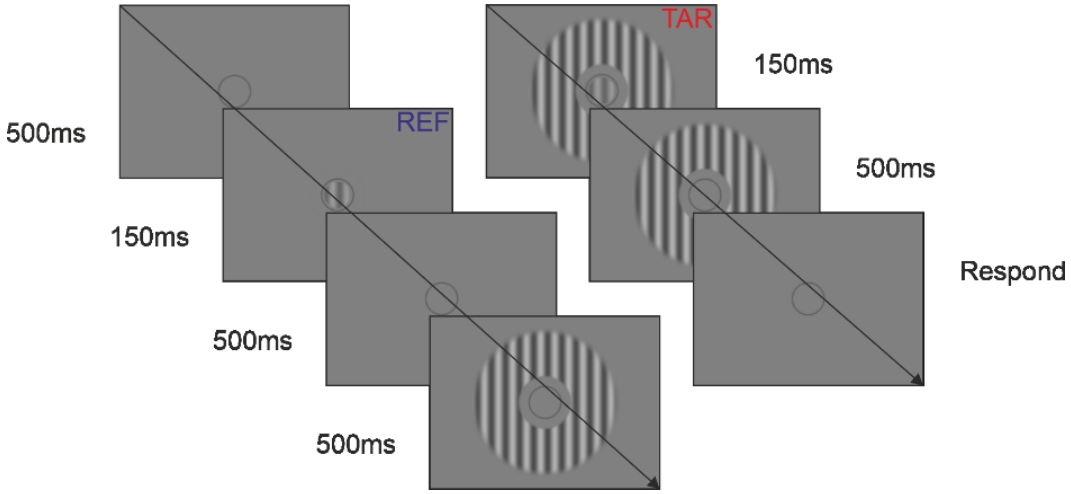

Figure 2.2: Contrast matching procedure
Timing of stimulus presentation. To be read from top-left to bottom-right, following diagonal arrows. The reference interval is represented by the blue "REF" text, while the target interval is represented by the red "TAR" text.

2.3.4 Model fitting and statistical analysis

Observers' final PSEs were estimated by fitting a logistic cumulative density function (CDF) to accuracy estimates. This was accomplished using the maximum likelihood fitting routines of the Palamedes Toolbox (http://www.palamedestoolbox.org/) via MATLAB R2018b. The threshold and slope of this psychometric function (PF) were allowed to vary as free parameters, while the guess-rate and lapse-rate were both fixed at 0%, such that the threshold parameter equates the contrast achieving an accuracy of 50% (the PSE). For each observer, and each condition, this fitting procedure was carried out on the accumulated accuracies from all staircase repetitions by summing the number of correct responses and total trials. The effect of age-group and surround configuration was investigated via a two-way mixed-model analysis of variance (ANOVA), with age-group (younger, older) used as the between subjects factor, and surround condition (no surround, collinear,

orthogonal) used as the within subjects factor. In younger observers, for the orthogonal condition, there was a violation of the assumption of normality of residuals (Shapiro-Wilk, $p < .001$) due to the presence of an outlier - how we dealt with this datapoint is detailed in the results section. This dataset satisfied Levene's test for the homogeneity of variance ($p > .05$), but failed Mauchly's test of sphericity ($\chi^2(2) = 29.03$, $p < .001$), so the Greenhouse-Geisser correction was applied. For completeness, we also report a comparison between younger and older observers' contrast detection thresholds. Due to the presence of two outliers (beyond 1st and 3rd quartiles -/+ 1.5 * IQR), this comparison was made using an independent samples Mann-Whitney U test (comparing medians). All statistical analyses were carried out using SPSS (https://www.ibm.com/products/spss-statistics).

2.4 Results

2.4.1 Elevated contrast detection thresholds in older observers

First, to demonstrate that all observers were able to perceive the suprathreshold stimulus, we report their contrast detection thresholds. Absolute detection thresholds are shown in Figure 2.3, and are all well below the contrast of our matching stimulus (20%). Though, there are two high-threshold outliers in the older age-group, and a single marginal outlier in the younger age-group. It is unclear whether these reflect genuine visual deficit, or whether they a misapprehension/artefact of the method of adjustment. Nevertheless, we have no reason to suspect that these data-points are artefactual, so we retained their thresholds and conducted a non-parametric comparison. We also retained these observers' supra-threshold contrast matching data, as the reference grating would still have been approximately twice their detection threshold. There is there is a clear trend towards younger observers having lower detection thresholds (median = 2.02%, IQR = .60%), than older observers (median = 3.55%, IQR = 1.74%), a difference in medians of 1.53% that was significant according to a Mann-Whitney U test (U = 28, z = -4.552, $p < .001$). This age-related increase in contrast detection thresholds is consistent with existing literature (Beard et al., 1994; Elliott, 1987; Mei et al., 2007; Owsley, 1983). Having confirmed that all observers could readily perceive our suprathreshold stimulus, we will now report the results of the contrast matching experiment.

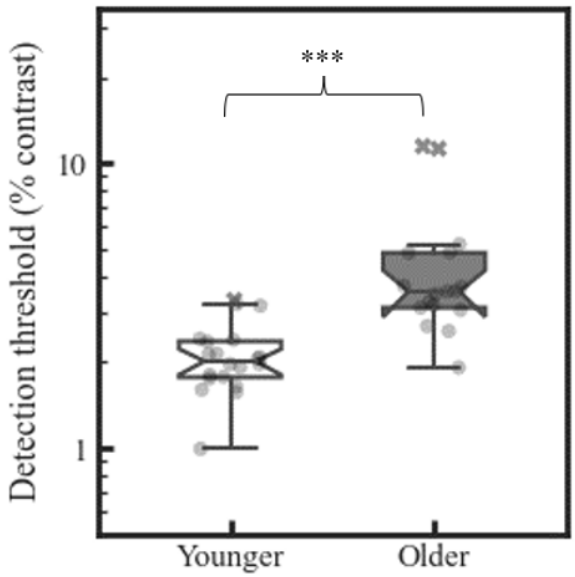

Figure 2.3: Detection thresholds compared across age-groups
*Contrast detection thresholds in younger and older adults. Represented as boxplots due to the use of non-parametric statistics to compare medians. Grey circles represent individual data points, crosses represent outliers. *** = p<.001. The upper whiskers of each boxplot extend to the last data-point less than the 3rd quartile + 1.5* the interquartile range (IQR), and the lower whiskers extend to the last data-point more than the 1st quartile – 1.5*IQR.*

2.4.2 Surround suppression does not change with age

In Figure 2.4, we present example psychometric function fits from a younger and older observer. Both have discrimination accuracy estimates that are well described by a logistic CDF, and the same is true of most observers (see Appendix A for all observers' fits), though two had chance-level accuracy at high target contrasts that were not consistent with their performance at lower contrasts. In these observers (highlighted in Appendix A), this discontinuity in performance was driven by several target intensities with only two trials, likely reflecting a brief excursion of the staircase into high target contrasts due to a series of finger errors. For this reason, we removed any target intensities with fewer than three trials prior to model fitting for these two observers.

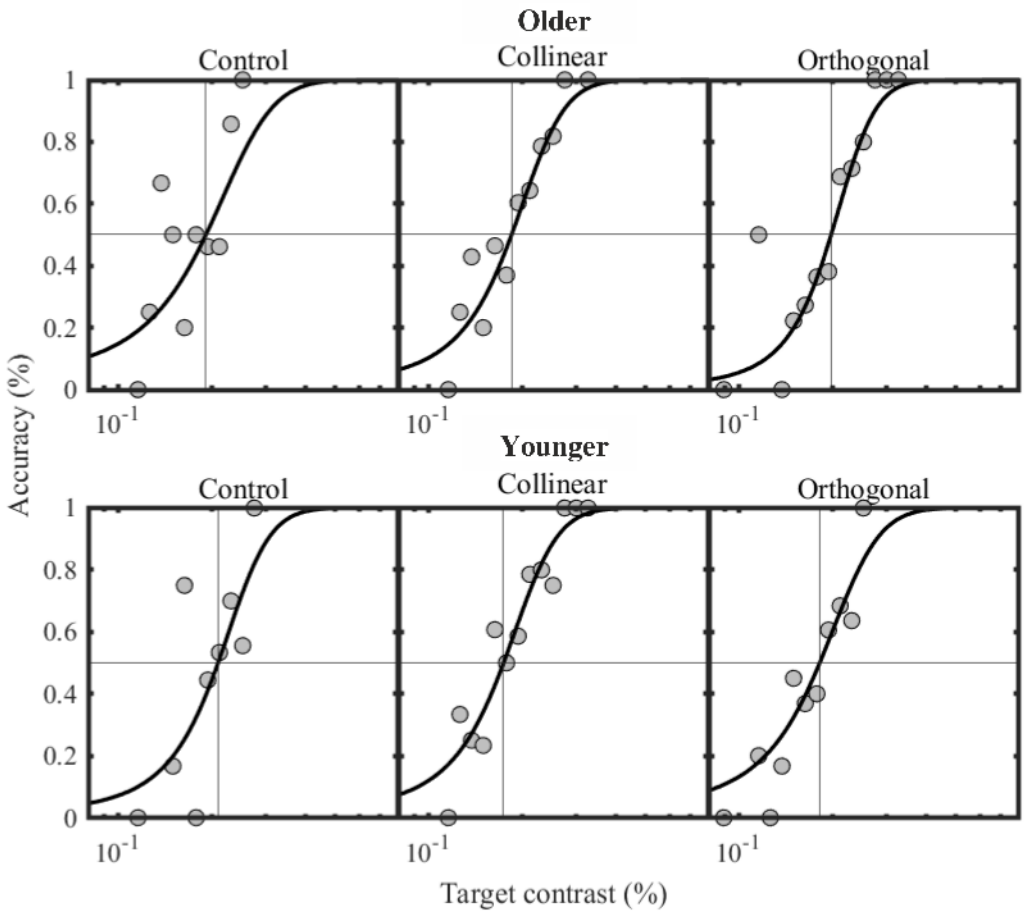

Figure 2.4: Example psychometric function fits
Psychometric function fits from two example observers in the three matching condition we used. Where the vertical line of the crosshair intercepts the abscissa denotes the point of subjective equality that we use for group-level comparisons.

In Figure 2.5, the mean of PSEs estimated by psychometric function fits for each surround condition are presented split by age-group. We analysed observers' PSEs with a mixed-model ANOVA, but examination of studentised model residuals (via quantile-quantile plots) showed that there were two clear outliers (a studentised residual of ± 3), one in the no-surround condition (older), and one in the orthogonal condition (younger). Further inspection revealed that the generative psychometric functions for these thresholds had sensible fits (these are highlighted in Appendix A). As we had no reason to suspect these data-points were artefactual, we elected to run the ANOVA with and without their data to examine whether the statistical outcome was contingent on their contribution. The interpretation of all p-values was the same in either case, so we elected to interpret results based on the full-dataset, but provide ANOVA tables both with and without outliers in Appendix B. In Figure 2.5, older and younger observers appear to have quite similar PSEs in all surround conditions, though

there is some indication that younger observers PSEs may be more sensitive to the orientation of the surround. Statistical analysis showed that the within-subjects comparison of surround condition was significant (F(1.280,42.240) = 26.545, p < = .001), while the between subjects comparison of age-group was non-significant (F(1,33) = .060, p = .808), and there was a non-significant interaction between age-group and surround condition (F(1.280,47.240) = .273, p = .662). These results indicate that PSEs were contingent on the configuration and/or presence of the surround, but do not support the presence of a broad age-related effect on surround suppression, nor any surround-configuration specific effect of age (i.e., due to differences in orientation tuning across age-groups, which has been found in monkey electrophysiology (Schmolesky et al., 2000)). This pattern of results is better conveyed in Figure 2.5B, where we have plotted surround-on PSEs as multiples of observers' surround-off PSEs, emphasising within-subjects differences relative to the surround-off control. Here, we can see that 95% confidence intervals on the mean (generated via bootstrapping with replacement) do not cross 1 for both surround configurations for both age-groups (indicative of a significant PSE reduction from control), but are almost entirely overlapping within each surround condition, implying a similar degree of surround suppression across age-groups, in disagreement with previous findings of increased central supra-threshold surround suppression in the elderly (Karas & McKendrick, 2009, 2011, 2012, 2015).

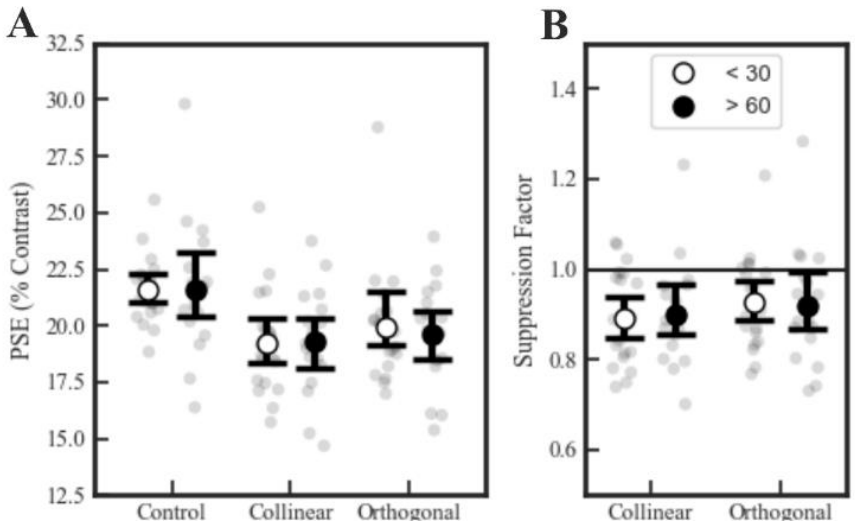

Figure 2.5: Points of subjective equivalence split by age-group

A: mean of PSEs split by age group & surround condition. Larger open circles represent the within condition mean of young participants, larger closed circles represent the same for older participants. Error bars represent 95% confidence intervals. B: Same format as A, but showing suppression rations. Thresholds below the solid reference line are indicative of suppression.

As the within-subjects effect of surround condition was significant, we conducted post-hoc paired t-tests across all observers, collapsed across age-groups (Figure 2.6A). Observers' PSEs were, on average, highest in the no-mask control condition (mean = 21.5%, 95%CI =[20.8, 22.3]), lowest in the collinear mask condition (mean = 19.2%, 95% CI [18.6, 19.9]) and intermediate in the orthogonal mask condition (mean = 19.7%; 95CI = [19.1, 20.6]). The mean reduction in PSEs between the control and collinear conditions was significant (mean Δ = -2.3%, 95% CI: [1%, 3%] $p < .001$), as was the reduction in PSEs between control and orthogonal conditions (mean Δ = -1.8%, 95% CI [.09%, 2.6%] $p < .001$), and the difference between orthogonal and collinear surround PSEs (mean Δ = 0.5%, [.03%, 1.3%], $p < .05$), though the latter was very small - less than the smallest step-size of our staircase procedure. As before, these PSEs are replotted as suppression factors in Figure 2.6B, where, on average, collinear mask PSEs were 90% (95% CI = [86, 93]) of the no mask control, increasing to 92% (95% CI = [89, 96]) with an orthogonal mask, representing a mild (20%) reduction in masking strength. Unlike Xing and Heeger (2000), who found suppression to be strongly attenuated when using orthogonal surrounds.

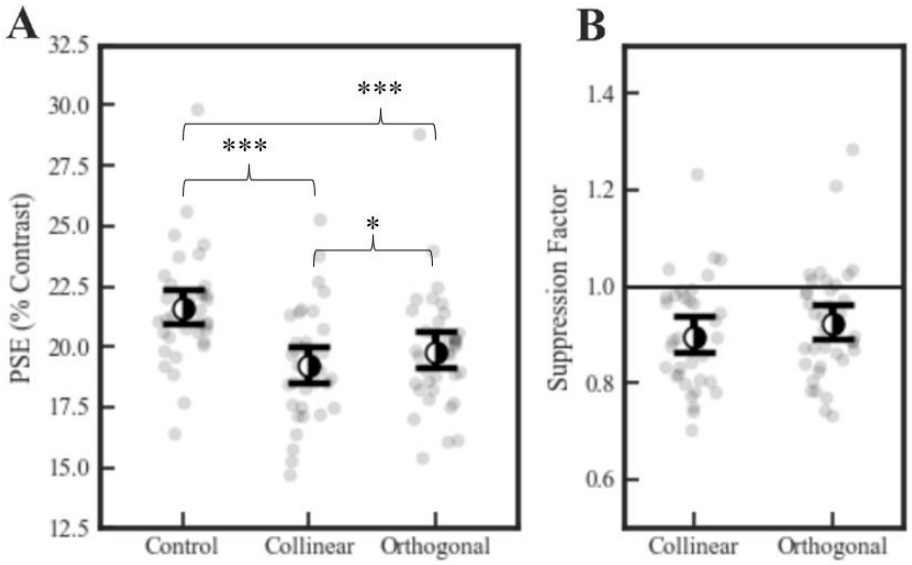

Figure 2.6: Points of subjective equivalence collapsed across all observers
*A: mean of PSEs collapsed across age groups, split by surround condition. Half-filled circles represent the within-condition mean across all observers. Smaller grey dots represent individual data points. Error bars represent 95% BCa CIs. B: mean of suppression factors, collapsed across age-groups. As before, points below the horizontal reference line are suggestive of suppression. * = p<.05, *** = p<.001.*

2.5 Discussion

Here, we asked whether age alters a fundamental component of early visual processing in humans: surround suppression. We used stimuli contrasts based on the work of Karas and McKendrick (2015), which they reported to elicit the strongest age-related difference. However, to minimise edge effects and overlay masking, we included a 1 stimulus cycle gap between the centre and surround (Petrov et al., 2005). Although we did observe some surround suppression, it was relatively weak and broadly spatially tuned compared to the effect routinely found in the periphery both at absolute detection and suprathreshold (Petrov et al., 2005; Shushruth et al., 2013; Xing & Heeger, 2000). Contrary to previous publications, we found no evidence for an age-related change in surround suppression. We will now discuss the meaning of our findings in the context of previous reports of surround suppression strengthening with age, and where our broadly tuned masking may fit into the existing literature on suprathreshold surround suppression in central vision.

2.5.1 An alternative explanation for age effects in contrast suppression

Previous work has indicated that surround suppression of contrast at suprathreshold intensities increases in magnitude with age (Karas & McKendrick, 2009, 2011, 2012). We failed to replicate this finding, and propose that the age-related alteration of a different suppressive mechanism (such as overlay masking) could explain the difference between our results. The idea that surround suppression is not the only explanation is compatible with several of McKendrick and colleagues current findings. First, their recent work has found that the age-related increase in central-field suppression is abolished by interocular presentation of the centre and surround (Pitchaimuthu, Nguyen, et al., 2017). This indicates that the inhibitory circuitry affected by age exists prior to the fusion of the monocular visual streams in V1. McKendrick an colleagues suggest that the mechanism affected by age could be a separate, ‘early’ surround-suppression component, such as that proposed by Webb et al (2005) (based on macaque V1 electrophysiology) which is more apparent at low contrasts, where the age effect has been shown to be most profound in humans (Karas & McKendrick, 2015). However, the same pattern of results could just as easily be explained by overlay masking, as it too has a broadly tuned monocular component (Baker, Meese, & Summers, 2007; Bonds, 1989; DeAngelis et al., 1992; Petrov et al.,

2005), and has been shown to occur prior to orientation tuned surround suppression (Petrov et al., 2005). Though, it is important to note that these findings for overlay suppression are at detection threshold, and it is uncertain how far they extend to suprathreshold contrast perception. Moreover, Nguyen and McKendrick (2016) demonstrated that a very narrow surrounding annulus presented in close proximity to a central grating (a gap of ~λ/3) is sufficient to produce increased suppression in the elderly – approximately 69% of the age-effect they reported when using spatially extensive directly abutting surrounds at the same contrasts in a previous experiment (Karas & McKendrick, 2015). This implies that the wide-field extent of the surround is of less importance relative to the inner boundary, and indeed, they found the inverse age-effect in the periphery when they happened to use an intervening gap in excess of 1 stimulus cycle. While this was interpreted as a foveal bias in their observed age-effect, any contribution from overlay suppression would likely have been severely reduced in their peripheral condition.

A competing explanation for the absence of an age-related strengthening of surround suppression in our results could be based on the timing of our stimuli. Our central targets were presented for 150ms, a duration shown to produce the age-effect previously (Nguyen & McKendrick, 2016). However, the onset of our surrounding annulus was not synchronous with the onset of the central target. In our experiment, surrounds appeared 500ms before the target, and were sustained until 500ms after the offset of the target, whereas McKendrick and colleagues always use targets and masks that are synchronous. When a surrounding mask's onset and offset occurs before or after the onset/offset of a central probe, suppression can still be produced (this is often referred to as forward and backwards masking). This masking is transient: relative to the central probe, suppression is much weaker when the surround's whole duration is offset by as little as 20ms (Petrov & McKee, 2009). However, this weak suppression is not to be confused with the suppression we have measured, as our surrounding masks still coincided with the timing of the central target (see Figure 2.2). This is an important distinction, as it means that we have measured the cumulative suppression of forward, simultaneous, and backwards masking in our observers. Could this explain our lack of an effect of age on surround suppression? Previous M/EEG experiments (Ohtani et al., 2002; Schallmo et al., 2018), have presented surrounds up to 2s before and after the target (to measure the separate responses to the centre and surround), and

still measure significant suppression of neuronal responses. In a control condition, Ohtani and colleagues even compared the suppression elicited by simultaneous and sustained early-onset surrounds, finding similar suppression in both cases. Moreover, in a psychophysical experiment, Schallmo and Murray (2016) presented surrounds 1-2s in advance of a probe until 500ms after the probe's offset, and still reported surround suppression at suprathreshold contrasts. Schallmo and colleagues' stimuli were presented in the perifovea (5.3°), as were those of Ohtani et al., so there remains the possibility that this pattern of results may change at fixation. Nevertheless, if the effect of age on surround suppression in central vision requires the simultaneous onset and/or offset of both the centre and surround, our stimuli would indeed have escaped it. As stronger surround suppression with age has been identified using centres and surround presented for as long as 500ms (Karas & McKendrick, 2015), we think this is unlikely, but investigating the temporal dependencies of this effect would be interesting line for enquiry.

Our findings of no effect of age on supra-threshold surround suppression in central vision coupled with alternative interpretations of McKendrick and colleagues' findings suggest that further investigation to identify the unique contributions of overlay masking and surround suppression to age-related changes in contrast suppression would be fruitful. Although, work from the same lab recently did not to find the age effect even when using directly abutting surrounds (Nguyen et al., 2020), suggesting that a general contribution from overlay masking may not be the sole explanation for the differences in our findings. In this experiment, the authors used stimuli at lower spatial frequencies than usual (1cpd), where a sizeable proportion of the surrounding annulus would be within a stimulus wavelength of the central target. That they did not find the effect of age in this case suggests that there may be some selectivity in the mechanism effected by age, such that it is absent at low spatial frequencies. If we accept that surround suppression changes in spatial tuning as a function of absolute stimulus contrast (Levitt & Lund, 1997; Webb et al., 2005), then a different uppressivee mechanisms may be dominant at low spatial frequencies which may escape the effect of age. Some temporal specificity is also implied, such that the increase in suppression with age is present for intermediate durations of 150 – 200ms (Karas & McKendrick, 2015; Pitchaimuthu, Nguyen, et al., 2017), but severely diminished at longer durations of 500ms (Karas & McKendrick, 2015) and

totally absent shorter durations of 40ms (Pitchaimuthu, Nguyen, et al., 2017). Indeed, that such minor alterations to the stimulus (including our own) abolish the age-effect suggests that it is generated by a suppressive mechanism (or a subset of mechanisms) that is spatially and temporally narrow-band, rather than a general feature of contrast gain-control processes throughout the visual system.

2.5.2 Broadly tuned central supra-threshold surround suppression

Surround suppression is an umbrella term for a number of distinct neural mechanisms with distinct spatiotemporal tunings. At the centre of the visual field, surround suppression is mostly absent at detection threshold (Kim et al., 2010; Petrov et al., 2005; Saarela & Herzog, 2008), but is present at supra-threshold intensities (Vanegas et al., 2015) and found to be orientation tuned (Cannon & Fullenkamp, 1991; Xing & Heeger, 2000). In fact, Xing & Heeger found an orthogonal surround to abolish suppression almost entirely when using a central stimulus. We too found significant supra-threshold suppression, but only found a small (thought significant) difference between collinear and orthogonal suppression indices. What might explain the broader tuning we have encountered? It is difficult to draw comparisons between Xing and Heeger's findings and our own, as they (like many others), use almost directly abutting surrounds, so it is impossible to fully dissociate contributions from overlay masking and surround suppression in their findings. Nevertheless, one explanation for our observation of broad spatial tuning is a difference the spatiotemporal parameters of our stimuli. Xing and Heeger used gratings that contrast-reversed at 8Hz while we used static gratings that were presented for 150ms, and which were therefore relatively long-lived. It is possible that Xing and Heeger's stimuli were detected by neurons tuned to higher temporal frequencies than the neurons subject to the suppressive mechanism we have observed. At detection threshold in the periphery, transient and sustained surround suppression are both found to be orientation tuned (Petrov & McKee, 2009), but it may be that not all mechanisms contributing to surround suppression in central vision share these tuning properties. In other words, it is possible that transient suprathreshold surround suppression is sharply orientation tuned in central vision, while the more sustained suppression that we may have isolated is not. Furthermore, when measuring orientation tuning, Xing and Heeger used much larger stimuli than we did (both centres and surrounds), at just over half our spatial frequency (we used 3.33cpd, they used 2cpd). Their surrounds

subtended 11° of visual angle, while ours only subtended 3.2°, so it is plausible that their stimuli recruited longer-range suppression that may be distinct in its tuning characteristics, although this would be counter to evidence suggesting that long-range surround suppression is less orientation tuned than the shorter-range suppression that we may have measured (Angelucci et al., 2017; Nurminen et al., 2010).

If we accept that Xing and Heeger's (2000) results may be explained via a difference in the spatiotemporal tuning of the suppressive mechanisms activated by our respective stimuli, and if we discount any intrusion of overlay masking into their design, then our findings provoke further exploration of how stimulus transience and spatial location interact to determine the strength of surround suppression. In human observers, previous work using stimuli presented at the near-periphery has identified that transient and sustained suppressive mechanisms are strongly orientation tuned (Petrov & McKee, 2009), and that the suppression enacted upon briefly presented stimuli (<100ms) is more profound. However, it remains to be seen whether the temporal invariance of tuning found at peripheral eccentricities is also present at fixation – it is possible that surround suppression in central vision has distinct spatiotemporal tuning characteristics. Tentatively: our findings of broad orientation tuning using sustained stimuli, taken together with the tighter tuning observed using transient stimuli (Xing & Heeger, 2000) are compatible with the orientation tuning of surround suppression in central vision being more temporally dependant than the suppression found in the periphery. Future experiments could directly address this possibility, and might also explore whether these spatiotemporal tuning characteristics vary as a function of absolute stimulus contrast. In Macaque V1, at perifoveal eccentricities, the orientation tuning of surround suppression appears to be strongest when using high contrast stimuli. Levitt and Lund (1997) found such contrast specificity in over half of the contrast-dependant neurons from which they measured responses, and Webb et al. (2005) noted that a low-contrast grating was similarly suppressed by a collinear grating annulus and one of uniform luminance. If the orientation tuning of surround suppression changes as a function of input contrast, perhaps the transience of its effect on perceived contrast does as well.

2.6 Conclusion

We have shown that the age-related increase in surround suppression reported in a recent series of publications is not a universal finding. Using stimuli designed to isolate surround suppression from overlay masking, we report no effect of age on supra-threshold contrast suppression, despite using contrasts previously reported to yield the maximum age-related difference. We propose that the lack of agreement between our results and previous reports could be unintentional contributions from overlay masking. The suppression we did find is weak, and only broadly tuned to the collinearity of the target and the surround, unlike previous reports of suprathreshold masking in central vision. We propose that future investigation of the strength of overlay masking relative to surround suppression across the human lifespan would be fruitful.

Chapter 3. Human sensitivity to global form violates Weber's law

3.1 Abstract

Decades of psychophysical experiments have shown that the perception of luminance contrast violates Weber's law: contrast discrimination is best at low (but non-zero) pedestal intensities. More recently, this "dipper effect" has been found in more complex visual parameter spaces, suggesting that it may be the consequence of a "canonical" computation in visual processing. Here, we ask whether the mid-level perception of global form elicited by Glass patterns also violates Weber's law, and whether the pattern of discrimination thresholds reveals a bias towards a certain axis of global form alignment (translational, radial, concentric). We parameterised global form by sampling Glass pattern dipole orientations from a Gaussian distribution, such that the variance of the sampling distribution controls the coherence of the pattern. In an online psychophysical experiment, we examined observers' sensitivity to global form across several pedestal coherence levels. By fitting the resultant detection and discrimination thresholds with a differentiated neuronal transducer function, we find evidence of mild dipper effect: threshold facilitation, followed by suppression proportional to Weber's law. Unlike previous reports, we find little evidence for a difference in sensitivity across pattern axes. Our work shows, for the first time, that the perception of global form elicited by Glass patterns may be subject to a dipper effect.

3.2 Introduction

Our senses provide graded, rather than binary responses to external stimuli. This means that we can usually distinguish or discriminate between two stimuli of the same type but different magnitudes. This discrimination is limited both by the noise in the measurement system and the difference in neuronal response that the two stimuli elicit. Over the past 100 years, a key finding from sensory psychophysics is that the discriminability of two stimuli depends on their absolute intensity. Two stimuli, with a set additive *difference* in intensity, are often harder to discriminate when their absolute intensities are high. One of the first (likely *the* first) attempts to make a

general model of the 'just noticeable differences' (JNDs) in stimulus intensity comes in the form of Weber's law (Fechner, 1860). For a given sensory parameter space, for any stimulus intensity I, Weber's law states that the JND in stimulus intensity (ΔC) is a fixed proportion (k) of the original magnitude such that $\Delta C = k \cdot C$. In other words, the difference in stimulus intensity necessary for sensory discrimination is a *multiplicative* constant, rather than *additive* one, such that the difference in intensity required is proportional to the absolute intensity of the stimulus. We now know Weber's law to be an approximation that can fail at very low and high stimulus magnitudes. At these extremes, JNDs often follow a different monotonic rule, such as a square root law ($\Delta C = \sqrt{C}$). Examples come from light intensity (Barlow, 1957), loudness (Parker & Schneider, 1980) and pure tones (McGill & Goldberg, 1968– - all show a small but reliable violation of Weber's proportionality. The perception of luminance contrast, however, has proven to be more problematic for Weber's law.

In 1974, Jacob Nachmias & Richard Sansbury found an intriguing violation of Weber's law. At very low pedestal contrasts, close to absolute detection, discrimination thresholds improve (are "facilitated") before showing a Weber-like trend of proportional increases (or "suppression"). This early facilitation of contrast (often called the "dipper" or "pedestal" effect) has been reliably observed for decades (Baker, Meese, & Georgeson, 2007; Foley, 1994; Foley & Legge, 1981; Legge & Foley, 1980; Nachmias & Sansbury, 1974), and has encouraged the development of theoretical explanations for nonmonotonicity in contrast discrimination. There are competing explanations from nonlinear transduction (Legge & Foley, 1980; Nachmias & Sansbury, 1974), and uncertainty reduction (Pelli, 1985). While the explanation for the dipper effect is still debated in some circles (see Solomon (2009)), recent experimental work has indicated that the "dipper effect" can be found in much more complex forms of visual perception. Morgan et al. (2008) presented observers with fields of vertically oriented Gabor patches and measured the increase in orientation variability required for the discrimination of two different Gabor fields. They found that the increase in standard deviation required for discrimination was subject to early facilitation followed by suppression. More recently, Gray et al. (2020) found evidence for a dipper effect in a facial emotion parameter space, where sensitivity to more subtle facial expressions is improved beyond the threshold for expression detection. Similar behavioural findings extend to the visual perception of perceived speed (Gori et al.,

2011) and 2nd order luminance patterns (P.-C. Huang & Chen, 2014). That the "dipper effect" can be found in several forms of visual perception – even as far as the perception of facial expressions - suggests that it may be the consequence of a "canonical computation": a fundamental neuronal operation that is applied at multiple stages of visual processing. For example, a "half-squaring" threshold nonlinearity has been found to predict the accelerating responses of neurons at low contrasts (Heeger, 1992b), and has been incorporated into models of divisive gain control (Carandini & Heeger, 2011). It is possible that a similar transformation (exponentiation at low stimulus intensities) is applied throughout the visual processing hierarchy, producing "dipper effects" in both rudimentary and more complex visual parameter spaces. In this chapter, we further examine the ubiquity of the "dipper effect" by exploring observers' sensitivity using a putatively mid-level visual stimulus that is an intermediary of luminance contrast and facial perception: the global form percept elicited by Glass patterns.

Glass patterns (Glass, 1969; Glass & Pérez, 1973; Glass & Switkes, 1976) are a global form stimulus comprised of orientated dot-pairs (dipoles) generated by drawing a field of randomly placed individual dots, applying a geometric transformation to their coordinates (i.e., dilation, rotation, translation) and redrawing them. The output is a field of dipoles that are locally distinct as separate pairs, but geometrically correlated across the whole pattern. This correlation yields the percept of form along the transformation axis that was applied, though no individual element is innately oriented. The perception of Glass patterns is thought to rely on at least two distinct pooling stages (H. R. Wilson & Wilkinson, 1998; H. R. Wilson et al., 1997): one that can pool dots into oriented dipoles, likely supported by the small oriented receptive fields of V1 and V2 (Movshon et al., 2010; Smith et al., 2007); and another that can pool local orientation signals into detectible global structure, likely located in extrastriate processing regions like V4 and the lateral occipital cortex (Gallant et al., 1996; Mannion et al., 2013; Ostwald et al., 2008). Glass patterns allow experimenters to manipulate the local information (carried by each dipole) and the global form information (carried across dipoles) independently. They are therefore an ideal stimulus for studying intermediary spatial pooling processes. Most experiments investigating Glass pattern global form focus on probing absolute detection thresholds by manipulating the number of locally corruptive unpaired dots in the dipole field

(Maloney et al., 1987), or more often by altering the proportion of signal and noise dipoles (Dakin & Bex, 2002; Seu & Ferrera, 2001; H. R. Wilson & Wilkinson, 1998). Recent experiments using the latter manipulation asked whether the human visual system is optimised to perceive global form along specific axes, finding a general enhancement for radial and concentric form over translational form (P.-C. Huang & Chen, 2016; Seu & Ferrera, 2001; H. R. Wilson & Wilkinson, 1998; H. R. Wilson et al., 1997).

Unlike absolute detection, the pattern of global form discrimination thresholds has scarcely been investigated, nor compared across pattern axes. Nankoo et al. (2012) indirectly assessed this by measuring global form sensitivity under signal detection theory (fitting a psychometric function to values of d' as a function of pattern coherence), but did not report or compare any parameters of these fits beyond the d' threshold (they also found increased sensitivity for concentric and radial patterns). For luminance contrast, discrimination thresholds can be predicted based on zero baseline d' values (Nachmias & Sansbury, 1974; Pelli, 1985), but this prediction assumes that observers' performance is limited by an additive source of noise (with a nonlinear transducer) which is still debated for contrast (Pelli, 1985; Solomon, 2009), and unknown for the discrimination of Glass pattern global form. To parameterise global form in the present study, we take an approach similar to that of Dakin (1997) and Morgan et al. (2008) by sampling dipole orientations from a Gaussian distribution. Though unlike Morgan et al. and Dakin, who effectively treat a perfectly aligned pattern as "zero" intensity (as in, zero variance), we do the opposite, assessing discrimination starting from random orientation noise (dipole orientations sampled from a uniform distribution). In an online experiment, we use dipole fields of increasing specificity around their principal axis to measure global form discrimination thresholds at several pedestal levels – separately for translational, radial, and concentric patterns. We have two aims. First, to assess the degree to which global form discrimination obeys Weber's law. Second, to investigate whether there is enhanced supra-threshold sensitivity to radial and concentric form over translational form, as in existing reports at detection threshold. To accomplish these objectives, we fit group-average discrimination thresholds with both a power-law (which describes only threshold suppression/a saturating nonlinearity) and the first derivative of a sigmoidal neuronal transducer that can simultaneously describe threshold facilitation

and suppression. We compare the descriptive merit of these models, and then contrast the bootstrapped fit parameters of descriptive model fits across Glass pattern alignment axes.

3.3 Methods

3.3.1 Participants

Forty-two observers were recruited for this experiment via Prolific (www.prolific.co), an online participant recruitment platform; 25 males, 16 females with a mean age of 27 years (SD = 7.8). One observer did not report their sex, and 3 did not report their age. Using Prolific's database of screening questions, recruitment was targeted at individuals who had self-reported normal or corrected-to-normal visual acuity; no disease or disorder of the eye or visual system; no severe hearing impairments; no personal history of epilepsy, cognitive impairment, or dementia; and usage of a computer, not a smartphone or tablet device, for participation. Satisfaction of these criteria was recapitulated via forced-choice questions on the consent form, presented to participants via Qualtrics (www.qualtrics.com), where their age and gender were also acquired. Before the point of participation, further screening was performed using custom JavaScript (JS) code; if a smartphone, tablet device, or unstable frame-timings were detected, participants were not able to continue. Participants were advised of all screening criteria in advance via the experiment's landing page, information sheet and consent form. Participants who satisfied these criteria and completed all 3 sessions of the experiment were paid £9 for their time. Ethical approval for the study was provided by the Department of Psychology at the University of York.

3.3.2 Online design and calibration

The experiment was conducted online using Pavlovia (www.pavlovia.org), a web-browser based experimentation platform. As such, full control of display and room conditions was not possible, but measures were taken to control for observers' viewing distance using the "virtual chinrest" protocol set out by Li et al. (2020). Briefly: for each observer's display, a pixels-per-centimetre value was obtained by asking the observer to scale a rectangle to the size of a standard reference (a bank

card). Viewing distance was then calibrated by identifying the distance from fixation to the nasal border of the left-eye optic disk blind-spot. This distance (in centimetres) was treated as equivalent to 13.5 degrees of visual angle. This equivocation was found to optimal by Li and colleagues, and we used it to maintain (approximately) the size of our stimuli across observers' retinae. Observers were asked to minimise movement throughout the experiment, and to try and sit at the same distance again if they pause the experiment for any reason. Remote psychophysical gamma correction via motion nulling (Ledgeway & Smith, 1994) was not possible at the time of data-collection due to web-browser limitations.

3.3.3 Visual Stimuli

The programming language Python (https://www.python.org/) was used to generate stimuli frames, which were stored as a set of high-resolution rasterised images for online presentation. The spatial parameters of the dipole-fields used are illustrated in Figure 3.1.

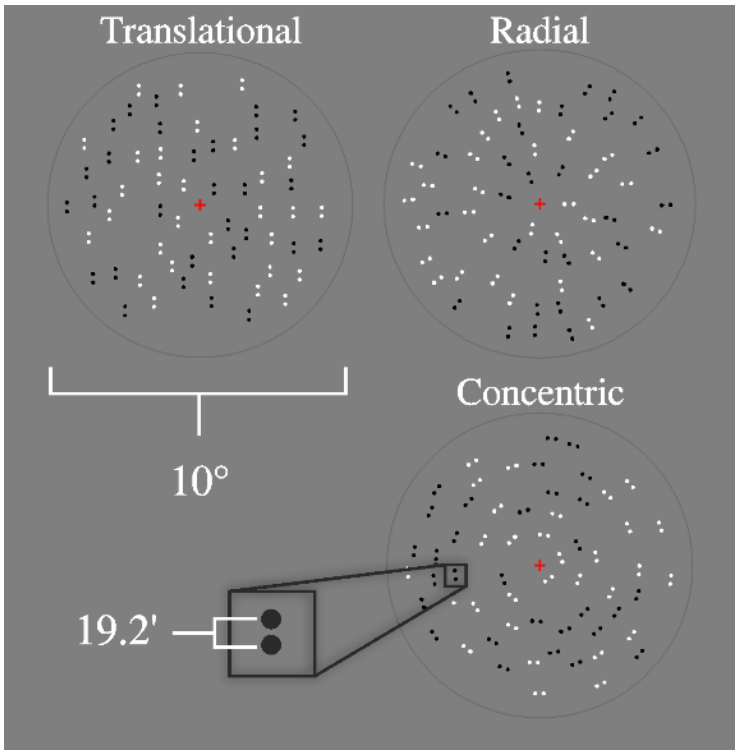

Figure 3.1: Spatial configuration of Glass patterns
An example of the Glass patterns used in this experiment; Translational, Radial, and Concentric transformations were applied to random dot coordinates.

Each stimulus frame was a field of 62 full contrast dot-pairs (dipoles), with a 50/50 split of black and white pairs. Assuming the accurate remote calibration of

observer distance and display scale, these dipoles were presented within a circular field subtending 10° of visual angle, with an areal density of 3.9% (1.6 pairs/deg^2). Within a dipole, each dot had a diameter of 10.02 arcmin and a centre-to-centre separation of 19.2 arcmin. The distance between dipole centres was at least 45.6 arcmin. This was ensured for each dipole field by continuously generating random dipole-centre coordinates and keeping values that were lawfully spaced and rejecting those that were not until a sufficient number of coordinates were generated (a method often referred to as rejection sampling). A thin grey ring was presented around the dipole fields for the duration of the experiment to eliminate any spatial uncertainty about where the stimuli would appear. Dipole-fields with different degrees of global form coherence were generated for translational, radial, and concentric pattern axes – see Figure 3.1 for examples of these patterns at maximum coherence.

$$vM_{pdf}(x \mid \kappa, \mu) = \frac{e^{\kappa \cos(x-\mu)}}{2\pi I_0(\kappa)} \quad \textbf{(1)}$$

To generate dipole fields of different global form coherence, the angular orientation of each dipole in a field was sampled from a semi-circular von-Mises distribution. The von-Mises probability density function is provided in equation 1. Here, *kappa (κ)* is a specificity parameter – it is analogous to the reciprocal of the variance. A *kappa* of zero produces a uniform distribution, while increasingly positive values increase the probability density symmetrically around the mean angle *mu (μ)* – which we set to zero. Sampling from this distribution returns values spanning 2π, which are halved, wrapping them on the semi-circle instead ($-\pi/2 > \pi/2$). This was done as each dipole is oriented along a line, but has no sense/direction (i.e., for a radial dipole, there is no distinction between an inwards and outwards heading). In the end, this produces an array of 62 semi-circular angles/arc lengths, centred on zero, with an angular specificity dependant on the value of *kappa*. This array is then used to displace a 2nd array of 62 angles that are perfectly aligned along the desired global form axis. Therefore, as *kappa* increases, so does the global orientation alignment of the dipole field.

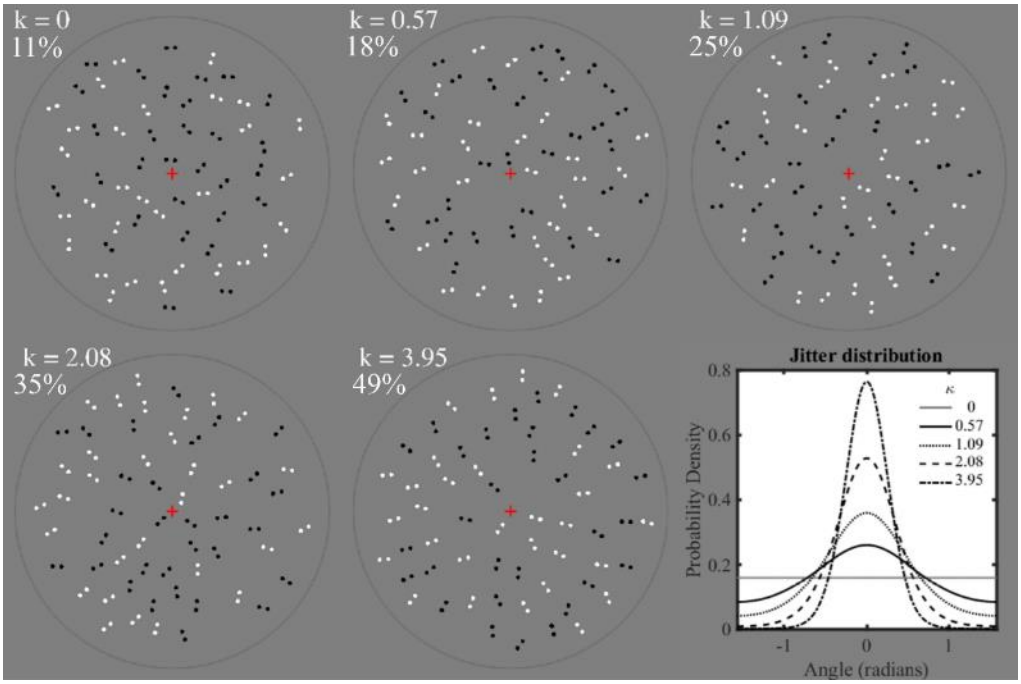

Figure 3.2: Parameterisation of global form

The pedestal levels of von-Mises concentration (k) used in this experiment shown using a radial Glass pattern. The bottom-right figure shows the jitter distribution for each level of kappa used as a pedestal. The percentages under k values are the output of equation 2 for each value of k.

To estimate observers' global form discrimination thresholds psychophysically, the von-Mises distribution was used to generate patterns across a range of *kappa* values, illustrated in Figure 3.2 for a radial pattern. For translational, radial, and concentric pattern axes, dipole orientation arrays were generated at five pedestal levels of *kappa,* a range we refer to as κ_p, where *p* is the rank of the pedestal. The first pedestal level (κ_1) was zero, representing detection of global form against random noise. Subsequent levels ($\kappa_2 > \kappa_5$) were spaced in 5.6dB increments, starting from $\kappa = 0.57$. Then, for each value of κ_p, a separate list of $\kappa_{p + increment}$ values was generated allow the measurement of discrimination thresholds. There were five increment values used, more widely spaced (by 8dB) to accommodate cross-observer variance in discrimination thresholds. For pedestals κ_1, κ_2, and κ_3, the 8db increments started at 0.150, but for κ_4 and κ_5, they started at 0.600. This decision was made based on pilots suggesting that a range of increments starting from 0.15 would mostly capture chance-level performance for κ_4 and κ_5. Dipole-field images were then generated for each unique value of κ_p and $\kappa_{p + increment}$. For every unique *kappa* value used, we produced fifteen pre-rasterised exemplars drawn from the same probability distribution. The order of exemplars used was pseudo-randomised such that the same exemplar could never be presented twice in immediate succession. Multiple exemplars were used to

reduce the probability of pattern discrimination purely on the basis of recognising a previously seen image. As the number of dipoles used in this experiment was quite low, each distribution generated was assessed using a Kolmogorov-Smirnov test (against $\alpha = 0.05$) for the equality of the randomly sampled variates with the theoretical von-Mises PDF. Any sample distributions that failed this test were resampled until a legal distribution was obtained.

$$Pedestal_{\%} = \left(\int_{-10^{\circ}}^{10^{\circ}} vM_{pdf}(x, pedestal_k, 0) \right) \quad \textbf{(2)}$$

$$Increment_{\%} = \left(\int_{-10^{\circ}}^{10^{\circ}} vM_{pdf}(x, pedestal_k + increment_k, 0) \right) - Pedestal_{\%}$$

To represent our parameterisation of global form in finite units that can related to neuronal orientation tuning profiles, pedestal and increment kappa values were converted to a bounded range that we refer to as "global form coherence". These transformed values were obtained by integrating the von-Mises PDF at all values of κ_p and $\kappa_{p + increment}$ across a 20° range centred on zero, and treating the value yielded by this integration as a measure of global form intensity. This value describes the probability of sampled dipole orientations falling within 10° of perfect alignment, and we treat it as though it faithfully equates this proportion, which is reliable assumption due to the use of a K-S tests to filter out aberrant distributions. The 20° integration window was chosen based on primate literature indicating that this is the lower-bound on the full-width half-maximum bandwidth of neurons in V1 with a receptive field size that covers a dipole width (Gur et al., 2004). This approximates the integration range reported by Dakin (1997), who found the judgement of mean orientation of translational Glass patterns to be stable until a local orientation standard deviation of approximately 9°. That our transformation (equation 2) assumes that all dipoles within the 20° integration range will equally affect the neuronal response is an oversimplification, but we believe it is sufficient for the relative comparisons performed in the present experiment, and for the comparisons performed in Chapter 4. This transformation was performed using equation 2, where the difference between the transformed values of κ_p and $\kappa_{p + increment}$ becomes the increment values to which we fit a psychometric function for each observer to obtain JNDs. The transformation between kappa and coherence defined in this way is illustrated in Figure 3.3.

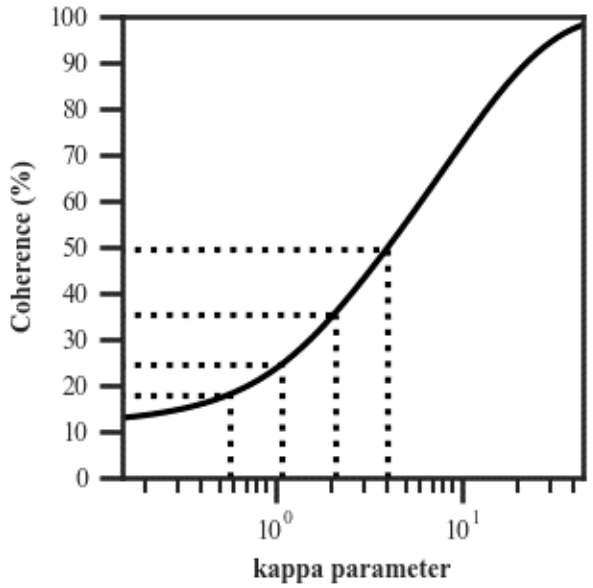

Figure 3.3: Transformation of kappa to coherence
The function of equation 2 used to convert the von-Mises concentration parameter (k) to a percentage based on existing tuning bandwidth estimates of neurons in primate visual cortex. The dotted lines represent the pedestal levels used in this experiment.

3.3.4 Experimental procedure

In a between-subjects design, each observer was allocated to either the translational, radial, or concentric pattern condition. While a within-subjects comparison between pattern axes would have been preferable, this would have required 3x the time commitment from each observer, and we were concerned that online participants may not return data for all three conditions. For each observer, the whole experiment took place across three identical sessions, for which the following procedure was carried out. Initially, observers carried out the calibration procedure as detailed previously. Then, they were shown short animations and written instructions explaining their task, the pattern axes they would see; (i.e., what is meant by "concentric", "radial", and "vertical"); and images conveying our operationalisation of global form alignment (similar to Figure 3.2, but without the distribution curves). Then, using the method of constant stimuli (MOCS), observers performed a two-interval-forced-choice (2IFC) procedure to obtain discrimination accuracy estimates for each increment coherence, at each pedestal. The timing of a single 2IFC trial is illustrated in Figure 3.4. Note, all blank intervals still contained the grey annulus and fixation cross. Observers saw a 500ms empty period, 100ms of pattern interval one, a

750ms empty period, followed by 100ms of pattern interval two. Then, they had unlimited time to make their response. Each interval contained a dipole-field generated either from a pedestal or pedestal + increment value of *kappa*, and the order of pedestal and target presentation was randomised on every trial. Observers were told that their task was to select the pattern that appeared to be most aligned along the pattern axis that had been explained to them. In each session, for each of five pedestal *kappa* levels, 20 trials were collected per increment *kappa* level, totalling 100 trials per session. To stabilise observers' response criteria prior to data collection, they carried out an abbreviated version of the experiment before the real task began, which contained 1 trial for each increment value, for every pedestal (a total of 25 practice trials before each session).

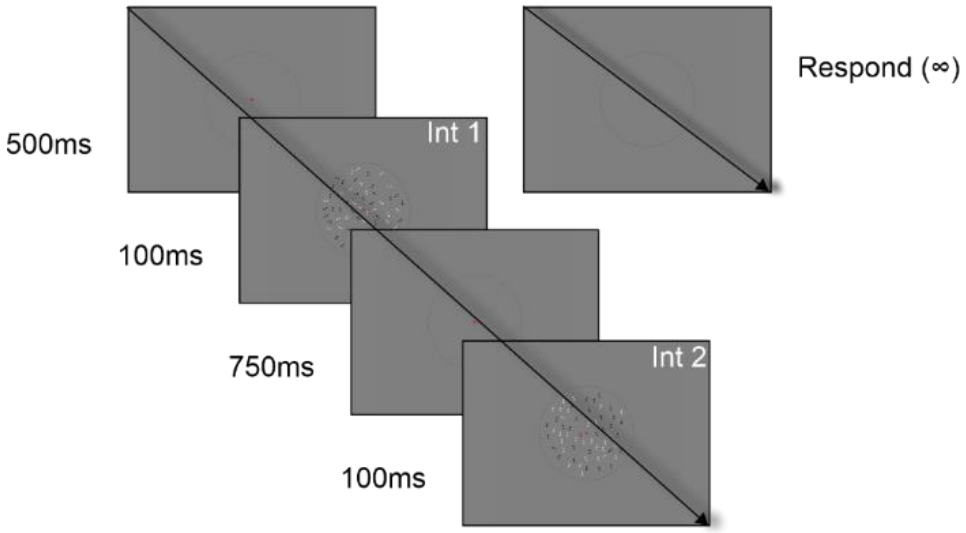

Figure 3.4: Stimulus presentation
Timing of a single 2AFC trial of the MOCS procedure used in this experiment. Observers perceived two 100ms intervals separated by 750ms. One interval contained the pedestal patterns, the other contained the target pattern.

3.3.5 Model fitting

To extract JNDs at each pedestal coherence, each observer's response accuracy for each increment coherence was fit with a Weibull function using the psignifit-4 toolbox (Schütt et al., 2016) in Python. The lapse rate was fixed to 1%, the lower asymptote of the function was fixed at the 2IFC chance level of 50%, and the threshold and slope parameters were allowed to vary freely. This fit was then used to extract a coherence value equating the 75% response accuracy. For each pedestal level, all 300 trials across an observer's three sessions were pooled and treated as if from a single

session. This means that that each observer had a single Weibull fit (and JND) for each pedestal level. Observers' discrimination thresholds were normalised to their respective detection thresholds prior to further model fitting, a process which will be elaborated upon in the results section. We compared normalised discrimination thresholds to the median threshold for detection using estimation statistics (Claridge-Chang & Assam, 2016), whereby significance is assessed by generating 95% confidence intervals on the normalised values. Where these confidence intervals do not contain the median detection threshold, a significant difference can be concluded.

$$Response(c \mid n, c_{50}, rMax) = rMax * \frac{c^n}{(c^n + c_{50}^n)} \quad \textbf{(3)}$$

To provide descriptive parameter estimates on the pattern of detection and discrimination thresholds across observers, median thresholds were fit with the first derivative of a hyperbolic ratio function, a neuronal transducer curve which is often used to model contrast responses (Albrecht & Geisler, 1991; Albrecht & Hamilton, 1982; Boynton et al., 1999; Williford & Maunsell, 2006). This function (eq. 3) models the input-output transfer function of neurons whose responses encode a sensory parameter space, and its 1st derivative can be used to predict behavioural thresholds. This function has the useful quality of being able to describe a range of response profiles with only three parameters,: *rMax*, the *c50*, and *n-exponent*. The *rMax* is a scaling coefficient reflecting the intensity at which a neuron's (or population of neurons') response reaches its maximum value - it is defined in arbitrary units when fitting behavioural thresholds. Changing the *rMax* parameter vertically translates the behavioural thresholds predicted by the model, and has previously been related to attentional modulation (L. Huang & Dobkins, 2005). The *c50* describes the stimulus value that produces and output equal to half the value of *rMax*. Changes in the c50 diagonally translate the thresholds predicted by the model, such that increasing the c50 shifts the discrimination function to operate over higher pedestal and threshold values. The exponent, *n*, describes the form of nonlinearities occurring around the *c50* value. When *n > 1,* this simultaneously predicts facilitation at low intensities, and suppression at high intensities (the form of suppression may or may not resemble Weber's law). Exponents at or below one mean that the input-output function is never expansive, only compressive, reflecting behavioural threshold suppression but an

absence of facilitation. Note, a *c50* beyond 100% is possible (in that it describes a legal function), but the resultant hyperbolic ratio function may be a less than optimal descriptor for the underlying neuronal response function. We compared the fit quality of the differentiated hyperbolic ratio function with that of a power-law $R(C\ |y\ m) = y\ \cdot C^m$, where an exponent (m) of 1 equates to the proportional increase predicted by Weber's law. For both functions, the parameters that best fit the median of the normalised JNDs were found using a Nelder-Mead simplex search algorithm that minimised the sum of squared error between the model and the median thresholds. A bootstrapping procedure was then used to generate confidence intervals on fit parameters and the fit residuals. In a single iteration of this bootstrap: pedestal JNDs were resampled with replacement across all observers. Then, the median of these values was fit with the desired model, and the parameters saved. This was repeated 10^4 times, resulting in 10^4 resampled estimates on the model fit parameters. Contrasts of the fit parameters between the pattern axes (translational, radial, concentric) were then performed whereby the distributions of fit parameters were differenced, and confidence intervals on these differences calculated as the 2.5^{th} and 97.5^{th} percentile values of the distribution of differences. Where this confidence interval does not contain zero, we can conclude with a high degree of certainty that there is a difference in a given fit parameter between conditions.

3.4 Results

3.4.1 Example psychometric functions

Examples of psychometric functions fit to each pedestal level for a single observer are presented in Figure 3.5. Each plot shows the response accuracy estimates for a single pedestal at each increment value, and the best fitting psychometric function. For this observer, at all pedestal levels, the data are well described by a Weibull CDF. The same is true of most observers, except for three who's data were rejected on the grounds of consistent chance-level performance (highlighted in Appendix C). Additionally, some observers' Weibull fits to the highest pedestal level had poor support above the JND threshold, meaning they were unable to achieve 75% accuracy at any increment value used in the experiment for this pedestal. This was the case for three observers in the concentric condition, and one observer in the translational

condition. For these three observers, for the highest pedestal, the JND thresholds generated were not entered into the bootstrapping procedure, but thresholds from all lower pedestals were kept. Overall, bootstrapped model fits were conducted on data from 12, 14, and 13 observers for the translational, radial, and concentric conditions, respectively. See Appendix C all observers' psychometric fits with rejected thresholds highlighted.

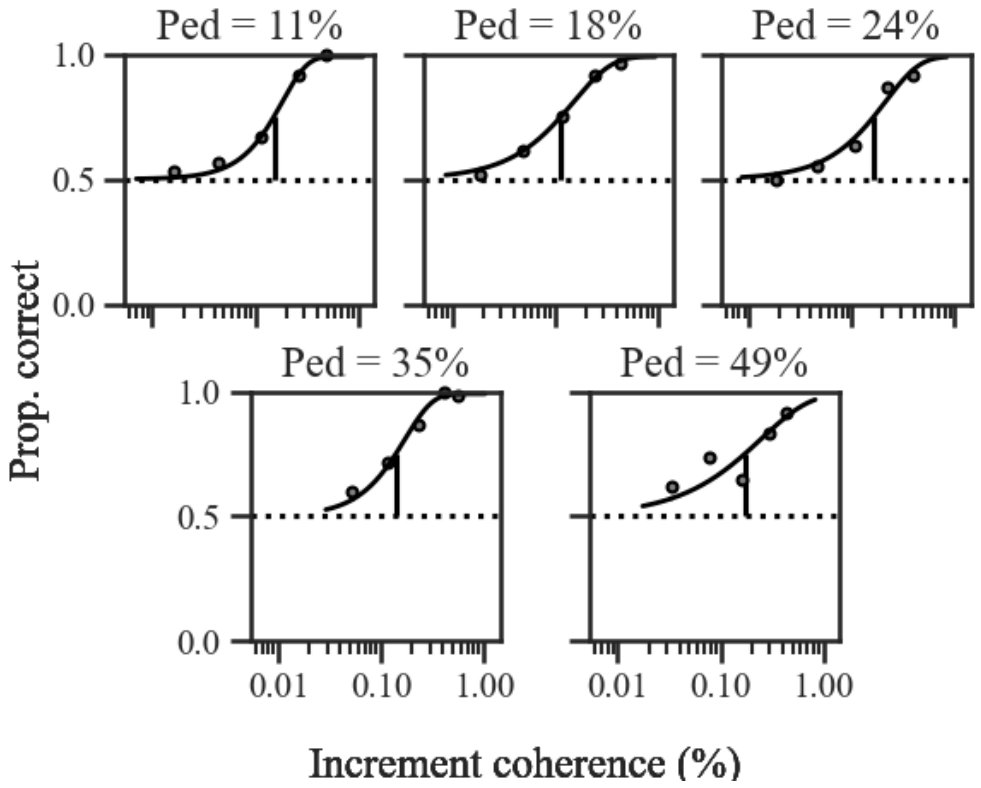

Figure 3.5: Example global form psychometric function fits

Psychometric function fits for a single observer. Each subplot shows the Weibull function fit to the accuracy estimates for a single pedestal level (see headings). Where the vertical reference line intercepts the curve marks the 75% JND, while the horizontal dotted line shows the 2AFC chance level of 50% accuracy.

3.4.2 Normalisation to variable detection thresholds

Global form coherence thresholds prior to normalisation are shown in Figure 3.6A. Qualitatively, these thresholds show some indication of facilitation at the first pedestal for the concentric condition, and at the 2nd pedestal for the radial condition. Note, however, that the variance in detection thresholds (the 95% CI on the open circles) includes the discrimination thresholds at all but the final pedestal. Although this variance in detection thresholds was present in each pattern condition, the spread of detection thresholds was quite similar across conditions. Figure 3.6B shows bootstrapped confidence intervals on the difference between the median detection

threshold for each combination of conditions, where there was a slight bias towards lower detection thresholds in the translational condition, but the confidence interval indicates that this bias was non-significant.

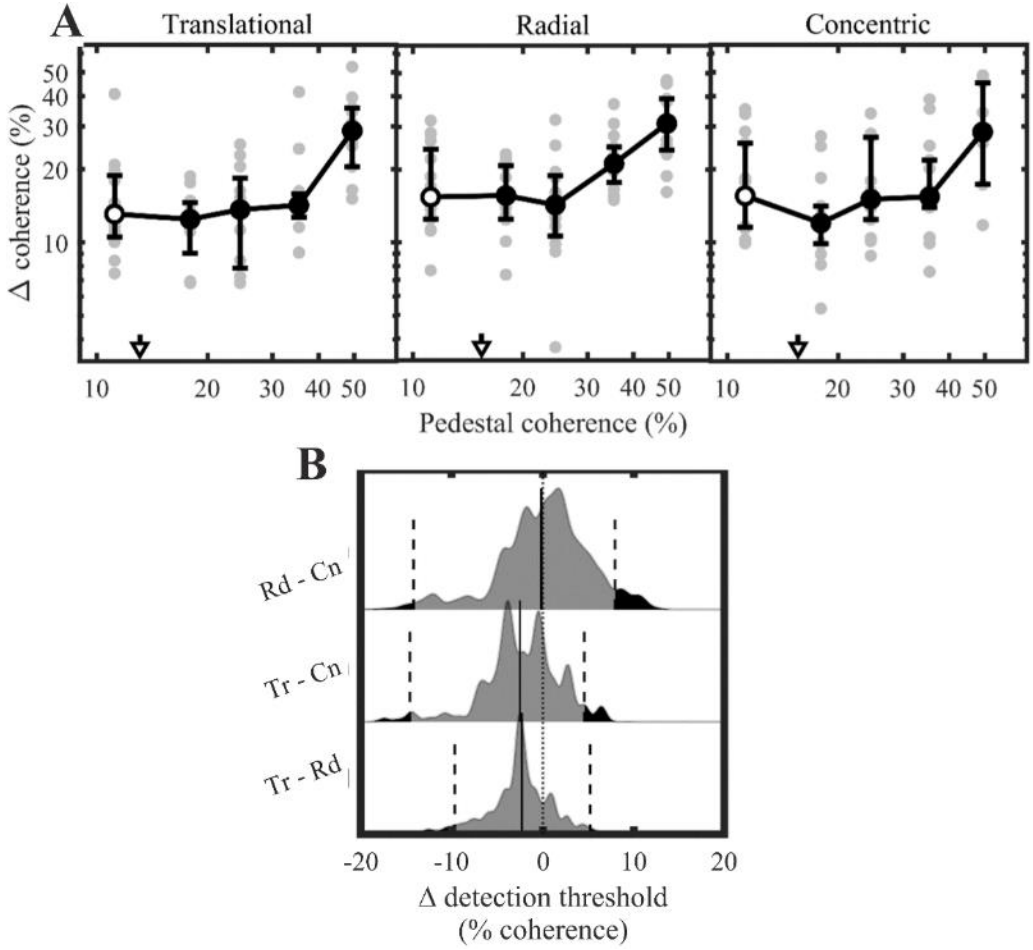

Figure 3.6: Coherence discrimination and detection thresholds

A: Coherence threshold estimates for all pedestal levels across pattern conditions. Within each plot, the filled circles show median discrimination thresholds, while the open circle is the median detection threshold. The small grey circles show individual data points, and the small arrows on each abscissa show the median detection threshold for a given condition. B: kernel density estimates on the bootstrapped difference in detection threshold between each pattern condition. Vertical solid lines denote the median, while vertical dashed lines denote the 95% CI on the median difference. Where this CI does not contain zero (vertical dotted line), a significant difference can be concluded.

To account for any observer-specific vertical translation of discrimination thresholds (caused, for example, by slightly different levels of internal noise), the data in Figure 3.6A were normalised to each observer's threshold for detection, such that $F_c = \Delta C / C_{det}$, where F_c is a change-from-detection coefficient, ΔC is an observer's discrimination threshold for a given pedestal, and C_{det} is their threshold for global form detection. These coefficients were then transformed back into the global form coherence parameter space via $\Delta C_{norm} = F_c \cdot \widetilde{C_{det}}$, where the normalised

discrimination thresholds for each observer are the product of the group median detection threshold and their individual change-from-detection coefficient. This effectively equates absolute detection thresholds across observers. The result of this normalisation is shown Figure 3.7, where thresholds above or below the dashed horizontal line suggest suppression or facilitation, respectively. Confidence intervals that do not cross this line are indicative of a significant group-level difference from detection for a given pedestal.

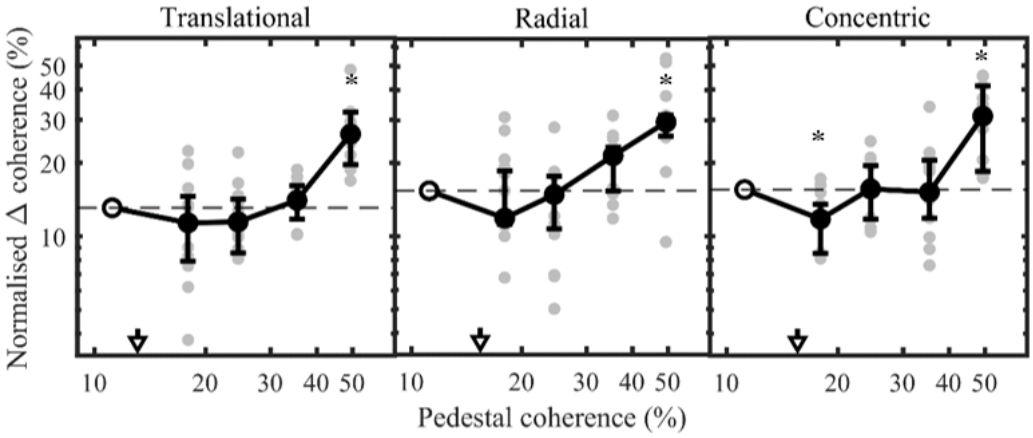

Figure 3.7: Normalised coherence discrimination threshold
Normalised discrimination thresholds for each pattern condition. The horizontal dashed line denotes the median detection threshold, as does the open circle. Thresholds with a star above them deviate significantly from median detection.

In all pattern conditions, the median discrimination thresholds fall below the detection line, but this facilitation is only quantitatively supported in the concentric condition at the first pedestal level. For all conditions, there is evidence for significant suppression of thresholds relative to detection at the highest pedestal level. Overall, initial assessment of normalised data suggests significant suppression for all conditions at the highest pedestal, and significant facilitation for concentric Glass patterns. However, interpreting the threshold values in isolation of the overall pattern of thresholds does not capture the full extent of these results. So, to describe the pattern of sensitivity across discrimination thresholds, and compare this pattern across conditions, we now present the results of our model fitting procedure.

3.4.3 Violation of Weber's law at low global form coherence

To assess the degree to which global form coherence thresholds adhere to Weber's law, we fit normalised thresholds with the first derivative of a hyperbolic ratio function and a power-law and bootstrapped their fit parameters. To determine the

optimal model, we compare the distribution of root-mean-square error returned by each model generated by our bootstrapping procedure. As fits were performed on thresholds normalised to absolute detection, the residuals being minimised by the simplex algorithm were weighted such that fits were forced to pass through the median detection threshold (the open circles). The fits of the hyperbolic ratio and power law are shown in Figure 3.8A and Figure 3.8B, respectively.

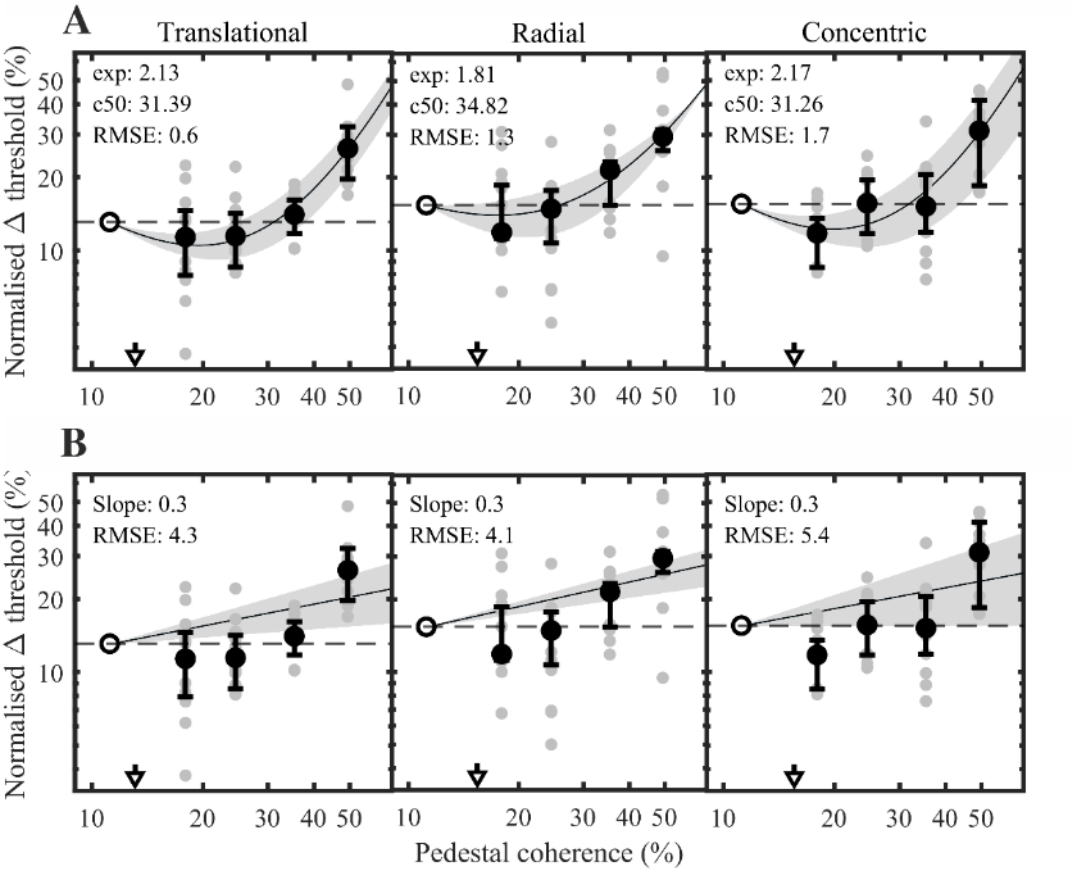

Figure 3.8: Model fits to normalised thresholds

A: fit of a hyperbolic ratio function. Black markers show the median thresholds with a 95% confidence interval. The solid line shows the best fit of the model to the median thresholds, and the grey shaded area shows the 95% CI on the model parameters derived via our bootstrapping procedure. Fit parameters of the solid line are shown in the inlaid text. The dashed reference line represents the median detection threshold. B: Same as A, but for the power law model.

The hyperbolic ratio fits generated sensible values for all conditions; for radial and concentric very few bootstraps resamples yielded incompatible fits with c50s beyond 100% coherence ($10/10^4$ and $11/10^4$, respectively). For the translational condition, all fits were compatible. As the proportion of incompatible fits was so low (< 0.01%), we proceeded with further comparisons, but removed the incompatible fits from the parameter distribution yielded by the bootstrapping procedure. It is notable that the predictive power of the power-law fits were significantly worse than that of

the hyperbolic ratio. The residuals of fits to the actual medians are shown inlaid in Figure 3.8, where the RMSE of the power-law is more than three times that of the hyperbolic ratio (>7x for translational). This reduction in error for the hyperbolic ratio fits was significant, as shown in the contrasts in Figure 3.9. Across all bootstraps, this represents a median improvement in RMSE of 3.5% (95% CI [1.8, 5.0], 3.2% (95% CI [1.4, 5.4]), and 3.7% (95% CI [0.7, 7.1]) for the translational, radial, and concentric conditions, respectively.

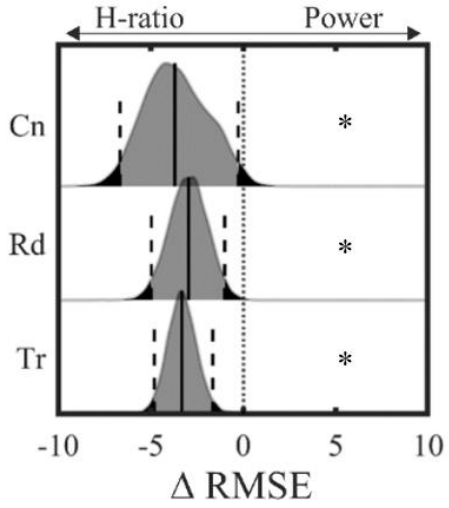

Figure 3.9: Contrast of model residuals
Kernel density estimates showing the difference in residuals between the bootstrapped fits of the hyperbolic ratio and power-law functions. A shift left of the zero-line indicates lower error for the hyperbolic ratio, a rightward shift indicates a better fit for the power-law. Tr: translational, Rd: radial, Cn: concentric. Contrasts with a star besides them are indicative of a significant difference between comparators.

As the fit from the hyperbolic ratio was the better descriptor of observers' thresholds across the board, we now report the bootstrapped exponent parameter of this model, which describes the degree of suppression and facilitation across the pattern of thresholds. The distribution of exponent values for each condition are shown in Figure 3.10, where the median and lower confidence intervals on the exponent were greater than 1 for the translational (median = 2.15, 95% CI [1.75, 2.56]), radial (median = 1.86, 95% CI [1.42, 2.39]), and concentric (median = 2.22, 95% CI [1.57, 2.72]) conditions. Exponents greater than 1 suggest that the pattern across thresholds is best explained by a model containing both the facilitation and suppression, possibly due to an accelerating and saturating nonlinearity. This indicates that discrimination of the global form percept elicited by Glass patterns does not obey Weber's law at low levels of form coherence. In other words, when the pattern of thresholds across all

pedestals is considered, we find evidence suggestive of a "dipper effect" in the perception of global form coherence for all pattern axes.

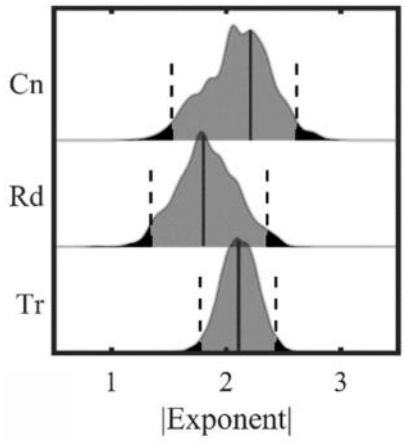

Figure 3.10: Hyperbolic-ratio exponent parameter
Exponent parameter KDEs from the bootstrapped fit of a hyperbolic ratio function to normalised thresholds. Vertical dashed lines and black regions of the KDE represent the 2.5th and 97.5th percentiles of the distribution, while the solid line represents the median.(Tr: translational, Rd: radial, Cn: concentric.

To test the extent to which the suppressive "handle" of the dipper function is consistent with Weber's law, we repeated the bootstrapped fit of the power-law model, but constrained it to only consider the discrimination thresholds. The results of this fit, and the contrast of its residuals to the hyperbolic ratio fits, are shown in Figure 3.11A and B, respectively. Here, the fit of the power law is vastly improved, and the contrast of the power-law and hyperbolic ratio residual distributions demonstrates that both models are similar in the quality of their fit, still with a slight bias towards the hyperbolic ratio, but this does not reach significance. In Figure 3.11C, the bootstrapped estimates on the exponent of the power-law function are provided, which cluster around a median of 1 for translational and radial, but are more variable in the concentric condition (likely owing to the discontinuity at the 3rd pedestal level). Overall, these fits indicate that supra-threshold discrimination of global form in Glass patterns is equally described by both a hyperbolic ratio and a power law, and that the form of this function is, on average, in agreement with Weber's law. However, when detection thresholds are taken into consideration, the hyperbolic ratio function is clearly the better predictor of the data between these models.

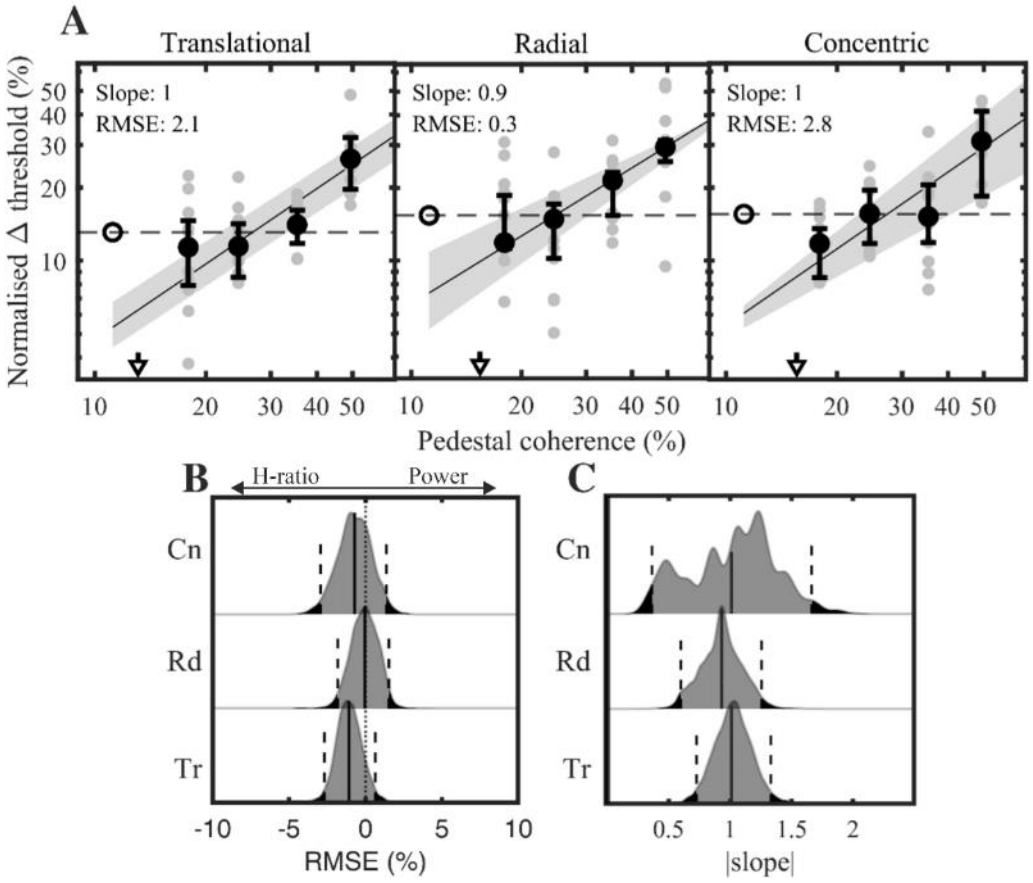

Figure 3.11: Fit of a power law suprathreshold discrimination
A: Normalised thresholds fit with a power-law function, ignoring absolute detection thresholds. B: contrast of the power-law residuals with the residuals of the hyperbolic fit. C: the exponent parameter distribution of the power-law. Tr: translational, Rd: radial, Cn: concentric

3.4.4 No evidence for enhanced sensitivity to polar form

Next, we investigated whether observers demonstrated enhanced sensitivity to a specific global form axis, as has been found in previous reports (Seu & Ferrera, 2001; H. R. Wilson & Wilkinson, 1998). We compared the bootstrapped parameter estimates of the hyperbolic ratio function across pattern conditions. Figure 3.12A shows the distribution of parameter estimates for each condition, while Figure 3.12B shows contrasts performed between parameter distributions for each condition. In Figure 3.12A, the median *c50* appears quite similar across pattern types; but is highest for radial (median: 33%, 95% CI [29.4, 42.5]), lowest for concentric (median: 30.8%, 95% CI [27, 56.1]) and intermediate for the translational condition (median: 31.3%, 95% CI [28, 36.4]). For the values of the exponent, the concentric (median: 2.12, 95% CI [1.75, 2.55]) and translational (median: 2.22, 95% CI [1.56, 2.7]) conditions have similar distributions, with a leftwards shift in the radial condition towards a lower

exponent (median: 1.86, [1.42, 2.38]), responsible for the shallower negative deflection in Figure 3.8A for this condition. For the *rMax*, which denotes the vertical translation of the curve, the medians of the radial (3.80, 95% CI [3.50, 4.65]) and concentric (3.90, 95% CI [3.31, 6.93]) condition are similar, while the translational condition shows a slight rightwards shift in the distribution (4.38, 95% CI [3.90, 5.15]). Note that both the *c50* and *rMax* distributions of the concentric condition are heavily skewed relative to the other conditions. As before, this is likely due to the discontinuity present at the 3rd pedestal level. To assess the confidence that can be placed on these differences, fit parameter contrasts across conditions are shown Figure 3.12B. Here, there is little evidence for a significant difference in any of the hyperbolic ratio fit parameters for translational, radial, or concentric Glass patterns as all confidence intervals, for the difference in all parameters, cross zero. This suggests, when discrimination thresholds are considered in aggregate, observers' sensitivity to the global form percept elicited by Glass patterns is not uniquely enhanced for radial or concentric polar form.

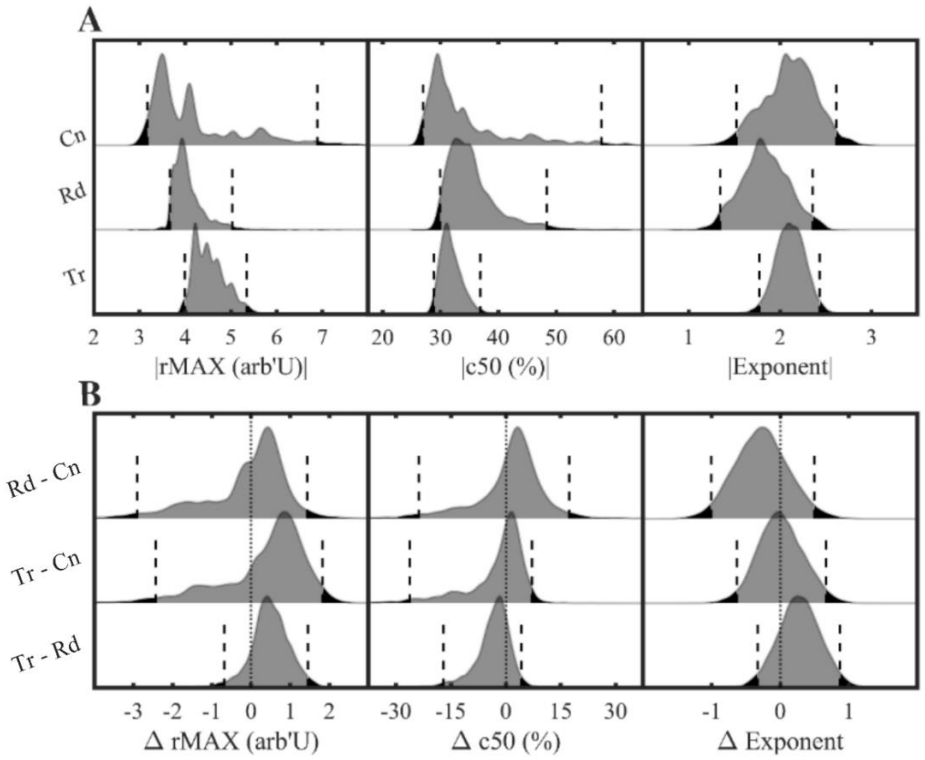

Figure 3.12: Cross-condition contrast of model fit parameters

A: The distribution of parameter estimates retrieved from the bootstrapped model fit for each pattern condition. Each sub-plot represents the distributions for a single fit parameter, and each ridgeline shows the kernel density estimate of the distribution's histogram. The vertical dashed lines show the percentile 95% CI, any area of the KDE outside of this range is shaded black. B: The condition-wise contrasts of the distributions in A. Where the vertical dashed lines do not contain zero a significant difference can be concluded. Tr: translational, Rd: radial, Cn: concentric

3.5 Discussion

This experiment had two objectives: to explore the pattern of behavioural sensitivity underlying the perception of Glass pattern global form; and to investigate whether there is any evidence for a difference in suprathreshold sensitivity between three often used pattern axes (translational, radial, and concentric). Our findings indicated that Glass pattern detection and discrimination thresholds cannot be explained by Weber's law, and are significantly better described by a model based on a neuronal response function that includes an accelerating and saturating nonlinearity, consistent with a "dipper effect". When hyperbolic ratio parameters were compared across pattern types, there was little evidence for a consistent enhancement of sensitivity specific to a certain global form axis. We will now discuss where our results fit in with previous work on Glass pattern sensitivity and competing explanations of observers' sensitivity to changes in visual stimuli magnitude.

3.5.1 A "dipper effect" for global form coherence?

Our modelling results, which summarised the shape of the discrimination function across thresholds, reported parameters consistent with a mild dipper effect for the percept of global form elicited by Glass patterns. However, it is true that there was a great deal of variance in the horizontal translation (conveyed via the *c50)* of the modelled discrimination functions, particularly for radial and concentric patterns. What might explain this variation? One possibility is that the range of pedestals at which facilitation occurs is unstable, perhaps due to top-down attentional processes. Indeed, significant suppression of thresholds (relative to detection) was only apparent at the final pedestal coherence, with many individual thresholds still less than detection for the preceding pedestal. In the contrast domain, Huang and Dobkins (2005) demonstrated that shifts in the c50 can be induced by introducing a secondary task that occupies observers' attention. As our experiment was conducted online (and was thus unsupervised) varying degrees of observer attention should perhaps be expected. Since we did not directly measure observers' attention, and have insufficient data to accurately estimate the lapse rate of psychometric function fits, we cannot rule out this possibility. However, if we accept that increased reaction times are a corollary of inattention, we can say that they do not appear to correlate with the *c50s* estimated by our bootstrapping procedure, as shown in Figure 3.13. It could also be argued that the

lack of control over the luminance of the displays observers' used might affect the horizonal location of thresholds. We doubt this is the case, however, as the psychophysical results of Burton et al. (2016) demonstrate that global form density thresholds are quite stable across photopic and mesopic viewing conditions. Moreover, the perception of Glass patterns is most dependant on the similarity between the dots comprising a dipole (Kovaćs & Julesz, 1992; J. A. Wilson et al., 2004), as opposed changes in the luminance of all dots producing the pattern.

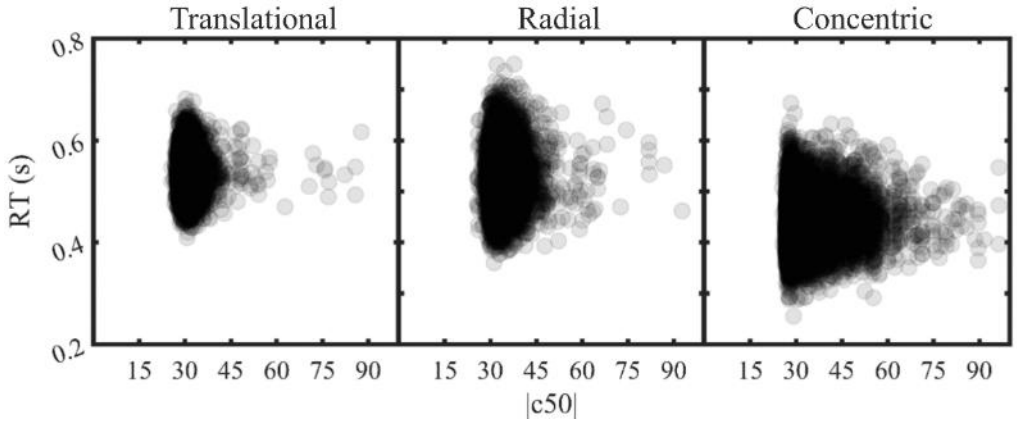

Figure 3.13: Reaction times as function of c50
The bootstrapped distribution of c50s plotted against the median reaction time of observers within each resample, for each pattern condition. Individual circles represent the values from a single resample.

3.5.2 Nonlinear transduction or uncertainty reduction?

Our model fitting results are consistent with a "dipper effect" in the perception of global form. The most-often cited explanation for the dipper effect is a nonlinear/sigmoidal neuronal transducer (Boynton et al., 1999; Foley & Legge, 1981; Nachmias & Sansbury, 1974), where the factor determining discrimination thresholds is the form the transfer function possessed by neurons encoding a given parameter space. Indeed, by fitting our behavioural thresholds with the differentiated neuronal transducer function, it would appear that we are implicitly supporting the validity of this account. However, the dipper effect can, in theory, be explained without the use of any neuronal nonlinearity. A competing explanation for threshold facilitation is uncertainly reduction (Pelli, 1985), an approach that models the visual system as a collection of channels, with each channel describing some parameter in the visual world. Here, the dipper effect arises because of uncertainty about the optimal channel(s) to use as the basis for a behavioural decision. When measuring absolute detection thresholds via a pedestal experiment, one of the comparators is the absence

of stimulation (a pedestal of zero). This means the perceptual cues indicating informative channels are severely reduced, elevating channel uncertainty. This uncertainty means the observer may be attending to a mixture of informative and uninformative channels, thus the threshold for absolute detection is elevated relative to discrimination at low pedestal intensities. Could uncertainty reduction explain our results? It is unclear where uncertainty might have infiltrated our methodology. When the dot-field was not being presented, its location was constantly indicated by a thin grey annulus, so any uncertainty in the location of the stimulus was unlikely. It is possible that there was a degree of initial uncertainty in pattern axis that should be attended, but observers practiced with all pedestals prior to the task, giving them ample time to develop their response criterion. Moreover, all pedestals were tested in separate blocks, so the ideal channel to attend for a given pedestal was likely optimised within the first few trials. For these reasons, we suspect that an explanation from uncertainty reduction is unlikely, though we cannot definitively rule out some contribution from intrinsic uncertainty. One indication that our behavioural results may be explained by a nonlinear transducer is the steepness of the psychometric function underlying observers' discrimination thresholds. If the transducer is subject to an accelerating nonlinearity followed by a saturating nonlinearity, we should expect to see the slope of the Weibull psychometric function steepen where performance is best (proximal to the first pedestal after detection, in this case) which is qualitatively the case for our results for two of the pattern axes we tested (on average, as shown in Appendix D). However, it is also true that the slope parameter is quite variable in our results, and some observers show a reduction in slope at this and later pedestals, which would be consistent with an alternative explanation from uncertainty reduction (Pelli, 1985). Indeed, these explanations are difficult to distinguish psychophysically, and they likely both play a role in determining observer sensitivity. However, an alternative way to explore the role of nonlinear transduction in relative isolation is to measure neuronal responses to increasing stimulus magnitude, as Boynton et al. (1999) did for luminance contrast using fMRI. If the form of the transducer (when differentiated), reflects the pattern of detection thresholds we have observed, this would be quite convincing evidence for an explanation from nonlinear transduction. Making this comparison is the aim of the next chapter of this thesis.

3.5.3 Equal sensitivity along different global form axes

Previous work has indicated that concentric and radial patterns are more readily perceived than translational patterns (Anderson & Swettenham, 2006; Kelly et al., 2001; Nankoo et al., 2012; Seu & Ferrera, 2001; H. R. Wilson & Wilkinson, 1998). We found very little evidence to this effect, in that the fit parameters of both the hyperbolic ratio function and power-law were similar across pattern axes. Dakin and Bex (2002) argued that the enhanced sensitivity to concentric Glass patterns reported by Wilson and colleagues is purely an artefact of the circular aperture within which patterns are often presented, and that any enhancement is abolished by the use of a square aperture. While our results agree with notion of uniformity in the sensitivity to different pattern axes, we used a circular aperture, so why did we not find evidence for consistent enhancement of sensitivity to concentric form? Ours would not be the first to find evidence contradicting Dakin & Bex's proposal; an earlier experiment by Kelly et al. (2001) examined human observers' sensitivity to different Glass pattern axes using a square aperture, and still found a concentric and radial enhancement relative to horizontal and vertical. Like many existing experiments showing concentric and radial enhancement, they parameterised global form coherence using a proportion-aligned approach. Here, the salience of global form is manipulated such that a coherence of 50% means half of the dipoles in the pattern field are obeying (perfectly) the geometric rule, while the remaining dipoles are randomly oriented. Perhaps the alternative parameterisation of global form used in our experiment is responsible for the absence of polar form enhancement? One explanation could be separate contributions to global form perception from neural populations with narrow and broad local orientation pooling. Some sub-populations of neurons encoding global form along polar axes may have very narrow local orientation bandwidths, essentially encoding the degree of co-circularity within an image (elements tangent and cotangent to the centre of the dipole field). When using the proportion aligned approach as the parameterisation of global form, increasing the intensity of the stimulus would innervate these tightly tuned sub-populations. Conversely, in the present experiment, perfectly aligned dipoles are quite rare as their orientations are sampled from a Gaussian distribution. This forces observers to make behavioural responses based on more broadly tuned sub-populations, which may not possess a preference for polar

form. Future work may be directed towards exploring the possibility of distinct contributions from tightly and broadly tuned global form detectors.

3.6 Conclusion

We have shown that the detection and discrimination of Glass pattern global form does not obey Weber's law at low coherences. As coherence increases further beyond the threshold for detection, discrimination thresholds increase in a proportion resembling Weber's law. Overall, these results are consistent with a "dipper effect" in the perception of global form. This is compatible with the idea of an accelerating and saturating neuronal response function in higher-level parameter spaces beyond luminance contrast. The role of a sigmoidal neuronal transducer in the thresholds we have observed could be further explored by measuring population neuronal response across the same range of global form coherences. Unlike previous reports, we found very little evidence for an enhancement of behavioural sensitivity to any specific axis of global form, a finding which may be related to the parameterisation of global form we elected to use.

Chapter 4. Evidence of nonlinear transduction in the perception of global form

4.1 Abstract

In the previous chapter, we demonstrated that the percept of global form elicited by Glass patterns is subject to a "dipper effect". In this chapter, we ask whether the population neuronal responses evoked by the same range of Glass pattern coherences show a profile that is predictive of the facilitation and suppression we observed in behavioural thresholds. By fitting the same transducer function to the neuronal responses evoked by translational, radial, and concentric global form, we find evidence for an accelerating nonlinearity that saturates as global form coherence increases. However, the region of peak sensitivity predicted by the transducer we found does not align with psychophysical estimates, in that it occurs over a significantly higher range of coherence values. We offer explanations for this horizontal shift via the inherent limitations of EEG and from attentional modulation. We find little evidence for increased sensitivity to polar form over translational form, consistent with our psychophysical predictions.

4.2 Introduction

In the previous chapter, we demonstrated that the coherence of the global form percept elicited by Glass patterns may be subject to a "dipper" effect. Relative to the threshold for absolute detection, discrimination thresholds are facilitated at low (but detectable) pattern coherences, and suppressed as pattern coherence increases. However, the neuronal circuitry that may explain this pattern of global form discrimination thresholds is unknown, thus gaining insight into its neuronal underpinnings is the purpose of the present chapter. Hypothetical explanations for the dipper effect we have observed can be made by turning to a more thoroughly investigated visual parameter space. In the domain of luminance contrast, there are competing explanations for both the early facilitation and late suppression described by the dipper effect. The most enduring, which accounts for both facilitation and suppression, is the explanation from nonlinear transduction (Albrecht & Hamilton, 1982; Legge & Foley, 1980; Nachmias & Sansbury, 1974). This explanation proposes that contrast detection and discrimination thresholds are determined almost purely by

the shape of the input-output transducer possessed by neurons encoding luminance contrast. This explanation assumes that any contribution from internal noise (within the sensory system) is additive, invariant to the intensity of the stimulus, and applied at a post-transduction processing stage. At low contrasts, the transducer is accelerating along its output axis, meaning the increase in stimulus intensity (the input) required to produce a discernible internal response is reduced relative to detection – producing threshold facilitation. As contrast increases, the transducer begins to flatten (or saturate), so relatively large increments in contrast are required to produce a response that is equally as discernible (threshold suppression).

Importantly, alternative explanations for the dipper effect that do not require a nonlinear transducer have been proposed. For threshold facilitation, a reduction in decision uncertainty could just as easily explain why observers appear to be more sensitive at low (but non-zero) contrasts (Pelli, 1985). Here, the visual system is modelled as a collection of information channels, with each channel encoding some aspect of visual perception (size, contrast, spatial frequency, position, etc.). Keeping with the example of luminance contrast, this explanation proposes that the information available to observers indicating which channels are relevant to stimulus detection is diminished due to one of the comparators always being invisible. When discriminating two supra-threshold stimuli, all ancillary stimulus information (where the stimulus is, how big it is, etc.) is readily available, so the observer can devote resources to monitoring the channel representing luminance contrast, producing threshold facilitation. An alternative account for the suppression of thresholds relative to detection (the “handle” of the “dipper” function) argues that a saturating transducer with additive noise is equivalent to a linear transducer with noise that is proportional to the stimulus intensity (Kontsevich et al., 2002).

Ultimately, it is difficult to disentangle competing explanations using a pedestal vs threshold experiment (as we used in the previous chapter), as they predict the same behaviour - just via different models. A different approach to this problem is to measure neural responses in human observers (via M/EEG or fMRI) across a range of stimulus intensities and examine the degree to which the measured response function predicts sensory discrimination. If we accept assumptions on the coupling between neural response magnitude and behaviour, this is equivalent to measuring the neuronal

transducer itself, instead of its behavioural derivative. For luminance contrast, Boynton et al. (1999) performed this experiment using fMRI, and found a pattern of BOLD responses in primary visual cortex that mirrored contrast discrimination thresholds obtained from the same observers, offering convincing evidence that the contrast dipper effect is due to transducer nonlinearity. The present experiment applies a similar methodology, but to the facilitation and suppression we uncovered in the previous chapter for the global form coherence of Glass patterns. Using electroencephalography (EEG), we aim to investigate (for the first time to our knowledge) whether this non-monotonicity in global form discrimination thresholds is reflected by an accelerating and saturating nonlinearity in neuronal responses.

Previous work investigating neuronal responses to Glass patterns has mainly focused on answering two questions: where is global form information represented in the human visual system, and do global form selective regions show a preference for a certain pattern axis? The fMRI work conducted by Ostwald et al. (2008) was amongst the first to address these questions in human observers by exploring blood-oxygen level dependant (BOLD) responses to translational, radial, and concentric Glass patterns. Using multivariate pattern analysis, Ostwald and colleagues demonstrated that low to mid-level retinotopic regions (V1 – V4) and lateral occipital cortex (LOC) all carry information that can accurately predict the presence of Glass patterns, a finding also supported by the later fMRI work of Mannion et al. (2010). Ostwald et al. also found a general bias in early retinotopic areas towards improved classification of radial patterns over other types. More importantly, they found that the LOC was the only region to show a bias towards the classification of concentric form. The LOC also demonstrated a significant difference in classification accuracy when comparing local and global shifts in dipole orientation, a finding that could provide a neurological basis for the behavioural enhancement often found for concentric Glass patterns (P.-C. Huang & Chen, 2016; Seu & Ferrera, 2001; H. R. Wilson & Wilkinson, 1998; H. R. Wilson et al., 1997).

Generally, the finding of enhanced responses to radial and concentric form over translational form is common in the Glass pattern literature (Pei et al., 2005; Rampone & Makin, 2020; Swettenham et al., 2010). However, we propose that the tendency for existing research to measure responses to maximally coherent patterns and Glass

pattern detection thresholds precludes a holistic understanding of human observers sensitivity to global form. Indeed, the discrimination of supra-threshold intensities is just as fundamental an operation as detecting the absolute presence of a stimulus. As in the psychophysical domain, very little work has been done to investigate the gradient of neuronal responses that encodes degrees of Glass pattern global form, and how such gradients may differ across different pattern alignment axes. One experiment by Mannion et al. (2013) investigated neural responses to a range of Glass pattern global form intensities, but did not distinguish between pattern types, or make any behavioural comparisons. As with most of the preceding psychophysical work (apart from Dakin (1997) and ours), Mannion et al. parameterised global form as the proportion of perfectly aligned dipoles to randomly oriented dipoles. They found a mostly linear gradient of responses for increasing alignment proportions in V3, and in other mid-level Glass pattern responsive regions, though they did not measure any behavioural thresholds for comparison. To our knowledge, no published work has measured neural responses to a gradient of global form coherences and compared them with psychophysically obtained detection and discrimination thresholds. Therefore, the present experiment aims to extend our psychophysical work and test the predictive power of psychophysically obtained thresholds for neural responses obtained via EEG. We do so by measuring steady-state visually evoked potentials (SSVEPs) to translational, radial, and concentric Glass patterns across the same range of coherences we have tested psychophysically. By fitting the same transducer model to neuronal responses as we did to detection and discrimination thresholds (in differentiated form), we investigated the degree to which their parameter estimates correspond.

4.3 Methods

4.3.1 Participants

Twenty-one participants were recruited for this experiment (10 females, 11 males, mean age = 28, standard deviation = 9) from the Department of Psychology's participant pool at the University of York, UK. Participants had normal or corrected-to-normal vision, no known disease or disorder of the eye or visual system, no severe hearing impairments, and no personal history of epilepsy, cognitive impairment, or dementia. All experimental protocols were approved by the Department of Psychology ethics committee at the University of York.

4.3.2 Visual Stimuli

4.3.2.1 Spatial parameters

Stimuli were generated in the Python (https://www.python.org/) programming language using routines from the PsychoPy toolbox (https://www.psychopy.org/). Stimuli were presented to observers in a dark room via a ViewPixx 3DLite LED display (1920x1080, 120Hz) using an Apple Mac Pro running OSX 3.5. To ensure precise frame timings, graphical processing overhead was minimised by rendering all frames at an internal resolution of 1280x720, but upscaling them to the native resolution of the display. The display was gamma-corrected via a luminance look-up-table acquired using a Minolta LS100 photometer and had a (post-correction) mean luminance of 41cd/m^2. The visual stimuli, illustrated in Figure 4.1, were a field of 62 full contrast dot-pairs (dipoles), with a 50/50 split of black and white pairs, and all spatial parameters matched to our earlier psychophysical work (Chapter 3, Methods). Dipole fields were presented to observers at a distance of 100cm within a circular field subtending 10° of visual angle, with an areal density of 3.9% (1.6 pairs/deg^2). Within a dipole, each dot had a diameter of 10.02 arcmin and a centre-to-centre separation of 19.2 arcmin. To match our psychophysical experiment, a thin grey ring was always present around the dipole field to reduce spatial uncertainty. The distance between dipole centres was at least 45.6 arcmin, which was ensured for each unique field by continuously generating random dipole-centre coordinates and keeping values that were lawfully spaced and rejecting those that were not until a sufficient number of coordinates were generated (a method often referred to as rejection sampling).

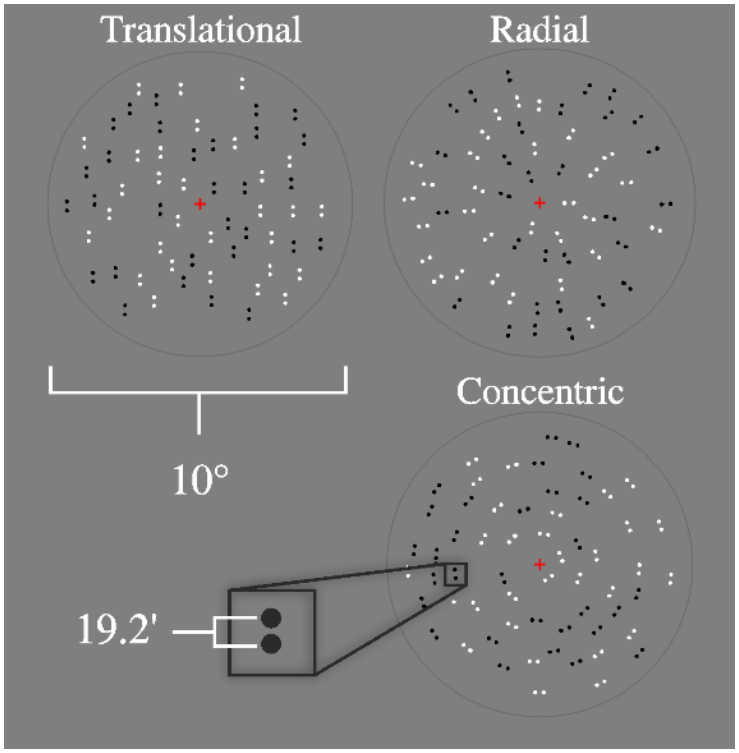

Figure 4.1: Glass pattern spatial specification
An example of the Glass patterns used in this experiment; Translational, Radial, and Concentric. Spatial parameters are annotated where relevant.

4.3.2.2 Parameterisation of global form

To generate dipole fields of different global form coherences, the angular orientation of each dipole in a field was sampled from a semi-circular von-Mises distribution, which produces a Gaussian probability distribution function (PDF), defined in radians.

$$vM_{pdf}(x \mid \kappa, \mu) = \frac{e^{\kappa \cos(x-\mu)}}{2\pi I_0(\kappa)} \quad \textbf{(4)}$$

The equation for the PDF of a von-Mises distribution is provided in equation 4. Here, *kappa* is a specificity parameter, analogous to the reciprocal of the variance. A *kappa* of zero produces a uniform distribution, while progressively positive values symmetrically increase the probability density around the angle *mu*, which is set to zero. Sampling from this distribution returns values spanning 2π, which are halved, wrapping them on the semi-circle instead ($-\pi/2 > \pi/2$). This half-wrapping was performed because each dipole is oriented along a line, but has no direction (i.e., for a radial dipole, there is no distinction between an inwards and outwards heading). In the end, this produces an array of 62 semi-circular angles (one per dipole), centred on zero, with an angular specificity dependant on the value of *kappa*. This array is then

used to displace a 2nd array of 62 dipole angles that are perfectly aligned along the desired axis. Therefore, as *kappa* increases, so does the global orientation alignment of the dipole field. Dipole fields were generated at 8 pre-defined *kappa* levels spanning the range of our psychophysically sampled discrimination thresholds; zero (chance-level coherence), six values spaced in 5dB steps, starting from and including 0.30, and maximally aligned patterns (i.e., $k = \infty$). As the number of dipoles in the patterns used in this experiment was relatively low, all patterns generated were assessed using a Kolmogorov-Smirnov test (against $\alpha = 0.05$) for the equality of the sampled orientations with the theoretical von-Mises PDF. Any distributions failing this test were rejected until a valid distribution was achieved.

$$Global\ coherence = \left(\int_{-10^{\circ}}^{10^{\circ}} vM_{pdf}(x, kappa, 0) \right) \quad \textbf{(5)}$$

To allow for direct comparison, we represent global form in in the same units as our psychophysical work by converting *kappa* values to a bounded range that we refer to as "global form coherence". These transformed values were obtained by integrating the von-Mises PDF at all values of *kappa* across a 20° range centred on zero, and treating the finite probability yielded by this integration as a measure of global form intensity. This value describes the probability of sampled dipole orientations falling within 10° of perfect alignment, and we treat it as though it faithfully equates this proportion, which is reliable assumption due to the use of a K-S tests to filter out aberrant distributions. The 20° integration window was chosen based on non-human primate literature indicating that this is the lower-bound on the full-width half-maximum bandwidth of neurons in V1 with a receptive field size that covers a dipole width (Gur et al., 2004). This calculation assumes that all dipoles within the 20° integration range will equally affect the neuronal response. This is an oversimplification, but we believe it is sufficient for the relative comparisons performed here. This transformation was performed using equation 5, whose kappa-coherence transfer function is provided in the previous chapter (Figure 3.3).

4.3.2.3 Temporal parameters

Dipole fields were presented to observers as dynamic Glass patterns (dGPs), which are distinct from static Glass patterns in that they are temporally modulated at two different frequencies. The first is a global form switch between a field of dipoles

with orientations drawn from a uniform distribution (i.e., chance-level coherence), and a field of dipoles with orientations drawn from a distribution with above-chance coherence (using the levels defined in the previous section of this chapter). The second frequency of modulation is a constant (and faster) resampling of these two distributions. In theory, the global form switching will elicit responses from neuronal populations encoding global form correlations, while the local switching will capture responses that encode local parameters, such as dipole positions and their local orientations. In the present experiment, we adopt the temporal design of Palomares et al. (2012), such that our global-form switching occurred at 0.83Hz, and distributions were continuously locally resampled at 30Hz. From now on, we will use the following convention when referring to these input frequencies: "1F1" refers to 0.83Hz global form switching, and "1F2" refers to 30Hz local resampling. When referring to harmonics of stimulation frequencies, the order of the harmonic in reference prefixes the "F", such that 1F1 and 2F1 refer to the fundamental frequency and the 2nd harmonic of the global switching, respectively. Both 1F1 (global) and 1F2 (local) were square wave modulations with a 50% duty cycle, and the phase of 1F1 was set such that the observer saw random orientations in the first half-cycle of global form switching. To present pattern switching with as little computational overhead as possible, dipole coordinates and orientation distributions for each stimulus frame were pre-generated and loaded into memory before any stimuli were shown to the observer.

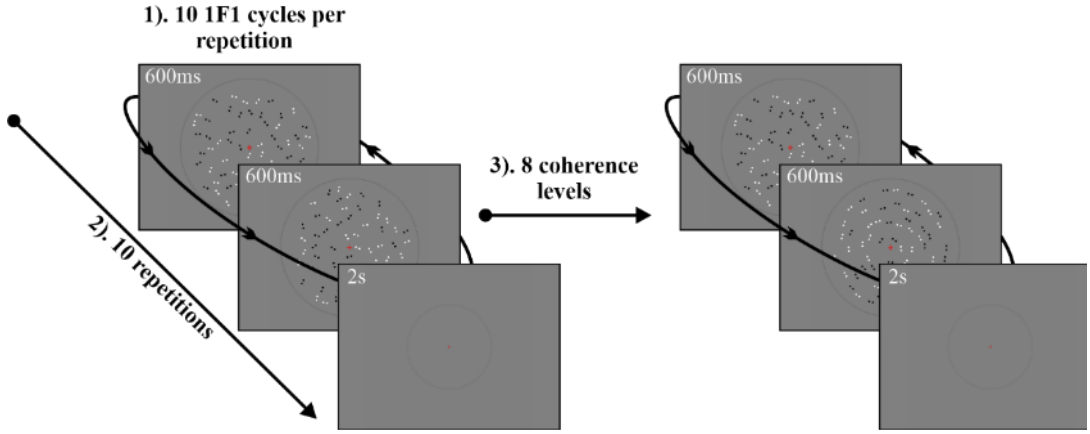

Figure 4.2: Presentation protocol for a single pattern condition
An illustration of the experimental design for a single pattern alignment axis (concentric in this example). Within each repetition, observers saw 10 1F1 stimulation cycles. There were 10 repetitions per coherence level.

4.3.3 Experimental procedure

Upon their arrival, observers were given an overview of the experiment, and a brief demonstration of the stimuli. Once settled, they were fit with an appropriately sized EEG cap that was aligned to their nasion, inion, and pre-auricular points according to the 10-20 system (Klem et al., 1999). For the duration of the experiment, observers sat at a distance of 1 metre from the display on a comfortable chair - no chinrest was used. In separate blocks for translational, radial, and concentric pattern axes, observers viewed dGPs across all pre-defined global form coherence levels. The procedure carried out for a single axis condition is illustrated in Figure 4.2. For a single coherence level, there were ten repetitions, with each repetition containing ten 1F1 cycles (12s of stimulation), with each repetition separated by a 2 second break period. Therefore, across all coherence levels observers experienced eighty 12s stimulation trains for each pattern axis – 240 for the whole experiment. The "pypixxlib" toolbox from ViewPixx was used to send transistor-transistor logic (TTL) triggers to the recording software immediately after the first and last frame of each 12s repetition, which were later used to split continuous records into epochs. Observers were encouraged to avoid blinking, jaw clenching, swallowing, and head movement during visual stimulation, and to restrict these activities to the 2s break period. Every 20 repetitions, and in-between each pattern condition, observers were given a longer self-paced break to allow them to stretch, rest their eyes, or take a drink. The order in which coherence levels were presented was randomised within each pattern axis, and the order of pattern axes was counterbalanced across observers. For most individuals, the experiment lasted for approximately an hour and fifteen minutes. Due to COVID-19 restrictions, the experimenter could not stay in the room during data acquisition, but was monitoring the observer via a video/audio link from the next room, checking-in with them after each condition, and whenever they required assistance.

4.3.4 Data acquisition and pre-processing

Electroencephalograms were recorded at a sampling rate of 1kHz using an antNeuro "asalab" system: a 64-channel amplifier (+ bipolar channels) and "waveguard" EEG cap with 64 electrodes (https://www.ant-neuro.com/). Cap electrodes were monopolar, using electrode FCz as the online reference, and vertical/horizontal bipolar electrooculograms (EOGs) were recorded to later aid offline artifact rejection. All pre-processing was carried out offline using custom code

and/or routines from the EEGLAB toolbox where specified (Delorme & Makeig, 2004) in MATLAB version 2018b. The following process was repeated for all observers. First, to allow the entire dataset to fit into computer memory, all continuous timeseries were down-sampled to 250Hz (factor 4 decimation), and then further filtered using a 0.1Hz high-pass filter and a 45Hz low-pass filter. As a first pass of quality control, the power-spectrum of each channel was calculated and manually inspected to reveal any records with channel-wide data quality issues (i.e., high power relative to other channels with a smooth roll-off). Three observers were rejected at this stage due to a presumed reference fault that produced extremely high power at all channels across a broad range of frequencies.

Further artifact rejection was carried out using EEGLAB's implementation of infomax independent components analysis (ICA), guided by the semi-automated selection of ICA components toolbox (SASICA - see Chaumon et al. (2015)). Initially, ICA weights were trained on the continuous EEGs, and the routines of SASICA were used to identify high-ranking single-channel components that exhibited a low correlation with neighbouring channels. The generative channel for these components were rejected as artefactual, as plausible EEG signals are necessarily correlated with nearby channels (due to volume conduction). Rejected electrodes were then replaced using EEGLab's spherical spline interpolation method (Ferree, 2006), and ICA weights were then recalculated on the remaining channels (with rank adjusted accordingly). The remaining components were then screened for high correlation with vertical and horizonal EOG signals (again using SASICA). Usually, the first few components were highly correlated with the EOGs, and had a topography consistent with eye blinking and eye movements. However, any component that was highly correlated with the EOGs was removed from the data.

The resultant "clean" records were then split into epochs. Using the TTL trigger timings, the first 1.2s of each 12s stimulation repetition were removed, and the remaining samples broken down into 3.6s epochs (yielding three epochs per repetition). We removed the first 1.2s of stimulation to avoid the contamination of steady-state responses by transient pattern onset responses and eye-blinks or movement that occurred just prior to pattern onset. Then, a discrete Fourier transform was performed on the samples of each individual 3.6s epoch, yielding a Fourier series

with frequency bins 0.27Hz (1F1 / 3) in width. This series was clipped at the Nyquist limit (125Hz). In the end, this produced a 450 x 63 x 240 x 8 array (Fourier coefficients x channels x epochs x coherence level) for each observer, and for each pattern axis condition. To ensure we are fitting models to meaningful signals, we then applied a final screening using the cross-channel average of complex valued Fourier coefficients. Any observers who did not show a significant 1F1 response at maximum pattern coherence across conditions were removed from further consideration. This screening was performed based on the outcome of a t2-circ test (Victor & Mast, 1991) against an alpha level of 0.05, which removed 5 observers from the experiment. The timeseries from the remaining thirteen observers were then submitted for further analysis.

4.3.5 Reliable components analysis

To produce a channel-wise spatial filter that describes cross-observer responses to global form switching and local resampling, the Fourier coefficients of frequencies of interest for each epoch, channel, and observer were submitted to frequency-domain reliable components analysis (RCA – Dmochowski et al., (2015)). Like principal components analysis, frequency-domain RCA reduces high dimension data (in this case 63 EEG channels) to linear channel combinations, with each combination ranked in terms of the value of an objective function. However, where PCA would conventionally find a spatial filter that maximises within-epoch covariance, RCA instead produces a filter that maximises the ratio of cross-epoch covariance to within-epoch covariance, exploiting the assumption that the amplitude and phase of SSVEPs are consistent across trials due to the periodicity of the stimulus. We used RCA to generate a group-level channel spatial filter by submitting the within and across epoch spatial covariance matrices accumulated across observers to the RCA algorithm (this accumulation was performed using the MATLAB toolbox for RCA; https://github.com/svndl/rcaBase). A global form response filter was calculated for each pattern axis condition separately, and a local resampling filter was calculated across all pattern conditions. The global filter was generated used the Fourier coefficients of multiple harmonic frequencies (1F1, 2F1, 3F1, 4F1 and 5F1), from 7 above-chance global form coherence values, while the local filter used only the coefficients of 1F2 – but was trained across all 8 global coherences values, including the zero baseline (as the 1F2 switching is present in these epochs). In summary,

multiple group-level RCA computations were used to produce three separate global component hierarchies (one for each pattern axis), and a single local component hierarchy (across all pattern axes). These "reliable components" (RCs) can then be visualised in the original electrode space by projecting the within-epoch covariance matrices through an RCs spatial filter weighting. The scalp topography produced represents the activation expected were only the neural generators of a given RC to be present in the sensor records (Haufe et al., 2014; Parra et al., 2005). For transducer fitting, individual observers' time-domain sensor data for each coherence level were projected through the group-level filters, yielding a weighted sum of sample amplitudes at each coherence level, for each reliable component. These values were then used to determine the signal to noise ratio (SNR) at each global form coherence level, for each pattern axis.

4.3.6 Signal-to-noise ratio calculation

To control for cross-observer variance in the baseline noise level, we transformed RCA filtered time-domain data to SNR. For each observer, coherence level, and pattern axis, RCA filtered time-domain data were Fourier transformed, the complex valued Fourier coefficients at frequencies of interest extracted, averaged across all 3.6s epochs, and their moduli transformed to amplitudes (via magnitude/DFT length). Then, to produce a pooled measure that captures all harmonic responses to 1F1 global form switching, the root-mean-square (RMS) amplitude across the harmonics used to train RCA filters (1F1 > 1F5) was calculated (Cottereau et al., 2011).

$$SNR = \frac{RMS_{signal}}{RMS_{baseline}} \tag{5}$$

To control for cross-observer variance in baseline noise, a signal-to-noise ratio was calculated at the individual observer level by dividing the RMS amplitude of above-chance global form coherence ($k > 0$) by the RMS amplitude of responses to chance-level global coherence ($k = 0$), shown in equation 5. A similar process was carried out using data projected through the cross-condition local resampling RCA filter, but only 1F2 was used, so no harmonic pooling was necessary. Unlike global form responses, the denominator (i.e., noise) used when calculating the SNR of local responses was the RMS of the five frequency bins adjacent to 1F2.

Prior to model fitting, we directly compared the SNR values of different components using estimation statistics (Claridge-Chang & Assam, 2016), whereby significance is assessed by generating 95% confidence intervals on the differences between median component SNR values. For a given comparison, confidence intervals on the difference in medians are generated by resampling (with replacement) SNR values from each comparator 10000 times (using "bootci" in MATLAB), and each time taking the median of each comparator, and differencing those medians. Unless otherwise specified, intervals represent the 95% bias-corrected accelerated (BCa) bootstrap confidence interval (DiCiccio & Efron, 1996). Where this confidence interval on the difference between two conditions does not contain zero, a significant difference can be concluded.

4.3.7 Model fitting

The median of observers' SNRs for each global form coherence were fit with the same hyperbolic ratio function which we used to fit global form discrimination thresholds in the previous chapter, allowing for the direct comparison of their parameter estimates.

$$Response = rMax * \frac{c^k}{\left(c^k + c_{50}^k\right)} + rMin \quad (6)$$

Unlike in Chapter 3, the *rMin* parameter was specified, and fixed to an SNR of 1 (i.e., the noise floor), as shown in equation 6. As before, the *rMax*, *c50*, and *k-exponent* parameters were optimised using a loosely constrained Nelder-Mead simplex search algorithm that minimised the sum of squared error (SSE) between the model and median SNR values. This fit was then entered into a bootstrapping procedure where, for each iteration, the cross-observer SNRs for each coherence level were resampled with replacement, the median taken, and a new hyperbolic ratio fit produced. The resultant 10000 x 3 matrix of bootstrapped fit parameters was then used to generate a 10000 x 8 matrix of SNRs. The 2.5th at 97.5th percentiles of this final matrix were used to produce upper and lower-bound curves that constitute 95% confidence intervals on the median model fit. As a brief reminder, the *rMax* parameter describes the maximum response of the neural population, the *c50* describes the coherence level at half of *rMax*, while the *k-exponent* describes the shape of the transducer function around the *c50*. Exponents above 1 are indicative of an

accelerating and saturating nonlinearity (the theoretical generator of the "dipper effect"), while exponents at or below 1 are indicative of only a saturating nonlinearity. A *c50* beyond 100% coherence is possible, and it indicates that the neural response showed little evidence of response saturation. In this situation, it's possible that a different function may be a better descriptor of the data. The bootstrapped distributions of model fit parameters were also compared using estimation statistics (Claridge-Chang & Assam, 2016).

4.4 Results

4.4.1 Grand-average responses to global and local switching

We first present an overview of grand-average responses prior to RCA filtering, as this cross-channel average informed the number of 1F1 harmonics to be used for RCA. These data are presented in Figure 4.3, which shows the timeseries and amplitude spectrum of full coherence 1F1 global form switching averaged across all conditions, observers, and channels. Figure 4.3A shows this data in the time-domain, where there is a clear negative peak approximately 150ms after the switch to a full coherence Glass pattern (the dashed square-wave shows the 1F1 stimulation model). This signal is markedly asymmetric, resembling a positively clipped sine wave. As such, both odd and even harmonics of 1F1 are present in the Fourier representation of this signal, as evidenced in Figure 4.3B, where both global and local switching harmonics are denoted by the black bars and inlaid text, the latter after the axis break. We calculated the T2circ statistic (Victor & Mast, 1991) on the Fourier coefficients of harmonics 1F1 through 8F1, and found significant responses up to and including 5F1 ($p < .05$, FWE corrected; Benjamini and Yekutieli (2001)), and at 1F2 (the local resampling frequency). Evidently, responses to global and local form switching have been captured at these harmonics, so they were selected as the input frequencies for RCA dimension reduction.

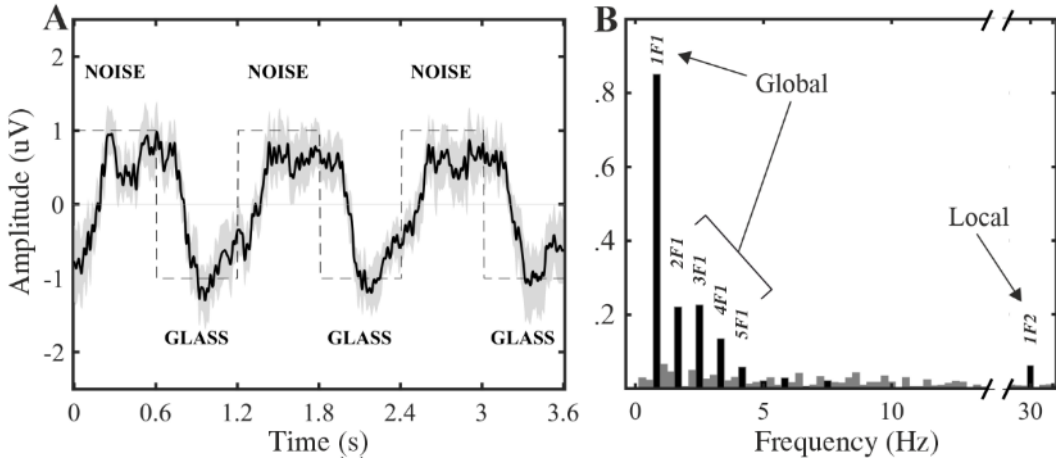

Figure 4.3: Grand average sensor timeseries and spectrum

A: Grand average timeseries. The solid line represents the grand-mean waveform, while the grey shading represents the bootstrapped 95% CI on the mean (n=10000). The dashed square-wave overlaid represents the temporal profile of 1F1 global form switching stimulus. B: The Fourier series of A, with stimulus harmonics labelled appropriately. The abscissa has been broken to show the response at 30Hz.

4.4.2 Spatially extensive responses to global form

To allow group-level assessment of a single scalar value that collectively describes scalp responses at channels showing consistent phase-locked responses, group-level RCA global form switching filters were calculated for each axis of pattern alignment, and separately for the 1F2 dipole resampling rate across pattern axes. The proportion of the total trial-to-trial covariance recovered by the first six components are presented in Figure 4.4. For global form RCs, the first component accounts for 79 - 81% of the cross-trial covariance, the 2nd component only explains approximately 12 – 14%, while the third component accounts for very little (3 – 6%). However, the local component does not show such a steep decline, with the 1^{st,} 2nd, and 3rd components capturing 58%, 28%, and 19%, respectively. Across all filters, the proportion of covariance recovered is uniformly low by the 4th component – a maximum of 6% for the 1F2 filter. Therefore, we focus our analysis on the first three components of each RCA hierarchy.

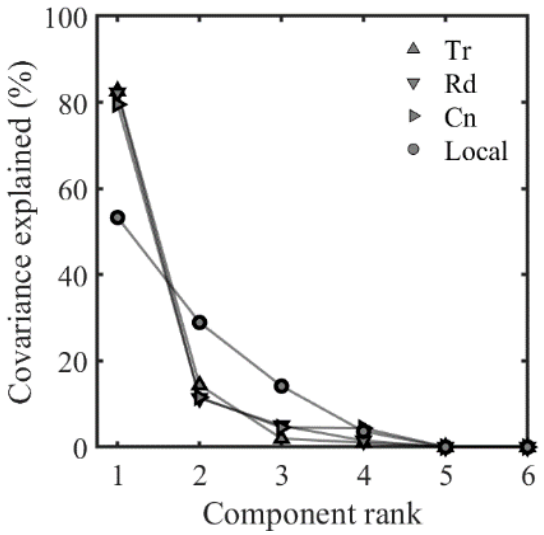

Figure 4.4: Hierarchy of reliable components

The proportion of trial-to-trial covariance in Fourier coefficients recovered by each RCA component (i.e., each spatial filter) trained separately on global and local frequency harmonics. Markers represent the condition that was submitted to calculate RCA filters; 1F1 – 5F1 global form harmonics (Tr: translational, Rd: radial, Cn: concentric) and 1F2 local resampling.

To visualise the components recovered by RCA, we present their projected scalp topographies in Figure 4.5. For the first component of filters trained on 1F2 local resampling, we find a slightly left-lateralised occipital topography with peak responses at electrodes O1 and PO7, and with convex contour lines encapsulating parieto-occipital electrodes. In comparison, for the first component of all global form switching filters, we find a topography that includes that of local resampling, but with peak responses that extend to more anterior parietal electrodes with more concave contour lines, suggestive of more lateralised responses. These topographies are consistent with those found by Palomares et al. (2012), and the notion that information pertaining to global form correlations is represented in early and mid-level processing regions (Mannion et al., 2010; Ostwald et al., 2008). For global form, there is still an RC2 topography over parieto-central and parieto-occipital electrode sites that is consistent across pattern conditions. Although it does not account for a great deal of covariance, this topography resembles the 2nd harmonic topography found by Pei et al. (2005). While we do not distinguish harmonics in the present experiment, it is interesting to note that these authors found that the activation of this topography was uniform across pattern axes and randomly oriented dipole fields. By RC3, cross pattern-axis consistency is lost for global form filters, suggesting that they reflect noise unrelated to the stimulus, especially when coupled with the low proportion of covariance recovered by this component. For the local resampling projections, RC2

and RC3 still explain a significant proportion of cross-trial covariance and are heavily lateralised. It is possible that these topographies represent some contribution from areas sensitive to the illusory motion percept elicited by dynamic Glass patterns, such as V5/hMT (Krekelberg et al., 2005). For both local and global RCs, we compare the SNR of each component across coherence levels in the next section.

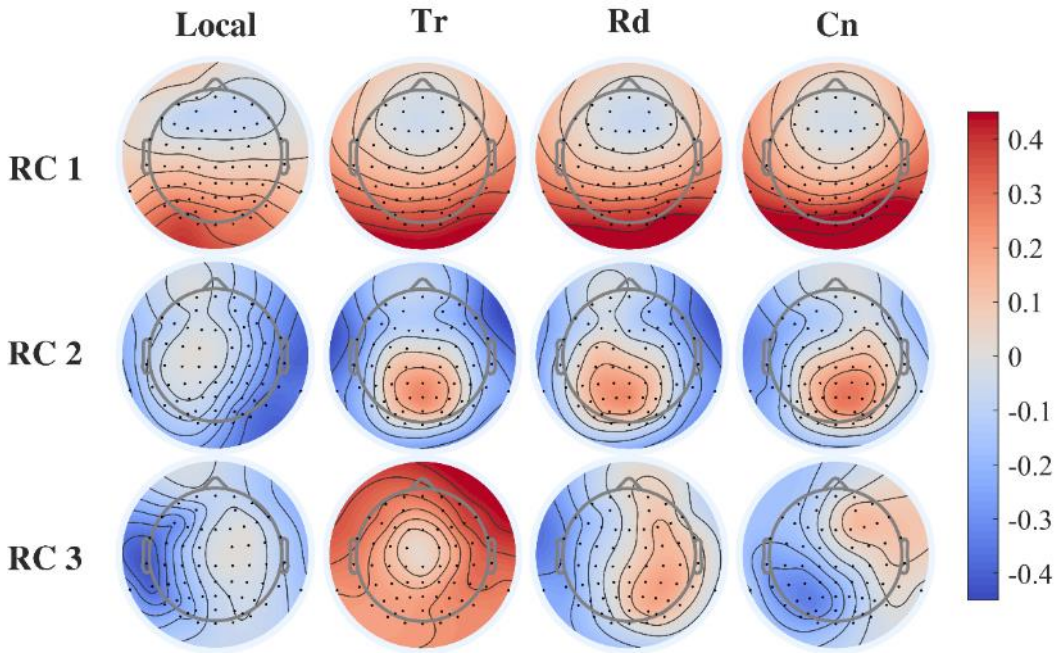

Figure 4.5: Topographies of reliable components
Scalp projection of RCs one, two, and three for the local and global stimulation harmonics. Local RCs are based on all conditions, global RCs are generated on a per-condition basis. Tr = translational, Rd = radial, and Cn = concentric. The heatmap is in the arbitrary units of the RCA forward model, and uses the "coolwarm" diverging colourmap (Moreland, 2009).

4.4.3 Responses to local resampling

Observers' sensor data was projected through the weightings of the group-level spatial filter for the first three RCs, enabling the assessment of component signal-to-noise ratios. We first describe the projected values generated from the local filters, and then shift focus solely to the global filters, as they are of prime interest. As shown in Figure 4.6A, all coherence levels, for all local resampling RCs had 95% confidence intervals that did not contain baseline noise, indicative of significant 1F2 responses. As expected, the SNR for the local switching (30Hz) response did not vary with global form coherence for any RC, as evidenced by bootstrapped contrasts carried out between each unique combination of coherence levels, for each component (see Appendix E). Averaging across coherence levels, the more occipital local RC1 filter showed the highest SNR (median: 4.1, 95% CI: [3.8, 4.3]), followed by RC2 (median: 3.0, 95% CI: [2.7, 3.7]), and RC3 (median: 1.8; 95% CI: [1.5, 2.0]). In Figure 4.6B,

contrasts on the distribution of SNRs across coherence levels are shown, indicating that the SNRs of all RCs were distinct from one-another, as all 95% confidence intervals on the differences in medians did not contain zero, with the largest difference being between RC1 and RC3 (median Δ: 2.3, 95%CI: [2.0, 2.6]). Overall, these results suggest that RCs one through three all show significant 1F2 activity, that the more occipital RC1 was most responsive, and that 1F2 responses were static across coherence levels.

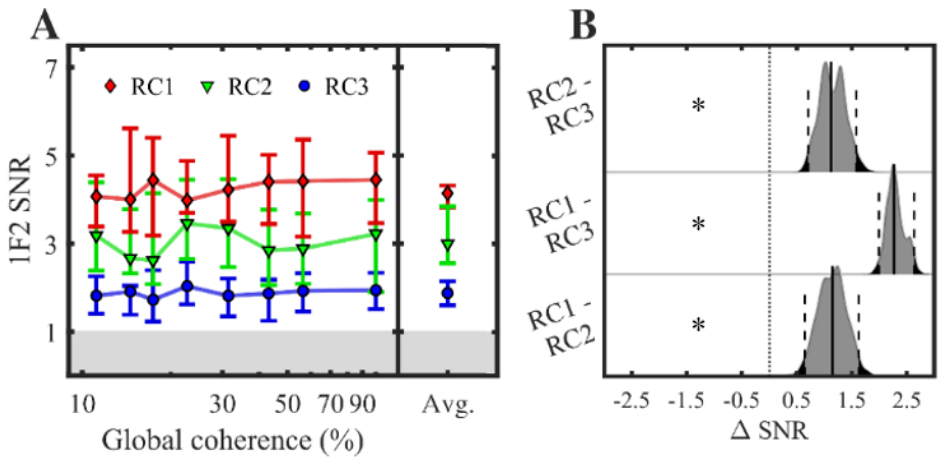

Figure 4.6: SNR of local RCs across global form coherence levels
A: Local component SNR as a function of global form coherence collapsed across all pattern axes. Line and marker colour denotes the reliable component (RC). The grey shaded region illustrates the noise floor (SNR = 1). The median across all coherence levels is provided in the right-most section of A, and is contrasted in B. B: Contrasts on the median SNR across coherence levels. Shaded curves show kernel density estimates generated by bootstrapping the difference in medians. The vertical solid line shows the empirical difference in medians, while the dashed lines show the bootstrapped confidence interval (CI) on this difference. Where CIs do not cross the zero a significant difference is implied (marked by stars) .

4.4.4 Lateralised responses scale with global form coherence

Having established that peak local responses were static across coherences levels, we now report responses to global form switching for each RC. Now, SNR is defined as detailed in equation 5. In Figure 4.7A, the median SNR, and 95% CIs on those medians, for each RC are shown for each pattern axis. The markers inlaid at the top of each panel denote SNR values with confidence intervals that do not contain baseline noise. A significant change from baseline is implied where 95% CIs do not cross the baseline threshold (SNR=1, grey region). RC1 appears to show a response that grows with global form coherence, as does RC2 to some extent in the translational condition, but with a general tendency towards much lower SNR. Some RC2

activation is not unexpected, as RCA components are not necessarily orthogonal, and RC1 shares some topography with RC2.

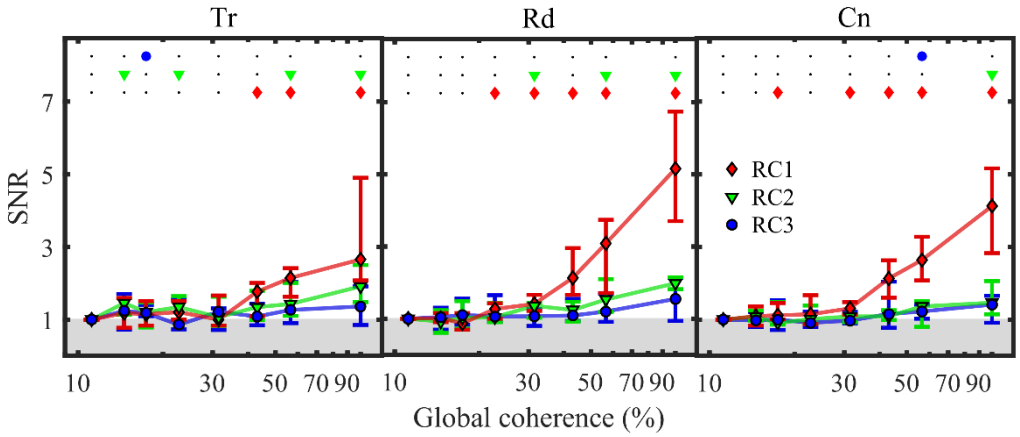

Figure 4.7: SNR of global RCs across global form coherence levels and pattern axes

For each subplot, different line colours represent the three reliable components (RCs) we compared. Error bars represent 95% confidence intervals. The shaded grey region represents the noise floor (an SNR of 1). The smaller markers at the top of each subplot indicate where the SNR was significantly larger than 1 (black dot = N.S.). Tr: translational, Rd: radial, Cn: concentric

In Figure 4.8 we illustrate the unique contribution from RC1 by performing contrasts between the SNR values of each RC, for each coherence level. These comparisons indicate that, when comparing RC1 with RC2 and RC3, a relative increase in RC1 SNR starts at 30 – 40% coherence, which continues to grow up to 100% coherence for all pattern conditions. However, when comparing RC2 and RC3, no such difference is encountered, having SNR values that are quite similar despite their distinct topographies. As RC1 shows a response that scales with global form, and shows SNR values that significantly differ from other RCs, we focus our model fitting analysis on sensor records that have been projected through this filter.

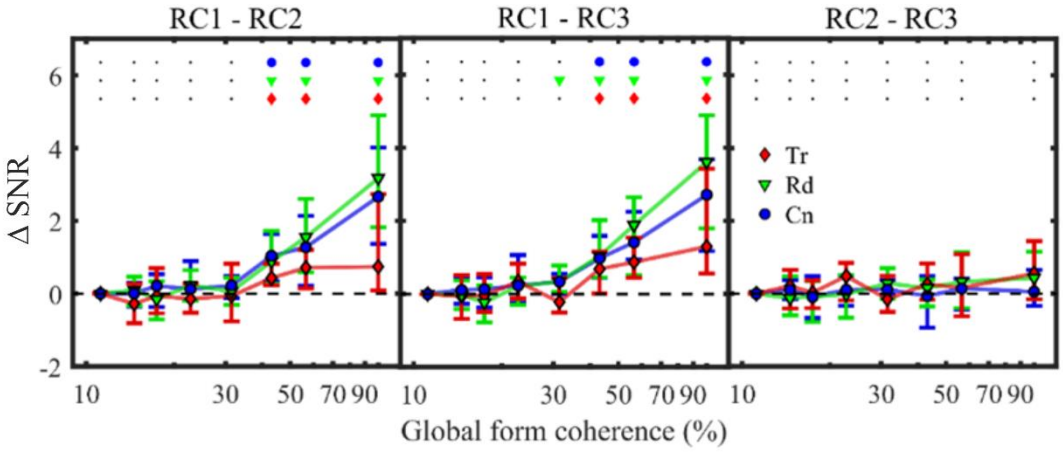

Figure 4.8: Contrasts of SNR between each RC, across coherence and pattern axes

Each panel represents a reliable component contrast (see headings). For each panel, different line colours represent different pattern axes, and the dashed horizontal zero line represents a difference of zero between conditions. Where error bars (95%CIs) do not contain this reference line, a significant difference in SNR between RCs is implied, and is represented by the smaller markers inlaid at the top of each plot (a black dot = N.S.). Tr: translational, Rd: radial, Cn: concentric

4.4.5 An accelerating nonlinearity independent of pattern axis

In Figure 4.9, global RC1 median SNR is plotted as a function of global form coherence. We have changed the abscissa to a linear scale, to illustrate any sigmoidal (accelerating/saturating) nonlinearities more clearly. Also shown is the fit of a hyperbolic ratio function, optimising the *c50*, *rMax*, and *exponent* parameters. Qualitatively, responses appear to be dynamic over a wide range of coherence values, showing some indication of mild response saturation, though a horizontal asymptote is never clearly achieved in the radial and concentric condition. In the translational condition, this saturation is more distinct, and has responses to maximum coherence that are generally lower than that of the other conditions. While some saturation is apparent in the median SNR of all conditions, some bootstrapped fits returned a c50 parameter beyond 100% coherence. This indicates that some combinations of observers were best fit by a function that shows an absence of saturation, evidenced by the shape of the upper limit on the confidence interval in all conditions. The proportion of fits with a c50 > 100% was 22% for the translational and radial condition, and 15% for the concentric. This suggests that a saturating nonlinearity can describe the pattern of responses in the majority of observers, but that some may be better described by a different (perhaps linear) form of transducer.

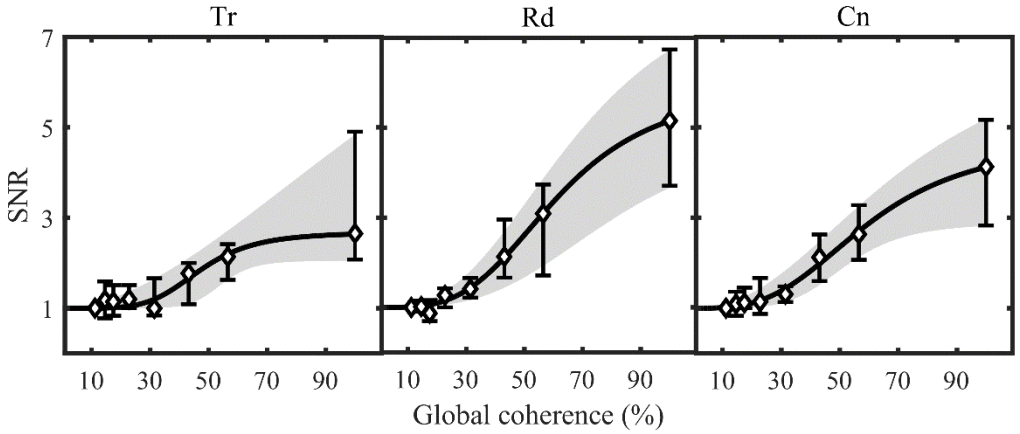

Figure 4.9: Bootstrapped model fits to median SNR

Median SNR from baseline for data projected through RC1 for each pattern condition. The solid line is the best fit of a hyperbolic ratio function to the median SNR, while the shaded grey region around the line is the 95% CI on this model fit. Tr: translational, Rd: radial, Cn: concentric. Individual datapoints are omitted for clarity, but are available in Appendix F.

When the c50 is beyond 100%, the value of the other parameters describes a function that is beyond the possible range of stimulus values, so we have removed bootstraps that contain c50s > 100 from consideration. This means that all forthcoming parameter contrasts are performed based on the lowest number of bootstrapped hyperbolic ratio fits with c50s < 100% (n = 7784). The distribution of remaining bootstrap parameter values is shown in Figure 4.10 for each pattern condition. First, that the median *exponent* values are all above 1 (and that they do not cross 1 in their confidence intervals) is consistent with an accelerating nonlinearity followed by saturation in the majority of observers. The exponents are quite similar for radial and concentric responses, showing median values of 3.45 (95% CI [2.5, 5.8]) and 3.15 (95% CI [3.2, 8.7]), respectively. However, the translational condition shows a great deal of variation in the exponent parameter, having an upper confidence interval that reaches the simplex fitting constraint (median: 5, 95% CI: [1.6, 10]). Such exponent values occur when responses traverse the *rMin* to *rMax* delta over a very brief range of coherence values, and these extreme fits are visible in the lower CI of the model fit for this condition (Figure 4.9). The differences in exponent values between pattern conditions (Figure 4.10B, right-most panel) were within the margin for error in all cases, though it is clear that the variability of translational exponent fit is likely to mask any subtle differences when making comparisons with this condition. Extreme exponent values do not preclude the assessment of other parameters, where we continue to see a trend towards the translational condition showing exceptional estimates. The c50 distributions are quite similar for radial and concentric responses,

showing median values of 60% (95% CI [43, 95]) and 62% (95% CI [53, 90]), while the translational condition reaches half-saturation at a lower coherence, on average (median c50: 47%, 95% CI [40, 83]). However, as shown in Figure 4.10B, cross-condition contrasts of the *c50* indicate that this difference was not significant, as its confidence intervals all cross zero.

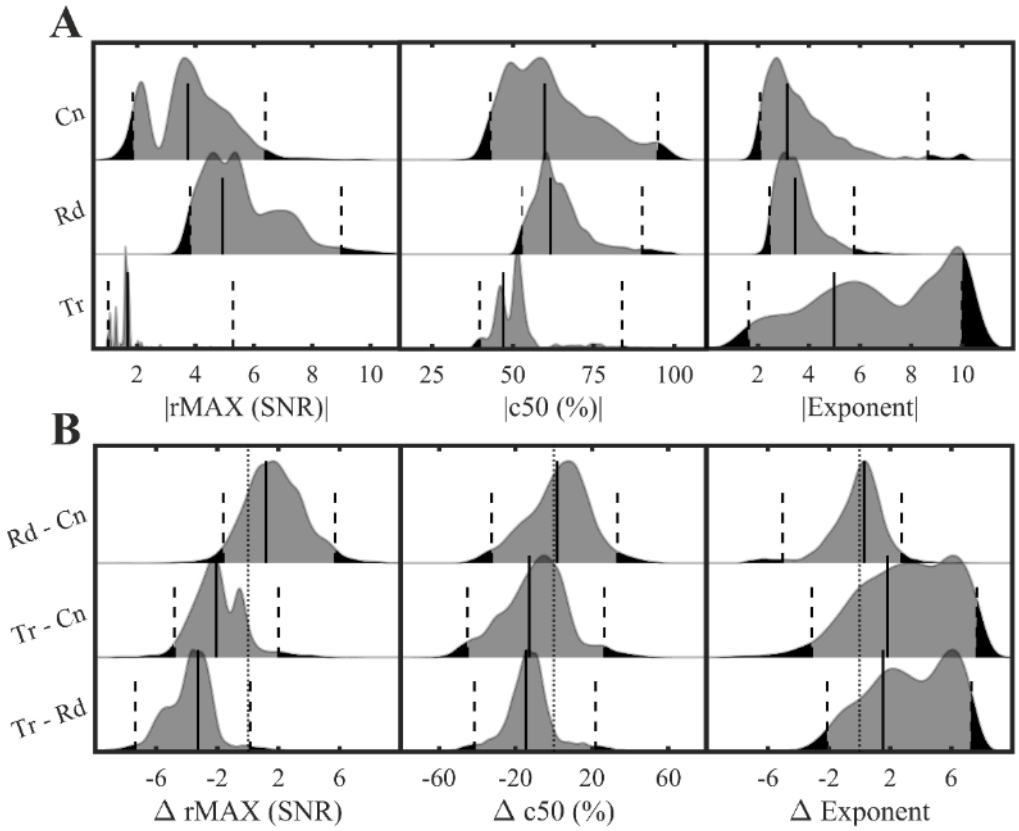

Figure 4.10: Comparison of fit parameters between Glass pattern axes
A: Distribution of parameter estimates recovered from the model fit bootstrapping procedure. Vertical dashed lines and black regions show the 2.5th and 97.5th percentiles. B: Distribution of differences in parameter estimates. When a confidence interval (vertical dashed lines) does not contain the zero reference (dotted) line, a significant difference is implied. Translational, radial, and concentric are abbreviated to Tr, Rd, and Cn (respectively).

The final fit parameter - the rMax - also shows a trend towards lower values in the translational condition (median: 1.8, 95% CI [1.0, 5.29]) when compared to the radial (median: 4.9. 95% CI [3.8, 9.0]) and concentric (median: 3.7, 95% CI [1.8, 6.4]). Despite the magnitude of this reduction, it does not prove to be significant. Indeed, some observers show maximum responses in the translational condition on-par with the other pattern conditions, and the highest overall SNR was actually from one observer in the translational condition. Overall, when comparing the hyperbolic ratio fit parameters across pattern types, we find little evidence for a reliable difference in

parameters across pattern types, consistent with the fits we have previously performed on psychophysically obtained discrimination thresholds. Although, the confidence interval on the difference between translational and radial form only marginally crosses zero, and the ubiquity of reports finding increased responses to radial patterns prompted further investigation. As we have directly measured the responses to 100% coherence, and because the rMax parameter is affected by responses at all levels of coherence, we ran a supplementary analysis contrasting observers' SNR values to 100% coherence between pattern axes. The results of this contrast are shown in Figure 4.11, where we find that SNR values were significantly higher in the radial condition than in the translational condition, in-line with previous research. This indicates that some observers had increased responses to maximally aligned radial patterns that were not fully captured by the hyperbolic ratio transducer model.

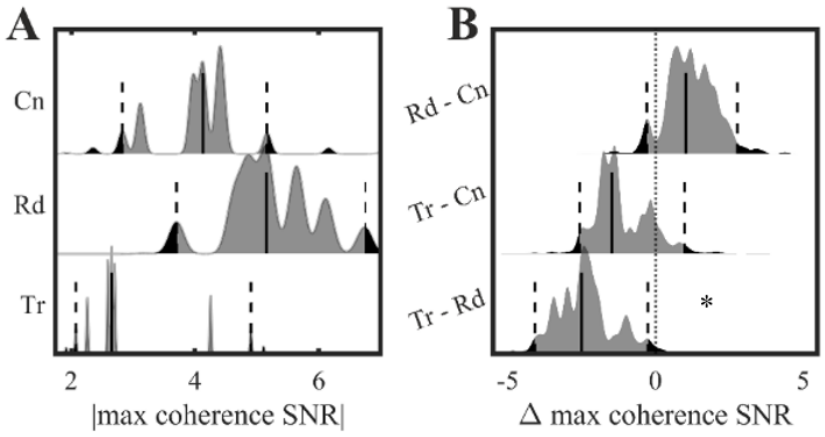

Figure 4.11: Comparison of SNR at maximum coherence across pattern axes
A: Bootstrap distribution of SNR at maximum pattern coherence. B: Cross-condition contrasts of responses to maximum coherence. Contrasts with a star besides them are indicative of a significant difference between comparators.

4.4.6 Response dynamic range is higher than behavioural predictions

Having established the general form of the transducer recovered by our SSVEP methodology for global form RC1, we now directly compare the parameter estimates of the neuronal transducer fits with those predicted by our earlier psychophysical work. In Figure 4.12, the hyperbolic ratio function that best fit behavioural thresholds is overlayed with those of the present experiment (this was accomplished by substituting the SSVEP c50 and exponent values with those obtained psychophysically, but still using the rMAX value of the SSVEP). Initial inspection suggests that psychophysical thresholds predict peak sensitivity over a much lower range of coherence levels, and had a more pronounced response saturation. This is

illustrated by the leftwards shift of the dotted curve relative to the curves derived from EEG responses in all three conditions.

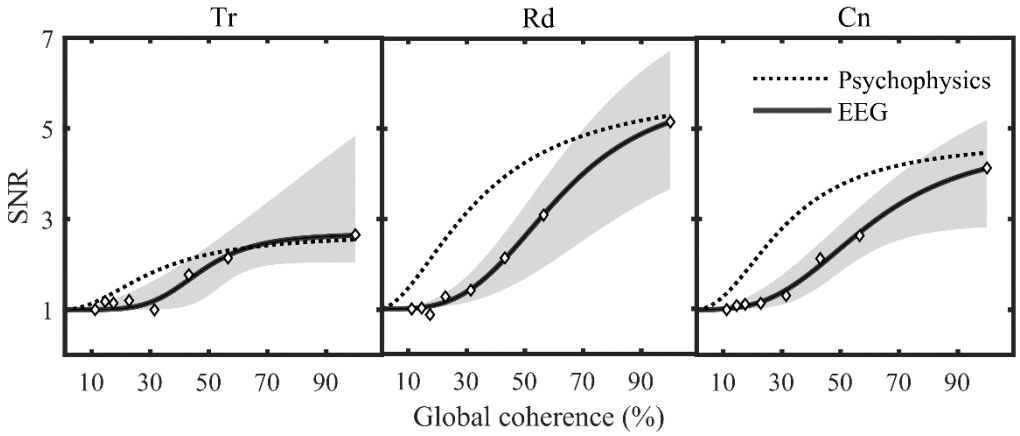

Figure 4.12: Comparison of SSVEP models with psychophysical predictions
Comparison of hyperbolic ratio fits from the psychophysical experiment (dotted line) and those based on electrophysiological recordings (solid). The bootstrapped confidence interval on the psychophysical fits is not shown here, but is instead represented in the distributions in Figure 4.13. Individual datapoints and CIs on the median are omitted for clarity. Tr: translational, Rd: radial, Cn: concentric

In Figure 4.13, the distribution of parameter estimates relating to sensitivity (c50s and exponents) are compared across psychophysical and electrophysiological estimates. There is quite a striking difference between experiments in the distribution of c50s, with the median value from the present experiment being almost double that of our psychophysical study. This reduction was significant for translational (median *Δc50*: 15.7%, 95% CI [7.5, 51.6]) and radial (median *Δc50*: 26.8%, 95% CI [11.2, 55.0]), but was not significant for the concentric condition (median *Δc50*: 28.6%, 95% CI [-3.5, 62.2]), with the latter showing the most *c50* variation in both experiments. For the *exponent* parameter, confidence intervals cross zero for the translational and concentric condition, but do not for the radial condition, where the EEG values show a median increase of 1.6 (95% CI [.45,3.9]). Overall, in comparing the distribution of *c50* and *exponent* fit parameters across experiments, we find significant differences in the radial condition for both, and in the c50 only for the translational condition. For the concentric condition, confidence intervals on the difference between parameters cross zero for both the *c50* and *exponent*, though for the former this is likely due to the heavily skewed distribution from the psychophysical experiment.

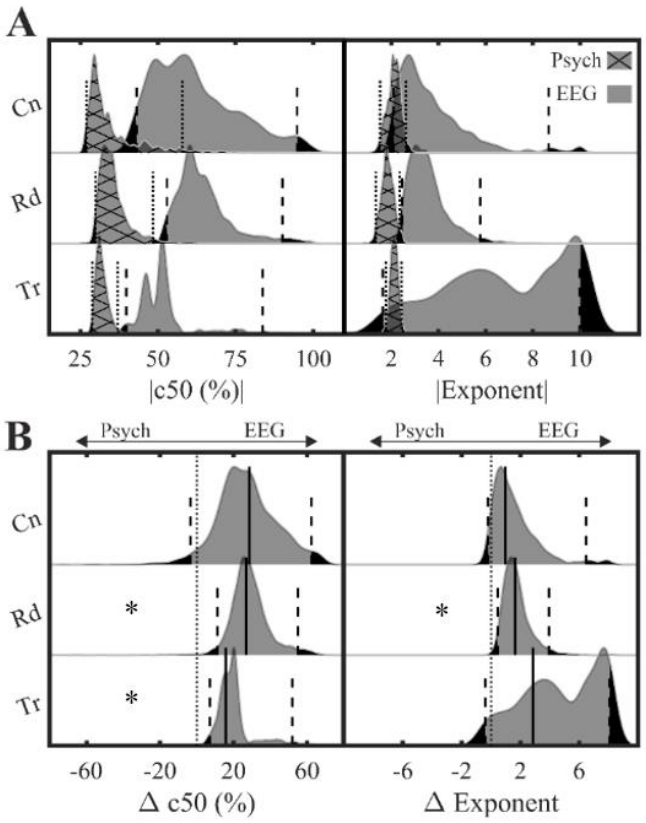

Figure 4.13: Contrast of SSVEP and psychophysical fit parameters
A: The distribution of absolute fit parameter values, with the psychophysical distribution denoted by the hatched pattern. B: the distributions of differences in parameter values between the EEG and psychophysical data. Where the confidence intervals (vertical dashed lines) do not contain zero (vertical dotted line), a significant difference is implied (marked by star symbols). Distributions to the right of the dotted line imply a higher parameter value for the SSVEP model fits. Tr: translational, Rd: radial, Cn: concentric

4.5 Discussion

This experiment aimed to explore the shape of the population neuronal transducer for global form, and to investigate the degree to which the gradient of neuronal responses predicts the pattern of discrimination thresholds we previously obtained psychophysically (Chapter 3). Our findings indicate that both behaviour and electrophysiology are indicative of threshold facilitation and suppression, in that they both can be described by an accelerating and saturating transducer function. However, the range of global form coherence values over which peak sensitivity is predicted (the steepest part of the transducer) is significantly elevated in electrophysiological recordings. Indeed, the psychophysical model fits predicted a transducer that was steepest over a lower range of coherence values, such that a pattern of lower coherence is still considered to be globally correlated. Nevertheless, response nonlinearity was

observed for the majority of observers, although our bootstrapping analysis revealed that some observer combinations may have been better fit by a function that lacks a saturating nonlinearity (those with a c50 > 100%). In bootstraps that did show saturation (the majority of them did), there was an absence of polar form enhancement. The hyperbolic ratio parameters that describe the shape and diagonal translation of responses (the exponent and the c50) did not significantly differ across pattern conditions, which is in line with our psychophysical predictions.

4.5.1 Relating global form responses to behavioural thresholds

While we did find evidence for an accelerating and mildly saturating nonlinearity in the perception of global form (as would be predicted by our psychophysical findings), we found that its accelerating regime occurred beyond the range of peak sensitivity predicted by psychophysical thresholds. There are several possible explanations for this difference. First, while RCA produces a filter that maximises the contribution from electrodes showing consistent responses at the global form switching frequency, this does not guarantee that the generative neuronal populations are the same population that is used to make an actual decision. It is probable that the neuronal population responsible for a behavioural response is just a small fraction of the responsive population, the selection of which may be driven by attentional processes (Verghese et al., 2012). In other words, our behavioural data are likely determined by the most sensitive global form encoding mechanism, rather than the average of the entire responsive population (which is what is measured by EEG). Indeed, previous fMRI experiments investigating responses to maximally aligned Glass patterns have shown that some degree of global form selectivity is broadly distributed along low to mid-level visual processing regions, but that only a small number of subregions (LOC, V4) encode differences between pattern types (Mannion et al., 2010; Ostwald et al., 2008). If global form information is represented throughout the visual system, but a relatively small sub-population are behaviourally decisive, our findings are limited by the low spatial resolution/volume conduction inherent to EEG, and we should perhaps expect that this average sensitivity differs from behavioural performance.

While we consider the limitations of EEG to explain the majority of the difference between our psychophysical and electrophysiological data, there are some

further points for consideration. One possibility is that observers' attention affects the range of global form values over which they are most sensitive. In the domain of luminance contrast, introducing an attentionally loading task to observers while measuring contrast discrimination is known to horizontally shift thresholds, such that attention acts as a form of input gain control (L. Huang & Dobkins, 2005; Reynolds & Heeger, 2009). When measuring discrimination thresholds psychophysically, a minimum level of attention is necessary to engage in the task and to provide accuracy estimates that do not reflect chance-level performance. However, in a pure SSVEP paradigm (as we have used), a more severe deterioration of attention is almost guaranteed as observers have no task other than to visually fixate, which may alter the sensitivity of neural populations that are modulated by attention. Currently, there is scant evidence to support the existence of a mechanism analogous to contrast gain control for global form stimuli, though Palomares et al., (2012) found that diminished attention reduces SSVEP responses to maximally aligned dynamic Glass patterns. This may be equivalent to changing the *rMAX* parameter of the present experiment, though this reflects a vertical translation of responses (output gain control), not the horizontal translation that would unify our psychophysics and electrophysiology. The psychophysical work of Pavan et al. (2019) demonstrated that introducing an attentional load during adaptation to maximally aligned Glass patterns increased the coherence at which a pattern is considered to be composed of randomly oriented dipoles. In other words, with an attentional load, observers were less sensitive to Glass pattern global form, and required patterns of higher coherence before considering them to be geometrically correlated. If this result is interpreted as attentional load elevating Glass pattern detection thresholds, then this could reflect an attentionally modulated input gain mechanism. This is because increasing the *c50* of the transducer diagonally translates both detection and discrimination thresholds, such that the threshold for detection also increases. Thus, there is some evidence that attention can modulate global form responses, but it remains unknown whether a mechanism akin to contrast gain exists for the perception of global form. Overall, that we did not find significant responses at global form coherence levels that are known to be behaviourally discernible limits the interpretation of our results. At the very least, we have found strong evidence for a nonlinear profile of responses to global form coherence in most

observers, but its relationship to behavioural sensitivity rests on an explanation for why the transducer may be horizontally translated.

4.5.2 Little evidence for polar form enhancement

Increased sensitivity to concentric form has been identified by earlier psychophysical reports that focus on detection thresholds of Glass patterns in noise (Kelly et al., 2001; Seu & Ferrera, 2001; H. R. Wilson & Wilkinson, 1998; H. R. Wilson et al., 1997). As with our psychophysical experiment, we found little evidence for a difference in hyperbolic ratio parameter estimates between Glass pattern alignment axes. However, it is evident that the translational condition was exceptional in the variability of its parameter estimates. The *rMax* and *c50* parameters for this condition had quite narrow distributions relative to other pattern types, while the exponent varied over quite an extreme range of values. To a first approximation, extreme exponents in the model fits for some observers could be taken as suggestive of increased sensitivity to translational form (that is, a sharper acceleration and saturation). A finding of increased sensitivity to translational form would be exceptional with respect to the existing psychophysical literature, and is not predicted in our own psychophysical experimentation, prompting alternative explanations. A consistent finding from M/EEG and fMRI experiments that investigated responses to maximally aligned Glass patterns is that translational patterns evoke a significantly weaker responses than concentric and radial patterns (Ohla et al., 2005; Pei et al., 2005; Rampone & Makin, 2020; Swettenham et al., 2010). In the present experiment we also find this distinction between radial and translational patterns when directly comparing maximum-coherence SNR values. If translational patterns, for no reason related to behavioural thresholds, yield reduced SNR when measured with EEG, then a higher global form coherence would be required to reveal significant responses. In our own SSVEP measurements, this would mean that a larger proportion of the transducer is masked by physiological noise, perhaps only becoming measurable when it is well into its accelerating regime. When setting baseline noise (*rMin = SNR of 1*) as the lower asymptote of the response function (as we have), a more severe exponent parameter is required to describe the sudden elevation of responses above the noise floor. For this reason, we would advise caution when interpreting the variability of translational exponent, as we suspect that it is not related to behavioural sensitivity,

but that it may reflect an inherent limitation of measuring the responses from a population of neurons using EEG.

4.6 Conclusion

In this experiment we demonstrated that activity over parieto-occipital electrodes demonstrate a response that scales with the global form coherence of translational, radial, and concentric Glass patterns. In the majority of observers, this gradient was well described by an accelerating nonlinearity followed by a mild saturation. When compared with predictions obtained psychophysically, the SSVEP response function predicted peak sensitivity over a significantly higher range of global form coherence values. We suspect that this horizontal shift is most likely to be due to the lack of spatial selectivity of electroencephalography and the method we used to extract signal-to-noise estimates, but also provide an alternative explanation from attentional modulation. Consistent with our psychophysical work, we find little evidence for a bias in supra-threshold sensitivity to radial or concentric patterns, but we do find that radial patterns evoke the largest responses overall, consistent with the finding of other experiments.

Chapter 5. General discussion

5.1 Summary of experimental work

The goal of this thesis was to explore nonlinearity and signal transduction in low and mid-level human visual perception, focusing on surround suppression of contrast, a divisive nonlinearity, and the transduction of global form. More specifically, we aimed to understand the extent to which surround suppression of contrast is dependent on advancing age, to further our understanding of *suprathreshold* surround suppression in central vision, and to explore observers' sensitivity to differences in global form and the extent to which psychophysically obtained thresholds can be predicted from population neuronal responses. In the following section, a brief overview of the findings from each chapter of this thesis is presented.

In Chapter 2, we investigated the possibility that recent reports of an age-related increase in surround suppression at supra-threshold contrasts (Karas & McKendrick, 2009, 2011, 2015) could be due to contamination from overlay masking, a separate, possibly pre-cortical form of contrast normalisation. Using stimuli designed to preclude contributions from overlay masking, we measured suppression strength in older (> 60) and younger (< 30) observers. Despite using centre and surround contrasts reported to produce the maximal age-related increase in the suppression of perceived contrast by previous reports, there was little evidence of an age-related increase in suppression magnitude. This implies that short-range, pre-cortical overlay masking could be the mechanism affected by age in existing experiments. The surround suppression we did find was weak, but significant. It was also mostly independent of the relative orientation of the centre and surround gratings, unlike the surround suppression routinely found in the periphery.

In Chapter 3, we explored behavioural sensitivity to differences in the global form percept elicited by Glass patterns. In particular, we were interested to see if the "dipper effect" (a paradoxical facilitation of discrimination thresholds at low pedestal levels) can be observed in a mid-level parameter space, as it has been for the higher-level perception of facial expression (Gray et al., 2020). A confirmation of a Glass pattern "dipper" would strengthen the notion that the consequences of nonlinear neuronal signal transduction can be found throughout the human visual system, and

perhaps throughout the whole brain. This also provided the opportunity to investigate previous reports of increased sensitivity to concentric and radial global form over translational global form, though while measuring across the entire range of discrimination thresholds, rather than the detection thresholds that have previously been investigated (Seu & Ferrera, 2001; H. R. Wilson & Wilkinson, 1998; H. R. Wilson et al., 1997). The results of our experiments suggest that there is evidence for both threshold facilitation and suppression, and that they could be the consequence of nonlinear transduction in the perception of global form. Unlike previous studies, we found little evidence for enhanced sensitivity to polar form over translational form.

In Chapter 4, we investigated the extent to which the thresholds obtained in our online psychophysical experiment could be explained by a nonlinear neuronal 'global form' transducer. We measured steady-state visually evoked potentials for dynamic Glass patterns across the same range of global form coherences that we examined psychophysically, and fit responses with the same transducer model used to fit discrimination thresholds. Our results demonstrated that responses to our parameterisation of global form accelerate and saturate (reflective of a dipper effect), but that the accelerating and saturating regimes occurred over a higher level of global form coherences than would be predicted by our psychophysical thresholds. For the most part, we found little difference in the steepness and location of the response function across pattern types, but radial patterns did show considerably higher responses at maximum pattern coherence, as has been identified in previous reports (Ostwald et al., 2008; Pei et al., 2005; Swettenham et al., 2010).

5.2 Future directions

The experiments presented in this thesis suggest several future lines of enquiry. In Chapter 2, the fact that we did not find any evidence for an age-related effect of surround suppression when contributions from overlay masking were attenuated suggests that surround suppression is not affected by age, contrary to previous reports (Karas & McKendrick, 2009, 2011, 2015; Pitchaimuthu, Nguyen, et al., 2017). It is possible that overlay masking is altered by age, and future work should specifically compare overlay masking using spatially superimposed stimuli and surround suppression between older and younger individuals.

Finding an age-related change that specifically affects one form of suppression would not only inform our understanding of the changes in vision that accompany age, but might also provide clues about their dependency on the concentrations of excitatory and inhibitory neurotransmitters. There is evidence from the study of schizophrenia that a reduction of the concentration of inhibitory neurotransmitter GABA (gamma-aminobutyric acid) in the visual cortex is a corollary of weakened orientation tuned surround suppression (Yoon et al., 2010). However, the effect of healthy aging on GABAergic signalling has received only limited investigation in humans. A study of post-mortem tissue samples found evidence for an age-related reduction in GABA compounds in primary visual cortex (Pinto et al., 2010). Conversely, recent evidence from magnetic resonance spectroscopy (a method that allows the non-invasive measurement of neurotransmitter concentration in living humans) has found the concentration of GABA in the visual cortex to increase with age (Pitchaimuthu, Wu, et al., 2017). If the strength of surround suppression is broadly contingent on the action of GABA, these reports predict contradictory effects. That we have not found an effect of age on surround suppression may indicate that GABA does not directly govern its magnitude.

The suppression we did find in Chapter 2 was only broadly orientation tuned, though other groups (Cannon & Fullenkamp, 1991; Xing & Heeger, 2000) have reported stronger orientation tuning when using centrally presented suprathreshold stimuli. Xing and Heeger used long-duration stimuli that flickered at 8Hz, and the stimuli of Cannon and Fullenkamp also had a relatively long duration. It is possible that some of the differences between our results are due to the presence of long-term contrast normalisation in Xing and Heeger's paradigm (releasing orientation-tuned mechanisms from saturation), or perhaps some artefact of exploratory eye movements in the case of the extended probes used by Cannon and Fullenkamp. Moreover, both studies used closely abutting surrounds when measuring orientation tuning which might have generated contamination from overlay masking (Petrov et al., 2005) or edge segmentation effects (Appelbaum et al., 2008). Future work exploring the orientation tuning of foveal suprathreshold surround suppression could explore these issues by parametrically altering the duration, flicker frequency and separation of the central probe region.

In Chapter 3, we found evidence for a "dipper effect" in global form detection and discrimination thresholds, and in Chapter 4 we found evidence for an accelerating and saturating nonlinearity in the evoked responses to global form. The semi-saturation points (*c50s*) computed from the electrophysiology were higher for two out of the three pattern axes we investigated, such that the psychophysical thresholds predicted high sensitivity over a lower range of global form coherences. Future work might be directed towards further exploring this inconsistency, as previous work in the contrast domain has been successful in relating discrimination to functional activation (Boynton et al., 1999). An explanation we provided was some contribution from a top-down attentional process. When making a behavioural decision, attention likely acts to select the most sensitive neuronal sub-population (Verghese et al., 2012), while EEG can only measure the whole population response. Therefore, the transducer that we fit to evoked scalp responses may be contaminated by response populations that are not behaviourally decisive. Alternatively, attention has been previously shown translate threshold vs contrast curves (L. Huang & Dobkins, 2005; Reynolds & Heeger, 2009). An attention driven gain analogous to contrast gain has not been investigated for global form, though it may explain our results. Indeed, the attentional load observers are subject to when performing the behavioural discrimination task would be higher than the passive viewing required for a steady-state VEP experiment. Future work might explore possibility of attentional modulation of global form discrimination. A final possibility is that the population responses to subtle global form may have simply been below the EEG noise floor in our measurements. Although, if this were the case, we should still expect to see a similar saturation of responses across both experiments, which was qualitatively not the case. Nevertheless, the higher signal-to-noise ratio of a method like magnetoencephalography (Goldenholz et al., 2009) may be useful in resolving responses to low levels of global form coherence.

We are the first to establish the shape of the Glass pattern global form discrimination function using local orientations sampled from a Gaussian distribution. As such, our pedestal levels were chosen relatively sparsely, and we did not sample discrimination thresholds at or below absolute detection. This is particularly important because the facilitatory regime of the dipper effect for contrast is usually most prominent proximal to the threshold for absolute detection (Nachmias & Sansbury,

1974). Future experiments may be directed towards a more granular assessment of global form discrimination near to the absolute detection thresholds we have reported. Moreover, future work would likely have the benefit of vastly improved experimental control, as our work was inherently uncontrolled due to the use of online experimentation.

Finally, it is curious that we observed little evidence of enhanced sensitivity to concentric global form, as is often reported at detection threshold (Seu & Ferrera, 2001; H. R. Wilson & Wilkinson, 1998). Dakin (2002) argued that previous reports of enhanced sensitivity to concentric form are solely due to the presence of edge cues introduced by using a circular aperture. Our experiments used a circular aperture, yet in Chapter 3 even the threshold for global form detection was similar across all pattern axes we examined. Why do previous reports find concentric enhancement, but we have not? A fundamental difference between previous parameterisations of global form and our own is the presence of dipoles that are perfectly aligned co-circularly or radially to the centre of the pattern field. Previous parameterisations increased the salience of global form by increasing the proportion of dipoles that followed the geometric rule, such that a "coherence" of 50% means half the dipoles are perfectly aligned, while the other half are randomly oriented. In our parameterisation, however, perfectly aligned dipoles are relatively rare because their orientation is drawn from a Gaussian distribution. One explanation for our results could be the presence of a population of neurons with tightly tuned preferences for circular or radial global form. If such populations exists, they may also be the generators of concentric bias identified in previous reports, and would be scarcely innervated by our stimuli at low and intermediary levels of global form coherence. Future work could explore the possibility of such a narrow-band global form mechanism via a psychophysical adaptation paradigm. For example, observers could be adapted to a perfectly aligned concentric Glass pattern (with regular "top-ups"), and their detection thresholds for concentric global form obtained both using our parameterisation of global form, and (separately) for the proportion-aligned parameterisation. Finding a reduction in thresholds for the latter parameterisation, but relatively stable thresholds in the former would be indicative of a unique contribution from a highly tuned population

5.3 Final conclusions

The present thesis explored questions relating to nonlinearities in visual processing, specifically the spatial modulation of surround suppression and the behavioural sensitivity and transduction of mid-level global form stimuli. Our main findings are summarised below:

1. Suprathreshold surround suppression of contrast in central vision is relatively stable across the lifespan and is only broadly orientation tuned.
2. There is threshold facilitation and suppression in the behavioural discrimination of the global form percept produced by Glass patterns, consistent with a “dipper effect”.
3. The gradient of neuronal responses to increasing global form coherence shows evidence of both acceleration and saturation. While this is also consistent with a “dipper effect”, further work is required to reconcile the difference in the sensitivity predicted by our psychophysical and electroencephalographical experiments
4. The psychophysical sensitivity to global form, as we have defined it, is similar across polar and translational Glass pattern alignment axes.
5. Steady-state evoked responses are also similar across polar and cartesian form, although, at maximum global form coherence, the response to radial patterns is significantly higher than the response translational patterns, consistent with previous findings.

These findings extend our understanding of the characteristics of suprathreshold spatial modulation of contrast and how it changes across the lifespan, and adds to the evidence that nonlinear neuronal transduction is present throughout the human brain.

0

References

Albrecht, D. G., & Geisler, W. S. (1991). Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Visual Neuroscience*, *7*(6), 531–546. https://doi.org/10.1017/s0952523800010336

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/10.1152/jn.1982.48.1.217

Anderson, S. J., & Swettenham, J. B. (2006). Neuroimaging in human amblyopia. *Strabismus*, *14*(1), 21–35. https://doi.org/10.1080/09273970500538082

Angelucci, A., Bijanzadeh, M., Nurminen, L., Federer, F., Merlin, S., & Bressloff, P. C. (2017). Circuits and Mechanisms for Surround Modulation in Visual Cortex. *Annual Review of Neuroscience*, *40*, 425–451. https://doi.org/10.1146/annurev-neuro-072116-031418

Anstis, S. M. (1970). Phi movement as a subtraction process. *Vision Research*, *10*(12), 1411–1430. https://doi.org/10.1016/0042-6989(70)90092-1

Appelbaum, L. G., Wade, A. R., Pettet, M. W., Vildavski, V. Y., & Norcia, A. M. (2008). Figure-ground interaction in the human visual cortex. *Journal of Vision*, *8*(9), 8.1-19. https://doi.org/10.1167/8.9.8

Arundale, K. (1978). An investigation into the variation of human contrast sensitivity with age and ocular pathology. *The British Journal of Ophthalmology*, *62*(4), 213–215. https://doi.org/10.1136/bjo.62.4.213

Baker, D. H., Meese, T. S., & Georgeson, M. A. (2007). Binocular interaction: contrast matching and contrast discrimination are predicted by the same

model. *Spatial Vision*, *20*(5), 397–413. https://doi.org/10.1163/156856807781503622

Baker, D. H., Meese, T. S., & Summers, R. J. (2007). Psychophysical evidence for two routes to suppression before binocular summation of signals in human vision. *Neuroscience*, *146*(1), 435–448. https://doi.org/10.1016/j.neuroscience.2007.01.030

Baker, D. H., Vilidaite, G., & Wade, A. R. (2021). Steady-state measures of visual suppression. *PLoS Computational Biology*, *17*(10), e1009507. https://doi.org/10.1371/journal.pcbi.1009507

Barlow, H. B. (1957). Increment thresholds at low intensities considered as signal/noise discriminations. *The Journal of Physiology*, *136*(3), 469–488. https://doi.org/10.1113/jphysiol.1957.sp005774

Beard, B. L., Yager, D., & Neufeld, S. (1994). Contrast detection and discrimination in young and older adults. *Optometry and Vision Science: Official Publication of the American Academy of Optometry*, *71*(12), 783–791. https://www.ncbi.nlm.nih.gov/pubmed/7898886

Benjamini, Y., & Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, *29*(4), 1165–1188. https://doi.org/10.1214/aos/1013699998

Betts, L. R., Taylor, C. P., Sekuler, A. B., & Bennett, P. J. (2005). Aging reduces center-surround antagonism in visual motion processing. *Neuron*, *45*(3), 361–366. https://doi.org/10.1016/j.neuron.2004.12.041

Bird, C. M., Henning, G. B., & Wichmann, F. A. (2002). Contrast discrimination with sinusoidal gratings of different spatial frequency. *JOSA A*, *19*(7), 1267–1273. https://doi.org/10.1364/JOSAA.19.001267

Bonds, A. B. (1989). Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Visual Neuroscience*, *2*(1), 41–55. https://doi.org/10.1017/s0952523800004314

Bonin, V., Mante, V., & Carandini, M. (2005). The suppressive field of neurons in lateral geniculate nucleus. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(47), 10844–10856. https://doi.org/10.1523/JNEUROSCI.3562-05.2005

Boynton, G. M., Demb, J. B., Glover, G. H., & Heeger, D. J. (1999). Neuronal basis of contrast discrimination. *Vision Research*, *39*(2), 257–269. https://doi.org/10.1016/S0042-6989(98)00113-8

Boynton, G. M., & Foley, J. M. (1999). Temporal sensitivity of human luminance pattern mechanisms determined by masking with temporally modulated stimuli. *Vision Research*, *39*(9), 1641–1656. https://doi.org/10.1016/S0042-6989(98)00199-0

Brewer, A. A., & Barton, B. (2012). Effects of healthy aging on human primary visual cortex. *Health*, *4*(9A), 695–702. https://doi.org/10.4236/health.2012.429109

Burton, E., Wattam-Bell, J., Rubin, G. S., Atkinson, J., Braddick, O., & Nardini, M. (2016). Cortical processing of global form, motion and biological motion under low light levels. *Vision Research*, *121*, 39–49. https://doi.org/10.1016/j.visres.2016.01.008

Burton, G. J. (1981). Contrast Discrimination by the human visual system. *Biological Cybernetics*, *40*(1), 27–38. https://doi.org/10.1007/BF00326678

Cai, Y., Zhou, T., & Chen, L. (2008). Effects of binocular suppression on surround suppression. *Journal of Vision*, *8*(9), 9.1-10. https://doi.org/10.1167/8.9.9

Cannon, M. W., & Fullenkamp, S. C. (1991). Spatial interactions in apparent contrast: inhibitory effects among grating patterns of different spatial frequencies, spatial positions and orientations. *Vision Research*, *31*(11), 1985–1998. https://doi.org/10.1016/0042-6989(91)90193-9

Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., Gallant, J. L., & Rust, N. C. (2005). Do we know what the early visual system does? *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(46), 10577–10597. https://doi.org/10.1523/JNEUROSCI.3726-05.2005

Carandini, M., & Heeger, D. J. (2011). Normalization as a canonical neural computation. *Nature Reviews. Neuroscience*, *13*(1), 51–62. https://doi.org/10.1038/nrn3136

Carandini, M., Heeger, D. J., & Anthony Movshon, J. (1997). Linearity and Normalization in Simple Cells of the Macaque Primary Visual Cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *17*(21), 8621–8644. https://doi.org/10.1523/JNEUROSCI.17-21-08621.1997

Cavanaugh, J. R., Bair, W., & Movshon, J. A. (2002). Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. *Journal of Neurophysiology*, *88*(5), 2530–2546. https://doi.org/10.1152/jn.00692.2001

Chaumon, M., Bishop, D. V. M., & Busch, N. A. (2015). A practical guide to the selection of independent components of the electroencephalogram for artifact correction. *Journal of Neuroscience Methods*, *250*, 47–63. https://doi.org/10.1016/j.jneumeth.2015.02.025

Chen, Y., Ko, H., Zemelman, B. V., Seidemann, E., & Nauhaus, I. (2020). Uniform spatial pooling explains topographic organization and deviation from receptive-field scale invariance in primate V1. *Nature Communications*, *11*(1), 6390. https://doi.org/10.1038/s41467-020-19954-9

Chubb, C., Sperling, G., & Solomon, J. A. (1989). Texture interactions determine perceived contrast. *Proceedings of the National Academy of Sciences of the United States of America*, *86*(23), 9631–9635. https://doi.org/10.1073/pnas.86.23.9631

Claridge-Chang, A., & Assam, P. N. (2016). Estimation statistics should replace significance testing. *Nature Methods*, *13*(2), 108–109. https://doi.org/10.1038/nmeth.3729

Cottereau, B. R., McKee, S. P., Ales, J. M., & Norcia, A. M. (2011). Disparity-tuned population responses from human visual cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *31*(3), 954–965. https://doi.org/10.1523/JNEUROSCI.3795-10.2011

Dakin, S. C. (1997). The detection of structure in glass patterns: psychophysics and computational models. *Vision Research*, *37*(16), 2227–2246. https://doi.org/10.1016/s0042-6989(97)00038-2

Dakin, S. C., & Bex, P. J. (2002). Summation of concentric orientation structure: seeing the Glass or the window? *Vision Research*, *42*(16), 2013–2020. https://doi.org/10.1016/s0042-6989(02)00057-3

DeAngelis, G. C., Robson, J. G., Ohzawa, I., & Freeman, R. D. (1992). Organization of suppression in receptive fields of neurons in cat visual cortex. *Journal of Neurophysiology*, *68*(1), 144–163. https://doi.org/10.1152/jn.1992.68.1.144

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

DiCiccio, T. J., & Efron, B. (1996). Bootstrap confidence intervals. *Schweizerische Monatsschrift Fur Zahnheilkunde = Revue Mensuelle Suisse d'odonto-Stomatologie / SSO*, *11*(3), 189–228. https://doi.org/10.1214/ss/1032280214

Dmochowski, J. P., Greaves, A. S., & Norcia, A. M. (2015). Maximally reliable spatial filtering of steady state visual evoked potentials. *NeuroImage*, *109*, 63–72. https://doi.org/10.1016/j.neuroimage.2014.12.078

Durand, S., Freeman, T. C. B., & Carandini, M. (2007). Temporal properties of surround suppression in cat primary visual cortex. *Visual Neuroscience*, *24*(5), 679–690. https://doi.org/10.1017/S0952523807070563

Elliott, D. B. (1987). Contrast sensitivity decline with ageing: a neural or optical phenomenon? *Ophthalmic & Physiological Optics: The Journal of the British College of Ophthalmic Opticians* , *7*(4), 415–419. https://www.ncbi.nlm.nih.gov/pubmed/3454919

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

Fechner, G. T. (1860). *Elemente der Psychophysik*. Breitkopf und Härtel. https://play.google.com/store/books/details?id=l5ui0NKWXyAC

Ferree, T. C. (2006). Spherical splines and average referencing in scalp electroencephalography. *Brain Topography*, *19*(1–2), 43–52. https://doi.org/10.1007/s10548-006-0011-0

Foley, J. M. (1994). Human luminance pattern-vision mechanisms: masking experiments require a new model. *Journal of the Optical Society of America. A, Optics, Image Science, and Vision*, *11*(6), 1710–1719. https://doi.org/10.1364/josaa.11.001710

Foley, J. M., & Legge, G. E. (1981). Contrast detection and near-threshold discrimination in human vision. *Vision Research*, *21*(7), 1041–1053. https://doi.org/10.1016/0042-6989(81)90009-2

Freeman, T. C. B., Durand, S., Kiper, D. C., & Carandini, M. (2002). Suppression without inhibition in visual cortex. *Neuron*, *35*(4), 759–771. https://doi.org/10.1016/s0896-6273(02)00819-x

Fu, Y., Wang, X. S., Wang, Y. C., Zhang, J., Liang, Z., Zhou, Y. F., & Ma, Y. Y. (2010). The effects of aging on the strength of surround suppression of receptive field of V1 cells in monkeys. *Neuroscience*, *169*(2), 874–881. https://doi.org/10.1016/j.neuroscience.2010.05.015

Gallant, J. L., Connor, C. E., Rakshit, S., Lewis, J. W., & Van Essen, D. C. (1996). Neural responses to polar, hyperbolic, and Cartesian gratings in area V4 of the macaque monkey. *Journal of Neurophysiology*, *76*(4), 2718–2739. https://doi.org/10.1152/jn.1996.76.4.2718

Glass, L. (1969). Moiré effect from random dots. *Nature*, *223*(5206), 578–580. https://doi.org/10.1038/223578a0

Glass, L., & Pérez, R. (1973). Perception of random dot interference patterns. *Nature*, *246*(5432), 360–362. https://doi.org/10.1038/246360a0

Glass, L., & Switkes, E. (1976). Pattern recognition in humans: correlations which cannot be perceived. *Perception*, *5*(1), 67–72. https://doi.org/10.1068/p050067

Goldenholz, D. M., Ahlfors, S. P., Hämäläinen, M. S., Sharon, D., Ishitobi, M., Vaina, L. M., & Stufflebeam, S. M. (2009). Mapping the signal-to-noise-ratios of cortical sources in magnetoencephalography and electroencephalography. *Human Brain Mapping*, *30*(4), 1077–1086. https://doi.org/10.1002/hbm.20571

Gori, M., Mazzilli, G., Sandini, G., & Burr, D. (2011). Cross-Sensory Facilitation Reveals Neural Interactions between Visual and Tactile Motion in Humans. *Frontiers in Psychology*, *2*, 55. https://doi.org/10.3389/fpsyg.2011.00055

Gray, K. L. H., Flack, T. R., Yu, M., Lygo, F. A., & Baker, D. H. (2020). Nonlinear transduction of emotional facial expression. *Vision Research*, *170*, 1–11. https://doi.org/10.1016/j.visres.2020.03.004

Gur, M., Kagan, I., & Snodderly, D. M. (2004). Orientation and Direction Selectivity of Neurons in V1 of Alert Monkeys: Functional Relationships and Laminar Distributions. *Cerebral Cortex* , *15*(8), 1207–1221. https://doi.org/10.1093/cercor/bhi003

Haufe, S., Meinecke, F., Görgen, K., Dähne, S., Haynes, J.-D., Blankertz, B., & Bießmann, F. (2014). On the interpretation of weight vectors of linear models in multivariate neuroimaging. *NeuroImage*, *87*, 96–110. https://doi.org/10.1016/j.neuroimage.2013.10.067

Hawken, M. J., Blakemore, C., & Morley, J. W. (1997). Development of contrast sensitivity and temporal-frequency selectivity in primate lateral geniculate nucleus. *Experimental Brain Research. Experimentelle Hirnforschung. Experimentation Cerebrale*, *114*(1), 86–98. https://doi.org/10.1007/pl00005626

Heeger, D. J. (1992a). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, *9*(2), 181–197. https://doi.org/10.1017/s0952523800009640

Heeger, D. J. (1992b). Half-squaring in responses of cat striate cells. *Visual Neuroscience*, *9*(5), 427–443. https://doi.org/10.1017/s095252380001124x

Huang, L., & Dobkins, K. R. (2005). Attentional effects on contrast discrimination in humans: evidence for both contrast gain and response gain. *Vision Research*, *45*(9), 1201–1212. https://doi.org/10.1016/j.visres.2004.10.024

Huang, P.-C., & Chen, C.-C. (2014). A comparison of pedestal effects in first- and second-order patterns. *Journal of Vision*, *14*(1). https://doi.org/10.1167/14.1.9

Huang, P.-C., & Chen, C.-C. (2016). Contrast Gain Control in Plaid Pattern Detection. *PloS One*, *11*(10), e0164171. https://doi.org/10.1371/journal.pone.0164171

Hubel, D. H., & Wiesel, T. N. (1960). Receptive fields of optic nerve fibres in the spider monkey. *The Journal of Physiology*, *154*, 572–580. https://doi.org/10.1113/jphysiol.1960.sp006596

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, *195*(1), 215–243. https://doi.org/10.1113/jphysiol.1968.sp008455

Hutchinson, C. V., Ledgeway, T., & Allen, H. A. (2014). The ups and downs of global motion perception: a paradoxical advantage for smaller stimuli in the aging visual system. *Frontiers in Aging Neuroscience*, *6*, 199. https://doi.org/10.3389/fnagi.2014.00199

Karas, R., & McKendrick, A. M. (2009). Aging alters surround modulation of perceived contrast. *Journal of Vision*, *9*(5), 11.1-9. https://doi.org/10.1167/9.5.11

Karas, R., & McKendrick, A. M. (2011). Increased surround modulation of perceived contrast in the elderly. *Optometry and Vision Science: Official Publication of the American Academy of Optometry*, *88*(11), 1298–1308. https://doi.org/10.1097/OPX.0b013e31822f4d51

Karas, R., & McKendrick, A. M. (2012). Age related changes to perceptual surround suppression of moving stimuli. *Seeing and Perceiving*, *25*(5), 409–424. https://doi.org/10.1163/187847611X595873

Karas, R., & McKendrick, A. M. (2015). Contrast and stimulus duration dependence of perceptual surround suppression in older adults. *Vision Research*, *110*(Pt A), 7–14. https://doi.org/10.1016/j.visres.2015.02.016

Kelly, D. M., Bischof, W. F., Wong-Wylie, D. R., & Spetch, M. L. (2001). Detection of glass patterns by pigeons and humans: implications for differences in higher-level processing. *Psychological Science*, *12*(4), 338–342. https://doi.org/10.1111/1467-9280.00362

Kim, Y. J., Haun, A. M., & Essock, E. A. (2010). The horizontal effect in suppression: Anisotropic overlay and surround suppression at high and low speeds. *Vision Research*, *50*(9), 838–849. https://doi.org/10.1016/j.visres.2010.01.020

Klem, G. H., Lüders, H. O., Jasper, H. H., & Elger, C. (1999). The ten-twenty electrode system of the International Federation. The International Federation of Clinical Neurophysiology. *Electroencephalography and Clinical*

Neurophysiology. Supplement, *52*, 3–6. https://www.ncbi.nlm.nih.gov/pubmed/10590970

Kontsevich, L. L., Chen, C.-C., & Tyler, C. W. (2002). Separating the effects of response nonlinearity and internal noise psychophysically. *Vision Research*, *42*(14), 1771–1784. https://doi.org/10.1016/s0042-6989(02)00091-3

Kovaćs, I., & Julesz, B. (1992). Depth, motion, and static-flow perception at metaisoluminant color contrast. *Proceedings of the National Academy of Sciences of the United States of America*, *89*(21), 10390–10394. https://doi.org/10.1073/pnas.89.21.10390

Krekelberg, B., Vatakis, A., & Kourtzi, Z. (2005). Implied motion from form in the human visual cortex. *Journal of Neurophysiology*, *94*(6), 4373–4386. https://doi.org/10.1152/jn.00690.2005

Kuffler, S. W. (1953). Discharge patterns and functional organization of mammalian retina. *Journal of Neurophysiology*, *16*(1), 37–68. https://doi.org/10.1152/jn.1953.16.1.37

Ledgeway, T., & Smith, A. T. (1994). Evidence for separate motion-detecting mechanisms for first- and second-order motion in human vision. *Vision Research*, *34*(20), 2727–2740. https://doi.org/10.1016/0042-6989(94)90229-1

Legge, G. E., & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America*, *70*(12), 1458–1471. https://www.ncbi.nlm.nih.gov/pubmed/7463185

Levitt, J. B., & Lund, J. S. (1997). Contrast dependence of contextual effects in primate visual cortex. *Nature*, *387*(6628), 73–76. https://doi.org/10.1038/387073a0

Li, Q., Joo, S. J., Yeatman, J. D., & Reinecke, K. (2020). Controlling for Participants' Viewing Distance in Large-Scale, Psychophysical Online Experiments Using a Virtual Chinrest. *Scientific Reports*, *10*(1), 904. https://doi.org/10.1038/s41598-019-57204-1

Lin, Y.-S., Cho, P.-C., & Chen, C.-C. (2017). Contrast gain control determines global form percept in tripole Glass patterns. *Journal of Vision*, *17*(5), 2. https://doi.org/10.1167/17.5.2

Maloney, R. K., Mitchison, G. J., & Barlow, H. B. (1987). Limit to the detection of Glass patterns in the presence of noise. *Journal of the Optical Society of America. A, Optics and Image Science*, *4*(12), 2336–2341. https://doi.org/10.1364/josaa.4.002336

Mannion, D. J., Kersten, D. J., & Olman, C. A. (2013). Consequences of polar form coherence for fMRI responses in human visual cortex. *NeuroImage*, *78*, 152–158. https://doi.org/10.1016/j.neuroimage.2013.04.036

Mannion, D. J., McDonald, J. S., & Clifford, C. W. G. (2009). Discrimination of the local orientation structure of spiral Glass patterns early in human visual cortex. *NeuroImage*, *46*(2), 511–515. https://doi.org/10.1016/j.neuroimage.2009.01.052

Mannion, D. J., McDonald, J. S., & Clifford, C. W. G. (2010). Orientation anisotropies in human visual cortex. *Journal of Neurophysiology*, *103*(6), 3465–3471. https://doi.org/10.1152/jn.00190.2010

McGill, W. J., & Goldberg, J. P. (1968). A study of the near-miss involving Weber's law and pure-tone intensity discrimination. *Perception & Psychophysics*, *4*(2), 105–109. https://doi.org/10.3758/BF03209518

Meese, T. S., & Baker, D. H. (2009). Cross-orientation masking is speed invariant between ocular pathways but speed dependent within them. *Journal of Vision*, *9*(5), 2. https://doi.org/10.1167/9.5.2

Meese, T. S., & Hess, R. F. (2004). Low spatial frequencies are suppressively masked across spatial scale, orientation, field position, and eye of origin. *Journal of Vision*, *4*(10), 843–859. https://doi.org/10.1167/4.10.2

Mei, M., Leat, S. J., & Hovis, J. (2007). Supra-threshold contrast matching and the effects of contrast threshold and age. *Clinical & Experimental Optometry: Journal of the Australian Optometrical Association*, *90*(4), 272–281. https://doi.org/10.1111/j.1444-0938.2007.00162.x

Meier, L., & Carandini, M. (2002). Masking by fast gratings. *Journal of Vision*, *2*(4), 2–2. https://doi.org/10.1167/2.4.2

Moreland, K. (2009). Diverging Color Maps for Scientific Visualization. *Advances in Visual Computing*, 92–103. https://doi.org/10.1007/978-3-642-10520-3_9

Morgan, M., Chubb, C., & Solomon, J. A. (2008). A 'dipper' function for texture discrimination based on orientation variance. *Journal of Vision*, *8*(11), 9–9. https://doi.org/10.1167/8.11.9

Movshon, J. A., Smith, M. A., & Kohn, A. (2010). Responses to glass patterns in macaque V1 and V2. *Journal of Vision*, *3*(9), 151–151. https://doi.org/10.1167/3.9.151

Nachmias, J., & Sansbury, R. V. (1974). Letter: Grating contrast: discrimination may be better than detection. *Vision Research*, *14*(10), 1039–1042. https://doi.org/10.1016/0042-6989(74)90175-8

Naka, K. I., & Rushton, W. A. (1966). S-potentials from luminosity units in the retina of fish (Cyprinidae). *The Journal of Physiology*, *185*(3), 587–599. https://doi.org/10.1113/jphysiol.1966.sp008003

Nankoo, J.-F., Madan, C. R., Spetch, M. L., & Wylie, D. R. (2012). Perception of dynamic glass patterns. *Vision Research*, *72*, 55–62. https://doi.org/10.1016/j.visres.2012.09.008

Nguyen, B. N., Chan, Y. M., Bode, S., & McKendrick, A. M. (2020). Orientation-dependency of perceptual surround suppression and orientation decoding of centre-surround stimuli are preserved with healthy ageing. *Vision Research*, *176*, 72–79. https://doi.org/10.1016/j.visres.2020.07.015

Nguyen, B. N., & McKendrick, A. M. (2016). Foveal and parafoveal contrast suppression are different: Mechanisms revealed by the study of healthy aging. *Journal of Vision*, *16*(3), 10. https://doi.org/10.1167/16.3.10

Nurminen, L., Peromaa, T., & Laurinen, P. (2010). Surround suppression and facilitation in the fovea: very long-range spatial interactions in contrast perception. *Journal of Vision*, *10*(13), 9. https://doi.org/10.1167/10.13.9

Ohla, K., Busch, N. A., Dahlem, M. A., & Herrmann, C. S. (2005). Circles are different: the perception of Glass patterns modulates early event-related potentials. *Vision Research*, *45*(20), 2668–2676. https://doi.org/10.1016/j.visres.2005.03.015

Ohtani, Y., Okamura, S., Yoshida, Y., Toyama, K., & Ejima, Y. (2002). Surround suppression in the human visual cortex: an analysis using magnetoencephalography. *Vision Research*, *42*(15), 1825–1835. https://doi.org/10.1016/S0042-6989(02)00103-7

Orban, G. A., Van Essen, D., & Vanduffel, W. (2004). Comparative mapping of higher visual areas in monkeys and humans. *Trends in Cognitive Sciences*, *8*(7), 315–324. https://doi.org/10.1016/j.tics.2004.05.009

Ostwald, D., Lam, J. M., Li, S., & Kourtzi, Z. (2008). Neural coding of global form in the human visual cortex. *Journal of Neurophysiology*, *99*(5), 2456–2469. https://doi.org/10.1152/jn.01307.2007

Owsley, C. (2011). Aging and vision. *Vision Research*, *51*(13), 1610–1622. https://doi.org/10.1016/j.visres.2010.10.020

Owsley, C., Sekuler, R., & Siemsen, D. (1983). Contrast sensitivity throughout adulthood. *Vision Research*, *23*(7), 689–699. https://www.ncbi.nlm.nih.gov/pubmed/6613011

Palomares, M., Ales, J. M., Wade, A. R., Cottereau, B. R., & Norcia, A. M. (2012). Distinct effects of attention on the neural responses to form and motion processing: a SSVEP source-imaging study. *Journal of Vision*, *12*(10), 15. https://doi.org/10.1167/12.10.15

Parker, S., & Schneider, B. (1980). Loudness and loudness discrimination. *Perception & Psychophysics*, *28*(5), 398–406. https://doi.org/10.3758/BF03204883

Parra, L. C., Spence, C. D., Gerson, A. D., & Sajda, P. (2005). Recipes for the linear analysis of EEG. *NeuroImage*, *28*(2), 326–341. https://doi.org/10.1016/j.neuroimage.2005.05.032

Pavan, A., Contillo, A., Ghin, F., Foxwell, M. J., & Mather, G. (2019). Limited Attention Diminishes Spatial Suppression From Large Field Glass Patterns. *Perception*, *48*(4), 286–315. https://doi.org/10.1177/0301006619835457

Payne, B. R. (1993). Evidence for visual cortical area homologs in cat and macaque monkey. *Cerebral Cortex* , *3*(1), 1–25. https://doi.org/10.1093/cercor/3.1.1

Pei, F., Pettet, M. W., Vildavski, V. Y., & Norcia, A. M. (2005). Event-related potentials show configural specificity of global form processing. *Neuroreport*, *16*(13), 1427–1430. https://doi.org/10.1097/01.wnr.0000177003.12322.9b

Peirce, J. W. (2007). The potential importance of saturating and supersaturating contrast response functions in visual cortex. *Journal of Vision*, *7*(6), 13. https://doi.org/10.1167/7.6.13

Pelli, D. G. (1985). Uncertainty explains many aspects of visual contrast detection and discrimination. *JOSA A*, *2*(9), 1508–1532. https://doi.org/10.1364/JOSAA.2.001508

Petrov, Y., Carandini, M., & McKee, S. (2005). Two distinct mechanisms of suppression in human vision. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(38), 8704–8707. https://doi.org/10.1523/JNEUROSCI.2871-05.2005

Petrov, Y., & McKee, S. P. (2006). The effect of spatial configuration on surround suppression of contrast sensitivity. *Journal of Vision*, *6*(3), 224–238. https://doi.org/10.1167/6.3.4

Petrov, Y., & McKee, S. P. (2009). The time course of contrast masking reveals two distinct mechanisms of human surround suppression. *Journal of Vision*, *9*(1), 21.1-11. https://doi.org/10.1167/9.1.21

Petrov, Y., Verghese, P., & McKee, S. P. (2006). Collinear facilitation is largely uncertainty reduction. *Journal of Vision*, *6*(2), 170–178. https://doi.org/10.1167/6.2.8

Pinto, J. G. A., Hornby, K. R., Jones, D. G., & Murphy, K. M. (2010). Developmental changes in GABAergic mechanisms in human visual cortex across the lifespan. *Frontiers in Cellular Neuroscience*, *4*, 16. https://doi.org/10.3389/fncel.2010.00016

Pitchaimuthu, K., Nguyen, B. N., & McKendrick, A. M. (2017). Aging alters intraocular but not interocular foveal center surround contrast suppression. *Journal of Vision*, *17*(1), 16. https://doi.org/10.1167/17.1.16

Pitchaimuthu, K., Wu, Q.-Z., Carter, O., Nguyen, B. N., Ahn, S., Egan, G. F., & McKendrick, A. M. (2017). Occipital GABA levels in older adults and their relationship to visual perceptual suppression. *Scientific Reports*, *7*(1), 14231. https://doi.org/10.1038/s41598-017-14577-5

Prazdny, K. (1986). Some new phenomena in the perception of glass patterns. *Biological Cybernetics*, *53*(3), 153–158. https://doi.org/10.1007/BF00342883

Rampone, G., & Makin, A. D. J. (2020). Electrophysiological responses to regularity show specificity to global form: The case of Glass patterns. *The European Journal of Neuroscience*, *52*(3), 3032–3046. https://doi.org/10.1111/ejn.14709

Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, *61*(2), 168–185. https://doi.org/10.1016/j.neuron.2009.01.002

Ringach, D. L., Hawken, M. J., & Shapley, R. (1997). Dynamics of orientation tuning in macaque primary visual cortex. *Nature*, *387*(6630), 281–284. https://doi.org/10.1038/387281a0

Ringach, D. L., Shapley, R. M., & Hawken, M. J. (2002). Orientation Selectivity in Macaque V1: Diversity and Laminar Dependence. *The Journal of*

Neuroscience: The Official Journal of the Society for Neuroscience, *22*(13), 5639–5651. https://doi.org/10.1523/JNEUROSCI.22-13-05639.2002

Saarela, T. P., & Herzog, M. H. (2008). Time-course and surround modulation of contrast masking in human vision. *Journal of Vision*, *8*(3), 23.1-10. https://doi.org/10.1167/8.3.23

Sanborn, A. N., & Dayan, P. (2011). Optimal decisions for contrast discrimination. *Journal of Vision*, *11*(14). https://doi.org/10.1167/11.14.9

Sceniak, M. P., Ringach, D. L., Hawken, M. J., & Shapley, R. (1999). Contrast's effect on spatial summation by macaque V1 neurons. *Nature Neuroscience*, *2*(8), 733–739. https://doi.org/10.1038/11197

Schallmo, M.-P., Kale, A. M., & Murray, S. O. (2018). The time course of different surround suppression mechanisms. In *bioRxiv* (p. 432773). https://doi.org/10.1101/432773

Schallmo, M.-P., Kolodny, T., Kale, A. M., Millin, R., Flevaris, A. V., Edden, R. A. E., Gerdts, J., Bernier, R. A., & Murray, S. O. (2020). Weaker neural suppression in autism. *Nature Communications*, *11*(1), 2675. https://doi.org/10.1038/s41467-020-16495-z

Schallmo, M.-P., & Murray, S. O. (2016). Identifying separate components of surround suppression. *Journal of Vision*, *16*(1), 2. https://doi.org/10.1167/16.1.2

Schmolesky, M. T., Wang, Y., Pu, M., & Leventhal, A. G. (2000). Degradation of stimulus selectivity of visual cortical cells in senescent rhesus monkeys. *Nature Neuroscience*, *3*(4), 384–390. https://doi.org/10.1038/73957

Scholl, B., Tan, A. Y. Y., Corey, J., & Priebe, N. J. (2013). Emergence of orientation selectivity in the Mammalian visual pathway. *The Journal of Neuroscience:*

The Official Journal of the Society for Neuroscience, *33*(26), 10616–10624. https://doi.org/10.1523/JNEUROSCI.0404-13.2013

Schütt, H. H., Harmeling, S., Macke, J. H., & Wichmann, F. A. (2016). Painfree and accurate Bayesian estimation of psychometric functions for (potentially) overdispersed data. *Vision Research*, *122*, 105–123. https://doi.org/10.1016/j.visres.2016.02.002

Sclar, G., & Freeman, R. D. (1982). Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Experimental Brain Research. Experimentelle Hirnforschung. Experimentation Cerebrale*, *46*(3), 457–461. https://doi.org/10.1007/BF00238641

Sengpiel, F., & Vorobyov, V. (2005). Intracortical origins of interocular suppression in the visual cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(27), 6394–6400. https://doi.org/10.1523/JNEUROSCI.0862-05.2005

Serrano-Pedraza, I., Romero-Ferreiro, V., Read, J. C. A., Diéguez-Risco, T., Bagney, A., Caballero-González, M., Rodríguez-Torresano, J., & Rodriguez-Jimenez, R. (2014). Reduced visual surround suppression in schizophrenia shown by measuring contrast detection thresholds. *Frontiers in Psychology*, *5*, 1431. https://doi.org/10.3389/fpsyg.2014.01431

Seu, L., & Ferrera, V. P. (2001). Detection thresholds for spiral Glass patterns. *Vision Research*, *41*(28), 3785–3790. https://doi.org/10.1016/s0042-6989(01)00235-8

Shushruth, S., Nurminen, L., Bijanzadeh, M., Ichida, J. M., Vanni, S., & Angelucci, A. (2013). Different orientation tuning of near- and far-surround suppression in macaque primary visual cortex mirrors their tuning in human perception.

The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, *33*(1), 106–119. https://doi.org/10.1523/JNEUROSCI.2518-12.2013

Simpson, W. A., & Finsten, B. A. (1995). Pedestal effect in visual motion discrimination. *Journal of the Optical Society of America. A, Optics, Image Science, and Vision*, *12*(12), 2555–2563. https://doi.org/10.1364/josaa.12.002555

Skottun, B. C., Bradley, A., Sclar, G., Ohzawa, I., & Freeman, R. D. (1987). The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. *Journal of Neurophysiology*, *57*(3), 773–786. https://doi.org/10.1152/jn.1987.57.3.773

Smith, M. A., Bair, W., & Movshon, J. A. (2002). Signals in macaque striate cortical neurons that support the perception of glass patterns. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *22*(18), 8334–8345. https://www.ncbi.nlm.nih.gov/pubmed/12223588

Smith, M. A., Kohn, A., & Movshon, J. A. (2007). Glass pattern responses in macaque V2 neurons. *Journal of Vision*, *7*(3), 5. https://doi.org/10.1167/7.3.5

Snowden, R. J., & Hammett, S. T. (1998). The effects of surround contrast on contrast thresholds, perceived contrast and contrast discrimination. *Vision Research*, *38*(13), 1935–1945. https://doi.org/10.1016/S0042-6989(97)00379-9

Solomon, J. A. (2009). The history of dipper functions. *Attention, Perception & Psychophysics*, *71*(3), 435–443. https://doi.org/10.3758/APP.71.3.435

Somers, D. C., Todorov, E. V., Siapas, A. G., Toth, L. J., Kim, D. S., & Sur, M. (1998). A local circuit approach to understanding integration of long-range

inputs in primary visual cortex. *Cerebral Cortex , 8*(3), 204–217. https://doi.org/10.1093/cercor/8.3.204

Spear, P. D. (1993). Neural bases of visual deficits during aging. *Vision Research*, *33*(18), 2589–2609. https://doi.org/10.1016/0042-6989(93)90218-L

Swettenham, J. B., Anderson, S. J., & Thai, N. J. (2010). MEG responses to the perception of global structure within glass patterns. *PloS One*, *5*(11), e13865. https://doi.org/10.1371/journal.pone.0013865

Tadin, D., Lappin, J. S., Gilroy, L. A., & Blake, R. (2003). Perceptual consequences of centre-surround antagonism in visual motion processing. *Nature*, *424*(6946), 312–315. https://doi.org/10.1038/nature01800

Tootell, R. B., Hamilton, S. L., & Switkes, E. (1988). Functional anatomy of macaque striate cortex. IV. Contrast and magno-parvo streams. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *8*(5), 1594–1609. https://www.ncbi.nlm.nih.gov/pubmed/3367212

Vanegas, M. I., Blangero, A., & Kelly, S. P. (2015). Electrophysiological indices of surround suppression in humans. *Journal of Neurophysiology*, *113*(4), 1100–1109. https://doi.org/10.1152/jn.00774.2014

Verghese, P., Kim, Y.-J., & Wade, A. R. (2012). Attention selects informative neural populations in human V1. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *32*(46), 16379–16390. https://doi.org/10.1523/JNEUROSCI.1174-12.2012

Victor, J. D., & Mast, J. (1991). A new statistic for steady-state evoked potentials. *Electroencephalography and Clinical Neurophysiology*, *78*(5), 378–388. https://doi.org/10.1016/0013-4694(91)90099-p

Webb, B. S., Dhruv, N. T., Solomon, S. G., Tailby, C., & Lennie, P. (2005). Early and late mechanisms of surround suppression in striate cortex of macaque. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(50), 11666–11675. https://doi.org/10.1523/JNEUROSCI.3414-05.2005

Wiesel, T. N., & Hubel, D. H. (1966). Spatial and chromatic interactions in the lateral geniculate body of the rhesus monkey. *Journal of Neurophysiology*, *29*(6), 1115–1156. https://doi.org/10.1152/jn.1966.29.6.1115

Williford, T., & Maunsell, J. H. R. (2006). Effects of spatial attention on contrast response functions in macaque area V4. *Journal of Neurophysiology*, *96*(1), 40–54. https://doi.org/10.1152/jn.01207.2005

Wilson, H. R., & Wilkinson, F. (1998). Detection of global structure in Glass patterns: implications for form vision. *Vision Research*, *38*(19), 2933–2947. https://doi.org/10.1016/s0042-6989(98)00109-6

Wilson, H. R., & Wilkinson, F. (2003). Further evidence for global orientation processing in circular Glass patterns [Review of *Further evidence for global orientation processing in circular Glass patterns*]. *Vision Research*, *43*(5), 563–564; author reply 565-6. https://doi.org/10.1016/s0042-6989(02)00651-x

Wilson, H. R., Wilkinson, F., & Asaad, W. (1997). Concentric orientation summation in human form vision. *Vision Research*, *37*(17), 2325–2330. https://doi.org/10.1016/s0042-6989(97)00104-1

Wilson, J. A., Switkes, E., & De Valois, R. L. (2004). Glass pattern studies of local and global processing of contrast variations. *Vision Research*, *44*(22), 2629–2641. https://doi.org/10.1016/j.visres.2003.06.001

Xing, J., & Heeger, D. J. (2000). Center-surround interactions in foveal and peripheral vision. *Vision Research*, *40*(22), 3065–3072. https://doi.org/10.1016/s0042-6989(00)00152-8

Xing, J., & Heeger, D. J. (2001). Measurement and modeling of center-surround suppression and enhancement. *Vision Research*, *41*(5), 571–583. https://doi.org/10.1016/s0042-6989(00)00270-4

Xu, X., Ichida, J., Shostak, Y., Bonds, A. B., & Casagrande, V. A. (2002). Are primate lateral geniculate nucleus (LGN) cells really sensitive to orientation or direction? *Visual Neuroscience*, *19*(1), 97–108. https://doi.org/10.1017/s0952523802191097

Yoon, J. H., Maddock, R. J., Rokem, A., Silver, M. A., Minzenberg, M. J., Ragland, J. D., & Carter, C. S. (2010). GABA concentration is reduced in visual cortex in schizophrenia and correlates with orientation-specific surround suppression. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *30*(10), 3777–3781. https://doi.org/10.1523/JNEUROSCI.6158-09.2010

Yu, S., Wang, Y., Li, X., Zhou, Y., & Leventhal, A. G. (2006). Functional degradation of extrastriate visual cortex in senescent rhesus monkeys. *Neuroscience*, *140*(3), 1023–1029. https://doi.org/10.1016/j.neuroscience.2006.01.015

Zhuang, X., Chen, Y., Zhuang, X., Xing, T., Chen, T., Jiang, G., & Yang, X. (2016). Impaired Center-Surround Suppression in Patients with Alzheimer's Disease. *Journal of Alzheimer's Disease: JAD*, *55*(3), 1101–1108. https://doi.org/10.3233/JAD-160603

Appendices

Appendix A: Individual contrast-matching psychometric functions

For each observer (rows), in each subplot, the grey circles represent the accuracy estimates obtained from our staircase experiment. They are scaled according to the number of trials spent at a given target increment level. The solid lines represent the best fit of a Weibull function, and the vertical dashed line represents the PSE obtained from this fit.

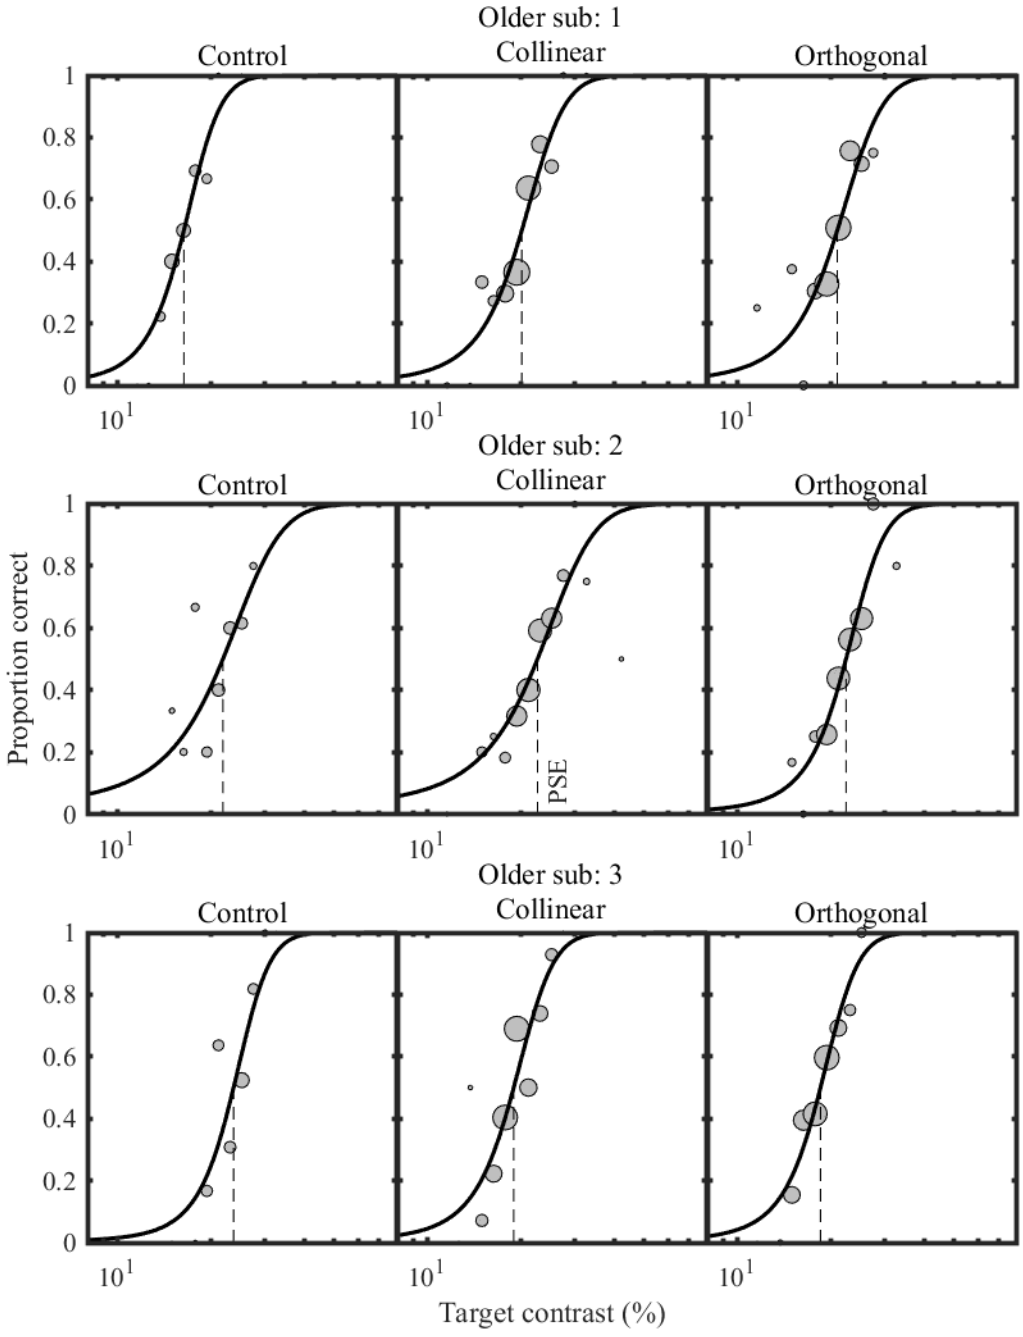

0

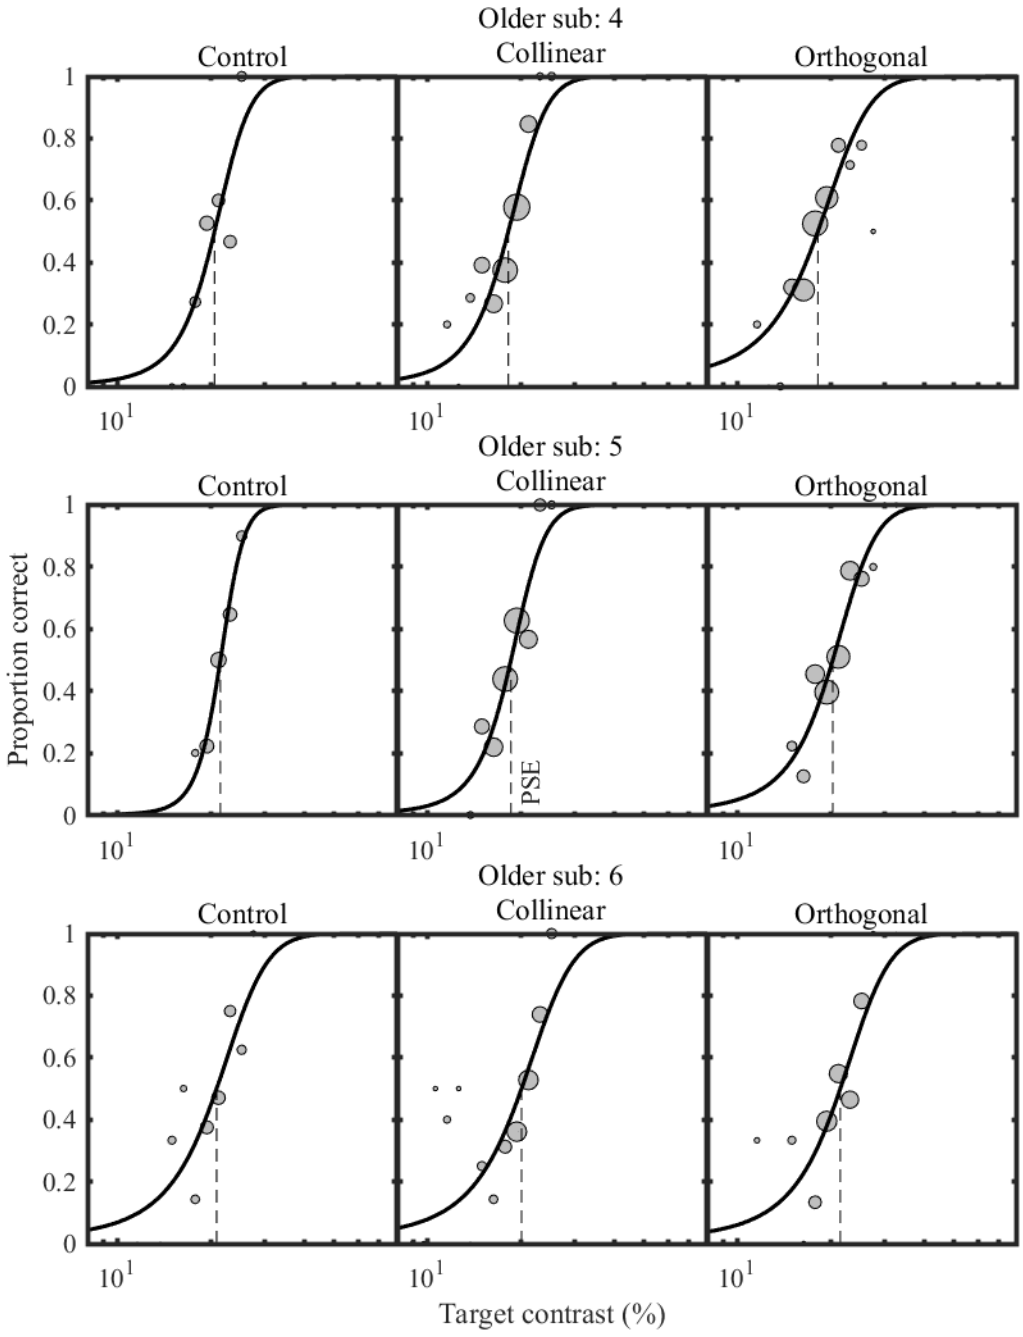
Older sub: 4
Control
Collinear
Orthogonal
Older sub: 5
Control
Collinear
Orthogonal
PSE
Older sub: 6
Control
Collinear
Orthogonal
Proportion correct
Target contrast (%)
0
0.2
0.4
0.6
0.8
1
10^1

0

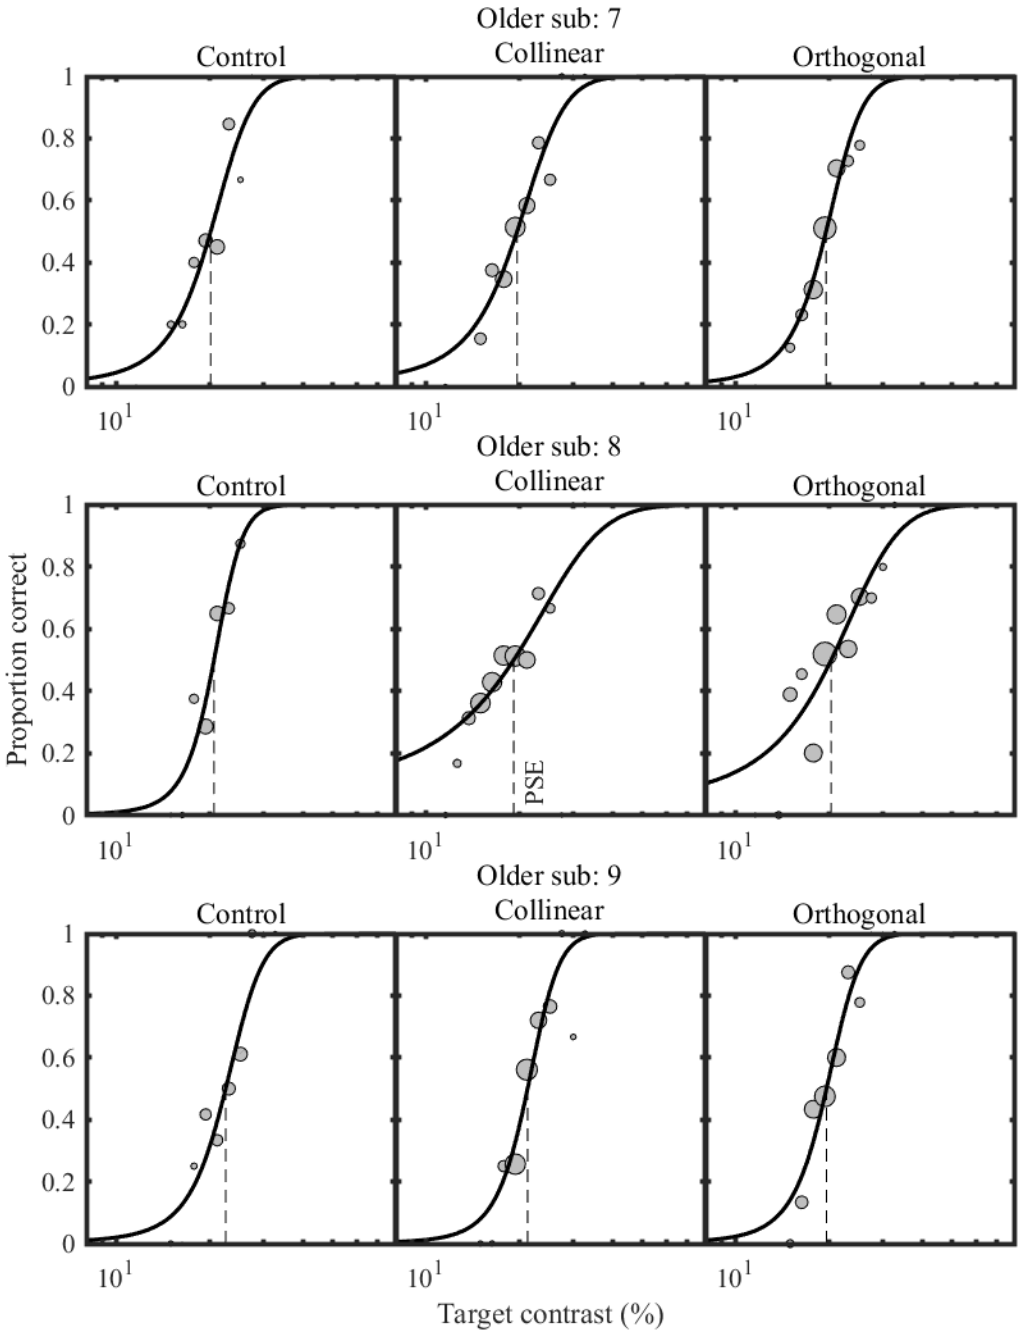
Older sub: 7
Control
Collinear
Orthogonal
Older sub: 8
Control
Collinear
Orthogonal
PSE
Older sub: 9
Control
Collinear
Orthogonal
Proportion correct
Target contrast (%)
0
0.2
0.4
0.6
0.8
1
10^1

0

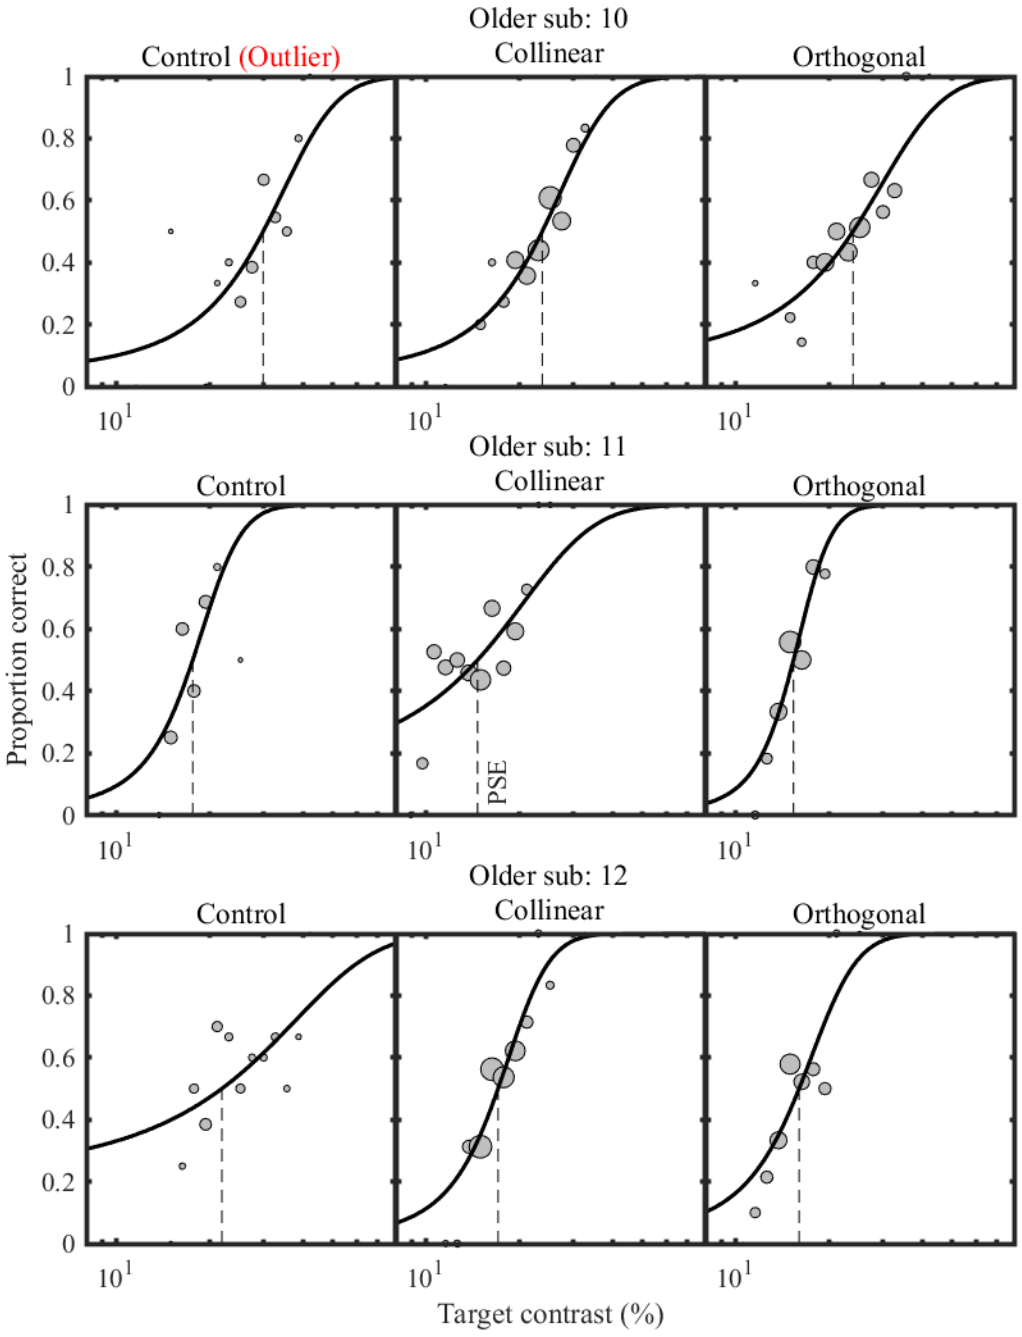
Older sub: 10
Control (Outlier)
Collinear
Orthogonal
Older sub: 11
Control
Collinear
Orthogonal
PSE
Older sub: 12
Control
Collinear
Orthogonal
Proportion correct
Target contrast (%)
0
0.2
0.4
0.6
0.8
1
10^1

0

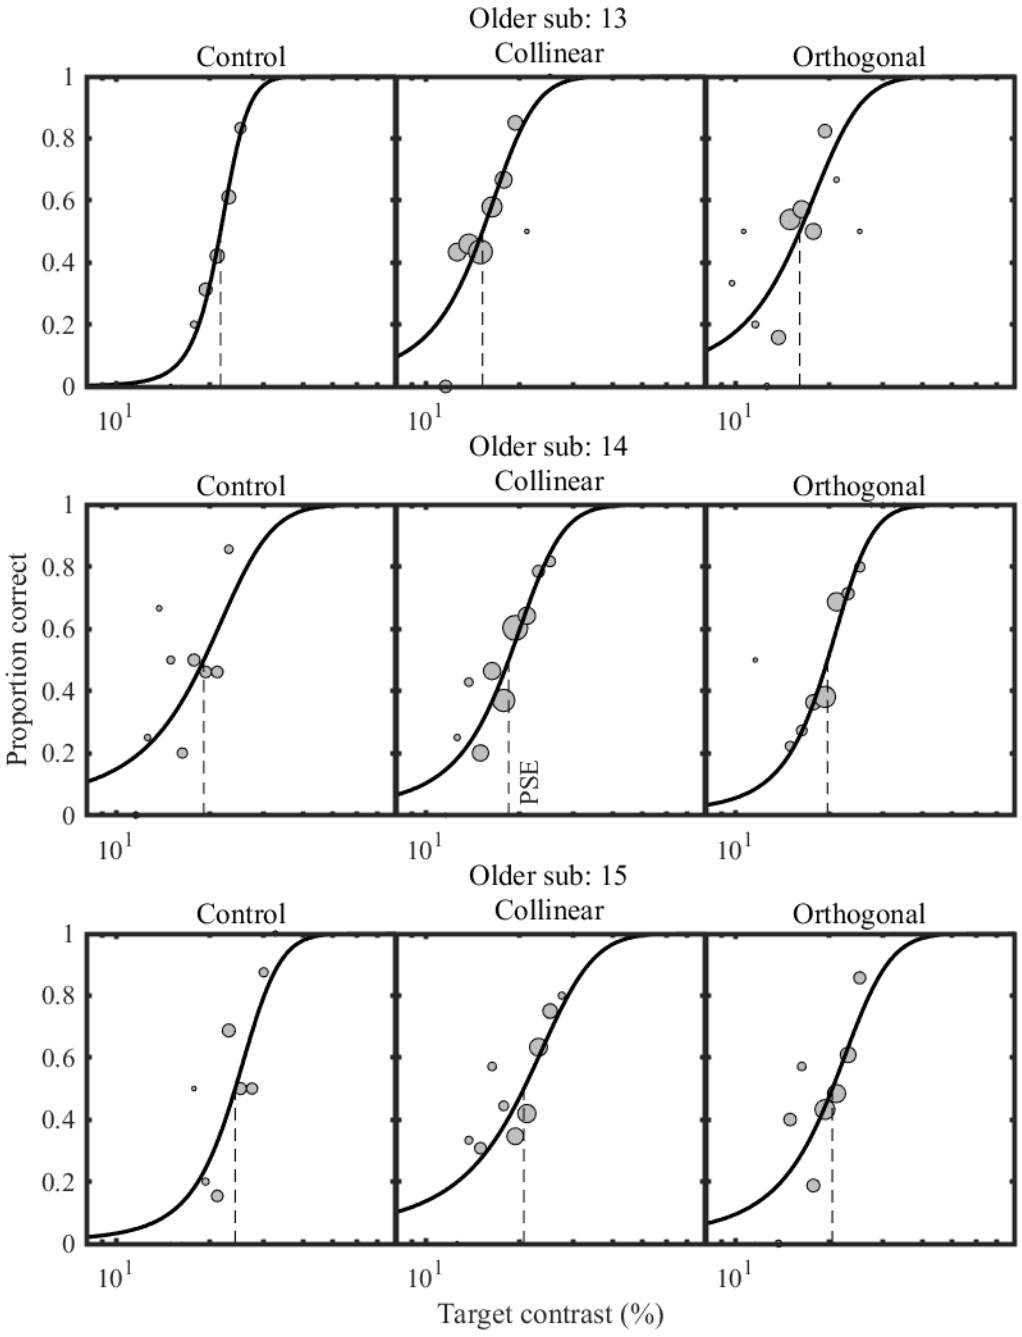
Older sub: 13
Control
Collinear
Orthogonal
Older sub: 14
Control
Collinear
Orthogonal
PSE
Older sub: 15
Control
Collinear
Orthogonal
Proportion correct
Target contrast (%)
0
0.2
0.4
0.6
0.8
1
10^1

0

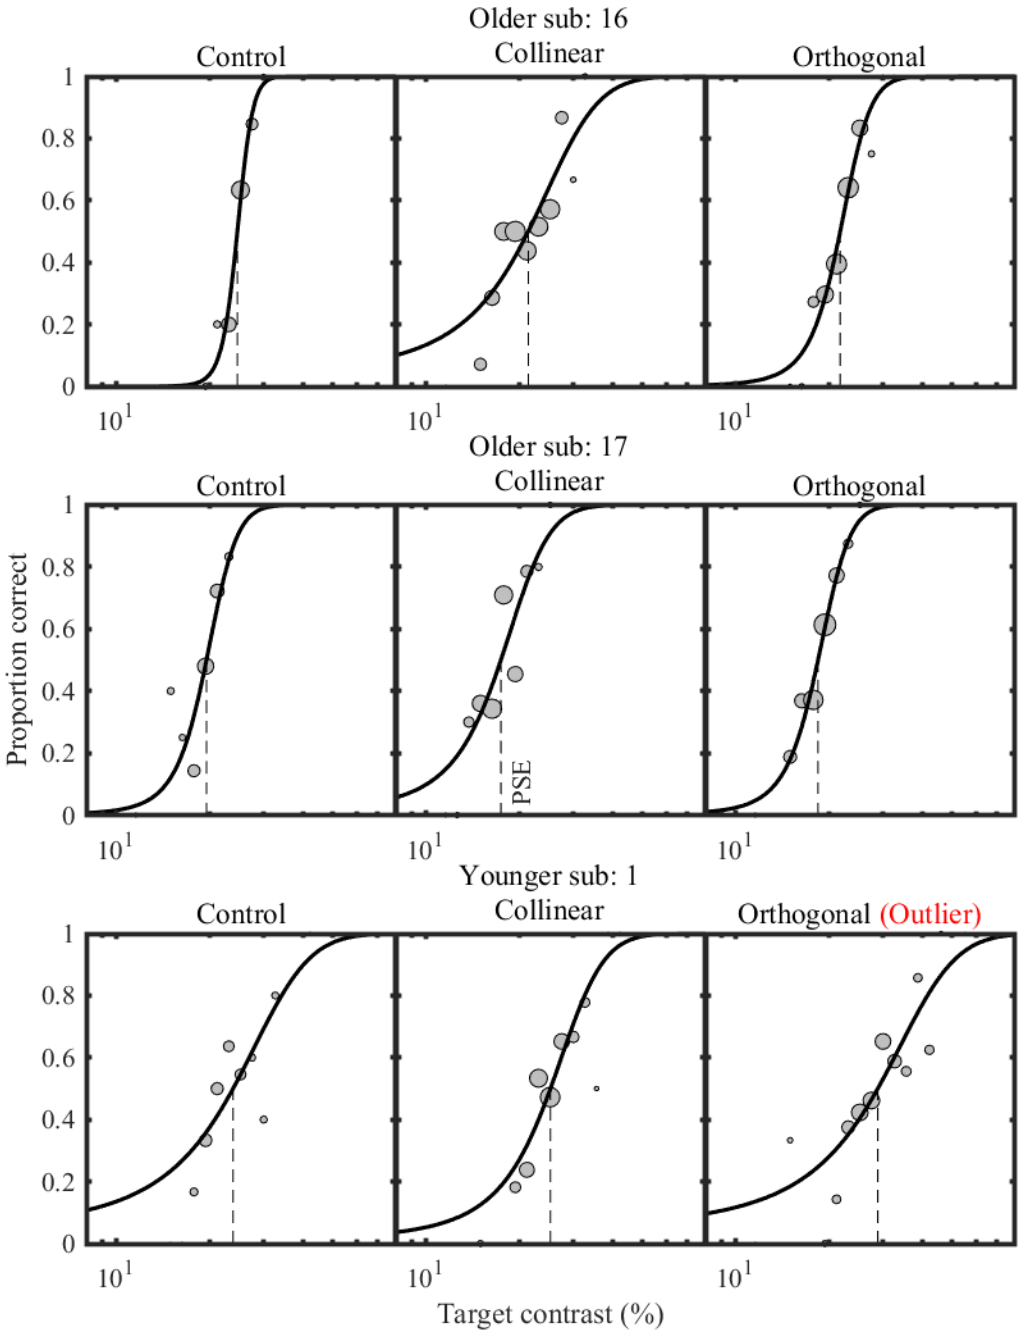
Older sub: 16
Control
Collinear
Orthogonal
Older sub: 17
Control
Collinear
Orthogonal
PSE
Younger sub: 1
Control
Collinear
Orthogonal (Outlier)
Proportion correct
Target contrast (%)

0

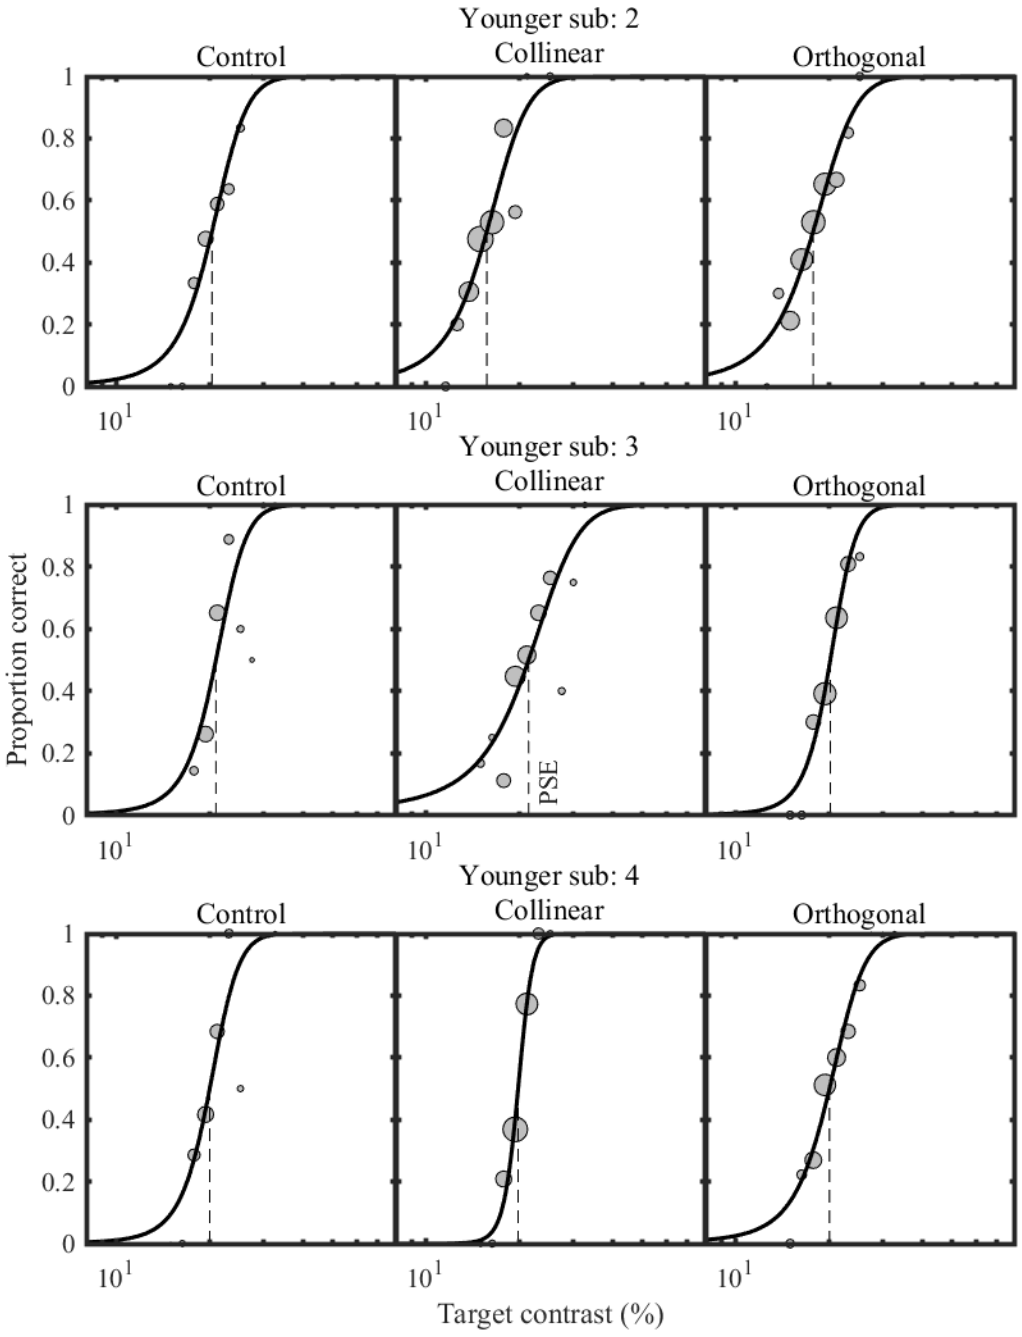
Younger sub: 2
Control
Collinear
Orthogonal
Younger sub: 3
Control
Collinear
Orthogonal
PSE
Younger sub: 4
Control
Collinear
Orthogonal
Proportion correct
Target contrast (%)

0

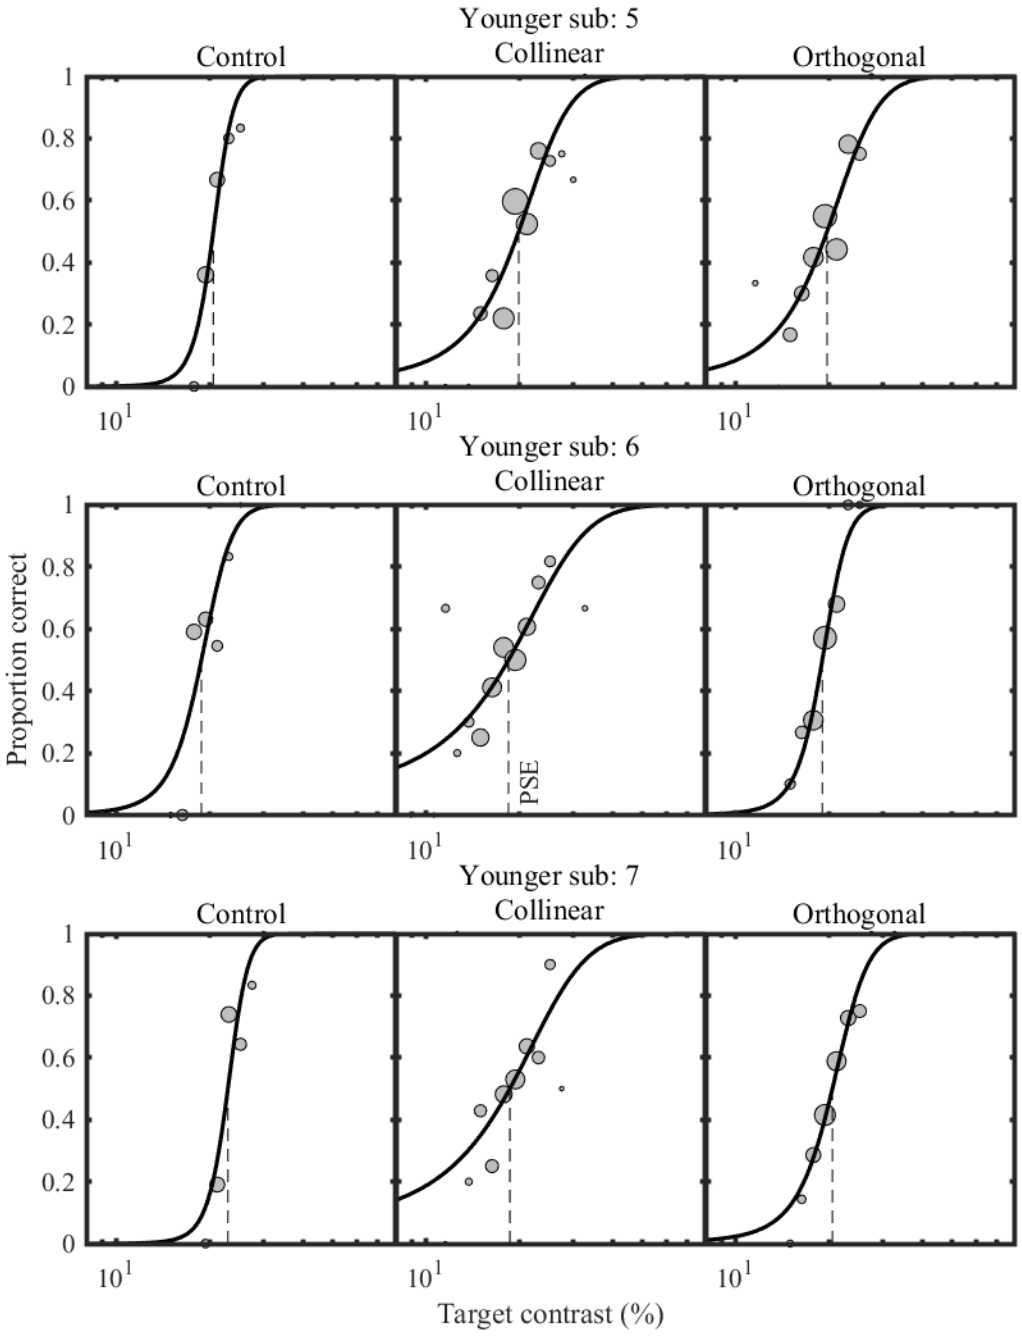
Younger sub: 5
Control
Collinear
Orthogonal
Younger sub: 6
Control
Collinear
Orthogonal
PSE
Younger sub: 7
Control
Collinear
Orthogonal
Proportion correct
Target contrast (%)

0

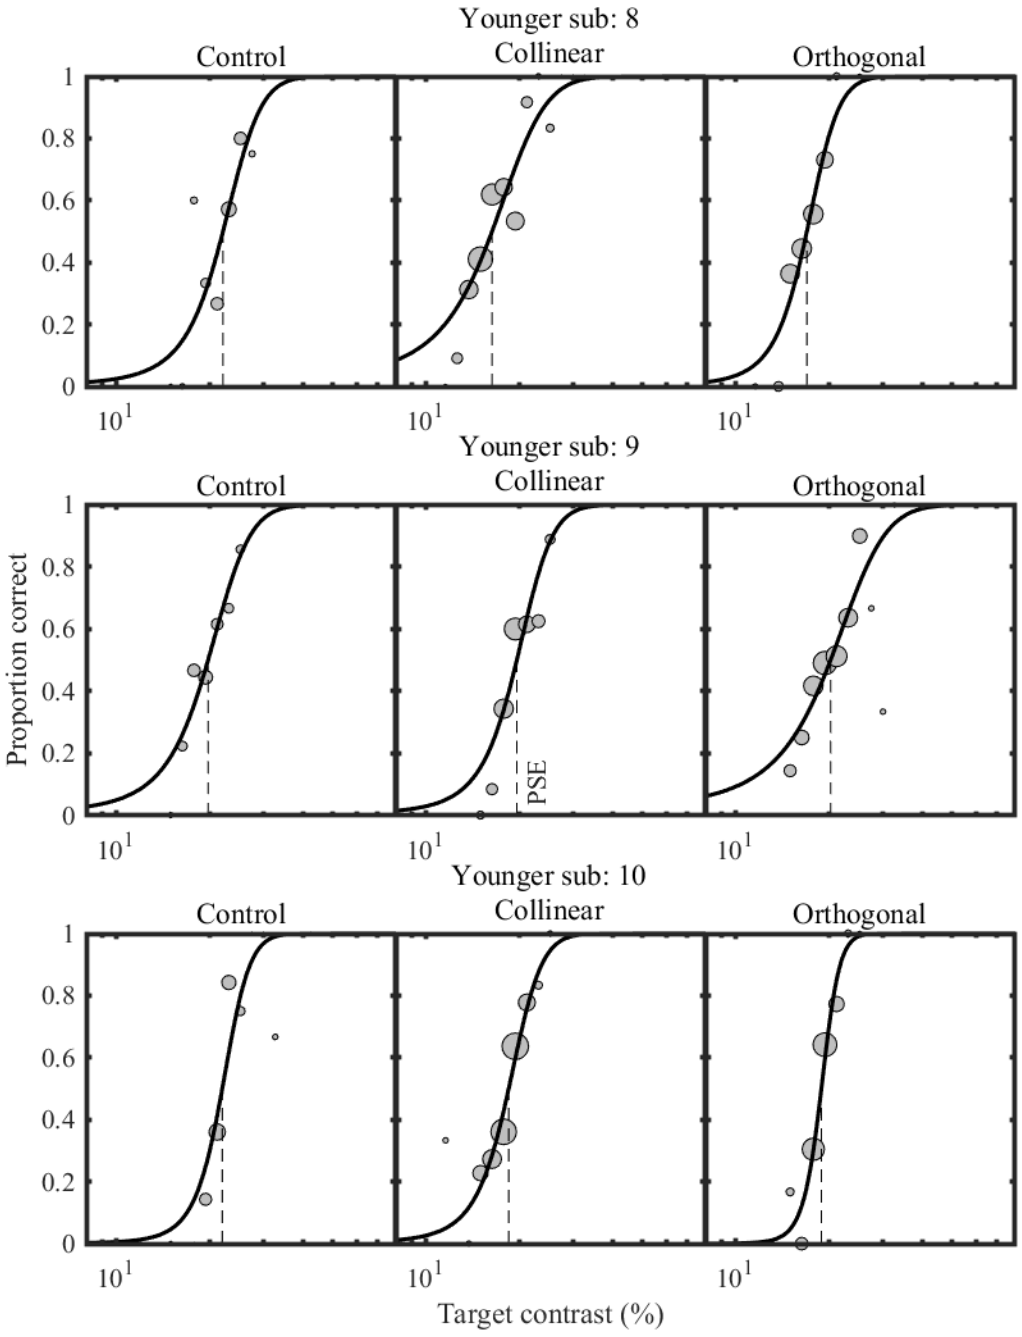
Younger sub: 8
Control
Collinear
Orthogonal
Younger sub: 9
Control
Collinear
Orthogonal
PSE
Younger sub: 10
Control
Collinear
Orthogonal
Proportion correct
Target contrast (%)

0

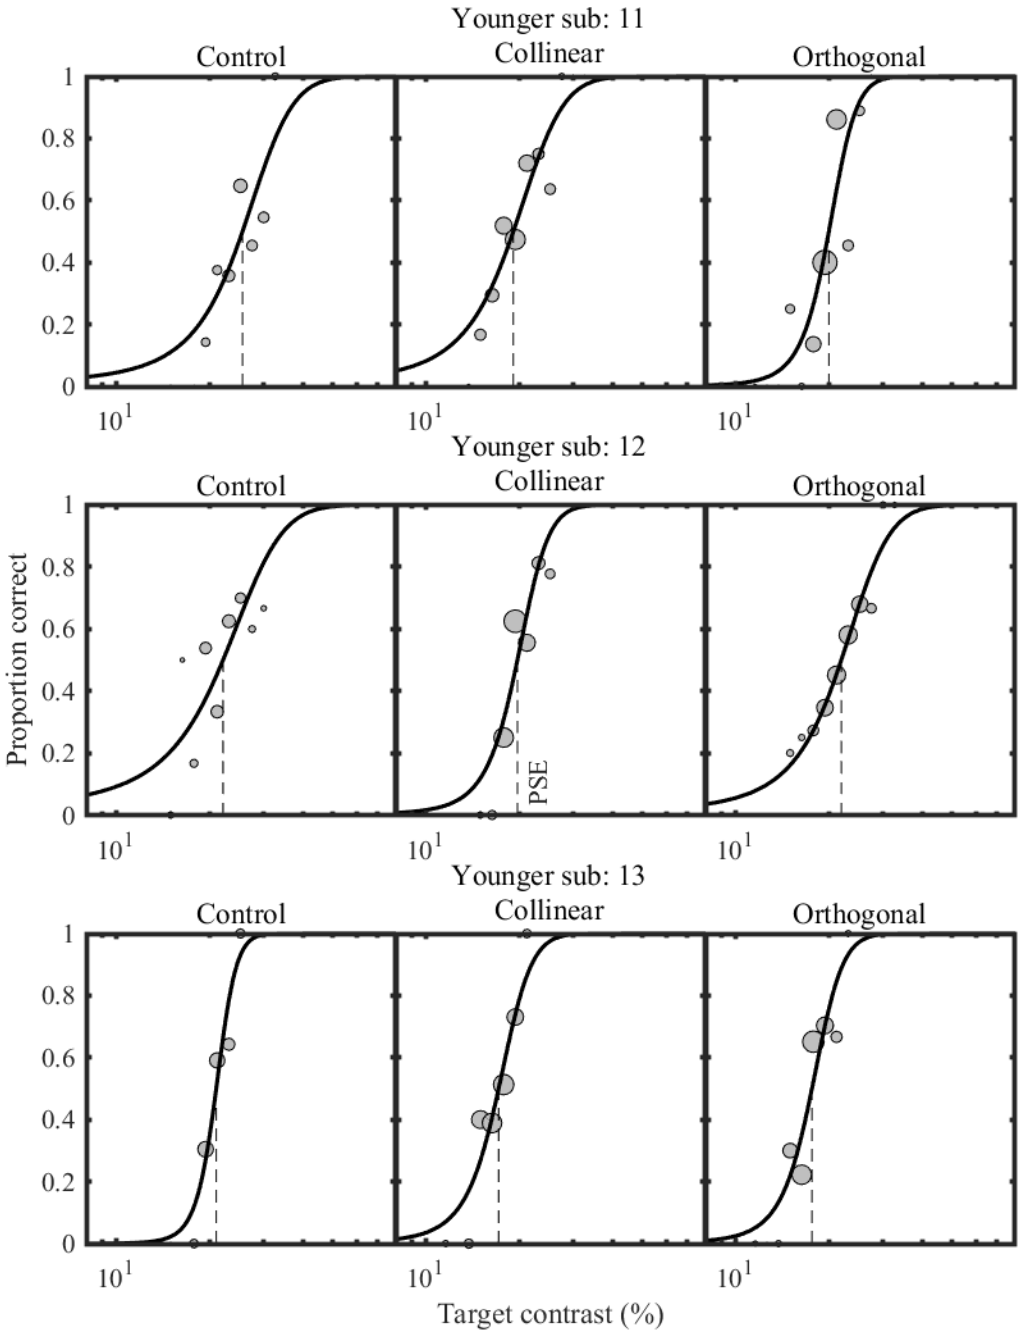
Younger sub: 11
Control
Collinear
Orthogonal
Younger sub: 12
Control
Collinear
Orthogonal
PSE
Younger sub: 13
Control
Collinear
Orthogonal
Proportion correct
Target contrast (%)

0

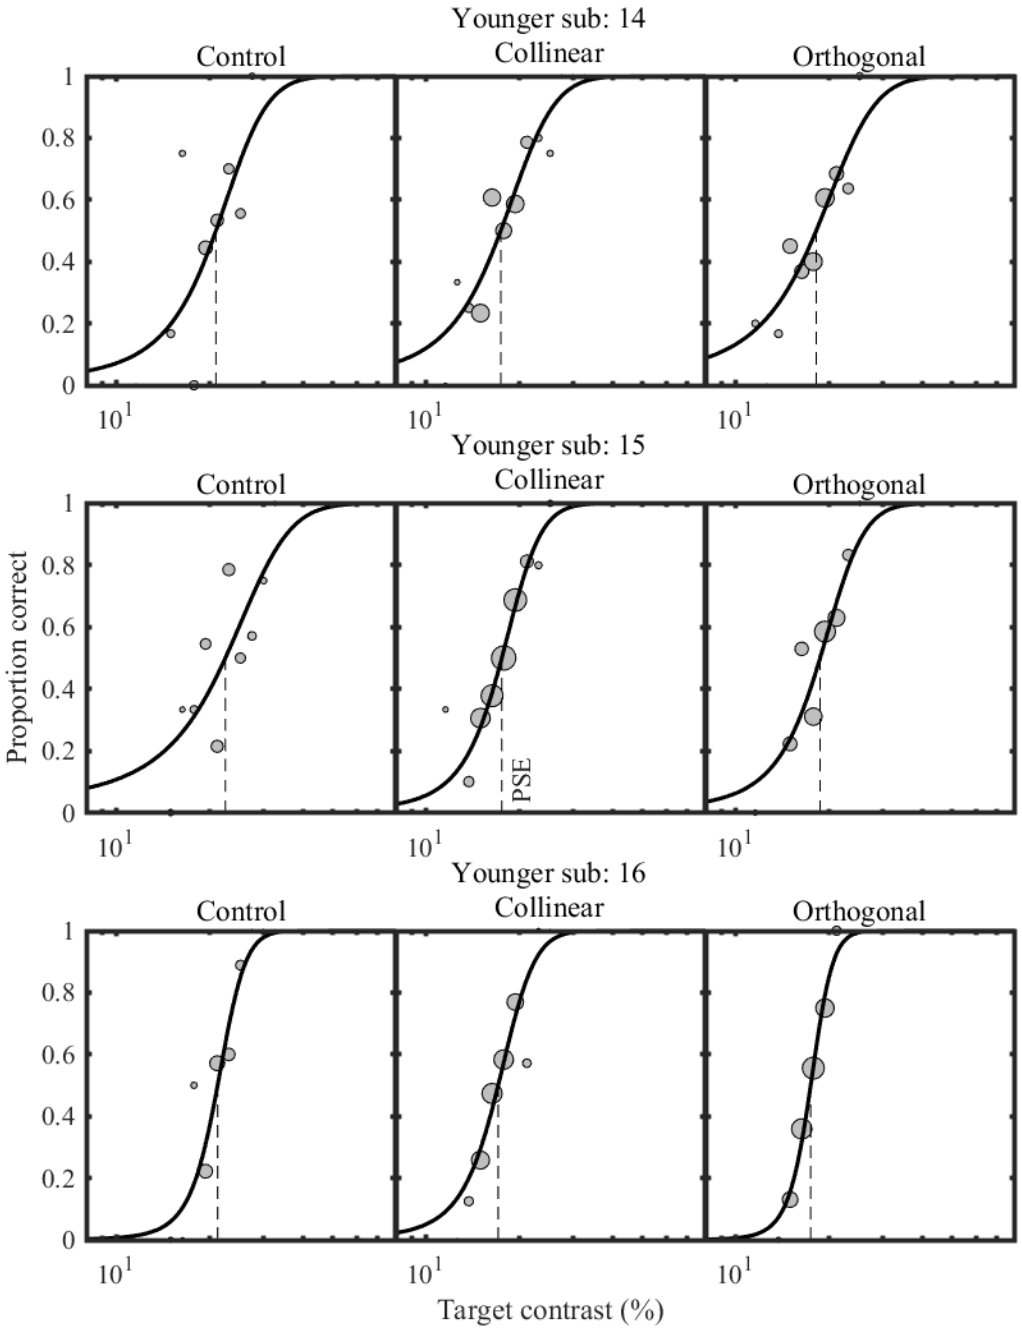
Younger sub: 14
Control
Collinear
Orthogonal
Younger sub: 15
Control
Collinear
Orthogonal
PSE
Younger sub: 16
Control
Collinear
Orthogonal
Proportion correct
Target contrast (%)

0

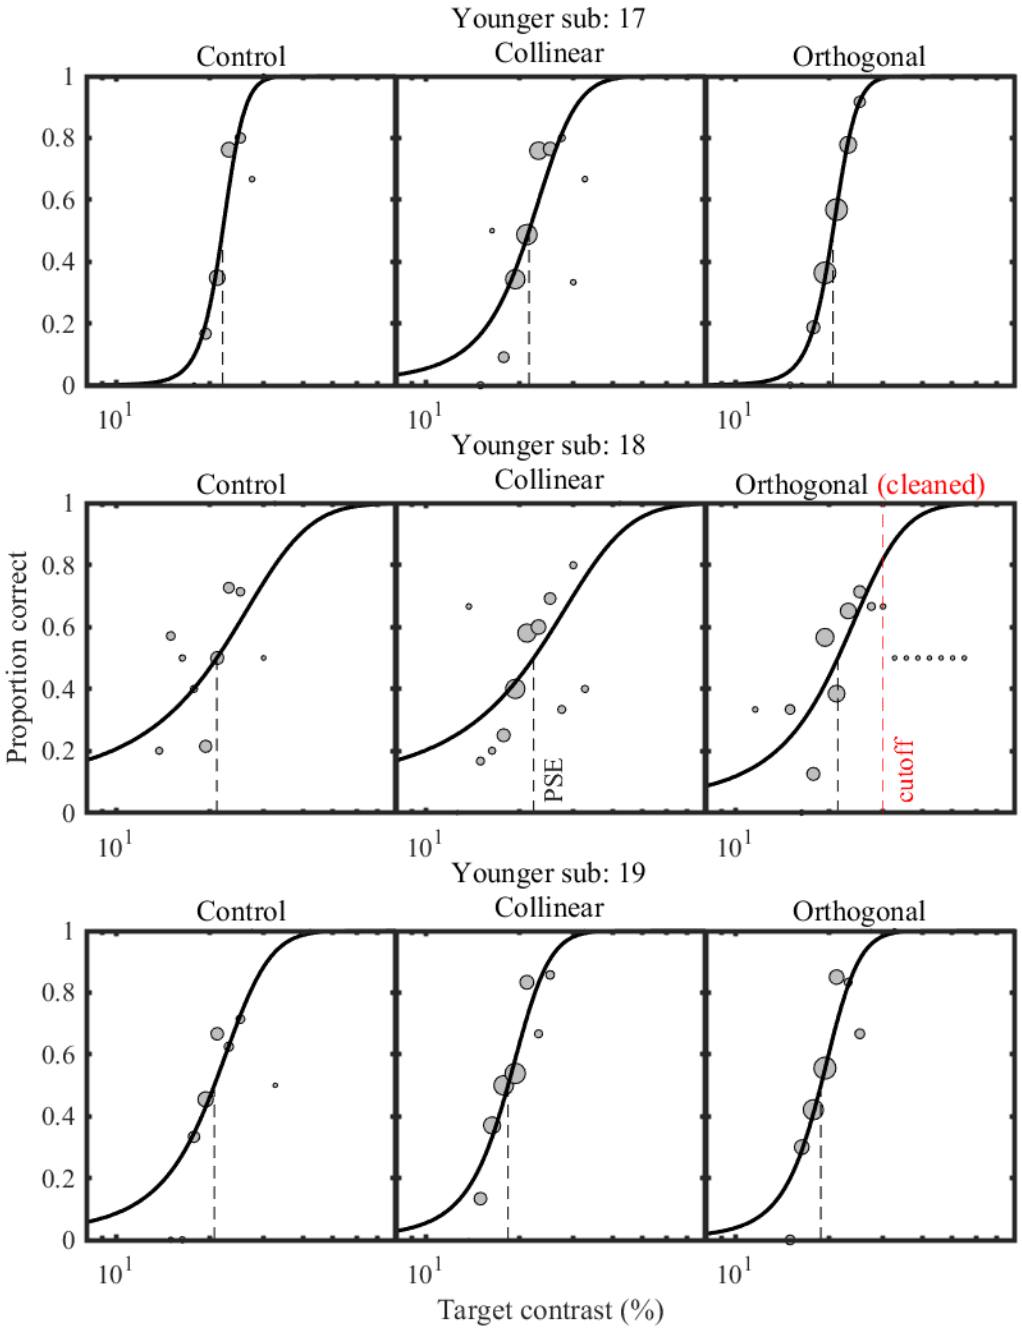
Younger sub: 17
Control
Collinear
Orthogonal
Younger sub: 18
Control
Collinear
Orthogonal (cleaned)
PSE
cutoff
Younger sub: 19
Control
Collinear
Orthogonal
Proportion correct
Target contrast (%)
1
0.8
0.6
0.4
0.2
0
10^1

Appendix B: ANOVA tables for Chapter 2 with and without outliers

ANOVA tables with outliers removed

Two subjects had thresholds that violated the assumption of normality of residuals. In these ANOVA tables, those outliers have been removed.

Tests of Within-Subjects Effects

Measure: supp

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
Surround	Sphericity Assumed	.010	2	.005	26.545	<.001
	Greenhouse-Geisser	.010	1.280	.008	26.545	<.001
	Huynh-Feldt	.010	1.349	.008	26.545	<.001
	Lower-bound	.010	1.000	.010	26.545	<.001
Surround * Age group	Sphericity Assumed	.000	2	5.329E-5	.273	.762
	Greenhouse-Geisser	.000	1.280	8.327E-5	.273	.662
	Huynh-Feldt	.000	1.349	7.900E-5	.273	.674
	Lower-bound	.000	1.000	.000	.273	.605
Error (Surround)	Sphericity Assumed	.013	66	.000		
	Greenhouse-Geisser	.013	42.240	.000		
	Huynh-Feldt	.013	44.527	.000		
	Lower-bound	.013	33.000	.000		

Tests of Between-Subjects Effects

Measure: supp

Transformed Variable: Average

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Intercept	4.119	1	4.119	6312.394	<.001
Age_group	3.913E-5	1	3.913E-5	.060	.808
Error	.022	33	.001		

0

ANOVA tables with outliers retained

Two subjects had thresholds that violated the assumption of normality of residuals. In these ANOVA tables, those outliers have been retained.

Tests of Within-Subjects Effects

Measure: supp

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
Surround	Sphericity Assumed	.011	2	.005	23.644	<.001
	Greenhouse-Geisser	.011	1.272	.009	23.644	<.001
	Huynh-Feldt	.011	1.336	.008	23.644	<.001
	Lower-bound	.011	1.000	.011	23.644	<.001
Surround * Age group	Sphericity Assumed	.000	2	5.165E-5	.223	.801
	Greenhouse-Geisser	.000	1.272	8.123E-5	.223	.697
	Huynh-Feldt	.000	1.336	7.733E-5	.223	.709
	Lower-bound	.000	1.000	.000	.223	.640
Error (Surround)	Sphericity Assumed	.016	70	.000		
	Greenhouse-Geisser	.016	44.512	.000		
	Huynh-Feldt	.016	46.755	.000		
	Lower-bound	.016	35.000	.000		

Tests of Between-Subjects Effects

Measure: supp

Transformed Variable: Average

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Intercept	4.502	1	4.502	3744.758	<.001
Age group	1.327E-5	1	1.327E-5	.011	.917
Error	.042	35	.001		

Appendix C: Global form discrimination psychometric functions for individual observers

For each observer, in each subplot, the small circles represent the accuracy estimates obtained from our method of constant stimuli experiment. The solid lines represent the best fit of a Weibull function, and the intercept of the vertical and horizontal solid lines represents the JND obtained from this fit. Plots shaded red did not contribute to the bootstrapped fit of a hyperbolic ratio function.

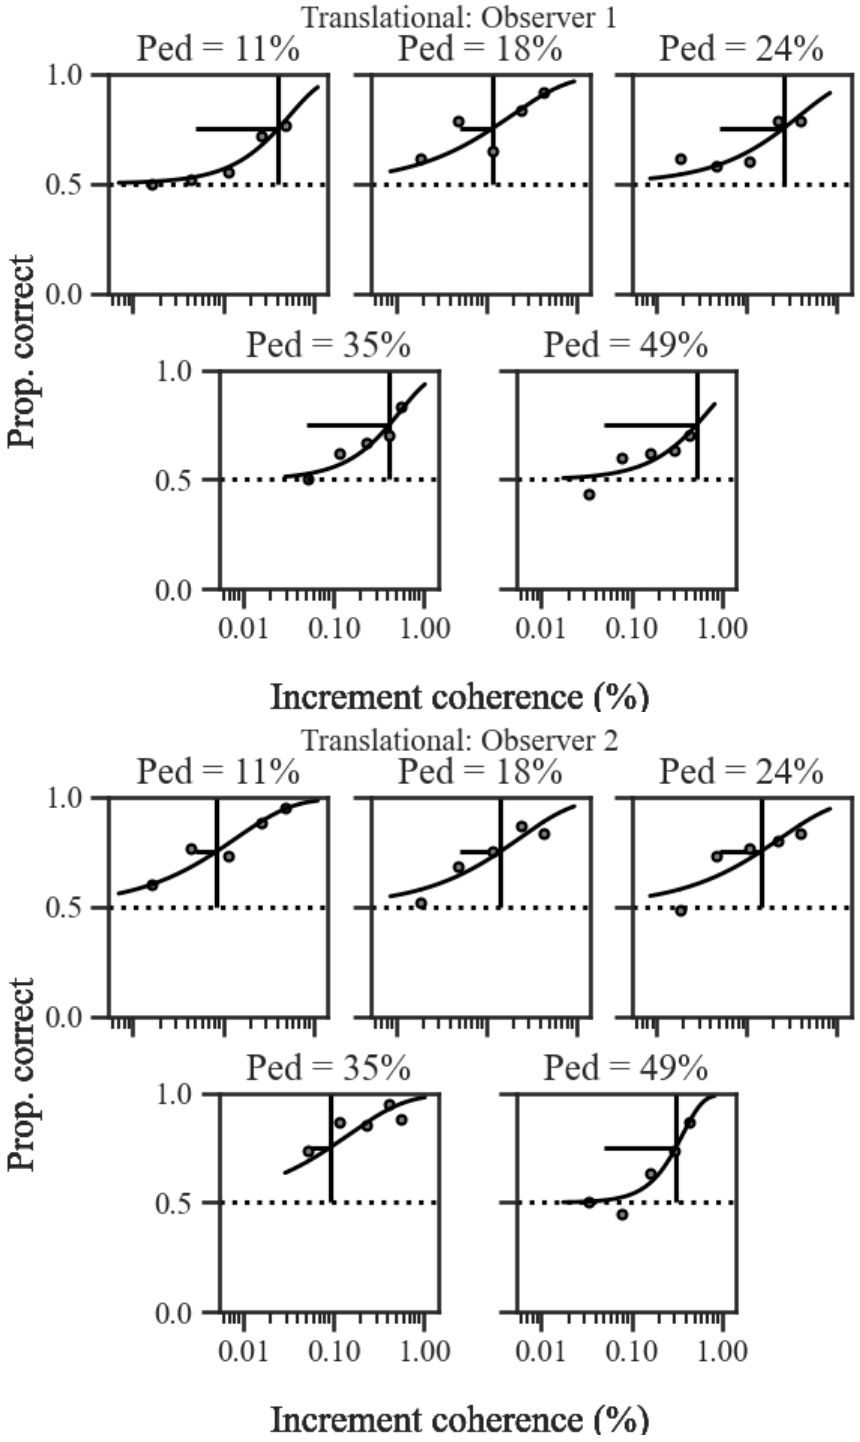

0

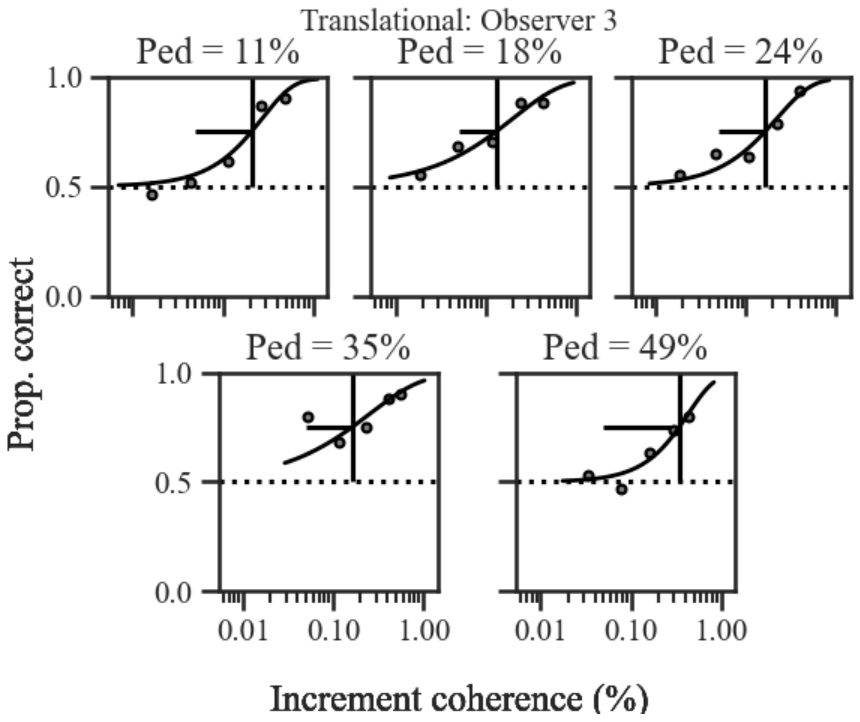
Translational: Observer 3
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

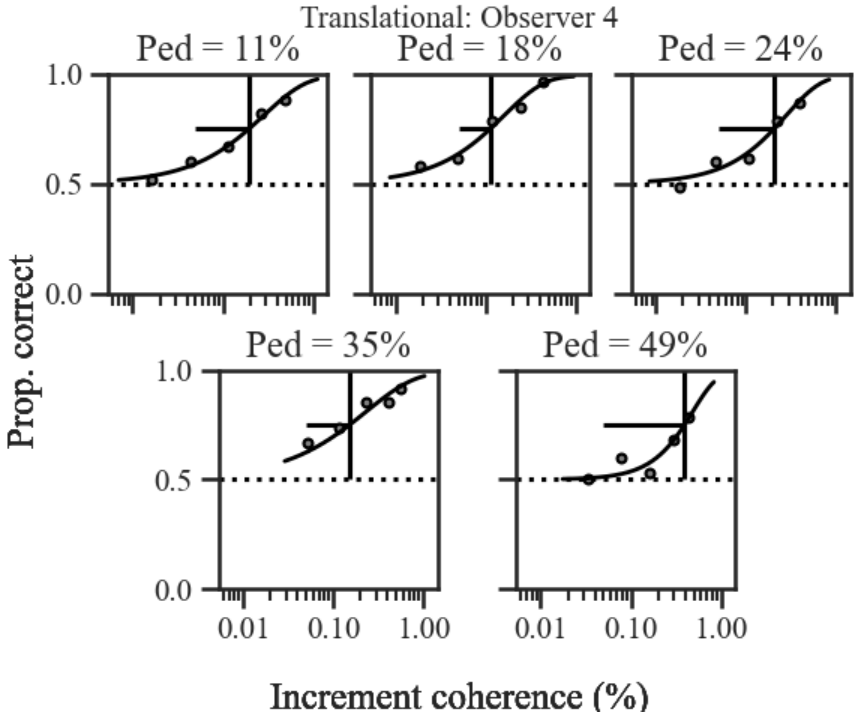
Translational: Observer 4
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

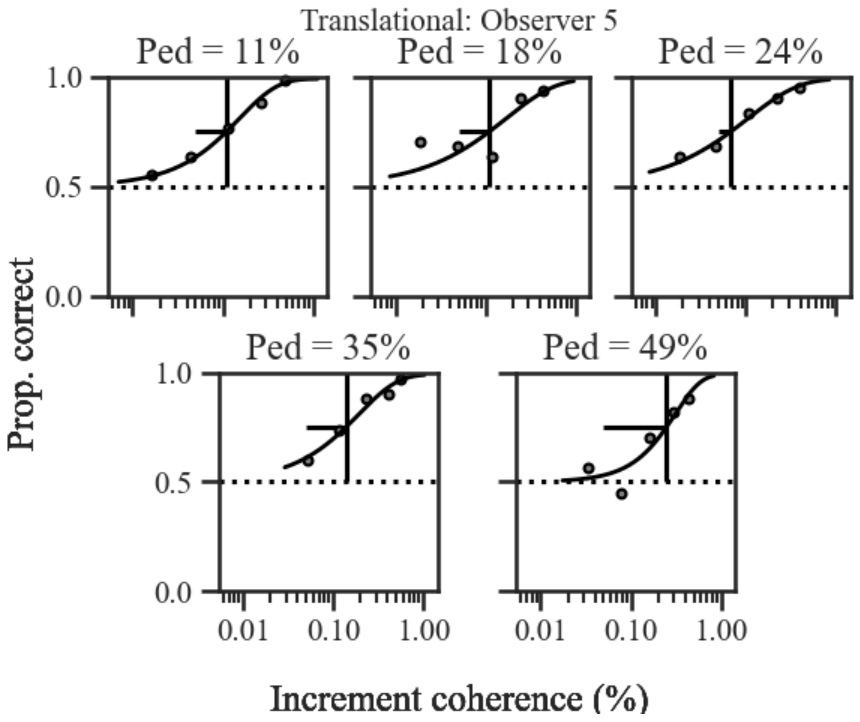
Translational: Observer 5
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

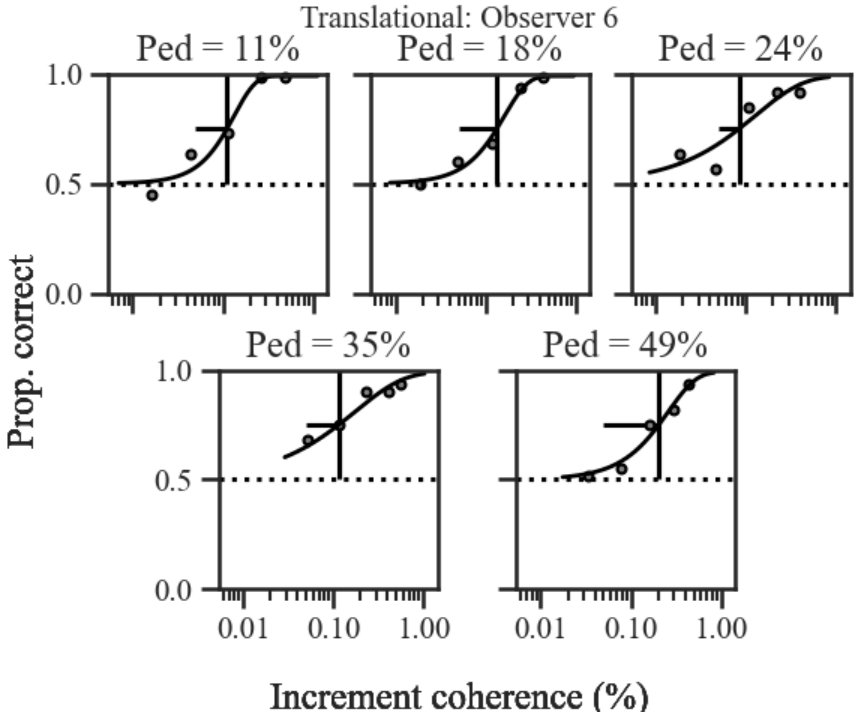
Translational: Observer 6
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

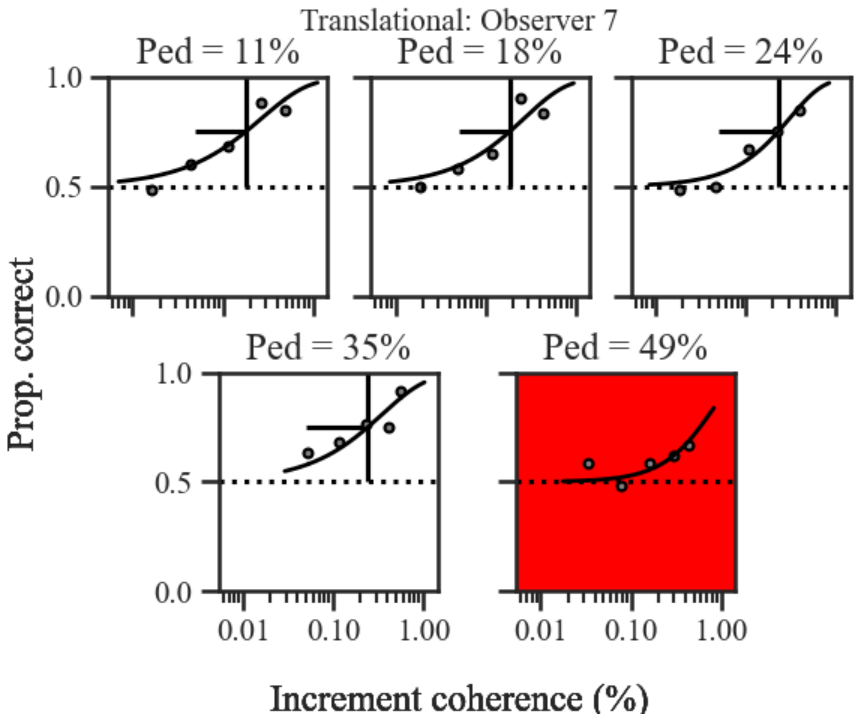
Translational: Observer 7
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

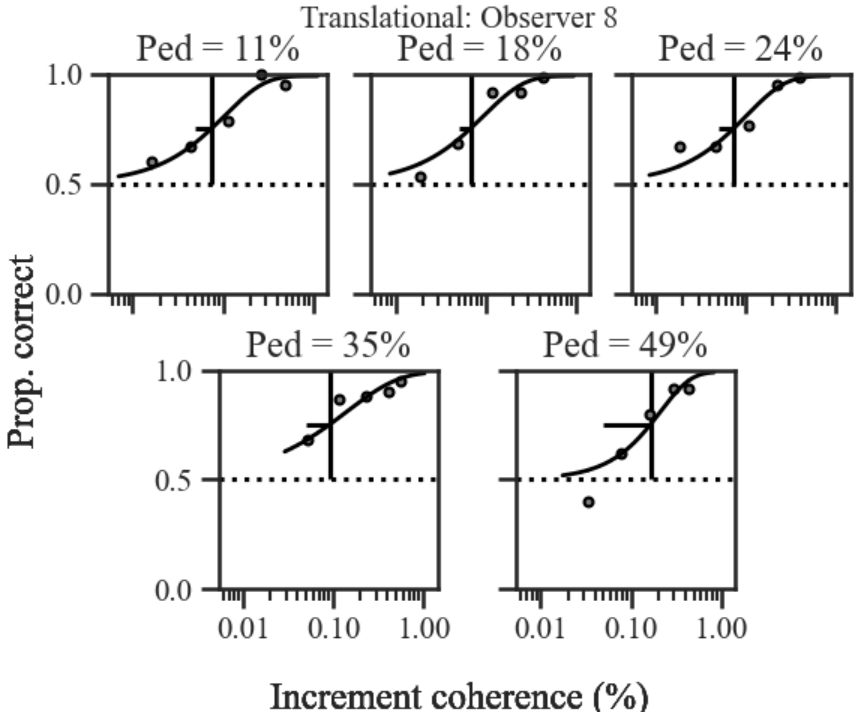
Translational: Observer 8
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

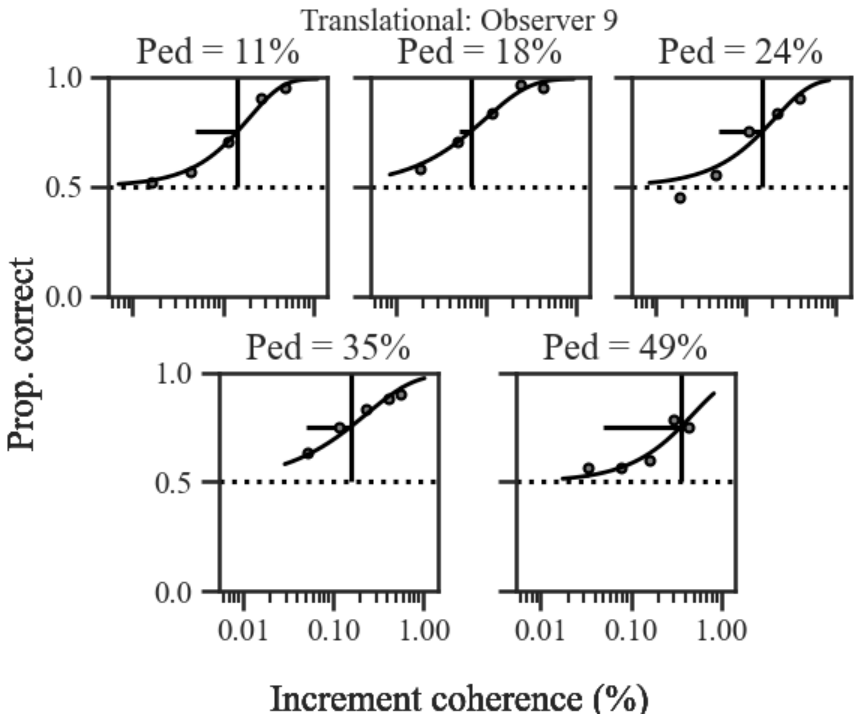
Translational: Observer 9
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
1.0
0.5
0.0
Prop. correct
0.01
0.10
1.00
Increment coherence (%)

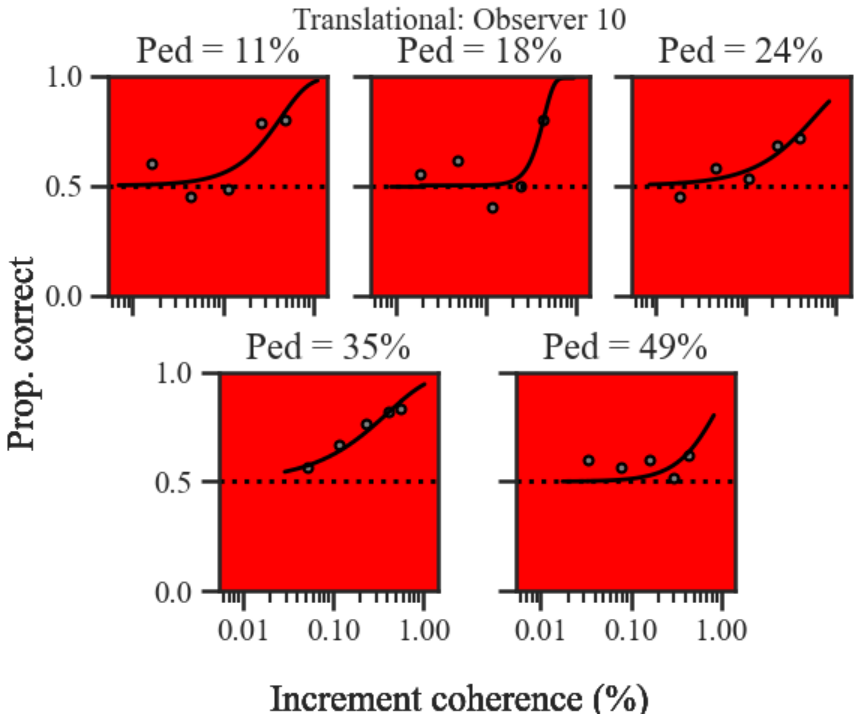
Translational: Observer 10
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
1.0
0.5
0.0
Prop. correct
0.01
0.10
1.00
Increment coherence (%)

0

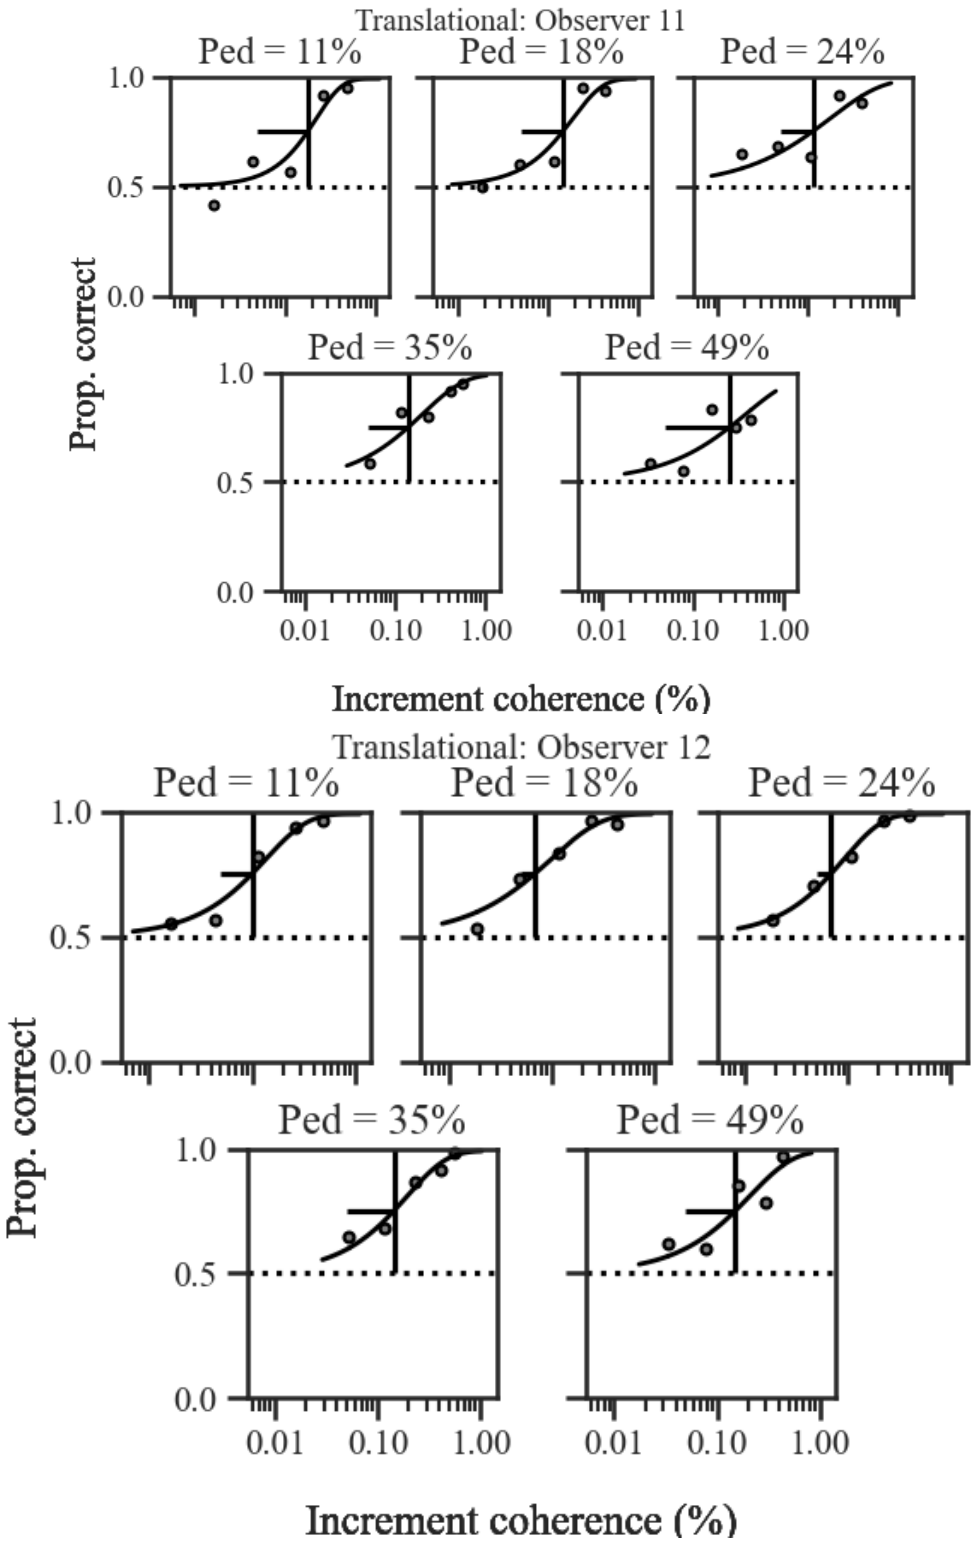
Translational: Observer 11
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
1.0
0.5
0.0
Prop. correct
0.01
0.10
1.00
Increment coherence (%)
Translational: Observer 12
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
1.0
0.5
0.0
Prop. correct
0.01
0.10
1.00
Increment coherence (%)

0

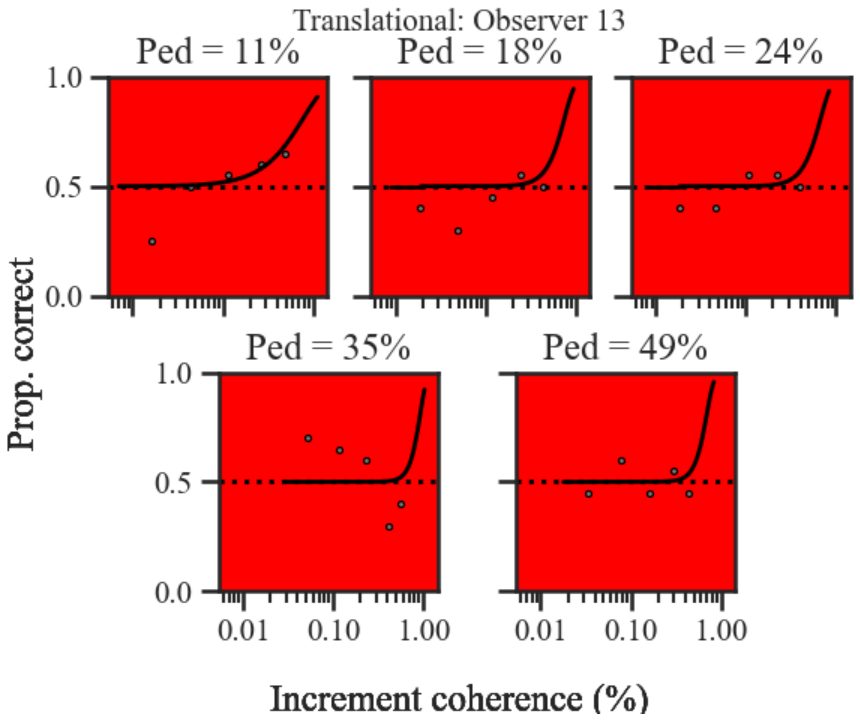
Translational: Observer 13
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

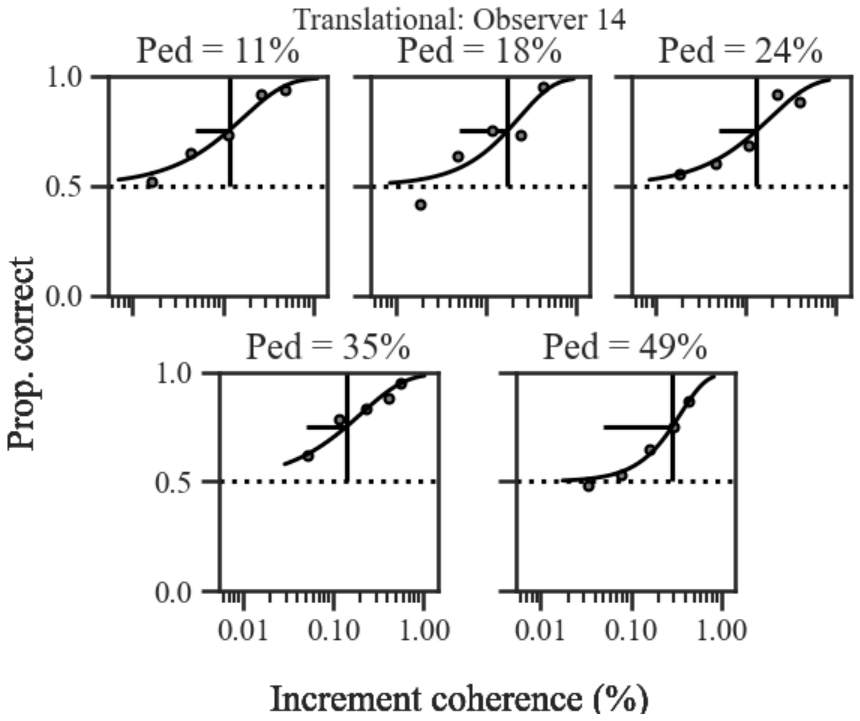
Translational: Observer 14
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

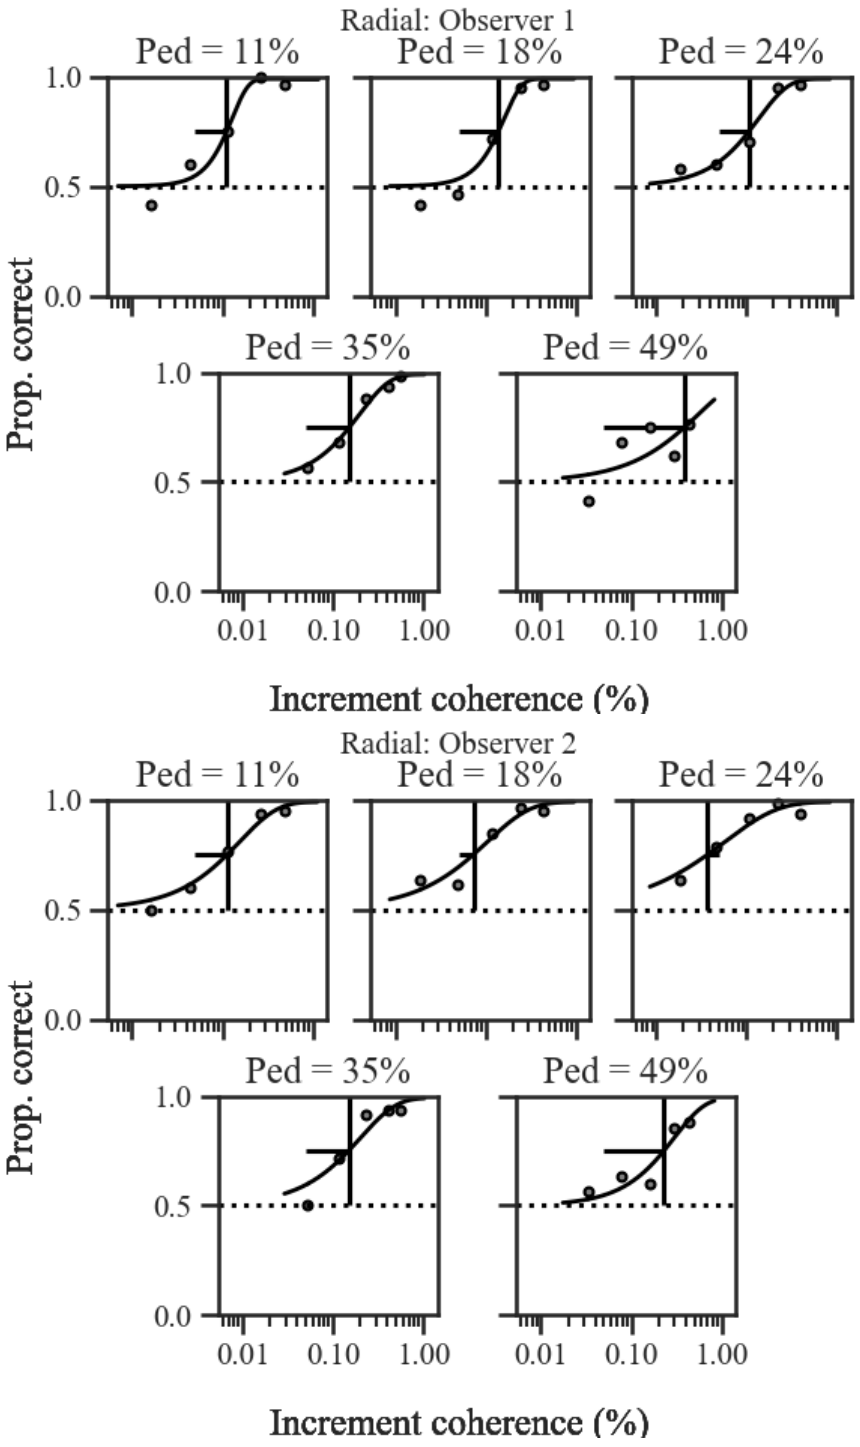
Radial: Observer 1
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)
Radial: Observer 2
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

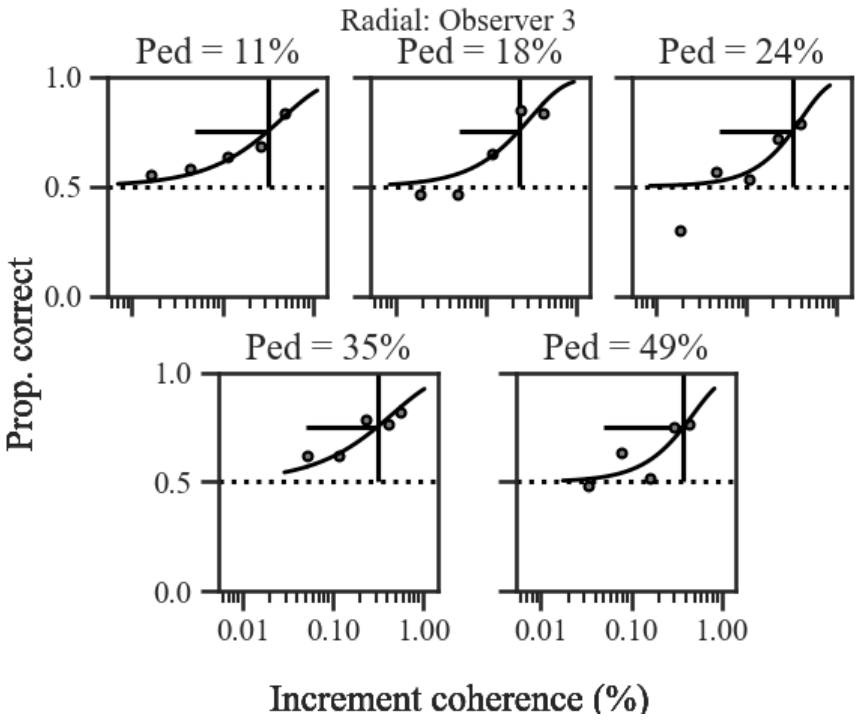
Radial: Observer 3
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

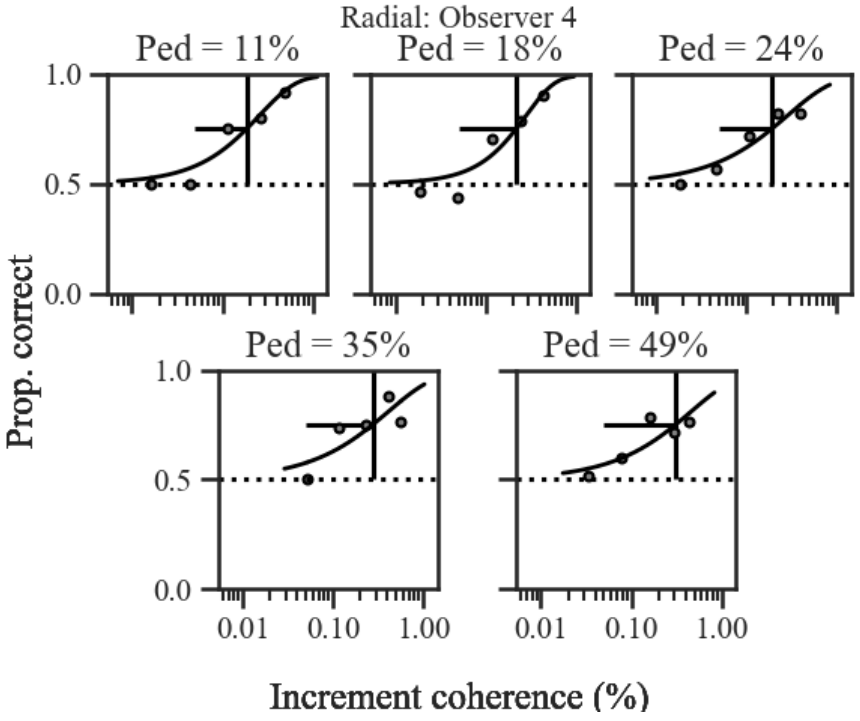
Radial: Observer 4
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

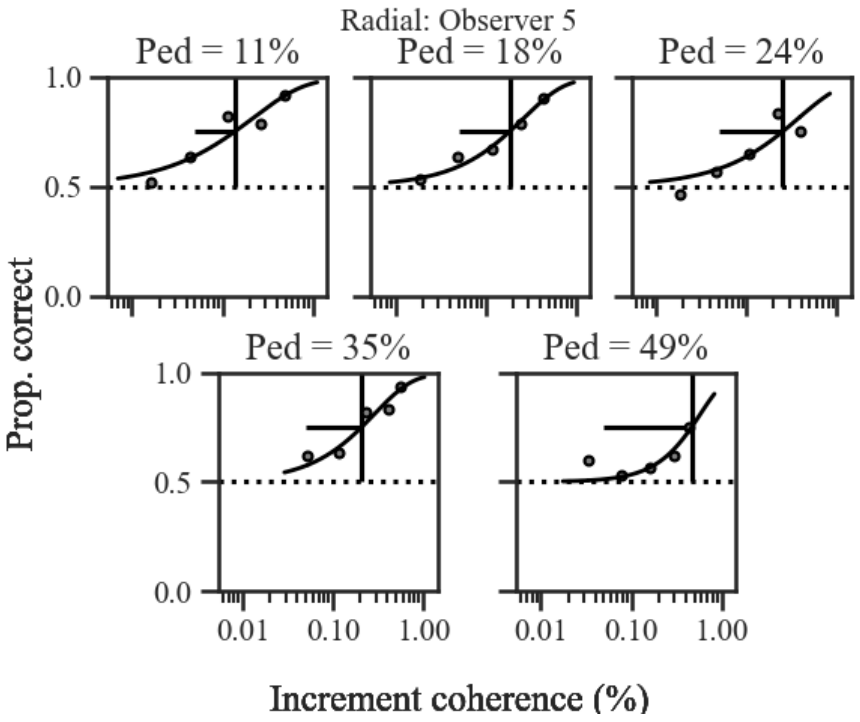
Radial: Observer 5
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

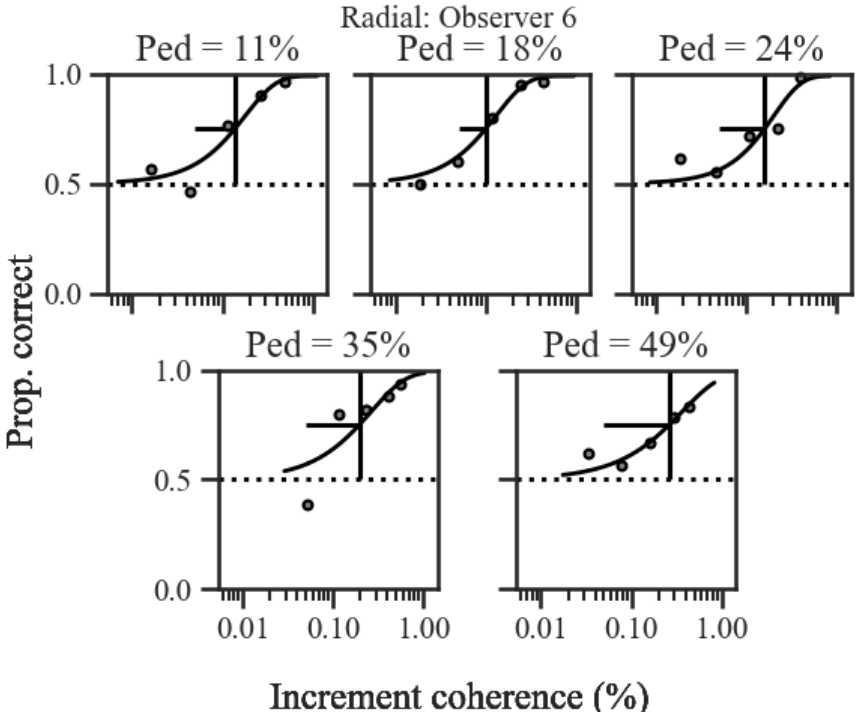
Radial: Observer 6
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

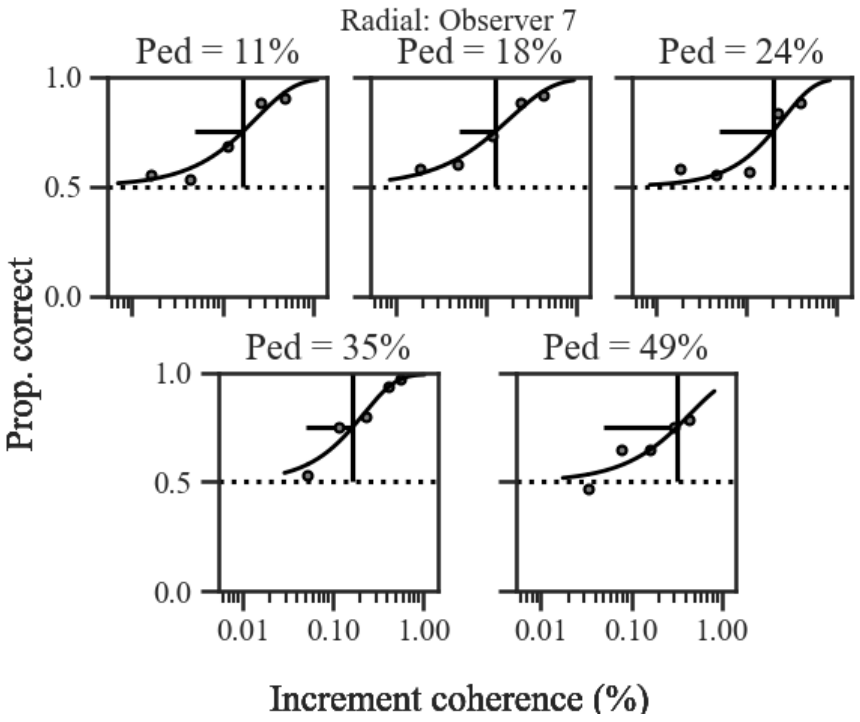
Radial: Observer 7
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

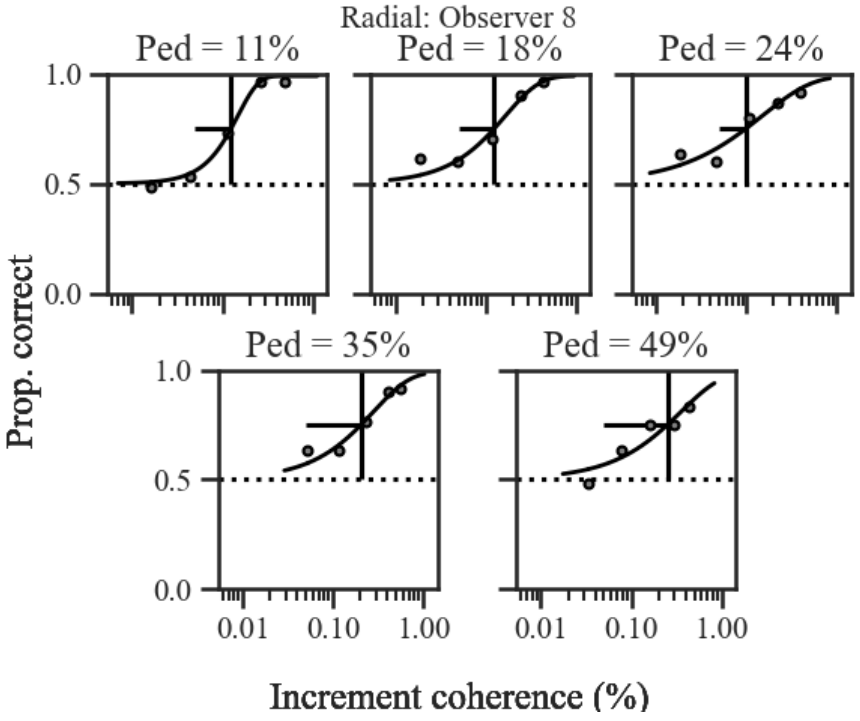
Radial: Observer 8
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

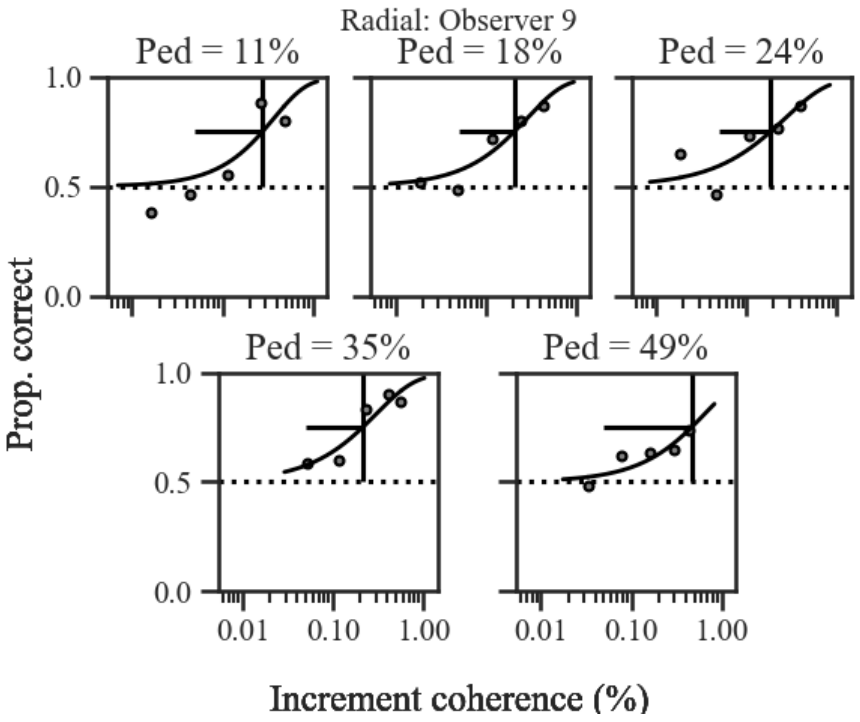
Radial: Observer 9
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

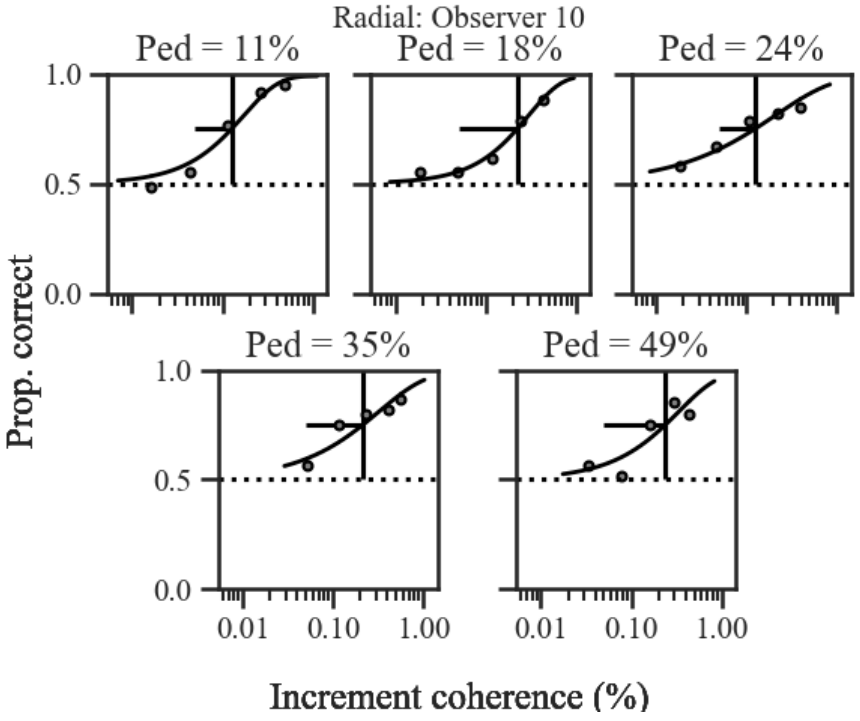
Radial: Observer 10
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

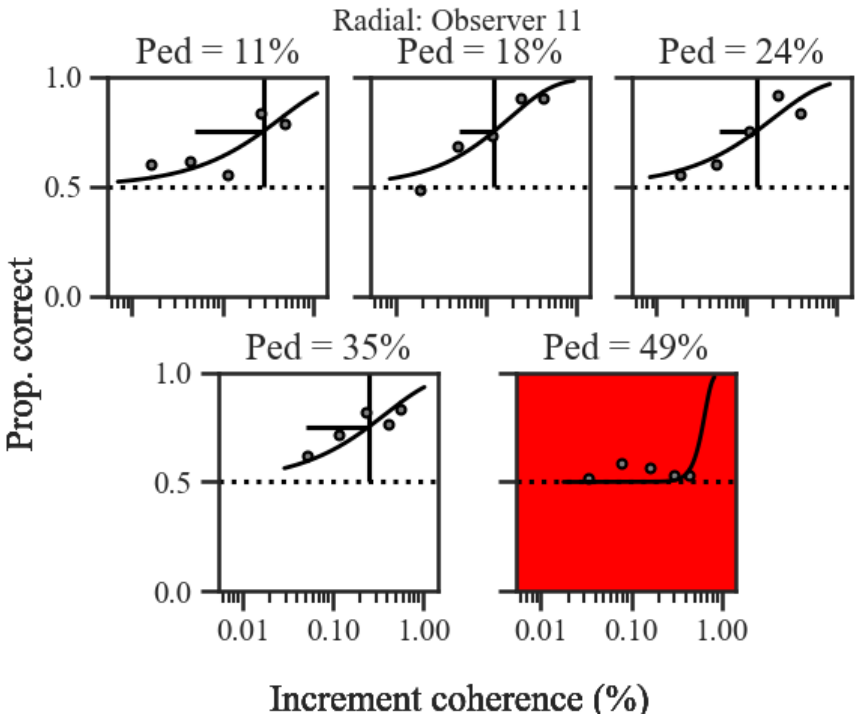
Radial: Observer 11
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

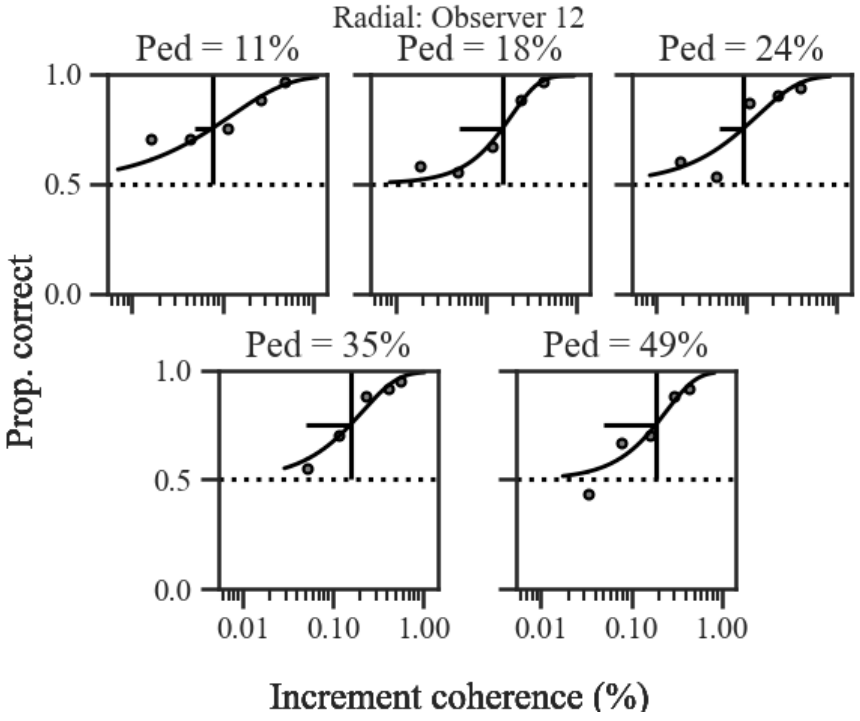
Radial: Observer 12
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

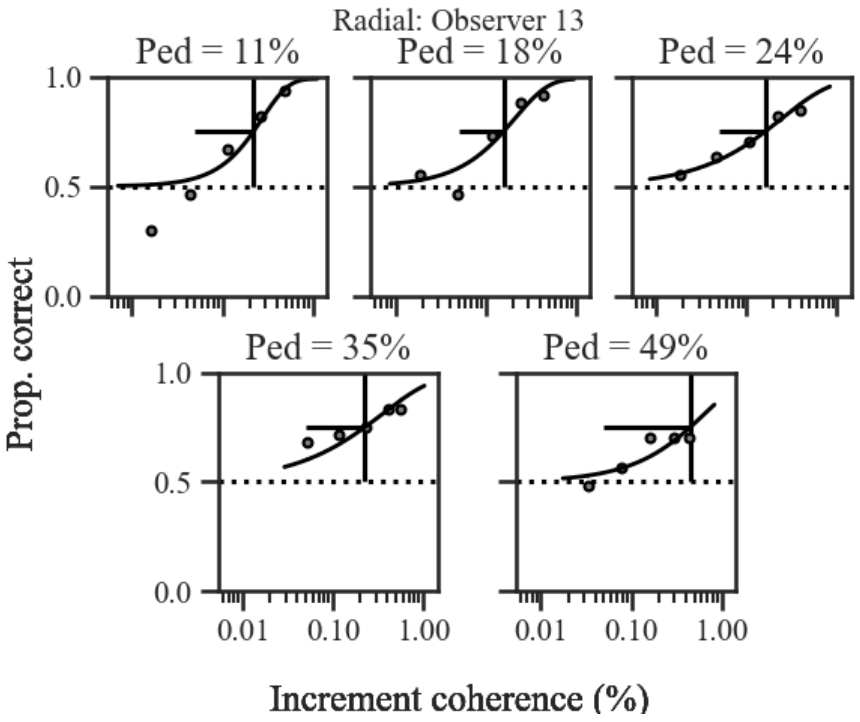
Radial: Observer 13
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

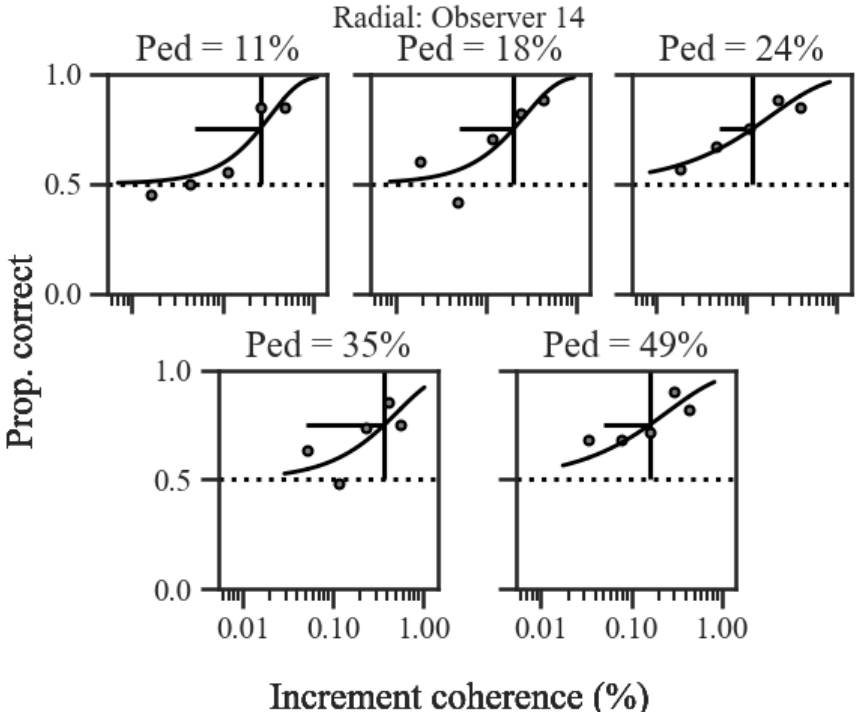
Radial: Observer 14
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

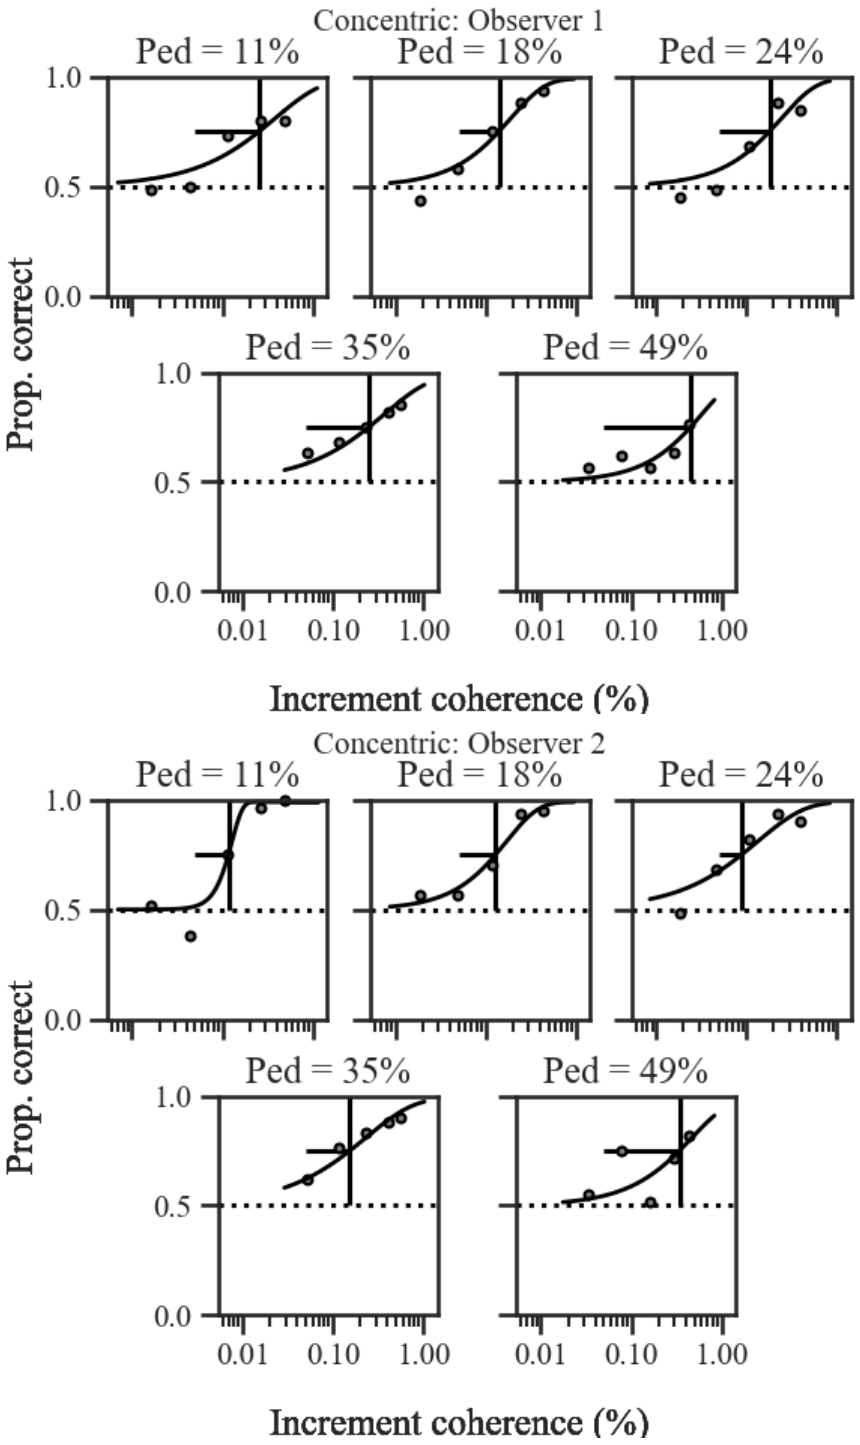
Concentric: Observer 1
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)
Concentric: Observer 2
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

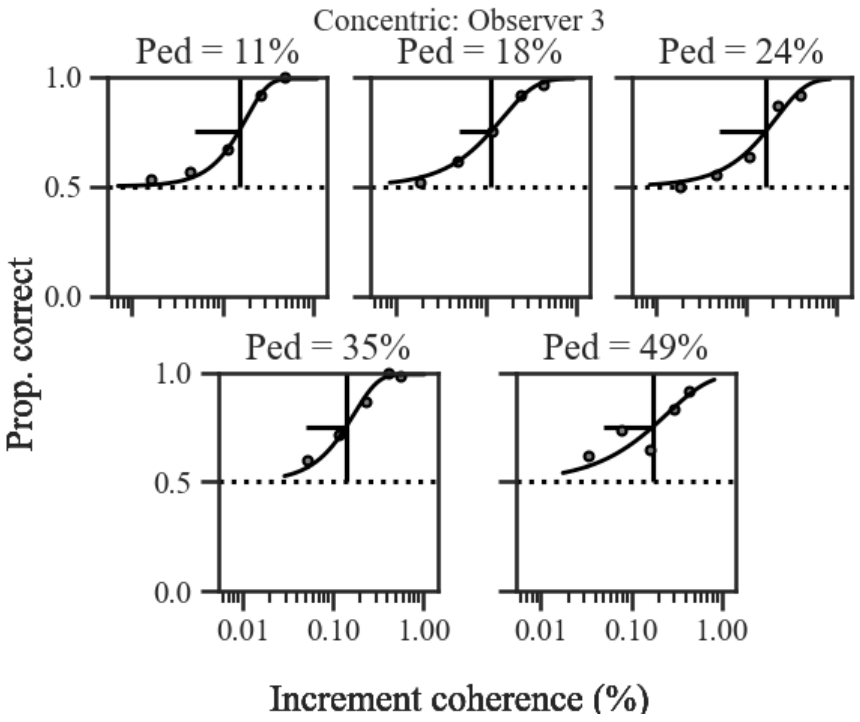
Concentric: Observer 3
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

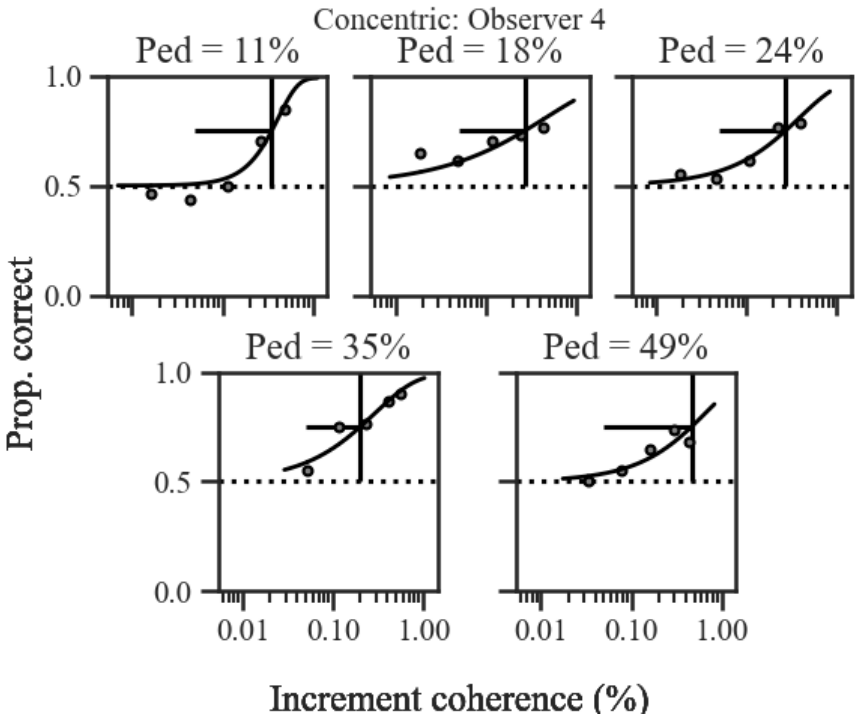
Concentric: Observer 4
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

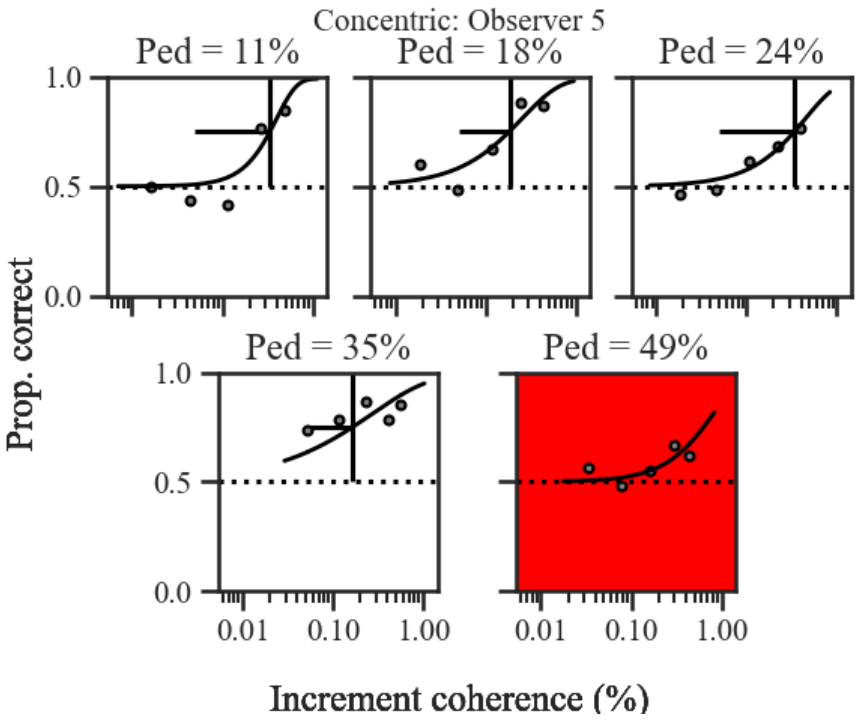
Concentric: Observer 5
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

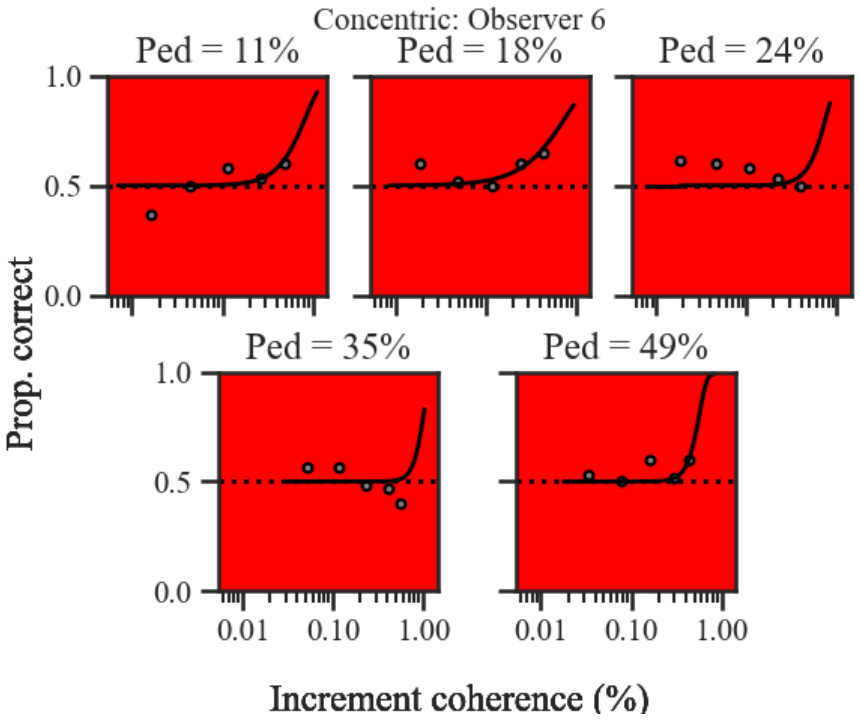
Concentric: Observer 6
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

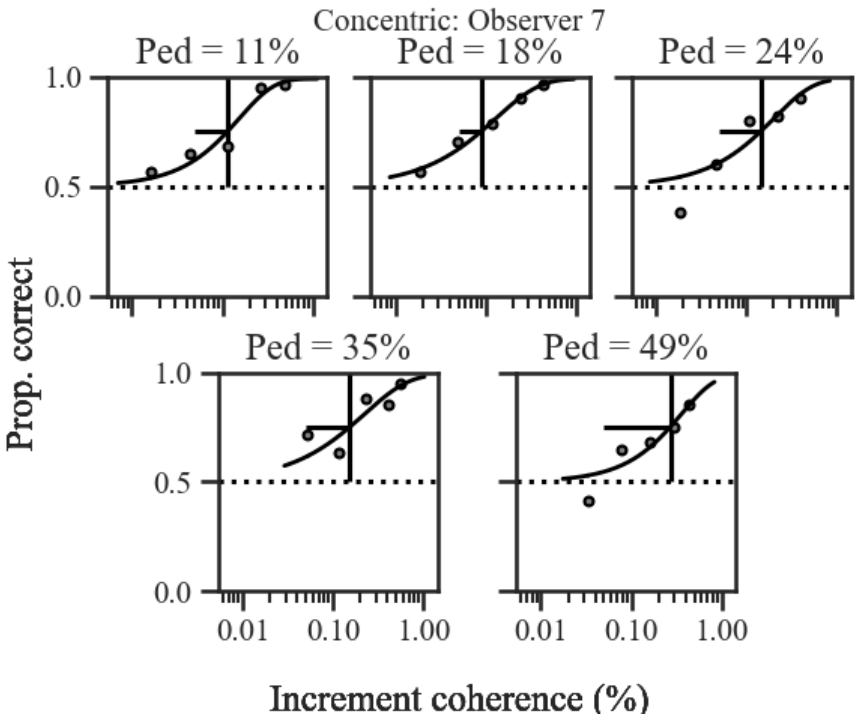
Concentric: Observer 7
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

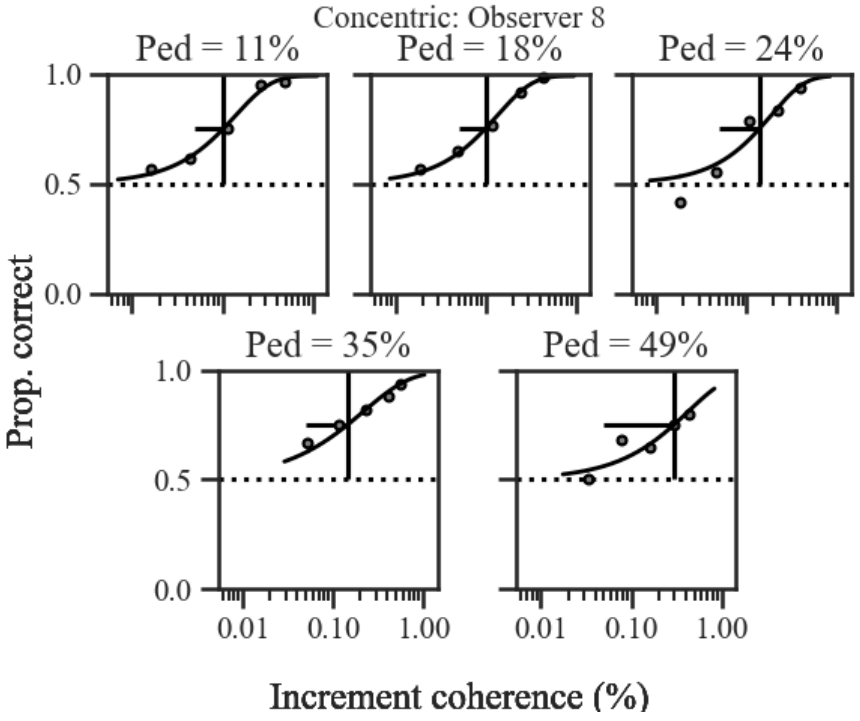
Concentric: Observer 8
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

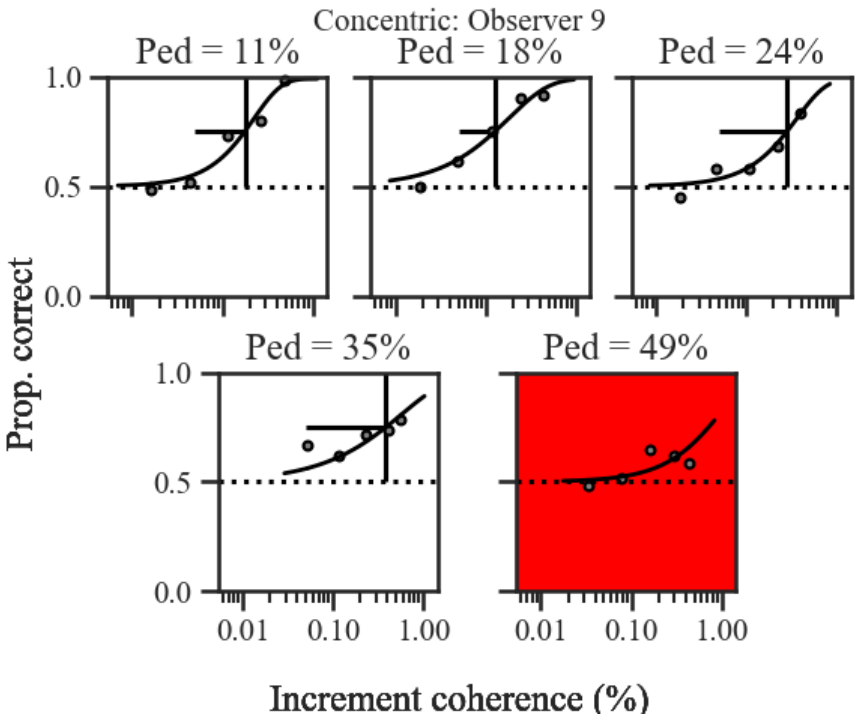
Concentric: Observer 9
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

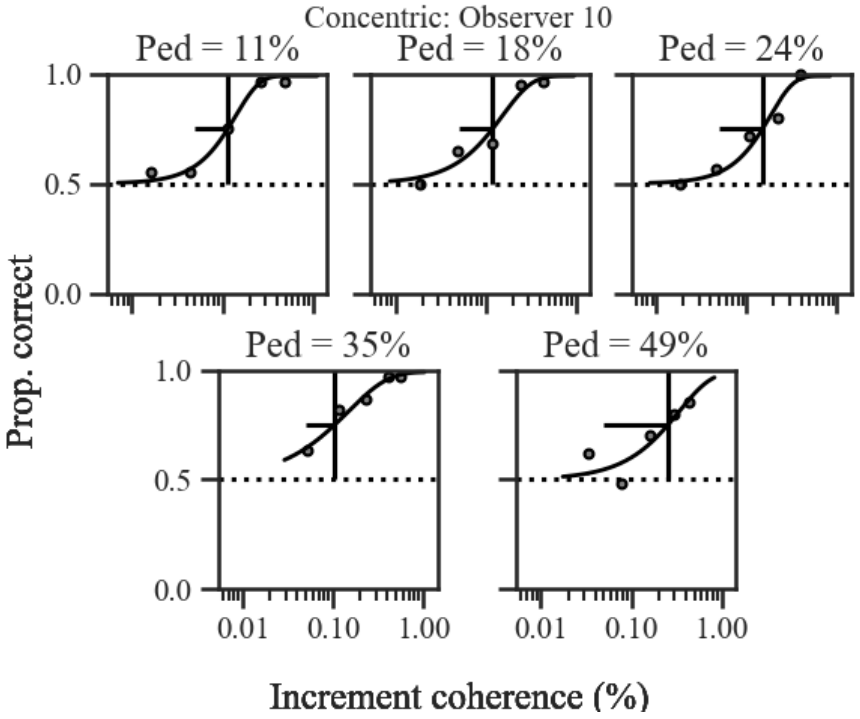
Concentric: Observer 10
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

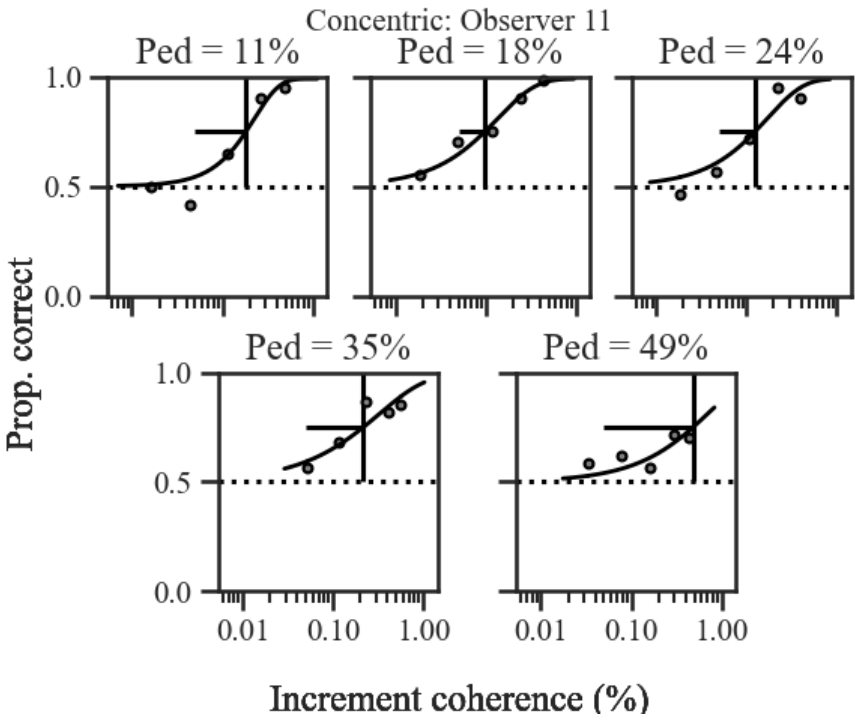
Concentric: Observer 11
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

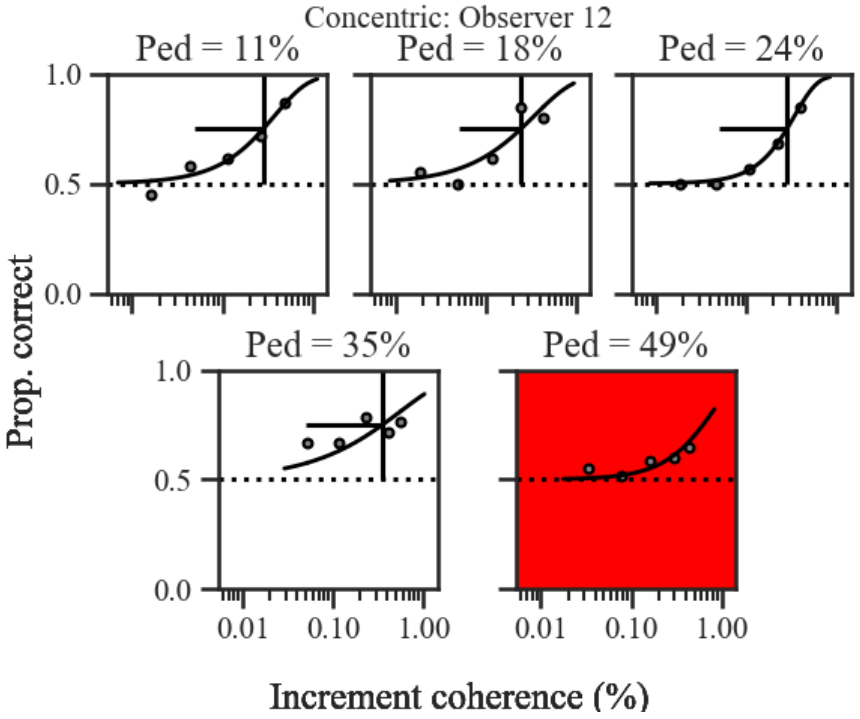
Concentric: Observer 12
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

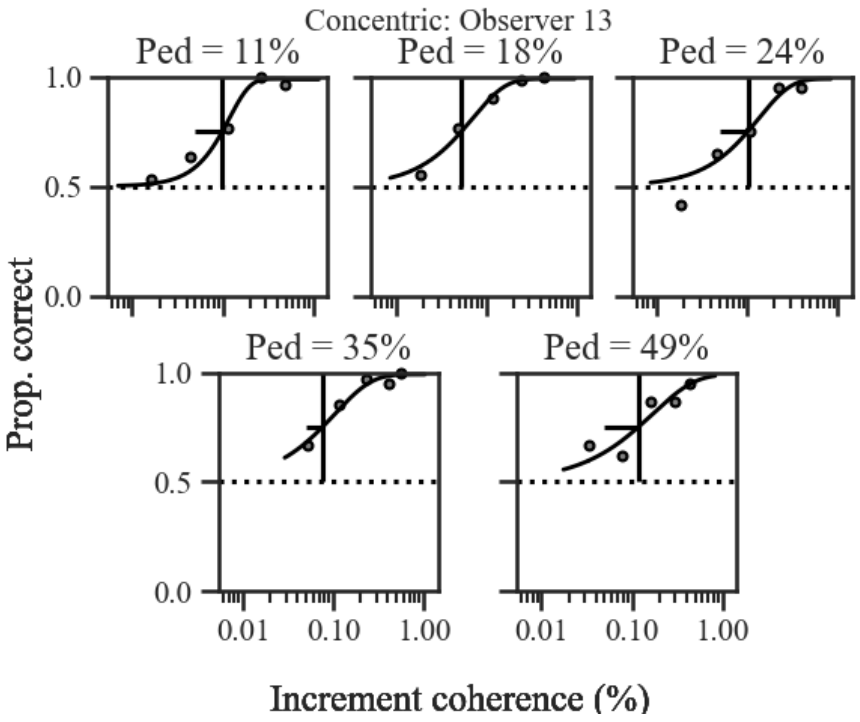
Concentric: Observer 13
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

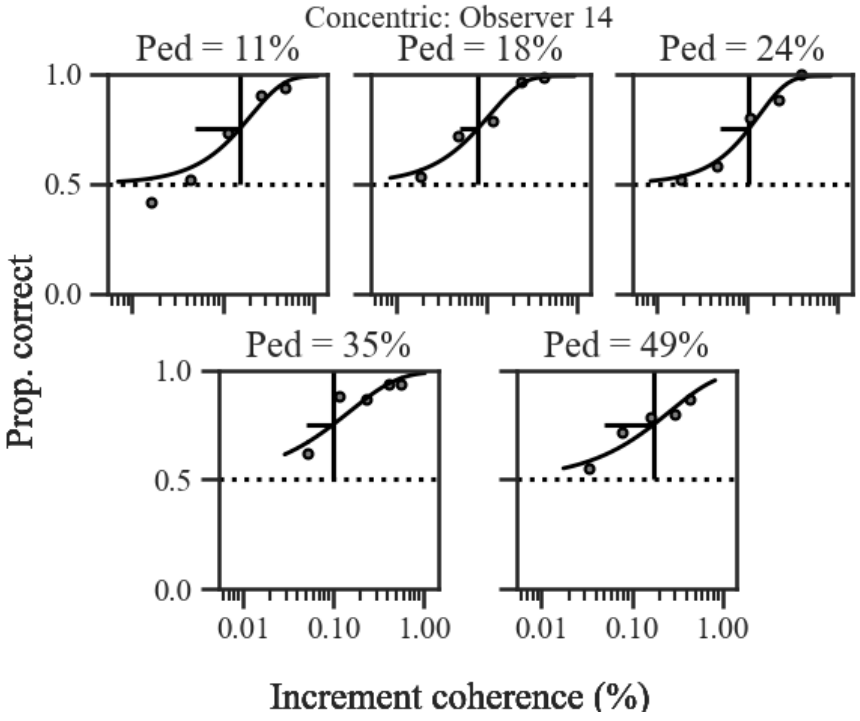
Concentric: Observer 14
Ped = 11%
Ped = 18%
Ped = 24%
Ped = 35%
Ped = 49%
Prop. correct
1.0
0.5
0.0
0.01
0.10
1.00
Increment coherence (%)

0

Appendix D: Psychometric function slopes for global form thresholds

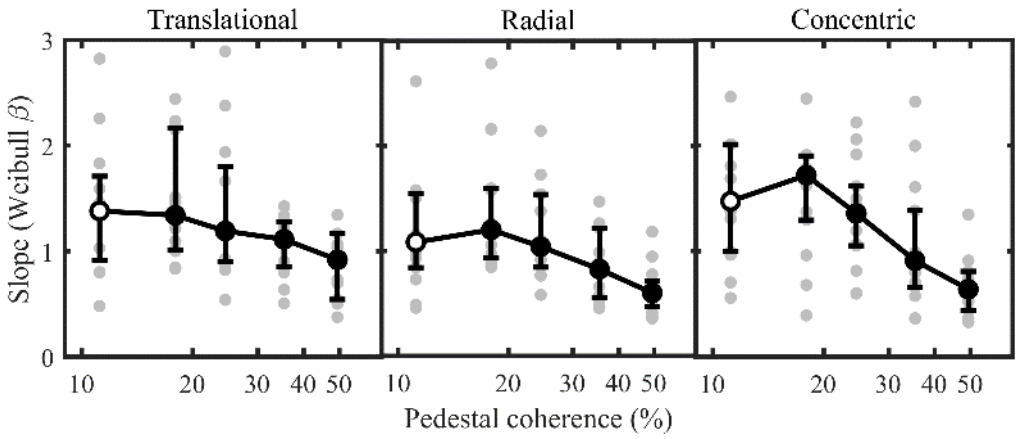

Psychometric function slopes for global form discrimination

For each pattern axis, the fitted value of Weibull-B (slope) is presented as a function of pedestal global form coherence. The open circle represents the median slope at detection threshold, while the solid circles represent the median for discrimination thresholds. Error bars represent 95% BCa confidence intervals from 10000 resamples of the median.

Appendix E: Contrasts of local component SNR across global form coherence levels

For each reliable component recovered for local form processing (RCs 1, 2, and 3) the projected 1F2 SNR has been contrasted between each global form coherence level to reveal any changes in local SNR as a function of global form coherence. Each RC is represented by a separate figure. Within each figure, contrasts are represented as ridgeline kernel-density estimates on the ordinate axis, such that a line for "1-2" represents the contrast of 1F2 SNR for the 1st and 2nd global form intensities. Where the confidence intervals on the difference do not contain zero (indicative of a significant difference), the KDE is coloured red (there were no significant differences, so all KDEs are grey).

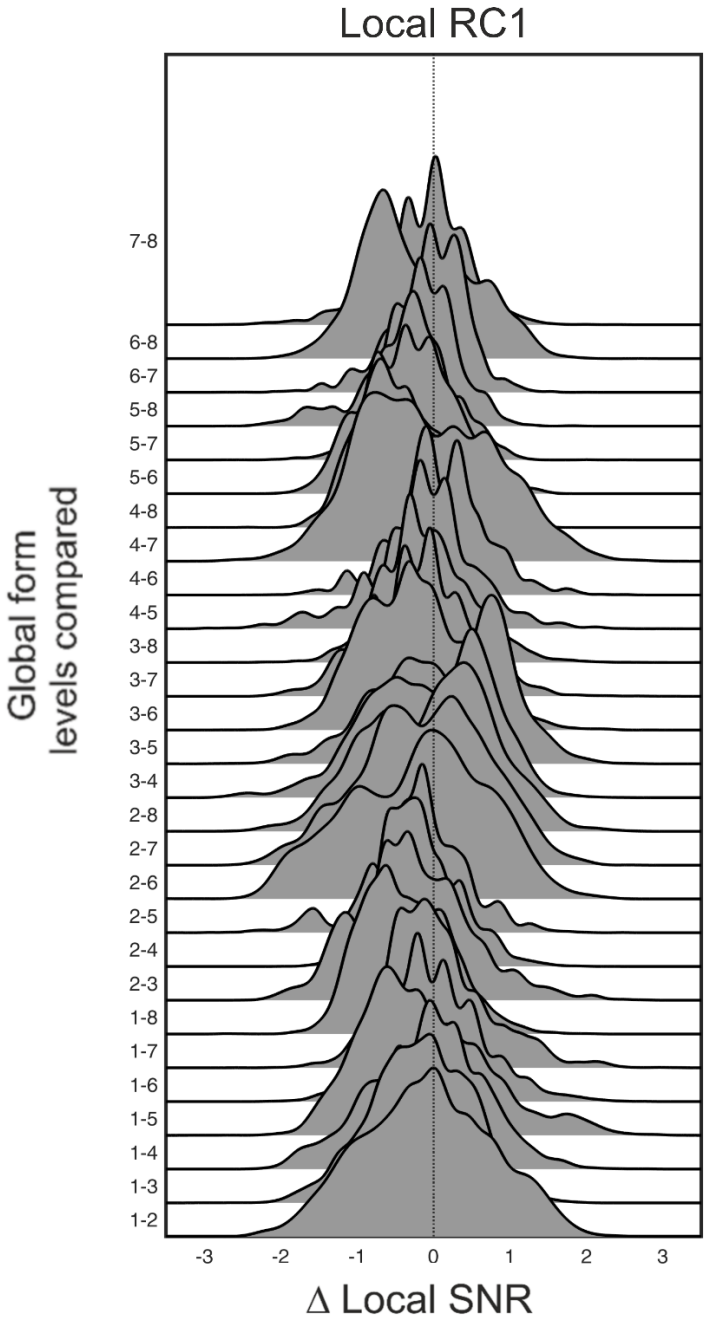

0

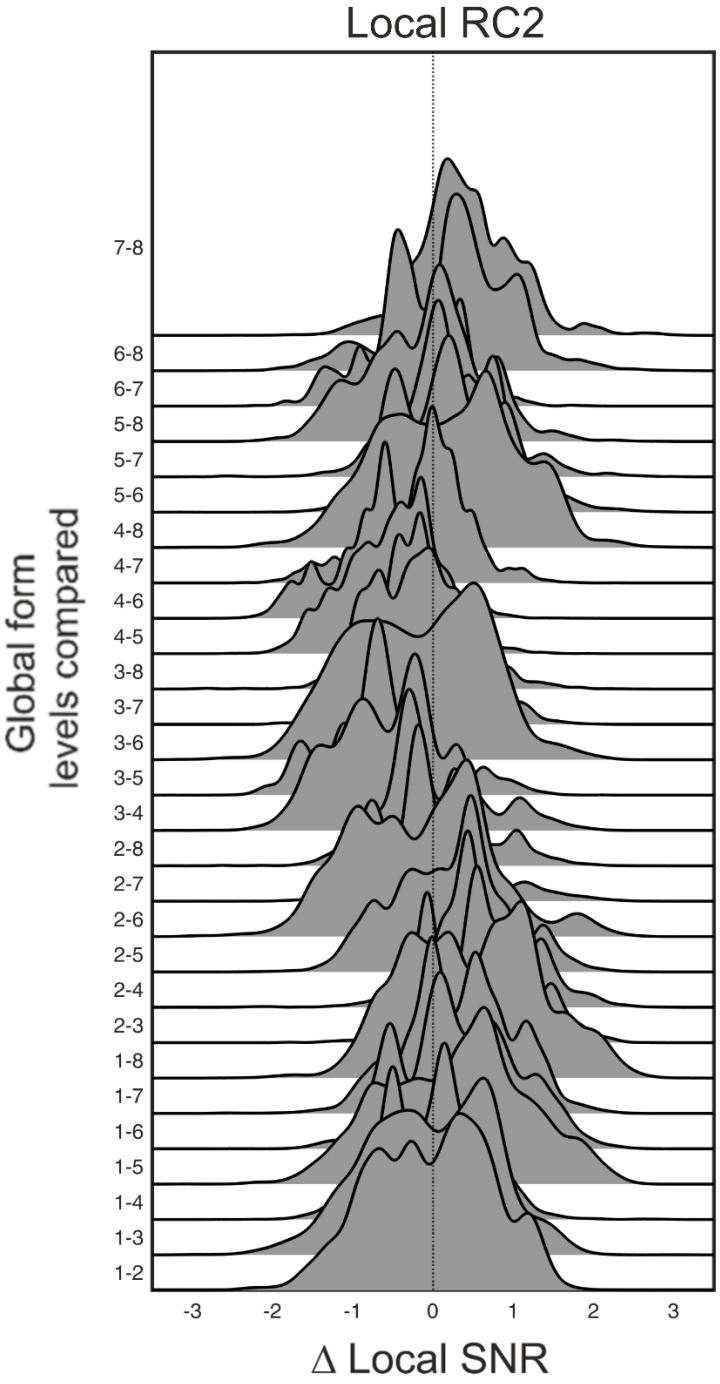
Local RC2
Global form
levels compared
7-8
6-8
6-7
5-8
5-7
5-6
4-8
4-7
4-6
4-5
3-8
3-7
3-6
3-5
3-4
2-8
2-7
2-6
2-5
2-4
2-3
1-8
1-7
1-6
1-5
1-4
1-3
1-2
-3
-2
-1
0
1
2
3
Δ Local SNR

0

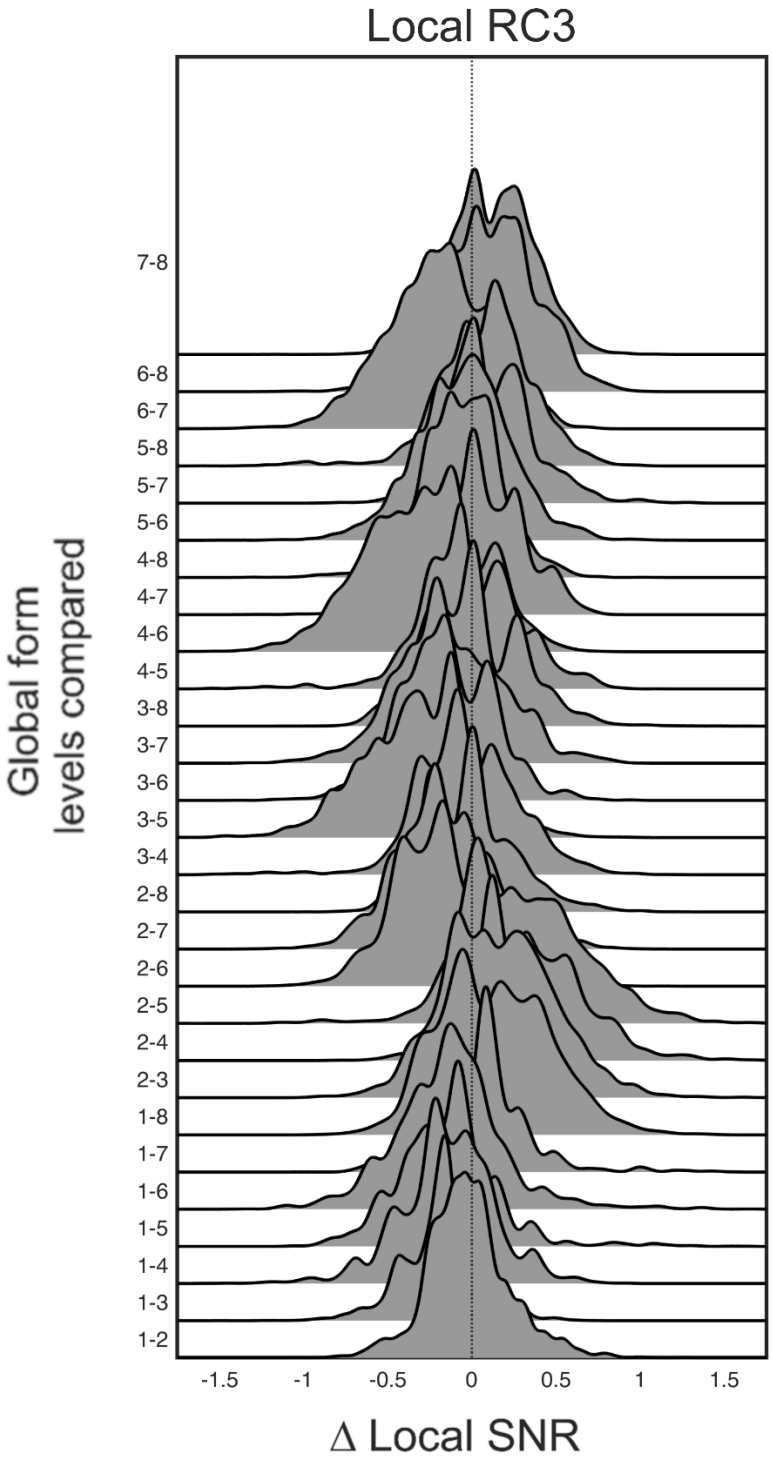
Local RC3
Global form
levels compared
7-8
6-8
6-7
5-8
5-7
5-6
4-8
4-7
4-6
4-5
3-8
3-7
3-6
3-5
3-4
2-8
2-7
2-6
2-5
2-4
2-3
1-8
1-7
1-6
1-5
1-4
1-3
1-2
-1.5
-1
-0.5
0
0.5
1
1.5
Δ Local SNR

Appendix F: SSVEP model fits with individual datapoints

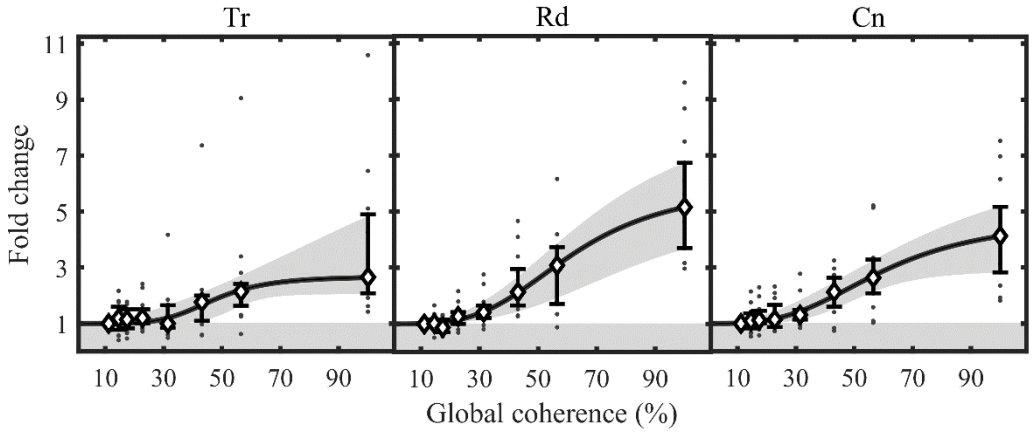

SSVEP model fits with individual datapoints

Median SNR from baseline for data projected through RC1 for each pattern condition. The solid line is the best fit of a hyperbolic ratio function to the median SNR, while the shaded grey region around the line is the 95% CI on this model fit. Tr: translational, Rd: radial, Cn: concentric. Individual datapoints are included here unlike the figure in the main text.