

Measuring nonlinear signal combinations using EEG

Darren Gavin Matthew Cunningham, MRes.

Thesis submitted to the University of Nottingham for the Degree of

Doctor of Philosophy

September 2017

Abstract

This thesis has examined nonlinear signal summation using a combination of EEG and computational modelling. Nonlinearities are essential to many perceptual phenomena, but remain poorly understood beyond the earliest levels of the sensory pathways.

Many nonlinear physiological phenomena, such as cross-orientation suppression (XOS), can be readily described by models of normalisation for neuronal gain control in primary visual cortex (V1). However, there are several nonlinearities that normalisation cannot fully explain. For example, super saturation – which can occur in around 17% of V1 and 25% of V2 neurons in macaque (Peirce, 2007b) – would be considered metabolically wasteful within a framework of normalisation: an over-exertion of the normalisation pool upon the excitatory response of a neuron. It seems unlikely that this non-monotonic nonlinearity does not serve a purpose. Considering this, gain control may not be the only function served by nonlinearities in the visual system (and beyond).

Peirce (2007b, 2011, 2013) proposed that nonlinearities in V1 could also be used by neurons in mid-level vision to detect signal conjunctions for combinations of stimuli. This kind of signal summation would make possible neurons with more complex receptive field preferences than are commonly observed in V1. For example, neurons that are sensitive to multiple orientations and a narrow bandwidth of spatial frequencies would be useful for detecting patterns coherent plaids.

However, at any one point in time, several different nonlinearities can occur in response to a stimulus. Being able to distinguish one from the other is more difficult than it might at first seem. The experiments described throughout this thesis aimed to disentangle nonlinearities, identify those that were selective for specific stimulus combinations and characterise them.

In Chapter 3 we used transient electroencephalography (EEG) to measure the earliest component – C1 – of visual evoked potentials (VEPs) to brief presentations of gratings and their combinations into coherent and non-coherent plaids. By comparing the C1 response to gratings and plaids, we aimed to measure the degree of nonlinear summation taking place for coherent and non-coherent grating combinations. The

outcome was inconclusive; there was limited evidence to suggest the involvement of extra nonlinearities in the processing of coherent plaids that were not involved in processing non-coherent plaids. This might be an inherent problem with the transient EEG approach. Although it produces a rich time course of data following the presentation of a stimulus, the response is the sum of many nonlinearities.

To overcome this, we took an alternative approach in Chapter 4 and used the two-frequency method of steady-state EEG. This allows you to tag each of the grating components forming a plaid, as well as directly measure nonlinearities at intermodulation frequencies. We found a plaid-selective intermodulation response, which was larger for coherent plaids than it was for non-coherent plaids. In support of this representing an additional nonlinearity beyond normalisation, the degree of component suppression did not differ between coherent and non-coherent plaids for any of the grating components used.

We generated a simple two-layered computational model of signal summation to try and explain the complexity of responses generated in to combinations of gratings. The model took the form of a logical AND gate, allowing it to respond selectively to conjunctions of signals. It appears that this kind of mechanism can represent well the responses we observed using EEG.

It is not clear how a mechanism that makes use of saturating nonlinearities to perform selective signal summation would behave across contrast. At lower contrast levels, before many neurons reach the rising slope of their dynamic range, it might be that the mechanism fails altogether. Using a similar paradigm to Chapter 4, we measured intermodulation responses across a wide range of contrast levels in Chapter 5. We again found a selective intermodulation response that was larger for coherent plaids. However, this difference only occurred at the highest component contrast level that we used – 32%.

Having found a nonlinearity in the visual system that appeared to selective for particular combinations of grating stimuli, we wanted to investigate whether similar nonlinearities are put to use in other brain regions. In Chapter 6 we generated auditory stimuli – three pure tones – that were combined to form a consonant and a dissonant chord. Substantial component suppression was observed for one of the

components. However, no intermodulation responses or component-based harmonic responses were observed.

Bringing these findings together, the transient approach to measuring nonlinear responses is somewhat limited, and provided only hints at what might be the presence of 'conjunction' detectors in mid-level vision. On the other hand, it appears that the two-frequency approach is extremely useful for measuring and disentangling multiple nonlinear responses. Here – in the visual system, at least – this was useful for distinguishing responses relating to lateral inhibition caused by the presence of multiple stimulus components from responses relating to the combination of responses relating to those stimulus components in the brain. Conjunction detectors that operate at moderate to high contrast levels appear to be present in mid-level vision. In the one auditory study that we conducted, no clear pattern of results were observed, making it difficult to assess the usefulness of the two-frequency approach in that domain.

Acknowledgements

I would firstly like to thank my supervisor and friend Jon Peirce. The value added to my experience, my attitude, my work ethic and my possibilities by his hand is immeasurable. He has been patient when it mattered, direct when it mattered, but most strikingly for me, I think, was that his infectious and engaging curiosity encouraged the same in me. This thesis and the career I'm now in would not have been possible without him.

Secondly, my fiancé John. You have been a saint by my side throughout this process. I know how much you've given up over the last year to make it work – to keep me grounded and keep us going – and I can't tell you how much I appreciate it. Here's to things to come!

Finally, I'd like to thank my family and the rest of my PhD cohort for all the phone calls, the drinks and keeping me sane for the last four years.

Contents page

Chapter 1: Literature Review	1
The visual system	1
Nonlinearities are prevalent.....	1
Normalisation in V1.....	4
Nonlinear responses beyond V1	6
Detecting nonlinear summation using EEG	14
ERPs: what do they represent?	15
Using ERPs to study nonlinear responses.....	18
Frequency-tagging is useful for separating nonlinearities	19
Using frequency-tagging to measure nonlinear signal summation.....	21
Summary	25
Chapter 2: General Methods.....	26
Recruitment.....	26
EEG Set-up.....	26
Gamma correction	26
Stimulus configuration and experimental procedure	27
Maintaining visual attention: the green dot task	28
Making plaids	28
Choosing spatial frequencies.....	30
Analysing SSVEP data	31
Chapter 3: Visual evoked responses to coherent and non-coherent plaid stimuli	33
Rationale	33
Methods	34
Results	39
C1 Responses.....	39
P1 Responses.....	41
Discussion.....	43
Amplitude effects	43
LI scores all less than linear	43
LI scores only different for AA, not BB, compared to AB	45
Chapter 4: Measuring nonlinear signal combinations using intermodulation responses	46
Rationale	46
Methods	47
Results	53
General overview	53
Response to components.....	55
Response to plaids.....	57
Modelling.	64
Discussion.....	67
Chapter 5: Measuring the contrast response function of intermodulation responses	71
Rationale	71
Methods	72
Results	74

Response to components.....	74
Response to plaids.....	80
Discussion.....	83
Chapter 6: Measuring intermodulation responses in the auditory domain.....	86
Rationale	86
Methods.....	87
Results	91
General overview	91
Response to tone components	92
Response to chords	94
Discussion.....	94
Chapter 7: General Discussion	97
Summary of findings	97
Implications.....	99
Relationship to the previous literature	100
XOS complexity.....	100
Intermodulation responses	102
Conclusion	105
References.....	106
Appendix one – supplementary figures from Chapter 3.....	1
Appendix two – supplementary figures from Chapter 5.....	16

Chapter 1: Literature Review

The visual system

Nonlinearities are prevalent

The visual system can be thought of as a functional hierarchy. It is composed of distinct and increasingly complex neural populations, each combining signals from the previous region to form more complex representations. Receptive field (RF) size and complexity increases up through the cortical regions of the visual stream (Serre, Oliva, & Poggio, 2007). In primary visual cortex (V1), neurons have elongated RFs that can be derived by the summation and subtraction of geniculo-cortical projections (Figure 1.1) (Hubel & Wiesel, 1959, 1962, 1963). There, many neurons respond selectively to stimulus features like size (Greenlee & Magnussen, 1988; Xing, Shapely, Hawken, & Ringach, 2005), orientation (Blakemore & Campbell, 1969; Campbell & Maffei, 1970; Hubel & Wiesel, 1962) and spatial frequency (Albrecht & DeValois, 1981; Campbell, Cooper, & Enroth-Cugell, 1969; Campbell & Robson, 1968; Greenlee & Magnussen, 1988; Maffei & Fiorentini, 1973), while neurons in higher visual areas respond selectively to complex stimuli such as faces (Kanwisher, McDermott, & Chun, 1997; McCarthy, Puce, Gore, & Allison, 1997).

Few neurons within this hierarchy respond linearly across contrast; they display a limited dynamic range of contrasts over which the majority of their response occurs followed by saturation (Albrecht & Hamilton, 1982; Peirce, 2007b). The beginning of this process involves transforming the external visual scene into the retinal image; a representation of the external visual scene on the retina. It is created from the transduction process that occurs when light from the external environment is projected through the lens of the eye and onto the mosaic of rods and cones across the retina, known as photoreceptors. The transformation from light energy to neural electrical impulses results in projections to bipolar cells and retinal ganglion cells. Opponent arrangements in retinal ganglion cells allow them to detect local contrast intensities in the visual scene. For example, centre-surround classical receptive fields

(CRF) – the limited range of visual space within which a neuron can detect changes in luminance (Enroth-Cugell & Robson, 1966; Kuffler, 1953).

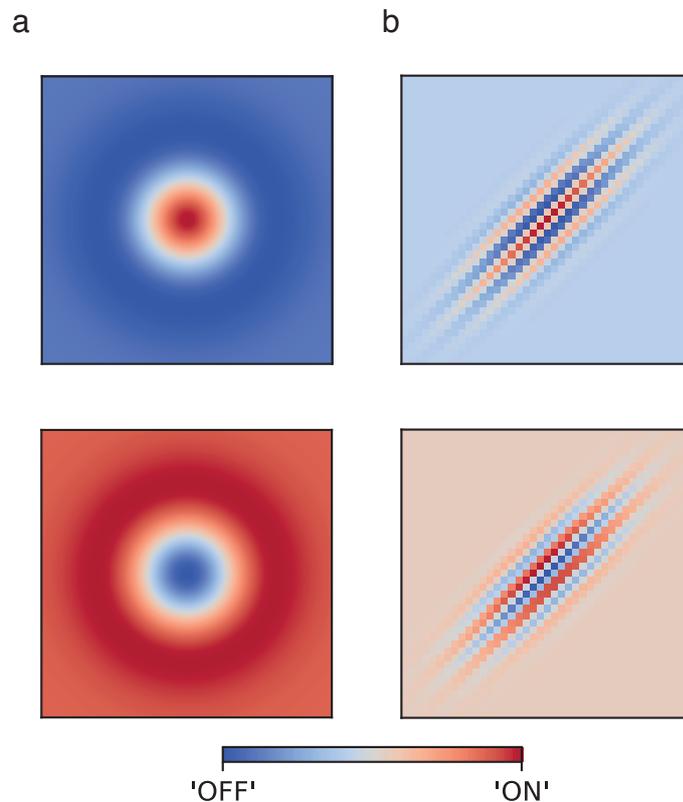


Figure 1.1 CRF structures in a) centre-surround LGN neurons and b) V1 simple cells, based on Hubel and Wiesel (1959, 1962). Red regions represent 'ON' regions of the receptive field, where high luminance regions of a stimulus provides optimal stimulation. Blue areas represent 'OFF' regions of the receptive field, where low luminance regions of stimulus provides optimal stimulation. The LGN CRFs shown here can be described as the difference of two Gaussian functions, while the V1 simple cells can be described as the product of a sinusoidal function and a Gaussian function. A diverging blue-white-red colormap was used instead of a greyscale colormap to better highlight the spatial organisation of the receptive fields in 2D space.

RFs allow retinal ganglion, LGN and V1 neurons to act as spatiotemporal filters, removing enough redundant information from the retinal image to encode the spatial distribution of luminance and contrast energy in the visual scene (Lennie, 2003). Basic 'linear' models these receptive fields involve projections first passing through a linear spatial filter, where the local intensities of a retinal image are multiplied by the local weights of the receptive field and summed together. Being linear, the filter can be used to describe the selectivity of a neuron; images that are similar to the filter will produce larger responses than those that are not. For example, localised areas of high

luminance intensity will optimally excite 'ON' regions of a retinal ganglion cell's centre-surround RF, while 'OFF' regions will be optimally excited by localised areas of low (or no) luminance intensity. The centre-surround spatial organisation LGN neuron filters selectivity (Kaplan, Marcus, & So, 1979; Shapely & Lennie, 1985) and the elongated filters of V1 simple cells (Movshon, Thompson, & Tolhurst, 1978) are able to describe their spatial frequency, while their temporal weighting functions predict the temporal frequency selectivity of LGN neurons (Benardete & Kaplan, 1999; Kremers, Weiss, & Zrenner, 1997) to arbitrary stimuli, such as white noise. However, when it comes to generating firing rate responses (impulse spikes) from the linearly summed signal, nonlinearities are required. Receptive fields make use of several static nonlinearities at their output, such as half-wave rectification (firing rate can never be negative) and saturation (the response ceiling of a neuron due to refractory periods) (Kuffler, Fitzhugh, & Barlow, 1957; Passaglia, Enroth-cugell, & Troy, 2001; Troy & Robson, 1992).

Potentially, this saturation is a simple biophysical artefact that neurons are unable to respond at increasingly high rates (e.g. due to a lack of resources). This seems unlikely due to the relative heterogeneity of neural responsiveness, which varies from neuron to neuron, and even within a given neuron responsiveness can change. It appears that this nonlinear behaviour results from an active process of suppression. Two major nonlinearities that demonstrate this are luminance gain control ('light adaptation') and contrast gain control. Gain, or physiological sensitivity, is defined as the magnitude of a stimulus that is required to generate a response. Luminance gain control starts in retinal ganglion cells and places the accelerating dynamic range of neurons to match the recent average light intensity of the visual scene. In areas where light intensity is high, light adaptation reduces the gain, but increases it in areas where light intensity is low. By holding the mean luminance across the retinal image relatively constant, luminance gain control allows the retina to adjust to local variations in contrast. Contrast gain control also begins in the retina (Baccus & Meister, 2002; Victor, 1987). With mean luminance adjusted for, contrast gain control adjusts response gain using local variations in luminance (i.e. RMS contrast), such that gain is reduced where contrast is high and increased where contrast is low. This is then strengthened as it is projected to the LGN and cortex

(Kaplan, Purpura, & Shapley, 1987; Sclar, Maunsell, & Libnie, 1990). Mechanisms of gain control, collectively referred to as ‘normalisation’ (Carandini & Heeger, 1994, 2012; Carandini, Heeger, & Movshon, 1997), allows neurons to maintain optimal response sensitivity across contrast. It also reduces redundancy in neural populations, for example by reducing statistical dependence between neighbouring neurons with similar tuning properties (Albrecht & Geisler, 1991; Albrecht & Hamilton, 1982; Bonin, Mante, & Carandini, 2005; Heeger, 1992; Schwartz & Simoncelli, 2001; Shou, Li, Zhou, & Hu, 1996).

Normalisation in V1

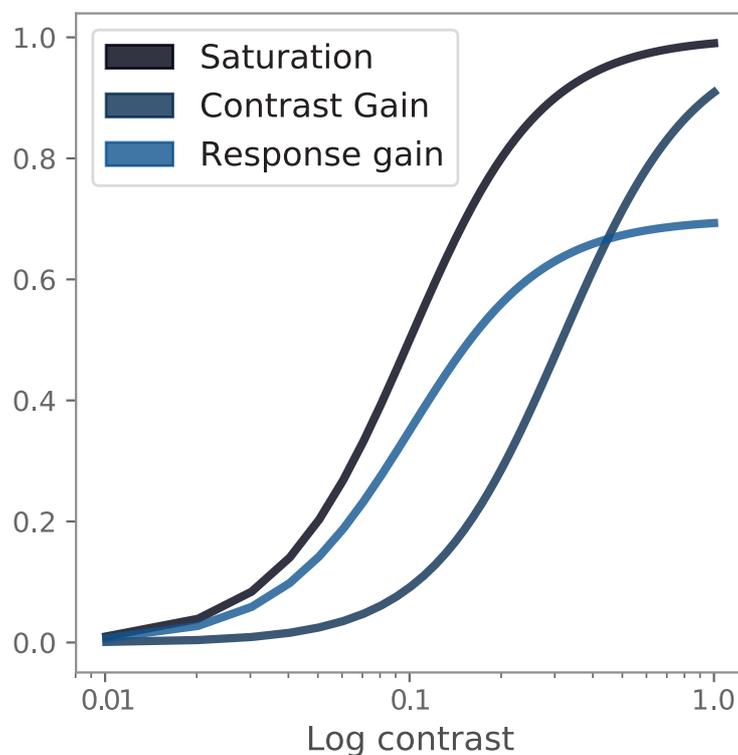


Figure 1.2 Contrast response functions for a linear response and nonlinear responses. The black function represents the output of a saturating nonlinearity. The dark blue function represents response gain control (created by reducing RMax to 0.7) and the lighter blue function contrast gain control (created by adding a masker at a standing contrast of 0.3). Compared to a linear response, gain control and normalisation processes extend the neuron’s dynamic range at low-to-mid level contrasts.

Take cross-orientation suppression (XOS) as an example. A grating whose orientation falls outside of the tuning bandwidth of a neuron will generate very little (if any) response when presented alone. When spatially superimposed on top of the

preferred grating however, it typically results in a reduction in response to the preferred grating. A divisive process of lateral inhibition between neurons explains this well. For example, DeAngelis, Freeman and Ohzawa (1994) and Petrov, Carandini and McKee (2005) found that XOS was broadly tuned for spatial frequency (up to a factor of ~ 4 difference between target and mask SF: Petrov et al., 2005). As a non-preferred grating generates very little response, the neuron might instead receive inhibitive input from a wide pool of other cortical neurons with different spatial frequency preferences, resulting in the observed broadly-tuned response inhibition. Several recent studies have provided evidence that this process may occur in two stages; before and after binocular summation takes place, the latter taking place in cortex (e.g. D. H. Baker, Meese, & Summers, 2007; Cass, Stuit, Bex, & Alais, 2009; Meese & Holmes, 2007, 2010; Petrov et al., 2005; Viswanathan, Jayakumar, & Vidyasagar, 2011).

Normalisation is described well by the hyperbolic Naka-Rushton function:

$$resp = RMax \left(\frac{C_a^n}{C_{50}^n + C_a^n + C_b^n} \right)$$

where $RMax$ is the maximum response rate of the neuron, C_a^n the contrast intensity of a stimulus preferred by the neuron, C_b^n the contrast intensity of a masker, for example, C_{50} the 'semi-saturation point' (the contrast required to produce half of the neuron's maximum response), and n an exponent that determines the slope of the response function (Albrecht & Hamilton, 1982; Naka & Rushton, 1966; Peirce, 2007b, 2011). The numerator in the equation determines neural excitation, while the denominator determines neural inhibition.

Without considering C_b^n , this first equation represents contrast saturation (Figure 1.2). As contrast intensity increases, the neuron's responsiveness reaches the accelerating dynamic range of its function, provided that exponent $n > 1$. Acting as a counter-weight to excitation, contrast intensity proportionally increases in the denominator. Following the semi-saturation point of the response function, this balance between numerator and denominator results in a compression of response

function. When a masker is also present at a standing contrast, the C_b^n term in the denominator results in a lateral shift in the response function (Figure 1.2).

Though many neurons display nonlinearities that are useful for adjusting the gain of neurons, there are other nonlinearities that are difficult to explain in terms of gain control. For example, some neurons display supersaturation – a contrast response function with a characteristic inverted u-function (Peirce, 2007b). This type of activity is redundant within a framework of normalisation; an exaggerated suppressive input from the normalisation pool. However, so much attention has been placed on normalisation and gain control in V1 over the last 40-50 years that a great deal has yet to be understood about how the rest of the visual system operates.

Nonlinear responses beyond V1

Further combinations beyond V1 necessarily contribute to our visual experience. Multiple extra-striate regions respond to stimuli more complex than oriented Gabors or gratings; stimuli like contours, plaids, shapes and faces (Anzai, Peng, & Van Essen, 2007; Desimone, 1991; Desimone, Albright, Gross, & Bruce, 1984; Gallant, Connor, Rakshit, Lewis, & Van Essen, 1996; Pasupathy & Connor, 1999, 2002; Willmore, Prenger, & Gallant, 2010). It seems unlikely that processes such as supersaturation are redundant. Instead, nonlinearities might serve to do more than adjust response gain. For example, supersaturation may instead reflect mechanisms with very selective tuning properties (Peirce, 2007b, 2011). It's of clear importance to understand whether this is the case, and what nonlinearities might be useful for encoding stimuli more complex than simple gratings (for a review, see Peirce, 2015).

The adaptation approach presents an opportunity for targeting specific mechanisms carrying out these combinations. As mentioned earlier, neurons are able to adjust their response gain to a particular stimulus given the average luminance (light adaptation in the retina) and local variations in luminance (contrast adaptation) across the surface of a reflecting stimulus. The degree to which this happens can be manipulated by prolonged exposure to a stimulus. Prolonged exposure to a high contrast grating, for example, results in stronger adaptation and in turn a reduced response to test stimuli at lower contrast levels that share the same features, such as

spatial frequency or orientation. By moving the neuron's dynamic range to higher contrasts, this active process increases the neuron's sensitivity to higher-contrast stimuli (Blakemore & Campbell, 1969; Greenlee & Magnussen, 1988; Manahilov & Vassilev, 1986; Ohzawa, Sclar, & Freeman, 1985; Sclar, Lennie, & Depriest, 1989; Shou et al., 1996; Snowden & Hammett, 1992). Therefore, by presenting an adapter stimulus with certain features, such as spatial frequency and orientation, you can target the responsiveness of neurons tuned to those features. Consider the following simple experiment. You determine a participant's contrast threshold – the contrast required to detect a stimulus – for a sinusoidal grating of a certain spatial frequency. You then flicker a high contrast grating identical to the original for an adaptation period. Following a brief interval, a target grating of the same spatial frequency as the initial grating is presented in the same position. If the target grating is matched in spatial frequency with the adapter grating, there will likely be an elevation in contrast threshold (greater contrast would be required) for detecting the target relative to the initial contrast threshold measurement – a contrast after-effect of adaptation. This suggests that the response gain of a particular mechanism that is sensitive to the features within the adapter stimulus has been decreased, providing a way of identifying and manipulating certain mechanisms within the visual system.

A range of adaptation after-effects have also been identified, such as the tilt aftereffect (TAE) (Gibson, 1937; Gibson & Radner, 1937). Like contrast detection thresholds, these can be used to study feature-selective mechanisms within the visual system. Adaptation aftereffects describe when the appearance of a stimulus is changed in some way following adaptation to another stimulus that may share similar features. In the TAE, prolonged adaptation to a line or grating causes a change in the perception of another oriented stimulus' orientation away from the adapter (Mitchell & Muir, 1976; Paradiso, Shimojo, & Nakayama, 1989; Peter Wenderoth & Johnstone, 1988). Contrast adaptation to a stimulus of one orientation results in the temporary strengthening of inhibitive lateral activity between populations of broadly-tuned orientation-selective neurons (Tolhurst & Thompson, 1974). The broadly-tuned adaptation results in neurons – whose orientation tuning is further away from the orientation of the stimulus – responding more strongly, which results in the tilted perception of the test stimuli. Roach and colleagues (Roach & Webb, 2013; Roach,

Webb, & McGraw, 2008) identified a form of TAE induced by circular and radial global patterns that was insensitive to spatial frequency but sensitive to orientation (relative to the local structure of the pattern). Interestingly, adaptation seemed to occur in spatially remote locations that did not receive input from the adapter patterns. Following adaptation, they found TAEs still occurred for grating patterns displayed in these spatially remote locations, irrespective of differences in spatial frequency between adapter and test. This suggests the involvement of a neural population sensitive to the global orientation structure of the circular/radial pattern recurrently inhibiting the local orientation channels that they receive their input from. Tan et al. (Tan, Bowden, Dickinson, & Badcock, 2015; Tan, Dickinson, & Asaad, 2016; Tan, Dickinson, & Badcock, 2016) found that shape information within texture-defined radial frequency patterns is globally processed, indicating the involvement of global

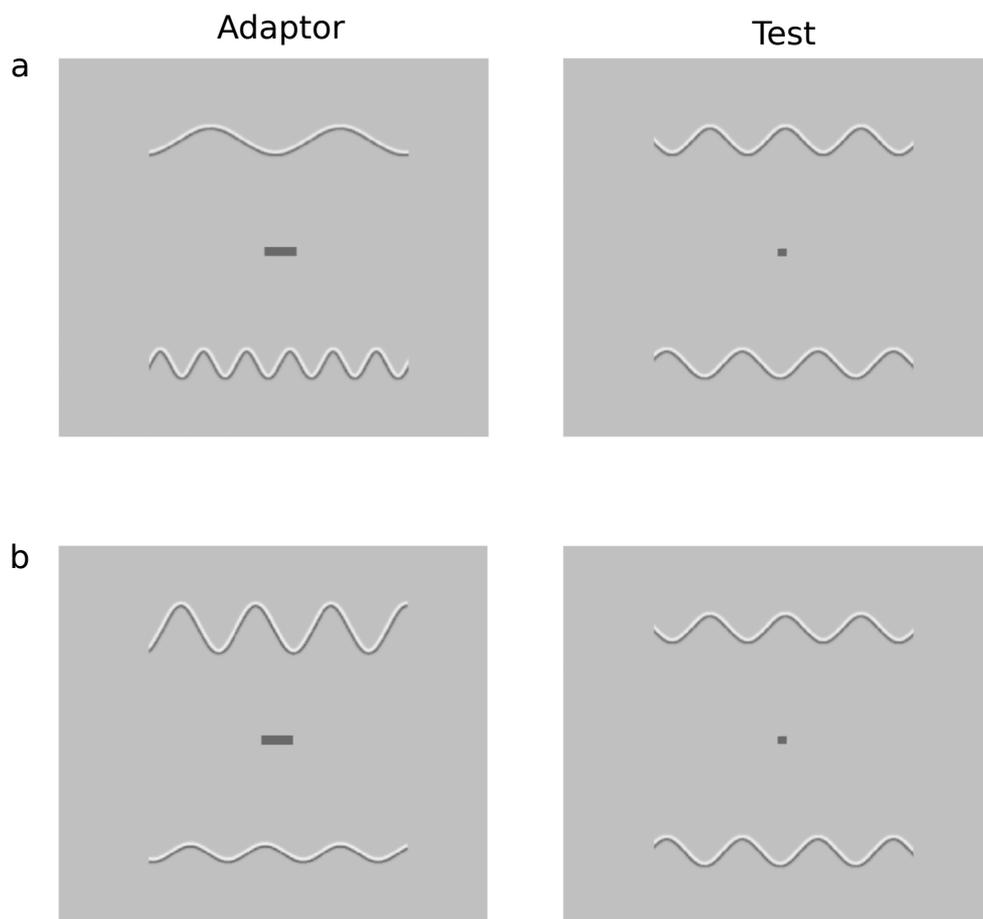


Figure 1.2. Examples of a) SFAE and b) SAAE stimuli taken from Gheorghiu and Kingdom (2009).

mechanisms sensitive to annular shape (however, see Baldwin, Schmidtman, Kingdom, & Hess, 2015; Kingdom, Baldwin, & Schmidtman, 2015).

To more easily understand how V1 outputs are combined in mid-level vision, though, simpler compound stimuli whose components are predictably and optimally encoded in V1 might be more useful (Coen-Cagli & Schwartz, 2013). To this end, a number of novel adaptation aftereffects have been found for simple contours (Bell, Sampasivam, McGovern, Meso, & Kingdom, 2014; Gheorghiu & Kingdom, 2009; Hancock, McGovern, & Peirce, 2010; Hancock & Peirce, 2008; McGovern, Hancock, & Peirce, 2011; Peirce, McGovern, & Hancock, 2009; Suzuki, 2001, 2003). Contour integration and detection is crucial for object perception. It allows us to separate objects in space from other objects and from their background (e.g. Gheorghiu & Kingdom, 2012; Sassi, Vancleef, Machilsen, Panis, & Wagemans, 2010; Vancleef et al., 2013). Gheorghiu and Kingdom (2009) used two adaptation aftereffects; the shape-frequency aftereffect (SFAE) and the shape-amplitude aftereffect (SAAE) (Figure 1.2). Adaptation to two 2D sine-wave adaptor contours, one above a fixation point and one below, was induced with a 90s adaptation period, followed by a repeated 0.5s test period with a 2.5s top-up adaptation period. One of the adapter contours either had a greater peak frequency or greater peak amplitude than the other. During the test period, the adapter contours were replaced with two test contours, both of which were identical in peak amplitude and frequency. They found that, where the frequency or the amplitude of adapter contours was different, participants also perceived a difference in that feature between the test contours (SFAE and SAAE, respectively). For example, in the SFAE condition, following adaptation participants would perceive one of the test contours to have more cycles from end-to-end than the other, even though they were physically identical. A model of contour integration by a curvature detection mechanism explained their findings well. The authors suggested that this can be achieved by an operation in which outputs of sub-unit RFs along the chord of the contours, each like orientation-tuned V1 filters, are combined by multiplication. If just one of the sub-unit RFs do not receive input, for example if there is a long-enough gap in the contour, then the AND-like operation does not take place. However, it has been debated repeatedly whether, in fact, summation of nonlinear V1 responses could perform this job (May & Zhaoping, 2011, 2013, Peirce, 2007b, 2011, 2013).

The SFAE and SAAE methods are difficult to interpret in terms of local versus global processing because the adapter stimuli vary in orientation structure at both levels. Peirce and Taylor (2006) developed a novel compound adaptation procedure to get around this problem and disentangle responses to compound stimuli – in this particular case to plaids – from responses to the components of the compound. They used two pairs of spatial frequency-matched gratings: two had a spatial frequency of 1.5cpd (gratings A and B), and the other pair shared a spatial frequency of 2.5cpd (gratings C and D). Each grating in the pair was orthogonal in orientation to the other: gratings A and B had orientations of 0° and 90°, while gratings C and D had orientations of 45° and 135°. To form compound plaid patterns, each of the grating components were linearly combined with two of the other gratings, which resulted in two spatial-frequency matched, orthogonally oriented ‘coherent’ plaids (AB and CD) and two non-matched, ‘non-coherent’ plaids (AC and BD). Initial adaptation took place for 30s in two retinally independent locations on either side of a fixation spot. On one side, the side of ‘compound adaptation’, the two coherent plaid patterns (AB and CD) were counter-phase modulated at a rate of 2Hz and were interleaved with each other every 1s. On the other side, the side of ‘component adaptation’, the two non-coherent plaids (AC, BD) did the same. This was followed by a 200ms inter-stimulus-interval (ISI), a 200ms probe, 2s top up adaptation, another 200ms ISI and then a second probe. The probes presented on either side of fixation were the same as each other, for example both were plaid CD or both were plaid AD. By presenting all of the grating components in both locations, but in different configurations, the authors hypothesised that if a mechanism exists that is selective for coherent plaids, then the magnitude of contrast adaptation will be larger on the side of compound adaptation. On the other hand, if there is no such mechanism, then contrast adaptation will occur equally as strong on both sides. Their results supported the former hypothesis: they found that contrast adaptation to the plaids was selective for plaid coherence.

Hancock et al., (2010; 2008) adapted this method to identify a curvature aftereffect (CAE). Participants were adapted to flickering compound contours made of two oriented grating components in one hemi-field, and to the flickering grating components alternating 180° out of temporal phase (so that one was always present) in the other hemi-field. The perceived curvature of a straight test pattern was greater

following compound adaptation than following component adaptation, but was abolished when the components were spatially separated (Hancock & Peirce, 2008) or when the spatial frequency of each component was mismatched by more than 2.14 octaves (Hancock et al., 2010). A mechanism sensitive to the compound pattern, but not to either of its components presented in isolation, can explain these results. The abolition of the effect by spatial separation or spatial frequency mismatching suggests that, rather than just indiscriminately combining output from spatially proximal RFs, the mechanism requires them to have similar spatial frequency tuning.

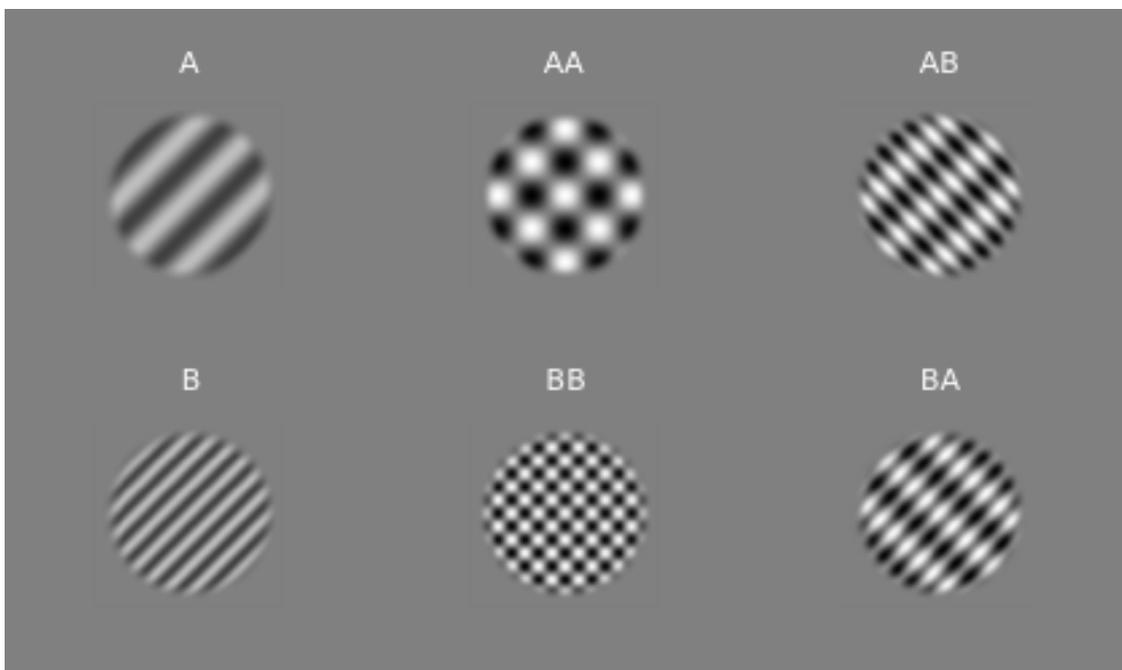


Figure 1.3. Combining gratings to make different kinds of plaids. Grating components (left), coherent plaids (middle) and non-coherent plaids (right). Grating A can be orthogonally superimposed on top of another grating A to form coherent plaid AA, or on top of a grating B to form non-coherent plaid AB.

Plaids present a simple but compelling case of ‘phenomenal coherence’ where the perception of a compound stimulus is greater than the sum of its parts. The combination of superimposed, oriented grating components in certain configurations results in the perception of a distinct chequerboard pattern (a ‘coherent’ plaid). Deviations from these configurations, however, result in the perception of the superimposed, semi-transparent grating components (‘non-coherent’ plaid) (Figure 1.3). This has been observed for plaids in motion (Adelson & Movshon, 1982; Kim & Wilson, 1993) and static plaids (Georgeson, 1992; Meese & Freeman, 1995; Olzak &

Thomas, 1991), though mechanisms responding to drifting plaids likely encode combinations of motion signals while those responding to static plaids encode the form of the plaid (Huk & Heeger, 2002; Movshon & Newsome, 1996; Rust, Mante, Simoncelli, & Movshon, 2006).

Building on Peirce and Taylor's (2006) findings, Hancock, McGovern and Peirce (2010) measured adaptation to compound plaid patterns across a range of component spatial frequencies. They found a contrast aftereffect that was again greatest in magnitude when plaid components matched in spatial frequency (i.e. their combination formed a coherent plaid), with a tuning bandwidth of around ~ 2.72 octaves. This held for a relatively wide range of coherent plaid spatial frequencies; from 0.4c/deg to 6.4 c/deg and displayed different temporal dependencies from the CAE (McGovern et al., 2011). A similar effect has been found for chromatic plaid adaptation (Robinson & MacLeod, 2011). Perceptual plaid coherence breaks altogether (Olzak & Thomas, 1991), or requires higher contrast thresholds as a function of spatial frequency mismatching (Georgeson, 1998). Non-coherent plaids are also picked in a more sequential manner in visual search tasks while coherent plaids 'pop-out' of the search arrays, indicative of pre-attentive mechanisms for perceiving coherent plaids (Nam, Solomon, Morgan, Wright, & Chubb, 2009). Together, these findings point to the existence of mechanisms, distinct from those responsible for the CAE, that selectively combine plaid components based on their relative spatial frequency. Mechanisms that 'match' by spatial frequency might serve to limit the total number of signal combinations that the visual system must carry out, and help it identify object features like edges that are more likely to originate from the same object.

The perceived contrast of plaids also differs from that of the grating components that form them. Likely due to XOS, they have a lower perceived contrast than either of their components presented at the same contrast (Georgeson & Shackleton, 1994). For example, a component presented alone at 32% Michelson contrast would be perceived as being slightly higher in contrast than a plaid composed of two 16% contrast grating components. When presented at low contrast levels plaids fail to cohere (Meese & Georgeson, 1996), and show selective adaptation at higher contrast levels, while gratings show the opposite pattern (McGovern & Peirce,

2010). McGovern and Peirce (2010) also found that contrast adaptation to plaids was spatially phase-invariant, whereas contrast adaptation to the grating components was sensitive to spatial phase. As the spatial phase of a probe grating stimulus relative to the adapter increased, a smaller magnitude in adaptation was observed, while for a compound plaid stimulus this shift in spatial phase of the probe did not affect the magnitude of adaptation. Together, the evidence of spatial frequency selectivity of plaid perception, phase invariance and differences in contrast coding suggests that a selective mechanism beyond V1 encoding of simple edges is involved in the processing of plaids.

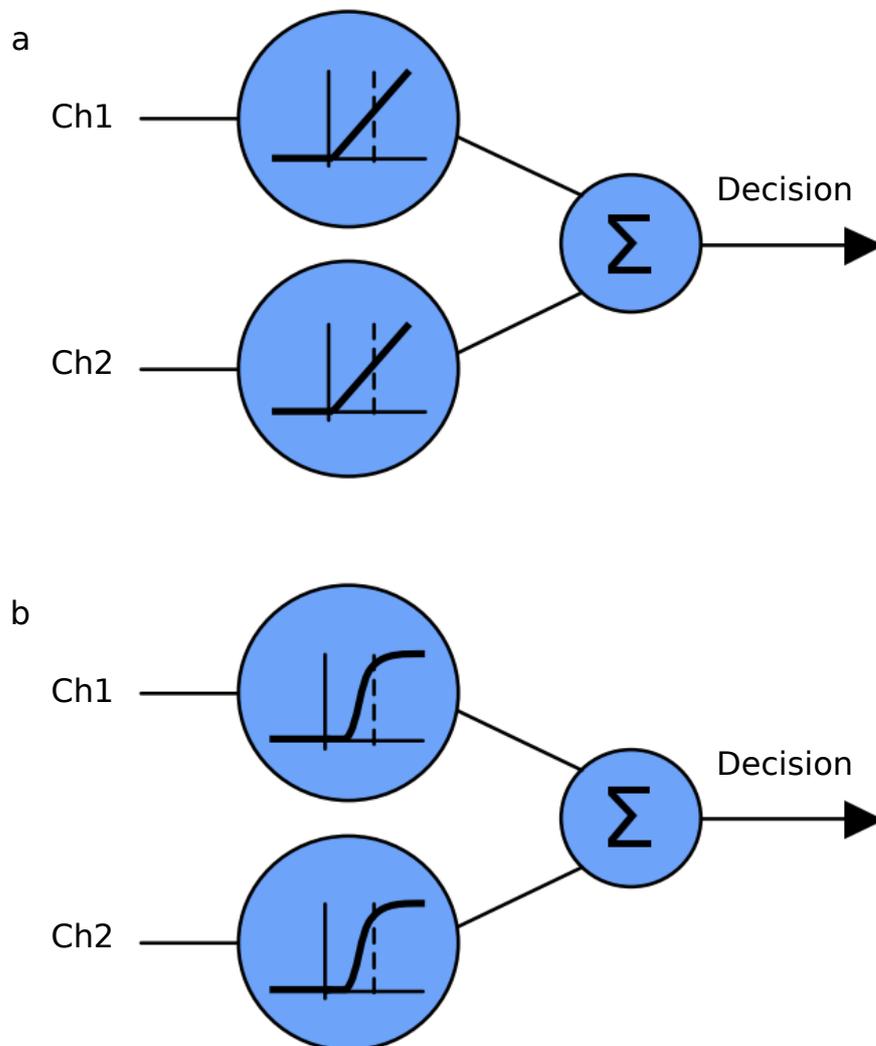


Figure 1.4. Taken from Peirce (2007b). A summing circuit that sums linear outputs (a) cannot differentiate between 100% contrast stimulation to one channel alone and 50% contrast stimulation to two channels. A summing circuit that makes use of compressive nonlinearities (b), such as contrast gain and normalisation, can.

Inspired by the findings of Peirce and Taylor (Peirce & Taylor, 2006) and Olzak and Thomas (Olzak & Thomas, 1999), Peirce (Peirce, 2007b, 2011, 2013) suggested that nonlinearities might be useful for more than just optimising gain control and normalisation. They might in fact be essential for detecting signal conjunctions. Logical AND-gates are summing circuits that can take advantage of saturating nonlinearities in V1 to discriminate between a compound pattern and its components (Figure 1.4) By nonlinearly summing the compressed output of V1 channels, greater stimulation can be provided by the circuit than any one channel could generate on its own. This would allow a mechanism to encode the combination of signals as more than the sum of their parts. For example, imagine a summing circuit composed of two V1 channels that both provide input to a 'conjunction' neuron higher in the visual system. If this circuit sums the output of both V1 channels linearly, the conjunction neuron cannot differentiate between 100% contrast stimulation to one V1 channel alone and 50% contrast stimulation to both channels. On the other hand, if the maximum response achievable is limited by saturating responses, a summing circuit that makes use of compressive nonlinearities can.

Detecting nonlinear summation using EEG

The difficulty of understanding what nonlinearities might drive selective mechanisms in mid-level vision lies in being able to disentangle one nonlinearity from another. Those giving rise to normalisation processes clearly play a large role in the detection of features in the visual scene. It is also clear that for detecting conjunctions of information, other nonlinearities must be involved. Following XOS, an additional nonlinearity, such as a logical AND-gate, may then process the plaid form by selectively combining outputs from V1 channels.

So how might one disentangle the multiple nonlinear processes of mid-level vision? The spatio-temporal resolution of fMRI seems too coarse to distinguish the different potential neural components that contribute to the overall response. Although EEG has a poorer spatial resolution, its superior temporal precision might be useful in this endeavour. For instance, if the additive responses, the normalisation

nonlinearities and the combination nonlinearities operate on different timescales, but in similar cortical regions then EEG might be able to distinguish them.

ERPs: what do they represent?

Transient visual evoked potentials (VEPs, but also referred to as visually evoked responses (VERs) e.g. Vassilev, Manahilov, & Mitov, 1983) are ERPs time-locked to the onset of a brief visual stimulus. Owing to the ease with which sensory systems can be examined relative to higher cognitive functions, transient VEPs were among the first ERPs to be studied. They provide an incredibly rich data set with excellent temporal precision, allowing researchers to study the time course of characteristic peaks and troughs in amplitude, known as response 'components', in depth. Those which occur earlier in the time course, that is to say those with an earlier latency, tend to be more strongly influenced by low-level stimulus properties than later components (Manahilov & Vassilev, 1986; Smith & Jeffreys, 1978).

The earliest of these components is referred to as the C1 (Fuxe et al., 2008; Jeffreys & Axford, 1972a), as well as other names including the N75 and N1 (Fortune & Hood, 2003; e.g. Kevin, Doug, Matthias, & Gerhard, 2008; Manahilov & Vassilev, 1986; Parker & Salzen, 1977). It reflects feedforward signals in early visual cortex that can be observed consistently around posterior occipital channels (Elleberg, Hammarrenger, Lennie, Roy, & Guillemot, 2001; Hansen, Haun, Johnson, & Elleberg, 2016; Miller, Shapiro, & Lovegrove, 2015; Reed, Marx, & May, 1984). The typical onset latency of the C1 is around 40-60ms and reaches peak amplitude anywhere between 50-110ms post-stimulus (Clark, Fan, & Hillyard, 1995; Di Russo, Martínez, Sereno, Pitzalis, & Hillyard, 2002; Elleberg et al., 2001; Hansen et al., 2016; Rokszin, Gyori-Dani, Nyúl, & Csifcsák, 2015).

Unlike later components such as the N170 or P300, the polarity and topography of the C1 component is known to vary in polarity. It has been suggested that this is consistent with the retinotopic organisation of V1 (e.g. Clark et al., 1995; Di Russo et al., 2002), and in-line with the cruciform model of striate cortex (Butler et al., 1987; Holmes, 1945; Jeffreys & Axford, 1972a, 1972b; Kelly, Vanegas, Schroeder, & Lalor, 2013; Vanegas, Blangero, & Kelly, 2013). This proposes that the polarity

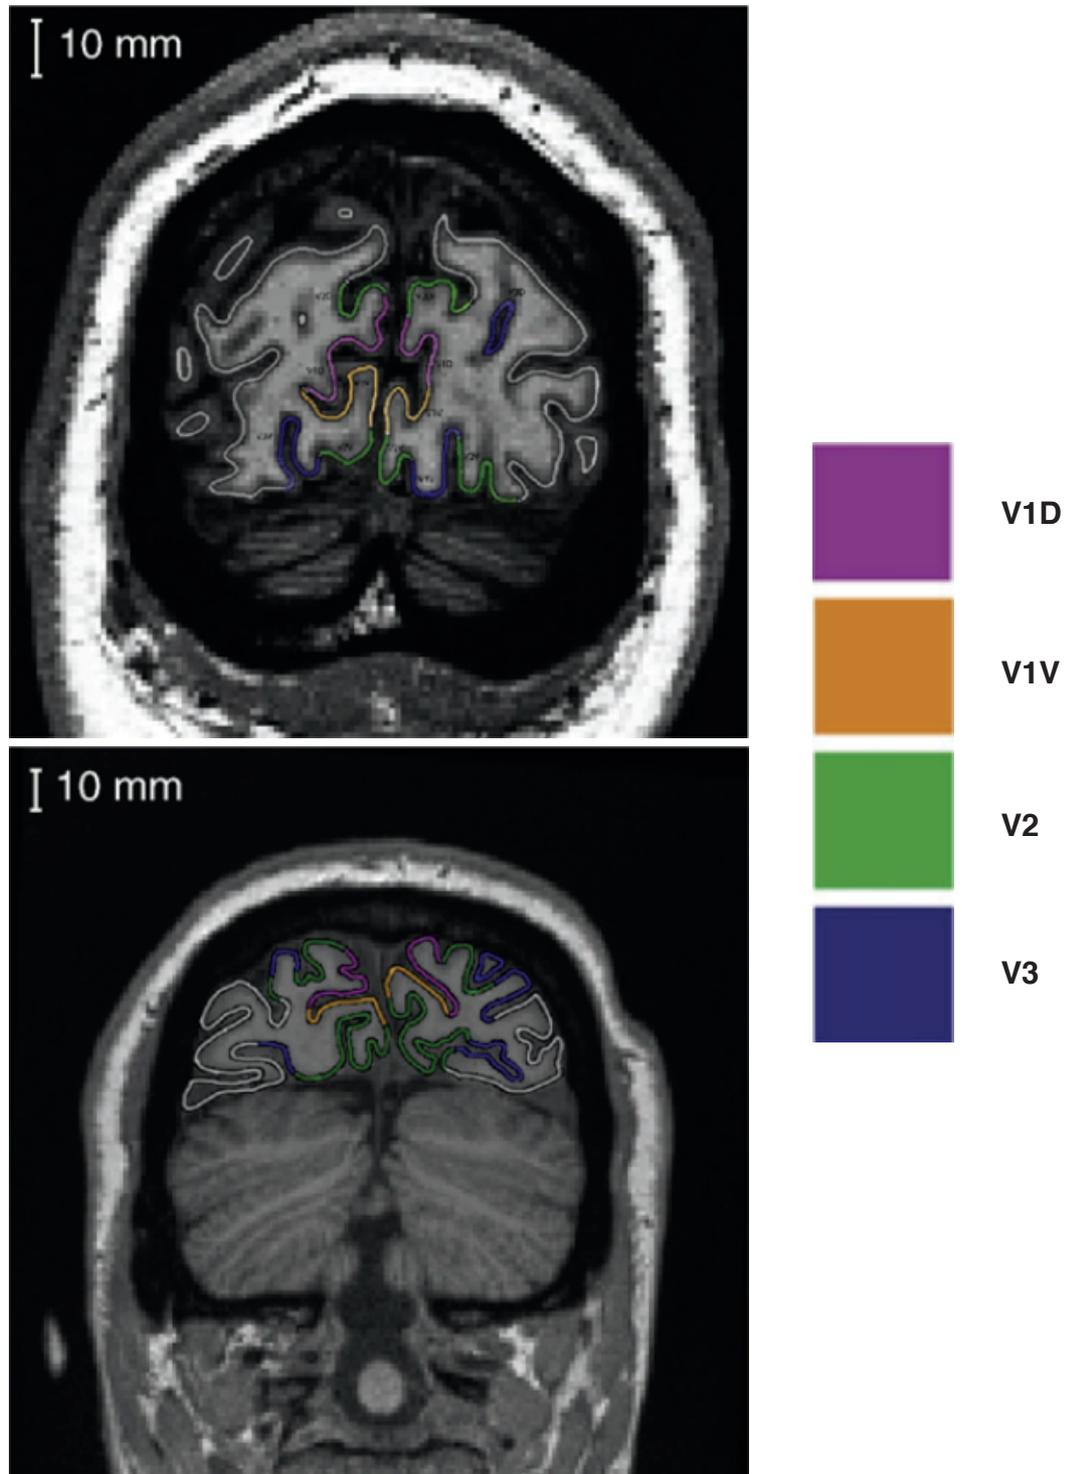


Figure 1.6. Taken from Ales, Yates and Norcia (2010). MRI scans from two participants with visual areas V1, V2 and V3 labelled. Neither participant conforms directly to the cruciform model of the calcarine sulcus. For example, the bottom participant resembles the cruciform model in terms of having well defined upper and lower banks of V1, but there is also complex cortical folding in areas V2 and V3.

reversal observed for C1 is caused by cortical folding around the calcarine fissure. Stimuli presented to the upper right quadrant of visual space, for example, will be

projected to the lower left bank of the calcarine fissure, resulting in a negative scalp potential. Stimuli presented in the lower right quadrant of visual space, on the other hand, will project to the upper left bank of the calcarine fissure and result in a positive scalp potential. Using this model, recent studies have attempted to 'map' retinotopic sensitivities onto the visual field on an individual participant basis with some success (Kelly, Gomez-Ramirez, & Foxe, 2008; Vanegas et al., 2013) to minimise inter-individual variation caused by inherent differences in cortical folding (e.g. Rademacher, Caviness, Steinmetz, & Galaburda, 1993). However, it remains unclear whether the cruciform model is a good candidate model for identifying activity originating in V1. Ales, Yates and Norcia (2010, 2013) found that in some cases, no polarity inversion was found in V1 sources, but was for areas V2 and V3. They displayed that the cortical folding assumed by the cruciform model often does not occur (see Figure 1.6), and suggested that multiple sources contribute to the C1 response beyond V1.

What is clear, however, is that foveal presentations of a stimulus tend to generate a potential with different characteristics than following peripheral presentations. The VEP following foveal presentations changes drastically across spatial frequency. At lower spatial frequencies the time course is dominated by a positive-going deflection (P1) with a peak latency around 120msec, but at higher spatial frequencies it is dominated by a negative-going deflection between 70-110ms (Baas, Kenemans, & Mangun, 2002; e.g. Elleberg et al., 2001; Hansen et al., 2016; J. G. May & Lovegrove, 1987; Proverbio, Zani, & Avella, 1996; Reed et al., 1984; Tobimatsu, Tomoda, & Kato, 1996). Hansen et al (2016) found that this initial negative component to centrally placed stimuli, which they referred to as fC1 (foveal C1), was high pass for spatial frequency and was generated at more posterior occipital channels, while the first major component generated by peripheral stimuli (pC1), was bandpass for spatial frequency and observed at slightly more anterior occipital channels. This, combined with a latency difference of ~32msec, led them to suggest that the fC1 component generated to centrally placed stimuli may in fact reflect the C2 component, which typically follows the pC1 referred to by Jeffreys and Axford (1972a, 1972b).

Using ERPs to study nonlinear responses

VEP components can be used to measure specific nonlinear responses. In a study examining how the human VEP is changed by adaptation, Manahilov and Vassilev (1986) measured the first and second negative-going deflections of the VEP, which I will refer to as the C1 and N1. They label them differently in their paper, but as I mentioned earlier there is inconsistency in component labelling in the literature. They found that the N1 component showed a general adaptation irrespective of stimulus features. The C1, on the other hand, showed suppression only when the spatial frequency of the test and adapter gratings were matched, indicating that it reflected a spatial-frequency selective response.

More recently, Miller, Shapiro and Luck (2015) measured C1, C2 and lateral occipital P1 responses to checker 'wedges' in the upper or lower visual field quadrants that were presented in isolation from each other (e.g. one wedge in the top left and another in the top right quadrant), presented together (e.g. two wedges in the top left and another two in the top right quadrant) but spatially separated by 2° of visual angle, or proximal to one another with a gap of 0.16° of visual angle. They hypothesised that if stimuli presented simultaneously are processed relatively independently, then the VEP measured should roughly equate to the linear sum of responses to each component. On the other hand, if they are processed by a common mechanism then the VEP response to multiple components together should be less than the linear (i.e. a normalised) sum of responses to each of the components, which they referred to as 'competition'. They found that the C1 response showed moderate levels of normalisation when the check wedges were closer together than when further apart, showing a less than linear response that was still larger than the response to a single stimulus. The C2 component, which Hansen et al., (2016) suggested is the C1 observed for central stimuli, was represented by an average response to the components than a linear sum. Similarly, Chen, Yu, Zhu, Peng and Fang (2016) found that early components (C1, P1) showed differing degrees of nonlinear spatial summation to proximal grating components. They found that C1 to three proximal grating stimuli was sub-additive only when participants actively attended to the stimuli. This affect

disappeared when participants did not attend and when the gratings were further apart from each other.

Clearly, the early VEP response components are highly, and reliably, sensitive to low-level stimulus features. As demonstrated by more recent studies, they may also be suitable for studying nonlinear summation of information in the visual system. The nonlinear summation involved in coherent plaid processing has not yet been assessed using the VEP approach.

Frequency-tagging is useful for separating nonlinearities

The 'steady state' ERP (Cobb, Morton, & Ettliger, 1967; Compston, 2010; D. Regan, 1966) or steady state visual evoked potential (SSVEP) also provides a useful opportunity for studying visual system nonlinearities. The method involves modulating the intensity ('flickering') or position of a visual stimulus at a fixed rate (F Hz, where F is any real number >0) over an extended period. The result is a train of synaptic potential responses observable as a waveform that achieves a 'steady state'; it shares a stable, periodic relationship to the stimulus in amplitude and phase across the presentation time.

The human visual EEG can be decomposed into many superimposed sinusoidal functions (D. Regan, 1966), something which the steady-state approach can take advantage of in combination with Fourier analyses. An SSVEP to one stimulus, sinusoidally modulating in intensity between 0 and full intensity at a rate of $F1$, will not generally display a sinusoidal response only at $f1$ (Figure 1.5a). To do so would require a neural population to respond linearly to the full stimulus signal. As outlined earlier, neurons in the visual system tend to display asymmetric responses; signal nonlinearities indicative saturation and sensitivities to things like orientation and spatial frequency. The result of these signal transformations is additional 'harmonic' responses at multiples of the stimulus input frequency, at ' nf ', where ' n ' is a small integer value greater than 1 (see Figure 1.5b). A large body of work, particularly clinical research, utilises 'pattern-reversal' stimulus modulation (or 'contrast-reversal'), where stimulus intensity sinusoidally modulates between 1 and -1, effectively reversing the spatial phase of the stimulus half way through the stimulus period (Cobb

et al., 1967). The result of this is a reliable ‘fundamental’ response at $2f_1$, not f_1 , because the phase-reversal results in double the number of synaptic potentials being generated within a single stimulus cycle.

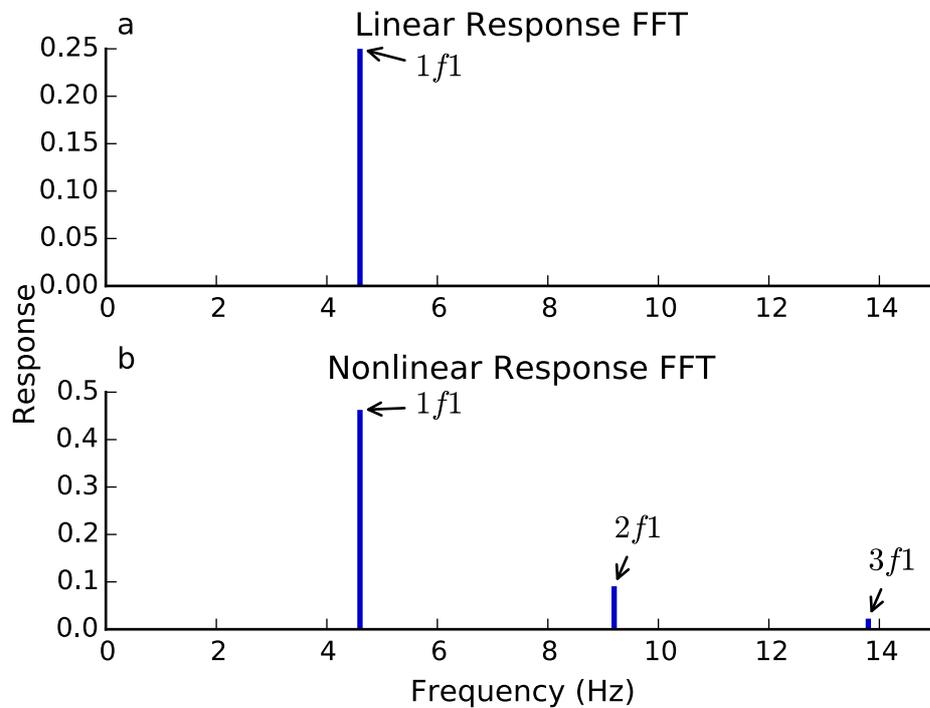


Figure 1.5. Examples of a) linear and b) nonlinear SSVEP response. Both responses were generated to the presentation of a stimulus sinusoidally modulating in Michelson contrast between 0 and 1 at a rate of 4.6Hz. A neuron using a contrast saturation operation was simulated to generate the nonlinear response in b). This resulted in harmonic responses at 9.2Hz ($2f_1$) and 13.8Hz ($3f_1$).

Victor (1979) and Victor and Shapley (1980) made single-unit recordings in cat retinal Y-cells in response to stimuli modulated in intensity as a function of the sum of 8 periodic sinusoidal functions. Similar in principle to measuring the contrast response function to investigate the response nonlinearity of a mechanism to a particular stimulus, using this ‘sum-of-sinusoids’ approach allowed them to investigate the order of nonlinearity of the cell under study at nf_i , where ‘i’ is an integer between 1 and 8. Additional frequencies, not present in the input signal, at harmonics and at $nf_i \pm mf_j$ (known as ‘intermodulation frequencies’), where ‘m’ is another integer and ‘j’ an integer value such that $1 \leq j \leq 8$, are present in the resulting SSVEP.

Using frequency-tagging to measure nonlinear signal summation

Modulating stimuli at many different temporal frequencies, however, can be problematic. The signal-to-noise ratio (SNR) at higher order nonlinear responses is quite low; by summing multiple narrow-band signals, attenuation occurs for higher-order responses with low absolute amplitude through overlap with other higher-order frequency responses. A wide range of stimulus frequencies can be used to avoid this, but is not necessarily a useful option for investigating visual system nonlinearities where relatively low driving frequencies are preferable for the generation of higher-order responses (e.g. Alonso-Prieto, Belle, Liu-Shuang, Norcia, & Rossion, 2013; Boremanse, Norcia, & Rossion, 2013). Regan and Regan (1988) suggested that the properties of higher-order responses to a particular stimulus may represent a ‘fingerprint’ of nonlinear response for the mechanisms involved in processing it. The attenuation of higher-order responses resulting from the sum-of-sinusoids approach is therefore a major disadvantage for investigating specific nonlinear mechanisms. A more frequency-selective technique would be preferable to understand what nonlinear responses correspond to, and to be able to disentangle lateral nonlinearities from other nonlinearities.

The ‘two-frequency’ or ‘frequency-tagging’ approach involves simultaneously modulating the components of a two-component stimulus at slightly different temporal frequencies to one another (see Figure 1.6). This makes it much simpler to pick stimulus frequencies that are temporally incommensurate; very little (if any) overlap occurs between their harmonics and intermodulation responses. It provides a rich dataset that can be used to look at, for example, the responses at the two fundamental frequencies to examine competition effects between the tagged stimuli (Andersen, Müller, & Hillyard, 2015; Appelbaum, Wade, Vildavski, Pettet, & Norcia, 2006; Keitel, Andersen, Quigley, & Müller, 2013; D. Regan & Heron, 1969). Of interest to the present thesis, however, is the narrow-band quality of the evoked response. It makes it easier to examine the higher-order nonlinearities involved in the combination of the two stimuli at nf_1 , nf_2 and $nf_1 \pm nf_2$. Responses at intermodulation frequencies now represent a nonlinearity at or after the point of signal combination between the two stimulus components (Spekreijse & Oosting, 1970a; Zemon & Ratliff, 1984). They

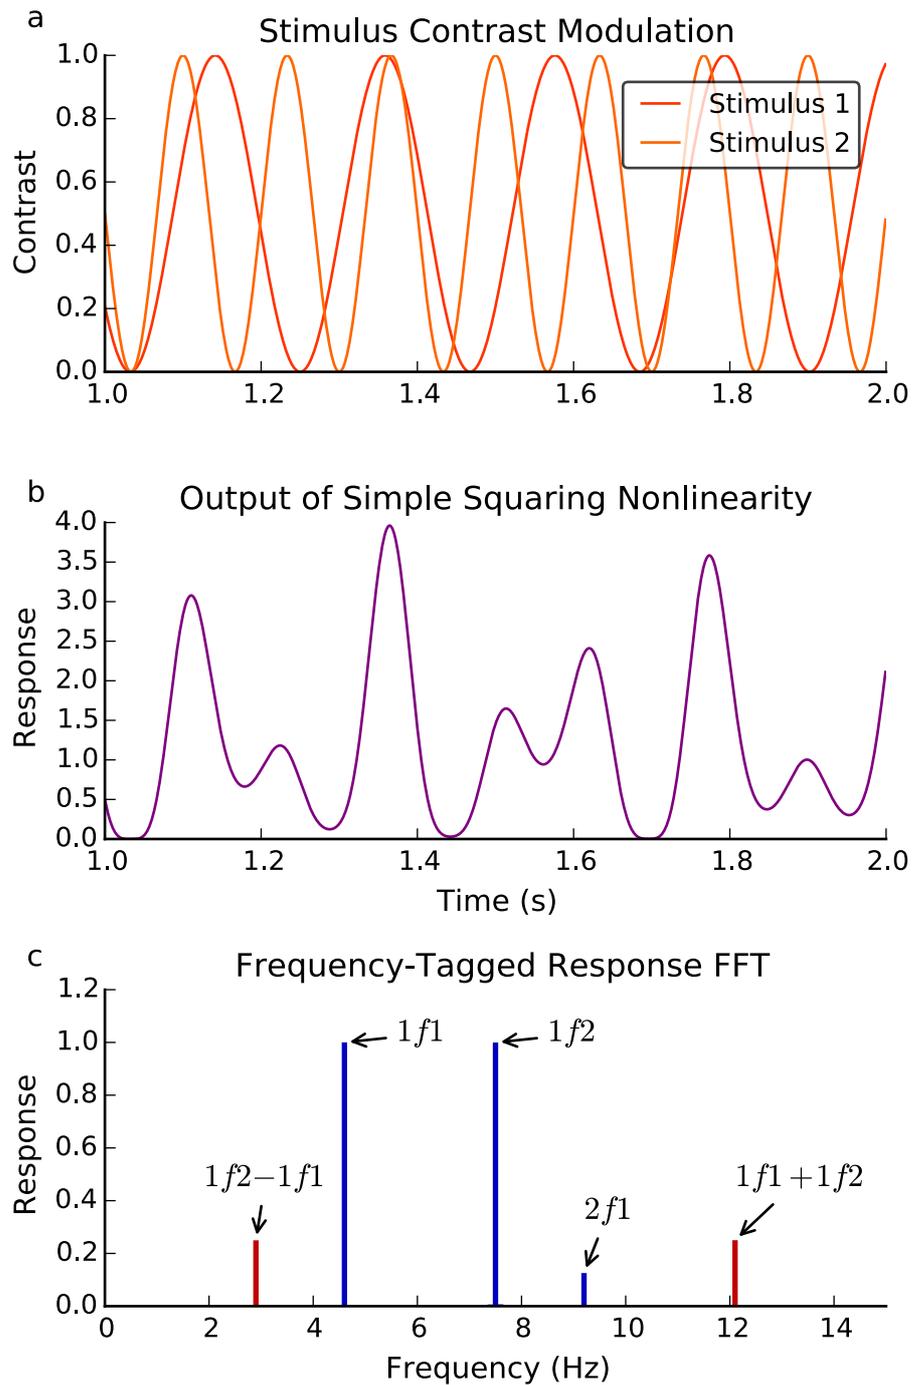


Figure 1.6. Example of frequency-tagged stimuli and the output of a nonlinear operation. a) 1 second window of independent contrast modulation of two stimuli at 4.6Hz (red) and 7.5Hz (orange). b) The output produced by first summing the contrast modulations shown in a) and then squaring this sum to generate a nonlinear response. c) This represents the FFT of this squared 'response'. Marked in blue are fundamental and harmonic responses relating directly to each stimulus component signal. Marked in red are the intermodulation responses generated by the nonlinear (squaring) summation of the two signals

cannot result from independent channels, nor from simple linear combinations of

signals. On the other hand, they may represent lateral, suppressive nonlinearities and, critically, spike-generating nonlinearities by neurons whose receptive fields make them sensitive to the combination of the two stimuli (see M. P. Regan & Regan, 1988 for a mathematical explanation).

This technique has for several decades been used to examine multiple lateral suppressive nonlinearities in the visual system and the tuning functions of early visual system mechanisms (Baith & Levi, 1988; T. J. Baker, Norcia, & Candy, 2011; Candy, Skoczenski, & Norcia, 2001; Suter, Perrier, Parker, Fox, & Roessler, 1996; Tsai, Wade, & Norcia, 2012). Norcia and colleagues (T. J. Baker, Norcia, & Candy, 2011; Candy et al., 2001; Tsai et al., 2012) have used frequency-tagged pattern-reversal gratings to examine the development of monocular and binocular XOS and orientation tuning bandwidths. For example, nonlinear responses measured at $nf1 \pm nf2$ in combination with examination of responses at $2f1$ and $2f2$ have allowed them to identify immaturities specific to XOS (non-selective masking at $2f1$ and $2f2$) but not binocular integration ($nf1 \pm nf2$ responses still present, indicative of binocular integration) in infants. Relatively few studies, however, have used the frequency-tagging technique to investigate responses to compound stimuli in mid-level vision.

Compound stimuli that are perceived as more than the sum of their parts are pervasive perceptual phenomena in human vision; the visual system must somehow segment irrelevant information and bind that which is relevant to identify patterns, shapes, surfaces and objects in the visual environment. The nonlinear combinations involved may be fundamental operations performed by mid-level visual system mechanisms and could result in intermodulation responses. Boremanse, Norcia and Rossion (2013) presented half-face stimuli at two different frequencies and measured the sum and difference intermodulation responses. These were larger when two face-halves were presented as a whole face compared to when they were horizontally separated or vertically misaligned. They attribute this to a nonlinear mechanism sensitive to the whole, rather than the parts, of the face. Similarly, intermodulation responses have been found when parts of a Kanizsa-type illusory stimulus (Kanizsa, 1979) are presented at different rates but not when the parts are rotated so as to 'break' the illusory percept (Alp, Kogo, Van Belle, Wagemans, & Rossion, 2016;

Gundlach & Müller, 2013), when an object is distinct from its background (i.e. figure-ground interactions: (Appelbaum, Wade, Pettet, Vildavski, & Norcia, 2008b), for contour integration in Vernier stimuli when the contours are collinear (Victor & Conte, 2000), and to visual form-motion binding when the abutting sides of a square are moved closer together (Aissani, Cottureau, Dumas, Paradis, & Lorenceau, 2011). Mechanisms that perform selective signal combinations might then operate throughout cortex.

As electrophysiological evidence for selective plaid responses based on plaid coherence has yet to be provided in the literature, the frequency-tagging SSVEP technique presents a unique and data-rich approach for doing so. Likewise, as relatively little is known about how intermodulation responses are generated by specific mechanisms, the dependence on coherence for the putative plaid detector presents a relatively simple way to distinguish the nonlinearities feeding the intermodulation response, and therefore develop new insights into how they occur. If special mechanisms for plaids exist and are partially driving the intermodulation responses, then it would be expected that they would differ according to plaid coherence. Conversely if the intermodulation response is driven only by a normalisation pool for XOS between the orthogonal grating components, which is thought to be very broadly tuned to spatial frequency (DeAngelis et al., 1994; Petrov et al., 2005), then the intermodulation response should be the same for different forms of plaid (or should at least be related to the strength of suppression directly measured at component frequencies). Surround-suppression may also contribute to the intermodulation response for spatially extensive stimuli. Using a contrast sensitivity task, Petrov, Carandini and McKee (2005) demonstrated for the first time in humans that XOS and surround suppression are two distinct mechanisms. XOS is more broadly tuned for orientation and spatial frequency and occurs earlier than surround suppression. Also, the spatial extent of XOS covered both foveal and peripheral areas of stimuli, while surround suppression was present only in peripheral areas. They suggested that the mechanism producing the sharp orientation and spatial frequency tuning observed in the periphery for surround suppression actively enhances cortical response specificity, rather than contributing to normalisation between responses to the target and mask by XOS. The use of orthogonally oriented grating components,

however, should make contributions from surround suppression to intermodulation responses relatively weak. Petrov et al (2005) found that an orientation difference of around 60° was enough to abolish surround suppression.

The response asymmetries that the frequency-tagging approach is able to take advantage of are not a property of visual system neurons alone, but are prevalent throughout the different sensory systems (M. P. Regan & Regan, 1988). Likewise, the need to understand combinatorial mechanisms is not a visual-domain specific problem (e.g. Colon, Legrain, & Mouraux, 2014; Wile & Balaban, 2007). Nozaradan, Zerouali, Peretz and Mouraux (2015), for example, recently used the frequency-tagging approach to study cross-modal mechanisms that combine auditory and somatosensory information.

The technique of measuring responses to multiple frequency-tagged stimuli and using the responses to understand a variety of nonlinear combinations would seem to be very powerful technique, given that it's likely that nonlinear signal combinations take place throughout the brain.

Summary

This thesis will examine a range of possible nonlinear combinations by using plaids. It will focus on nonlinear interactions in mid-level vision using both transient ERPs and frequency tagging techniques. It will also study signal combinations in the auditory system by combining tones to form consonant and dissonant chords; similar mechanisms may be operationalised in systems other than the visual system.

Chapter 2: General Methods

Recruitment

All participants had normal or corrected-to-normal vision (for vision experiments) and normal or corrected-to-normal hearing (for the auditory experiment), and gave informed consent to participate. The ethics board at the School of Psychology, University of Nottingham, granted ethical approval. The work described was conducted in accordance with the 2008 version of the Declaration of Helsinki.

EEG Set-up

A DBPA-1 Sensorium bio-amplifier (Sensorium Inc., Charlotte, VT, USA) was used for EEG recording at a sampling rate of 1000Hz. Voltage responses were recorded from 122 electrode channels (silver/silver-chloride) on a set of customized whole-head caps with twisted and fixed electrode cables (EasyCap, Munich, Germany). This included a ground electrode placed on the forehead, reference electrode at the left mastoid, EOG electrodes (RHE, LHE and LIO) and 117 scalp electrodes. Caps were centered on electrode Cz, halfway between the nasion and inion. Where possible, impedances were brought below 25 kilohms ($k\Omega$) before the experiment began (below 50 $k\Omega$ if this could not be achieved). Note that on many commercial systems this might seem like high impedance, but is the suggested setting for the Sensorium DBPA-1 amplifier.

A parallel port from the stimulus computer was used to indicate when a stimulus onset occurred, sending a trigger signal time-locked to the screen refresh. We confirmed in advance that this was precise on our hardware, using a photometer connected to the amplifier via a StimTracker (Cedrus Light Sensor, Cedrus Corporation, San Pedro, CA, USA).

Gamma correction

It was critical that we did not introduce nonlinearities into the measurements erroneously by using uncalibrated monitors. In the transient experiment (Chapter 3), a computer-controlled monitor (Iiyama Prolite X2472HD, Iiyama, Hoofddorp, The Netherlands) with a screen resolution of 1920 x 1080 pixels, mean luminance of

133cd/m² and a refresh rate of 60 Hz was used for stimulus presentation. Gamma correction was carried out using a colorimeter (ColorCal MKII, Cambridge Research Systems, Kent, UK). The intermodulation experiment reported in Chapter 4 used the same monitor, but the gamma function of the screen was linearized using a spectroradiometer (SpetraScan PR-655, Photo Research Inc., Chatsworth, CA, USA). Chapter 5's experiment used a computer-controlled LCD monitor (Display++, Cambridge Research Systems Ltd., Kent, UK) with a screen resolution of 1920x1080, mean luminance of 120cd/m² and a refresh rate of 120Hz was used for stimulus presentation. The gamma function of the screen was linearized using the pre-set gamma table stored within the monitor, and this was checked using a spectroradiometer (SpetraScan PR-655, Photo Research Inc., Chatsworth, CA, USA).

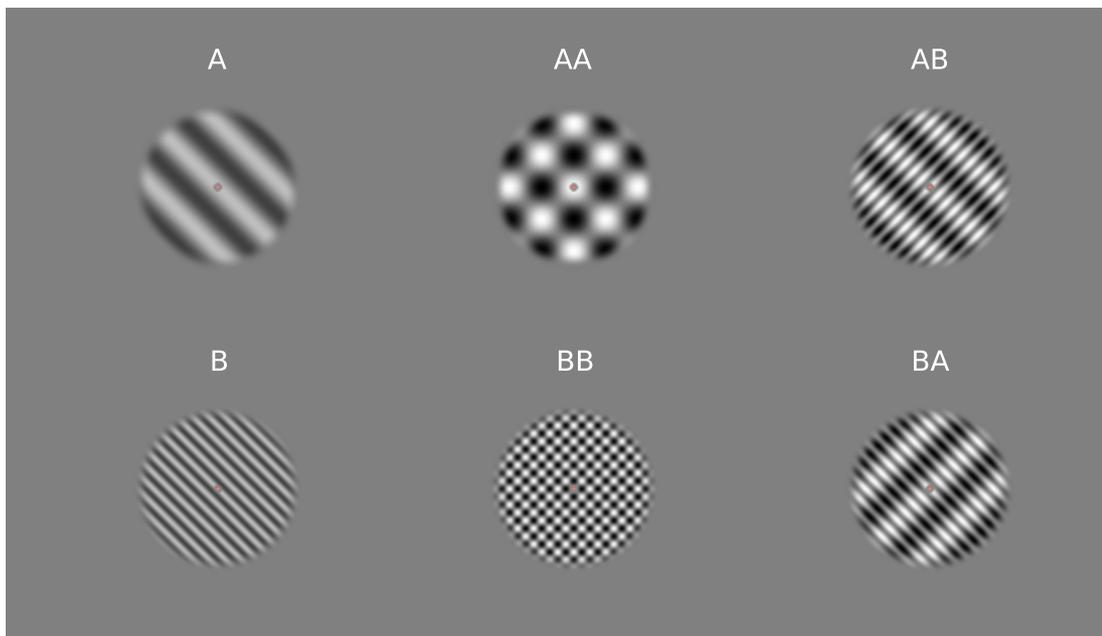


Figure 2.1. Combining gratings to make different kinds of plaids. Grating components (left), coherent plaids (middle) and non-coherent plaids (right). Grating A can be orthogonally superimposed on top of another grating A to form coherent plaid AA, or on top of a grating B to form non-coherent plaid AB.

Stimulus configuration and experimental procedure

Stimuli (see **Figure 2.1** for a general example) were presented within a circular window that had a raised cosine edge profile (width=1.5°, referring to the width of the blurred edge of the raised cosine mask that was overlaid on the grating). A further circle was placed in the centre of the stimulus with the same mid-grey colour as in the

background. This also had a raised cosine edge profile (width = 0.08°), and helped accentuate the fixation dot – a red circle in the centre of the screen. In the transient experiment, the stimulus circle subtended 10° in diameter, the central circle 0.6° and the fixation dot 0.2° . In the intermodulation experiments the stimulus circle subtended 7.5° in diameter, the central circle 0.4° , and the fixation dot 0.175° .

Maintaining visual attention: the green dot task

In the visual experiments that I go on to outline, participants were asked to fixate on the central fixation dot described in the General Configuration subsection, but spread their visual attention to detect and respond as quickly as possible to the appearance of a green dot on the surface of the stimulus on each trial. To collect responses, a button box (Cedrus RB840) was placed in front of participants, where participants could easily rest their hand throughout the experiment. The use of an attentional task, independent of the stimulus conditions, is common practice in fMRI studies since the finding that differential attention to different stimuli can provide an important confound (Huk, Ress, & Heeger, 2001). The task dot, subtending 0.3° of visual angle in the transient experiment and 0.25° in the intermodulation experiments, appeared in a random radial location half-way between the edge of the central circle and the edge of the stimulus circle. The timing of the dot's appearance was set to be unpredictable to the participant to make the task challenging. The timing of this varied for each experiment, and will be described in each chapter separately.

Making plaids

Plaid coherence is achieved by spatially superimposing similar orthogonal sinusoidal gratings, each presented synchronously in time at similar contrast levels. This is well-suited for measuring transient VEPs, which are generated in response to brief presentations of a stimulus, and is the approach used in the transient experiment discussed in this thesis.

The principle of spatially superimposing orthogonal sinusoidal gratings remained unchanged in the intermodulation experiments. However, to generate the intermodulation responses in plaid conditions, the two grating components had to

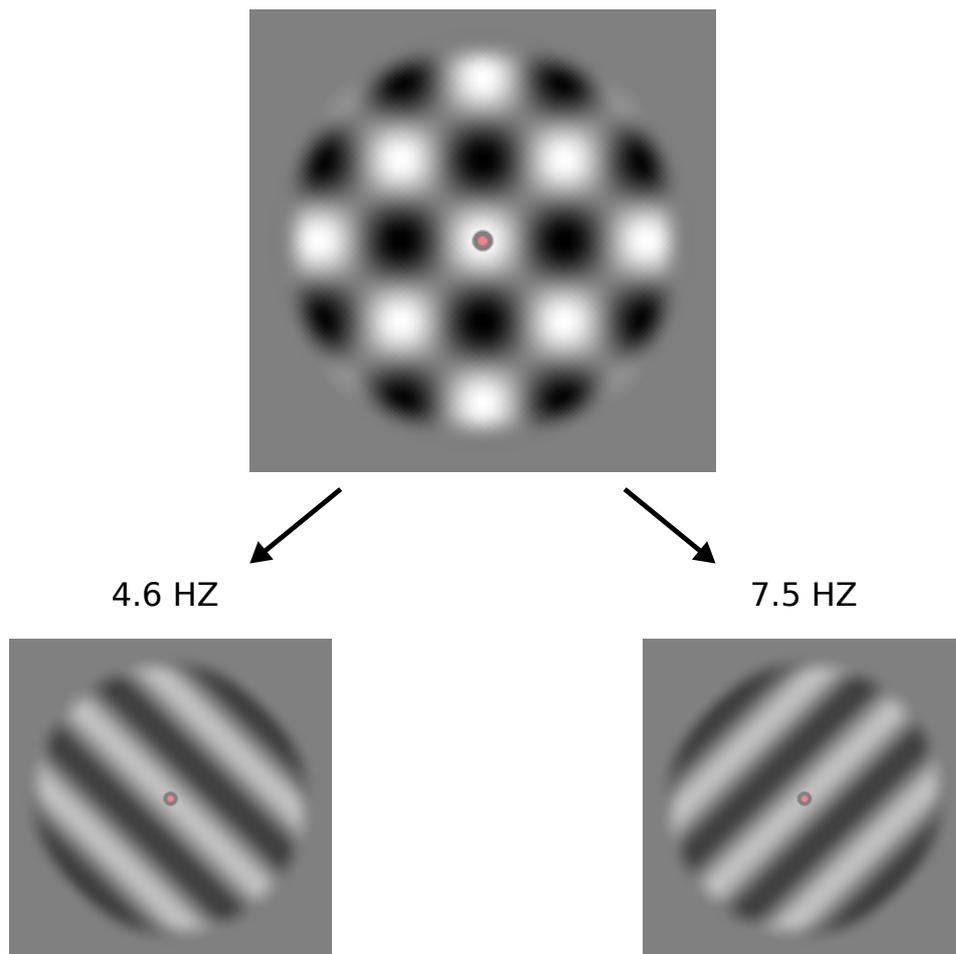


Figure 2.2. Illustration of the separate frequency-tagging of plaid components.

have their contrast intensity simultaneously modulated at slightly different temporal frequencies (**Figure 2.2**). The contrast of any one component varied sinusoidally in time between 0 and a given maximum contrast relative to the full Michelson contrast of the monitor (0.99). The plaid patterns were formed from the signed sum (black is negative, white is positive) of the pixel values resulting from the two components such that, when the components were both at the peak of their sinusoidal modulation in time, a standard plaid pattern was physically present momentarily, whereas when either component was at its trough only the other component was physically present, as a simple grating. By smoothly (sinusoidally) ramping up and down the stimulus intensity between 0 and a maximum contrast, the frequencies chosen to modulate stimuli did not have to be tied to the monitor refresh rate. The sampling rate of the EEG system and the length of presentation time were both critical for measuring at the correct frequency resolution. By accounting for these, the measured responses

were smooth responses at our frequencies with the exception a few frequencies in Chapter 4 (as described below).

The frequencies chosen for contrast modulation were selected on the basis that they allowed the response frequencies to be temporally incommensurate; the fundamental frequencies and their harmonics, and the intermodulation terms and harmonics would not overlap. Further, we aimed to avoid overlap between the typical alpha band (8-12Hz) and the sum intermodulation response frequency as in Boremanse, Norcia and Rossion (Boremanse et al., 2013) to boost signal-to-noise. The experiment reported in chapter 4 used contrast modulation frequencies of 2.3Hz and 3.75Hz, and 4.6Hz and 7.5Hz were used in chapter 5. Chapter 4 was supposed to use 4.6Hz and 7.5Hz, too, but stimuli were erroneously presented at half the presentation frequency. The recorded data did not suffer for this mistake; as will be outlined in chapter 4, we found significant results that went on to be published.

An initial concern was whether a coherent plaid still appears when you have two components that are infrequently matched in contrast across time. Due, presumably, to the reasonably high rate of contrast modulation, the percept for the observer is not of two gratings changing contrast gradually. Instead, the percept is of a plaid that alternates with gratings at a high and unpredictable rate.

Choosing spatial frequencies

Transient. Gratings A and B were two sinusoidal grating patterns. Grating A had a spatial frequency of 2.5cpd and grating B a spatial frequency of 5.5cpd. These spatial frequencies were chosen to ensure that C1 responses driven by centrally placed stimuli would be generated (Elleberg et al., 2001; Hansen et al., 2016). They were also chosen so that there was a big enough separation between A and B to generate perceptually distinct coherent and non-coherent plaids but keep spatial frequency low enough to not affect visual acuity.

Intermodulation. As in the transient experiment, gratings A and B were two sinusoidal grating patterns. However, grating A had a spatial frequency of 1cpd and grating B a spatial frequency of 3cpd. This was because the C1 component response was no longer the measurement of interest. Having spatial frequencies ~ 4 cpd was no longer

necessary, and I wanted to make the checks in the plaid patterns more salient by making them larger.

Analysing SSVEP data

Data were band-pass filtered between 0.1Hz and 100Hz. They were then epoched according to stimulus onset, with the first second of data removed to exclude onset transients from the analysis, resulting in a 10-second epoch for each trial. Trials were time-averaged by condition for each participant, averaging out activity that is not phase-locked to the stimulus presentation, such as the prominent occipital cortex alpha-wave response. Fast Fourier transforms (FFTs) were then conducted on these average waveforms to bring the data into frequency space, resulting in amplitude responses (μV) at discrete frequencies (for a 10s stimulus the FFT has a resolution of 0.1Hz) between 0.1 and 100Hz.

The amplitude response at each frequency at each electrode site was converted into a measure of signal-to-noise ratio (SNR) by dividing the amplitude at the frequency of interest by the average amplitude of the surrounding 12 frequency bins. The choice of channels from which data was extracted in each of the three SSVEP experiments was based on group topographies showing peak responses for all stimuli occurring there and observations from previous literature.

To determine whether SNRs were significantly above background noise (SNR=1), a series of one-sample t-tests were conducted separately for each set of component SNRs, difference intermodulation SNRs, and sum intermodulation SNRs. Each series of t-tests were corrected using the ranked Bonferroni-Holm method to control for Type 1 errors (Holm, 1979). The extent to which a stimulus type (e.g. grating/plaid) predicted response SNRs was examined using linear mixed-effects modelling. The analytical model was generated using the `mixed` function of the `Afex` package in \mathbb{R} (Singmann et al., 2016). In Chapters 4 and 6, Stimulus Type was the only predictor with random slopes as a function of 'participant' using a maximal random effects structure (as recommended by Barr, Levy, Scheepers, & Tily, 2013). In Chapter 5, Stimulus Type and Contrasts were set as predictors, again with random slopes as a function of 'participant' using a maximal random effects structure. This model was

applied to both component and intermodulation SNRs. The `lsmeans` function in R for examining pairwise comparisons from linear mixed effects model structures was used when a significant main effect of pattern was found, and comparisons were corrected using the Tukey HSD method for multiple comparisons (Russell, 2016).

Chapter 3: Visual evoked responses to coherent and non-coherent plaid stimuli

Rationale

The most natural way to try and understand the nonlinear responses to plaids and gratings is simply to measure the Visual Evoked Potential (VEP). VEPs are simple waveforms, time-locked to the onset of the stimulus presentation comprising multiple response components, each of which may indicate a different neural substrate.

While many have studied the neural responses to low-level image features using VEPs (e.g. Butler et al., 1987; Jeffreys & Axford, 1972a; Kelly et al., 2008; J. G. May & Lovegrove, 1987; Shigeto, Tobimatsu, Yamamoto, Kobayashi, & Kato, 1998), only a few have examined how processes such as normalisation (Chen, Yu, Zhu, Peng, & Fang, 2016; Miller et al., 2015) are represented within the early VEP components. None have directly assessed the nonlinearities involved summation of individual grating components involved in plaid processing.

Other methodologies have been used to try and understand nonlinear summation. McDonald, Mannion and Clifford (2012) used fMRI to investigate the summation of orthogonally oriented grating components that formed a coherent plaid. They found that BOLD responses to the plaid were greater than to the grating components alone, but less than would be predicted by linear summation of individual component responses. This is in agreement with animal studies (Allison, Smith, & Bonds, 2001; Brouwer & Heeger, 2011a; DeAngelis, Robson, Ohzawa, & Freeman, 1992; Grossberg & Hong, 2006; Heeger, 1992) and psychophysical studies (Georgeson & Shackleton, 1994) of plaid processing. However, these have not compared responses to coherent and non-coherent plaids, instead framing their results with the context of normalisation.

We aimed to investigate the nonlinearities that result from combinations of coherent and non-coherent gratings using EEG, to try and differentiate the contribution from suppressive, versus expansive, nonlinearities. Here, the first reliable VEP component (C1), and the first positive (P1) deflection that follows the C1 were measured in response to centrally-presented gratings and their combination into both

coherent and non-coherent plaids. Based on observations of nonlinear spatial summation being represented in these early VEP components (Chen et al., 2016; Miller et al., 2015), and because cross-orientation suppression is largely un-tuned for spatial frequency (DeAngelis et al., 1994; Petrov et al., 2005), it was expected that if an additional nonlinearity is required for processing coherent plaids but not non-coherent plaids, this additional nonlinearity would be present in early VEP components.

Methods

Participants. Twenty-one participants were recruited. Upon completion of testing, two participants reported difficulty with the task due to undisclosed visual impairments. As a result, their data was removed from analyses. All other participants had normal or corrected-to-normal vision. Informed consent to participate in the study was given from all participants. The ethics board at the School of Psychology, University of Nottingham, granted ethical approval.

Gamma correction and materials. A computer-controlled monitor (Iiyama ProLite X2472HD, Iiyama, Hoofddorp, The Netherlands) with a screen resolution of 1920 x 1080 pixels, mean luminance of 133.17cd/m² and a refresh rate of 60 Hz was used for stimulus presentation. Gamma correction was carried out using a photometer (ColorCal MKII, Cambridge Research Systems, Kent, UK). No chin rest was used, so as to maximise participant comfort, albeit reducing the precision with which the stimulus size/position could be calculated. Participants sat at a screen distance of ~100-120cm (measured for each participant and used to adjust screen properties at the beginning of the session). The PsychoPy stimulus generation library (Peirce, 2007a) was used for stimulus presentation and collecting participant responses to a simple detection task. A button box (Cedrus RB-x30 Response Pad, Cedrus Corporation, San Pedro, CA, USA) was used for participants to make these responses.

Stimuli. Gratings were all based around two sinusoidal grating patterns ('A':2.5cpd and 'B':5.5cpd). They were sometimes combined with another spatial-frequency-

matched grating to form two coherent plaids (i.e. 'AA' and 'BB') or combined with the other grating to form a non-coherent plaid (i.e. 'AB'). While components were randomly oriented on each trial, they always overlapped at a plaid angle of 90° . Stimuli were presented at Michelson contrasts of 20%, 40% and 80%, resulting in 15 stimulus conditions including the blank (no stimulus) condition.

Stimuli were presented within a 10deg raised cosine annular envelope. The blank inner portion of the annulus subtended 0.6deg of visual angle leaving the outer portion with a radius of 4.7deg. A red fixation dot was placed in the centre of the screen and subtended 0.2deg. A green dot of 0.3deg that appeared in a random location at the half-radius of the outer annulus, was employed in a simple task to maintain attention to the stimulation area.

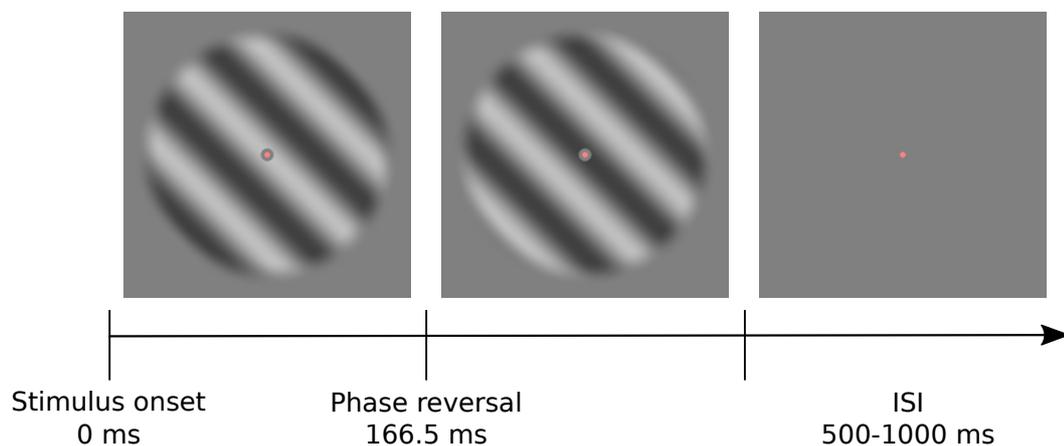


Figure 3.1. Example of an experimental trial for a grating stimulus.

Experimental Procedure. When seated in the recording booth, participants were instructed to maintain fixation on the central red fixation point. They were then instructed to respond on the button box as quickly as possible to the presence of the green task dot (without looking towards it) and that this would appear in a pseudo-random order throughout the experiment. They were then given a brief presentation of the experiment to ensure they understood what they were being asked to do. Following this, they were presented with a blank fixation screen consisting of a grey background and the red fixation dot until they indicated that they were ready to begin.

Stimuli were presented for 333ms with an inter-stimulus-interval of between 500ms and 1000ms (**Figure 3.1**). To prevent after-images, stimuli were phase-reversed

halfway through presentation (at 166.5ms). Each stimulus type was randomly interleaved and repeated 20 times in a run, with 5 runs in total. Participants were given short breaks between runs. Upon completion, participants were thanked, debriefed and given the opportunity to ask any questions.

Analytical Procedure. Data were band-pass filtered between 1Hz and 100Hz to remove slow drifts, and epoched between -100ms and 500ms from stimulus onset. Minimal artefact rejection was used on a participant-by-participant basis. Peak-to-peak fluctuations greater than 100 μ V were identified on a channel-by-channel basis and corrected using the average of the nearest neighbouring channels. If this occurred at more than 16 channels, the entire trial was removed from analysis. This process removed an average of 8.55% (SEM: \pm 8.54) of trials across participants. This was calculated across participants, rather than the total of number of runs, to account for inter-individual variability. Grand average global field power (GFP) was used to determine the approximate peak latency of the C1 and P1 components and generate a window around this for analysis (see General Methods for details). The latency of particular components can be variable between participant and even between conditions within a single participant. To measure the amplitude and latency of each component we determined a relatively wide window, the median location of the component in GFP \pm 30ms for that condition and across participants. A window of 60-120ms was used for the C1 component and a window of 107-167ms for the P1 component. The selection of these windows allowed us to avoid the impact of transient phase reversal, which occurred as a result of the spatial phase reversal of stimuli at 166ms, on the time course of the response.

Grand average topographies (see Figures 3.2 to 3.16 in Appendix One) with discrete 10ms time-windows between 0ms and 170ms were then used to determine the electrodes from which to extract data. On a first pass of this procedure, four electrodes (AF1, AF4, TPP7h and PPO5h) were observed as being particularly noisy and were removed from analysis. The GFPs and topographies were then re-calculated. Using this procedure, electrode Oz was deemed the most appropriate site for both C1 and P1 peak responses to our stimuli.

Data were averaged by condition for each participant. Peak component responses for each condition were identified in Python using Scipy's `signal.argrelextrema` package. A criterion window of 14ms (7 points on either side) was set to identify points in the response function greater (positive or negative) in magnitude than their surround. For the C1 component, the most negative point identified within the C1 window of 60-120ms was picked, and for the P1 component the most positive point during its 107-167ms time window was picked. If no peak was identified using the above criteria, i.e. in the case where the VEP was very flat and no reliable peak could be identified, the average amplitude during the time window was stored instead.

To assess differences in signal summation between coherent and non-coherent plaids, a ratio of amplitude response 'linearity' was calculated as:

$$LI = \frac{pResp}{c1Resp + c2Resp}$$

where *LI* is 'linearity index', *pResp* is the absolute response to a plaid stimulus and *c1Resp* and *c2Resp* are the separate absolute responses to the grating components forming that plaid. For example, coherent plaid AA at 40% contrast was formed simply by the orthogonal superimposition of two 20% contrast grating A's. By adding the response of grating A at 20% contrast to itself, a hypothetical 'linear' response for plaid AA at 40% contrast can be generated. An *LI* of 1 would reflect linear summation; the magnitude of response for that plaid was equivalent to linearly summing the response to each of the components forming that plaid alone. On the other hand, an *LI* of 0.5 would indicate a response equal to the average of the component responses.

The extent to which grating/plaid pattern predicted peak amplitudes was examined using linear mixed-effects modelling. The analytical models were generated using the `mixed` function of the `Afex` package in R (Singmann et al., 2016). Grating/plaid pattern (A, B, AA, BB and AB) and contrast (20, 40 and 80%) were used as predictors with random slopes as a function of 'Participant' using a maximal random effects structure (as recommended by Barr et al., 2013). These models were applied to peak C1 and peak P1 amplitudes, separately. Another model was then generated for fitting to *LI* scores. This was again done separately for C1 and P1, and included

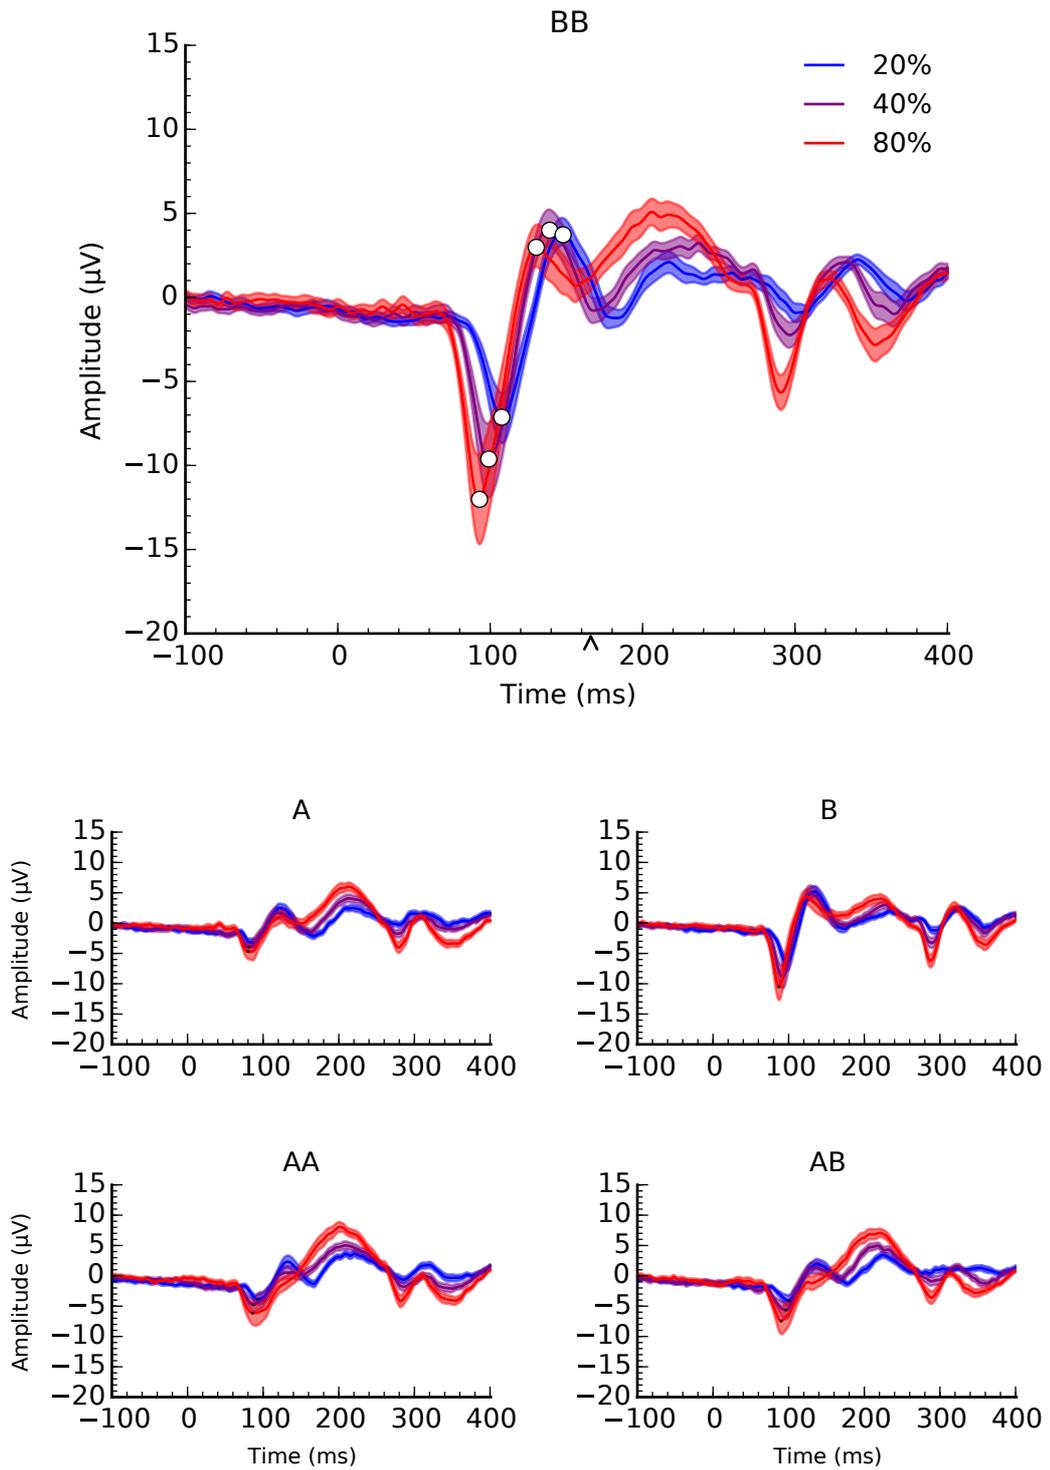


Figure 3.17. Grand average waveforms for each condition. Top: Coherent plaid BB provided as an example because very clear peak responses were observed for this condition. Bottom: all other conditions (A, B, AA and AB). The blue waveform represents the grand average response at 20% Michelson contrast, purple to 40% and red to 80%. In the top plot, white circles mark the peak of the C1 (first negative) and P1 (first positive) responses for each contrast presentation of coherent plaid BB. The solid coloured line represents the mean response, and the transparent shading around of the same colour represents the SEM. The small arrow marked on the x-axis line at 166ms in the top plot represents when stimulus phase-reversal occurred.

pattern (AA, BB and AB) and contrast (40, 80) as predictors with random slopes as a function of 'Participant', as in the previous models. The `lsmeans` function in R for examining pairwise comparisons from linear mixed effects model structures was used when a significant main effect of pattern or contrast was found, and comparisons were corrected using the Tukey HSD method for multiple comparisons (Russell, 2016). Correction was made assuming the full possible set of 105 comparisons even though many comparisons were not actually of interest *a priori* and have not been presented or examined.

Results

C1 Responses

First, the C1 peak amplitudes were analysed to test for effects of stimulus contrast and pattern (see Figure 3.18). Contrast significantly increased C1 response magnitude for all stimuli except grating A. The higher-spatial-frequency grating and coherent plaid (B and BB) resulted in generally larger C1 peak amplitudes. At equivalent contrast levels, plaid responses were similar in magnitude to their component responses. For

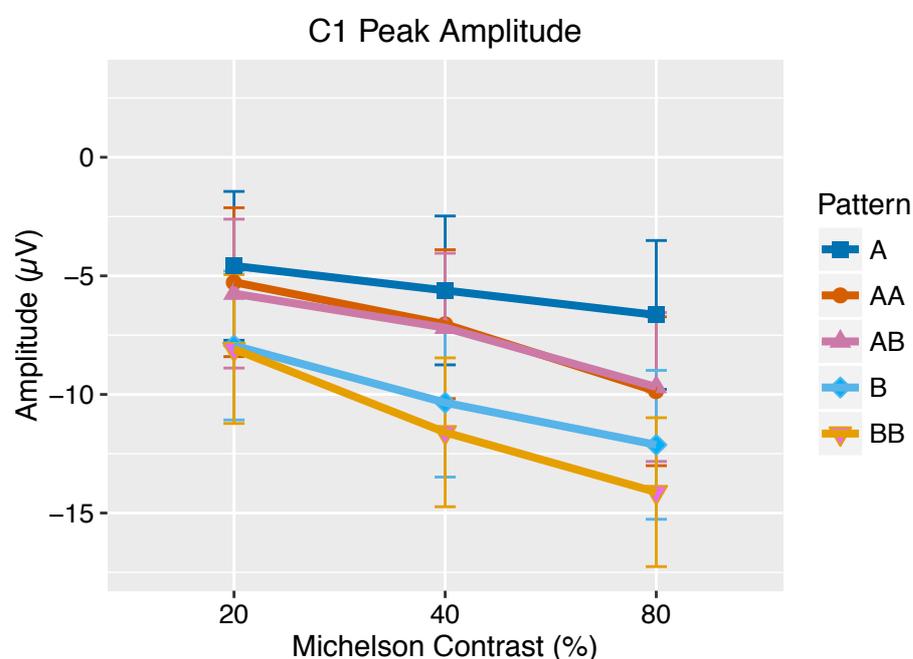


Figure 3.18. C1 peak amplitudes. Each point represents the mean across participants for that condition and error bars represent 95% CIs.

example, the response to plaid BB at 80% contrast was similar in magnitude to that of grating B at 80% contrast.

A significant main effect of Contrast ($F(2,252)= 45.03, p<.001$) and Pattern ($F(4, 252)= 32.32, p<.001$) was found. The interaction between Contrast and Pattern was non-significant ($F(8,252)= 1.21, p>.05$). Post-hoc tests showed that peak C1 amplitudes decreased as a function of Contrast for all patterns except grating A, which only showed a very slight (non-significant) decrease in response amplitude. The 20% contrast response amplitude was significantly smaller than the 80% contrast response amplitude (at least $p<.05$, at most $p<.001$). Only for coherent plaid BB was the peak response amplitude at 20% contrast significantly less than that at 40% contrast ($p<.05$). At 40% and 80% contrast, peak C1 response amplitudes were smaller for grating A than grating B ($p<.001$), and for coherent plaids AA and AB than coherent plaid BB ($p<.001$).

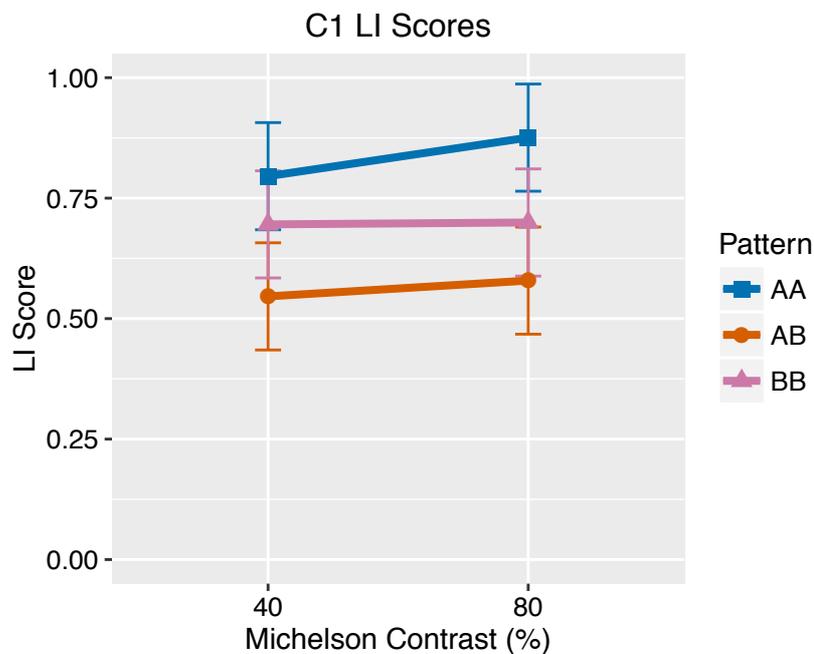


Figure 3.19. C1 *LI* scores. Each point represents the group average, and error bars 95% CIs.

Next, the question of whether there was a difference in nonlinear spatial summation between coherent and non-coherent plaids was addressed. *LI* scores were calculated for each plaid. These were generated by dividing the absolute peak amplitude measured at 40% or 80% contrast for a plaid pattern, and dividing this by the sum of the unsigned response to the grating components forming that plaid (at 20% and 40% contrast, respectively). C1 *LI* scores (Figure 3.19) were less than linear and the average score was smaller for non-coherent plaid AB than for both coherent

plaids. One participant was excluded from analysis because of an extreme outlying ratio value (>4). A significant main effect of Pattern ($F(2,85)= 13.60, p<.001$) but not Contrast ($F(1,85)= 0.83, p=.37$) was found. The interaction between the two was also non-significant ($F(2,85)= 0.27, p=.76$). Post-hoc tests revealed that the LI score for pattern AA at 40% contrast was larger than AB at 40% contrast ($p<.05$), and the 80% contrast LI score for AA was larger than AB ($p<.01$). In summary, other than for grating A, clear contrast response functions were observed for C1 peak amplitudes; they became larger in magnitude with increasing contrast. C1 peak amplitudes were also larger for the higher spatial frequency patterns (grating B and coherent plaid BB) than all other patterns.

P1 Responses

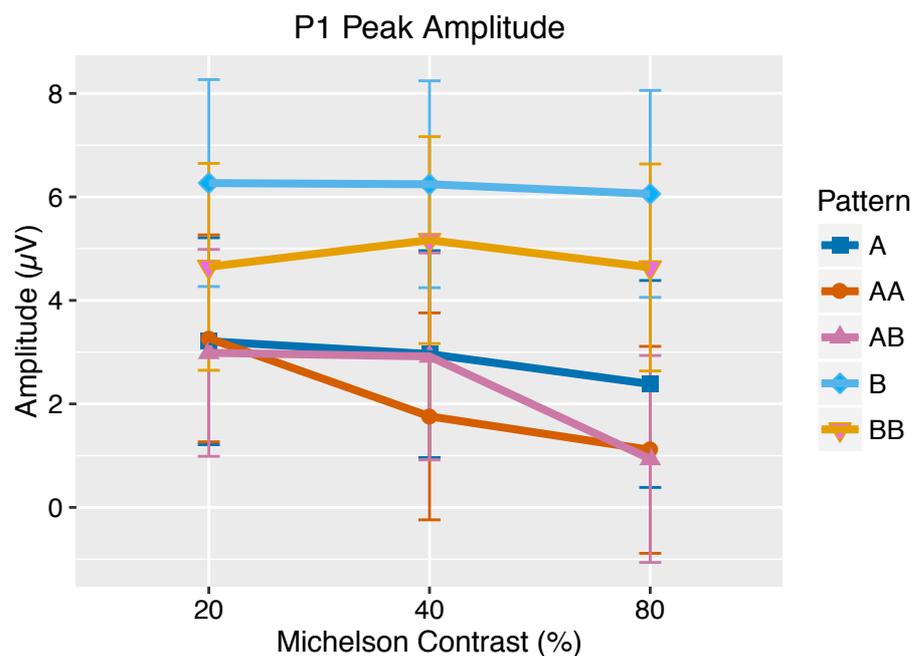


Figure 3.20. P1 peak amplitudes. Each point represents the mean across participants for that condition and error bars represent 95% CIs.

For the P1 peak component, the higher spatial frequency grating B and coherent plaid BB resulted in significantly larger P1 peak amplitudes compared to all other conditions. A significant main effect of Contrast ($F(2,252)= 6.37, p<.01$) and Pattern ($F(2,252)= 41.34, p<.001$) was found for P1 peak amplitudes (Figure 3.20). The

interaction between Contrast and Pattern was non-significant ($F(8,252)= 1.46$, $p=0.17$). Contrary to the main effect of contrast, no post-hoc tests for any one stimulus pattern comparing across contrast levels revealed a significant difference. Larger P1 amplitudes were observed for grating B than grating A at all contrast levels ($p<.01$ at 20% and $p<.001$ at 40% and 80%). At 40% contrast, a smaller P1 peak amplitude was observed for coherent plaid AA than plaid BB ($p<.001$), and again at 80% contrast ($p<.001$), along with a smaller P1 for non-coherent plaid AB than for coherent plaid BB ($p<.001$).

As for component C1, P1 *LI* scores (Figure 3.21) were less than linear and the average score was smaller for non-coherent plaid AB than for both coherent plaids.

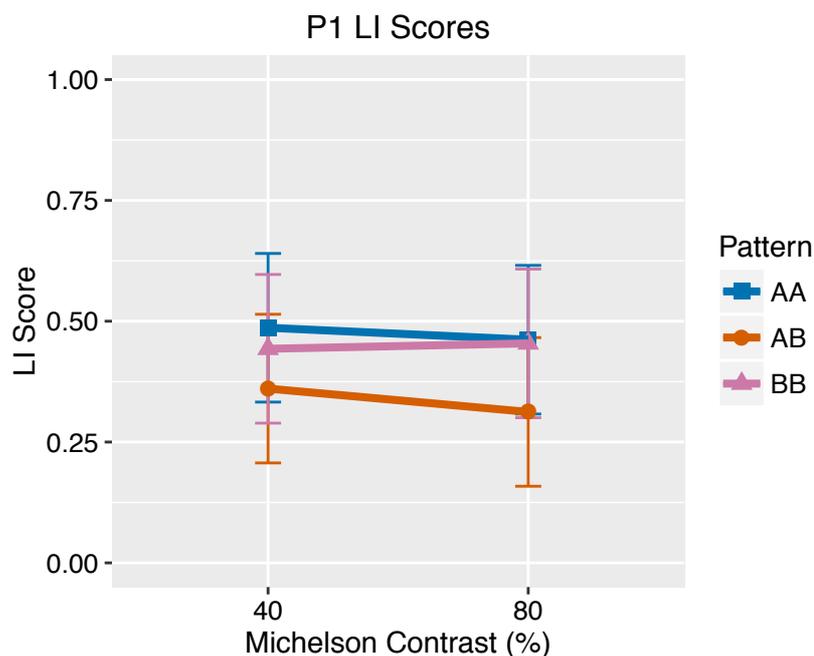


Figure 3.21. P1 *LI* scores. Each point represents the group average, and error bars 95% CIs.

Two participants were excluded from analysis because of extreme outlying values (>6 and >3 , respectively). The main effects of Pattern and Contrast were both non-significant ($F(2,80)= 1.81$, $p=.17$; $F(1,80)= .11$, $p=.74$), as was the interaction between them ($F(2,80)= 0.08$, $p=.93$).

Discussion

The peak amplitudes of the first major negative (C1) and positive (P1) VEP components were measured in response to presentations of sinusoidal grating patterns presented alone, and when combined to form plaid patterns. This was to assess the nonlinear combination of signals to the gratings when forming different kinds of plaid pattern. These early VEP components have been shown to indicate differences in nonlinear spatial summation (Chen et al., 2016; Miller et al., 2015). Here, coherent plaid AA resulted in significantly larger C1 *LI* scores than non-coherent plaid AB, but the same was not true for coherent plaid BB compared to the non-coherent plaid.

Amplitude effects

An increase in response amplitude with increasing stimulus contrast was observed for the C1 component, particularly for the 5.5 cpd stimuli (grating B and plaid BB). Based on the latency and the high-pass spatial frequency characteristic of the C1 component in response to centrally placed stimuli vs. the bandpass C1 response to peripherally presented stimuli, Hansen et al., (2016) suggested that the former may represent the C2 component response, not the C1. We, too, measured a central C1 response that was high-pass for spatial frequency, though how the response might have differed if stimuli were presented peripherally was not the focus of the study and is therefore not known. The P1 peak amplitude did not vary with contrast for any stimulus, and like the C1 component was larger in magnitude for higher spatial frequency than lower spatial frequency stimuli. The spatial-frequency tuning, but contrast invariance, suggests a generation site beyond striate cortex, and several studies have suggested that the P1 component is extra-striate in origin (Di Russo et al., 2002; Martinez et al., 1999).

LI scores all less than linear

It's not so surprising that the response to a plaid is less than a linear summation of responses to components alone. Presumably due to XOS, single cell studies and human investigations using fMRI (Bonds, 1989; Busse, Wade, & Carandini, 2009;

DeAngelis et al., 1992; McDonald et al., 2012) have also found that the summed response to a plaid is less than a linear summation of responses to its components. Plaids also have a lower perceived contrast relative to sinusoidal gratings presented at the same contrast (Georgeson & Shackleton, 1994), though McDonald, Mannion and Clifford (2012) and the present study found larger responses for plaids compared to one of their component gratings of equivalent Michelson contrast alone. This was specific to coherent plaids in the current instance, and might reflect a difference in processing for coherent plaids vs. non-coherent plaids. That assertion is difficult to make, though, because the same comparison of non-coherent plaid vs. component responses cannot be made; the different components (A and B) forming the non-coherent plaid generate responses of different magnitudes.

Gheorghiu and Kingdom (2009) suggested that a multiplication of V1 outputs was involved in the processing of simple contour stimuli. This is an appealing solution because of its simplicity; it would allow a mechanism detecting contours to respond only to a combination of contour components but neither of the components alone. However, this does not explain the less-than-linear summation observed for plaid stimuli – a greater than linear response would be expected from multiplication. Perceptual studies of plaids and gratings have found that plaids are perceived at lower contrast compared to gratings of the same contrast (Georgeson & Shackleton, 1994). Though it would have been useful to have measured the perceived contrast of our stimuli, our finding that plaids produced larger responses to gratings at all contrast levels suggests that contrast alone is not enough to determine the magnitude of the response to the stimulus, and that an extra factor, such as a plaid-response, resulted in the difference between coherent plaid and grating responses. This supports Peirce's suggestion (2007b, 2011) that, in combination with component suppression, a logical AND-gated mechanism in mid-level vision could provide greater stimulation than any one component could alone. This additional response could be what McDonald, Mannion and Clifford (2012) observed as a 'release from suppression' to plaid stimuli. They only used coherent plaid stimuli so had no comparison to non-coherent plaids, whereas we have shown that this 'release' only happens for coherent plaids.

LI scores only different for AA, not BB, compared to AB

LI scores were calculated to compare the linearity of summation between the different plaid patterns. XOS is largely un-tuned for spatial frequency (DeAngelis et al., 1994; Petrov et al., 2005), so if similar amounts of XOS occurred for each plaid and no extra summation was taking place, it would be expected that no difference would be found in LI score between any of the plaid conditions. A significant difference was only found between coherent plaid AA and non-coherent plaid AB; in all other cases the data fell in the expected direction but differences were not significant. It is difficult therefore to conclude that there was a difference in summation between coherent and non-coherent plaids.

The perception of plaids has been shown, by several methods, to depend on matched spatial frequencies. When spatial frequencies differ in the two gratings being combined, observers perceive a pair of semi-transparent gratings sliding past each other (Adelson & Movshon, 1982) whereas a plaid with matched spatial frequency components appears as a single coherent checkerboard pattern with a single direction of motion. Similarly, selective adaptation to plaids decreases when components are unmatched (Hancock et al., 2010) and the pop-out effects of visual search disappears when plaid targets have unmatched components (Nam et al., 2009). Despite this, I found limited evidence for the involvement of extra nonlinearities in processing coherent plaids relative to non-coherent plaids in the current experiment. Perhaps VEPs are not sensitive enough to differentiate between the different nonlinear interactions that give rise to selective plaid processing.

Chapter 4: Measuring nonlinear signal combinations using intermodulation responses

Rationale

The previous chapter examined nonlinear signal summation by measuring transient VEP responses. While there were observable patterns within the data in the direction that we hypothesised, these were not very clear. The transient EEG approach brings with it several inherent limitations. The first is the somewhat arbitrary application of a time window within which to pull out peak responses. This is done based largely on the previous literature to limit any experimenter bias in analysing results, but brings with it the danger that genuine peak responses are missed, confounding the results. Secondly, and critically for the purpose of this thesis, the transient EEG approach measures mass evoked potentials. This made it extremely difficult to pick apart different nonlinearities within the response, and could explain why we observed only very subtle effects in the hypothesised direction.

Measuring amplitude responses at intermodulation terms to different kinds of plaid pattern might present a more effective way of distinguishing between response nonlinearities. Steady state VEP's are more objective in that no time window of interest has to be applied to the analysis (save for epoching) – a response is either generated at a frequency of interest or it is not. Further, the method allows you to measure direct responses to the components of a compound stimulus at f_1 and f_2 and their harmonics at $2f_1$ and $2f_2$, as well as responses at intermodulation terms like $f_2 \pm f_1$. This allows you to examine nonlinearities between components directly at component frequencies as well as at intermodulation terms.

One possibility is that mechanisms for plaids exist that contribute to intermodulation responses. If this is the case, then I would expect such responses to differ according to plaid coherence. Conversely if the intermodulation response is driven by a normalisation pool for cross-orientation suppression (XOS), which is thought to be relatively un-tuned to spatial frequency (DeAngelis et al., 1994; Petrov et al., 2005), then the intermodulation response should be similar for different forms of plaid. An exception to this might be if the intermodulation response is also driven

by surround suppression, which unlike XOS is tightly tuned for spatial frequency (Petrov et al., 2005). However, surround suppression is also tightly tuned to the same orientation as the excitatory region of neurons, so I would expect contributions from this to be weak for orthogonally arranged grating components.

Methods

Participants. Fifteen participants (7 females, 8 males) with normal or corrected-to-normal vision gave informed consent to participate in the study. The ethics board at the School of Psychology, University of Nottingham, granted ethical approval. The work was conducted in accordance with the 2008 version of the Declaration of Helsinki.

Stimuli and Experimental Procedure. Stimuli comprised of two sinusoidal gratings (denoted as 'A' and 'B') and various combinations thereof. Grating A had a spatial frequency of 1cpd and grating B a spatial frequency of 3cpd. These could then be combined with a second, spatially orthogonal, grating to form plaid patterns that were either coherent ('AA', 'BB') or non-coherent ('AB', 'BA') as shown in Figure 1. On each trial the overall orientation of the stimulus, either grating or plaid, was randomly assigned but the orthogonal configuration of the grating components that formed a plaid was maintained.

Components were presented within a 7.5° diameter circular window with a raised cosine edge profile (width= 1.5°). In the center of the stimulus a further circle (diameter 0.4°) was placed with the same mid-grey color as in the background, also with a raised cosine edge profile (0.08°). This was done to accentuate the fixation dot, a red circle subtending 0.175° in the center of the screen. A further green dot (0.25° diameter) appeared occasionally and briefly, located at a radius of 3.75° from the centrally located fixation point but a random radial angle. This was used as part of the attentional control task that participants were asked to perform.

When seated in the recording booth, participants were instructed to maintain fixation on the red fixation point. They conducted the green dot attentional task, independent of the stimulus presentation, to ensure that attention was constant. The

green dot appeared at a random time within any trial number divisible by 10 or 6 (6, 10, 12, 18 etc.). This resulted in rare events, unpredictable to the participants, that required their attention to detect, preventing them from assigning greater or lesser attention to the actual stimulus. Although diverting attention away from the stimulus might decrease our chances of measuring strong signals Appelbaum et al., (2009) we considered it more important to keep attention constant. It remains very possible, however, that we underestimate the size of responses in our data because of this. The percentage of task occurrences across each condition relative to the grand total number of occurrences shows no bias towards any particular condition: A1-13.33%; A2-13.07%; B1-14.67%; B2-10.67%; A1A2-12.27%; B1B2-13.33%; A1B2-9.60%; B1A2-13.07%). It is therefore unlikely that the dot task was a factor in determining any systematic effects observed between conditions. Participants responded to 76.34% (SEM: 2.22) of dot occurrences, indicating that the task was neither too easy nor too difficult.

A brief presentation of the experiment was provided to ensure the participants understood what they were being asked to do. Following this, they were presented with a blank fixation screen consisting of a grey background and red fixation dot until they indicated that they were ready to begin.

In the literature review and General Methods, I covered the need to independently modulate plaid components at slightly different frequencies to generate intermodulation responses. Here, the frequencies chosen were 2.3Hz and 3.75Hz on the basis that they allowed the response frequencies to be temporally incommensurate; the fundamental frequencies (2.3Hz, 3.75Hz), their harmonics ($2f_1$: 4.6Hz, $2f_2$: 7.5Hz) and the intermodulation terms (f_2-f_1 : 1.45Hz, f_1+f_2 : 6.05Hz) and harmonics ($2f_2-2f_1$: 2.9Hz, $2f_1+2f_2$: 12.1Hz) would not overlap. Further, we aimed to avoid overlap between the typical alpha band (8-12Hz) and the sum intermodulation response frequency (f_1+f_2 : 6.05Hz) as in Boremanse, Norcia and Rossion (Boremanse et al., 2013) to boost signal-to-noise. The contrast of any one component varied sinusoidally in time between 0 and 50% of the maximal contrast of the monitor (which had a maximal Michelson contrast of 0.99).

For the remainder this thesis, in the context of intermodulation stimuli only, '1' denotes a grating component that was flickered at 2.3Hz and '2' denotes a

component flickering at 3.75Hz. For example, grating A flickered at 2.3Hz will now be referred to as A1, while a coherent plaid with grating A components will be referred to as A1A2, where one component was flickered at 2.3Hz and the other at 3.75Hz. The full set of components and compounds resulted in 8 stimuli: four grating components (A1, A2, B1, and B2), two coherent plaids (A1A2 and B1B2) and two non-coherent plaids (A1B2, and B1A2). A trial consisted of an 11 second presentation of a flickering grating or two simultaneously flickering, superimposed grating components, followed by a 7-9 second inter-stimulus-interval. Each of the 8 stimuli was presented 3 times in a run (lasting 8 minutes), and participants completed 5 runs, with short breaks between each. Upon completion, participants were thanked, debriefed and given the opportunity to ask any questions.

Analytical Procedure. The amplitude response at each frequency at each electrode site was converted into a measure of signal-to-noise ratio (SNR) by dividing the amplitude at the frequency of interest by the average amplitude of the surrounding 12 frequency bins. The data used in the following analyses were taken from electrode Oz. This was based on group topographies showing peak responses for all stimuli occurring there, and is consistent with measurements in the vicinity of primary visual cortex. We also conducted the analysis using a cluster of electrodes around Oz, but this made no difference to the conclusions from the analyses.

Initially, the frequencies of interest for analysis were at f_1 (2.3Hz), f_2 (3.75Hz), f_2-f_1 (1.45Hz), and f_1+f_2 (6.05Hz). Upon inspection of the data, responses at the $2f_2-2f_1$ (2.9Hz) and $2f_1+2f_2$ (12.1Hz) frequencies warranted further analysis. To determine whether SNRs were significantly above background noise (SNR=1), a series of one-sample t-tests were conducted separately for each set of component SNRs (i.e. A1, A1A2 and A1B2; A2, A1A2 and B1A2; B1, B1B2, B1A2; B2, B1B2 and A1B2), difference intermodulation SNRs (at f_2-f_1 and $2f_2-2f_1$: A1, A2, B1, B2, A1A2 and A1B2, B1B2 and B1A2), and sum intermodulation SNRs (at f_1+f_2 and $2f_1+2f_2$: A1, A2, B1, B2, A1A2, A1B2, B1B2 and B1A2). Each series of t-tests were corrected using the ranked Bonferroni-Holm method to control for Type 1 errors (Holm, 1979).The

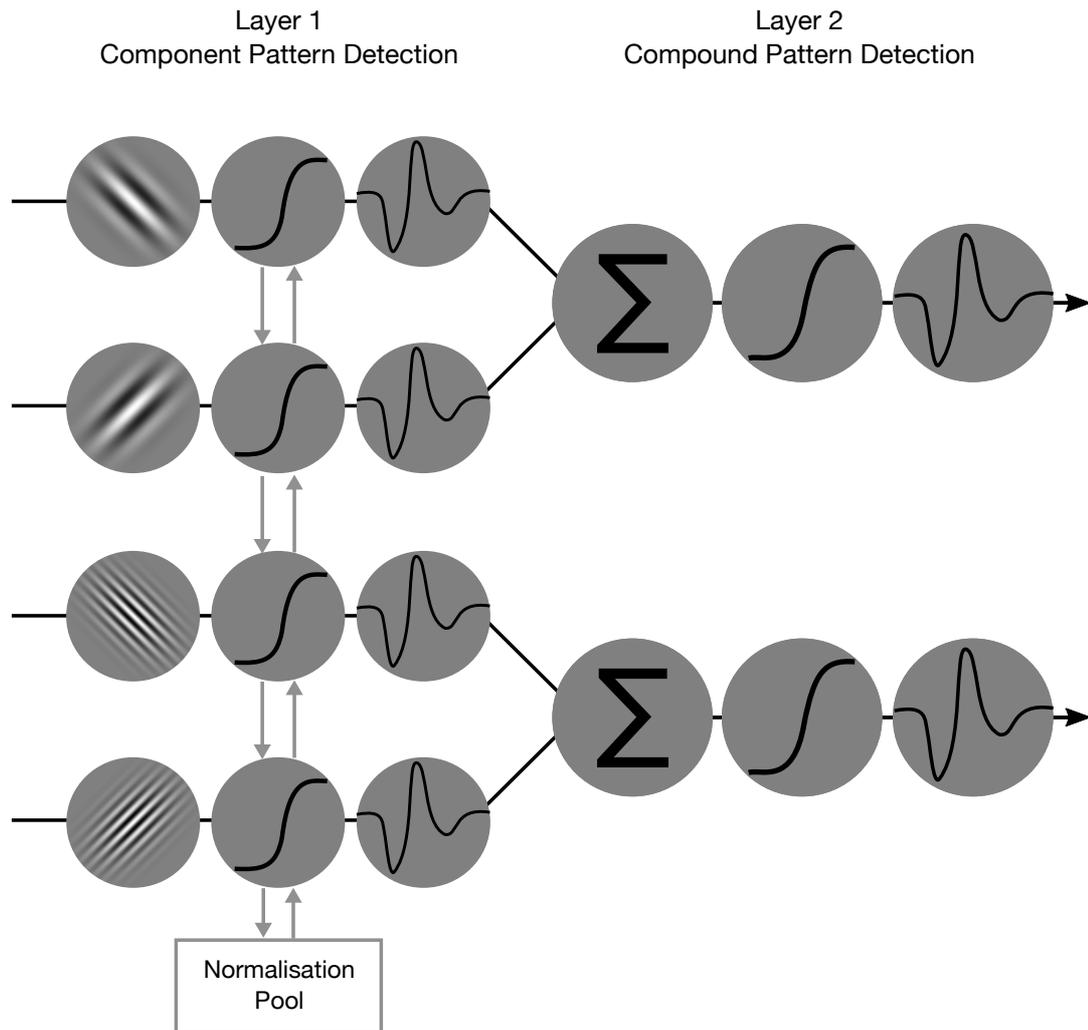


Figure 4.1. Model schematic displaying how contrast input was transformed within each layer. Though the Layer 1 channels were un-tuned in this study, the Gabor patches differing in orientation and spatial frequency are visual representations used to indicate how the channels, within the framework that we propose, might be tuned. These are followed by a static sigmoidal nonlinearity and then a temporal filter (if appropriate to the simulation). The normalisation pool exerted its influence (for XOS) as part of the static nonlinearity. A similar configuration (without normalisation) was used to transform input into Layer 2 channels.

extent to which grating/plaid pattern predicted response SNRs was examined using linear mixed-effects modelling.

Modelling. A two-layer network model was generated based on Peirce (2007b, 2011; see **Figure 4.1**) in which static sigmoidal nonlinearities were applied to the outputs of each channel (as expected by any model of V1 outputs) and summed by a “Layer 2” mechanism, which also has a sigmoidal nonlinearity. Such a model can be used to investigate the relationship between XOS, nonlinear additive summation and the generation of intermodulation responses to plaids. Contrast input to the model was

generated in the same way that contrast was modulated in the stimuli presented to participants:

$$contr = C \cdot (\sin(t \cdot f \cdot 2\pi) \cdot 0.5 + 0.5),$$

where t represents a point in time between 0 and 11 seconds (in steps of 0.01s), f the temporal frequency for that component and C the maximum Michelson contrast (set to 0.5). The multiplication and addition by 0.5 scaled the minimum and maximum contrast to be 0 and 1 initially and the further multiplication by C reduced the maximum contrast to 0.5.

Layer 1 involved four channels with static nonlinearities (Naka & Rushton, 1966) with an extra term in their denominators to account for XOS (Carandini & Heeger, 2012) and is here assumed to represent grating component responses. The structure of a Layer 1 channel was as follows:

$$compResp = rMax \cdot \left(\frac{I_{at}^n}{C_{50} + I_{at}^n + NP} \right),$$

where $rMax$ is the maximum response of the channel, I_{at} the input contrast of the stimulus component being encoded at that channel at time point t , n an exponent and C_{50} the semi-saturation point. NP , used as the normalisation pool, is the sum of three extra terms corresponding to the input to the other three Layer 1 channels. The value of each of these terms was determined in a similar fashion to the I_{at} term:

$$NP = \sum_{i=3} I_{it}^n$$

Two channels corresponded to ‘detectors’ for components A1 and A2, and their output was summed (CSa) before being passed to Layer 2. The same was done for the other two channels, (‘B1’ and ‘B2’) and their sum referred to as CSb .

The second layer can be thought of as two additional “channels” with a similar static nonlinearity (without the additional XOS term on the denominator), which would respond selectively to the presence of a plaid:

$$plaidResp = rMax \cdot \left(\frac{CS_t^n}{C_{50} + CS_t^n} \right),$$

where CS (either a or b) represents the linearly summed component responses (as described above), n an exponent and C_{50} the semi-saturation point (the same used in Layer 1). For both channel layers, $rMax$ was held constant at 1, n at 2 and C_{50} at 0.2. These were reasonable values to set parameters with, given similar values reported in

studies that aimed to model similar phenomena (D. H. Baker & Wade, 2017; Sit, Chen, Geisler, Miikkulainen, & Seidemann, 2009; Tsai et al., 2012). The overall output of the model was then calculated as the linearly summed response of *compResp* and *plaidResp*, simulating the population response measured with EEG.

It has been suggested that the temporal processing of signals by the mechanisms generating intermodulation responses may be key to their almost-always asymmetric response patterns (e.g. Alp et al., 2016; Boremanse et al., 2013). Thus, we

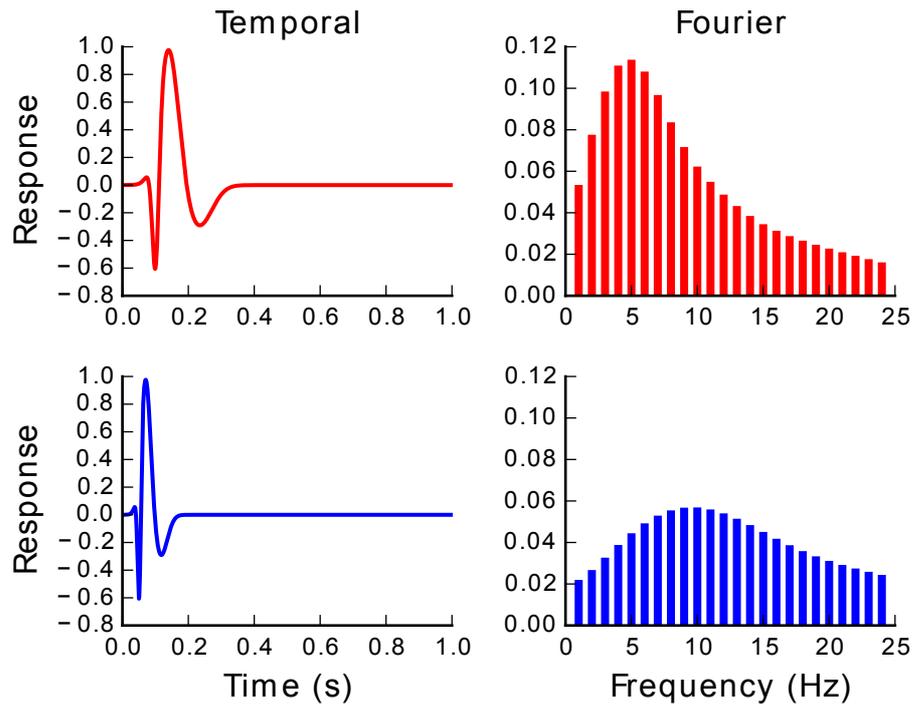


Figure 4.2. Temporal and Fourier responses of the bandpass and higher-pass temporal filters. The left figures represent temporal responses and the right Fourier responses. Red indicates responses of the bandpass filter and blue the higher-pass filter.

wanted to evaluate the importance of various temporal filters, modeling the neural impulse response functions (NIRs) on the performance of the model. Further, it has been suggested that later mechanisms in the visual pathway may differ in their temporal filtering properties compared to earlier mechanisms. We therefore conducted simulations where NIRs were generated and used as temporal filters on each channel's output.

We tested two variants of filter (**Figure 4.2**). The first variant was generated by summing three Gaussian distributions. It had a bandpass Fourier response peaking at ~5Hz with a slow decay towards 20Hz. The second form of linear filter was generated simply by halving the peaks and widths of the band-pass filter described above,

resulting in a high(er)-pass Fourier response peaking at $\sim 10\text{Hz}$. These filters were combined in three ways; a) the bandpass filter for both Layer 1 and Layer 2, b) the bandpass filter for Layer 1 and higher frequency filter for Layer 2, and c) no temporal filtering at either level. This allowed us to assess the importance of the temporal integration functions in the resulting responses, and whether there is any evidence for these differing in early- and late-stages of the model.

Following this, noise was added to the output of the model. A vector of 10,000 values that was the average of 1000 iterations of random sampling from a normal distribution was filtered with a 3rd order bandpass Butterworth filter between 0.1 and 100Hz. This resulted in a $1/f$ decay like that observed in the EEG data. The mean of the random distribution allowed us to scale the overall magnitude of the $1/f$ noise and the *SD* the degree of randomisation (mean=10 and *SD*=0.5 for no temporal filter, *m*=190 and *SD*=9.5 for application of temporal filter).

FFTs were performed on the model output with the added noise. The magnitude of model responses was normalized between 0 and 1 by using the minimum and maximum response for the condition being simulated. This controlled for the arbitrary scaling of model responses introduced by the temporal filter. In all cases, this maximum was found at the $1/f$ noise magnitude at 0.1Hz, rather than at a signal bin related to the stimulus input. The same approach that was used for the EEG data was used to calculate SNRs at each frequency, and these were then used to perform model fitting.

Results

General overview

We wanted to measure fundamental responses to grating components, both alone and when forming part of a plaid. We also wanted to measure intermodulation responses at the difference and sum of the fundamental frequencies used to modulate the contrast of each component forming a plaid. Clear component-based responses at f_1 (2.3Hz) and f_2 (3.75Hz) were observed at posterior occipital sites for gratings presented alone as well as when they were presented along with another grating component in the plaid conditions. Further, clear intermodulation responses

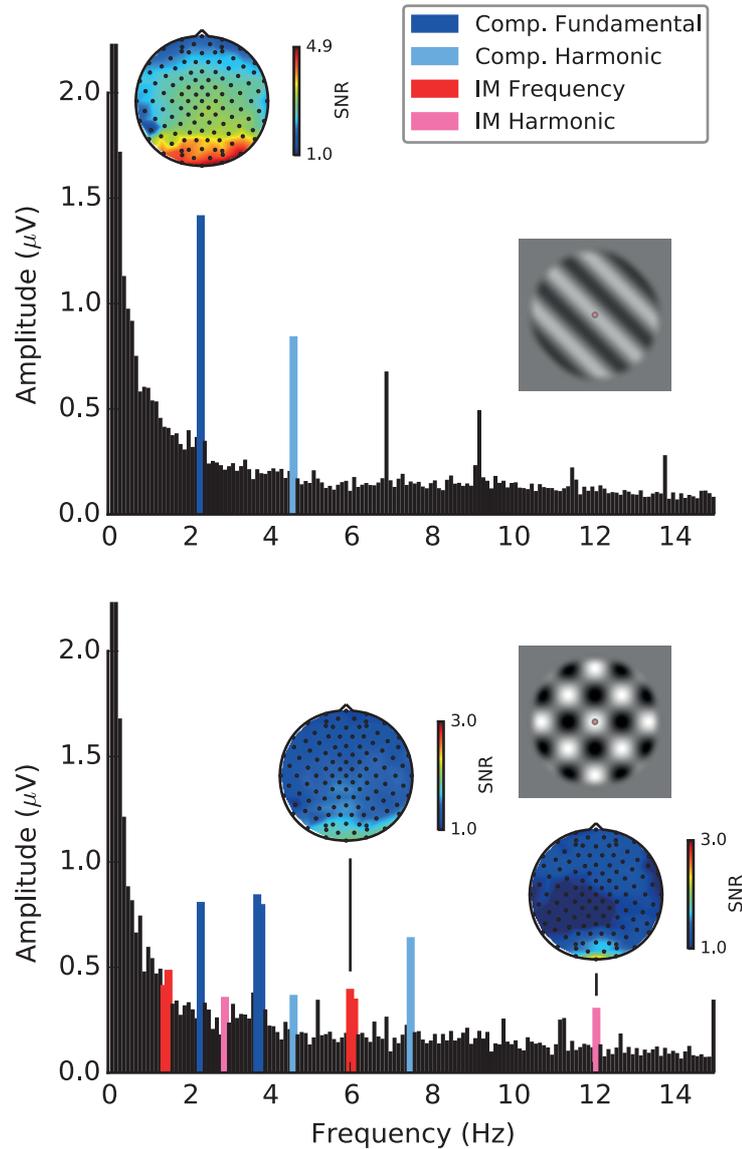


Figure 4.3. Example Fourier amplitude spectra with SNR topographies. Top: grating component A1 alone with fundamental SNR topography. Bottom: Coherent plaid A1A2 with f_1+f_2 and $2f_1+2f_2$ intermodulation SNR topographies. A component alone resulted in fundamental responses at the driving frequency as well as harmonic responses. Simultaneously presenting two components resulted in fundamental and harmonic responses for each component frequency. Additional nonlinearities were observed at the intermodulation frequencies. The reason that a few responses appear to straddle multiple bins (the f_2 , f_2-f_1 and f_1+f_2 responses) was that the stimulus frequencies fell exactly between two bins in these cases; the FFT bins had a resolution of 0.1Hz.

were observed in all plaid conditions at f_1+f_2 (6.05Hz), but only in coherent plaid conditions at $2f_1+2f_2$ (12.1Hz). It should be noted that in Figure 4.3 there are a few responses that straddle multiple bins (the f_2 , f_2-f_1 and f_1+f_2 responses). This was a result of the coding error mentioned in the General Methods section that resulted in

stimuli being presented at half the intended driving-frequency. This resulted in frequency responses that fell exactly between two bins, where each bin had a resolution of 0.1Hz. For example, a driving response of 3.75Hz straddled the 3.7Hz and 3.8Hz bins.

Response to components.

To test the well-documented effects of XOS (DeAngelis et al., 1994; Meese & Holmes, 2007; Petrov et al., 2005) we could compare directly the response to each frequency-tagged grating component in isolation and in the presence of a second grating. The SNR measures of these responses to each component – calculated by dividing the amplitude at the frequency of interest by the average amplitude of the surrounding 12 frequencies (with signal bins excluded) – can be seen in Figure 4.4. Substantial suppression was observed for components presented at 2.3Hz in the presence of another grating component at 3.75Hz (Figure 4.5). This was not the case for components presented at 3.75Hz.

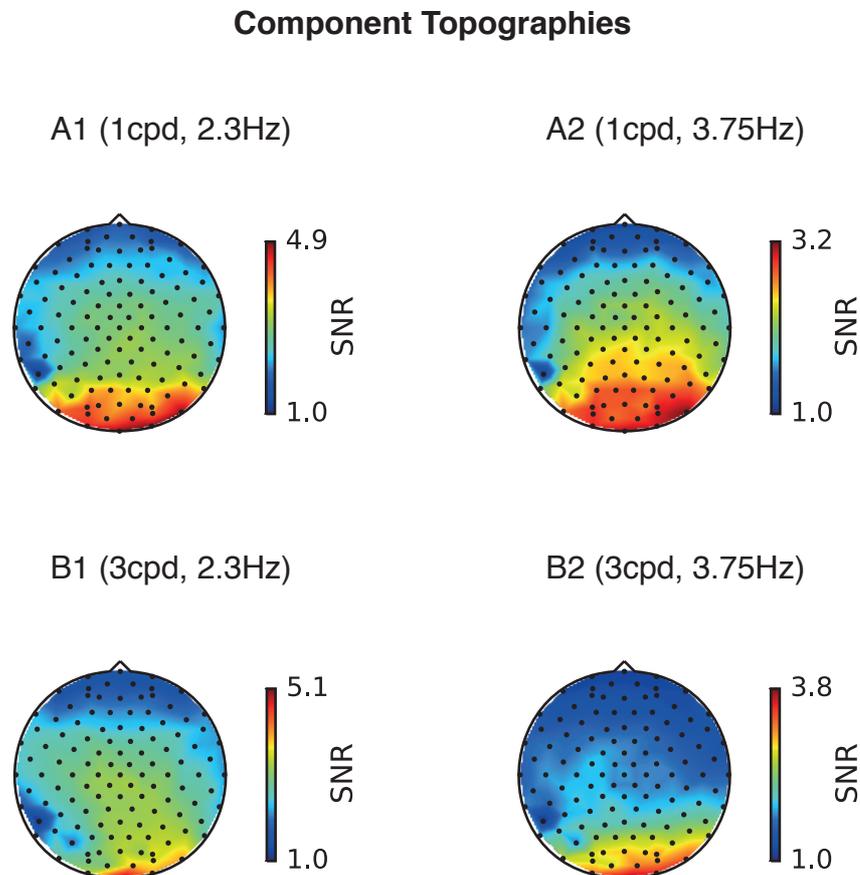


Figure 4.4. SNR topographies for grating components presented alone.

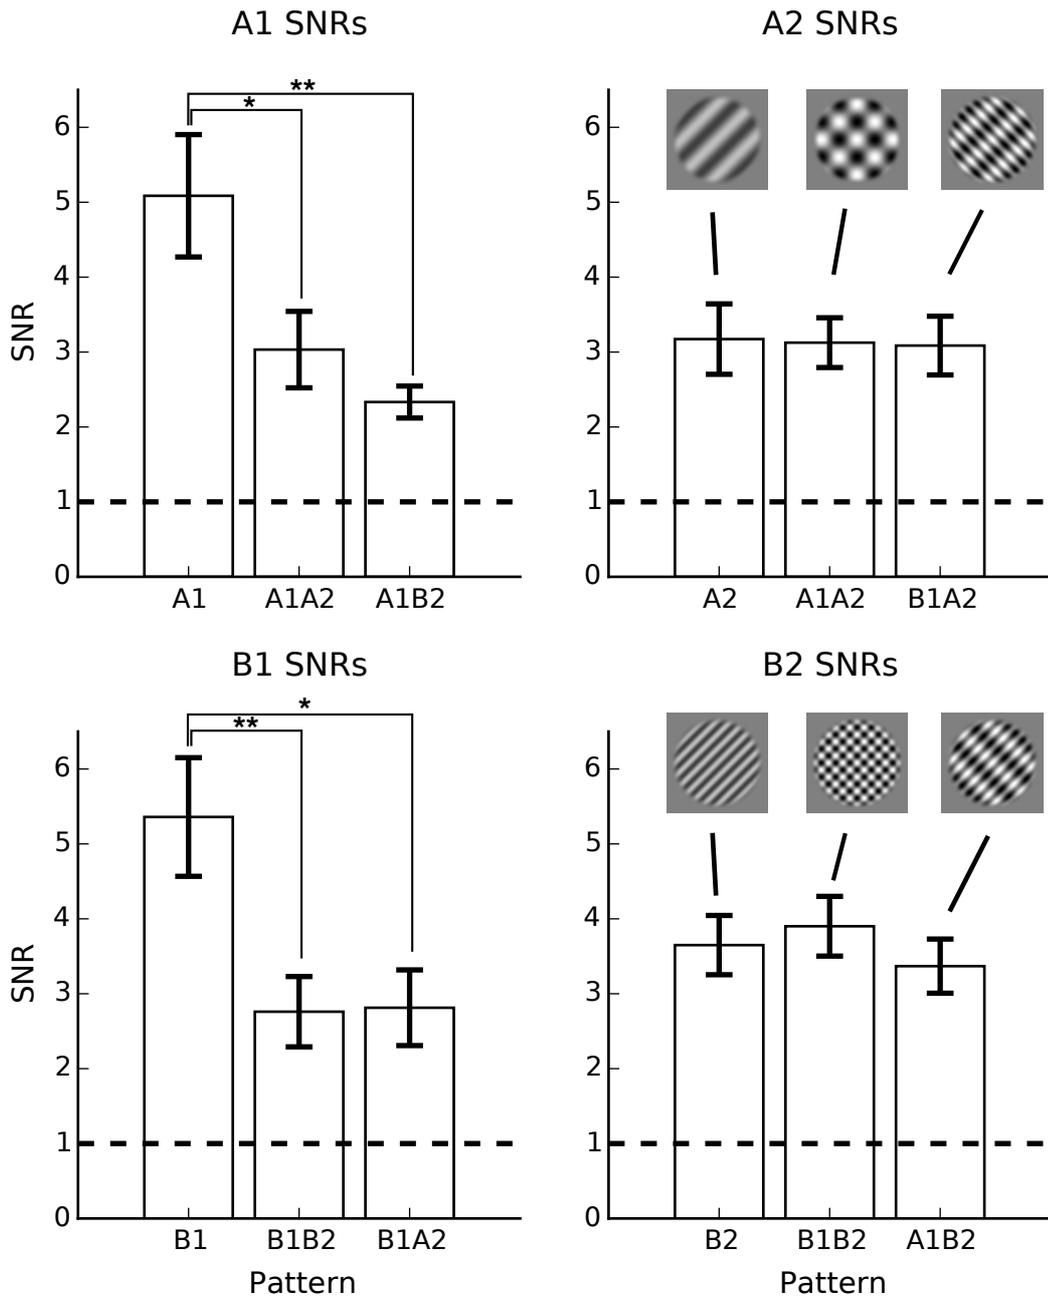


Figure 4.5. Bar plots displaying fundamental frequency response SNRs to component A1, component A2, component B1 and component B2. As indicated by the stimulus insets, the leftmost bar of each plot represents the fundamental frequency SNR when that component was presented alone, the middle when it was part of a coherent plaid, and the right as part of a non-coherent plaid. * represents a significant difference at $p < .05$ and ** at $p < .01$.

All fundamental component responses, whether presented alone or in the presence of another grating, were significantly above background noise ($p < .01$ in all cases). For the 2.3Hz response to grating A1 there was a significant main effect of stimulus pattern ($F_{2,28} = 8.26$, $p = .002$). Post-hoc pairwise comparisons showed that component A1 SNRs were reduced for both the A1A2 plaid ($t_{28} = 2.92$, $p = .018$) and the

A1B2 plaid ($t_{28} = 3.91$, $p = .002$) but there was no significant difference between the two plaid conditions ($t_{28} = 0.99$, $p = 0.587$). A similar effect was observed for 2.3Hz responses to B1 ($F_{2,28} = 6.64$, $p = .004$), again driven by significant differences between the grating alone and each plaid condition (B1B2: $t_{28} = 3.12$, $p = .009$; B1A2: $t_{28} = 3.12$, $p = .011$). There was no significant difference between the response to coherent and non-coherent plaids ($t_{28} = -.06$, $p = .998$).

For the higher temporal frequency components, at 3.75Hz the levels of suppression were less pronounced and did not result in significant main effects of stimulus type (A2: $F_{2,28} = .02$, $p = .98$; B2: $F_{2,28} = .65$, $p = .528$). This pattern of suppression may reflect an additional complex interaction between spatial and temporal frequency tuning of cross-orientation normalisation processes (Cass & Alais, 2006; Meese & Holmes, 2007, 2010).

In summary, although different patterns of suppression were observed between components, levels of cross-orientation suppression were similar for both coherent and non-coherent plaids for all conditions. This indicates that suppressive effects were not spatial-frequency tuned.

Response to plaids.

Intermodulation frequencies were used to assess responses to the conjunction of grating components (see Figures 4.6 to 4.9 for plaid topographies). Responses at the difference frequency ($f_2 - f_1$: 1.45Hz) were not prominent compared to background noise for all conditions (Figure 4.10). A main effect of pattern was found ($F_{6,84} = 3.09$, $p = .009$), though no significant differences were observed in post-hoc pairwise comparisons between plaid conditions. At the harmonic of the intermodulation difference frequency ($2f_2 - 2f_1$: 2.9Hz), responses in all conditions were not prominent compared to background noise (Figure 4.11), and the main effect of stimulus pattern on SNRs was non-significant ($F_{6,84} = 2.07$, $p = .066$).

A1A2

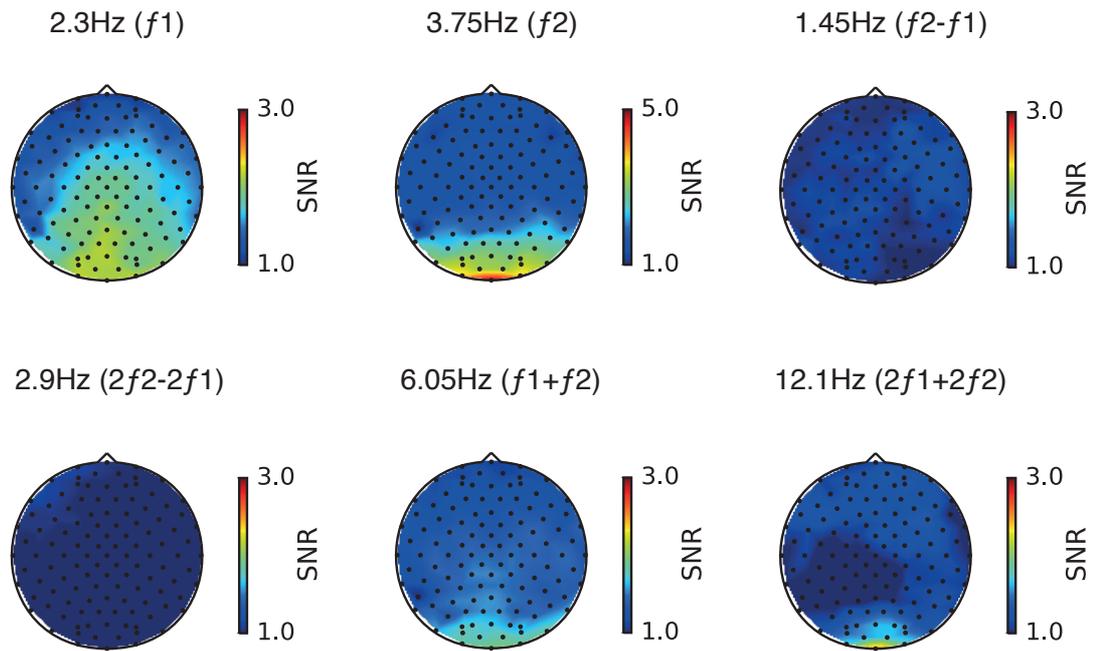


Figure 4.3. SNR average topographies at f_1 , f_2 , $f_2 - f_1$, $2f_2 - 2f_1$, $f_1 + f_2$ and $2f_1 + 2f_2$ for coherent plaid A1A2.

B1B2

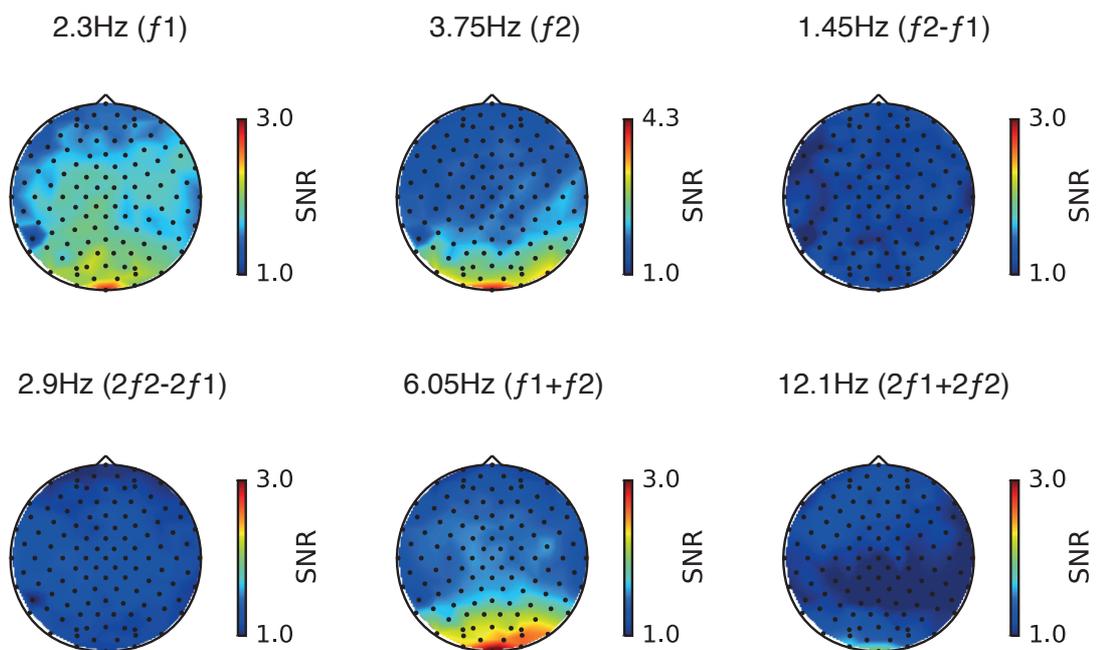


Figure 4.4. SNR average topographies at f_1 , f_2 , $f_2 - f_1$, $2f_2 - 2f_1$, $f_1 + f_2$ and $2f_1 + 2f_2$ for coherent plaid B1B2.

A1B2

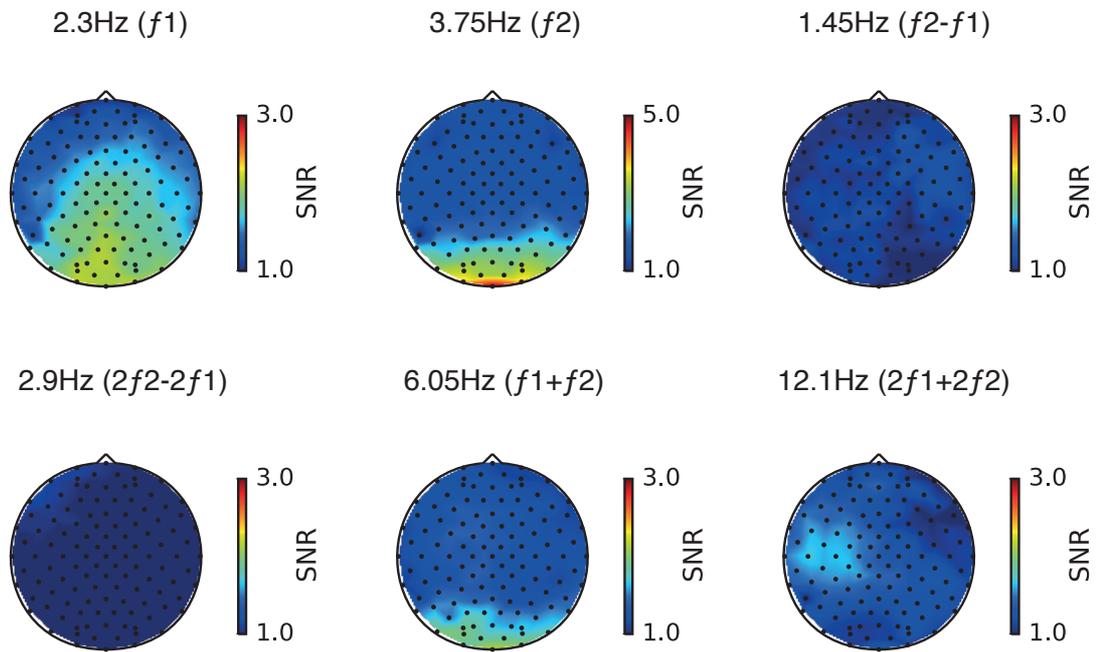


Figure 4.6. SNR average topographies at f_1 , f_2 , $f_2 - f_1$, $2f_2 - 2f_1$, $f_1 + f_2$ and $2f_1 + 2f_2$ for non-coherent plaid A1B2.

B1A2

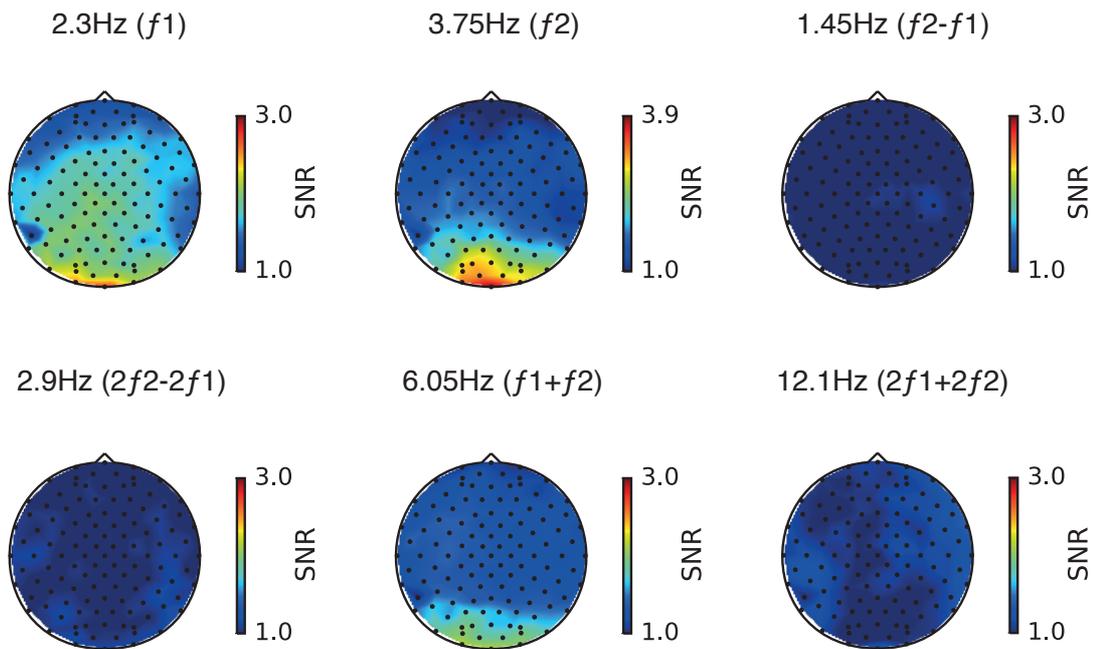


Figure 4.5. SNR average topographies at f_1 , f_2 , $f_2 - f_1$, $2f_2 - 2f_1$, $f_1 + f_2$ and $2f_1 + 2f_2$ for non-coherent plaid B1A2.

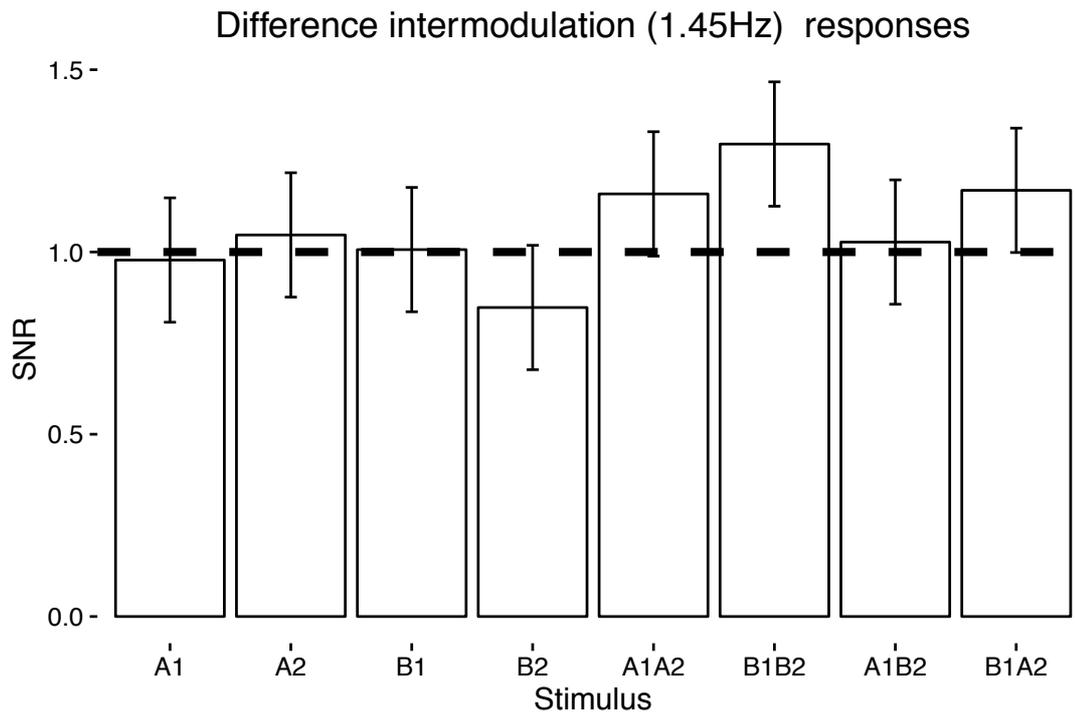


Figure 4.7. Response SNRs at 1.45Hz (f_2-f_1) for each grating component and plaid condition. The bars represent the SNR, the error bars 95% CIs, and the black dashed line represents background noise at an SNR of 1.

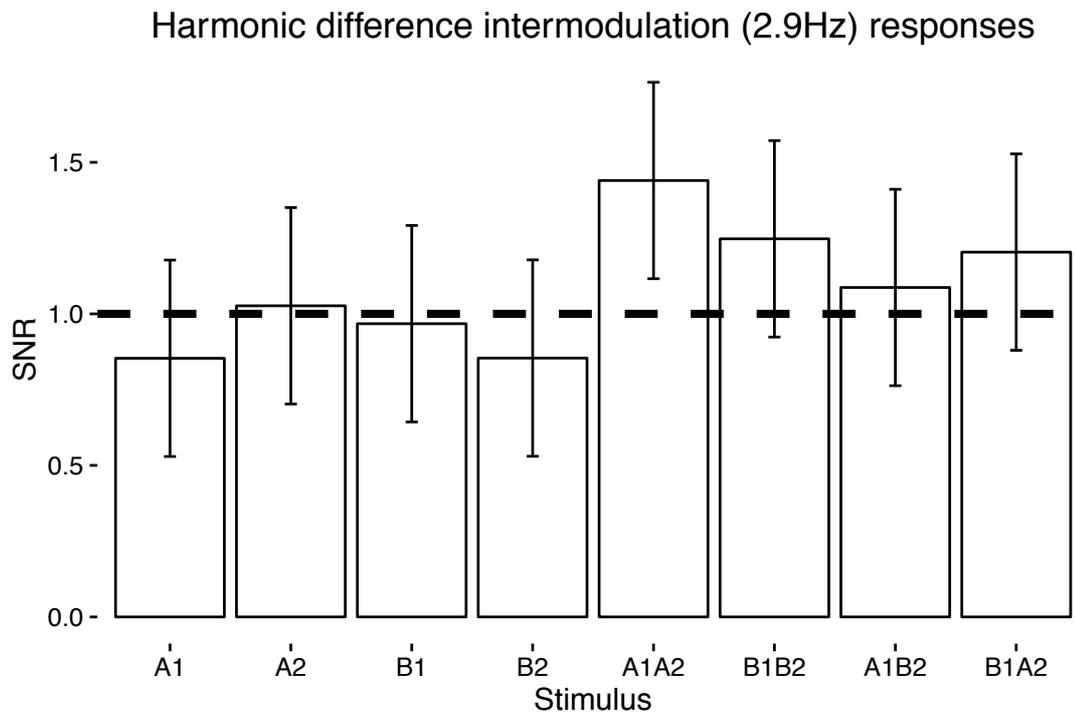


Figure 4.11. Response SNRs at 2.9Hz ($2f_2-2f_1$) for each grating component and plaid condition. The bars represent the SNR, the error bars 95% CIs, and the black dashed line represents background noise at an SNR of 1.

At the sum intermodulation frequency (f_1+f_2 : 6.05Hz) prominent responses were observed (Figure 4.12). One-sample t-tests comparing each condition's SNR to background noise revealed significant differences for both coherent plaid and non-coherent plaid conditions (A1A2: $t_{14}= 5.07$, $p=.001$, 95% CI= [1.62, 2.53]; B1B2: $t_{14}= 7.91$, $p<.001$, 95% CI= [2.70, 3.97]; A1B2: $t_{14}= 4.45$, $p=.003$, 95% CI= [1.66, 2.88]; B1A2: $t_{14}= 3.38$, $p<.023$, 95% CI= [1.41, 2.84]), as well as for component A2 ($t_{14}= 3.24$, $p<.024$, 95% CI= [1.13, 1.65]). The latter finding was unexpected as only one component (i.e. only one flickering stimulus) was presented in that condition; mathematically, an intermodulation response should not take place. This finding might be a statistical false alarm. A significant main effect of stimulus pattern was found ($F_{6,84}= 24.23$, $p<.001$). Post-hoc pairwise comparisons showed that response SNRs at 6.05Hz were significantly larger in response to plaid B1B2 than for plaid A1A2 ($t_{84}= -4.22$, $p<.001$) and the mean 6.05Hz response SNR to the non-coherent plaids (the mean of A1B2 and B1A2 responses; $t_{84}= -3.00$, $p=.009$).

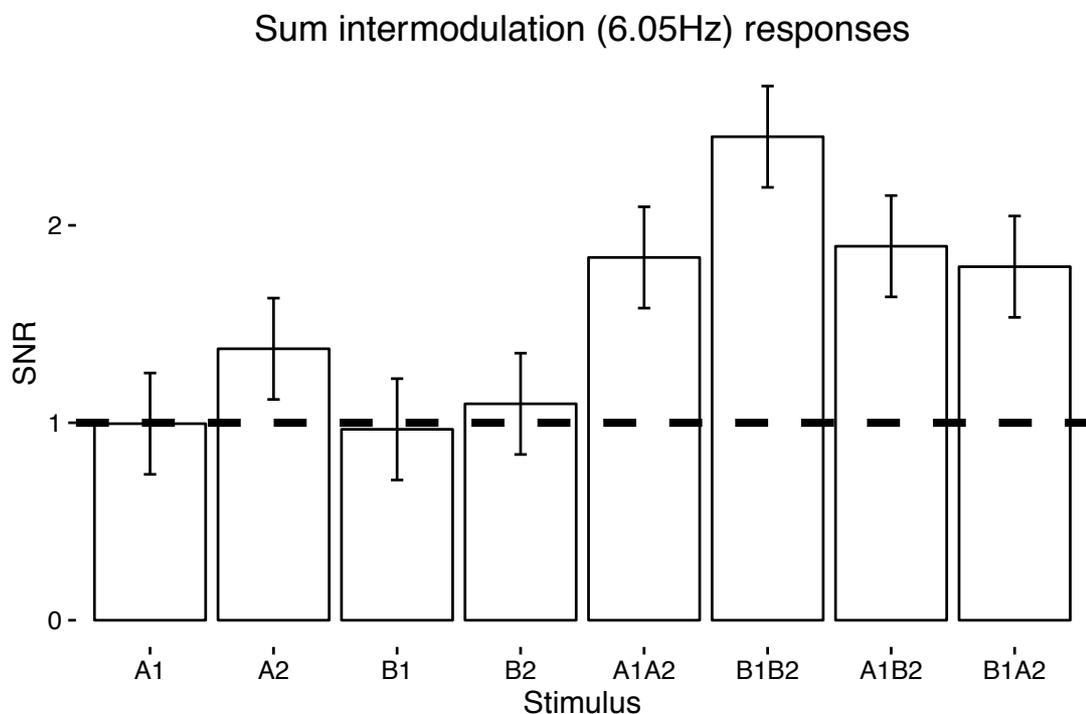


Figure 4.12. Response SNRs at 6.05Hz (f_1+f_2) for each grating component and plaid condition. The bars represent the SNR, the error bars 95% CIs, and the black dashed line represents background noise at an SNR of 1.

Harmonic sum intermodulation (12.1Hz) responses

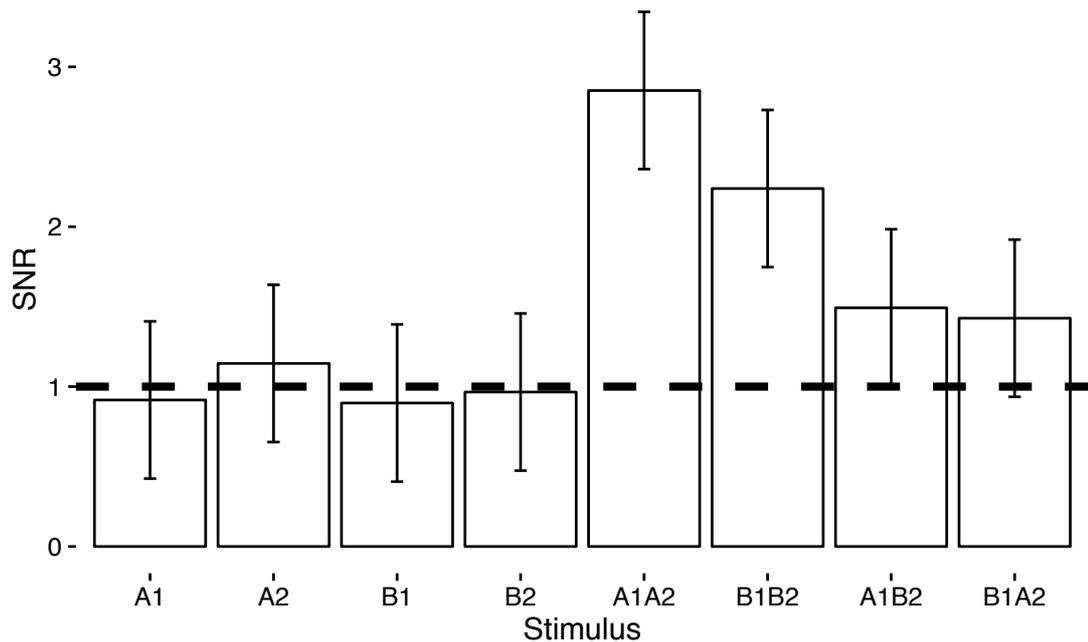


Figure 4.13. Response SNRs at 12.1Hz ($2f_1+2f_2$) for each grating component and plaid condition. The bars represent the SNR, the error bars 95% CIs, and the black dashed line represents background noise at an SNR of 1.

At the harmonic of the intermodulation sum frequency ($2f_1+2f_2$: 12.1Hz; see Figure 4.13) only responses to coherent plaids were significantly above background noise (A1A2: $t_{14}= 3.74$, $p=.017$, 95% CI= [1.71, 3.60]; B1B2: $t_{14}= 4.00$, $p= .017$, 95% CI= [1.48, 2.82]). There was a significant main effect of stimulus pattern on response SNRs ($F_{6,84}= 11.61$, $p< .001$). Significant differences were observed in post-hoc pairwise comparisons between both coherent plaids and the mean non-coherent plaid response (A1A2: $t_{84}= 4.69$, $p< .001$; B1B2: $t_{84}= 2.62$, $p= .030$).

Summary. The response at the sum intermodulation term (f_1+f_2 : 6.05Hz) was significantly above background noise for all plaid conditions, and was greater for coherent plaid B1B2 than any other condition. That all plaid responses were significantly above background noise and that suppression for coherent and non-coherent plaid responses was similar, may indicate that the 6.05Hz response primarily reflected XOS. At the harmonic of the sum intermodulation term (12.1Hz), however, there was a degree of selectivity; the response was larger for coherent than non-coherent plaids, and was significantly above background noise for coherent plaids but

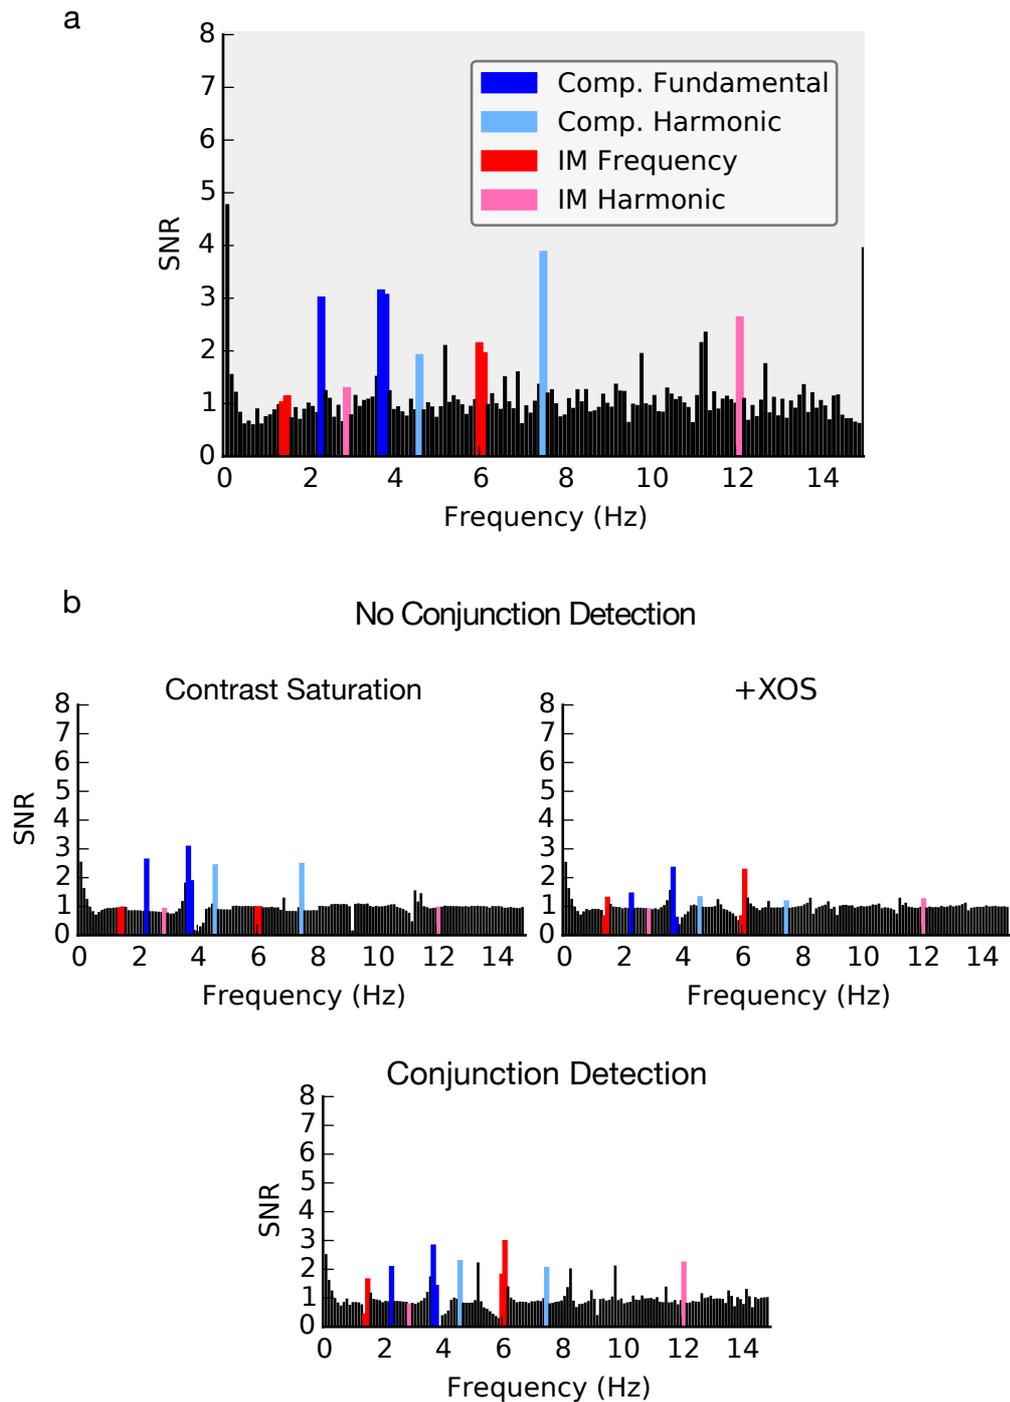


Figure 4.14. SNR Spectra for a) EEG data and b) model data without an NIR being applied. Highlighted in dark blue are fundamental component-based SNRs (f_1 and f_2 : 2.3Hz and 3.75Hz), light blue the component harmonic SNRs ($2f_1$ and $2f_2$: 4.6Hz and 7.5Hz), red the intermodulation responses (f_2-f_1 and f_1+f_2 : 1.45Hz and 6.05Hz) and magenta the intermodulation harmonic responses ($2f_2-2f_1$ and $2f_1+2f_2$: 2.9Hz and 12.1Hz). The contrast saturation model was the simplest, followed by the inclusion of XOS (normalisation pool) and then the conjunction detection model.

not for non-coherent plaids, suggesting that it was spatial-frequency tuned. This appears more consistent with a nonlinearity resulting from plaid-selective mechanisms (which are expected to be spatial-frequency tuned) rather than a XOS mechanism (which is not). Computational modeling is needed to understand whether that intuitive explanation fits with the data, however.

Modelling.

Simulations of several candidate models, including the conjunction detection model outlined in the methods section, were generated. The candidate models with no conjunction detection were each composed of a channel bank of 4 V1 neurons (like Layer 1 of the conjunction model). In the interest of understanding how different component-based nonlinearities contribute to the summed nonlinearities observed in EEG data, input signals underwent only contrast saturation in one model, and both contrast saturation and XOS in another. The output of these simulations is shown in the top rows of Figure 4.14b and Figure 4.15b. In both figures, the model output is being compared to the EEG data we collected for coherent plaid A1A2. The EEG data is shown in section 'a' of both figures, and the various model outputs are shown in section 'b'. The same noise as described earlier was injected into the output of these. Two versions of each were run; one where no temporal filter was applied to the channel output (Figure 4.14b) and another where the temporal filters that were outlined earlier were applied (Figure 4.15b). It should be noted that the application of the second temporal filter (whether bandpass or high-pass) was applied to Layer 2 in the conjunction detection model only (shown in the bottom row of Figure 4.14b and Figure 4.15b).

In the first candidate model, each channel had a stage of contrast saturation upon contrast input (Figure 4.14b top left; Figure 4.15b top left). Fundamental responses and their harmonics were produced. No intermodulation responses were produced because the operation performed by each channel was independent of the other channels; the different temporal signals could not combine. This is dissimilar to the observed EEG data in sections a) of both figures, and clearly not the case for much of the visual system. In V1 substantial XOS is usually observed in response to

superimposed grating stimuli (e.g. Bonds, 1989; Brouwer & Heeger, 2011b; Burr & Morrone, 1987).

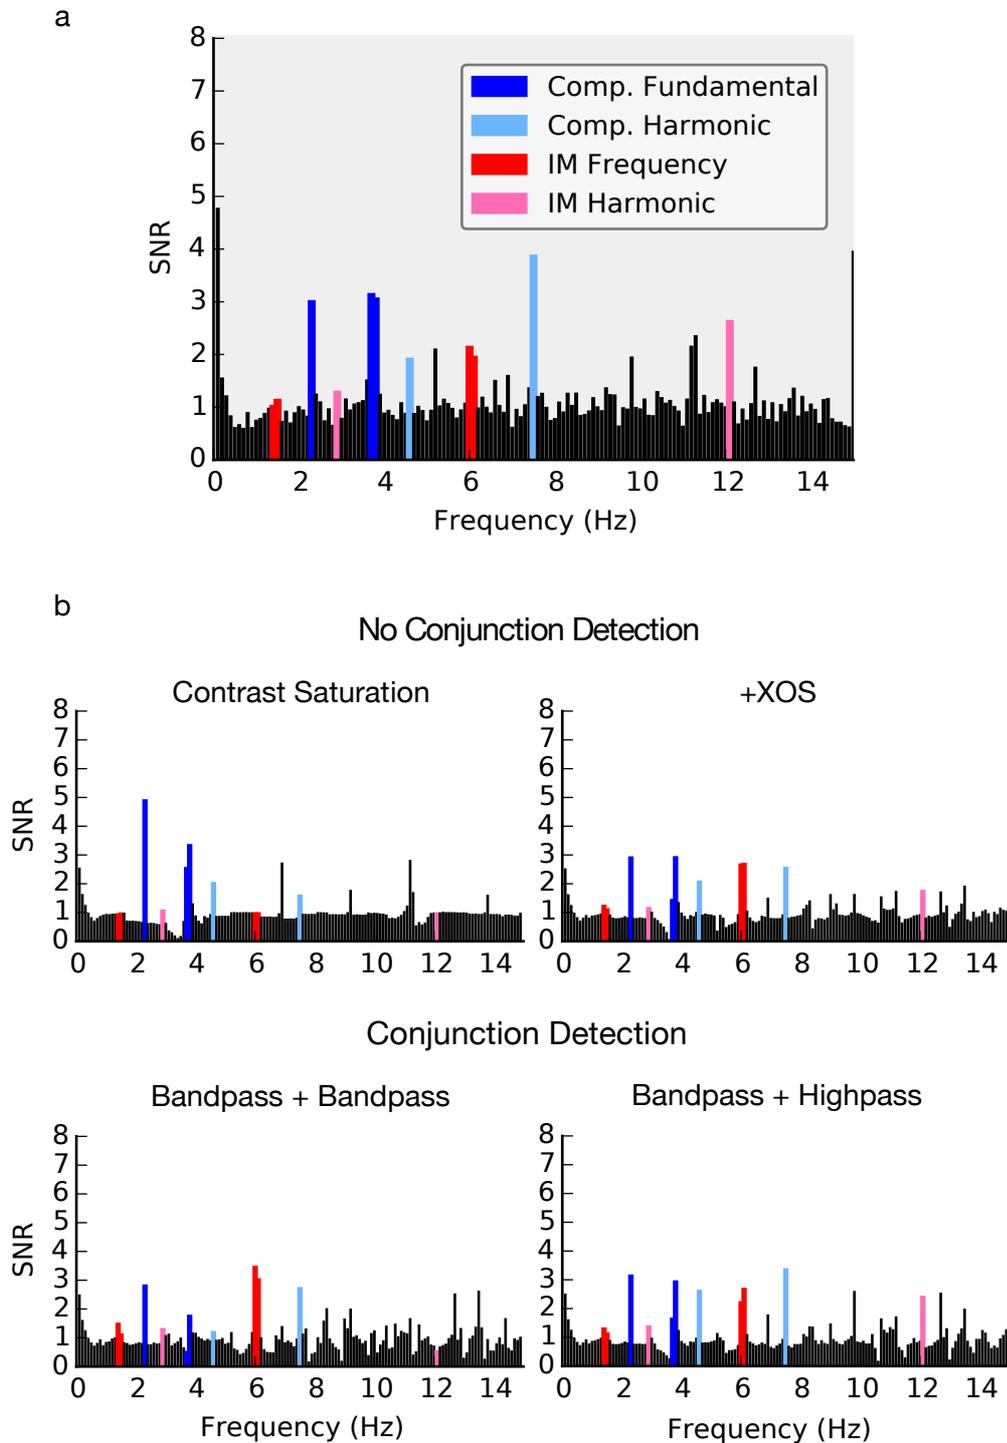


Figure 4.15. SNR Spectra for a) EEG data and b) model data with NIRs applied. Highlighted in dark blue are fundamental component-based SNRs (f_1 and f_2 : 2.3Hz and 3.75Hz), light blue the component harmonic SNRs ($2f_1$ and $2f_2$: 4.6Hz and 7.5Hz), red the intermodulation responses (f_2-f_1 and f_1+f_2 : 1.45Hz and 6.05Hz) and magenta the intermodulation harmonic responses ($2f_2-2f_1$ and $2f_1+2f_2$: 2.9Hz and 12.1Hz). The contrast saturation model was the simplest, followed by the inclusion of XOS (normalisation pool) and then the conjunction detection models.

The second candidate model performed contrast saturation and received input from a normalisation pool (outlined in methods section); i.e. a model of contrast gain to account for contrast saturation and XOS (Figure 4.14b top right; Figure 4.15b top right). When no temporal filter was applied, suppression was observed at both the fundamental and harmonic component-based SNRs. With the application of the temporal filter to the output of each channel, suppression was observed at the fundamental response frequencies, but the harmonic SNRs showed a slight increase. In both cases, a substantial intermodulation response SNR was observed at 6.05Hz (f_1+f_2) and a slight intermodulation response was observed at 12.1Hz ($2f_1+2f_2$). On the surface, this supports the finding that XOS drives substantial intermodulation responses (e.g. T. J. Baker, Norcia, & Rowan Candy, 2011; Candy et al., 2001). However, the output produced by these models is still dissimilar to the raw data in showing relatively weak harmonic component responses (and essentially no harmonic *intermodulation* response).

When no temporal filter was applied, the additional nonlinearity of the conjunction detector raised the SNR at the fundamental, harmonic and intermodulation terms compared to the contrast gain (contrast saturation + XOS) only model. In contrast, the addition of the bandpass temporal filter at Layer 1 and the higher-pass temporal filter at Layer 2 resulted in larger SNRs only at the harmonic responses (4.6Hz: $2f_1$, 7.5Hz: $2f_2$ and 12.1Hz: $2f_1+2f_2$). Applying the bandpass temporal filter to both Layer 1 and Layer 2 of the conjunction model decreased fundamental component responses, the 4.6Hz response and the 12.1Hz response, but notably increased the 6.05Hz response. This latter model was the most dissimilar of the conjunction detector models to the observed EEG data. The model with no temporal filter and the model with different temporal filters at stage 1 and stage 2 both resulted in larger 12.1Hz SNRs; the frequency at which a significant difference between coherent and non-coherent plaids was found in the EEG data. They also display several additional frequency combination responses that are present within the EEG data spectra, such as the response at 9.8Hz (f_1+2f_2). Both clearly capture some of the subtleties involved in the signal combinations taking place in the EEG data, though the larger 12.1Hz response produced by the bandpass + higher-pass model is closer to the EEG findings.

In summary, we generated several simple fixed models to account for different kinds of signal summation in response to frequency-tagged flickering stimuli. Rather than comparing models using measures of fit quality and (somewhat arbitrary) penalties for numbers of parameters, we have compared models in terms of whether they produced responses at the expected frequencies. The only model variants capable of producing the full range of responses observed in the data were the models including a 2nd layer non-linear step; the model with only cross-orientation suppression (normalisation) does not show the observed intermodulation responses. The use of a bandpass temporal filter at Layer 1 and a higher-pass temporal filter at Layer 2 of the conjunction model appears to slightly better account for our data than without the application of any model of the neural impulse response function, and much better than when a different bandpass temporal filter is applied to Layers 1 and 2 of the model. This presumably indicates different temporal integration windows at different stages of the visual hierarchy.

Discussion

We have measured neural responses to sinusoidal grating patterns presented alone and combined as coherent and non-coherent plaids to assess the nonlinear combination of neural responses to the gratings. To do this we measured EEG responses at intermodulation frequencies, which have previously been shown to indicate a nonlinearity at or after the point of summation (D. Regan, 1983; D. Regan & Heron, 1969; M. P. Regan & Regan, 1988; Spekreijse & Oosting, 1970b; Spekreijse & Reits, 1982; Zemon & Ratliff, 1984). We found, for compound stimuli (plaids), there was a non-selective nonlinear response at 6.05Hz, and a plaid-selective response at 12.1Hz ($2f_1+2f_2$), when the combination formed a *coherent* plaid.

Previous measures of intermodulation responses to plaids have suggested that they reflect lateral suppressive effects such as XOS (T. J. Baker, Norcia, & Rowan Candy, 2011; Candy et al., 2001). The frequency-tagging technique we have used allows a more direct measure of XOS by examining the effect of each component (e.g. presented at 3.75Hz) on the response to the other (at 2.3Hz) so, to test whether the

intermodulation response could be caused by XOS we could compare it with the pattern of suppression we measured directly.

We did indeed observe substantial suppression; for example, when a component (either A1 or B1) was presented alone at 2.3Hz, a greater 2.3Hz response was observed than when the component was presented in combination with an orthogonal grating. Critically, however, the 12.1Hz response was very much dependent on the spatial frequencies being matched in the two components. Conversely, and in keeping with previous findings that XOS is largely un-tuned for spatial frequency (DeAngelis et al., 1994; Petrov et al., 2005), the reductions in component responses that we measured occurred equally for any combination of spatial frequencies. It seems unlikely that the observed intermodulation responses resulted purely from XOS, given that one is tuned for spatial frequency and the other is not.

The perception of moving plaids has been shown to depend on matched spatial frequencies in several ways. When spatial frequencies differ in the two gratings being combined, observers perceive a pair of semi-transparent gratings sliding past each other (Adelson & Movshon, 1982) whereas a plaid with matched spatial frequency components appears as a single coherent checkerboard pattern with a single direction of motion. Similarly, selective adaptation to static plaids decreases when components are unmatched (Hancock et al., 2010) and the pop-out effect in visual search disappears when plaid targets have unmatched components (Nam et al., 2009). A mechanism selective for coherent plaids appears to explain better the nonlinear intermodulation responses we have measured. In support of this, we found that the addition of ‘conjunction detector’ channels beyond contrast gain operations resulted in model output more like the EEG data.

Studies of fMRI (McDonald et al., 2012) and positron emission topography (PET) (P Wenderoth, Watson, Egan, Tochon-Danguy, & O’keefe, 1999) have shown similar responses to plaid and grating stimuli of equivalent contrast. McDonald et al. (McDonald et al., 2012) used fMRI to study the summation of signals to gratings and plaids. However, this method measures a ‘mixed’ signal; it cannot distinguish changes in the responses of a set of neurons from changes in the number of neurons responding. For instance, a greater response in V1 could be caused either by reduced

suppression between neurons or by recruitment of an additional mechanism. The inherently superior temporal resolution of EEG, combined with the frequency-tagging technique, allows us to separate responses to different components within the stimulus.

Previous measurements using the frequency tagging technique with combinations of oriented gratings have shown intermodulation responses (T. J. Baker, Norcia, & Rowan Candy, 2011; Candy et al., 2001). These were attributed to suppression between channels; coherent and non-coherent plaids were not compared in terms of spatial frequency matching. These studies also used counterphase flicker, whereas we modulated between low and high contrast intensities and measured fundamental responses at the driving frequencies. The harmonic responses at $2f_1$ (4.6Hz) and $2f_2$ (7.5Hz) - not shown here - displayed the same effects as the fundamental responses in terms of component suppression; all components displayed the same amount of harmonic component suppression irrespective of plaid coherence.

The intermodulation responses we measured were predominantly observed at the sum intermodulation frequencies, and not at the difference. This has also been the case for several other intermodulation studies (Aissani et al., 2011; Alp et al., 2016; Appelbaum, Wade, Pettet, Vildavski, & Norcia, 2008a; Gundlach & Müller, 2013) but the converse was true for Boremanse et al. (2013). One interpretation for differences between the intermodulation terms suggested by Boremanse et al. (2013) is that responses at both frequencies reflect parallel nonlinearities but the sum intermodulation response output may be a temporally band-pass or higher-pass nonlinearity and signal early local spatial interactions, whereas the difference intermodulation response may be generated by a temporally low-pass nonlinearity and generated by signal integration to higher-level (global) stimuli, such as their face-part stimuli, which require longer to process (Alonso-Prieto et al., 2013). Here we found that the application of a bandpass temporal filter at Layer 1 and a higher frequency filter at Layer 2 of our model resulted in better model output than by applying a bandpass filter at both layers. Rapid local combinations would certainly fit in with our EEG results as we used simple sinusoidal gratings that presumably were being combined across many receptive fields to encode another pattern (the plaid).

Further, sum intermodulation responses were strongest around Oz , placed approximately over the occipital pole, consistent with activity relatively early in visual cortex.

Alp et al., (2016) suggested that the temporal resonance properties of different neural mechanisms may influence the varied response at the difference and sum intermodulation frequencies. These resonances may depend on specific synaptic connections to and from the mechanisms, feedback connectivity and the relative complexity of the receptive field within the visual hierarchy (e.g. sensitive to compound plaids or sensitive to faces). The effects of such differences in temporal integration have not been applied quantitatively in a computational model (e.g. to explain differential responses at sum- and difference-intermodulation terms). Here we used a simple approach to model the temporal properties of mechanisms and found that using different neural impulse response functions (temporal filters) at early and late layers was sufficient to explain a wide range of features in the data. In the case studied here (plaid combinations) it appeared that responses at 6.05Hz (f_1+f_2) in the present data primarily represented XOS mechanisms, while responses at 12.1Hz ($2f_1+2f_2$) reflected plaid-selective mechanisms.

The responses to gratings and plaids appear similar in terms of their topography. Although they may result from different neurons, we would expect these to be anatomically proximal. For example, if the response to plaids originated in V2, it would be hard to distinguish from the V1 response to gratings using EEG. Furthermore, although they are reliable, the intermodulation responses do not have large amplitudes, which compounds the difficulty in localizing them.

In summary, we have shown a nonlinear response to a compound of gratings (plaid) that does not arise purely from contrast normalisation between spatial frequency channels. The data are in keeping with a mechanism for detecting conjunctions of visual features, as might result from a logical AND operation. The frequency-tagging technique provides a useful tool to investigate AND-gates in a wide variety of neural mechanisms.

Chapter 5: Measuring the contrast response function of intermodulation responses

Rationale

In the previous chapter I reported a study in which we measured SSVEP responses to high-contrast plaid stimuli. The frequency-tagging approach critically allowed us to disentangle nonlinearities resulting in XOS from other nonlinearities, like those indicative of a putative “plaid detector”. We found a plaid-selective intermodulation response for coherent plaids; when the plaid components were matched in spatial frequency. A simple two-layer model of nonlinear summation was generated that described the data well. The mechanism was structured as a logical-AND gated summing circuit that performs selective signal combinations.

The study reported in the previous chapter used only one high contrast level. We found a plaid selective response, but this tells us nothing about how the mechanism generating that response depends on contrast. Critical for the mechanism is the use of saturating nonlinearities in V1 (Peirce, 2007b, 2011). A computational consequence of the suggested mechanism is that at very low component contrasts – before the expansive rising slope of a neuron’s contrast response function – the mechanism cannot detect a compound pattern over and above its components. This caused May and Zhaoping (2011) to suggest that multiplication is needed for AND gates, instead of the literal “nonlinear summation”, suggested by Peirce (2007). While Peirce (2011) agrees that the mechanism does not detect plaids at low contrast, he points to the psychophysical literature indicating that plaids, indeed, do not have any special status at low contrasts; they appear simply as a pair of non-coherent gratings.

McGovern and Peirce (2010) found that adaptation to plaid components was strongest at low probe contrasts, but adaptation to plaid compounds was strongest at higher probe contrasts. The perception of plaids also changes from a distinct checkerboard pattern to that of its two superimposed grating components at lower contrast levels (Meese & Freeman, 1995). This suggests that, measured psychophysically, plaid-detecting mechanisms operate at mid-to-high contrast levels, but fail at low contrasts.

The aim here was to measure the contrast response function of plaid-selective responses to determine whether the effects noted in the previous chapter are contrast-dependent. As before, using the intermodulation approach allows us to examine the effect that the presence of one component has on the response to another, as well as to disentangle response nonlinearities for XOS and plaid selectivity. Here, we measured intermodulation responses to different combinations of gratings and plaids over a wide range of contrast levels.

As contrast increases, component-based responses should show similar levels of suppression irrespective of the plaid they are a part of. A plaid-selective intermodulation response should occur for the coherent plaid, and should be visible at moderate-to-high contrast levels.

Methods

Participants. Fifteen participants with normal or corrected-to-normal vision gave informed consent to participate in the study. The ethics board at the School of Psychology, University of Nottingham, granted ethical approval.

Stimuli and Experimental Procedure. Stimuli comprised of two sinusoidal gratings (gratings 'A' and 'B') and various combinations thereof. Grating A had a spatial frequency of 1cpd and grating B a spatial frequency of 3cpd. These could then be combined with a second, spatially orthogonal, grating to form plaid patterns that were either coherent ('AA') or non-coherent ('AB'). On each trial the overall orientation of the stimulus, either grating or plaid, was randomly assigned but the orthogonal configuration of the gratings was maintained.

Participants were instructed to maintain fixation on the central red fixation point and to respond as quickly as possible when they detected the appearance of the green task dot (without looking towards it) by pressing the central button on the response box. This green dot appeared at a random time during each trial. Participants were presented with a blank fixation screen consisting of a grey background and red fixation dot until they indicated that they were ready to begin.

The frequencies chosen to generate intermodulation responses were 4.6Hz (f_1) and 7.5Hz (f_2), again on the basis that they allowed the components to be temporally incommensurate. These frequencies are higher than those used in the previous chapter for reasons outlined in the General Methods. We aimed to avoid overlap between the low-frequency $1/f$ decay noise and the difference IM response frequency (2.9Hz), as well as overlap between the typical alpha band (8-12Hz) and the sum IM response frequency (12.1Hz) as in Boremanse, Norcia and Rossion (Boremanse et al., 2013) to boost signal-to-noise.

The component contrast varied sinusoidally in time between 0 and a given maximum contrast for the duration of a trial. The range of maximum component contrast levels used were 2%, 4%, 8%, 16% and 32% Michelson contrast. The contrast of the plaid then was formed by the addition of these components, which varied sinusoidally in time such that the maximal contrast a plaid could have was 64% (32% x 2). The result was 25 stimulus configuration x contrast conditions: three grating components (A1, A2 and B2), one coherent plaid (A1A2) and one non-coherent plaid (A1B2), each displayed at the individual contrast levels. A trial consisted of an 11 second presentation of a flickering grating or two simultaneously flickering, superimposed grating components, followed by a 7-9 second inter-stimulus-interval. Each of the 25 stimuli was presented 1 time in a run (lasting around 8 minutes), with 8 runs in total split across two recording sessions. Participants were given short breaks between runs and were thanked, debriefed and given the opportunity to ask any questions upon completion.

Analytical procedure. The amplitude response at each frequency at each electrode site was converted into a measure of signal-to-noise ratio (SNR) by dividing the amplitude at the frequency of interest by the average amplitude of the surrounding 12 frequency bins. The data used in the following analyses were taken from electrode Oz. This was based on group topographies showing peak responses for all stimuli occurring there, and is consistent with measurements in the vicinity of primary visual cortex.

The frequencies of interest for analysis were at f_1 (4.6Hz), f_2 (7.5Hz), $2f_1$ (9.2Hz), $2f_2$ (15Hz), f_2-f_1 (2.9Hz), $2f_2-2f_1$ (5.8Hz), f_1+f_2 (12.1Hz), and $2f_1+2f_2$ (24.2Hz). To

determine whether SNRs were significantly above background noise (SNR=1), a series of one-sample t-tests were conducted separately for each set of component SNRs (i.e. A1, A1A2 and A1B2; A2, A1A2 and B1A2; B2, B1B2 and A1B2), difference intermodulation SNRs (at f_2-f_1 and $2f_2-2f_1$: A1, A2, B2, A1A2 and A1B2), and sum intermodulation SNRs (at f_1+f_2 and $2f_1+2f_2$: A1, A2, B2, A1A2 and A1B2) at each contrast level. Each series of t-tests were corrected using the ranked Bonferroni-Holm method to control for Type 1 errors (Holm, 1979). The extent to which grating/plaid pattern predicted response SNRs was examined using linear mixed-effects modelling.

Results

Response to components

Most the component-based SNRs were significantly above a background noise level of 1. The only exceptions were component A1 at 2% and 4% contrast, when presented alone and in plaids A1A2 and A1B2, and component A2 at 2% presented alone and 2% and 4% in plaid A1A2. See Figure 5.1 for component-based topographies when presented alone.

To measure directly the effect of suppression for different stimulus configurations the component-based SNR was measured, for each component, when presented alone to when it was presented as part of a plaid. By doing so we could determine (i) whether the component response was suppressed by the presence of another grating component in the plaid conditions and (ii) whether the amount of cross-orientation suppression observed for a given grating component differed as a function of plaid coherence.

A1. Figure 5.2a shows the increasing contrast response functions for responses at the 4.6Hz fundamental frequency for component A1 when presented alone, within plaid A1A2 and within plaid A1B2. By fitting the SNR data for component A1 to the LMEM, a significant main effect of Pattern ($F(2,196)= 10.61, p<.0001$) and Contrast ($F(2,196)= 31.73, p<.0001$) was found. There was also a significant interaction between them ($F(8,196)= 2.89, p<.01$).

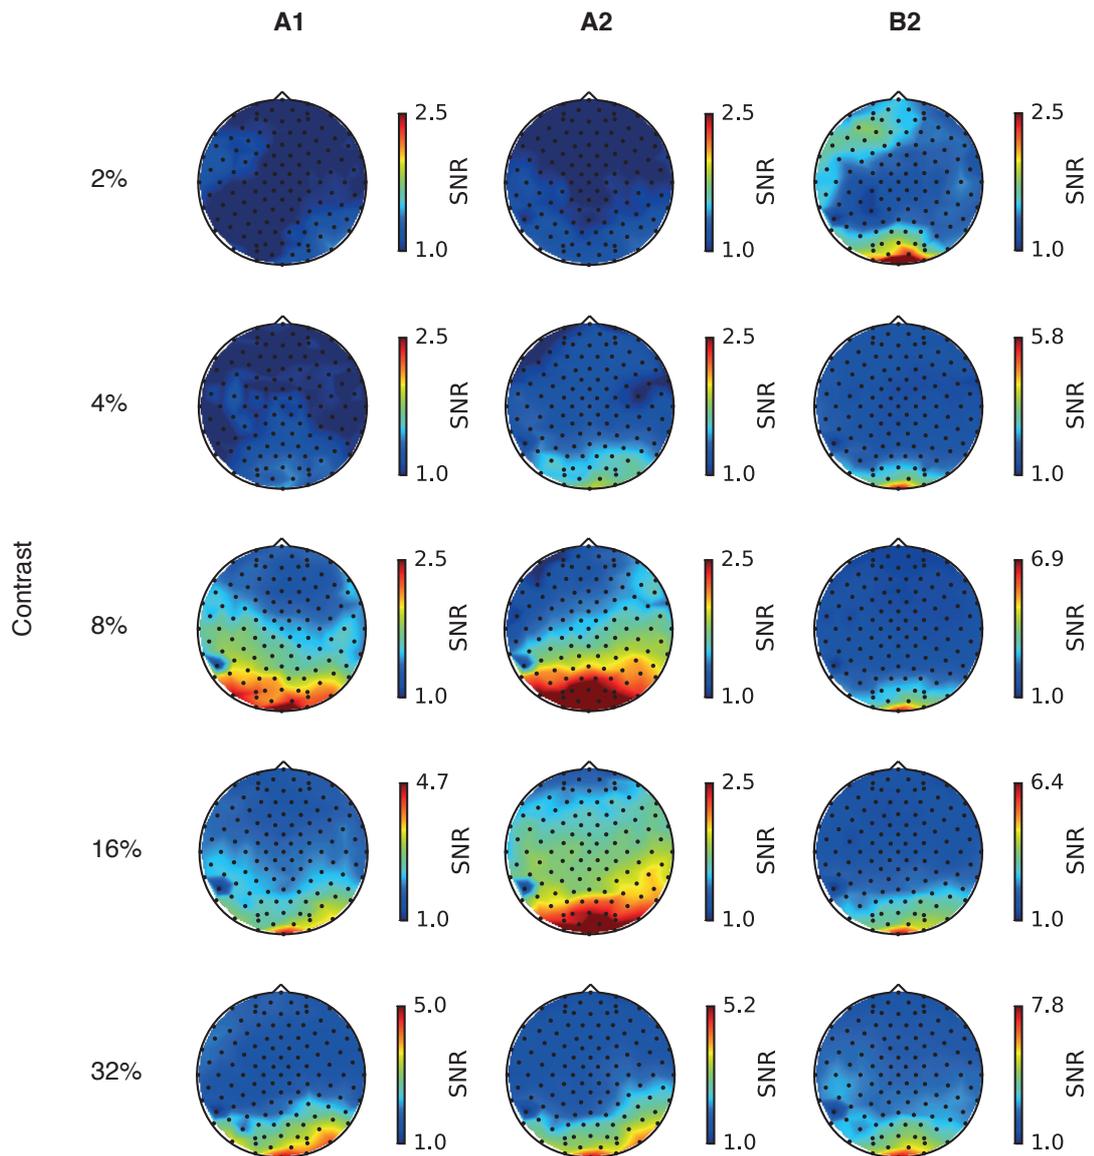


Figure 5.1. Component-based SNR topographies for A1 (left), A2 (middle) and B2 (right). Contrast increases from the top to the bottom of the figure, from 2% Michelson through to 32% Michelson component contrast. Colour maps were scaled so that if the topography's maximum value was less than 2.5, the maximum of the topography was set to 2.5. This helped avoid noisy topographical maps as SNR approached 1.

Pairwise comparisons for component A1 between A1, A1A2 and A1B2 across the five contrast levels were then carried out. There were no significant differences between any grating and plaid patterns at 2%, 4%, and 8% contrast. At 16% contrast, the response to component A1 presented alone was larger than when presented in plaid A1B2 ($t(196) = 3.91, p < .05$). At 32% contrast, responses to component A1 presented in plaid A1B2 were smaller than A1 presented alone ($t(196) = -4.53, p < .001$) and presented in plaid A1A2 ($t(196) = 3.81, p < .05$). This suggests that, at a component

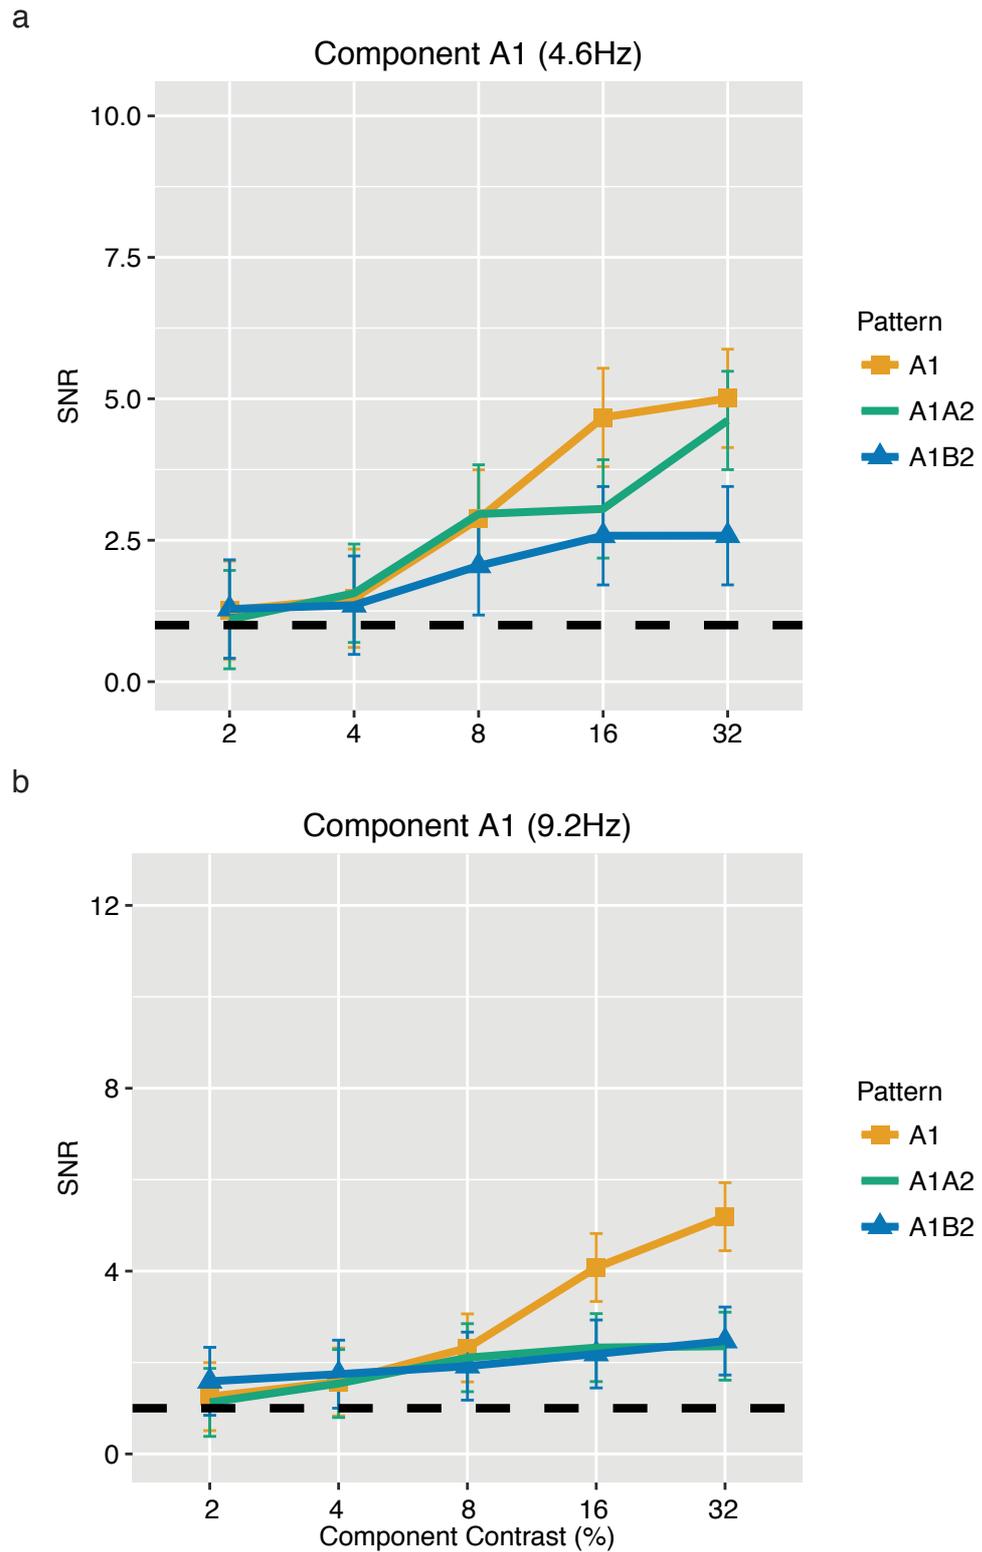


Figure 5.2. Contrast response functions of a) fundamental component-based SNRs at 4.6Hz and at b) harmonic SNRs at 9.2Hz to component A1. As indicated by the key, each figure displays the contrast response function of component A1 presented alone vs. when presented as part of a plaid. The black dashed bar represents an SNR of 1, and error bars represent 95% CIs.

contrast of 32%, there was a difference in component suppression for component A1 between plaids A1A2 and A1B2, with less suppression observed for plaid A1A2.

Figure 5.2b shows the contrast response functions for responses at the harmonic 9.2Hz frequency to component A1 presented alone and as part of plaids A1A2 and A1B2. A significant main effect of Pattern ($F(2,196)= 14.95$, $p<.0001$) and Contrast ($F(4,196)= 21.15$, $p<.0001$) was found, as well as a significant interaction between the two ($F(8,196)= 5.96$, $p<.0001$).

Pairwise comparisons revealed that the 9.2Hz harmonic response increased significantly across contrast (2%, 4%, and 8% all less than 16% and 32% at $p<.05$) only when component A1 was presented alone. A1 presented alone also resulted in larger 9.2Hz response SNRs than when presented as part of plaids A1A2 and A1B2 at 16% ($p<.05$ and $p<.01$, respectively) and 32% ($p<.0001$) contrast. This suggests that substantial suppression occurred to a similar degree at 9.2Hz for responses to component A1 when presented in plaids A1A2 and A1B2.

A2. An increasing contrast response function was observed at the fundamental frequency of 7.5Hz for component A2 SNRs when presented alone and in plaid A1A2 (Figure 5.3a). A main effect of stimulus pattern affected ($F(1, 129)= 4.68$, $p<.05$), and component contrast ($F(4,126)= 23.74$, $p<.0001$) was found. No interaction was observed between the two.

Pairwise comparisons for component A2 between A2 and A1A2 across the five contrast levels were then carried out. SNR's to A2 presented alone and in plaid A1A2 were similar across contrast, suggesting that there was very little suppression of component A2 when combined with component A1. The significant effect of stimulus pattern upon the SNRs is likely due to the change in function steepness for A2 between 4% and 16% component contrast, which resulted in a local peak in the function at 8%.

Figure 5.3b displays harmonic response SNR functions at 15Hz for component A2. As suggested by the figure, only a significant main effect of Contrast was found ($F(4,126)= 55.37$, $p<.0001$). Further post-hoc analyses were not warranted for this data.

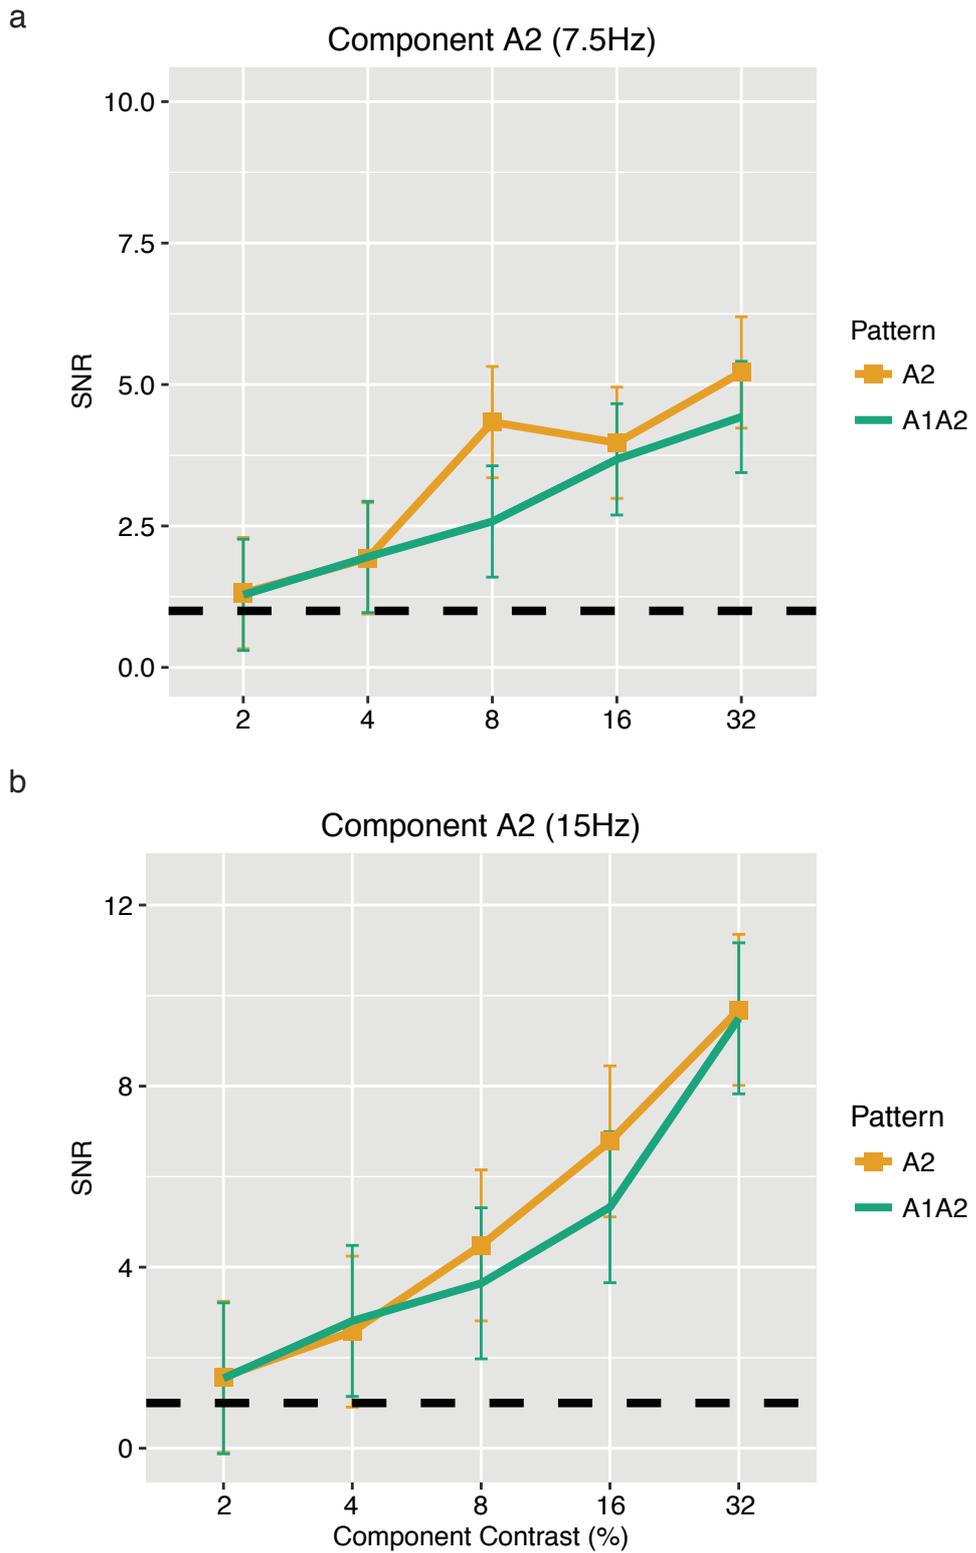


Figure 5.3 Contrast response functions of a) fundamental component-based SNRs at 7.5Hz and at b) harmonic SNRs at 15Hz to component A2. As indicated by the key, each figure displays the contrast response function of component A1 presented alone vs. when presented as part of a plaid. The black dashed bar represents an SNR of 1, and error bars represent 95% CIs.

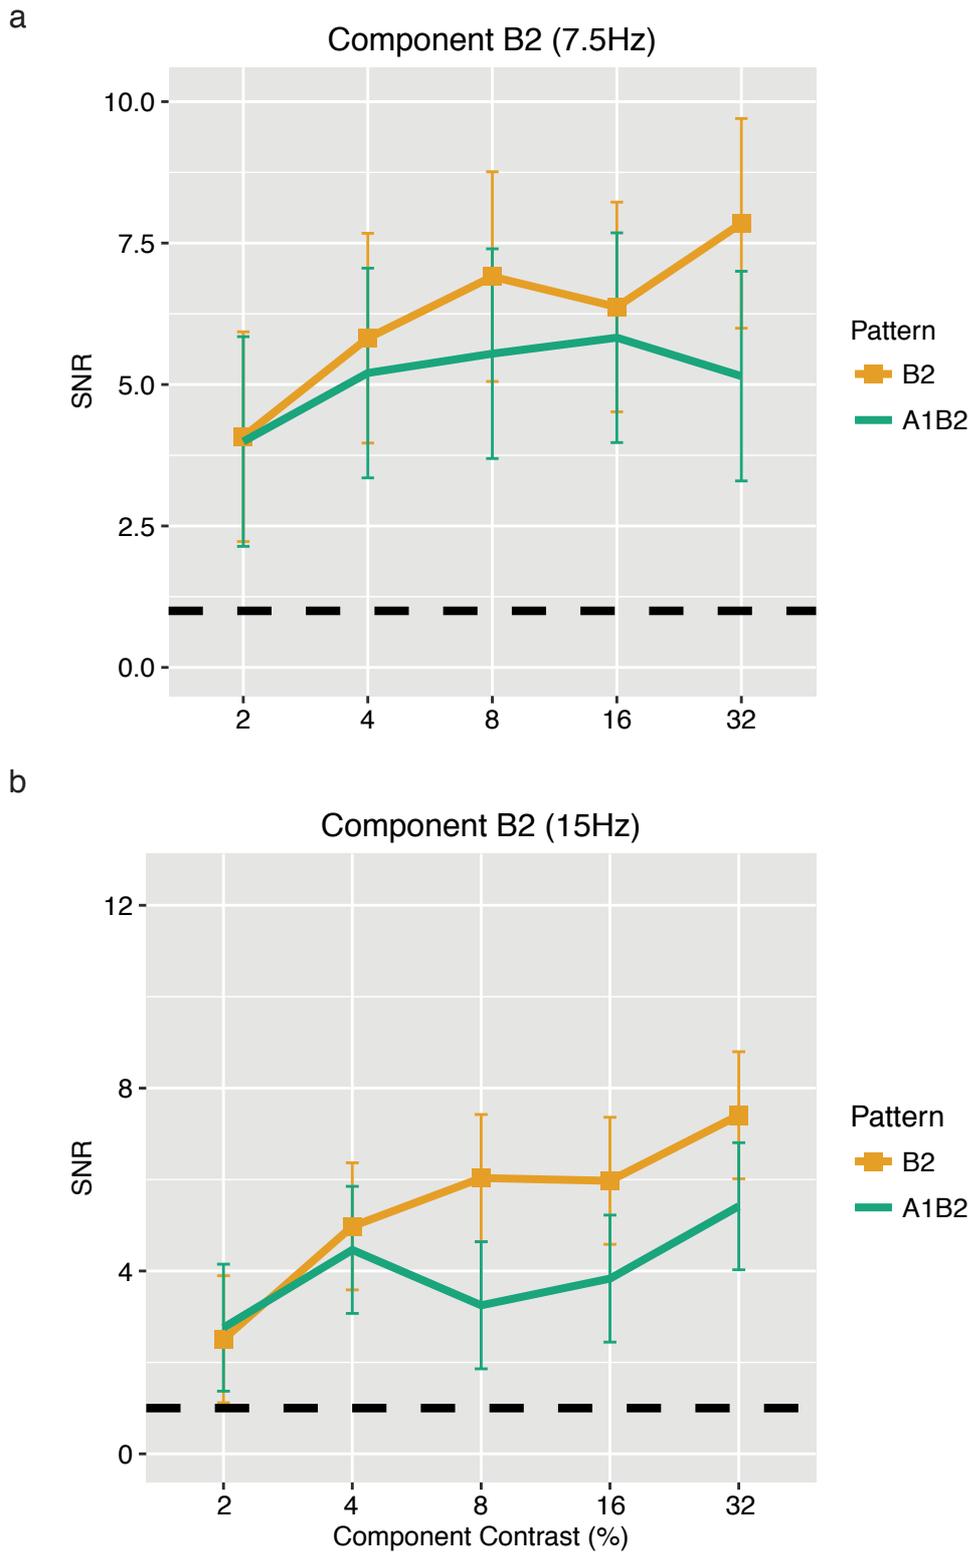


Figure 5.4. Contrast response functions of a) fundamental component-based SNRs at 7.5Hz and at b) harmonic SNRs at 15Hz to component B2. As indicated by the key, each figure displays the contrast response function of component A1 presented alone vs. when presented as part of a plaid. The black dashed bar represents an SNR of 1, and error bars represent 95% CIs.

B2. The contrast response function for B2 presented alone and in plaid A1B2 is shown in Figure 5.4c. Stimulus pattern affected SNRs ($F(1,126)= 6.69, p=.01$), as did component contrast ($F(4,126)= 4.6, p<.01$). There was no statistical interdependence between stimulus pattern and component contrast ($F(4, 126)= 1.24, p=.30$). Despite the significant main effects just described, the only significant post-hoc test results were between B2 at 2% contrast and B2 at 32% contrast ($t(126)= -4.11, p<.01$), and between A1B2 at 2% contrast and B2 at 32% contrast ($t(126)= -4.2, p<.01$). This suggests that, in general, SNRs for component B2 presented alone and in plaid A1B2 were similar.

Figure 5.4c displays harmonic response SNR functions at 15Hz for component B2. A significant main effect of Pattern ($F(1, 126)= 15.91, p<.001$) and Contrast ($F(4,126)= 11.18, p<.0001$) was found, as well as a just-significant interaction ($F(4, 126)= 2.44, p=.05$). Despite this, post-hocs revealed only that B2 presented alone increased significantly in contrast between steps in contrast (2% less than 8%, 16% and 32% component contrast at $p<.01$), as did A1B2 (2% less than 32% component contrast at $p<.05$). At 8% component contrast, B2 presented alone resulted in a significantly larger SNR than it did when presented in plaid A1B2, indicative of some suppression.

In summary, component-based fundamental and harmonic responses displayed increasing contrast response functions. The only difference in suppression between plaids was at 32% component contrast for A1 presented in plaids A1A2 and A1B2; the fundamental 4.6Hz response SNR in A1B2 was lower than in A1A2.

Response to plaids

Only responses at the difference intermodulation frequency were substantially above background noise (see Figure 5.5 for SNR topographies, and Figure 5.6 for the analysed response SNRs). Topographies and responses for the difference harmonic at 5.8Hz, the sum intermodulation response at 12.1Hz and its harmonic at 24.2Hz can be seen in Figures 5.7 to 5.12 in Appendix Two.

Difference Frequency. At 2.9Hz, an increasing contrast response function that was steepest between 16 and 32% contrast was observed for coherent plaid A1A2. A substantial response at 2.9Hz was observed for coherent plaid A1A2 at 32% contrast ($t(14)= 4.48, p<.05$). No other condition resulted in an intermodulation response significantly different from an SNR of 1.

A significant main effect of Pattern ($F(4,336)= 7.26, p<.0001$) and Contrast ($F(4,336)=5.81, p<.001$), as well as a significant interaction between them ($F(16,336)= 3.27, p<.0001$) was found . Paired-samples t-tests were conducted between plaid patterns A1A2 and A1B2 across the five component contrast levels. For plaid A1A2,

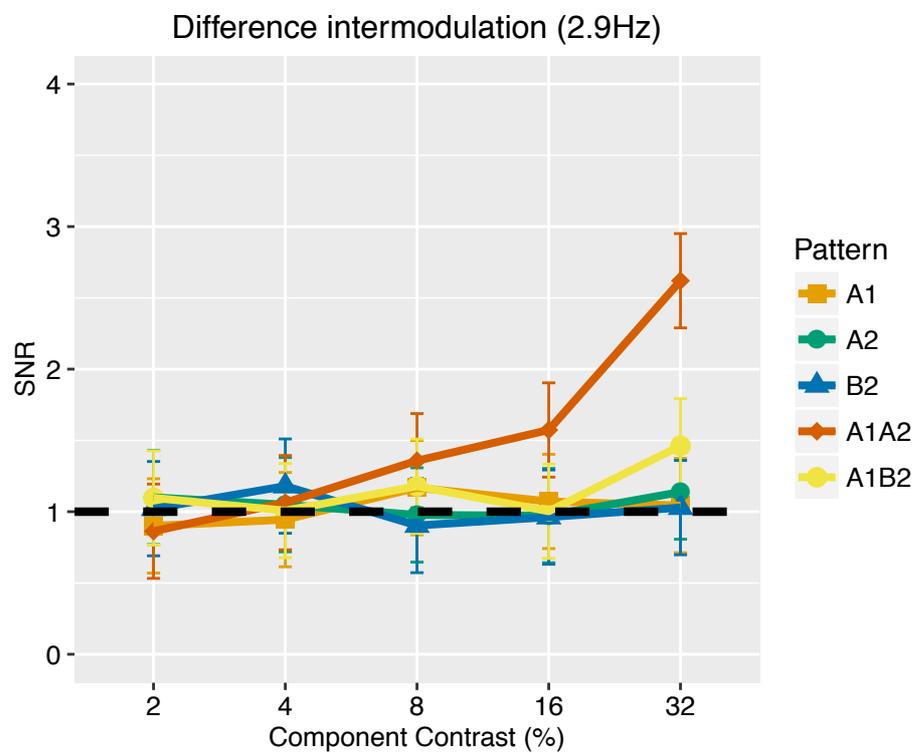


Figure 5.5. Contrast response functions for SNRs at the difference intermodulation frequency (2.9Hz) for each condition. The black dashed line represents an SNR of 1, and error bars represent 95% CIs.

responses at contrast levels 2%, 4%, 8% and 16% were non-significantly different from each other, but were all significantly smaller than responses at 32% contrast (all at $p<.001$). On the other hand, there was no significant increase across contrast for non-coherent plaid A1B2. A significant difference at 32% contrast between plaids A1A2 and A1B2 ($t(336)=4.88, p=.0001$) was found. The interaction between Pattern and Contrast is likely due to the SNRs for A1A2 increasing with component contrast

(significant from 8% contrast when not corrected for multiple comparisons) and A1B2's function remaining flat.

To summarise, only the difference intermodulation response displayed robust responses. This was found for coherent plaid A1A2 presented at 32% component contrast, and this was larger than the 2.9Hz response to plaid A1B2.

Discussion

We have measured neural responses to sinusoidal grating patterns presented alone and combined as coherent and non-coherent plaids at a range of component contrast levels. This was with the aim of measuring the contrast response functions of different nonlinear combinations of neural responses to the grating components. To do this we measured responses at intermodulation frequencies, which have previously been shown to indicate a nonlinearity at or after the point of summation (D. Regan, 1983; D. Regan & Heron, 1969; M. P. Regan & Regan, 1988; Spekreijse & Oosting, 1970b; Spekreijse & Reits, 1982; Zemon & Ratliff, 1984). The main finding was that for compound stimuli (plaids) we found reliable responses at 2.9Hz (f_2-f_1), but only when the combination formed a coherent plaid and its components were presented at 32% contrast. This is essentially a conceptual replication of the previous chapter, with the difference that the contrast in the previous chapter was higher (50% component contrast) and the component modulation frequencies lower (2.3Hz and 3.75Hz).

As discussed in the previous chapter, this intermodulation response indicates a nonlinear combination of signals; it does not tell us the nature of that nonlinearity. It could be the case that it just reflects lateral suppressive effects like XOS (T. J. Baker, Norcia, & Rowan Candy, 2011; Candy et al., 2001). The frequency-tagging technique we have used allows us to investigate that hypothesis directly by examining the effect of each component (e.g. presented at 7.5Hz) on the response to the other (at 4.6Hz). Here, we observed suppression of component A1, but only as part of plaid A1B2 at 16% and 32% component contrast. Further, this suppression only differed from plaid A1A2 at 32% component contrast.

As in the previous chapter, it is difficult to explain the intermodulation response (this time at 2.9Hz) for plaid A1A2 in terms of XOS, given that this was

directly measured only for A1B2. XOS and the 2.9Hz response appear to be oppositely tuned. Instead, a plaid-selective mechanism combining component-based responses might be contributing to the larger intermodulation response. This is in line with previous findings of perceptual plaid coherence, compound adaptation and the pop-out effect in visual search tasks being dependent on spatial frequency matching (Adelson & Movshon, 1982; Hancock et al., 2010; Nam et al., 2009).

The difference intermodulation response appears to increase continuously across component contrast until 16%-32% where the function slope became steeper. Tying back into the conjunction-detection model, these findings could suggest that the response threshold of the mechanism is higher than most of the contrast range that can be used on a stimulus like a plaid. The contrasts of a plaid's components are interdependent because of their spatial superposition. This seems unlikely, though, considering that the perception of plaid coherence doesn't break until below 4% contrast (Meese & Freeman, 1995). Based on my own observations, it seems more likely that responses at the intermodulation frequencies are intrinsically low in magnitude – certainly much lower than component-based responses – so get obscured by background noise for mid contrast stimuli. This background noise could be due to high phase variability between participants that does not get cancelled out by the incoherent averaging (i.e. independent of phase) across participants that was used here.

To the best of my knowledge, only two other studies have measured intermodulation responses across contrast, and both used similar frequency-pairs to the present study. Brown, Candy and Norcia (1999) and Candy, Skoczenski and Norcia (2001) used, respectively: 5Hz-7.5Hz; 3.3Hz-8.3Hz and 5.5Hz-8.3Hz. Surprisingly, they found neither f_1+f_2 nor f_2-f_1 responses across contrast for orthogonally superimposed test and mask gratings, but did find robust responses at f_2-f_1 and $2f_2-2f_1$ for parallel test-mask grating configurations. A later study also found low (if any) intermodulation responses for orthogonal test and mask grating stimuli at a standing component contrast of 40% (T. J. Baker, Norcia, & Rowan Candy, 2011). Lateral suppression brought about by XOS is highly nonlinear, and would be expected to produce intermodulation responses as observed in the present and previous study (Cunningham, Baker, & Peirce, 2017).

Three key differences between my work and those mentioned above are that they used a square-wave profile for contrast modulation, the modulations were counterphase (i.e. between 1 and -1) and they used contrast sweeps of test stimuli to measure suppression across contrast. Particularly to the former difference, using sinusoidal contrast modulation is more specific than square-wave contrast modulation; it only introduces those frequency components to the response signal relating to the stimuli and their combination, whereas square-wave contrast modulation introduces responses at additional frequencies. These additional frequencies could have overlapped at intermodulation frequencies, meaning there is a risk that they could be enhanced or attenuated responses at those frequencies.

Here we found a difference in XOS between plaids A1A2 and A1B2 for component A1. Though XOS is largely un-tuned for spatial frequency (DeAngelis et al., 1994; Petrov et al., 2005), recent research has highlighted that the tuning of XOS can change depending on the interaction between the spatial and temporal frequency of the stimuli being presented (Cass et al., 2009; Meese & Holmes, 2007). Different from the present study, the adaptation studies mentioned above used stimuli where both components (the test and mask, as they were being used) were presented at the same temporal frequency. Brown, Candy and Norcia (1999) and Candy, Skoczenski and Norcia (2001) used the intermodulation approach and found that, for test and mask gratings that were orthogonally superimposed, suppression of components followed a similarly shaped function to that found here. Though they used contrast sweeps to measure contrast response functions, this would not be expected to alter the shape of the response functions for XOS, just the magnitude.

We have again shown a nonlinear response to a compound of gratings (plaid) that does not appear to be determined purely from contrast normalisation between spatial frequency channels. This response became prominent at the highest component contrast level used (32%). The data support the idea of logical AND operations being used by mechanisms for detecting conjunctions of visual features.

Chapter 6: Measuring intermodulation responses in the auditory domain

Rationale

In the previous two chapters I used the two-frequency SSVEP method to study signal combinations in the human visual system. In both chapters I observed intermodulation responses that appeared to be tuned to compound plaid stimuli that were selective for plaid coherence. When the components forming a plaid matched in spatial frequency, a greater intermodulation response was observed that could not be fully explained by normalisation brought about by lateral inhibition. The main implications of this are two-fold. First, this suggests that the human visual system is equipped with mechanisms that are sensitive to certain combinations of information, in this case the combination of grating components within a plaid. Second, these studies demonstrate the usefulness of the two-frequency method for investigating nonlinear responses in the human brain.

The brain might employ nonlinear combinations like these beyond just the visual cortex. Some studies of multisensory integration have employed the two-frequency method to understand how multisensory stimulus combinations interact (e.g. Colon et al., 2014; Nozaradan et al., 2015). For example, Nozaradan, Zerouali, Peretz and Mouraux (2015) examined audio-tactile interactions, and found that intermodulation responses were indicative of sensorimotor integration when participants tapped to a beat. The auditory system, like the visual system, is hierarchically organised (Wessinger, VanMeter, Tian, Pekar, & Rauschecker, 2001). Neural populations in primary auditory cortex (A1) also make use of nonlinearities (Ahrens, Linden, & Sahani, 2008; David, Mesgarani, Fritz, & Shamma, 2009; Sutter & Loftus, 2003). For example, A1 responses increase and saturate with spectro-temporal auditory contrast (analogous to spatiotemporal visual contrast), and an A1 neuron's response can be reduced by noise masking (Phillips, 1990) – both of which can be represented well by the normalisation equation (Carandini & Heeger, 2012; Rabinowitz, Willmore, Schnupp, & King, 2011).

The same techniques as we have used in previous chapters are all able to be applied in the same way to study auditory cortex. Previous studies have used EEG to show responses to amplitude-modulated sounds in the same way that we used contrast-modulated visual stimuli. The combination of pure tones to form chords could be used as an auditory analogue to the combination of sinusoidal gratings to form plaids in the visual system. Chords can be consonant or dissonant; the former describing a combination of pure tones that results in a clean or coherent sound, while the latter results in a combination that is perceptually jarring or non-coherent. Several studies have used EEG to understand how the auditory system processes consonance and dissonance (Bidelman & Grall, 2014; Fishman et al., 2001; Kung et al., 2014; Park, Park, Kim, & Park, 2010). For example, Park et al., (2010) found significantly higher induced gamma-band activity when listening to consonant chords compared to listening to dissonant chords. In a number of cases, gamma-band activity has been related to perceptual binding, for example to Kanizsa triangles (Tallon-Baudry, Bertrand, Delpuech, & Pernier, 1996), though this is still contested (Palanca & DeAngelis, 2005). Park et al suggested that the increase gamma-band activity for consonant chords is indicative of perceptual binding between the notes forming the consonant chord. In this chapter, we wanted to implement the two-frequency method to determine whether we find any evidence for nonlinear signal combinations in the auditory ssVEP responses, and whether they show the same selectivity for *coherent* compounds as we have found in visual cortices.

Methods

Participants. Fifteen participants with normal hearing and normal or corrected-to-normal vision gave informed consent to participate in the study. The ethics board at the School of Psychology, University of Nottingham, granted ethical approval. The work was conducted in accordance with the 2008 version of the Declaration of Helsinki.

Stimuli and Experimental Procedure. The stimuli were comprised of six frequency-modulated pure tones – notes F (349.23Hz) A (440Hz), C (523.25Hz), E (659.25Hz), F[#]

(369.99Hz) and C[#] (554.37Hz). These were all generated at a sample rate of 44KHz (see Figure 6.1 for an example). They were combined, as a first step, into pairs of tones (FA, CE, F[#]C[#]). These were combined with each other to form chords; consonant chord FACE and dissonant chord FF[#]AC[#] (see Figure 6.2). On each trial the stimuli were played dichotically.

When seated in the recording booth, participants were instructed to maintain their attention on a nature documentary that was playing on screen. The volume for

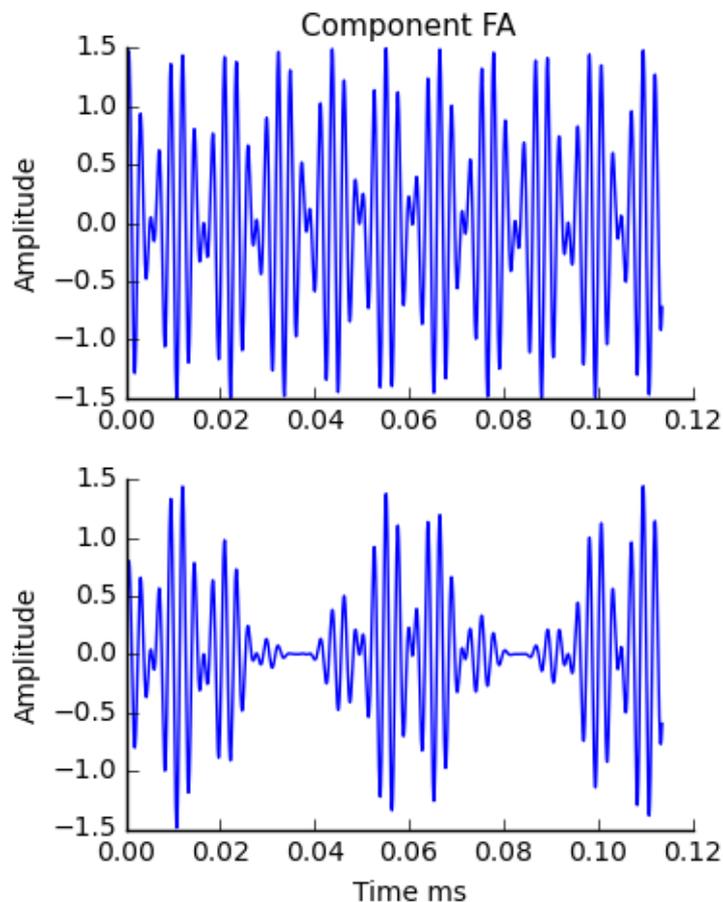


Figure 6.1. Component tone FA signal before amplitude modulation at 21Hz (top) and after (bottom).

the documentary was muted and subtitles were displayed. A brief presentation of the experiment was provided to ensure the participants understood what they were being asked to do. Following this, the documentary began playing and participants were instructed that the experiment was going to start.

To generate intermodulation responses, the pure tone components were amplitude modulated. Component tone FA was amplitude modulated at 21Hz, and CE

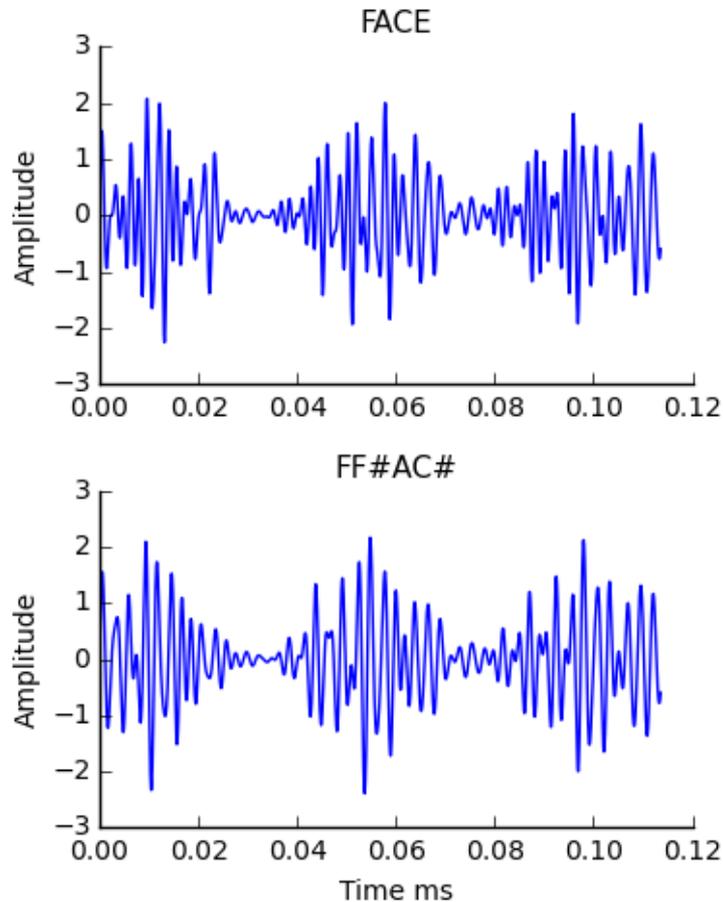


Figure 6.2. Chord FACE (top) and FF#AC# (bottom) signals. These represent a mixture of FA and CE signals (top) and FA and F#C# signals (bottom).

and F#C# were amplitude modulated at 24Hz. These frequencies were chosen on the basis that they allowed the response frequencies to be temporally incommensurate; the fundamental frequencies (21Hz, 24Hz), their harmonics (2f1: 42Hz, 2f2: 48Hz) and the intermodulation terms (f2-f1: 3Hz, f1+f2: 45Hz) would not overlap. Further, we aimed to avoid overlap between line noise at 50Hz and the sum intermodulation response frequency (f1+f2: 45Hz). The amplitude was modulated between 0 and 75% of the computer's maximal volume (which had a maximal Michelson contrast of 0.99).

A trial consisted of an 11 second presentation of an amplitude modulating tone or two simultaneously modulating, superimposed tones, followed by a 7-9 second inter-stimulus-interval. Each of the 5 stimuli was presented 10 times in a run, with 3

runs in total. Participants were given short breaks between runs. Upon completion, participants were thanked, debriefed and given the opportunity to ask any questions.

Component Responses

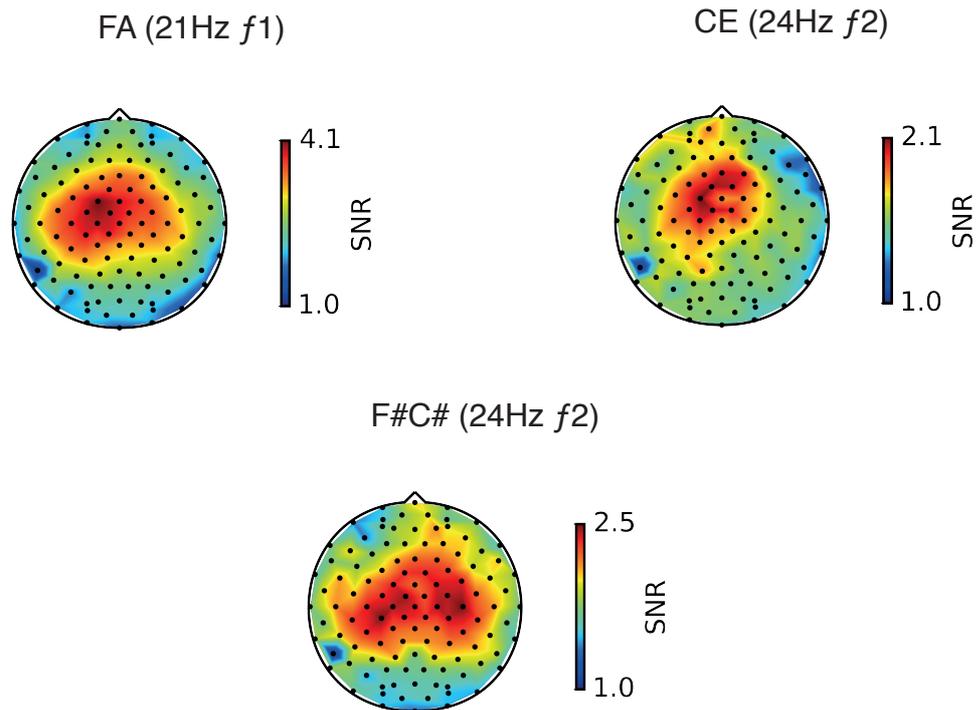


Figure 6.3. SNR topographies for components FA (upper left), CE (upper right) and F#C# (bottom).

Analytical Procedure. The data were band-pass filtered between 0.1Hz and 100Hz. The data were then epoched according to stimulus onset and the first second of data removed to exclude onset transients from the analysis, resulting in 10-second epoch for each trial. Trials were then time-averaged by condition for each participant, averaging out activity that is not time-locked to the stimulus presentation, such as the alpha-wave response. Fast Fourier transforms (FFTs) were then conducted on these average waveforms to bring the data into frequency space, resulting in amplitude responses (μV) at discrete frequencies (for a 10s stimulus the FFT has a resolution of 0.1Hz) between 0.1 and 100Hz.

The amplitude response at each frequency of interest at each electrode site was converted into a measure of signal-to-noise ratio (SNR) by dividing the amplitude

at the frequency of interest by the average amplitude of the surrounding 12 frequency bins. Based on group topographies (see Figure 6.3), the data used in the following analyses were taken from a cluster of electrode sites (Fz, FFC1h, FFC2h, FC1, FCz, FC2, FCC1h and FCC2h).

To determine whether SNRs were significantly above background noise (SNR=1), a series of one-sample t-tests were conducted separately for each set of component SNRs (i.e. FA, FA(CE) and FA(F[#]C[#]); CE and (FA)CE; F[#]C[#] and FF[#]AC[#]), difference intermodulation SNRs (at f_2-f_1 : FA, CE, F[#]C[#], FACE and FF[#]AC[#]), and sum intermodulation SNRs (at f_1+f_2 : FA, CE, F[#]C[#], FACE and FF[#]AC[#]). Each series of t-tests were corrected using the ranked Bonferroni-Holm method to control for Type 1 errors (Holm, 1979).

The extent to which pure tone/chord predicted response SNRs was examined using linear mixed-effects modelling. The analytical model was generated using the mixed function of the Afex package in R (Singmann et al., 2016). 'Stimulus' (levels: FA, CE, F[#]C[#], FACE and FF[#]AC[#]) was the only predictor with random slopes as a function of 'participant' using a maximal random effects structure (as recommended by Barr et al., 2013). This model was applied to both component SNRs and intermodulation SNRs. The lsmeans function in R for examining pairwise comparisons from linear mixed effects model structures was used when a significant main effect of pattern was found, and comparisons were corrected using the Tukey HSD method for multiple comparisons (Russell, 2016).

Results

General overview

We wanted to measure fundamental responses to the tone components, both alone and when forming part of a chord. We also wanted to measure intermodulation responses at the difference and sum of the fundamental frequencies used to modulate the amplitude of each component forming a chord. Clear component-based responses at f_1 (21Hz) and f_2 (24Hz) were observed for tones presented alone as well as when they were presented along with another tone component in the chord conditions. None of the chord stimuli resulted in robust intermodulation responses.

Response to tone components

As in the visual system, the two-frequency method meant that I could compare directly the response to each frequency-tagged tone component in isolation and in the presence of a tone. The SNR measures of these responses to each component – calculated by dividing the amplitude at the frequency of interest by the average amplitude of the surrounding 12 frequencies (with signal bins excluded) – can be seen

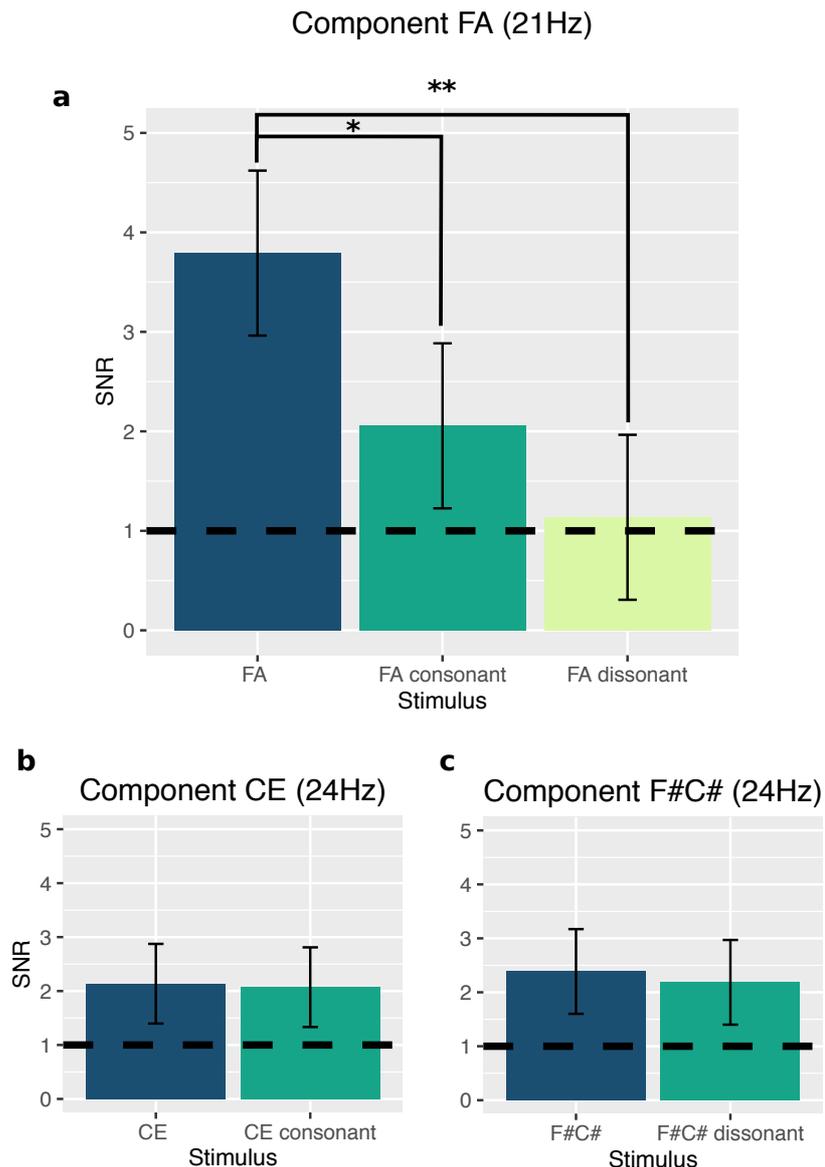


Figure 6.4. Component response SNRs for a) component FA, b) component CE, and c) component F#C#. In a) the leftmost bar represents SNR when FA was presented alone, the middle when part of consonant chord FACE, and the rightmost when part of dissonant chord FF#AC#. In b) and c) the left bar represents the SNR when that component was presented alone, and the right when presented as part of b) consonant chord FACE and c) when part of dissonant chord FF#AC#. Error bars represent 95% CIs, and the dashed line at 1 represents a response equivalent to background noise.

in Figure 6.4. Substantial suppression was observed for component FA presented at 21Hz in the presence of both CE (24Hz) and F[#]C[#] (24Hz). This was not the case for both CE and F[#]C[#].

Except for the 21Hz response to tone FA in chord FF[#]AC[#], all fundamental responses were significantly above background noise (at either $p < .05$ or $p < .01$). For the 21Hz response to tone FA, a significant main effect of Stimulus was found ($F_{2,30} = 13.79$, $p < .0001$). Post-hoc pairwise comparisons showed that tone FA alone resulted in significantly larger SNRs than when presented within chords FACE ($t_{30} = 3.38$, $p < .01$) and FF[#]AC[#] ($t_{30} = 5.17$, $p < .0001$). The 21Hz response SNR to FA when presented in chord FACE did not differ significantly from when presented in FF[#]AC[#]. There was no observable suppression for tones CE and F[#]C[#] when combined with tone FA (CE: $F_{1,15} = .03$, $p > .05$; F[#]C[#]: $F_{1,15} = .32$, $p > .05$).

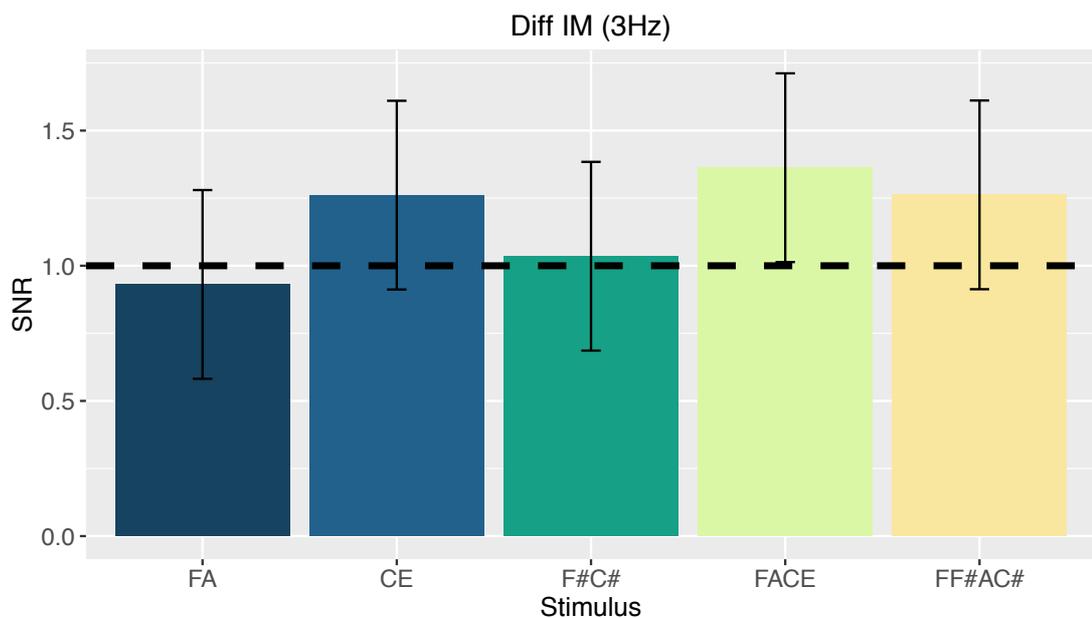


Figure 6.5. Response SNRs at the difference intermodulation term, 3Hz. The black dashed line represents background noise at an SNR of 1. Error bars represent 95% CIs.

In summary, response suppression was observed only for component FA when combined with the other two tone components. This suppression did not differ between the coherent and non-coherent tone combinations, indicating that the suppression was not frequency-tuned.

Response to chords

Intermodulation frequencies were used to assess responses to the conjunction of tone components. Responses at the difference intermodulation frequency (f_2-f_1 : 3Hz, Figure 6.5) and sum intermodulation frequency (f_1+f_2 : 45Hz, **Figure 6.6**) were not prominent compared to background noise for any stimulus. Likewise, no significant main effects were found for both (3Hz: $F_{4,60} = 1.05$, $p > .05$; 45Hz: $F_{4,60} = 1.52$, $p > .05$). These results suggest that very little, if any intermodulation response occurred to the pairs of amplitude modulated tones.

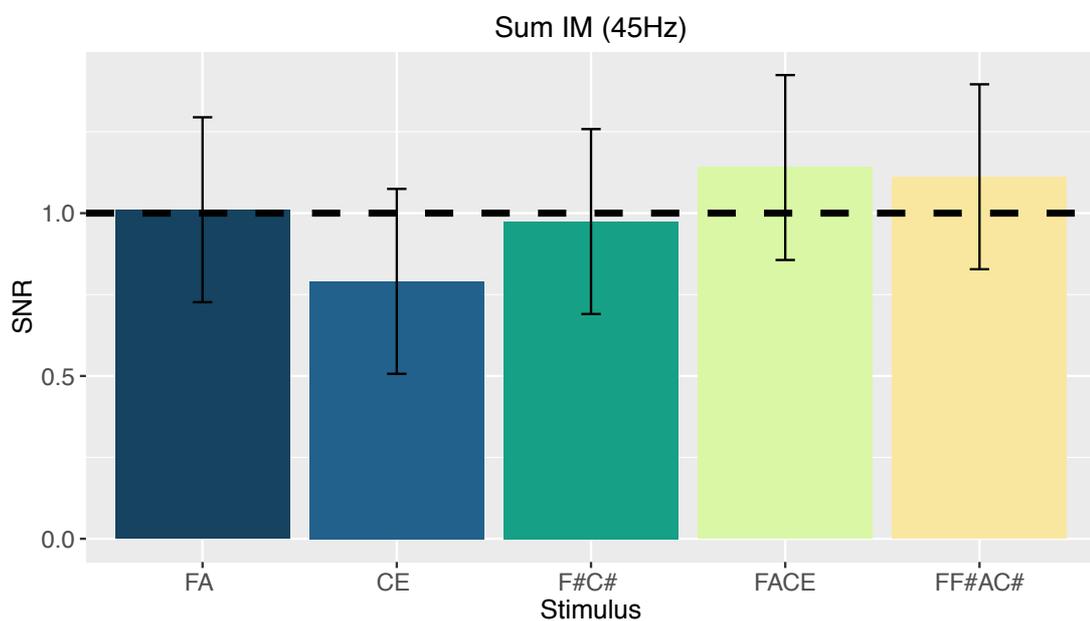


Figure 6.6. **Response SNRs at the sum intermodulation term, 45Hz.** The black dashed line represents background noise at an SNR of 1. Error bars represent 95% CIs.

Discussion

I have measured neural responses to amplitude-modulated tones presented as simple 2-tone combinations alone and then combined to form a coherent (FACE) and non-coherent (FF[#]AC[#]) chord. This was done to assess the nonlinear combination of neural responses to the tones. EEG responses were measured at intermodulation frequencies to specifically investigate nonlinear combinations. Contrary to the previous studies described in this thesis that examined nonlinearities in the visual system, robust intermodulation responses were not observed here in the auditory domain. However,

substantial component suppression was observed for tone FA, and this did not depend on the chord coherence.

The auditory masking is not a surprising finding; it has been previously established that auditory response masking occurs (Abbas & Sachs, 1976; de la Rocha, Marchetti, Schiff, & Reyes, 2008; Phillips, 1990; Sachs, S Kiang, Sac, & S, 1968), and can be described well by the normalisation equation (Carandini & Heeger, 2012; Rabinowitz et al., 2011). Substantial suppression was observed only for component FA, and this did not differ in magnitude as a function of compound coherence. Given the observed response suppression, it was surprising to find no intermodulation response. Intermodulation responses are generated at the point of or just after a nonlinear signal combination has taken place (Spekreijse & Oosting, 1970a; Zemon & Ratliff, 1984), such as lateral suppression between neural populations tuned to different frequencies (Baitch & Levi, 1988; T. J. Baker, Norcia, & Candy, 2011; Candy et al., 2001; Suter et al., 1996; Tsai et al., 2012).

Intermodulation responses may simply have a lower signal-to-noise ratio in the auditory system than in the visual system, and mass-measurements of neural responses through EEG might not be sensitive enough to capture them. We have consistently observed low signal-to-noise ratios for intermodulation responses in the visual system, certainly smaller than component-based responses to the gratings forming plaids. This is also a common observation in the literature (e.g. Appelbaum et al., 2008b; Boremanse et al., 2013; Boremanse, Norcia, & Rossion, 2014; Sutoyo & Srinivasan, 2009).

Alternatively, the auditory system might also not express nonlinearities in the same way as the visual system does. There were no noticeable component-based harmonic responses in this study, but we observed large harmonic responses in the visual system in the previous chapters. We wondered whether the sounds were simply not loud enough, such that the summation was occurring in the linear part of the response curve (moderate volumes had been used in the data collected here). To test this, we re-ran the study at 100% volume but, again, no intermodulation responses nor harmonic responses were observed.

A potential confound of using amplitude-modulating tones to study coherent chords is that the amplitude modulation process actually adds power at additional

frequencies. For instance, by taking an A tone (440Hz) and modulating it at 24Hz results in a spectrum that includes power, not only at 440Hz, but also 416 and 464Hz. As a result the amplitude-modulated chords did not sound as clearly coherent, or as discordant, as was intended. So, though plaid coherence appears to persist in the visual system, despite the temporal modulation that is necessary to make this method work, the same might not be true of chord coherence in the auditory system.

We attempted to delineate responses to component tones from compound responses at intermodulation frequencies to chord stimuli. In this instance, we have found that the method used in the two previous chapters does not translate directly, and more work is needed to refine the stimuli and experimental procedure used to examine nonlinear combinations in the auditory cortex in the same way as in the visual cortex.

Chapter 7: General Discussion

This thesis has examined nonlinear signal summation using a combination of EEG and computational modelling. Nonlinearities are essential to many perceptual phenomena, but remain poorly understood beyond the earliest levels of the sensory pathways.

Many nonlinear physiological phenomena, such as XOS, can be readily described by models of normalisation for neuronal gain control in V1. However, there are several nonlinearities that normalisation cannot fully explain. For example, super saturation – which can occur in around 17% of V1 and 25% of V2 neurons in macaque (Peirce, 2007b) – would be considered metabolically wasteful within a framework of normalisation: an over-exertion of the normalisation pool upon the excitatory response of a neuron. It seems unlikely that this non-monotonic nonlinearity does not serve a purpose. Considering this, gain control may not be the only function served by nonlinearities in the visual system (and beyond).

Peirce (2007b, 2011, 2013) proposed that nonlinearities in V1 could also be used by neurons in mid-level vision to detect signal conjunctions for combinations of stimuli. This kind of signal summation would make possible neurons with more complex receptive field preferences than are commonly observed in V1. For example, neurons that are sensitive to multiple orientations and a narrow bandwidth of spatial frequencies would be useful for detecting patterns coherent plaids.

However, at any one point in time, several different nonlinearities can occur in response to a stimulus. Being able to distinguish one from the other is more difficult than it might at first seem. The experiments described throughout this thesis aimed to disentangle nonlinearities, identify those that were selective for specific stimulus combinations and characterise them.

Summary of findings

In Chapter 3 we used transient EEG to measure the earliest component – C1 – of VEPs to brief presentations of gratings and their combinations into coherent and non-coherent plaids. By comparing the C1 response to gratings and plaids, we aimed to measure the degree of nonlinear summation taking place for coherent and non-

coherent grating combinations. The outcome was inconclusive; there was limited evidence to suggest the involvement of extra nonlinearities in the processing of coherent plaids that were not involved in processing non-coherent plaids. This might be an inherent problem with the transient EEG approach. Although it produces a rich time course of data following the presentation of a stimulus, the response is the sum of many nonlinearities.

To overcome this, we took an alternative approach in Chapter 4 and used the two-frequency method of steady-state EEG. This allows you to tag each of the grating components forming a plaid, as well as directly measure nonlinearities at intermodulation frequencies. We found a plaid-selective intermodulation response, which was larger for coherent plaids than it was for non-coherent plaids. In support of this representing an additional nonlinearity beyond normalisation, the degree of component suppression did not differ between coherent and non-coherent plaids for any of the grating components used.

We generated a simple two-layered computational model of signal summation to examine the complexity of responses generated in to combinations of gratings. The channels were not made selective to spatial frequency or orientation, but the model was structured to operate like a logical AND gate. It appears that this kind of mechanism can represent well the responses we observed using EEG.

It is not clear how a mechanism that makes use of saturating nonlinearities to perform selective signal summation would behave across contrast. At lower contrast levels, before many neurons reach the rising slope of their dynamic range, it might be that the mechanism fails altogether. Using a similar paradigm to Chapter 4, we measured intermodulation responses across a wide range of contrast levels in Chapter 5. We again found a selective intermodulation response that was larger for coherent plaids. However, this difference only occurred at the highest component contrast level that we used – 32%.

Having found a nonlinearity in the visual system that appeared to be selective for particular combinations of grating stimuli, we wanted to investigate whether similar nonlinearities are put to use in other brain regions. In Chapter 6 we generated auditory stimuli – three pure tones – that were combined to form a consonant and a dissonant chord. Substantial component suppression was observed for one of the

components. However, no intermodulation responses or component-based harmonic responses were observed.

Implications

The two-frequency method is extremely useful for disentangling different nonlinearities in the visual system. As discussed in Chapter 2, methods like transient EEG recording, fMRI and psychophysics measure the sum of many nonlinearities across many hundreds of thousands of neurons. This makes it difficult to attribute those response measures to specific processes, especially when additional nonlinearities might represent a very small proportion of the summed response (i.e. in the case where only small populations of neurons are performing additional selective signal summations). Here, the two-frequency method allowed us to measure responses to each of the components in compound plaids, giving us a direct measure of how the presence of one affected the response to the other. The technique also enabled us to measure intermodulation responses, and we were therefore able to distinguish XOS from any additional nonlinearities that might be taking place. Taking this a step further, Sutoyo and Srinivasan (2009) flickered four separate components simultaneously to analyse nonlinearities to the various combinations that took place monocularly and interocularly.

Certainly, a major role that nonlinearities in the early visual system play is to adjust the gain of neuronal responses, but as suggested by the results in Chapters 4 and 5 that may not be their only purpose. Twice we showed a nonlinearity that was larger when the grating components forming a plaid were matched in spatial frequency (coherent) than when they were not (non-coherent). The degree of component suppression was similar for all components in Chapter 4, irrespective of whether they formed a coherent or non-coherent plaid. Therefore, XOS alone could not readily explain the larger intermodulation response for coherent plaids. In Chapter 5, there was more component suppression observed for a component in the non-coherent plaid condition, rather than in the coherent plaid condition. Again, this is difficult to interpret within a framework of XOS alone. However, it could be explained by an additional nonlinearity involved in summing responses to the plaid components,

in line with the model output in Chapter 4. Since in both cases the results could not be explained by differences in component suppression, the combined results of Chapters 4 and 5 suggest that nonlinear summation at least mostly contributes to intermodulation responses and contributes very little to component suppression. Further, Chapter 5 suggests that the mechanism performing nonlinear summation of plaid components does not operate at lower component contrast levels. This is in-line with the concept of logical AND gates making use of saturating nonlinearities. The input threshold to the summing mechanism would need to be higher than that which any one component can produce on its own.

Based on the results in Chapter 6, this does not appear to take place in the auditory system. Our results would suggest that auditory system nonlinearities behave differently from those in the visual system. For example, we observed clear component suppression for component FA when combined with either of the other pure tone components, but saw no intermodulation response (or even component-based harmonic responses). It seems more likely, however, that the stimuli that we generated confounded the measurements made. By amplitude-modulating frequency-modulated signals (i.e. the tones), new frequencies are introduced to the signal that can alter the sound of the tone. So while the percept of the coherent plaid persists when you contrast modulate its components, the same is not necessarily true for amplitude modulating the components of a consonant chord.

Relationship to the previous literature

XOS complexity

We observed two different patterns of component suppression in Chapters 4 and 5. The components that were contrast modulated at 2.3Hz in Chapter 4 displayed suppression, while those modulated at 3.75Hz did not. This suggests that, irrespective of spatial frequency, higher temporal frequency components exhibit stronger inhibition on responses to lower temporal frequency components. Again, in Chapter 5 where the stimulus components were modulated at 4.6Hz and 7.5Hz, the higher frequency (7.5Hz) components showed no suppression. However, no suppression was observed for component A1 when combined with component A2, but substantial

suppression was observed when combined with B2. Contrary to Chapter 4, this would suggest that when maskers are higher in spatiotemporal frequency relative to the other stimulus component, XOS becomes spatial frequency tuned.

DeAngelis, Freeman and Ohzawa (1994) and Petrov, Carandini and McKee (2005) found that XOS was broadly tuned for spatial frequency (up to a factor of ~ 4 difference between target and mask SF: Petrov et al., 2005), indicating that the mechanism involved receives spatial frequency-related inhibition from a wide broadly-tuned pool of cortical neurons. However, several studies (Cass et al., 2009; Meese & Holmes, 2007, 2010; Snowden, 1992) have found that XOM occurs more strongly when superimposed grating components are of a lower spatial frequency (< 2 cpd: Meese & Holmes, 2007) and presented at a higher temporal frequency (> 4 Hz: Meese & Holmes, 2007). This might reflect different routes to suppression in the visual system (e.g. D. H. Baker et al., 2007; Cass et al., 2009; Meese & Holmes, 2007, 2010; Petrov et al., 2005; Viswanathan et al., 2011). Before full binocular summation takes place in cortex, there is evidence of a subcortical component of XOS broadly tuned for orientation that operates primarily at lower spatial frequencies and higher temporal frequencies. This propagates to cortex where a second component with the same sharp tuning as neuronal excitation strengthens the overall suppressive effect.

This is almost in line with the pattern of suppression observed for components A1 and B1 in Chapter 4, which had spatial frequencies of 1cpd and 3cpd. Our results would suggest that what mattered for suppression in Chapter 4 was the temporal frequency of the component acting as the masker. The suppression observed in Chapter 5 might reflect a complex interaction between the spatial and temporal frequency tuning of XOS. However, it is difficult to interpret the effect that temporal frequency played in our results relative to the findings of Meese and Holmes (2007) and Cass et al., (2009) because they presented both components at the same temporal frequency. Further, our stimuli were presented binocularly and foveally, whereas they had multiple viewing conditions and Cass et al., (2009) so understanding at which point different forms of suppression impacted the responses that we observed is difficult.

The two-frequency EEG method would be well-suited to investigating this further. Using a wide array of conditions across temporal and spatial frequency, and

monocular, dichoptic and binocular viewing conditions in a style similar to Sutoyo and Srinivasan (2009), one could carry out a thorough investigation of different routes to suppression (D. H. Baker et al., 2007; Priebe & Ferster, 2006) and their dependence on orientation, spatial frequency and temporal frequency.

Intermodulation responses

Our results from Chapters 4 and 5 are in general agreement with most studies measuring intermodulation responses to compound stimulus configurations; they found larger intermodulation responses for both spatially coincidental and non-coincidental ‘compound’ conditions than ‘component’ conditions (Aissani et al., 2011; Alp et al., 2016; Boremanse et al., 2013, 2014; Gundlach & Müller, 2013; Hou, Pettet, Sampath, Candy, & Norcia, 2003; Sutoyo & Srinivasan, 2009). Alp et al., (2016) examined the perception of illusory surfaces using Kanizsa stimuli – a spatially separated stimulus configuration (i.e. the components are in different regions of the visual scene – and found larger intermodulation responses at f_1+f_1 and $2f_1+f_2$ when the four Kanizsa components formed an illusory square. Similarly, Gundlach and Müller (2013) found larger intermodulation responses at f_1+f_2 when Kanizsa stimuli formed an illusory bar. Aissani et al., (2011) found that when moving bars arranged in a square shape were manipulated to appear ‘bounded’ (i.e. to form the percept of a single object), larger intermodulation responses occurred at $2f_1+2f_2$. Boremanse, Norcia and Rossion (2013, 2014) found that responses at different intermodulation frequencies like f_2-f_1 and $2f_2-2f_1$ were larger when face-halves were vertically and horizontally aligned such that they formed a whole face, and when both halves belonged to the same person. It would be interesting to use the intermodulation approach to examine surround suppression using and add to Petrov et al’s (2005) contribution to our understanding of the functionality of surround suppression.

Combined, this provides support for the use of two-frequency method to measure responses to compound stimuli, and for the idea that the visual system is equipped with mechanisms that are tuned to selectively encode patterns like coherent plaids (Adelson & Movshon, 1982; Hancock et al., 2010; Nam et al., 2009). A mechanism selective for coherent plaids in combinations with mechanisms of XOS

appears to explain better the nonlinear intermodulation responses we have measured than XOS alone. In Chapter 4 we found that the addition of ‘conjunction detector’ channels beyond contrast gain operations resulted in model output more like the EEG data. Additionally, in Chapter 5 we found that responses to the coherent compound plaid at f_2-f_1 largely followed what would be expected of a mechanism making use of saturating nonlinearities to selectively detect signal conjunctions. A logical AND operation performing nonlinear summation in combination with XOS cannot be ruled out as a candidate mechanism for driving plaid-selective responses at intermodulation frequencies. Using a model that operated like a logical AND gate in Chapter 4, we captured differences in amplitude between difference and sum intermodulation frequencies that were also observed in the measured EEG data and are widely observed in the literature. A simple nonlinear combination, such as a squaring nonlinearity on the linear sum of V1 signals, produces equal amplitude at the difference and sum frequencies.

Another way in which signals might be nonlinearly combined is by multiplication of output signals (Gheorghiu & Kingdom, 2009). There is evidence of multiplicative combinations of signals in the fly (Gabbiani, Krapp, Koch, & Laurent, 2002) and rabbit (Taylor, He, Levick, & Vaney, 2000), but not in primate. It would seem easier to make use of saturating nonlinearities that already exist in the visual system to perform nonlinear summation as opposed to multiplication. Pierce (2011) points out that though multiplication and nonlinear summation seem computationally different, they in effect lead to the same neuronal decision to fire (or not). Further, though we did not observe clear plaid-selective responses, the results from Chapter 3 did not show a large increase in response to coherent plaids that might be expected from a mechanism performing multiplication. Instead we saw responses that were slightly less-than-linear for coherent plaids, and responses that resembled the average response to both components for the non-coherent plaid.

It is intriguing that the two-frequency studies referenced above have observed intermodulation responses at either the difference or sum intermodulation terms for a particular type of stimulus. One interpretation for differences between the intermodulation terms suggested by Boremanse et al. (2013) is that responses at both frequencies reflect parallel nonlinearities, but the sum intermodulation response

output may be a temporally band- or high-pass nonlinearity and signal early local spatial interactions. On the other hand, the difference intermodulation response may be generated by a temporally low-pass nonlinearity and generated by signal integration to higher-level (global) stimuli, such as their face-part stimuli, which require longer to process (Alonso-Prieto et al., 2013). In Chapter 4 we found that the application of a bandpass temporal filter at Layer 1 and a higher frequency filter at Layer 2 of our model resulted in better model output than by applying a bandpass filter at both layers. Rapid local combinations would certainly fit in with our EEG results for Chapter 4 as we used simple sinusoidal gratings that presumably were being combined across many receptive fields to encode another pattern (the plaid). However, in Chapter 5 when the temporal frequencies used were doubled to 4.6Hz and 7.5Hz, our main finding was larger responses to coherent plaids at the difference intermodulation term f_2-f_1 .

Alp et al., (2016) suggested that the temporal resonance properties of different neural mechanisms may influence the varied response at the difference and sum intermodulation frequencies. These resonances may depend on specific synaptic connections to and from the mechanisms, feedback connectivity and the relative complexity of the receptive field within the visual hierarchy (e.g. sensitive to compound plaids or sensitive to faces). The effects of such differences in temporal integration have not been applied quantitatively in a computational model (e.g. to explain differential responses at sum- and difference-intermodulation terms). In Chapter 4 we used a simple approach to model the temporal properties of mechanisms and found that using different neural impulse response functions (temporal filters) at early and late layers was sufficient to explain a wide range of features in the data. It appeared that responses at 6.05Hz (f_1+f_2) in Chapter 4 primarily represented XOS mechanisms, while responses at 12.1Hz ($2f_1+2f_2$) reflected plaid-selective mechanisms. However, in Chapter 5 using plaid stimuli again we observed robust responses only at the difference intermodulation term f_2-f_1 .

It is therefore difficult to reconcile what might be influencing the prominence of responses at either the difference or the sum terms. What is an exciting prospect is that the many nonlinearities that appear when you use the two-frequency method might not simply represent phenomenological noise, but meaningful responses. By

carrying out an in-depth study into how the frequency range (e.g. using a range of component frequencies between 1Hz and 8Hz) and frequency distance (i.e. the distance between f_1 and f_2 in Hz) affect responses to compound stimuli like plaids, one might come closer to understanding the complex temporal dependencies involved in generating these responses. If what Boremanse et al., (2013, 2014) suggested is accurate, then the resulting Fourier spectrums might be different depending on the complexity of the compound stimulus being presented.

Conclusion

The visual system makes use of nonlinearities for more than just contrast gain control. Mechanisms exist that selectively detect signal conjunctions for certain stimulus combinations, such as coherent plaids, by making use of the saturating nonlinearities that are prevalent throughout V1. These mechanisms can be measured using the two-frequency steady-state EEG approach.

The potential of this approach to understanding nonlinearities is vast. Developing paradigms around the two-frequency methodology will prove invaluable to understanding what kind of mechanisms exist in mid-level vision and how they operate, as well as to further our understanding of the complexities involved in encoding multiple stimulus components at earlier levels of the visual system.

References

- Abbas, P. J., & Sachs, M. B. (1976). Two-tone suppression in auditory-nerve fibers: Extension of a stimulus-response relationship Two-tone suppression in auditory-nerve fibers: Extension of a stimulus-response relationship*. *The Journal of the Acoustical Society of America*, *59*(112).
<http://doi.org/10.1121/1.380841>
- Adelson, E. H., & Movshon, J. A. (1982). Phenomenal coherence of moving visual patterns. *Nature*, *300*(5892), 523–525.
- Ahrens, M. B., Linden, J. F., & Sahani, M. (2008). Nonlinearities and contextual influences in auditory cortical responses modeled with multilinear spectrotemporal methods. *The Journal of Neuroscience*, *28*(8), 1929–1942.
<http://doi.org/10.1523/JNEUROSCI.3377-07.2008>
- Aissani, C., Cottureau, B. R., Dumas, G., Paradis, A.-L., & Lorenceau, J. (2011). Magnetoencephalographic signatures of visual form and motion binding. *Brain Research*, *1408*, 27–40. <http://doi.org/10.1016/j.brainres.2011.05.051>
- Albrecht, D. G., & DeValois, R. L. (1981). Striate cortex responses to periodic patterns with and without the fundamental harmonics. *The Journal of Physiology*, *319*(1), 497–514. <http://doi.org/10.1113/jphysiol.1981.sp013922>
- 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.
<http://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), 217–37. Retrieved from <http://jn.physiology.org/content/48/1/217.abstract>
- Ales, J. M., Yates, J. L., & Norcia, A. M. (2010). V1 is not uniquely identified by polarity reversals of responses to upper and lower visual field stimuli. *NeuroImage*, *52*(4), 1401–1409.
<http://doi.org/10.1016/j.neuroimage.2010.05.016>
- Ales, J. M., Yates, J. L., & Norcia, A. M. (2013). On determining the intracranial

sources of visual evoked potentials from scalp topography: A reply to Kelly et al. (this issue). *NeuroImage*. <http://doi.org/10.1016/j.neuroimage.2012.09.009>

Allison, J. D., Smith, K. R., & Bonds, A. B. (2001). Temporal-frequency tuning of cross-orientation suppression in the cat striate cortex. *Visual Neuroscience*, *18*(6), 941–948. <http://doi.org/null>

Alonso-Prieto, E., Belle, G. Van, Liu-Shuang, J., Norcia, A. M., & Rossion, B. (2013). The 6Hz fundamental stimulation frequency rate for individual face discrimination in the right occipito-temporal cortex. *Neuropsychologia*, *51*(13), 2863–2875. <http://doi.org/10.1016/j.neuropsychologia.2013.08.018>

Alp, N., Kogo, N., Van Belle, G., Wagemans, J., & Rossion, B. (2016). Frequency tagging yields an objective neural signature of Gestalt formation. *Brain and Cognition*, *104*, 15–24. <http://doi.org/10.1016/j.bandc.2016.01.008>

Andersen, S. K., Müller, M. M., & Hillyard, S. A. (2015). Attentional Selection of Feature Conjunctions Is Accomplished by Parallel and Independent Selection of Single Features. *Journal of Neuroscience*, *35*(27), 9912–9919. <http://doi.org/10.1523/JNEUROSCI.5268-14.2015>

Anzai, A., Peng, X., & Van Essen, D. C. (2007). Neurons in monkey visual area V2 encode combinations of orientations. *Nature Neuroscience*, *10*(10), 1313–21. <http://doi.org/10.1038/nn1975>

Appelbaum, L. G., & Norcia, A. M. (2009). Attentive and pre-attentive aspects of figural processing. *Journal of Vision*, *9*(11), 18.1-12. <http://doi.org/10.1167/9.11.18>

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

Appelbaum, L. G., Wade, A. R., Pettet, M. W., Vildavski, V. Y., & Norcia, A. M. (2008b). Figure – ground interaction in the human visual cortex. *Journal of Vision*, *8*(9), 1–19. <http://doi.org/10.1167/8.9.8.Introduction>

Appelbaum, L. G., Wade, A. R., Vildavski, V. Y., Pettet, M. W., & Norcia, A. M. (2006). Cue-Invariant Networks for Figure and Background Processing in Human Visual Cortex. *The Journal of Neuroscience*, *26*(45), 11695–11708. <http://doi.org/10.1523/JNEUROSCI.2741-06.2006>

- Baas, J. M. P., Kenemans, J. L., & Mangun, G. R. (2002). Selective attention to spatial frequency: An ERP and source localization analysis. *Clinical Neurophysiology*, *113*(11), 1840–1854. [http://doi.org/10.1016/S1388-2457\(02\)00269-9](http://doi.org/10.1016/S1388-2457(02)00269-9)
- Baccus, S. A., & Meister, M. (2002). Fast and Slow Contrast Adaptation in Retinal Circuitry. *Neuron*, *36*, 909–919.
- Baitch, L. W., & Levi, D. M. (1988). Evidence for nonlinear binocular interactions in human visual cortex. *Vision Research*, *28*(10), 1139–1143. [http://doi.org/10.1016/0042-6989\(88\)90140-X](http://doi.org/10.1016/0042-6989(88)90140-X)
- 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), 435–48. <http://doi.org/10.1016/j.neuroscience.2007.01.030>
- Baker, D. H., & Wade, A. R. (2017). Evidence for an Optimal Algorithm Underlying Signal Combination in Human Visual Cortex. *Cerebral Cortex*, (January 2017), 254–264. <http://doi.org/10.1093/cercor/bhw395>
- Baker, T. J., Norcia, A. M., & Candy, T. R. (2011). Orientation tuning in the visual cortex of 3-month-old human infants. *Vision Research*, *51*(5), 470–8. <http://doi.org/10.1016/j.visres.2011.01.003>
- Baker, T. J., Norcia, A. M., & Rowan Candy, T. (2011). Orientation tuning in the visual cortex of 3-month-old human infants. *Vision Research*, *51*(5), 470–478. <http://doi.org/10.1016/j.visres.2011.01.003>
- Baldwin, A. S., Schmidtman, G., Kingdom, F. A. A., & Hess, R. F. (2015). Rejecting probability summation for RF patterns, not so Quick!, *5*(1), 1997. <http://doi.org/10.1016/j.visres.2016.03.003>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278. <http://doi.org/10.1016/j.jml.2012.11.001>
- Bell, J., Sampasivam, S., McGovern, D. P., Meso, A. I., & Kingdom, F. A. A. (2014). Contour inflections are adaptable features. *Journal of Vision*, *14*(7), 1–14. <http://doi.org/10.1167/14.7.2.doi>
- Benardete, E. A., & Kaplan, E. (1999). The dynamics of primate M retinal ganglion

cells. *Visual Neuroscience*, 16(2), 355–368.
<http://doi.org/10.1017/S0952523899162151>

Bidelman, G. M., & Grall, J. (2014). Functional organization for musical consonance and tonal pitch hierarchy in human auditory cortex.
<http://doi.org/10.1016/j.neuroimage.2014.07.005>

Blakemore, C., & Campbell, F. W. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology*, 203, 237–260.

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.
<http://doi.org/10.1017/S0952523800004314>

Bonin, V., Mante, V., & Carandini, M. (2005). The Suppressive Field of Neurons in Lateral Geniculate Nucleus. *Journal of Neuroscience*, 25(47), 10844–10856.
<http://doi.org/10.1523/JNEUROSCI.3562-05.2005>

Boremanse, A., Norcia, A. M., & Rossion, B. (2013). An objective signature for visual binding of face parts in the human brain. *Journal of Vision*, 13(11), 6.
<http://doi.org/10.1167/13.11.6>

Boremanse, A., Norcia, A. M., & Rossion, B. (2014). Dissociation of part-based and integrated neural responses to faces by means of electroencephalographic frequency tagging. *The European Journal of Neuroscience*, 40(November 2013), 1–12. <http://doi.org/10.1111/ejn.12663>

Brouwer, G. J., & Heeger, D. J. (2011a). Cross-orientation suppression in human visual cortex. *Journal of Neurophysiology*, 106(5), 2108–2119.
<http://doi.org/10.1152/jn.00540.2011>

Brouwer, G. J., & Heeger, D. J. (2011b). Cross-orientation suppression in human visual cortex. *Journal of Neurophysiology*, 106(5), 2108–2119.
<http://doi.org/10.1152/jn.00540.2011>

Brown, R. J., Candy, T. R., & Norcia, A. M. (1999). Development of rivalry and dichoptic masking in human infants. *Investigative Ophthalmology and Visual Science*, 40(13), 3324–3333.

Burr, D. C., & Morrone, C. (1987). Inhibitory interactions in the human vision system

- revealed in pattern-evoked potentials. *The Journal of Physiology*, 389(1), 1–21.
<http://doi.org/10.1113/jphysiol.1987.sp016643>
- Busse, L., Wade, A. R., & Carandini, M. (2009). Representation of concurrent stimuli by population activity in visual cortex. *Neuron*, 64(6), 931–42.
<http://doi.org/10.1016/j.neuron.2009.11.004>
- Butler, S. R., Georgiou, G. A., Glass, A., Hancox, R. J., Hopper, J. M., & Smith, K. R. (1987). Cortical generators of the CI component of the pattern-onset visual evoked potential. *Electroencephalography and Clinical Neurophysiology/ Evoked Potentials*, 68(4), 256–267. [http://doi.org/10.1016/0168-5597\(87\)90046-3](http://doi.org/10.1016/0168-5597(87)90046-3)
- Campbell, F. W., Cooper, G. F., & Enroth-Cugell, C. (1969). The spatial selectivity of the visual cells of the cat. *Journal of Physiology*, 203, 223–235.
- Campbell, F. W., & Maffei, L. (1970). Electrophysiological evidence for the existence of orientation and size detectors in the human visual system. *The Journal of Physiology*, 207(3), 635–652. <http://doi.org/10.1113/jphysiol.1970.sp009085>
- Campbell, F. W., & Robson, J. G. (1968). Application of Fourier analysis to the visibility of gratings. *Journal of Physiology*, 197, 551–566.
- Candy, T. R., Skoczenski, A. M., & Norcia, A. M. (2001). Normalization Models Applied to Orientation Masking in the Human Infant. *The Journal of Neuroscience*, 21(12), 4530–4541.
- Carandini, M., & Heeger, D. J. (1994). Summation and division by neurons in primate visual cortex. *Science*, 264(5163), 1333–1336.
<http://doi.org/10.1126/science.8191289>
- Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1), 51–62.
<http://doi.org/10.1038/nrn3136>
- Carandini, M., Heeger, D. J., & Movshon, J. A. (1997). Linearity and Normalization in Simple Cells of the Macaque Primary Visual Cortex. *The Journal of Neuroscience*, 17(21), 8621–8644. Retrieved from
<http://www.jneurosci.org/content/17/21/8621>
- Cass, J., & Alais, D. (2006). Evidence for two interacting temporal channels in human visual processing. *Vision Research*, 46(18), 2859–2868.

<http://doi.org/10.1016/j.visres.2006.02.015>

- Cass, J., Stuit, S., Bex, P. J., & Alais, D. (2009). Orientation bandwidths are invariant across spatiotemporal frequency after isotropic components are removed. *Journal of Vision*, *9*(12), 17.1-14. <http://doi.org/10.1167/9.12.17>
- Chen, J., Yu, Q., Zhu, Z., Peng, Y., & Fang, F. (2016). Spatial summation revealed in the earliest visual evoked component C1 and the effect of attention on its linearity. *J Neurophysiol*, *115*, 500–509. <http://doi.org/10.1152/jn.00044.2015>
- Clark, V. P., Fan, S., & Hillyard, S. A. (1995). Identification of early visual evoked potential generators by retinotopic and topographic analyses. *Human Brain Mapping*, *2*(3), 170–187. <http://doi.org/10.1002/hbm.460020306>
- Cobb, W. A., Morton, H. B., & Ettliger, G. (1967). Cerebral Potentials evoked by Pattern Reversal and their Suppression in Visual Rivalry. *Nature*, *216*(5120), 1123–1125. <http://doi.org/10.1038/2161123b0>
- Coen-Cagli, R., & Schwartz, O. (2013). The impact on midlevel vision of statistically optimal divisive normalization in V1. *Journal of Vision*, *13*(8), 1–20. <http://doi.org/10.1167/13.8.13>
- Colon, E., Legrain, V., & Mouraux, A. (2014). EEG Frequency Tagging to Dissociate the Cortical Responses to Nociceptive and Nonnociceptive Stimuli. *Journal of Cognitive Neuroscience*, *26*(10), 2262–2274. http://doi.org/10.1162/jocn_a_00648
- Compston, A. (2010). The Berger rhythm: potential changes from the occipital lobes in man, by E.D. Adrian and B.H.C. Matthews (From the Physiological Laboratory, Cambridge). *Brain* 1934: 57; 355-385. *Brain*, *133*(1), 3–6. <http://doi.org/10.1093/brain/awp324>
- Cunningham, D. G. M., Baker, D. H., & Peirce, J. W. (2017). Measuring nonlinear signal combination using EEG. *Journal of Vision*, *17*(5), 1–14. <http://doi.org/10.1167/17.5.10>
- David, S. V., Mesgarani, N., Fritz, J. B., & Shamma, S. A. (2009). Rapid Synaptic Depression Explains Nonlinear Modulation of Spectro-Temporal Tuning in Primary Auditory Cortex by Natural Stimuli. *Journal of Neuroscience*, *29*(11). Retrieved from <http://www.jneurosci.org/content/29/11/3374.short>

- de la Rocha, J., Marchetti, C., Schiff, M., & Reyes, A. D. (2008). Linking the response properties of cells in auditory cortex with network architecture: cotuning versus lateral inhibition. *J Neurosci*, *28*(37), 9151–9163. <http://doi.org/28/37/9151> [pii]r10.1523/JNEUROSCI.1789-08.2008
- DeAngelis, G. C., Freeman, R. D., & Ohzawa, I. (1994). Length and width tuning of neurons in the cat's primary visual cortex. *J Neurophysiol*, *71*(1), 347–374.
- 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), 144–63. Retrieved from <http://jn.physiology.org/content/68/1/144.short>
- Desimone, R. (1991). Face-Selective Cells in the Temporal Cortex of Monkeys. *Journal of Cognitive Neuroscience*, *3*(1), 1–8. <http://doi.org/10.1162/jocn.1991.3.1.1>
- Desimone, R., Albright, T. D., Gross, C. G., & Bruce, C. (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. *The Journal of Neuroscience*, *4*(8), 2051–2062. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/6470767>
- Di Russo, F., Martínez, A., Sereno, M. I., Pitzalis, S., & Hillyard, S. A. (2002). Cortical sources of the early components of the visual evoked potential. *Human Brain Mapping*, *15*(2), 95–111. <http://doi.org/10.1002/hbm.10010>
- Ellemberg, D., Hammarrenger, B., Lennie, P., Roy, M.-S., & Guillemot, J. P. (2001). Contrast dependency of VEPs as a function of spatial frequency: the parvocellular and magnocellular contributions to human VEPs. *Spatial Vision*, *15*(1), 99–111.
- Enroth-Cugell, C., & Robson, J. G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *Journal of Physiology*, *187*, 517–552.
- Fishman, Y. I., Volkov, I. O., Noh, M. D., Garell, P. C., Bakken, H., Arezzo, J. C., ... Steinschneider, M. (2001). Consonance and Dissonance of Musical Chords: Neural Correlates in Auditory Cortex of Monkeys and Humans. *Journal of Neurophysiology*, *86*(6), 2761–2788. <http://doi.org/10.1152/jn.2001.86.6.2761>
- Fortune, B., & Hood, D. C. (2003). Conventional pattern-reversal VEPs are not equivalent to summed multifocal VEPs. *Investigative Ophthalmology and Visual*

Science, 44(3), 1364–1375. <http://doi.org/10.1167/iovs.02-0441>

Foxe, J. J., Strugstad, E. C., Sehatpour, P., Molholm, S., Pasiaka, W., Schroeder, C. E., & McCourt, M. E. (2008). Parvocellular and magnocellular contributions to the initial generators of the visual evoked potential: High-density electrical mapping of the “C1” component. *Brain Topography*, 21(1), 11–21.
<http://doi.org/10.1007/s10548-008-0063-4>

Gabbiani, F., Krapp, H. G., Koch, C., & Laurent, G. (2002). Multiplicative computation in a visual neuron sensitive to looming. *Nature*, 420(6913), 320–324.
<http://doi.org/10.1038/nature01190>

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.

Georgeson, M. A. (1992). Human vision combines oriented filters to compute edges. *Proceedings. Biological Sciences / The Royal Society*, 249(1326), 235–245.
<http://doi.org/10.1098/rspb.1992.0110>

Georgeson, M. A. (1998). Edge-finding in human vision: a multi-stage model based on the perceived structure of plaids'. *Image and Vision Computing*, 16, 389–405.

Georgeson, M. A., & Shackleton, T. M. (1994). Perceived contrast of gratings and plaids: Non-linear summation across oriented filters. *Vision Research*, 34(8), 1061–1075. [http://doi.org/10.1016/0042-6989\(94\)90010-8](http://doi.org/10.1016/0042-6989(94)90010-8)

Gheorghiu, E., & Kingdom, F. A. A. (2009). Multiplication in curvature processing. *Journal of Vision*, 9(2), 23–23. <http://doi.org/10.1167/9.2.23>

Gheorghiu, E., & Kingdom, F. A. A. (2012). Local and global components of texture-surround suppression of contour-shape coding. *Journal of Vision*, 12(6), 20–20. <http://doi.org/10.1167/12.6.20>

Gibson, J. J. (1937). Adaptation, after-effect, and contrast in the perception of tilted lines. II. Simultaneous contrast and the areal restriction of the after-effect. *Journal of Experimental Psychology*, 20(6), 553–569.
<http://doi.org/10.1037/h0057585>

Gibson, J. J., & Radner, M. (1937). Adaptation, after-effect and contrast in the

- perception of tilted lines. I. Quantitative studies. *Journal of Experimental Psychology*, 20(5), 453–467. <http://doi.org/10.1037/h0059826>
- Greenlee, M. W., & Magnussen, S. (1988). Interactions among spatial frequency and orientation channels adapted concurrently. *Vision Research*, 28(12), 1303–1310. [http://doi.org/10.1016/0042-6989\(88\)90061-2](http://doi.org/10.1016/0042-6989(88)90061-2)
- Grossberg, S., & Hong, S. (2006). A neural model of surface perception: lightness, anchoring, and filling-in. *Spatial Vision*, 19(2–4), 263–321. <http://doi.org/10.1163/156856806776923399>
- Gundlach, C., & Müller, M. M. (2013). Perception of illusory contours forms intermodulation responses of steady state visual evoked potentials as a neural signature of spatial integration. *Biological Psychology*, 94(1), 55–60. <http://doi.org/10.1016/j.biopsycho.2013.04.014>
- Hancock, S., McGovern, D. P., & Peirce, J. W. (2010). Ameliorating the combinatorial explosion with spatial frequency-matched combinations of V1 outputs. *Journal of Vision*, 10(8), 7. <http://doi.org/10.1167/10.8.7>
- Hancock, S., & Peirce, J. W. (2008). Selective mechanisms for simple contours revealed by compound adaptation. *Journal of Vision*, 8(7), 11. <http://doi.org/10.1167/8.7.11>
- Hansen, B. C., Haun, A. M., Johnson, A. P., & Elleberg, D. (2016). On the Differentiation of Foveal and Peripheral Early Visual Evoked Potentials. *Brain Topography*, 29(4), 506–514. <http://doi.org/10.1007/s10548-016-0475-5>
- Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, 9(2), 181–197. <http://doi.org/10.1017/S0952523800009640>
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70.
- Holmes, G. (1945). The organization of the visual cortex in man. *Proceedings of the Royal Society of Britain*, 132(869), 348–361. <http://doi.org/10.1098/rstb.2006.1832>
- Hou, C., Pettet, M. W., Sampath, V., Candy, T. R., & Norcia, A. M. (2003). Development of the spatial organization and dynamics of lateral interactions in the human visual system. *The Journal of Neuroscience : The Official Journal of*

the Society for Neuroscience, 23(25), 8630–8640. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/14507962>

Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology*, 148(3), 574–591.

<http://doi.org/10.1113/jphysiol.1959.sp006308>

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160, 106–54.

Hubel, D. H., & Wiesel, T. N. (1963). Shape and arrangement of columns in cat's striate cortex. *The Journal of Physiology*, 165(3), 559–568.

<http://doi.org/10.1113/jphysiol.1963.sp007079>

Huk, A. C., & Heeger, D. J. (2002). Pattern-motion responses in human visual cortex.

Nature Neuroscience, 5(1), 72–75. <http://doi.org/10.1038/nn774>

Huk, A. C., Ress, D., & Heeger, D. J. (2001). Neuronal Basis of the Motion Aftereffect Reconsidered. *Neuron*, 32, 161–172.

Jeffreys, D. A., & Axford, J. G. (1972a). Source Locations of Pattern-Specific Components of Human Visual Evoked Potentials. I. Component of Striate Cortical Origin. *Exp. Brain Res*, 16, 1–21.

Jeffreys, D. A., & Axford, J. G. (1972b). Source locations of pattern-specific components of human visual evoked potentials. II. Component of extrastriate cortical origin. *Experimental Brain Research*, 16(1), 22–40.

<http://doi.org/10.1007/BF00233372>

Kanizsa, G. (1979). *Organization in vision: essays on gestalt perception*. New York, New York State: Praeger.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 17(11), 4302–11.

Kaplan, E., Marcus, S., & So, Y. T. (1979). Effects of dark adaptation on spatial and temporal properties of receptive fields in cat lateral geniculate nucleus. *The Journal of Physiology*, 294(1), 561–580.

<http://doi.org/10.1113/jphysiol.1979.sp012946>

Kaplan, E., Purpura, K., & Shapley, R. M. (1987). Contrast affects the transmission of visual information through the mammalian lateral geniculate nucleus. *The Journal of Physiology*, *391*(1), 267–288.

<http://doi.org/10.1113/jphysiol.1987.sp016737>

Keitel, C., Andersen, S. K., Quigley, C., & Müller, M. M. (2013). Independent Effects of Attentional Gain Control and Competitive Interactions on Visual Stimulus Processing. *Cerebral Cortex*, *23*(4), 940–946.

<http://doi.org/10.1093/cercor/bhs084>

Kelly, S. P., Gomez-Ramirez, M., & Foxe, J. J. (2008). Spatial Attention Modulates Initial Afferent Activity in Human Primary Visual Cortex. *Cerebral Cortex* November, *18*, 2629–2636. <http://doi.org/10.1093/cercor/bhn022>

Kelly, S. P., Vanegas, M. I., Schroeder, C. E., & Lalor, E. C. (2013). The cruciform model of striate generation of the early VEP, re-illustrated, not revoked: A reply to Ales et al. (2013). *NeuroImage*, *82*, 154–159.

<http://doi.org/10.1016/j.neuroimage.2013.05.112>

Kevin, W., Doug, W., Matthias, S., & Gerhard, S. (2008). Correspondence of visual evoked potentials with fMRI signals in human visual cortex. *Brain Topography*, *21*(2), 86–92. <http://doi.org/10.1007/s10548-008-0069-y>

Kim, J., & Wilson, H. R. (1993). Dependence of Plaid Motion Coherence on Component Grating Directions. *Vision Res*, *33*(17), 2479–2489.

Kingdom, F. A. A., Baldwin, A. S., & Schmidtman, G. (2015). Modeling probability and additive summation for detection across multiple mechanisms under the assumptions of signal detection theory. *Journal of Vision*, *15*:5(1), 1–16.

<http://doi.org/10.1167/15.5.1>

Kremers, J., Weiss, S., & Zrenner, E. (1997). Temporal Properties of Marmoset Lateral Geniculate Cells. *Vision Res*, *37*(19), 2649–2660.

Kuffler, S. W. (1953). Discharge Patterns and Functional Organization of Mammalian Retina. *Journal of Neurophysiology*. Retrieved from

<http://jn.physiology.org/content/16/1/37.short>

Kuffler, S. W., Fitzhugh, R., & Barlow, H. B. (1957). Maintained activity in the cat's

retina in light and darkness. *The Journal of General Physiology*, 40(5), 683–702. <http://doi.org/10.1085/jgp.40.5.683>

Kung, C.-C., Hsieh, T.-H., Liou, J.-Y., Lin, K.-J., Shaw, F.-Z., & Liang, S.-F. (2014). Musicians and non-musicians's™ different reliance of features in consonance perception: A behavioral and ERP study. <http://doi.org/10.1016/j.clinph.2013.10.016>

Lennie, P. (2003). Receptive fields. *Current Biology : CB*, 13(6), R216-9. [http://doi.org/10.1016/S0960-9822\(03\)00153-2](http://doi.org/10.1016/S0960-9822(03)00153-2)

Maffei, L., & Fiorentini, A. (1973). The visual cortex as a spatial frequency analyser. *Vision Research*, 13(7), 1255–1267. [http://doi.org/10.1016/0042-6989\(73\)90201-0](http://doi.org/10.1016/0042-6989(73)90201-0)

Manahilov, V., & Vassilev, A. (1986). The effect of standing contrast and pattern adaptation on the human visually evoked potential. *Electroencephalography and Clinical Neurophysiology*, 64(6), 536–545. [http://doi.org/10.1016/0013-4694\(86\)90192-6](http://doi.org/10.1016/0013-4694(86)90192-6)

Martinez, A., Anllo-Vento, L., Sereno, M. I., Frank, L. R., Buxton, R. B., Dubowitz, D. J., ... Hillyard, S. A. (1999). Involvement of striate and extrastriate visual cortical areas in spatial attention. *Nat Neurosci*, 2(4), 364–369. Retrieved from <http://dx.doi.org/10.1038/7274>

May, J. G., & Lovegrove, W. J. (1987). The effects of grating complexity on transient evoked potentials. *Electroencephalography and Clinical Neurophysiology*, 66(6), 521–528. [http://doi.org/10.1016/0013-4694\(87\)90099-X](http://doi.org/10.1016/0013-4694(87)90099-X)

May, K. A., & Zhaoping, L. (2011). Exploring the roles of saturating and supersaturating contrast-response functions in conjunction detection and contrast coding. *Journal of Vision*, 11(9), 11. <http://doi.org/10.1167/11.9.11>

May, K. A., & Zhaoping, L. (2013). The potential roles of saturating and supersaturating contrast-response functions in conjunction detection: Reply to Peirce. *Journal of Vision*, 13(4), 22. <http://doi.org/10.1167/13.4.22>

McCarthy, G., Puce, A., Gore, J. C., & Allison, T. (1997). Face-Specific Processing in the Human Fusiform Gyrus. <http://dx.doi.org/10.1162/jocn.1997.9.5.605>, 9(5), 605–610. <http://doi.org/10.1162/jocn.1997.9.5.605>

- McDonald, J. S., Mannion, D. J., & Clifford, C. W. G. (2012). Gain control in the response of human visual cortex to plaids. *Journal of Neurophysiology*, *107*(9), 2570–2580. <http://doi.org/10.1152/jn.00616.2011>
- McGovern, D. P., Hancock, S., & Peirce, J. W. (2011). The timing of binding and segregation of two compound aftereffects. *Vision Research*, *51*(9), 1047–1057. <http://doi.org/10.1016/j.visres.2011.02.017>
- McGovern, D. P., & Peirce, J. W. (2010). The spatial characteristics of plaid-form-selective mechanisms. *Vision Research*, *50*(8), 796–804. <http://doi.org/10.1016/j.visres.2010.01.018>
- Meese, T. S., & Freeman, T. C. (1995). Edge computation in human vision: anisotropy in the combining of oriented filters. *Perception*, *24*(6), 603–622. <http://doi.org/10.1068/p240603>
- Meese, T. S., & Georgeson, M. A. (1996). The tilt aftereffect in plaids and gratings: channel codes, local signs and “patchwise” transforms. *Vision Research*, *36*(10), 1421–1437. [http://doi.org/10.1016/0042-6989\(95\)00212-X](http://doi.org/10.1016/0042-6989(95)00212-X)
- Meese, T. S., & Holmes, D. J. (2007). Spatial and temporal dependencies of cross-orientation suppression in human vision. *Proceedings of the Royal Society of London B: Biological Sciences*, *274*(1606), 127–136. <http://doi.org/10.1098/rspb.2006.3697>
- Meese, T. S., & Holmes, D. J. (2010). Orientation masking and cross-orientation suppression (XOS): Implications for estimates of filter bandwidth. *Journal of Vision*, *10*(12), 9. <http://doi.org/10.1167/10.12.9>
- Miller, C. E., Shapiro, K. L., & Lovegrove, W. J. (2015). Electrophysiological measurement of the effect of inter-stimulus competition on early cortical stages of human vision. *NeuroImage*, *105*, 229–237. <http://doi.org/10.1016/j.neuroimage.2014.10.033>
- Mitchell, D. E., & Muir, D. W. (1976). Does the tilt after-effect occur in the oblique meridian? *Vision Research*, *16*(6), 609–613. [http://doi.org/10.1016/0042-6989\(76\)90007-9](http://doi.org/10.1016/0042-6989(76)90007-9)
- Movshon, J. A., & Newsome, W. T. (1996). Visual response properties of striate cortical neurons projecting to area MT in macaque monkeys. *Journal of Neuroscience*, *16*(23), 7733–7741. <http://doi.org/10.1167/9.8.752>

- Movshon, J. A., Thompson, I. D., & Tolhurst, D. J. (1978). Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *The Journal of Physiology*, *283*(1), 53–77. <http://doi.org/10.1113/jphysiol.1978.sp012488>
- Naka, K. I., & Rushton, W. A. H. (1966). S-potentials from luminosity units in the retina of fish (Cyprinidae). *The Journal of Physiology*, *185*(3), 587–599. <http://doi.org/10.1113/jphysiol.1966.sp008003>
- Nam, J.-H., Solomon, J. A., Morgan, M. J., Wright, C. E., & Chubb, C. (2009). Coherent plaids are preattentively more than the sum of their parts. *Attention, Perception, & Psychophysics*, *71*(7), 1469–1477. <http://doi.org/10.3758/APP.71.7.1469>
- Nozaradan, S., Zerouali, Y., Peretz, I., & Mouraux, A. (2015). Capturing with EEG the neural entrainment and coupling underlying sensorimotor synchronization to the beat. *Cerebral Cortex*, *25*(3), 736–747. <http://doi.org/10.1093/cercor/bht261>
- Ohzawa, I., Sclar, G., & Freeman, R. D. (1985). Contrast gain control in the cat's visual system. *Journal of Neurophysiology*, *54*(3), 651–67. <http://doi.org/10.1038/298266a0>
- Olzak, L. A., & Thomas, J. P. (1991). When orthogonal orientations are not processed independently. *Vision Research*, *31*(1), 51–57. [http://doi.org/10.1016/0042-6989\(91\)90073-E](http://doi.org/10.1016/0042-6989(91)90073-E)
- Olzak, L. A., & Thomas, J. P. (1999). Neural recoding in human pattern vision: model and mechanisms. *Vision Research*, *39*(2), 231–256. [http://doi.org/10.1016/S0042-6989\(98\)00122-9](http://doi.org/10.1016/S0042-6989(98)00122-9)
- Palanca, B. J. A., & DeAngelis, G. C. (2005). Does neuronal synchrony underlie visual feature grouping? *Neuron*, *46*(2), 333–346. <http://doi.org/10.1016/j.neuron.2005.03.002>
- Paradiso, M. A., Shimojo, S., & Nakayama, K. (1989). Subjective contours, tilt aftereffects, and visual cortical organization. *Vision Research*, *29*(9), 1205–1213. [http://doi.org/10.1016/0042-6989\(89\)90066-7](http://doi.org/10.1016/0042-6989(89)90066-7)
- Park, J. Y., Park, H., Kim, J., & Park, H.-J. (2010). Consonant chords stimulate higher EEG gamma activity than dissonant chords. *Neuroscience Letters*, *488*, 101–105. <http://doi.org/10.1016/j.neulet.2010.11.011>

- Parker, D. M., & Salzen, E. A. (1977). The spatial selectivity of early and late waves within the human visual evoked response. *Perception, 6*, 85–95.
- Passaglia, C. L., Enroth-cugell, C., & Troy, J. B. (2001). Effects of Remote Stimulation on the Mean Firing Rate of Cat Retinal Ganglion Cells. *The Journal of Physiology, 21*(15), 5794–5803. <http://doi.org/10.1113/jphysiol.1977.sp011898>
- Pasupathy, A., & Connor, C. E. (1999). Responses to contour features in macaque area V4. *Journal of Neurophysiology, 82*(5), 2490–502. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10561421>
- Pasupathy, A., & Connor, C. E. (2002). Population coding of shape in area V4. *Nature Neuroscience, 5*(12), 1332–1338. <http://doi.org/10.1038/972>
- Peirce, J. W. (2007a). PsychoPy--Psychophysics software in Python. *Journal of Neuroscience Methods, 162*(1–2), 8–13. <http://doi.org/10.1016/j.jneumeth.2006.11.017>
- Peirce, J. W. (2007b). The potential importance of saturating and supersaturating contrast response functions in visual cortex. *Journal of Vision, 7*(6), 13. <http://doi.org/10.1167/7.6.13>
- Peirce, J. W. (2011). Nonlinear summation really can be used to perform AND operations: Reply to May and Zhaoping. *Journal of Vision, 11*(9), 18–18. <http://doi.org/10.1167/11.9.18>
- Peirce, J. W. (2013). Nonlinear response functions can still be used to build conjunction detectors: reply to May and Zhaoping. *Journal of Vision, 13*(2013), 23. <http://doi.org/10.1167/13.4.23>
- Peirce, J. W. (2015). Understanding mid-level representations in visual processing. *Journal of Vision, 15*(7), 5. <http://doi.org/10.1167/15.7.5>
- Peirce, J. W., McGovern, D. P., & Hancock, S. (2009). Does plaid-selective adaptation arise from the same mechanism as the curvature aftereffect? *Journal of Vision, 9*(8), 890–890. <http://doi.org/10.1167/9.8.890>
- Peirce, J. W., & Taylor, L. J. (2006). Selective mechanisms for complex visual patterns revealed by adaptation. *Neuroscience, 141*(1), 15–18. <http://doi.org/10.1016/j.neuroscience.2006.04.039>

- Petrov, Y., Carandini, M., & McKee, S. P. (2005). Two distinct mechanisms of suppression in human vision. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 25(38), 8704–7. <http://doi.org/10.1523/JNEUROSCI.2871-05.2005>
- Phillips, D. P. (1990). Neural representation of sound amplitude in the auditory cortex: effects of noise masking. *Behavioural Brain Research*, 37, 197–214.
- Priebe, N. J., & Ferster, D. (2006). Mechanisms underlying cross-orientation suppression in cat visual cortex. *Nature Neuroscience*, 9(4), 552–561. <http://doi.org/10.1038/nn1660>
- Proverbio, A. M., Zani, A., & Avella, C. (1996). Differential activation of multiple current sources of foveal VEPs as a function of spatial frequency. *Brain Topography*, 9(1), 59–68. <http://doi.org/10.1007/BF01191643>
- Rabinowitz, N. C., Willmore, B. D. B., Schnupp, J. W. H., & King, A. J. (2011). Contrast Gain Control in Auditory Cortex. <http://doi.org/10.1016/j.neuron.2011.04.030>
- Rademacher, J., Caviness, V. S., Steinmetz, H., & Galaburda, A. M. (1993). Topographical Variation of the Human Primary Cortices: Implications for Neuroimaging, Brain Mapping, and Neurobiology. *Cerebral Cortex*, 3(4), 313–329. <http://doi.org/10.1093/cercor/3.4.313>
- Reed, J. L., Marx, M. S., & May, J. G. (1984). Spatial frequency tuning in the visual evoked potential elicited by sine-wave gratings. *Vision Research*, 24(9), 1057–1062. [http://doi.org/10.1016/0042-6989\(84\)90083-X](http://doi.org/10.1016/0042-6989(84)90083-X)
- Regan, D. (1966). Some characteristics of average steady-state and transient responses evoked by modulated light. *Electroencephalography and Clinical Neurophysiology*, 20(3), 238–248. [http://doi.org/10.1016/0013-4694\(66\)90088-5](http://doi.org/10.1016/0013-4694(66)90088-5)
- Regan, D. (1983). Spatial frequency mechanisms in human vision investigated by evoked potential recording. *Vision Research*, 23(12), 1401–1407. [http://doi.org/10.1016/0042-6989\(83\)90151-7](http://doi.org/10.1016/0042-6989(83)90151-7)
- Regan, D., & Heron, J. R. (1969). Clinical investigation of lesions of the visual pathway: a new objective technique. *Journal of Neurology, Neurosurgery, and Psychiatry*, 32(5), 479–83.

- Regan, M. P., & Regan, D. (1988). A frequency domain technique for characterizing nonlinearities in biological systems. *Journal of Theoretical Biology*, *133*(3), 293–317. [http://doi.org/10.1016/S0022-5193\(88\)80323-0](http://doi.org/10.1016/S0022-5193(88)80323-0)
- Roach, N. W., & Webb, B. S. (2013). Adaptation to implied tilt: extensive spatial extrapolation of orientation gradients. *Frontiers in Psychology*, *4*, 438. <http://doi.org/10.3389/fpsyg.2013.00438>
- Roach, N. W., Webb, B. S., & McGraw, P. V. (2008). Adaptation to global structure induces spatially remote distortions of perceived orientation. *Journal of Vision*, *8*(3), 31.1-12. <http://doi.org/10.1167/8.3.31>
- Robinson, A., & MacLeod, D. (2011). The McCollough effect with plaids and gratings: Evidence for a plaid-selective visual mechanism. *Journal of Vision*, *11*(1), 26. <http://doi.org/10.1167/11.1.26>
- Rokszin, A. A., Gyori-Dani, D., Nyúl, L. G., & Csifcsák, G. (2015). Electrophysiological correlates of top-down effects facilitating natural image categorization are disrupted by the attenuation of low spatial frequency information. *International Journal of Psychophysiology*, pp. 19–27. <http://doi.org/10.1016/j.ijpsycho.2015.12.006>
- Russell, L. V. (2016). Least-squares means: The R package lsmeans, *69*(1), 1–33. <http://doi.org/10.18637/jss.v069.i01>
- Rust, N. C., Mante, V., Simoncelli, E. P., & Movshon, J. A. (2006). How MT cells analyze the motion of visual patterns. *Nature Neuroscience*, *9*(11), 1421–1431. <http://doi.org/10.1038/nn1786>
- Sachs, M. B., S Kiang, N. Y., Sac, B., & S, B. Y. (1968). Two-Tone Inhibition in Auditory-Nerve Fibers Rate versus level functions for auditory-nerve fibers in cats: tone-burst stimuli. *The Journal of the Acoustical Society of America*, *43*(56), 1120–1835. <http://doi.org/10.1121/1.1903521>
- Sassi, M., Vancleef, K., Machilsen, B., Panis, S., & Wagemans, J. (2010). Identification of everyday objects on the basis of Gaborized outline versions. *I-Perception*, *1*(3), 121–42. <http://doi.org/10.1068/i0384>
- Schwartz, O., & Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nature Neuroscience*, *4*(8), 819–25. <http://doi.org/10.1038/90526>

- Sclar, G., Lennie, P., & Depriest, D. D. (1989). CONTRAST ADAPTATION IN STRIATE CORTEX OF MACAQUE. *Vision Res*, 29(7), 147–155.
- Sclar, G., Maunsell, J. H. R., & Libnie, P. (1990). CODING OF IMAGE CONTRAST IN CENTRAL VISUAL PATHWAYS OF THE MACAQUE MONKEY. *Vision Research*, 30(1), 1–10.
- Serre, T., Oliva, A., & Poggio, T. (2007). A feedforward architecture accounts for rapid categorization. *Proceedings of the National Academy of Sciences*, 104(15), 6424–6429. <http://doi.org/10.1073/pnas.0700622104>
- Shapely, R. M., & Lennie, P. (1985). Spatial frequency analysis in the visual system. *Annual Review of Neuroscience*, 8, 547–583. <http://doi.org/10.1146/annurev.ne.08.030185.002555>
- Shigeto, H., Tobimatsu, S., Yamamoto, T., Kobayashi, T., & Kato, M. (1998). Visual evoked cortical magnetic responses to checkerboard pattern reversal stimulation: A study on the neural generators of N75, P100 and N145. *Journal of the Neurological Sciences*, 156(2), 186–194. [http://doi.org/10.1016/S0022-510X\(98\)00026-4](http://doi.org/10.1016/S0022-510X(98)00026-4)
- Shou, T., Li, X., Zhou, Y., & Hu, B. (1996). Adaptation of visually evoked responses of relay cells in the dorsal lateral geniculate nucleus of the cat following prolonged exposure to drifting gratings. *Visual Neuroscience*, 13(4), 605–13. <http://doi.org/10.1017/S0952523800008518>
- Singmann, H., Bolker, B., Westfall, J., Aust, F., Højsgaard, S., Fox, J., ... Mertens, U. (2016). afex: Analysis of factorial experiments. Comprehensive R Archive Network (CRAN). Retrieved from <https://github.com/singmann/afex>
- Sit, Y. F., Chen, Y., Geisler, W. S., Miikkulainen, R., & Seidemann, E. (2009). Complex Dynamics of V1 Population Responses Explained by a Simple Gain-Control Model. *Neuron*, 64(6), 943–956. <http://doi.org/10.1016/j.neuron.2009.08.041>
- Smith, A. T., & Jeffreys, D. A. (1978). Evoked potential evidence of adaptation to spatial Fourier components in human vision. *Nature*, 274(5667), 156–158. <http://doi.org/10.1038/274156a0>
- Snowden, R. J. (1992). Orientation bandwidth: the effect of spatial and temporal frequency. *Vision Research*, 32(10), 1965–1974. [http://doi.org/10.1016/0042-6989\(92\)90056-O](http://doi.org/10.1016/0042-6989(92)90056-O)

- Snowden, R. J., & Hammett, S. T. (1992). Subtractive and divisive adaptation in the human visual system. *Nature*, *355*(6357), 248–50.
<http://doi.org/10.1038/355248a0>
- Spekreijse, H., & Oosting, H. (1970a). Linearizing: a method for analysing and synthesizing nonlinear systems. *Kybernetik*, *7*(1), 22–31.
<http://doi.org/10.1007/BF00270331>
- Spekreijse, H., & Oosting, H. (1970b). Linearizing: A method for analysing and synthesizing nonlinear systems. *Kybernetik*, *7*(1), 22–31.
<http://doi.org/10.1007/BF00270331>
- Spekreijse, H., & Reits, D. (1982). SEQUENTIAL ANALYSIS OF THE VISUAL EVOKED POTENTIAL SYSTEM IN MAN; NONLINEAR ANALYSIS OF A SANDWICH SYSTEM. *Annals of the New York Academy of Sciences*, *388*(1), 72–97.
<http://doi.org/10.1111/j.1749-6632.1982.tb50785.x>
- Suter, S., Perrier, D. T., Parker, K. L., Fox, J. A., & Roessler, J. S. (1996). Differentiation of VEP intermodulation and second harmonic components by dichoptic, monocular, and binocular stimulation. *Visual Neuroscience*, *13*(6), 1157–1166.
<http://doi.org/10.1017/S0952523800007793>
- Sutoyo, D., & Srinivasan, R. (2009). Nonlinear SSVEP responses are sensitive to the perceptual binding of visual hemifields during conventional “eye” rivalry and interocular “percept” rivalry. *Brain Research*, *1251*, 245–55.
<http://doi.org/10.1016/j.brainres.2008.09.086>
- Sutter, M. L., & Loftus, W. C. (2003). Excitatory and inhibitory intensity tuning in auditory cortex: evidence for multiple inhibitory mechanisms, 2629–2647.
- Suzuki, S. (2001). Attention-dependent brief adaptation to contour orientation: a high-level aftereffect for convexity? *Vision Research*, *41*(28), 3883–3902.
[http://doi.org/10.1016/S0042-6989\(01\)00249-8](http://doi.org/10.1016/S0042-6989(01)00249-8)
- Suzuki, S. (2003). Attentional selection of overlapped shapes: a study using brief shape aftereffects. *Vision Research*, *43*(5), 549–561.
[http://doi.org/10.1016/S0042-6989\(02\)00683-1](http://doi.org/10.1016/S0042-6989(02)00683-1)
- Tallon-Baudry, C., Bertrand, O., Delpuech, C., & Pernier, J. (1996). Stimulus specificity of phase-locked and non-phase-locked 40 Hz visual responses in human. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*,

16(13), 4240–9. <http://doi.org/10.1016/j.neuropsychologia.2011.02.038>

- Tan, K. W. S., Bowden, V. K., Dickinson, J. E., & Badcock, D. R. (2015). Modulated textures with shape structures implied by a closed flow are processed globally. *Journal of Vision*, 15(3), 17-. <http://doi.org/10.1167/15.3.17>
- Tan, K. W. S., Dickinson, J. E., & Asaad, W. (2016). Discrete annular regions of texture contribute independently to the analysis of shape from texture. *Journal of Vision*, 16(11), 10. <http://doi.org/10.1167/16.11.10>
- Tan, K. W. S., Dickinson, J. E., & Badcock, D. R. (2016). In unison: First- and second-order information combine for integration of shape information. *Journal of Vision*, 16(11)(9), 1–14. <http://doi.org/10.1167/16.11.9>
- Taylor, W. R., He, S., Levick, W. R., & Vaney, D. I. (2000). Dendritic computation of direction selectivity by retinal ganglion cells. *Science (New York, N.Y.)*, 289(5488), 2347–50. <http://doi.org/10.1126/SCIENCE.289.5488.2347>
- Tobimatsu, S., Tomoda, H., & Kato, M. (1996). Normal variability of the amplitude and phase of steady-state VEPs. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, 100(3), 171–176. [http://doi.org/10.1016/0168-5597\(95\)00279-0](http://doi.org/10.1016/0168-5597(95)00279-0)
- Tolhurst, D. J., & Thompson, P. G. (1974). ORIENTATION ILLUSIONS AND AFTER-EFFECTS: INHIBITION BETWEEN CHANNELS.
- Troy, J. B., & Robson, J. G. (1992). Steady discharges of X and Y retinal ganglion cells of cat under photopic illuminance. *Visual Neuroscience*, 9(6), 535–553. <http://doi.org/10.1017/S0952523800001784>
- Tsai, J. J., Wade, A. R., & Norcia, A. M. (2012). Dynamics of Normalization Underlying Masking in Human Visual Cortex. *The Journal of Neuroscience*, 32(8), 2783–2789. <http://doi.org/10.1523/JNEUROSCI.4485-11.2012>
- Vancleef, K., Putzeys, T., Gheorghiu, E., Sassi, M., Machilsen, B., & Wagemans, J. (2013). Spatial arrangement in texture discrimination and texture segregation. *i-Perception*, 4(1), 36–52. <http://doi.org/10.1068/i0515>
- Vanegas, M. I., Blangero, A., & Kelly, S. P. (2013). Exploiting individual primary visual cortex geometry to boost steady state visual evoked potentials. *Journal of Neural Engineering*, 10(3), 36003. <http://doi.org/10.1088/1741->

- Vassilev, A., Manahilov, V., & Mitov, D. (1983). Spatial frequency and the pattern onset-offset response. *Vision Research*, 23(12), 1417–1422. [http://doi.org/10.1016/0042-6989\(83\)90153-0](http://doi.org/10.1016/0042-6989(83)90153-0)
- Victor, J. D. (1979). Nonlinear systems analysis: comparison of white noise and sum of sinusoids in a biological system. *Proceedings of the National Academy of Sciences of the United States of America*, 76(2), 996–8. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/284424>
- Victor, J. D. (1987). The dynamics of the cat retinal X cell centre. *The Journal of Physiology*, 386(1), 219–246. <http://doi.org/10.1113/jphysiol.1987.sp016531>
- Victor, J. D., & Conte, M. M. (2000). Two-frequency analysis of interactions elicited by Vernier stimuli. *Visual Neuroscience*, 17(6), 959–973.
- Victor, J. D., & Shapely, R. M. (1980). A method of nonlinear analysis in the frequency domain. *Biophysical Journal*, 29(3), 459–483. [http://doi.org/10.1016/S0006-3495\(80\)85146-0](http://doi.org/10.1016/S0006-3495(80)85146-0)
- Viswanathan, S., Jayakumar, J., & Vidyasagar, T. R. (2011). Role of feedforward geniculate inputs in the generation of orientation selectivity in the cat's primary visual cortex. *The Journal of Physiology*, 589(9), 2349–2361. <http://doi.org/10.1113/jphysiol.2010.202317>
- Wenderoth, P., & Johnstone, S. (1988). The different mechanisms of the direct and indirect tilt illusions. *Vision Research*, 28(2), 301–312. [http://doi.org/10.1016/0042-6989\(88\)90158-7](http://doi.org/10.1016/0042-6989(88)90158-7)
- Wenderoth, P., Watson, J. D., Egan, G. F., Tochon-Danguy, H. J., & O'keefe, G. J. (1999). Second order components of moving plaids activate extrastriate cortex: a positron emission tomography study. *NeuroImage*, 9(2), 227–34. <http://doi.org/10.1006/nimg.1998.0398>
- Wessinger, C. M., VanMeter, J., Tian, B., Pekar, J., & Rauschecker, J. P. (2001). Hierarchical organization of the human auditory cortex revealed by functional magnetic resonance imaging, 13, 1–7. Retrieved from <http://www.mitpressjournals.org/doi/pdf/10.1162/089892901564108>
- Wile, D., & Balaban, E. (2007). An auditory neural correlate suggests a mechanism

underlying holistic pitch perception. *PLoS ONE*, 2(4), e369.
<http://doi.org/10.1371/journal.pone.0000369>

Willmore, B. D. B., Prenger, R. J., & Gallant, J. L. (2010). Neural Representation of Natural Images in Visual Area V2. *The Journal of Neuroscience*, 30(6), 2102–2114. <http://doi.org/10.1523/JNEUROSCI.4099-09.2010>

Xing, D., Shapely, R. M., Hawken, M. J., & Ringach, D. L. (2005). Effect of stimulus size on the dynamics of orientation selectivity in Macaque V1. *Journal of Neurophysiology*, 94(1), 799–812. <http://doi.org/10.1152/jn.01139.2004>

Zemon, V., & Ratliff, F. (1984). Intermodulation components of the visual evoked potential: Responses to lateral and superimposed stimuli. *Biological Cybernetics*, 50(6), 401–408. <http://doi.org/10.1007/BF00335197>

Appendix one – supplementary figures from Chapter 3

Component A at 20% Michelson Contrast, 0-160ms

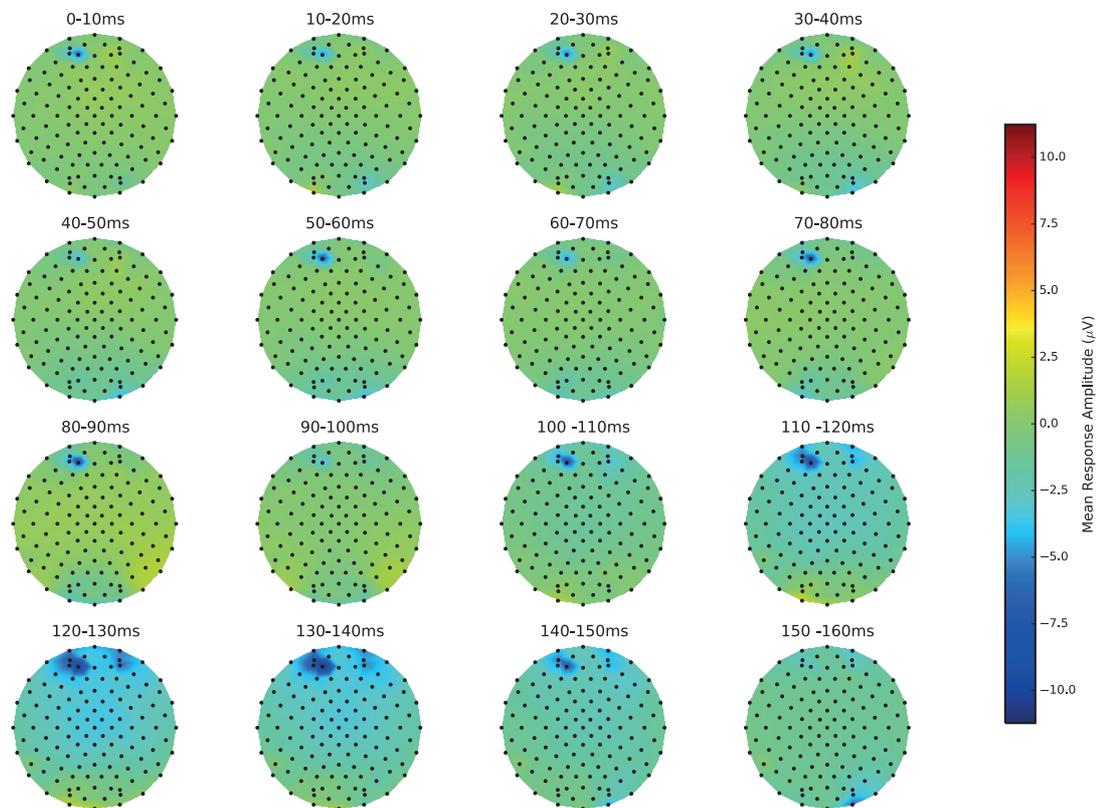


Figure 3. 2. Grand average topographies in 10ms averaged chunks for grating component A at 20% contrast.

Component A at 40% Michelson Contrast, 0-160ms

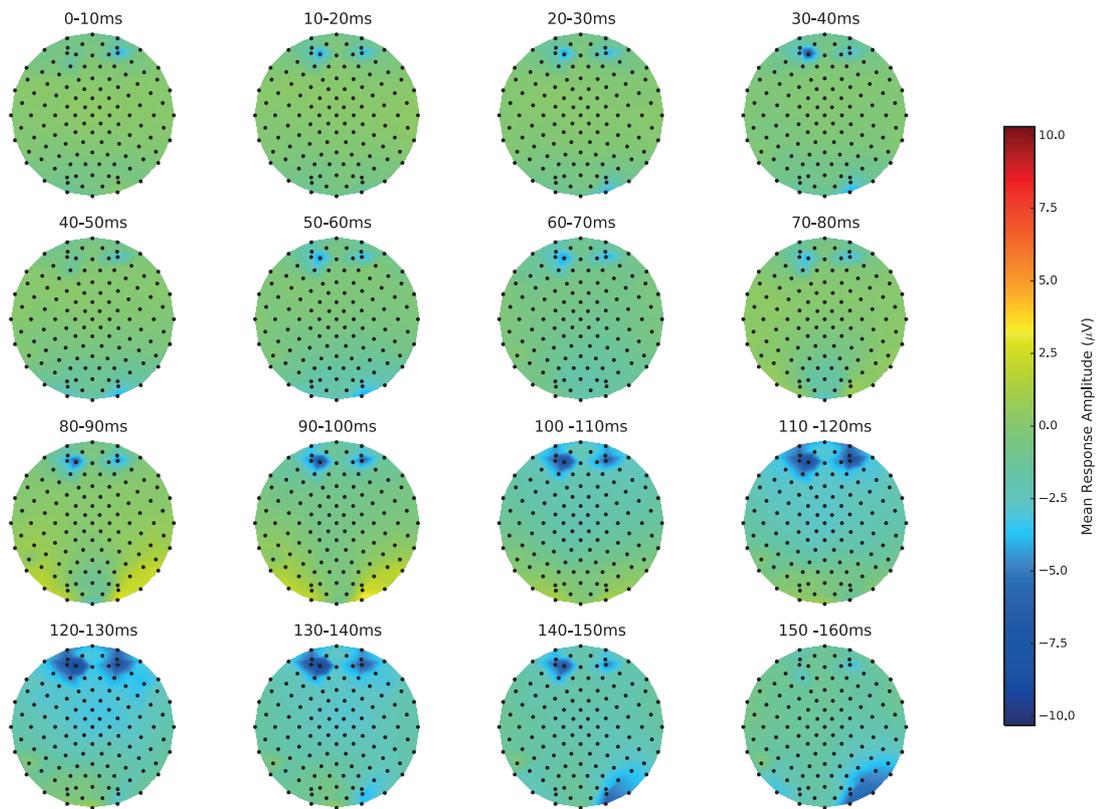


Figure 3.3. Grand average topographies in 10ms averaged chunks for grating component A at 40% contrast.

Component A at 80% Michelson Contrast, 0-160ms

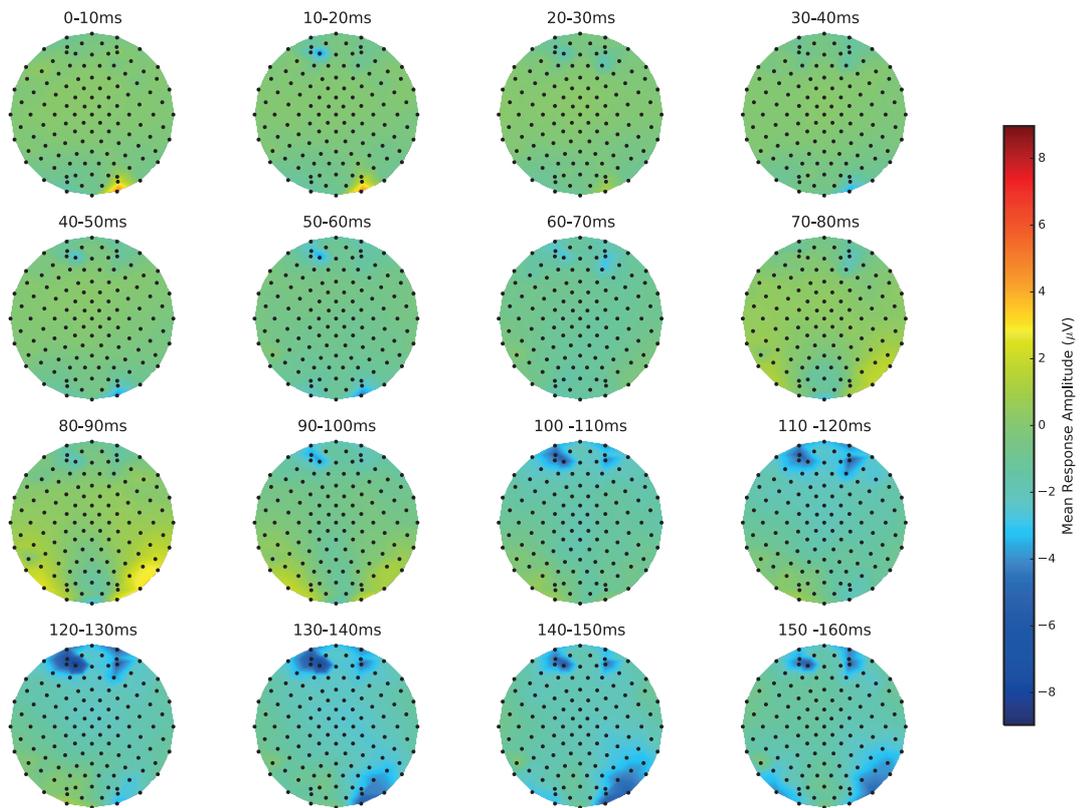


Figure 3.4. Grand average topographies in 10ms averaged chunks for grating component A at 80% contrast.

Component B at 20% Michelson Contrast, 0-160ms

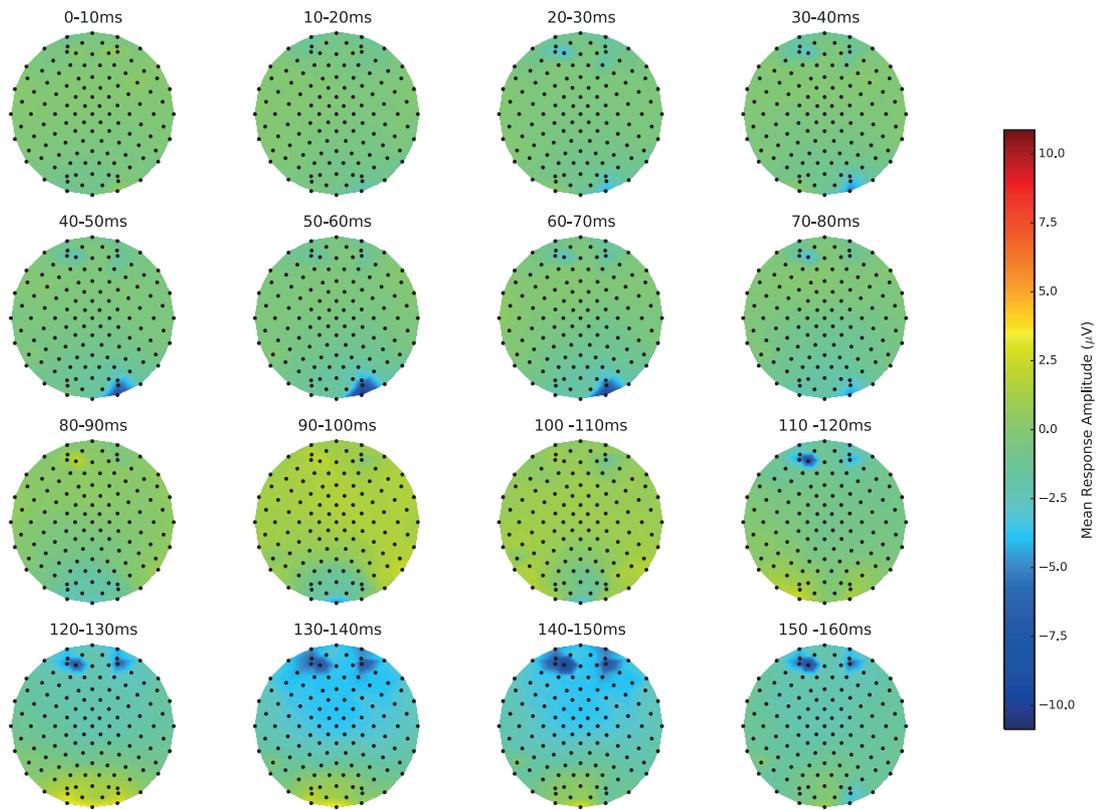


Figure 3.5. Grand average topographies in 10ms averaged chunks for grating component B at 20% contrast.

Component B at 40% Michelson Contrast, 0-160ms

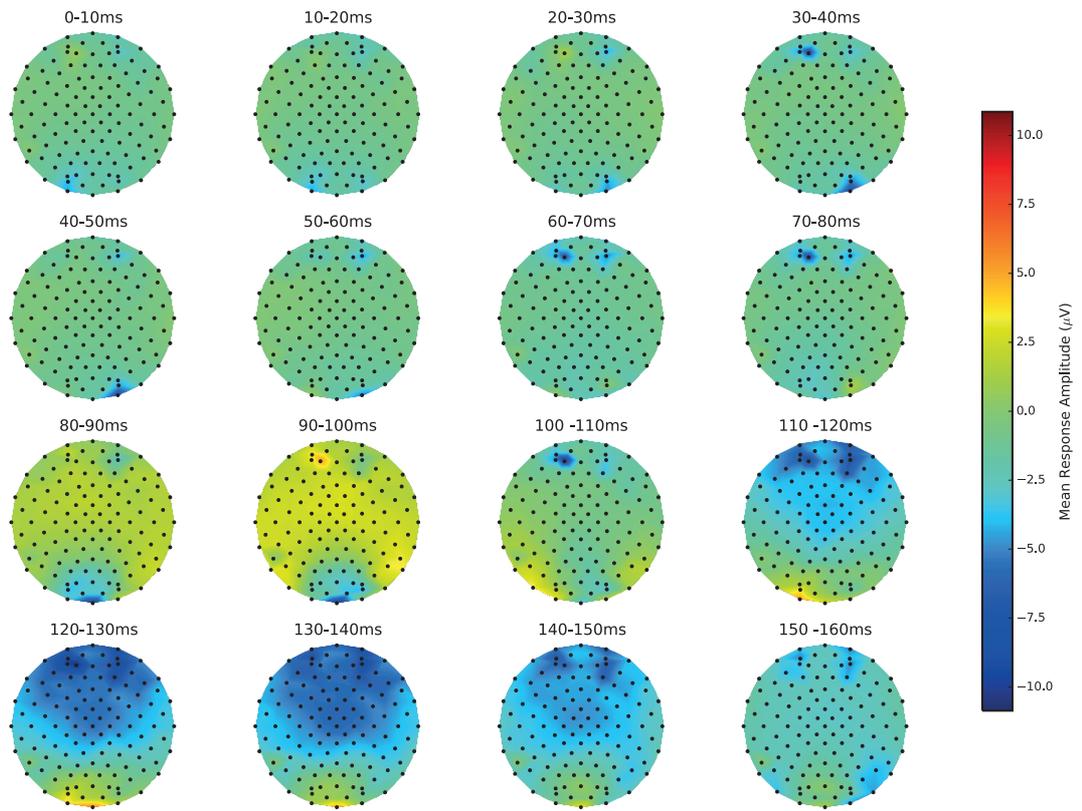


Figure 3.6. Grand average topographies in 10ms averaged chunks for grating component B at 40% contrast.

Component B at 80% Michelson Contrast, 0-160ms

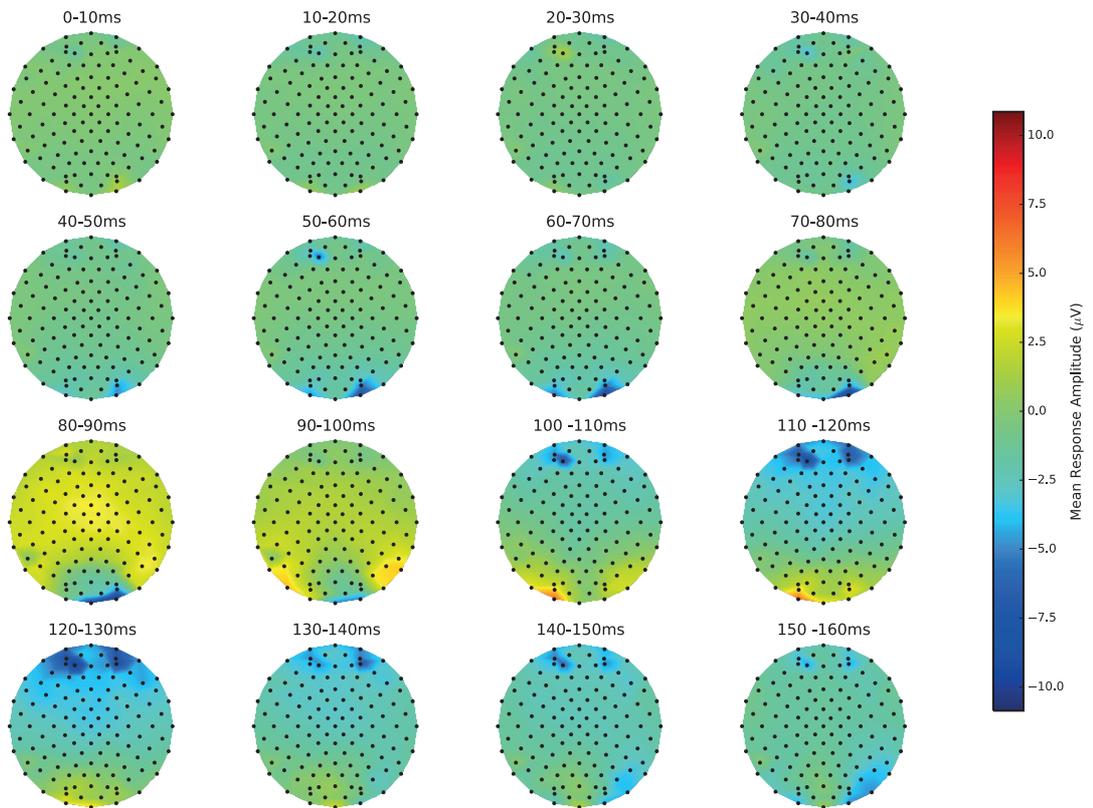


Figure 3.7. Grand average topographies in 10ms averaged chunks for grating component B at 80% contrast.

Plaid AA at 20% Michelson Contrast, 0-160ms

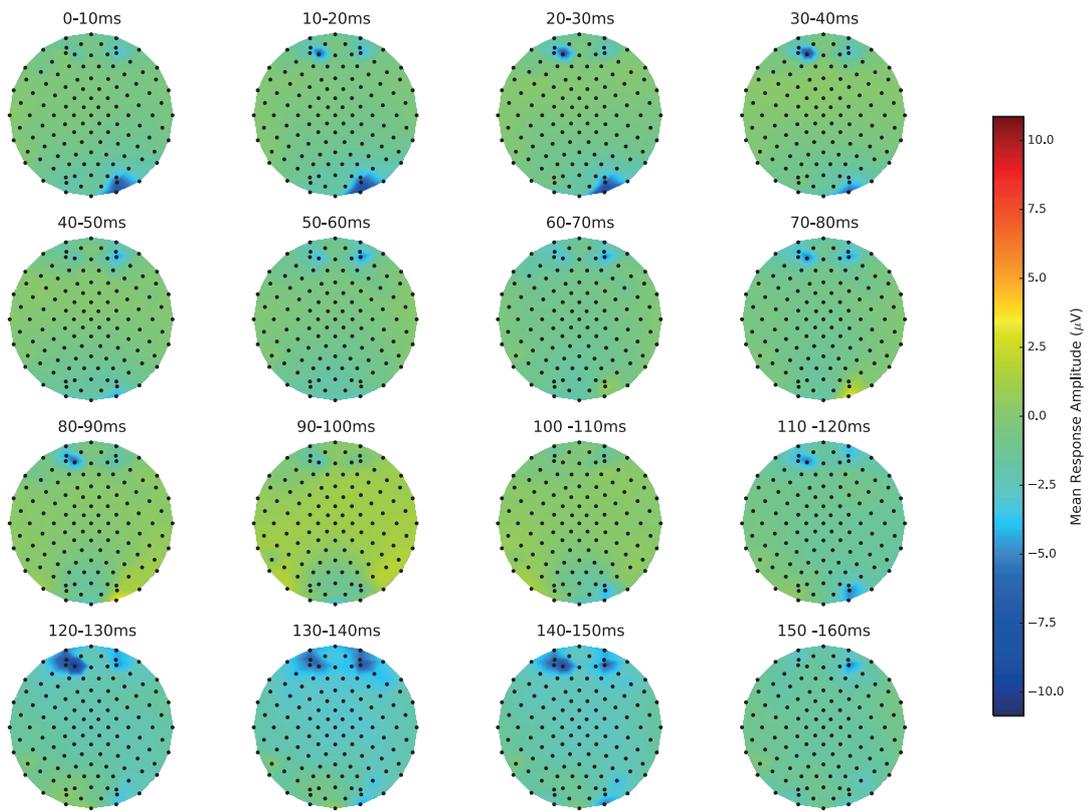


Figure 3.8. Grand average topographies in 10ms averaged chunks for grating plaid AA at 20% contrast (10% component contrast).

Plaid AA at 40% Michelson Contrast, 0-160ms

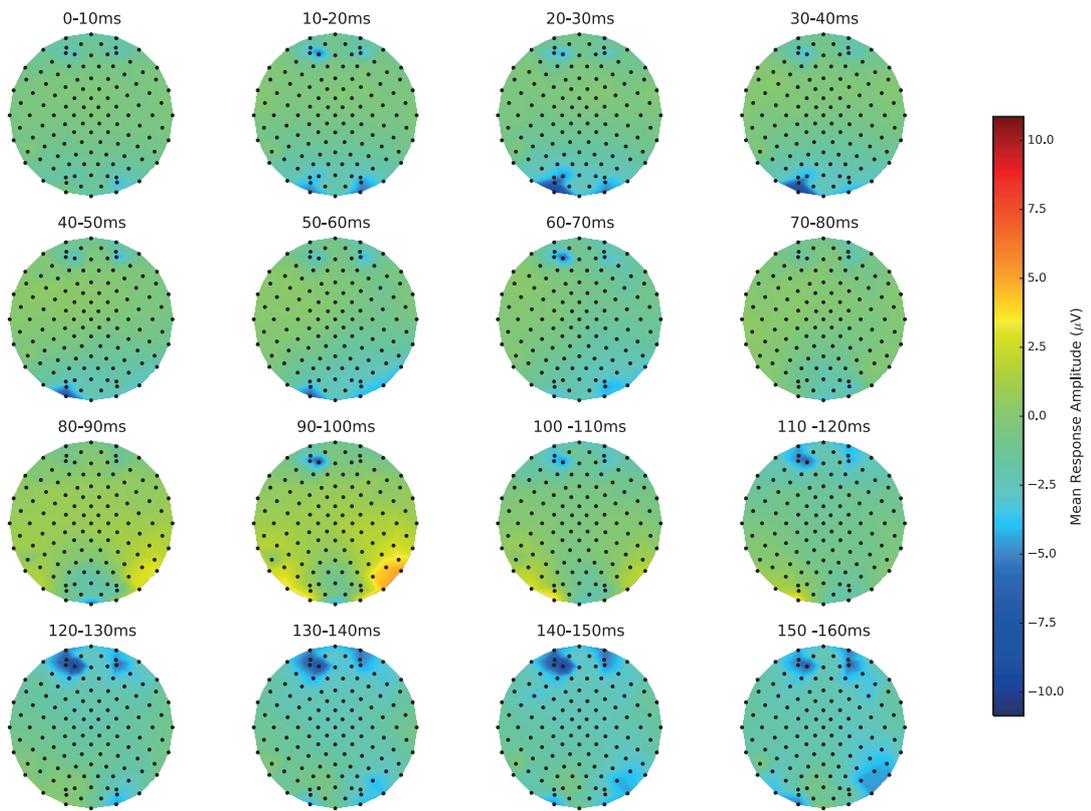


Figure 3.9. Grand average topographies in 10ms averaged chunks for grating plaid AA at 40% contrast (20% component contrast).

Plaid AA at 80% Michelson Contrast, 0-160ms

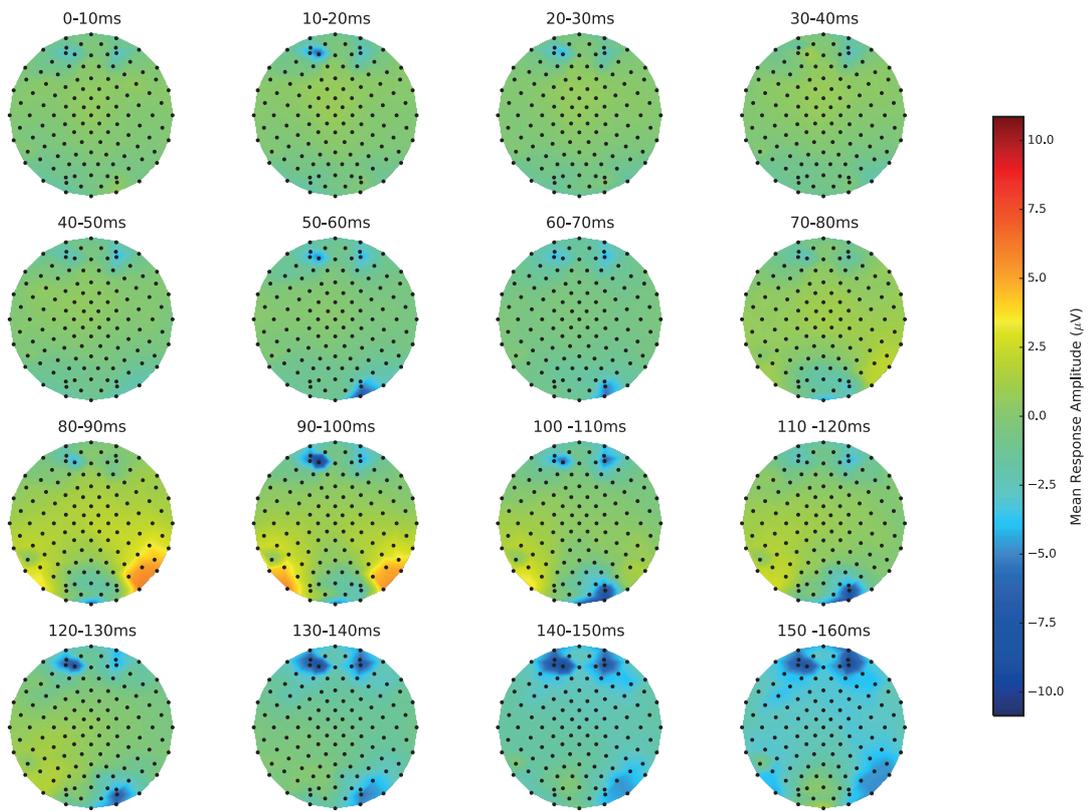


Figure 3.10. Grand average topographies in 10ms averaged chunks for grating plaid AA at 80% contrast (40% component contrast).

Plaid BB at 20% Michelson Contrast, 0-160ms

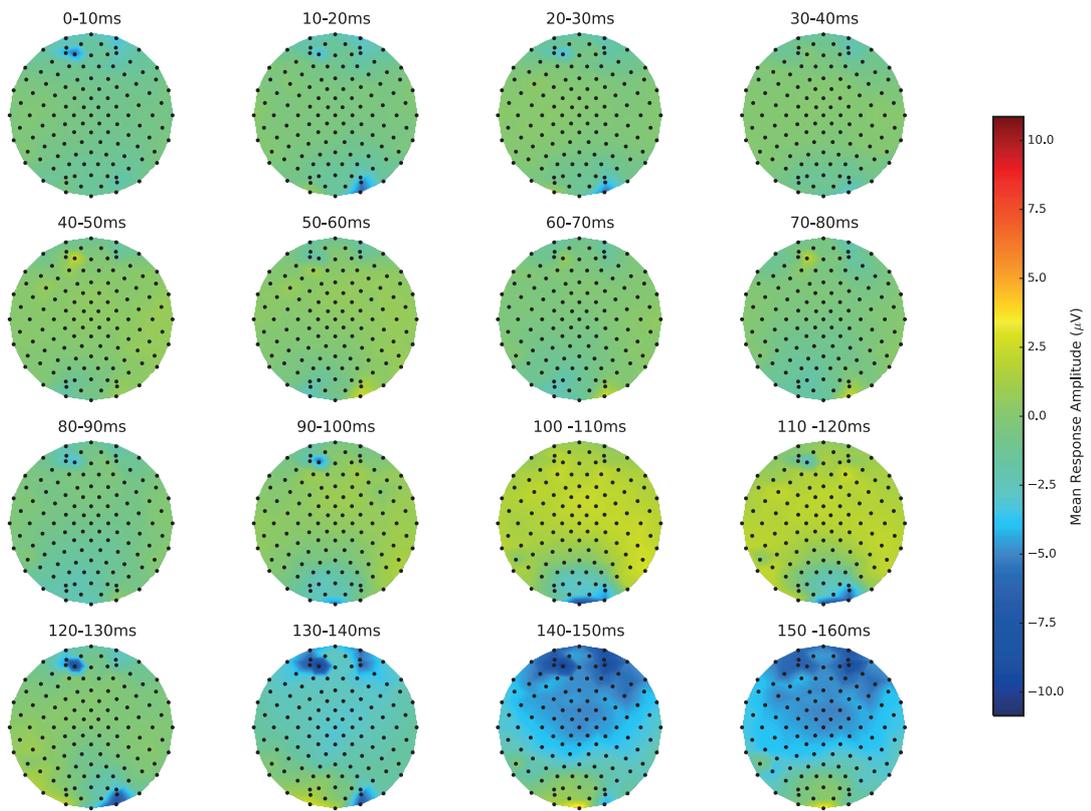


Figure 3.11. Grand average topographies in 10ms averaged chunks for grating plaid BB at 20% contrast (10% component contrast).

Plaid BB at 40% Michelson Contrast, 0-160ms

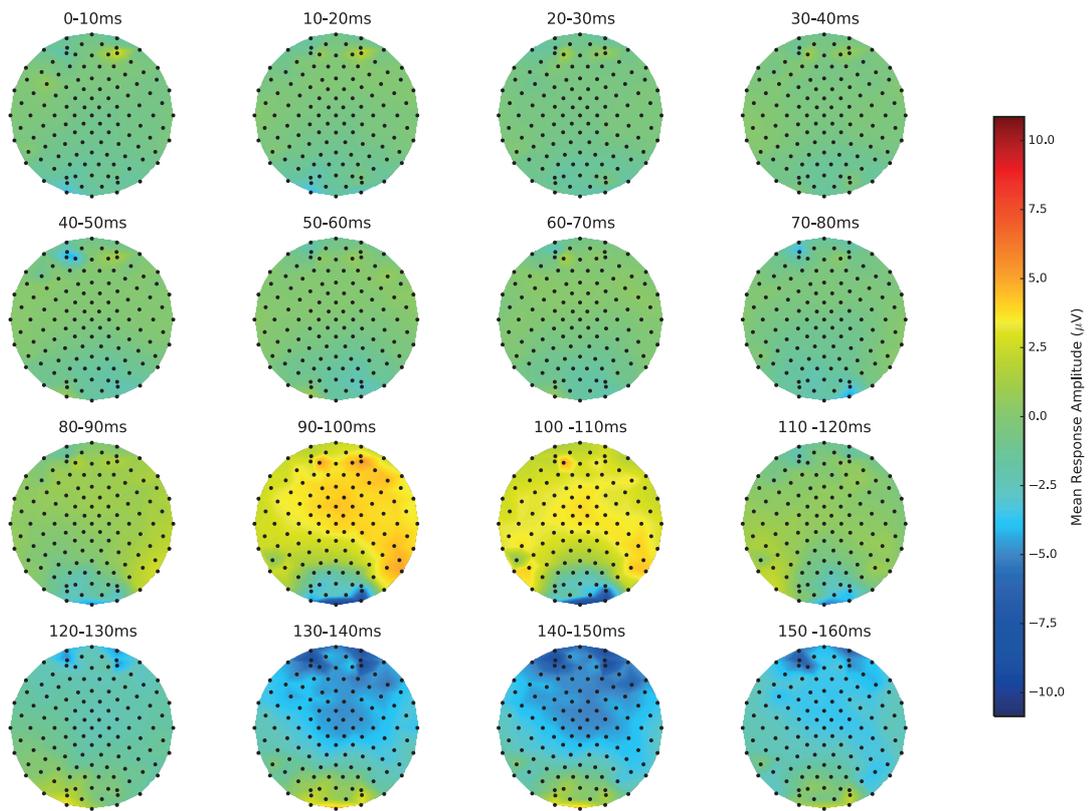


Figure 3.12. Grand average topographies in 10ms averaged chunks for grating plaid BB at 40% contrast (20% component contrast).

Plaid BB at 80% Michelson Contrast, 0-160ms

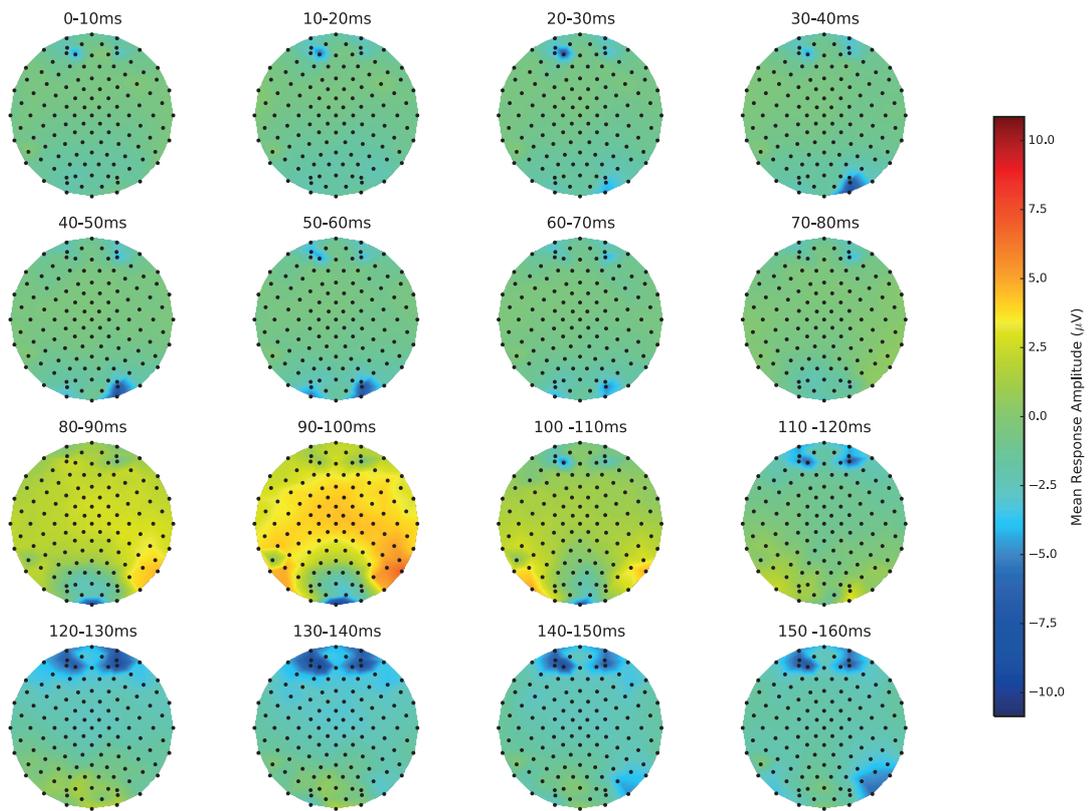


Figure 3.13. Grand average topographies in 10ms averaged chunks for grating plaid BB at 80% contrast (40% component contrast).

Plaid AB at 20% Michelson Contrast, 0-160ms

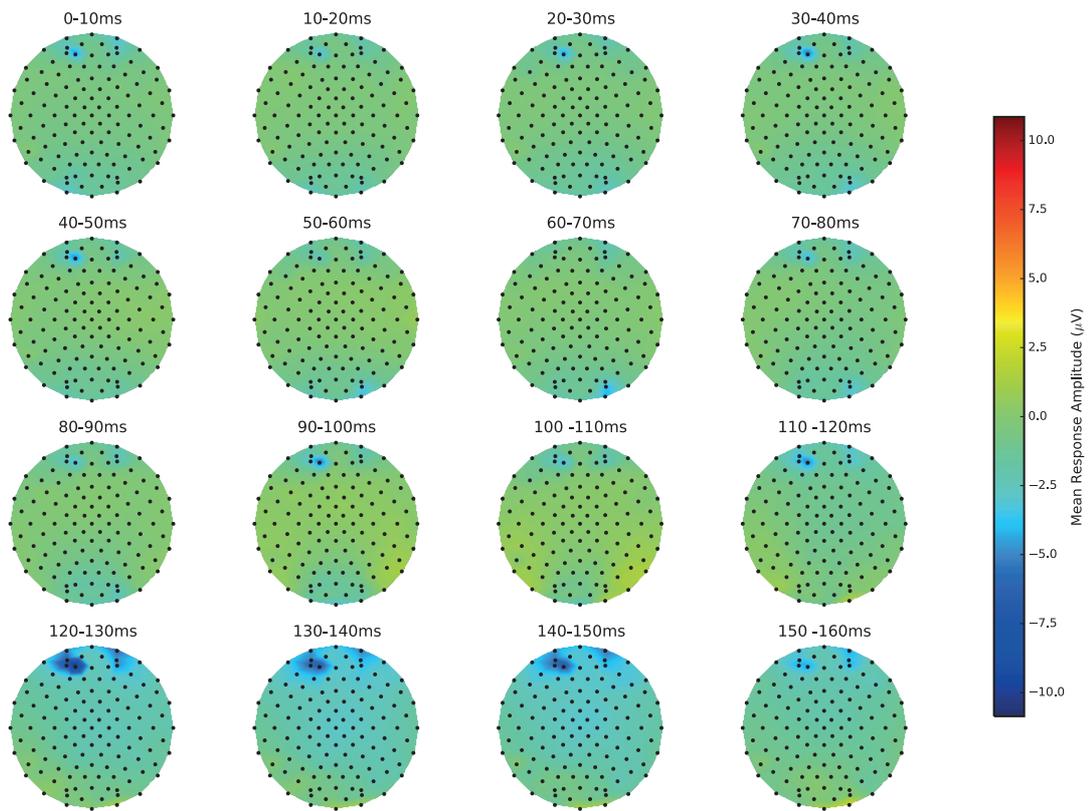


Figure 3.14. Grand average topographies in 10ms averaged chunks for grating plaid AB at 20% contrast (10% component contrast).

Plaid AB at 40% Michelson Contrast, 0-160ms

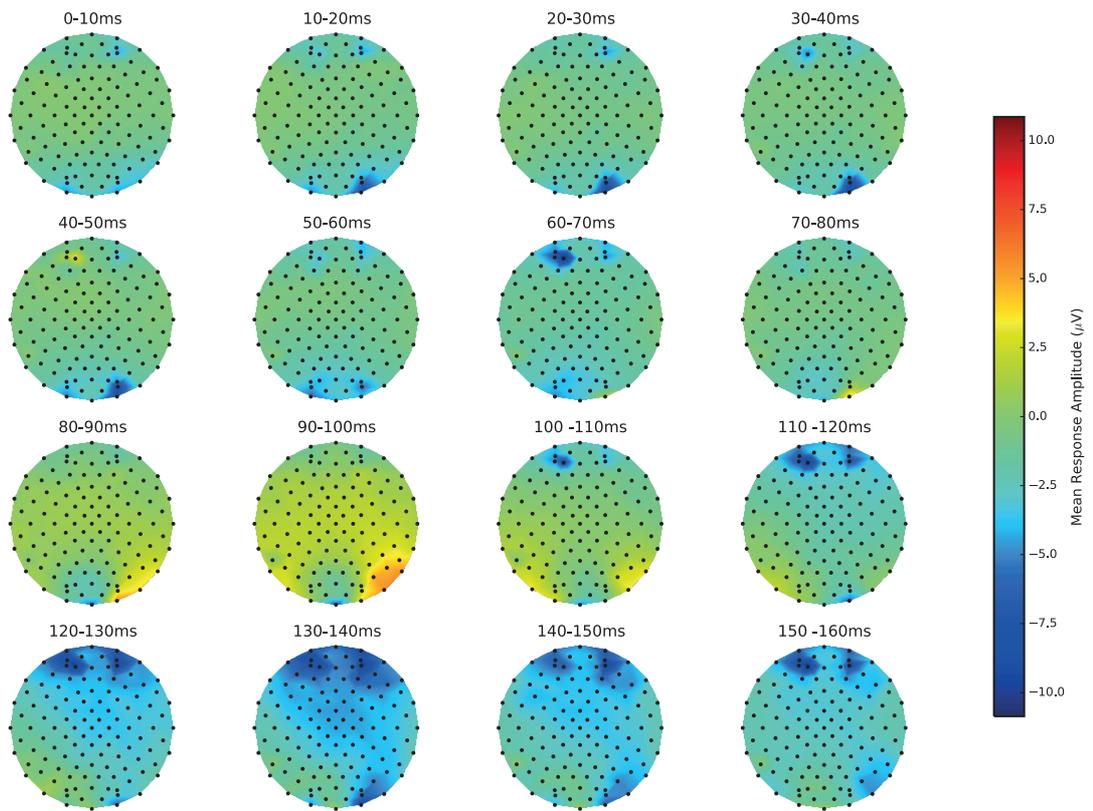


Figure 3.15. Grand average topographies in 10ms averaged chunks for grating plaid AB at 40% contrast (20% component contrast).

Plaid AB at 80% Michelson Contrast, 0-160ms

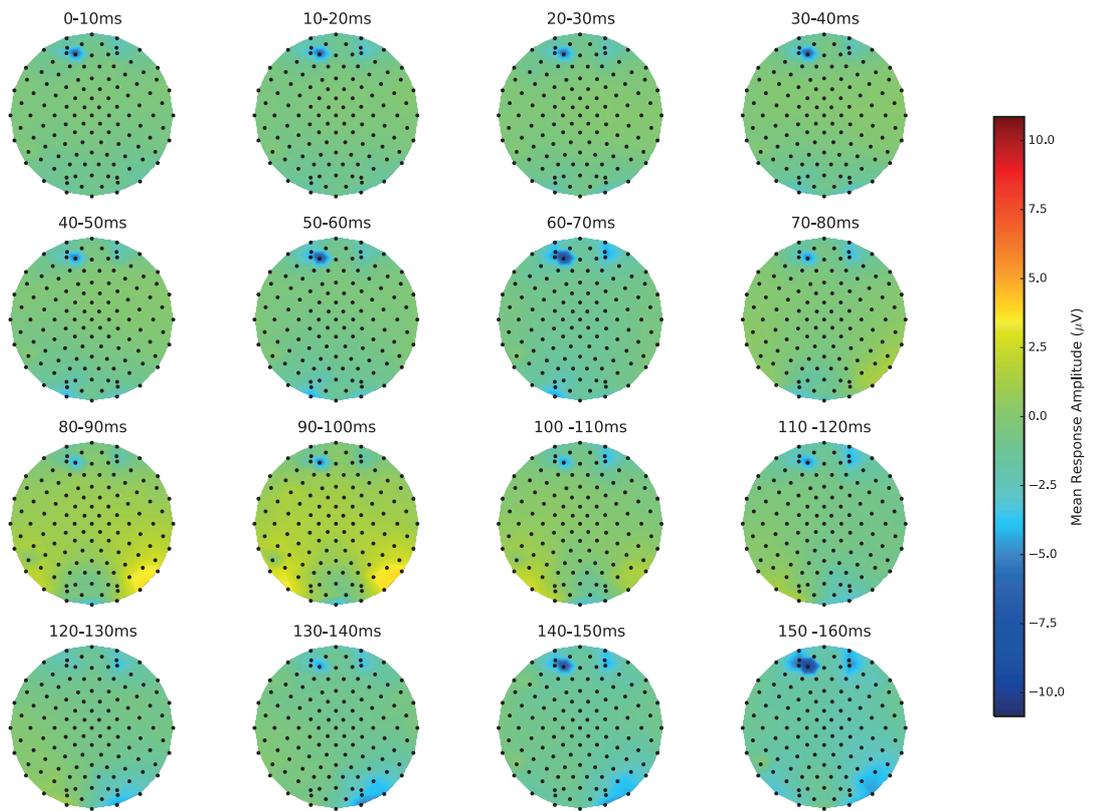


Figure 3.16. Grand average topographies in 10ms averaged chunks for grating plaid AB at 80% contrast (40% component contrast).

Appendix two – supplementary figures from Chapter 5

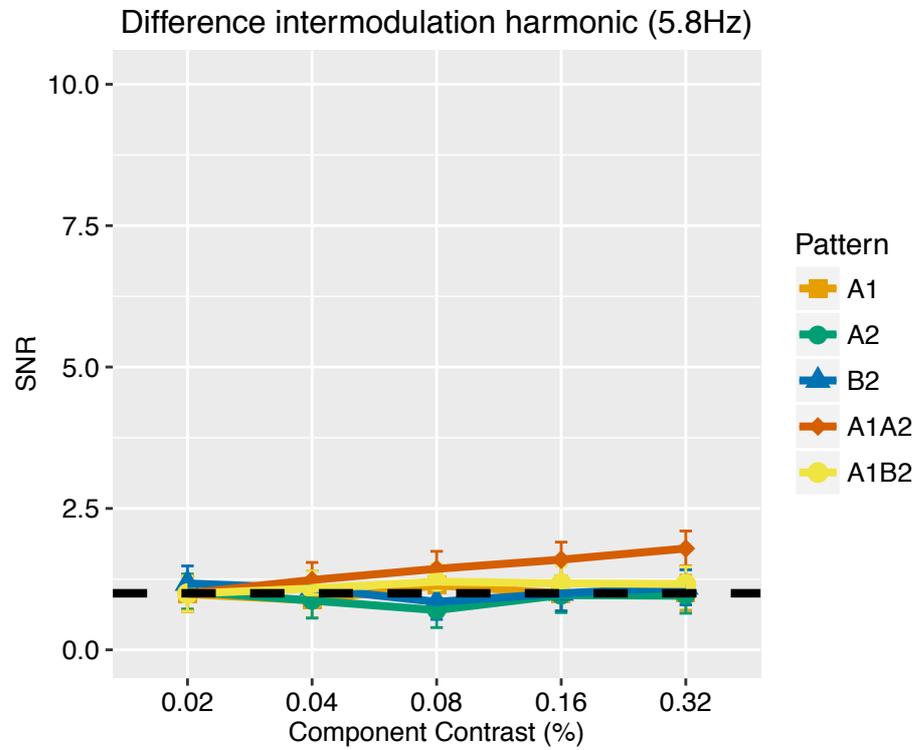


Figure 5.7. Contrast response functions for SNRs at the harmonic difference intermodulation frequency (5.8Hz) for each condition. The black dashed line represents an SNR of 1, and error bars represent 95% CIs.

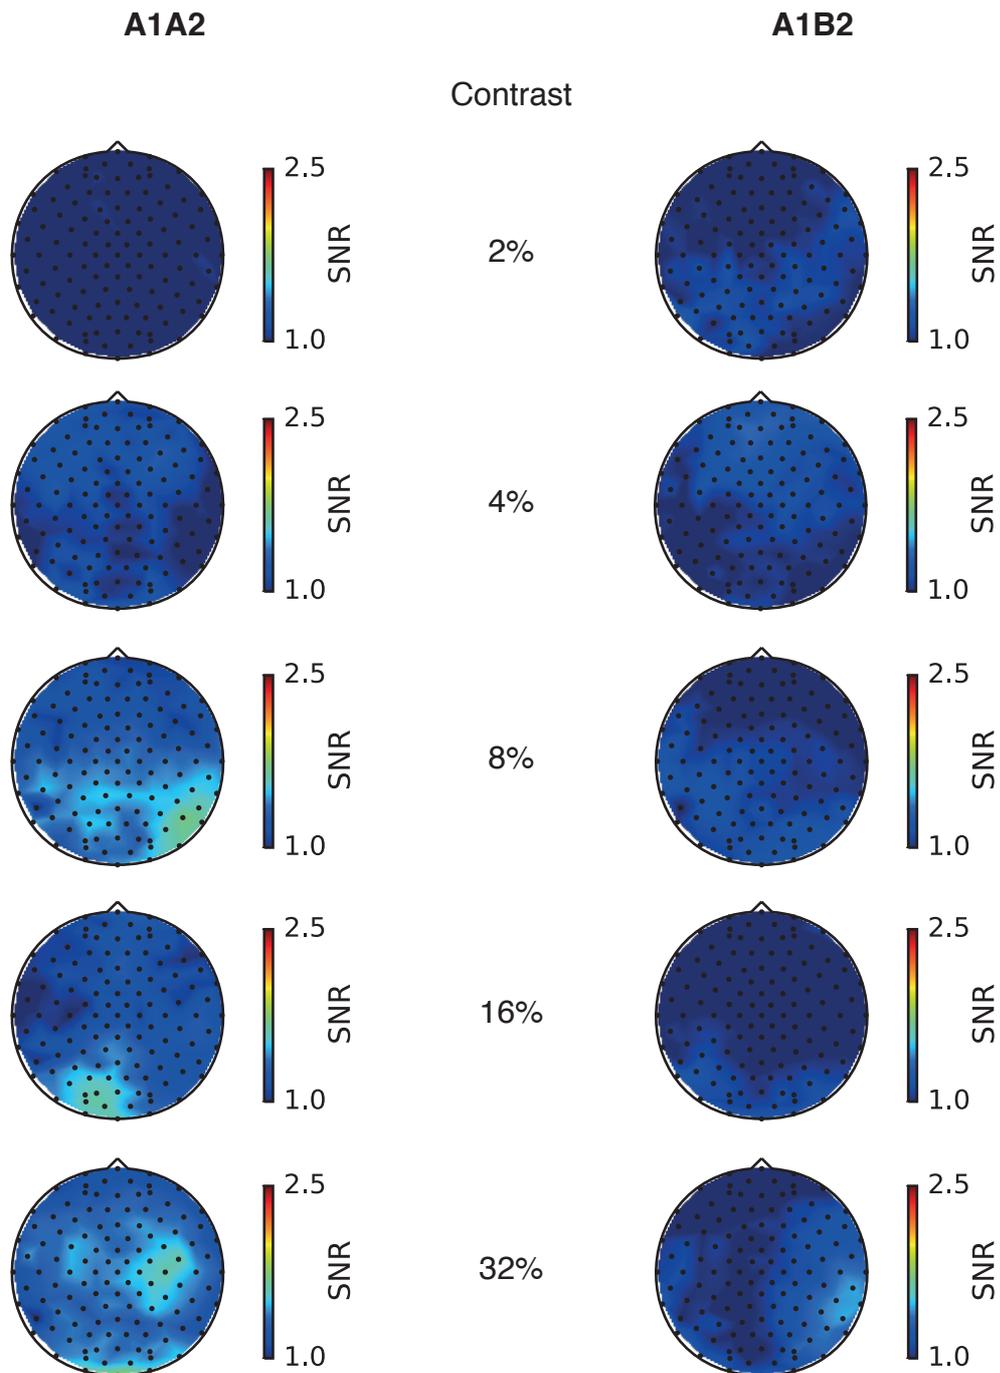


Figure 5.8. SNR topographies at 5.8Hz ($2f_2-2f_1$) for coherent plaid A1A2 and non-coherent plaid A1B2. Contrast increases from the top to the bottom of the figure, from 2% Michelson through to 32% Michelson component contrast. Colour maps were scaled so that if the topography's maximum value was less than 2.5, the maximum of the topography was set to 2.5. This helped avoid noisy topographical maps as SNR approached 1.

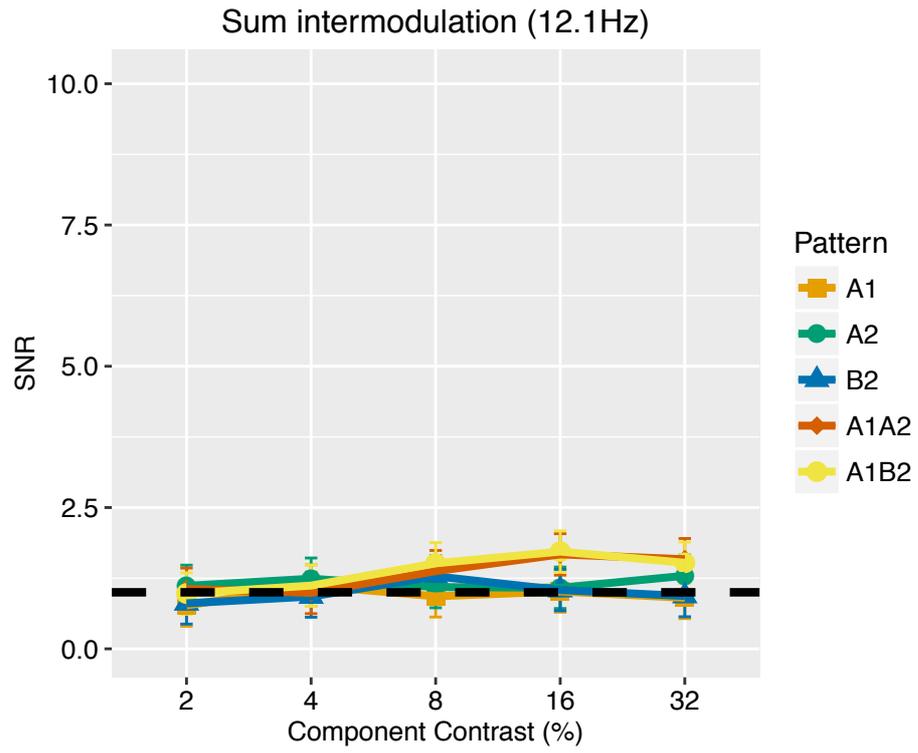


Figure 5.9 Contrast response functions for SNRs at the sum intermodulation frequency (12.1Hz) for each condition. The black dashed line represents an SNR of 1, and error bars represent 95% CIs.

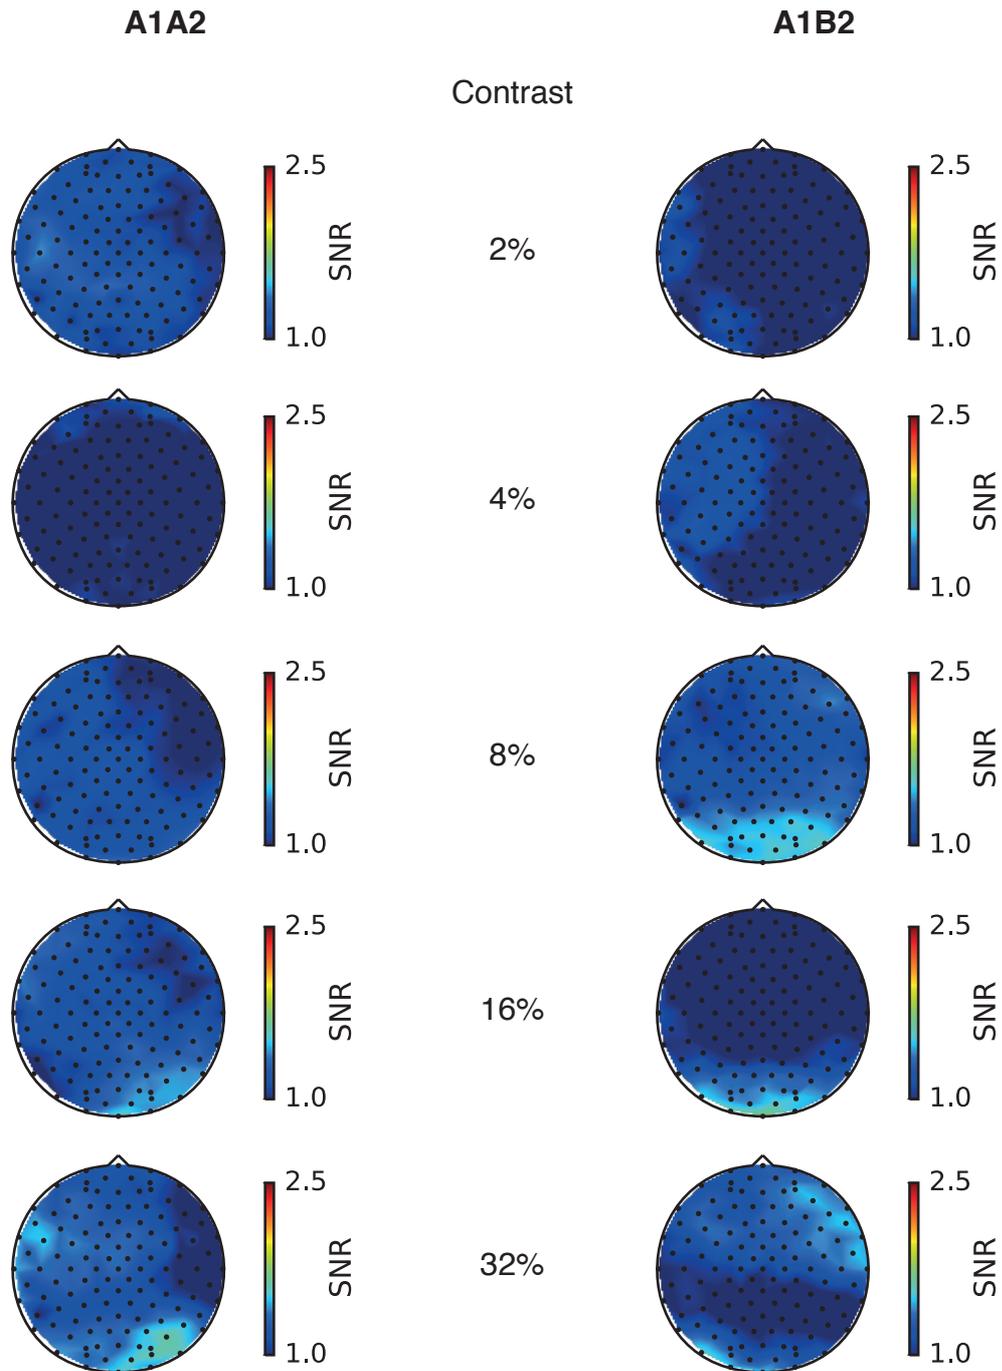


Figure 5.10. SNR topographies at 12.1Hz (f_1+f_2) for coherent plaid A1A2 and non-coherent plaid A1B2. Contrast increases from the top to the bottom of the figure, from 2% Michelson through to 32% Michelson component contrast. Colour maps were scaled so that if the topography's maximum value was less than 2.5, the maximum of the topography was set to 2.5. This helped avoid noisy topographical maps as SNR approached 1.

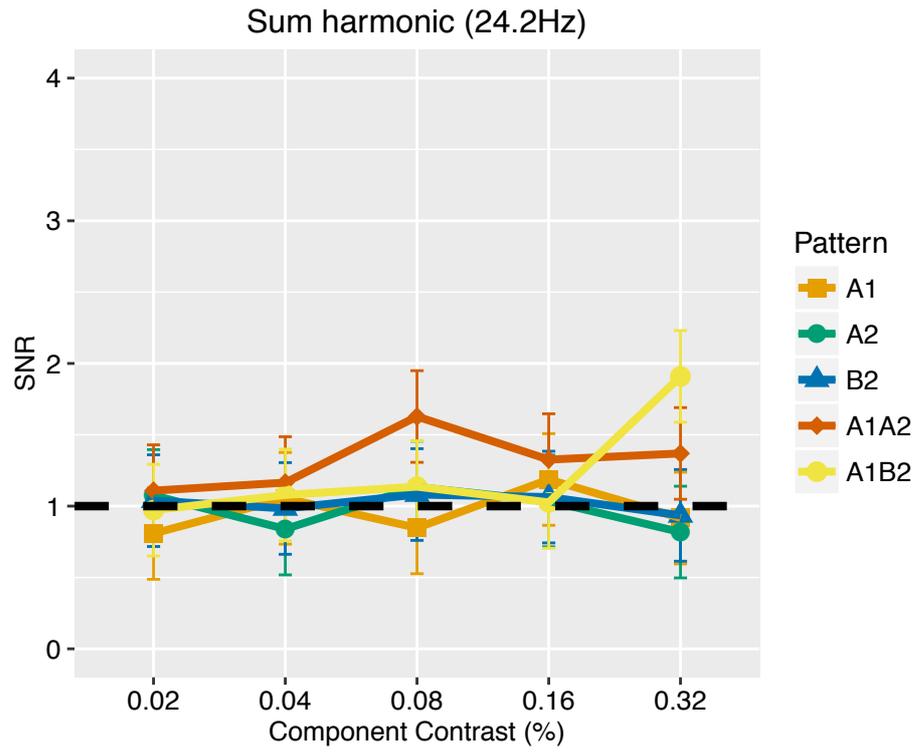


Figure 5.11. Contrast response functions for SNRs at the harmonic sum intermodulation frequency (24.2Hz) for each condition. The black dashed line represents an SNR of 1, and error bars represent 95% CIs.

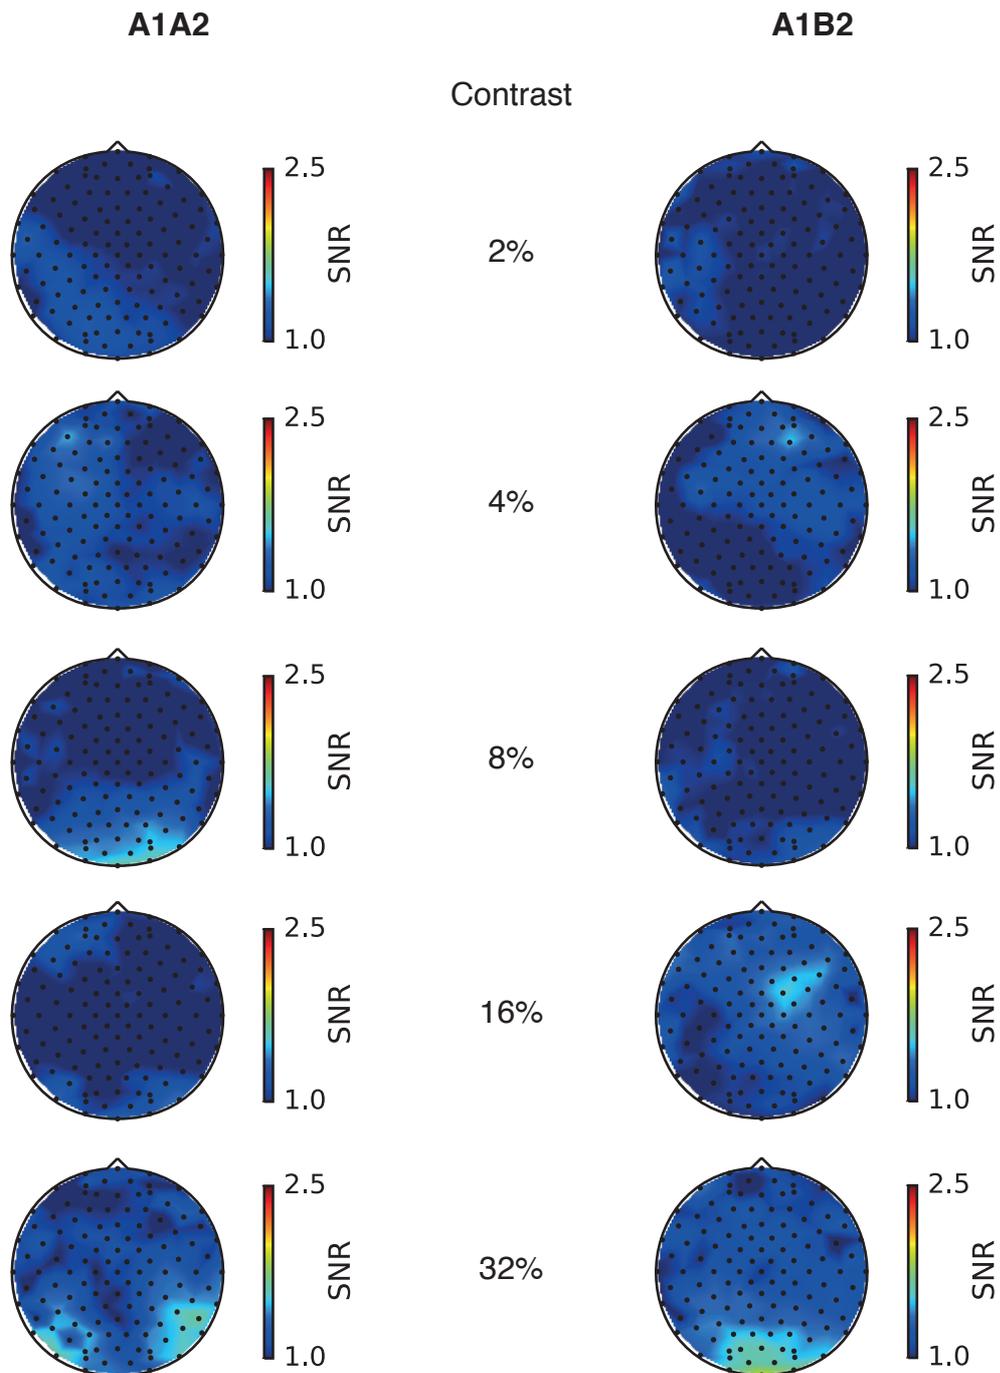


Figure 5.12. SNR topographies at 24.2Hz ($2f_1+2f_2$) for coherent plaid A1A2 and non-coherent plaid A1B2. Contrast increases from the top to the bottom of the figure, from 2% Michelson through to 32% Michelson component contrast. Colour maps were scaled so that if the topography's maximum value was less than 2.5, the maximum of the topography was set to 2.5. This helped avoid noisy topographical maps as SNR approached 1.