Differences in image properties across facial expressions: effects on perceived contrast, stimulus efficacy and oculomotor responses

Abigail L. M. Webb

Supervisor: Prof. Paul B. Hibbard
Dr. Rick O'Gorman

Department of Psychology
University of Essex

This dissertation is submitted for the degree of
Doctor of Philosophy

July 2019

Acknowledgements

"(...) change comes like a little wind that ruffles the curtains at dawn, and it comes like the stealthy perfume of wildflowers hidden in the grass."

- Steinbeck, J. (2008). *Sweet thursday*. Penguin.

I would like to express my gratitude to family and friends for the endless support over recent years, without which this journey would not have been possible. I feel incredibly fortunate to be surrounded by so many wonderful individuals who have inspired optimism and motivation, and provided reassurance when navigating personal doubts.

And thank you to Sam, for this same support, but also for his endless optimism and patience, and the continual reminders that there is always goodness.

Thank you also to office 4.738, for the sweetest gestures of support, and for the unwavering satire that made all the deadlines and late nights bearable.

I especially want to thank both supervisors, Prof. Paul Hibbard and Dr. Rick O'Gorman, for being such fantastic supervisors and supporting the progress of my PhD to this point. I am especially thankful to Paul, for all the patience and reassurance, but also for demonstrating that it is possible to be both of everything. I could not have had a better supervisor or mentor.

And to the small and interested, freckled girl; thank you for keeping me on my toes. I love how we have grown together =^._.^=/

Abstract

Studies within visual psychophysics suggest that fearful facial expressions possess a special status within the visual system. This is evidenced across a range of experimental paradigms that together suggest perceptual biases for fearful expressions are driven by their low spatial frequency content. Unaddressed by this account, however, is the extent that expressions of fear compared to other faces differ in terms of the statistical image properties that define facial stimuli (their Fourier amplitude spectra), and the extent that such natural differences in image properties of expressions are influenced under experimental conditions. The present thesis contributes findings to the broader literature, demonstrating that fear expressions in particular are not naturally higher in stimulus properties known to influence salience. Broadband fear expressions instead contain lower levels of RMS contrast compared to other expressions, where this effect becomes most pronounced as their spatial frequency content increases. These expression-related differences in low-level image properties are discussed in relation to contrast normalisation and differences in perceived image contrast. Findings from Experiment 2 emphasise that facial stimuli normalised for physical contrast do not necessarily have the same perceived contrast, and that this depends on the contrast metric used to normalise faces. A fear advantage for perceived salience is significantly influenced by RMS contrast normalisation and high spatial frequency filtering. A behavioural investigation of Hedger, Adams and Garner (2015) shows that these same effects of contrast normalisation may, to some degree, inadvertently enhance the extent that fearful expressions exploit the contrast sensitivity function. Additionally, Experiments 4 and 5 explore these effects using continuous flash suppression (CFS) and saccadic latency (SL). Findings showed that CFs in particular is vulnerable to faces' contrast content and spatial frequency content, and supports evidence of a high frequency-dependent bias for fear expressions. This contrasts with SL data where no such fear advantage, nor one that is influenced by contrast normalisation, is found when using a SL paradigm. Together, these findings are discussed in relation to the proposed perceptual bias for fearful face expressions, and the extent to which these reflect a genuine adaptive phenomena, or an inadvertent effect incurred from stimulus normalisation.

Table of contents

List of figures

List of tables

Chapter 1

Introduction to the threat bias for fearful faces

Brand New. (2017). Science Fiction. Artwork copyright to Brand New (2017).

1.1 Introduction to threat bias theory

Oatley and Johnson-Laird (1987) were among the first to propose that humans are especially sensitive to perceiving local, threat-relevant information in the environment. This is regarded as a phylogenetically old perceptual behaviour that is shared by all primates, thought to have been selected by nature for its benefits in threat avoidance (Lang et al., 1997; LeDoux, 1994; Öhman and Mineka, 2001a). The ability to rapidly and efficiently detect the presence of a possible threat facilitates the production of the appropriate reflexive responses necessary to maximise an individual's chances of survival (LeDoux and Phelps, 1993; Öhman and Mineka, 2001a). This includes responses such as traditional flight or fight motor reactions, the orienting of visual attentional resources toward a location of interest, and the biases that enable the ability to discern and prioritise the processing of threatening information compared to other environmental stimuli (Öhman and Mineka, 2001a; Phelps et al., 2006). An important feature of such a perceptual bias is its theoretical status as an adaptive visual function that occurs at the universal level, such that it is considered to be present in all human individuals (LeDoux and Phelps, 1993; Oatley and Johnson-Laird, 1987; Öhman and Mineka, 2001a). Across the psychological literature, discussions of such perceptual biases call upon a variety of visual stimuli. This includes perceptual biases for threatening objects including various weapons (Koster et al., 2004; Yiend and Mathews, 2001), emotional words (Hunt et al., 2006; MacLeod and Mathews, 1988; Stormark et al., 1995), offensive imagery (Koster et al., 2007), angry facial expressions (Armony and Dolan, 2002; Mogg and Bradley, 2002), and enhanced fear responses within sub-clinical anxiety populations (Fox et al., 2001, 2002; Georgiou et al., 2005). Because the definition of what constitutes threat-relevant visual information is so broadly defined across the literature, it is important to clarify here that this thesis is concerned only with a perceptual threat bias for fearful facial expressions. For clarity, Table 1.1 includes a list of those findings that are often citated as evidence of a threat bias, but that which is not reflective of a specific bias for fearful facial stimuli; the kind that this thesis is concerned with. This focus on the role of fearful facial expressions, as opposed to other forms of threat-relevant stimuli, relies on the notion that fearful facial expressions play a functionally unique role in the visual system (Öhman and Mineka, 2001a). Fear is an automatic, aversive emotional state that occurs as part of an automated response to cues of a possible threat (Ekman, 1999; Öhman and Mineka, 2001a). A typical fear response in humans is associated with multiple muscular actions in the face of the expressor, including vertical widening of the mouth and retraction of the lips, and the extension of the upper eyelid such that the eyes are visibly widened, illustrated in Figure 1.2. (Ekman and Friesen, 1975; Lee et al., 2013). These physical characteristics that categorise a fearful face expression are associated with functional advantages during potentially threatening situations, relevant to both expressor and observer (Susskind et al., 2008). For the expressor, this relates to the sensory benefits gained from raised eyelids and an increased openness of the eyes that increases the visual field by 9.4% into the available periphery (compared to that associated with neutral faces) (Lee et al., 2013). On the other hand, fearful facial expressions also benefit an observer; because the presence of a fearful face signals the possibility of a local threat shared by both expressor and observer, observation of a fearful face elicits escape and avoidance characteristics in the fear system of the observer (Öhman and Mineka, 2001a; Phelps et al., 2006). Perceptual biases for fearful faces are therefore considered a unique aspect of humans' threat-avoidance system that rely on social-environmental cues from conspecifics.

1.2 Perceptual biases for fearful faces

Findings across a range of experimental paradigms evidence perceptual biases for fearful faces compared to other face expressions. In studies that measure differences in the detectability and explicit recognition of emotional facial expressions, fearful expressions are associated with better detection, but poorer recognition compared to happy and neutral faces (Elsherif et al., 2017; Smith and Rossit, 2018). Findings from neuroimaging studies also evidence preferential responses from the amygdala to images of fearful expressions compared to neutral faces, and innocuous objects including images of houses. This is true both

Fig. 1.2 An example of the physical configuration belonging to a fearful facial expression, including the vertical widening of the mouth and eyes, and extension of the upper eyelid. Copyright to IMDb (2018)

when the fearful face is consciously and unconsciously perceived under conditions of visual suppression (Williams et al., 2004).The continuity of this effect under conditions where observers report no awareness of having seen a face supports the notion that such threat-relevant information in fearful faces may access subcortical processing (Bayle et al., 2009, 2011; Williams et al., 2004). Techniques including visual masking and rapid presentation durations (20 milliseconds) allow visual stimuli to be presented outside of observers' conscious awareness. Visually masked fearful expressions, compared to neutral counterparts, facilitate the orienting of visual attention towards a given target stimulus (Carlson and Reinke, 2008). When presented for rapid durations (20 milliseconds), fearful expressions both capture and maintain visual attention for longer periods of time compared to neutral faces. (Bannerman et al., 2009a, 2010). Such studies show that fearful faces benefit from rapid allocation of attentional resources (Bayle et al., 2011; Carlson and Reinke, 2008; Holmes et al., 2005; Yang et al., 2007); even when they are presented in peripheral vision (Bannerman et al., 2012). Indeed, findings of a bias for fearful faces compared to happy and neutral counterparts when presented in the periphery is not uncommon (Bannerman et al., 2012; Bayle et al., 2011; Phelps et al., 2006), where a fearful face can be detected up to 40 degrees of eccentricity (Bayle et al., 2011). Images of fearful faces also emerge faster compared to neutral (Stein et al., 2014), and both angry and happy expressions when faces compete for visual awareness against a salient mask stimulus (Gray et al., 2013).

Chapter 9 of this thesis provides a detailed discussion of the different techniques that are used to measure perceptual biases for fearful faces, and the findings that these different measures yield.

Authors (year)	Reason for disqualification from thesis
Mogg et al. (2004)	Angry faces as threat stimuli
Koster et al. (2004)	Knives, weapons as threat stimuli
Koster et al. (2007)	Images of mutilation
Yiend and Mathews (2001)	Animals and weapons for stimuli. Clinical population.
Fox et al. (2002)	No inclusion of fearful faces
Fox et al. (2001)	Clinical population
Rigoulot et al. (2011)	Panoramic display
Lee et al. (2013)	Schematic facial stimuli
MacLeod and Mathews (1988)	Clinical population
Hunt et al. (2006)	Emotional words as threat stimuli
Alpers and Gerdes (2007)	Schematic facial stimuli
Stormark et al. (1995)	Word stimuli
Eastwood and Smilek (2005)	Schematic facial stimuli, no fear
Georgiou et al. (2005)	Clinical population
Armony and Dolan (2002)	Angry facial stimuli, no fearful

Table 1.1 List of findings not included within this thesis on the grounds that they do not evidence, or measure, a general and universal threat bias for fearful facial stimuli. Included for clarity.

1.3 Visual pathways for processing threat-relevant information

Traditional theories of the threat bias tend to share two features: that the evolutionary relevance of fearful expressions enables these expressions to undergo specialised visual processing via rapid and dedicated visual pathways, and that their function in threat avoidance dictates that such rapid processing is likely to take place independently of conscious awareness (LeDoux, 1994; Oatley and Johnson-Laird, 1987; Öhman and Mineka, 2001a; Tamietto and De Gelder, 2010). The specificity of such threat-responsive visual pathways is most evident when compared to the retino-LGN-PVC pathway; the central visual pathway where a large component of visual processing takes place.

1.3.1 The Retino-LGN-PVC pathway

Under normal conditions for visual processing, information received by the retina is projected to the thamalas and on toward the primary visual cortex (PVC). This is the primary visual afferent, often referred to as the retino-LGN-PVC pathway, or the geniculo-calcorine-pathway (Johns, 2014). In this central visual pathway, visual information from the retina is received by relay stations within the thalamas; a region of the brain designated as the ‘gateway’ to cortical and conscious processing (Johns, 2014). Once information arrives at the thalamas, it is processed via the lateral geniculate nucleus (LGN), located within the pulvinar region of the thalamas. Informaton is then projected from thalamic relay stations to the primary visual cortex via optic radiations (Johns, 2014). Once the information is received by the primary visual cortex, or area V1, extensive processing of the stimulus extends to other layers of the PVC, other retinotopically mapped visual areas, and then on towards other cortical areas such as occipito-parietal regions for egocentric-spatial processing, and also to subcortical structures such as the superior colliculus and amygdala, where the emotional relevance of the stimulus can be registered (Johns, 2014). An alternative visual pathway, however, provides a route for visual information received by the retina to bypass cortical areas, granting it direct access to subcortical areas renowned for rapid processing, including the superior colliculus and amygdala (Johns, 2014; Vuilleumier et al., 2003). Threat-relevant information is thought to undergo processing via such a channel, specifically so that it may benefit from direct access to subcortical areas (Tamietto and De Gelder, 2010). This alternate route, thought to be unique to threat processing, is often referred to as the standard hypothesis, or ‘quick and dirty’ visual pathway (LeDoux, 2012; LeDoux and Phelps, 1993).

1.3.2 The standard hypothesis

Subcortical processing channels are thought to be the mechanisms that allow fear expressions direct access to the amygdala during processing (Tamietto and De Gelder, 2010). Namely, the retino-colliculuar-pulvinar, or extrastriate pathway, was initially proposed as the neural mechanism underpinning the ‘low road’ for processing threatening visual information (LeDoux, 2012; LeDoux and Phelps, 1993; Tamietto and De Gelder, 2010). This standard hypothesis (Tamietto and De Gelder, 2010) is reminiscent of the well-recognised ‘quick and dirty’ visual pathway proposed by LeDoux and Phelps (1993), where threat-relevant information receives prioritised access to visual processing via distinct and phylogenetically old visual mechanisms located within subcortical regions of the brain. Such subcortical processing streams are considered to operate within regions of the brain associated with non-conscious processing of stimuli, and are sometimes referred to as the evoluationarily-old functions of the brain, thought to be shared by all primates (LeDoux, 1994; Tamietto and De Gelder, 2010). In terms of processing time, this retino-collicular-pulvinar pathway is dwarfed by the lengthier processing associated with the retino-pulvinar-PVC afferent. Visual stimuli can take hundreds of milliseconds to achieve conscious visual awareness, and so a mechanism responsible for rapid threat detection are likely to be selected for their ability to forego the temporal costs associated with visual processing at the level of conscious perception (Koch et al., 2005). It is argued that it is therefore beneficial for threat-relevant stimuli to undergo pre-attentive processing in as rapid a manner as possible, in a way that bypasses conscious perceptual processes (Tamietto and De Gelder, 2010).Figure 1.3 (a) depicts *a schematic view of* the functional and anatomical distinction between normal visual processing that takes place via the retino-LGN-PVC pathway, and the specificity of visual processing that operates via extrastriate pathways that selectively process threat-relevant input.

1.3.3 The multiple-waves model

An alternative version of the standard hypothesis model proposes that rapid processing of fear expressions need not exclusively occur at the subcortical level. Instead, Pessoa and Adolphs (2010) propose the multiple-waves model, whereby the importance of cortical regions during the processing of fear expressions is argued to be equal to the role played by the amygdala. Pessoa and Adolphs (2010) propose that information in fear expressions undergoes processing via multiple ‘short cuts’ between visual areas, such as the afferents that form direct connections between the LGN, V2, and inferior temporal regions (TEO), where information here is projected via long-range short cuts accessing ventro-lateral-prefrontal areas directly to the amygdala. Figure 1.3 (b) illustrates such short cuts thought to be implicated as the mechanisms for fear face processing by the multiple-waves model. This model, compared to the standard hypothesis, does not restrict visual information to hierarchical processes such as retino-pulvinar-PVC or retino-pulvinar-collicular pathways (illustrated in Figure 1.3 (a)). Instead, rapid latencies associated with fear expression processing can occur at the subcortical level, but also via processing short cuts located within cortical routes (Pessoa and Adolphs, 2010). Rapid detection of threat-relevant information can therefore operate via projections that both send and receive information to and from cortical regions, including prefrontal cortex areas, the anterior geniculate and medial frontal lobe (Johns, 2014; Pessoa and Adolphs, 2010). The multiple-waves model is upheld by evidence of Bubbles studies; a technique whereby face images are segmented and presented as such to observers during facial expression discrimination tasks. This is discussed in more detail in Chapter 10.

Together, both theoretical and empirical contributions establish the consensus within the literature that fearful expressions have an advantage for rapid visual processing. The general consensus is *that* both subcortical processes and cortical short cuts between visual areas, and is thought to occur with little, if any, recruitment of conscious perceptual functions related to deliberation and evaluation (Morris et al., 1998; Tamietto and De Gelder, 2010). Indeed, these notions are supported by evidence from neuroimaging studies. Neuroimaging data show selective involvement of extrastriate pathways during the generation of reflexive responses to threatening visual stimuli (Pourtois et al., 2006). Namely, this has been evidenced by selective activation of regions such as the temporo-parietal and temporo-occipital junction, and the ventrolateral prefrontal cortex, during reflexive redirection of attentional resources towards new, behaviourally relevant stimuli, including a fearful face expression (Pourtois et al., 2006). This effect was specific to fearful faces, and was not found for expressions of happiness, while both happy and fearful expressions were found to elicit shared activation in fronto-parietal areas including those implicated in conscious deliberation, such as calculated and voluntary control of spatial attention (Armony and Dolan, 2002; Pourtois et al., 2006). Pessoa and Adolphs (2010) also cite neuroimaging evidence in favour of cortical activity during perceptual discrimination of fearful face images. An example is Smith et al. (2005), who include a categorisation task whereby observers are required to discriminate and categorise elements of facial stimuli according to emotional expression.

These discussions relate to the growing consensus in emotion processing theories that lean towards the presence of a general emotion network (Jastorff et al., 2015). Evidence for a general emotion network suggests that emotion processing operates within a shared and expansive neural network that extends between both cortical and subcortical regions. Emotion processing is therefore less of a distinct and emotion-specific isolated processes, but rather a larger network with sub-levels of subtle and intertwined activation patterns that have potential to deploy preferential responses to different emotions. Similar to the multiple-waves model of emotion processing, a general emotion network may accommodate a threat bias that can operate via a shared network between cortical and subcortical areas. Using functional resonance imaging techniques, Jastorff et al. (2015) presented observers with expressions of neutral, happy, angry, sad and fearful emotions, displayed in the form of dynamic (animated) avatars whose gait and posture varied according to the emotional expression. Jastorff et al. (2015) showed activation patterns between several overarching neutral networks, including default and salience areas. Activation patterns in response to emotionally-oriented avatar gaits were observed in lateral parietal and prefrontal cortex regions, and medial temporal lobe areas. Such activation patterns from default network areas implicate the role of empathy and affective processes. Activation patterns from salience network areas included that of the anterior cingulate, frontal-insular cortices and long-range connections between subcortical and limbic structures including the amygdala and putamen. Activation from such salience areas implicates some role of processes related to socio-emotional and involuntary assessment of relevant information. Such a general emotion network that spans between several sub-networks, and one which contains shared and subtle activation patterns for emotion processing, is theoretically compatible with the multiple-waves model of threat processing.

It is perhaps important to note here that the multiple-waves model does not address the extent to which emotional discrimination process refer to the same construct as that of rapid and crude detection of a threat-relevant stimuli, including fearful facial expressions (Pessoa and Adolphs, 2010). It may therefore be that cortical involvement is implicated during distinction of such facial stimuli, but that detection relies more on the registration of crude signals that take place at the subcortical level (Tamietto and De Gelder, 2010). Further evidence is discussed in more detail in the following chapters.

1.4 The visual-based hypothesis

Any given visual stimulus, including a human face, can be represented and understood in terms of both its low-level image properties and higher-level content.

A recent focus within the psychophysical literature, and the approach that this thesis is concerned with, is the extent to which perceptual biases for fearful expressions are driven by *low-level* compared to *high-level* factors (Gray et al., 2013). The first refers to constituent physical properties that together form the image, and this information is extracted in early processing regions of the visual cortex; the latter refers to increasingly complex information related to the image's meaning, processed in temporal-visual regions associated with object recognition (Stojanoski and Cusack, 2014). These approaches are often cited as high-level versus low-level accounts of the threat bias, or the affective versus visual-based hypotheses of the bias for fearful faces (Elsherif et al., 2017; Gray et al., 2013; Hedger et al., 2016). The higher-level approach, or affective hypothesis, implies that fearful faces are salient because of the way in which their semantic and affective content is contextually relevant (Elsherif et al., 2017; Öhman and Mineka, 2001a), allowing them to undergo processing via pathways that recruit input from cortical functions related to cognitive and semantic evaluation (Lang et al.,

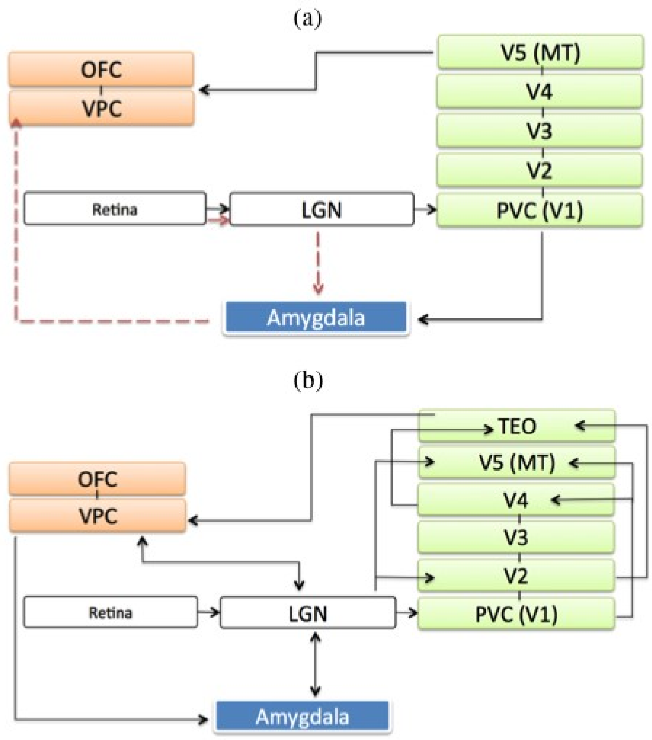

Fig. 1.3 (a) Representations of visual processing via the primary visual afferent (retino-LGN-PVC pathway), compared to that of the extrastriate pathway implicated in the standard hypothesis model (Tamietto and De Gelder, 2010), compared to visual processing associated with the multiple-waves model (b) proposed by (Pessoa and Adolphs, 2010). (a) According to the standard hypothesis model (Tamietto and De Gelder, 2010). Normally formatted connections (–>) illustrate typical visual processing via a retino-LGN-PVC pathway. Red connections (- - -) demonstrate an extrastraite visual pathway, otherwise referred to as the retino-LGN-collicular pathway, whereby certain visual information benefits from direct access to subcortical regions for rapid encoding. *Later stages of processing may relay this information to orbitofrontal (OFC) and ventromedial (VPC) frontal areas of the cortex.* (b) According to the multiple-waves model (Pessoa and Adolphs, 2010), connections between visual areas are not strictly hierarchical, rather, V1 independent 'short cuts' allow information to project between regions. Shortcuts between V2 and inferior temporal areas (TEO) allow visual information access to prefrontal regions *(such as OFC, VPC)*, which in turn send afferents directly to subcortical regions including the amygdala. Information processed here is done so via afferent and efferent pathways between both cortical and subcortical regions.

2000; Stojanoski and Cusack, 2014). For the threat bias, this means that fearful faces achieve perceptual dominance because of the way in which their threat content is evaluated in cortical regions associated with evaluative judgements (Doi and Shinohara, 2013). Such semantic evaluation of fearful faces is thought to be computationally costly and time-consuming due to input from higher order cortical regions, such as the medial prefrontal cortex and orbitofrontal cortex - areas of the brain associated with emotional evaluation and conscious deliberation (Adolphs, 2003). The temporal costs associated with such a high-level pathway for processing fear have been criticised for their inconsistency with a fear-processing pathway that is rapid and automatic; primary and foundational characteristics of the automatic pathway proposed by LeDoux and Phelps (1993) (Gray et al., 2013). Alternatively, the low-level approach, or visual-based hypothesis, suggests that the bias for fearful faces is driven by low-level factors, specifically image properties known to influence image salience (Gray et al., 2013; Hedger et al., 2015, 2016). These properties include an image's contrast and spatial frequency content. It is important to note that these properties do not relate directly to the semantic information that allows, for example, the recognition of the identity, gender or expression of the face. This approach implicates the role of bottom-up visual processes that begin with sensitivity to low-level image properties associated with fearful faces (Elsherif et al., 2017; Gray et al., 2013). The role of low-level stimulus information here is analogous to that proposed by ecological models of stimulus detection, where the salience associated with a given stimulus is driven not by the meaning that it conveys, but by the physical signal, or image attributes, that it creates (Westheimer, 2001). There are two proposed low-level possibilities that may account for biases for fearful faces: the first, that humans have evolved specific visual mechanisms that are sensitive to low-level image properties present in these expressions. During the competition for encoding, these mechanisms allow fearful faces to be rapidly detected and encoded, and so may have been selected by nature for their usefulness in threat avoidance (Oatley and Johnson-Laird, 1987). A slightly different but equally low-level approach suggests that it is not specific mechanisms that have evolved to detect fearful faces, rather, it is the appearance of fearful expressions that have evolved to exploit already sensitive low-level visual functions. Hedger et al. (2015) propose that the physical configuration of fearful expressions have adapted in such a way that they are able to tune into already rapid and sensitive processing pathways in early vision; namely that their physical composition allows them access to processing via the contrast sensitivity function. In either case, the central tenet of the low-level approach to the threat bias is that salient effects associated with fearful expressions occur wholly because of the low-level image properties, or signals, that they contain.

1.5 Summary of Chapter 1

This chapter introduces the notion that fearful face expressions are uniquely salient stimuli. Because of the social and contextually relevant cues emitted by facial expressions, fearful expressions are considered to possess a special status within the visual system because they indicate local threats in the environment. This notion is upheld by evidence from a range of experimental paradigms within cognitive psychology. Recent psychophysical studies propose that these threat bias effects are driven by the physical signal associated with fearful face stimuli, rather than an evaluation of their emotional content. Fearful faces are therefore prioritised during visual processing due to the low-level image properties that they contain, not because of the way in which their semantic content is evaluated during processing. Low-level approaches provide an understanding of the way in which a fearful face, as a physical stimulus, can be broken down and understood in terms of their image properties, or characteristics. These image properties are therefore considered to be the possible mechanisms responsible for perceptual biases for fearful facial expressions. These shall be discussed in Chapter 2.

Chapter 2

Statistical properties of human faces

Sufjan Stevens. (2005). Illinois. Artwork copyright to Asthmatic Kitty (2005).

Chapter 1 introduces the notion that fearful facial expressions are salient simply because of the low-level image properties that they contain, as opposed to an evaluation of their "semantic", or meaningful content. The low-level approach, or visual-based hypothesis, allows visual stimuli to be broken down and understood in terms of their basic image properties and features. The present chapter introduces ways in which such properties can be defined and expressed mathematically, and how they can be used to understand the physical nature of a human facial expression.

2.1 Fourier theory and visual psychophysics

The basic principles of Fourier theory, when applied to visual psychophysics, provide an understanding of the way in which the human visual system extracts and processes information during the perception of a complex visual scene; just as a piece of music may be expressed and understood in terms of many individual frequencies, so can the appearance of a visual scene be decomposed and understood in terms of its basic image properties. In this way, Fourier theory is one method that can be used to identify the physical features and properties that together form the stimulus, or retinal image, of what is being observed.

This technique can be applied under experimental conditions, to establish the individual elements of an image that may be demodulated during perception, and encoded via specific or dedicated visual processing pathways. In this way, Fourier theory has been a valuable tool for understanding human vision and perception as an optical instrument, whose performance suggests functional relationships between subjective observation and the physical world (Campbell and Robson, 1968; Westheimer, 2001).

2.1.1 Local Luminance

The smallest scale of information present in an image relates to the number of pixels that it contains, and their properties. Any image is composed of a given number of pixels, where each pixel is defined as a single unit or area within the image. Pixels can be defined in terms of their luminance; an objective quality that is determined by how much light is emitted, or reflected, from each individual 'point' (Prins et al., 2016). Mathematically, this luminance is defined as a photometric measure of luminous intensity, which is a measure of the radiant energy over a given area, weighted by the spectral sensitivity of the visual system. For example, a photograph of a human face might contain 69,000 pixels that together form a complex face stimulus. This face image can be decomposed and converted into individual elements of information so that it becomes an array of information. Within this array, each pixel from the image can be regarded as an element; 69,000 elements of information, where each value refers to a unit of luminance belonging to each pixel. These elements, or individual pixel luminance intensities, are typically scaled between 0 and 255, relating the 8-bit resolution with which luminance information tends to be stored in digital images. A pixel value of 0 refers to the absolute darkest pixels within an image, and 255 to the very brightest. Extracting this kind of local information provides a detailed and full understanding of an image (Prins et al., 2016). An image's average luminance is also limited in what it tells us about the percept, in that it does not inform of the way in which such information is distributed across the image. Measures that tell of the relationships between these points of information, and how they are distributed within two-dimensional space of an image, can be represented in different ways. Examples are luminance contrast and spatial frequency. These are discussed below.

2.1.2 Luminance contrast

Luminance refers to the properties of individual points within an image, as outlined above. The degree to which these points differ in luminance across the entire image refers to the image's contrast; the aggregate measure of the magnitude of difference between the lightest and darkest points within the image (Campbell and Robson, 1968). This is referred to as an images' global contrast, and there are several different metrics that can be used to measure and quantify this. The most popularly used contrast metrics are Michelson, Weber and Root Mean Squared (RMS) contrast. Michelson contrast is defined as:

$$C_{Michelson} = \frac{L_{max} - L_{min}}{L_{max} + L_{min}} \tag{2.1}$$

Here, the absolute maximum (L_{max}) and minimum (L_{min}) luminance is measured for the whole image. This is perhaps one of the simplest contrast metrics, in that it provides succinct information about the range of luminance within an image (the disparity between the lightest and darkest points). However, because this metric is concerned only with absolute luminance values at the "brightest" and "darkest" polarities, it provides little information about variations in luminance that occur *between* these two extremes (Meese et al., 2017). Such variations that occur within these maximum and minimum locations may have a significant effect on how salient the image appears to an observer, but won't be represented in the image's calculated Michelson contrast because they fall within the two absolute points of interest. Because of this, Michelson contrast is less likely to match, or be equal to, the apparent, or perceived, contrast of a complex image (Meese et al., 2017). Weber contrast, or the Weber fraction, is defined as:

$$C_W = \frac{I - I_b}{I_b} \tag{2.2}$$

Here, the Weber fraction takes the luminance value from both the feature (I) and the image background (I_b). For Weber contrast, contrast energy increases with the size of the image in a way that is unlike the nature of human vision (Meese et al., 2017). Because of this, Weber contrast is often employed as a measure of the relative brightness of a primary feature against a uniform background (Meese et al., 2017; Prins et al., 2016). Root Mean Squared (RMS) contrast refers to measurement of the standard deviation of aggregate local luminance values from each pixel, which is therefore normalised by the luminance mean, defined as:

$$C_{RMS} = \sqrt{\frac{1}{MN}\sum_{i=1}^{M}\sum_{j=1}^{N}(I_{ij} - \bar{I})^2} \tag{2.3}$$

Here, I_{ij} is the luminance of the pixel at position (x, y) in the MxN pixel image, $\bar{I}$ is the mean luminance of the image and the contrast values are normalised such that 0 is the lowest possible value and 1 the highest possible value. Rather than simply capturing the range of luminance values in an image, RMS captures the variability across the whole set of pixels that the image contains. Michelson, Weber and RMS contrast all refer to an image's physical contrast content. An image's apparent contrast, on the other hand, refers to its subjectively perceived contrast, or brightness, that is associated with the images' physical strength (Haun and Peli, 2013; Peli, 1990). When 2 visual stimuli are assigned the same contrast value from a given contrast metric, though they are matched at the physical level, they will not necessarily possess the same apparent contrast (O'Hare and Hibbard, 2011). In other words, two face images that are physically matched for physical contrast may still differ in terms of how salient an observer perceives them to be. Psychophysical studies show that the degree of this disparity between physical and apparent contrast varies as a function of certain features that images contain (O'Hare and Hibbard, 2011). This point is returned to in Chapter 4.

2.1.3 Spatial frequency

Spatial frequency content refers to the number of variations in light intensity, or luminance, that occur across an image. In the Fourier transform, these variations that occur across the space of an image are defined in terms of sine-wave patterns; spatial patterns that vary in terms of their amplitude and orientation. Their amplitude indexes the magnitude of difference between the lightest and darkest points within its waveform pattern, while frequency and orientation refers to the way in which luminance intensity varies as a function of the orientation and distance between pixel locations within an image, i.e. their horizontal and vertical coordinates (Humphreys and Bruce, 1989). The frequency informs how many cycles of pattern there are per unit of space. Different ranges of spatial frequency information carry different information about the qualities and features of an image. Low spatial frequency components are recognised as image attributes that convey a coarse and "rough" representation of the changes in light intensity that occur across the image. Conversely, it is the high frequency components that provide a finer representation of these variations, providing the visual information necessary for identifying local features such as fine lines and details in an image (Humphreys and Bruce, 1989; Milner and Goodale, 2006). Low spatial frequency information is also thought to benefit from faster processing through the visual system compared to higher spatial frequencies (Bullier, 2001). Using the Fourier transform, it is possible to extract from an image a particular "band" of frequency information, and to recreate the image using only this range of spatial frequency content. This filtering technique makes it possible to isolate a particular band of frequency information, and to measure its influence on the appearance of the image, and other aspects of its perception. An example of this is demonstrated in Figure 2.2, where a normal face image containing intact broad spatial frequency information (both low and high frequencies) has been filtered to contain either its high or low frequency components. Here, it is apparent that the high frequency components convey finer information, such as the edges and expressional lines within the face that correspond to abrupt spatial changes within an image (Bar, 2004), but provide less information about larger-scale intensity changes across the face, such as the difference in luminance between the fact and hair. In comparison, the face filtered to contain only low frequency components reveals more information about the images' configuration, the proportion of facial features, and changes in intensity across the face, but appears as a very indistinct and crude impression of the expression that the face displays (Bar, 2004).

The spatial frequency content of an image will influence its perceived, apparent contrast. Differences between images' physical and apparent contrast are most pronounced when images contain different spectral ranges of spatial frequency information (O'Hare and Hibbard, 2011). It has been argued that, when visual stimuli are presented at suprathreshold, such that they are presented above their associated detection threshold, their perceived contrast stabilises to a constant level that is independent of spatial frequency content (Peli et al., 1996). However, other studies have identified clear differences in apparent contrast, in broadband images, that depend on their spatial and temporal frequency content (Bex and Langley, 2007; O'Hare and Hibbard, 2011), and mirror differences found in contrast sensitivity. Additionally, the relationship between an image's physical and perceived contrast is not necessarily linear. Although contrast sensitivity tends to fall as spatial frequency increases, contrast constancy refers to the point at which the perceived contrast of a grating becomes constant regardless of its physical contrast, or associated threshold (Georgeson and Sullivan, 1975).

Fig. 2.2 A sample of filtered facial stimuli from Vuilleumier et al. (2003). Left: a raw face image containing broadband frequency information. Centre: the same face image filtered to contain only its high frequency components. Right: the same face when filtered to contain only low frequency components.

2.1.4 Amplitude and phase spectrum

The Fourier transform allows the information that forms a complex percept to be decomposed into sinusoidal components, characterising images in terms of their orientation, contrast and spatial frequency content. The association between an image's contrast and spatial frequency scales is defined as the image's Fourier amplitude spectrum (Baker and Graf, 2009; Humphreys and Bruce, 1989). In a static image, the Fourier amplitude spectrum therefore describes the degree of contrast energy at each spatial frequency and orientation (Bieniek et al., 2012). An image's Fourier phase spectrum, on the other hand, determines how this energy is disseminated across the image; providing information about local properties and features within the image, such as the way energy at high frequencies converges for contours and edges (Baker and Graf, 2009; Bieniek et al., 2012). The process of phase scrambling preserves an image's amplitude spectrum, such that this physical characteristic remains the same, while disrupting its phase spectrum such that the distribution of information is no longer representative of the features that define the image. This technique is a useful tool for distorting the configural information in an image while preserving image properties such as contrast and spatial frequency. The outcome is an image whose content cannot be explicitly recognised, but whose contrast and spatial frequency are identical to those of its original form (Baker and Graf, 2009). An example is shown in Figure 2.3, where the phase scrambled version of a face image renders it unrecognisable, due to the way in which the distribution of information has been randomised, thus disrupting the configural information necessary for recognition (Farah et al., 1995). During analyses of images' physical composition, Fourier transforms describe images in terms of the sum of the components that they contain. This provides an understanding of the way that information is *arranged* within a given visual stimulus, as opposed to extracting units of information at the local level, such as what is achieved from measures of image luminance. Theoretically, however, the ability to dissociate between images' amplitude and phase spectrum (and to filter images accordingly) is a process that allows us to measure the extent that a given aspect of face perception relies on constituent image features, such as spatial frequency and orientation, or the higher-level abstraction of these properties that occur during image categorisation, i.e. evaluation of the image's meaning (Stojanoski and Cusack, 2014).

2.2 The contrast sensitivity function

Aspects of human vision have been shaped by evolution for efficient encoding of basic image attributes that are consistent across viewing conditions, such as contrast and spatial frequency information. The contrast sensitivity function is one example of such an adaptation, referring to the basic visual function that is responsible for detecting changes in contrast (Campbell and Robson, 1968). In everyday viewing conditions, the contrast sensitivity function modulates observers' ability to distinguish and discern the image properties that compose an object compared to its neighbouring information (Santos et al., 2009). Under experimental conditions, contrast sensitivity is indexed by an observer's ability to detect a sinusoidal grating stimulus composed of certain spatial and temporal frequency and orientation information. This is defined in terms of visual contrast thresholds; evidenced by the minimum level of intensity difference between these light and dark components required to detect a barely

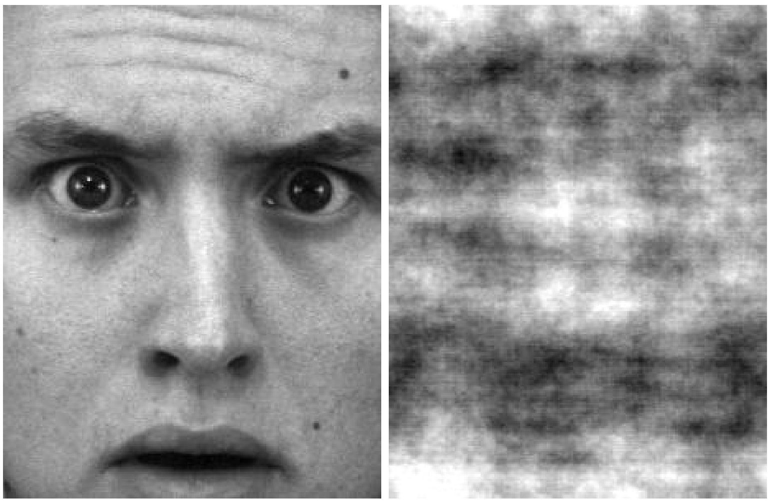

Fig. 2.3 A fearful face (Lundqvist et al., 1998), shown in its normal format (left), and the same image when it is subjected to Fourier phase scrambling (right). Both images are identical in terms of their amplitude spectrum, but the randomised phase spectrum assigned to the manipulated image renders it visibly unrecognisable. This is analogous to the removal of configural information that renders face images unrecognisable (Farah et al., 1995)

visible stimulus (Lee et al., 2014; Milner and Goodale, 2006). The contrast of a sinusoidal grating is determined by its maximum and minimum luminance whose values range from 0 to 1 (Milner and Goodale, 2006), where the contrast sensitivity function is generated from differences in thresholds for discerning contrast at different spatial frequencies. Psychophysical studies show that the human visual system is most sensitive to detecting changes in contrast occurring around a certain range of spatial frequency information; around 4 cycles per degree (cycles/deg) (Sunness et al., 1997). Contrast sensitivity declines for spatial frequencies lower or higher than this band of information, creating the typical U shape of contrast threshold by spatial frequency plots (Campbell and Robson, 1968). Contrast sensitivity is therefore an early dimension of vision, optimised for detecting changes in *luminance* around 4 cycles/degree, as a compromise between accommodating the typical amplitude spectrum of natural images, such as landscapes, and preserving the signal-to-noise ratio of the encoded information (Atick and Redlich, 1992). This category of stimuli show certain regularities regarding the way in which their information is distributed. The way in which contrast energy decreases as a function of increasing spatial frequency information accounts for the characteristic 1/f function and slope belonging to these images (Field, 1987; Tolhurst et al., 1992; Tversky and Hemenway, 1983). These physical regularities in ecological viewing scenes account for the contrast sensitivity function.

2.3 Image properties of human faces

Every day the human visual system is presented with an array of visual information, a lot of which can be classified as belonging to specific image categories. An example of this is the 1/f amplitude spectrum that defines images of natural scenes. Exemplars within this image category, such as different real world landscapes, are regular in terms of both their appearance and their physical composition (Costen et al., 1996). Physical differences in the amplitude spectrum occur *between* image categories (Torralba and Oliva, 2003). While studies concerned with how such information differs according to image categories have tended to focus on natural scenes, such physical regularities and variation are also characteristic of images of human faces. Redies et al. (2007) have been among the first to demonstrate that the Fourier amplitude spectrum associated with photographs of human faces is different from that associated with natural viewing scenes. When scaled and plotted logarithmically, the curve denoting the association between spectral power and spatial frequency is visibly steeper for face stimuli than it is for natural scenes. This amplitude by spatial frequency slope shows that, in comparison to natural landscape images, information in face images declines at a faster rate as spatial frequency increases, accounting for the unique steepness of face's amplitude spectra slope when compared that of natural images. In other words, human faces contain less information around higher spatial frequencies when compared to natural viewing scenes. Figure 2.4 shows data from Redies et al. (2007). This study analysed and compared the Fourier spectra belonging to photographs of faces from the Yale and AR face, artistic depictions of human faces, and natural scenes. The steepness associated with the slopes of face photographs deviates significantly to that of artist's depictions, that appear to share greater physical similarity to natural scenes. Importantly, Redies et al. (2007) note that the presence of person information surrounding the face, such as the shoulders and neck, increases the complexity of the image, and reduces the steepness of the slope. Normalising the face stimuli such that eye locations are fixed and centralised across all faces shows that the associated steepness of the slope occurs because of the face itself, rather than because of neighbouring information around the face (Redies et al., 2007). Another physical regularity observed in photographs of human faces is the regions of sharpness that occur across the face. Pixel intensities for human faces are more regularly distributed across the image than they are for other object categories, such as chairs, and natural scenes (Torralba and Oliva, 2003). It is proposed that this is due to the potential for variation in complex visual patterns that can occur within natural scenes, depending on, for example, whether they are man-made, or contain specific objects and people. Normalised face stimuli, on the other hand, are more constrained in terms of the physical complexities that they might contain. Because of this, the distribution of sharpness information across faces is more consistent (Torralba and Oliva, 2003).

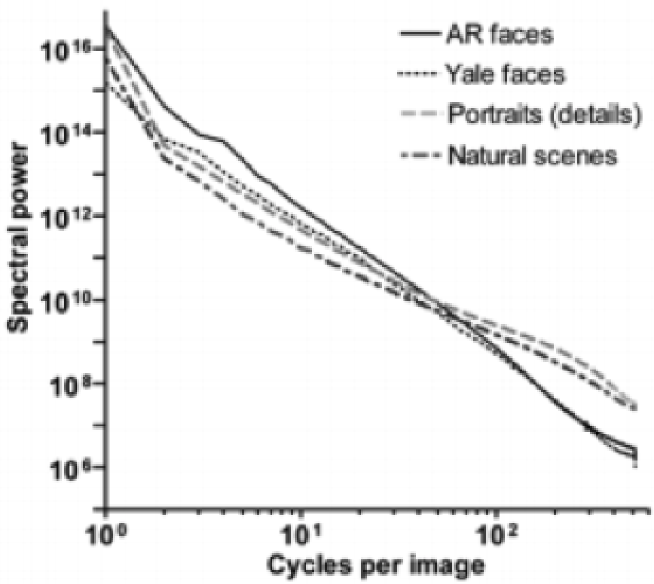

Fig. 2.4 Data from Redies et al. (2007), depicting the logarithmic average spectral power curve for each of the 3 image categories. The logarithmic average for photographs of human faces (AR and YALE databases) is steeper than for face portraits and natural scenes, demonstrating that faces contain less information as spatial frequency information increases. Reprinted from Network: Computation in Neural Systems, 18(3), Redies, C., Hänisch, J., Blickhan, M., Denzler, J. (2007). Artists portray human faces with the Fourier statistics of complex natural scenes, 235-248. Copyright (2007), with permission from Taylor Francis.

2.3.1 Distinct roles of spatial frequency in face perception

Identifying the physical composition of face images is a necessary prerequisite for understanding the way in which certain aspects of face perception are functionally dependent on low level stimulus attributes. Psychophysical studies show that certain bands, or ranges, of spatial frequency information are used depending on the nature of the task (Crouzet and Thorpe, 2011). For example, low spatial frequency components of around 2 to 8 cycles/face provide a coarse representation of a face, with little information about the finer details that it contains, but enough information to convey crude emotional cues (Calder et al., 2000; Costen et al., 1994; Schyns and Oliva, 1999). High spatial frequencies of around 8 to 16 cycles/face provide a finer representation of nuanced facial details, such as age and and expression-related lines (Fiorentini et al., 1983; Liu et al., 2000; Schyns and Oliva, 1999). Theses differences were previously shown in Figure 2.2.

The notion that there may exist specialised visual mechanisms responsible for processing this information is a familiar concept within cognitive psychology, following that specific mechanisms may have been selected for their efficiency when processing physical regularities in the environment; one such regularity being the presence of facial stimuli (Attneave, 1954; Baddeley, 2003). An example comes from evidence that suggests a crucial range of spatial frequency information that is particularly weighted during accurate facial identification. Costen et al. (1996) tasked observers to learn the identity of 6 human faces when these face images were filtered to contain a range of frequency information between 4.5 and 22 cycles/face. Observers' response times to correctly identify low frequency filtered faces increased as the range of information was gradually reduced such that higher frequency information was gradually removed from the face. This effect of frequency range was not found for high frequency faces, where the pruning of the spatial frequency range had no adverse effects on observers' response times to identify faces. Their findings suggest that successful facial recognition relies on spatial frequencies between 8 and 16 cycles/face width (Costen et al., 1996). A similar study by the same authors showed that faces whose spatial frequency content was decreased from 22 to 10.5 cycles per face did not result in adverse effects on observers' accuracy or response times, but that when this was further decreased to 5.5 cycles per face accuracy and timings for identifying faces deteriorated (Costen et al., 1994). Fiorentini et al. (1983) found that accurate identification of faces composed of less than 5 cycles/face was significantly worse compared to identification of faces containing spatial frequency information greater than 5 cycles/face. Importantly, Fiorentini et al. (1983) presented faces at a viewing distance that inhibited the visibility of spatial frequencies greater than 15 cycles/face, suggesting that the observed critical range of frequency information for accurate face identification must occur between 5 and 15 cycles/face. A study concerned with oculomotor responses to face images showed that eye movements toward the location of a learned face were disrupted when the same intermediate band of frequency information was removed; resulting in longer search times and number of eye fixations for faces (Ojanpää and Näsänen, 2003).

A consensus within the current literature is that different spatial frequency ranges within a face convey information whose relevance varies depending on the need for global compared to local facial information. For example, establishing second-order relations between facial features, characterised by low spatial frequencies, underpins global and holistic face processing (Goffaux and Rossion, 2006). Findings show that spatial frequencies of 2-8 cycles/face convey the global information necessary to form a coarse representation of facial structure relationships (Costen et al., 1996; Näsänen, 1999). For face identification in particularly this may be used for assessing information about cardinal facial features such as the mouth, eyes and nose. Conversely, spatial frequency components between 8-16 cycles/face are thought to provide finer-grained, local details, such as nuanced information relating to expression and age-related facial lines (Fiorentini et al., 1983; Hayes et al., 1986; Schyns and Oliva, 1999).

2.4 Summary of Chapter 2

Images of human faces show certain regularities in their low level properties. These unique image properties associated with face stimuli are the physical 'signatures' that distinguish them from other categories of visual stimuli, such as man-made objects and natural scenes. The information that defines a face as a face, and not, for example, a landscape image, is determined by components associated with its phase and amplitude spectrum. Psychophysical studies show that image qualities such as these convey different information about a face, and because of this may be selectively processed during certain tasks of face perception. Specialised visual mechanisms have evolved for their efficiency in processing commonly experienced visual stimuli; this notion may account for the perceptual biases observed for fearful faces.

Chapter 3

Low spatial frequencies in fearful faces

Manchester Orchestra. (2017). A Black Mile to the Surface. Artwork copyright by Loma Vista Recordings (2017).

Chapter 1 introduces the notion that fearful facial expressions are salient to the visual system because of the low-level image properties that they contain. This approach is termed the visual-based hypothesis, or the low-level approach to the threat bias. The different kinds of image attributes that could account for these effects are introduced in Chapter 2. This chapter introduces the currently accepted notion that the threat bias for fearful expressions is driven specifically by the low frequency components within these faces.

3.1 Revisiting the visual-based hypothesis

Recent discussions of the threat bias explore the extent to which visual biases for fearful expressions are driven by early visual mechanisms, or processes, that may be especially sensitive to the low-level characteristics of fearful expressions. This is the approach adopted by the visual-based hypothesis, or low-level account, that places a particular focus on the role of low level image properties that are already known to mediate image salience, including the mean luminance, contrast and spatial frequency content of the image (Bannerman et al., 2012; Gray et al., 2013; Hedger et al., 2015). In other words, fearful faces are considered to be salient because of their physical composition, rather than because of the emotional and semantic meaning that they convey. Some of the first studies to argue for this low level effect adopt the use of inverted face stimuli that allow a separation of face images' low level and semantic content, by manipulating the faces to disrupt observers' ability to readily interpret the images. Manipulated versions of face stimuli are created by rotating them by 180° and reversing their luminance polarity, such that their darkest components become their brightest, and vice versa. Manipulating facial stimuli using inversion disrupts configural information necessary for successful recognition but at the same time preserves their low-level image composition, including their amplitude spectra (Gray et al., 2013). The result is two versions of the same face stimulus, with identical spatial frequency content and contrast, but with the expression of the manipulated image being unrecognisable (Farah et al., 1995; Gray et al., 2013; Itier and Taylor, 2002). Inversion can be a valuable tool for assessing the extent to which perceptual biases rely on low-level versus higher-level content in faces. If effects found for naturally presented (non-manipulated) faces are also found when the configural information is disrupted -such as in inverted and manipulated faces- it is likely that the effect relies on the low level image properties that are preserved under conditions of manipulation. This method is similar to that of phase scrambling, outlined in 2.1.4. Studies of face detection that use such stimuli show that not only are fearful expressions, compared with neutral and happy faces, more quickly detected when slowly released from visual masking, but importantly, that this effect remains true for manipulated fearful faces (Yang et al., 2007). That perceptual biases for fearful faces remain despite the absence of configural information supports the notion that it is their low-level stimulus properties that determine their salience, as opposed to the evaluation of their emotional significance (Gray et al., 2013; Stein et al., 2014; Yang et al., 2007).

Attributes known to modulate stimulus salience include spatial frequency content and contrast (Hedger et al., 2015). Recent findings from psychophysical and neuroimaging studies generally converge on the notion that it is the low spatial frequency components within fearful expressions, in particular, that drive the threat bias. Low spatial frequency information provides a coarse representation of the changes in contrast that occur within an image (Kaplan and Shapley, 1986), thus building only a global and crude representation of a face, but one that is sufficient enough to convey cues of emotional expression without the need for analysis of fine facial details (see Wang et al. (2015) for review). On this basis, a recent focus has been the role of low spatial frequency information in determining the salience associated with fearful expressions (Bannerman et al., 2012; Holmes et al., 2005; Vlamings et al., 2009; Vuilleumier et al., 2003).

3.2 Neural correlates: how low spatial frequencies in fear expressions are processed in the brain

A central feature of threat bias theories is the key role that the amygdala plays in threat detection and avoidance (Davis and Whalen, 2001; LeDoux, 2012; LeDoux and Phelps, 1993; Öhman and Mineka, 2001b; Vuilleumier et al., 2003). The amygdala is considered to influence attentional processing of emotionally laden stimuli, responsible for eliciting appropriate behavioural responses that will maximise an individual's chances of survival (LeDoux and Phelps, 1993; Phelps et al., 2006). One way in which this may be achieved is through modulating cortical processing of visual information via feedback connections to the visual cortex, allowing the prioritisation of threat-relevant information (Kapp et al., 1994). A thalamas-amygdala sensory pathway allows the amygdala to receive direct input from subcortical regions including the superior colliculus and pulvinar; often referred to as the automatic 'low road' for processing threat-relevant stimuli (Davis and Whalen, 2001; Day-Brown et al., 2010; LeDoux, 2012). Coarse visual cues in fearful facial stimuli are considered to access to these subcortical regions that are sensitive to low spatial frequencies (Vuilleumier et al., 2003). These projections act as the processing gateways to subcortical regions including the superior colliculus, and therefore the amygdala. Low-frequency-tuned subcortical pathways operate via preconscious processes, with little to no input from cortical visual processes (Davis and Whalen, 2001; LeDoux, 2012). Information processed here will therefore benefit from bypassing the temporal costs associated with recruitment of cortical processes (Berson, 1988; Vuilleumier et al., 2003). This proposed dedicated subcortical pathway for processing threat-relevant over other information is physically and functionally distinct from the 'high road' (LeDoux, 2012), whereby visual information undergoes higher level processing via channels projecting from the thalamas to the primary visual cortex, on towards regions such as the extrastriate visual cortex and fusiform cortex; areas associated with cortical evaluation (De Gelder et al., 2005; Vuilleumier et al., 2003). This high road is thought to make particular use of high spatial frequency information, and deals with information relating to perceptually complex facial information, including expression and age-related lines and wrinkles (Schyns and Oliva, 1999). A fearful face processed in this way can therefore be expected to receive access to subcortical regions after they have undergone projections to cortical areas via parvocellular channels (Livingstone and Hubel, 1988; Méndez-Bértolo et al., 2016; Vuilleumier et al., 2003). Notions of both traditional subcortical (such as the 'low road') and low frequency-tuned magnocellular routes for encoding fearful expressions are similar in nature, and in recent discussions the two are fused, such that they can refer to the same function (Méndez-Bértolo et al., 2016; Vuilleumier et al., 2003).

3.2.1 Evidence from neurophysiology and neuroimaging

Findings from studies of neurophysiology show that while cortical visual areas such as the fusiform gyrus respond to both high and low spatial frequencies in fearful faces, activity in this region is less responsive to low frequency information and more responsive to high frequency information (Vuilleumier et al., 2003). Subcortical regions such as the amygdala, however, display selective responses only to the low frequency components in fearful faces (Pessoa et al., 2002; Vuilleumier et al., 2003). Electrophysiological studies have shown fast, selective responses from the lateral amygdala to low spatial frequency components in fearful faces. Here, faster amygdala latencies of around 70ms are associated with low frequency cues in fearful faces, where this selectivity is not observed for natural or happy faces regardless of their frequency content. Rapid responses for fear faces were not observed in the visual cortex, whereas happy faces were associated with longer latencies of around 130ms when composed of higher frequencies (Vuilleumier et al., 2003). The authors emphasise that rapid latency periods for low frequency fearful expressions occurred 30ms *before* responses from face-sensitive visual cortical areas, demonstrating the precedence of selective amygdala responses to low frequency components in fearful faces that are distinct from longer processing routes which include recruitment of visual cortical areas (Méndez-Bértolo et al., 2016; Vuilleumier et al., 2003). Similar evidence comes from studies of event related potentials (ERPs). The P1 is an ERP component associated with changes in visual attentional processing that dictate the allocation of attentional resources to emotionally-laden stimuli (Eimer and Holmes, 2002; Pourtois et al., 2005). Vlamings et al. (2009) observed an association between enhanced P1 amplitudes during the presentation of fearful faces that had been filtered to contain low frequency components. This effect was located in the right hemisphere, suggesting that the presence of a fearful faces is associated with enhanced visual attentional responses in areas of the brain known for facial expression processing (Halgren et al., 2000; Vlamings et al., 2009). Together, these findings suggest that coarse visual cues, or low frequency components in fearful faces, are associated with subcortical routes that process input from magnocellular layers of the lateral geniculate nucleus. The result is that fearful faces, compared to other expressions, benefit from prioritised access to subcortical regions implicated in the generation of automated fear responses (Bayle et al., 2011; Milner and Goodale, 2006; Robinson and Petersen, 1992). These findings from neuroimaging studies are compatible with theories arguing that such fear responses operate at very low levels of cognition (De Gelder et al., 1999; LeDoux and Phelps, 1993; Morris et al., 2001; Ohman and Soares, 1994).

3.2.2 Evidence from behavioural studies

Studies of visual attention investigate the extent to which fearful faces receive preferential allocation of attentional resources. Bannerman et al. (2012) measured reflexive eye movements towards fearful, happy and neutral faces presented for 20ms in the periphery. Broadband faces were filtered to create versions of the faces that only contained high or low frequency components. When faces were composed of broadband frequencies, saccadic eye movements showed biases for both happy and fearful faces compared to neutral expressions. However, when faces were composed only of low frequency information, faster saccades were more strongly associated with fearful than for happy and neutral faces. This effect was not observed when faces were composed of high frequency information, where no differences in saccadic eye movements were observed (Bannerman et al., 2012). Faster reflexive eye movements occurring only in response to low frequency fear faces suggests that coarse visual cues in fearful faces modulate rapid orientation of visual attention.

The observation of a fearful expression also facilitates low spatial frequency processing. Bocanegra et al. (2012) presented observers with masked fearful or neutral expressions, measuring their effect on observers' performance when indicating the location of low and high frequency gratings. Their findings showed faster responses for locating low frequency gratings, compared with slower responses for high frequency gratings. The authors suggest that the process of encoding fearful facial stimuli stimulates coarse processing streams that are necessary for encoding information related to motion, depth, and global cues; important information for navigating a threatening situation (Bocanegra et al., 2012).

However, it is important to note here that findings from Stein et al. (2014) suggest that high frequency components in fearful faces allow them to achieve perceptual dominance faster than neutral and happy faces; an effect that Stein et al. (2014) argue is in agreement with two other studies that decompose facial stimuli using Gaussian windows (Smith and Schyns, 2009; Smith et al., 2005). According to Stein et al. (2014), these data support the view that a bias for fear face expressions operates via rapid cortical 'short cut' connections, such as those outlined in the multiple-waves model (Pessoa and Adolphs, 2010) outlined in Chapter 1. Indeed, these findings are compatible with the Multiple Waves model in that they associate the *discrimination* of fearful expressions with high frequency image components. However, the extent that rapid stimulus detection -that which was measured under binocular rivalry (Stein et al., 2014)- is equal to facial discrimination is not addressed by Stein et al. (2014). This is addressed in the following chapter.

3.3 Summary of Chapter 3

Chapter 3 presents findings that suggest low spatial frequency content plays a key role in determining the perceptual biases for fearful facial expressions. This effect is thought to be driven by the way in which these low level image properties are processed via subcortical processing channels that allow fearful expressions direct access to the amygdala, that is not mediated by cortical processing. Such subcortical processing is therefore rapid, and thought to occur independently of conscious awareness. However, shortcomings that remain unaddressed by these low level approaches raise questions about the exact nature of low-level properties in fearful faces, and how these might be differently measures across experimental designs. These issues are the premises for Experiments 1 and 2, and are introduced in Chapter 4.

Chapter 4

Implications of contrast normalisation: the premise for Experiment 1 and 2

QOSTA. (2017). Villains. Artwork copyright to Boneface (2017).

Chapters 1-3 unpack the currently accepted view that fearful facial expressions are especially salient to the human visual system, and that this salience is accounted for by the way that their low frequency components gain access to early visual processing (Bannerman et al., 2012; Gray et al., 2013; Hedger et al., 2016). The majority of studies converge on this finding. However, a sub-category within the threat bias literature rejects this approach, arguing that higher frequencies are responsible for fearful face biases (Adolphs et al., 2005; Smith and Schyns, 2009; Stein et al., 2014). One possible explanation for such inconsistencies, which remains unaddressed in the current literature, is the effects of contrast normalisation in experimental studies, and deviations between physical and apparent contrast in face images. It is possible that these factors may significantly influence the salience of facial stimuli; an effect that could in turn influence responses from perceptual behaviours. Understanding these effects is the motivation for Experiments 1 and 2. By way of introducing this premise for these first 2 studies, these issues are discussed in more detail below.

4.1 Inconsistencies in findings: evidence from Bubbles studies

The currently accepted account is that low frequency information in fearful expressions is responsible for findings of the fearful face advantage (See Chapter 3 for overview). A small body of evidence opposes this notion, arguing that information from higher spatial-scales determines the salience associated with fearful expressions (Smith and Schyns, 2009). Findings from studies that use a Bubbles technique suggest an important role of high frequencies in fear discrimination. In face perception, this technique involves presenting observers with of a combination of fragments, or bubbles, of a facial expression, but never the whole image (Murray and Gold, 2004). Each fragment, or bubble, consists of a sample from a particular spatial region and frequency range. By repeated presentation of many images, each containing a different combination of samples, the combinations of position and scale required to identify the expression can be determined (Murray and Gold, 2004). An example is the experimental procedure used by Smith and Schyns (2009), where observers were presented with randomly sampled face information and instructed to categorise and label stimuli according to the expression collectively portrayed by the bubbles, or segments of faces. Studies using this technique show that discriminating facial expressions relies on both high and low frequency information, but that discrimination of fearful faces relies more on higher spatial frequencies compared to other expressions, with a particularly important role played by the eye region (Smith and Rossit, 2018; Smith et al., 2005). It is noteworthy to consider here whether the discrimination of a face stimulus, in which the observer is required to judge and explicitly identify the expression, relies on the same information required for rapid detection, in which the observer is merely asked to detect the *presence* of the face. If perceptual discrimination demands more of identification processes, we might expect that it require analysis of detailed information from a face, while rapid detection may require only coarse visual cues. Using the continuous flash suppression (CFS) technique, Stein et al. (2014) measured response times to detect neutral and fearful faces that were filtered to contain high or low frequency spatial content. Their findings showed response times to detect faces were faster for fearful expressions when they were composed of high frequency ($> 6_{cpd}$) rather than low frequency information ($< 2_{cpd}$), contrary to findings from the wider literature (Bannerman et al., 2012; Méndez-Bértolo et al., 2016; Vlamings et al., 2009; Vuilleumier et al., 2003). Stein et al. (2014) interpret these findings as evidence against the presently accepted low-frequency-sensitive magnocellular pathways for fear encoding, arguing the importance of higher spatial information in faces for generating threat bias responses. Here, they propose that the amygdala performs less of a rudimentary response to fearful expressions and more of a secondary response that takes place after expressions have been evaluated for their valence and relevance. This is thought to occur via a cerebral cortex route, responsible for processing faces' high frequency content, such as that outlined by the multiple-waves model (Pessoa and Adolphs, 2010). Stein et al. (2014) do not directly address the inconsistencies between their data and the wider body of literature, but do argue that rapid detection times for fearful expressions do not equate to subcortical processing, on the basis that visual cortical areas are equally capable of such short latency responses.

Short latencies in subcortical regions for fearful expressions are thus not sufficient evidence to assume a low-frequency-magnocellular route for processing fearful expressions. However, it is noteworthy here to re-visit studies demonstrating significant differences in temporal and spatial responses to fear faces between subcortical and cortical visual areas (Méndez-Bértolo et al., 2016; Pessoa et al., 2002; Vuilleumier et al., 2003). Therefore, although visual cortices and subcortical regions are both capable of short latencies, as Stein et al. (2014) argue, the visual cortex shows selective responses to high frequency fearful faces, whereas faster responses to low frequency fear faces are exclusive to subcortical areas (Vuilleumier et al., 2003), supporting the currently accepted notion of distinct low and high roads for processing fearful faces. It may therefore be that like facial discrimination tasks, the use of CFS measures only higher-level *conscious* processing of visual stimuli, but not that which is expected to occur at the subcortical level (Stein and Sterzer, 2014). Indeed, findings from Stein et al. (2014), Smith and Schyns (2009) and Smith et al. (2005) do suggest that high frequency components in facial stimuli and responses from higher-level visual processes. Stein et al. (2014) do also note that the effect of high frequency content on faster detection for fearful compared to neutral faces may reflect differences when identifying expressive compared to non-expressive faces, rather than a threat bias for fearful faces that is driven by high frequencies.

It may be that inconsistencies in biases for fear expressions are in part due to the nature of the task used. Stein et al. (2014) use methods of conscious stimulus detection under conditions of intraocular suppression, and refer to studies of facial discrimination to support the notion that the bias for fear expressions relies on high frequency information. But it is important to note here that both CFS and Bubbles techniques are the centre of unresolved debates regarding their efficacy (Gosselin and Schyns, 2004; Murray and Gold, 2004; Stein et al., 2011; Stein and Sterzer, 2014; Yang and Blake, 2012). A notable argument here is that CFS in particular is praised for its suitability as a measure of higher-level visual and semantic processing (Stein and Sterzer, 2014); features of processing that, by definition, conflict with low-level and rapid detection of fear faces proposed by traditional threat bias theories (LeDoux, 2012; Öhman and Mineka, 2001b; Tamietto and De Gelder, 2010). This is discussed in more detail in Chapter 10.

It is worth including here that some studies produce inconsistent findings regarding the face inversion effect. Studies that both evidence perceptual biases for fear faces and include manipulated versions of face stimuli are consistent in showing that biases for fear faces are preserved under conditions of manipulation (Bayle et al., 2011; Gray et al., 2013; Hedger et al., 2015; Stein et al., 2014; Yang et al., 2007), supporting the notion that such effects are driven by low level image properties that are preserved when configural content is disrupted. However, other studies evidencing the same perceptual advantages in fear perception find this effect for upright faces but *not* inverted faces, where such findings are interpreted as evidence of an emotion effect (Bannerman et al., 2012, 2009a; Holmes et al., 2005; Phelps et al., 2006). It is unclear why this inversion effect is observed in some cases and not others, particularly when data from such cases converge on the role of low frequency factors for determining fear biases. These inconsistencies, and their possible causes, are not acknowledged nor discussed in the current literature.

Chapters 10-12 address the extent that differences in findings are task relevant, but another possible cause is the use of contrast normalisation, addressed in the following section.

4.2 Effects of contrast normalisation

Contrast normalisation, or contrast equalisation, is a technique used to ensure that a given set of visual stimuli share the same physical contrast. This can be done by taking an overall average measure of physical contrast for a given set of stimuli, and then attributing this average to each individual stimulus, thus normalising, or equating, stimuli in terms of their contrast. The general purpose of doing so is to confirm that physical contrast- an image feature known to influence stimulus salience- remains the same across all images. This ensures that observed effects can be unambiguously attributed to the semantic content of the image as opposed to differences in their low level signals. In studies of the bias for fearful expressions, contrast normalisation is often used in conjunction with spatial filtering techniques. The motivation here is to manipulate effects of spatial scale information while stabilising and controlling for those associated with differences in physical contrast.

In theory this is a plausible technique for standardising stimuli at the physical level, however several issues remain unclear that are particularly relevant for facial stimuli: to what extent does normalising faces for contrast remove natural differences in physical contrast that are key determinants of their composition? Does this have an inhibitory or bolstering effect on their salience? Do face stimuli matched for contrast still appear the same in terms of their apparent contrast? and which out of several contrast metrics is the most appropriate given the nature of face images? These issues are not necessarily mutually exclusive, but are intertwined, as is discussed in detail below.

4.2.1 Does contrast normalisation mask natural expression-related differences in physical contrast?

Chapter 2 introduced the notion that image categories, including faces and natural scenes, differ in their typical Fourier amplitude spectra (Redies et al., 2007). However, it remains unknown whether physical differences relating to images' amplitude spectrum occur within facial stimuli, and specifically, whether such differences exist between facial expressions: if the physical configurations of face expressions are distinct because they correspond to unique evolutionary functions (Ekman and Cordaro, 2011), then we might expect differences between expressions' physical features to play key role during visually categorisation and processing of such stimuli (Smith et al., 2005). If these differences are inherent characteristics of faces, then contrast normalisation may mask this variation. Establishing pre-existing differences between expressions' physical properties, such as their contrast and amplitude spectra, is thus necessary to establish how natural properties of expressions respond to contrast normalisation. Specifically for threat bias research, if fearful expressions have evolved to be especially salient to the visual system, then we might expect unique Fourier amplitude spectra for these faces compared to other expressions. Contrast normalisation may mask or otherwise influence the key image properties unique to fear expressions.

4.2.2 Differences in apparent contrast for physically matched stimuli

The purpose of contrast normalisation is to exclude differences in image properties that may otherwise affect stimulus salience. Although this is successful when normalising natural scenes for RMS contrast (Bex and Makous, 2002), not all visual stimuli matched for physical contrast will necessarily have the same apparent contrast. O'Hare and Hibbard (2011) demonstrated this dissociation using filtered random noise stimuli, showing that images' amplitude spectra determine differences in their perceived contrast, even when they have been matched for RMS content. Mid range frequency information in particular is generally perceived as having higher contrast compared to higher and lower ranges of information (O'Hare and Hibbard, 2011). Data from O'Hare and Hibbard (2011) is displayed in Figure 4.2. These findings emphasise the importance of first understanding the amplitude spectra, or contrast by frequency profile, associated with images before they undergo contrast normalisation.

4.2.3 Which contrast metric is most suitable for normalising faces?

Chapter 2 shows that there are several possible contrast metrics that can be selected for normalisation and for interpreting psychophysical results, and that the appropriate metric depends on the nature of visual stimulus. RMS is a commonly used contrast metric, particularly when normalising the contrast of images of both natural scenes and face images. The high degree of consistency that remains between natural images' physical and apparent, perceived contrast after they undergo contrast normalisation suggests that RMS contrast is most representative of the human visual system (Bex and Makous, 2002; Peli, 1990). However, findings show that face images have a significantly different amplitude spectrum compared to that of natural images (Redies et al., 2007), and given that we may also expect expression-related differences in statistical properties *within* face categories, it is reasonable to question whether faces normalised for RMS are guaranteed this same consistency between their physical and apparent contrast. There is evidence that, in broadband stimuli, apparent contrast is affected by the amplitude spectrum of the image, as well as its contrast (O'Hare and Hibbard, 2011). This is shown in Figure 4.2. In other words, while images of natural scenes with similar amplitude spectra, once normalised for RMS contrast, are likely to be matched for perceived contrast, the same may not be true for facial stimuli given that the physical composition of faces and natural scenes are not the same, and that the amplitude spectrum may also differ across facial expressions.

4.2.4 Equivocal use of contrast normalisation

It is unclear whether contrast normalisation affects the perceived salience of facial expressions, and equally whether these effects vary as a function of spatial filtering of faces. This is in part because it is an under-addressed topic, but also because the process by which contrast normalisation is administered varies across different studies. It is therefore difficult to identify possible outcomes of the technique. The point at which stimuli are normalised (i.e. before or after spatial filtering) is often not specified (Bannerman et al., 2012; Hedger et al., 2015; Vlamings et al., 2009; Williams et al., 2004; Yang et al., 2007), or in some studies takes place before *and* after filtering (Stein et al., 2014). In some cases, only contrast is normalised (Bayle et al., 2011; Hedger et al., 2015; Yang et al., 2007), while in others both luminance and contrast are normalised (Gray et al., 2013; Williams et al., 2004). In others, it is not specified whether or not face stimuli underwent any normalisation at all (Bannerman et al., 2010; Carlson and Reinke, 2008; Holmes et al., 2005; Phelps et al., 2006; Pourtois et al., 2006; Schupp et al., 2004; Smith et al., 2005; Whalen et al., 2001). Few studies include statistical tests measuring the effects of normalisation on face stimuli (Vlamings et al., 2009), or to confirm that differences in raw face stimuli were not significantly large enough to equalise stimuli through normalisation. For example, Vuilleumier et al. (2003) measured no differences between faces' luminance. However, none of these studies includes measures of differences in images' perceived contrast before or after stimuli are prepared. The various ways normalisation can be administered, including those listed above, emphasise the absence of a consensus for the way in which such process ought to take place, such as the point during stimulus preparation this ought to occur, and indeed whether contrast normalisation is necessary at all. To best illustrate some of the inconsistencies regarding the standardisation, or normalisation, of facial stimuli, Figures 4.3, 4.4, and 4.5 provide an overview experimental parameters of behavioural and neuroimaging studies that have concluded specialised processing of fear expressions. Each Figure summarises the following information: Authors of individual studies and their associated paradigm and/or experimental task; Overarching Face information column includes: face Database and a list of the emotions included within the study (N= neutral, F= fear, H= happy, D= disgust, S= sad, S=surprise, A= anger); Cropping style of facial stimuli, and inclusion of internal features only, if stated; Size of facial stimuli denoted in degrees of visual angle, unless otherwise stated; and indication of a facial control condition used by study i.e. inversion and/or reversal of luminance polarity (LP). Column for normalisation indicates whether face images were normalised for contrast or luminance. Spatial filtering (SF) range column indicates the range of frequencies used, and associated method. Conclusion column includes a brief summary of effects. All hyphens in table imply that relative information was not evident from the manuscript.

4.3 Summary of Chapter 4

Contrast normalisation is a commonly used technique, employed for its benefits in stabilising differences in image contrast in order to isolate and observe effects of other factors. In face perception, this equalises the physical contrast of face stimuli such that, in theory, they appear the same in terms of their perceived contrast. Contrast normalisation may mask natural physical differences between expressions, which as a result could render image salience non-ecological. It also remains unclear whether faces matched for physical contrast appear the same in their apparent contrast, and if so, which contrast metric most ensures this consistency. To investigate the effects of contrast normalisation on fearful face biases, such as whether they inhibit or bolster effects, it is necessary to first measure natural differences in contrast between expressions, and second, their associated perceived contrast the point at which they are physically identical. These two questions form the premise for Experiments 1 and 2.

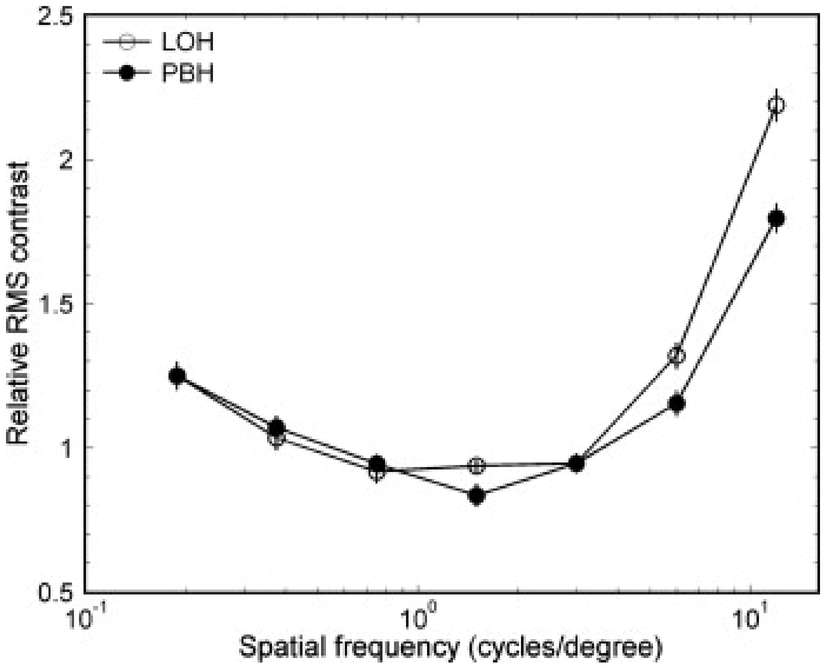

Fig. 4.2 Data from (O'Hare and Hibbard, 2011), illustrating the relative amount of RMS contrast required for perceived contrast to be identical, plotted as a function of images' spatial frequency content. Reprinted from Vision research, 51(15), O'Hare, L., & Hibbard, P. B. (2011), Spatial frequency and visual discomfort, Pages No., 1767-1777. Copyright (2011), with permission from Elsevier.

Authors (year)	Paradigm/task	Face info.				Normalisation		SF filtering		Conclusion
		Database (emotion)	*Crop*	*Size*	*Control cond.*	*Contrast*	*Luminance*	*Range*	*Method*	
Whalen et al. (1998)	fMRI/observational/ backward masking	Ekman & Friesen '76 (fhn)	-	-	None	-	-	None	None	Amygdala responds to masked fear (not happy) faces. Substantia innominata response to happy and fear.
Whalen et al. (2001)	fMRI/observational	Ekman & Friesen '76 (fan)	-	-	None	-	-	None	None	Dorsal amygdala activation greater for seen fear (compared to angry) faces. Ventral amygdala response to both compared to neutral.
Whalen et al. (2004)	fMRI/masking/ observational	-	Eye-whites	-	None	-	-	None	None	Ventral amygdala activation greater for fear than happy eye-whites.
Williams et al. (2004)	fRMI/binocular rivalry	Ekman & Friesen '67 (fhn)	Rectangle	2.2x2.9°	None	Yes	Yes	None	None	Amygdala responds to masked fear and happy faces compared to neutral faces
Schupp et al. (2004)	ERP/observational	KDEF (hsfdsan)	-	-	None	-	-	None	None	Augmented EPN amplitudes for threat compared to neutral faces pronounced 200-280ms post stimulus onset.
Pourtois et al. (2006)	fMRI/covert attentional cueing/dot-probe	Ekman & Friesen '76 (fhn)	-	-	None	-	-	None	None	Compared to happy faces, fear expressions associated with faster RTs for locating target, regardless of valid or invalid trial. Specific areas for fear-valid and fear-invalid trials, compared to happy equivalents.
Carlson & Reinke (2008)	Covert attention/dot-probe /backward masking	3D faces Gur et l. '02 (fn)	Oval, internal	5x7°	None	-	-	None	None	Masked fear expressions on congruent trials facilitate attention, compared to neutral
Bannerman et al. (2010)	Saccadic latency\attentional cueing	Ekman & Friesen '76 (fn)	-	7.5x11.2°	None	-	-	None	None	Fear expressions capture and inhibit disengagement at rapid (20ms) durations, compared to neutral faces. Manual RTs not affected by expression at this duration.
Bannerman et al. (2009)	Saccadic latency	Ekman & Friesen '76 (fn)	Oval, Internal	7.5x11.2°	Inversion	-	-	None	None	Saccades deployed faster for fear compared to neutral faces at rapid (20ms) presentation durations. This effect true for manual responses, but only at 500ms presentation durations. Effect not preserved for inverted faces.

Fig. 4.3 Table 1/3: Summary of experimental parameters for threat bias studies.

Authors (year)	**Paradigm/task**	**Face info.**				**Normalisation**		**SF filtering**		**Conclusion**
		Database (emotion)	*Crop*	*Size*	*Control cond.*	*Contrast*	*Luminance*	*Range*	*Method*	
Phelps, Ling & Carrasco (2006)	Attentional cueing/contrast sensitivity	Ekman & Friesen '76 (fn)	Rectangle, head	5x6.6°	Inversion	-	-	None	None	Contrast sensitivity task preceded by a fearful compared to neutral face results in lower contrast thresholds (Experiment 1). Transient fear-cues in periphery also enhances contrast sensitivity. Effects not preserved for inverted faces (Experiment 2).
Bayle et al. (2011)	Response time detection/peripheral vision	NimStim (fdn)+ Ekman & Friesen '76 +own images	Oval, internal	7.5x10.5°	None	-	Yes	None	None	Fear expressions maintained detectability up to highest eccentricity (40°).
Gray et al. (2015) *Experiment 3*	b. Continuous flash suppression	NimStim (fhan)	Oval, internal	2.1x2.8°	Inversion + LP reversal	RMS	Yes	None	None	Fear expressions break suppression faster compared to happy, angry, neutral faces. Effects were preserved for control counterparts (inverted + LP reversal).
Bannerman et al. (2012)	Saccadic latency	KDEF (fhn)		6.9x10.4°	Inversion	No	Yes	Low(<0.8$_{cpd}$) High(>3.3$_{cpd}$)	SO Butter-worth	Saccades deployed faster for fear compared to neutral and happy faces. General emotion preference for broad faces, but a fear bias for low frequency fear faces compared to neutral and happy. No effects for high frequency faces (Experiment 1). Effects not preserved for inverted faces (Experiment 2).
Vuilleumier et al. (2003)	fMRI/observational	KDEF + own images (fn)	Rectangle, internal	-	None	-	No	Low(<6$_{cpi}$) High(>24$_{cpi}$)	-	Fear associated with preferential activation in bilateral amygdala, insula, temporal areas. Bilateral amygdala also preferentially responds to Broad and low frequency content. No specific responses to high frequency fear faces.
Stein et al. (2013)	b. Continuous flash suppression	Ekman & Friesen '76 + NimStim (fn)	Rectangle, internal	3.2x3.8°	Inversion	RMS	Yes	Low(2$_{cpd}$) High(6$_{cpd}$)	SO Butter-worth	Broad and high frequency fear expressions break suppression faster than neutral faces (Experiment 1 and supported by Experiment 2). Effect preserved for inverted faces.
Vlamings, Goffaux & Kemner (2009)	ERP/manual response time	NimStim (fn)	Oval, internal	6.3°		RMS	Yes	Low(<12$_{cpi}$) High(>36$_{cpi}$)	Gaussi-an filters	P1 amplitude specifically to low frequency fearful faces.

Fig. 4.4 Table 2/3: Summary of experimental parameters for threat bias studies.

Authors (year)	Paradigm/task	Face info.				Normalisation		SF filtering		Conclusion
		Database (emotion)	*Crop*	*Size*	*Control cond.*	*Contrast*	*Luminance*	*Range*	*Method*	
Yang, Zald & Blake (2007) *Experiment 1,2*	b. Continuous flash suppression	Ekman & Friesen '76 (fhn)	Square, internal	1.9x1.9°	Inversion	25% RMS	-	None	None	Fearful faces break suppression faster than neutral and happy faces. Effect were preserved for inverted versions of faces.
Yang, Zald & Blake (2007) *Experiment 3*	b. Continuous flash suppression	Ekman & Friesen '76 (fhn)	Eye-region	.5x1.8°	Inversion	25% RMS	-	None	None	Eye-region of fearful faces break suppression faster than those of neutral and happy faces.
Holmes, Green & Vuilleumier (2005)	Probe detection task	Ekman & Friesen '76 (fn)	Rectangle, internal	8.1x10.9°	None	-	-	Low(<2_{cpd}) High(>8_{cpd})	-	Low frequency fearful compared to neutral faces cue attention towards congruently-located probes. No emotion effects found for high frequency faces (Experiment 1). Effects were not preserved for inverted versions of faces (Experiment 2).
Hedger et al. (2015) *Experiment 2*	b. Continuous flash suppression	NimStim (fhan)	Oval, internal	6.2x4.1°	Inversion + LP reversal	-	-	None	None	Fear faces have overall visibility advantage over happy, neutral and angry faces. Effects preserved for inverted faces.
Hedger et al. (2015) *Experiment 3*	Backward masking	NimStim (fhan)	Oval, internal	6.2x4.1°	Inversion + LP reversal	-	-	None	None	Fear faces have overall visibility advantage over happy, neutral and angry faces. Effects preserved for inverted faces.
Smith et al. (2005)	Explicit categorization task/Gaussian windows "bubbles"	California database (fhadss)	Gaussian windows	na	None	-	-	120-60; 60-30; 30-15;15-17.5; 7.5-3.8_{cpd} samples	Gaussian windows	Role of high frequency ranges when discriminating/categorising fearful expressions.
Pessoa et al. (2002)	fMRI	Ekman & Friesen '76 (fhn)	Oval, internal	-	-	-	-	None	None	Selective responses from fusiform gyrus for fearful faces.
Bocanegra, Huijding & Zeelenberg (2012)	Spatial attention/dot-probe task	Ekman & Friesen '76 (fn)	Oval, internal	7° diameter	None	-	-	None	None	Fearful facial cues facilitate responses to low frequency Gabors and inhibit rapid responses to high frequency Gabors.

Fig. 4.5 Table 3/3: Summary of experimental parameters for threat bias studies.

Chapter 5

Experiment 1: differences in RMS contrast and Fourier amplitude spectra between facial expressions

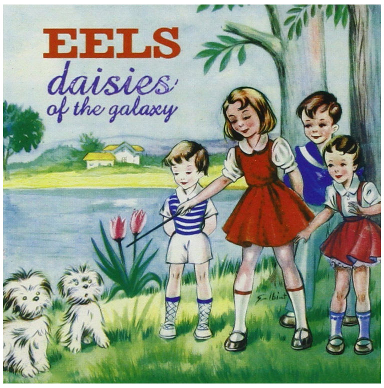

Eels. (2000). Daisies of the Galaxy. Artwork copyright to DreamWorks (2000).

5.1 Introduction

Experiment 1 investigates whether facial expressions naturally differ in terms of their global contrast, in the absence of any contrast normalisation. Faces used in this experiment were analysed in their original form, and also following low, mid-range and high spatial frequency filtering, in order to establish whether the physical contrast associated with these faces varies with spatial frequency. Face images used were a 140-face image sample extracted from the Karolinska Directed Emotional Faces set (Lundqvist et al., 1998). Measures of differences in the Fourier amplitude spectra between natural, raw face expressions are also analysed. Analyses address the question of whether such commonly used face expression images naturally differ from each other at the physical level, and if so, how such differences are influenced when images are spatial-frequency filtered.

5.2 RMS contrast between face expressions

5.2.1 Methods

Stimuli were grayscale front-view face pictures of 140 individuals (70 male, 70 female) extracted from the Karolinska Directed Emotional Faces (KDEF) set (Lundqvist et al., 1998). Each individual portrays 1 of 5 expressions: fear, anger, happiness, disgust or neutral. The numbers of 5 expressions were selected to provide a measure of differences between expressions that is broader compared to a fear-neutral comparison, or a negative (fearful) and positive (happy) comparison often adopted. Faces were cropped to included internal features only, and their dimensions were 300 (height) x 230 (width), measured in pixels. Assuming a hypothetical viewing distance of 65cm, the stimulus size of faces was 7.09 degrees. Spatial frequency versions of faces were created using a second-order Butterworth filter in MATLAB, defined as follows:

$$G(f) = \frac{1}{1 + (\frac{f}{f_0})^{2n}} \tag{5.1}$$

Where *G(f)* refers to the gain function; *f* is the spatial frequency; f_0 refers to the cut-off frequency, and *2n* refers to the order (second-order). This created low (LSF), mid-range (MSF) and high spatial frequency (HSF) versions for all face stimuli, in addition to the original broad spatial frequency (BSF) versions. The cut-off frequencies were f<1_{cpd} for low frequency faces, 1<f<6_{cpd} for mid-range frequency faces, and f>6_{cpd} for high-pass cut-off faces. Low and high bandpass cut-offs were consistent with those used by Stein et al. (2014) and Vlamings et al. (2009). Face stimuli were not normalised for contrast content, to ensure that any naturally occurring variation in physical contrast would be preserved. Examples of these faces are shown in Figure 5.2. The RMS contrast belonging to each face expression was was measured across 140 KDEF actors, and performed in MATLAB.

5.2.2 Results

Differences in RMS contrast across expression, for each of the broad-, low-, mid-range, and high-frequency stimuli are summarised respectively under the Data Tables section of this chapter. Average RMS contrast values for each expression at each spatial frequency are illustrated in Figure 5.3. Separate Analysis of Variance (ANOVA) analyses and Sidak comparisons were performed for each of the 4 spatial frequency categories of facial stimuli. Sidak comparisons compare differences in RMS contrast between all face expressions, rather than exclusively between fearful and counterpart faces, because this experiment was interested in overall expression-related differences between faces.

Broad spatial frequency faces

A repeated measures Analysis of Variance (ANOVA) showed a significant effect of expression (F(4, 556) = 11.25, p<.001, ηp_2 .07). Sidak-corrected pairwise comparisons showed that

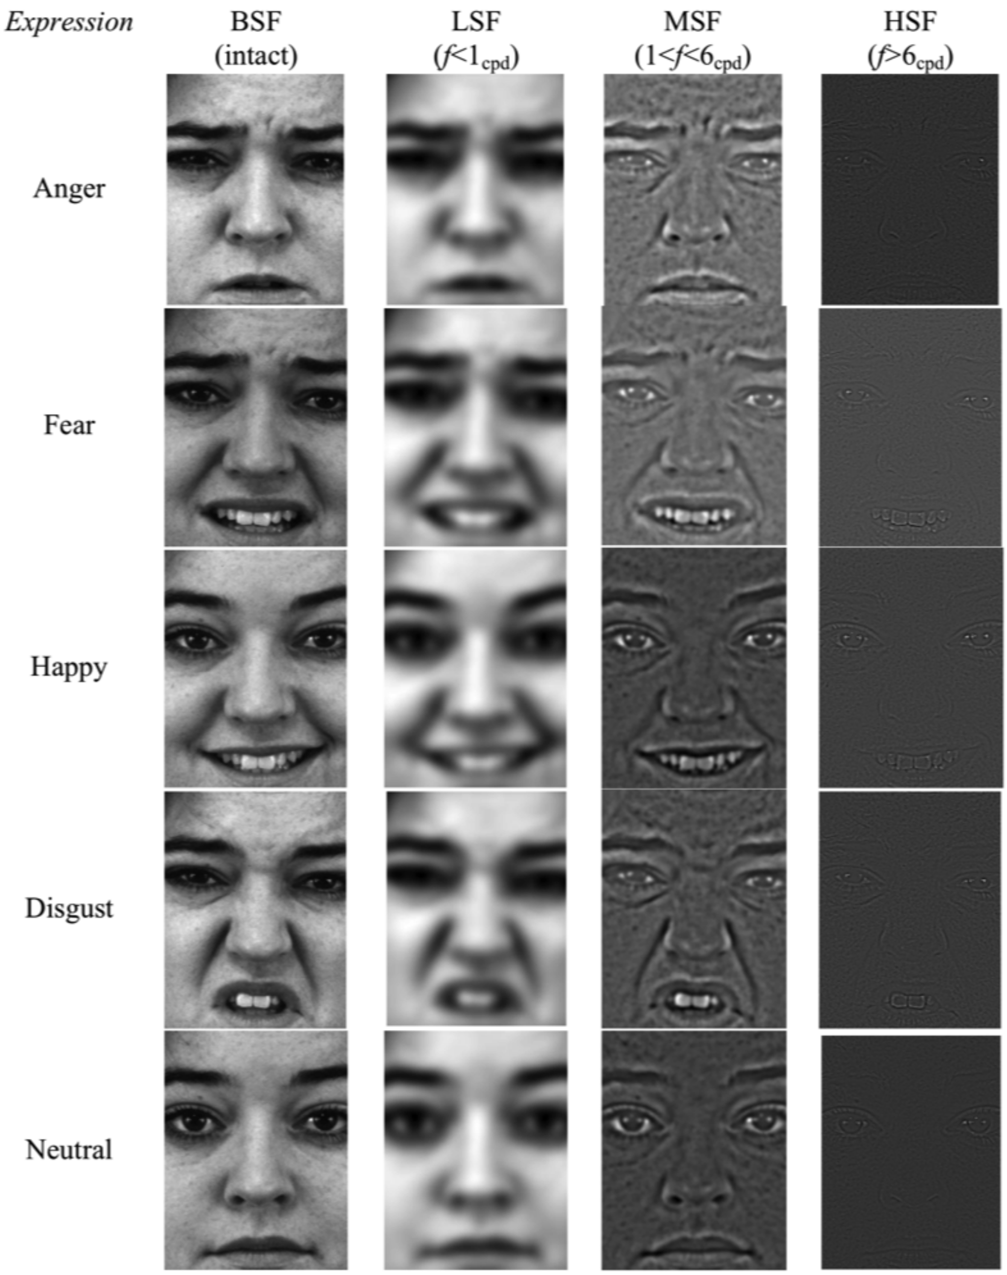

Fig. 5.2 An actor extracted from the KDEF database, portraying a fearful facial expression (Lundqvist et al., 1998). The image is shown in its normal, unfiltered (broadband) form, and when it is LSF, MSF, and HSF filtered.

for broadband faces, RMS contrast values were highest for angry expressions and lowest for fearful faces. RMS contrast was significantly (p <.05) higher in angry compared to fearful and happy expressions; lower in fearful compared to disgusted and angry faces; lower in happy compared to angry and disgusted faces; and higher in disgusted compared to happy and fearful expressions. Data are summarised in Table 5.1.

The Fourier amplitude spectrum was also measured for all 140 broad faces, shown in Figures 5.4 and 5.5. Average slope values for each face expression, pooled across 140 actors, illustrated the association between faces' contrast energy and spatial frequency. A repeated measures ANOVA showed a significant effect of expression, ($F(4, 556) = 22.63$, p<.001, ηp_2 .14). Sidak pairwise comparisons revealed differences in the Fourier amplitude spectrum of faces that varied according to facial expression. In particular, broadband fearful face images possess a steeper Fourier amplitude slope compared to neutral and angry facial expressions. Data are summarised in Table 5.2.

Low spatial frequency faces

A repeated measures ANOVA showed no significant effect of expression ($F(4, 556) = 1.68$, p.15, ηp_2 .01), showing that RMS contrast does not differ between expressions when they are low-frequency filtered.

Mid-range spatial frequency faces

A repeated measures ANOVA showed a significant effect of expression ($F(4, 556) = 22.86$, p<.001, ηp_2 .14). Sidak-corrected pairwise comparisons showed that for mid-range filtered faces, RMS contrast values were lowest for angry and highest for disgusted expressions. RMS contrast was significantly (p<.001) higher in disgusted compared to all other expressions; lower in angry compared to happy and disgusted faces; lower in fear compared to disgusted faces; and higher in happy compared to angry faces. Data are summarised in Table 5.3.

High spatial frequency faces

A repeated measures ANOVA showed a significant effect of expression ($F(4, 556) = 41.93$, p<.001, ηp_2 .23). Sidak-corrected pairwise comparisons showed that for high frequency filtered faces, RMS contrast values were highest for disgusted and lowest for fearful expressions. RMS contrast was significantly (p<.001) higher in disgusted compared to neutral, fearful and happy expressions; higher in angry compared to all but disgusted faces; and lower in happy compared to angry and disgust expressions. Fearful faces were significantly lower in contrast compared to all but neutral faces, where RMS did not differ between neutral and fear. Data are summarised in Table 5.4.

a

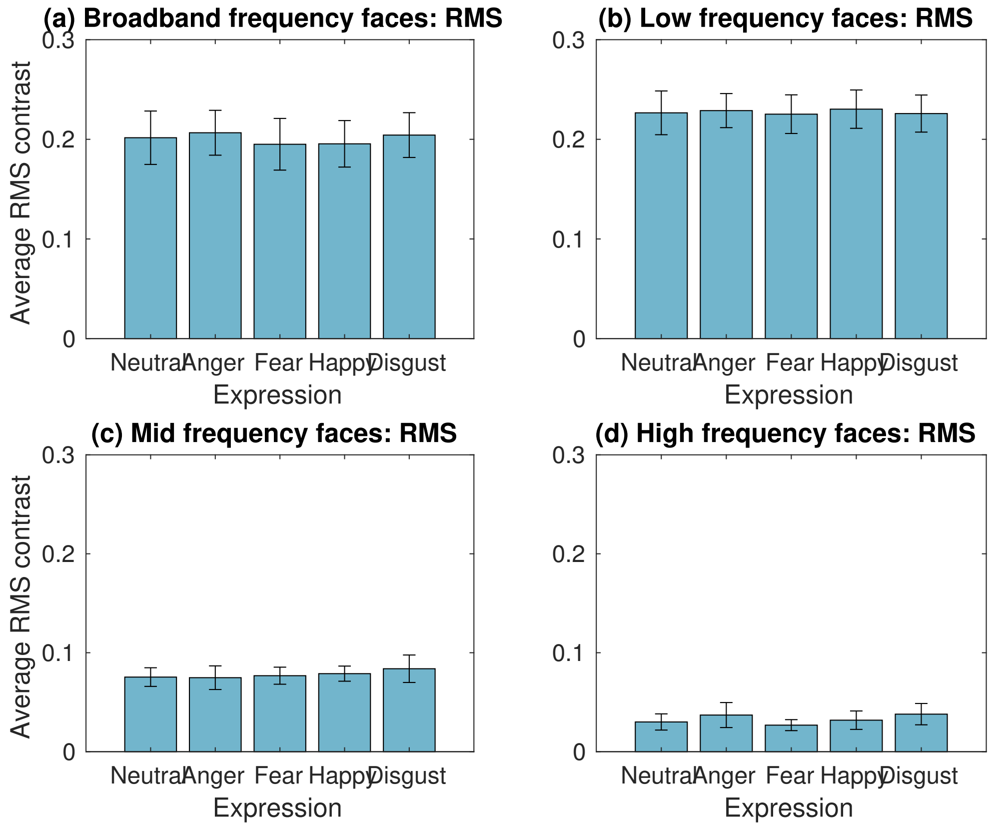

Fig. 5.3 Mean RMS contrast calculated for 140 broadband faces (a), and low (b), mid-range (c), and high (d) frequency filtered versions of these faces. Error bars depict the associated standard deviation. All face stimuli were unmatched for physical contrast. The RMS contrast between facial expressions differs most when these faces are filtered to contain a higher range of spatial frequency content, where fearful facial expressions are consistently lower in RMS contrast compared to neutral and other emotional expressions. Error bars show associated standard deviations.

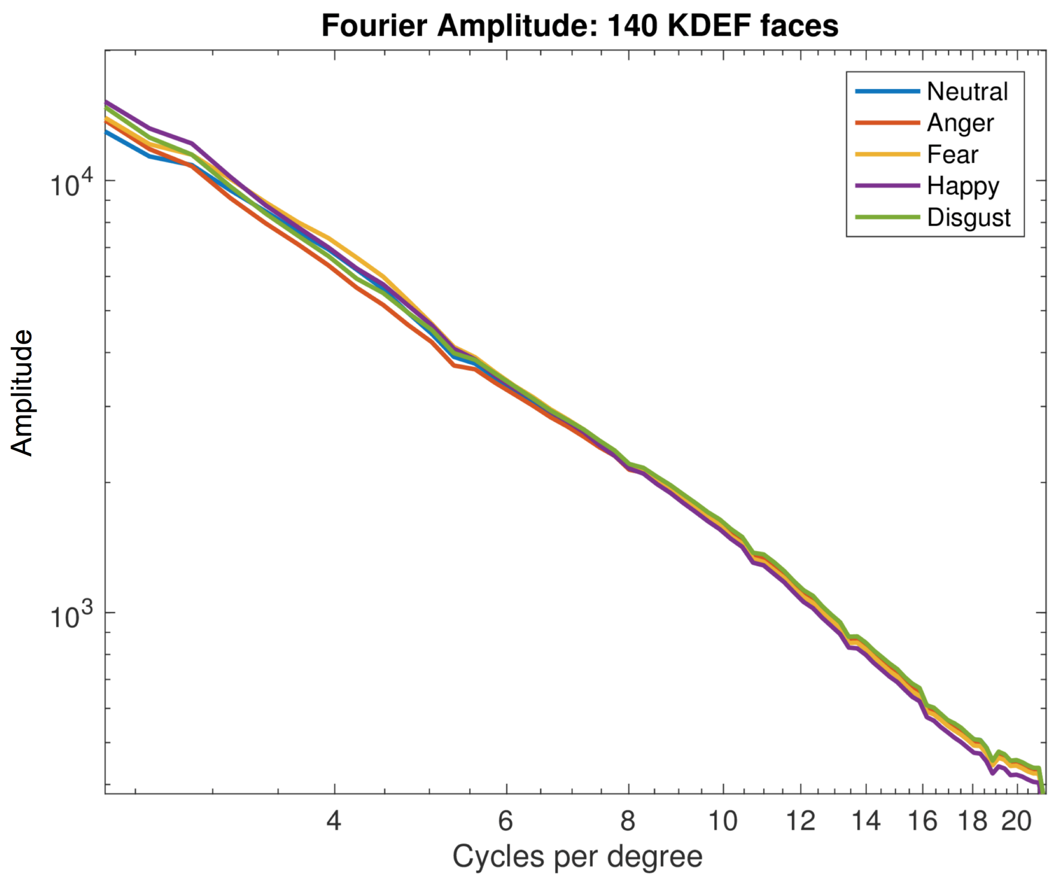

Fig. 5.4 The Fourier amplitude spectrum for 140 KDEF faces, each portraying 1 of the 5 face expressions. Average amplitude slopes plot the association between contrast energy and spatial frequency content. Face images analyses were done so in their natural format, such that they were not normalised for physical contrast, or subjected to spatial filtering.

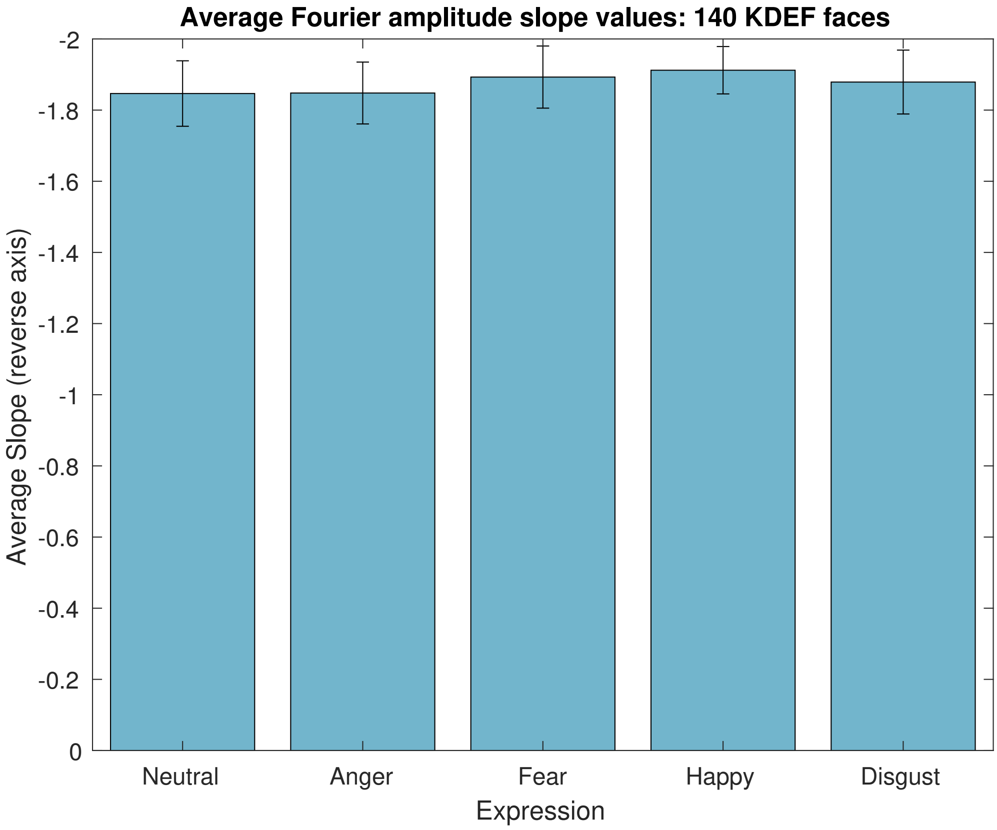

Fig. 5.5 Average Fourier amplitude slopes, presented in bar-graph format, for neutral, angry, fearful, happy and disgust faces, presented against a reversed y-axis. ’Larger’ average slope values represent a steeper amplitude slope, denoted by more extreme negative values. Error bars show associated standard deviations.

5.3 Data Tables

Expression comparisons	Sig	CI (95%)
Neutral		
Anger	.59	-0.013, 0.003
Fear	.21	-0.002, 0.015
Happy	.30	-0.002, 0.014
Disgust	.98	-0.011, 0.005
Anger		
Neutral	.59	-0.003, 0.013
Fear	.001	0.003, 0.020
Happy	.001	0.003, 0.019
Disgust	.99	-0.006, 0.010
Fear		
Neutral	.21	-0.015, 0.002
Anger	.001	-0.020, -0.003
Happy	1.0	-0.009, 0.008
Disgust	.01	-0.017, -0.001
Happy		
Neutral	.30	-0.014, 0.002
Anger	.001	-0.019, -0.003
Fear	1.0	-0.008, 0.009
Disgust	.02	-0.017, -0.001
Disgust		
Neutral	.98	-0.005, 0.011
Anger	.99	-0.010, 0.006
Fear	.01	0.001, 0.017
Happy	.02	0.001, 0.017

Table 5.1 RMS contrast differences between broadband expressions. Multiple pairwise and Sidak comparisons of RMS contrast between 5 unfiltered facial expressions.

Expression comparisons	Sig	CI (95%)
Neutral		
Anger	1.0	-0.027, 0.030
Fear	<.001	0.018, 0.075
Happy	<.001	0.037, 0.094
Disgust	.01	0.004, 0.061
Anger		
Neutral	1.0	-0.030, 0.027
Fear	<.001	0.016, 0.073
Happy	<.001	0.036, 0.093
Disgust	.02	0.002, 0.059
Fear		
Neutral	<.001	-0.075, -0.018
Anger	<.001	-0.073, -0.016
Happy	.45	-0.009, 0.048
Disgust	.84	-0.043, 0.015
Happy		
Neutral	<.001	-0.094, -0.037
Anger	<.001	-0.093, -0.036
Fear	.45	-0.048, 0.009
Disgust	.01	-0.062, -0.005
Disgust		
Neutral	.01	-0.061, -0.004
Anger	.02	-0.059, -0.002
Fear	.84	-0.015, 0.043
Happy	.01	0.005, 0.062

Table 5.2 Differences in the average Fourier amplitude slope associated with 5 different broadband facial expressions. Multiple pairwise and Sidak-corrected comparisons between the average slope value for each expression.

Expression comparisons	Sig	CI (95%)
Neutral		
Anger	1.0	-.002, .004
Fear	.95	-.004, .002
Happy	.06	-.007, 8.6e-5
Disgust	<.001	-.011, -.004
Anger		
Neutral	1.0	-.004, .002
Fear	.69	-.005, .001
Happy	.01	-.007, -5.33e-4
Disgust	<.001	-.012, -.005
Fear		
Neutral	.95	-.002, .004
Anger	.69	-.001, .005
Happy	.65	-.005, .001
Disgust	<.001	-.010, -.003
Happy		
Neutral	.06	8.6e-5, .007
Anger	.01	5.33e-4, .007
Fear	.65	-.001, .005
Disgust	.001	-.008, -.001
Disgust		
Neutral	<.001	.004, .011
Anger	<.001	.005, .012
Fear	<.001	.003, .010
Happy	.001	.001, .008

Table 5.3 RMS contrast differences between mid-range filtered expressions. Multiple pairwise and Sidak-corrected comparisons of RMS contrast between 5 unfiltered facial expressions.

Expression comparisons	Sig	CI (95%)
Neutral		
Anger	<.001	-.010, -.003
Fear	.05,	-3e-6, .006
Happy	.71	-.005, .001
Disgust	<.001	-.011, -.004
Anger		
Neutral	<.001	.003, .010
Fear	<.001	.006, .013
Happy	<.001	.001, .008
Disgust	.99	-.004, .002
Fear		
Neutral	.05	-.006, -3e-6
Anger	<.001	-.013, -.006
Happy	<.001	-.008, -.001
Disgust	<.001	-.014, -.007
Happy		
Neutral	.71	-.001, .005
Anger	<.001	-.008, -.001
Fear	<.001	.001, .008
Disgust	<.001	-.009, -.002
Disgust		
Neutral	<.001	.004, .011
Anger	.99	-.002, .004
Fear	<.001	.007, .014
Happy	<.001	.002, .009

Table 5.4 RMS contrast differences between high filtered expressions. Multiple pairwise and Sidak-corrected comparisons of RMS contrast between 5 unfiltered facial expressions.

5.4 Conclusion

The findings from Experiment 1 are best discussed in relation to those from Experiment 2, as both studies were similarly motivated and composed 2 stages of an investigation of the extent to which facial expressions differ at the physical and perceptual level. A detailed discussion of the present findings are therefore provided in the following chapter, under the Discussion for Experiment 2.

However, a brief summary of the present findings is that facial expressions *do* naturally differ in terms of their RMS contrast content and Fourier amplitude spectrum. In terms of RMS contrast, expression-related differences become more pronounced as the images are filtered to contain higher frequency information. For broadband stimuli, fearful faces are significantly lower in RMS contrast compared to other expressions except from neutral. When the same faces are filtered to contain only low frequency information, findings revealed no effect of facial expression. Mid-range frequency fearful faces are lower in RMS contrast compared to disgust expressions. When high frequency filtered, the effect is most pronounced, and fearful faces are significantly lower in RMS contrast compared to all emotional expression except neutral faces. Fourier analysis data showed that facial expressions naturally differ in their Fourier amplitude spectra. This was true for 140 KDEF faces analysed here, where these images were not normalised for contrast or spatial frequency filtered. Fear expressions in particular, compared to neutral and angry expressions, contain less information as spatial frequency content increases, though they do not differ significantly when compared to disgust or happy expressions.

Chapter 6

Experiment 2: expression-related differences in perceived image salience when faces are physically matched for contrast

ASIWYFA. (2017). The Endless Shimmering. Artwork copyright to ASIWYFA (2017).

6.1 Introduction

Experiment 2 investigates whether normalising images of facial expressions for their physical contrast content negates any remaining differences in their perceived image salience, or apparent contrast. Subjective ratings of apparent contrast are obtained for 5 facial expressions, calculated using Michelson and RMS contrast. In order to establish the extent that expression-related differences in apparent contrast vary as a function of the spatial range used to filter images, these analyses are also performed for low, mid-range and high spatial frequency versions of the faces. These analyses were performed to determine the effects of contrast normalisation on the salience of facial stimuli.

6.2 Methods

6.2.1 Participants

A total of 32 individuals took part in the study. All participated in the experiment as part of a credited research module assessment. The first session of data collection obtained data from 19 participants: measuring responses for broad-, low-, and high-frequency stimuli. A second session of data collection obtained data from an additional 19 participants, in order to include an experimental condition for mid-range frequency stimuli. All participants had normal or corrected to normal vision. The number of participants was determined by previous studies of a similar nature (Gray et al., 2013; Peli, 1997).

6.2.2 Stimuli and apparatus

Stimuli were presented using a VIEWPIXX 3D monitor, viewed from a distance of 96cm. For the duration of the experiment, participants used a chin rest to maintain this viewing distance. The monitor screen was 52cm wide and 29cm tall. The screen resolution was 1920x1080 pixels, with a refresh rate of 120Hz and an average luminance of 50 cd/m^2. Each pixel subtended 1.0 arc min. Stimuli were presented at 10 bit resolution. The luminance responses of the monitor was linearised using gamma correction based on measurements taken using a Minolta LS-110 photometer. Stimuli were generated and presented using MATLAB and the Psychophysics Toolbox extensions (Brainard and Vision, 1997; Peli, 1997). Stimuli were grayscale and front-view face pictures of 16 individuals (8 male, 8 female) extracted at random from the Karolinska Directed Emotional Faces set (Lundqvist et al., 1998). These included internal features only, and each individual portrayed 5 emotions: fearful, anger, happy, disgust or neutral. These 80 (16 identifies x 5 emotions) face pictures were composed of intact broad spatial frequencies. Filtering the images using a second-order Butterworth filter in MATLAB created low (LSF), mid-range (MSF) and high spatial frequency (HSF) versions for all face stimuli. The cut-off frequencies were $f<1_{cpd}$ for low frequency faces, $1<f<6_{cpd}$ for mid-range frequency faces, and $f>6_{cpd}$ for high-pass cut-off faces. Low and high bandpass cut-offs were consistent with those used by Stein et al. (2014) and Vlamings et al. (2009). A sample is shown in the previous chapter, Figure 5.1. All stimuli were presented in their normal form (upright with retained luminance polarity), and in a manipulated form (rotated by 180° with inverted luminance polarity). This manipulated condition was included for its consistency with previous studies (Bannerman et al., 2010; Gray et al., 2013); the process of inverting and negating the luminance of stimuli prevents the use of configural information needed for accurate emotional recognition, but preserves low-level image properties such as contrast and spatial frequency content (Gray et al., 2013; Tanaka and Farah, 2003).

6.2.3 Procedure

Participants were tested individually and informed prior to the experiment that the study was concerned with face perception. The University of Essex University Ethics Committee approved the study and all participants gave written, informed consent.

Each trial commenced with the target face on the left or right side of the screen, chosen at random. A small red reference mark above and below the image was used to indicate which of the two stimuli was the target. Target face images were presented within a range of randomly selected contrast values between 10 and 20% Michelson contrast, and were presented one at a time adjacent to the reference face portraying a broadband, neutral expression, and against a grey uniform background. This reference face always portrayed a neutral expression with a fixed 10% Michelson contrast. This same reference was present for all trials, and participants were informed of its role as the reference stimulus to which they would perceptually match target stimuli. Using left and right arrow keys, participants adjusted the contrast of the target face until it appeared to have the same contrast as the reference. The study was separated into 2 sessions. The first part collected data from 19 participants in a single block of trials for broad, low and high frequency faces.

A second session of data collection was added for an additional 19 participants in order to include a condition for only mid-range filtered faces. Because of this, mid-range faces were presented as part of a single block.

6.3 Results

Data tables are displayed under the Data Tables section of this chapter.

For each stimulus, we recorded both the Michelson contrast at which each face appeared to have the same contrast as the neutral reference stimulus, and also calculated the RMS contrast of each stimulus at this setting. Low physical contrast settings reflect less physical contrast necessary for observers to perceptually match a 10% Michelson reference face stimulus. In other words, the stimulus appeared to have a relatively high contrast, while high settings mean that it had a relatively low apparent contrast, since the settings represent the observers' compensation for differences in apparent contrast. Data were analysed using separate 2 way ANOVAs for each of the 4 frequency conditions: broadband, low, mid-range and high frequency face stimuli. Within each frequency condition, Sidakc-corrected comparisons were used to explore differences in apparent contrast between fear expressions and each counterpart emotion, including neutral faces.

6.3.1 Unfiltered broadband faces

Mean contrast settings for broadband faces are shown in Figure 6.2. For set contrast calculated using RMS contrast, a two-way Expression (neutral, angry, fearful, happy, disgust) x Manipulation (normal, manipulated) repeated measures ANOVA showed significant main effects of expression and manipulation ($F(4, 72) = 7.72$, $p<.001$, ηp_2 .30; $F(1, 18) = 9.89$, $p<.001$, ηp_2 .35, respectively), but no significant interaction ($F(4, 72) = 2.17$, $p<.001$, ηp_2

.10). An effect of manipulation showed that for manipulated faces (inverted and with reversed luminance polarity), participants tended to make lower contrast settings to match the non-manipulated standard. Lower contrast settings mean that manipulated faces were perceived as higher in contrast compared to normally presented face stimuli. This is consistent with findings that demonstrate an effect of stimulus' contrast polarity on image salience; when judging contrast, darker (or negative-appearing) regions of an image, such as those that characterise the appearance of manipulated faces, are associated with responses associated with those of high contrast and therefore greater salience (Haun and Peli, 2013). A significant effect of expression showed that broadband expressions differ in terms of RMS contrast at the point at which they are perceptually matched. Paired comparisons between fear and expression counterparts were explored using Sidak corrections (α= 0.0063, corrected according to 8 comparisons) showed that unfiltered fear expressions require significantly less RMS contrast in order to perceptually match a 10% reference stimulus, compared to neutral, happy and disgust faces. When images were manipulated, fearful expressions still required significantly less RMS contrast compared to neutral and angry manipulated faces. No other significant differences were observed. Data are summarised in Table 6.1, and illustrated in Figure 6.2 (a).

Analyses were repeated for apparent contrast data that was calculated using Michelson contrast, to provide a comparison of how apparent contrast might vary as a function of this contrast metric. A two-way Expression (neutral, angry, fearful, happy, disgust) x 2 Manipulation (normal, manipulated) repeated measures ANOVA showed a significant effect of expression and manipulation (F(4, 72) = 3.56, p.01, ηp_2.16; F(1, 18) = 10.39, p.005, ηp_2 .36, respectively), but no significant interaction (F(4, 72) = 2.18, p.07, ηp_2.10). Manipulated faces require significantly less Michelson contrast compared to their normally presented face counterparts in order to match a reference stimulus. Sidak-corrected paired comparisons (α= 0.0063) revealed no significant differences in apparent contrast between fearful faces and their face counterparts. Data are summarised in Table 6.2, and illustrated in Figure 6.2 (b).

6.3.2 Low frequency filtered faces

Mean contrast settings for low frequency faces are shown in Figure 6.3. For set contrast calculated using RMS contrast, a two-way Expression (neutral, anger, fearful, happy, disgust) X Manipulation (normal, manipulated) repeated measures ANOVA showed no significant effect of expression, and a significant effect of manipulation (F(4, 72) = 1.47, p.21, ηp_2 .07; F(1, 18) = 41.82, p<.001, ηp_2 .69, respectively). No significant interaction was observed (F(4, 72) = .03, p.99, ηp_2 .002). Manipulated faces were perceived as more salient compared to their normally presented counterparts, and thus required less RMS contrast to match the reference stimulus.. As detailed in the analysis of broadband face stimuli, this is due to the association between the negated regions of images and enhanced perceived salience (Haun and Peli, 2013). Data are illustrated in Figure 6.3, (a).

Analyses were repeated for apparent contrast data that was calculated using Michelson contrast. A two-way Expression (neutral, angry, fearful, happy, disgust) X 2 Manipulation (normal, manipulated) repeated measures ANOVA showed a significant effect of expression

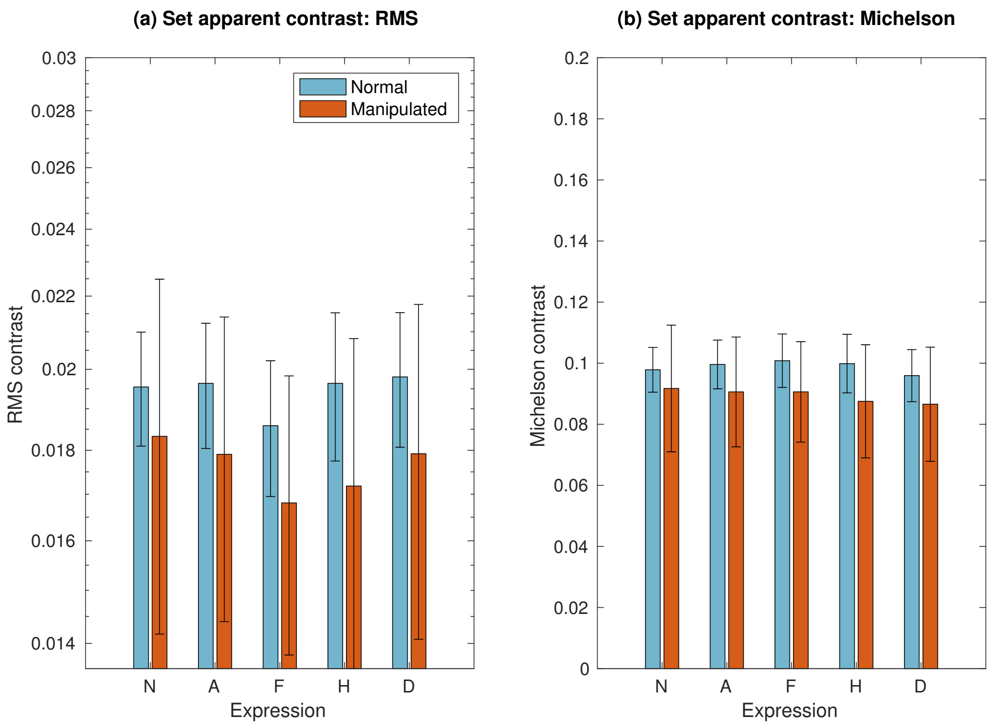

Fig. 6.2 Contrast settings for broadband face expressions, when apparent contrast is calculated using (a) RMS contrast or (b) Michelson contrast. Y-axes represent the degree of physical contrast (RMS or Michelson) required for a face to perceptually match that of a reference (set at 10% Michelson). All error bars represent associated standard deviations.

and manipulation (F(4, 72) = 2.58, p.04, ηp_2 .12; F(1,18) = 42.73, p<.001, ηp_2 .70, respectively) , but no significant interaction (F(4, 72) = .08, p.98, ηp_2 .004). Manipulated faces were perceived as more salient compared to their normally presented counterparts, and thus required less Michelson contrast to match the reference stimulus. Sidak-corrected paired comparisons (α= 0.0063) revealed no significant differences in apparent contrast between fearful faces and their face counterparts. Data are summarised in Table 6.3, and illustrated in Figure 6.3, (b).

6.3.3 Mid frequency filtered faces

Mean contrast settings for mid-range frequency faces are shown in Figure 6.4. For set contrast calculated using RMS contrast, a two-way Expression (neutral, anger, fearful, happy, disgust) X Manipulation (normal, manipulated) repeated measures ANOVA showed an effect of expression and manipulation (F(4, 72) = 11.26, p<.001, ηp_2 .38; F(1,18) = 6.34, p.02, ηp_2 .26, respectively), but no significant interaction (F(4,72) = 1.66, p.16, ηp_2 .08). Unlike the effect of manipulation observed for broadband and low-frequency faces, manipulated mid-range frequency faces required *more* contrast in order to appear the same as a reference stimulus, suggesting that they initially appeared less salient compared to their normally presented face counterparts. A significant effect of expression shows that when filtered to contain only mid-range frequency components, expressions differ in terms of their RMS contrast when perceptually matched. Sidak-corrected paired comparisons (α= 0.0063) revealed that mid-range frequency filtered fear expressions require significantly less RMS contrast compared to disgust faces in order to match a reference stimulus. This was also true when the two faces were manipulated. No other significant differences were found. Data are summarised in Table 6.4, and illustrated in Figure 6.4, (a).

Analyses were repeated for apparent contrast data that was calculated using Michelson contrast. A two-way Expression (neutral, angry, fearful, happy, disgust) by Manipulation (normal, manipulated) repeated measures ANOVA showed significant effects of expression and manipulation (F(4, 72) = 13.91, p<.001, ηp_2 .43; F(1,18) = 6.46, p.02, ηp_2 .26, respectively), and a significant interaction (F(4,72) = 2.55, p.04, ηp_2 .12). Again, manipulated face images require significantly more Michelson contrast in order to appear the same as a reference stimulus, compared to their normally presented face counterparts. Sidak-corrected paired comparisons (α= 0.0063) showed that fear expressions require significantly less Michelson contrast compared to neural and angry expressions in order to match a reference stimulus. These differences remained true between manipulated versions of fear and angry faces. No other significant differences were observed. Data are summarised in Table 6.5, and illustrated in Figure 6.4, (b).

6.3.4 High frequency filtered faces

Mean contrast settings for high frequency filtered faces are shown in Figure 6.5. A two-way Expression (neutral, anger, fearful, happy, disgust) by Manipulation (normal, manipulated) repeated measures ANOVA showed a significant effect of expression and manipulation (F(4,

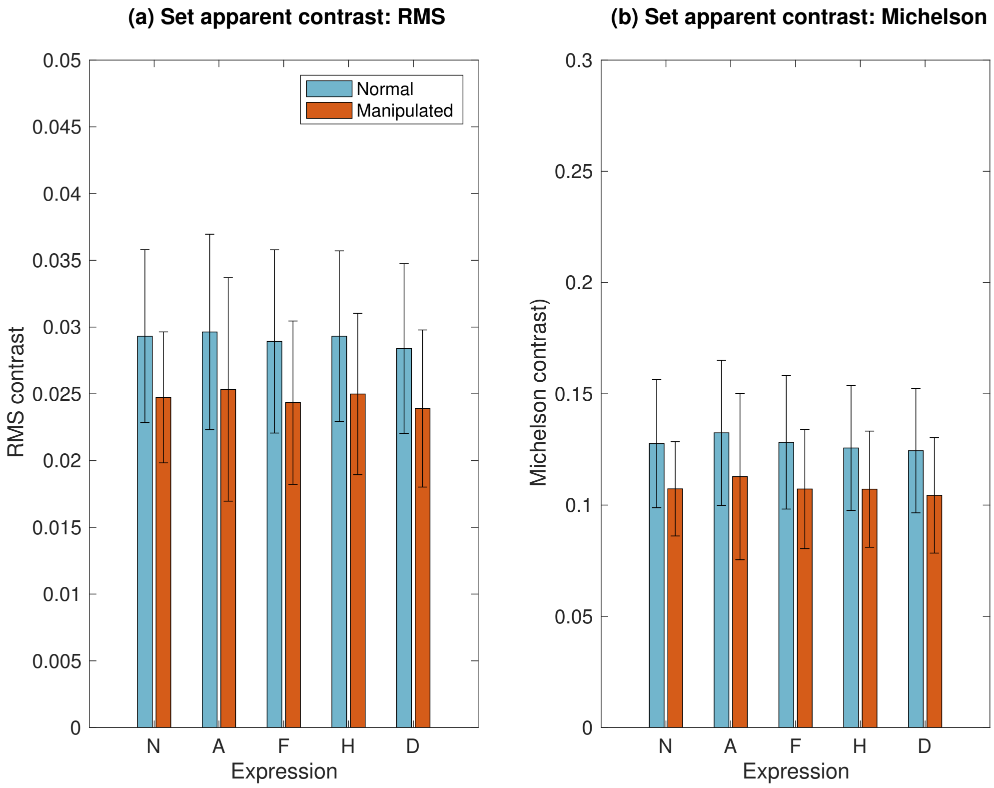

Fig. 6.3 Contrast settings for low frequency filtered face expressions, when apparent contrast is calculated using (a) RMS contrast or (b) Michelson contrast. Y-axes represent the degree of physical contrast (RMS or Michelson) required for a face to perceptually match that of a reference (set at 10% Michelson) (c) Demonstrates the physical composition of the same faces, when measured for their naturally occurring RMS content. All error bars represent associated standard deviations.

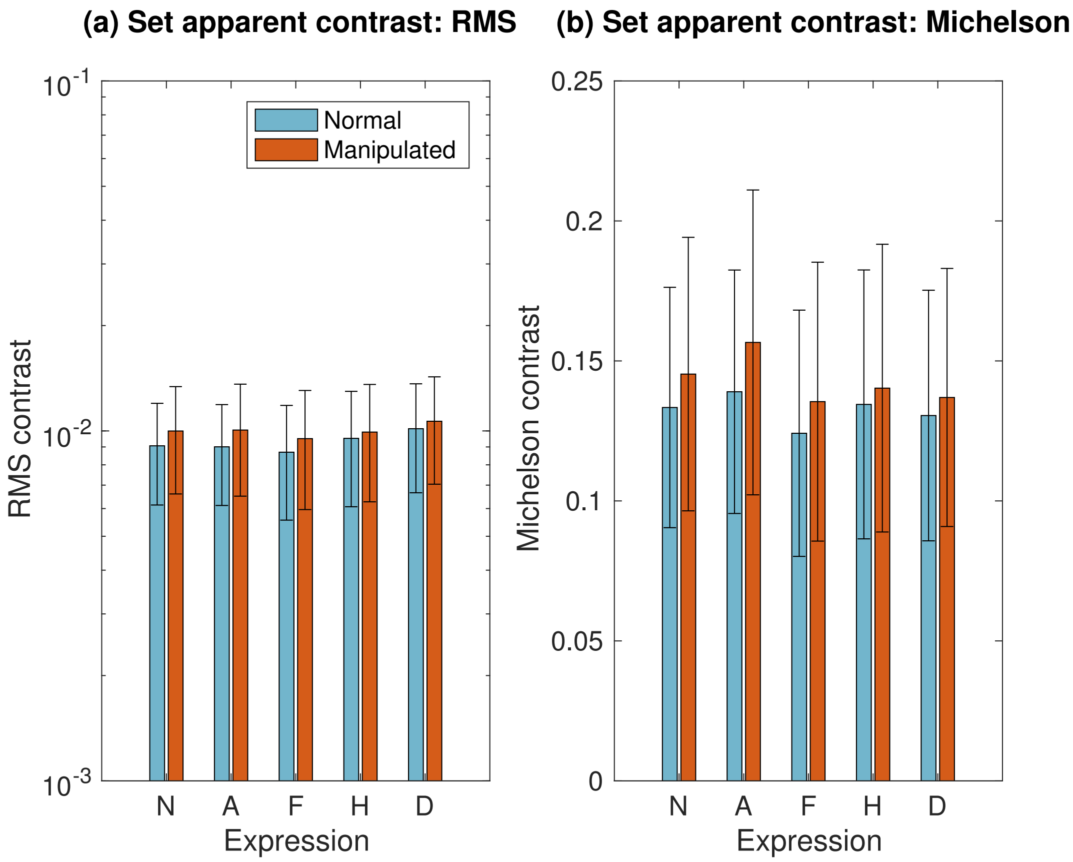

Fig. 6.4 Contrast settings for mid-range frequency filtered face expressions, when apparent contrast is calculated using (a) RMS contrast or (b) Michelson contrast. Y-axes represent the degree of physical contrast (RMS or Michelson) required for a face to perceptually match that of a reference (set at 10% Michelson) (c) Demonstrates the physical composition of the same faces, when measured for their naturally occurring RMS content. All error bars represent associated standard deviations.

72) = 26.04, p<.001, ηp_2 .59; F(1,18) = 7.33, p.01, ηp_2 .28, respectively), but no significant interaction (F(4, 72) = 2.11, p.08, ηp_2 .10). An effect of manipulation shows that, similar to the effect observed for mid-range frequency faces, high frequency manipulated faces also require more RMS contrast compared to their normally presented face counterparts in order to appear the same as a reference stimulus. A significant effect of expression shows that when filtered to contain only high frequency components, expressions differ in terms of their RMS contrast when perceptually matched. Sidak-corrected paired comparisons (α= 0.0063) showed that high frequency fearful expressions require significantly less RMS contrast compared to all but happy faces in order to match a reference stimulus. These differences between fear and their counterpart faces remained true when faces were manipulated. Data are summarised in Table 6.6, and illustrated in Figure 6.5 (a).

Analyses were repeated for apparent contrast data that was calculated using Michelson contrast. A two-way Expression (neutral, angry, fearful, happy, disgust) by Manipulation (normal, manipulated) repeated measures ANOVA showed a significant effect of expression and manipulation (F(4, 72) = 10.49, p<.001, ηp_2 .36; F(1,18) = 8.05, p.01, ηp_2 .30, respectively), but no significant interaction (F(4, 72) = 1.63, p.17, ηp_2.08. Similar to the effect of manipulation found for mid-range frequency faces, high frequency manipulated faces also require significantly more Michelson contrast compared to their normally presented face counterparts in order to match a reference stimulus. Sidak-corrected paired comparisons (α= 0.0063) showed that, unlike the effects observed for other frequency filtering conditions, high frequency fear expressions require significantly *more* Michelson contrast compared to happy and disgust counterparts in order to match a reference stimulus. These differences remained true for manipulated versions of faces. No other significant differences were observed. Data are summarised in Table 6.7 and illustrated in Figure 6.5, (b).

6.3.5 RMS contrast content of experimental stimuli

For the 16 KDEF face stimuli used in this study, 4 repeated measures ANOVA tests were conducted to confirm that expressions followed the same pattern of RMS contrast differences as those observed in a 140 KDEF face sample (Experiment 1). A repeated measures ANOVA revealed a significant effect of expression for broadband (F(4, 60) = 3.55, p.01, ηp_2.91), mid-range (F(4, 60) = 6.27, p<.001, ηp_2.29), and high frequency filtered faces (F(4, 60) = 8.63, p<.001, ηp_2 .36), but not for low frequency filtered faces (F(4, 60) = 1.58, p.19, ηp_2.09). Sidak-corrected pairwise comparisons (α= 0.0127, accounting for 4 comparisons) explored differences in RMS contrast between fearful expressions and their emotion counterparts, at each of the 3 levels of spatial frequency where an effect of expression was observed (broadband, mid-range, and high frequency filtered faces). The pattern of differences in RMS contrast between fearful and other expressions generally appeared similar in nature to the original 140 KDEF face analysis (Experiment 1), however, statistically significant differences included: broadband fear expressions were lower in RMS contrast compared to disgust expressions; high frequency fear expressions were lower in RMS contrast compared to angry and disgust expressions. No other significant differences were found. Data are summarised in Table 6.8, and illustrated in Figure 6.6.

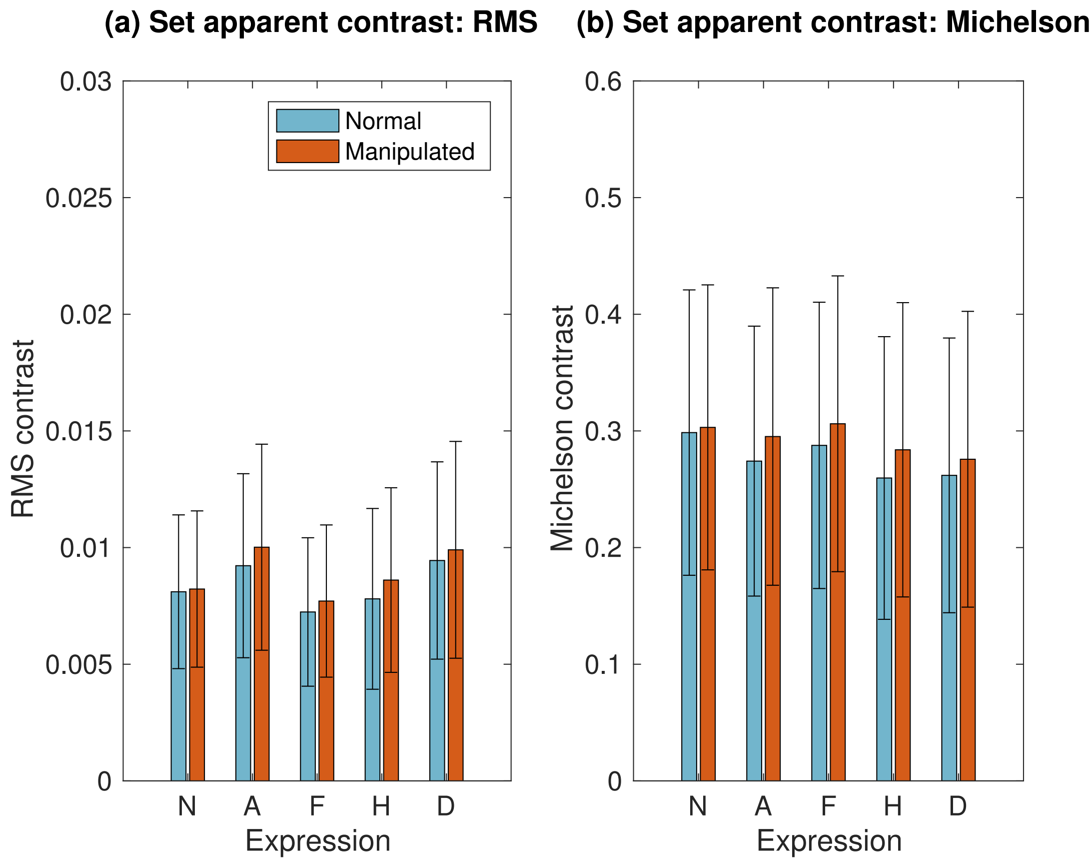

Fig. 6.5 Contrast settings for high frequency filtered face expressions, when apparent contrast is calculated using (a) RMS contrast or (b) Michelson contrast. Y-axes represent the degree of physical contrast (RMS or Michelson) required for a face to perceptually match that of a reference (set at 10% Michelson) (c) Demonstrates the physical composition of the same faces, when measured for their naturally occurring RMS content. All error bars represent associated standard deviations.

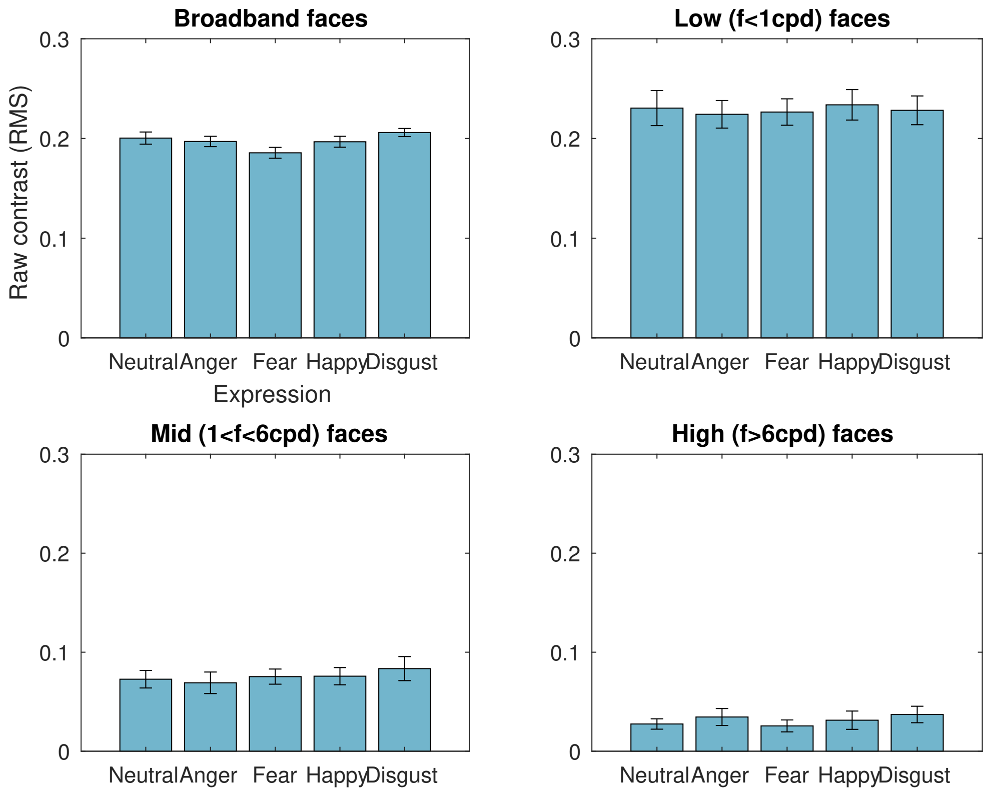

Fig. 6.6 The physical RMS contrast for broad, low, mid-range and high frequency filtered face images used as the experimental stimuli for Experiment 2. Faces were 16 KDEF (Lundqvist et al., 1998) face images. All error bars represent associated standard deviations.

6.4 Data tables

Expression comparisons	T	CI	Sig.
Fear			
Neutral	-3.44	-.001, -3e-4	.003
Anger	1.02	-7e-4, .002	.32
Happy	-3.34	-.001, -3e-4	.004
Disgust	-3.55	-.001, -4e-4	.002
Manipulated faces			
Fear			
Neutral	-3.51	-.002, -6e-4	.002
Anger	-3.49	-.001, -4e-4	.003
Happy	-1.27	-9e-4, 2e-4	.22
Disgust	-2.26	-.002, -1e-4	.03

Table 6.1 Apparent contrast for broadband faces: Sidak corrected comparisons (α= 0.0063, corrected according to 8 comparisons) exploring differences in the amount of RMS contrast required for faces to appear the same as a reference face image. Fear expressions are used as a reference comparison. Degrees of Freedom= 18 for all comparisons.

Expression comparisons	T	CI	Sig.
Fear			
Neutral	1.91	-2e-4, .006	.07
Anger	.71	-.002, .004	.48
Happy	.60	-.002, .004	.55
Disgust	2.73	.001, .008	.01
Manipulated faces			
Fear			
Neutral	-.52	-.005, .003	.60
Anger	-.003	-.003, .003	.99
Happy	2.02	-1e-4, .006	.05
Disgust	1.64	-.001, .009	.11

Table 6.2 Apparent contrast for broadband faces: Sidak corrected comparisons (α= 0.0063, corrected according to 8 comparisons) exploring differences in the amount of Michelson contrast required for faces to appear the same as a reference face image. Fear expressions are used as a reference comparison. Degrees of Freedom= 18 for all comparisons.

Expression comparisons	T	CI	Sig.
Fear			
Neutral	.17	-.006, .007	.86
Anger	-1.33	-.010, .002	.19
Happy	.91	-.003, .008	.37
Disgust	1.40	-.001, .009	.17
Manipulated faces			
Fear			
Neutral	-.02	-.006, .006	.98
Anger	-1.46	-.013, .002	.16
Happy	.02	-.006, .006	.98
Disgust	.79	-.004, .010	.43

Table 6.3 Apparent contrast for low frequency faces: Sidak corrected comparisons (α= 0.0063, corrected according to 8 comparisons) exploring differences in the amount of Michelson contrast required for low-frequency faces to appear the same as a reference face image. Fear expressions are used as a reference comparison. Degrees of Freedom= 18 for all comparisons.

Expression comparisons	T	CI	Sig.
Fear			
Neutral	-1.25	-7e-4, 1e-4	.22
Anger	-.71	-6e-4, 3e-4	.48
Happy	-2.79	-.001, -1e-4	.01
Disgust	-6.01	-.001, -8e-4	<.001
Manipulated faces			
Fear			
Neutral	-1.30	-8e-4, 1e-4	.20
Anger	-1.77	-.001, 1e-4	.09
Happy	-1.60	-8e-4, 1e-4	.12
Disgust	-5.65	-.001, -6e-4	<.001

Table 6.4 Apparent contrast for mid-range frequency faces: Sidak corrected comparisons (α= 0.0063, corrected according to 8 comparisons) exploring differences in the amount of RMS contrast required for mid-range frequency faces to appear the same as a reference face image. Fear expressions are used as a reference comparison. Degrees of Freedom= 18 for all comparisons.

Expression comparisons	T	CI	Sig.
Fear			
Neutral	-3.28	-.015, -.003	.004
Anger	-4.78	-.021, -.008	<.001
Happy	-2.79	-.018, -.002	.01
Disgust	-2.42	-.011, -8e-4	.02
Manipulated faces			
Fear			
Neutral	-3.08	-.016, -.003	.006
Anger	-5.21	-.029, -.012	<.001
Happy	-1.60	-.011, .001	.12
Disgust	-.57	-.006, .003	.57

Table 6.5 Apparent contrast for mid-range frequency faces: Sidak corrected comparisons (α= 0.0063, corrected according to 8 comparisons) exploring differences in the amount of Michelson contrast required for mid-range frequency faces to appear the same as a reference face image. Fear expressions are used as a reference comparison. Degrees of Freedom= 18 for all comparisons.

Expression comparisons	T	CI	Sig.
Fear			
Neutral	-4.64	-.001, -4e-4	<.001
Anger	-7.66	-.002, -.001	<.001
Happy	-2.14	-.001, -1e-4	.04
Disgust	-6.38	-.002, -.001	<.001
Manipulated faces			
Fear			
Neutral	-3.54	-8e-4, -2e-4	.002
Anger	-6.85	-.003, -.001	<.001
Happy	-3.32	-.001, -3e-4	.004
Disgust	-5.62	-.003, -.001	<.001

Table 6.6 Apparent contrast for high frequency faces: Sidak corrected comparisons (α= 0.0063, corrected according to 8 comparisons) exploring differences in the amount of RMS contrast required for mid-range frequency faces to appear the same as a reference face image. Fear expressions are used as a reference comparison. Degrees of Freedom= 18 for all comparisons.

Expression comparisons	T	CI	Sig.
Fear			
Neutral	-1.86	-.023, .001	.07
Anger	2.28	.001, .025	.03
Happy	4.13	.013, .042	.001
Disgust	3.33	.009, .041	.004
Manipulated faces			
Fear			
Neutral	5.73	-.008, .014	.57
Anger	1.42	-.005, .027	.17
Happy	3.28	.008, 0.36	.004
Disgust	4.15	.015, .045	.001

Table 6.7 Apparent contrast for high frequency faces: Sidak corrected comparisons (α= 0.0063, corrected according to 8 comparisons) exploring differences in the amount of Michelson contrast required for mid-range frequency faces to appear the same as a reference face image. Fear expressions are used as a reference comparison. Degrees of Freedom= 18 for all comparisons.

Expression comparisons: **BSF** faces	T	CI	Sig.
Fear			
Neutral	-2.42	-.027, -.001	.02
Anger	-1.56	-.026, .004	.13
Happy	-2.15	-.022, -1e-4	.04
Disgust	-3.46	-.032, -.007	.003
Expression comparisons: **MSF** faces	T	CI	Sig.
Fear			
Neutral	1.54	-9e-4, .006	.14
Anger	1.97	-4e-4, 0.13	.06
Happy	-.27	-.003, .003	.78
Disgust	-2.19	-.015, -2e-4	.04
Expression comparisons: **HSF** faces	T	CI	Sig.
Fear			
Neutral	-.90	-.006, .002	.38
Anger	-3.61	-.014, -.003	.003
Happy	-2.83	-.010, -.001	.013
Disgust	-6.15	-.015, -.007	<.001

Table 6.8 Differences in RMS contrast between 16 KDEF faces for broad (BSF), mid-range (MSF), and high frequency (HSF) faces. Fear expressions are used as a reference comparison. (α= 0.0127, corrected according to 4 comparisons). Degrees of Freedom= 15 for all comparisons.

6.5 Conclusion

Expression-related differences in perceived (apparent) contrast were measured by the physical contrast set by observers for different face expressions in order for them to appear the same as a reference face (whose contrast was set to 10% Michelson), and the extent to which this is influenced by the range of spatial frequency filtering. Intact broadband faces portraying anger, fearful, happy, disgusted and neutral expressions were filtered to create low $f < 1_{cpd}$, mid-range $1 < f < 6_{cpd}$, and high $f > 6_{cpd}$ frequency versions. The data shows that apparent contrast does differ between facial expressions, that the degree of physical contrast necessary to perceptually match a facial expression is influenced by the spatial range used to filter the image, and that the degree of difference between perceived and physical contrast depends on the physical contrast metric used to normalise face images. These effects can be summarised as follows:

6.5.1 Unfiltered, broadband stimuli

Fear expressions are perceived as having higher contrast than neutral, happy and disgusted expressions. This is true when faces' apparent contrast is calculated using RMS contrast; but no differences were evident when Michelson contrast was used. In other words, unfiltered and broadband fearful faces require significantly less RMS contrast in order to appear the same as a given reference. Conversely, when the same faces are matched for Michelson contrast, no significant differences between expressions' apparent contrast exist. It could then be argued that observers matched face images according to Michelson contrast, rather than RMS.

6.5.2 Low frequency filtered stimuli

Overall, findings showed that the perceived contrast of low frequency filtered expressions is affected very little by contrast metric. Apparent contrast does not differ between expressions when RMS contrast is used as the contrast measure. An effect for Michelson contrast is small, but does, however, show that low frequency fearful faces are likely to require less Michelson contrast compared to disgust faces in order to be perceptually matched.

6.5.3 Mid frequency filtered stimuli

Apparent contrast did not differ between mid-range frequency expressions when the contrast metric used was RMS contrast. However, the same fearful facial expressions do appear as having higher contrast compared to neutral and angry expressions when the contrast metric used is Michelson contrast. These findings imply that matching mid-range faces for RMS contrast may limit differences in percevied contrast more than when normalising the same faces for Michelson contrast, but only marginally.

6.5.4 High frequency filtered stimuli

High frequency filtered fear expressions are perceived as more salient, and therefore require less RMS contrast in order to be perceptually matched, compared to neutral, angry and disgust faces when the contrast metric used is RMS contrast. Interestingly, when Michelson contrast is the metric of interest, fear expressions require *more* Michelson contrast compared to happy and disgust faces in order to be perceptually matched. These effects are consistent with those of O'Hare and Hibbard (2011), where according to contrast constancy, high frequency filtered versions of stimuli are lower in perceived contrast. However, findings also highlight the importance of understanding the implications of the contrast metric that is used to normalise facial stimuli.

To summarise, for normal face images that are not spatially filtered, findings suggest that matching face expressions for RMS contrast does not guarantee that they are normalised at the perceptual level. Fear expressions in particular require less physical contrast to be appear perceptually matched to other expressions. Some of these effects were also found for manipulated face stimuli, showing that when substantial configuration content is removed and low-level content preserved, fearful faces will still appear more salient compared to other expressions when they are matched for their RMS contrast. This finding is consistent with previous studies (Gray et al., 2013), suggesting that low level stimulus properties may determine a specific salience associated with fearful facial expressions (Gray et al., 2013; Yang and Blake, 2012). Because no differences in Michelson contrast were observed when the same broadband faces were perceptually matched, it may be that Michelson contrast assures a great degree of consistent between faces physical and perceived contrast compared to RMS contrast. Importantly, they show that when composed of a natural variation of spatial frequencies, and matched for physical contrast, fearful faces are salient, in that they initially appear higher in contrast compared to other face expressions.

Findings suggest that the perceived salience of low frequency filtered facial expressions is relatively uninfluenced by contrast metric. However, when filtered to contain mid-range and high frequency content, effects of salience associated with fearful expressions re-emerge, suggesting that they are driven by the higher-frequency $f > 1_{cpd}$ components of the images. This suggests that the use of filtering techniques may either diminish (in the case of low frequencies) or enhance (in the case of high frequencies) the salience of fearful faces. Together, these results show that the salience of fearful faces associated with the threat bias is evident in the apparent contrast of fearful faces that are filtered to contain high spatial frequency information, but that the strength of this advantage depends very much on the contrast metric.

6.6 Discussion

Experiment 1 asked whether images of face expressions naturally differ from each other at the physical level. To follow on from this, the present study (Experiment 2) measured differences in perceived salience even when the same face stimuli are matched for physical contrast.

Recent psychophysical studies suggest that fearful faces have a significant status within the visual system, and that this effect is determined by their low spatial frequency content (Bannerman et al., 2012; Vuilleumier et al., 2003). Such coarse, low-level information is thought to undergo rapid subcortical visual processing (Bannerman et al., 2012; Vuilleumier et al., 2003). This approach to understanding the threat bias for fearful faces is theoretically compatible with notions of a low-road for processing threat-relevant information, including fearful faces (LeDoux, 2012; Öhman and Mineka, 2001b; Vuilleumier et al., 2003), which depends primarily on low-lever image properties (Hedger et al., 2016). Experiment 1 and 2 therefore address two related questions: to what extent do face expressions differ in contrast and spatial frequency content before they are physically matched for contrast (Experiment 1), and whether normalisation processes assure that no differences remain between expressions' perceived, or subjective salience (Experiment 2). The overarching conclusions are detailed below.

6.6.1 Facial expressions naturally differ at the physical level

A preliminary question was whether or not facial expressions differ in terms of their low-level image properties. Images of faces differ from those of other natural scenes in their Fourier amplitude spectra (Redies et al., 2007). However, as far we are aware, it remains unknown whether there are expression-related differences in image properties that occur within facial stimuli. Acknowledging natural differences such as these is a necessary prerequisite for understanding effects of contrast normalisation on the physical composition of faces, and also the extent to which differences remain between these physical attributes of faces and their perceived image salience. We measured the physical contrast in 140 KDEF face expressions, and versions of these faces filtered to contain different spatial frequency ranges. The RMS contrast and Fourier amplitude spectra were calculated for these images. Unfiltered (broadband) images were found to differ across facial expressions in their RMS content and amplitude spectra. Differences in physical contrast between expressions occur across spatial filtering conditions, but fearful facial expressions in particular are naturally lower in physical RMS contrast compared to other emotion counterparts. This effect is more enhanced as spatial frequency increases. These findings contribute to the broader literature that demonstrates characteristic physical properties associated with face stimuli (Redies et al., 2007), adding that physical differences also exist between face stimuli, and vary as a function of facial expression and spatial frequency range.

6.6.2 Differences between expressions' perceived salience after contrast normalisation

The currently accepted view is that normalising natural images for RMS contrast ensures continuity between their physical and perceived contrast, but it remains unclear whether the same is true for faces as well as other images. Given the unique characteristic amplitude spectrum associated with face images (Redies et al., 2007), and the possibility of natural expression-related differences between image properties, we might expect that facial expressions could be differently affected by contrast normalisation than, for example, images of natural scenes. We measured the physical contrast (RMS and Michelson) necessary for facial expressions to appear perceptually matched for contrast. Previous findings show that perceived contrast of stimuli, when matched for RMS contrast (O'Hare and Hibbard, 2011), varies depending on spatial frequency content. A consistent finding from Experiment 2 was that fearful expressions tend to require less physical contrast in order to provide a perceptual match with other expressions, and that this effect is most pronounced as the spatial frequency content of faces is increased. These findings show that contrast normalisation using RMS and Michelson contrast does not ensure subjective similarity for images of facial expressions. Rather, fearful faces will appear to have a higher contrast than other expressions when matched for RMS contrast.

6.6.3 Relevance of the present data to broader literature

The perception of facial displays of emotion is not exclusively studied by visual psychophysicists, but is a broadly relevant topic within cognitive psychology and others areas of social sciences. Our data highlight the importance of understanding the effects of contrast normalisation on the salience of facial stimuli. In particular, the present study raises the question of whether contrast normalisation is a necessary component for a given experimental design, and if so, which contrast metric is most appropriate given the nature of the experimental stimuli. The technique is widely used, but for facial stimuli specifically can alter the contrast content in a way that could lead to misconstrued perceptual biases for given stimuli. In terms of the threat bias literature, the present data suggests perceptual biases for fear expressions are likely driven to some degree by effects of contrast normalisation. Experiment 1 showed that fearful expressions are naturally lower in physical contrast compared to other expressions; an effect that increased linearly with their spatial content. Therefore, during contrast normalisation, raw broadband fear faces likely benefit from an artificial boost in their contrast content. Data from Experiment 2 shows that fearful faces require less contrast to perceptually match other expressions, but that this may only be the case when faces are normalised for RMS contrast. This suggests that RMS contrast normalisation of faces may boost the perceived salience of fearful expressions. This alone could account for perceptual biases to these expressions. Bolstering effects become more pronounced as expressions' spatial content increases, suggesting that highpass filtered expressions (namely fear expressions) are most susceptible to these effects after they have been normalised for contrast. This may account for findings from Stein et al. (2014) of a link between high spatial content and rapid detection of fear faces. However, findings do suggest that for normal, unfiltered facial stimuli, Michelson contrast may ensure that expressions matched for Michelson contrast do not differ at the perceived level. Importantly, however, this effect does not remain true as spatial frequency content of images increases, where for high frequency fear expressions, normalising faces for Michelson contrast may in fact reduce the salience of fear expressions compared to other faces.

Very minimal expression-related differences in apparent contrast were found when stimuli were low pass filtered. At first glance this finding implies that low frequency filtered expressions are not susceptible to the bolstering effects of contrast normalisation proposed above, suggesting studies that evidence perceptual biases for low frequency fear may reflect a genuine bias, rather than a bias in response to an increase in image's apparent contrast. However, it is important to clarify here that 'low' frequency was defined as $f<1_{cpd}$. Studies reporting an effect of low frequency on fear expression salience use ranges that vary closely above this bandpass cutoff: $f<2_{cpd}$ (Holmes et al., 2005; Vlamings et al., 2009). Such cutoffs are consistent with our mid-range frequency condition, defined as $1f<6_{cpd}$ a range associated with normalisation effects on faces' apparent contrast. Equally, the present findings cannot account for studies reporting the same low-frequency finding whose bandpass cutoffs lay outside of the affected band: $f<.08_{cpd}$ (Bannerman et al., 2012; Vuilleumier et al., 2003). Nor can they account for why threat bias effects in other studies were not found in higher frequency conditions, given that bandpass cutoffs used are more likely to match those that are associated with normalisation, and therefore apparent contrast boosting. This suggests that effects of contrast normalisation are not wholly responsible for findings of threat biases for fearful faces, and highlight the importance of emphasising the exact band of information an effect is associated with, rather than a statement of whether this reflects frequencies categorised as "high" or "low". Chapter 4 provides a brief discussion of such deviations in how spatial frequency content of filtered images is defined.

In terms of which contrast metric is most appropriate for normalising facial stimuli, data from the present study shows that this is determined by several factors: the facial expressions used in the study, as some do not differ in apparent contrast after normalisation, and also the range of spatial frequency used for image filtering. Low frequency faces are better matched for perceived salience when normalised for RMS (compared to Michelson) contrast, whereas studies using unfiltered (broad) face stimuli might prefer normalisation for Michelson contrast given that differences were more limited than when these faces are normalised for RMS content. Especially important is the influence of contrast normalisation on the salience of fear expressions when these are high frequency filtered. High frequency fear faces will require less physical contrast to perceptually match other expressions if they are normalised for RMS contrast. If normalised for Michelson contrast, however, they require more contrast. As differences between images' apparent and physical contrast vary according to multiple factors, it may be reasonable to question whether normalisation is necessary at all, especially if the focus is on the role of the properties of naturally occurring stimuli. In other cases, such as those interested in the role of frequency information, it may be that normalising facial stimuli in terms of their perceived contrast is the most appropriate contrast metric. This technique would at least ensure consistency across faces perceived contrast, in a way that is not assured through normalising images' RMS or Michelson contrast.

Our data show that contrast normalisation of face stimuli does not ensure consistency between the apparent, or perceived contrast, of facial expressions. For studies outside of visual psychophysics that are not concerned with spatial filtering of images, normalising face expressions for RMS contrast is likely to artificially boost the physical contrast of expressions over and above the amount necessary for them to *appear* the same as a desired reference. Unfiltered, broadband face images naturally differ at the physical level, whereby fearful expressions are notably lower in physical contrast. Although after normalisation face expressions will be physically identical, fear expressions will appear more salient, or higher in contrast, compared to other faces. A reasonable inference here is that such faces will likely be associated with observer responses relating to visual sensitivity and salience. An example is data from Hedger et al. (2016), whereby face images of neutral, angry, happy and fear expressions were normalised for RMS contrast, and presented to observers under conditions of intraocular suppression. Observers task was to indicate the point at which a face became visible while masked from awareness by a noise stimulus, used as an index of faces relative 'visibility'. Hedger et al. (2016) found that these RMS matched fear expressions were more visible compared to counterpart faces. This effect is consistent with our own data for broadband face images, whereby broadband fear faces matched for RMS contrast receive a boost in perceived contrast after their physical RMS contrast is increased through the normalisation procedure. In this case, it is reasonable to argue that effects observed by Hedger et al. (2016) reflect inadvertent consequences incurred form the standard normalisation of experimental stimuli, rather than a genuine perceptual bias for natural low-level properties inherent to this face.

6.6.4 Addressing inconsistencies in the wider literature

The present study provides an explanation of the high frequency effects observed by (Stein et al., 2014). They argue that their data are consistent with those of other studies (Smith and Schyns, 2009; Smith et al., 2005), in that the perceptual biases for fear expressions operate via a visual cortex-amygdala pathway that relies on high spatial scales of information. We suggest that the findings from (Stein et al., 2014) might be explained by effects of contrast normalisation on apparent contrast, and findings of a high frequency effect appear to be exclusive to the Bubbles technique in particular. The implications of both CFS and Bubbles techniques on findings are still unclear (see Murray and Gold (2004) and Gosselin and Schyns (2004) for Bubbles; see Stein et al. (2011) for CFS). Data regarding the crucial spatial range necessary to drive biases for fear expressions are perhaps not so inconsistent. A recent meta-analysis by (Hedger et al., 2016) argues that evidence to support the threat bias is relatively weak, and is influenced by several experimental conditions making it more or less observable (see Hedger et al. (2016) for more detail). This highlights the importance of conceptualising the threat bias as a composition of several different subcategories of threat-sensitive perceptual behaviours, such as those related to threatening objects (Koster et al., 2004, 2007; Stormark et al., 1995), as opposed to an all-purpose visual function (LeDoux, 2012; Öhman and Mineka, 2001a). It seems to be that biases for fearful expressions are one of such subcategories, where studies concerned with their mechanisms are, for the most part, consistent in their findings (Bannerman et al., 2012; Bayle et al., 2011; Carlson and Reinke, 2008; Gray et al., 2013; Holmes et al., 2005; Phelps et al., 2006; Pourtois et al., 2006; Schupp et al., 2004; Vlamings et al., 2009; Whalen et al., 2001; Yang et al., 2007).

Prior chapters highlight some methodological inconsistencies, particularly those related to face perception. An example is the unclear efficacy of contrast normalisation; whether it is necessary, and if so, which contrast metric is most appropriate. Michelson and RMS contrast are commonly used, but other metrics are available and perhaps more appropriate

(See Meese et al. (2017) for review). The most appropriate metric for a given task is an under-addressed topic within the literature. A similar issue relates to variable ways to categorise spatial frequency. Much like the variety of contrast metrics, measures and definitions of spatial frequency also vary across studies. These can be expressed in terms of cycles/degree, cycles/face, cycles/image, or cycles/face-width, as Chapter 4 (4.2.4). Cycles/degree is perhaps most the most common measure, and most relevant if one is interested in the low-level sensitivity of the visual system. Under this metric, images will vary as the distance between observer and stimulus, whereas cycles/face width is argued to be a more ideal measure because it remains unchanged as this distance varies. Fiorentini et al. (1983) argue that the nature of face perception tasks does not require distance, and spatial frequency should be defined consistently as a single dimension, where cycles/face width an ideal frequency metric (Näsänen, 1999). However, it could be argued that certain visual functions, relating to face perception, such as those relating to the underlying sensitivity of the visual system, are not independent of distance. The ability to detect a fearful face, for example, needs to operate beyond that associated with social interaction. It is also important to recognise arbitrary subdivision of the spatial frequency spectrum into 'low' and 'high' bands of information. This is shown above, where 'low' frequency can refer to information below $f{<}1_{cpd}$, or below $f{<}2_{cpd}$. These variations in definitions impose an overlap between what should be two distinct bands of information, and can make it difficult to discern precisely where on the spatial frequency spectrum an effect might be expected to occur. This is particularly important for defining 'high' frequencies, and how often they abide by upper limit of frequency information: in photographs of faces, for example, contrast dissipates by around 27 cycles/face-width, but is also relatively low below this limit (Hayes et al., 1986; Näsänen, 1999). Just as with contrast measures, these inconsistencies in methodology are addressed very little in the literature.

6.7 Summary of Chapter 6

The present study showed that raw images of facial expressions do naturally differ at the physical level. Expression-related differences in physical contrast are masked by the contrast normalisation technique, such that faces naturally lower in contrast are artificially enhanced. Contrast normalisation thus matches facial expressions at the physical level, but inadvertently bolsters differences in their perceived contrast. Fearful facial expressions are more susceptible to these effects than other facial expressions, and the likelihood that these effects occur increases with spatial frequency. Findings raise questions regarding the use of contrast normalisation, and the extent that findings of perceptual biases occur due to experimentally driven effects of apparent contrast.

Chapter 7

Are fear expressions salient because of their stimulus efficacy?

Belle and Sebastian. (2006). The Life Pursuit. Artwork copyright to Matador Records (2006).

Much of the threat bias literature agrees that in the competition for processing, fearful facial expressions are selected for their relevance for survival via specialised subcortical visual mechanisms (De Gelder et al., 1999; LeDoux and Phelps, 1993; Morris et al., 2001; Ohman and Soares, 1994). The subcortical nature of these mechanisms is widely believed to operate independently of conscious awareness. Chapter 3 introduces the proposed low-frequency-sensitive subcortical pathway as one such mechanism, but an alternative approach proposed by Hedger et al. (2015) argues that perceptual biases for fear expressions are better predicted by their stimulus efficacy. This recently proposed account of the threat bias incorporates key aspects of the sensory bias hypothesis, but is yet to be tested at the behavioural level. The present chapter introduces this approach, and in doing so presents the motivation for Experiment 3.

7.1 Stimulus efficacy

A visual stimulus might be preferentially processed for two reasons: because it is semantically and meaningfully relevant, or because its configuration is somehow congruent with low-level mechanisms in early vision that allow for it to be rapidly and effectively processed. Low-frequency-sensitive subcortical theories propose that the human visual system has evolved specific visual neural mechanisms that enable rapid identification of fearful expressions because they possess a unique status in the threat avoidance system (Bayle et al., 2011; Milner and Goodale, 2006; Vuilleumier et al., 2003). An alternative low-level, but directionally different view argues that it is the stimulus properties associated with fear expressions that allow them effective processing by already sensitive visual processing pathways (Gray et al., 2013; Hedger et al., 2015). In other words, fearful expressions are salient to the visual system due to their stimulus efficacy. Both Hedger et al. (2015) and Gray et al. (2013) refer to the sensory bias hypothesis which states that perceptual biases for certain facial expressions rely not on selective attentional processes, but on the detection of physical image attributes contained within faces that are available to fundamental aspects of early vision (Horstmann and Bauland, 2006). Horstmann and Bauland (2006) use the following example of a visual search task: if a red disk stimulus presented within an array of green disks is consistently detected at fixed latency of 1,000 ms, despite an increasing number of green disks, it can be assumed that the red disk-bias occurs due to the efficiency of its stimulus properties that are accessed by the cognitive system prior to attention allocation. If latencies for detection times increase linearly with set size, or the number of green disks, this implies the role of attentional processes that transition between stimuli until the target stimulus is detected. Here, Horstmann and Bauland (2006) highlight the distinction between stimulus detection that arises from non-efficient but attentional processes, and that which occurs pre-attentively, and is therefore efficient by definition. Hedger et al. (2015) incorporate these notions of stimulus efficacy and efficient information processing in their approach to the biases for fearful face expressions. Their alternative approach suggests that the physical configuration characterising a fear expression has evolved to maximise its stimulus signal, such that this information may be readily and preattentively accessed by fundamental visual functions. This position contrasts with other theories within the literature that suggest the existence of visual mechanisms specifically for fearful face detection (De Gelder et al., 1999; Ohman and Soares, 1994). Instead, physical attributes are able to undergo 'general purpose' processes that underpin basic and rapid stimulus detection, namely, according to Hedger et al. (2015), via the contrast sensitivity function.

7.2 Does the fear expression bias operate via the contrast sensitivity function?

According to the efficacy-explanation by Hedger et al. (2015), image attributes characteristic of fearful faces benefit from preattentive processing via already early and rapid visual channels. Specifically, they implicate the role of the contrast sensitivity function, noting that the human visual system is most sensitive to contrast with a certain spatial frequency range of 1-5 cpd as outlined in Section 2.2 (Campbell and Robson, 1968; Sunness et al., 1997). The extent that the Fourier amplitude spectra of fear faces is well matched to the contrast sensitivity function, relative to other images, is thought to determine their stimulus efficacy, thus accounting for the associated perceptual biases. This approach is consistent with previous literature in terms of specialised fear processing that takes place independent of awareness. Both Hedger et al. (2015) and traditional notions argue the necessity for a threat bias to elicit appropriate behavioural responses that can be generated with maximum rapidity (De Gelder et al., 1999; Lee et al., 2014; Öhman and Mineka, 2001b; Ohman and Soares, 1994). Hedger et al. (2015) calculated the amplitude spectra for images of fearful and neutral face expressions. Measures of the expressions' effective contrast were then obtained by multiplying their amplitude spectra by a standard model of the contrast sensitivity function, based on the Modelfest dataset (Watson and Ahumada, 2005). Their findings showed that fearful compared to neutral faces contained more contrast energy at the spatial range associated with the peak of the contrast sensitivity function, suggesting that fear expressions are prioritised for processing purely due to the efficacy of their stimulus properties given the contrast sensitivity function. Consistency of the sensory advantages associated with fear expressions was found across 5 commonly employed face databases. As part of the same investigation the degree of effective contrast associated with anger, fear, happy, and neutral expressions was shown to be a predictor of their visibility under interocular masking conditions (Hedger et al., 2015). Face expressions of fear were most visible compared to other face stimuli, an effect that was preserved when their configural content was disrupted using stimulus manipulation. These findings were attributed to the effective contrast, or stimulus efficacy, unique to fear expressions. It is important to include here that the data from (Hedger et al., 2015) does not necessarily negate the currently accepted accounts implicating the importance of low frequency components in detecting fear expressions. This is because the critical spatial frequency range associated with effective contrast does not necessarily deviate substantially from that associated with the contrast sensitivity function. Indeed, previous psychophysical studies concerned with face identification question the extent to which perceptual biases for filtered faces are facilitated due to the relative closeness of their spatial range to that associated with the peak of the contrast sensitivity function (Costen et al., 1996). During spatial filtering, certain bands of frequency information are removed while the spatial range of interest is preserved. However, as outlined in Chapter 2, in images of faces, as with other scenes, contrast amplitude decreases with spatial frequency, such that lower frequencies contain more contrast content. Costen et al. (1996) suggest that faces filtered using low frequency bandpass filters may preserve additional contrast that accompany this frequency range. Their suggestion was proposed as an explanation for some of the empirical inconsistencies for crucial spatial range required for face identification (for more information see, (Costen et al., 1996)), but extrapolation of this notion raises the question of whether the same is also true for studies that low frequency filter images of fear expressions. It may be that when fear expressions are filtered to contain low frequency content they become similarly matched to optimal contrast sensitivity, therefore boosting their image salience. Hedger et al. (2015) acknowledge this possible effect, highlighting the likelihood that lower spatial frequencies would naturally have been included in the measurement of faces' effective contrast. The authors do, however, emphasise that their findings highlight the importance of contrast at these scales, rather than specifically the evaluation of information at this location.

7.3 The premise for Experiment 3

The argument that fear expressions are particularly well tuned for early sensory processes is consistent with the wider body of threat bias literature, and is upheld by data from an image analysis and behavioural data from a visual suppression paradigm, both performed by Hedger et al. (2015). However, there remain several elements of this approach that are not addressed by the authors. The first relates to the behavioural evidence in support of the efficacy-account. The role of preconscious processing of threat-information is behaviourally evidenced using an intraoculor visual suppression paradigm, but does not directly measure expression-related effects on contrast sensitivity processing. If it is the case that the fearful faces are well matched to the human contrast sensitivity function, we should expect to observe an increase in contrast sensitivity, reflected by decreased contrast thresholds, for fearful faces compared to other facial expressions. The second issue is that prior to the transformation of face images,(Hedger et al., 2015) normalised these faces for their luminance and RMS contrast, such that they were identical at the physical level. As demonstrated by data obtained from Experiment 1 and 2, the process of attributing the aggregate physical contrast to all facial stimuli matches face stimuli at the physical level, but it does not guarantee the same consistency across their perceived image salience. It is therefore possible that normalising faces for physical contrast prior to analysis may inadvertently include effects of their variable apparent contrast in subsequent image transformations. For example, if contrast normalisation artificially boosts the apparent contrast of fearful expressions above the necessary amount to appear identical to other expressions, as findings from Experiment 2 suggest, this might drive effects of stimulus salience such as those observed by Hedger et al. (2015).

Experiment 3 addresses these questions. Here, we conduct a replication of the image analyses conducted by Hedger et al. (2015), in order to include the same set of facial stimuli that are physically matched for contrast, but with the additional of faces that remain physically unmatched, such that they contain natural differences in both physical and apparent contrast. Furthermore, we conduct a traditional contrast sensitivity task in order to psychophysically measure predictions from image analyses by Hedger et al. (2015). Here, we employ facial expressions as opposed to sinusoidal grating stimuli to measure expression-related differences in contrast sensitivity. An important feature of this later study is that it directly addresses the association between face expression and contrast sensitivity at the behavioural level, using an experimental paradigm that is traditionally used when measuring this phenomenon; as opposed to unrelated tasks such as those associated with intraocular suppression.

Chapter 8

Experiment 3: The effect of facial expression on contrast sensitivity: a behavioural investigation and extension of Hedger, Adams & Garner (2015).

Weezer. (1994). The Blue Album. Artwork copyright to DGC Records (1994).

8.1 Introduction

The present experiment conducts a behavioural investigation of image analyses performed by Hedger et al. (2015), in the form of a traditional contrast sensitivity task, employing facial expressions as opposed to sinusoidal grating stimuli to measure expression-related differences in contrast sensitivity. In addition to this, image analyses of faces' effective contrast is repeated for the same face samples used by Hedger et al. (2015), and this is extended to include versions of face images that have *not* been normalised for physical RMS contrast.

8.2 Methods

8.2.1 Participants

Eighteen (15 female, 3 male) participants took part in the study. All participants were informed of the nature of the study and provided written informed consent prior to the study beginning. The University of Essex Ethics Committee approved the experimental procedures. All participated in the experiment as part of a credited research module assessment, or in exchange for monetary reward. All participants had normal or corrected to normal vision.

8.2.2 Stimuli and apparatus

Stimuli were grayscale images of 16 individuals, 8 male and 8 females, taken from the Karolinska Directed Emotional Faces set (Lundqvist et al., 1998). Face images included internal features only, and included 4 emotional expressions of neutral, fear, anger and happiness. Though Hedger et al. (2015) refer to only fearful and neutral expressions, we included an additional two expressions so as to include positively and negatively-valenced comparisons. All individual faces were presented in their normal, upright form, and in a phase scrambled format. Phase scrambled versions of the face images were used as a control measure, providing versions of faces whose configural content was disrupted but low level statistical properties preserved. An example is shown in Figure 8.2. Phase scrambling was performed using MATLAB fast Fourier transform functions. Contrast thresholds were determined using an adaptive staircase technique (see under Procedure, below). Stimuli were presented using a VIEWPIXX 3D monitor (52cm X 29cm), viewed from a distance of 65 cm. The stimulus size of faces was 5.5 degrees. The screen resolution was 1920x1080 pixels, with a refresh rate of 120Hz and an average luminance of 50 cdm-2. Each pixel subtended 1.43 arc min. Stimuli were presented at 10-bit resolution. Participants' responses were recorded using the RESPONSEPixx response box. Stimuli were generated and presented using MATLAB and the Psychophysics Tool box extensions (Brainard and Vision, 1997).

8.2.3 Procedure

Participants were tested individually in a quiet room and informed prior to the experiment that the study was concerned with face perception. As a 2AFC location task, the participants' objective was to indicate, using 1 of 2 buttons on a RESPONSEPixx response box, whether the target face image appeared to the left or right of centre. The beginning of each trial commenced with the face stimulus on the left or right side of the screen. Participant responses determined the onset of the next trial. The proportion of times that the participant correctly indicated the location of the stimulus was recorded for all face stimuli. The adaptive staircase method was used to establish the Michelson contrast required for correct detection (75% of the time) for each expression stimulus. Here, the starting contrast level for each expressions' staircase began at 0.01 Michelson contrast. According to an up-down rule (García-Pérez et al., 2011), Michelson contrast was increased by one initial step of 0.005 proceeding 1 incorrect observer response, thus boosting stimulus visibility. Conversely, 3 correct observer

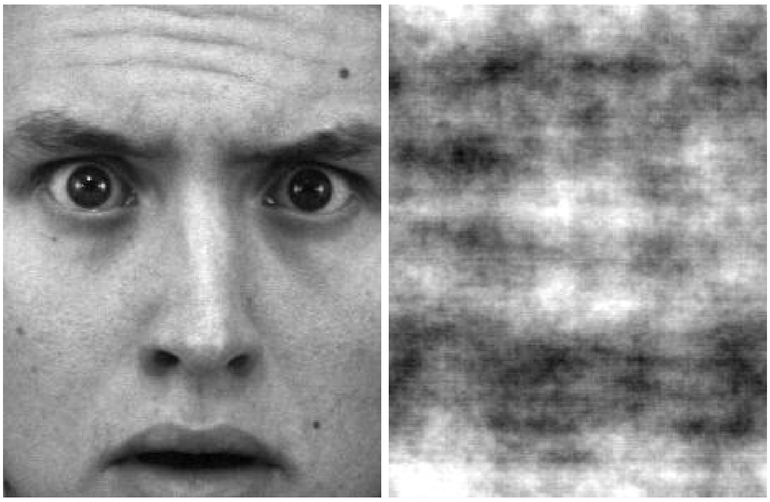

Fig. 8.2 A fearful face (Lundqvist et al., 1998), shown in its normal format (left), and the same image when it is subjected to Fourier phase scrambling (right). Both images are identical in terms of their amplitude spectrum, but the randomised phase spectrum assigned to the manipulated image renders it visibly unrecognisable.

responses triggered a decrease in Michelson contrast, initially by 0.005. The overall staircase length was 70 trials, where the initial step size (0.005 Michelson) halved after 17, 35 and 52 trials. 4 experimental blocks were completed, and the 280 trials for each combination of expression and phase scrambling were combined to create a single psychometric function.

8.3 Results

The proportions of participants' correct responses for each expression, at each contrast level, were used to create a psychometric function. A cumulative Gaussian function was fit to this data using the Palemedes toolbox (Prins et al., 2016) and used to determine a contrast detection threshold for each expression in its normal and manipulated (scrambled) formats. This 75% contrast detection threshold was defined as the contrast required for the participant to correctly identify the location of the face stimulus on 75% of trials. These results are plotted in Figure 8.3. A 4 Emotion (neutral, anger, fear, happy) x 2 Manipulation (normal, scrambled) within subjects ANOVA revealed no significant effects of expression (F(3, 51)= .26, p=.85, ηp_2=.01), or manipulation (F(1, 17)= .13, p= .72, ηp_2 =.01), and no significant expression x manipulation interaction (F(3, 51)= 1.20, p=.32, ηp_2 =.06). Analyses were repeated for contrast thresholds that were calculated using the RMS contrast of face stimuli. A 4 Emotion (neutral, anger, fear, happy) x 2 Manipulation (normal, scrambled) within subjects ANOVA revealed no significant effect of expression (F(3, 51)= .42, p=.73, ηp_2=.02), or manipulation (F(1, 17)= .07, p= .79 ηp_2=.004), and no significant expression x manipulation interaction (F(3, 51)= 1.23, p=.31, ηp_2 =.07). These findings show that visual contrast thresholds do not vary between face expressions, nor are these findings different according to the two contrast metrics used here (Michelson and RMS). The absence of

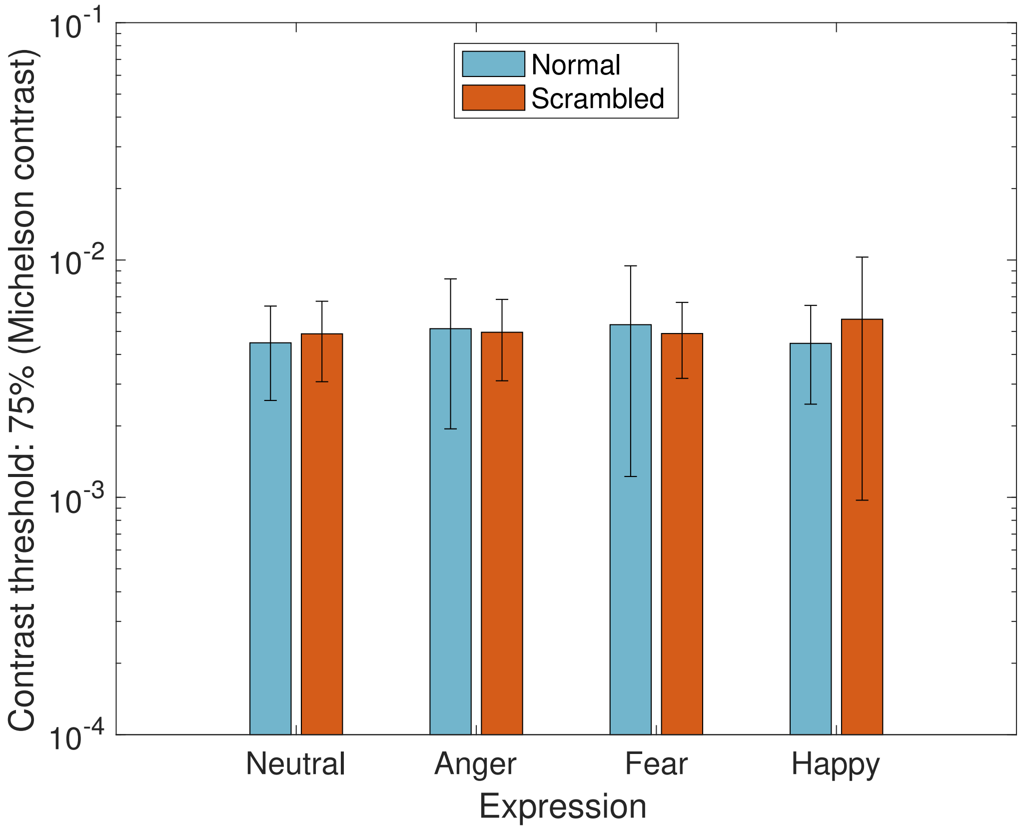

Fig. 8.3 Visual contrast thresholds for neutral, fearful, angry and happy facial expressions. Error bars show standard error values. Faces are unfiltered. Fearful face expressions are not associated with lower visual contrast thresholds, contrary to what might be predicted from Hedger et al. (2015). Error bars represent ±1 standard error of the mean.

an expression-related effect on contrast sensitivity provides evidence against Hedger et al. (2015)'s original claim that fear expressions (compared to neutral faces) exploit the contrast sensitivity function. In an attempt to understand the inconsistency between the present behavioural data, and that generated from image analyses, we conducted the same measure of faces' effective contrast as that performed by Hedger et al. (2015) and extended this to include expressions of anger, happiness and disgust, including a condition where all face images had been either normalised for RMS contrast (as was the procedure for Hedger et al. (2015)) or non-normalised, such that face images were analysed in their raw format, containing possible natural variations in physical contrast.

8.4 Image analyses: effective contrast

Hedger et al. (2015) calculated the effective contrast for face images extracted from 5 face databases: Nimstim (Tottenham et al., 2009), KDEF (Lundqvist et al., 1998), Radboud

(Langner et al., 2010), Montreal (Beaupré and Hess, 2005) and Ekman and Friesen (Ekman, 1976) face sets. Stimuli were cropped to include internal features only and normalised for RMS contrast prior to analyses. Here, we calculate effective contrast for the 16 KDEF face images used in our experimental study, referring to the same procedure described by Hedger et al. (2015). First, Fourier amplitude spectra were calculated for each face image. From the ModelFest dataset (Watson and Ahumada, 2005), we extracted visual contrast thresholds for 10 stimulus parameters. These corresponded to Gabor stimuli, ranging from 1.12-30 cycles per degree. A smooth curve was fit to the average threshold (over 4 repetitions and all observers in the ModelFest data set) using a cubic spline. The resulting contrast sensitive distribution was then multiplied by the Fourier amplitude spectrum for each face image to establish each face's effective contrast. Figure 8.4 shows an example of the procedure for calculating effective contrast for the 16 face images used in the present contrast sensitivity study. To extend our analysis, effective contrast was measured for face images across 4 of the face databases employed by Hedger et al. (2015), with the exception of the Ekman Friesen face set (Ekman, 1976). As outlined by Hedger et al. (2015) the overall estimate of effective contrast for each face image was obtained by summing contrast across spatial frequency after application of the contrast sensitivity model. All face images were analysed in two conditions: after they had been normalised for RMS contrast (according to Hedger et al. (2015)), and also in their raw form, such that no contrast normalisation had taken place. In the RMS-matched analysis, the RMS contrast of each face was set to be equal to that of the image with the lowest contrast in each set. It is for this reason that the RMS-matched stimuli have an overall lower effective contrast. All face images depict forward-facing actors displaying one of 5 expressions (neutral, anger, fear, happy or disgust), cropped to include internal features only. The average effective contrast for each facial expression, compared across the 5 face image samples, including the experimental stimuli for the present contrast sensitivity study, is displayed in Figure 8.5.

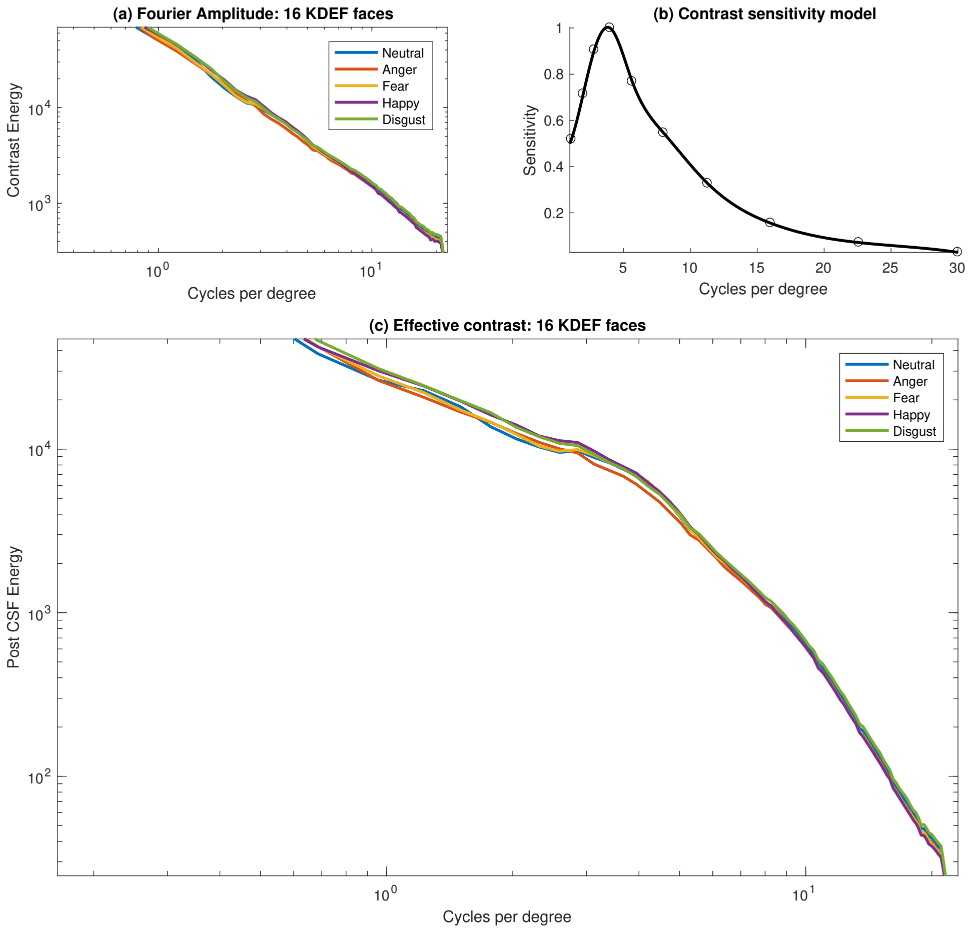

Fig. 8.4 (a) The mean amplitude spectrum for each of the five expressions. (a) The contrast sensitivity function based on the ModelFest data (Watson and Ahumada, 2005). (c) The effective contrast, obtained by multiplying the original amplitude function by the contrast sensitivity function. This method for calculating effective contrast replicates that used by Hedger et al. (2015)

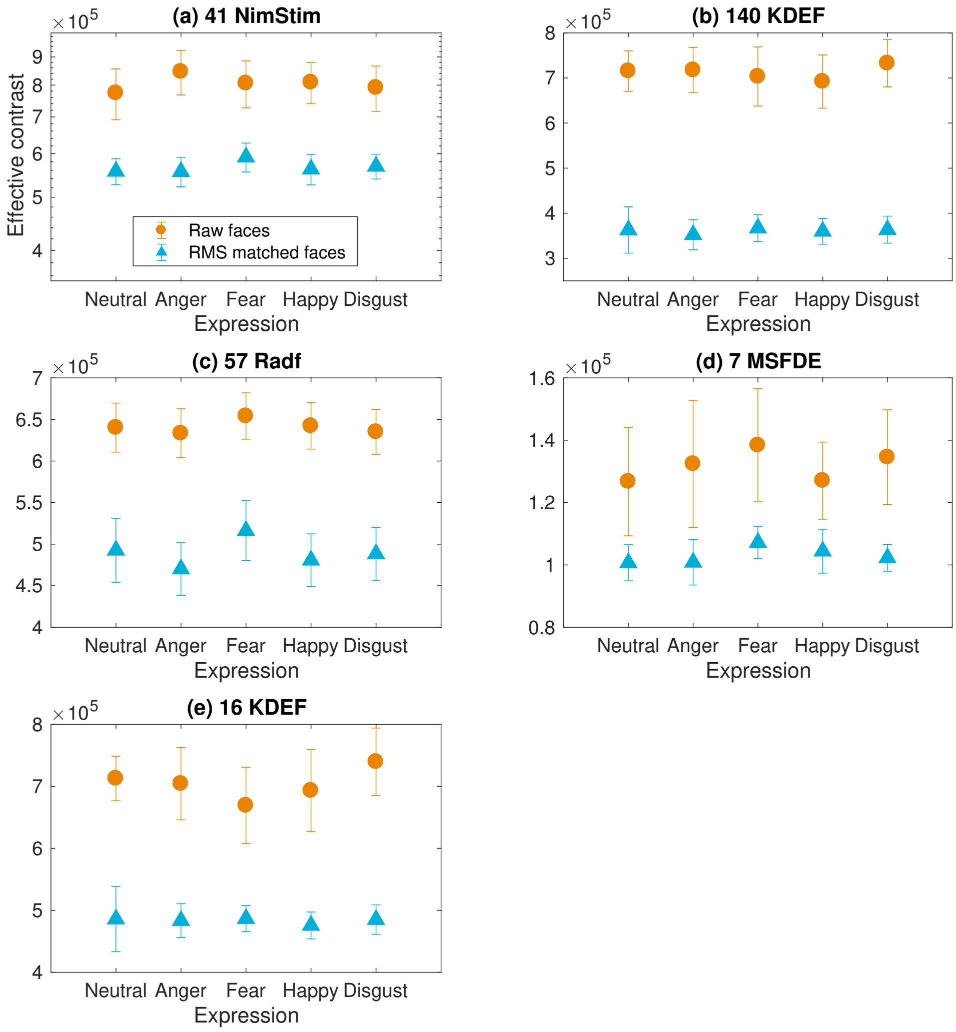

Fig. 8.5 Effective contrast for neutral faces, and anger, fear, happiness and disgust expressions, measured for raw faces (circle data) and the same faces normalised for RMS contrast (triangle data). Effective contrast measures were performed across 4 samples of face images, including the NimStim (a), KDEF (b), Radboud (c), MSFDE (d), face sets employed by Hedger et al. (2015), and for the 16 KDEF faces used in the present contrast sensitivity study (e). Error bars represent associated standard deviations.

42 NimStim face images

Effective contrast for neutral, angry, fearful, happy and disgust NimStim faces are shown in Figure 8.5 (a), and summarised in Table 8.1 (a). Sidak-corrected paired comparisons explored differences in effective contrast between fear expressions and neutral, anger, happy and disgust counterparts (α= 0.0127, accounting for 4 comparisons). When faces had been normalised for RMS contrast, NimStim fear expressions were significantly higher in effective contrast compared to all other expressions, including neutral faces. This finding replicates that observed by Hedger et al. (2015), whereby RMS-normalised fear expressions were found to be significantly higher in effective contrast compared to neutral expressions. Alternatively, when faces aren't normalised for RMS contrast, NimStim fear expressions are significantly higher in effective contrast compared to neutral faces, and lower in effective contrast compared to angry faces. No other significant differences were observed.

For 42 raw (not normalised) NimStim faces, RMS contrast was calculated across the 5 face expressions in order to explore how natural differences in physical contrast compare with expression-related differences in effective contrast that vary according to whether or not they have been contrast normalised. Sidak-corrected comparisons compared RMS contrast between fear expressions and each of their face counterparts, including neutral faces (α= 0.0127). Fearful NimStim faces naturally contained significantly less RMS contrast compared to anger and happy expressions. No other significant differences were observed. These data are illustrated in Figure 8.6, and summarised in Table 8.2 (a).

To summarise effects found for NimStim faces: fearful NimStim faces are not naturally higher in RMS contrast, rather, they are in fact lower in RMS contrast compared to angry and happy NimStim faces. In terms of effective contrast, raw fear expressions are only higher in effective contrast compared to happy faces. However, when RMS contrast normalised, the effective contrast for fear expressions becomes higher than that of all other expressions, including neutral faces.

140 KDEF face images

Effective contrast for neutral, angry, fearful, happy and disgust KDEF faces are shown in Figure 8.5 (b), and summarised in Table 8.1 (b). Sidak-corrected paired comparisons explored differences in effective contrast between fear expressions and neutral, anger, happy and disgust counterparts (α= 0.0127). For KDEF faces normalised for RMS contrast, fear expressions are significantly higher in effective contrast compared to angry faces. Alternatively, when the same KDEF faces are not normalised for RMS contrast, fear expressions are significantly lower in effective contrast compared to disgust expressions. No other significant differences were observed.

For 140 raw (not normalised) KDEF faces, naturally-occurring and expression-related differences in RMS contrast were explored using Sidak-corrected comparisons (α= 0.0127). Fearful KDEF faces naturally contain significantly less RMS contrast compared to neutral, angry and disgust expressions. No other significant differences were observed. These data are illustrated in Figure 8.6, and summarised in Table 8.2 (b).

To summarise effects found for KDEF faces: fearful KDEF faces are naturally lower in RMS contrast compared to all but happy expressions. The effective contrast of raw fearful faces only differs compared to expressions of disgust, where effective contrast for fear expressions is lower compared to that associated with disgust expressions. However, when these KDEF faces are normalised for RMS contrast, the effective contrast for fear expressions becomes higher, but only compared to that of angry faces.

57 Radboud face images

Effective contrast for neutral, angry, fearful, happy and disgust Radboud faces are shown in Figure 8.5 (c), and summarised in Table 8.1 (c). Sidak-corrected paired comparisons explored differences in effective contrast between fear expressions and neutral, anger, happy and disgust counterparts (α= 0.0127). For Radboud face images normalised for RMS contrast, fear expressions are significantly higher in effective contrast compared to all other expressions, including neutral. Alternatively, when the same Radboud faces are not normalised for contrast, the same effect is true; raw fear expressions are significantly higher in effective contrast compared to all other expressions, including neutral. These findings in particular require further discussion, presented in the following section.

For 57 raw Radboud faces, naturally-occurring and expression-related differences in RMS contrast were explored using Sidak-corrected comparisons (α= 0.0127). Fearful Radboud faces naturally contained significantly less RMS contrast compared to all other face expressions. These data are illustrated in Figure 8.6, and summarised in Table 8.2 (c).

To summarise effects found for Radboud faces: fearful Radboud faces are naturally lower in RMS contrast compared to all other expressions, including neutral faces. Interestingly, effective contrast for these fear expressions, both in their raw and RMS normalised format, are higher in effective contrast compared to their other face counterparts.

7 Montreal face images

Effective contrast for neutral, angry, fearful, happy and disgust Montreal faces are shown in Figure 8.5 (d), and summarised in Table 8.1 (d). Sidak-corrected paired comparisons explored differences in effective contrast between fear expressions and neutral, anger, happy and disgust counterparts (α= 0.0127). For Montreal faces that are normalised for RMS contrast, fear expressions are significantly higher in effective contrast compared to neutral faces. Alternatively, when the same Montreal faces are not normalised for RMS contrast, effective contrast in fear expressions does not differ significantly compared to any other face. No other significant differences were observed.

For 7 raw Montreal faces, naturally-occurring and expression-related differences in RMS contrast were explored using Sidak-corrected comparisons (α= 0.0127). Fearful Montreal faces do not naturally differ in terms of RMS contrast compared to any other face expression, including neutral faces. These data are illustrated in Figure 8.6, and summarised in Table 8.2 (d).

To summarise effects found for Montreal faces: fearful Montreal do not naturally differ in terms of their RMS contrast content compared to any of their face counterparts. Effective contrast also does not differ between fear and other expressions for these same, raw faces. However, when these faces are normalised for RMS contrast, effective contrast in fearful faces increases, but only to the extend that it differs, and is higher than, that associated with neutral faces.

16 KDEF face images (experimental stimuli)

Effective contrasts for the experimental face stimuli used in our contrast sensitivity study are shown in Figure 8.5 (e), and summarised in Table 8.1 (e). Sidak-corrected paired comparisons explored differences in effective contrast between fear expressions and neutral, anger, happy and disgust counterparts (α= 0.0127). For the 16 KDEF (experimental stimuli) that were normalised for RMS contrast, no differences in effective contrast were observed between fear and any other expression, including neutral. Alternatively, when the same faces were not normalised for RMS contrast, fear expressions are significantly lower in effective contrast compared to both neutral and disgust faces. No other significant differences were observed.

For the 16 raw KDEF faces used in the present contrast sensitivity study, naturally-occurring and expression-related differences in RMS contrast were explored using Sidak-corrected comparisons (α= 0.0127). Experimental fearful expressions were lower in RMS contrast compared to disgust. No other significant differences were observed. These data are illustrated in Figure 8.6, and summarised in Table 8.2 (e).

To summarise effects found for 16 KDEF experimental stimuli: natural differences in RMS contrast were observed between fear and disgust, were fearful faces contained naturally lesser amounts of RMS contrast than disgust faces. These same raw faces were lower in effective contrast compared to both neutral and disgust faces. However, when they were normalised for RMS contrast, effective contrast did not differ between fear and any other face counterpart, including neutral faces.

Together, data from the present contrast sensitivity study showed that visual contrast thresholds are not influenced by differences between images of facial expressions. Namely, fearful expressions portrayed by face images did not enhance observers' contrast sensitivity, in contrast to predictions following findings from Hedger et al. (2015). Fearful expressions, according to image analyses by Hedger et al. (2015) are higher in effective contrast, and thus well-tuned to contrast sensitivity processing. This proposal was driven by data from image analyses measuring differences in effective contrast between fear and neutral face images that had been normalised for RMS contrast. The stimuli used in the present contrast sensitivity study were raw face images that were not normalised for physical contrast in any way; perhaps a difference in experimental design between the present experiment and that of Hedger et al. (2015) that may account for a discrepancy between our findings. We replicate measures of effective contrast used by Hedger et al. (2015) to establish the extent that CSF advantages exclusive to fear expressions might be driven by the effect of contrast normalisation on the effective contrast of faces. A general, but not robustly consistent, trend across the present image analyses is that greater effective contrast in fear expressions is an effect that is most pronounced when images have undergone contrast normalisation (RMS). This was the case for the KDEF database; a set of facial stimuli used as both the experimental stimuli in the present CSF study, and that which was included in image analyses conducted by Hedger et al. (2015). Importantly, although there is a fair pattern of effects that favour effective contrast in RMS normalised, but not raw, fear expressions, it is important to note that this was not true across all analyses for all face databases.

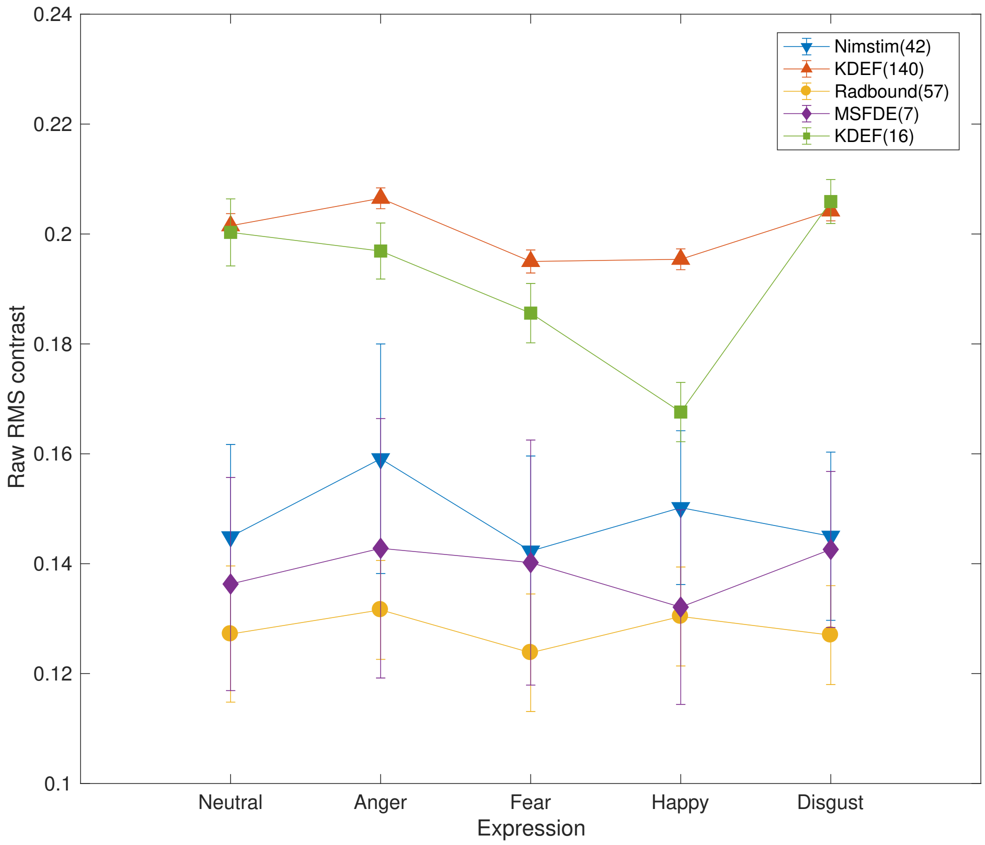

Fig. 8.6 RMS contrast for face expressions *before* faces are subjected to contrast normalisation i.e. when kept in raw format. RMS contrast for 5 expressions is measured across the 5-database face samples used to calculate faces' effective contrast. Error bars represent associated standard deviations.

8.5 Data tables

Table 8.1 Effective contrast compared between fear and counterpart expressions. Sidak-corrected comparisons (α= 0.0127, accounting for 4 comparisons) are made between raw, and thus not normalised faces, and also when they are matched for RMS contrast. Measures are performed for 4 databases (a-d), and experimental stimuli used in the present behavioural study (e).

(a) 42 NimStim faces

Not normalised	t	df	CI	Sig.
Fear-neutral	3.97	41	15887.84, 48742.20	<.001
Fear-anger	-3.50	41	-3337.05, -16987.44	.001
Fear-happy	-.36	41	-21229.94, 14786.41	.72
Fear-disgust	1.27	41	-8667.98, 38236.24	.21

RMS normalised	t	df	CI	Sig.
Fear-neutral	9.75	41	27128.97, 41299.98	<.001
Fear-anger	6.58	41	24173.68, 45546.27	<.001
Fear-happy	7.79	41	21670.22, 36397.82	<.001
Fear-disgust	5.62	41	14337.81, 30382.54	<.001

(b) 140 KDEF faces

Not normalised	t	df	CI	Sig.
Fear-neutral	-2.40	139	-21956.07, -2130.43	.018
Fear-anger	-2.47	139	-25622.43, -2854.71	.015
Fear-happy	1.81	139	-1011.76, 23548.61	.075
Fear-disgust	-4.96	139	-40878.60, -17581.18	<.001

RMS normalised	t	df	CI	Sig.
Fear-neutral	.96	139	-4403.07, 12759.73	.33
Fear-anger	4.96	139	8864.86, 20593.09	<.001
Fear-happy	2.49	139	1473.99, 12835.65	.014
Fear-disgust	1.27	139	-1958.42, 8989.18	.20

(c) 57 Radboud faces

Not normalised	t	df	CI	Sig.
Fear-neutral	6.05	56	9302.71, 18501.62	<.001
Fear-anger	8.01	56	15594.01, 25979.80	<.001
Fear-happy	3.92	56	5821.49, 17985.59	<.001
Fear-disgust	9.00	56	14834.82, 23323.09	<.001

RMS normalised	t	df	CI	Sig.
Fear-neutral	10.18	56	18898.93, 28150.01	<.001
Fear-anger	18.90	56	41225.73, 50999.21	<.001

Fear-happy	14.01	56	30423.72, 40576.68	<.001
Fear-disgust	12.10	56	23278.46, 32513.48	<.001

(d) 7 Montreal (MSFDE) faces

Not normalised	t	df	CI	Sig.
Fear-neutral	1.95	6	-2924.24, 26224.08	.09
Fear-anger	.60	6	-18156.74, 30104.36	.56
Fear-happy	1.78	6	-4255.20, 26954.75	.12
Fear-disgust	.73	6	-8982.68, 16669.87	.49

RMS normalised	t	df	CI	Sig.
Fear-neutral	4.65	6	3087.98, 9943.33	.003
Fear-anger	2.82	6	843.11, 11834.65	.03
Fear-happy	1.82	6	-955.55, 6513.98	.18
Fear-disgust	2.67	6	424.47, 9425.20	.037

(e) 16 KDEF faces

Not normalised	t	df	CI	Sig.
Fear-neutral	-3.49	15	-70194.13, -16964.54	.003
Fear-anger	-1.86	15	-75544.43, 5119.73	.08
Fear-happy	-1.59	15	-56093.38, 8083.37	.13
Fear-disgust	-3.65	15	-111632.35, -29367.94	.002

RMS normalised	t	df	CI	Sig.
Fear-neutral	.05	15	-29304.24, 30754.92	.96
Fear-anger	.43	15	-12714.96, 19164.70	.67
Fear-happy	2.53	15	1753.64, 20077.86	.02
Fear-disgust	2.77	15	-10552.70, 13706.44	.78

Table 8.2 Differences between RMS contrast in raw fear expressions and 4 emotion counterparts (α= 0.0127). Fear comparisons are measured across all 4 databases (a-d), and also for the experimental stimuli used in the present contrast sensitivity study (e).

Raw RMS:

(a) 42 Nimstim faces

	t	df	CI	Sig.
Fear-neutral	-2.27	41	-.006, .001	.21
Fear-anger	5.46	41	-.023, -.010	<.001
Fear-happy	-4.09	41	-.011, -.003	<.001
Fear-disgust	1.06	41	-.007, .002	.29

(b) 140 KDEF faces

	t	df	CI (95%)	Sig (2 tailed)
Fear-neutral	-2.67	139	-.011, -.001	.008
Fear-anger	-5.09	139	-.016, -.007	<.001
Fear-happy	-.19	139	-.005, .004	.85
Fear-disgust	-4.27	139	-.013, -.004	<.001

(c) 57 Radboud faces

	t	df	CI (95%)	Sig (2 tailed)
Fear-neutral	-4.37	56	-.005, -.001	<.001
Fear-anger	-10.24	56	-.009, -.006	<.001
Fear-happy	-7.05	56	-.008, -.004	<.001
Fear-disgust	-4.55	56	-.004, -.001	<.001

(d) 7 Montreal (MSFDE) faces

	t	df	CI (95%)	Sig (2 tailed)
Fear-neutral	.70	6	-.009, .017	.50
Fear-anger	-.26	6	-.025, .020	.79
Fear-happy	1.17	6	-.001, .025	.28
Fear-disgust	-.45	6	-.015, .010	.66

(e) 16 KDEF (experimental stimuli) faces

	t	df	CI (95%)	Sig (2 tailed)
Fear-neutral	-2.42	15	-.027, -.001	.02
Fear-anger	-1.56	15	-.026, .004	.13
Fear-happy	-2.15	15	-.022, .000	.04
Fear-disgust	-3.46	15	-.032, -.007	.003

8.6 Discussion

A widely accepted view in the threat bias literature is that fearful face expressions possess a special status in the human visual system, due to their low level image content (Bayle et al., 2011; Carlson and Reinke, 2008; Gray et al., 2013; Hedger et al., 2015; Holmes et al., 2005; Phelps et al., 2006). Hedger et al. (2015) recently showed that visibility, or salience, associated with fear expressions is predicted by their effective contrast content; the extent that the Fourier amplitude spectrum of fear expressions, compared to neutral faces, exploits the contrast sensitivity function. In the present study, we conducted a traditional contrast sensitivity task to test whether higher effective contrast purported for fear expressions is associated with lower visual contrast thresholds at the behavioural level. Here, we measured contrast sensitivity for facial stimuli of 5 raw face expressions. No expression-related differences were observed across visual thresholds, as was predicted based on data from Hedger et al. (2015). Specifically, a decrease in visual thresholds for fearful expressions was not observed. Hedger et al. (2015) showed greater effective contrast unique to fear expressions (when compared to neutral faces) was found only for face images that had been normalised for RMS contrast. In order to investigate whether the use of contrast normalisation in Hedger et al. (2015) may have driven effective contrast effects that in its absence were not replicated by our contrast sensitivity study, we repeated calculations of effective contrast using the same procedure employed by Hedger et al. (2015). Here, effective contrast was calculated for images of face expressions both when they were normalised for RMS contrast, as was performed by Hedger et al. (2015), but also when the same faces had not been normalised for physical contrast. These analyses were performed for the NimStim, KDEF, Montreal (MSFDE) and Radboud face sets used by Hedger et al. (2015), and also for the 16 KDEF face images used as the experimental stimuli in the present contrast sensitivity study. Importantly, our findings replicate those of Hedger et al. (2015), showing that fear expressions normalised for RMS contrast are significantly higher in effective contrast than neutral counterparts. We extend this finding to show that this is also true when fearful faces are compared to other face expressions. This advantage was observed for NimStim, KDEF and Radboud face databases. However, when the same faces were analysed in their raw form (i.e. when they were not normalised for physical contrast), this effect of fear diminishes for Nimstim and KDEF face databases. These findings indicate that the process of normalising face stimuli, to some extent, significantly increases the effective contrast in fearful face expressions, where naturally (not normalised) these faces tend not to differ in effective contrast compared to other facial expressions, or indeed are more likely to be lower in effective contrast. An important finding to discuss here is the absence of this contrast normalisation effect for face images taken from the Radboud face database. Here, we observed that Radboud fear expressions normalised for contrast were significantly higher in effective contrast compared to neutral faces, as well as other expressions; an effect that is consistent with that observed by Hedger et al. (2015). However, this effect did not diminish when images were not normalised for contrast; an effect that was not found for other face samples. Radboud face images were included in the present study on the basis that they were included in the original study by Hedger et al. (2015). Details of the image processing used to create and standardise these actor photographs includes white-balance correction (Langner et al., 2010). This process adjusts raw image data in order to remove certain unrealistic and biased appearances, such as those incurred under different lightning conditions (Innocent, 2011; Yanof et al., 2007). It is important to note that database production information for KDEF and NimStim face sets do not refer to any image processing related to white-balance correction, or contrast normalisation (Lundqvist et al., 1998; Tottenham et al., 2009). No information about image processing is provided for the Montreal image database (Beaupré and Hess, 2005). It may therefore be that contrast and luminance information in 'raw' Radboud face images had already been subjected to some degree of normalisation, or standardisation, when they were created. In sum, the present study performed a traditional contrast sensitivity task to address the proposal that fearful faces exploit the contrast sensitivity function, and as a result undergoes efficient visual processing (Hedger et al., 2015). Together, these findings suggest that contrast normalisation –a standard procedure in psychophysical studies- significantly influences the physical composition of face stimuli in a way that can be expected, in some cases, to influence their perceived salience under both experimental and neurophysiological conditions.

Chapter 9

An overview of the techniques used to investigate perceptual biases for fearful expressions

Nathaniel Rateliff and the Night Sweats. (2015). Artwork copyright to N. Rateliff and the Night Sweats (2015).

The overarching theme of Chapters 1-8 is, for the most part, concerned with the physical composition of facial expressions, with a particular focus on fearful facial expressions. The focus of the present and following chapters changes somewhat to the behavioural techniques used to measure perceptual biases for fear expressions. Part 1 of the present chapter presents an overview of some of the most commonly used experimental paradigms. Part 2 focuses specifically on the continuous flash suppression and saccadic latency techniques. Understanding the benefits and costs associated with these techniques, in both theory and practice as means for assessing the threat bias, forms the motivation for Experiments 4 and 5.

9.1 Part 1: Techniques in visual attention and conscious perception

Between cognitive and psychophysical studies, the range of experimental techniques used to measure perceptual biases for face stimuli is vast. Many of those concerned with perceptual biases for fear expressions place a particular focus on the role of processes that operate within visual attention and conscious perception (Carlson and Reinke, 2008; Lang et al., 1997; Posner et al., 1980; Pourtois et al., 2006; Stein et al., 2014; Whalen et al., 2001; Yang et al., 2007). This is due to similar capacity limitations in both visual attention and conscious visual awareness. The debate regarding the extent to which visual attentional and awareness processes are dissociable is on-going (e.g., Koch and Tsuchiya (2007)), and so naturally studies that measure one of these visual functions often also measure the other, and findings cannot always be unambiguously attributed to just one of these processes. In terms of visual attention, because it is not possible to attend to all aspects of a visual scene at any given time, attentional processes must selectively identify and extract from an array of visual information that which may be of particular importance, or require higher-acuity detailed analysis (Dehaene and Changeux, 2011; Lavie et al., 2004). Similar capacity limitations also occur in conscious visual awareness, such that only a subset of the visual information available to the visual system can reach visual consciousness at any one time. Here, the time for a given visual stimulus to compete and achieve conscious perception over that of neighbouring information is thought to be determined by the extent to which the visual system regards it as significant (Gray et al., 2013; Stein et al., 2014). Experimental paradigms that focus on constructs associated with visual attention and conscious perception are useful tools for measuring the degree to which a certain visual stimulus may be preferentially selected and prioritised during perception (Carlson and Reinke, 2008; Dehaene and Changeux, 2011; Lavie et al., 2004). In terms of the threat bias literature, these techniques have been used to address a number of questions: to what extent do fearful expressions undergo processing independently of conscious visual perception; can this be indexed by preferential allocation of attentional resources to fearful stimuli; and does the presence of an unconsciously perceived fear expression maintain the capacity to influence cognitive functioning and perceptual behaviours in a way that we might expect if such fear expressions really do possess a unique function within the threat-avoidance system (Hedger et al., 2016; Öhman and Mineka, 2001b). These questions are not mutually exclusive.

9.1.1 Masked visual probe paradigm

Visual probe tasks measure the effects of exogenously elicited spatial attention, and the degree to which this is influenced by threat-relevant stimuli; an effect that translates to the relevance of a threatening stimulus (MacLeod and Mathews, 1988). In other words, if a given visual stimulus is evolutionarily (threat) relevant, does it facilitate spatial attention in a way that might be expected if its function is to maximise chances of survival? Dot probe tasks address whether a fearful expression in the visual field is better at capturing visual attention, compared to a non-threatening face image. Such a trial typically begins with a central fixation point, followed by 2 different stimulus images appearing simultaneously at opposite visual fields from the centre. One of these images will be a fearful face, and the other a counterpart expression (such as a neutral face). When these faces are ‘masked’, both face images appear in synchrony for around 40ms, such that they are visually registered but not consciously ‘seen’ by participants. The presentation of the stimulus is therefore consciously noticed by the observer, but the stimulus content, i.e. the expression of the face, is unidentifiable. A single target dot then appears at one side of the visual field. The dot’s location will be either spatially congruent with the position of the fearful face, where they have both appeared in the same spatial location, or it will be spatially incongruent with the fearful face, where the dot and fearful face have occupied opposite spatial locations (Carlson and Reinke, 2008). The observers’ task is to indicate the location of the target dot, usually in a 2AFC task. Faster reaction times to detect a congruently located dot indicate that visual attentional resources are allocated to the area of visual space previously occupied by a fearful face and, because of this, are readily available to detect a target dot. Conversely, slower reaction times to locate incongruently located dots denote difficulty to disengage attention from spatial regions previously occurred by fear expressions (Carlson and Reinke, 2008; Lang et al., 1997; Posner et al., 1980; Pourtois et al., 2006). The question addressed by the masked visual probe paradigm is straightforward: are fearful expressions capable of capturing and directing visual attention, even at the point where they have been inhibited from reaching visual conscious awareness? A primary benefit of this particular version of the dot probe task is the way that an observer’s capacity to attend to a fear expression is artificially attenuated. Rapid presentations of face images prevent the degree of temporal integration of neural responses required for it to reach conscious visual attention. The result is that the face stimulus is transient; it is not explicitly ‘seen’ by observers but is present long enough to influence spatial attentional functioning. This design is thus thought to measure two perceptual functions: attentional capture and disengagement (or redirection).

Example studies: Using a masked dot-probe task Carlson and Reinke (2008) showed that the presence of a fear expression, although masked from conscious perception, modulates the spatial distribution of visual attention. Their task began with a central fixation point, followed by 2 face images appearing simultaneously at opposite sides from centre. One of these faces was a fearful expression, and the other a neutral face counterpart. Both face images were presented for 33ms, such that they were ‘masked’ from observers’ conscious awareness. Such brief presentation times (<40ms) of images limit the progression of sensory processing required to achieve conscious recognition. After the presentation of a face image, a single dot-probe was presented at a location either spatially congruent or incongruent with the fearful face. In a 2-alternative-forced-choice task (2AFC), observers manually indicated the location of the dot-probe, providing response times for the trials. Results from the study showed that dot-probes were more quickly and accurately indicated when preceded by a masked fearful expression (i.e. on congruent trials), compared to their neutral counterparts. These findings suggest that visual attention has been allocated and engaged to the area of visual space previously occupied by a fearful face, facilitating rapidity to detect a dot-probe. Conversely, slower reaction times for an incongruently located dot-probes suggest that fearful faces are associated with inhibition to disengage attention from the location of a fearful face such that it may be reoriented to the accurate location of a dot-probe (Carlson and Reinke, 2008). These findings were also compatible with models of visual attention and emotion that posit the propensity for threat-relevant visual cues to capture and maintain attention for longer (Öhman and Mineka, 2001b).

Bocanegra and Zeelenberg (2009) used a psychophysical variation of the traditional dot probe paradigm to investigate effects of transient presentations of fear and neutral expressions on sensitivity to different spatial frequency information. Their experimental design is like that outlined by Carlson and Reinke (2008), the primary difference being that instead of using a 2AFC task to indicate the location of a target dot, a 2AFC task was used to indicate the orientation of a grating -a stimulus with a narrow spatial frequency range- tilted left or right from centre. Data from Bocanegra et al. (2012) showed an association between masked fear expressions and better performance to detect the orientation of gratings with a low frequency content. This effect was not found for high frequency gratings, where the presentation of unseen fear expressions produced a deficit effect on observers' ability to detect the orientation of such gratings. Interestingly, the emotion portrayed by face stimuli did not influence observers' ability to locate the spatial location of gratings i.e. the side of screen occupied by the grating (Bocanegra et al., 2012). These dissociable effects of a fear expression on grating detection and grating location suggests that fearful faces influence contrast sensitivity processing (indexed by grating detection), but not overt spatial location (indexed by grating location). This data is consistent with that of Phelps et al. (2006), where transient fear expressions are associated with lower visual contrast thresholds, and therefore for relatively low spatial ranges. It is argued that a function of the threat bias may be to stimulate and improve processing for coarse features conveyed by low frequency information, in a way that does not discriminate the spatial location of such target stimuli (Bocanegra et al., 2012). This is thought to be achieved via coarse processing streams by fear expressions, mobilising visual processes implicated in the encoding of information related to motion, depth, and global cues; necessary cues for maximising successful threat avoidance (Bocanegra et al., 2012; Hedger et al., 2015).

It is important to include here that two studies often cited as producing such inconsistencies are Carlson and Reinke (2008)'s study that showed an association between masked fearful faces and enhanced attentional capture, but not disengagement, and data from Koster et al. (2004) who reported findings of both attentional capture and disengagement effects associated with threatening stimuli. It has been suggested that the effect of disengagement may be unique to certain experimental parameters. Evidence of this disengagement effect are cited by studies of sub-clinical participant populations, whose participants include high-trait anxiety individuals (Fox et al., 2001, 2002; Yiend and Mathews, 2001). It may be that sensitivity to threat-relevant information is enhanced in such populations, to the extent that the absence of disengagemrnt of attention may be the result of adaptive mechanisms that have gone 'awry' in cases of high-trait anxiety (MacNamara et al., 2013). If this is the case, it may be that fearful expressions only produce the difficulty-to-disengage effect within certain populations, rather than at the universal level. Of equal importance is that Koster et al. (2004) employ visual stimuli defined as threatening objects, such as knives, and therefore do not refer specifically to fearful expressions. The inconsistency of findings between Carlson and Reinke

(2008) and Koster et al. (2004) suggests that fear expressions capture but do not maintain visual attention, while threatening objects influence both perceptual functions. Theoretically, we might expect this to be the case. An effective mechanism for threat avoidance requires fear expressions to be both rapidly detectable but equally simple to disengage from, such that attention can readily able to reorient to a location inhabited by an environmental threat. In contrast, the ability to maintain visual attention for stimuli such as knives is more likely to be an object-specific aspect of the threat bias. Acknowledging differences in experimental techniques and methodology is therefore necessary to disentangle mixed evidence of a bias exclusive to fearful faces.

9.1.2 Attention in the periphery

Other studies of visual attention measure the degree to which fear expressions are preferentially attended to as their eccentricity increases, presenting stimuli up to 40° outside of foveal vision. Bayle et al. (2011) presented observers with images of fear and neutral expressions, as well as faces of different genders. Face stimuli were presented at 7 different eccentricities, ranging from 10° to 40° into the peripheral visual field. In a detection task, participants were required to indicate the point at which they perceived the face, and its location. Observers were most efficient at detecting fear compared to neutral expressions up to 40° of eccentricity. Bayle et al. (2011) interpret these findings as evidence that the visual system's ability to preferentially detect fearful faces remains intact even when such stimuli deviates from central vision, and that it does so in a way that is not true for neutral and disgust face images.

9.1.3 Neuroimaging data

Data from neuroimaging studies implicate the role of limbic regions such as the amygdala during processing of emotionally laden facial stimuli, such as fear expressions. Often, the observer's task during these studies is simply to passively observe presentations of facial stimuli. In some cases, passive observation includes observers being presented with images that are masked from visual awareness. Data from fMRI studies show increases of oxygenated blood flow to the amygdala while observers are presented with rapid displays of fearful facial expressions (Breiter et al., 1996; Morris et al., 1999; Whalen et al., 1998). Backward-masking techniques involve the rapid presentation of a target stimulus followed immediately by second image intended to attenuate, or mask, conscious registration of the former stimulus (Breitmeyer et al., 2006). Whalen et al. (1998) presented observers with images of fearful and happy facial expressions for 33 ms, followed immediately by a neutral face expression to "mask" the faces from visual awareness. Despite observer reports of not having consciously registered emotional faces, amygdala activity remained significantly greater in response to fearful masked faces compared to happy masked faces. These data suggest that neural registration occurs for fearful facial expressions prior to their conscious appraisal, and that this method is exclusive to fearful expressions, rather than occurring as a general emotion effect (Morris et al., 1998, 1999). It is important to acknowledge the mixed evidence from neuroimaging studies. These inconsistencies are primarily found for different levels of activation in response to fearful facial stimuli, and ambiguity regarding the extent that effects are unique to fear expressions presented outside of visual conscious awareness. Though, inconsistencies may be accounted for by limitations of the temporal and spatial sensitivity associated with certain neuroimaging techniques (for more information, see Phillips et al. (2004) and Palermo and Rhodes (2007).

9.1.4 Saccadic latency

Behaviourally, visual attention operates via orienting of the eyes toward a focal point of interest. For larger viewing scenes this involves a synchronised shift of both the head and the oculolomotor muscles that determine gaze orientation. At smaller scales, such as the perception of a single image, visual attention is indexed by rapid shifts, or step-like rotations of the eye towards a single point of fixation (Carlson and Reinke, 2008; Hughes, 1991; Posner et al., 1980). The time that lapses between points of fixation is not associated with significant visual processing. Instead, encoding takes place at fixation points where the eyes remain stationary at the newly landed focal point. These shifts refer to saccadic eye movements and are modulated by selective visual attentional mechanisms (Hughes, 1991). Saccadic eye movements can be explicitly directed towards an endogenous stimulus, or occur as implicit reflexive responses to unexpected, transient stimuli (Carlson and Reinke, 2008). Under experimental conditions, visual features that successfully guide saccades to their location, or encourage or inhibit the redirection of saccades, reflect selective allocation of visual attentional resources to the given stimulus. An important feature of saccadic eye movements is their independence from conscious awareness; they occur as reflexive and automated responses to external visual stimuli. For this reason, they are useful measures of implicit perceptual biases. In studies of the threat bias, monitoring eye movements toward a visual target indexes its associated capacity to obtain visual attentional resources (Bannerman et al., 2009a; Quaia et al., 1999). In other words, the smaller the time delay (measured in milliseconds) between the onset of a visual stimulus and the first successful saccade toward its location suggests that it receives preferential and automatic allocation of attention resources. It is important to note here that metrics associated with saccadic eye movements are multidimensional. Saccadic eye movements can be measured in terms of the time taken to initiate the first successful orientation toward a target, the magnitude and direction between successive saccades, and the trajectory of these movements across a percept. This thesis is concerned only with the first saccadic metric. Initial orienting saccade, or saccadic latency response time, is a crude measure of the point at which a target stimulus is automatically detected and attended to.

Example study: Bannerman et al. (2010) presented observers with fearful or neutral faces in an exogenous attentional cueing paradigm. Faces were presented for durations of either 20 or 100ms, and observer responses were measured in the form of both manual and saccadic responses, counterbalanced across trials. Saccadic response data showed that fearful face cues enhanced attentional capture and impaired attentional disengagement, but only when these faces were presented for short (20ms) durations. Findings from the manual response data showed that when observers were required to manually indicate the location of a face, an emotional modulation effect for inhibiting attentional disengagement was only found for longer (100ms) cue durations. Bannerman et al. (2010) concluded that the fear expression effect for 20ms cue durations occurred because fearful faces are able to capture and maintain attentional resources in vision, where this effect is most detectable when measuring implicit oculomotor behaviours compared to calculated manual responses.

9.1.5 Binocular rivalry and Continuous Flash Suppression

Under normal viewing conditions, each eye receives a slightly different image of the world. Binocular fusion is the natural process by which the two representations are brought together to form a single percept, while allowing the differences between the images to be used as a cue to three-dimensional shape (Blake and Logothetis, 2002). When these 2 inputs deviate significantly from each other binocular fusion cannot be achieved, resulting in perception that switches interchangeably between the two inputs, or representations (Hedger et al., 2016; Tsuchiya and Koch, 2005). The process by which information presented to one eye must compete for conscious appraisal with that received by the other eye is termed binocular rivalry. Here, each input progresses through stages of the visual processing hierarchy while simultaneously exerting inhibition upon its competitor (Tsuchiya and Koch, 2005). The representation that exerts the most inhibition during the competition for visual processing successfully achieves conscious perception, where this is experienced as the perceptual shift from one image to the other (Alais, 2012; Alpers and Gerdes, 2007; Hardcastle, 2003). Under experimental conditions, binocular rivalry can be synthesised using dichoptic presentation, where two different images are presented to each eye, typically via the use of a mirror stereoscope or LCD glasses. Dichoptic presentation of visual stimuli is used to measure perceptually the degree of dominance associated with a certain percept and is a central feature of traditional binocular rivalry paradigms (Hedger et al., 2016). This has been a commonly employed experimental technique for investigations of specialised and pre-conscious processing of certain facial stimuli. In terms of measuring a bias for fearful faces, prioritisation for processing is denoted by the expression's effectiveness to achieve the status of being the 'first percept' (Carter and Cavanagh, 2007). In other words, if fear expressions are detected faster, or seen first, compared to counterpart stimuli, it can be inferred that this is due to selective visual processes that grant such stimuli prioritised access to visual awareness (Ritchie et al., 2013).

Continuous flash suppression (CFS) is a variant of the traditional binocular rivalry paradigm that induces a strong intraocular suppression against the stimulus of interest (Lin and He, 2009). Continuous flash suppression thus also incorporates dichoptic stimulation to induce 2 different visual stimuli, but instead of using static images, as are used in binocular rivalry paradigms, CFS relies on spatiotemporal dynamic masking to attenuate the visibility of a target stimulus. This dynamic mask is usually composed of a high contrast noise, or Mondrian-style pattern, and usually has a refresh rate of 10hz; changing, or appearing to flash, every 100ms. A dynamic mask is presented to one eye where it will repeatedly register for conscious processing, while the other eye is presented with the static face image, such as a fear expression. Dichoptic presentation of both a dynamic noise mask and a target facial expression prevents the facial stimulus from achieving visual awareness, therefore rendering it invisible for relatively long periods of time (Tsuchiya and Koch, 2005). Observers provide manual responses to indicate the presence or location of target stimuli, indexing intraocular suppression duration for a given stimulus. The associated interocular suppression duration for a stimulus translates to the point at which is 'breaks' suppression. Faster response times imply that a stimulus undergoes rapid and prioritised visual processing; an indication that the associated stimulus is regarded as important by the visual system (Gayet and Stein, 2017; Gray et al., 2013; Tsuchiya and Koch, 2005). CFS has become popular for investigating the extent to which fearful faces might gain preferential access to unconscious processing and prioritised access to perceptual awareness (Gray et al., 2013; Stein et al., 2014). Such CFS studies evidence an advantage for fearful faces to break through visual suppression faster compared to other facial expressions (Gray et al., 2013; Jiang et al., 2007; Stein et al., 2014; Yang et al., 2007). These findings further support the notion that selective processes operate in unconscious perception that may preferentially process threat-relevant cues in facial stimuli.

Example study: Yang et al. (2007) presented observers with images of fearful, happy or neutral facial expressions. Monocular presentations of face images were accompanied by dynamic mask stimuli presented to the other eye. As part of a 4AFC task, observers indicated as fast as possible 1 of 4 quadrants that the face stimulus was assigned to. Fearful faces were associated with faster response times compared to neutral and happy expressions. This finding was also true when faces were inverted, such that the bias for conscious registration of fear expressions remained even when their configural content was disrupted. These findings suggest that fear expressions may achieve conscious visual perception at a faster rate than other facial displays of emotion, perhaps because they possess a special status within the visual system that is signalled by their low-level stimulus properties (Yang et al., 2007).

9.1.6 Paradigm differences

The main commonality shared by visual attentional paradigms is the shared aim to understand unconscious visual processing of fear expressions by way of disrupting normal visual processing. However, the designs used to achieve this can vary significantly between studies (Hedger et al., 2016). Masked visual probe experiments, for example, have been used to investigate threat-relevant images in a broader sense rather than for fear expressions specifically. This includes dot probe tasks that evidence biases for threatening objects such as knives (Koster et al., 2004), emotionally-laden words (Hunt et al., 2006), or the use of the sub-clinical anxiety populations (Fox et al., 2001). This diminishes the extent that threat bias findings from the masked probe paradigm in general can be cited as evidence of the bias for fear expressions. Data outlined above from Carlson and Reinke (2008) and Bocanegra and Zeelenberg (2009) appear to be the only two studies whose masked probe paradigms provide evidence of a fear expression bias, but because both studies measure different dimensions of perception it is difficult to compare their findings. Carlson and Reinke (2008) measure effects of fear expression during the allocation of attentional resources, whereas Bocanegra and Zeelenberg (2009) measure whether the same expressions better activate early visual functioning. Both conclude a fear advantage effect, but one that appears to happen in different modalities -attention and sensitivity- such that the possible mechanism is unclear. There is also a methodological appeal of the process by which such masking paradigms render a stimulus invisible to an observer. However, it has been argued by Bannerman et al. (2010) that the use of rapid stimulus duration necessitates a sensitive measure of implicit responses, rather than of manual responses measured by almost all studies outlined in this chapter. It therefore seems to be that the masked probe paradigm is intuitively appealing, but that the evidence it provides of a fear expression bias is small and varied (Bocanegra and Zeelenberg, 2009; Carlson and Reinke, 2008). A recent review by Bannerman et al. (2010) emphasises the benefits associated with the saccadic latency paradigm, particularly when compared to continuous flash suppression. These two paradigms have become popular tools for measuring specific biases for fear expressions and are discussed in more detail in Part 2 of this chapter.

9.2 Part 2: CFS versus saccadic latency

In the last decade CFS has become an especially popular tool for measuring the degree to which certain visual stimuli reach the level of visual conscious recognition. The technique was introduced in 2004 by Tsuchiya and Koch (2004), and has been cited multiple times by studies ranging from the investigation of face perception (Gray et al., 2013; Hedger et al., 2015), unconscious perception of word stimuli (Heyman and Moors, 2014), and to unconscious processing mechanisms responsible for abstract mathematical computations (Sklar et al., 2012). Specifically, it has been used a small number of times to study perceptual biases for fear expressions and has shown that fearful faces break interocular suppression faster compared to other facial expressions (Gray et al., 2013; Jiang et al., 2007; Stein et al., 2014; Yang et al., 2007). To date, however, Stein et al. (2014) have conducted the only investigation concerned directly with the spatial frequency content of faces that might underpin the bias for fearful expressions. Interestingly, and counter to the currently accepted low-frequency account (introduced in Chapter 3), Stein et al. (2014) report that salient effects associated with fear expressions are determined by high spatial frequency components. The authors also report that fear expressions only break suppression faster when they are upright, compared to when they are inverted; another finding that contrasts with the low-level account, and findings from Gray et al. (2013) who show that under CFS fear expressions continue to break suppression rapidly regardless of stimulus inversion. In an attempt to understand why data from Stein et al. (2014) deviate from the broader literature, the growing uncertainties regarding the technique became apparent within the literature. An example is a review by Bannerman et al. (2010) who discuss features of the technique that are of important consideration if it is to be used as a measure of the threat bias. Recent methodological discussions also raise questions regarding its reliability in studies of visual psychophysics, primarily because of the way in which the response times generated under CFS are sensitive to changes in design features (Bannerman et al., 2010; Lunghi et al., 2016; Yang and Blake, 2012). The following sections detail the argument that CFS by definition cannot be an ecologically valid measure of implicit perceptual biases (Bannerman et al., 2010), and present data demonstrating the nuances associated with CFS experimental features.

9.2.1 Ecological validity

Perhaps the most prominent argument against the use of CFS as a measure of visual biases is the uncertainty over its ecological validity. This topic has been both discussed at length, and behaviourally investigated by Bannerman et al. (2010). Here, Bannerman et al. (2010) review practical and theoretical differences between the saccadic latency and CFS technique as measures of visual biases. Central to their discussion is the reliance of both measures on observers' response times, when the task is to visually locate a fearful face stimulus. The primary difference is that CFS requires observers to manually indicate the point at which they detect the face, and its associated location. By contrast, a saccadic latency technique bypasses the use of such manual response times, measuring saccadic latency responses instead; the time taken to initiate the first saccadic shift after stimulus onset. The difference between the use of manual response times and saccadic response times forms the crux of Bannerman et al. (2010)'s argument that SL is the more suitable and effective measure of perceptual biases. In an attentional cueing paradigm, Bannerman et al. (2010) measured response times for participants to orient attention towards an exogenously presented fearful expression. Face stimuli were presented for short (20ms) or long (100ms) durations and were followed by the appearance of a target (cross) at either a valid or invalid location to the face. In a saccade condition, observers were required only to 'look' toward target stimuli as soon as they appeared, compared to manual conditions where button presses were used to indicate the location of a face. Manual responses showed that fearful faces were better at maintaining visual attention (disengagement effect), but only when they were presented for longer durations (100ms). Saccade responses, however, showed that fearful faces were associated with both capturing and maintaining visual attention, but only when they were presented for brief durations (20ms). Bannerman et al. (2010) argue that at the reflexive and automatic level, the visual system is quick to orient toward a threatening stimulus within a visual scene (a fearful facial expression), and to allocate attentional resources accordingly. Such early and implicit behavioural responses are not represented by manual response times, including those recorded under CFS by Stein et al. (2014). Bannerman et al. (2010) argue that such responses are the result of consciously calculated decisions, and indeed Stein et al. (2014) acknowledge the focus of their study on stimulus detectability occurring at the conscious level, as opposed to unconscious and early aspects of perception. It is this explicit detection of stimuli under the CFS paradigm that is argued by Stein et al. (2014) to account for the high frequency specificity for the fearful face bias observed under CFS. Here, Stein et al. (2014) argue that high frequency fear expressions in their CFS study were detected faster because the process of explicit detection of facial expressions requires high frequency information in a similar way to explicit facial emotion recognition (Stein et al., 2014). This interpretation of findings from Stein et al. (2014) provides further support for the argument that the CFS technique measures very different aspects of visual processing compared to those measured by saccadic latency (Bannerman et al., 2010).

A primary advantage of saccadic latency, absent in the CFS paradigm, is its ecological validity and compatibility with the threat bias theory. The relative speed of saccades suggests that less information is required for saccadic initiation than for the execution of a decision-generated motor response (Bannerman et al., 2010; Hunt et al., 2010). Sensitive to low thresholds of visual input, saccades occur as early as 11-110 ms in response to a face stimulus, compared to manual response times to detect a face that vary between 150-300 ms following onset (Crouzet et al., 2010; Kirchner and Thorpe, 2006; Rousselet et al., 2003). It is also argued that information sampled early after stimulus onset is implicated in responses associated with signal detection, whereas information extracted at longer durations is necessary for detailed analysis. This is evidenced by a manual response advantage for emotion discrimination when face stimuli are presented for longer durations, where information sampled at 500ms stimulus durations is thought to be elaborate enough for discrimination decisions (Bannerman et al., 2009a). It therefore seems to be that saccadic and manual response modes related to two separate and functionally distinct response systems (Bannerman et al., 2010). Threat bias theory states that biases for fear expressions operate not via overt discrimination processes, but with rapid and likely non-conscious responses to low-level stimulus signals (Vuilleumier et al., 2003). It therefore follows that saccadic latency, as a reliable measure of reflexive oculomotor behaviours, is perhaps the most appropriate measure of the bias for fear expressions.

9.2.2 Individual differences across trials

Central to some of the debates regarding the reliability of CFS is the substantial variability across observer response times. Reaction times to detect target stimuli suppressed by dynamic masks differ significantly compared to those associated with both traditional binocular rivalry and detection tasks that do not involve intraocular suppression (Hedger et al., 2016). Reaction times also differ across experimental conditions, within participant performance across trials, and overall between participants (Gayet and Stein, 2017). Observers who exhibit overall slower reaction times across trials tend to show larger response time differences between conditions (Gayet and Stein, 2017), an effect that can be replicated by artificially lengthening the duration of a trial, or by reducing the contrast and thus visibility of a target stimulus. Such variability may be due to the frequency of observer spontaneous eyeblinks during the CFS task; behaviours that are, in themselves, little understood (Van Opstal et al., 2016). Spontaneous eyeblinks are thought to be closely related to conscious awareness, indicated by the temporal relationship observed between their frequency and breaks that occur during information processing. Van Opstal et al. (2016) found that although observers often refrain from blinking during a CFS trial, the occurrence of a spontaneous eyeblink is associated with a significant increase in observer response times. Such increases in response times may be interpreted as evidence of mask and/or stimulus strength, when they may in fact be irrelevant factors that require experimental control. This association between spontaneous eyeblink frequency and suppression duration may, in part, account for some of the variable suppression duration across trials (Van Opstal et al., 2016). A similar and equally notable issue is the extent to which, under experimental conditions, observers become aware of the nature of the CFS task. Whether stimuli are delivered using goggles or mirror stereoscopes, the process of closing one eye makes it possible to determine on which side of the visual field a stimulus is presented. During these tasks, participants are required to indicate the location of a stimulus, which can often be difficult to discern given the strength of the mask. Misunderstanding the purpose of the task, or due to simple demand characteristics, observers may therefore use the ability to close one eye in order to achieve better performance on trials. Such effects can undoubtedly influence reaction times, due to the removal of the suppression mask from view altogether. In a recent discussion regarding variable response times under CFS, Gayet and Stein (2017) propose that response data be normalised to extract such extraneous variance. Here, Gayet and Stein (2017) argue that it should be common practice to examine differences in reaction times between conditions, and whether these condition-related effects correlate with participants' average reaction time scores.

Saccadic latency responses are not an exception, and also demonstrate variable response times both within and between observers. Compared to those associated with CFS, however, variability associated with saccadic latency data appears to be better understood, naturally-driven, and minimised during experimental design. With the exception of subclinical populations, average latencies do vary between observers and also experimental trials (Quaia et al., 1999). An example is the 'gap effect'. Varying the temporal relationship between the central fixation point and the point of target appearance shows that a temporal gap between the former disappearing and the latter appearing is linked with shorter latencies. When the temporal gap is closed, such that the disappearance of the central fixation point and appearance of the target stimulus occur at the same time, saccadic latencies are longer (Quaia et al., 1999). Ensuring continuity of this temporal relationship across trials should control for any related differences in saccadic responses. It is also recognised that in order for an exogenously presented visual stimulus to influence saccade behaviours, its onset must occur at least 70ms before the saccade is initiated. For attentional cueing paradigms, this means that effects of a transient stimulus may not be captured outside of this time window. Though it is recognised that saccades are not entirely immune to transient stimuli appearing outside of this window, and do show slight deviations, or curvatures, when progressing toward the target stimulus (Quaia et al., 1999). Similarly, anticipatory saccades may occur if the location of a stimulus is expected, whereby a saccade is initiated as an anticipatory response to a target stimulus that has not yet onset. This may mean that saccade responses are less likely to reflect implicit attentional processes, instead indexing a degree of anticipation. Variability occurring because of this can also be avoided during experimental design by ensuring a central fixation point is calibrated for observer gaze before trail onset begins. It is also recognised that microsaccades are necessary perceptual behaviours that persist in the form of miniscule step-like tremor movements, occurring regardless of explicit instructions to remain fixated on a single point (Findlay and Walker, 2012). Despite variability in saccade behaviours according to these different factors, certain regularities are associated with such eye movements; so much so that the term main sequence has been used to describe these consistencies (Quaia et al., 1999). For example, shortest latencies of around 20ms are designated the smallest movements, while largest possible movement is associated with an absolute maximum of 100ms. A primary and evident benefit of saccadic latency measures is their relation to implicit behavioural responses, rather than manual, motor responses such as those associated with experimental paradigms, discussed above.

9.2.3 Mechanisms of control

Binocular rivalry techniques, including CFS, do not necessarily render stimuli completely invisible; rather stimuli are available to both eyes, but must compete for perceptual dominance such that one and not the other becomes the dominant percept. This contrasts with masked probe paradigms, where neural representation of visual stimuli is attenuated via rapid presentation durations. Here, a stimulus is not continuously available to the observer, as in CFS, but is presentated for rapid durations such that the degree of visual processing that can take place is limited. Because the CFS paradigm allows stimuli to remain available to observers, it is not clear whether they are still registered in early vision or attended to at the subconscious level. It seems to be that the use of CFS has foregone an understanding of the way in which suppressed stimuli are dealt with by the visual system. Such uncertainties include the locus of dominance competition; whether it is the eyes or the stimulus of interest that are suppressed under CFS (see Ritchie et al. (2013) for review), and the extent that stimuli suppressed by visual masks are processed via unconscious visual attentional processes or those that operate within conscious perception, and the degree to which CFS is a suitable measure for this (see Beck and Clark (1997); Gayet and Stein (2017); Kiefer and Brendel (2006). Similar discussions also debate the extent that visual processing under CFS differentially recruits input from dorsal and ventral processing pathways (see Ludwig and Hesselmann (2015); Lunghi et al. (2016). Recent discussions also introduce the notion that the mechanisms of visual processing under CFS may rely on working memory; whereby images of face expressions break suppression faster when congruent with emotionally-relevant information already held in working memory (for a review, see Gayet et al. (2014); Pan et al. (2014)). The cognitive and perceptual constructs accessed under CFS therefore remain relatively unclear, which in turn limits the ability to form a functional understanding of the way that visual stimuli are processed under CFS. Saccadic eye movements, by comparison, are well-understood perceptual behaviours. Findings have shown that they operate within early stages of visual processing, and are likely to respond to certain low-level image properties (discussed below). The association between such physical image characteristics and stimulus detection under CFS is less understood, where the literature so far suggests that responses under CFS are subject to characteristics of experimental masks and stimuli.

9.2.4 Relevance to the low-level approach

It is reasonable to question the extent to which each technique is a suitable measure to be used for a low-level investigation of the fearful face bias. Saccadic latencies, for example, appear not to be sensitive to the phase spectrum associated with facial stimuli: shorter latency periods–the time lapsing between stimulus onset and first orienting saccade- that are specific to face stimuli are preserved even when the configural content in faces is substantially removed (i.e. for phase scrambled versions of faces) (Honey et al., 2008). This finding suggests that face-specific eye movements for detecting a facial stimulus respond to the basic and coarse low-level features of the image. This supports the notion that saccadic latencies reflect responses to low level stimulus properties; the kind of processing occurring before higher-level visual representation is achieved (Honey et al., 2008). Bompas and Sumner

(2008) compared manual versus saccadic responses to stimuli of different luminance levels, showing that saccades were faster to respond to visual signal information, including that thought to determine image salience. Indeed, the accepted view appears to be that saccadic eye movements reflect naturalistic responses to the physically salient aspects of a visual scene, sampling information during early visual processing that is not available to or reflected by motor responses (Bannerman et al., 2010; Bompas and Sumner, 2008; Theeuwes et al., 1998). Bompas and Sumner (2008) also note the shared neutral correlates associated with both saccadic eye movements and models of rapid threat detection. Subcortical regions, including the pulvinar regions and supervisor colliculus are both implicated in rapid threat-relevant information processing, and the generation of saccadic eye movements (Bannerman et al., 2010; Bompas and Sumner, 2008). However, the suitability of saccadic responses for investigating the low-level approach may depend on the nature of the experimental design. For example, some attentional cueing paradigms measure the effect of fear expressions on attentional disengagement. Anti-saccade techniques require an observer to suppress an incomplete saccade towards a target stimulus in order to redirect this toward a stimulus located directly opposite. Failure to suppress a saccade toward a target stimulus is refers to as pro-saccade errors, where the average latency is around 250-350ms, but is associated with both large individual differences in error rates and practice effects, whereby the ability to suppress saccades increases linearly with trials (Blaukopf and DiGirolamo, 2005; Quaia et al., 1999). Although exogenous cueing paradigms are theoretically appealing, application of this design within a saccadic latency framework may not be suitable. Here, an advantage could be argued for the use of CFS instead, where observers are comparatively less likely to report false responses, as next trial onset can only be elicited by manual observer responses.

9.2.5 Selectivity of CFS mask strength

A primary issue exclusive to the CFS paradigm is the reliability of CFS masks to suppress a visual stimulus from conscious awareness. Under CFS, one eye is continuously presented with a highly salient masking stimulus. The purpose of a CFS mask is to present a salient (high in contrast) noise pattern whose physical contrast changes sequentially at a fixed rate (Tsuchiya and Koch, 2005). The dynamic feature of the mask allows it to be a novel stimulus, such that it can continually compete for perceptual dominance and therefore 'mask' a target image presented to the other eye from visual awareness. Although theoretically the suppressing mask is thought to be an effective way of attenuating visual consciousness of a stimulus, recent data suggests that the suppression strength of CSF masks varies according to the features of the target stimulus of interest. Under CFS, Yang and Blake (2012) measured contrast thresholds for face images composed of different spatial frequency ranges. Their CFS mask, or display, consisted of multiple grayscale rectangles normalised for luminance and RMS contrast, changing every 100 ms. Contrast thresholds were significantly higher when face stimuli were composed of low spatial information (0.75 cpd) than when the same faces were composed of higher frequency content (6 cpd). In other words, observers' ability to detect the location of a face under CFS is significantly worse if the face image is filtered to contain low ranges of spatial information (0.75 cpd). These findings suggest that the suppression strength of a given CFS mask is not necessarily a universal effect, rather that this strength is dependent on the spatial content of experimental face stimuli (Yang and Blake, 2012). If it is the case that low frequency face images are better suppressed by CFS masks compared to their high frequency versions, as data from Yang and Blake (2012) suggests, these data may account for the higher suppression durations associated with low frequency compared to high frequency faces observed by Stein et al. (2014). Yang and Blake (2012) show that this selectivity of mask strength is driven by its spectral properties; greater contrast energy at lower spatial frequencies contained within their mask display account for its propensity to suppress face stimuli composed of the same low frequency content. Furthermore, they show that when CFS masks are bandpass filtered to contain low frequency content, suppression strength is greatest for low frequency target stimuli, but that when masks are high bandpass filtered, the magnitude of this suppression effect is significantly smaller, such that high frequency CFS are comparatively weaker to supress high frequency images and even more so for low frequency images. These findings show that the strength of CFS masks depends on two factors: the spectral properties of the mask display itself, and those associated with the suppressed stimuli. It is therefore not possible to draw unambiguous conclusions regarding the effects of low-level image properties on the visibility of stimuli under CFS conditions (Lunghi et al., 2016; Yang and Blake, 2012). This selectivity of CFS mask suppression strength is upheld across other studies, highlighting the need for a better understanding of the way in which low-level image properties are determining factors of stimulus strength under CFS (Fahle, 1982; Levelt, 1965; Lunghi et al., 2016; Mueller and Blake, 1989). To account for this mask selectivity, Stein et al. (2014) (Experiment 4b) employ hybrid CFS masks composed of contrast energy equal at both low and high spatial ranges. The same advantage for high frequency fear faces to break suppression faster than low frequency faces was still found, suggesting that effects may not have been driven by the selectivity of mask strength. However, it is worth mentioning here that although Stein et al. (2014) replicate effects using hybrid CFS masks, face stimuli were still normalised for RMS contrast and luminance. Findings from Experiment 2 provide a detailed account of how this process of normalisation may artificially drive effects observed by Stein et al. (2014).

Findings from a similar investigation also show that the associated strength of masks are equally sensitive to, and dependent on, spatiotemporal properties of the stimulus: the rate at which the dynamic pattern sequences that compose the mask are updated (Lunghi et al., 2016). Lunghi et al. (2016) show that suppression strength, indexed by higher suppression duration, peaks when the temporal frequency of masks is 1Hz. Temporal frequencies either side of this peak show a marked decline in mask strength, where masks composed of temporal frequencies below .22 Hz and above 4 Hz are performed worse compared to static masks. Importantly, peak suppression for 1 Hz is true both when stimuli are composed of low (1 cpd) and high (10 cpd) spatial frequency content. However, suppression for high frequency masking stimuli with a temporal frequency of 1 Hz is markedly more effective than the same stimuli composed of high frequency information. For low frequency stimuli presented at 1 Hz, suppression strength was almost comparable to that of a static mask image. These findings show that the strength of a CFS mask varies significantly according to the design properties of the mask, including its temporal and spatial frequency features. It seems to be that 1 Hz is the ideal temporal frequency for a mask to produce maximum suppression effects, but that these effects are significantly stronger when visual stimuli are composed of high compared to low frequency information. Lunghi et al. (2016) argue that typical masks often composed of 10-12 Hz likely produce very weak suppression effects, and that this further demonstrates reason to doubt the extent that reliable conclusions can be drawn from CFS studies.

9.2.6 CFS yields inconsistent findings

Studies of CFS concerned with perceptual biases for fear expressions yield inconsistent findings, and the reason for this is unclear. A recent meta-analysis compared the strength of evidence for prioritised processing of visual threats across experiments using masked visual probe, binocular rivalry and continuous flash suppression techniques (Hedger et al., 2016). Their findings showed small and medium effects associated with masked visual probe and binocular rivalry designs (respectively), but that studies of continuous flash suppression tended to produce findings that, while significant, consisted of biases for positive over negative visual stimuli; findings that directly contrast with the central feature of threat bias theory. However, it is important to note here that this analysis included studies that employed the use of not only fearful faces, but other visual threats including dominant faces, negatively-valanced words, and threatening objects such as knives. It might therefore be that the presence of a CFS effect associated with fearful faces could have been weakened by the inclusion of broadly defined visual threats. However, relatively speaking, the same polarised bias effects were not found for masked visual probe and binocular rivalry paradigms. This suggests that CFS may characteristically yield different and inconsistent results compared to different experimental paradigms. Equally unclear is the extent that an advantage for a given face image to break suppression relies on its low-level image properties, or higher-level content. Yang et al. (2007) were among the first to observe a fear advantage under CFS, noting that because the effect persisted when the same faces were inverted, the bias must be driven by low level image content preserved within inverted face images. This inversion effect was also found by Gray et al. (2013). However, both Stein et al. (2014) and Alpers and Gerdes (2007) observe a fear expression advantage to break suppression, but one that only occurs when face images are upright and therefore configural content preserved. At the theoretical level, CFS possesses intuitive appeal that is reinforced by its propensity to render stimuli consciously invisible for long periods of time. However, findings clearly demonstrate the many subtleties surrounding the technique that must be considered. While it has been suggested that these shortcomings inherent to CFS may be alleviated by controlling for factors such as eyeblinking and demand characteristics, and the use of normalisation techniques to rid data of unnecessary variability, the consensus discussed in a review by Hedger et al. (2016) is that our understanding of the technique is limited, and provides insufficient confidence of its use to effectively measure the perceptual biases for threat-relevant stimuli (Hedger et al., 2016).

9.3 Premise for Experiments 4 and 5

9.3.1 Does contrast normalisation influence fear expressions' visibility under CFS?

Amidst the uncertainties surrounding the CFS paradigm, it is clear there are several experimental factors and features that may determine the outcome of results. This is two-fold, as it is evident that properties of both the CFS mask and the stimuli of interest can influence the outcome of CFS data. Experiments 1-3 of this thesis show that both natural differences in image properties and effects of contrast normalisation will influence the the perceived salience of a facial expression. In particular, findings from Experiment 1 show that fear expressions are in fact naturally lower in physical RMS contrast compared to other expressions; Experiment 2 shows that the perceived salience of fearful expressions is influenced by contrast normalisation, whereby the physical contrast of fear expressions is artificially boosted above and beyond that which is necessary to match images for apparent contrast; and Experiment 3 shows that the salience of fear expressions under CFS is likely driven by RMS normalisation of face images that bolsters the effective contrast of fear expressions in particular. It is therefore important to understand how each of these factors may individually influence fear biases under CFS; an argument similar to that of Gayet and Stein (2017) who highlight the importance of developing a better understanding of the specific experimental factors that influence the outcome of CFS response data. Stein et al. (2014) observe that a fear expression advantage under CFS relies on faces' high frequency content; a finding that directly contrasts with the currently accepted low-frequency mechanism outlined in Chapter 3. We might expect, given that face images used by Stein et al. (2014) were normalised for RMS contrast both before and after spatial filtering, that their findings are at least in part driven by contrast normalisation. Findings from Experiment 2 and 3 would suggest that the process by which their stimuli were normalised for RMS contrast would provide fear expressions with an artificial boost in perceived, or effective contrast under CFS. If this is the case, we might expect that RMS contrast normalisation of facial images boosts the visibility of high frequency filtered fear expressions under CFS. Experiment 4 therefore provides a measure of stimulus visibility for face images in two scenarios: when faces are normalised for physical RMS contrast, or when they are psychophysically matched for apparent contrast. These effects are also measured across spatial frequency conditions, matching the frequency cutoffs used by Stein et al. (2014).

9.3.2 Are responses under CFS and SL equally influenced by effects of contrast normalisation?

Findings from Bannerman et al. (2010) strongly suggest that compared to CFS, saccadic latency is a more effective and reliable tool for observing implicit and reflectively-driven perceptual biases for fearful facial expressions. This is primarily because of the automatic oculomotor responses captured using saccadic latency that do not require any motor-driven responses from observers (Bannerman et al., 2010). It is also clear that the CFS paradigm engenders biases for suppressing certain types of stimulus information compared to others. However, it remains unclear to what extent this is also the case within the context of saccadic latency. Data from Bannerman et al. (2012) shows that saccades are selectively executed to low frequency ($< .08_{cpd}$) fear expressions compared to when the same faces are composed of broadband frequencies, or filtered to contain high frequency information ($> 3.3_{cpd}$). These effects were observed by Bannerman et al. (2012) even when face stimuli had *not* been normalised for RMS contrast (on account of contrast normalising reducing the perceived visibility of low frequency images). However, in the study conducted by Bannerman et al. (2012) no comparison was made for face images that were normalised for RMS contrast, or apparent contrast. Face stimuli also only consisted of neutral, fear and happy expressions. Experiment 5 therefore conducts a replication of this study that is extended to include a condition for expressions of anger, and face stimuli presented either in their raw form, such that they are not normalised for contrast; when they are psychophysically normalised for apparent contrast; and when the are physically matched for RMS contrast.

9.3.3 CFS VS SL: which is the most appropriate measure for the fear bias?

Bannerman et al. (2010) show that saccadic latency responses are naturalistic measures of crude perceptual biases for fear expressions compared to motor response times generated under the CFS technique. Here, we address the same question, but also include a measure of the extent to which image properties of faces that are differently influenced by contrast normalisation can be expected to generate different response data between CFS and saccadic latency techniques. We therefore ask: do spatially filtered faces normalised for RMS contrast break suppression under CFS faster compared to those normalised for perceived contrast? And is this pattern of effects also true for response data generated from saccadic latency techniques?

9.4 Summary of Chapter 9

This chapter provides an overview of findings that show distinct variability in the data generated by CFS studies. This appears to be driven, in part, by image properties of target stimuli. Data from Experiments 1-3 show that perceptual biases for certain face expressions are significantly influenced by whether or not these images have been normalised for physical contrast. Here, we introduce the question of whether this is one of the little-understood factors that may drive differences in CFS response data. Bannerman et al. (2010) show that compared to CFS, saccadic latency is a reliable and naturalistic tool for investigating perceptual biases. We address whether this technique is also influenced by stimulus properties related to contrast normalisation, and if so, the extent that saccadic latency and CFS responses are similarly affected by such stimulus differences.

Chapter 10

Experiment 4: Response times to detect facial expressions under CFS: effects of physical and perceived contrast, and spatial frequency content

ASIWYFA. (2009). ASIWYFA. Artwork copyright to Smalltown America (2009).

10.1 Introduction

Experiment 4 investigates the visibility -or the time taken to break suppression- of 5 facial expressions that have been filtered to contain broad, low, mid-range, or high spatial frequency content. Stein et al. (2014) show that the advantage for fear expressions to break suppression faster compared to neutral expressions is driven by their high frequency content; an effect that does not support the generally-accepted low frequency approach to the fear bias (Bannerman et al., 2012; Vuilleumier et al., 2003). Notably, stimuli used by Stein et al. (2014) had been normalised for RMS contrast. Findings from Experiment 2 (and indeed, 3), suggest that this use of contrast normalisation is likely to have artifically boosted perceived salience of fearful expressions. Here, we address whether responses to spatially filtered face expressions, including fearful faces, vary under CFS according to whether these faces have been normalised for physical (RMS) contrast, or normalised psychophysically for their apparent contrast, such that physical differences between faces' RMS content reflects the degree that is necessary for each face to perceptually match a reference stimulus (composed of 10% Michelson contrast). In addition, Stein et al. (2014) used mask stimuli that were a hybrid of low and high frequency content. This makes it difficult to disentangle the origin of the masking stimuli, since it is possible for example that high frequency content in the target faces might have been masked by low frequency components of the dynamic mask. In order to directly match the low level attributes of our target and mask, we filtered both stimuli in the same way, so that both contained information in the same frequency range.

10.2 Methods

10.2.1 Participants

Data collection took place across 4 sessions, corresponding to each spatial frequency condition. Twenty-nine observers took part in the broadband condition, and 17 observers for the low, mid-range, and high frequency conditions. All participated in the experiment as part of a credited research module assessment. All participants had normal or corrected to normal vision.

10.2.2 Stimuli and apparatus

Stimuli were presented using a VIEWPIXX 3D monitor, viewed from a distance of 80cm. For the duration of the experiment, participants used a chin rest to maintain this viewing distance. The monitor screen was 52cm wide and 29cm tall. The screen resolution was 1920x1080 pixels, with a refresh rate of 120Hz and an average luminance of 50 cdm^2. Each pixel subtended 1.0 arc min. Stimuli were presented at 10-bit resolution. Stimuli were generated and presented using MATLAB and the Psychophysics Toolbox extensions (Brainard and Vision, 1997; Kleiner et al., 2007; Pelli, 1997).

Face stimuli

Stimuli were grayscale front-view face pictures of 16 individuals (8 male, 8 female) extracted from the Karolinska Directed Emotional Faces set (Lundqvist et al., 1998). Faces were cropped to included internal features only. Each individual portrayed 1 of 5 expressions: fear, anger, happiness, disgust or neutral. The size of each face image was 4.5°(face-width). Five expressions were used in order to provide a wider understanding of how differences might occur between expressions more broadly, as opposed to how fearful expressions differ only from neutral or a single positively-valenced face. Broadband face images were filtered using a second-order Butterworth filter in MATLAB. This included low (LSF), mid-range (MSF) and high spatial frequency (HSF) versions of all faces, and broadband face images were included such that there were 4 frequency conditions in total. The cut-off frequencies were $f < 1_{cpd}$

for LSF faces , $1 < f < 6_{cpd}$ for MSF faces, and $f > 6_{cpd}$ for HSF faces. Spatial frequency for face images therefore varied between 4.5 and 27 cycles per face-width. Low and high bandpass cut-offs were consistent with those used by Stein et al. (2014) and Vlamings et al. (2009). Two contrast conditions were used: one where all faces were normalised for physical root mean squared (RMS) contrast, and another where the same faces were normalised for perceived contrast, such that each possessed the physical contrast necessary for them to appear the same. RMS contrast represented the standard deviation of normalised pixel intensities, indexing variation in luminance values in face images (Peli, 1990), defined as:

$$C_{RMS} = \sqrt{\frac{1}{MN}\sum_{i=1}^{M}\sum_{j=1}^{N}(I_{ij} - \bar{I})^2} \tag{10.1}$$

Perceived-contrast-matched face stimuli were generated from behavioural data generated from Experiment 2 (Chapter 6). To achieve this, each raw face image was assigned the associated Michelson contrast level required for it to appear as equal in contrast to the 10% Michelson contrast reference. Apparent contrast values for face stimuli, generated by data from Experiment 2, are displayed in Figure 10.2. It is noteworthy to include here that because behavioural data from Experiment 2 were used to normalise face images for apparent contrast in the present study, and because Experiment 2 and the present study did not share the same participant pool, the resulting apparent contrast values may only have embodied idiosyncrasies in perceived contrast for those individuals who took part in Experiment 2. The advantage of this is that the same stimuli were used for Experiments 2 and 4; the disadvantage is that, conversely, contrast normalisation of faces' perceived contrast was not tailored to each individual who took part in Experiment 4. However, previous work has shown good agreement between participants in contrast matching for complex stimuli (O'Hare and Hibbard, 2011).

All stimuli were presented in their normal, upright form, and in manipulated form, where faces were rotated by 180° and the luminance polarity inverted. Using these techniques to manipulate stimuli reduces the visibility of configural information for successful recognition of facial expressions, but preserves images' low-level properties including contrast and spatial frequency content (Gray et al., 2013; Tanaka and Farah, 2003).

Mask stimuli

The same second order Butterworth filters used to filter face images were also applied to CFS masks composed of randomly positioned rectangles. This meant that the spatial frequency content of the mask was matched to that of the face image for each trial. The CFS mask display flashed at a rate of 10 Hz, and was present for the duration of each trial. An example is shown in Figure 10.3.

10.2.3 Procedure

The University of Essex University Ethics Committee approved the study and all participants gave written, informed consent after they were informed that the study was concerned with

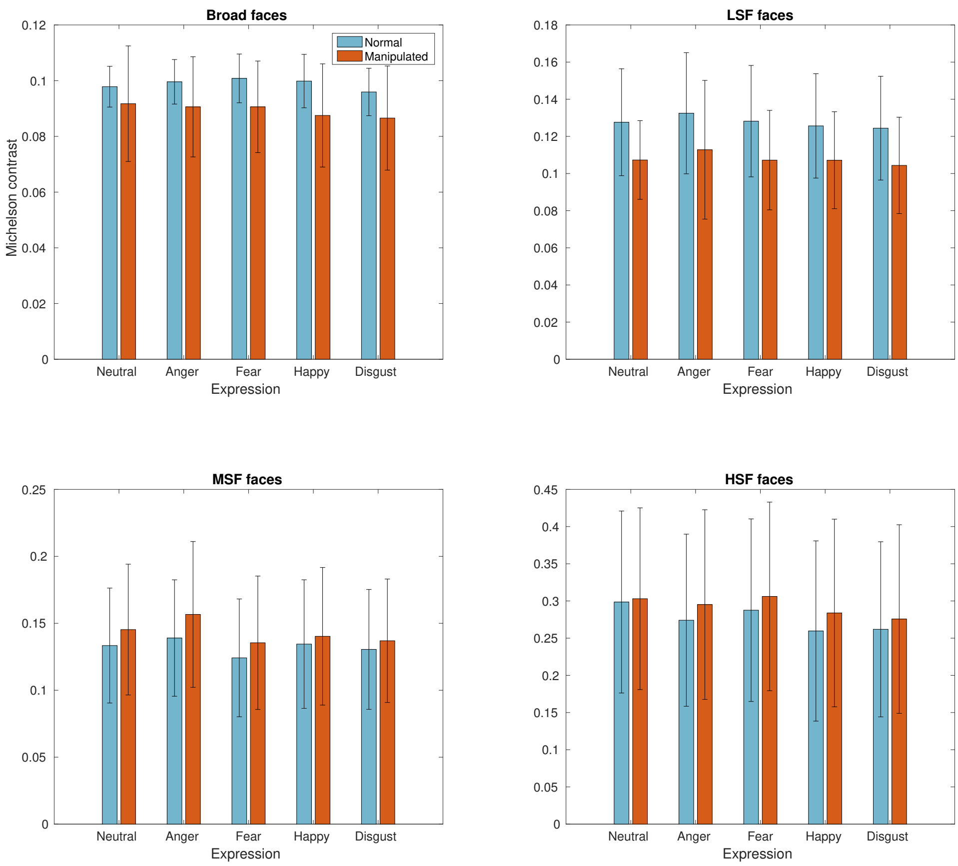

Fig. 10.2 Apparent contrast-matched faces -the degree of Michelson contrast required for images to appear the same as the 10% Michelson contrast reference - for face expressions composed of broad, low, mid-range, and high spatial frequency information. Error bars represent associated standard deviations.

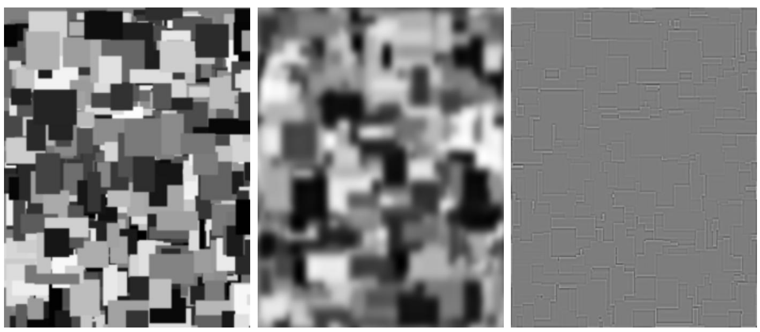

Fig. 10.3 From left to right: a broadband CFS mask stimulus filtered to contain low or high spatial frequency information ($f < 1_{cpd}$, and $f > 6_{cpd}$, respectively)

face perception.

Participants were tested individually in a quiet room, and wore NVIDIA 3D vision liquid-crystal shutter goggles for the duration of the study. Face stimuli were presented against a uniform background and located in 1 of 4 quadrants on the monitor screen. The 3D goggles delivered different images to the two eyes, simultaneously. One image was the face stimulus, the other a dynamic and high contrast noise-pattern, or mask. Face stimuli reached full Michelson contrast at 1 second after stimulus onset. As part of a four-alternative forced choice procedure observers responded by indicating in which 1 of 4 quadrants the face was located. Manual responses were recorded using the RESPONSEPixx response box, so that response times reflect the point at which stimuli emerged from intraocular suppression. Next trial onset was triggered by observers' motor response using the response box. After 7 seconds, absence of a button press also elicited onset of the next trial. These procedural parameters were consistent across the 4 sessions of data collection, with the only deviation being the structure and number of trials. These were as follows:

Session 1 (Broadband condition)

Each of the 29 observers completed 320 trials. This included 16 actors x 5 expressions (neutral, fearful, happy, angry, disgusted) x 2 manipulations (normal vs. inverted and reversed polarity) x 2 contrast matches (RMS vs. psychophysically matched). Trials were dispersed into 8 blocks, such that each block contained 40 trials.

Session 2 (LSF condition)

Each of the 17 observers completed 320 trials. This included 16 actors x 5 expressions (neutral, fearful, happy, angry, disgusted) x 2 manipulations (normal vs. inverted and reversed polarity) x 2 contrast matches (RMS vs. psychophysically matched). Trials were dispersed into 8 blocks, such that each block contained 40 trials.

Session 3 (MSF condition)

Each of the 17 observers completed 320 trials. This included 16 actors x 5 expressions (neutral, fearful, happy, angry, disgusted) x 2 manipulations (normal vs. inverted and reversed polarity) x 2 contrast matches (RMS vs. psychophysically matched). Trials were dispersed into 8 blocks, such that each block contained 40 trials.

Session 4 (HSF condition)

Each of the 17 observers completed 320 trials. This included 16 actors x 5 expressions (neutral, fearful, happy, angry, disgusted) x 2 manipulations (normal vs. inverted and reversed polarity) x 2 contrast matches (RMS vs. psychophysically matched). Trials were dispersed into 8 blocks, such that each block contained 40 trials.

10.3 Results

Suppression durations, or observer response times, for faces at each frequency condition (broad-, low-, mid-, and high) were analysed using a 5 (Expression) x 2 (Contrast) x 2 (Manipulation) repeated measures ANOVA. Suppression durations for each frequency condition are presented (at the end of this section) in Figure 10.4 , and are separated into a left and right column according to the RMS and apparent contrast metrics, respectively.

10.3.1 Response times for unfiltered, broadband frequency faces

Response times for broadband face stimuli when matched for either RMS or apparent contrast are plotted in Figure 10.4 (a) and (e), respectively. The results of the ANOVA are summarised in Table 10.1 . No significant expression-by-manipulation-by-contrast interaction was observed. A non-significant effect of contrast metric showed that suppression duration of face stimuli does not vary according to whether faces are physically (RMS) or perceptually (apparent) matched for contrast. A significant effect of manipulation shows that manipulated faces (inverted by 180° with retained luminance polarity) break suppression faster than normally presented faces. This finding is consistent with previous literature that suggests luminance polarity inversion increases images' perceived contrast, or salience (Haun and Peli, 2013). A significant effect of expression showed that broadband face expressions differ in terms of their associated suppression duration. A non-significant contrast-by-expression interaction shows that these differences are not differently influenced by contrast metric. Findings from this ANOVA are summarised in Table 10.1, at the end of this section. Sidak-corrected comparisons explored differences in average response times to detect expressions of fear compared to other faces. Comparisons included data for manipulated versions of faces (α= 0.0063, corrected according to 8 comparisons). These response time comparisons were performed independently for each contrast condition. Findings showed that when normalised for RMS contrast, broadband fearful expressions break suppression significantly faster compared to angry faces, where this effect is preserved when faces are manipulated. No other significant differences were observed. Alternatively, this is also true when the same face expressions are normalised psychophysically for apparent contrast, where broadband fear expressions break suppression significantly faster compared to angry expressions; however this effect is not preserved when faces are manipulated. These data are summarised in Table 10.1.

10.3.2 Response times for low frequency filtered faces

Response times for low frequency filtered facial stimuli ($f < 1_{cpd}$), when matched for either RMS or apparent contrast, are plotted in Figure 10.4 (b) and (f), respectively. The results of the ANOVA are summarised in Table 10.2. No significant expression-by-manipulation-by-contrast interaction was observed. A non-significant effect of stimulus contrast showed that suppression durations, or visibility of faces, do not vary between the contrast metric used to normalise images (physical RMS, or apparent contrast). An effect of manipulation showed that manipulated versions of low frequency faces are faster to break suppression and are therefore more visible compared to upright counterparts. This does not vary as a function of the contrast metric used for normalisation. A significant effect of expression shows that low frequency filtered face stimuli differ relative to the emotional expression of the face. A non significant contrast-by-expression interaction shows that differences in stimulus visibility between expressions is not determined by whether images are normalised for RMS compared to apparent contrast. Sidak-corrected comparisons explored differences in average response times to detect expressions of fear compared to other faces. Comparisons included data for manipulated versions of faces (α= 0.0063, corrected according to 8 comparisons).

These response time comparisons were performed independently for each contrast condition. Findings showed that regardless of contrast metric used to normalise faces, low frequency filtered fearful expressions do not differ in terms of response time compared to any other face expression. This is also true when faces are manipulated.These data are summarised in Table 10.2.

10.3.3 Response times for mid frequency filtered faces

Response times for mid-range frequency filtered facial stimuli ($1 < f < 6_{cpd}$), when matched for either RMS or apparent contrast, are plotted in Figure 10.4 (c) and (g), respectively. The results of the ANOVA are summarised in Table 10.3. No significant expression-by-manipulation-by-contrast interaction was observed. A significant effect of stimulus contrast showed that faces break suppression faster when they are normalised for RMS rather than apparent contrast , denoted by lower suppression durations associated for RMS normalised stimuli. An effect of manipulation showed that overall, manipulated versions of mid-range frequency faces break suppression faster and are therefore more visible compared to upright counterparts. A significant manipulation-by-contrast interaction also shows that manipulated versions of faces also break suppression faster when they are normalised for RMS contrast, and are slower to break suppression when psychophysically matched. A significant effect of facial expression shows that mid-range frequency filtered face stimuli differ according to the emotional expression of the face, and a significant contrast-by-expression interaction demonstrates that this differences vary as a function of the contrast metric used to normalise face images. Sidak-corrected comparisons explored differences in average response times to detect expressions of fear compared to other faces (α= 0.0063, corrected according to 8 comparisons). These response time comparisons were performed independently for each contrast condition. Findings showed that when mid-range facial stimuli are normalised for RMS contrast, fearful facial expressions break suppression faster compared to angry faces. This effect was not preserved for manipulated faces, and no other significant differences were found. Alternatively, when mid-range faces are normalised psychophysically for apparent contrast, suppression durations only differ between faces when they are manipulated. Here, happy faces break suppression faster compared to fearful expressions. No other significant differences were observed. These data are summarised in Table 10.3.

10.3.4 Response times for high frequency filtered faces

Response times for high frequency filtered facial stimuli ($f > 6_{cpd}$), when matched for either RMS or apparent contrast, are plotted in Figure 10.4 (d) and (h), respectively. The results of the ANOVA are summarised in Table 10.4. A significant expression-by-manipulation-by-contrast interaction shows that the effects of the 3 factors on suppression durations and therefore stimulus visibility are not independent of each other. A significant effect of stimulus contrast showed that faces break suppression faster when they are normalised for RMS contrast rather than apparent contrast. A significant expression-by-contrast effect also shows that the effects of contrast normalisation on the salience of facial expressions varies as a function of the contrast metric used to normalise images. A significant expression-by-manipulation interaction shows that visibility of facial expressions does vary as a function of the degree to which they are manipulated. A non-significant effect of manipulation, however, shows that unlike findings for the 3 preceding frequency conditions, the visibility of manipulated and normally presented face images do not differ overall.

Sidak-corrected comparisons explored differences in average response times to detect expressions of fear compared to other faces. Comparisons included data for manipulated versions of faces (α= 0.0063, corrected according to 8 comparisons). These response time comparisons were performed independently for each contrast condition. Findings showed that when high frequency faces are normalised for RMS contrast, fearful faces break suppression faster than both angry and happy faces. This differences between fear and anger was preserved when faces were manipulated. Alternatively, when the same high frequency faces were normalised psychophysically for perceived contrast, manipulated fear expressions were slower to break suppression compared to manipulated neutral faces. No other significant differences were observed. These data are summarised in Table 10.4.

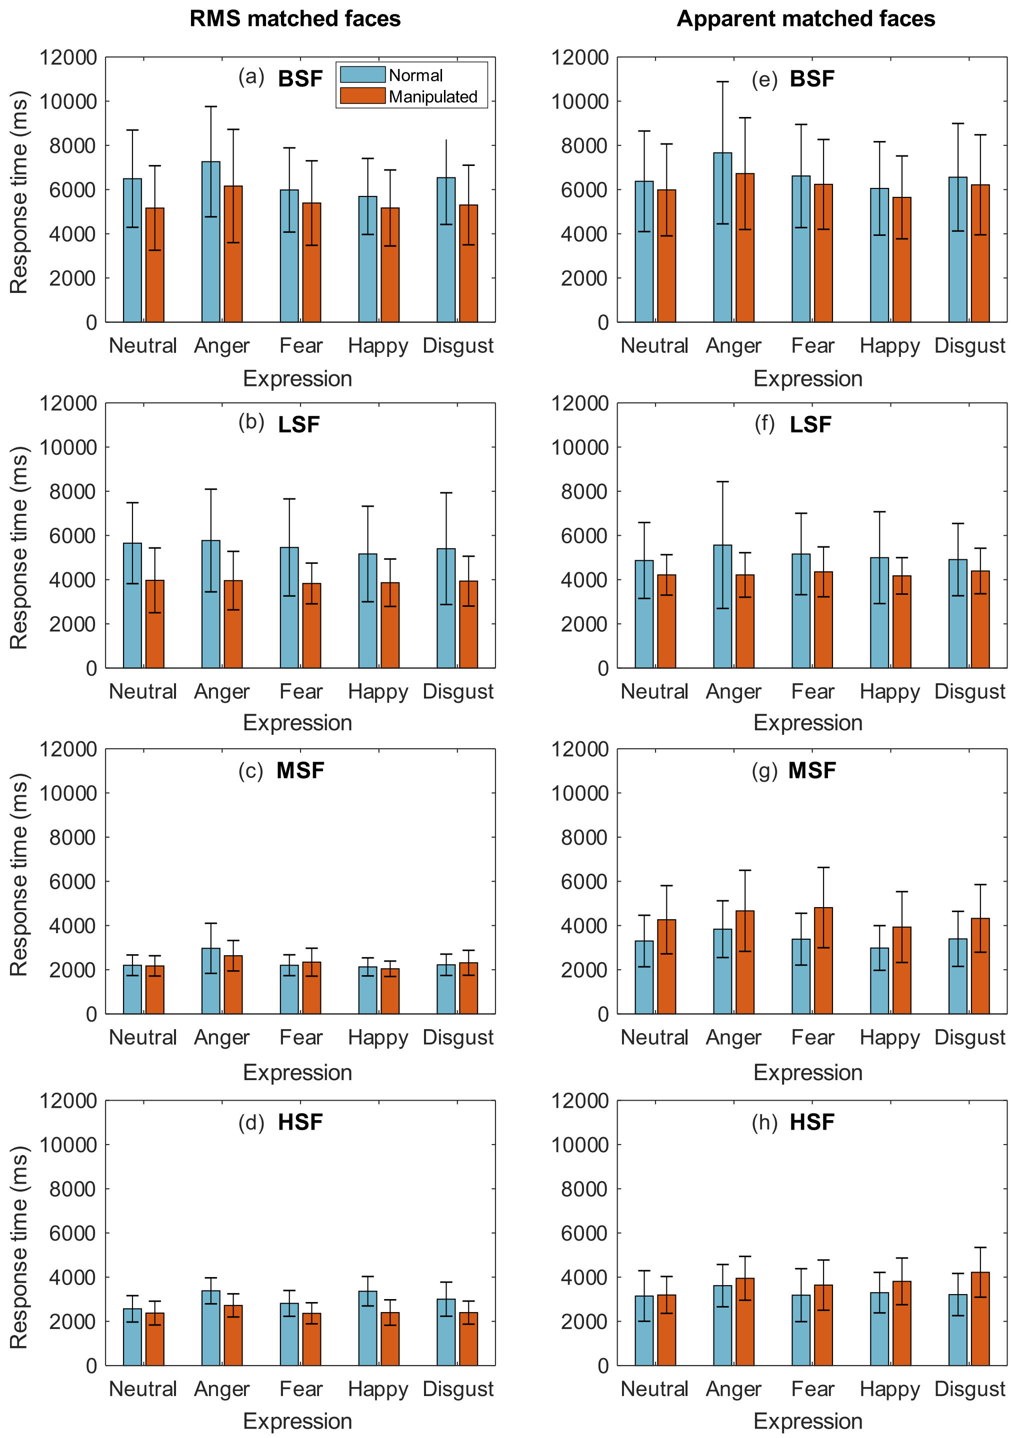

Fig. 10.4 Average response times to detect spatially filtered face images when normalised for RMS or apparent (perceived) contrast. Figures (a)-(d) show response times for 5 expressions when RMS normalised when faces are composed of broadband (a), low (b), mid-range (c), or high (d) spatial frequency content. Figures (e)-(h) show response times for the same face images when psychophysically matched, composed of broadband (e), low (f), mid-range (g), or high (h) spatial frequency content. Error bars represent associated standard deviations.

10.4 Data tables

Table 10.1 Response times to detect broadband faces, denoting the associated visibility of a stimulus. Data include an overall ANOVA for all broadband frequency stimuli, and Sidak-corrected comparisons between normal and manipulated fearful faces and their face counterparts (α= 0.0063, corrected according to 8 comparisons). Sub-tables show analyses for both contrast conditions.

Broadband faces (ANOVA)

	df	F	*p*	ηp_2
Contrast metric	1, 28	2.10	.15	.07
Expression	4, 112	22.59	<.001	.44
Manipulation	1, 28	22.10	<.001	.44
Contrast*expression	4, 112	.80	.52	.02
Manipulation*expression	4, 112	2.18	.07	.07
Contrast*Manipulation	1, 28	4.74	.03	.14
3 way	4, 112	1.62	.17	.05

Expression comparisons (RMS)

	t	df	CI	*p*
Fear-neutral	-2.45	28	-933.12, -84.68	.02
Fear-anger	-5.59	28	-1748.20, -811.27	<.001
Fear-happy	1.70	28	-59.44, 646.64	.10
Fear-disgust	-2.27	28	-1053, -56.20	.03
Manipulated versions:	t	df	CI	*p*
Fear-neutral	.91	28	-282.66, 736.11	.37
Fear-anger	-2.98	28	-1298.65, -240.71	.006
Fear-happy	.92	28	-272.62, 719.46	.36
Fear-disgust	.47	28	-296.50, 476.39	.63

Expression comparisons (Apparent)

	t	df	CI	*p*
Fear-neutral	.97	28	-267.30, 748.63	.34
Fear-anger	-3.19	28	-1724.45, -377.41	.003
Fear-happy	1.80	28	-76.80, 1204.39	.08
Fear-disgust	.18	28	-572.64, 684.56	.85
Manipulated versions:	t	df	CI	*p*
Fear-neutral	1.07	28	-227.53, 728.68	.29
Fear-anger	-1.69	28	-1073.34, 100.64	.10
Fear-happy	2.54	28	115.63, 1063.52	.01
Fear-disgust	.09	28	-414.32, 454.26	.92

Table 10.2 Response times to detect low frequency filtered faces, denoting the associated visibility of a stimulus. Data include an overall ANOVA for all low frequency stimuli, and Sidak-corrected comparisons between normal and manipulated fearful faces and their face counterparts (α= 0.0063, corrected according to 8 comparisons). Sub-tables show analyses for both contrast conditions.

LSF faces (ANOVA)

	df	F	*p*	ηp_2
Contrast metric	1, 17	.005	.94	.00
Expression	4, 68	3.52	.01	.17
Manipulation	1, 17	18.06	.001	.51
Contrast*expression	4, 68	1.16	.33	.06
Manipulation*expression	4, 68	1.47	.21	.08
Contrast*Manipulation	1, 17	12.50	.003	.42
3 way	4, 68	.80	.52	.04

Expression comparisons (RMS)

	t	df	CI	*p*
Fear-neutral	-.83	17	-680.02, 293.68	.41
Fear-anger	-1.17	17	-880.67, 251.04	.25
Fear-happy	2.73	17	67.62, 523.57	.01
Fear-disgust	.25	17	-396.67, 507.55	.79
Manipulated versions:	t	df	CI	*p*
Fear-neutral	-.92	17	-461.08, 179.83	.36
Fear-anger	-1.06	17	-374.90, 122.81	.30
Fear-happy	-.20	17	-374.81, 308.84	.84
Fear-disgust	-1.34	17	-268.08, 59.74	.19

Expression comparisons (Apparent)

	t	df	CI	*p*
Fear-neutral	1.38	17	-153.44, 741.18	.18
Fear-anger	-1.02	17	-1256.64, 447.38	.33
Fear-happy	.78	17	-276.66, 605.13	.44
Fear-disgust	1.59	17	-82.50, 587.59	.13
Manipulated versions:	t	df	CI	*p*
Fear-neutral	1.28	17	-90.02, 369.18	.21
Fear-anger	1.16	17	-113.88, 394.90	.26
Fear-happy	1.27	17	-118.23, 481.88	.21
Fear-disgust	-.29	17	-311.23, 235.08	.77

Table 10.3 Response times to detect mid-range frequency filtered faces, denoting the associated visibility of a stimulus. Data include an overall ANOVA for all mid-range frequency stimuli, and Sidak-corrected comparisons between normal and manipulated fearful faces and their face counterparts (α= 0.0063, corrected according to 8 comparisons). Sub-tables show analyses for both contrast conditions.

MSF faces (ANOVA)

	df	F	*p*	ηp_2
Contrast metric	1, 16	49.81	<.001	.75
Expression	4, 64	15.52	<.001	.49
Manipulation	1, 16	6.28	.02	.28
Contrast*expression	4, 64	2.74	.03	.15
Manipulation*expression	4, 64	3.70	.009	.18
Contrast*Manipulation	1, 16	7.44	.01	.31
3 way	4, 64	.75	.56	.04

Expression comparisons (RMS)

	t	df	CI	*p*
Fear-neutral	-.003	16	-198.17, 197.68	.99
Fear-anger	-3.29	16	-1259.86, -273.95	.005
Fear-happy	.93	16	-94.84, 244.35	.36
Fear-disgust	-.14	16	-307.50, 268.65	.88
Manipulated versions:	t	df	CI	*p*
Fear-neutral	2.57	16	30.11, 308.49	.02
Fear-anger	-2.21	16	-567.99, -12.89	.04
Fear-happy	2.12	16	.52, 596.90	.05
Fear-disgust	.33	16	-162.68, 222.97	.74

Expression comparisons (Apparent)

	t	df	CI	*p*
Fear-neutral	.48	16	-277.19, 442.14	.63
Fear-anger	-2.79	16	-796.04, -109.47	.01
Fear-happy	2.90	16	107.34, 690.07	.01
Fear-disgust	-.07	16	-410.74, 383.78	.94
Manipulated versions:	t	df	CI	*p*
Fear-neutral	2.96	16	157.16, 943.08	.009
Fear-anger	.79	16	-243.47, 536.11	.43
Fear-happy	4.61	16	476.03, 1284.62	<.001
Fear-disgust	2.68	16	101.94, 872.93	.01

Table 10.4 Response times to detect high frequency filtered faces, denoting the associated visibility of a stimulus. Data include an overall ANOVA for all high frequency stimuli, and Sidak-corrected comparisons between normal and manipulated fearful faces and their face counterparts (α= 0.0063, corrected according to 8 comparisons). Sub-tables show analyses for both contrast conditions.

HSF faces (ANOVA)

	df	F	*p*	ηp_2
Contrast metric	1, 16	36.07	<.001	.69
Expression	4, 64	22.92	<.001	.58
Manipulation	1, 16	1.28	.27	.07
Contrast*expression	4, 64	2.93	.02	.15
Manipulation*expression	4, 64	3.91	.007	.19
Contrast*Manipulation	1, 16	48.12	<.001	.75
3 way	4, 64	13.94	<.001	.46

Expression comparisons (RMS)

	t	df	CI	*p*
Fear-neutral	2.25	16	11.65, 480.13	.04
Fear-anger	-7.71	16	-725.77, -412.94	<.001
Fear-happy	-5.40	16	-766.13, -334.60	<.001
Fear-disgust	-1.95	16	-397.50, 16.13	.06
Manipulated versions:	t	df	CI	*p*
Fear-neutral	-.12	16	-207.67, 184.51	.90
Fear-anger	-4.66	16	-520.24, -194.94	<.001
Fear-happy	-.33	16	-251.02, 183.38	.74
Fear-disgust	-.81	16	-116.49, 51.66	.42

Expression comparisons (Apparent)

	t	df	CI	*p*
Fear-neutral	.28	16	-249.76, 326.35	.78
Fear-anger	-3.13	16	-720.03, -138.66	.006
Fear-happy	-.74	16	-432.89, 207.52	.46
Fear-disgust	-.227	16	-283.04, 228.26	.82
Manipulated versions:	t	df	CI	*p*
Fear-neutral	3.84	16	119.61, 690.45	.001
Fear-anger	-2.53	16	-568.04, -50.45	.02
Fear-happy	-1.25	16	-455.39, 117.53	.229
Fear-disgust	-2.94	16	-997.56, -163.21	.009

10.5 Discussion

A widely accepted view in the psychophysical literature is that perceptual biases for fearful face expressions are driven by the low spatial frequency properties contained within these faces (Bannerman et al., 2012; Bocanegra and Zeelenberg, 2009; Vuilleumier et al., 2003). However, a CFS study conducted by Stein et al. (2014) showed that the advantage for fearful facial expressions to break suppression quicker than neutral faces relies on the high frequency components of faces. Stimuli used by Stein et al. (2014) was, however, first normalised for RMS contrast. Here, we conduct a CFS study using the same frequency cut-offs as those used by Stein et al. (2014). Two contrast metric conditions were used. In one, all face images were normalised for their physical RMS content. In the other, face images are normalised psychophysically, such that they contain the Michelson contrast necessary for each expression to appear subjectively the same as a 10% reference image (data from Experiment 2). Findings showed that, similar to data reported by Stein et al. (2014), a threat bias for fear expressions is observed when fear expressions are filtered to contain high frequency content and normalised for RMS contrast. High frequency fear expressions were faster to break suppression compared to angry and happy expressions. Stein et al. (2014) reports that their observed threat bias effect was preserved under conditions of manipulation, i.e. when their face images were inverted by 180°. This same effect was also observed for broadband facial stimuli. Stein et al. (2014) also observed that overall, upright faces were faster to break suppression, and therefore more salient, compared to their inverted counterparts. In the present study, the threat bias observed for high frequency fear expressions was preserved under conditions of manipulation, but only between angry and fear expressions. Findings overall also showed that manipulated faces normalised for RMS contrast were generally faster to break suppression and therefore more salient compared to their upright, normal counterparts. This finding is not consistent with those reported by Stein et al. (2014) and Yang et al. (2007). It is important to note here, however, that differences of this effect between existing data and the present findings might be driven by the fact that manipulated images used in this study were not only inverted by 180°, but they were also subjected to inverted luminance polarity. This meant that manipulated versions of faces were inverted, but also had a negated effect. Findings from Haun and Peli (2013) show that judgements of an image's brightness are more informed by the darker rather than the lighter regions of an image; thought to occur due to the density of contrast-encoding neurons innervated within primary visual cortex areas (Haun and Peli, 2013). It may therefore be that under CFS, normally presented, upright face images do break suppression faster compared to inverted versions (Stein et al., 2014; Yang et al., 2007), but that reversing the luminance polarity of these faces under the same experimental conditions boosts the salience, or visibility of these faces. In any case, the present findings do show a fear advantage to break suppression under CFS when faces are normalised for RMS contrast. When normalised for apparent, perceived contrast, a small fear advantage was observed, but only for broadband faces between fear and angry expressions. At other levels of spatial frequency filtering, fear expressions perceptually matched for contrast either do not differ compared to other expressions, or have a *disadvantage* when overcoming suppression. These findings suggest that a threat bias effect whereby fear expressions break suppression faster than their expression counterparts is in part reliant on facial stimuli having been normalised for physical RMS contrast. Data from Experiments 1 and 2 suggest that the process by which these expressions are normalised for RMS contrast may provide fear expressions in particular with an artificial boost in perceived image contrast, such that we might expect them to have been granted an advantage in breaking suppression under CFS. This interpretation is further evidenced by the diminished threat bias effect observed for face stimuli normalised psychophysically for apparent contrast: these fear expressions had not been granted the boost in physical contrast necessary to enhance their perceived salience, and instead assigned the amount required to perceptually match them to other faces. This may account for why, overall, perceptually matched fear expressions did not appear significantly more salient than other expressions. The present findings do, however, suggest an interesting difference in the visibility between fear and angry expressions that is almost consistently preserved even when the configural content in these faces is disrupted via stimulus manipulation, or indeed one that is not influenced by the contrast metric used to normalise broadband face images. This is returned to in the following section.

Overall, in terms of the present findings in relation to those observed by Stein et al. (2014), it is visible that threat bias effects associated for high frequency fear expressions are driven, at least in part, by the process of physical contrast normalisation. This allows the higher frequency versions of fear expressions to appear more visible than other faces, when this would not be the case under circumstances where face images were not physically matched. Important, too, is the proposal by Stein et al. (2014) that fear expressions undergo rapid visual processing via high frequency-selective cortical pathways, such as the short cuts between visual areas outlined in Pessoa and Adolphs (2010)'s Multiple Waves Model. Here, Stein et al. (2014) argue that rapid detection of high frequency fear expressions under CFS is reflective of short latency cortical pathways that allow a readiness to establish the presence of a fearful face. Their argument includes reference to findings from 'Bubbles' studies Smith et al. (2005) and Smith and Schyns (2009), whose findings show a high frequency specificity during the discrimination of fear compared to other face expressions. At the semantic level it is important here to question the extent that discrimination processes are functionally equivalent to detection processes. It seems reasonable to argue that visual detection, such as that which takes place under CFS conditions, refers to an automatic readiness to detect crude elements within a visual scene. Discrimination processes of the same stimuli, however, implies input from processes related to evaluation and explicit differentiation. Findings from the study therefore suggest that data from Stein et al. (2014) are influenced by effects related to contrast normalisation, rather than genuine stimulus properties inherent to expressions of fear, but also support the currently accepted notion that the perceptual threat bias for fear expressions is driven by the low frequency contents of these faces.

10.5.1 Relevance of the present data to the broader literature

The present findings show that normalising facial stimuli for RMS contrast increases their overall stimulus visibility compared to when they are normalised for perceived contrast. This is true for broadband face images that are not spatially filtered in any way, and also for the same images filtered to contain higher ranges of spatial frequency content. This linear effect of contrast normalisation on the spatial frequency content of images is not found for low frequency filtered images. For low frequency faces, differences in the contrast metric used to normalise images does not have an overall visibility effect on low frequency face stimuli. These low frequency filtered facial expressions are also indifferent to effects of contrast normalisation, in that they do not differ in terms of visibility under CFS according to contrast normalisation. These results further highlight the importance of understanding the interaction between contrast metric used during normalisation and the spatial frequency content of experimental stimuli. Additionally, these findings also raise the question of why it seems to be that a general threat bias effect associated with low-frequency components of faces is not observed when CFS is the experimental paradigm.

Chapter 11

Experiment 5: Saccadic responses to facial expressions: effects of physical and perceived contrast, and spatial frequency content

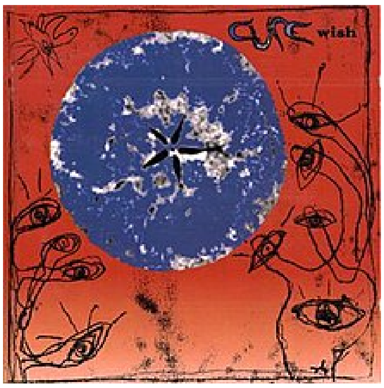

The Cure. (1992). Wish. Artwork copyright to Elektra Records (1992).

11.1 Introduction

Experiment 5 is motivated by findings from Bannerman et al. (2012) that showed selective execution of saccadic eye movements towards low frequency fearful expressions. In addition, Stein et al. (2014) observed advantages for fear expressions to break interocular suppression when they were composed of high frequency components, rather than low. Bannerman et al. (2012) did not normalise their face images for physical contrast, while in contrast Stein et al. (2014) used face images that were normalised for RMS contrast. Experiment 5 investigates the extent that the different frequency effects observed between the 2 studies might be accounted for by the degree to which facial stimuli are normalised for contrast, and if so, whether this effect varies according to the contrast metric used to normalise images. Experiment 5 was also conducted as a comparison study, in order to better understand how the threat bias effect on observer responses under saccadic latency differs compared to detection reaction times generated under CFS conditions. If it is the case that saccadic responses index automated and reflective oculomotor responses, while CFS requires a conscious and manually dictated awareness of a stimulus, then we might expect a possible threat bias to manifest itself differently between the two response time data. Again, it is also important to understand whether observer responses yielded by both paradigms are equally influenced by the contrast and spatial frequency content of facial stimuli.

11.2 Methods

11.2.1 Participants

Twenty-one participants took part in the study. The University of Essex University Ethics Committee approved the study and all participants gave written, informed consent after they were informed that the study was concerned with face perception.

11.2.2 Apparatus

Stimuli were presented using a Dell E2417H monitor, viewed from a distance of 80 cm. For the duration of the experiment, participants used a chin rest to maintain this viewing distance. The monitor screen was 48 cm wide and 27 cm tall. The screen resolution was 1920x1080 pixels, with a refresh rate of 60 Hz and average luminance of 127 cd/m^2. Each pixel subtended 0.02° arc min, measured horizontally. The luminance response of the monitor was linearised by gamma-correcting using a Spyder 5 Elite (Datacolor).

Using monocular recording, eye movements were monitored and recorded using an EyeLink 1000 eye-tracking device (SR Research Ltd). Observer responses were recorded in the form of saccade starts, recorded using the EyeLink ‘saccade start’ function. Saccade start is defined as the point within a trail at which the saccade begins, measured in milliseconds. The sampling rate was 1000 Hz. For 19 out of 21 observers, eye movements were recorded from the right eye.

11.2.3 Stimuli

Stimuli were grayscale front-view face pictures of 10 individuals (5 male, 5 female) extracted from the Karolinska Directed Emotional Faces set (Lundqvist et al., 1998). Faces were cropped to include internal features only. Each individual portrayed 1 of 4 expressions: fear, anger, happiness or neutral. The 10 actors used here were selected from the 16 KDEF sample used in Experiments 1-4 were used for the present study. Broadband face images were filtered using a second-order Butterworth filter in MATLAB. This included low (LSF), mid-range (MSF) and high spatial frequency (HSF) versions of all faces, and broadband face images were included such that there were 4 frequency conditions in total. The cut-off frequencies were $f < 1_{cpd}$ for LSF faces, and $f > 6_{cpd}$ for HSF faces. These were the same cutoffs used in all previous experiments, consistent with those used by Stein et al. (2014) and Vlamings et al. (2009). Each face image subtended 5.4°. The spatial content of each face was either broadband, or varied between 5 and 33 cycles per face-width. Three contrast versions were made for all face images: face images normalised for RMS contrast, the same faces normalised psychophysically for perceived (apparent) contrast, and face images that were not normalised at all, such that any natural differences in physical contrast were preserved in these faces. All stimuli were presented in their normal, upright form, and in manipulated form, where faces were rotated by 180° and the luminance polarity inverted. Using these techniques to manipulate stimuli reduces the visibility of configural information for successful recognition of facial expressions, but preserves images' low-level properties including contrast and spatial frequency content (Gray et al., 2013; Tanaka and Farah, 2003). The procedure for normalising face images for RMS matched and psychophysically matched versions are described in Chapter 10. For all frequency versions of faces, average luminance was set to $.127 cd/m^2$. Stimuli were generated and presented using MATLAB and the Psychophysics Toolbox extensions (Brainard and Vision, 1997; Kleiner et al., 2007; Pelli, 1997). This totalled a maximum of 720 trials: 10(actor) x 4(expression) x 3(frequency) x 3(contrast) X 2(manipulation).

11.2.4 Procedure

All observers began the study with 9-point calibration steps, according to the default EyeLink 1000 calibration program. Each trial commenced with a single fixation cross appearing at the centre of the screen for a pseudorandom time interval between 250-1,000 ms. In order to ensure that observers' gaze location remained as close to central as possible before the appearance of a face, fixation crosses were immediately followed by face images, without the use of a 200 ms gap period (as was included by Bannerman et al. (2012). Each face image was located either to the left or right of centre. All faces appeared for 200 ms. The distance between the centre of a face image and the centre of the screen was 7° degrees. The observers task was to "look towards" the face image, as quickly as possible. Each face was followed by a 200 ms gap after which the trial ended. Before the next trial onset, a default EyeLink 1000 drift correction, or recalibration function, was used to ensure that each observer started a trial at the point at which their gaze was central to a drift correction target located at the centre of the screen. This drift correction target was present between each trial. Each observer completed 360 trials that were separated into 3 blocks of 120 trials each, according to the spatial frequency content of images (broad, low and high). The 120 trials within each frequency block were defined as follows: 10 actors x 4 expressions x 3 contrast conditions, with manipulated and upright versions of these combinations interleaved throughout each block. Trials were randomised within each block, and the order of block presentation was randomised between observers.

11.3 Results

Data tables are displayed in section 11.3 (Data Tables) of this section. Saccade responses were measured in milliseconds. Faster saccades denote rapid detection of a visual stimulus. Data were removed for trials where no saccade response was recorded. For trials were more than one saccade was made, the first saccade was extracted for data analysis.

Data were analysed using a generalised linear mixed effects model. Both observer and the actor of the face image (1-10) were designated random effects (with intercepts only) due to possible variability related to these factors. This allowed analyses to establish generalisable estimates regarding the effect of contrast, expression and image manipulation on saccadic response times. For each model, categorical fixed effects of expression, contrast and manipulation were included, plus contrast-by-expression and manipulation-by-expression interactions, in order to investigate the effects of contrast normalisation and image manipulation on SRTs to different face expressions, respectively. Separate models were fit according to each frequency block, such that in total there were 3 models; one each for broadband, low, and high frequency conditions. For all models, fear expression set as the default reference and intercept in each case. Alpha (α) was Sidak-corrected to 0.0085, according to the 6 comparisons that compared fear with neutral, angry and happy counterparts at *both* levels of manipulation (upright vs. manipulated).

11.3.1 Saccadic responses to broadband faces

Average saccadic response times (SRTs) to detect broadband facial stimuli are displayed in Figure 11.2. Average saccade times are plotted in milliseconds, for broadband face images when they are (a) raw, and therefore not normalised for contrast in any way, (b) when they are normliased for physical RMS contrast, and (c) when they are psychophysically matched for apparent contrast. Data from a generalised linear mixed effects model are summarised in Table 11.1. No significant effect of face expression was observed, showing that the speed of saccadic eye movements do not differ according to the facial expression portrayed by a face. In terms of contrast effects, SRTs did not significantly differ according to whether broadband facial stimuli had been normalised for physical or apparent contrast, or not normalised at all. No overall effect of stimulus manipulation was found, such that SRTs did not significantly differ between normally presented upright faces and their manipulated counterparts.

No significant contrasts were found for contrast-by-expression effects, such that SRTs for facial stimuli did not significantly differ according to an interaction of faces contrast content or facial expression. Additionally, no significant manipulation-by-expression contrasts were found, such that SRTs for broadband facial expressions did not significantly differ according to whether they were presented to observers in their normal or manipulated form. These data are summarised in Table 11.1, and illustrated in Figure 11.2.

11.3.2 Saccadic responses to low frequency faces

Average SRTs to detect low frequency filtered facial stimuli are displayed in Figure 11.3. Average Saccade times for low frequency faces are plotted in milliseconds, when (a) raw, and therefore not normalised for contrast in any way, (b) when they are normalised for physical RMS contrast, and (c) when they are psychophysically matched for apparent contrast. Data from a generalised linear mixed effects model are summarised in Table 11.2. No overall effect of facial expression, contrast metric, or manipulation was observed for low frequency filtered face images. In other words, no significant differences in SRTs were found between facial expressions, or conditions whereby face images were normalised (or not normalised) for different contrast metrics, nor were any overall differences found for upright, normally presented faces compared to their manipulated counterparts.

Contrasts for contrast-by-expression effects showed no significant differences (α= 0.0085) in SRTs between facial expressions, nor were expressions found to vary as a function of whether or not they had been normalised for contrast. Additionally, no significant manipulation-by-expression contrasts were found, such that SRTs for low frequency filtered facial expressions did not significantly differ according to whether they were presented to observers in their normal or manipulated form. These data are summarised in Table 11.2, and illustrated in Figure 11.3.

11.3.3 Saccadic responses to high frequency faces

Average SRTs to detect high frequency filtered facial stimuli are displayed in Figure 11.4. Average saccade times for high frequency faces are plotted in milliseconds, when (a) raw, and therefore not normalised for contrast in any way, (b) when they are normalised for physical RMS contrast, and (c) when they are psychophysically matched for apparent contrast. Data from a generalised linear mixed effects model are summarised in Table 11.3. No overall effect of facial expression was observed, such that the expression portrayed by face images did not have an overall effect on SRTs for different faces images. A significant contrast effect showed that face images psychophysically normalised for apparent contrast received faster saccadic responses (18.38 ms) compared to face images that were not normalised for contrast at all (raw faces). No significant effect of image manipulation was found, such that SRTs did not differ significantly between faces presented in their normal, upright form, and their manipulated counterparts.

Contrasts for contrast-by-expression effects showed no significant (α= 0.0085) differences in response times to detect facial expressions, nor were they differences in SRTs for face expressions found to vary as a function of whether or not they had been normalised for contrast. Additionally, no significant manipulation-by-expression contrasts were found, such that SRTs for high frequency filtered facial expressions did not significantly differ according to whether they were presented to observers in their normal or manipulated form. These data are summarised in Table 11.3, and illustrated in Figure 11.4.

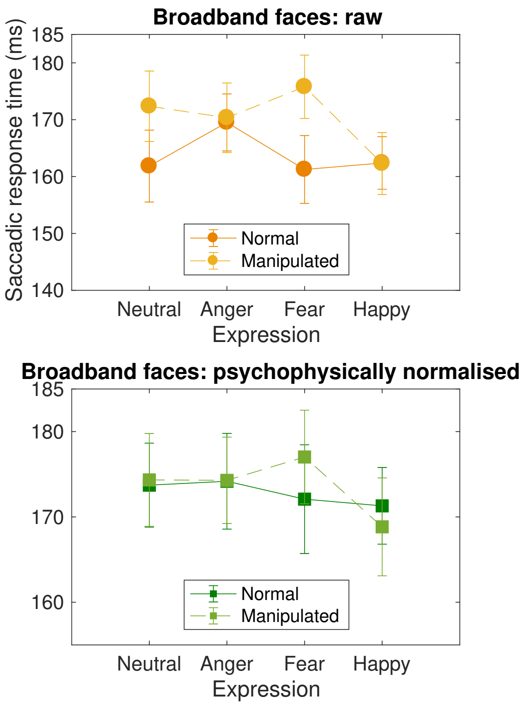

Fig. 11.2 Average SRTs for broadband facial expressions as a function of the contrast metric to normalise images: raw face images have not been contrast normalised in any way, RMS normalised faces have been normalised for physical RMS contrast, and apparent normalised faces have all been psychophysically normalised for apparent contrast according to data from Experiment 2. Separate lines reflect conditions of manipulation, where face images were presented to observers in normal, upright format, or inverted by 180° with retained luminance polarity. All error bars represent associated standard deviations.

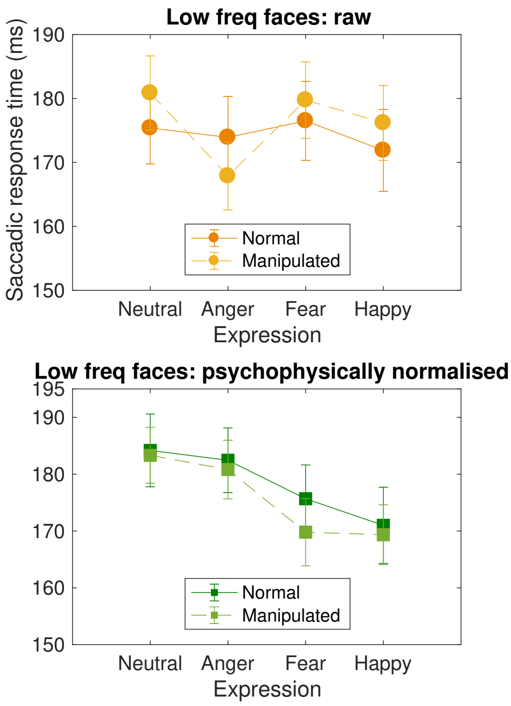

Fig. 11.3 Average SRTs for low frequency filtered facial expressions as a function of the contrast metric to normalise images: raw face images have not been contrast normalised in any way, RMS normalised faces have been normalised for physical RMS contrast, and apparent normalised faces have all been psychophysically normalised for apparent contrast according to data from Experiment 2. Separate lines reflect conditions of manipulation, where face images were presented to observers in normal, upright format, or inverted by 180° with retained luminance polarity. All error bars represent All error bars represent associated standard deviations.

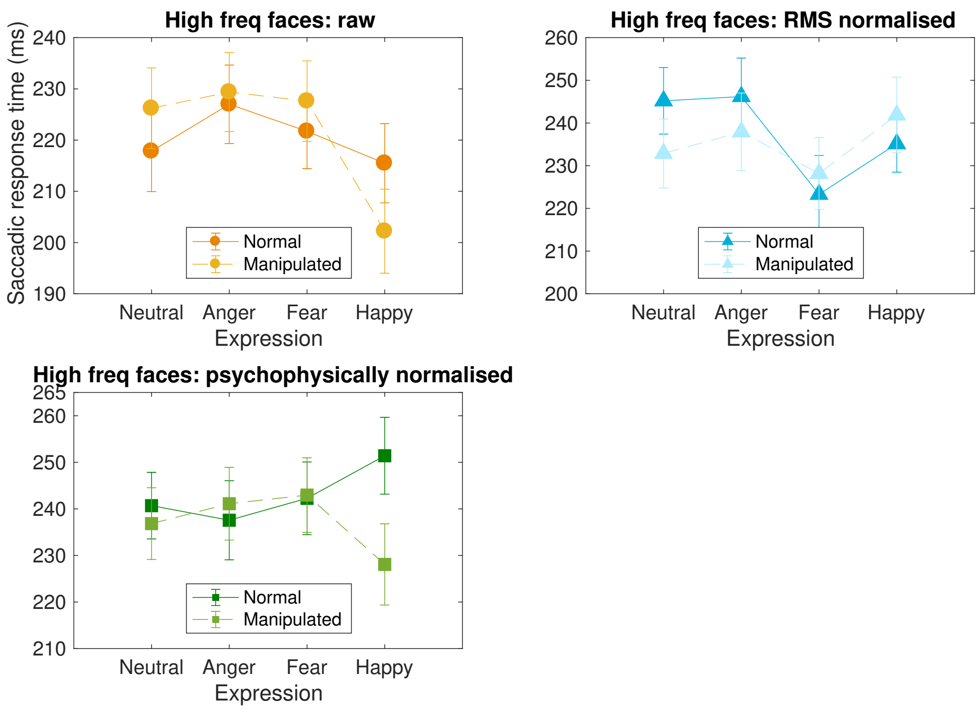

Fig. 11.4 Average SRTs for high frequency filtered facial expressions as a function of the contrast metric to normalise images: raw face images have not been contrast normalised in any way, RMS normalised faces have been normalised for physical RMS contrast, and apparent normalised faces have all been psychophysically normalised for apparent contrast according to data from Experiment 2. Separate lines reflect conditions of manipulation, where face images were presented to observers in normal, upright format, or inverted by 180° with retained luminance polarity. All error bars represent All error bars represent associated standard deviations.

11.4 Data tables

Table 11.1 MLE model and Sidak-corrected comparisons for broadband facial expressions.Analyses include 3 contrast levels, 5 emotion levels, and 2 manipulation levels.

SRTs to broad expressions

Expression effect	SRT Estimate (ms)	(SE)	tStat	df	*p*	CI (95%)
Intercept	164.42	5.90	27.61	2500	<.001	152.7, 176.1
Fear-neutral	1.86	6.09	0.30	2500	.75	-10.09, 13.82
Fear-anger	4.03	6.05	0.71	2500	.47	-7.56, 16.17
Fear-happy	-3.32	5.82	-.57	2500	.56	-14.76, 8.10
Contrast effect	SRT Estimate (ms)	(SE)	tStat	df	*p*	CI (95%)
Raw-RMS normalised	-0.12	5.09	-0.02	2500	0.98	-10.10, 9.86
Raw-apparent normalised	6.69	5.14	1.35	2500	0.17	-3.12, 17.04
Manipulation effect	SRT Estimate	SE)	tStat	df	*p*	CI (95%)
Normal-manipulated	8.17	4.17	1.95	2500	0.05	-0.02, 16.36
Contrast-by-expression	SRT Estimate	SE)	tStat	df	*p*	CI (95%)
Neutral:RMS	5.10	7.27	0.70	2500	0.48	-9.16, 19.37
Anger:RMS	-4.10	7.23	-0.56	2500	0.57	-18.28, 10.07
Happy:RMS	4.33	7.13	0.60	2500	0.54	-9.64, 18.32
Neutral:apparent	-2.22	7.39	-0.30	2500	0.76	-16.73, 12.28
Anger:apparent	-3.47	7.36	-0.47	2500	0.63	-17.91, 10.96
Happy-apparent	0.01	7.23	0.001	2500	0.99	-14.17, 14.20
Manipulation-by-expression	SRT Estimate	SE)	tStat	df	*p*	CI (95%)
Neutral:manipulated	-5.54	5.98	-0.92	2500	0.35	-17.28, 6.19
Anger:manipulated	-6.27	5.94	-1.05	2500	0.29	-17.93, 5.38
Happy:manipulated	-5.28	5.88	-0.89	2500	0.36	-16.81, 6.24

Table 11.2 MLE model and Sidak-corrected comparisons for low-frequency facial expressions.Analyses include 3 contrast levels, 5 emotion levels, and 2 manipulation levels.

SRTs to low frequency expressions						
Expression effect	SRT Estimate (ms)	(SE)	tStat	df	*p*	CI (95%)
Intercept	177.33	7.02	25.25	2498	<.001	163.57, 191.1
Fear-neutral	-0.05	6.06	-0.001	2498	0.99	-11.94, 11.83
Fear-anger	-6.77	6.01	-1.12	2498	0.26	-18.57, 5.02
Fear-happy	-5.23	6.09	-0.85	2498	0.39	-17.17, 6.71
Contrast effect	SRT Estimate (ms)	(SE)	tStat	df	*p*	CI (95%)
Raw-RMS normalised	4.73	5.36	0.88	2498	0.37	-5.77, 15.25
Raw-apparent normalised	-4.23	5.22	-0.81	2498	0.41	-14.83, 6.01
Manipulation effect	SRT Estimate	SE)	tStat	df	*p*	CI (95%)
Normal-manipulated	.088	4.33	0.20	2498	0.83	-7.61, 9.37
Contrast-by-expression	SRT Estimate	SE)	tStat	df	*p*	CI (95%)
Neutral:RMS	-11.80	7.63	-1.54	2498	0.12	-26.77, 3.16
Anger:RMS	8.22	7.44	1.10	2498	0.26	-6.37 , 22.18
Happy:RMS	-3.99	7.60	-0.52	2498	0.59	-18.90, 10.91
Neutral:apparent	9.83	7.44	1.32	2498	0.18	-4.76, 24.44
Anger:apparent	15.17	7.43	2.04	2498	0.04	0.58, 29.76
Happy-apparent	0.75	7.51	0.100	2498	0.91	-13.98, 15.49
Manipulation-by-expression	SRT Estimate	SE)	tStat	df	*p*	CI (95%)
Neutral:manipulated	2.03	6.16	0.33	2498	0.74	-10.05, 14.12
Anger:manipulated	.05	6.07	0.08	2498	0.93	-11.40, 12.43
Happy:manipulated	0.98	6.18	0.15	2498	0.87	-11.14, 13.12

Table 11.3 MLE model and Sidak-corrected comparisons for high frequency filtered facial expressions.Analyses include 3 contrast levels, 5 emotion levels, and 2 manipulation levels.

SRTs to high frequency expressions

Expression effect	SRT Estimate (ms)	(SE)	tStat	df	*p*	CI (95%)
Intercept	223.15	7.59	29.37	2470	<.001	208.25, 238.05
Fear-neutral	-0.29	8.58	-0.03	2470	0.97	-17.65, 17.05
Fear-anger	4.04	8.83	0.45	2470	0.64	-13.27, 21.36
Fear-happy	-10.04	9.02	-1.11	2470	0.26	-27.73, 7.64
Contrast effect	SRT Estimate (ms)	(SE)	tStat	df	*p*	CI (95%)
Raw-RMS normalised	2.12	7.94	0.26	2470	0.78	-13.45, 17.69
Raw-apparent normalised	18.38	7.65	2.39	2470	0.01	3.36, 33.39
Manipulation effect	SRT Estimate	SE)	tStat	df	*p*	CI (95%)
Normal-manipulated	3.66	6.38	0.57	2470	0.56	-8.85, 16.19
Contrast-by-expression	SRT Estimate	SE)	tStat	df	*p*	CI (95%)
Neutral:RMS	16.01	11.14	1.43	2470	0.15	-5.83, 37.86
Anger:RMS	13.67	11.11	1.22	2470	0.21	-8.12, 35.47
Happy:RMS	28.02	11.33	2.47	2470	0.01	5.79, 50.24
Neutral:apparent	-2.95	10.91	-0.27	2470	0.78	-24.35, 18.45
Anger:apparent	-6.15	10.84	-0.56	2470	0.57	-27.42, 15.10
Happy-apparent	14.45	10.93	1.32	2470	0.18	-6.93, 35.90
Manipulation-by-expression	SRT Estimate	SE)	tStat	df	*p*	CI (95%)
Neutral:manipulated	-4.98	9.04	-0.55	2470	0.58	-22.71, 12.73
Anger:manipulated	-3.88	9.01	-0.43	2470	0.66	-21.56, 13.79
Happy:manipulated	14.61	9.09	-1.61	2470	0.10	-32.44, 3.21

11.5 Conclusion

Experiment 5 measures saccadic response times to neutral, angry, fearful and happy expressions, as a function of the degree to which face images were normalised for physical or perceived contrast, and their spatial frequency content. Data showed that for broadband face images (i.e. that are not spatially filtered), saccadic response times do not vary according to whether or not face images had been normalised for contrast or presented in their normal compared to their manipulated form. Broadband fear expressions were not associated with an advantage for saccadic response times (see Table 11.1). When filtered to contain only low frequency components, the expression portrayed by faces, their contrast content, and their presentation format (upright or manipulated) did not have a significant overall effect on SRTs to faces (see Table 11.2). Again, no fear advantage for saccadic response times were observed for low-frequency faces. When filtered to contain only high frequency components, the expression of faces and their presentation format (upright vs manipulated) did not significantly influence overall SRTs to faces. However, an overall effect of contrast at high frequencies showed that high frequency faces normalised for apparent, perceived contrast were overall faster to detect compared to when the same faces that were presented in their raw contrast form. Finally, no fear advantage was observed at this high frequency condition either (see Table 11.3).

11.6 Discussion

Experiment 5 sought to investigate effects of face expression, contrast normalisation, and spatial frequency content on the rapidity of eye movements towards facial stimuli. The motivation for Experiment 5 came from studies of saccadic latency and response time biases for fearful facial expressions (Bannerman et al., 2012, 2010), where the objective of the present experiment was to establish a similar experimental design and to include measures of effects related to the contrast content of facial stimuli.

11.6.1 Summary of findings from Experiment 5

Findings from Experiment 5 showed few effects related to facial expression and contrast normalisation effects. No evident threat bias was observed at broadband frequencies, as was expected based on findings from Bannerman et al. (2010) and Bannerman et al. (2012). The present findings did not show an advantage for faster saccadic responses when responding to broadband fear expressions, or indeed at any other level of spatial filtering. In terms of contrast metric effects, findings from Experiment 4 showed that under CFS, threat bias effects are more likely to occur when face images have been normalised for RMS contrast. This effect was not observed for Experiment 5, in that the time taken to saccade towards broadband face stimuli did not differ according to contrast normalisation, nor as a function of the contrast metric used to normalise facial expressions. We might have expected this to be the case for high frequency filtered faces, in the form of a fear bias for RMS normalised faces for high frequency filtered stimuli. However, this was not the case, and in fact high frequency faces were overall better detected when normalised psychophysically for perceived contrast, rather than RMS contrast, or in the absence of any normalisation at all. Findings from the present experiment and Experiment 4, taken together, suggest that effects of contrast normalisation differently influence the salience of face images according to the nature of the experimental paradigm used.

11.6.2 The present findings compared to those of Bannerman et al. (2012)

Findings from Bannerman et al. (2010) showed that rapidly presented (20 ms) unfiltered broadband fear and neutral expressions were associated with a fear advantage for capturing attention -denoted by faster SRTs- and an inhibitory effect on attentional disengagement when faces were presented as cues to a given target location. Enhanced attentional capture and affected disengagement for target location was only found for fear expressions, and this effect diminished when face expression cues were presented for longer durations of 100 ms. The findings were interpreted by Bannerman et al. (2010) as showing that saccades, compared to manual responses, access stimulus information available at earlier stages of processing, such that fear biases in the form of saccadic eye movements are likely to be exclusive to short presentation durations of 20 ms, as opposed to >100 ms. However, it is important to note here that Bannerman et al. (2010)'s experimental task related to rapid identification of targets that had been preceded by rapid presentations of fear or neutral expressions. The aspects of visual attention measured here were therefore related to the attentional facilitatory and inhibitory effects associated with fear expressions, rather than initial and first saccade movements alone for fear expressions. Bannerman et al. (2012) conducted an adaptation of Bannerman et al. (2010), presenting observers with neutral, fearful and happy faces that had been spatially filtered to include low ($f < .08_{cpd}$) and high ($f > 3.3_{cpd}$) frequency information. Stimulus duration matched that of Bannerman et al. (2010), where face stimuli were presented for rapid 20 ms durations, where no attentional task was involved. Bannerman et al. (2012) showed rapid saccade biases for fear expressions only when they were composed of low frequency information. Information about stimulus contrast normalisation is not included within Bannerman et al. (2010), and Bannerman et al. (2012) note that their experimental stimuli were not normalised for physical contrast. The present study only measured saccadic responses to 200 ms presentations of facial stimuli, as opposed to shorter presentation durations of 20 ms such as those used by Bannerman et al. (2010) and Bannerman et al. (2012). This was because the aim of the present study was not concerned with concealing facial expressions from conscious visual awareness, rather, the variable of interest was simply how rapid saccades are deployed for spatially filtered face expressions in general. Effects of fear expressions as cues to target locations were therefore not relevant here. Additionally, different spatial frequency cut offs to those used by Bannerman et al. (2012) were used in order to maintain consistency between stimuli used for Experiments 1-4, and also to match cut-offs employed by Stein et al. (2014) and Vlamings et al. (2009).

In summary, Experiment 5 conducted a saccadic latency study similar in nature to that of Bannerman et al. (2012), with the exception of stimulus presentation duration and the precise spatial frequency content of images. These deviations in the experimental design between the present study of and that Bannerman et al. (2012) may account for why the present data do not support that generated by Bannerman et al. (2012). Here, no threat bias effect was found in any spatial frequency condition: faster SRTs were not observed for fear compared to any other expression when faces were broadband, or low compared to high frequency filered. Importantly, no such effect of low frequency for fearful faces was observed, as would have been predicted by the current low frequency account. Additionally, findings did not continue the expected pattern of effects based on those generated from Experiments 2 and 4, whereby high frequency fear expressions receive a boost in their salience that is driven by normalisation of faces' RMS contrast. In the present experiment, no such bias effect was observed for RMS normalised high frequency faces. Instead, at high spatial frequencies, normalising faces for apparent, perceived contrast facilitated their detection. Overall, the present findings do not support the generally-accepted view, and indeed, pre-existing evidence that the visual system preferentially detects the low spatial frequency components in fearful facial expressions. To confirm, there was also no evidence to suggest that under saccadic latency the probability of detecting a threat bias may be inadvertently facilitated by the use of RMS contrast normalisation. One possible explanation for this is that if saccadic responses sample information early during processing, such that biases for fear expressions are 'short lived' (Bannerman et al., 2010), the presentation duration of 200 ms used in the present study may have exceeded the presentation time necessary to capture fear effects. Indeed, data from Bannerman et al. (2010) showed that attentional advantages associated with fear expressions were only observed for saccadic responses to fearful faces presented for 20 ms but not when the same faces were presented for 100 ms. However, the objective of the present study was not to render visual stimuli explicitly invisible to observers, and so stimulus duration was not determined by very rapid presentation windows. It may be that under normal viewing conditions, where stimuli are not concealed from conscious visual awareness, that saccadic eye movements simply do not prioritise threat-relevant facial information. A re-design of Experiment 5 might therefore be to include additional design factors including a second stimulus duration condition. This would reflect those used by Bannerman et al. (2012), and an additional spatial frequency condition for mid-range information, as there is a degree of overlap between the cutoffs used by Experiments 1-5, Stein et al. (2014), Vlamings et al. (2009) and those used by Bannerman et al. (2012). This point is returned to in the Discussion of the present thesis.

Chapter 12

Discussion

Manchester Orchestra. (2018). I Know How To Speak. Artwork copyright to Manchester Orchestra (2018).

12.1 A summary of Experiments 1-5

Experiment 1: image analyses

Experiment 1 adopted the use of image analysis techniques to investigate differences in low-level image properties between facial expressions of anger, fear, happiness, disgust, and neutral faces. This included a measure of raw faces' physical RMS contrast, and Fourier amplitude spectra; the contrast energy by spatial frequency relationship within each image. Findings, taken together, show that expression-related differences in naturally-occurring RMS contrast become more pronounced as faces are filtered to contain higher spatial frequency information. For raw faces -not normalised for contrast or spatially filtered in any way- fearful facial expressions naturally contain lower RMS contrast content compared to expressions of disgust and anger. When the same faces were low frequency filtered, no expression-related differences in RMS contrast were observed at all. Moreover, mid-range frequency filtered fear expressions are only lower in RMS contrast compared to expressions of disgust, but at the level of high frequency filtering, fear expressions become significantly lower in RMS contrast compared to all but neutral facial expressions. In terms of the Fourier amplitude spectra for raw and unfiltered facial expressions, fearful faces are associated with a steeper amplitude slope compared to neutral and angry expressions. Consistent with the analyses of RMS contrast for fear expressions, these findings show that normal, broadband fearful expressions contain naturally less information as spatial frequency increases, compared to the rate at which this occurs for neutral and angry faces. Fearful faces may thus be characterised as containing lower contrast at high frequencies compared to some other expressions.

Experiment 2: contrast matching

Experiment 2 conducted a behavioural experiment using a contrast matching task, in order to establish differences between perceived image salience for different expressions; the extent that expressions differ in the amount of physical contrast they each require in order to appear the same as a reference face composed of 10% Michelson contrast. Apparent contrast values reflected the degree of physical contrast attributed to faces by observers such that they matched the reference stimulus. Apparent contrast was calculated based on both the RMS and the Michelson contrast of faces; the degree of each metric required to match faces at the perceived level. Findings showed that expressions naturally differ in their apparent, perceived contrast, but also that these expression-related differences in perceived image salience vary according to both the spatial frequency content of images, and the contrast metric used to calculate apparent contrast. For natural broadband face images, fearful expressions required significantly less RMS contrast compared to all but angry faces in order to perceptually match a reference stimulus. That some of these effects remained consistent under conditions of manipulation suggests that the perceived salience of broadband face expressions is driven by low-level image properties, rather than configural emotion-relevant content in faces. Interestingly, Michelson contrast was not found to differ at all between fear and other faces at the point at which they were perceptually matched, suggesting some degree of continuity may be possible between the perceived contrast of expressions if these are to be normalised for Michelson contrast. When the same faces are filtered to contain mid-range spatial frequency content, perceived contrast differs very little between fear and other faces. The only difference observed here was that low frequency fear expressions require less Michelson contrast compared to disgust in order to be perceptually matched. This implies that the perceived contrast of mid-range spatial filtered faces is largely uninfluenced by the contrast metric used for normalisation. However, when faces are high spatial frequency filtered, differences in perceived salience between fear and other emotions are most pronounced. Importantly, the direction of this effect i.e. whether fear is lower or high in perceived contrast compared to other faces, depends on contrast metric. High frequency fear expressions require less RMS contrast than all but happy faces but *more* Michelson contrast compared to happy and disgust faces in order to be perceptually matched. Findings suggest that the salience of fear expressions is significantly influenced by physical contrast metric, but also the degree to which faces have been spatial frequency filtered. In particular, the general message here seems to be that high frequency fear expressions normalised for RMS contrast may receive a boost in physical contrast above and beyond that which is necessary for them to be perceptually matched to a reference stimulus.

Experiment 3: contrast sensitivity and analyses of effective contrast

Experiment 3 was formed of two parts; the first conducted a behavioural investigation of Hedger et al. (2015)'s proposal that effective contrast is higher for fearful expressions than it is for neutral faces. A contrast sensitivity task measured contrast thresholds for facial expressions of fear, anger and happiness and neutral faces. Findings showed that visual contrast thresholds did not differ as a function of the facial expression portrayed by faces, therefore providing no evidence to suggest that fear expressions may be especially well detected as a result of being well matched to the contrast sensitivity function. A second stage to this investigation was to re-conduct the procedure used by Hedger et al. (2015) to measure faces' effective contrast; the extent to which their statistical properties were matched to undergo effective processing via the contrast sensitivity function. An extension of this procedure gave measures of effective contrast when faces were raw, and therefore not normalised for overall contrast, but also when they had been normalised for RMS contrast prior to analyses. Consistent with data from Hedger et al. (2015), faces that were normalised for RMS contrast were, in many cases, higher in effective contrast compared to other expressions. This was true across 3 out of 5 face image samples analysed. Importantly, however, some of these patterns of a fear advantage for effective contrast diminished when the same faces had not been normalised for RMS contrast. Together, findings from Experiment 3 showed that advantages for fear expressions to exploit the contrast sensitivity function may rely in part on the process by which they are first normalised for physical RMS contrast. Analysis of the raw RMS contrast naturally contained within the face samples used in these image analyses revealed a fairly consistent, but not always significant, pattern of lower RMS contrast associated with fear expressions, demonstrating that expressions of fear do not always contain naturally higher levels of physical contrast compared to other face expressions; an effect that was consistent and extended findings from Experiment 1.

Experiment 4: detection of consciously suppressed face expressions under CFS

Experiment 4 conducted a behavioural study measuring the effects related to the spatial frequency content of face expressions on observers' ability to rapidly detect a face under conditions of intraocular suppression (CFS). The extent to which detection times for spatially filtered face expressions was also measured as a function of whether facial stimuli had been normalised for physical RMS contrast, or normalised psychophysically for apparent contrast. Findings showed a relatively small fear bias for breaking suppression under CFS, but that this expression effect is not influenced by the contrast metric used to normalise faces. Here, broadband fearful expressions are detected faster compared to angry faces, regardless of whether faces are normalised for RMS or apparent contrast. The detection of low frequency faces was unaffected by both facial expression and contrast metric, although manipulated versions of faces did have an overall advantage in breaking suppression. As the spatial frequency content of faces increased, so did the visibility of these faces under CFS when they were normalised for RMS contrast. Normalising faces that contain higher ranges of spatial frequency information appears to reduce the overall salience of faces, but also reduces the probability of detecting a fear advantage. A fear advantage for such high frequency stimuli is most evident when these faces are normalised for RMS contrast: RMS normalised high frequency fear expressions are more break suppression faster compared to both angry and happy faces. These findings suggest that effects of normalising stimuli on the salience of facial images increases to some degree with their spatial frequency content, provided that the contrast metric used is RMS contrast.

Experiment 5: oculomotor responses to facial expressions using saccadic latency

Experiment 5 conducted a saccadic latency study, measuring first orienting eye movements towards spatially filtered facial expressions that were normalised for physical compared to apparent contrast, or not normalised for contrast in any way at all. No evidence was found to support a fear advantage in receiving faster saccadic response times, regardless of faces' spatial frequency content. Unlike findings from Experiment 4, response times to saccade towards facial stimuli were not influenced by whether or not they had been normalised for contrast, or the contrast metric used to achieve this. Instead, perceptually matching faces for perceived contrast was associated with a general advantage for looking towards high frequency faces, compared to when these were RMS normalised. Together, findings did not provide evidence in support of an threat bias for fear faces, nor findings of a low frequency driven advantage for fearful faces evidenced by Bannerman et al. (2010) and Bannerman et al. (2012).

12.2 Relevance to the broader literature

12.2.1 Facial expressions naturally differ at the physical level

Theories of face expression perception posit that the physical configurations of facial expressions are uniquely distinct from each other (Ekman and Cordaro, 2011; Susskind et al., 2008). This is because they correspond to distinct evolutionary functions, where the perceptual categorisation of each expression is thought to rely on the physical signal emitted from such expression-related differences between faces' physical configuration (Smith et al., 2005; Susskind et al., 2008). However, as far as we are aware, the physical characteristics that define these faces, and the way in which they differ *between* faces remains relatively unknown. Redies et al. (2007) show, for example, that image categories such as images of natural scenes display certain regularities in terms of the statistical image properties that they contain, defined as the $1/f$ Fourier amplitude slope of natural scenes. Redies et al. (2007) show that images of faces, by comparison, can be categorised according to their Fourier amplitude slope that is visibly steeper compared to that of natural images, such that facial expressions contain less information at higher spatial frequencies. Following the argument that the contrast sensitivity in humans has adapted to optimise visual processing of information commonly experienced in natural scenes (Campbell and Robson, 1968), we might also expect a similar adaptation to have been selected by nature for processing different kinds of facial information, such as that related to facial expression perception (Susskind et al., 2008). Findings from Experiment 1 extend and contribute findings to the current literature that evidence naturally occurring expression-related differences in basic image properties of facial stimuli. When face images are not spatially filtered or normalised for physical contrast, physical RMS contrast differs between expressions, showing that different expressions differ in terms of characteristics that modulate image salience. In particular, RMS contrast is naturally lower for expressions of fear than it is for expressions of disgust and anger. Experiment 1 also measured the Fourier amplitude spectrum for different facial expressions. Fearful faces, compared to neutral, angry and disgust faces, are associated with a steeper amplitude slope compared to that associated with neutral and angry faces. Importantly, when these images are subjected to spatial filtering, differences in physical contrast between faces become increasingly apparent as the spatial frequency content in faces is increased. For example, naturally lower RMS contrast present in broadband expressions of fear diminishes when faces are filtered to contain only low frequency information. This effect re-emerges for mid-range frequency versions of faces, where RMS contrast is lower for fear than for disgust expressions, and at the level of high frequency filtered, RMS contrast is lower for fear expressions compared to all but neutral faces. These findings contribute to the current consensus that certain facial expression images naturally posses a steeper amplitude slope compared to natural image categories Redies et al. (2007), adding that such differences also occur *within* face categories and *between* facial displays of emotion. This information is a prerequisite for understanding the way in which certain experimental techniques, including spatial filtering and contrast normalisation, may mask or in some way influence such naturally occurring low-level differences.

12.2.2 Faces normalised for physical contrast are not necessarily normalised at the subjective level

Contrast normalisation is a commonly employed technique within both psychophysical studies, but also studies within other areas of cognitive psychology and social sciences. The objective here is to match, or average, physical contrast differences that might exist between stimuli. Theoretically, this controls for any possible low-level image differences that may result in some images being perceived as more salient than others. Observed salience differences that remain between stimuli after normalisation of stimulus contrast can therefore be attributed to the *content* of the image, rather than image features that may be irrelevant. Another reason particularly relevant to studies of visual psychophysics is to balance differences that occur when stimuli are spatially filtered. Controlling for effects of contrast-driven differences in stimulus salience is a useful tool when the feature of interest is the relative effects of spatial frequency on image salience, but also, low frequency filtered images will naturally contain more contrast compared to high frequency versions. Normalising images for contrast therefore helps to limit some of this variability. The currently accepted view is that normalising images of natural scenes for RMS contrast is a beneficial procedure because doing so maintains an ideal degree of consistency between the resulting images' physical and perceived contrast (Meese et al., 2017; Redies et al., 2007). In other words, natural images normalised for RMS contrast will be identical at the physical level, but also consistent in terms of their perceived image salience. However, the issue here is that the statistical image properties that characterise natural scenes are different to those that characterise face stimuli (Redies et al., 2007). Findings from Experiment 1 show naturally occurring expression-related differences in such statistical properties of facial stimuli.

Indeed, findings were consistent with those observed by O'Hare and Hibbard (2011), where the perceived contrast of broadband images varies as a function of the spatial frequency content of their frequency filtered versions. Normalising broadband face expressions for RMS contrast, as is the procedure for natural images (Redies et al., 2007), does not ensure that facial stimuli will be consistent in terms of their perceived image salience. For typical broadband face stimuli that have not been spatially filtered, fearful face expressions naturally contain *less* physical contrast compared to that required for other face expressions to appear subjectively the same, in terms of contrast, to a given reference stimulus. A finding that is important for studies using typical broadband face images, as opposed to spatial frequency filtered stimuli, differences in perceived contrast between expressions depend on the contrast metric used to normalise images. Indeed findings from Experiment 2 showed that broadband facial expressions may be better normalised for Michelson contrast, if the objective of such normalisation is to reduce differences in stimulus salience. Results from the same study *suggest* that the result of normalising broadband faces for RMS contrast is a fear advantage whereby fearful faces appear higher in terms of stimulus salience. However, Experiment 4 showed that when the same faces are normalised psychophysically for apparent contrast, such that they contain relative amounts of physical contrast necessary for each expression to match a given reference image, this fear bias effect diminishes. This finding suggests that normalising broadband face expressions for RMS contrast boosts the perceived salience for fear expressions compared to other faces. This is not the case when normalising faces for perceived contrast. It could then also be argued that apparent contrast may be a useful candidate for contrast normalisation. Moreover, Experiment 2 showed that differences in salience between low frequency filtered facial expressions are relatively indifferent to contrast normalisation. This finding shows that low frequency versions of face stimuli ($f < 1_{cpd}$) do not require different levels of physical contrast in order to appear the same. But this is not true for all frequency versions of stimuli, as the magnitude of differences in perceived salience between expressions after contrast normalisation increases with the spatial frequency content of spatially filtered faces. Fearful face expressions still require less physical contrast to appear the same as disgust faces when these images are mid-range frequency filtered, but this is provided that they are normalised for RMS contrast, and less Michelson contrast than angry faces. Facial expressions filtered to contain high frequency information, however, are not guaranteed any consistency of perceived salience regardless of contrast metric.

These findings show that normalising facial stimuli does not guarantee that images will be consistent in terms of their associated perceived contrast. Rather, ensuring that faces are normalised at both the physical and subjective level depends on several factors, including the spatial frequency content of images, and the contrast metric used to normalise stimuli.

12.2.3 Contrast normalisation can artificially enhance expressions' perceived salience

Combined findings from Experiments 1 and 2 show that: fearful expressions do not naturally contain higher amounts of physical contrast, but actually contain relatively lesser amounts of RMS contrast compared to other expressions, where this effect becomes more pronounced as faces are high frequency filtered (Experiment 1); fearful expressions also require significantly less physical contrast than other expressions in order to be equal at the perceptual level, an effect that again increases with faces' frequency content (Experiment 2); and finally, that different contrast metrics used to normalise stimuli can reduce residual differences between expressions' salience (Experiment 2 and 4). The importance of these findings are twofold; in the first instance they tell us about natural physical and perceptual differences that exist between expression stimuli. In the second instance they suggest *a priori* that the process by which fear expressions are normalised for physical contrast inadvertently boosts their perceived salience. Naturally, these faces contain low levels of contrast, but increasing this via normalisation is likely to do so above and beyond the amount required for these expressions to match other faces in terms of their subjective salience. In other words, normalising fearful faces for physical contrast artificially enhances their perceived contrast so that they appear more salient compared to other facial expressions. These findings demonstrate that normalising facial stimuli for physical contrast is likely not only to leave outstanding differences in perceived contrast between images, but that it will also allow fear expressions in particular to benefit from an artificial boost in apparent contrast. The magnitude of this normalisation effect increases linearly with the spatial frequency content of facial stimuli. Such that the salience of fear expressions is most bolstered the more that these images are filtered to contain higher frequency content.

12.2.4 To what extent is the fearful face bias a contrast normalisation effect?

Findings from Experiments 1 and 2 suggest that salience effects associated with fear are substantially driven by contrast normalisation. Fearful expressions that contain higher frequency components are more likely to benefit from an artificial boost in apparent contrast. Indeed, Experiment 3 shows that the effective contrast of fearful expressions -a predictor of their visibility under CFS (Hedger et al., 2015)- is facilitated by the normalisation of faces for RMS contrast. Effective contrast measured for fearful faces in their natural format shows that fear expressions are often, but not always, no better matched to contrast sensitivity processing than other expressions. In some cases, fear expressions are actually lower in effective contrast than their counterparts when they are not contrast normalised. Findings from Experiment 4 also show that the same broadband fear faces are more likely to break suppression faster compared to other expressions when they have been normalised for RMS contrast as opposed to perceived contrast. This is likely due to the way in which psychophysically normalised faces have not received a boost in perceived salience. The strength of this effect was small and almost non-existent for low frequency filtered faces, and increased with faces' frequency content. Under CFS, visibility does not differ between low frequency fear expressions and other faces. At the level of high frequency filtering, a fear bias emerges when these faces have been normalised for RMS, rather than apparent, percevied contrast. In fact, normalising faces for RMS contrast under CFS had an overall effect on the visibility of faces, but specifically when the spatial frequency content of images was increased. RMS normalised mid-range and high frequency faces were overall more visible than they were when normalised psychophysically for perceived contrast. This overall effect of contrast metric used to normalise images was not found for broadband or low frequency filtered faces, suggesting that effects of contrast normalisation that enhance both the visibility of fear expressions and overall stimulus salience operates most at higher spatial frequency scales. Combined, these findings suggest that fear biases are likely driven at least in part by the use of contrast normalisation; particularly when the contrast metric used to achieve this is RMS contrast. In particular, we expect then that high frequency faces are most vulnerable to artificial boosts in perceived contrast compared to their expression counterparts, such that threat bias effects become more reliant on contrast normalisation as the spatial frequency content of faces increases.

Experiments 1, 2 and 4 show that contrast normalisation influences the subjective appearance of facial stimuli in a way that can be expected to drive threat bias effects for fearful facial expressions in particular. Data from Experiment 3 shows how these effects might operate at the physical level. Measuring the Fourier amplitude spectrum of 5 facial expressions from 5 commonly used face databases after these had been attenuated by a model for the human contrast sensitivity function generated measures of effective contrast for each face expression. Higher effective contrast was shown by Hedger et al. (2015) to be uniquely associated with fearful expressions, and was a reliable predictor of a fear advantage to break suppression under CFS conditions. Experiment 3 showed the same general pattern that effective contrast *is* higher for expressions of fear compared to other faces. However, findings also show that this effect is positively influenced by face stimuli having first been normalised for RMS contrast. Many faces that underwent analyses in absense of RMS contrast normalisation were not associated with differences in effective contrast. These findings, when applied to those from Experiment 4, show that expression-related differences in image salience under CFS are more frequent when faces are normalised for RMS contrast, rather than when they are normalised perceptually for apparent contrast. Importantly, however, Experiment 4 did not include a condition whereby faces were presented to observers in a non-normalised format, and so it is not possible to extrapolate from this that RMS normalised faces under CFS are salient wholly because of RMS normalisation. This point is returned to under 'Avenues for future research'. Experiment 5 measured the time taken for saccadic eye movements to deploy in response to spatial filtered face expressions, including a condition whereby faces were not normalised for contrast in anyway, such that they were presented in their raw form. Consistent with findings so far, a fear bias effect was again found for high frequency fear expressions, but only when these faces were normalised for RMS contrast. Interestingly, the pattern of effects observed in Experiments 1-4 were not consistent here. No fear advantage for receiving first orienting saccades was observed for fear expressions, regardless of their frequency content, nor whether or not they had been normalised for contrast.

It is important to note here that these conclusions cannot be strictly generalised to those generated by using different experimental techniques. The is because these findings measure the threat bias only using paradigms related to contrast matching, contrast sensitivity, CFS and saccadic latency; only a sub-section of the available techniques used to investigate perceptual biases, as discussed in the techniques overview in Chapter 9.

12.2.5 To what extent does the fearful face bias rely on low-frequency information?

Bolstering effects of contrast normalisation on images' perceived salience varies according to the spatial frequency content of expressions, where the degree of physical contrast necessary to achieve perceptual matching was greater or lesser depending on the contrast metric (RMS or Michelson) used to determine apparent contrast. However, this was not true for low frequency filtered faces, where perceived contrast for upright and low frequency fearful faces was no different to other expressions (Experiment 2). Normalising the same faces for RMS contrast under CFS was also not associated with any fearful face advantage. This finding in particular suggests that low frequency fearful faces are not as vulnerable to effects of contrast normalisation that boost the perceived contrast of such faces, provided that contrast normalisation occurs after frequency filtering. Because effects of contrast normalisation appear to operate at higher spatial frequency scales, a reasonable conclusion is that perceptual biases found for high frequency fear expressions are likely driven by contrast normalisation. This notion is upheld by findings from both Experiment 2 and 4, and could also account for the high frequency-dependent fear advantage observed by Stein et al. (2014), where faces normalised for RMS contrast were faster to break suppression compared to low frequency counterparts. Findings from Experiments 1-4 suggest that these high frequency faces likely received a boost in physical contrast that exceeded the degree required to perceptually match these images, such that they appeared considerably more salient compared to their expression counterparts. It therefore seems to be that contrast normalisation effects operate within high frequency ranges, where high frequency fearful faces receive an artificial boost in perceived salience where low frequency versions do not. This gives little reason to question the body of evidence that demonstrates a fear advantage that is driven by low frequency components. Despite this, findings from both Experiment 4 and 5 provided no evidence in support of the low-frequency account of the fear bias. These findings have two important implications: that there are no substantial fear bias effect for low frequency face images under conditions of CFS and saccadic latency, but also, that low frequency filtered face images are far less likely compared to high frequency faces to generate artificial threat bias effects. It therefore remains unclear why some studies continue to find a low frequency dependent threat bias effect when the present studies did not.

In summary, findings from the present studies do not provide sufficient evidence to suggest that the currently accepted low-frequency account of the fear bias is simply an effect of contrast normalisation (Vuilleumier et al., 2003), rather than a reflection of a genuine perceptual bias for information naturally contained within fearful expressions.

12.2.6 Is the fear bias task dependent?

Bannerman et al. (2010) argue that saccadic latency is an ideal and reliable measure of implicit biases for facial expressions. This is primarily on the basis that this paradigm relies on oculomotor responses that are reflexively-driven, and therefore less likely to reflect complex and higher level stages of processing such as those related to manually driven responses to consciously perceived stimuli (Bannerman et al., 2010). Bannerman et al. (2010) therefore argue that saccadic latency is a preferred measure of the threat bias compared to CFS. Based on this argument, we might expect that a threat bias is more likely to be detectable under conditions of saccadic latency as opposed to CFS, as rapid and automated eye movements to visual stimuli are more naturalistic and ecologically relevant measures of perceptual biases. It is therefore surprising that no fear bias effects were observed by Experiment 5. This is especially true given that findings from the CFS study conducted for Experiment 4 display several different effects related to fear biases, each differently influenced by factors including spatial frequency and contrast normalisation of stimuli. It therefore remains unclear whether saccadic latency is simply not as vulnerable as other experimental paradigms when it comes to effects of stimulus normalisation, or whether threat bias effects are simply less likely to operate at such levels of perceptual behaviours.

A notable finding here, too, is the visibility advantage found for fearful over angry faces under CFS (Experiment 4). This finding was fairly consistent between broadband and higher spatial frequency conditions. This implies an interesting difference driving salient effects for fear over angry faces, but particularly one that is robustly detected under CFS conditions but not when studied using saccadic latency.

The issue of task-relevance and a threat bias for fear expressions is summarised in a recent meta-analyses by Hedger et al. (2016). Hedger et al. (2016) reviewed overall threat bias findings generated from masked visual probe, binocular rivalry and CFS paradigms. Each paradigm claims to measure unconscious registration of threat-relevant information, yet achieves this in very different ways, contributing mixed effect sizes of threat biases. Isolating effects only from studies using fear expressions as their experimental stimuli substantially increases the size and consistency of these effects, but the fact that variability remains in the outcomes and sizes of findings implies that effects at least to some extent task-relevant.

12.2.7 To what extent do the present findings explain inconsistencies in the wider literature?

In terms of the present findings in relation to those observed by Stein et al. (2014), it is evident that threat bias effects associated for high frequency fear expressions are driven, at least in part, by the process of physical contrast normalisation. This allows the higher frequency versions of fear expressions to appear more visible than other faces, when this would not be the case under circumstances where face images were not physically matched. Important, too, is the proposal by Stein et al. (2014) that fear expressions undergo rapid visual processing via high frequency-selective cortical pathways, such as the short cuts between visual areas outlined in Pessoa and Adolphs (2010)'s multiple waves model. Here, Stein et al. (2014) argue that rapid detection of high frequency fear expressions under CFS is reflective of short latency cortical pathways that allow a readiness to establish the presence of a fearful expression. Stein et al. (2014)'s argument is supported by evidence from Gaussian bubbles studies, including Smith et al. (2005) and Smith and Schyns (2009), whose findings show a high frequency specificity during the discrimination of fear compared to other face expressions. In order to address this small body of evidence that opposes the low-frequency fear account, it is important to question at the semantic level the extent that *discrimination* processes are functionally and practically equivalent to *detection* processes. Observers' task for Smith et al. (2005) was to "judge" and categorise segments of facial expressions according to the emotion they portrayed. This is not the same perceptual ability as that which is measured under CFS, where observers are simply required to press a button at the point at which they are consciously aware of faces' presence. *Discrimination* processes of the same stimuli, however, implies input from processes related to evaluation and explicit differentiation. A reasonable argument against inconsistent findings of a high frequency-dependent fear bias is that perceptual discrimination tasks using the Gaussian Bubbles technique is reflective of a different process compared to that required to detect faces under CFS, and should be interpreted as such. Additionally, findings from the present thesis strongly suggest that data from Stein et al. (2014) is driven by artificial salience effects inadvertently driven by contrast normalisation of stimuli, rather than a genuine fear bias that reflects high frequency fear advantages proposed by the multiple waves model (Pessoa and Adolphs, 2010; Stein et al., 2014). The constructs measured via Gaussian Bubbles studies, on the other hand, remain unequivocally related to configural versus local face processing, where high frequency effects for a fear bias are exclusive to such experimental designs (see Murray and Gold (2004) for a review).

Hedger et al. (2016) argue that inconsistencies within the literature are significantly influenced by differences in experimental design and methodology that vary between studies. This includes the importance of conceptualising the threat bias as a composition of several different subcategories of threat-sensitive perceptual behaviours, such as those related to threatening objects (Koster et al., 2004; Stormark et al., 1995), instead of an all-purpose visual function. A perceptual bias for fearful expressions is one of such subcategories, where studies concerned with their mechanisms in attention are for the most part consistent (Bannerman et al., 2012; Bayle et al., 2011; Carlson and Reinke, 2008; Gray et al., 2013; Phelps et al., 2006; Pourtois et al., 2006; Tsuchiya and Koch, 2005; Vlamings et al., 2009; Whalen et al., 1998, 2001; Yang et al., 2007). It therefore seems to be that evidence of a fearful face bias is more rosbust than a general, overall threat-bias effect for all threat-relevant visual information (Hedger et al., 2016). However, it is also evident that the magnitude of these effects are likely influenced by the experimental paradigm used to investigate an attentional bias effect (Hedger et al., 2016), and as the present findings show, also by factors related to physical nature of experimental stimuli.

A fear-specific, or overall emotion effect?

A consistent argument raised across several studies is the extent to which perceptual biases for fearful expressions reflect a genuine prioritisation of fearful facial stimuli over and above all other expressions, or, whether they occur simply as an overall preference for emotional stimuli when presented with fearful compared to neutral faces (Bannerman et al., 2012; Gray et al., 2013). It is difficult to address this question, as many studies within the literature refer only to 1 or 2 emotion comparisons when investigating fear effects. For example, Bannerman et al. (2009a); Carlson and Reinke (2008); Hedger et al. (2015); Holmes et al. (2005); Phelps et al. (2006); Stein et al. (2014); Vlamings et al. (2009); Vuilleumier et al. (2003) and Bannerman et al. (2010) all include only neutral faces as a comparison for fear expressions, compared to Bannerman et al. (2012); Pourtois et al. (2006); Schupp et al. (2004); Whalen et al. (1998, 2001); Williams et al. (2004); Yang et al. (2007) and Bayle et al. (2011) who include comparisons of both a neutral face and another emotional expression; often expressions of happiness or anger. Notably, though, all studies cited here report bias effects for fear expressions, weakening the argument that fear biases are entirely driven by an overall preference or prioritisation of emotional over neutral stimuli. However, a clear effect exclusive to expressions of fear is further clouded by evidence from studies that demonstrate advantages for detection times in response to happy compared to fearful and anger expressions (Killgore and Yurgelun-Todd, 2004; Whalen et al., 1998; Williams et al., 2004). These findings highlight the importance of including measures of other emotional stimuli, in order for findings to be comparable across facial expressions, therefore ruling out possibilities of overall emotion-sensitive responses in favour of fear-discriminative ones. They also introduce issues related to inconsistent bias effects found in response to expressions of anger, and the way in which these findings are accounted for by the current threat bias account.

Another important (and also relevant) consideration is biases that have been found not only for fearful faces, but also for facial expressions of anger (Williams et al., 2004); to the extent that such findings have been dubbed the ‘anger superiority effect’ (Hansen and Hansen, 1988). For example, studies that observer biases for suppressed fearful faces also reported the same effects associated with angry expressions (Gray et al., 2013; Hansen and Hansen, 1988)). The way in which we consider these findings in relation to those associated with fearful faces raises important questions about how we define the threat bias. For example, whether the definition of the threat bias is exclusive to expressions of fear, or whether there may also be similar or different mechanisms that are also sensitive to angry faces. If this is the case, then we might expect similar bias responses to angry faces. This anger superiority effect is also challenged by findings that demonstrate a disadvantage for expressions of anger under certain conditions. Indeed, findings from Experiment 4 show that angry faces can be consistently slower to break suppression compared to fearful faces, suggesting that the differences that distinguish between expressions of fear and anger are more evident under certain experimental conditions than others. This effect was almost always true even when faces were manipulated, suggesting that the difference may be driven by differences in their low level image properties. Accounts of these a reverse anger superiority effects discuss the relevance of angry faces in representing unambiguous and direct threats, compared with fear expressions that ambiguously denote possible threats (for a review, see Stewart et al. (2012).

Inconsistent fear biases for inverted (and manipulated) faces

Studies that question the extent to which perceptual biases are driven by low-level stimulus information rather than emotional content of faces often include a condition whereby faces are inverted, or manipulated. This includes the inversion of faces by 180° and retaining luminance polarity, such that faces are identical in terms of contrast and spatial frequency content compared to their upright counterparts, but are rendered unrecognisable due to disruption of their configural content (Gray et al., 2013). Perceptual biases for upright fearful expressions that are also preserved when the same faces are manipulated therefore suggests that the effect relies substantially on the low-level image properties preserved under manipulation. However, not all studies that demonstrate fear biases also measure these at the inverted, or manipulated level, such that it can be difficult to extrapolate from findings the extent that effects are driven by stimulus properties contained within faces rather than their meaning, or emotional content. Bayle et al. (2011); Pourtois et al. (2006); Schupp et al. (2004); Vlamings et al. (2009); Vuilleumier et al. (2003); Whalen et al. (1998, 2001); Williams et al. (2004) and Carlson and Reinke (2008) are all examples of studies reporting some form of perceptual bias effect for fear expressions, but that do not include a measure of the same faces when inverted and/or manipulated. On the other hand, studies employing both upright and inverted and, or manipulated versions of stimuli report perceptual biases for fearful expressions at both stimulus orientations, such as Gray et al. (2013); Stein et al. (2014); Yang et al. (2007). Interestingly, these studies all employ the use of CFS techniques, suggesting another possible effect of task-relevance; though Jiang et al. (2007) show that these effects were not preserved for inverted faces under CFS. Finally, studies that include both upright and inverted or manipulated facial stimuli, but that do *not* find fear biases consistent under conditions of manipulation, include those of saccadic latency (Bannerman et al., 2012, 2009b) and overt (Holmes et al., 2005) and covert (Phelps et al., 2006) spatial attention.

There are therefore outstanding inconsistencies across studies regarding the extent to which perceptual biases for fearful expressions are driven by low level image information rather than the content of the percept. Findings from the present studies do not provide any clear explanation for such inconsistencies between inverted and upright findings. What they do show, however, is that the process by which the configural content in facial stimuli is disrupted has different effects on image salience depending on whether faces are simply inverted, or also subjected to retained luminance polarity. This is because face images subjected to retained luminance polarity are assigned a negated appearance, due to the way in which the light components of the face are converted to dark and vice versa. Haun and Peli (2013) suggest that this effect enhances the apparent salience associated with darker regions, therefore facilitating biases for these versions of faces compared to their upright, or indeed perhaps their inverted-only versions.

12.2.8 Other experimental features to consider

Findings from the present studies highlight the importance of transparency when reporting the experimental conditions and process by which stimuli have been normalised and standardised for experimental use. Nuanced differences between experimental design may result in significant outcomes regarding behavioural effects, as has been evidenced by the effects observed as a result of contrast normalisation. Such transparency is necessary if we are to understand the experimental conditions where a threat bias for fearful faces might be found. It is also important to consider the extent that a threat bias in perception that is exclusive to fear expressions occurs as a genuine adaptive response, or, whether it is reflective of an overall emotion bias.

Definitions of spatial frequency

Across psychophysical studies of the threat bias, conditions for filtering facial stimuli according to "high" and "low" frequency ranges are inconsistently defined. This was illustrated in Section 4.3, Chapter 4. This is related to the viewing distance and therefore the retinal size of facial stimuli. Because image statistics including spatial frequency vary relative to the viewing distance of the observer (Torralba and Oliva, 2003), it has been argued that a standard viewing distance that mimics that of every day social interactions ought to be a standard feature of experimental design (Fiorentini et al., 1983; Smith et al., 2005; Torralba and Oliva, 2003). Additionally, because spatial filtering effects of faces will vary according to viewing distance, it has also been argued that a standard metric for frequency content termed cycles per face-width be used as the default reference to faces' spatial content, as opposed to cycles per degree; the value for which depends on the size and viewing distance of the stimulus, and the orientation of measurement (Fiorentini et al., 1983). This would also reduce the way in which the terms cycles per face-width and cycles per degree are used interchangeably between studies. It is perhaps worth arguing here, however, that a threat bias for fear expressions ought to be effective across a range of viewing distances, and so this should be represented across frequency ranges, as opposed to a threat bias that has poor distal resolution. However, if the threat bias is related to the contrast sensitivity function, which is defined in angular rather than face units, then its effect will depend on viewing distance.

Related to this issue is the way in which spatial frequency is subjectively subdivided according to 'low' and 'high' ranges of information. An example is the way in which low spatial frequency is defined by Vuilleumier et al. (2003) as $f < 6_{cpd}$ but $f < .08_{cpd}$ by Bannerman et al. (2012), where this low frequency cut off used by Vuilleumier et al. (2003) is more closely matched to the high frequency cut off used by Stein et al. (2014) ($f > 6_{cpd}$. Such overlaps in the definition of what constitutes and differentiates low and high spatial frequencies can make it difficult to establish precisely where on the frequency spectrum an effect is expected to occur. Equally, it could be argued that there naturally exists a standard 'upper-limit' for high frequency information in face stimuli. For example, contrsat dissipates by around 27 cycles per face-width (Hayes et al., 1986; Näsänen, 1999).

Face databases

Related to the topic of methodological transparency is the importance of understanding pre-existing differences between facial expressions extracted from different face databases. Often it is the case that validation papers for such image databases include a description of the way in which raw photographs were subjected to image processing. This includes procedures typical for image standardisation such as cropping images to include internal features only, but also factors such as whether whiteness-correction was used, and whether images were normalised for luminance or contrast content. Findings from Experiment 3 demonstrate limitations for interpreting findings that can occur as a result of this information being absent. The fact that patterns of effective contrast differences between expressions were different according to the face database to which faces belonged suggests that different procedures employed during face stimulus creation can significantly influence the physical composition of these images. Such differences are visible simply by glancing at the Fourier amplitude slopes for expressions that differ between databases. Obtaining information about the way in which face stimuli are created, and stating the precise processes by which these are adjusted under experimental conditions is therefore a useful way to ensure that image properties contained within faces are understood and controlled for.

Threshold versus suprathreshold

Hedger et al. (2016) recently proposed that mixed findings of a fear expression might be influenced by the extent that different studies present fearful face stimuli at variable degrees of visibility that may be more or less likely to detect a threat bias. For example, certain experimental procedures such as the contrast sensitivity task employed for Experiment 3 present stimuli only at threshold level. It is possible that such lower contrast levels for face stimuli may weaken signals crucial for eliciting bias effects at the behavioural level. It may be that facial stimuli require a degree of stimulus salience in order to pose as realistic threat-information (Hedger et al., 2016). This may also differ according to the experimental paradigm used. For example, contrast thresholds are naturally lower for *detection* compared to *recognition* tasks (Näsänen, 1999). It is therefore important to consider the extent that a given bias effect might be expected at a detection versus recognition level, and according to this, which is most ideal experimental measure. Thresholds for detecting biases may be lower for certain participant populations compared to others, where threat bias effects may be more detectable for sub-clinical anxiety populations compared to those representing general non-anxious populations (Fox et al., 2002).

Artefacts associated with "raw" facial stimuli

It is important to note here that it is not possible to know with absolute certainty the exact properties that compose faces constituted as “raw” facial stimuli in the present studies. Such “raw” facial stimuli were defined as such based on the absence of contrast normalisation, spatial filtering and assumed absence of whiteness and colour correction. However, the amount known about the exact properties of these face photographs is limited to the amount of information provided on their associated database confirmation and/or creation manuscripts. This very issue was highlighted in Experiment 4, where it was not possible to determine whether or not images from each database had been subjected to the same image standardisation procedures. Such information is also related to the parameters and calibration camera equipment, not just the procedures that raw images are subjected to during standardisation. In the absence of a full account of such information, it is not possible to determine with precise certainty the extent that these faces are truly "raw" images at the point at which they undergo image analyses. Also noteworthy here are the limitations of fast Fourier transformations (FFT) of stimuli, including images of faces. Because a FFT algorithm relies on faces being a size that is the power of 2 (i.e. 256/128 pixels), for images not of this size padding methods can impose pixels with properties of 0. The outcome of FFT of such images may distort conclusions about the genuine properties of an image, including those related to their Fourier amplitude spectra. Finally, another relevant factor to consider is the extent that Gamma-corrected facial stimuli experienced by observers can be considered identical, at the stimulus level, to those that underwent image analyses. However, it is not necessarily the case that this limits the degree that findings from analyses of raw faces can be applied to an observer's perception of the same face. Gamma correction used in Experiment 5 is an example of how this method is used to transform images in a way that ensures they are presented to observers in their "true", physical format i.e. constituting the same properties that underwent image analyses, including FFTs.

12.3 Avenues for future research

12.3.1 Local information in fearful expressions

Another way to investigate mechanisms that might underpin the perception of fearful facial expressions is principle component analysis (PCA). PCA refers to a type of image-based analysis used to understand the way in which facial expressions are structurally encoded during cognition (Calder et al., 2000). This method has been applied to several aspects of face perception, including facial identity recognition (Burton et al., 1999; Calder et al., 2000), and also the coding of different facial expressions (Calder et al., 2001, 2000). This includes the identification of certain regularities in terms of the image features that define and vary between facial expressions, such as shape information related to feature positions within faces (Calder et al., 2000). An important question here is the extent that structural encoding of facial expressions can be investigated in terms of part-based versus whole-face approaches to facial expression processing. This is related to ongoing questions regarding the role of global and focal areas of faces during the cognitive categorisation of such stimuli, or other words, the extent that facial expressions are processed according to holistic over categorical processing. Many studies discuss the notion that the physical configuration of facial expressions suggest that distinct facial regions are key to facial emotion processing (Ellison and Massaro, 1997; Susskind et al., 2008). A particular focus here has been on local regions of the mouth and eyes that vary substantially between facial expressions. For fearful expressions in particular, the eye region has been designated an especially important source of information (Ellison and Massaro, 1997; Elsherif et al., 2017; Yang et al., 2007). Studies that present only the eye region of fearful facial expressions, for example, show that these segments of face stimuli are more quickly detected compared to the same regions from happy and neutral faces (Yang et al., 2007), and also receive selective central amygdala responses when backward masked from observers conscious visual awareness (Whalen et al., 2004).

Such findings have been attributed to the increased surface area of the sclera associated with fearful eyes; an effect that has been hypothesised as being driven by high contrast unique to the eye region (Gray et al., 2013; Yang et al., 2007). Consistent with this argument, Elsherif et al. (2017) argue that perceptual biases for the eye regions of fearful expressions may also be driven by increased low spatial frequency power that is highest for this region of the face. These factors are thought to facilitate rapid detection of fear expressions in the periphery, simply due to enhanced signals emitted from the sclera and iris (Keil, 2008; Lee et al., 2013). Expression-related differences between facial regions' statistical image properties may therefore be an interesting scope for future research.

12.3.2 Adaptations of Experiments 4 and 5

Experiment 4 shows that the visibility of expressions is related to their spatial frequency content but also the contrast metric used to normalise faces. Broadband stimuli were associated with bias effects for detecting fear expressions under CFS but only when these are first normalised for RMS contrast. When normalised psychophysically for perceived contrast, this fear bias effect diminishes. It is, however, not possible from results of these experimental conditions alone to unequivocally argue that threat bias effects under CFS are determined by physical contrast normalisation. This is on the basis that normalising stimuli for RMS contrast allows an understanding of how artificial changes in physical and perceived contrast affects faces' visibility, while normalising faces psychophysically allows an understanding of the extent that visibility differences between expressions remain when these artificial enhancements are removed; neither condition allow a measure of expression-related differences in visibility that occur when faces simply contain natural variations of physical contrast. In other words, to directly answer the question of the extent that fear biases are driven by contrast normalisation effects under CFS, it is necessary to understand how differences in visibility behave between expressions when faces are presented in their natural and raw format, such that their contrast content has not been influenced in any way. It would therefore be useful to re-conduct a single condition of Experiment 4, in which response times to detect faces under intraocular suppression are measured for faces when these are presented in normal broadband, low, mid-range and high frequency form, but importantly when their contrast content has not been adjusted accordingly.

Experiment 5 showed a consistent high frequency fear bias effect that was depedent on face image first having been normalised for RMS stimuli. However, an unexpected finding here, or lack thereof, was the absense of an effct at broad and low frequency conditions. Bannerman et al. (2012) demonstrated evidence in support of the low-frequency fear bias account, where low frequency fear expressions not normalised for contrast received faster saccadic responses compared to happy and neutral counterparts. Both Bannerman et al. (2010) and Bannerman et al. (2012) claim that such low frequency fear biases are short lived phenomena, such that they are detected only at rapid stimulus exposures of less than 100 ms. This notion is also upheld by findings from Holmes et al. (2005), where fear bias effects were again shown to dissipate at longer presentation durations exceeding 500 to 1,000 ms. A reasonable follow-on experiment to Experiment 5 would therefore be to include two conditions for stimulus duration, corresponding to the same 20 and 100 ms exposures employed by Bannerman et al. (2012). This is a particularly interesting effect, as if it is the case that attentional responses to information present in fear expressions operate at at very early stages of attentional allocation (Bannerman et al., 2012, 2010; Holmes et al., 2005), it is not evident as to why, particularly for saccadic response data, these behaviours are not detectable within the time frame where eye movements are recorded.

12.4 Summary of the present thesis

Findings from the present thesis contribute to the current consensus that face category images naturally posses a steeper amplitude slope compared to natural image categories Redies et al. (2007), adding that such differences also occur *within* face categories and *between* facial displays of emotion. This includes differences in physical contrast, and the relationship between such contrast energy and spatial frequency profile that differ between facial expressions. Findings also add that such physical differences are influenced by, and vary according to spatial filtering. These findings are necessary and preliminary information for understanding the way in which certain experimental techniques, including spatial filtering and contrast normalisation, may mask or in some way influence such naturally occurring low-level differences. Findings from Experiment 2 show that facial expressions naturally differ in terms of their perceived image salience, and demonstrate that the process by which they are normalised for physical contrast does not necessarily guarantee a result whereby face expressions do not differ in terms of their perceived image salience. These findings are particularly relevant to effects of contrast normalisation on the salience of facial stimuli. Findings from Experiment 3 demonstrate that visual contrast thresholds do not differ significantly according to facial expression, and that effective contrast in fear expressions is driven primarily by artificial increases in RMS contrast for these faces, and that this is broadly consistent across several face databases. Findings from Experiments 4 and 5 both suggest that threat bias effects manifest differently according to experimental condition, but importantly demonstrate that CFS in particular appears vulnerable to effects of contrast normalisation that artificially drive CFS fear biases. Together, these findings provide new and relevant evidence to the current understanding of the physical properties that define and differ between face expressions, and the way in which these can be expected to vary according to specific experimental and methodological factors.

References

Adolphs, R. (2003). Cognitive neuroscience: Cognitive neuroscience of human social behaviour. *Nature Reviews Neuroscience*, 4(3):165–178.

Adolphs, R., Gosselin, F., Buchanan, T. W., Tranel, D., Schyns, P., and Damasio, A. R. (2005). A mechanism for impaired fear recognition after amygdala damage. *Nature*, 433(7021):68–72.

Alais, D. (2012). Binocular rivalry: competition and inhibition in visual perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, 3(1):87–103.

Alpers, G. W. and Gerdes, A. (2007). Here is looking at you: emotional faces predominate in binocular rivalry. *Emotion*, 7(3):495–506.

Armony, J. L. and Dolan, R. J. (2002). Modulation of spatial attention by fear-conditioned stimuli: an event-related fmri study. *Neuropsychologia*, 40(7):817–826.

Atick, J. J. and Redlich, A. N. (1992). What does the retina know about natural scenes? *Neural computation*, 4(2):196–210.

Attneave, F. (1954). Some informational aspects of visual perception. *Psychological review*, 61(3):183–193.

Baddeley, A. (2003). Working memory and language: An overview. *Journal of communication disorders*, 36(3):189–208.

Baker, D. H. and Graf, E. W. (2009). Natural images dominate in binocular rivalry. *Proceedings of the National Academy of Sciences*, 106(13):5436–5441.

Bannerman, R. L., Hibbard, P. B., Chalmers, K., and Sahraie, A. (2012). Saccadic latency is modulated by emotional content of spatially filtered face stimuli. *Emotion*, 12(6):1384.

Bannerman, R. L., Milders, M., De Gelder, B., and Sahraie, A. (2009a). Orienting to threat: faster localization of fearful facial expressions and body postures revealed by saccadic eye movements. *Proceedings of the Royal Society of London B: Biological Sciences*, pages 1635–1641.

Bannerman, R. L., Milders, M., De Gelder, B., and Sahraie, A. (2009b). Orienting to threat: faster localization of fearful facial expressions and body postures revealed by saccadic eye movements. *Proceedings of the Royal Society of London B: Biological Sciences*, pages 1635–1641.

Bannerman, R. L., Milders, M., and Sahraie, A. (2010). Attentional bias to brief threat-related faces revealed by saccadic eye movements. *Emotion*, 10(5):733–738.

Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, 5(8):617–629.

Bayle, D. J., Henaff, M.-A., and Krolak-Salmon, P. (2009). Unconsciously perceived fear in peripheral vision alerts the limbic system: a meg study. *PloS one*, 4(12):e8207.

Bayle, D. J., Schoendorff, B., Hénaff, M.-A., and Krolak-Salmon, P. (2011). Emotional facial expression detection in the peripheral visual field. *PloS one*, 6(6):e21584.

Beaupré, M. G. and Hess, U. (2005). Cross-cultural emotion recognition among canadian ethnic groups. *Journal of cross-cultural psychology*, 36(3):355–370.

Beck, A. T. and Clark, D. A. (1997). An information processing model of anxiety: Automatic and strategic processes. *Behaviour research and therapy*, 35(1):49–58.

Berson, D. M. (1988). Retinal and cortical inputs to cat superior colliculus: composition, convergence and laminar specificity. In *Progress in brain research*, volume 75, pages 17–26. Elsevier.

Bex, P. J. and Langley, K. (2007). The perception of suprathreshold contrast and fast adaptive filtering. *Journal of Vision*, 7(12):1, 1–23.

Bex, P. J. and Makous, W. (2002). Spatial frequency, phase, and the contrast of natural images. *Journal of the Optical Society of America A*, 19(6):1096–1106.

Bieniek, M. M., Pernet, C. R., and Rousselet, G. A. (2012). Early erps to faces and objects are driven by phase, not amplitude spectrum information: evidence from parametric, test-retest, single-subject analyses. *Journal of vision*, 12(13):12, 1–24.

Blake, R. and Logothetis, N. K. (2002). Visual competition. *Nature Reviews Neuroscience*, 3(1):13–21.

Blaukopf, C. L. and DiGirolamo, G. J. (2005). The automatic extraction and use of information from cues and go signals in an antisaccade task. *Experimental brain research*, 167(4):654–659.

Bocanegra, B. R., Huijding, J., and Zeelenberg, R. (2012). Beyond attentional bias: A perceptual bias in a dot-probe task. *Emotion*, 12(6):1362–1366.

Bocanegra, B. R. and Zeelenberg, R. (2009). Emotion improves and impairs early vision. *Psychological science*, 20(6):707–713.

Bompas, A. and Sumner, P. (2008). Sensory sluggishness dissociates saccadic, manual, and perceptual responses: An s-cone study. *Journal of Vision*, 8(8):10–10.

Brainard, D. H. and Vision, S. (1997). The psychophysics toolbox. *Spatial vision*, 10:433–436.

Breiter, H. C., Etcoff, N. L., Whalen, P. J., Kennedy, W. A., Rauch, S. L., Buckner, R. L., Strauss, M. M., Hyman, S. E., and Rosen, B. R. (1996). Response and habituation of the human amygdala during visual processing of facial expression. *Neuron*, 17(5):875–887.

Breitmeyer, B., Ogmen, H., and Öğmen, H. (2006). *Visual masking: Time slices through conscious and unconscious vision*. Number 41. Oxford University Press.

Bullier, J. (2001). Integrated model of visual processing. *Brain research reviews*, 36(2-3):96–107.

Burton, A. M., Bruce, V., and Hancock, P. J. (1999). From pixels to people: A model of familiar face recognition. *Cognitive Science*, 23(1):1–31.

Calder, A. J., Burton, A. M., Miller, P., Young, A. W., and Akamatsu, S. (2001). A principal component analysis of facial expressions. *Vision research*, 41(9):1179–1208.

Calder, A. J., Young, A. W., Keane, J., and Dean, M. (2000). Configural information in facial expression perception. *Journal of Experimental Psychology: Human perception and performance*, 26(2):527.

Campbell, F. W. and Robson, J. (1968). Application of fourier analysis to the visibility of gratings. *The Journal of physiology*, 197(3):551–566.

Carlson, J. M. and Reinke, K. S. (2008). Masked fearful faces modulate the orienting of covert spatial attention. *Emotion*, 8(4):522–529.

Carter, O. and Cavanagh, P. (2007). Onset rivalry: Brief presentation isolates an early independent phase of perceptual competition. *PloS one*, 2(4):e343.

Costen, N. P., Parker, D. M., and Craw, I. (1994). Spatial content and spatial quantisation effects in face recognition. *Perception*, 23(2):129–146.

Costen, N. P., Parker, D. M., and Craw, I. (1996). Effects of high-pass and low-pass spatial filtering on face identification. *Perception & Psychophysics*, 58(4):602–612.

Crouzet, S. M., Kirchner, H., and Thorpe, S. J. (2010). Fast saccades toward faces: face detection in just 100 ms. *Journal of vision*, 10(4):16, 1–17.

Crouzet, S. M. and Thorpe, S. J. (2011). Low-level cues and ultra-fast face detection. *Frontiers in psychology*, 2:342.

Davis, M. and Whalen, P. J. (2001). The amygdala: vigilance and emotion. *Molecular psychiatry*, 6(1):13–34.

Day-Brown, J. D., Wei, H., Chomsung, R. D., Petry, H. M., and Bickford, M. E. (2010). Pulvinar projections to the striatum and amygdala in the tree shrew. *Frontiers in neuroanatomy*, 4:143.

De Gelder, B., Morris, J. S., and Dolan, R. J. (2005). Unconscious fear influences emotional awareness of faces and voices. *Proceedings of the National Academy of Sciences*, 102(51):18682–18687.

De Gelder, B., Vroomen, J., Pourtois, G., and Weiskrantz, L. (1999). Non-conscious recognition of affect in the absence of striate cortex. *Neuroreport*, 10(18):3759–3763.

Dehaene, S. and Changeux, J.-P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron*, 70(2):200–227.

Doi, H. and Shinohara, K. (2013). Unconscious presentation of fearful face modulates electrophysiological responses to emotional prosody. *Cerebral Cortex*, 25(3):817–832.

Eastwood, J. D. and Smilek, D. (2005). Functional consequences of perceiving facial expressions of emotion without awareness. *Consciousness and cognition*, 14(3):565–584.

Eimer, M. and Holmes, A. (2002). An erp study on the time course of emotional face processing. *Neuroreport*, 13(4):427–431.

Ekman, P. (1976). Pictures of facial affect. *Consulting Psychologists Press*.

Ekman, P. (1999). Facial expressions. *Handbook of cognition and emotion*, 16:301–320.

Ekman, P. and Cordaro, D. (2011). What is meant by calling emotions basic. *Emotion review*, 3(4):364–370.

Ekman, P. and Friesen, W. V. (1975). Unmasking the face: A guide to recognizing emotions from facial cues.

Ellison, J. W. and Massaro, D. W. (1997). Featural evaluation, integration, and judgment of facial affect. *Journal of Experimental Psychology: Human Perception and Performance*, 23(1):213–226.

Elsherif, M. M., Saban, M. I., and Rotshtein, P. (2017). The perceptual saliency of fearful eyes and smiles: A signal detection study. *PloS one*, 12(3):e0173199.

Fahle, M. (1982). Binocular rivalry: Suppression depends on orientation and spatial frequency. *Vision Research*, 22(7):787–800.

Farah, M. J., Tanaka, J. W., and Drain, H. M. (1995). What causes the face inversion effect? *Journal of Experimental Psychology: Human perception and performance*, 21(3):628–634.

Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, 4(12):2379–2394.

Findlay, J. and Walker, R. (2012). Human saccadic eye movements. *Scholarpedia*, 7(7):5095.

Fiorentini, A., Maffei, L., and Sandini, G. (1983). The role of high spatial frequencies in face perception. *Perception*, 12(2):195–201.

Fox, E., Russo, R., Bowles, R., and Dutton, K. (2001). Do threatening stimuli draw or hold visual attention in subclinical anxiety? *Journal of Experimental Psychology: General*, 130(4):681–700.

Fox, E., Russo, R., and Dutton, K. (2002). Attentional bias for threat: Evidence for delayed disengagement from emotional faces. *Cognition & emotion*, 16(3):355–379.

García-Pérez, M. A., Alcalá-Quintana, R., Woods, R. L., and Peli, E. (2011). Psychometric functions for detection and discrimination with and without flankers. *Attention, Perception, & Psychophysics*, 73(3):829–853.

Gayet, S. and Stein, T. (2017). Between-subject variability in the breaking continuous flash suppression paradigm: potential causes, consequences, and solutions. *Frontiers in psychology*, 8:437.

Gayet, S., Van der Stigchel, S., and Paffen, C. L. (2014). Breaking continuous flash suppression: competing for consciousness on the pre-semantic battlefield. *Frontiers in Psychology*, 5:460.

Georgeson, M. and Sullivan, G. (1975). Contrast constancy: deblurring in human vision by spatial frequency channels. *The Journal of Physiology*, 252(3):627–656.

Georgiou, G., Bleakley, C., Hayward, J., Russo, R., Dutton, K., Eltiti, S., and Fox, E. (2005). Focusing on fear: Attentional disengagement from emotional faces. *Visual cognition*, 12(1):145–158.

Goffaux, V. and Rossion, B. (2006). Faces are" spatial"–holistic face perception is supported by low spatial frequencies. *Journal of Experimental Psychology: Human Perception and Performance*, 32(4):1023.

Gosselin, F. and Schyns, P. G. (2004). No troubles with bubbles: a reply to murray and gold. *Vision Research*, 44(5):471–477.

Gray, K. L., Adams, W. J., Hedger, N., Newton, K. E., and Garner, M. (2013). Faces and awareness: Low-level, not emotional factors determine perceptual dominance. *Emotion*, 13(3):537–544.

Halgren, E., Raij, T., Marinkovic, K., Jousmäki, V., and Hari, R. (2000). Cognitive response profile of the human fusiform face area as determined by meg. *Cerebral cortex*, 10(1):69–81.

Hansen, C. H. and Hansen, R. D. (1988). Finding the face in the crowd: an anger superiority effect. *Journal of personality and social psychology*, 54(6):917–924.

Hardcastle, V. G. (2003). Attention versus consciousness. *Neural basis of consciousness*, 49:105–120.

Haun, A. M. and Peli, E. (2013). Perceived contrast in complex images. *Journal of vision*, 13(13):3, 1–21.

Hayes, T., Morrone, M. C., and Burr, D. C. (1986). Recognition of positive and negative bandpass-filtered images. *Perception*, 15(5):595–602.

Hedger, N., Adams, W. J., and Garner, M. (2015). Fearful faces have a sensory advantage in the competition for awareness. *Journal of Experimental Psychology: Human Perception and Performance*, 41(6):1748–1757.

Hedger, N., Gray, K. L., Garner, M., and Adams, W. J. (2016). Are visual threats prioritized without awareness? a critical review and meta-analysis involving 3 behavioral paradigms and 2696 observers. *Psychological bulletin*, 142(9):934–968.

Heyman, T. and Moors, P. (2014). Frequent words do not break continuous flash suppression differently from infrequent or nonexistent words: implications for semantic processing of words in the absence of awareness. *Plos One*, 9(8):e104719.

Holmes, A., Green, S., and Vuilleumier, P. (2005). The involvement of distinct visual channels in rapid attention towards fearful facial expressions. *Cognition & Emotion*, 19(6):899–922.

Honey, C., Kirchner, H., and VanRullen, R. (2008). Faces in the cloud: Fourier power spectrum biases ultrarapid face detection. *Journal of vision*, 8(12):9, 1–13.

Horstmann, G. and Bauland, A. (2006). Search asymmetries with real faces: Testing the anger-superiority effect. *Emotion*, 6(2):193–207.

Hughes, M. S. (1991). Dictionary of eye terminology. *Archives of Ophthalmology*, 109(9):1208–1208.

Humphreys, G. W. and Bruce, V. (1989). *Visual cognition: Computational, experimental, and neuropsychological perspectives*. Psychology Press.

Hunt, A. R., Van Zoest, W., Kingstone, A., Nobre, A., and Coull, J. (2010). Attending to emerging representations: The importance of task context and time of response. *Attention and time*, pages 3–15.

Hunt, C., Keogh, E., and French, C. C. (2006). Anxiety sensitivity: The role of conscious awareness and selective attentional bias to physical threat. *Emotion*, 6(3):418–428.

Innocent, M. (2011). White balance correction using illuminant estimation. US Patent 8,049,789.

Itier, R. J. and Taylor, M. J. (2002). Inversion and contrast polarity reversal affect both encoding and recognition processes of unfamiliar faces: a repetition study using erps. *Neuroimage*, 15(2):353–372.

Jastorff, J., Huang, Y.-A., Giese, M. A., and Vandenbulcke, M. (2015). Common neural correlates of emotion perception in humans. *Human brain mapping*, 36(10):4184–4201.

Jiang, Y., Costello, P., and He, S. (2007). Processing of invisible stimuli: Advantage of upright faces and recognizable words in overcoming interocular suppression. *Psychological science*, 18(4):349–355.

Johns, P. (2014). *Clinical Neuroscience E-Book*. Elsevier Health Sciences.

Kaplan, E. and Shapley, R. M. (1986). The primate retina contains two types of ganglion cells, with high and low contrast sensitivity. *Proceedings of the National Academy of Sciences*, 83(8):2755–2757.

Kapp, B. S., Supple, W. F., and Whalen, P. J. (1994). Effects of electrical stimulation of the amygdaloid central nucleus on neocortical arousal in the rabbit. *Behavioral neuroscience*, 108(1):81–93.

Keil, M. S. (2008). Does face image statistics predict a preferred spatial frequency for human face processing? *Proceedings of the Royal Society of London B: Biological Sciences*, 275(1647):2095–2100.

Kiefer, M. and Brendel, D. (2006). Attentional modulation of unconscious “automatic” processes: Evidence from event-related potentials in a masked priming paradigm. *Journal of Cognitive Neuroscience*, 18(2):184–198.

Killgore, W. D. and Yurgelun-Todd, D. A. (2004). Activation of the amygdala and anterior cingulate during nonconscious processing of sad versus happy faces. *Neuroimage*, 21(4):1215–1223.

Kirchner, H. and Thorpe, S. J. (2006). Ultra-rapid object detection with saccadic eye movements: Visual processing speed revisited. *Vision research*, 46(11):1762–1776.

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., Broussard, C., et al. (2007). What’s new in psychtoolbox-3. *Perception*, 36(14):1.

Koch, C. and Tsuchiya, N. (2007). Attention and consciousness: two distinct brain processes. *Trends in cognitive sciences*, 11(1):16–22.

Koch, G., Oliveri, M., Torriero, S., and Caltagirone, C. (2005). Modulation of excitatory and inhibitory circuits for visual awareness in the human right parietal cortex. *Experimental brain research*, 160(4):510–516.

Koster, E. H., Crombez, G., Verschuere, B., and De Houwer, J. (2004). Selective attention to threat in the dot probe paradigm: Differentiating vigilance and difficulty to disengage. *Behaviour research and therapy*, 42(10):1183–1192.

Koster, E. H., Crombez, G., Verschuere, B., Vanvolsem, P., and De Houwer, J. (2007). A time-course analysis of attentional cueing by threatening scenes. *Experimental psychology*, 54(2):161–171.

Lang, P. J., Bradley, M. M., Cuthbert, B. N., et al. (1997). Motivated attention: Affect, activation, and action. *Attention and orienting: Sensory and motivational processes*, 97:135.

Lang, P. J., Davis, M., and Öhman, A. (2000). Fear and anxiety: animal models and human cognitive psychophysiology. *Journal of affective disorders*, 61(3):137–159.

Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D. H., Hawk, S. T., and Van Knippenberg, A. (2010). Presentation and validation of the radboud faces database. *Cognition and emotion*, 24(8):1377–1388.

Lavie, N., Hirst, A., De Fockert, J. W., and Viding, E. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology: General*, 133(3):339–354.

LeDoux, J. E. (1994). Emotion, memory and the brain. *Scientific American*, 270(6):50–57.

LeDoux, J. E. (2012). *Evolution of human emotion: a view through fear*, volume 195. Elsevier.

LeDoux, J. E. and Phelps, E. A. (1993). Emotional networks in the brain. *Handbook of emotions*, 109:157–172.

Lee, D. H., Susskind, J. M., and Anderson, A. K. (2013). Social transmission of the sensory benefits of eye widening in fear expressions. *Psychological Science*, 24(6):957–965.

Lee, T.-H., Baek, J., Lu, Z.-L., and Mather, M. (2014). How arousal modulates the visual contrast sensitivity function. *Emotion*, 14(5):978–984.

Levelt, W. J. (1965). *On binocular rivalry*. PhD thesis, Van Gorcum Assen.

Lin, Z. and He, S. (2009). Seeing the invisible: The scope and limits of unconscious processing in binocular rivalry. *Progress in neurobiology*, 87(4):195–211.

Liu, C. H., Collin, C. A., Rainville, S. J., and Chaudhuri, A. (2000). The effects of spatial frequency overlap on face recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 26(3):956–979.

Livingstone, M. and Hubel, D. (1988). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science*, 240(4853):740–749.

Ludwig, K. and Hesselmann, G. (2015). Weighing the evidence for a dorsal processing bias under continuous flash suppression. *Consciousness and Cognition*, 35:251–259.

Lundqvist, D., Flykt, A., and Öhman, A. (1998). The karolinska directed emotional faces (kdef). *CD ROM from Department of Clinical Neuroscience, Psychology section, Karolinska Institutet*, 91:630.

Lunghi, C., Alais, D., et al. (2016). The temporal frequency tuning of continuous flash suppression reveals peak suppression at very low frequencies. *Scientific reports*, 6:35723.

MacLeod, C. and Mathews, A. (1988). Anxiety and the allocation of attention to threat. *The Quarterly Journal of Experimental Psychology Section A*, 40(4):653–670.

MacNamara, A., Kappenman, E. S., Black, S. R., Bress, J. N., and Hajcak, G. (2013). Integrating behavioral and electrocortical measures of attentional bias toward threat. *Handbook of self-regulatory processes in development: New directions and international perspectives*, pages 215–242.

Meese, T. S., Baker, D. H., and Summers, R. J. (2017). Perception of global image contrast involves transparent spatial filtering and the integration and suppression of local contrasts (not rms contrast). *Royal Society open science*, 4(9):170285.

Méndez-Bértolo, C., Moratti, S., Toledano, R., Lopez-Sosa, F., Martínez-Alvarez, R., Mah, Y. H., Vuilleumier, P., Gil-Nagel, A., and Strange, B. A. (2016). A fast pathway for fear in human amygdala. *Nature neuroscience*, 19(8):1041–1049.

Milner, D. and Goodale, M. (2006). *The visual brain in action*. Oxford University Press.

Mogg, K. and Bradley, B. P. (2002). Selective orienting of attention to masked threat faces in social anxiety. *Behaviour research and therapy*, 40(12):1403–1414.

Mogg, K., Philippot, P., and Bradley, B. P. (2004). Selective attention to angry faces in clinical social phobia. *Journal of abnormal psychology*, 113(1):160–165.

Morris, J. S., DeGelder, B., Weiskrantz, L., and Dolan, R. J. (2001). Differential extrageniculostriate and amygdala responses to presentation of emotional faces in a cortically blind field. *Brain*, 124(6):1241–1252.

Morris, J. S., Öhman, A., and Dolan, R. J. (1998). Conscious and unconscious emotional learning in the human amygdala. *Nature*, 393(6684):467–470.

Morris, J. S., Öhman, A., and Dolan, R. J. (1999). A subcortical pathway to the right amygdala mediating “unseen” fear. *Proceedings of the National Academy of Sciences*, 96(4):1680–1685.

Mueller, T. and Blake, R. (1989). A fresh look at the temporal dynamics of binocular rivalry. *Biological cybernetics*, 61(3):223–232.

Murray, R. F. and Gold, J. M. (2004). Troubles with bubbles. *Vision Research*, 44(5):461–470.

Näsänen, R. (1999). Spatial frequency bandwidth used in the recognition of facial images. *Vision research*, 39(23):3824–3833.

Oatley, K. and Johnson-Laird, P. N. (1987). Towards a cognitive theory of emotions. *Cognition and emotion*, 1(1):29–50.

Öhman, A. and Mineka, S. (2001a). Fears, phobias, and preparedness: toward an evolved module of fear and fear learning. *Psychological review*, 108(3):483–522.

Öhman, A. and Mineka, S. (2001b). Fears, phobias, and preparedness: toward an evolved module of fear and fear learning. *Psychological review*, 108(3):483–522.

Ohman, A. and Soares, J. (1994). " unconscious anxiety": phobic responses to masked stimuli. *Journal of Abnormal Psychology*, 103(2):231–240.

Ojanpää, H. and Näsänen, R. (2003). Utilisation of spatial frequency information in face search. *Vision Research*, 43(24):2505–2515.

O’Hare, L. and Hibbard, P. B. (2011). Spatial frequency and visual discomfort. *Vision research*, 51(15):1767–1777.

Palermo, R. and Rhodes, G. (2007). Are you always on my mind? a review of how face perception and attention interact. *Neuropsychologia*, 45(1):75–92.

Pan, Y., Lin, B., Zhao, Y., and Soto, D. (2014). Working memory biasing of visual perception without awareness. *Attention, Perception, & Psychophysics*, 76(7):2051–2062.

Peli, E. (1990). Contrast in complex images. *Journal of the Optical Society of America A*, 7(10):2032–2040.

Peli, E. (1997). In search of a contrast metric: Matching the perceived contrast of gabor patches at different phases and bandwidths. *Vision Research*, 37(23):3217–3224.

Peli, E., Arend, L., and Labianca, A. T. (1996). Contrast perception across changes in luminance and spatial frequency. *Journal of the Optical Society of America A*, 13(10):1953–1959.

Pelli, D. G. (1997). The videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial vision*, 10(4):437–442.

Pessoa, L. and Adolphs, R. (2010). Emotion processing and the amygdala: from a'low road'to'many roads' of evaluating biological significance. *Nature reviews neuroscience*, 11(11):773–783.

Pessoa, L., McKenna, M., Gutierrez, E., and Ungerleider, L. (2002). Neural processing of emotional faces requires attention. *Proceedings of the National Academy of Sciences*, 99(17):11458–11463.

Phelps, E. A., Ling, S., and Carrasco, M. (2006). Emotion facilitates perception and potentiates the perceptual benefits of attention. *Psychological science*, 17(4):292–299.

Phillips, M. L., Williams, L. M., Heining, M., Herba, C. M., Russell, T., Andrew, C., Bullmore, E. T., Brammer, M. J., Williams, S. C., Morgan, M., et al. (2004). Differential neural responses to overt and covert presentations of facial expressions of fear and disgust. *Neuroimage*, 21(4):1484–1496.

Posner, M. I., Snyder, C. R., and Davidson, B. J. (1980). Attention and the detection of signals. *Journal of experimental psychology: General*, 109(2):160–174.

Pourtois, G., Schwartz, S., Seghier, M. L., Lazeyras, F., and Vuilleumier, P. (2005). Portraits or people? distinct representations of face identity in the human visual cortex. *Journal of Cognitive Neuroscience*, 17(7):1043–1057.

Pourtois, G., Schwartz, S., Seghier, M. L., Lazeyras, F., and Vuilleumier, P. (2006). Neural systems for orienting attention to the location of threat signals: an event-related fmri study. *Neuroimage*, 31(2):920–933.

Prins, N. et al. (2016). *Psychophysics: a practical introduction*. Academic Press.

Quaia, C., Lefèvre, P., and Optican, L. M. (1999). Model of the control of saccades by superior colliculus and cerebellum. *Journal of Neurophysiology*, 82(2):999–1018.

Redies, C., Hänisch, J., Blickhan, M., and Denzler, J. (2007). Artists portray human faces with the fourier statistics of complex natural scenes. *Network: Computation in Neural Systems*, 18(3):235–248.

Rigoulot, S., D'hondt, F., Defoort-Dhellemmes, S., Despretz, P., Honoré, J., and Sequeira, H. (2011). Fearful faces impact in peripheral vision: behavioral and neural evidence. *Neuropsychologia*, 49(7):2013–2021.

Ritchie, K. L., Bannerman, R. L., Turk, D. J., and Sahraie, A. (2013). Eye rivalry and object rivalry in the intact and split-brain. *Vision research*, 91:102–107.

Robinson, D. L. and Petersen, S. E. (1992). The pulvinar and visual salience. *Trends in neurosciences*, 15(4):127–132.

Rousselet, G. A., Macé, M. J.-M., and Fabre-Thorpe, M. (2003). Is it an animal? is it a human face? fast processing in upright and inverted natural scenes. *Journal of vision*, 3(6):5, 440–455.

Santos, N. A. d., Alencar, C. C. G., and Dias, Y. H. N. (2009). Contrast sensitivity function of sine-wave gratings in children with acute malnutrition. *Psychology & neuroscience*, 2(1):11–15.

Schupp, H. T., Öhman, A., Junghöfer, M., Weike, A. I., Stockburger, J., and Hamm, A. O. (2004). The facilitated processing of threatening faces: an erp analysis. *Emotion*, 4(2):189–200.

Schyns, P. G. and Oliva, A. (1999). Dr. angry and mr. smile: When categorization flexibly modifies the perception of faces in rapid visual presentations. *Cognition*, 69(3):243–265.

Sklar, A. Y., Levy, N., Goldstein, A., Mandel, R., Maril, A., and Hassin, R. R. (2012). Reading and doing arithmetic nonconsciously. *Proceedings of the National Academy of Sciences*, 109(48):19614–19619.

Smith, F. W. and Rossit, S. (2018). Identifying and detecting facial expressions of emotion in peripheral vision. *PloS one*, 13(5):e0197160.

Smith, F. W. and Schyns, P. G. (2009). Smile through your fear and sadness: Transmitting and identifying facial expression signals over a range of viewing distances. *Psychological Science*, 20(10):1202–1208.

Smith, M. L., Cottrell, G. W., Gosselin, F., and Schyns, P. G. (2005). Transmitting and decoding facial expressions. *Psychological science*, 16(3):184–189.

Stein, T., Hebart, M. N., and Sterzer, P. (2011). Breaking continuous flash suppression: a new measure of unconscious processing during interocular suppression? *Frontiers in human neuroscience*, 5:167.

Stein, T., Seymour, K., Hebart, M. N., and Sterzer, P. (2014). Rapid fear detection relies on high spatial frequencies. *Psychological science*, 25(2):566–574.

Stein, T. and Sterzer, P. (2014). Unconscious processing under interocular suppression: getting the right measure. *Frontiers in Psychology*, 5:387.

Stewart, L. H., Ajina, S., Getov, S., Bahrami, B., Todorov, A., and Rees, G. (2012). Unconscious evaluation of faces on social dimensions. *Journal of Experimental Psychology: General*, 141(4):715–727.

Stojanoski, B. and Cusack, R. (2014). Time to wave good-bye to phase scrambling: creating controlled scrambled images using diffeomorphic transformations. *Journal of vision*, 14(12):6, 1–16.

Stormark, K. M., Nordby, H., and Hugdahl, K. (1995). Attentional shifts to emotionally charged cues: Behavioural and erp data. *Cognition & Emotion*, 9(5):507–523.

Sunness, J. S., Rubin, G. S., Applegate, C. A., Bressler, N. M., Marsh, M. J., Hawkins, B. S., and Haselwood, D. (1997). Visual function abnormalities and prognosis in eyes with age-related geographic atrophy of the macula and good visual acuity. *Ophthalmology*, 104(10):1677–1691.

Susskind, J. M., Lee, D. H., Cusi, A., Feiman, R., Grabski, W., and Anderson, A. K. (2008). Expressing fear enhances sensory acquisition. *Nature neuroscience*, 11(7):843–850.

Tamietto, M. and De Gelder, B. (2010). Neural bases of the non-conscious perception of emotional signals. *Nature Reviews Neuroscience*, 11(10):697–709.

Tanaka, J. W. and Farah, M. J. (2003). The holistic representation of faces. *Perception of faces, objects, and scenes: Analytic and holistic processes*, pages 53–74.

Theeuwes, J., Kramer, A. F., Hahn, S., and Irwin, D. E. (1998). Our eyes do not always go where we want them to go: Capture of the eyes by new objects. *Psychological Science*, 9(5):379–385.

Tolhurst, D., Tadmor, Y., and Chao, T. (1992). Amplitude spectra of natural images. *Ophthalmic and Physiological Optics*, 12(2):229–232.

Torralba, A. and Oliva, A. (2003). Statistics of natural image categories. *Network: computation in neural systems*, 14(3):391–412.

Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., Marcus, D. J., Westerlund, A., Casey, B., and Nelson, C. (2009). The nimstim set of facial expressions: judgments from untrained research participants. *Psychiatry research*, 168(3):242–249.

Tsuchiya, N. and Koch, C. (2004). Continuous flash suppression. *Journal of Vision*, 4(8):61–61.

Tsuchiya, N. and Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature neuroscience*, 8(8):1096–1101.

Tversky, B. and Hemenway, K. (1983). Categories of environmental scenes. *Cognitive psychology*, 15(1):121–149.

Van Opstal, F., De Loof, E., Verguts, T., and Cleeremans, A. (2016). Spontaneous eyeblinks during breaking continuous flash suppression are associated with increased detection times. *Journal of vision*, 16(14):21, 1–10.

Vlamings, P. H., Goffaux, V., and Kemner, C. (2009). Is the early modulation of brain activity by fearful facial expressions primarily mediated by coarse low spatial frequency information? *Journal of vision*, 9(5):12, 1–3.

Vuilleumier, P., Armony, J. L., Driver, J., and Dolan, R. J. (2003). Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nature neuroscience*, 6(6):624–631.

Wang, S., Eccleston, C., and Keogh, E. (2015). The role of spatial frequency information in the recognition of facial expressions of pain. *Pain*, 156(9):1670–1682.

Watson, A. B. and Ahumada, A. J. (2005). A standard model for foveal detection of spatial contrast. *Journal of vision*, 5(9):717–740.

Westheimer, G. (2001). The fourier theory of vision. *Perception*, 30(5):531–541.

Whalen, P. J., Kagan, J., Cook, R. G., Davis, F. C., Kim, H., Polis, S., McLaren, D. G., Somerville, L. H., McLean, A. A., Maxwell, J. S., et al. (2004). Human amygdala responsivity to masked fearful eye whites. *Science*, 306(5704):2061–2061.

Whalen, P. J., Rauch, S. L., Etcoff, N. L., McInerney, S. C., Lee, M. B., and Jenike, M. A. (1998). Masked presentations of emotional facial expressions modulate amygdala activity without explicit knowledge. *Journal of Neuroscience*, 18(1):411–418.

Whalen, P. J., Shin, L. M., McInerney, S. C., Fischer, H., Wright, C. I., and Rauch, S. L. (2001). A functional mri study of human amygdala responses to facial expressions of fear versus anger. *Emotion*, 1(1):70–83.

Williams, M. A., Morris, A. P., McGlone, F., Abbott, D. F., and Mattingley, J. B. (2004). Amygdala responses to fearful and happy facial expressions under conditions of binocular suppression. *Journal of Neuroscience*, 24(12):2898–2904.

Yang, E. and Blake, R. (2012). Deconstructing continuous flash suppression. *Journal of vision*, 12(3):8, 1–14.

Yang, E., Zald, D. H., and Blake, R. (2007). Fearful expressions gain preferential access to awareness during continuous flash suppression. *Emotion*, 7(4):882–886.

Yanof, A., Skow, M., Tran, H., and Adby, A. (2007). Digital image processing using white balance and gamma correction. US Patent 7,236,190.

Yiend, J. and Mathews, A. (2001). Anxiety and attention to threatening pictures. *The Quarterly Journal of Experimental Psychology: Section A*, 54(3):665–681.