

Variations in face perception across the visual field

A thesis presented for the degree of

Doctor of Philosophy

by

Anisa Yasmina Morsi

BSc (Hons) Cognitive Neuroscience and Psychology

Supervised by:

Dr John Greenwood

Dr Tessa Dekker

BBSRC London Interdisciplinary Biosciences PhD Programme

University College London

June 2023

Declaration of Authorship

I, Anisa Yasmina Morsi, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Face recognition is widely considered to be “special”, carried out in dedicated brain regions with distinct mechanisms compared to the recognition of other objects. Typically, low-level vision varies across the visual field in characteristic patterns, with our ability to see fine detail better along the horizontal vs. vertical meridian and in the lower vs. upper visual field. Given that faces appear to be processed uniquely, does face recognition vary across peripheral visual field locations in the same way? Using behavioural psychophysics, I uncover a clear advantage for face perception along the horizontal vs. vertical meridian, and a smaller-but-consistent advantage in the lower vs. upper field. Therefore, location influences face perception as it does in low-level vision. I then measured the retinotopic properties of three face-selective brain regions (the occipital face area, pFus and mFus), to determine whether visual field sampling within these regions could explain the variations in face perception. In all three areas there was a greater number of population receptive fields (pRFs) and better visual field coverage along the horizontal vs. vertical meridian and in the lower vs. upper field. These patterns resemble those observed in early visual cortex. Lastly, the process of localising and delineating face-selective areas of the brain was examined. Using novel functional localiser stimuli, I show that the delineation of face-selective brain areas, and the retinotopic properties subsequently measured from them, varies according to whether localiser stimuli were foveal or peripheral. Altogether, these findings demonstrate that the visual field anisotropies of low-level vision also affect high-level face recognition, with similar variations in retinotopic properties. This supports a hierarchical model of vision whereby spatial selectivity in higher-level areas is built upon the selectivity of lower regions, even within specialised face processing.

Impact Statement

We are experts at face recognition. A plethora of research indicates that this is because relative to general object recognition, faces are treated as “special” within the visual system, processed using unique mechanisms and in distinct parts of the brain. The central aim of this thesis was to investigate how the specialised processing of faces is linked to more basic visual abilities, by examining variations in face perception using behavioural and neuroimaging measures and comparing these patterns to those of low-level vision.

A major challenge in determining how specialised processes operate within the visual system comes from the considerably different methodology used to examine low-level vision and high-level face processing. To address this issue, a novel face acuity test was designed (Chapter 2), which measured the spatial resolution of face perception across the visual field, similar to methods used to assess low-level visual resolution. Using this approach revealed that our *performance* at judging face gender varies across the visual field in a similar way to low-level visual abilities. This has implications for everyday life by indicating that our ability to accurately perceive faces varies according to where they appear – for example, we are worse at recognising faces if they appear in our upper rather than lower visual field. These findings also highlight a previously uncovered link between spatial selectivity in low- and high-level vision, suggesting that variations in spatial vision may be inherited throughout the visual system. In general, the acuity test in Chapter 2 provides a useful framework for comparing variations in low- and high-level processing, as a tool to gain further clarity about how specialisation is ingrained within our visual system.

Chapter 3 measured the neural correlates of the variations in face perception. These findings show that the shared spatial selectivity observed between low-level vision and face perception arises due to common patterns of retinotopy in early visual cortex and face-selective brain regions. While face-selective cortices have previously been considered to have relatively distinct retinotopy, these results indicate that certain retinotopic principles operate throughout the visual system, even in high-level, specialised regions.

Chapter 4 showed that the spatial properties of the stimuli used to functionally localise category-selective brain regions can impact not only their delineation, but the properties subsequently measured from them. The retinotopic properties of face-selective brain regions were found to vary according to whether localiser stimuli contained foveal/central or peripheral objects. This indicates that the accuracy of retinotopic measurements may be affected by the initial localisation methods, revealing the importance of considering how functionally defined brain regions are identified, and whether the localisation approach might bias subsequent analyses. This has particular implications for face perception research, as face-selective brain regions are typically localised using only foveal stimuli. Altogether, the findings in this thesis provide further insight into the links between low-level vision and face perception, with even high-level, specialised face processing systems inheriting the spatial selectivity of earlier vision.

Acknowledgements

First and foremost, I would like to thank my primary supervisor Dr John Greenwood. Your support and guidance has been invaluable in this journey and I am deeply grateful. Thank you for always challenging me to think in new ways, and always making time to patiently answer my questions, however big or small. Any student would be truly lucky to have you as a supervisor.

I am also grateful to my second supervisor, Dr Tessa Dekker, for the excellent advice you provided throughout my PhD. You pushed me to keep the bigger picture in mind, and your encouraging words were a great source of motivation.

I would also like to thank the other academics who have provided me with invaluable knowledge and feedback along the way, including Dr Valérie Goffaux and Dr Sam Schwarzkopf. A huge thanks also goes to Hugo, for all your help with scanning, and always making time to help me with my analyses. Your help has made my project possible and I am truly grateful. Vijay, thank you for sharing an office with me and providing essential emotional support (and sometimes, snacks).

Charlie, thank you for the unending love and encouragement (and silliness) throughout the many ups and downs of this process. I am so lucky to have had you by my side.

Vicky, thank you for the most supportive friendship. Your positive outlook on life has without a doubt made me who I am today, and helped me not give up when the going gets rough.

Alex and Billy, thank you for the general entertainment and of course, the emotional support burrata.

To my grandparents and Mum, thank you for encouraging me throughout my academic endeavours. Things have not always been easy, but I hope I can make you proud.

Lastly, thank you to Laila for always being a furry source of joy.

Contents

Declaration of Authorship	2
Abstract.....	3
Impact Statement	4
Acknowledgements.....	6
Contents.....	7
Chapter 1 General Introduction	11
1.1 Faces engage specialised processing mechanisms	12
1.1.1 Configural processing	12
1.1.1.1 Inversion effects.....	12
1.1.1.2 Composite effects	15
1.1.1.3 Part-whole effects	16
1.1.2 “Face-space” suggests another unique encoding strategy	17
1.2 Faces are processed in dedicated brain regions.....	18
1.3 Is face perception a domain specific process?.....	22
1.3.1 Face recognition can be selectively impaired	22
1.3.1.1 Impaired face recognition is linked to configural processing deficits .	24
1.3.2 A domain general account of face recognition	25
1.4 Low- and high-level visual processing.....	27
1.4.1 General properties of the visual system.....	27
1.4.2 Category versus spatial selectivity in higher-level regions.....	29
1.4.3 Sensitivity to low-level information is different for faces.....	30
1.5 Visual field anisotropies reveal insights into processing	31
1.5.1 Low-level vision varies systematically	32
1.5.2 Face perception may vary uniquely or idiosyncratically	35
1.5.3 Retinotopic sensitivity is linked to function in category-selective cortices.	39
1.6 Thesis outline	41
Chapter 2 Measuring acuity for facial gender across the visual field	44
2.1 Introduction.....	44
2.2 Experiment 1	47

2.2.1	Method	47
2.2.1.1	Participants	47
2.2.1.2	Apparatus.....	47
2.2.1.3	Stimuli	47
2.2.1.4	Procedure	48
2.2.1.5	Analyses	49
2.2.2	Results	50
2.3	Experiment 2	53
2.3.1	Method	53
2.3.1.1	Participants	53
2.3.1.2	Apparatus.....	54
2.3.1.3	Stimuli	54
2.3.1.4	Procedure	54
2.3.2	Results	55
2.4	Experiment 3	58
2.4.1	Method	58
2.4.1.1	Participants	58
2.4.1.2	Stimuli	58
2.4.1.3	Procedure	58
2.4.2	Results	59
2.5	Discussion	60
Chapter 3 Measuring the spatial properties of face-selective brain regions		65
3.1	Introduction.....	65
3.2	Method	70
3.2.1	Participants	70
3.2.2	Apparatus.....	71
3.2.3	Stimuli	71
3.2.4	Procedure.....	73
3.2.5	Localisation of face-selective ROIs	74
3.2.6	MRI data acquisition.....	75
3.2.7	MRI data preprocessing	76

3.2.8	pRF fitting.....	76
3.2.9	pRF model analysis.....	78
3.2.10	Delineation of early visual cortex	79
3.2.11	Vertex selection.....	79
3.2.12	Location analyses	80
3.2.13	Visual field coverage	80
3.2.14	Statistical analyses.....	81
3.3	Results	83
3.3.1	Behavioural results.....	83
3.3.2	Neuroimaging results	83
3.3.2.1	Eccentricity and inversion	86
3.3.2.1.1	pRF size.....	87
3.3.2.1.2	pRF number.....	89
3.3.2.1.3	Visual field coverage.....	91
3.3.2.2	Summary of eccentricity and inversion effects.....	94
3.3.2.3	Horizontal-vertical difference	94
3.3.2.3.1	pRF size.....	96
3.3.2.3.2	pRF number.....	97
3.3.2.3.3	Visual field coverage.....	100
3.3.2.4	Summary of horizontal-vertical differences.....	103
3.3.2.5	Upper-lower difference.....	104
3.3.2.5.1	pRF size.....	105
3.3.2.5.2	pRF number.....	106
3.3.2.5.3	Visual field coverage.....	109
3.3.2.6	Summary of upper-lower differences	112
3.4	Discussion	112
Chapter 4 Localising face-selective parts of the brain		121
4.1	Introduction.....	121
4.2	Method	127
4.2.1	Participants	127
4.2.2	Stimuli	128
4.2.3	‘fLoc’ functional localiser	130

4.2.4	MRI data acquisition.....	132
4.2.5	MRI data preprocessing	132
4.2.6	Region of interest (ROI) definition.....	133
4.2.7	Statistical analyses.....	135
4.3	Results	136
4.3.1	ROI size	138
4.3.2	ROI locations.....	140
4.3.3	Effects on pRF measurements.....	141
4.3.4	Face-selectivity vs. R^2	148
4.4	Discussion	149
Chapter 5 General Discussion		155
5.1	Overview of thesis findings.....	155
5.2	Face perception varies systematically across the visual field	159
5.2.1	pRF number explains visual field variations more consistently than size	160
5.2.2	Could acuity for face perception be linked to spatial integration?	162
5.2.3	The horizontal-vertical and upper-lower anisotropies differ in magnitude	164
5.2.4	Spatial selectivity follows similar patterns despite inversion	166
5.2.5	Shared patterns of retinotopy reveal links between face processing and other visual abilities	167
5.3	Featural vs. spatial selectivity within face processing systems	169
5.3.1	Inversion effects indicate configural processing across the visual field ..	169
5.3.2	Retinotopy is less clearly linked to featural selectivity.....	170
5.3.3	Can retinotopy explain what makes face processing special?.....	173
5.4	Localisation technique can affect retinotopic measurements.....	176
5.5	Conclusions.....	177
Appendices		179
Appendix A Supplementary Information for Chapter 3.....		179
Appendix B Supplementary Information for Chapter 4.....		190
References		194

Chapter 1

General Introduction

Face perception is crucial in our everyday lives – we use it to recognise our family, friends, and colleagues, as well to extract important information such as gender and emotional expressions, which help us navigate social situations. Most of us have a seemingly unlimited memory for faces and are able to recognise hundreds if not thousands of faces by a young age (Freire & Lee, 2001). Face recognition is no easy task for the visual system, however, and is instead a complex, multi-stage process, involving various static (identity, gender) and dynamic (emotion, facial speech) aspects (Bruce & Young, 1986). Faces typically share the same basic configuration (two eyes, a nose and a mouth) and differ only in their individual features, meaning that we need to make fine perceptual discriminations to recognise individuals. Yet, most of us process faces effortlessly and within a fraction of a second. How is this impressive recognition ability built into the visual system?

Presumably due to the complexity and social importance of face recognition, the human visual system appears to have developed *specialised* mechanisms for processing faces. Despite a plethora of research on face perception, there is still much debate as to the extent of this specialisation, and how exactly it operates. The question of specialisation is central to this thesis. How “special” really is face perception? And how might the visual system subserve distinct processes? Firstly, I will discuss the evidence that faces are processed uniquely within the visual system, with specific cognitive strategies and in dedicated brain regions. Then, I will delve into the neural properties of parts of the brain which are selectively involved in face perception, which suggest that processing within these regions has distinct characteristics. Finally, I will outline how these distinct neural properties may be linked to variations in face perception across the visual field, which have been found to differ considerably compared to other aspects of vision.

1.1 Faces engage specialised processing mechanisms

1.1.1 Configural processing

The visual system appears to have developed specialised processing strategies for faces. Arguably the most distinctive aspect of face recognition is that we appear to process faces holistically¹ (Tanaka & Farah, 1993), where faces are not only perceived in terms of their individual features but as a unified percept. This process involves encoding face features as well as *configural* information within the face (Maurer et al., 2002). Configural information may be divided into first-order relations, which refers to the structural configuration of the eyes, nose and mouth, and second-order relations, which refers to the spatial relations between features, such as the distance between the eyes and nose (Piepers & Robbins, 2012). This type of configural processing provides richer information about the face, and as such, faces may be perceived as more than the sum of their parts (Galton, 1879). Three main paradigms have been used to demonstrate the importance of configural face processing: *inversion*, *composite* and *part-whole* effects.

1.1.1.1 Inversion effects

Evidence for the configural processing of faces comes from the *face inversion effect*, first coined by Yin (1969). The detrimental effect of inverting a face (showing it upside down) is remarkably demonstrated using “Thatcherised” faces, where the features of a face have been inverted (Thompson, 1980). Figure 1.1 illustrates that when faces are inverted, it is difficult to tell which one has been Thatcherised. However, if the observer rotates the images so that the faces are upright it becomes immediately apparent, as its inverted features give it a grotesque appearance. Our inability to detect that features are in the wrong orientation when a face is inverted has been attributed to the disruption of configural mechanisms, where information about the spatial configuration of features within a face is lost (Bartlett & Searcy, 1993; Lewis, 2001).

Yin’s (1969) original study examined how performance on face recognition tasks may be disrupted by inversion. Participants were presented with a set of images of faces

¹ It should be noted that some researchers disagree with the term “holistic” to describe the specific nature of face processing, for example because holistic mechanisms are also present in other areas of perception, and because it may not fully capture the relational aspects of configural processing. These are indeed valid points. In this thesis, however, holistic processing may be used as an umbrella term which includes configural processing.

and other objects which are also usually seen in one orientation, such as houses. They were then presented with image pairs and asked to indicate which of the images had appeared in the original set. While participants were slower to recognise all objects when they were shown upside-down, they were *disproportionately* slower for inverted faces. This suggests that the inversion effect was mainly driven by processes specific to faces. There were inconsistencies in the stimuli used for this experiment, such as photographs used in one category and cartoon drawings in another, which could have made certain categories of object easier to recognise regardless of their orientation. However, these findings were later replicated using photographs for both face and house categories, and using different photographs during testing than the ones shown in the original set, to prevent participants performing the task by remembering the specific images themselves (Valentine & Bruce, 1986a).




Figure 1.1. The Thatcher illusion (Thompson, 1980). When faces are inverted, it is difficult to tell which face has had its features inverted. When faces are rotated so that they are upright, the manipulated face becomes immediately obvious, as its inverted features makes it appear grotesque.

Larger inversion effects for faces compared to other objects have since been demonstrated across various tasks, making it perhaps the most robust finding in the face perception literature (for review, see Rossion & Gauthier, 2002). For example, performance on a face recognition *memory* test was disproportionately impaired by inversion, compared to the recognition of upside-down dogs (Robbins & McKone, 2007).

Another study found worse performance for inverted faces when participants were required to *match* faces based on identity (Meinhardt-Injac et al., 2010). Similarly, detecting differences in face gender and emotion becomes harder when faces are upside-down, showing that inversion impairs the perception of both the changeable and unchangeable aspects of faces (Pallett & Meng, 2015; Prkachin, 2003). Face inversion effects have been consistently reported for many aspects of face perception and memory.

How exactly is inversion thought to disrupt holistic processing? As mentioned previously, researchers have proposed that when faces are upside-down, our ability to



Figure 1.2. Examples of face stimuli used by Le Grand et al. (2001), which were either manipulated **(a)** configurally, by altering the spacing between face features, or **(b)** featurally, by replacing the eyes and mouth.

process the configural information within a face is reduced, leading to poorer recognition (Bartlett & Searcy, 1993). This was investigated in more depth in a study where participants judged whether a test face was the same or different than a briefly presented target face (Le Grand et al., 2001). Performance for upright faces was found to be significantly worse when faces had been manipulated configurally, by changing the spacing between features, than when faces were altered featurally, where the eyes or mouths were replaced (Figure 1.2). When faces were inverted, however, configural changes did not disproportionately impair performance. This finding suggests that inversion impairs face recognition by disrupting our ability to encode the spatial relationships between features.

Tanaka and Sengco (1997) also manipulated the position of the eyes within faces to demonstrate a similar effect. Participants showed poorer recognition of individual face features if the features were tested in a different configuration (e.g. the eyes were further apart) to when participants had initially viewed them. Not only was the recognition of the eyes affected, but participants were worse at making judgements about other face features (noses and mouths) which were not directly manipulated, showing that the manipulation of one face feature changed the configuration of the entire face. This effect

disappeared when faces were inverted, with participants insensitive to configural changes. Similar configural and inversion effects were not found for houses, suggesting that they are specific to face perception. Altogether this suggests that although the low-level characteristics (e.g. spatial frequency; Willenbockel et al., 2010) of a face remain the same regardless of its orientation, inversion may cause our visual system to revert to a part-based strategy, where the spatial relationships between face features are not processed.

Face inversion effects have indeed led some researchers to argue that upright and inverted faces elicit qualitatively different processing strategies, where inverted faces are processed more like non-face objects (Rossion, 2008). However, configural processing has been observed for inverted faces, with horizontal displacements between face features detected better within inverted faces than vertical displacements (Goffaux & Rossion, 2007). Other findings showed that observers relied on similar, local regions within upright and inverted faces during various face discrimination tasks (Sekuler et al., 2004). This suggests that similar processing strategies may be engaged regardless of face orientation, but more efficiently for upright faces. Inversion effects could therefore reflect a quantitative rather than a qualitative change in face recognition. Either way, the fact that face inversion effects are found consistently throughout the literature suggests that faces are typically processed holistically, with inversion disrupting our ability to extract important configural information (McKone et al., 2007).

1.1.1.2 Composite effects

Other behavioural tasks have provided further evidence for holistic face processing. The *composite effect* demonstrates that recognising the individual identities of two different face halves is more difficult when the halves are aligned, compared to when they are misaligned (Figure 1.3; Young et al., 1987). This effect emerges because when the face halves are aligned, our visual system combines them to create a whole (illusory) face, impairing the




Figure 1.3. Composite faces taken from McKone et al. (2013). When two different face halves are aligned (left), it is harder to recognise either the top (Barack Obama) or bottom (Will Smith) identity, as the illusion of a new identity is created. When face halves are misaligned (right), this effect disappears.

recognition of each face half. The composite effect has been shown to disappear when faces were inverted, consistent with the idea that inversion disrupts configural processing (Le Grand et al., 2001; Maurer et al., 2002; Tanaka & Sengco, 1997). Similar composite effects have been reported for facial gender and emotion, with poorer recognition when face halves were aligned (Calder & Jansen, 2005; Calder et al., 2000; Liu et al., 2020). Like inversion effects, composite effects suggest that upright face recognition is superior due to the more efficient engagement of configural mechanisms, which are impaired by inversion.

1.1.1.3 Part-whole effects

Lastly, the *part-whole effect* also points toward configural face processing. Tanaka and Farah (1993) showed that the features of upright faces were better recognised if they were originally presented as part of a whole face as opposed to in isolation. This suggests that encoding a face feature *and* its spatial relationships to other features could benefit subsequent recognition. This effect was markedly reduced for scrambled faces and inverted faces, as well as non-face objects such as houses. In fact, isolated face parts produced better recognition scores for scrambled faces, consistent with the idea that when faces are not in their normal configuration, we switch from a holistic to a feature-based processing strategy. The whole-face advantage has been found not only for whole faces vs. isolated features but for faces in their original configuration vs. a new configuration, where the spatial relationships between features have been altered (Tanaka & Sengco, 1997). Some have used this to argue that the part-whole effect is due to an “encoding specificity” principle (Tulving & Thomson, 1973) whereby better recognition is reported for the whole face simply because that is how the part was originally encoded (Gauthier et al., 2009). However, this would predict a part-whole effect for inverted faces and houses too, which is not the case (Tanaka & Farah, 1993). These findings add to the view that our visual system does not just process the individual features of faces but integrates them holistically, a process which is largely specific to faces.

So far, this section has described different measures of configural processing (the inversion, composite and part-whole effects). These measures employ different approaches but all point towards configural processing strategies being specifically engaged for face perception (Boutet et al., 2021). What else might make the processing

of faces “special”? A multidimensional “face-space” has been put forward as another way that faces are uniquely encoded.

1.1.2 “Face-space” suggests another unique encoding strategy

Researchers have proposed the concept of a multidimensional “face-space”, where faces are encoded along multiple dimensions according to their features (Valentine, 1991; Valentine et al., 2016). These dimensions could include the shape of the face, the colour of the eyes, the distance between the eyes and the nose, age, ethnicity, and so on. Each individual identity occupies its own location within the psychological space – as no two faces are exactly the same – with similar faces represented nearer to each other. As the dimensions are assumed to follow a normal distribution, the faces closer to the centre of the space are more typical in appearance, while more distinctive faces are located further away. This could explain why distinctive faces are easier to recognise (Valentine & Bruce, 1986b, 1986c) – there are fewer faces located close by to them within face-space, so the recognition of their identity is less prone to error. Typical faces, on the other hand, may be easier to misidentify as there are more identities represented within close proximity.

Support for face-space has come from the study of adaptation effects. Adaptation effects follow the constant presentation of a certain stimulus, where perception is biased in the opposite way for a short period of time afterwards. An example of this is the tilt aftereffect, where if observers are adapted to lines oriented towards one direction, they will perceive subsequently presented lines as tilted more towards the opposite direction (Gibson & Radner, 1937; He & MacLeod, 2001). Similar effects of adaptation have been measured for faces that were generated using a morphing continuum along several feature dimensions (Leopold et al., 2001). Each face had an “anti-face”, created along the same identity trajectory but in the opposite direction. Participants were found to be better at recognising the identity of a face when they had first been adapted to its anti-face, rather than a face which did not vary along the dimensions in an opposing manner. In other words, the beneficial effects of adaptation were specific to faces that were encoded along the same-but-opposing trajectory within the multidimensional space. This is consistent with the view that individual identities are represented at different locations within a psychological space (Valentine et al., 2016).

The idea of face-space has been used to explain why inversion effects occur. Originally, it was suggested that presenting faces upside-down may lead to encoding errors, with inverted faces not represented as precisely or efficiently within face-space (Valentine, 1991). In Valentine's (1991) study, inversion effects were not as strong for distinctive faces, suggesting that the encoding errors introduced by inversion have less of an effect if there are fewer identities nearby that could interfere with recognition. Alternatively, while adaptation effects have been demonstrated for both upright and inverted faces, Leopold et al. (2001) found that the effects did not transfer across orientation. This suggests that there could be different spaces for upright and inverted faces, with inversion effects arising due to the inverted face-space being less well-developed (as we usually perceive faces in their upright orientation). Either way, these findings support the idea of faces being processed within a dedicated psychological space(s).

So far, then, we have seen that faces appear to engage specific cognitive processing mechanisms, most widely demonstrated by the inversion effect (Yin, 1969). There may also be a dedicated "face-space" within the visual system, which highlights another way that faces might be processed uniquely (Valentine et al., 2016). Next, I will discuss how specialised face perception might be subserved by distinct brain regions.

1.2 Faces are processed in dedicated brain regions

Our visual system is organised hierarchically. Information first enters through ganglion cells in the retina and travels through the lateral geniculate nucleus (LGN), before reaching the visual cortex. The visual cortex is comprised of several layers, which are tuned to different stimulus properties. Primary and secondary visual cortex (V1 and V2) encode the most basic or "low-level" elements of scenes, such as edges and orientation information (Hubel & Wiesel, 1962; Tootell et al., 1998). As information becomes more complex it is processed in successive stages of the visual system, with "high-level" processing such as object and face recognition carried out in comparatively later stages within the visual hierarchy (Riesenhuber & Poggio, 1999). Processing therefore becomes more specialised, and directed towards certain categories of object, higher in the hierarchy (Kanwisher & Dilks, 2013).

Functional magnetic resonance imaging (fMRI) has revealed a network of brain regions specialised for processing faces, which have been linked to different face processing streams (Grill-Spector et al., 2017; Haxby et al., 2000). The ventral or “core” stream consists of regions in ventral temporal cortex: the occipital face area (OFA) on the occipitotemporal gyrus, and two anatomically and functionally separate regions of the fusiform face area (FFA) referred to as pFus and mFus, on the posterior and medial fusiform gyrus, respectively (Figure 1.4). The ventral stream is thought to encode the fixed structural properties of faces, subserving tasks such as identity recognition. The “distributed” (or “extended”) face processing network consists of regions in the dorsal stream, such as the superior temporal sulcus (STS). The dorsal stream is associated with the processing of dynamically changing parts of faces, such as eye gaze and emotional expression. The distributed neural model of face perception also includes regions of the brain implicated in more general visual and cognitive abilities, such as parietal regions that are involved in spatial processing, or the amygdala which is important for processing emotion (Haxby et al., 2000). Multiple brain regions therefore work together to not only




Figure 1.4. Representation of brain regions in early visual cortex (V1-V3) and the core face processing network (OFA, pFus and mFus), displayed on the ventral surface of the brain. Together, pFus and mFus form the fusiform face area (FFA). Visual regions are organised hierarchically, from posterior to anterior locations, with regions higher in the hierarchy processing more complex information.

identify fixed aspects of faces (e.g. identity), but to extract meaning from faces – using various cues such as expression and eye gaze – which is crucial for social interaction. While face perception undoubtedly involves various brain regions and cognitive abilities, this thesis will primarily focus on the ventral or “core” face processing stream.

The FFA was the first face-selective part of the brain to be identified, with now-classic fMRI studies showing that it selectively responds to faces over other objects, measured through an increased blood oxygen dependent (BOLD) signal (Kanwisher et al., 1997; McCarthy et al., 1997). The FFA also responds more

strongly to upright compared to inverted faces, suggesting that it is most active when processing faces in their normal, upright configuration and providing a neural basis for the face inversion effect (Kanwisher et al., 1998). Some researchers have argued that inverted faces predominantly activate object-selective cortices (e.g. those that process houses), adding to the view that they are processed in a more object-like manner (Haxby et al., 1999). However, inverted faces do still activate the FFA, just to a smaller extent than upright faces, which suggests that the FFA is inherently face-selective regardless of orientation (Goffaux et al., 2016; Kanwisher et al., 1998). Interestingly, if faces were both inverted and converted to two-tone Mooney (1957) images – which make faces within the images much harder to perceive – BOLD signal in the FFA was markedly reduced compared to when faces were inverted only (Kanwisher et al., 1998). This suggests that the less “face-like” stimuli become, the less the FFA responds (Tong et al., 2000).

The FFA has been shown to be particularly sensitive to holistic face percepts, demonstrated using the composite face illusion and an adaptation repetition paradigm where participants were instructed to attend to the top face half (Schiltz et al., 2006). When face halves were aligned, the FFA was activated more strongly if the bottom face half was different during the second presentation, compared to when face halves were of the same identity in both presentations. Importantly, there were no differences in FFA activation when face halves were misaligned or the composite faces were inverted. This provides neuroimaging evidence that the FFA is involved in the holistic processing of faces.

Research has investigated this further, finding that there was stronger functional connectivity between the FFA and areas of parietal cortex involved in spatial processing when participants detected configural changes within a face, as opposed to featural changes (Zachariou et al., 2017). This suggests that the FFA is indeed involved in processing the spatial relations between face features, due to its increased communication with brain regions specialised for spatial processing. However, other studies have found a similar FFA response for configural and feature-based manipulations, suggesting that the FFA may be equally interested in face features and their configuration (Liu et al., 2010; Yovel & Kanwisher, 2004). Either way, the FFA appears to facilitate holistic processing.

On the other hand, fMRI suggests that the OFA is specifically concerned with encoding face features. While there were increased BOLD responses in the OFA when face parts appeared, the activation was similar regardless of whether parts appeared in a normal or scrambled configuration (Liu et al., 2010). Transcranial magnetic stimulation (TMS) over the right OFA has also been shown to disrupt performance on a feature-based face task, while performance on a configural task remained intact (Pitcher et al., 2007). With its preference towards face features in any configuration, the OFA may represent an early stage in the ventral face processing pathway, with the integration of face features occurring later in the hierarchy, within the FFA (Pitcher et al., 2011).

While the FFA was originally identified as one region, as mentioned previously it can also be split into two anatomically and functionally distinct regions. First, an area on the posterior fusiform gyrus referred to as pFus, and another located more anteriorly on the mid fusiform sulcus, labelled as mFus (Grill-Spector et al., 2017; Pinsk et al., 2009; Weiner & Grill-Spector, 2010). These studies have suggested that the more anterior a face-selective region is, the more selective for faces it is, with mFus showing higher selectivity for faces compared to pFus, which in turn is more selective for faces than OFA (although see Chen et al., 2022). Even more anteriorly, a face-selective region in the anterior temporal lobe has been found. This region has been linked to the final stages of face recognition and may serve as the interface between face perception and face memory (Collins & Olson, 2014). While earlier face-selective regions may focus on processing specific perceptual features, later regions may encode a more abstract face representation.

Other brain regions have been associated with the processing of more dynamic and changeable aspects of faces, which as mentioned previously, may be carried out in separate streams (Bruce & Young, 1986; Haxby et al., 2000; although see Fisher et al., 2016a). For example, research suggests that facial emotion is encoded in the superior temporal sulcus (STS). When participants were instructed to match face stimuli by emotional expression there was greater BOLD signal in the right STS, compared to when performing an identity-matching task (Narumoto et al., 2001). The role of the STS in processing facial emotion is directly evidenced using transcranial magnetic stimulation (TMS), which is a technique that uses magnetic fields to temporarily disrupt activity within certain parts of the brain. While TMS applied to a ventral part of the brain caused deficits

in recognising face identity, TMS applied to a more dorsal region reduced the ability to judge facial expression (Pitcher, 2014). The STS also shows preference for eye gaze (Engell & Haxby, 2007), demonstrating its role in various changeable aspects of faces.

Altogether these findings suggest that there are specific parts of the brain which are selectively activated for faces (Kanwisher et al., 1997). A network of brain regions is dedicated to processing the static and dynamic aspects of faces, with certain modules implicated in face-specific configural mechanisms (Grill-Spector et al., 2017). While various other brain regions may also be required in order to extract complex meaning from faces (Haxby et al., 2000), the above findings generally support the idea of domain specificity within the brain, with face perception carried out in separable processing streams to other objects (Kanwisher, 2000). Next, evidence for and against a modular theory of face processing will be examined in more detail.

1.3 Is face perception a domain specific process?

Evidence that faces are processed using unique cognitive mechanisms and within dedicated brain regions have fuelled the argument that domain specificity exists within the visual system (Fodor, 1983; Kanwisher, 2000; McKone et al., 2007). In this view, certain cognitive strategies and brain areas have developed to be specific for faces and are either innate (developed over human evolution) or develop very early on in life. We may indeed have innate mechanisms geared towards face perception, with newborn babies preferentially orienting towards faces and stimuli with face-like configurations (Johnson et al., 1991; Turati et al., 2002; Valenza et al., 1996). While this preference in newborns includes non-face objects that are top-heavy in their features (Simion et al., 2002), indicating that experience may be needed to shape functional specialisation, this suggests that our visual system is predisposed to support specialised face processing. Innate or not, is face perception special because its underlying mechanisms are domain specific, involving separable processes to other types of vision? Some of the most compelling evidence for this view comes from cases of prosopagnosia.

1.3.1 Face recognition can be selectively impaired

Prosopagnosic individuals offer further insight into the “special” nature of face processing. Prosopagnosia, also known as “face blindness”, describes a *selective* deficit

for face perception while other visual abilities remain intact (Bodamer, 1947). The condition can be developmental, where it is present from birth in the absence of any clear brain abnormality (Behrmann & Avidan, 2005). It can also be acquired due to a brain injury, usually involving right occipitotemporal lesions which often include the fusiform gyrus (Barton et al., 2002; Sergent et al., 1992). Prosopagnosia usually refers to a deficit in recognising face identity (Fisher et al., 2017), although some individuals are also impaired at perceiving gender or emotional expression (Humphreys et al., 2007; Rezlescu et al., 2012). In severe cases, prosopagnosia means that people cannot identify family members, or recognise that two identical photographs contain the same individual (Sergent & Signoret, 1992).

While individuals with prosopagnosia typically perform poorly on face recognition tasks, they often show intact *object* recognition (Behrmann & Avidan, 2005). For example, LH developed prosopagnosia following a brain injury which included bilateral occipitotemporal damage (Farah et al., 1995). Although LH could not recognise faces, he could accurately discriminate similar-looking objects of the same category, such as forks or eyeglasses. Similar findings exist for individuals with developmental prosopagnosia, who performed poorly on face memory tests but could successfully recognise non-face objects such as cars or houses (Duchaine & Nakayama, 2006). These findings show that face processing deficits can exist separately from general object agnosia and are not caused by a problem with within-category discrimination.

Individuals can also experience selective deficits in object recognition while face perception remains intact, highlighting a double dissociation between the two processes. After a brain injury, CK was no longer able to efficiently recognise non-face objects, such as a guitar or tennis racquet (Behrmann & Moscovitch, 2001; Moscovitch et al., 1997). However, his face recognition ability remained intact; he could recognise new faces as well as individuals he knew prior to his injury, and performed well on face matching tasks. CK showed larger-than-average face inversion effects, supporting the view that inverted faces engage object rather than face processing mechanisms (Haxby et al., 1999). Interestingly, when shown faces that were made up of objects, he could detect the face but could not identify the objects it was comprised of, indicating that his deficit is indeed specific to non-face objects. CK's case provides strong evidence for separable object and

face processing mechanisms, with selective damage to object processing modules leaving face perception intact.

1.3.1.1 Impaired face recognition is linked to configural processing deficits

Problems with face recognition have been associated with an inability to process configural face information. Individuals with developmental or acquired prosopagnosia often lack the expected face inversion effect of better recognition for upright compared to inverted faces, along with reduced composite and part-whole effects (Behrmann & Avidan, 2005; Busigny et al., 2010; de Gelder & Rouw, 2000; Ramon et al., 2010; Shah et al., 2015). After damage to their right occipital lobe and fusiform gyrus, patient GG could no longer recognise the identities of faces and did not display face inversion effects (Busigny et al., 2010). However, GG could accurately detect the distances between features of other objects of similar complexity. This suggests that face processing deficits in prosopagnosia may be linked to a specific inability to process faces (but not other objects) configurally.

In some cases, those with prosopagnosia can even show *opposite* face inversion effects, with better accuracy or faster reaction times for inverted compared to upright faces (Behrmann & Avidan, 2005; Farah et al., 1995). One such study examined JM, an individual with developmental prosopagnosia who showed a general overreliance on local features when processing objects, including faces (Schmalzl et al., 2009). JM was significantly worse at recognising upright compared to inverted faces, demonstrating a reversed face inversion effect. This suggests that while he could recognise inverted faces by processing individual face features, upright face recognition was impaired as his visual system attempted to engage configural processes which do not work properly. While configural processing usually enables upright faces to be perceived efficiently, these mechanisms may be damaged in prosopagnosia.

Electroencephalogram (EEG) research has suggested that prosopagnosia can also involve deficits in processing face features. The N170 is an electrophysiological marker of face encoding which can be reliably measured from occipitotemporal electrodes 140-200 ms after the presentation of a face (Eimer, 2000; Rossion et al., 2000). In individuals with developmental prosopagnosia, N170 amplitudes reduced less than in control subjects when the contrast of the eyes within face stimuli was reversed (Fisher et

al., 2016b). Prosopagnosia has therefore been linked to a reduced ability to process face features, particularly the eyes, which have been shown to be especially important for face perception (Brown & Perrett, 1993; Schyns et al., 2002; Yamaguchi et al., 2013). However, participants with prosopagnosia showed a reduced N170 attenuation compared to controls when they were instructed to fixate the mouth rather than the eyes, suggesting that they were generally less sensitive to holistic configuration (Fisher et al., 2016b). Altogether it seems that prosopagnosia can involve impairments in perceiving face features *and* integrating them holistically, with deficits in configural processing commonly found.

Altogether, these studies tell a convincing story of domain specificity within the visual system, with prosopagnosia demonstrating the selective impairment of face perception (Sergent & Signoret, 1992; Towler et al., 2017). However, others have argued that the apparently “special” nature of face recognition may instead arise from expertise. These theories oppose the idea of domain specificity within the brain, instead putting forward a domain general view of face perception.

1.3.2 A domain general account of face recognition

On the other side of the coin, researchers have proposed that the processes involved in face recognition are domain general, developing for any objects of expertise that require within-category discrimination (Gauthier, 1998; Gauthier & Tarr, 1997). Some studies have found inversion effects for objects of expertise, such as dogs in dog experts that have many years’ experience (Diamond & Carey, 1986). However, others have failed to replicate this finding in dog experts, pointing out that Diamond and Carey’s (1986) findings may have been due to previous familiarity with the specific photos of the dogs, which would have been in the upright configuration (Robbins & McKone, 2007). Instead, Robbins and McKone (2007) show that although expertise for certain objects may produce small inversion effects, they remain disproportionately large for faces.

Other evidence for a domain general view has involved training participants to become experts at recognising “Greebles”, which are novel objects considered to require similar within-category discrimination to faces (Gauthier, 1998; Gauthier & Tarr, 1997). Participants in these studies showed sensitivity to configural changes within upright but not inverted Greebles, suggesting that objects of expertise could trigger similar processing

strategies to faces. However, inversion effects were found for reaction time but not accuracy, along with variable part-whole effects, with parts sometimes being better recognised in isolation (Gauthier, 1998). Similarly, Greebles have been shown to lack a similar composite effects to faces (Robbins & McKone, 2007). Although another study did find composite effects for Greebles, they only occurred if the initial learning phase included aligned Greebles, and if the test quickly followed the learning phase (Richler 2009). In comparison, composite effects for faces were present regardless of these factors. This suggests that although configural mechanisms may be elicited by objects of expertise in some situations, they remain disproportionately involved in face processing. This is consistent with the idea of a continuum within holistic processing – while holistic mechanisms may be employed in some cases to aid object recognition, depending on various factors such as the specific features of the objects or difficulty of the task, face perception exists at the far end of the spectrum, with holistic processing almost always elicited if possible (Tanaka & Farah, 2003).

Individuals with prosopagnosia can also become experts at recognising Greebles despite being impaired on face recognition tasks (Rezlescu et al., 2014). One study examined Edward, who performed poorly on a multitude of face recognition tasks including famous face recognition and unfamiliar face memory (Duchaine et al., 2004). His performance was also similar for upright and inverted faces, showing a lack of face inversion effect. Edward was able to accurately recognise and name Greebles learnt during recent training just as well as (or in some cases, better than) control subjects. His ability to discriminate Greebles adds to the view that face processing mechanisms are largely separable from those employed for other objects, even those that require similar within-category judgements.

Perhaps more compelling evidence for a domain general account comes from neuroimaging. fMRI has revealed that car and bird experts show increased BOLD responses in the FFA in response to their objects of expertise (Gauthier et al., 2000). This increase in FFA activation was correlated with behavioural measures of expertise. However, other researchers in favour of a domain specific hypothesis have argued that FFA responses were at least twice as high for faces than for the birds or cars (Kanwisher, 2000). They also point out that the activation in Gauthier et al.'s (2000) study extended into other category-selective parts of cortex that are not involved in face perception,

suggesting that the increased BOLD responses were not specific to the FFA but may reflect more general attentional or object-selective processes. Overall there is more evidence for the FFA's specific involvement in face processing rather than general within-category discriminations (Grill-Spector et al., 2004; Kanwisher et al., 1997; Kanwisher et al., 1998).

What can we conclude from this exploration into domain specificity? The role of experience undoubtedly plays a role in face perception, and indeed may be required for domain specificity to develop – monkeys that were deprived of seeing faces from birth were shown to lack face-selective parts of the brain (Arcaro et al., 2017). Yet, these distinct brain regions develop shortly after birth given typical visual experience (Livingstone et al., 2017), with studies showing that they are more consistently involved in face processing than other objects of similar complexity or expertise (Grill-Spector et al., 2004; Kanwisher, 2000). Furthermore, the selective deficit for face perception observed in prosopagnosia indicates separable mechanisms for face and object recognition (Sergent & Signoret, 1992; Towler et al., 2017). On the whole, the evidence seems more in favour of a domain specific theory of face perception (at least within the “core” face recognition system as opposed to other brain regions that may be recruited during face perception, which are involved in more general cognitive functions; Haxby et al., 2000). How might this category specific processing be built into the brain? First, we must take a step back to consider some general properties of the visual system, and the differences between low- and high-level vision.

1.4 Low- and high-level visual processing

1.4.1 General properties of the visual system

Visual neurons essentially take “snapshots” of the world in front of us, detecting information that falls within their *receptive field*, the region of visual space that each neuron responds to. As mentioned previously, our visual system is organised hierarchically. Primary and secondary visual cortex (V1 and V2) encode the most basic or “low-level” elements of scenes, such as edges and orientation information (Hubel & Wiesel, 1962; Tootell et al., 1998). V1 neurons fire if they detect an edge within their receptive fields, at different magnitudes depending on various factors such as the length or orientation of the edge(s). As information becomes more complex it is processed in successive stages of

the visual system, with “high-level” processing such as object and face recognition carried out in comparatively later stages within the visual hierarchy (Riesenhuber & Poggio, 1999). Each stage of the visual system has its own set of response properties, shaped by the stimulus preferences of receptive fields. While receptive fields in early visual cortex fire differently depending on the orientation of simple edges, receptive fields in other parts of the brain exhibit selectivity for other aspects of visual scenes, such as colour in V4 (Schein & Desimone, 1990) or direction of motion in MT (the middle temporal visual area; Albright, 1984; Dubner & Zeki, 1971).

A major property of the visual cortex is that it is organised retinotopically, with neurons next to each other in cortex responding to locations next to each other in visual space (Wandell et al., 2007). Due to this retinotopic organisation, research has uncovered certain principles that appear to be universal within the visual system. One of these is the cortical magnification of the fovea, which refers to the increase in cortical area dedicated to processing the central visual field (Cowey & Rolls, 1974; Daniel & Whitteridge, 1961). In V1, for example, cortical magnification is highest at the fovea and decreases with eccentricity (distance from fixation; Duncan & Boynton, 2003). This likely stems from differences at the level of the retina, where the fovea is processed by a greater proportion of retinal ganglion cells than the periphery, explaining its enhanced resolution (Anstis, 1998; Rosenholtz, 2016). The oversampling of the fovea is an inherent property of the visual system that has been observed throughout the brain, in low- and high-level areas (Amano et al., 2009; Harvey & Dumoulin, 2011; Kay et al., 2015; Silson et al., 2016).

To overcome the comparatively fewer neural resources dedicated to sampling peripheral vision, receptive fields are typically larger in the periphery (Freeman & Simoncelli, 2011). In V1, receptive fields are smallest at the fovea and increase linearly in size with eccentricity (Dumoulin & Wandell, 2008). As such, there is typically an inverse relationship between receptive field size and cortical magnification factor (Harvey & Dumoulin, 2011). The linear increases in receptive field size have been specifically linked with the decreasing acuity gradient from the fovea to the periphery in V1 (Duncan & Boynton, 2003). In other words, resolution decreases as receptive fields become larger in the periphery. This is another integral property of the visual system, with similar increases in receptive field size with eccentricity observed throughout the brain (Dumoulin & Wandell, 2008; Winawer et al., 2010).

At the same time, overall receptive field size differs between low- and high-level visual cortex (Kay et al., 2015; Winawer et al., 2010). Receptive fields are smallest in lower brain regions, enabling early visual cortex to sample the visual field with high spatial resolution, and increase in size in successive levels of the hierarchy, which is thought to allow the pooling of information to support progressively more complex, higher-level processing (Freeman & Simoncelli, 2011; Riesenhuber & Poggio, 1999). Research has suggested a shift from the encoding of information with high spatial sensitivity in lower areas, supported by smaller receptive fields, to the encoding of category information with large receptive fields in higher areas (Groen et al., 2022). The next sections will discuss the distinction between spatial and category sensitivity in more detail.

1.4.2 Category versus spatial selectivity in higher-level regions

Early visual cortex encodes visual information with high spatial precision (Kamitani & Tong, 2005). On the other hand, object-selective parts of cortex have been thought to encode category representations regardless of location (Barlow, 2009; Logothetis & Sheinberg, 1996). *Object invariance* is considered an inherent property of object perception, allowing us to recognise objects despite changes in viewpoint or lighting conditions, which alters their appearance on the retina (DiCarlo & Cox, 2007). This invariance was previously thought to occur at the level of individual neurons (Logothetis & Sheinberg, 1996; Tanaka, 1996). For example, research has studied the properties of neurons in monkey inferotemporal (IT) cortex, an area which contains face-selective cells similar to face-selective brain regions in humans (Tsao et al., 2006). Single cell recordings have revealed that compared to early visual cortex, IT neurons have much larger receptive fields (Gross et al., 1972), similar to the large receptive fields found in human face-selective brain regions (Kay et al., 2015). The responses of monkey IT neurons were found to be invariant to position, firing as long as stimuli appeared within their receptive fields (Ito et al., 1995). IT neurons can be tuned to specific face viewpoints and identities, suggesting that these neurons are sensitive to the properties of faces, rather than their position within the visual field (Nam et al., 2021).

However, face-selective IT neurons have also been shown to respond more to preferred retinal locations (DiCarlo & Maunsell, 2003). In this study, rhesus monkeys were shown simple shapes either at fixation or at 1.5° in the left or right visual field. Some IT neurons were found not only to be selective for specific shapes but would only respond if

those shapes fell at a certain location. Full receptive field mapping was not performed, so it is possible that some neurons had very small receptive fields which only covered one of the three locations. However, many neurons responded to all three locations but at different strengths, demonstrating that their receptive fields covered all three locations and that they did have preferred stimulus locations. This aligns with other research showing that although many IT neurons have large receptive fields, they still have only one region of high sensitivity within their receptive fields (Op De Beeck & Vogels, 2000).

Human fMRI studies provide further evidence that object-selective brain regions can also contain position information, and that location and category can be encoded independently (Kravitz et al., 2010). In both the OFA and FFA, BOLD responses were found to be significantly higher in response to faces appearing in the lower vs. the upper visual field (Schwarzlose et al., 2008). Other research has revealed that the quadrant of the visual field that faces appeared in could be decoded from FFA activity with above chance accuracy (37%, with chance being 25%; Carlson et al., 2011). Although this is not a particularly high accuracy level, it nonetheless shows that the FFA encodes location information which can be later extracted. Interestingly, in Schwarzlose et al.'s (2008) study there was more location information in the OFA than the FFA, suggesting that some face processing areas may be more concerned with location while others are more category-focused. Consistent with this, opposing biases of spatial and category selectivity have been found in more posterior (OFA) and anterior (FFA) face-selective regions, respectively (Silson et al., 2022). Variations in position and category encoding have also been found between the ventral and dorsal face processing streams, with the dorsal stream linked to greater spatial precision in order to process the dynamically changing aspects of faces, such as expression (Freiwald et al., 2016; Pitcher, 2014). Altogether, these studies suggest that like early visual cortex, face processing brain regions respond differently depending on location.

1.4.3 Sensitivity to low-level information is different for faces

Besides location, does the low-level featural² content of faces affect their perception? Research suggests that it can – with contrast levels affecting face recognition,

² While the term “featural” is often used to refer to face features (e.g. the eyes, nose, mouth) in the face perception literature, it is also used in a more general sense to describe the featural content of faces, for example in terms of low-level image characteristics such as spatial frequency, orientation, and contrast.

for example (Fisher et al., 2016b; Schyns et al., 2002; Yue et al., 2011) – but in unique ways. In earlier brain regions, selectivity for low-level features tends to occur due to retinotopic organisation. For example, receptive fields in early visual cortex respond preferentially to radial as opposed to tangential orientations, known as the radial bias (Rovamo et al., 1982; Sasaki et al., 2006). As mentioned previously, however, horizontal orientations are important for various aspects of face perception, including identity recognition and holistic processing (Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux & Greenwood, 2016). Horizontal spatial frequencies have also been found to elicit the strongest BOLD responses in the FFA (Goffaux et al., 2016). This suggests that selectivity for featural content also follows category as opposed to spatially based principles within face perception.

As with orientation, spatial frequency preference in early visual cortex ties in with retinotopic organisation. V1 receptive fields are tuned towards higher spatial frequencies near the fovea and progressively lower ones in the periphery (Henriksson et al., 2008). In contrast, the holistic processing of faces has been shown to rely on lower spatial frequencies (Awasthi et al., 2011; Collishaw & Hole, 2000; Goffaux et al., 2005). Inversion effects were equally strong when faces were filtered to only contain lower frequency information, but reduced (indicating decreased configural processing) when only higher spatial frequencies were left intact (Goffaux & Rossion, 2006). This suggests an overall benefit of low spatial frequencies for face perception – the holistic aspect, at least – which is not tied to retinotopic location. These variations in sensitivity to low-level information – which are differentially associated with retinotopy and category – highlight further discrepancies between low- and high-level vision.

1.5 Visual field anisotropies reveal insights into processing

So far, we have seen that face perception involves specialised cognitive processes that are subserved by distinct brain regions. The retinotopy of these face-selective regions appears to differ considerably from earlier brain areas, highlighting differences between low- and high-level visual processing. Further evidence for this distinction comes from research investigating how visual abilities vary according to location. As we will see,

In this thesis, it will mainly be used in the more general sense to refer to the featural *content* of faces (which of course encompasses the face features themselves, but does not necessarily refer to them specifically).

measuring how perceptual abilities differ across the visual field in this way can provide insights about the underlying functionality of the visual system. The rest of this chapter will focus on the visual field variations that have been found for low-level vision and face recognition, and how they may be linked to the way that neurons sample the visual field.

1.5.1 Low-level vision varies systematically

Our ability to perceive simple, low-level stimuli – such as lines or letters – varies across the visual field. These abilities are often assessed by measuring the smallest size or the minimum spatial resolution required to make judgements about simple stimuli. The most obvious example of how vision varies across the visual field is that acuity – our ability to resolve fine detail – is highest in the fovea (1-2° of visual angle at the centre of vision) and declines with eccentricity (distance from fixation; Rosenholtz, 2016).




Figure 1.5. Horizontal (pink) and vertical (purple) meridians within the visual field. The size of the letters and contrast of the gratings represent visual field anisotropies in low-level vision, with smaller letters and lower contrast indicating better acuity.

The perception of low-level stimuli also varies according to *angular location* around fixation, with eccentricity kept constant. Acuity is typically better along the horizontal meridian compared to the vertical, reflecting a *horizontal-vertical* anisotropy, and in the lower compared to the upper visual field, highlighting an *upper-lower* difference (Figure 1.5; Carrasco et al., 2001; Greenwood et al., 2017; Westheimer, 2003). For example, at 10° in the periphery participants could more accurately judge the orientation of gratings that appeared along the horizontal than the vertical meridian, and in the lower

versus upper field (Barbot et al., 2021; Carrasco et al., 2023). These horizontal-vertical and upper-lower anisotropies consistently emerge for many elements of vision, including spatial frequency, contrast sensitivity and crowding (Abrams et al., 2012; Barbot et al., 2020; Benson et al., 2002; Himmelberg et al., 2020; Rubin et al., 1996), indicating that low-level visual perception is systematically influenced by location.

Variations in low-level vision may be explained by how neurons in early visual cortex sample the visual field. *Retinotopic maps* exist within each region of early visual cortex – adjacent neurons encode adjacent regions of the visual field, resulting in a clear topological representation of visual space (Wandell et al., 2007). In other words, there are clearly observable gradients of polar angle and eccentricity preferences within V1-V3 (Figure 1.6). This well-defined retinotopic organisation has allowed researchers to identify regions of the visual field that are processed by a greater proportion of visual cortex. As covered earlier, cortical magnification factor in V1 is highest at the fovea and decreases with eccentricity (Duncan & Boynton, 2003). This likely stems from differences at the level of the retina, where the fovea is processed by a greater proportion of neurons than the periphery, explaining its enhanced resolution (Anstis, 1998; Rosenholtz, 2016). V1 also has smaller receptive fields at the fovea, which increase in size with eccentricity (Dumoulin & Wandell, 2008). As such, an inverse relationship between receptive field size and cortical magnification factor is considered an inherent property of the visual system (Harvey & Dumoulin, 2011). Smaller receptive fields at the fovea have been linked with increased acuity in V1 (Duncan & Boynton, 2003). Therefore, differences in acuity between the fovea and periphery can be explained by neuronal *density* and *receptive field size* in early visual cortex.




Figure 1.6. Polar angle and eccentricity maps showing retinotopic organisation within early visual cortex (V1-V3). UVF = upper visual field, LVF = lower visual field, d = dorsal, v = ventral. Figure reproduced from Groen et al. (2022).

These properties can also explain variations in low-level vision according to polar angle (Abrams et al., 2012; Barbot et al., 2021; Carrasco et al., 2023; Carrasco et al., 2001; Himmelberg et al., 2020). Anatomical studies report higher densities of retinal ganglion cells along the horizontal compared to the vertical meridian (Curcio & Allen, 1990; Perry & Cowey, 1985). Consistent with this, research has found a greater cortical

magnification factor along the horizontal vs. the vertical meridian and in the lower vs. upper field in V1-V3 (Silva et al., 2018), which correlates with better acuity and orientation discrimination (Duncan & Boynton, 2003). Adults also have a greater surface area in V1 along the horizontal vs. vertical meridian and in the lower vs. upper field, although children have a similar representation of the upper and lower locations, which could explain why the upper-lower difference does not emerge until adolescence (Himmelberg et al., 2023).

Population receptive field (pRF) mapping is a technique that takes advantage of the retinotopic organisation of the visual system, using fMRI to identify populations of neurons which respond to specific parts of the visual field, as well as properties such as pRF location (the centre of the pRF itself) and size (Dumoulin & Wandell, 2008). Consistent with variations in cortical magnification and surface area, research has shown that there are a greater number of pRFs located along the horizontal vs. the vertical meridian in V1 (Amano et al., 2009). Visual field coverage represents how well parts of the visual field are sampled, based on the number and size of pRFs in each region. Due to the increased neuronal density along the horizontal meridian, Amano et al. (2009) also report better visual coverage along the horizontal meridian. These findings suggest that variations in low-level vision could be linked to the *number* of neurons in early visual cortex which sample each part of the visual field.

Smaller pRFs have also been identified along the horizontal vs. vertical meridian in V1-V3 (Silva et al., 2018), which again correlates with acuity (Duncan & Boynton, 2003). In V1, this horizontal-vertical difference in pRF size increased with eccentricity, meaning that retinotopic properties could explain why behavioural anisotropies become more pronounced further into the periphery (Greenwood et al., 2017). Similarly, pRFs in early visual cortex that sample the lower field have been found to be smaller than those sampling the upper field, which could explain the lower field advantage for low-level stimuli (Silson et al., 2018; Silva et al., 2018). Interestingly, in Silva et al.'s (2018) study the upper-lower difference was significant in V1 but not in V2 and V3, despite a trend in the same direction. This is consistent with behavioural reports that the upper-lower difference is smaller in magnitude and harder to measure than the horizontal-vertical difference (Barbot et al., 2021; Kurzwski et al., 2021). These findings also link variations in low-level vision to differential sampling of the visual field, with *smaller receptive fields* beneficial for perception.

1.5.2 Face perception may vary uniquely or idiosyncratically

Like low-level vision, face recognition declines in peripheral vision compared to the fovea (McKone, 2004). However, we can still accurately recognise faces that appear in the periphery. McKone et al. (2004) demonstrated that facial identity could be accurately recognised at various locations in peripheral vision, with above chance performance even for faces that appeared 21° from fixation. Importantly, inversion effects were found at all eccentricities, indicating that configural face processing mechanisms (Maurer et al., 2002) were engaged across the visual field. Other studies also report the accurate recognition of faces and face inversion effects in peripheral vision (Kalpadakis-Smith et al., 2018; Roux-Sibilon et al., 2023).

Unlike low-level vision, however, face perception appears to *differ* across the visual field in distinct ways. Studies have suggested that faces are better perceived in the left visual field (Ellis & Shepherd, 1975; Harrison & Strother, 2021; McKone, 2004). At 10° in the periphery, discriminating the identity of synthetic faces was found to be better in the left visual field compared to the right, upper, and lower field (Schmidtman et al., 2015). This is consistent with research suggesting that the right hemisphere is more specialised for face processing (Grill-Spector et al., 2017; Kanwisher et al., 1997; Rangarajan et al., 2014), and with lesions in the right occipitotemporal cortex being sufficient to cause prosopagnosia (Barton et al., 2002; Sergent et al., 1992). However, it highlights a departure from low-level vision, with no horizontal-vertical or upper-lower differences found for face perception.

There have also been suggestions of an upper field advantage for face perception, with the gender of faces found to be recognised faster in the upper vs. lower visual field (Quek & Finkbeiner, 2014, 2016). This could be linked to findings that face features are better discriminated when they appear at their typical visual field locations (assuming central fixation), with the eyes – which are particularly important for gender recognition (Brown & Perrett, 1993; Schyns et al., 2002; Yamaguchi et al., 2013) – better recognised in the upper field (de Haas et al., 2016). While there is evidence of a general upper field bias in temporal processing (Honda & Findlay, 1992), if faces are resolved more accurately in the upper visual field, this would indicate an opposite upper-lower difference to low-level vision.

Other studies report that face perception varies across location with no systematic pattern at all. For example, biases in the apparent gender and age of morphed faces were found to vary idiosyncratically across the visual field; where one participant could perceive faces as more female in the left field and more male in the right, another participant could show the opposite pattern (Afraz et al., 2010). As faces appeared only 3° from fixation, it is possible that they were not shown far enough into the periphery to observe systematic anisotropies, which typically become more pronounced with eccentricity (Greenwood et al., 2017). However, similar idiosyncratic variation was found for perceiving the identity of morphed faces presented at 7° eccentricity (Visconti di Oleggio Castello et al., 2018). In both these studies, the perceptual biases were stable across time within individuals. This highlights a potentially dramatic departure from low-level vision, with no systematic variation for face perception.

What could be the neural underpinnings of these unique variations? While face-selective parts of the brain may not have retinotopic maps that are as fully formed as those in early visual cortex, they show characteristics of retinotopic organisation, with biases towards certain regions of the visual field (Groen et al., 2022; Silson et al., 2016; Silson et al., 2022)³. While there is cortical magnification of the fovea throughout the visual cortex (Dekker et al., 2019), pRFs in V1 are still distributed relatively evenly across the visual field (Amano et al., 2009; Arcaro et al., 2009; Wandell & Winawer, 2015). In contrast, pRFs have been found to cluster near the fovea across multiple face-selective brain regions (Gomez et al., 2018; Kay et al., 2015; Poltoratski et al., 2021), with around 80% of pRFs located less than 5° from fixation (Finzi et al., 2021). This *exaggerated* cortical magnification of the fovea increases even further as face-selective areas become more anterior in the ventral face processing stream, from OFA, to pFus, to mFus (Kay et al., 2015). The foveal bias also appears to increase during development, with pRFs in pFus located more centrally in adults compared to children (Gomez et al., 2018). As face recognition abilities increase over childhood and become adult-like by adolescence (Bruce et al., 2000; Carey et al., 1980), improvements in face perception – which are typically measured in the fovea – could be linked to a stronger representation of the central visual

³ While the term “retinotopy” is used in a general sense throughout this thesis, a distinction should be highlighted between the strict retinotopic maps which are observable in V1-V3, and visual field biases within higher-level category-selective cortex, which may occur in the absence of a fully formed retinotopic map (Groen et al., 2022).

field in parts of the brain that process faces. pRFs in the FFA have also been shown to be located more foveally in the right than the left hemisphere, consistent with research suggesting that faces are perceived better in the left visual field (Butler & Harvey, 2005; Ellis & Shepherd, 1975; Harrison & Strother, 2021; McKone, 2004). Altogether these findings suggest that face-selective parts of the brain predominantly represent the fovea at the expense of the periphery, highlighting a dissociation from early visual cortex.

Recently, however, pRFs positioned more peripherally have been reported in some face-selective areas. Like Kay et al. (2015), Finzi et al. (2021) found the same pattern of increasing cortical magnification at the fovea in the ventral face processing stream; the foveal bias increased from OFA, to pFus, to mFus, with the majority of pRFs located in the central 5° of the visual field. However, pRFs in two lateral face-selective areas, the posterior (pSTS) and mid superior temporal sulcus (mSTS) were more spread out across the visual field; 60% of pRFs in pSTS and 80% in mSTS covered the visual field at 20° eccentricity, compared to only around 30% in mFus (Finzi et al., 2021). In this study the researchers also carried out diffusion-weighted MRI, which reveals white matter tracts that connect visual areas with each other. While the majority of white matter tracts in the ventral face-selective areas came from the central 10° of early visual cortex, tracts in lateral face processing regions originated from more peripheral eccentricities, suggesting that the position of pRFs in face-selective areas may depend on the input that they receive from earlier visual regions. Even so, this still suggests that representations of the visual field within the ventral or “core” face processing network (Grill-Spector et al., 2017) are heavily biased towards the fovea.

As discussed earlier, receptive fields increase in size higher in the visual hierarchy (Dumoulin & Wandell, 2008; Freeman & Simoncelli, 2011). Consequently, another major way that face-selective brain regions differ from early visual cortex is their much larger receptive fields. While pRFs in ventral face-selective areas (OFA, pFus and mFus) are also smallest at the fovea, they are generally much larger than in V1-V3, and increase in size more rapidly with eccentricity (Finzi et al., 2021; Gomez et al., 2018; Kay et al., 2015; Poltoratski et al., 2021). For example, at 3° eccentricity, the median size of pRFs in V1 was just below 0.5°, but approximately 5° in mFus (Kay et al., 2015). In this study, pRF size also increased from posterior to anterior regions in ventral temporal cortex, with smaller pRFs in OFA (posterior) and the largest in mFus (anterior). pRFs in face-selective

regions – but not in V1 – also became larger when attention was directed towards faces, suggesting that larger receptive fields are beneficial for face recognition. However, as pRFs increased in size even for a task which simply required participants to detect when a dot appeared on the faces, instead of being inherently face-based, this finding could be a general property of object-selective cortices instead of being specific to face recognition.

Other studies have also linked larger receptive fields to better face recognition. One in particular found smaller pRFs in ventral face-selective regions – but not in V1 – for *inverted* compared to upright faces (Poltoratski et al., 2021). In response to inverted faces, pRFs not only decreased in size but also shifted downwards, away from the fovea. Overall visual field coverage in the face-selective areas was therefore reduced for inverted faces. Interestingly, this shift in coverage was correlated with an increased behavioural face inversion effect in the lower left visual field. In this paper, the authors propose that larger receptive fields aid the configural processing of faces by enabling the spatial integration of face features over a larger area. This could explain why the recognition of inverted faces is worse, as smaller receptive fields are less able to integrate information. It is not fully clear why a larger inversion effect was found in the bottom left hemifield, though, as this specific anisotropy has not been consistently reported for inversion effects.

Smaller pRFs have also been found within the face-selective brain regions – but not V1-V3 – of individuals with developmental prosopagnosia (Witthoft et al., 2016). As these individuals also showed a strong foveal bias, the small size of pRFs within their face-selective areas meant that visual field coverage was even more restricted to the fovea than it was in controls, suggesting that large receptive fields may be needed to provide enough coverage for face recognition. Overall, research highlights a puzzling dissociation between low- and high-level vision, where smaller receptive fields are associated with better acuity in early visual cortex (Duncan & Boynton, 2003; Silson et al., 2018; Silva et al., 2018; Silva et al., 2021) yet *larger receptive fields* appear to be beneficial for face processing (Kay et al., 2015; Poltoratski et al., 2021; Witthoft et al., 2016).

Overall, there appear to be substantial differences in how early visual cortex and face processing brain areas sample the visual field (Figure 1.7). Both show cortical magnification of the fovea and receptive fields that increase in size with eccentricity (Dumoulin & Wandell, 2008; Kay et al., 2015). Early visual cortex samples the visual field

relatively uniformly (Amano et al., 2009; Arcaro et al., 2009), with smaller receptive fields linked to better acuity (Duncan & Boynton, 2003; Silva et al., 2021). On the other hand, face-selective regions show an exaggerated foveal bias, with studies suggesting that larger receptive fields are beneficial for face perception and that receptive field properties are tied to (Finzi et al., 2021; Gomez et al., 2018; Poltoratski et al., 2021; Witthoft et al., 2016). Could this dissociation in receptive field properties stem from spatial versus category selectivity within the brain?




Figure 1.7. Differences in receptive field size and position between primary visual cortex (V1) and the fusiform face area (FFA). Receptive fields in V1 tile the visual field, with small receptive fields linked to high acuity at the fovea. In contrast, research suggests that receptive fields in the FFA are large and cluster at the fovea, with poor coverage of the periphery. Note that this does not illustrate polar angle anisotropies, nor receptive field number. Size is represented in an abstract sense and is not meant to convey exact physiological estimates.

1.5.3 Retinotopic sensitivity is linked to function in category-selective cortices

Like earlier visual regions, category-selective brain areas show characteristics of retinotopic organisation, such as a contralateral visual field bias in object, scene, and face-selective cortices (Silson et al., 2016; Silson et al., 2022). The retinotopic biases of category-selective brain regions appear to vary according to their functional objectives. For example, scene or place perception requires us to process information across the visual field. Consistent with this, studies have shown that peripheral vision plays an important role in scene analysis (Wang & Cottrell, 2017), to the extent that patients with

central visual field loss can successfully categorise scenes (Thibaut et al., 2014). fMRI has been used to identify the parahippocampal place area (PPA), an area which responds selectively to scenes (Epstein & Kanwisher, 1998). As the PPA responds similarly to rooms with furniture and empty rooms, this suggests that it encodes the general spatial layout of the environment as opposed to the objects within it. In line with this, research indicates that pRFs in the PPA are biased towards peripheral eccentricities (Silson et al., 2015), with two areas in the parahippocampal cortex found to have pRFs which were largely restricted to the fovea and eccentricities ranging from 7.5-15° (Arcaro et al., 2009). These findings suggest that receptive field properties in place-selective brain areas allow the peripheral visual field to be efficiently sampled, which supports scene processing.

On the other hand, when faces appear in our visual field, we usually fixate them directly to recognise them (de Haas et al., 2019). As such, we are worse at recognising faces in peripheral compared to central vision (McKone, 2004). In particular, fixating the eyes – or the upper part of the face just below the eyes – has been found to be diagnostic of identity (Orban de Xivry et al., 2008) and gender (de Lissa et al., 2014). Face features themselves are more accurately discriminated when they appear in their expected locations within the visual field, assuming central fixation; eyes in the upper field, mouths in the lower (de Haas et al., 2016). Adults also tend to fixate the central part of faces more than children do, which coincides with adults' generally better face recognition abilities (Gomez et al., 2018). Individuals with central visual field loss often perform poorly at face recognition, as they are unable to look directly at faces (although they experience difficulties with other visual tasks too; Glen et al., 2012). As fixating faces is clearly beneficial for their recognition, it makes sense for face processing neurons to preferentially sample the fovea.

Direct fixation is not a behaviour specific to faces, however. Most visual abilities are performed much better if we look directly at the stimuli involved (de Haas et al., 2019; Rayner, 1977). Yet, cortical magnification at the fovea seems *disproportionately* enlarged for face recognition, to the extent that face-selective areas appear to have limited coverage of the periphery. How, then, can we still accurately recognise faces that appear in peripheral vision (Kalpadakis-Smith et al., 2018; McKone, 2004; Roux-Sibilon et al., 2023)? A puzzling distinction arises; while retinotopic mapping suggests that there is

limited coverage of the periphery in regions of the brain which process faces, behavioural research shows that face recognition can be performed well in peripheral vision.

To summarise these sections, evidence suggests that low-level vision varies predictably across the visual field, with *horizontal-vertical* and *upper-lower* anisotropies (Abrams et al., 2012; Barbot et al., 2021; Benson et al., 2021; Carrasco et al., 2023; Carrasco et al., 2001; Greenwood et al., 2017; Himmelberg et al., 2020). On the other hand, face recognition does not appear to vary in the same way. While some findings suggest that face perception is better in the left (Schmidtman et al., 2015) or upper visual field (Quek & Finkbeiner, 2016), others report entirely idiosyncratic variations (Afraz et al., 2010; Visconti di Oleggio Castello et al., 2018). This disconnect between low- and high-level anisotropies could stem from variations in how early visual cortex and face-selective brain regions sample the visual field (Kay et al., 2015), and a difference between spatial and category based processing (Groen et al., 2022). Although these distinct behavioural and neuronal variations are consistent with the “special” nature of face perception, it is unclear how they emerge within the visual system.

1.6 Thesis outline

The general aim of the thesis was to investigate the distinct nature of face perception, and how this fits into the wider visual system. Converging evidence indeed suggests that faces are processed in a specialised way. Faces engage distinct cognitive mechanisms (McKone et al., 2007) and activate specific regions of the brain (Grill-Spector et al., 2017; Kanwisher et al., 1997; Kanwisher et al., 1998), with cases of prosopagnosia demonstrating that face recognition can be selectively impaired while the recognition of other objects remains intact (Behrmann & Avidan, 2005; Busigny et al., 2010; Sergent & Signoret, 1992). Face perception also seems to vary across the visual field in distinct ways compared to low-level vision (Afraz et al., 2010; Quek & Finkbeiner, 2016; Schmidtman et al., 2015; Visconti di Oleggio Castello et al., 2018), with face-selective parts of the brain sampling the visual field differently to earlier visual regions (Figure 1.6; Kay et al., 2015; Poltoratski et al., 2021). These sampling differences may be linked to spatial versus category-based encoding within the brain (Groen et al., 2022).

How can we accurately assess how “special” face processing really is? A major challenge in doing so revolves around the different methodology used to investigate low-

level vision and higher-level face perception. While low-level visual abilities are often examined using spatial properties such as acuity (e.g. Barbot et al., 2021; Carrasco et al., 2001), face perception studies typically involve subtle judgements of appearance, using faces of the same size (e.g. Afraz et al., 2010). This makes it difficult to compare face recognition abilities to other types of vision. It was therefore essential to align the methodology used to measure variations in low- and high-level visual processing. By examining how face perception varies across the visual field when it has been measured in a similar way to low-level vision, this would help us better understand the links between low- and high-level vision.

The first specific aim of the thesis was therefore to assess the apparent uniqueness of face recognition by measuring the spatial resolution of face perception across the visual field. Chapter 2 describes a novel face acuity test, which measured acuity for judging face gender (upright and inverted) at various locations in peripheral vision. Because this is a similar approach to acuity measures used for low-level vision, this allows low- and high-level anisotropies to be more easily compared. If there is systematic variation to be found for face perception, there should be similar anisotropies to low-level vision, such as a horizontal-vertical and upper-lower difference (Carrasco et al., 2001). There could also be systematic yet unique variation for faces, such as better acuity in the left (Schmidtman et al., 2015) or upper visual field (Quek & Finkbeiner, 2016). Alternatively, acuity variations for faces could vary in an idiosyncratic manner, with no common variation across individuals (Afraz et al., 2010; Visconti di Oleggio Castello et al., 2018). Investigating the spatial resolution of face perception should not only provide further insight into the distinct nature of face processing, but also how these specialised mechanisms relate to other facets of vision, as a further tool to examine the “specialness”.

The second specific aim was two-fold: first, to determine whether the spatial properties of face-selective brain regions could explain the variations in face perception measured in Chapter 2, and second, whether these spatial properties align with – or diverge from – those in early visual cortex. In Chapter 3, the retinotopic properties of face-selective brain areas (OFA, pFus and mFus) and of early visual cortex (V1-V3) were investigated using population receptive field (pRF) mapping (Dumoulin & Wandell, 2008) within central and peripheral vision. Here, I examine whether the retinotopic properties of face-selective areas differ according to visual field location, to determine whether the

behavioural anisotropies found for faces (Chapter 2) could arise from variations in visual field sampling. I also analyse whether these properties vary according to face inversion more generally, to assess the links between featural selectivity and sampling within face-selective regions. Importantly, patterns of retinotopy are compared between early visual cortex and face-selective areas, to determine whether there is shared spatial selectivity between parts of the brain involved in low-level vision and face perception. These direct comparisons of sampling variations provide further insight into the links between low- and high-level vision.

The third specific aim was to determine whether the retinotopy measured within face-selective brain regions would be influenced by the methods used to originally identify the regions. Face-selective areas are typically identified using functional localisers that present face stimuli foveally, which could result in inaccurate measurements during subsequent analyses of these regions, such as an overestimated foveal bias. To investigate this, Chapter 4 describes a novel method of functionally localising face-selective parts of the brain, which involves foveal *and peripheral* face stimuli shown using a large field of view. First, I explore whether the cortical location and size of the face-selective areas differs according to whether they were delineated using foveal or peripheral stimuli. Then, retinotopic analyses were performed using the different delineations, which revealed that certain pRF measures varied according to the specific stimuli used. This indicates that the spatial properties of the stimuli used during localisation can affect subsequent analyses within functionally defined brain regions, and reveals further insight into whether the retinotopy in face-selective areas really is distinct. Altogether these experiments investigate the spatial properties of face perception, and how these spatial properties relate to those of low-level vision, to address the overarching question of whether faces really are “special”.

Chapter 2

Measuring acuity for facial gender across the visual field

2.1 Introduction

As discussed in the previous chapter, vision varies across the visual field. For the recognition of simple low-level stimuli ranging from lines to letters, these variations are often assessed using measures of spatial resolution or acuity (e.g. the smallest size needed for accurate recognition). Our perception of simple stimuli is systematically influenced by both eccentricity (becoming worse as distance from the fovea increases; Rosenholtz, 2016) and angular location around fixation (e.g. worse performance in the upper vs. lower visual field; Abrams et al., 2012). In contrast, the perception of high-level stimuli such as faces has been found to vary across the visual field in a unique or even entirely idiosyncratic fashion (Afraz et al., 2010; Quek & Finkbeiner, 2016; Schmidtmann et al., 2015). These distinct variations are consistent with the view that faces are “special” in the visual system, in that they are processed by unique mechanisms which are disproportionately disrupted by inversion (Robbins & McKone, 2007; Rossion, 2008). However, this dissociation between low- and high-level perceptual variations means that the extent to which specialised face processing systems rely on earlier visual processing is unclear – variations in face recognition may arise solely in high-level brain areas, independently of the variations in low-level areas. This possibility is difficult to examine because so far, the methodology used to measure variations in low-level vision and face perception has been considerably different. Consequently, I sought to measure the spatial resolution of face perception and its variation around the visual field.

As above, low-level properties such as visual acuity not only decline with eccentricity (Rosenholtz, 2016) but also vary by location, even with eccentricity held constant. Acuity is typically better along the horizontal meridian compared to the vertical (Figure 2.1A; Barbot et al., 2021; Carrasco et al., 2001; Greenwood et al., 2017; Westheimer, 2005). Along the vertical meridian, acuity is better in the lower compared to the upper visual field (Carrasco et al., 2001; Greenwood et al., 2017). These two

anisotropies – horizontal-vertical and upper-lower – consistently emerge for many elements of vision, including bisection acuity, contrast sensitivity and crowding (Abrams et al., 2012; Barbot et al., 2021; Benson et al., 2021; Himmelberg et al., 2020; Rubin et al., 1996), demonstrating that low-level visual perception is fundamentally influenced by location.

These anisotropies have been linked with the retinotopic organisation of the visual system. Multiple retinotopic maps exist throughout the visual hierarchy, with adjacent locations in the visual field encoded by anatomically adjacent neurons (Arcaro et al., 2009; Wandell et al., 2007; Wandell & Winawer, 2011). Anatomical studies report higher densities of retinal ganglion cells along the horizontal vs. the vertical meridian (Curcio & Allen, 1990; Perry & Cowey, 1985). Similarly, in early visual cortex (V1-V3), smaller population receptive field (pRF) sizes have been found along the horizontal vs. the vertical meridian and the lower vs. the upper field, highlighting variations in sampling across the visual field (Silson et al., 2018; Silva et al., 2018). Higher cell densities and smaller pRF sizes have been linked with better acuity (Duncan & Boynton, 2003), which could explain variations in low-level information processed by early visual cortex.

Like low-level vision, face recognition declines in peripheral vision compared to the fovea (McKone, 2004). However, variations in face recognition according to angular location may differ substantially. For example, the identity of synthetic-contour faces was more accurately discriminated in the left visual field, with no significant horizontal-vertical or upper-lower differences (Schmidtman et al., 2015). This is consistent with a left hemifield bias within face perception (Ellis & Shepherd, 1975; Harrison & Strother, 2021; McKone, 2004). Other studies have found that the gender of faces was recognised more quickly in the upper vs. lower field (Quek & Finkbeiner, 2014, 2016), suggesting an opposite upper-lower difference to low-level vision. Biases in the apparent gender and age of morphed faces have also been found to vary idiosyncratically across the visual field; where one participant perceived morphed faces as more female in the left field and more male in the right, another participant showed the opposite pattern (Afraz et al., 2010). Similar idiosyncrasies were found for biases in the perceived identity of morphed faces (Visconti di Oleggio Castello et al., 2018). These distinct patterns suggest a dissociation in the mechanisms driving the visual field variations in low-level vision and face perception.

This dissociation may not be surprising given evidence that faces undergo distinct forms of cognitive processing (Robbins & McKone, 2007). For instance, relative to other objects, face recognition is disproportionately impaired for upside-down vs upright faces (Yin, 1969). This *inversion effect* is driven by increased sensitivity to the spatial relationships between features (configural processing) within upright faces (Le Grand et al., 2001; Maurer et al., 2002; Piepers & Robbins, 2012; Rossion, 2008). Functional magnetic resonance imaging (fMRI) has identified a network of ventral occipitotemporal brain regions dedicated to face processing, such as the fusiform face area (FFA) which shows greater activation for upright vs inverted faces (Yovel & Kanwisher, 2004, 2005). Like early visual cortex, higher-level face-selective regions also show retinotopy, with smaller receptive field sizes in the fovea vs. the periphery (Kay et al., 2015; Poltoratski et al., 2021). It is unclear how the retinotopy in face-selective areas is linked to earlier regions, however. In low-level vision, neural selectivity derives from earlier stages in the visual hierarchy, with orientation selectivity in V1 simple cells arising from specific combinations of circular-symmetric LGN cells, for example (Hubel & Wiesel, 1962). The dissociations between low-level vision and face perception suggest that neuronal selectivity in face-selective regions may not have the same dependence on inputs from earlier visual cortex.

A major challenge in comparing the above variations in low- and high-level vision derives from differences in methodological approach. While measurements of low-level vision tend to focus on spatial properties such as acuity or grating resolution, face recognition experiments often measure the percent-correct recognition of facial characteristics such as gender at a fixed face size (Afraz et al., 2010; Quek & Finkbeiner, 2016; Schmidtmann et al., 2015). To align the methodology used to assess variations in low-level vision and face perception, I developed an acuity test for faces, which measures the smallest size necessary to judge gender at a given visual field location. Upright and inverted faces were included to determine whether the paradigm engaged configural mechanisms, and to assess whether visual field variations would differ according to inversion. If face processing systems inherit the spatial properties of early visual cortex, anisotropies similar to those found for low-level vision should emerge for faces. Alternatively, there could be idiosyncratic variation (Afraz et al., 2010; Visconti di Oleggio Castello et al., 2018) or a systematic but unique pattern of anisotropy (Quek & Finkbeiner,

2016; Schmidtman et al., 2015). The latter outcomes would suggest that face recognition involves distinct mechanisms that do not inherit the spatial properties of earlier brain regions, with spatial selectivity arising within face-selective cortices themselves.

2.2 Experiment 1

2.2.1 Method

2.2.1.1 Participants

14 participants (13 female, one male, $M_{age} = 24.9$ years) took part, including myself and one supervisor; the rest were naïve. All had normal or corrected-to-normal vision of at least 20/20, assessed using a Snellen chart viewed at central fixation. Nine were right-eye dominant, determined using the Crider ring test (Crider, 1944). This sample size was derived from previous studies with similar designs (e.g. Abrams et al., 2012). All experiments were approved by the Research Ethics Committee for Experimental Psychology at University College London and all participants gave written informed consent before testing began.

2.2.1.2 Apparatus

The experiment was programmed in MATLAB (MathWorks, Inc) and conducted on an Apple iMac running PsychToolbox (Brainard, 1997; Kleiner et al., 2007; Pelli & Vision, 1997). Stimuli were viewed binocularly on a Cambridge Research Systems Display++ monitor with 2560 x 1440 resolution and 120 Hz refresh rate. The monitor was gamma corrected and linearised through software to have a minimum luminance of 0.16 cd/m² and a maximum of 143 cd/m². Participants were seated at a 50cm viewing distance, with head movements minimised using forehead and chin rests. The experiment took place in a dark room, and responses were recorded with a keypad.

2.2.1.3 Stimuli

Eight male and eight female faces were selected from a bank of faces created by researchers at the UCLouvain (as in Laguesse et al., 2012), on the basis that they had received ratings of more than eight out of ten for either maleness or femaleness in a

separate study. Consequently, the task would involve a binary judgement⁴ to enable the fitting of psychometric functions, but would avoid effects tied to specific identities. All faces were grayscale, front-facing, and had a neutral expression (Figure 2.1B). Using Adobe Photoshop CS6, each face was edited into an egg-shaped aperture measuring 657 x 877 pixels (at its widest and highest point, respectively) so that the only differences between images were due to internal features and not outer face shape (e.g. jawline). The faces were set to have the same mean luminance as the monitor, with matched root-mean square (RMS) contrast values of 0.68. This ensured that overall luminance or contrast values could not be used as cues to gender.

2.2.1.4 Procedure

Participants were instructed to fixate on a white two-dimensional Gaussian element (standard deviation of 13.8 minutes of arc) in the centre of the screen. During each trial, a face was presented for 500 ms, with the image centre located at 10° eccentricity and at one of eight possible angles (Figure 2.1). The presentation length of the stimulus was




Figure 2.1. **A.** The eight polar angles tested, starting with 0° in the right visual field and preceding counterclockwise in 45° increments. The horizontal and vertical meridians are represented by pink and purple dashed lines, respectively. **B.** Examples of female (left) and male (right) face stimuli. **C.** Experimental paradigm. A Gaussian fixation point first appeared, then a face was presented for 500 ms at one of the eight possible locations (shown here at 180°). Each face was followed by a mask which remained on screen until a keyboard response was made. Faces varied in size from trial to trial according to an adaptive QUEST procedure.

⁴ Although gender itself is not binary, a binary judgement was required for psychometric function fitting, hence the choice to use faces at each end of the gender spectrum.

chosen to ensure that the task would not be too difficult to perform in peripheral vision. The face was immediately followed by a 1/f noise egg-shaped mask, which broadly matches the spatial frequency content of faces and natural scenes (Párraga et al., 2000). The size of the mask varied trial-by-trial to match the size of the face just shown, remaining on screen until participants made their response. A single interval two-alternative forced choice (2AFC) response method was used, with participants reporting the face as either male or female using a numeric keypad. Audio feedback was provided after each response.

Before experimental trials began, participants completed a shorter set of 72 practice trials to become accustomed to the task. For the practice trials, faces were presented at fixed sizes of 600, 400 and 200 pixels (face size refers to vertical height, with width scaled proportionately), with nine trials at each location. Participants were required to be at least 90% correct in order to continue. During the experimental trials an adaptive QUEST procedure (Watson & Pelli, 1983) varied face sizes presented at each location according to the participant's responses, set to converge on the size at which 75% of responses were correct. Within each block of trials, QUEST estimates were computed separately for each location. Faces were presented at sizes within $\pm 1/3$ of the QUEST threshold estimate on each trial (minimum five and maximum 640 pixels). This "jitter" allowed data to be collected for a range of sizes, which improved the subsequent fitting of psychometric functions to the data (Tailor et al., 2021).

Each experimental block contained 50 faces shown at each of the eight locations (with independent QUEST procedures) to give 400 trials in total. Each face was shown an equal number of times, in a randomised order, with the location it appeared at also randomised. Upright and inverted faces were presented in alternate blocks. The experiment consisted of one or two practice blocks, followed by eight experimental blocks (four repeats for both upright and inverted faces) to give 3200 trials in total. During analysis, psychometric functions were fit to the combined data from these four repeats (separately for each location and inversion condition).

2.2.1.5 Analyses

Responses were first sorted by face size (in pixels) and collated in 20-pixel bins (e.g. faces of eight, 15 and 18 pixels would fall in the same bin). The proportion of correct

responses was then calculated for each face-size bin. Cumulative Gaussian functions (Figure 2.2) were fit to these data using three free parameters for the mean, variance and lapse rate (Greenwood, 2023). For some participants with particularly noisy data, a maximum allowable lapse rate of either 0.05 or 0.1 was also applied in order to improve curve fitting. Gender acuity thresholds for each location were taken as the size at which 75% accuracy was reached, then converted from pixels to degrees of visual angle.

Statistical analyses were carried out using a 3-way mixed effects analysis of variance (ANOVA), with participant as a between-subjects random factor and inversion (upright, inverted) and location (0, 35, 90, 135, 180, 225, 270, 315°) as within-subjects fixed factors. A priori and post-hoc comparisons took the form of repeated-measures t-tests, comparing thresholds between the horizontal (0 and 180°) vs. vertical (90 and 270°) meridians, the upper (90°) and lower (270°) field, and the left (180°) and right (0°) locations, for upright and inverted faces separately.




Figure 2.2. Psychometric functions for a single participant, showing the proportion of correct gender judgements for different sized upright faces at each of the eight visual field locations (labelled at the top of each graph). Performance improves monotonically as a function of face size. Dashed grey lines plot thresholds for gender acuity (the size at which 75% accuracy was reached).

2.2.2 Results

Mean gender acuity thresholds across participants are plotted as both a bar chart and according to the polar angle of each of the eight locations in Figure 2.3. Smaller values




Figure 2.3. Mean gender acuity thresholds (in degrees of visual angle) measured in Experiment 1, plotted in two ways. Firstly, as a bar graph (A) with each angular location indicated via colour (see legend). Dots represent individual thresholds for each participant. Secondly, as a polar angle plot (B), where 0° is at the right and angles proceed counterclockwise in 45° jumps. Upright faces are shown in purple and inverted in pink. Shaded regions denote ± 1 SEM.

represent better gender acuity. Mean gender acuity thresholds were worse for inverted compared to upright faces overall, indicating that inversion disrupted the ability to judge gender at all locations. There was a sizeable difference in gender acuity according to location; for upright faces, there was a range of almost 2° of visual angle between the smallest threshold value in the lower field (270°) and the largest in the upper field (90°). Thresholds were smaller along the horizontal (0, 180°) as opposed to the vertical (90, 270°) meridian. Thresholds at the diagonal locations varied inconsistently.

The ANOVA revealed a main effect of location, $F(7,91) = 3.41$, $p = .003$, $d = 0.21$, confirming that the location of faces in the visual field influenced gender perception. Planned contrasts revealed that thresholds were significantly smaller along the horizontal (0° and 180° averaged) compared to the vertical (90° and 270° averaged) meridian for both upright, $t(13) = -2.84$, $p = .014$, and inverted faces, $t(13) = -2.21$, $p = .046$. Thresholds were also smaller in the lower compared to the upper field for upright faces, $t(13) = 2.68$, $p = .019$, although not for inverted faces, $t(13) = -0.10$, $p = .923$. There was no difference between thresholds at the left and right locations for upright, $t(13) = -1.79$, $p = .096$, or inverted faces, $t(13) = -0.61$, $p = .551$. In other words, both horizontal-vertical and upper-

lower anisotropies were observed for gender acuity, though performance did not differ between left and right hemifields.

Figure 2.3 indicates that there was considerable between-participants variability in gender acuity thresholds, with the ANOVA showing a main effect of participant, $F(7,91) = 2.98$, $p = .029$, $d = 0.75$. However, there was no interaction between location and participant, $F(91,91) = 1.07$, $p = .383$, $d = 0.52$, indicating that individuals varied in their overall threshold magnitude rather than exhibiting idiosyncratic variations across the eight locations. Of particular note are a subset of participants who showed thresholds that were considerably smaller (with values averaging $\sim 0.5^\circ$) than the rest of the group.

The presence of an inversion effect was supported by a main effect of inversion, $F(1,13) = 6.61$, $p = .023$, $d = 0.34$, showing that thresholds were significantly higher for inverted compared to upright faces. Therefore, configural information appears to have




Figure 2.4. Face inversion effects for each participant in Experiment 1, calculated as the mean gender acuity thresholds (in degrees of visual angle) for inverted faces minus the mean thresholds for upright faces. Each participant is plotted separately (see colour legend). The red dashed line represents the mean face inversion effect across participants.

benefitted performance in the task. There was no interaction between inversion and location, $F(7,91) = 0.91$, $p = .506$, $d = 0.07$, indicating that inversion disrupted gender perception to a similar extent across the visual field. There was, however, a significant interaction between inversion and participant, $F(13,91) = 10.55$, $p = <.001$, $d = 0.60$.

To investigate individual inversion effects further, mean face inversion effect (FIE) values were calculated across participants by subtracting upright from inverted thresholds (Figure 2.4). Upon closer analysis the inversion-participant interaction appeared to be driven by a couple of participants with very large FIEs, and indeed I found that removing their data from the analysis eliminated the interaction.

2.3 Experiment 2

Experiment 1 demonstrates that acuity for face perception varies in the same way as low-level vision, with both horizontal-vertical and upper-lower anisotropies. Experiment 2 was designed to test the validity of these findings. Firstly, in Experiment 1, the thresholds of several participants suggested that they could correctly judge the gender of faces as small as 0.5° in size. Although it is not impossible that faces that small could be recognised at 10° eccentricity, the divergence from the group average suggested that they may have been directly fixating the faces. Second, although the upper-lower difference was significant for upright faces, the face stimuli were centred on the image centre. This meant that for upright faces the eyes appeared closer to fixation in the lower vs. the upper field, and vice versa for inverted faces. This could potentially have driven the upper-lower difference, given the importance of the eye region for gender perception (Brown & Perrett, 1993; Schyns et al., 2002; Yamaguchi et al., 2013). In Experiment 2 I therefore added eye-tracking and controlled for eye position by centring the locations of face stimuli on the eyes.

2.3.1 Method

2.3.1.1 Participants

14 participants (12 female, two male, $M_{\text{age}} = 23.6$ years) took part, including myself; the rest were naïve and newly recruited. As before, all had normal or corrected-to-normal vision. Seven were right-eye dominant.

2.3.1.2 Apparatus

In addition to the setup in Experiment 1, an EyeLink 1000 (SR Research, Mississauga, ON, Canada) was used to monitor fixation during trials.

2.3.1.3 Stimuli

The same face stimuli from Experiment 1 were used. The position of each face in the egg aperture was shifted and/or rotated slightly if needed, so that the position of the eyes within the egg aperture was as similar as possible across faces. Stimulus locations were now centred on the eyes themselves, so that the centre of the eyes was always 10° from fixation regardless of face size, angular location or inversion (Figure 2.5).

2.3.1.4 Procedure

Following calibration of the EyeLink to track their left eye, participants were required to fixate the Gaussian element (with an allowable error of 1.5° radius) in order for each trial to start. Trials in which fixation diverged from this region were cancelled and repeated




Figure 2.5. In Experiment 1 (left panels), faces were presented 10° from fixation according to the centre of the face image. In Experiment 2 (right panels), the centre of the eyes was always 10° from fixation regardless of face size, angular location or inversion.

at the end of the block. Participants first completed the practice block(s), as in Experiment 1. Experimental blocks were split according to whether faces were upright or inverted and locations were cardinal (0, 90, 180, 270°) or diagonal (45, 135, 225, 315°) angles, resulting in four blocks: upright cardinal, upright diagonal, inverted cardinal and inverted diagonal. These split blocks were introduced so that the blocks would not be too long, as the eye tracking increased the duration of data collection. Data were collected over four hour-long sessions, with each of the four conditions repeated once per session. This gave a total of 16 blocks and 4096 trials.

2.3.2 Results

Mean gender acuity thresholds are plotted in Figure 2.6. Again, smaller values represent better gender acuity. Compared to Experiment 1, thresholds in Experiment 2 were higher overall and had reduced variability across participants (particularly in the inverted condition), suggesting that eye tracking successfully stopped participants from looking directly at faces. Indeed, the smallest-measured threshold in Experiment 2 was 1.31°, compared with 0.27° in Experiment 1.

The ANOVA revealed a main effect of location, $F(7,91) = 5.55$, $p < .001$, $d = 0.30$, indicating that gender acuity was influenced by the location of faces. There was a clear horizontal-vertical difference, with planned comparisons revealing that thresholds were significantly smaller along the horizontal (0° and 180° averaged) compared to the vertical (90° and 270° averaged) meridian, for both upright, $t(13) = -6.16$, $p < .001$, and inverted faces, $t(13) = -3.00$, $p = .010$. However, although thresholds were smaller in the lower compared to the upper field, the difference was not significant for either upright, $t(13) = 1.19$, $p = .256$, or inverted faces, $t(13) = 0.38$, $p = .713$. Similarly, thresholds did not differ between the left and right locations for either upright, $t(13) = -0.10$, $p = .926$, or inverted faces, $t(13) = 0.39$, $p = .704$.

An overall inversion effect can be seen in the mean data (Figure 2.6A and B), with higher thresholds and therefore reduced ability to perceive gender (i.e. larger faces needed) for inverted compared to upright faces. This was confirmed by a significant main effect of inversion, $F(1,13) = 17.93$, $p = .001$, $d = 0.58$. There was no interaction between inversion and location, $F(7,91) = 1.29$, $p = .265$, $d = 0.09$, indicating that inversion effects did not differ significantly across the visual field.




Figure 2.6. Mean gender acuity thresholds from Experiment 2, first shown as a bar graph (**A**) with each location indicated via colour (see legend). Dots represent individual thresholds for each participant. Mean thresholds are also visualised in a polar angle plot (**B**), with 0° at the right and angles increasing counterclockwise by 45° each time. Upright faces are shown in purple and inverted in pink. Shaded regions represent ± 1 SEM. (**C**) Bar charts comparing the horizontal-vertical difference and upper-lower difference in Experiments 1 and 2. Data are plotted for upright faces only. Horizontal refers to thresholds averaged across 0° and 180° locations, with vertical the average of 90° and 270°. Upper represents 90° and lower 270°. Dots represent individual thresholds for each participant. Significant differences are marked with an asterisk.

Overall gender-recognition abilities again differed between individuals, which was confirmed by a main effect of participant, $F(7,91) = 11.52$, $p < .001$, $d = 0.93$. There was however no interaction between location and participant, $F(91,91) = 0.94$, $p = .623$, $d = 0.48$, indicating that location-based variations in gender perception are not wholly specific to the individual. In other words, face perception differed across the visual field in a characteristic pattern, shared across participants.




Figure 2.7. Face inversion effects for each participant in Experiment 1, calculated as the mean gender acuity thresholds (in degrees of visual angle) for inverted faces minus the mean thresholds for upright faces. Each participant is plotted separately (see colour legend). The red dashed line represents the mean face inversion effect across participants.

Unlike the previous experiment there was no interaction between inversion and participant, $F(13,91) = 1.60$, $p = .100$, $d = 0.19$, indicating that individuals did not vary substantially in their inversion effects (Figure 2.7). This suggests that the significant interaction in Experiment 1 may have been caused by a subset of participants looking at the faces – accordingly, these individuals had low thresholds for both upright and inverted faces (i.e. little to no inversion effect), likely driven by them fixating the faces in both conditions.

To compare anisotropies more clearly, bar charts displaying the horizontal-vertical difference and upper-lower difference in both experiments are shown in Figure 2.6C. Only data for upright faces are included. The charts on the left-hand side show that the horizontal-vertical difference was consistent across both experiments, with significantly lower thresholds (better gender acuity) for faces at horizontal compared to vertical

locations. However, the upper-lower difference was only significant in Experiment 1, where effects could have been driven by variations in eye position within face stimuli. Although a trend in the same direction persisted when these factors were controlled for in Experiment 2 – showing that the upper-lower difference cannot be attributed to these factors alone – the difference was no longer significant.

To summarise, gender acuity was better along the horizontal vs. vertical meridian but did not differ significantly in the lower vs. upper field. However, even in Experiment 1 the upper-lower difference was smaller than the horizontal-vertical difference, suggesting that it may simply be harder to measure. This possibility was examined in Experiment 3.

2.4 Experiment 3

Given the trend towards better gender acuity in the lower vs. upper field in Experiment 2, I next conducted further measurements at these locations to determine whether a significant difference would emerge with a greater number of trials.

2.4.1 Method

2.4.1.1 Participants

14 participants (11 female, two male, one non-binary, $M_{\text{age}} = 22.1$ years) took part, including myself; the rest were naïve and newly recruited. Again, all had normal or corrected-to-normal vision. Seven were right-eye dominant.

2.4.1.2 Stimuli

Stimuli were as in Experiment 2, with faces shown at the upper (90°) and lower (270°) locations only.

2.4.1.3 Procedure

Blocks were split according to whether faces were upright or inverted, with 128 trials in each block. Data were collected over two hour-long testing sessions, with each of the two conditions (upright/inverted) repeated eight times per session (with an extra block completed at the start of the first session, which acted as a practice block and was not included in data analysis). This gave a total of 16 experimental blocks and 2048 trials over

the experiment (doubling the number of trials per location compared to Experiment 2). Remaining parameters were as in Experiment 2.

2.4.2 Results

Mean gender acuity thresholds are plotted in Figure 2.8, with smaller values representing better gender acuity. The ANOVA revealed a main effect of location, $F(1,13) = 22.97$, $p < .001$, $d = 0.91$, indicating that gender acuity differed between the upper and lower fields. Gender acuity thresholds were significantly smaller in the lower field compared to the upper for upright faces, $t(13) = 3.82$, $p = .002$, and approached significance for inverted faces, $t(13) = 2.07$, $p = .059$.




Figure 2.8. Mean gender acuity thresholds for the upper (90°; pink) versus lower (270°; purple) visual field, for both upright and inverted faces. Dots represent individual thresholds for each participant. The asterisk represents a significant difference.

On an individual level, 12 of 14 participants showed better acuity in the lower vs. the upper field. These results highlight the presence of an upper-lower difference in face recognition, with better gender acuity in the lower half of the visual field.

Like the previous two experiments, there was a main effect of orientation, $F(1,13) = 8.38$, $p = .013$, $d = 0.39$, indicating an overall inversion effect whereby gender acuity thresholds were larger for inverted compared to upright faces, $t(27) = -2.92$, $p = .007$. There was also a main effect of participant, $F(1,13) = 31.12$, $p = .032$, $d = 1.00$, again highlighting overall differences in gender acuity between individuals. There was no significant interaction between location and participant, $F(1,13) = 0.58$, $p = .829$, $d = 0.37$, suggesting that there was a common pattern of gender acuity across individuals. Interactions were similarly non-significant for location and orientation, $F(1,13) = 0.91$, $p =$

.358, $d = 0.07$, and orientation and participant, $F(1, 13) = 1.02$, $p = .484$, $d = 0.51$, indicating that gender acuity patterns were similar for upright and inverted faces, and that inversion effects did not vary significantly between individuals.

2.5 Discussion

This chapter demonstrates that face perception varies across the visual field in a systematic pattern, analogous to that of low-level vision. Across three experiments, acuity for judging facial gender showed a horizontal-vertical anisotropy – accurate recognition was possible with smaller faces on the horizontal vs. the vertical meridian. Overall, a small-but-reliable upper-lower difference was also present, with better acuity in the lower vs. the upper field. This upper-lower difference persisted when controlling for factors such as participants' eye movements and the eye position within face stimuli, though its measurement required an increase in the number of trials. The presence of both these anisotropies, and the smaller magnitude of the upper-lower than the horizontal-vertical difference, matches the patterns found in low-level vision (Barbot et al., 2021; Kurzawski et al., 2021). That is, instead of varying uniquely or idiosyncratically (Afriz et al., 2010; Quek & Finkbeiner, 2016; Schmidtman et al., 2015), at least one aspect of face perception – judging gender – varies predictably across the visual field in the same fashion as lower-level abilities (e.g. Abrams et al., 2012; Benson et al., 2021; Carrasco et al., 2001). This suggests that spatial properties are preserved through the visual hierarchy and inherited by face processing systems.

Why were systematic variations in face perception found, rather than purely idiosyncratic variations? Other studies measured judgements of appearance (Afriz et al., 2010; Visconti di Oleggio Castello et al., 2018), which can show pronounced perceptual idiosyncrasies in low-level vision (Moutsiana et al., 2016). Similarly, perceived object position has been shown to vary across participants due to individual differences in localised distortion across the visual field (Kosovicheva & Whitney, 2017). Using judgements of appearance to measure subtle perceptual biases in face perception may introduce similar individual-dependent distortions. Indeed, idiosyncrasies in identity judgements for specific face pairs did not correlate with other face pairs, suggesting that perceptual biases were tied to particular face stimuli (Visconti di Oleggio Castello et al., 2018). This could be driven by variations in factors like heightened contrast in face

features like the eyebrows or lips, which would skew judgements towards ‘female’ (Russell, 2009). The use of faces morphed between one male and one female prototype (Afraz et al., 2010) could have similarly increased the susceptibility to stimulus-based and individual-based distortion effects. In contrast, I used categorical judgements with multiple unambiguously gendered faces, measuring variations in performance instead of judgments of appearance. This method would be less susceptible to smaller individual-dependent variations, allowing me to uncover the larger systematic anisotropies observed.

Why was gender acuity not found to differ between the left and right visual fields? Others have reported that the identity of synthetic-contour faces was recognised more accurately in the left visual field, compared to the right, upper and lower locations (Schmidtman et al., 2015). These faces had been filtered to only contain a narrow band of intermediate spatial frequencies, which removes details that may usually be used during face perception, such as skin texture and contrast around the features (Wilson et al., 2002). As such, the left hemifield advantage measured for these cartoon-like faces may not be representative of the full spectrum of face processing (Schmidtman et al., 2015). For example, while the intermediate spatial frequencies may have enabled subtle changes in identity to be detected, low spatial frequencies have been shown to contribute considerably to holistic processing (Goffaux et al., 2005; Goffaux & Rossion, 2006). Other findings suggest that lateralisation effects are not always present for face perception and may indeed depend on the task involved (Bourne et al., 2009; Kovacs et al., 2017). The faster recognition of face gender previously found in the upper vs. lower field (Quek & Finkbeiner, 2014, 2016) could also be linked to methodological approach, reflecting a more general upper field bias for temporal processing that has been observed during lower-level tasks (Abegg et al., 2015; Honda & Findlay, 1992). The findings in this chapter show that when the spatial resolution of face perception is measured, similar visual field anisotropies to low-level vision emerge.

In Experiments 2 and 3, the upper-lower difference was reduced when the position of the eyes within face stimuli was matched across locations, relative to Experiment 1 where faces were centred on the nose. This effect of eye position highlights a role for specific face features in gender perception, and confirms a particular importance of the eyes (Brown & Perrett, 1993; Schyns et al., 2002; Yamaguchi et al., 2013). Across all

experiments, acuity thresholds were nonetheless consistently lower for upright compared to inverted faces, which may reflect an added benefit of configural processing when faces were upright. Consistent with prior work (Kovacs et al., 2017; McKone, 2004), this inversion effect suggests that the configural processing of upright faces is maintained across both foveal and peripheral vision. This suggests that the gender acuity task was sufficient to engage these face-specific processes.

Inversion effects of this nature are often argued to represent a qualitative change in face processing between upright and inverted faces (Rossion, 2008). While inversion effects are typically measured as percent-correct performance using faces of a constant size (Yin, 1969), here I highlight a spatial component to face recognition abilities. Namely, in the face acuity task, percent-correct performance for upright and inverted faces could be matched by increasing the size of inverted faces. The requirement that inverted faces be larger than upright faces could reflect the added benefit of configural processing – or in a quantitative view, more efficient engagement (Sekuler et al., 2004) – of configural processing, which may allow the gender of upright faces to be perceived at smaller sizes. In line with this, fMRI evidence shows that face inversion reduces visual field coverage driven by smaller population receptive field (pRF) sizes (Poltoratski et al., 2021). Impaired perception of inverted faces could therefore reflect sampling by fewer neurons, with this more localised processing reduced by increasing face size.

While inversion effects are often attributed to differences in configural processing between upright and inverted faces (Bartlett & Searcy, 1993; Rossion, 2008), it is also possible that the improved acuity for upright faces resulted from non-face-specific processes, such as a general effect of greater experience with upright rather than inverted face features. To confirm whether the inversion effects observed can indeed be attributed to differences in configural processing, it would be desirable to modify the acuity task to include both intact faces and individual face features. If there are large inversion effects for intact faces but not face features, this would indicate that there is better gender acuity for upright vs. inverted faces due to the added benefit of configural information.

Along similar lines, what might be driving the anisotropies observed for face recognition around the visual field? One possibility is that they are driven by variations in featural selectivity. Contrast sensitivity and orientation discrimination are better for

peripheral stimuli that are radially as opposed to tangentially oriented, known as the *radial bias* (Rovamo et al., 1982; Westheimer, 2003). This could explain why larger inversion effects have been found along the horizontal vs. vertical meridian (Roux-Sibilon et al., 2023), as horizontal information is particularly diagnostic of various aspects of face perception, including identity recognition and holistic processing (Dakin & Watt, 2009; Goffaux & Greenwood, 2016). Given the hierarchical, feedforward nature of the visual system, face processing systems could inherit featural selectivity such as the radial bias through the passive pooling of low-level information (Loffler et al., 2003; Riesenhuber & Poggio, 1999). Alternatively, the process could be more active; given that horizontal spatial frequencies are especially important for face perception (Dakin & Watt, 2009; Goffaux & Greenwood, 2016) and elicit stronger responses in the FFA (Goffaux et al., 2016), a boost in the selection of these orientations could facilitate face recognition. However, though these factors could contribute to the horizontal-vertical anisotropy, radial variations of this kind would be matched in the upper and lower visual fields, making them unlikely to explain the upper-lower anisotropy.

Anisotropies for face recognition could also occur through variations in *spatial selectivity* around the visual field. In low-level vision, better acuity across location has been linked to smaller receptive field size and a larger cortical magnification factor (Benson et al., 2021; Duncan & Boynton, 2003; Silson et al., 2018; Silva et al., 2018). Similar to above, improved acuity for faces could arise simply because higher levels *passively* inherit the enhanced low-level input from these locations (Riesenhuber & Poggio, 1999). Unlike low-level vision, however, better face perception has been linked with larger pRF sizes and the resulting increase in visual field coverage within face-selective regions (Poltoratski et al., 2021; Witthoft et al., 2016). The anisotropies I measured for face perception could therefore stem from the way that face-selective neurons *actively* sample the visual field, with variations in coverage across different locations. Given that the same pattern of anisotropies was found for upright and inverted faces, the findings in this chapter either suggest that these spatial properties vary in the same way for neurons sensitive to upright and inverted faces, or as above that these effects are driven by the quality of the information inherited from lower-level stages of processing. Either way, the similar spatial selectivity uncovered suggests that the retinotopy within face-selective regions is not entirely distinct from earlier brain areas.

The findings in this chapter suggest that face processing systems inherit the spatial selectivity of earlier visual regions, by showing that the ability to judge face gender varies systematically across the visual field in the same patterns as low-level visual abilities. However, face perception also involves other judgements, such as the recognition of identity or emotional expression. Determining whether similar patterns of spatial selectivity are also inherited for the recognition of face identity and expression would require further exploration, with some modifications to the face acuity task. For identity recognition this could involve a face matching task whereby on each trial, a face in the periphery is judged as the same or different as a face briefly shown at fixation – this would avoid idiosyncrasies tied to specific face stimuli, as multiple faces could be used within the identity recognition task. Finding similar horizontal-vertical and upper-lower anisotropies for identity recognition would suggest that for faces, these variations are not specific to gender recognition. Instead, patterns of spatial selectivity may be inherited by face processing systems more generally, and maintained within the different aspects of face perception.

In conclusion, face gender perception varies around the visual field with both horizontal-vertical and upper-lower anisotropies, matching patterns consistently found for low-level vision (Abrams et al., 2012; Barbot et al., 2021; Benson et al., 2021; Carrasco et al., 2001) and contrary to suggestions that face perception varies in a unique or entirely idiosyncratic manner (Afraz et al., 2010; Quek & Finkbeiner, 2016; Schmidtman et al., 2015). These results are consistent with a hierarchical model of face processing whereby spatial selectivity for faces is built on the selectivity of earlier levels. The resulting variations in high-level vision could be driven by variations in featural and/or spatial selectivity, and inherited passively or actively. Although the mechanisms underlying these variations may differ from those of low-level vision, the finding that these face recognition anisotropies match those of low-level vision, both for upright and inverted faces, points to the possibility of a common basis. Ultimately, the experiments in this chapter demonstrate that spatial variations are found throughout the visual system, causing location to influence face perception.

Chapter 3

Measuring the spatial properties of face-selective brain regions

3.1 Introduction

As discussed in the previous chapter, low-level vision varies across the visual field in characteristic patterns. At the same eccentricity, aspects of vision such as acuity and contrast sensitivity are reliably better along the horizontal vs. the vertical meridian (horizontal-vertical difference) and in the lower vs. upper visual field (upper-lower difference; Abrams et al., 2012; Barbot et al., 2021; Carrasco et al., 2001; Himmelberg et al., 2020). Across three behavioural experiments in Chapter 2, I uncovered a similar pattern of systematic variation for face perception, with a clear advantage for judging face gender along the horizontal vs. vertical meridian and a smaller-yet-consistent advantage in the lower vs. upper field. What is the neural basis of these anisotropies found for faces? In this chapter, I use retinotopic mapping to investigate whether the spatial properties of face-selective brain regions can explain the systematic variation found for face perception.

As covered in Chapter 1, anisotropies in low-level vision have been linked to the way that neurons in early visual cortex differentially sample the visual field. Population receptive field (pRF) mapping – a technique used to identify populations of neurons which respond to each region of visual space (Dumoulin & Wandell, 2008) – has revealed that there are a greater number of pRFs located along the horizontal vs. the vertical meridian in V1 (Amano et al., 2009; Arcaro et al., 2009). Variations in neuronal density also affect visual field coverage – an estimate of how well each part of the visual field is sampled, based on the amount of pRFs in each region and their size – which was accordingly found to be better along the horizontal vs. vertical meridian (Amano et al., 2009). These spatial variations in pRF properties are consistent with recent research showing that V1 has a greater surface area along the horizontal vs. vertical meridian and in the lower vs. upper field (Himmelberg et al.,

2023). Similarly, a greater cortical magnification factor – the amount of cortex dedicated to processing each part of the visual field (Daniel & Whitteridge, 1961) – has been found along the horizontal vs. the vertical meridian and in the lower vs. upper field in V1-V3 (Silva et al., 2018), which correlates with better acuity (Duncan & Boynton, 2003). Together, these findings suggest that variations in low-level vision could be linked to the number of neurons in early visual cortex which sample each region of the visual field.

Better acuity has also been linked to smaller receptive field size in V1 (Duncan & Boynton, 2003; Silva et al., 2021). Accordingly, smaller pRFs have been identified along the horizontal vs. vertical meridian in V1-V3 (Silva et al., 2018). In V1, the magnitude of the horizontal-vertical difference in pRF size increased with eccentricity, which could explain why behavioural anisotropies in low-level vision become more pronounced further into the periphery (Greenwood et al., 2017). Similarly, pRFs in early visual cortex that sample the lower field have been found to be smaller and less elliptical (therefore covering a smaller area) than those sampling the upper field, which could explain the lower field advantage for low-level stimuli (Silson et al., 2018; Silva et al., 2018). This again links variations in low-level vision to differential sampling of the visual field, with smaller receptive fields being beneficial for perception.

Could the anisotropies that were found for face perception in Chapter 2 also arise due to sampling variations? Face-selective parts of the brain also show retinotopic sensitivity (Silson et al., 2016; Silson et al., 2022), but appear to sample the visual field differently from earlier regions. While the fovea is better represented than the periphery throughout the visual cortex (Harvey & Dumoulin, 2011), pRFs in V1 are still relatively spread out across the visual field, providing coverage in peripheral vision (Arcaro et al., 2009; Kay et al., 2015). In comparison, the bias towards a better representation of the central visual field is considerably magnified within face-selective brain regions, with the vast majority of receptive fields positioned near the fovea (Silson et al., 2022). Studies show that in mFus, which is part of the fusiform face area (FFA), around 80% of pRFs were located with their centres less than 5° from fixation, with few pRF centres beyond 10° (Finzi et al., 2021; Kay et al., 2015). A similar exaggeration has been observed for visual field coverage – while approximately 15% of pRFs were

found to be positioned in the fovea in V1, there were more than 80% in mFus (Poltoratski et al., 2021). While most of these studies only performed retinotopic mapping at a maximum eccentricity of around 5-7° (Gomez et al., 2018; Kay et al., 2015; Poltoratski et al., 2021), they suggest that receptive fields within face-selective cortex are heavily centred on the fovea, with limited coverage of peripheral vision.

The relationship between pRF size and perceptual abilities may differ considerably in face processing systems compared to low-level vision. pRFs are much larger in face-selective brain regions than they are in V1, and increase in size more rapidly with eccentricity (Finzi et al., 2021; Kay et al., 2015; Poltoratski et al., 2021). Larger receptive fields have been linked to better face recognition, with smaller pRFs for *inverted* compared to upright faces in the FFA (Poltoratski et al., 2021). When inverted faces were shown, pRFs were smaller in size and shifted downwards, away from the fovea. Overall visual field coverage in the face-selective areas was therefore reduced for inverted faces. The authors propose that larger receptive fields aid the configural processing of faces by enabling the spatial integration of face features over a larger area, and that this explains the poorer recognition of inverted compared to upright faces (Yin, 1969), with smaller receptive fields less able to integrate information. Poltoratski et al.'s (2021) findings suggest that links between neuronal properties and face perception may be dependent not only on spatial selectivity within face-selective brain regions, but selectivity based on the featural *content* of faces. Although pRFs in V1 can shift position based on attention (Klein et al., 2014), they generally show consistent spatial selectivity regardless of content (Kay et al., 2015) – on the other hand, face-selective areas may exhibit featural selectivity in addition to, or in place of, spatial selectivity.

Further evidence for an unusual relationship between pRF size and face perception comes from prosopagnosia. Smaller pRFs have been found within the face-selective brain regions – but not V1-V3 – of individuals with developmental prosopagnosia (Witthoft et al., 2016). As these individuals also showed a strong foveal bias, the small size of pRFs within their face-selective areas meant that visual field coverage was even more restricted to the centre of vision than it was in controls, with pRFs only covering the central few degrees. These findings are consistent with studies

reporting that those with prosopagnosia are less able to process configural face information, such as the spatial relationships between face features (Behrmann & Avidan, 2005; Busigny et al., 2010; Towler et al., 2017). Overall, research highlights a puzzling dissociation between low- and high-level vision, where smaller receptive fields are associated with better acuity in early visual cortex (Duncan & Boynton, 2003; Silson et al., 2018; Silva et al., 2018; Silva et al., 2021) yet larger receptive fields appear to be beneficial for face processing (Kay et al., 2015; Poltoratski et al., 2021; Witthoft et al., 2016).

Altogether, while the above findings suggest that face-selective brain regions share some hallmarks of retinotopy with early visual cortex, such as increases in receptive field size with eccentricity, they also show distinct differences (Kay et al., 2015). This raises the question of how the anisotropies found for face perception in Chapter 2 – which closely resembled those of low-level vision – are supported by visual field sampling in face-selective areas. Firstly, the cortical magnification of the fovea is dramatically enhanced within face-selective regions relative to V1, suggesting that they have an impoverished representation of the periphery (Finzi et al., 2021; Gomez et al., 2018; Poltoratski et al., 2021). Yet, the results in Chapter 2 along with previous research show that we can successfully recognise faces in peripheral vision, with similar inversion effects – a hallmark of face processing – found in the periphery (Kalpadakis-Smith et al., 2018; Kovacs et al., 2017; McKone, 2004; Roux-Sibilon et al., 2023). How do these peripheral face recognition abilities arise, if face-selective regions sample the periphery so poorly? Secondly, face-selective areas sample the visual field with much larger receptive fields than earlier areas, and while smaller receptive fields benefit acuity in V1 (Duncan & Boynton, 2003; Silva et al., 2021), they have been associated with *worse* face recognition (Poltoratski et al., 2021; Witthoft et al., 2016). Does this mean that better face perception along the horizontal meridian and in the lower field would be linked to *larger* pRFs, contrary to patterns found in V1?

To investigate these questions I carried out retinotopic mapping of early visual cortex (V1-V3) and three face-selective regions within ventral temporal cortex: the occipital face area (OFA) in the inferior occipital gyrus, and two regions in the posterior (pFus) and medial (mFus) fusiform gyrus which comprise the FFA. As mentioned

earlier, most other studies have mapped face-selective areas to a maximum of around 5-7° eccentricity (Gomez et al., 2018; Kay et al., 2015; Poltoratski et al., 2021). Evidence suggests that the properties of pRFs which only receive partial stimulation – due to sampling an area of visual space closer to the edge of the display – are estimated less accurately (Alvarez et al., 2015). This would be a particular concern within face-selective parts of the brain, due to their large receptive fields (Kay et al., 2015; Poltoratski et al., 2021). So that we could accurately map peripheral visual field locations, a large field of view (21.65° eccentricity) was used. The spatial properties of the three face-selective areas, and of V1-V3, were assessed through three measures: pRF size, pRF number (quantity), and visual field coverage. I was specifically interested in whether these measures varied according to visual field location and between upright and inverted faces.

As smaller receptive fields have been linked to better low-level acuity, I expected that pRFs in V1-V3 would be smaller along the horizontal vs. vertical meridian and in the lower vs. upper field (Duncan & Boynton, 2003; Silva et al., 2018; Silva et al., 2021). There should also be a greater number of pRFs and better visual field coverage (Amano et al., 2009; Arcaro et al., 2009; Benson et al., 2021; Himmelberg et al., 2023; Silva et al., 2018) along the horizontal vs. vertical meridian and in the lower vs. upper field. pRF properties were not expected to differ for upright and inverted faces in V1-V3, as these areas do not carry out face-specific processing (Kay et al., 2015; Sayres et al., 2010).

The spatial properties of face processing neurons appear to vary according to the perceptual characteristics of faces (e.g. upright vs inverted), with large receptive fields benefitting face perception (Poltoratski et al., 2021; Witthoft et al., 2016). If this is indeed the case, within the face-selective regions there should be smaller pRFs for inverted faces across the visual field, as similar inversion effects have been found in central and peripheral vision (Kovacs et al., 2017; McKone, 2004; Roux-Sibilon et al., 2023). If there is a similar relationship for spatial as well as featural selectivity, face-selective areas should have larger pRFs at locations where better acuity for faces was found (Chapter 2), along the horizontal vs. vertical meridian and lower vs. upper field. This would suggest that there is indeed a dissociation within the visual system, where

smaller receptive fields provide better resolution in low-level vision but larger receptive fields are linked to better acuity for face recognition.

Alternatively, if the dependence on smaller receptive fields for acuity is a common property throughout low- *and* high-level brain regions, then smaller receptive fields should provide better acuity in face-selective areas (as they do in V1; Duncan & Boynton, 2003). In this case, within face-selective cortex there should be smaller pRFs along the horizontal vs. vertical meridian, lower vs. upper field, and for upright vs. inverted faces.

As better low-level acuity is associated with increased neuronal density and visual field coverage in early visual cortex (Amano et al., 2009; Arcaro et al., 2009), I expected that better acuity for face perception would be linked to these properties in a similar manner. In all three face-selective parts of cortex, there should be a greater number of pRFs and increased visual field coverage along the horizontal vs. vertical meridian, and in the lower vs. upper field. As similar behavioural anisotropies were found for upright and inverted faces in Chapter 2, I reasoned that these retinotopic properties would vary across the visual field in a similar way irrespective of face inversion. However, I expected that the face-selective regions would show an overall increase in pRF number and visual field coverage for upright vs. inverted faces, linking these measures to featural selectivity within face perception. Ultimately, my aim was to investigate whether the spatial properties of face processing neurons could explain the behavioural anisotropies I found for faces, and to examine how unique the retinotopy within face processing brain regions really is, compared to earlier visual areas.

3.2 Method

3.2.1 Participants

Ten participants (six female, four male, $M_{\text{age}} = 29.1$ years) took part, all of whom had normal or corrected-to-normal vision. The experiment was approved by the Research Ethics Committee for Experimental Psychology at University College London and all participants gave written informed consent before testing began.

3.2.2 Apparatus

Stimuli were displayed on a back-projection screen in the bore of the magnet using an EPSON EB-L1100U projector that had a maximum luminance of 502 cd/m². The screen size was 27 cm x 27 cm with a resolution of 1920 x 1200 pixels, and stimuli were displayed at 1200 x 1200 pixels. The monitor had a refresh rate of 60 Hz. Participants viewed the screen through a mirror attached to the head coil at a viewing distance of 34 cm, giving a maximum field of view of 43.3° (21.65° eccentricity). Gamma correction was performed so that the grey background of the experimental screen matched the mean luminance of the projector (251 cd/m²).

3.2.3 Stimuli

Stimuli were programmed using MATLAB (MathWorks, Inc) and PsychToolbox (Brainard, 1997; Kleiner et al., 2007; Pelli & Vision, 1997). 15 male and 15 female identities were selected from the Radboud Face Database (Langner et al., 2010). Images were in colour and faces had a neutral expression. To maximise face-selective activation by minimising adaptation effects to faces of the same viewpoint (Fang et al., 2007; Henson, 2016), faces of three viewpoints were used, according to the view of the model: front- (90°), left- (135°) and right-facing (45°). This resulted in a total of 90 face images, which had their background removed and were resized to be 332 x 450 pixels using Adobe Photoshop CS6.


Bars – which have been shown to map eccentricities outside of the fovea more accurately than wedge and ring stimuli (Linhardt et al., 2021) – covered the full field of view in length (43.3°; Figure 3.1A). Some retinotopic mapping studies have used stimuli which are scaled for eccentricity according to V1 cortical magnification factor (e.g. Alvarez et al., 2015), however this would have been insufficient for face-selective regions, where cortical magnification is enhanced (as discussed in Chapter 1; Kay et al., 2015). As misjudging this could affect pRF properties, three different bar thicknesses were used, covering 5.30°, 6.97° and 10.06° in width. Face size was determined by the width of the bar so that faces would fit within the bars as opposed to being cut off. Bars of different thicknesses therefore contained faces of different sizes, so that smaller faces would target smaller pRFs (presumably nearer to the

fovea) and larger faces would target larger pRFs (in the periphery especially). For each bar width, orientation (horizontal and vertical) and face condition (upright or inverted), five bars containing male faces and ten bars containing female faces were generated (male bars appeared less frequently). This made for a total of 180 bars across the


A



B



C



D




Figure 3.1. **A.** Examples of retinotopic mapping stimuli. A blue fixation cross appeared at the centre of the display while bars containing either male or female faces traversed the screen in one of four directions: 0° (rightwards), 90° (upwards), 180° (leftwards) and 270° (downwards). Arrows represent the direction of motion. Different bar thicknesses are shown in different panels. Bars appeared at one location per TR (one second), with various crops shown in quick succession to induce the illusion of motion. **B.** Bar conditions throughout the experiment. For illustration purposes the screen is shown as a circle, with each circle representing one sweep (15 TRs/locations) across the screen. The smallest width bars were shown first, then the medium, then the largest width. Blank periods (15 TRs) occurred after every two sweeps. Arrows represent the bars' direction of movement. **C.** pRF model. The stimulus was converted to a binary aperture image and each pRF modelled as a 2D Gaussian before a static nonlinearity was applied using a compressive spatial summation parameter. The model output is the predicted BOLD response. **D.** Definition of pRF location and size in the compressive spatial summation model. Position is determined by x and y coordinates, while size is the standard deviation (σ) divided by the square root of the spatial summation exponent n . Figures C and D are adapted from Kay et al., (2013).

experiment. To minimise crowding and repetition effects (Fang et al., 2007; Henson, 2016), each bar contained faces of different viewpoints in a pseudo-randomised manner, such that faces of one viewpoint could not appear next to a face of the same viewpoint. To avoid adaptation effects tied to identity, each identity could not appear in the same bar twice (Natu et al., 2016). The position of the faces within the bars was offset along the width of the bar (the x axis for vertical bars and the y axis for horizontal bars) so that they could be moved closer together along the opposite axis, leaving less of the bar as blank space. The value of the offset varied according to the orientation and width of each bar (for example, a greater offset was required for thicker bars). So that faces could be moved along the bars during the experiment (explained in 3.2.4), bars were initially made longer than required and then cropped eight times along the longer axis, with each crop cutting out a bar of the desired length but starting at a different point at equal increments along the bar. This resulted in eight differently cropped versions of each bar. For the inverted face bars, faces were flipped along the vertical axis. The background of each bar matched the grey background of the experimental screen.

3.2.4 Procedure

Each run began with a blank screen for five seconds, containing only a fixation cross in the centre which subtended 0.95° of visual angle. Bars then traversed the screen in four directions: 0° (rightwards), 90° (upwards), 180° (leftwards) and 270° (downwards), appearing in one location per repetition time (TR), which lasted one second (Figure 3.1B). Each sweep across the screen contained 15 equal steps, meaning that the steps became smaller as the bar widths increased. The number of steps was kept the same as otherwise the pRF fitting may have been biased towards bar widths that had more TRs (by contributing more to the minimisation of the least-squared error between the observed data and model prediction). The thinnest bars were presented first (four sweeps, one per direction of motion) before moving on to the next thickness. As there were three bar thicknesses, each run had a total of 12 sweeps. Every second sweep (i.e. after 0° and 90° , then after 180° and 270°) was followed by a blank period of 15 TRs. Each run therefore comprised 275 TRs, lasting four minutes and 35 seconds.

Because the bars contained spaces between the faces, the faces were moved along the length of the bar so that the time-averaged bar for each TR would have contained faces in as much of the bar as possible. Each bar was rapidly presented from the first to the eighth crop and then back to the first, meaning that 15 crops of each bar were shown within quick succession (15 being a multiple of the monitor refresh rate, to ensure that each crop was shown for the exact same amount of time). As such, the faces within the bars appeared to move smoothly side to side (horizontal bars) or up and down (vertical bars) at each location. Most bars contained female faces, while bars consisting of male faces occurred with 0.075 probability. To ensure that attention was directed towards the bars but also encourage fixation, participants were instructed to maintain fixation and respond when a bar containing male faces appeared, or when the fixation cross changed from blue to purple (0.002 probability, and lasting 0.2 seconds). Responses were recorded via a button box. Participants did not receive feedback, however key presses were monitored throughout the experiment to ensure that participants were performing the task properly. Upright and inverted runs were interleaved in order to avoid effects like fatigue disproportionately affecting one condition.

3.2.5 Localisation of face-selective ROIs

Participants also took part in a functional localiser experiment (which is fully described and assessed in Chapter 4) in the same scan session, to identify face-selective regions of interest (ROIs) within the brain. This localiser was newly developed to ensure that voxels responding to locations in the peripheral visual field would be captured, as we suspected this might not be the case with existing localisers (e.g. Stigliani, 2015). Stimuli were displayed using the same maximum eccentricity (21.65°) as the retinotopic mapping stimuli, so images covered a 43.3° field of view. Images of faces, hands and instruments were shown in order to identify face-selective brain regions by contrasting the activation in response to faces against the other object categories (Kanwisher et al., 1997; Weiner & Grill-Spector, 2010). To maximise both foveal and peripheral stimulation, two configurations of stimuli were used: large, single faces centred on the fovea, and smaller faces which were tiled across the screen in a nine-by-nine grid (and the same for hands and instruments). Faces, hands and

instruments were shown in separate blocks lasting ten seconds each, interspersed by baseline (blank) periods. Within each block there were 20 stimuli displayed for 500 ms each. Single and tiled configurations were presented in the same run, in different blocks. Each run consisted of 51 blocks, lasting eight minutes and 35 seconds. Participants were instructed to maintain fixation and press a button when a phase-scrambled image appeared. Each participant completed two runs.

To analyse blood-oxygen-level-dependent (BOLD) fMRI responses to the different object categories a general linear model (GLM) was run using SPM12 software (Ashburner & Friston, 2014). Each of the stimulus conditions was entered into a design matrix, with the predicted fMRI time series modelled according to the onset times of each stimulus. Motion regression was accounted for by using the output from the motion correction step carried out in the preprocessing stage, with one regressor added for each of the six directions of movement (roll, pitch and yaw around both x and y axis). The observed fMRI time series was convolved with the canonical haemodynamic response function (HRF) and compared with the predicted time series generated by SPM. Statistical contrasts were then carried out, thresholded at the $t \geq 2$ level. This threshold was chosen to help maximise the number of pRFs remaining for further analyses after filtering by visual field location (described in 3.2.12). In nine participants I defined three face-selective areas (OFA, pFus and mFus) and two in the remaining participant (OFA and mFus) by contrasting faces against other objects (single faces, tiled faces > single hands, tiled hands, single instruments, tiled instruments). Statistical T maps were surface projected using Freesurfer (Fischl, 2012) and used as a visual guide during the delineation of face-selective ROIs. Large areas were initially drawn manually, before an automatic process defined the ROI by identifying the vertex with the peak T statistic in each region, along with any neighbouring vertices that were above the chosen T threshold ($t \geq 2$). This process is described in more detail in Chapter 4.

3.2.6 MRI data acquisition

Functional and anatomical scans were obtained using a Siemens Prisma 3T MRI scanner (Siemens, Erlangen, Germany). A 64-channel head coil was used, with cushions placed around participants' heads in order to minimise movement. A T1-

weighted anatomical magnetisation-prepared rapid acquisition with gradient echo (MPRAGE) image was acquired (TR = 2300 ms and TE = 2.98 ms, voxel size = 1 mm isotropic voxels). For the functional scans only the back of the head coil was used, leaving 42 channels. Functional T2-weighted multiband 2D echoplanar images were acquired (repetition time (TR) = 1000 ms, TE = 35.20 ms, voxel size = 2 mm isotropic voxels, 48 slices, flip angle = 60°, acceleration factor = 4). Each functional scan contained 270 volumes. A short 30 second localiser was carried out before the functional scans and again before the anatomical scan, after the front head coil was fitted. Fixation was monitored throughout the experiment using an Eyelink 1000, although we did not record fixation data. pRF runs were carried out before the functional localiser in all participants except author AYM, as although important for both paradigms, keeping precise fixation was particularly crucial for the pRF mapping.

3.2.7 MRI data preprocessing

For each participant, the T1 anatomical scan was automatically segmented and used to generate a 3D representation of the cortical surface using Freesurfer (Dale et al., 1999; Fischl, 2012; Fischl et al., 1999). Functional images were B0 distortion corrected and motion corrected using AFNI software (Cox, 1996). An alignment volume was created by finding the volume with the fewest voxel outliers across all runs, which all functional volumes were then aligned to. Using Freesurfer (Fischl, 2012) the alignment volume was co-registered to the structural image, and surface projection was performed.

3.2.8 pRF fitting

pRF analyses were carried out using the SamSrf 9.4 MATLAB toolbox (Schwarzkopf, 2022). Similar to the method described in Dumoulin and Wandell (2008), this was based around a forward modelling approach where the fMRI data and the position of the bars within the visual field were used to estimate pRF parameters. The four runs were concatenated before pRF estimation.

A compressive spatial summation (CSS) model was used within SamSrf 9.4 (Figure 3.1C), where each pRF was estimated as a two-dimensional Gaussian with a compressive non-linearity that was subsequently applied (Kay et al., 2013; Poltoratski

et al., 2021). This approach has been found to estimate pRF properties more accurately compared to a linear model (which I also assessed in 3.2.9), particularly in higher visual areas which show nonlinear spatial summation, i.e. where responses to visual stimuli sum in a subadditive manner rather than linearly. The CSS model involved four free parameters: x and y (the position of the pRF within the visual field, with x_0 and y_0 denoting the centre), σ (the standard deviation or spatial spread of the pRF, in degrees of visual angle) and n (the exponent of the compressive non-linearity; Figure 3.1D). Estimates of σ in the CSS model are affected by the compression exponent – pRFs with low exponent values (high compression) can produce strong responses to stimuli throughout the spatial extent of their receptive fields, compared to a Gaussian model which predicts strong responses in the centre and a sharper drop off at the edges (Kay et al., 2013). This may be linked to increased position invariance within category-selective regions (DiCarlo & Cox, 2007). Given the effects of this compression on the spread of the receptive field profile, during analyses pRF size was defined as σ divided by the square root of the exponent (Kay et al., 2013):

$$\frac{\sigma}{\sqrt{n}}$$

A stimulus aperture was created for each run which consisted of a binary mask representing the position (in x and y Cartesian coordinates) of the bar stimulus within the visual field at each time point. Apertures were averaged across the four runs, resulting in one aperture comprising 270 frames (one for each TR). Due to the motion of the bar stimuli (described in 3.2.4) and the averaging, apertures formed solid bars instead of containing space within the faces (Figure 3.1C). The linear overlap between the model estimation and the aperture was used to estimate the response of the underlying neural population at each vertex on the cortical surface.

To compare it to the observed fMRI time series, the model prediction was convolved with a canonical haemodynamic response function (HRF). Research has shown that the differences between using a canonical HRF and measuring the HRF in each individual are minimal, so data analysis should not have been affected by this choice (Dumoulin & Wandell, 2008; van Dijk et al., 2016).

The pRF fitting involved a coarse-to-fine approach. The coarse fit was carried out using an extensive multidimensional search space comprised of 35496 grid points, with different combinations of x , y and σ at each vertex. The parameters which provided the highest Pearson correlation between the predicted and observed time series were then used for the fine fit. The fine fit used the Nelder-Mead simplex-based method (Nelder & Mead, 1965) to reduce the residual sum of squares (RSS) between the predicted and observed time series, and determine optimal values for all four free parameters (x , y , σ and n).

3.2.9 pRF model analysis

As the compressive spatial summation (CSS) model was newly implemented within SamSrf 9, prior to analysing the pRF data I assessed its performance compared to a standard Gaussian model, which estimates pRFs as 2D Gaussians without any spatial summation. Firstly, I compared R^2 (the proportion of variance explained) values between the Gaussian and CSS model within the first three participants. Within these participants there was a consistent trend towards the CSS model explaining more of the variance (increased R^2) than the Gaussian model across all six ROIs (V1, V2, V3, OFA, pFus and mFus), with the face-selective regions benefitting most from the CSS model (for data from an example participant, see Figure A.1).

Next, I checked that the values of the exponent parameter in the CSS model (which estimate the amount of compression) were comparable to those reported by Kay et al. (2013). As expected, exponent values were highest in V1 (indicating the least compression) and decreased for higher-level brain regions, with the lowest values (the most compression) in mFus (Figure A.2). Upon investigating these values I found that some vertices had extremely high exponents, which were associated with poor fits (low R^2). To avoid erroneously high exponents affecting pRF estimates in a way that would not be representative of physiology and to improve fitting, the model was altered so that the exponent could be a maximum of 2.

I then compared estimates of pRF size across eccentricity between the standard Gaussian model and the new CSS implementation. While the Gaussian model estimates pRF size as σ (the standard deviation of the pRF, in degrees of visual angle),

the CSS model calculates it as $\sigma\sqrt{n}$ (where n is the exponent of the compressive non-linearity). The CSS model produced generally similar estimates of pRF size to the Gaussian model and to previous studies (including some which used a compressive exponent; Dumoulin & Wandell, 2008; Kay et al., 2015; Poltoratski et al., 2021), showing the expected increase in size with eccentricity within each ROI (Figure A.3). These checks confirmed that the CSS model was performing appropriately, so it was used for all pRF analyses. Mean beta and R^2 values for all ROIs are available in Figure A.4.

3.2.10 Delineation of early visual cortex

Prior to delineation, data were smoothed using a goodness-of-fit threshold of 0.1 and a smoothing kernel of 3 mm full width half maximum (FWHM). pRF locations (x and y) were used to project colour-coded polar angle and eccentricity maps onto the cortical surface. Visual areas V1-V3 were delineated by running an auto-delineation tool and then corrected manually using the SamSrf 9.4 toolbox (Schwarzkopf, 2022). This involved using standard criteria based on reversals in polar angle (DeYoe et al., 1994; Engel et al., 1997; Sereno et al., 1995), assisted by the eccentricity maps. Regions were delineated based on the maps generated for the upright face condition, before being checked and corrected (if needed) using the inverted maps.

3.2.11 Vertex selection

Vertices that had beta amplitudes of less than 0.01 or greater than 3 (z scores), sigma values of 0, or were located perfectly at the centre (x and y of exactly 0, which is indicative of fitting errors) were removed. To avoid noisy and unreliable vertices, those with a goodness-of-fit threshold (R^2) below 0.2 were also removed. In OFA and mFus, there were a few participants who had some vertices with very low pRF size estimates (almost 0) at high eccentricities. Upon closer analysis, these vertices had poor fits. To avoid these less reliable estimates affecting the main pattern of results, I adjusted the R^2 threshold within certain face-selective ROIs for some participants (OFA: four participants = 0.4, mFus: one participant = 0.4, one participant = 0.3).

3.2.12 Location analyses

To compare properties across visual field location, pRFs were filtered according to their centre position. Four wedges were defined, each including polar angle locations within $\pm 45^\circ$ on either side of the left horizontal, right horizontal, upper vertical and lower vertical meridians. Although behavioural research has suggested that visual field anisotropies decline at locations more than 30° away from the meridian (Abrams et al., 2012; Benson et al., 2021), fMRI studies that also used a wedge-based approach have shown that asymmetries in cortical surface were similar across different wedge widths (Himmelberg et al., 2023), and that anisotropies in pRF properties could be found using 45° wedges (Silva et al., 2018). Upon checking, I found that patterns of data were similar regardless of whether the wedge width was 30° or 45° . As a wider wedge width of 45° around the meridian considerably increased the number of pRFs remaining after filtering by ROI label, R^2 threshold *and* location (which otherwise reduced the amount of remaining pRFs more than was ideal), I chose to use 45° wedges for my location comparisons. For horizontal-vertical comparisons, the left and right horizontal wedges were combined to make the horizontal location, while the upper and lower vertical wedges were combined to make the vertical location.

3.2.13 Visual field coverage

There is no standard method of calculating coverage, with researchers adopting different approaches. Here, visual field coverage was determined by generating a Gaussian receptive field profile for each vertex based on its centre position (x, y), eccentricity and sigma (σ), and then raising the receptive field profile by the spatial summation exponent (n). Other approaches have calculated coverage using binary circles, which takes into account the position and size ($2\sigma\sqrt{n}$) of receptive fields, but not their Gaussian profiles (e.g. Poltoratski et al., 2021; Witthoft et al., 2016). Within the CSS model, pRFs with higher exponents (less compression) respond strongly to their centre position and relatively weakly towards the edges of the region they cover. On the other hand, pRFs with lower exponents (more compression) would have broader receptive field profiles, responding more equally across the spatial extent they cover. Two pRFs could therefore be positioned at the same location and have the same size but respond differently across the region of space they cover. By including

the exponent, I aimed to better account for *responsiveness* across the spatial extent of each receptive field profile when generating estimates of coverage.

Receptive field profiles were averaged across vertices to result in a mean coverage plot for each ROI, face orientation (upright/inverted), and participant. Because absolute values would differ based on various factors (e.g. number of vertices, sigma values, exponents applied), each mean coverage plot was divided by its maximum value to normalise the values to between 0 and 1. Using this specific method, coverage therefore represents the mean responsiveness of pRFs at each visual field location, relative to the population of pRFs within the same plot. Coverage values were extracted from these plots – according to eccentricity and polar angle location, using the wedges described in 3.2.12 – for further analyses.

3.2.14 Statistical analyses

Statistical analyses were carried out on estimates of three pRF measurements: size ($\sigma\sqrt{n}$), number (the total amount of pRFs after poorly fit vertices were removed; see 3.2.11) and visual field coverage (the values extracted from the coverage plots; see 3.2.13). Eccentricity bin widths of 1° were used for the statistical analyses of all three properties. The eccentricity range spanned from 0.5° (to exclude the region where the fixation point appeared) up to the maximum eccentricity of 21.65°, so there were 21 bins in total. Analyses were initially performed with ROI (V1, V2, V3, OFA, pFus, mFus) as a factor. As there were significant effects of ROI in all analyses (Table A.5-A.13), they were run within each ROI separately.

Linear mixed effects models were used to investigate whether location, inversion and eccentricity could predict pRF size. Because the location of each pRF was determined by its centre (described in 3.2.12), size could not be estimated if there were no pRF centres within that region. Linear mixed models were an appropriate choice as they could deal with these “missing” estimates and because pRF size demonstrates linear relationships (e.g. with eccentricity). Separate mixed effects models were run for each ROI (V1, V2, V3, OFA, pFus and mFus). The first set of analyses examined pRF size prior to location filtering, with fixed factors of eccentricity and inversion (upright/inverted), while the second set examined the effects of location,

with fixed factors of eccentricity, inversion (upright/inverted) and location (horizontal/vertical or upper/lower). Participant was specified as a random factor for the intercept as well as for each of the fixed factors, as the slope of the relationship between pRF size and eccentricity, location and/or inversion could vary across individuals. Participants could therefore have different intercepts and slopes for each of the factors. As one of the main hypotheses was to determine whether/how inversion affects pRF size across the visual field, following the linear mixed models t-tests were also performed to investigate whether differences were present at each eccentricity.

Mixed effects analyses of variance (ANOVAs) were used to assess the effects of eccentricity, inversion, location and participant on pRF number and visual field coverage. Separate ANOVAs were run for each ROI (V1, V2, V3, OFA, pFus and mFus). The first set of analyses looked at the effects before pRFs were filtered according to location (outlined in 3.2.12). The within-subjects fixed factors were therefore eccentricity and inversion (upright/inverted). In the second set of analyses I focused on the effects of location, with separate ANOVAs run for each location comparison. As such, the within-subjects fixed factors were eccentricity, inversion (upright/inverted) and location (horizontal/vertical or upper/lower). In all of the ANOVAs, participant was entered as a between-subjects random factor to account for individual variation within the model. As I was primarily interested in the group-level effects of location and inversion, and to improve the clarity of the results, the participant effects will not be reported. Following significant main effects or interactions, t-tests or Wilcoxon signed rank tests (if sphericity or homoscedasticity assumptions were violated) were carried out to explore the inversion and location differences in more depth.

3.3 Results

3.3.1 Behavioural results




Figure 3.2. Behavioural results from the gender task in the pRF experiment, showing the percentage of correctly identified male bars. Dots show individual data, with lines joining each participant's performance for upright and inverted faces.

Figure 3.2 displays the behavioural results from the gender task for nine participants (one individual's data did not save), where participants were required to press a button when they saw a bar containing male faces, which occurred with 0.075 probability (see 3.2.4). These results are displayed as the percentage of correctly identified male bars. Responses were counted as correct if there was a button press within two seconds of the onset of the bar. There was a

clear behavioural face inversion effect, with participants significantly better at recognising gender when the faces within the bars were upright rather than inverted, $t(8) = 12.19$, $p < .001$. This may reflect configural processing mechanisms being more efficiently engaged for upright compared to inverted faces (Rossion, 2008). Results from the fixation task – where participants detected the fixation cross changing colour – confirmed that participants were fixating and that there was no inversion effect for the non-face-based task (Figure A.5).

3.3.2 Neuroimaging results

Retinotopic maps of early visual cortex were successfully identified in all participants. Figure 3.3 displays polar angle, eccentricity, pRF size and R^2 maps for V1-V3 on the inflated cortical surface of one participant's occipital lobe. Figure 3.3 also shows the three face-selective regions delineated – OFA, pFus and mFus – on the ventral cortical surface within the same individual. Figure 3.4 displays representative predicted (by the pRF model) and observed time courses in OFA, pFus and mFus.





Figure 3.3. Retinotopic maps and ROIs within one participant. **A.** Polar angle map plotted on an inflated, spherical cortical surface (right hemisphere), with delineations of V1-V3 and OFA, pFus and mFus outlined. The colour wheel indicates polar angle coordinates (green for the lower visual field, blue around the horizontal meridian, red for the upper field). **B.** Eccentricity map. Purple represents central eccentricities and yellow the periphery. **C.** Sigma (σ) map, with purple representing small and blue representing large σ values. **D.** R^2 map, with black showing low and yellow showing high R^2 values. **E.** Three face-selective brain regions displayed on the ventral surface of the right hemisphere, from OFA in the occipital lobe (most posterior), to pFus and mFus (most anterior) along the fusiform gyrus.






Figure 3.4. Representative examples of the observed (blue) and predicted (red) time courses from selected vertices in OFA, pFus and mFus. x , y , sigma (σ), the spatial summation exponent (n), Beta (z) and R^2 values are displayed along the top of each time course.

3.3.2.1 Eccentricity and inversion

Prior to testing hypotheses about visual field anisotropies, I examined hallmarks of retinotopy for upright and inverted faces by looking at how pRF size, number and coverage varied across eccentricity. This was so that I could assess how the retinotopic properties of face-selective areas generally compared to V1-V3, across central and peripheral eccentricities (including at 10° , where better gender acuity was found for upright vs. inverted faces in the previous chapter). pRFs in the face-selective regions were expected to follow similar patterns across eccentricity as V1-V3, with larger sizes but reduced numbers and coverage as eccentricity increased.

In the face-selective areas only, pRFs were also predicted to vary according to inversion. Larger pRFs for upright vs. inverted faces would indicate distinct retinotopy,

while smaller sizes for upright faces would highlight similarities to early visual cortex. Increased pRF numbers and better visual field coverage for upright vs. inverted faces were expected, consistent with differences in behavioural performance for upright and inverted faces. These analyses allowed me to investigate how retinotopic properties are linked to *featural selectivity* within face perception.

3.3.2.1.1 pRF size

The first property to be examined was pRF size. In V1-V3 there were the expected increases in size with eccentricity, which were unaffected by inversion (Figure 3.5). Consistent with previous research, pRFs were larger and increased more sharply with eccentricity in successive visual areas (i.e. in V3 compared to V2, and V2 compared to V1; Dumoulin & Wandell, 2008). Linear mixed models confirmed these patterns, with significant effects of eccentricity but not of inversion in V1, V2 and V3 (Table 3.1).




Figure 3.5. Mean pRF size for upright (red) and inverted (blue) faces. At each eccentricity, size estimates were only plotted if they had been calculated from at least five vertices. The black lines indicate significant differences at certain eccentricities ($p < .05$). Note the different y axis scales between V1-V3 and face-selective regions.

pRFs were larger in the face-selective regions compared to early visual cortex (Figure 3.5), in line with previous research (Kay et al., 2015). As in V1-V3, significant main effects showed there were the expected increases in size with eccentricity in OFA, pFus and mFus (Table 3.1). Interestingly, these size increases were not as sharp in the FFA (pFus and mFus) compared to OFA. This highlights a difference in how the size-eccentricity relationship changed over the successive face-selective regions, compared to the progression between the earlier cortical regions, where size increases became more dramatic across eccentricity in successive areas (e.g. V3 vs. V1).

ROI	Fixed factors	β	p	CI
V1	Intercept	0.22		
	Inversion	-0.01	.910	-0.20, 0.18
	Eccentricity	0.08	< .001	0.06, 0.10
V2	Intercept	0.05		
	Inversion	0.08	.571	-0.19, 0.34
	Eccentricity	0.13	< .001	0.10, 0.16
V3	Intercept	0.53		
	Inversion	0.03	.856	-0.33, 0.40
	Eccentricity	0.20	< .001	0.14, 0.26
OFA	Intercept	1.71		
	Inversion	0.97	.003	0.32, 1.61
	Eccentricity	0.35	< .001	0.31, 0.39
pFus	Intercept	4.95		
	Inversion	0.32	.416	-0.46, 1.11
	Eccentricity	0.21	.004	0.06, 0.35
mFus	Intercept	7.36		
	Inversion	-0.58	.538	-2.45, 1.28
	Eccentricity	0.16	.002	0.06, 0.25

Table 3.1. Linear mixed model results comparing pRF size across inversion and eccentricity. Bold text indicates statistical significance ($p < .05$).

There was a main effect of inversion in OFA, with t-tests revealing significantly larger pRFs for inverted vs. upright faces – opposite to predictions based on Poltoratski et al.'s (2021) findings – at a few central and peripheral eccentricities (Table 3.1; Figure 3.5). Although pRFs looked to be larger for upright vs. inverted faces in mFus – which would be consistent with previous research (Poltoratski et al., 2021) – there were no

significant main effects of inversion in pFus or mFus. t-tests did not uncover consistent differences in the periphery (with no differences around 10°, where better gender acuity for upright vs. inverted faces was found in Chapter 2). As such, pRFs in the FFA did not reliably differ in size between upright and inverted faces.

3.3.2.1.2 pRF number

Next, pRF number (defined as the amount of vertices remaining at each eccentricity after filtering by R^2 , beta and sigma thresholds; see 3.2.11) was analysed. In V1-V3, main effects of eccentricity showed that pRF numbers were significantly higher towards the fovea⁵ and gradually decreased with eccentricity, apart from a small spike in the far periphery (Figure 3.6; Table 3.2). This gradual decrease meant that in early visual areas there was still a considerable amount of pRFs at peripheral eccentricities, highlighting their spatial spread across the visual field. As expected, there were no main effects of inversion in V1, V2, or V3, and t tests revealed only a few significant differences which varied inconsistently in terms of whether there were more pRFs for upright or inverted faces. Together these results show that there were similar amounts of pRFs for upright and inverted faces in early visual cortex.




Figure 3.6. Mean pRF number for upright (red) and inverted (blue) faces. The black lines indicate significant differences at certain eccentricities ($p < .05$). Note the different y axis scales between ROIs.

⁵ The fovea will refer to 0-2° eccentricity, the parafovea 2-5°, and the perifovea 5-10°.

Like early visual cortex, main effects of eccentricity indicated that all three face-selective regions had significantly more pRFs towards the fovea (Figure 3.6; Table 3.2). Consistent with previous research (e.g. Finzi et al., 2021), pRFs were more densely concentrated towards the fovea and decreased more dramatically in the periphery compared to V1-V3, indicating a magnified foveal bias in the face-selective areas.

ROI	Factors	<i>df</i>	<i>F</i>	<i>p</i>
V1	Inversion	1, 180	0.32	.583
	Eccentricity	20, 180	13.67	< .001
	Inversion*Eccentricity	20, 180	0.72	.802
V2	Inversion	1, 180	2.17	.175
	Eccentricity	20, 180	10.37	< .001
	Inversion*Eccentricity	20, 180	2.40	.001
V3	Inversion	1, 180	1.45	.259
	Eccentricity	20, 180	7.81	< .001
	Inversion*Eccentricity	20, 180	0.67	.852
OFA	Inversion	1, 180	0.24	.634
	Eccentricity	20, 180	15.10	< .001
	Inversion*Eccentricity	20, 180	2.69	< .001
pFus	Inversion	1, 160	2.48	.154
	Eccentricity	20, 160	5.66	< .001
	Inversion*Eccentricity	20, 160	1.97	.011
mFus	Inversion	1, 180	0.53	.484
	Eccentricity	20, 180	8.28	< .001
	Inversion*Eccentricity	20, 180	1.72	.034

Table 3.2. ANOVA results comparing pRF number across inversion and eccentricity. Bold text indicates statistical significance ($p < .05$).

Contrary to expectations, there were no main effects of inversion in any of the face-selective regions. However, there were significant interactions between inversion and eccentricity, and Wilcoxon tests revealed that in mFus, there were significantly more pRFs for upright vs. inverted faces at perifoveal and peripheral eccentricities. This includes a significant difference at 10°, where better gender acuity was found for upright vs. inverted faces in the previous chapter. These results indicate that the distribution of pRFs across the visual field varied according to inversion, with upright faces leading to more pRFs found in the periphery.

3.3.2.1.3 Visual field coverage

After pRF size and number, visual field coverage was assessed. Coverage plots for each ROI are displayed in Figure 3.7, representing the relative responsiveness of pRFs across the visual field, within the population of neurons that responded to upright or inverted faces (see 3.2.13 for more details). Coverage values were extracted from these maps and plotted across eccentricity in Figure 3.8. Upright and inverted coverage maps were normalised separately (divided by their own maximum value). To check whether this normalisation approach affected results, maps were also normalised jointly (both maps for each ROI divided by the maximum value in the upright map). This yielded generally similar estimates of coverage as those shown in Figure 3.8, demonstrating that the maximum values – a measure of absolute coverage – were not considerably different between upright and inverted maps (Figure A.6).

In V1-V3, main effects of eccentricity showed that coverage was significantly higher in the fovea and decreased in the periphery (Figure 3.8; Table 3.3). This bias towards greater responsiveness at the fovea increased in successive visual areas (from V1 to V2 to V3), in line with greater cortical magnification further up in the visual hierarchy. As expected in early visual cortex, coverage was not significantly affected by inversion.

Main effects showed that the face-selective regions also had significantly better coverage in the fovea than the periphery (Figure 3.8; Table 3.3). While these central biases were considerably exaggerated compared to early visual cortex, they also increased in magnitude in successive regions, from OFA to pFus to mFus. This simultaneously highlights differences and commonalities between low- and high-level retinotopy. Unexpectedly, there were no main effects of inversion or interactions between inversion and eccentricity in any of the face-selective areas. In mFus, however, Figure 3.7Figure 3.8 show a trend towards better coverage across the visual field for upright vs. inverted faces, with t-tests showing that this difference was significant in the fovea (although not at 10°, where better gender acuity was found for upright vs. inverted faces in Chapter 2). This highlights subtle differences in coverage for upright and inverted faces, suggesting that mFus had increased responsiveness in

the fovea for upright faces, with smaller-yet-consistent increases across the rest of the visual field.




Figure 3.7. Mean visual field coverage for upright (top) and inverted (bottom) faces (see 3.2.13 for how plots were generated). Coordinates represent eccentricity in degrees of visual angle, with negative values for the left and positive values for the right visual field. Values were converted to log scale before plotting, for visualisation purposes (see colour bar). Dots represent pRF centres from all participants. For plots without pRF centres imposed, see Figure A.7).




Figure 3.8. Mean visual field coverage values for upright (red) and inverted (blue) faces. Black lines indicate significant differences at certain eccentricities ($p < .05$).

ROI	Factors	<i>df</i>	<i>F</i>	<i>p</i>
V1	Inversion	1, 180	0.02	.882
	Eccentricity	20, 180	27.67	< .001
	Inversion*Eccentricity	20, 180	11.18	.965
V2	Inversion	1, 180	0.59	.462
	Eccentricity	20, 180	35.22	< .001
	Inversion*Eccentricity	20, 180	1.77	.027
V3	Inversion	1, 180	0.03	.862
	Eccentricity	20, 180	76.84	< .001
	Inversion*Eccentricity	20, 180	0.87	.625
OFA	Inversion	1, 180	0.02	.884
	Eccentricity	20, 180	263.79	< .001
	Inversion*Eccentricity	20, 180	0.97	.496
pFus	Inversion	1, 160	0.13	.731
	Eccentricity	20, 160	586.69	< .001
	Inversion*Eccentricity	20, 160	0.31	.998
mFus	Inversion	1, 180	2.18	.174
	Eccentricity	20, 180	329.15	< .001
	Inversion*Eccentricity	20, 180	1.00	.460

Table 3.3. ANOVA results comparing visual field coverage across inversion and eccentricity. Bold text indicates statistical significance ($p < .05$).

3.3.2.2 Summary of eccentricity and inversion effects

In all ROIs, pRF size increased while number and coverage decreased across eccentricity, as expected. This highlights shared patterns of retinotopy between low- and high-level visual cortex. While inversion was expected to modulate retinotopic measurements within face-selective areas, this was found not to be the case for pRF size. However, there were significantly more pRFs in the periphery for upright vs. inverted faces, and better visual field coverage for upright faces at the fovea. Although these effects of inversion were smaller than expected (as will be discussed), they could point to slightly better sampling of the visual field for upright vs. inverted faces, which may contribute to the better recognition of upright faces.

3.3.2.3 Horizontal-vertical difference

Having examined how retinotopic properties differed according to eccentricity and inversion (featural selectivity), I then considered their variations according to location. This allowed me to examine whether *spatial selectivity* within face-selective regions could explain the behavioural variations in the previous chapter, and whether patterns of spatial selectivity were similar to, or diverged from, earlier visual areas.

First, retinotopy was compared between the horizontal and vertical meridians (location filtering is described in section 3.2.12). In the face-selective regions, increases in pRF number and visual field coverage were expected along the horizontal vs. vertical meridian, which would explain the horizontal-vertical difference in gender acuity found in Chapter 2. As similar patterns were expected in V1-V3 (Amano et al., 2009; Arcaro et al., 2009), this would also highlight commonalities between low- and high-level sampling. Whether pRF size varied between the horizontal and vertical meridians in the face-selective regions would provide further insight into their links with early visual cortex. Along the horizontal meridian (where there was better gender acuity), smaller pRFs would demonstrate similarities with low-level sampling (Silva et al., 2018), while larger pRFs would point to more unique retinotopy within face-selective areas.

The main effects of eccentricity were significant in all tests comparing pRF size, number and coverage between the horizontal and vertical meridian (Table 3.4, Table

3.5 and Table 3.6), showing that the expected retinotopic characteristics across eccentricity (larger pRFs but decreased numbers and visual field coverage in the periphery) were also present in these analyses. Inversion was included as a factor because I was interested in whether the horizontal-vertical differences were similar for upright and inverted faces. As the effects of inversion have already been explored directly – where effects would be more robust, as the location filtering in the following sections meant that there were fewer data within the inversion comparisons – they will mainly be focused on in terms of their interaction with location.

ROI	Fixed factors	β	p	CI
V1	Intercept	-0.00		
	Location	0.13	.373	-0.16, 0.43
	Inversion	0.01	.949	0.19, 0.20
	Eccentricity	0.08	< .001	0.07, 0.10
V2	Intercept	0.10		
	Location	-0.01	.953	-0.38, 0.36
	Inversion	0.05	.704	-0.23, 0.34
	Eccentricity	0.13	< .001	0.12, 0.16
V3	Intercept	-0.99		
	Location	0.79	.012	0.17, 1.41
	Inversion	0.19	.422	-0.27, 0.65
	Eccentricity	0.23	< .001	0.17, 0.29
OFA	Intercept	1.69		
	Location	-0.48	.551	-2.05, 1.10
	Inversion	1.04	.003	0.35, 1.72
	Eccentricity	0.40	< .001	0.35, 0.45
pFus	Intercept	4.06		
	Location	0.24	.734	-1.15, 1.63
	Inversion	0.67	.206	-0.37, 1.71
	Eccentricity	0.21	.007	0.06, 0.36
mFus	Intercept	6.02		
	Location	1.22	.070	-0.10, 2.54
	Inversion	-0.56	.578	-2.53, 1.41
	Eccentricity	0.17	< .001	0.07, 0.27

Table 3.4. Linear mixed model results for pRF size across horizontal and vertical locations. Bold text indicates statistical significance ($p < .05$).

3.3.2.3.1 pRF size

Again, pRF size was the first property to be assessed. Figure 3.9 displays the average pRF size across eccentricity for the horizontal and vertical meridians (if the size estimate at each eccentricity was calculated from at least five vertices – this




Figure 3.9. Mean pRF size across the horizontal and vertical meridians, for upright (top) and inverted (bottom) faces. At each eccentricity, size estimates were only plotted if they were averaged from at least five vertices. The black lines indicate significant differences at certain eccentricities ($p < .05$). Note the different y axis scales between the ROIs.

number was chosen based on the observation that there were very few vertices in some locations). In V1 and V2, there were no main effects of location, showing that pRF size did not significantly differ between the horizontal and vertical meridians (Table 3.4). There was a significant effect of location in V3, with t-tests showing that pRFs were larger along the vertical meridian in the periphery and periphery, especially for inverted faces.

Because the face-selective regions had so few pRFs in the periphery along the vertical meridian, size estimates were often unavailable (Figure 3.9). While there was a trend towards more dramatic increases in pRF size along the vertical than horizontal meridian in all three face-selective areas, there were no main effects of location (Table 3.4). The t-tests yielded inconsistent differences (in terms of whether pRFs were larger along the horizontal or vertical meridian) at only a couple of eccentricities in OFA and pFus, and none in mFus. This shows that contrary to expectations, pRF size was not reliably modulated by location in the face-selective regions.

3.3.2.3.2 pRF number

Next, the number of pRFs along the two meridians was assessed. In V1-V3, differences in pRF number between the horizontal and vertical meridians emerged outside of the fovea, with main effects of location and interactions between location and eccentricity in all three regions (Figure 3.10; Table 3.5). t-tests showed that apart from a greater amount of pRFs along the vertical meridian at some eccentricities in V2, all three early visual areas had significantly more pRFs along the peripheral horizontal meridian. These differences covered a greater proportion of the visual field in V3 compared to V1 and V2, suggesting that cortical magnification of the horizontal meridian may increase further up in the visual hierarchy. These patterns were generally unaffected by inversion. As such, the horizontal-vertical difference in pRF number increased with eccentricity in early visual cortex, with similar patterns for upright and inverted faces.

The horizontal-vertical differences found in early visual cortex were magnified in face-selective regions. Main effects showed that in all three face-selective areas there were significantly more pRFs along the horizontal vs. vertical meridian, with t-

tests indicating that these differences were significant across the visual field (Figure 3.10; Table 3.5). There were also interactions between location and eccentricity in all three regions, reflecting that the horizontal-vertical differences in pRF number were greater in the fovea compared to the periphery. This is the opposite pattern to V1-V3 (where the differences increased in the periphery), which is presumably linked to there being considerably more pRFs in the fovea in face-selective regions overall. In mFus




Figure 3.10. Mean pRF number across the horizontal and vertical meridians, for upright (top) and inverted (bottom) faces. The black lines indicate significant differences at certain eccentricities ($p < .05$). Note the different y axis scales between ROIs.

ROI	Factors	df	F	p
V1	Location	1, 180	26.83	< .001
	Inversion	1, 180	0.32	.583
	Eccentricity	20, 180	13.67	< .001
	Location*Inversion	1, 180	3.86	.050
	Location*Eccentricity	20, 180	14.64	< .001
	Inversion*Eccentricity	20, 180	0.53	.957
V2	Location	1, 180	0.01	.911
	Inversion	1, 180	2.17	.175
	Eccentricity	20, 180	10.37	< .001
	Location*Inversion	1, 180	0.78	.377
	Location*Eccentricity	20, 180	16.54	< .001
	Inversion*Eccentricity	20, 180	1.84	.015
V3	Location	1, 180	224.16	< .001
	Inversion	1, 180	1.45	.259
	Eccentricity	20, 180	7.81	< .001
	Location*Inversion	1, 180	4.27	.039
	Location*Eccentricity	20, 180	7.94	< .001
	Inversion*Eccentricity	20, 180	0.53	.953
OFA	Location	1, 180	16.49	.003
	Inversion	1, 180	0.19	.700
	Eccentricity	20, 180	14.47	< .001
	Location*Inversion	1, 180	0.23	.630
	Location*Eccentricity	20, 180	15.86	< .001
	Inversion*Eccentricity	20, 180	1.02	.441
pFus	Location	1, 160	15.92	.004
	Inversion	1, 160	2.48	.154
	Eccentricity	20, 160	5.66	< .001
	Location*Inversion	1, 160	2.47	.117
	Location*Eccentricity	20, 160	7.81	< .001
	Inversion*Eccentricity	20, 160	0.49	.969
mFus	Location	1, 180	7.27	.025
	Inversion	1, 180	0.53	.484
	Eccentricity	20, 180	8.28	< .001
	Location*Inversion	1, 180	8.03	.005
	Location*Eccentricity	20, 180	6.53	< .001
	Inversion*Eccentricity	20, 180	0.99	.468

Table 3.5. ANOVA results for pRF number across horizontal and vertical locations. Bold text indicates statistical significance ($p < .05$).

there was also an interaction between location and inversion, reflecting that there was a clear horizontal-vertical anisotropy found across the visual field for upright faces, but fewer significant differences for inverted faces. This is likely related to there being fewer pRFs in the periphery for inverted faces overall (shown in 3.3.2.1.2), which would reduce the magnitude of the location differences. All in all, the expected horizontal-vertical differences in pRF number were found across the visual field within all three face-selective regions, with generally similar variation regardless of inversion.

3.3.2.3.3 Visual field coverage

Visual field coverage was then compared across horizontal and vertical locations. Main effects of location showed that in V1 and V3, there was increased coverage along the horizontal compared to the vertical meridian (Figure 3.11; Table 3.6). t-tests showed that these differences were significant in the parafovea and periphery in V1, and most of the visual field in V3. Although there was no main effect of location in V2, there was an interaction between location and eccentricity, with t-tests indicating that while there was significantly better coverage along the vertical meridian near the fovea, the expected pattern was present at more peripheral eccentricities, with better coverage along the horizontal meridian (as in V1 and V3). There were also significant interactions between location and eccentricity in V1 and V3, reflecting that across early visual cortex, the magnitude of the horizontal-vertical difference was larger at central and mid-eccentricities as opposed to the periphery, where coverage was significantly reduced overall. Although there was an interaction between location and inversion in V3, t-tests showed that the horizontal-vertical differences were similar for upright and inverted faces across the visual field. Altogether, coverage was generally better along the horizontal than vertical meridian in early visual cortex, with similar patterns irrespective of inversion.

The face-selective brain regions all showed horizontal-vertical anisotropies in visual field coverage, although these differences were smaller in magnitude than in early visual cortex. Main effects of location indicated that in all three face-selective areas, coverage was better along the horizontal vs. vertical meridian (Figure 3.11; Table 3.6). In OFA and pFus there were significant interactions between location and eccentricity, with the t-tests revealing that the horizontal-vertical difference was largest




Figure 3.11. Mean visual field coverage values across the horizontal and vertical meridians, for upright (top) and inverted (bottom) faces. Black lines indicate significant differences at certain eccentricities ($p < .05$).

at mid-eccentricities and reduced in the periphery, with similar patterns regardless of inversion. Interestingly, the horizontal-vertical difference was smallest in mFus (of the three face-selective regions). There was a significant interaction between location and inversion, with t-tests indicating that there was better coverage along the horizontal vs. vertical meridian across the visual field for upright but not inverted faces. In other

ROI	Factors	<i>df</i>	<i>F</i>	<i>p</i>
V1	Location	1, 180	8.36	.018
	Inversion	1, 180	0.02	.882
	Eccentricity	20, 180	27.48	< .001
	Location*Inversion	1, 180	1.30	.255
	Location*Eccentricity	20, 180	30.78	< .001
	Inversion*Eccentricity	20, 180	0.42	.988
V2	Location	1, 180	9.32	.014
	Inversion	1, 180	0.58	.466
	Eccentricity	20, 180	35.10	< .001
	Location*Inversion	1, 180	5.44	.020
	Location*Eccentricity	20, 180	31.70	< .001
	Inversion*Eccentricity	20, 180	2.26	.002
V3	Location	1, 180	46.57	< .001
	Inversion	1, 180	0.03	.856
	Eccentricity	20, 180	76.68	< .001
	Location*Inversion	1, 180	0.07	.788
	Location*Eccentricity	20, 180	59.97	< .001
	Inversion*Eccentricity	20, 180	0.84	.662
OFA	Location	1, 180	11.84	.007
	Inversion	1, 180	0.02	.883
	Eccentricity	20, 180	264.09	< .001
	Location*Inversion	1, 180	0.85	.357
	Location*Eccentricity	20, 180	3.17	< .001
	Inversion*Eccentricity	20, 180	2.21	.002
pFus	Location	1, 160	7.20	.028
	Inversion	1, 160	0.13	.731
	Eccentricity	20, 160	585.87	< .001
	Location*Inversion	1, 160	1.59	.209
	Location*Eccentricity	20, 160	8.67	< .001
	Inversion*Eccentricity	20, 160	0.38	.994
mFus	Location	1, 180	14.94	.004
	Inversion	1, 180	2.18	.174
	Eccentricity	20, 180	329.22	< .001
	Location*Inversion	1, 180	6.28	.013
	Location*Eccentricity	20, 180	0.27	.999
	Inversion*Eccentricity	20, 180	3.07	< .001

Table 3.6. ANOVA results for visual field coverage across horizontal and vertical locations. Bold text indicates statistical significance ($p < .05$).

words, coverage showed horizontal-vertical anisotropies in all three face-selective regions, although only for upright faces in mFus.

3.3.2.4 Summary of horizontal-vertical differences

While V3 had larger pRFs along the vertical vs. horizontal meridian, pRF size was not significantly affected by location in any of the other brain regions. However, V1-V3 and all three face-selective areas had significantly more pRFs and better visual field coverage along the horizontal compared to vertical meridian. These variations in pRF number and coverage are consistent with better gender acuity along the horizontal vs. vertical meridian (Chapter 2), and highlight commonalities in how face-selective regions and early visual cortex sample the visual field, irrespective of face inversion.

3.3.2.5 Upper-lower difference

After comparing the horizontal and vertical meridians, I assessed whether retinotopic properties differed between the upper and lower visual field (location filtering is described in 3.2.12). Similar to the horizontal-vertical results discussed above, in the face-selective regions increases in pRF number and visual field coverage were expected in the lower vs. upper visual field. This would tie in with the behavioural upper-lower difference measured in the previous chapter. Similar patterns were expected in V1-V3, which would highlight commonalities between low- and high-level sampling. Also as before, differences in pRF size between the upper and lower fields

ROI	Fixed factors	β	p	CI
V1	Intercept	-0.68		
	Location	0.66	.006	0.20, 1.13
	Inversion	-0.00	.985	-0.19, 0.18
	Eccentricity	0.07	< .001	0.05, 0.08
V2	Intercept	-1.11		
	Location	0.88	.017	0.16, 1.61
	Inversion	-0.02	.864	-0.31, 0.26
	Eccentricity	0.12	< .001	0.08, 0.16
V3	Intercept	-1.58		
	Location	0.84	.035	0.06, 1.62
	Inversion	0.36	.178	-0.17, 0.89
	Eccentricity	0.29	< .001	0.22, 0.36
OFA	Intercept	-4.56		
	Location	2.30	.001	0.94, 3.67
	Inversion	1.67	.023	0.23, 3.11
	Eccentricity	0.33	.042	0.01, 0.66
pFus	Intercept	-1.47		
	Location	-0.72	.689	-4.27, 2.83
	Inversion	3.92	< .001	2.33, 5.50
	Eccentricity	0.63	.029	0.07, 1.19
mFus	Intercept	8.52		
	Location	-1.72	.096	-3.76, 0.31
	Inversion	-0.02	.985	-2.50, 2.45
	Eccentricity	0.66	< .001	0.29, 1.03

Table 3.7. Linear mixed model results for pRF size across upper and lower locations. Bold text indicates statistical significance ($p < .05$).

in the face-selective regions would shed further light into the links with early visual cortex. In the lower field (where there was better gender acuity), smaller pRFs would indicate similarities with low-level sampling, while larger pRFs would highlight distinct sampling characteristics within face-selective areas.

The main effects of eccentricity were significant in all tests comparing pRF size, number and coverage between the upper and lower locations (Table 3.7, Table 3.8Table 3.9), again showing that the expected patterns of increased pRF size with eccentricity, and decreased pRF number and visual field coverage with eccentricity, were present in these analyses. As with the horizontal-vertical anisotropies, inversion was included as a factor so that I could assess whether the upper-lower differences varied between upright and inverted faces. As the effects of inversion have been explored in more depth previously, they will be focused on in terms of their interaction with location. It should be noted that the main effects of inversion may be even less reliable in the following analyses due to data reductions caused by the location filtering (this was particularly apparent along the vertical meridian, which comprises the upper and lower locations; Figure 3.10).

3.3.2.5.1 pRF size

Firstly, pRF size was examined. Figure 3.12 displays the average pRF size over eccentricity for the upper and lower locations (if the size estimate at each eccentricity was calculated from at least five vertices). Contrary to expectations, main effects of location showed that pRFs were larger in the lower vs. upper field across early visual cortex (Table 3.7), with t-tests uncovering significant differences across varied eccentricities. These patterns were similar for upright and inverted faces, with no effects of inversion in V1-V3.

Along the vertical meridian the face-selective regions had very few pRFs in the periphery, meaning that size estimates were unavailable at far eccentricities (Figure 3.12). There was an effect of location in OFA, showing that pRFs were larger in the lower vs. upper field. An effect of inversion along with the t-tests indicated that this upper-lower difference was significant in the perifovea, but only for upright faces (Table

3.7). There were no effects of location in pFus or mFus, highlighting that pRF size did not differ between the upper and lower locations within the FFA.




Figure 3.12. Mean pRF size in the upper and lower visual field, for upright (top) and inverted (bottom) faces. At each eccentricity, size estimates were only plotted if they were averaged from at least five vertices. Note the different y axis scales between the ROIs.

3.3.2.5.2 pRF number

Next, the number of pRFs was examined between the upper and lower locations (Figure 3.13). Main effects of location in V1-V3 confirmed that there were a greater

number of pRFs found in the lower vs. upper field, as expected (Table 3.8). Interactions between location and eccentricity, along with the Wilcoxon tests, showed that the upper-lower difference varied across the visual field but was generally greater at mid-eccentricities. These patterns were generally similar irrespective of inversion, with Wilcoxon tests showing that while there was an




Figure 3.13. Mean pRF number in the upper and lower visual field, for upright (top) and inverted (bottom) faces. Black lines indicate significant differences at certain eccentricities ($p < .05$). Note the different y axis scales between the ROIs.

ROI	Factors	<i>df</i>	<i>F</i>	<i>p</i>
V1	Location	1, 180	36.19	< .001
	Inversion	1, 180	0.60	.458
	Eccentricity	20, 180	16.65	< .001
	Location*Inversion	1, 180	2.57	.110
	Location*Eccentricity	20, 180	3.86	< .001
	Inversion*Eccentricity	20, 180	0.74	.782
V2	Location	1, 180	34.59	< .001
	Inversion	1, 180	0.11	.746
	Eccentricity	20, 180	17.99	< .001
	Location*Inversion	1, 180	4.58	.033
	Location*Eccentricity	20, 180	7.45	< .001
	Inversion*Eccentricity	20, 180	1.09	.356
V3	Location	1, 180	10.37	.011
	Inversion	1, 180	0.22	.654
	Eccentricity	20, 180	9.70	< .001
	Location*Inversion	1, 180	2.02	.156
	Location*Eccentricity	20, 180	2.68	< .001
	Inversion*Eccentricity	20, 180	0.52	.959
OFA	Location	1, 180	7.87	.021
	Inversion	1, 180	0.86	.379
	Eccentricity	20, 180	6.08	< .001
	Location*Inversion	1, 180	1.85	.174
	Location*Eccentricity	20, 180	5.05	< .001
	Inversion*Eccentricity	20, 180	1.03	.427
pFus	Location	1, 160	2.93	.125
	Inversion	1, 160	4.65	.063
	Eccentricity	20, 160	2.81	< .001
	Location*Inversion	1, 160	6.92	.009
	Location*Eccentricity	20, 160	2.92	< .001
	Inversion*Eccentricity	20, 160	1.53	.019
mFus	Location	1, 180	19.06	.002
	Inversion	1, 180	11.41	.008
	Eccentricity	20, 180	11.24	< .001
	Location*Inversion	1, 180	12.55	< .001
	Location*Eccentricity	20, 180	11.29	< .001
	Inversion*Eccentricity	20, 180	3.86	< .001

Table 3.8. ANOVA results for pRF number across upper and lower locations. Bold text indicates statistical significance ($p < .05$).

interaction between location and inversion in V2, there were similar significant differences found across the visual field.

Like with the horizontal-vertical anisotropy in pRF number, the upper-lower differences were even more pronounced in face-selective regions compared to early visual cortex. All three face-selective areas had strikingly few pRFs in the upper field. Unsurprisingly, in OFA and mFus there were main effects of location, with more pRFs in the lower than upper field, consistent with predictions (Figure 3.13; Table 3.8). There were also interactions between location and eccentricity, with Wilcoxon tests showing that the upper-lower differences were more pronounced near the fovea, where there were larger numbers of pRFs overall. In mFus there was an additional interaction between location and inversion, highlighting that these differences were only significant for the inverted faces. Although there was no main effect of location or significant differences found by the Wilcoxon tests in pFus, there was an interaction between location and eccentricity, suggesting that there were small upper-lower differences within the eccentricities tested.

3.3.2.5.3 Visual field coverage

Finally, visual field coverage was examined between the upper and lower locations. Main effects of location confirmed that across early visual cortex, coverage was higher in the lower vs. upper field, in line with expectations (Figure 3.14; Table 3.9). There were also significant interactions between location and eccentricity in all three regions, with t-tests showing that the upper-lower differences were more pronounced in central vision and reduced in the periphery. There were no interactions between location and inversion in V1-V3, highlighting that the differences in coverage between the upper and lower locations were similar for upright and inverted faces.

Like V1-V3, main effects of location showed that the face-selective regions had better visual field coverage in the lower vs. upper field, although the differences were smaller in magnitude (Figure 3.14; Table 3.9). t-tests revealed that the upper-lower anisotropy was present at central and mid-eccentricities across all face-selective




Figure 3.14. Mean visual field coverage values in the upper and lower visual field, for upright (top) and inverted (bottom) faces. Black lines indicate significant differences at certain eccentricities ($p < .05$).

regions, and extended furthest into the periphery in mFus. There were interactions between location and eccentricity in all three face-selective regions, again indicating that the upper-lower difference was more pronounced in the fovea and diminished in the periphery. Coverage patterns were generally similar for upright and inverted faces, although interactions between location and inversion were present in pFus and mFus, with t-tests revealing small variations according to inversion in the FFA.

ROI	Factors	<i>df</i>	<i>F</i>	<i>p</i>
V1	Location	1, 180	25.67	< .001
	Inversion	1, 180	0.07	.801
	Eccentricity	20, 180	28.89	< .001
	Location*Inversion	1, 180	2.35	.126
	Location*Eccentricity	20, 180	30.79	< .001
	Inversion*Eccentricity	20, 180	0.41	.990
V2	Location	1, 180	31.87	< .001
	Inversion	1, 180	0.17	.693
	Eccentricity	20, 180	39.03	< .001
	Location*Inversion	1, 180	0.03	.855
	Location*Eccentricity	20, 180	51.05	< .001
	Inversion*Eccentricity	20, 180	1.28	.184
V3	Location	1, 180	19.72	.002
	Inversion	1, 180	0.03	.868
	Eccentricity	20, 180	71.77	< .001
	Location*Inversion	1, 180	2.79	.095
	Location*Eccentricity	20, 180	32.37	< .001
	Inversion*Eccentricity	20, 180	0.28	.999
OFA	Location	1, 180	16.09	.003
	Inversion	1, 180	0.11	.747
	Eccentricity	20, 180	334.16	< .001
	Location*Inversion	1, 180	0.21	.643
	Location*Eccentricity	20, 180	41.28	< .001
	Inversion*Eccentricity	20, 180	1.07	.379
pFus	Location	1, 160	10.32	.012
	Inversion	1, 160	0.02	.880
	Eccentricity	20, 160	781.81	< .001
	Location*Inversion	1, 160	8.32	.004
	Location*Eccentricity	20, 160	13.21	< .001
	Inversion*Eccentricity	20, 160	0.20	1.000
mFus	Location	1, 180	24.80	< .001
	Inversion	1, 180	1.64	.232
	Eccentricity	20, 180	319.01	< .001
	Location*Inversion	1, 180	13.54	< .001
	Location*Eccentricity	20, 180	18.49	< .001
	Inversion*Eccentricity	20, 180	1.78	.020

Table 3.9. ANOVA results for visual field coverage across upper and lower locations. Bold text indicates statistical significance ($p < .05$).

3.3.2.6 Summary of upper-lower differences

In V1-V3, there were larger pRFs in the lower vs. upper visual field (opposite to expectations). Variations in pRF size were less apparent within the face-selective regions, with some evidence of larger pRFs in the lower vs. upper field in OFA, but no differences in pFus or mFus. Consistent with expectations, there were significantly more pRFs and better visual field coverage in the lower vs. upper field within all of the brain regions. As with the horizontal-vertical results described above, these variations in pRF number and coverage within the face-selective areas align with better gender acuity found in the lower vs. upper field (Chapter 2). They also reveal similarities in visual field sampling between face-selective regions and early visual cortex, regardless of face inversion.

3.4 Discussion

Do variations in face perception stem from the way that face-selective cortex samples visual space? In the previous chapter, face perception was found to vary systematically across the visual field, with better acuity for judging face gender along the horizontal vs. vertical meridian and in the lower vs. upper field. These *horizontal-vertical* and *upper-lower differences* mirror those found behaviourally for low-level vision (e.g. Carrasco et al., 2001), with the low-level anisotropies linked to the neuronal properties of early visual cortex (e.g. Silva et al., 2018). This chapter investigated the neural basis of the variations in face perception, by measuring population receptive field (pRF) properties in parts of the brain that preferentially process faces (OFA, pFus and mFus). There was a greater *number* of pRFs and better *visual field coverage* along the horizontal vs. vertical meridian and in the lower vs. upper field in both the early visual regions and face-selective areas. While there was evidence of larger pRFs along the vertical vs. horizontal meridian and (opposite to expectations) in the lower vs. upper field in early visual cortex, pRF size generally did not differ by location in the face-selective regions. Therefore, some – but not all – retinotopic properties of face-selective areas can provide a neural basis for the visual field anisotropies measured behaviourally for face perception (Chapter 2), with patterns of retinotopy resembling those in early visual cortex.

In early visual cortex, results were largely in line with previous research. A greater number of pRFs was found along the horizontal than vertical meridian and in the lower vs. upper field, replicating previous pRF findings (Amano et al., 2009) and consistent with greater cortical magnification and surface area in these regions (Himmelberg et al., 2023; Silva et al., 2018). Visual field coverage (an estimate of the *relative responsiveness* in each region of the visual field; see 3.3.2.3.3) was also better along the horizontal vs. vertical meridian and in the lower vs. upper field, similar to prior work (Amano et al., 2009). As for pRF size, V3 – but not V1 and V2 – showed the expected horizontal-vertical difference of larger pRFs along the vertical than horizontal meridian. This differs slightly from previous research, which found a horizontal-vertical difference across V1-V3 (Silva et al., 2018). Interestingly, pRFs were larger in the lower than upper field across early visual cortex, showing the opposite pattern to other studies (Silson et al., 2018; Silva et al., 2018). This was unexpected as smaller receptive fields are usually associated with better acuity in V1 (Duncan & Boynton, 2003). Therefore, only pRF number and visual field coverage showed the same horizontal-vertical and upper-lower anisotropies as the behavioural variations typically measured in low-level vision (e.g. Carrasco et al., 2001). Differences in pRF size were more variable and much less pronounced compared to pRF number, explaining why coverage patterns aligned more closely with the variations in pRF number rather than size. These findings suggest that even in V1, pRF size is not the only factor driving acuity variations, with receptive field number and resulting visual field coverage equally – if not more – important for acuity.

Some researchers have proposed that larger receptive fields result in better face perception (Poltoratski et al., 2021; Witthoft et al., 2016). However, while there were larger pRFs in the lower vs. upper field in OFA (the same pattern as V1-V3), pRF size did not significantly differ between locations in the FFA (pFus and mFus). Unlike early visual cortex, therefore, the behavioural variations that were previously found for face perception (Chapter 2) cannot be explained by differences in pRF size. This may reflect the fact that pRFs within face-selective regions were considerably larger than in V1-V3 (as has been consistently found in other studies; Finzi et al., 2021; Kay et al., 2015; Poltoratski et al., 2021), so did not vary as much over location. Despite research showing considerably poorer performance on face recognition tasks in the periphery

relative to the fovea (McKone, 2004), pRFs in the FFA did not increase in size as much with eccentricity compared to V1-V3 and OFA, suggesting that receptive field size may become less associated with acuity further in the ventral face processing stream. Research has also shown that pRF size is harder to reliably measure compared to other parameters such as position, being more susceptible to change due to various factors such as stimulus properties, attention, or fitting procedures (Alvarez et al., 2015; van Dijk et al., 2016). The lack of reliable effects for pRF size suggests that receptive field size is not the main factor driving differences in face perception across the visual field.

Instead, reliable effects were found according to the number of pRFs – and resulting visual field coverage – in each region of the visual field. All three face-selective brain regions had greater numbers of pRFs and better visual field coverage along the horizontal than vertical meridian. There was also an increased amount of pRFs and better coverage in the lower vs. upper field (although the difference in pRF number was not significant in pFus). These patterns tie in with the behavioural anisotropies found for faces in Chapter 2, suggesting that variations in face perception arise because face-selective cortex samples some regions of the visual field with a greater *number* of neurons than others. These sampling variations also align with those found in V1-V3, highlighting commonalities in retinotopic properties between early visual cortex and higher-level, face-selective parts of the brain.

The findings described above provide clues about how receptive field number and size contribute to visual field coverage within low- and high-level brain areas. The horizontal-vertical and upper-lower anisotropies in pRF number were magnified within the face-selective regions relative to early visual cortex, suggesting that variations in receptive field number across the visual field become exaggerated further up in the visual hierarchy. Interestingly, however, the anisotropies in visual field coverage were *less* pronounced in the face-selective areas than in V1-V3. This indicates that while differences in pRF number and their size affected overall coverage in V1-V3, coverage in the face-selective regions was predominantly driven by the amount of pRFs in each location. Consequently, this further implicates receptive field number as being

especially important for accurate face perception, suggesting that this factor drives overall coverage within face-selective cortex.

In the face-selective brain regions the horizontal-vertical differences were more pronounced than the upper-lower differences, particularly for pRF number. This could be related to the total amount of pRFs, as more were found along the horizontal than the vertical meridian (which comprised the upper and lower locations) overall. This variation is, however, consistent with studies showing that the behavioural upper-lower difference is smaller in magnitude and harder to measure in low-level vision (Barbot et al., 2021; Kurzawski et al., 2021). It also aligns with the results from Chapter 2, where a clear horizontal-vertical anisotropy was found across all three gender acuity experiments, but a greater number of trials was needed to reliably measure the upper-lower difference. Importantly, although the upper-lower anisotropy was smaller it is consistent with a general lower field bias within the visual system (Fortenbaugh et al., 2015; Greenwood et al., 2017; Rubin et al., 1996; Schmidtman et al., 2015), showing that common hallmarks of spatial vision emerge even within specialised face processing systems.

An exaggerated foveal bias was found in the face-selective regions compared to V1-V3, with the majority of pRF centres located within 5° from fixation, similar to previous findings (Finzi et al., 2021; Gomez et al., 2018; Poltoratski et al., 2021). As pRF number drives coverage estimates, coverage was also higher in the fovea in face-selective areas relative to V1-V3, and reduced more sharply in the periphery. Interestingly, when considering variations in pRF number, this central bias appeared to be even more pronounced in certain parts of the visual field. While pRFs were located in the central and peripheral visual field along the horizontal meridian, they were generally limited to central eccentricities along the vertical meridian. Within the vertical meridian, the upper field was particularly constricted, with almost no pRFs positioned in the periphery. As such, pRF size was unable to be estimated beyond the central few degrees in the upper field in face-selective regions. This highlights a distinct departure from early visual cortex, where pRFs were positioned across all eccentricities along both meridians.

Despite sharp reductions in pRF number in the periphery, face-selective areas showed coverage of the visual field up until around 15° eccentricity (similar to V1-V3), even along the upper vertical meridian. This clearly emphasises a role for large receptive fields that extend into peripheral vision, allowing us to recognise faces in peripheral vision (Kalpadakis-Smith et al., 2018; Kovacs et al., 2017; McKone, 2004; Roux-Sibilon et al., 2023) despite reduced pRF numbers. While large receptive fields may not be strongly linked to acuity, these results suggest that they are necessary for providing coverage in the periphery, given that a large proportion of receptive fields are positioned near the fovea. Large receptive fields do not appear to be enough to fully overcome the drop in pRF number, however, as gender acuity was worst (Chapter 2) in the upper field, which had the fewest pRFs.

The previous chapter suggested that variations in face perception could result from the passive pooling of information from earlier visual areas, with differing acuity for faces arising from differences in low-level input. Alternatively, anisotropies in face perception could be driven by an active process, where face-selective parts of the brain differentially sample visual field locations in the same way as earlier regions. These neuroimaging findings support the latter view, with the spatial selectivity of earlier brain areas inherited by face processing neurons, which themselves sample the visual field in a similarly biased manner. The horizontal-vertical and upper-lower differences in retinotopy (pRF number and coverage) were similar for upright and inverted faces, showing that spatial selectivity within face-selective cortex is similar regardless of the featural aspects of faces. This aligns with the results in Chapter 2, where the behavioural anisotropies were comparable in magnitude for upright and inverted faces (i.e. where performance was worse overall for inverted faces, but the *differences* in performance across the visual field followed a similar pattern for upright and inverted faces). Face-selective areas may apply the spatial properties they inherit universally, even though they process faces differently according to their featural content. While this would reflect an active process, it is also possible that similar spatial variations for upright and inverted faces occur more passively, with enhanced input from lower brain areas within certain visual field regions requiring increases in sampling (Riesenhuber & Poggio, 1999). Either way, while face-selective brain regions exhibit featural selectivity that is specific to faces, their spatial selectivity appears to be inherited from

other brain areas and *not* face-specific. This suggests that while the visual system employs specialised mechanisms for perceiving faces depending on their featural aspects, faces may not be special in terms of the way that they are processed spatially.

A clear behavioural face inversion effect was found, where participants were better at recognising the gender of upright compared to inverted faces. In line with previous suggestions (Le Grand et al., 2001; Rossion, 2008), this may reflect an added benefit of configural processing for upright faces, indicating that face-specific mechanisms were engaged during the pRF mapping task. Could this featural selectivity (upright faces processed better than inverted) be explained by the spatial properties (i.e. variations in pRF size) of neurons in face-selective cortex, despite spatial selectivity following similar patterns for upright and inverted faces? Although there were larger pRFs for inverted vs. upright faces in OFA, pRF size did not significantly differ in the FFA (pFus and mFus). Like with the location variations discussed above, this suggests that variations in pRF size are insufficient to explain featural as well as spatial selectivity.

Could featural selectivity be explained by the other neural properties measured? In mFus, there were more pRFs found for upright vs. inverted faces in the periphery, suggesting that an increased number of face-selective neurons sample the peripheral visual field when faces are upright. mFus also showed increased coverage in the fovea for upright faces, suggesting that there may be changes in responsiveness across the visual field depending on face orientation. However, these differences were smaller than expected – with no main effects of inversion in any of the face-selective regions – given the clear behavioural face inversion effect, and given that in Chapter 2, the variations in gender acuity according to inversion were of a similar magnitude to the variations across location. These smaller-than-expected neural inversion effects occurred despite there being considerably more data in the inversion compared to the location analyses⁶ (where effects of location emerged clearly). Altogether, the current findings suggest that configural processing – which relies on featural selectivity –

⁶ The inversion comparisons included data across the whole visual field, rather than being split by location.

cannot be fully explained by the receptive field properties of face-selective regions, as previously argued (Poltoratski et al., 2021; Witthoft et al., 2016).

Why, then, could there have been smaller pRFs for inverted vs. upright faces previously found in the FFA (Poltoratski et al., 2021)? This could be related to signal-to-noise variations, as the inverted faces in Poltoratski et al.'s (2021) study elicited weaker beta amplitudes (a measure of BOLD signal) than upright faces in the face-selective areas. Reduced beta amplitudes are typically associated with poorer goodness-of-fit (R^2) values (Figure A.4; Anderson et al., 2017; Schwarzkopf et al., 2014). This would affect the accuracy of size estimates, and indeed lower R^2 has been linked specifically to smaller pRF sizes (Alvarez et al., 2015). Similarly, Hughes et al. (2019) found that reductions in pRF size within parts of the brain that process motion could be attributed to lower stimulus visibility, with fewer voxels included in analyses due to their reduced responsivity. The results from the present experiment, which found evidence of configural face processing and observed similar beta amplitudes for upright and inverted faces in mFus (Figure A.4A), instead suggest that pRF size is not reliably modulated by face inversion.

Finding similar beta amplitudes for upright and inverted faces in mFus is perhaps unusual considering that other studies typically report weaker responses for inverted faces in the FFA (Goffaux et al., 2016; Kanwisher et al., 1998; Poltoratski et al., 2021; Yovel & Kanwisher, 2005). This could be linked to repetition suppression, with the repetition of face stimuli leading to reduced responses in the FFA (Dricot et al., 2008; Henson, 2016). These repetition effects have been linked to the high-level qualities (e.g. abstract identity) of faces, rather than their low-level characteristics or retinal position (Henson, 2016; Kovacs et al., 2012; Natu et al., 2016), and have been observed for upright but not inverted faces (Yovel & Kanwisher, 2005). This is presumably because upright faces are processed and perceived better – as was demonstrated in the current study, by the behavioural face inversion effect – making them more susceptible to effects of repetition. Increased repetition suppression for the upright faces could have resulted in their usually stronger neural responses being reduced to similar levels to the inverted faces.

The smaller-than-expected differences in coverage found in the present study highlight another distinction to previous research, where larger differences in coverage were found between upright and inverted faces in the FFA (Poltoratski et al., 2021). This could of course have been driven by the smaller pRFs found for inverted faces. It could also be related to the way that coverage was computed. When calculating coverage I included the exponent (the spatial summation parameter from the pRF model) to account for the Gaussian profile of receptive fields, as a better estimate of responsivity across the visual field (outlined in section 3.2.13). The addition of this extra parameter increased the need to apply some normalisation, especially as there were slightly lower exponents for inverted vs. upright faces (Figure A.2). The upright and inverted maps were normalised separately, meaning that the resulting estimates of coverage represented *relative* changes in responsiveness across the visual field. Coverage did not significantly differ between upright and inverted faces within the FFA when maps were jointly normalised, indicating that the absolute values were similar (Figure A.6). However, it is possible that a binary approach that does not account for the Gaussian shape of the receptive fields (Poltoratski et al., 2021) would have yielded larger differences in (absolute) coverage.

To summarise, these results show that face-selective parts of the brain sample locations across the visual field in a way that resembles early visual cortex, with a greater number of pRFs and better visual field coverage along the horizontal vs. vertical meridian and in the lower vs. upper field. These patterns explain the systematic behavioural variation found for face perception (Chapter 2), and reveal similar retinotopic properties between earlier visual areas and face-selective brain regions. pRF size was not reliably linked to acuity in face-selective regions, although large receptive fields may be necessary to support face perception in peripheral vision, given the magnified foveal bias compared to V1-V3. While there was some evidence that retinotopic measurements differed according to the featural content of faces more generally (e.g. increased foveal coverage for upright faces in mFus), the location variations were similar regardless of whether faces were upright or inverted, demonstrating that spatial selectivity within face-selective brain regions was not dependent on featural content. Together, these findings show that variations in face perception arise due to the way that face-selective parts of the brain differentially

sample the visual field, and uncover shared spatial selectivity between early visual cortex and specialised face-processing regions.

Chapter 4

Localising face-selective parts of the brain

4.1 Introduction

So far, I have demonstrated that face perception varies systematically across the visual field, with similar anisotropies to low-level vision. In the previous chapter, population receptive field (pRF) mapping (Dumoulin & Wandell, 2008) was used to show that these behavioural variations can be explained by the retinotopic profile of face-selective parts of the brain. Could the way that these face-selective regions were initially identified have altered the measurement of their retinotopic properties? To assess whether this was the case, I examined the process of localising and delineating face-selective brain areas, along with how variations in localisation approach affected retinotopic analyses. In this chapter, I will first discuss how regions of the brain that selectively respond to faces are typically identified. Then, I will describe how the methods involved could influence the measurement of retinotopic properties within these brain areas. Finally, a novel functional localiser will be introduced, which was designed to assess whether the delineation of face-selective brain regions depends on the spatial properties of the stimuli used for localisation.

In humans, cortical regions which selectively respond to faces rather than other objects were originally found by researchers wishing to map functional organisation within the brain. Using functional magnetic resonance imaging (fMRI), a region on the ventral surface of the right hemisphere – the fusiform face area (FFA) – was consistently found to respond more strongly to faces rather than other objects (Kanwisher et al., 1997; McCarthy et al., 1997). Using a similar approach, other studies quickly found that other parts of the brain could be selectively activated by presenting certain classes of stimuli, leading to the identification of body-selective and place-selective parts of cortex, for example (Downing et al., 2001; Epstein & Kanwisher, 1998). As these category-selective regions were located in similar anatomical locations across participants and could be reliably activated over time, this provided support for

the argument that the brain is organised into distinct modules that are specialised for certain types of processing (Fodor, 1983; Kanwisher, 2000).

A similar method of functional localisation is now commonly used to isolate face-selective brain areas, before exploring their properties. This typically involves separate fMRI analyses, where images of faces and other objects (e.g. houses), are presented in separate conditions (Poldrack, 2007; Saxe et al., 2006). Regions with high selectivity to faces are identified by contrasting the blood oxygen dependent (BOLD) responses to faces with BOLD responses to other (non-face) categories within each voxel of the brain and selecting voxels above a certain statistical threshold. This approach is useful as it provides greater statistical power compared to whole-brain analyses, by reducing the problem of multiple comparisons (Saxe et al., 2006). It also ensures that independent data are used to define and then analyse brain regions, to avoid spurious results which might otherwise arise from using the same data set for both, known as “double-dipping” (Kriegeskorte et al., 2009; Vul et al., 2009). Other analyses can then be applied to these regions to investigate their properties, such as pRF modelling.

The use of separate localisation analyses relies on the assumption that the localiser identifies fixed functional modules of the brain, which remain largely unchanged in location across both the localiser task and other tasks aiming to measure the modules’ properties of interest (e.g. retinotopy; Duncan et al., 2009; Friston et al., 2006). Yet, research shows that functionally defined regions can depend considerably on the localisation technique employed, such as the specific comparisons involved. For example, contrasting faces against houses led to stronger face-specific activation in the fusiform gyrus than contrasting faces with scrambled images (Berman et al., 2010). Similarly, researchers have shown that both the pattern and strength of face-specific activation in ventral temporal cortex depends on the number of categories that faces are contrasted with, again highlighting an effect of localisation methods on eventual delineation (Schwarz et al., 2019; Weiner & Grill-Spector, 2010). The statistical threshold used to define which voxels should be included also affects delineation; while a higher cut-off may identify face-selective regions with greater face-only activations, stringent thresholds can mean that category-selective regions are less reliably found (Duncan et al., 2009; Duncan & Devlin, 2011). Due to these

dependencies on the exact contrasts and statistical thresholds used, functional localisers may isolate brain regions which differ from those activated when the properties of interest are separately investigated.

The activation of category-selective brain areas clearly depends on the categories of object presented, and the statistical comparisons between them. But how much does localisation depend on the properties of the stimuli themselves (within each class of object)? This could be a particular issue for retinotopic mapping, which aims to measure the spatial properties of brain regions which have been localised using stimuli with their own spatial properties. Indeed, research shows that word-selective parts of ventral occipitotemporal cortex were larger in size when they were localised using large as opposed to small images (Le et al., 2017). Later pRF analyses of these word-selective areas showed that there was increased coverage of the peripheral visual field within the regions that had been isolated using the large localiser. Along similar lines, could the spatial properties of localiser stimuli affect the retinotopy measured within face-selective areas?

Throughout the visual system, the fovea is better represented – with increased neuronal density – than the periphery (Cowey & Rolls, 1974). Retinotopic mapping studies typically report an even more exaggerated representation of the central visual field in face-selective brain regions compared to V1-V3, with the majority of receptive fields in the FFA located less than 5° from fixation (Finzi et al., 2021; Gomez et al., 2018; Kay et al., 2015; Poltoratski et al., 2021). The extent of the central bias appears to increase from posterior to anterior areas within ventral face-selective cortex, with a greater foveal bias in the fusiform face area (FFA) than the occipital face area (Kay et al., 2015). Receptive fields are much larger in face-selective cortex than earlier visual areas, at several degrees of visual angle even at the centre of vision (Kay et al., 2015; Poltoratski et al., 2021). Despite these large receptive fields, these studies suggest that due to the foveal bias being so magnified, face-selective areas have an impoverished representation of the periphery (Finzi et al., 2021; Gomez et al., 2018; Kay et al., 2015; Poltoratski et al., 2021).

However, behavioural findings – including the three experiments in Chapter 2 – show that we can still accurately recognise faces that appear in peripheral vision

(McKone, 2004). Face inversion effects have also been observed in the periphery, suggesting that configural processing mechanisms are not just confined to the central visual field but occur in peripheral vision as well (Kalpadakis-Smith et al., 2018; Kovacs et al., 2017; Roux-Sibilon et al., 2023). This highlights a puzzling distinction – while existing retinotopic mapping studies suggest that face-selective parts of the brain have limited coverage in the periphery, behavioural research has shown that aspects of face processing thought to rely on face-selective mechanisms can be performed well in peripheral vision. Could this dissociation be linked to how face-selective parts of the brain were initially identified?

While retinotopic mapping studies present stimuli across the visual field, functional localisers typically show relatively small images at the centre of vision (Goffaux et al., 2016; Kanwisher et al., 1997; Stigliani, 2015; Weiner & Grill-Spector, 2010). One such localiser (Stigliani, 2015) is widely used in the literature to identify parts of the brain with selectivity for faces (Finzi et al., 2021; Gomez et al., 2018; Henderson et al., 2023; Kay et al., 2015; Natu et al., 2016; Poltoratski et al., 2021; Weiner & Grill-Spector, 2010; Witthoft et al., 2016). In this localiser (Stigliani, 2015), images were presented in five categories, each with two conditions (Figure 4.1): faces (adult, child), bodies (headless bodies, limbs), objects (cars, guitars), characters (pseudowords, numbers) and places (corridors, houses). Images varied slightly in their viewpoint, size and retinal position; one image could cover more of the lower left visual




Figure 4.1. Example stimuli from each of the categories in the functional localiser “fLoc”, described by Stigliani et al. (2015).

field and the next one more of the upper right, for instance. Importantly, however, images were always centred on the fovea. Stimuli were overlaid on phase-scrambled backgrounds which subtended the central 10.5° of the visual field (5.25° eccentricity). Notably, the stimuli themselves did not fill the whole extent of the image, with a considerable proportion of the scrambled background remaining visible across all image categories.

Presenting faces *only* foveally may preferentially activate receptive fields covering central vision, resulting in the delineation of face-selective regions which have a biased representation of the visual field. This may affect later analyses of retinotopy within these regions, by increasing the likelihood that pRFs will be centred near the fovea (Finzi et al., 2021; Gomez et al., 2018; Kay et al., 2015; Poltoratski et al., 2021). The importance of stimulus properties is illustrated by a pRF mapping study which used bars revealing different types of face stimuli (Sayres et al., 2010). Three ventral face-selective regions were originally localised by presenting stimuli foveally and at a considerably smaller eccentricity than the pRF experiment (Weiner & Grill-Spector, 2010), which could have biased them towards a central representation. Crucially, the position of pRFs in all three face-selective brain regions became more peripheral when bars (of the same size) revealed underlying images that contained large, single faces compared to multiple smaller faces (Sayres et al., 2010). In other words, altering the spatial properties – through changes in face size, quantity and positioning – of face stimuli shown during retinotopic mapping affected the position of pRFs, reducing the foveal bias. The spatial properties of *localiser* stimuli might therefore have a similar impact on pRF analyses.

Given the above issues, a functional localiser was developed which would examine whether altering the spatial properties – by using single and tiled configurations – of the stimuli involved would affect the delineation of face-selective brain regions. Ultimately, the goal was to define face-selective regions that represent both the central *and peripheral* visual field for later retinotopic analyses (described in Chapter 3). As such, stimulus images were displayed using a larger field of view (43.3° diameter, or 21.65° eccentricity) than previous studies, so that the images extended further into peripheral vision. Some of the images consisted of single, large, foveal

faces. Because it was not clear whether simply enlarging the single, foveal faces would sufficiently engage peripheral face processing (and because the single faces did not fill the stimulus background as much as desired when enlarged), smaller faces were also tiled across the visual field, so that faces were also *centred* in the periphery. As Chapter 2 indicated that the location of face features within the visual field affects face processing – with the position of the eyes linked to variations in gender acuity – including faces centred in the periphery was expected to drive peripheral face processing more optimally. To assess whether a combination of single and tiled faces would be beneficial, face-selective areas were also delineated using both stimuli configurations.

With this novel localiser, I expected to define three areas within the ventral face processing stream, as identified in other studies (Grill-Spector et al., 2017). Because the same inherently face-selective modules should be identified, the peak of the face-specific activation on the cortical surface was not expected to shift in location between single, tiled and combination configurations (Grill-Spector et al., 2017; Kanwisher et al., 1997; Silson et al., 2016). However, I hypothesised that compared to single (foveal only) faces, tiled (foveal *and peripheral*) faces would stimulate more of the peripheral visual field, resulting in larger face-selective areas. If the tiled faces activate additional peripheral neurons along with the *same* populations of neurons as the single faces, the face-selective areas identified for the tiled and combination localiser should be similar in size. Alternatively, if the single and tiled faces drive *different* populations of neurons there should be larger regions delineated for the combination version, compared to either the single or tiled stimuli alone.

In the previous chapter, pRF mapping was used to investigate how face-selective parts of the brain sample the visual field. I expected that the retinotopic properties measured during this experiment would differ depending on the localiser stimuli that had been used to identify the face-selective regions. Specifically, I hypothesised that there would be a greater number of pRFs and better visual field coverage in the periphery when face-selective regions were delineated using localiser stimuli that included peripheral (tiled and combination) vs. only foveal (single) faces. If different populations of neurons are activated by the localisers, there could also be

differences in pRF size subsequently measured from the regions, such as larger pRFs found for the single (large) faces, and smaller pRFs for the tiled (smaller) faces (Sayres et al., 2010). Otherwise, if the same populations of neurons are activated, pRF size should be unaffected by the localiser as it changes dynamically (Alvarez et al., 2015; Hughes et al., 2019; Kay et al., 2015) and should therefore depend on the retinotopic mapping stimuli in the previous chapter. Any differences uncovered during these analyses would confirm that the retinotopic properties of face processing brain regions are influenced by localisation methods.

Additionally, the relationship between face selectivity and variance explained (R^2) of the pRF model was examined. R^2 is not a direct measure of retinotopic sensitivity but instead quantifies how well a spatially localised pRF model (such as the one described in section 3.2.8) can explain the observed BOLD responses. It is also affected by factors such as BOLD response amplitude, with higher amplitudes (which would be expected from vertices that respond strongly to faces) associated with greater R^2 . However, as the pRF model is able to explain the most variance when the observed responses are spatially localised, higher R^2 values are generally associated with greater retinotopic sensitivity. Previous research suggests that there may be a negative relationship, with spatial sensitivity (R^2) decreasing as face-selectivity (T) increases (Silson et al., 2022). In contrast, a positive relationship would indicate that as face-selective increases, sensitivity to visual field location also increases. There could also be no relationship, suggesting that category selectivity is not linked to retinotopic sensitivity, but that the two principles are encoded independently (Kravitz et al., 2010).

4.2 Method

4.2.1 Participants

The same ten participants who participated in the retinotopic mapping experiment in Chapter 3 also took part here (six female, four male, $M_{\text{age}} = 29.1$ years, all Caucasian and with normal or corrected-to-normal vision). All ten participants took part in the functional localiser and pRF mapping experiment within one scanning session, while two also took part in an additional scanning session involving the

commonly used ‘fLoc’ localiser (Stigliani, 2015). The experiment was approved by the Research Ethics Committee for Experimental Psychology at University College London and all participants gave written informed consent before testing began.

4.2.2 Stimuli

Stimuli were programmed using MATLAB (MathWorks, Inc) and PsychToolbox (Brainard, 1997; Kleiner et al., 2007; Pelli & Vision, 1997) and displayed on a back-projection screen in the bore of the magnet using an EPSON EB-L1100U projector that had a maximum luminance of 502 cd/m². The screen size was 27 cm x 27 cm with a resolution of 1920 x 1200 pixels, and stimuli were displayed at 1200 x 1200 pixels. The monitor had a refresh rate of 60 Hz. Participants viewed the screen through a mirror attached to the head coil at a viewing distance of 34 cm, giving a maximum field of view of 43.3° (21.65° eccentricity).

To create my new functional localiser, I adapted code available from the fLoc functional localiser package (Stigliani, 2015). Localiser images were supplied by researchers at UCLouvain (Schuurmans et al., 2023), which were a subset of the images in Stigliani et al. (2015) and were unedited apart from having had their backgrounds removed. Images were greyscale and consisted of faces of various viewpoints, isolated hands in different positions, and stringed instruments (20 of each). This number of object categories was judged as sufficient to identify face-selective areas (Berman et al., 2010; Kanwisher et al., 1997), while allowing for additional manipulations of image configuration – single vs. tiled – within the time available, and so that the length remained similar to existing localisers (Peelen & Downing, 2005; Stigliani, 2015). Images were resized to 1200 x 1200 pixels in Adobe Photoshop CS6. If needed, the object within each image was made bigger so that it filled most of the total image dimensions, and so that the space filled was as uniform as possible across the different objects. Contrast normalisation was applied so that the images had a root mean square (RMS) contrast of 0.15 for faces, and 0.1 for hands and instruments. 20 tiled images for each object category were then created by randomly selecting nine images from the same category and combining them in a three-by-three grid. The same image could not appear twice within a tiled image. This resulted in *three object*

categories and two tiling conditions: single faces, tiled faces, single hands, tiled hands, single instruments and tiled instruments (Figure 4.2A).

Each image then underwent a fast Fourier transform (FFT) and was iteratively phase-scrambled, which consisted of scrambling the face (or hand/instrument) image, pasting the face back onto this scrambled image, scrambling again, and repeating this process 500 times (Petras et al., 2019). This ensured that image backgrounds had




Figure 4.2. **A.** Example stimuli from each of the six conditions in the new functional localiser. **B.** Experimental procedure. A red fixation dot appeared at the centre of the screen. Each block lasted ten seconds. Experimental blocks contained 20 stimuli per block, each lasting 500 ms. Scrambled images were inserted randomly in half of the experimental (not baseline) blocks. The localiser began with a baseline block, then 49 blocks were shown (seven for each of the six stimulus conditions, plus seven baseline) in a randomised order, followed by another baseline block to end. This comprised one localiser run, lasting eight minutes and 35 seconds in total. Image contrast has been increased slightly in this figure, to improve visibility.

similar spectral properties as the objects themselves and so would not drive visual responses differently, which would reduce the contrasted activation to each category of object (e.g. if faces appeared on backgrounds that had similar spectral properties to instruments, there would be less face-specific activation found with a face > instrument contrast than if the backgrounds contained similar spectral content to faces). In addition to the object images the final scrambled images (with no object superimposed on top) were saved, from which 20 were selected randomly – but with an equal distribution across object categories – as scrambled images to be used in the task. The projector was calibrated and these values used to gamma correct the images, so that their mean luminance was 251 cd/m².

Images from each of the six conditions were shown in separate blocks, along with baseline (blank) blocks (Figure 4.2B). Blocks lasted ten seconds each, which contained 20 stimuli each displayed for 500 ms. In each run there were 51 blocks (seven blocks for each of the six conditions plus baseline, plus an extra baseline block at both the start and the end of the run), lasting eight minutes and 35 seconds. Each participant completed two runs.

To ensure attention, participants were instructed to maintain fixation on a red dot in the centre of the screen and press a button in response to a scrambled image, which randomly appeared in half of the experimental (not baseline) blocks. Responses were recorded via a button box. Participants did not receive feedback, however performance was monitored throughout the experiment to ensure that the task was being performed properly. Behavioural responses were otherwise not analysed.

4.2.3 ‘fLoc’ functional localiser

To evaluate the effectiveness of the new localiser, two of the ten participants were also shown the commonly used ‘fLoc’ (Stigliani, 2015). As shown in Figure 4.1, there were five object categories, each with two subtypes: faces (adult and child), bodies (limbs and headless bodies), objects (cars and instruments), places (houses and indoor corridors), and characters (numbers and pseudowords). Images from each category were shown in separate blocks, interspersed by baseline (blank) blocks (Figure 4.3). Stimuli were split into two sets, with one set comprising adult faces,




Figure 4.3. Experimental procedure of the 'fLoc' functional localiser (Stigliani et al., 2015). A red fixation dot appeared at the centre of the screen. Each block lasted six seconds. Experimental blocks contained 12 stimuli, each lasting 500 ms. Scrambled images were inserted randomly in half of the experimental (not baseline) blocks. The fLoc began with a baseline block, then 36 blocks were shown (six for each of the five stimulus conditions, plus six baseline) in a randomised order, followed by another baseline block to end. This comprised one localiser run, lasting three minutes and 57 seconds in total.

headless bodies, pseudowords, cars and houses and the other set containing child faces, limbs, numbers, corridors and instruments. There were six blocks for each of the five categories in the stimulus set and the baseline, with one extra baseline block at the start and end, making for 38 blocks per run. Blocks lasted six seconds, containing twelve stimuli each presented for 500 ms. Apart from the baseline blocks at the start and end, the order of the blocks was randomised during each run. Each run lasted three minutes and 57 seconds.

Typically, previous studies have presented localiser images at a maximum eccentricity of around 5-7° (Silson et al., 2022; Stigliani, 2015; Weiner & Grill-Spector, 2010; Witthoft et al., 2016), although this can be hard to determine as the size of localiser stimuli is often not explicitly reported. To also test whether the overall size of localiser stimuli affects the properties of brain regions identified, the fLoc was displayed at both a small (5.44° eccentricity) and large (21.65° eccentricity) presentation size.

Participants completed four runs (two of each stimuli set) of the small version and four runs of the large version (a total of 15 minutes and 48 seconds for each size).

As with the new localiser, participants were instructed to maintain fixation on a red dot in the centre of the screen and press a button in response to a scrambled image, which randomly replaced one of the images in half of the experimental (not baseline) blocks. Responses were recorded via a button box. Participants did not receive feedback, however key presses were monitored throughout the experiment to ensure participants were performing the task properly. Behavioural responses were otherwise not analysed.

4.2.4 MRI data acquisition

Functional and anatomical scans were obtained using a Siemens Prisma 3T MRI scanner (Siemens, Erlangen, Germany). A 64-channel head coil was used, with cushions placed around participants' heads to minimise movement. A T1-weighted anatomical magnetisation-prepared rapid acquisition with gradient echo (MPRAGE) image was acquired (TR = 2300 ms and TE = 2.98 ms, voxel size = 1 mm isotropic voxels). For the functional scans only the back of the head coil was used, leaving 42 channels. Functional T2-weighted multiband 2D echoplanar images were acquired (repetition time (TR) = 1000 ms, TE = 35.20 ms, voxel size = 2 mm isotropic voxels, 48 slices, flip angle = 60°, acceleration factor = 4). Each functional scan contained 510 volumes. A short 30 second localiser was carried out before the functional scans and again before the anatomical scan, after the front head coil was fitted. Fixation was monitored live by the experimenter using an Eyelink 1000, although eye tracking data were not recorded. In all participants apart from author AYM, functional localiser runs were carried out after the pRF runs as I reasoned that although important for both paradigms, keeping precise fixation was particularly crucial for the pRF mapping.

4.2.5 MRI data preprocessing

For each participant, the T1 anatomical scan was automatically segmented and used to generate a 3D representation of the cortical surface using Freesurfer (Dale et al., 1999; Fischl, 2012; Fischl et al., 1999). Functional images were B0 distortion corrected and motion corrected using AFNI software (Cox, 1996; Cox & Hyde, 1997).

An alignment volume was created by finding the volume which had the least voxel outliers (relative to the trend of the time series) across all runs, which all functional volumes were then aligned to. Using Freesurfer (Fischl, 2012) the alignment volume was coregistered to the structural image, and surface projection was performed.

4.2.6 Region of interest (ROI) definition

To estimate blood oxygen level dependent (BOLD) responses to the different stimulus categories within each voxel, we ran a general linear model (GLM) for each participant using SPM12 software (Ashburner & Friston, 2014). I collapsed across object and tiling conditions and entered each into a design matrix. As such, there were six stimulus conditions entered: single faces, tiled faces, single hands, tiled hands, single instruments, tiled instruments. The scrambled images (used as targets to monitor attention during the object category blocks) were modelled as a regressor of non-interest, alongside the six stimulus conditions. This avoided the scrambled images being included as part of an object category. Motion regression was accounted for by using the output from the motion correction step carried out during preprocessing, with one regressor added for each of the six directions of movement (roll, pitch and yaw around both x and y axis). The predicted fMRI time series was then modelled according to the onset times of each stimulus, convolved with a canonical haemodynamic response function (HRF), and compared to the observed time series. Scans were concatenated and correction applied after the model estimation to account for the concatenation of different runs (for example, because the convolution of the HRF between the start and end of each scan would not be continuous between separate scans). This resulted in one GLM (per participant) which contrasts were then applied to.

Three types of statistical contrast were carried out to compare BOLD responses between the different face stimuli. The main contrast involved both image configurations combined, with single and tiled faces contrasted against other single and tiled objects (single faces, tiled faces > single hands, tiled hands, single instruments, tiled instruments). As I wished to assess whether using *peripheral* (tiled) or *only foveal* (single) faces during localisation would affect delineation, I ran two additional contrasts. The first consisted of single faces against other single objects

(single faces > single hands, single instruments), and the second involved tiled faces against other tiled objects (tiled faces > tiled hands, tiled instruments). I avoided intermixing the tiling conditions in these two contrasts so that the responses to the foveal and peripheral stimuli could be independently examined, without any influence of the other configuration. This resulted in *single*, *tiled* and *combination* contrasts for each participant.

Unthresholded statistical maps were generated in SPM, which consisted of a T statistic for each voxel which represented the difference in BOLD response between the contrasted stimuli (e.g. single and tiled faces > single and tiled hands and instruments). Separate maps were computed for each contrast. These maps were smoothed using a goodness-of-fit threshold of 0.1 and a smoothing kernel of 3 mm full width half maximum (FWHM), and surface projected using Freesurfer (Fischl, 2012). The contrast maps were used during delineation as a visual aid and to reveal the MNI coordinates of activated regions, which were converted to Talairach coordinates (Lacadie et al., 2008). There is no good consensus about which T (or sometimes p) value threshold should be used, with considerable variation across studies. I chose a threshold of $t \geq 2$ for delineating the ROIs, as more liberal thresholds have previously yielded more consistent regions (Duncan et al., 2009), and this threshold increased the number of vertices included in the pRF analyses in Chapter 3. All contrast maps were therefore thresholded at the $t \geq 2$ level, exported from SPM as NIFTI images, and surface projected using the SamSrf 9.4 MATLAB toolbox (Schwarzkopf, 2022), so that they could be viewed retinotopically alongside the V1-V3 delineations.

Delineations of face-selective ROIs were then carried out in SamSrf 9.4 (Schwarzkopf, 2022), using the thresholded T maps for each contrast. First, large areas were manually drawn around clusters of activation, based upon the expected anatomical locations (Weiner & Grill-Spector, 2010) and Talairach coordinates of face-selective areas (Kanwisher et al., 1997; Kung et al., 2007). The vertex with the peak T statistic within each large area was then automatically identified, with neighbouring vertices included in the resulting ROI if their T value exceeded the chosen threshold ($t \geq 2$). This automatic process meant that the resulting ROIs consisted of a continuous set of vertices surrounding the peak voxel, within the large area that was originally

drawn. In nine participants I delineated three face-selective regions: the occipital face area (OFA), the posterior (pFus) and the medial fusiform face area (mFus). In the remaining participant, pFus could not be identified and was most likely absent (as has previously been shown to be the case in some individuals; Chen et al., 2022). For each of the face-selective regions in each participant, a single, tiled and combination version was delineated. This resulted in nine ROIs per participant (three versions of each of the three face-selective regions, apart from one participant who had six, as only OFA and mFus were delineated).

4.2.7 Statistical analyses

To analyse whether the *size* of the face-selective regions delineated was affected by the type of localiser stimuli, t-tests were run to compare whether the number of vertices (averaged across participant) in each ROI differed depending on whether they had been defined using single, tiled or combination stimuli. This resulted in nine t-tests. As these comparisons were planned a priori and stated specifically in my hypotheses, I did not apply Bonferroni correction.

To determine whether the *position* of face-selective regions differed according to localiser stimuli, I extracted the vertex RAS (Right, Anterior, Superior) coordinates which represented the peak of the face-selective activation within each of the nine ROI labels (i.e. the most responsive vertex). These coordinates are individual specific, so it was not possible to average them across participant. Converting the coordinates to a standard brain template could have introduced mapping inaccuracies, so I chose to compare them using analyses of variance (ANOVAs), where the values could remain individual specific. As such I performed a four-way mixed effects ANOVA for each of the coordinate types, with localiser stimuli (single, tiled, combination), ROI (OFA, pFus, mFus) and hemisphere (left, right) as fixed factors and participant as a random factor, with either the Right, Anterior or Superior coordinates as the dependent variable.

After examining the size and location of the face-selective ROIs, their retinotopic properties were analysed. As there were significant main effects and/or interactions of ROI in all analyses (Table B.1-B.3), they were run within each ROI separately. Linear mixed models were used to determine whether localiser stimuli could predict pRF size

within the ROIs. Separate mixed effects models were run for each ROI (OFA, pFus and mFus). The model included fixed factors of localiser (single, tiled, combination), eccentricity, and inversion (upright/inverted). pRF data were included across the whole visual field, i.e. were not filtered by location. Participant was specified as a random factor for the intercept as well as for each of the fixed factors, as the slope of the relationship between pRF size and localiser, eccentricity, and/or inversion could vary across individuals. Participants could therefore have different intercepts and slopes for each of the factors.

Mixed effects ANOVAs were used to assess the effects of localiser, eccentricity, inversion and participant on pRF number and visual field coverage. Separate ANOVAs were run for each ROI (OFA, pFus and mFus). The within-subjects fixed factors were localiser, eccentricity and inversion (upright/inverted). In all of the ANOVAs, participant was entered as a between-subjects random factor to account for individual variation within the model, although the participant effects will not be reported as I was mainly interested in the group-level effects of the localiser stimuli. Following significant main effects, paired t-tests were performed to explore the localiser differences in more depth.

Finally, linear mixed models were performed to assess whether the variance explained (R^2) of the pRF model could predict face selectivity. As a measure of face-selectivity, T -statistics were taken from the SPM maps generated during delineation of face-selective ROIs using the combination localiser (single and tiled faces > single and tiled hands and instruments). These T values represent the strength of face-selective activation within each vertex. As there was a significant effect of ROI (Table B.4), separate analyses were run within each region (OFA, pFus and mFus). The model included a fixed factor of R^2 . Participant was specified as a random factor for R^2 and the intercept, to account for any individual variation in the slope of the relationship between face selectivity and R^2 .

4.3 Results

To ensure that the stimuli (4.2.2) and contrasts (4.2.6) used in the new localiser were sufficiently identifying face-selective regions, I compared the face-specific

activation obtained with the combination localiser (which included single and tiled faces, with a large presentation size) and the commonly used ‘fLoc’ (Stigliani, 2015), which two out of the ten participants also took part in. The fLoc was shown to both of these two participants at a small (as in Stigliani, 2015) and large (the same as the




Figure 4.4. **A.** ROI size displayed as the mean number of vertices for the small fLoc, large fLoc, and the new localiser (which used a large presentation size and a combination of single and tiled face stimuli), across two participants. Dots represent individual data points. **B.** Face-selective regions OFA, pFus and mFus on the ventral surface of the right hemisphere in one participant, for the small fLoc, large fLoc, and the new localiser.

combination localiser) presentation size. As only two participants took part in the fLoc, statistical analyses were not performed on these data, but patterns will be described. Firstly, the general locations of the face-selective areas were very similar between the combination localiser and the fLoc, suggesting that the new approach was identifying the same face-selective modules (Figure 4.4B). Secondly, there was generally a greater number of vertices – meaning that face-selective ROIs were larger – for the large vs. small fLoc (Figure 4.4A). This suggests that simply displaying the localiser at a larger size can lead to face-specific activation found across a greater region of the cortical surface (Figure 4.4B), consistent with previous findings for word-selective brain regions (Le et al., 2017). Lastly, there were trends towards the combination localiser identifying larger face-selective regions than both the small and large fLoc, particularly within the FFA (pFus and mFus). This indicates that the novel stimuli (i.e. with objects located peripherally as well as foveally, and with a larger field of view than usual) were sufficient – if not better – at identifying face-selective activity within the ventral face processing stream. Therefore, I moved on to comparing the three different versions of the stimuli within the new localiser (single, tiled and combination).

4.3.1 ROI size

Firstly, I examined whether the size of the face-selective ROIs was affected by the different localiser stimuli, in terms of the total number of vertices they contained. As outlined in the introduction, I expected to define larger face-selective regions when localiser stimuli included faces that were also located peripherally (tiled and combination faces) as opposed to only foveally (single faces). OFA was the largest face-selective area overall, averaging 912.9 vertices ($SD = 499.6$). pFus and mFus were around half the size of OFA, and were similar in size to each other, averaging 531.3 ($SD = 285.0$) and 467.0 ($SD = 310.6$) vertices respectively. Despite overall differences in size, the delineation of the three regions was similarly affected by the localiser stimuli (single faces, tiled faces or a combination of both). OFA, pFus and mFus all had the fewest vertices on average for single faces, then tiled faces, and the most when single and tiled faces were combined (Figure 4.5).

Paired t-tests revealed that the number of vertices did not significantly differ between single and tiled stimuli, for any of the three face-selective ROIs (Table 4.1). However, the combination stimuli resulted in a significantly higher number of vertices than either single or tiled faces alone. In other words, using both the single and tiled stimuli during localisation led to the delineation of significantly larger face-selective brain regions. This suggests that the single and tiled faces may have activated different vertices, with the combination stimuli resulting in significantly larger ROIs than either face type alone as both sets of vertices were sufficiently stimulated.




Figure 4.5. Mean number of vertices in each of the face-selective regions, delineated using localiser stimuli consisting of single faces, tiled faces, or a combination of both. Dots indicate the mean number of vertices for each individual, and error bars show the SEM. Asterisks denote statistically significant differences ($p < .05$).

4.3.2 ROI locations

After examining the size of the face-selective ROIs, their location on the cortical surface was assessed. As expected, Figure 4.6 shows that the location of the peak (i.e. most responsive) vertex of face-selective areas was similar regardless of the localiser stimuli. Three ANOVAs were run to analyse whether the RAS (right, anterior, superior) coordinates of the peak vertex varied between localiser stimuli (coordinates for each participant's right hemisphere are listed in Table B.4). There was no main effect of stimulus type (single, tiled or combination faces) for either the right ($F(2,36) = 0.01$, $p = .995$), anterior ($F(2, 36) = 0.04$, $p = .958$) or superior ($F(2, 36) = 0.00$, $p = .996$) coordinates, indicating that the peak of the face-selective activation did not change according to the localiser stimuli. Instead, Figure 4.6 shows that the localiser stimuli affected how the face-specific activation was spread out over the cortical surface. Specifically, the combination of both single and tiled faces resulted in a greater spread of activation over the cortical surface than either stimulus type alone. Upon visual inspection I noted that pFus and mFus tended to spread out more anteriorly and laterally, covering a greater region of the fusiform gyrus.

ROI	Comparison	M_{diff}	t	p	Cohen's d
OFA ($df = 9$)	single vs tiled	-136.6	-1.83	.101	-0.58
	combi vs single	316.8	6.56	< .001	2.1
	combi vs tiled	180.2	2.91	.017	0.92
pFus ($df = 8$)	single vs tiled	-48.9	1.12	.261	-0.40
	combi vs single	126.1	5.02	.001	1.67
	combi vs tiled	77.2	3.83	.005	1.28
mFus ($df = 9$)	single vs tiled	-36.0	-0.85	.419	-0.27
	combi vs single	167.4	4.87	< .001	1.54
	combi vs tiled	131.4	4.93	< .001	1.56

Table 4.1. Results of t-tests comparing the number of vertices as a measure of ROI size in three ventral face-selective regions, which were delineated using either single, tiled or combination localiser stimuli. Bold font indicates statistical significance ($p < .05$).

4.3.3 Effects on pRF measurements

The previous sections show that face-selective ROIs were larger when localiser stimuli contained faces that were located peripherally (the tiled and combination stimuli) as well as foveally. It was important to determine the properties of these additional vertices, and whether their inclusion or exclusion within face-selective ROIs affected the retinotopy measured within the regions. As such, I analysed how the retinotopic properties measured in Chapter 3 differed within ROIs delineated using single, tiled and combination localiser stimuli. Within each ROI version, I examined pRF size, number (the total amount of vertices remaining at each eccentricity after noise and artefact filtering; see 3.2.12) and visual field coverage (an estimate of how well each region of the visual field is sampled, given the number and size of pRFs in each area and their response profiles; see 3.2.13). Figure 4.7 displays these pRF data across eccentricity for upright faces. While the pRF experiment also involved inverted faces, patterns during the following analyses were similar regardless of inversion – with no statistical interactions between location and inversion (Table 4.3 and Table 4.4) – so for clarity, only upright face data are shown.

Firstly, I assessed pRF size. The different localiser stimuli all resulted in ROIs that showed the typical increase in pRF size with eccentricity (Figure 4.7A). Linear mixed models did not find any significant effects of the localiser on pRF size within any of the face-selective regions, indicating that pRFs were similar in size regardless of the spatial properties of localiser stimuli (Table 4.2).

Next, I looked at pRF number. Within all of the ROIs there was the expected decrease in pRFs across eccentricity, regardless of localiser stimuli (Figure 4.7B). In all three face-selective regions there were significant effects of the localiser on pRF number, with the fewest amount of pRFs found across the visual field when localisation was performed using single faces, then tiled faces, and the most pRFs found across eccentricities for the combination stimuli (Table 4.3). Significant interactions between localiser and eccentricity in all ROIs reflected that these differences were most apparent in the fovea, where there were the most pRFs overall. t-tests revealed that in all three face-selective regions, there were significantly more pRFs found across the




Figure 4.6. Face-selective regions OFA, pFus and mFus on the ventral surface of the right hemisphere in two participants (P1 and P2), defined using different localisation stimuli (single faces, tiled faces, or a combination of both). Each ROI version (single, tiled, combination) has the same general location, with the activation spread out over a different amount of the cortical surface depending on the localisation approach.




Figure 4.7. Mean pRF size (**A**), number (**B**), and visual field coverage (**C**) across eccentricity (in degrees of visual angle) for face-selective ROIs delineated using single, tiled or combination face stimuli. Only data from pRF mapping with upright faces are shown. For pRF number, statistically significant differences ($p < .05$) between combination and single stimuli (black line) and between the combination and tiled stimuli (grey line) are shown across the top. (**D**). The percentage of pRFs found across eccentricity using the single (blue) or tiled (orange) vs. the combination version of the localiser, which had the greatest number of pRFs overall. Lower percentages represent larger differences, i.e. a larger proportion of pRFs was found with the combination stimuli.

ROI	Fixed factors	β	p	CI
OFA	Intercept	1.70		
	Localiser	0.04	.643	-0.13, 0.20
	Inversion	1.02	.002	0.37, 1.67
	Eccentricity	0.34	< .001	0.30, 0.37
pFus	Intercept	5.32		
	Localiser	-0.05	.393	-0.16, 0.06
	Inversion	0.18	.632	-0.57, 0.94
	Eccentricity	0.20	.005	0.06, 0.35
mFus	Intercept	7.16		
	Localiser	-0.02	.922	-0.33, 0.30
	Inversion	-0.50	.582	-2.28, 1.28
	Eccentricity	0.18	< .001	0.08, 0.28

Table 4.2. Linear mixed model results comparing pRF size across localiser stimuli. Significant effects ($p < .05$) are indicated in bold.

visual field for the combination vs. single localiser, with mFus showing the greatest amount of significant differences (Figure 4.7B). In mFus, there were also more pRFs found across central and peripheral eccentricities for the combination compared to the tiled localiser. In OFA and pFus, while there were significantly more pRFs for combination vs. tiled stimuli in the central visual field, these differences did not extend as far into the periphery. mFus therefore appeared to have benefitted the most from a combination of single and tiled faces.

These findings show that the combination localiser identified significantly more pRFs across the visual field, but did these changes vary in magnitude between central and peripheral vision? To explore this, I calculated the amount of pRFs found with the single and tiled localisers as percentages of the number identified using the combination version, at each eccentricity (Figure 4.7D). There were generally smaller percentages for the single compared to the tiled faces, reflecting the larger differences in pRF number found between the combination and single, compared to the combination and tiled, localisers. Despite there being a significantly reduced amount of pRFs in the periphery within face-selective areas compared to the fovea, the differences in pRF number between the different localisers were proportionally similar

in central and peripheral vision. In fact, in mFus there was a trend towards the differences increasing in magnitude in the periphery.

ROI	Factors	df	F	p
OFA	Localiser	2, 180	16.37	< .001
	Inversion	1, 180	0.34	.573
	Eccentricity	20, 180	14.49	< .001
	Localiser*Inversion	1, 180	3.48	.052
	Localiser*Eccentricity	20, 180	6.10	< .001
	Inversion*Eccentricity	20, 180	3.21	< .001
pFus	Localiser	2, 160	6.70	.008
	Inversion	1, 160	2.95	.124
	Eccentricity	20, 160	5.46	< .001
	Localiser*Inversion	1, 160	1.22	.320
	Localiser*Eccentricity	20, 160	2.15	< .001
	Inversion*Eccentricity	20, 160	1.87	.018
mFus	Localiser	2, 180	12.17	< .001
	Inversion	1, 180	1.01	.341
	Eccentricity	20, 180	7.45	< .001
	Localiser*Inversion	1, 180	0.51	.609
	Localiser*Eccentricity	20, 180	5.84	< .001
	Inversion*Eccentricity	20, 180	1.53	.078

Table 4.3. ANOVA results comparing pRF number across localiser stimuli. Significant effects ($p < .05$) are indicated in bold.

Next, visual field coverage was analysed across the different localiser stimuli. Coverage plots were generated for each ROI and localiser using the centre position, size and Gaussian profile – determined by the compressive exponent – of pRFs, as described in 3.2.13. Because I was primarily interested in how the localiser affected measurements of the visual field representation within ROIs, the coverage plots were normalised separately between ROI versions, to provide an estimate of *relative responsiveness* across the visual field. Coverage should be increased in the periphery (i.e. a more homogenous response across the visual field) when localiser stimuli contained peripherally as well as foveally located faces. When only foveal faces were used, there should be a stronger foveal bias in coverage (i.e. a more imbalanced response between the fovea and periphery). This normalisation approach allowed me to examine the relative changes in coverage across eccentricity. However, even when

normalising the maps for each localiser within the same frame (dividing all maps by the largest value in the combination map) the coverage estimates were similar, particularly in the FFA (Figure B.1). This indicates that the absolute coverage values were comparable between localisers.




Figure 4.8. Mean visual field coverage for face-selective ROIs, delineated using different localiser stimuli (single faces, tiled faces, or a combination of both). Plots contain data from pRF runs containing upright faces (Chapter 3; see 3.2.13 for further detail on how plots were generated). Coordinates represent eccentricity in degrees of visual angle, with negative values for the left and positive values for the right visual field. Values were converted to log scale before plotting, for visualisation purposes (see colour bar). Dots represent pRF centres from all participants.

In line with typical retinotopy, coverage was significantly higher at the fovea and decreased in the periphery within all face-selective ROIs, for all localiser types (Figure 4.7C; Table 4.4). Figure 4.8 displays visual field coverage plots and reveals small variations between the localiser types. In OFA and mFus, higher (yellow) coverage estimates appear to reach slightly further into the peripheral visual field for the combination and tiled vs. the single localiser. Despite this there were no effects of the localiser in any of the face-selective regions (Table 4.4), indicating that coverage across the visual field was similar irrespective of localiser stimuli.

Clearer from the coverage plots, however, is the increased number of pRF centres in the peripheral visual field in OFA and mFus with the combination and tiled vs. single localiser. This emphasises that the changes in localiser stimuli had the largest impact on overall pRF number. While this variation is most visible in the periphery, where there are considerably fewer pRFs overall, including both single and tiled faces in the localiser stimuli increased the amount of pRFs found across the visual

ROI	Factors	<i>df</i>	<i>F</i>	<i>p</i>
OFA	Localiser	2, 180	0.17	.847
	Inversion	1, 180	0.06	.816
	Eccentricity	20, 180	330.93	< .001
	Localiser*Inversion	1, 180	3.45	.054
	Localiser*Eccentricity	20, 180	0.56	.986
	Inversion*Eccentricity	20, 180	1.05	.407
pFus	Localiser	2, 160	0.48	.627
	Inversion	1, 160	0.17	.691
	Eccentricity	20, 160	672.85	< .001
	Localiser*Inversion	1, 160	0.12	.889
	Localiser*Eccentricity	20, 160	0.40	1.00
	Inversion*Eccentricity	20, 160	0.27	.999
mFus	Localiser	2, 180	0.56	.579
	Inversion	1, 180	2.62	.140
	Eccentricity	20, 180	314.61	< .001
	Localiser*Inversion	1, 180	1.06	.366
	Localiser*Eccentricity	20, 180	1.24	.156
	Inversion*Eccentricity	20, 180	1.03	.429

Table 4.4. ANOVA results comparing visual field coverage across localiser stimuli. Significant effects ($p < .05$) are indicated in bold.

field. Altering the spatial properties of localiser stimuli did not therefore result in significant changes in coverage within the face-selective regions.

4.3.4 Face-selectivity vs. R^2

Lastly, the relationship between face selectivity (where higher T -statistics represent stronger face-selective activation) and R^2 (variance of the BOLD response explained by the pRF model, where higher values are generally consistent with greater retinotopic sensitivity) was examined. All three face-selective regions had a clear positive correlation between face selectivity and R^2 , with T values increasing as R^2 increased (Figure 4.9).




Figure 4.9. Face selectivity (T -statistics, with greater values indicating greater face-selective activation) plotted against R^2 (variance explained by the pRF model). Each data point represents a single vertex from one participant, with different colours for each participant.

Linear mixed models confirmed these positive correlations, with R^2 significantly predicting face selectivity in all three ROIs (Table 4.5). This suggests that as selectivity for faces increased, so did sensitivity to visual field location, as a spatially localised pRF model explained more of the variance in BOLD responses.

ROI	Fixed factors	β	p	CI
OFA	Intercept	2.35		
	R^2	5.52	< .001	3.63, 7.40
pFus	Intercept	0.53		
	R^2	9.83	.003	3.41, 16.24
mFus	Intercept	2.70		
	R^2	4.67	< .001	1.56, 4.14

Table 4.5. Linear mixed model results assessing whether R^2 (variance explained of the pRF model) predicts face-selectivity (T). Significant effects are represented in bold ($p < .05$).

4.4 Discussion

The aim of this chapter was to analyse whether the spatial properties of localiser stimuli would affect the delineation of, and retinotopic properties subsequently measured from, functionally defined brain regions. Three ventral face-selective parts of the brain (OFA, pFus and mFus) were localised using face stimuli that covered a large field of view (21.65° eccentricity). As expected, the location of the peak activation on the cortical surface was unaffected by localisation methods in all three face-selective regions (in line with Silson et al., 2016). However, a greater area of the cortical surface was delineated when the localiser stimuli included large, foveal faces *and* smaller faces that were tiled across the visual field, which included faces positioned in peripheral vision, compared to either single or tiled faces alone. Although the differences were not significant, all three face-selective areas were larger when they were delineated using the tiled (foveal and peripheral) as opposed to the single (foveal only) faces. Examination of their retinotopic properties revealed that the larger brain regions identified using a combination of single and tiled faces contained a greater number of pRFs across the visual field, while pRF size and visual field coverage were not significantly affected by the localisation method. Together, these results show that the retinotopy measured within face-selective brain regions was affected by the spatial properties of the stimuli used during localisation.

Localiser stimuli which included faces that were positioned in the periphery⁷ resulted in the delineation of larger face-selective brain regions, with activation spreading out over a larger region of the cortical surface, compared to when only foveal faces were presented. Retinotopic analyses showed that the larger ROIs contained a greater number of pRFs across the visual field, with these additional pRFs indicating a more accurate visual field representation. In other words, the larger ROIs identified by including peripherally located faces were specifically linked to a greater number of pRFs (rather than properties such as pRF size). Interestingly, when tiled faces were included, pRF numbers increased near the fovea as well as in the periphery. This could reflect the quantity of faces increasing, as each tiled image contained nine faces which

⁷Peripherally located faces were included in both the tiled and combination versions of the localiser. Where the tiled version is referred to in isolation, this will be made clear.

would each elicit a face-specific response in the brain, as opposed to only one when single faces were shown. While some studies have found suppressed neural responses for repeating faces (Henson, 2016; Kovacs et al., 2012), this suggests that repetition effects may not have played a significant role here. This could be because the new localiser included a number of different face viewpoints and identities (three and 20, respectively), which have been found to minimise repetition effects (Fang et al., 2007; Henson, 2016). This could also be related to there being twice the amount of data in the combination contrasts than in the single and tiled contrasts, which may have improved the detection of face-specific activation irrespective of the stimuli involved. In future studies it would be informative to disentangle the effect of stimulus type and amount of data by matching the number of experimental blocks included in each condition.

While changes in pRF number were found, localiser stimuli did not significantly impact pRF size. While there were small variations in coverage observed between the localiser stimuli, these effects were subtle and did not emerge as clearly as the differences in pRF number. Ultimately, incorporating peripherally located faces in localiser stimuli had the greatest effect on the overall number of pRFs stimulated, rather than their properties. This suggests that the different localiser approaches affected the number of neurons that were successfully stimulated, but did not significantly affect *how* they respond across the visual field.

While the localiser did not appear to significantly impact the selectivity of face-selective neurons, the differences in pRF number described above mean that the choice of localiser would still affect subsequent measurements of retinotopy. If fewer pRFs are successfully identified during retinotopic mapping due to the way that face-selective brain regions were initially identified, then the properties of the pRFs will be less accurately estimated. Essentially, my investigations of the spatial properties of face-selective brain regions in the previous chapter would have been less precise with fewer pRFs available. As I was particularly interested in the peripheral visual field, it was especially important to sufficiently stimulate the periphery, to avoid pRFs dropping out at far eccentricities. Otherwise, this would have led to inaccurate conclusions about how face-selective parts of the brain represent and sample the visual field. Specifically,

the horizontal-vertical and upper-lower differences in pRF number would have been smaller in magnitude in the previous chapter, if fewer pRFs were available for the comparisons. The present findings suggest that large face localisers with intermixed single and tiled stimuli are best for activating face-selective cortex, with benefits for subsequent analyses that investigate both foveal and peripheral space.

These findings have implications not only for the retinotopic measurements in the previous chapter but for the interpretation of previous research. Specifically, they suggest that the extent of the foveal bias within face-selective regions – which has been found to be magnified compared to early visual cortex, with even more neural resources directed to the fovea rather than the periphery (Finzi et al., 2021; Kay et al., 2015) – may be overestimated if the localiser images only presented single faces at the fovea. As the current results suggest that simply increasing the size of localiser images produces face-selective activation over a greater region of the cortical surface, this could be particularly true if the localiser images covered a smaller field of view than the stimuli in the main experiment. Because many studies do not report the size that localiser images were displayed at, this is difficult to assess, however.

As well as the effects of face position within localiser stimuli, these findings provide insight into manipulations of face size. As mentioned above, increasing the overall size of the localiser images stimulated a greater region of the cortical surface, which suggests that simply enlarging the face images goes some way towards stimulating the peripheral visual field. More interestingly, however, altering the size of the faces *within* the stimulus images may have affected delineation. While the tiled faces did activate a larger cortical region than the single faces, the most face-selective activation was found with a *combination* of single and tiled faces. This suggests that the large, single faces provided some benefit – otherwise, there should not have been any significant differences between the tiled and combination versions. However, it is also possible that the benefit stemmed from the larger amount of data available in the combined contrast, which will require further exploration.

Consistent with the idea of large faces being beneficial is recent research showing that face-selective inferotemporal (IT) neurons in monkeys responded most strongly to faces that had the same retinal size but that were perceived to be extremely

large, as 3D depth cues meant that they appeared to be further away (Khandhadia et al., 2023). Similarly, a retinotopic mapping study in humans found increased BOLD responses in ventral face-selective areas when bars (of a constant size) revealed segments of a large face, rather than containing multiple small faces (Sayres et al., 2010). While this suggests that large faces may drive face processing systems strongly, I found an increase in foveal activation when the localiser included *both* small and large faces (i.e. combination stimuli), and that separate vertices could be activated by the single and tiled localisers. This suggests that we may have populations of face-selective neurons that are tuned towards different face sizes, consistent with the finding that adaptation aftereffects for faces depend on the size of the adapting and test faces being similar (Zhao & Chubb, 2001). As the new localiser did not assess manipulations of face size separately to changes in face position, future work would be needed to disentangle these effects. Nevertheless, my findings suggest that a combination of face sizes within localiser stimuli could be beneficial in driving face-specific activation.

These findings are consistent with other studies showing that the specific image categories and contrasts used during localisation can affect the pattern and strength of activation within ventral category-selective regions (Schwarz et al., 2019; Weiner & Grill-Spector, 2010). Here, I show that the *properties* of the images within each category can also affect delineation. Some have raised concerns that the exact patterns of voxels identified within functionally defined brain regions can have low reliability over time, depending on factors such as the statistical threshold chosen (Friston et al., 2006; Poldrack, 2007; Stark et al., 2004). Using a more liberal statistical threshold can improve the reliability of regions delineated across separate sessions (Duncan et al., 2009; Duncan & Devlin, 2011). With the new localiser I used a slightly more liberal threshold of $t > 2$, compared to other studies which tend to use $t > 3$ or higher (Stigliani, 2015; Weiner & Grill-Spector, 2010, although see Gauthier et al., 2000; Poltoratski et al., 2021). I found that (within the same session) the same general pattern of activation was present within all of the face-selective areas for each of the localisation methods; the cortical area delineated was largest when single and tiled faces were combined, then tiled, then single. Although this does not address reliability over time, the consistency across brain regions indicates that functional localisers can

produce reliable activation patterns within and across participants, and are indeed useful tools for identifying category-selective brain regions (Saxe et al., 2006).

As highlighted previously, some have critiqued the theory behind functional localisation itself, arguing that separate localisers which define regions of interest based on functional selectivity do not necessarily capture the same regions that are activated during the main experiment (Friston et al., 2006). Consistent with this, my findings show that the spread of face-selective activation over the cortical surface was affected by the localisation approach. Research has demonstrated that category selective areas are not entirely discrete, but that regions of functional selectivity often overlap within ventral temporal cortex, with many voxels shown to respond to different categories – e.g. faces and body parts – to varying extents (Weiner & Grill-Spector, 2010). As an example, the new localiser may have been more effective at isolating peripheral vertices that are usually defined as body-selective – with body-selective parts of cortex found to be more peripherally biased than face-selective regions (Gregorek, 2022) – but preferentially respond to faces when they are shown in the periphery. This indeed suggests that the exact vertices included in category selective regions depends on how they have been functionally defined.

The more extreme side of Friston's (2006) argument questions the idea that the brain has reliable category selectivity at all, instead proposing that because functionality varies according to context, attempts to isolate these concrete regions may be inherently flawed. However, the position of the regions I delineated did not differ according to the localisation technique (e.g. contrasting stimuli with substantially different spatial properties). This demonstrates a degree of consistency in category selectivity irrespective of specific stimulus type, in line with at least some domain specificity within the brain (Fodor, 1983; Kanwisher, 2000). Separate localiser experiments may indeed be a quick and relatively easy way to identify such category-selective regions for further analyses, as long as the methodology is appropriate (Kanwisher, 2017; Saxe et al., 2006).

Lastly, localising face-selective parts of cortex presented an opportunity to investigate the link between category and spatial selectivity within the brain. As the pRF model is able to explain more variance in the observed BOLD responses when

the responses are spatially localised, higher R^2 values are associated with greater retinotopic sensitivity. In all three face-selective regions there was a positive relationship between face selectivity and R^2 , indicating that as selectivity for faces increased, so did sensitivity to retinotopic location. Contrary to previous reports (Silson et al., 2022), this suggests that face-selective regions may encode category and spatial information in a complementary rather than opposing manner, with more sensitivity to spatial location as category selectivity increases. This complementary encoding may occur actively within face-selective cortex, or could arise due to a more passive process, with face-selective neurons encoding more spatial information if they receive inputs from neurons in earlier visual areas that are more spatially localised (and so would have higher R^2). However, R^2 is not a direct measure of spatial selectivity, and may be confounded by other factors. For example, vertices that respond very strongly to faces would have higher beta amplitudes (due to their neurons being very visually responsive), which have been associated with higher R^2 . Further exploration will be needed in order to clarify the relationship between category and spatial selectivity within face-selective cortex.

To conclude, this chapter examined the process of localising and delineating face-selective parts of the brain. I introduced a novel functional localiser, which presented faces peripherally as well as foveally. Including peripherally located faces resulted in the delineation of larger face-selective brain regions, with a greater number of pRFs found across the visual field, which allowed for the more accurate measurement of retinotopic properties. These findings show that the localisation of functionally defined brain areas is affected not only by the choice of object categories, statistical contrasts, or overall image size, but by the spatial properties of the objects within the images. This demonstrates that while functional localisers can be a useful tool to identify category-selective brain regions, it is important to carefully design the stimuli used to localise them, considering the properties that will later be investigated within them.

Chapter 5

General Discussion

How distinct is face perception from other parts of vision, and how might specialised processing be built into the visual system? In this thesis, these overarching questions were addressed through a series of behavioural and functional magnetic resonance imaging (fMRI) experiments. Faces are widely considered to be “special”, processed using distinct mechanisms – shown by disproportionate inversion effects (McKone et al., 2007; Yin, 1969) – and with dedicated brain regions that form the “core” face recognition system (Grill-Spector et al., 2017; Kanwisher et al., 1997). Cases of prosopagnosia have supported this view, showing that face perception can be selectively impaired while other facets of vision, including object recognition, remain intact (Sergent & Signoret, 1992; Towler et al., 2017). Previous research has also suggested that while low-level vision differs predictably according to location (Abrams et al., 2012; Carrasco et al., 2001; Greenwood et al., 2017), face recognition abilities vary across the visual field in unique or idiosyncratic patterns (Afraz et al., 2010; Quek & Finkbeiner, 2016; Schmidtman et al., 2015; Visconti di Oleggio Castello et al., 2018). The extent to which high-level face perception interacts with and is built upon low-level vision therefore remained unclear. By investigating the spatial selectivity of face recognition in comparison to low-level spatial variations, I sought to examine this link more directly. The experiments in this thesis come together to shed further light on the apparent distinctness of face recognition, revealing that the mechanisms behind face perception are not as unique as previously thought, with even high-level, specialised processing showing similar spatial selectivity to low-level vision.

5.1 Overview of thesis findings

The first aim of the thesis was to measure exactly how face perception varies across the visual field. A plethora of studies have demonstrated that the perception of low-level stimuli – such as letters or lines – often follows systematic patterns, with better perception along the horizontal vs. vertical meridian and in the lower vs. upper visual field (Abrams et al., 2012; Barbot et al., 2021; Benson et al., 2021; Carrasco et

al., 2001; Greenwood et al., 2017; Himmelberg et al., 2020; Westheimer, 2003, 2005). In contrast, previous research suggests that face perception may vary across the visual field in unique or even wholly idiosyncratic patterns. Studies have reported better face perception in the upper (Quek & Finkbeiner, 2014, 2016) or left hemifield (Ellis & Shepherd, 1975; Harrison & Strother, 2021; McKone, 2004; Schmidtman et al., 2015; although see Bourne et al., 2009; Kovacs et al., 2017), with others suggesting that the perceived gender, identity and age of faces varies entirely idiosyncratically across location (Afraz et al., 2010; Visconti di Oleggio Castello et al., 2018). These distinct variations suggest a dissociation between low-level vision and face perception, leaving the link between the two unclear.




Figure 5.1. Schematic of the behavioural findings from Chapter 2. The size of the faces represents the anisotropies found for face perception, which match the systematic variation often measured for low-level vision (reflected by the size of the letters). Smaller faces and letters indicate better acuity.

whereby acuity for judging facial gender was clearly better along the horizontal vs. vertical meridian (Figure 5.1). I also found a smaller-but-consistent lower vs. upper visual field advantage for gender acuity. Therefore, the perception of face gender varied across the visual field in a systematic fashion, with patterns matching those of low-level vision. While face perception is indeed a multifaceted ability and involves

To investigate this dissociation, I sought to align the methodology used to measure variations in low- and high-level vision, by measuring the spatial resolution of face perception across the visual field. Specifically, a novel face acuity test was designed, which determined the smallest face size needed to accurately judge gender across various locations, 10° in the periphery. Using this method, I uncovered systematic variations

other judgements such as the recognition of identity and emotional expression, the findings in Chapter 2 suggest that face processing systems inherit the same patterns of spatial selectivity that are observed for low-level visual abilities (e.g. Carrasco et al., 2001).

The second aim was to use fMRI and population receptive field (pRF) modelling to investigate the neural underpinnings of face perception. In Chapter 3, I asked whether the retinotopic properties of three face-selective brain regions – OFA, pFus and mFus – could explain the systematic behavioural variations in face perception (Chapter 2), and whether these spatial properties align with, or deviate from, those of earlier brain regions (V1-V3). In all three face-selective brain regions there was a greater number of pRFs and better visual field coverage (a measure of relative responsiveness; see 3.2.13) along the horizontal vs. vertical meridian and in the lower vs. upper field, irrespective of face inversion. These patterns align with the behavioural anisotropies measured in Chapter 2, suggesting that better acuity for faces is linked to the way that face-selective neurons sample the visual field, with increased sampling of the horizontal meridian and lower field. While the retinotopy within face-selective areas varied considerably according to location, there were only small (neural) effects of inversion overall, suggesting that visual field sampling might be more similar for upright and inverted faces than previously thought (Poltoratski et al., 2021).

Similar horizontal-vertical and upper-lower variations in pRF number and coverage were found in V1-V3, highlighting commonalities in the way that the visual field is sampled by early visual cortex and face-selective parts of the brain. While receptive field size has previously been associated with perceptual performance (Duncan & Boynton, 2003; Poltoratski et al., 2021), pRF size was less reliably linked to acuity variations within the face-selective regions and in V1-V3, highlighting further similarities between low- and high-level visual cortex. Together, these results suggest that shared spatial selectivity between early vision and higher-level face processing occurs due to variations in neuronal density – and resulting coverage – that are maintained throughout the visual system. In other words, face-selective areas appear to inherit certain patterns of retinotopy from early visual cortex.

The third aim was to determine whether the measurement of the retinotopic properties described above would be affected by the spatial properties of the stimuli used to originally localise face-selective brain regions. For example, a considerably magnified foveal bias is typically observed in face-selective brain areas compared to early visual cortex, suggesting that the majority of neural resources are directed to the fovea at the expense of the periphery (Finzi et al., 2021; Kay et al., 2015). Could this exaggerated central bias be partly driven by localisation methods, which present face stimuli foveally (Stigliani, 2015; Weiner & Grill-Spector, 2010)? Previous research has shown that localisation technique can affect the delineation of category-selective brain regions, with patterns of functional activation differing depending on factors such as the specific categories of images or statistical contrasts used (Duncan et al., 2009; Duncan & Devlin, 2011; Friston et al., 2006; Weiner & Grill-Spector, 2010). I reasoned that the delineation of – and properties subsequently measured from – face-selective regions could also be affected by the spatial properties of the faces in the localiser images.

To investigate this possibility, in Chapter 4 I designed and implemented a novel functional localiser which contained either single, large (foveal only) faces, smaller faces which were tiled across the visual field (foveal *and peripherally* located faces), or a combination of the single and tiled faces. I found that all three face-selective regions identified – OFA, pFus and mFus – contained a significantly greater number of vertices and were therefore larger in size when localiser stimuli contained a combination of single and tiled faces, compared to either configuration alone. This suggests that the single and tiled faces may have been activating different sets of vertices, which were then both picked up using the combination localiser. Retinotopic analyses revealed that when both foveal and peripheral faces were included in the localiser, there were a greater *number* of pRFs found across the visual field, with an increase in peripheral pRFs particularly noticeable in OFA and mFus. These results show that the spatial properties of localiser stimuli can affect the delineation of functionally defined brain regions. Importantly, the approach used to initially identify face-selective parts of the brain can impact subsequent retinotopic measurements, such as the extent of the foveal bias, again suggesting that the retinotopy within face-selective areas may not be quite as distinct as previously thought.

5.2 Face perception varies systematically across the visual field

Chapter 2 revealed that the ability to judge facial gender varied across the visual field in a systematic fashion, with horizontal-vertical and upper-lower anisotropies that match those of low-level vision (Abrams et al., 2012; Barbot et al., 2021; Carrasco et al., 2001; Greenwood et al., 2017). This is contrary to suggestions that the spatial resolution of face perception varies uniquely (Quek & Finkbeiner, 2016; Schmidtman et al., 2015) or even entirely idiosyncratically (Afraz et al., 2010; Visconti di Oleggio Castello et al., 2018). Instead, the findings in Chapter 2 suggest that spatial selectivity is preserved throughout the visual system and inherited by higher-level vision. Chapter 3 provided a neural basis for these findings, showing that in all three face-selective brain regions there was a greater number of pRFs and better visual field coverage (a measure of relative responsiveness) along the horizontal vs. vertical meridian and in the lower vs. upper field. In other words, better gender recognition was linked to an increased number of pRFs and the resulting increases in visual field coverage. The systematic variations in face perception measured in Chapter 2 may therefore be explained by the differential sampling of the visual field within face-selective brain regions.

Chapter 3 aimed to not only determine the neural basis of the behavioural anisotropies found for face perception, but to compare retinotopic properties between early visual cortex and face-selective brain areas. Previously, research has linked better low-level acuity to small receptive fields and increased neuronal density in early visual cortex (V1-V3; Amano et al., 2009; Duncan & Boynton, 2003). Yet, studies have suggested that receptive field properties are considerably different within face-selective regions, with large receptive fields instead linked to better face recognition, and a comparatively impoverished representation of the periphery compared to the fovea (Finzi et al., 2021; Gomez et al., 2018; Kay et al., 2015; Poltoratski et al., 2021; Witthoft et al., 2016). This apparent dissociation in retinotopy left the link between low- and high-level vision unclear. In Chapter 3, however, I uncovered similar patterns for pRF number and coverage in face-selective areas and V1-V3. Rather than face-selective brain regions showing almost entirely distinct retinotopy, this highlights

commonalities in the way that the visual field is sampled by early visual cortex and face-selective parts of the brain.

5.2.1 pRF number explains visual field variations more consistently than size

Previous research has suggested that large receptive fields in face-selective areas aid face recognition, with smaller pRFs observed for inverted vs. upright faces within the FFA (Poltoratski et al., 2021). If there were similar variations in pRF size according to location, I should have found larger pRFs in regions of the visual field that had better acuity for face recognition. However, Chapter 3 revealed that apart from larger pRFs in the lower vs. upper field in OFA, pRF size did not significantly differ between visual field locations in the face-selective areas. pRF size was also similar between upright and inverted faces in the FFA (pFus and mFus), despite a clear behavioural face inversion effect where gender recognition was better for upright than inverted faces. These findings suggest that receptive field size is not reliably linked to acuity for face perception. In line with this, my results show that pRF size increased much less with eccentricity in the FFA compared to V1-V3, which would be consistent with a weaker link to acuity (which declines in the periphery; Rosenholtz, 2016). To determine whether these shallow pRF size-eccentricity slopes bear any relation to behavioural performance, future experiments could measure how gender acuity varies across eccentricity. Together, the current findings suggest that receptive field size is not the driving factor behind face perception abilities.

Further evidence for the importance of pRF number rather than size for face perception comes from examining their respective contributions to overall visual field coverage. Near the fovea, the horizontal-vertical and upper-lower differences in pRF number were generally more extreme within the face-selective regions compared to early visual cortex. However, the anisotropies for visual field coverage were *less* extreme in the face-selective areas compared to V1-V3. This suggests that receptive field size may have had a greater contribution to visual field coverage in early visual cortex, while coverage in the face-selective regions was predominantly defined by pRF number. This may again point to the dominance of receptive field density in defining coverage – and as a result, acuity – within face-selective areas. As receptive fields increase considerably in size further up in the visual hierarchy (Dumoulin & Wandell,

2008; Kay et al., 2015; Winawer et al., 2010), their size may not drive variations in coverage quite as strongly.

The functional localiser results in Chapter 4 further highlight the importance of receptive field number within face-selective brain regions. Manipulating the spatial properties of localiser stimuli resulted in changes in the size of face-selective areas identified on the cortical surface. Analysing the retinotopic properties of the differently defined regions revealed that these increases in size (of the regions delineated) were linked to changes in pRF number – but not size or visual field coverage⁸ – within OFA, pFus and mFus. Using localiser stimuli which included a combination of single, large (foveal only) faces and smaller, tiled (more peripheral) faces resulted in a greater density of pRFs found in all three face-selective regions. Essentially, the representation of the central *and* peripheral visual field was enhanced within face-selective regions through an increase in receptive field number, rather than changes in pRF properties such as size. Again, this points to receptive field number being the driving factor behind responses within face-selective parts of the brain.

Even in V1-V3, the link between pRF size and typical visual performance according to angular location (Carrasco et al., 2001) was variable. While other studies report smaller pRFs along the horizontal vs. vertical meridian in all three areas (Silva et al., 2018), I only found this pattern in V3, with no difference between the meridians in V1-V2. There was a consistent although opposite-to-expected pattern for the upper-lower difference, however, with larger pRFs in the lower vs. upper field across V1-V3. This suggests that even within early visual cortex, pRF number and resulting coverage estimates are more reliably linked to variations in visual ability than estimates of size. These findings tie in with research showing that estimates of pRF size can be more variable than other parameters, such as position (Lage-Castellanos et al., 2020). Studies have indeed shown that pRFs are more prone to changes in size depending on various factors, such as stimulus properties and visibility (Alvarez et al., 2015; Hughes et al., 2019), participant attention (Kay et al., 2015), and fitting procedures (Lerma-Usabiaga et al., 2020). As such, the results in Chapter 3 add to recent debate

⁸ Even when the coverage maps were jointly normalised, there were minimal differences in coverage between the localisers (Figure B.1).

around how we make conclusions about underlying changes in visual function based on pRF estimates, particularly size (Dumoulin & Knapen, 2018; Stoll et al., 2022).

What can be concluded about face processing mechanisms from the above findings? Altogether, I uncovered better coverage in some parts of the visual field – driven by an increased number of pRFs – which can explain variations in our ability to perceive face gender across different locations. This is still generally consistent with prior work (Poltoratski et al., 2021), indicating that variations in the way that face-selective brain regions sample the visual field are linked to face perception abilities. However, my findings suggest that these variations in sampling are primarily linked to differences in the *number* of face-selective neurons – or similar amounts that respond more strongly, which could also have resulted in a greater number of pRFs remaining during analyses⁹ – rather than changes in their properties, such as size.

5.2.2 Could acuity for face perception be linked to spatial integration?

Some researchers have argued that large receptive fields are beneficial for face recognition because they enable more efficient spatial integration, which aids holistic processing (Poltoratski et al., 2021; Witthoft et al., 2016). These studies report smaller pRFs in the FFA for inverted vs. upright faces (Poltoratski et al., 2021), and in the face-selective brain regions of individuals with developmental prosopagnosia (Witthoft et al., 2016). Yet, as discussed above, I did not find a strong link between pRF size and the ability to judge facial gender. Could the large receptive fields of face-selective neurons still serve another purpose, even though size itself seems to matter less than in earlier regions? Again, this links to the importance of receptive field density for face perception. As we usually fixate faces to recognise them (de Haas et al., 2019), it is desirable to have a large proportion of receptive fields sampling the fovea, demonstrated by the exaggerated central bias of face-selective parts of the brain (Finzi et al., 2021; Gomez et al., 2018; Kay et al., 2015; Poltoratski et al., 2021). The large receptive fields of face-selective areas may serve an important functional purpose by enabling as many receptive fields as possible to be centred near the fovea, while

⁹ If the neurons within a pRF respond more strongly, this increases the likelihood that it will remain after noise and goodness-of-fit filtering, due to improved signal-to-noise and parameter fitting, respectively.

simultaneously extending into the periphery. This way, face perception can still occur in peripheral vision (Kovacs et al., 2017; McKone, 2004; Roux-Sibilon et al., 2023; and as shown in Chapter 2), but the most neural resources – i.e. greatest amount of receptive fields, which Chapter 3 showed to be the defining factor driving acuity for faces – are directed towards the fovea. Indeed, visual field coverage in the face-selective regions was found to extend to a similar region in the periphery as V1-V3. The large receptive fields of face-selective neurons may be important not because their large size directly aids spatial integration but so that they can cluster within the central part of the visual field, which is most critical for face perception, while still providing coverage in the periphery.

Could enhanced spatial integration lead to better face perception due to increases in receptive field number instead? This could explain the variations in spatial selectivity, for example with an increased number of face-selective neurons along the horizontal meridian (Chapter 3) supporting more efficient spatial integration of face information in this location, resulting in better gender acuity (Chapter 2). However, it is less clear whether this explanation could be applied to featural selectivity, which Poltoratski et al.'s (2021) arguments were based upon. If differences in spatial integration were the main driver behind face perception abilities, there should have been clear differences in pRF measures according to inversion (e.g. main effects), in line with the clear behavioural inversion effects observed. There should also have been larger (behavioural and neural) inversion effects in regions with better gender acuity, if variations in spatial integration due to visual field sampling played a primary role in holistic processing.

Considering the above, an interesting avenue of research could be to compare retinotopic properties within the FFA when stimuli only include face features, in comparison to whole faces where holistic processing would be involved. (Research has found the FFA to be similarly sensitive to whole faces and face features (Liu et al., 2010; Yovel & Kanwisher, 2004), suggesting that signal-to-noise variations should not be a confounding factor). If more pRFs were found in response to the whole faces at certain locations – but this number then drops if the whole faces are inverted, as configural processing is disrupted – this could support the idea that better acuity is

achieved because an increased number of receptive fields leads to better spatial integration, rather than due to variations in properties such as their size.

5.2.3 The horizontal-vertical and upper-lower anisotropies differ in magnitude

While there was a clear horizontal-vertical anisotropy present across all three (behavioural) gender acuity experiments in Chapter 2, the magnitude of the upper-lower difference was smaller and more variable. The lower field advantage appeared to be partly driven by the position of the eyes within face stimuli in Experiment 1¹⁰, when faces were centred on the nose. This meant that they appeared further away from fixation in the upper compared to the lower field for upright faces, and vice versa for inverted faces. When this factor was controlled for in Experiment 2 by ensuring that the eyes always appeared at an equal distance from fixation, the upper-lower difference was smaller in magnitude and no longer significant. This confirms the importance of the eyes for judging gender (Brown & Perrett, 1993; Schyns et al., 2002; Yamaguchi et al., 2013). Due to there being a present but not significant upper-lower difference in Experiment 2, a further experiment was carried out where I focused on the upper and lower locations only, with a greater number of trials (Experiment 3). A clear upper-lower difference was subsequently found, confirming a lower field bias for face perception. The difference in magnitude between the horizontal-vertical and upper-lower anisotropies is consistent with other research which reports that for low-level vision, the upper-lower difference is smaller – and therefore harder to reliably measure – than the horizontal-vertical difference (Barbot et al., 2021; Kurzawski et al., 2021).

The neuroimaging findings in Chapters 3 and 4 provide a neural explanation of why the behavioural horizontal-vertical and upper-lower anisotropies differed in magnitude. Within the face-selective regions, the horizontal-vertical differences in pRF number were larger than the upper-lower differences. This likely reflects the much greater number of pRFs along the horizontal than the vertical meridian overall, meaning that the anisotropy between the two meridians was more exaggerated than between the upper and lower locations, which comprised the (comparatively sparser)

¹⁰ Experiments 1, 2 and 3 refer to the three gender acuity experiments in Chapter 2.

vertical meridian. Similarly, in Chapter 4 the additional pRFs found in response to certain localiser stimuli were mainly located along the horizontal rather than the vertical meridian. These findings speak to the strikingly different representation of the two meridians within the face-selective regions, with considerably fewer pRFs located along the vertical vs. the horizontal meridian. The behavioural upper-lower difference in face perception may therefore be less pronounced than the horizontal-vertical difference because they are driven by underlying variations in pRF number that differ in magnitude.

Interestingly, the horizontal-vertical anisotropy in pRF number became more exaggerated further along the ventral stream, from OFA to mFus. This suggests that as brain regions become more specialised for face processing, the representation of the horizontal meridian becomes increasingly prioritised over the vertical. Could this be related to face processing systems inheriting certain types of featural selectivity from lower areas? Horizontal spatial frequencies have been shown to be especially important for different aspects of face perception, with the disruption of horizontal information detrimental to identity recognition and holistic processing, for example (Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux et al., 2016). The radial bias describes the generally better perception of radially – i.e. pointing towards the fovea – oriented information within the visual system (Rovamo et al., 1982; Sasaki et al., 2006). A similar radial bias has been found for recognising face identity, with larger inversion effects – indicating greater disruption of face-specific mechanisms – along the horizontal vs. vertical meridian (Roux-Sibilon et al., 2023). The increasingly asymmetric representation of the horizontal and vertical meridians within more anterior face-selective parts of the brain suggests that the face-selective regions may develop mechanisms to improve sampling within locations where the most diagnostic facial information will be best processed (i.e. the horizontal meridian). In other words, they may *enhance* the radial bias inherited from earlier regions. This would point to a more active process behind how face-selective neurons differentially sample the visual field (although a radial bias in face recognition could also result from the passive pooling of information from lower areas, with successive pooling magnifying the biases within visual information).

Could the increased representation of the horizontal meridian as brain regions become more specialised for face processing be linked to visual experience? The recognition of individual face features has been shown to be tuned towards their typical visual field locations, with eyes better recognised in the upper field and mouths in the lower (de Haas et al., 2016). This could arise from a “faciotopy map” within OFA, where face-selective neurons that have preferences for certain face features are organised topologically, according to the typical configuration of a face (Henriksson et al., 2015). Along similar lines, retinotopic organisation within the FFA could be tuned towards locations where faces are more likely to appear within our visual field. Within everyday life, we may be exposed to more faces around eye-level (i.e. along the horizontal meridian) than significantly above or below, which could lead to a greater allocation of neural resources along the horizontal meridian. A way to investigate this indirectly could be to examine the retinotopic organisation of word-selective brain regions in English vs. Chinese speakers. Word-selective brain areas have been shown to exhibit increased coverage of the horizontal vs. vertical meridian in native English speakers (Le et al., 2017). Would this asymmetry be reduced within native Chinese speakers, as Chinese characters and words are represented less horizontally than English words, and would any reductions in retinotopic asymmetry be specific to word processing regions? This could shed further light on how visual experience shapes functional specialisation and retinotopy within category-selective brain regions.

5.2.4 Spatial selectivity follows similar patterns despite inversion

Chapter 2 showed that spatial variations in gender acuity were similar regardless of face inversion, with the behavioural horizontal-vertical and upper-lower anisotropies present for both upright and inverted faces. In other words, although there was worse performance overall for inverted faces, the visual field anisotropies were similar in magnitude for upright and inverted faces. This indicates that the *spatial* processing of faces was similar despite the *featural* change introduced by inversion. While another study found larger inversion effects along the horizontal vs. vertical meridian (Roux-Sibilon et al., 2023), this involved judgements of identity, which have been shown to be particularly reliant on horizontal information (Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux & Greenwood, 2016). As horizontal information is

processed better along the horizontal meridian throughout the visual cortex, due to the radial bias (Sasaki et al., 2006), this variation in inversion effects could stem from non-face-specific featural mechanisms. Using a gender judgement could have reduced susceptibility to these radial effects, allowing patterns of spatial selectivity to emerge over featural selectivity.

In line with this, Chapter 3 showed that in general, the horizontal-vertical and upper-lower differences in pRF number and coverage were similar for upright and inverted faces. This provides a neural explanation of the similar behavioural anisotropies measured for upright and inverted face perception (Chapter 2), and shows that the variations in sampling inherited by face processing systems are associated with the spatial location of faces, rather than their featural content. There could be an active process behind the inheritance, with face-selective parts of the brain applying their inherited spatial selectivity similarly, irrespective of face orientation. On the other hand, this could occur more passively, with variations in face representations – resulting in measurable differences in receptive field number and coverage – arising from differences in the amount and/or quality of information received from lower brain regions (Riesenhuber & Poggio, 1999). Either way, similar patterns of spatial selectivity – resembling those of low-level vision – were observed behaviourally and neurally for upright and inverted faces. This highlights a link between spatial and featural selectivity within face perception.

5.2.5 Shared patterns of retinotopy reveal links between face processing and other visual abilities

Retinotopy within face-selective regions has often been considered to be distinct from low-level vision. That is not entirely the case, with the apparently unique sampling characteristics of face-selective areas formed from an exaggeration of the retinotopy found within the visual system in general, which includes the retina, lateral geniculate nucleus (LGN) and early cortical regions. This is an important distinction to make; instead of face-selective regions exhibiting distinct retinotopy, their approach to sampling is built upon sampling within other brain regions. The most obvious example is the foveal bias that is often thought of as characteristic of face-selective areas, and an example of a difference in processing. However, cortical magnification is a common

property found throughout the visual cortex (Dekker et al., 2019; Duncan & Boynton, 2003; Harvey & Dumoulin, 2011) and even as early as the retina, with a higher concentration of retinal ganglion cells sampling the fovea compared to the periphery (Curcio & Allen, 1990). The central bias of the visual system is therefore *magnified* within face-selective parts of the brain.

Crucially, research suggests that the central bias is exaggerated to a similar extent within parts of ventral occipitotemporal cortex selective for words and letters, showing that this magnification is not unique to faces (Le et al., 2017). The same holds for other receptive field properties such as size. Receptive fields increase in size further up in the visual hierarchy (Dumoulin & Wandell, 2008; Winawer et al., 2010), becoming very large in brain regions selective not just for faces but for bodies, places, and objects (Gregorek, 2022; Silson et al., 2015). In particular, research suggests that receptive fields are similarly large in the visual word form area, where there was a similar foveal bias to face-selective regions (Le et al., 2017). In this sense, then, how special really is the retinotopy within face-selective areas? Together with previous findings, the results in this thesis suggest that the retinotopy is not so unique after all.

These shared patterns of retinotopy add further context to a modular view of face processing. The processing of faces has often been considered to be relatively distinct from other aspects of vision, and carried out in separable streams to other objects. The experiments in this thesis highlight that the “core” face recognition system shares similar spatial properties to earlier visual areas, which is consistent with the view that faces are not processed in a strictly modular sense, but that various distributed brain regions contribute to face perception (Haxby et al., 2000). For example, this could include parietal regions involved in directing spatial attention or inferring eye gaze, or limbic regions which help process facial emotion. Retinotopic maps have been found in regions of parietal and frontal cortex that are involved in more general cognitive processes (e.g. spatial attention and eye movements). Common principles of spatial encoding therefore appear to be present throughout the brain (Groen et al., 2022). The presence of multiple spatial maps is consistent with the broader idea that while there may be face-selective modules within the brain, these modules exist within a more distributed network of brain regions that work together to

produce our impressive face recognition abilities (Grill-Spector et al., 2017; Haxby et al., 2000).

5.3 Featural vs. spatial selectivity within face processing systems

5.3.1 Inversion effects indicate configural processing across the visual field

Behavioural face inversion effects were found in all three gender acuity experiments in Chapter 2, and during pRF mapping in Chapter 3, where participants were required to judge the gender of bars containing male or female faces. In both paradigms the ability to judge gender was worse for inverted compared to upright faces, suggesting that the tasks engaged face-specific mechanisms, with an added benefit of configural processing for upright compared to inverted faces (Le Grand et al., 2001; Rossion, 2008). This is consistent with other studies showing that holistic face processing occurs in peripheral as well as central vision (Kovacs et al., 2017; McKone, 2004; Roux-Sibilon et al., 2023). The gender acuity experiments highlight a spatial component to the face inversion effect, where the recognition of upright and inverted faces could be matched by simply altering face size. Presumably, the benefit of configural processing meant that the gender of smaller upright faces could be judged, while inverted faces had to be larger to achieve the same recognition accuracy. These behavioural inversion effects suggest that configural mechanisms were engaged across the visual field.

What are the cognitive mechanisms behind the differences in performance for upright and inverted faces? These inversion effects could reflect a qualitative change in face processing, with configural and feature-based processing employed for upright and inverted faces, respectively (Rossion & Gauthier, 2002). Alternatively, these results could stem from a quantitative change, where inversion effects represent a difference in how efficiently configural processes are engaged, rather than qualitatively different mechanisms (Sekuler et al., 2004). In Chapter 2, the psychometric functions used to determine gender acuity thresholds were generally similar in shape and simply shifted towards larger face sizes for the inverted faces. This is certainly consistent with a quantitative explanation of inversion effects, with increases in the engagement of configural processes as inverted faces became larger. However, it could reflect a

qualitative mechanism, with feature-based processing improving simply because increases in size meant that the features of the inverted faces were more visible. To determine which explanation is more likely, it would be desirable to modify the face acuity test to include configural manipulations, so that we could directly assess how configural processes would affect acuity variations for inverted faces. This would provide further clarification on the mechanisms that underlie face inversion effects, which remain hotly debated in the literature.

As described earlier, similar spatial variations were found behaviourally (Chapter 2) and neurally (Chapter 3) for upright and inverted faces. This emphasises the link between spatial and featural selectivity within face processing systems. How closely interlinked might these mechanisms be? As will be discussed over the following sections, the retinotopic mapping in Chapter 3 also uncovered distinctions between the two types of selectivity, consistent with the view that category and location information may be encoded independently – although not necessarily separately – within face processing systems (Schwarzlose et al., 2008).

5.3.2 Retinotopy is less clearly linked to featural selectivity

The findings discussed so far suggest that the retinotopy within face-selective regions shows similar patterns of spatial selectivity, irrespective of face inversion. Did the retinotopic properties of face-selective areas differ according to the featural content of faces more generally – that is, when upright and inverted faces were compared directly, rather than across location? I found some evidence for this in mFus, with an increased number of pRFs in the periphery for upright faces, and better visual field coverage (increased relative responsiveness) at the fovea for upright faces. This suggests that within face-selective areas, there is increased sampling of certain regions of the visual field when faces are upright rather than inverted. As touched on previously, however, these sampling variations were smaller than anticipated – with no significant main effects of inversion across any of the three pRF measures – given the clear behavioural inversion effect (in the pRF experiment). They are also smaller than expected given that in Chapter 2, the variations in gender acuity thresholds between upright and inverted faces were of a similar magnitude to the variations across location. The lack of main effects (of inversion) during retinotopic mapping occurred despite

there being a greater amount of data available for the inversion compared to the location comparisons, suggesting that the inversion variations were much smaller than the location differences, which emerged clearly despite the more limited data.

As with the location differences discussed earlier, pRF size did not reliably differ between upright and inverted faces. This highlights a discrepancy with previous research, which found smaller pRFs for inverted faces in the FFA (Poltoratski et al., 2021). However, Poltoratski et al.'s (2021) findings could be linked to the reduced beta amplitudes – a measure of blood oxygen dependent (BOLD) signal – that were observed for inverted compared to upright faces (as has been similarly observed in other studies; Kanwisher et al., 1998; Yovel & Kanwisher, 2004). Reduced beta amplitudes are typically associated with poorer pRF fits (Anderson et al., 2017; Schwarzkopf et al., 2014). This would affect the accuracy of size estimates, with poorer fits having been linked specifically to smaller pRF sizes (Alvarez et al., 2015). Whether or not this is the case, the lack of consistency between findings suggests that pRF size provides a less reliable explanation of perceptual performance than other retinotopic measurements (Lage-Castellanos et al., 2020), this time being less able to explain inversion effects as well as the visual field anisotropies. I also found that pRFs were most similar in size in central vision, despite behavioural inversion effects having been found consistently for centrally presented faces (Kovacs et al., 2017; McKone, 2004; McKone et al., 2007; Robbins & McKone, 2007; Rossion, 2008; Roux-Sibilon et al., 2023; Yin, 1969). If inversion effects arise because variations in pRF size affect holistic processing (Poltoratski et al., 2021), I should have found clear differences in size estimates in the fovea. Instead, this suggests that pRF size may not reliably drive variations in upright and inverted face perception.

The smaller-than-expected variations in coverage between upright and inverted faces (Chapter 3) highlight another difference to previous research, where larger increases in coverage were found for upright vs. inverted faces in the FFA (Poltoratski et al., 2021). This could of course be related to the issues described above, as smaller pRFs for the inverted faces would have resulted in lower estimates of coverage. There are also many different ways to estimate visual field coverage, and I chose to do so using a method that would account for the Gaussian shape of the receptive fields, as

a better estimate of responsiveness across location. As I normalised the upright and inverted coverage maps separately, this approach represents *relative* changes in sampling across the visual field. Coverage patterns were similar even when the maps were jointly normalised, suggesting that the absolute coverage values did not differ considerably for upright and inverted faces (Figure A.6). However, it is possible that a binary approach to calculating coverage, which accounts for the position and size of pRFs but does not factor in changes in their response profiles (Poltoratski et al., 2021), would have yielded larger differences in (absolute) coverage.

Could variations in neuronal sampling provide any clarification on the qualitative vs. quantitative argument? In mFus, I found an increased number of pRFs at more peripheral eccentricities for upright faces, and increased visual field coverage (responsiveness) at the fovea for upright faces. This suggests a quantitative difference in face processing, with face-selective parts of the brain sampling the visual field with a greater number of neurons – or similar amounts that respond more strongly, which could also result in a greater number of pRFs identified – when faces are upright rather than inverted. However, because these effects were smaller than expected given the clear behavioural inversion effect, and the upright vs. inverted coverage differences were only significant at the fovea, this suggests that quantitative sampling differences may not be the full story.

While there were small effects of inversion on pRF number and coverage, there were no reliable differences in pRF size within the FFA. This opposes previous research which found smaller pRFs for inverted vs. upright faces in the FFA (Poltoratski et al., 2021). As discussed previously, the size differences found in Poltoratski et al.'s (2021) study could have been linked to the reduced beta amplitudes – a measure of BOLD signal – that were observed for inverted compared to upright faces (as has been similarly observed in other studies; Kanwisher et al., 1998; Yovel & Kanwisher, 2004). Reduced BOLD signal in other studies could point to a quantitative change in a different way, with face-specific mechanisms less efficiently engaged for inverted faces simply because they do not stimulate face-selective neurons as effectively. In my study, beta amplitudes in mFus were similar for upright and inverted faces, so any differences – or lack of them – in pRF measures are less

likely to stem from variations in BOLD signal. Instead, the results from Chapter 3 suggest that if quantitative differences in neural processing play a role in behavioural face inversion effects, this may involve variations in the overall amount of pRFs stimulated, rather than properties such as their size (which may actually be more reflective of a qualitative difference, as variations in pRF size indicate changes in *how* neurons respond).

Whether they represent a quantitative or qualitative change in face processing, Chapter 3 demonstrates that there were smaller-than-expected neural inversion effects within face-selective brain regions, that reflect differences in overall pRF number rather than their spatial properties (pRF size). Only small neural differences were found despite the behavioural inversion effect showing that face-specific mechanisms were employed (and there being clear face inversion effects in all three gender acuity experiments in Chapter 2). Crucially, this suggests that although retinotopic properties may explain some of the differences in perception between upright and inverted faces – e.g. with more pRFs sampling the periphery for upright faces – they may not be the main factor driving configural processing. What, then, drives face-specific mechanisms, and what role does retinotopy play within these category-selective processes?

5.3.3 Can retinotopy explain what makes face processing special?

The configural processing of faces relies on featural selectivity, but how much does this depend on retinotopic location? Horizontal orientations are particularly important for different aspects of face perception, including identity recognition and holistic processing (Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux et al., 2016), and are preferentially processed within the FFA. Larger face inversion effects have previously been found along the horizontal vs. vertical meridian (Roux-Sibilon et al., 2023), which likely reflects horizontal face information being processed better along the horizontal meridian due to an inherited radial bias from earlier visual cortex (Rovamo et al., 1982; Sasaki et al., 2006). This indicates that featural selectivity can be associated with retinotopy (as was shown in Chapters 2 and 3, with similar patterns of spatial selectivity for upright and inverted faces). However, the findings described above suggest that face-selective regions exhibit general biases for certain types of

information, that do not inherently depend on location (although location may play a part). This would explain why the retinotopic properties measured in Chapter 3 showed clear variations across location, with this spatial selectivity following similar patterns for upright and inverted faces – yet at the same time, there were relatively small effects of inversion on retinotopy overall (i.e. little featural selectivity).

This distinction between featural and spatial encoding aligns with previous research which suggested that within category-selective brain regions, featural and spatial selectivity may be encoded independently – although not necessarily separately, with the same populations of neurons found to contain both category and location information (Carlson et al., 2011; Kravitz et al., 2010; Schwarzlose et al., 2008). Previously, separate and opposing gradients of featural (category bias) and spatial selectivity (bias for the contralateral visual field) have been found within face- and place-selective brain regions (Silson et al., 2022). The findings in this thesis tie in with previous research and suggest that while the mechanisms underlying featural and spatial variations are clearly interlinked, they may also be partially separable.

Chapter 4 presented an opportunity to investigate the link between spatial and featural selectivity further, by examining the relationship between selectivity for faces (greater responses to faces compared to other objects in general) and variance explained (R^2) by the pRF model. As R^2 quantifies how well a spatially localised pRF model can explain the observed BOLD responses, higher R^2 values are generally associated with greater spatial sensitivity. Interestingly, there was a positive relationship between face selectivity and R^2 in all three face-selective regions. This indicates that as selectivity for faces increased, so did sensitivity to retinotopic location. Therefore, contrary to the findings mentioned above (Silson et al., 2022), category and spatial selectivity may be encoded within face-selective cortex in a complimentary rather than competitive manner. In other words, face-selective neurons may encode both featural and spatial information (as shown by Carlson et al., 2011; Schwarzlose et al., 2008), without the two types of information being detrimental to each other.

While differences in spatial selectivity did not drive inversion effects – with similar variations across location found for upright and inverted faces in Chapters 2 and 3 – could the unique processing of faces be explained by other aspects of

retinotopy within face-selective areas? A magnified foveal bias (compared to early visual cortex) is considered an important functional specialisation of face-selective cortex, and has been associated with better face recognition (Gomez et al., 2018; Malach et al., 2002). Research has found a similarly strong central bias within word-selective areas, demonstrating that this retinotopic property is not unique to face processing (Gomez et al., 2018; Le et al., 2017). Yet, even though word- and face-selective parts of the brain may have similar retinotopic properties, faces still show larger inversion and part-whole effects compared to words (Farah et al., 1998; Ge et al., 2006). This suggests that while the retinotopic organisation of face-selective brain regions may certainly enable specialised face processing, with increased neural resources directed towards more critical locations (the fovea), the retinotopy itself cannot fully account for the differences in upright and inverted face perception.

What else could be driving the differences between upright and inverted face recognition, if variations in sampling are not the full story? Instead of holistic processing and inversion effects arising from the way that the visual field is sampled, they could stem from alternative mechanisms that shape how the information is *used*. Interestingly, fMRI research has shown that there was a greater degree of functional connectivity within regions of the face processing network – which included core (e.g. FFA) and extended (e.g. STS) regions – when faces were upright, rather than inverted (Wang et al., 2016). Similarly, others have found increased functional connectivity between the FFA and regions of parietal cortex implicated in spatial processing when faces were manipulated configurally, as opposed to having their features replaced (Zachariou et al., 2017). These findings suggest that rather than inversion effects arising primarily from the way that individual face-selective regions sample the visual field, the efficiency of configural mechanisms could be linked to how well face-selective areas *communicate* the information they sample. Configural mechanisms may rely not only on processes contained within the FFA, but on feedforward and/or feedback connections with the distributed face perception network, which includes brain regions involved in more general cognitive abilities (Grill-Spector et al., 2017; Haxby et al., 2000). More work is needed to uncover exactly how the specialised processing of faces is built into the wider visual system, and the specific role that featural selectivity plays within this.

5.4 Localisation technique can affect retinotopic measurements

Examining the retinotopy of category-selective brain regions can be a useful tool in determining how visual field sampling may be similar to, or differ from, other facets of vision. However, Chapter 4 showed that the way that category-selective brain areas are functionally defined can affect the subsequent measurement of their retinotopic properties. Presenting localiser stimuli that contained peripherally as well as foveally located faces (i.e. the combination and tiled versions) resulted in a greater number of pRFs found across the visual field in all three face-selective regions, compared to when only foveal faces were shown. In mFus, there was a trend towards larger increases in pRF number at peripheral eccentricities, suggesting that the periphery may have benefitted most (likely linked to the comparatively few peripheral pRFs overall). Many of the additional pRFs were located at peripheral eccentricities along the horizontal meridian, indicating that the localisation technique would have affected the magnitude of the horizontal-vertical differences in pRF number measured in Chapter 3. These findings have important implications for pRF mapping, showing that the spatial properties of the localiser images affected retinotopic measurements across the visual field. Furthermore, without sufficiently stimulating the peripheral visual field during localisation, the properties of peripheral pRFs may be particularly susceptible to inaccuracies.

The inclusion of peripheral faces in the localiser (i.e. the combination and tiled versions) resulted in a greater number of pRFs found near the fovea as well as in the periphery, relative to when only single faces were used. This added boost at central eccentricities could be the result of a greater quantity of faces providing increased activation within face-selective areas (i.e. nine faces per tiled image, compared to only one per single image). At the same time, the greatest benefit was observed for the combination stimuli, where there were more peripherally *and* centrally located pRFs. This suggests that the single and tiled faces may have activated different populations of neurons, with the single, large faces conferring their own kind of benefit. This could potentially be linked to the large size of the single faces, with very large faces shown to increase BOLD responses in the FFA compared to smaller ones (Sayres et al., 2010). As such, the results from Chapter 4 suggest that using a large localiser with a

variety of face sizes and tiling conditions could maximally activate face-selective cortex, and subsequently improve the accuracy of retinotopic measurements across the central and peripheral visual field.

This finding that pRF number varied across the visual field depending on the localisation method has implications for the interpretation of previous research. In particular, the extent of the foveal bias within face-selective regions – which has been found to be magnified compared to early visual cortex, with even more neural resources directed to the fovea rather than the periphery (Finzi et al., 2021; Kay et al., 2015) – may be overestimated if the localiser images only presented single faces at the fovea. As I found that simply increasing the size of the images could go some way towards stimulating the peripheral visual field – without manipulating the size of the faces within the images – this could be particularly true if studies displayed localiser images at a smaller size than the images in a separate paradigm which then investigates the properties of the functionally defined regions. This could involve pRF mapping (e.g. Finzi et al., 2021; Poltoratski et al., 2021), as well as other approaches used to make conclusions about central vs. peripheral face processing (e.g. Kamps et al., 2020; Levy et al., 2001; Schwarzlose et al., 2008). All in all, the results of Chapter 4 highlight that before examining the characteristics of functionally defined brain regions, the methodology used to localise them should be carefully considered.

5.5 Conclusions

The central aim of this thesis was to investigate the “special” nature of face perception. Using a novel face acuity test, I found that our ability to recognise the gender of faces varies across location in a similar manner to more basic aspects of vision, with better spatial resolution along the horizontal vs. vertical meridian and in the lower vs. upper visual field. These behavioural variations were matched by enhanced sampling of certain regions of the visual field – driven predominantly by increased pRF numbers – within face-selective parts of the brain. This indicates that variations in facial gender perception are driven by variations in neuronal sampling within face-selective brain regions. Importantly, these sampling variations resembled those of early visual cortex. As such, these behavioural and neuroimaging findings come together to

highlight shared spatial selectivity between low- and high-level visual processing. This is consistent with a hierarchical model of vision, where complex, high-level face representations are built upon the information received from lower stages (Riesenhuber & Poggio, 1999). Crucially, spatial properties appear to be preserved throughout the visual system and inherited by face processing systems.

Certain hallmarks of retinotopy were magnified within face-selective regions, such as the cortical magnification of the fovea, which was exaggerated even when face-selective parts of the brain were localised using peripherally as well as foveally located faces. While not unique to faces, with word-selective areas showing a similarly strong foveal bias, this highlights subtle variations in visual field sampling compared to low-level vision, within the inherently similar patterns. Certain principles of retinotopy within face-selective areas may therefore support face processing in general, with more neural resources directed towards the fovea as we usually fixate faces to recognise them. Similar behavioural and neural anisotropies were found across the visual field for upright and inverted faces, suggesting that spatial and featural selectivity for faces are linked. At the same time, the smaller-than-expected neural inversion effects despite clear behavioural inversion effects points to a separation between spatial and featural selectivity, suggesting that there is still a missing piece of the puzzle in terms of how the relatively unique – i.e. configural – processing of faces occurs. The specialised processing of faces could arise due to the way that face processing neurons *use* the information that they sample, rather than *how* they sample the information. Altogether, the findings in this thesis show that although the retinotopic characteristics of face-selective regions may certainly subserve specialised face processing – e.g. with an enhanced foveal bias – the retinotopy itself is not so unique after all. Instead, even high-level, category-selective parts of the brain adopt certain universal principles of the visual system, and show similar sampling variations across location.

Appendices

Appendix A Supplementary Information for Chapter 3




Figure A.1. Comparison of the variance explained (R^2) by the Gaussian and compressive spatial summation (CSS) models, in one participant. ROIs are displayed in different colours (see legend). R^2 values are higher for the CSS model, particularly for OFA, pFus and mFus. The black line indicates equal variance between the models.




Figure A.2. Median exponent values from the compressive spatial summation (CSS) pRF model for upright (red) and inverted (blue) faces in each of the ROIs. Error bars represent the SEM. The dashed line represents linear summation, with all values < 1 indicating compression. Exponent values are lower (indicating increased compression) in face-selective regions compared to V1-V3, as in Kay et al. (2013). The asterisk denotes a statistically significant difference ($p < .05$).




Figure A.3. Mean estimates of pRF size in one participant, for the Gaussian (top) and compressive spatial summation (CSS; bottom) models, for upright (red) and inverted (blue) faces. The Gaussian model estimates pRF size as σ (the standard deviation of the pRF, in degrees of visual angle), while the CSS model defines pRF size as $\sigma/\sqrt[n]{n}$, with n being the exponent of the compressive spatial summation parameter (Kay et al. 2013). Note the different axis scales between V1-V3 and the face-selective regions.




Figure A.4. Mean beta amplitudes **(A)** and R^2 values **(B)** for upright (red) and inverted (blue) faces in all ROIs. Asterisks denote statistically significant differences ($p < .05$).




Figure A.5. Behavioural results from the fixation task in the pRF experiment, showing the percentage of fixation cross colour changes that were correctly identified. Dots show individual data, with lines joining each participant's performance for upright and inverted faces. Four participants had 100% correct in both the upright and inverted runs.

Fixed factors	β	p	CI
Intercept	-2.95		
ROI	1.65	< .001	1.45, 1.85
Inversion	-0.01	.960	-0.54, 0.51
Eccentricity	0.17	< .001	0.12, 0.21

Table A.1. Linear mixed model results comparing pRF size across ROI, inversion and eccentricity. Bold text indicates statistical significance ($p < .05$).

Factors	df	F	p
ROI	5, 900	146.24	< .001
Inversion	1, 900	0.34	.577
Eccentricity	20, 900	27.01	< .001
ROI*Inversion	5, 900	1.10	.373
ROI*Eccentricity	100, 900	5.38	< .001
Inversion*Eccentricity	20, 900	2.81	< .001
ROI*Inversion*Eccentricity	100, 900	1.05	.362

Table A.2. ANOVA results comparing pRF number across ROI, inversion and eccentricity. Bold text indicates statistical significance ($p < .05$).

Factors	df	F	p
ROI	5, 900	8.19	< .001
Inversion	1, 900	0.24	.635
Eccentricity	20, 900	232.63	< .001
ROI*Inversion	5, 900	1.75	.143
ROI*Eccentricity	100, 900	16.46	< .001
Inversion*Eccentricity	20, 900	0.93	.555
ROI*Inversion*Eccentricity	100, 900	1.07	.299

Table A.3. ANOVA results comparing visual field coverage across ROI, inversion and eccentricity. Bold text indicates statistical significance ($p < .05$).

Fixed factors	β	p	CI
Intercept	-3.38		
ROI	1.66	< .001	1.42, 1.90
Location	0.13	.581	-0.33, 0.58
Inversion	0.07	.815	-0.50, 0.63
Eccentricity	0.17	< .001	0.14, 0.21

Table A.5. Linear mixed model results for pRF size across horizontal and vertical locations (with ROI as a predictor). Bold text indicates statistical significance ($p < .05$).

Factors	df	F	p
ROI	5, 180	144.24	< .001
Location	1, 180	94.25	< .001
Inversion	20, 180	0.32	.586
Eccentricity	5, 180	27.47	< .001
ROI*Location	100, 180	10.42	< .001
ROI*Inversion	20, 180	1.10	.373
ROI*Eccentricity	100, 180	5.18	< .001
Location*Inversion	1, 180	2.22	.170
Location*Eccentricity	20, 180	3.99	< .001
Inversion*Eccentricity	20, 180	2.71	< .001
ROI*Location*Inversion	5, 180	3.42	.004
ROI*Location*Eccentricity	100, 180	14.50	< .001
ROI*Inversion*Eccentricity	100, 180	0.67	.996
Location*Inversion*Eccentricity	20, 180	0.46	.980

Table A.4. ANOVA results for pRF number across horizontal and vertical locations (with ROI as a factor). Bold text indicates statistical significance ($p < .05$).

Fixed factors	β	p	CI
Intercept	-4.34		
ROI	1.48	< .001	1.12, 1.83
Location	0.91	< .001	0.46, 1.35
Inversion	0.13	.650	-0.42, 0.68
Eccentricity	0.17	< .001	0.13, 0.21

Table A.6. Linear mixed model results for pRF size across upper and lower locations (with ROI as a predictor). Bold text indicates statistical significance ($p < .05$).

Factors	df	F	p
ROI	5, 180	8.21	< .001
Location	1, 180	56.91	< .001
Inversion	20, 180	0.24	.633
Eccentricity	5, 180	232.37	< .001
ROI*Location	100, 180	6.62	< .001
ROI*Inversion	20, 180	1.74	.145
ROI*Eccentricity	100, 180	16.48	< .001
Location*Inversion	1, 180	0.13	.726
Location*Eccentricity	20, 180	50.77	< .001
Inversion*Eccentricity	20, 180	0.94	.543
ROI*Location*Inversion	5, 180	3.70	.003
ROI*Location*Eccentricity	100, 180	17.61	< .001
ROI*Inversion*Eccentricity	100, 180	1.55	< .001
Location*Inversion*Eccentricity	20, 180	0.04	1.00

Table A.7. ANOVA results for visual field coverage across horizontal and vertical locations (with ROI as a factor). Bold text indicates statistical significance ($p < .05$).

Factors	<i>df</i>	<i>F</i>	<i>p</i>
ROI	5, 180	123.48	< .001
Location	1, 180	57.71	< .001
Inversion	20, 180	1.26	.290
Eccentricity	5, 180	32.99	< .001
ROI*Location	100, 180	13.14	< .001
ROI*Inversion	20, 180	0.67	.647
ROI*Eccentricity	100, 180	8.69	< .001
Location*Inversion	1, 180	8.50	.017
Location*Eccentricity	20, 180	3.48	< .001
Inversion*Eccentricity	20, 180	3.06	< .001
ROI*Location*Inversion	5, 180	0.54	.744
ROI*Location*Eccentricity	100, 180	4.50	< .001
ROI*Inversion*Eccentricity	100, 180	0.67	.996
Location*Inversion*Eccentricity	20, 180	2.03	.004

Table A.9. ANOVA results for pRF number across upper and lower locations (with ROI as a factor). Bold text indicates statistical significance ($p < .05$).

Factors	<i>df</i>	<i>F</i>	<i>p</i>
ROI	5, 180	11.78	< .001
Location	1, 180	78.33	< .001
Inversion	20, 180	0.21	.658
Eccentricity	5, 180	228.11	< .001
ROI*Location	100, 180	2.91	.023
ROI*Inversion	20, 180	1.42	.235
ROI*Eccentricity	100, 180	13.59	< .001
Location*Inversion	1, 180	4.66	.059
Location*Eccentricity	20, 180	45.26	< .001
Inversion*Eccentricity	20, 180	1.04	.415
ROI*Location*Inversion	5, 180	1.97	.080
ROI*Location*Eccentricity	100, 180	6.90	< .001
ROI*Inversion*Eccentricity	100, 180	0.90	.745
Location*Inversion*Eccentricity	20, 180	1.62	.040

Table A.8. ANOVA results for visual field coverage across upper and lower locations (with ROI as a factor). Bold text indicates statistical significance ($p < .05$).




Figure A.6. Estimates of visual field coverage across eccentricity (in degrees of visual angle) derived from upright and inverted coverage maps that had been jointly rather than independently normalised. For each ROI, both the upright and inverted maps were divided by the maximum value of the upright map. As such, these values reflect absolute differences in coverage. Statistically significant differences between upright and inverted coverage values are indicated by the lines across the top ($p < .05$).




Figure A.7. Mean visual field coverage for upright (top) and inverted (bottom) faces. Coordinates represent eccentricity (in degrees of visual angle), with negative values for the left and positive values for the right visual field. Values were converted to log scale before plotting, for visualisation purposes (see colour bar).

Appendix B Supplementary Information for Chapter 4

Fixed factors	β	p	CI
Intercept	2.97		
ROI	0.94	< .001	0.54, 1.34
Localiser	-0.04	.488	-0.16, 0.08
Inversion	0.18	.706	-0.76, 1.13
Eccentricity	0.25	< .001	0.18, 0.33

Table B.1. Linear mixed model results comparing pRF size across localiser stimuli, with ROI as a predictor. Significant effects ($p < .05$) are indicated in bold.

Factors	df	F	p
ROI	2, 360	5.40	.015
Localiser	2, 360	21.22	< .001
Inversion	1, 360	0.17	.693
Eccentricity	20, 360	23.82	< .001
ROI*Localiser	4, 360	5.01	.003
ROI*Inversion	2, 360	0.96	.400
ROI*Eccentricity	40, 360	3.25	< .001
Localiser*Inversion	2, 360	1.55	.283
Localiser*Eccentricity	40, 360	9.97	< .001
Inversion*Eccentricity	20, 360	3.88	< .001
ROI*Localiser*Inversion	4, 360	0.65	.630
ROI*Localiser*Eccentricity	80, 360	1.05	.358
ROI*Inversion*Eccentricity	40, 360	7.40	< .001
Localiser*Inversion*Eccentricity	40, 360	0.24	1.00

Table B.2. ANOVA results comparing pRF number across localiser stimuli, with ROI as a factor. Significant effects ($p < .05$) are indicated in bold.

Factors	<i>df</i>	<i>F</i>	<i>p</i>
ROI	2, 360	1.02	.379
Localiser	2, 360	0.45	.642
Inversion	1, 360	0.96	.354
Eccentricity	20, 360	444.57	< .001
ROI*Localiser	4, 360	0.16	.959
ROI*Inversion	2, 360	2.47	.113
ROI*Eccentricity	40, 360	1.70	.007
Localiser*Inversion	2, 360	0.66	.527
Localiser*Eccentricity	40, 360	0.93	.588
Inversion*Eccentricity	20, 360	0.50	.962
ROI*Localiser*Inversion	4, 360	4.27	.002
ROI*Localiser*Eccentricity	80, 360	0.52	1.00
ROI*Inversion*Eccentricity	40, 360	5.07	< .001
Localiser*Inversion*Eccentricity	40, 360	0.14	1.00

Table B.3. ANOVA results comparing visual field coverage across localiser stimuli, with ROI as a factor. Significant effects ($p < .05$) are indicated in bold.

		OFA	pFus	mFus
P1	single	34.38 -48.46 -39.30	23.10 -41.66 -30.22	44.04 -21.69 -41.62
	tiled	34.38 -48.46 -39.30	23.69 -41.31 -30.48	44.20 -22.19 -41.46
	combi	35.16 -47.87 -38.31	22.29 -41.37 -30.51	44.04 -21.69 -41.62
P2	single	44.35 -55.06 -40.80	39.41 -51.49 -44.24	40.88 -34.66 -34.91
	tiled	44.35 -55.06 -40.80	39.41 -51.49 -44.24	40.88 -34.66 -34.91
	combi	44.35 -55.06 -40.80	39.86 -50.93 -43.97	40.88 -34.66 -34.91
P3	single	34.34 -56.99 -42.83	31.09 -47.92 -32.64	34.30 -25.33 -37.86
	tiled	33.60 -57.83 -41.89	30.95 -48.68 -32.53	34.30 -25.33 -37.86
	combi	36.62 -55.80 -43.25	31.08 -46.33 -32.63	34.52 -24.68 -38.38
P4	single	38.63 -62.22 -31.87	41.50 -35.60 -37.07	41.13 -20.56 -26.54
	tiled	38.63 -62.22 -31.87	39.14 -39.45 -36.33	39.84 -25.12 -31.08
	combi	38.15 -61.79 -30.94	41.50 -35.60 -37.07	36.43 -20.97 -31.75
P5	single	39.80 -49.13 -12.91	35.52 -30.74 -22.72	30.95 -19.76 -17.39
	tiled	39.80 -49.13 -12.91	35.52 -30.74 -22.72	30.95 -19.76 -17.39
	combi	39.80 -49.13 -12.91	35.52 -30.74 -22.72	30.95 -19.76 -17.39
P6	single	29.40 -65.88 -25.91	37.14 -50.04 -28.42	37.05 -28.06 -32.16
	tiled	29.40 -65.88 -25.91	37.14 -50.04 -28.42	37.05 -28.06 -32.16
	combi	29.40 -65.88 -25.91	36.88 -49.18 -28.28	37.17 -29.71 -32.06
P7	single	45.06 -58.95 -32.22	35.05 -41.96 -35.36	34.85 -26.27 -33.88
	tiled	45.06 -58.95 -32.22	35.05 -41.96 -35.36	35.06 -26.48 -34.65
	combi	45.84 -58.62 -32.37	35.05 -41.96 -35.36	34.85 -26.27 -33.88
P8	single	33.36 -55.58 -51.83	36.17 -39.88 -46.40	34.17 -23.59 -45.95
	tiled	34.11 -55.55 -51.96	27.19 -41.39 -49.28	34.17 -23.59 -45.95
	combi	33.36 -55.58 -51.83	36.17 -39.88 -46.40	37.79 -25.77 -47.10
P9	single	39.83 -69.83 -60.66	41.84 -50.19 -58.92	36.25 -24.77 -56.42
	tiled	48.31 -62.36 -53.66	41.84 -50.19 -58.92	36.25 -24.77 -56.42
	combi	39.83 -69.83 -60.66	41.84 -50.19 -58.92	35.77 -25.09 -55.69
P10	single	43.28 -47.62 -37.72	41.98 -20.91 -32.92	42.90 -5.18 -30.60
	tiled	43.90 -47.67 -38.69	41.98 -20.91 -32.92	42.90 -5.18 -30.60
	combi	43.28 -47.62 -37.72	37.04 -17.94 -28.21	42.90 -5.18 -30.60

Table B.4. Vertex RAS (Right, Anterior, Superior) coordinates of the vertex with the peak T statistic in each of the face-selective ROIs identified using the single, tiled and combination localiser stimuli. Coordinates are displayed for the right hemisphere of each participant (P1-10).




Figure B.1. Visual field coverage across eccentricity (in degrees of visual angle) for face-selective ROIs delineated using single, tiled or combination face stimuli. Coverage values were derived from coverage maps that had been jointly rather than independently normalised (all divided by the maximum value of the combination map, as opposed to their own maximum values). These plots therefore represent absolute differences in coverage. Only data from pRF mapping with upright faces are shown. Statistically significant differences between combination and single (blue lines) and combination and tiled (orange lines) are shown across the top ($p < .05$).

References

- Abegg, M., Pianezzi, D., & Barton, J. J. (2015). A vertical asymmetry in saccades. *Journal of Eye Movement Research*, 8(5). <https://doi.org/10.16910/jemr.8.5.3>
- Abrams, J., Nizam, A., & Carrasco, M. (2012). Isoeccentric locations are not equivalent: the extent of the vertical meridian asymmetry. *Vision Research*, 52(1), 70-78. <https://doi.org/10.1016/j.visres.2011.10.016>
- Afraz, A., Pashkam, M. V., & Cavanagh, P. (2010). Spatial heterogeneity in the perception of face and form attributes. *Current Biology*, 20(23), 2112-2116. <https://doi.org/10.1016/j.cub.2010.11.017>
- Albright, T. D. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. *Journal of Neurophysiology*, 52(6), 1106-1130. <https://doi.org/10.1152/jn.1984.52.6.1106>
- Alvarez, I., de Haas, B., Clark, C. A., Rees, G., & Schwarzkopf, D. S. (2015). Comparing different stimulus configurations for population receptive field mapping in human fMRI. *Frontiers in Human Neuroscience*, 9, 96. <https://doi.org/10.3389/fnhum.2015.00096>
- Amano, K., Wandell, B. A., & Dumoulin, S. O. (2009). Visual field maps, population receptive field sizes, and visual field coverage in the human MT+ complex. *Journal of Neurophysiology*, 102(5), 2704-2718. <https://doi.org/10.1152/jn.00102.2009>
- Anderson, E. J., Tibber, M. S., Schwarzkopf, D. S., Shergill, S. S., Fernandez-Egea, E., Rees, G., & Dakin, S. C. (2017). Visual Population Receptive Fields in People with Schizophrenia Have Reduced Inhibitory Surrounds. *The Journal of Neuroscience*, 37(6), 1546-1556. <https://doi.org/10.1523/JNEUROSCI.3620-15.2016>
- Anstis, S. (1998). Picturing peripheral acuity. *Perception*, 27(7), 817-825. <https://doi.org/10.1068/p270817>
- Arcaro, M. J., McMains, S. A., Singer, B. D., & Kastner, S. (2009). Retinotopic organization of human ventral visual cortex. *The Journal of Neuroscience*, 29(34), 10638-10652. <https://doi.org/10.1523/JNEUROSCI.2807-09.2009>
- Arcaro, M. J., Schade, P. F., Vincent, J. L., Ponce, C. R., & Livingstone, M. S. (2017). Seeing faces is necessary for face-domain formation. *Nature Neuroscience*, 20(10), 1404. <https://doi.org/10.1038/nn.4635>
- Ashburner, J., & Friston, K. J. (2014). *SPM12 Manual*.
- Awasthi, B., Friedman, J., & Williams, M. A. (2011). Faster, stronger, lateralized: Low spatial frequency information supports face processing. *Neuropsychologia*, 49(13), 3583-3590. <https://doi.org/10.1016/j.neuropsychologia.2011.08.027>

- Barbot, A., Xue, S., & Carrasco, M. (2021). Asymmetries in visual acuity around the visual field. *Journal of Vision*, 21(1), 2. <https://doi.org/10.1167/jov.21.1.2>
- Barlow. (2009). *Grandmother cells, symmetry, and invariance: how the term arose and what the facts suggest* (Vol. 4th Edn).
- Bartlett, J. C., & Searcy, J. (1993). Inversion and configuration of faces. *Cognitive Psychology*, 25(3), 281-316. <https://doi.org/10.1006/cogp.1993.1007>
- Barton, J. J. S., Press, D. Z., Keenan, J. P., & O'Connor, M. (2002). Lesions of the fusiform, face area impair perception of facial configuration in prosopagnosia. *Neurology*, 58(1), 71-78. <https://doi.org/10.1212/WNL.58.1.71>
- Behrmann, M., & Avidan, G. (2005). Congenital prosopagnosia: face-blind from birth. *Trends in Cognitive Sciences*, 9(4), 180-187. <https://doi.org/10.1016/j.tics.2005.02.011>
- Behrmann, M., & Moscovitch, M. (2001). Face recognition: Evidence from intact and impaired performance. *Handbook of Neuropsychology*, 4, 181-206.
- Benson, N. C., Kupers, E. R., Barbot, A., Carrasco, M., & Winawer, J. (2021). Cortical magnification in human visual cortex parallels task performance around the visual field. *eLife*, 10. <https://doi.org/10.7554/eLife.67685>
- Berman, M. G., Park, J., Gonzalez, R., Polk, T. A., Gehrke, A., Knaffla, S., & Jonides, J. (2010). Evaluating functional localizers: the case of the FFA. *NeuroImage*, 50(1), 56-71. <https://doi.org/10.1016/j.neuroimage.2009.12.024>
- Bodamer. (1947). Die Prosop-Agnosie. *Archiv für Psychiatrie und Nervenkrankheiten*, 179(6-53).
- Bourne, V. J., Vladeanu, M., & Hole, G. J. (2009). Lateralised repetition priming for featurally and configurally manipulated familiar faces: evidence for differentially lateralised processing mechanisms. *Laterality*, 14(3), 287-299. <https://doi.org/10.1080/13576500802383709>
- Boutet, I., Nelson, E. A., Watier, N., Cousineau, D., Beland, S., & Collin, C. A. (2021). Different measures of holistic face processing tap into distinct but partially overlapping mechanisms. *Attention Perception & Psychophysics*, 83(7), 2905-2923. <https://doi.org/10.3758/s13414-021-02337-7>
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10(4), 433-436.
- Brown, E., & Perrett, D. I. (1993). What gives a face its gender? *Perception*, 22(7), 829-840. <https://doi.org/10.1068/p220829>
- Bruce, V., Campbell, R. N., Doherty-Sneddon, G., Import, A., Langton, S., McAuley, S., & Wright, R. (2000). Testing face processing skills in children. *British Journal of Developmental Psychology*, 18, 319-333. <https://doi.org/10.1348/026151000165715>

- Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, 77 (Pt 3), 305-327. <https://doi.org/10.1111/j.2044-8295.1986.tb02199.x>
- Busigny, T., Joubert, S., Felician, O., Ceccaldi, M., & Rossion, B. (2010). Holistic perception of the individual face is specific and necessary: evidence from an extensive case study of acquired prosopagnosia. *Neuropsychologia*, 48(14), 4057-4092. <https://doi.org/10.1016/j.neuropsychologia.2010.09.017>
- Butler, S. H., & Harvey, M. (2005). Does inversion abolish the left chimeric face processing advantage? *Neuroreport*, 16(18), 1991-1993. <https://doi.org/10.1097/00001756-200512190-00004>
- Calder, A. J., & Jansen, J. (2005). Configural coding of facial expressions: The impact of inversion and photographic negative. *Visual Cognition*, 12(3), 495-518. <https://doi.org/10.1080/13506280444000418>
- Calder, A. J., Young, A. W., Keane, J., & Dean, M. (2000). Configural information in facial expression perception. *Journal of Experimental Psychology: Human Perception and Performance*, 26(2), 527-551. <https://doi.org/10.1037//0096-1523.26.2.527>
- Carey, Diamond, & Woods. (1980). Development of face recognition: A maturational component? *Developmental Psychology*, 16(4), 257-269. <https://doi.org/10.1037/0012-1649.16.4.257>
- Carlson, T., Hogendoorn, H., Fonteijn, H., & Verstraten, F. A. (2011). Spatial coding and invariance in object-selective cortex. *Cortex*, 47(1), 14-22. <https://doi.org/10.1016/j.cortex.2009.08.015>
- Carrasco, M., Myers, C., & Roberts, M. (2023). Visual field asymmetries vary between adolescents and adults. *bioRxiv*, 03. <https://doi.org/10.1101/2023.03.04.531124>
- Carrasco, M., Talgar, C. P., & Cameron, E. L. (2001). Characterizing visual performance fields: effects of transient covert attention, spatial frequency, eccentricity, task and set size. *Spatial Vision*, 15(1), 61-75. <https://doi.org/10.1163/15685680152692015>
- Chen, X., Liu, X., Parker, B. J., Zhen, Z., & Weiner, K. S. (2022). Functionally and structurally distinct fusiform face area(s) in over 1000 participants. *NeuroImage*, 265, 119765. <https://doi.org/10.1016/j.neuroimage.2022.119765>
- Collins, J. A., & Olson, I. R. (2014). Beyond the FFA: The role of the ventral anterior temporal lobes in face processing. *Neuropsychologia*, 61, 65-79. <https://doi.org/10.1016/j.neuropsychologia.2014.06.005>
- Collishaw, S. M., & Hole, G. J. (2000). Featural and configurational processes in the recognition of faces of different familiarity. *Perception*, 29(8), 893-909. <https://doi.org/10.1068/p2949>

- Cowey, A., & Rolls, E. T. (1974). Human cortical magnification factor and its relation to visual acuity. *Experimental Brain Research*, 21(5), 447-454. <https://doi.org/10.1007/BF00237163>
- Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, 29(3), 162-173. <https://doi.org/10.1006/cbmr.1996.0014>
- Cox, R. W., & Hyde, J. S. (1997). Software tools for analysis and visualization of fMRI data. *NMR in Biomedicine*, 10(4-5), 171-178. [https://doi.org/10.1002/\(SICI\)1099-1492\(199706/08\)10:4/5<171::AID-NBM453>3.0.CO;2-L](https://doi.org/10.1002/(SICI)1099-1492(199706/08)10:4/5<171::AID-NBM453>3.0.CO;2-L)
- Crider, B. (1944). A battery of tests for the dominant eye. *The Journal of General Psychology*, 31(2), 179-190. <https://doi.org/10.1080/00221309.1944.10543187>
- Curcio, C. A., & Allen, K. A. (1990). Topography of ganglion cells in human retina. *Journal of Comparative Neurology*, 300(1), 5-25. <https://doi.org/10.1002/cne.903000103>
- Dakin, S. C., & Watt. (2009). Biological “bar codes” in human faces. *Journal of Vision*, 9(4), 2-2. <https://doi.org/10.1167/9.4.2>
- Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical surface-based analysis. I. Segmentation and surface reconstruction. *NeuroImage*, 9(2), 179-194. <https://doi.org/10.1006/nimg.1998.0395>
- Daniel, P. M., & Whitteridge, D. (1961). The representation of the visual field on the cerebral cortex in monkeys. *Journal of Physiology*, 159(2), 203-221. <https://doi.org/10.1113/jphysiol.1961.sp006803>
- de Gelder, B., & Rouw, R. (2000). Configural face processes in acquired and developmental prosopagnosia: evidence for two separate face systems? *Neuroreport*, 11(14), 3145-3150. <https://doi.org/10.1097/00001756-200009280-00021>
- de Haas, B., Iakovidis, A. L., Schwarzkopf, D. S., & Gegenfurtner, K. R. (2019). Individual differences in visual salience vary along semantic dimensions. *Proceedings of the National Academy of Sciences*, 116(24), 11687-11692. <https://doi.org/10.1073/pnas.1820553116>
- de Haas, B., Schwarzkopf, D. S., Alvarez, I., Lawson, R. P., Henriksson, L., Kriegeskorte, N., & Rees, G. (2016). Perception and Processing of Faces in the Human Brain Is Tuned to Typical Feature Locations. *The Journal of Neuroscience*, 36(36), 9289-9302. <https://doi.org/10.1523/JNEUROSCI.4131-14.2016>
- de Lissa, P., McArthur, G., Hawelka, S., Palermo, R., Mahajan, Y., & Hutzler, F. (2014). Fixation location on upright and inverted faces modulates the N170.

- Dekker, T. M., Schwarzkopf, D. S., de Haas, B., Nardini, M., & Sereno, M. I. (2019). Population receptive field tuning properties of visual cortex during childhood. *Developmental Cognitive Neuroscience*, 37, 100614. <https://doi.org/10.1016/j.dcn.2019.01.001>
- DeYoe, E. A., Bandettini, P., Neitz, J., Miller, D., & Winans, P. (1994). Functional magnetic resonance imaging (fMRI) of the human brain. *Journal of Neuroscience Methods*, 54(2), 171-187. [https://doi.org/10.1016/0165-0270\(94\)90191-0](https://doi.org/10.1016/0165-0270(94)90191-0)
- Diamond, R., & Carey, S. (1986). Why faces are and are not special: an effect of expertise. *Journal of Experimental Psychology: General*, 115(2), 107-117. <https://doi.org/10.1037//0096-3445.115.2.107>
- DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, 11(8), 333-341. <https://doi.org/10.1016/j.tics.2007.06.010>
- DiCarlo, J. J., & Maunsell, J. H. (2003). Anterior inferotemporal neurons of monkeys engaged in object recognition can be highly sensitive to object retinal position. *Journal of Neurophysiology*, 89(6), 3264-3278. <https://doi.org/10.1152/jn.00358.2002>
- Downing, P. E., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A Cortical Area Selective for Visual Processing of the Human Body. *Science*, 293(5539), 2470-2473. <https://doi.org/doi:10.1126/science.1063414>
- Dricot, L., Sorger, B., Schiltz, C., Goebel, R., & Rossion, B. (2008). The roles of “face” and “non-face” areas during individual face perception: Evidence by fMRI adaptation in a brain-damaged prosopagnosic patient. *NeuroImage*, 40(1), 318-332. <https://doi.org/10.1016/j.neuroimage.2007.11.012>
- Dubner, R. F., & Zeki, S. M. (1971). Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Research*. [https://doi.org/10.1016/0006-8993\(71\)90494-X](https://doi.org/10.1016/0006-8993(71)90494-X)
- Duchaine, B. C., Dingle, K., Butterworth, E., & Nakayama, K. (2004). Normal greeble learning in a severe case of developmental prosopagnosia. *Neuron*, 43(4), 469-473. <https://doi.org/10.1016/j.neuron.2004.08.006>
- Duchaine, B. C., & Nakayama, K. (2006). Developmental prosopagnosia: a window to content-specific face processing. *Current Opinion in Neurobiology*, 16(2), 166-173. <https://doi.org/10.1016/j.conb.2006.03.003>
- Dumoulin, S. O., & Knapen, T. (2018). How Visual Cortical Organization Is Altered by Ophthalmologic and Neurologic Disorders. *Annual Review of Vision Science*, 4, 357-379. <https://doi.org/10.1146/annurev-vision-091517-033948>

- Dumoulin, S. O., & Wandell, B. A. (2008). Population receptive field estimates in human visual cortex. *NeuroImage*, 39(2), 647-660. <https://doi.org/10.1016/j.neuroimage.2007.09.034>
- Duncan, K. J., Pattamadilok, C., Knierim, I., & Devlin, J. T. (2009). Consistency and variability in functional localisers. *NeuroImage*, 46(4), 1018-1026. <https://doi.org/10.1016/j.neuroimage.2009.03.014>
- Duncan, K. J. K., & Devlin, J. T. (2011). Improving the reliability of functional localizers. *NeuroImage*, 57(3), 1022-1030. <https://doi.org/10.1016/j.neuroimage.2011.05.009>
- Duncan, R. O., & Boynton, G. M. (2003). Cortical magnification within human primary visual cortex correlates with acuity thresholds. *Neuron*, 38(4), 659-671. [https://doi.org/10.1016/S0896-6273\(03\)00265-4](https://doi.org/10.1016/S0896-6273(03)00265-4)
- Eimer, M. (2000). The face-specific N170 component reflects late stages in the structural encoding of faces. *Neuroreport*, 11(10), 2319-2324. <https://doi.org/10.1097/00001756-200007140-00050>
- Ellis, H. D., & Shepherd, J. W. (1975). Recognition of upright and inverted faces presented in the left and right visual fields. *Cortex*, 11(1), 3-7. [https://doi.org/10.1016/s0010-9452\(75\)80014-1](https://doi.org/10.1016/s0010-9452(75)80014-1)
- Engel, S. A., Glover, G. H., & Wandell, B. A. (1997). Retinotopic organisation in human visual cortex and the spatial precision of functional MRI. *Cerebral Cortex*, 7, 181-192. <https://doi.org/10.1093/cercor/7.2.181>
- Engell, A. D., & Haxby, J. V. (2007). Facial expression and gaze-direction in human superior temporal sulcus. *Neuropsychologia*, 45(14), 3234-3241. <https://doi.org/10.1016/j.neuropsychologia.2007.06.022>
- Epstein, & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, 392(6676), 598-601. <https://doi.org/10.1038/33402>
- Fang, F., Murray, S. O., & He, S. (2007). Duration-dependent fMRI adaptation and distributed viewer-centered face representation in human visual cortex. *Cerebral Cortex*, 17(6), 1402-1411. <https://doi.org/10.1093/cercor/bhl053>
- Farah, M. J., Wilson, K. D., Drain, H. M., & Tanaka, J. R. (1995). The Inverted Face Inversion Effect in Prosopagnosia - Evidence for Mandatory, Face-Specific Perceptual Mechanisms. *Vision Research*, 35(14), 2089-2093. [https://doi.org/10.1016/0042-6989\(94\)00273-O](https://doi.org/10.1016/0042-6989(94)00273-O)
- Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. N. (1998). What is "special" about face perception? *Psychological Review*, 105(3), 482-498. <https://doi.org/10.1037/0033-295x.105.3.482>
- Finzi, D., Gomez, J., Nordt, M., Rezai, A. A., Poltoratski, S., & Grill-Spector, K. (2021). Differential spatial computations in ventral and lateral face-selective regions are

- scaffolded by structural connections. *Nature Communications*, 12(1), 2278. <https://doi.org/10.1038/s41467-021-22524-2>
- Fischl, B. (2012). FreeSurfer. *NeuroImage*, 62(2), 774-781. <https://doi.org/10.1016/j.neuroimage.2012.01.021>
- Fischl, B., Sereno, M. I., & Dale, A. M. (1999). Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage*, 9(2), 195-207. <https://doi.org/10.1006/nimg.1998.0396>
- Fisher, K., Towler, J., & Eimer, M. (2016a). Facial identity and facial expression are initially integrated at visual perceptual stages of face processing. *Neuropsychologia*, 80, 115-125. <https://doi.org/10.1016/j.neuropsychologia.2015.11.011>
- Fisher, K., Towler, J., & Eimer, M. (2016b). Reduced sensitivity to contrast signals from the eye region in developmental prosopagnosia. *Cortex*, 81, 64-78. <https://doi.org/10.1016/j.cortex.2016.04.005>
- Fisher, K., Towler, J., & Eimer, M. (2017). Face identity matching is selectively impaired in developmental prosopagnosia. *Cortex*, 89, 11-27. <https://doi.org/10.1016/j.cortex.2017.01.003>
- Fodor, J. A. (1983). The modularity of mind: an essay on faculty psychology. Cambridge, MA: MIT.
- Fortenbaugh, F. C., Silver, M. A., & Robertson, L. C. (2015). Individual differences in visual field shape modulate the effects of attention on the lower visual field advantage in crowding. *Journal of Vision*, 15(2), 19-19. <https://doi.org/10.1167/15.2.19>
- Freeman, J., & Simoncelli, E. P. (2011). Metamers of the ventral stream. *Nature Neuroscience*, 14(9), 1195-1201. <https://doi.org/10.1038/nn.2889>
- Freire, A., & Lee, K. (2001). Face recognition in 4-to 7-year-olds: Processing of configural, featural, and paraphernalia information. *Journal of Experimental Child Psychology*, 80(4), 347-371. <https://doi.org/10.1006/jecp.2001.2639>
- Freiwald, W., Duchaine, B., & Yovel, G. (2016). Face processing systems: from neurons to real-world social perception. *Annual Review of Neuroscience*, 39, 325-346. <https://doi.org/10.1146/annurev-neuro-070815-013934>
- Friston, K. J., Rotshtein, P., Geng, J. J., Sterzer, P., & Henson, R. N. (2006). A critique of functional localisers. *NeuroImage*, 30(4), 1077-1087. <https://doi.org/10.1016/j.neuroimage.2005.08.012>
- Galton, F. (1879). Composite portraits, made by combining those of many different persons into a single resultant figure. *The Journal of the Anthropological Institute of Great Britain and Ireland*, 8, 132-144. <https://doi.org/10.2307/2841021>

- Gauthier, I. (1998). Training 'greeble' experts: a framework for studying expert object recognition processes. *Vision Research*, 38(15-16), 2401-2428. [https://doi.org/10.1016/S0042-6989\(97\)00442-2](https://doi.org/10.1016/S0042-6989(97)00442-2)
- Gauthier, I., Klaiman, C., & Schultz, R. T. (2009). Face composite effects reveal abnormal face processing in Autism spectrum disorders. *Vision Research*, 49(4), 470-478. <https://doi.org/10.1016/j.visres.2008.12.007>
- Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, 3(2), 191-197. <https://doi.org/10.1038/72140>
- Gauthier, I., & Tarr, M. J. (1997). Becoming a "Greeble" expert: exploring mechanisms for face recognition. *Vision Research*, 37(12), 1673-1682. [https://doi.org/10.1016/s0042-6989\(96\)00286-6](https://doi.org/10.1016/s0042-6989(96)00286-6)
- Ge, L., Wang, Z., McCleery, J. P., & Lee, K. (2006). Activation of face expertise and the inversion effect. *Psychological Science*, 17(1), 12-16. <https://doi.org/10.1111/j.1467-9280.2005.01658.x>
- Gibson, J. J., & Radner, M. (1937). Adaptation, after-effect and contrast in the perception of tilted lines. I. Quantitative studies. *Journal of Experimental Psychology*, 20(5), 453. <https://doi.org/10.1037/h0059826>
- Glen, F. C., Crabb, D. P., Smith, N. D., Burton, R., & Garway-Heath, D. F. (2012). Do patients with glaucoma have difficulty recognizing faces? *Investigative Ophthalmology and Visual Science*, 53(7), 3629-3637. <https://doi.org/10.1167/iovs.11-8538>
- Goffaux, V., & Dakin, S. C. (2010). Horizontal information drives the behavioral signatures of face processing. *Frontiers in Psychology*, 1, 143. <https://doi.org/10.3389/fpsyg.2010.00143>
- Goffaux, V., Duecker, F., Hausfeld, L., Schiltz, C., & Goebel, R. (2016). Horizontal tuning for faces originates in high-level Fusiform Face Area. *Neuropsychologia*, 81, 1-11. <https://doi.org/10.1016/j.neuropsychologia.2015.12.004>
- Goffaux, V., & Greenwood, J. A. (2016). The orientation selectivity of face identification. *Scientific Reports*, 6, 34204. <https://doi.org/10.1038/srep34204>
- Goffaux, V., Hault, B., Michel, C., Vuong, Q. C., & Rossion, B. (2005). The respective role of low and high spatial frequencies in supporting configural and featural processing of faces. *Perception*, 34(1), 77-86. <https://doi.org/10.1068/p5370>
- Goffaux, V., & 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), 1023-1039. <https://doi.org/10.1037/0096-1523.32.4.1023>

- Goffaux, V., & Rossion, B. (2007). Face inversion disproportionately impairs the perception of vertical but not horizontal relations between features. *Journal of Experimental Psychology: Human Perception and Performance*, 33(4), 995. <https://doi.org/10.1037/0096-1523.33.4.995>
- Gomez, J., Natu, V., Jeska, B., Barnett, M., & Grill-Spector, K. (2018). Development differentially sculpts receptive fields across early and high-level human visual cortex. *Nature Communications*, 9(1), 788. <https://doi.org/10.1038/s41467-018-03166-3>
- Greenwood, J. A. (2023). *Eccentric Vision Lab Github*. <https://www.github.com/eccentricvision>
- Greenwood, J. A., Szinte, M., Sayim, B., & Cavanagh, P. (2017). Variations in crowding, saccadic precision, and spatial localization reveal the shared topology of spatial vision. *Proceedings of the National Academy of Sciences*, 114(17), E3573-E3582. <https://doi.org/10.1073/pnas.1615504114>
- Gregorek, S. (2022). Population receptive field mapping reveals different spatial computations across body-and face-selective regions. <https://doi.org/http://arks.princeton.edu/ark:/88435/dsp01bn999991g>
- Grill-Spector, K., Knouf, N., & Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nature Neuroscience*, 7(5), 555-562.
- Grill-Spector, K., Weiner, K. S., Kay, K., & Gomez, J. (2017). The Functional Neuroanatomy of Human Face Perception. *Annual Review of Vision Science*, 3, 167-196. <https://doi.org/10.1146/annurev-vision-102016-061214>
- Groen, I. I. A., Dekker, T. M., Knapen, T., & Silson, E. H. (2022). Visuospatial coding as ubiquitous scaffolding for human cognition. *Trends in Cognitive Sciences*, 26(1), 81-96. <https://doi.org/10.1016/j.tics.2021.10.011>
- Gross, C. G., Rocha-Miranda, C. E., & Bender, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the Macaque. *Journal of Neurophysiology*, 35(1), 96-111. <https://doi.org/10.1152/jn.1972.35.1.96>
- Harrison, M., & Strother, L. (2021). Does face-selective cortex show a left visual field bias for centrally-viewed faces? *Neuropsychologia*, 159(107956). <https://doi.org/10.1016/j.neuropsychologia.2021.107956>
- Harvey, B. M., & Dumoulin, S. O. (2011). The relationship between cortical magnification factor and population receptive field size in human visual cortex: constancies in cortical architecture. *The Journal of Neuroscience*, 31(38), 13604-13612. <https://doi.org/10.1523/jneurosci.2572-11.2011>
- Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, 4(6), 223-233. [https://doi.org/10.1016/s1364-6613\(00\)01482-0](https://doi.org/10.1016/s1364-6613(00)01482-0)

- Haxby, J. V., Ungerleider, L. G., Clark, V. P., Schouten, J. L., Hoffman, E. A., & Martin, A. (1999). The effect of face inversion on activity in human neural systems for face and object perception. *Neuron*, 22(1), 189-199. [https://doi.org/10.1016/S0896-6273\(00\)80690-X](https://doi.org/10.1016/S0896-6273(00)80690-X)
- He, S., & MacLeod, D. I. A. (2001). Orientation-selective adaptation and tilt after-effect from invisible patterns. *Nature*, 411(6836), 473-476. <https://doi.org/10.1038/35078072>
- Henderson, M. M., Tarr, M. J., & Wehbe, L. (2023). Low-level tuning biases in higher visual cortex reflect the semantic informativeness of visual features. *Journal of Vision*, 23(4), 8-8. <https://doi.org/10.1167/jov.23.4.8>
- Henriksson, L., Mur, M., & Kriegeskorte, N. (2015). Faciotopy-A face-feature map with face-like topology in the human occipital face area. *Cortex*, 72, 156-167. <https://doi.org/10.1016/j.cortex.2015.06.030>
- Henriksson, L., Nurminen, L., Hyvärinen, A., & Vanni, S. (2008). Spatial frequency tuning in human retinotopic visual areas. *Journal of Vision*, 8(10), 5-5. <https://doi.org/10.1167/8.10.5>
- Henson, R. N. (2016). Repetition suppression to faces in the fusiform face area: A personal and dynamic journey. *Cortex*, 80, 174-184. <https://doi.org/10.1016/j.cortex.2015.09.012>
- Himmelberg, M. M., Tuncok, E., Gomez, J., Grill-Spector, K., Carrasco, M., & Winawer, J. (2023). Comparing retinotopic maps of children and adults reveals a late-stage change in how V1 samples the visual field. *Nature Communications*, 14(1), 1561. <https://doi.org/10.1038/s41467-023-37280-8>
- Himmelberg, M. M., Winawer, J., & Carrasco, M. (2020). Stimulus-dependent contrast sensitivity asymmetries around the visual field. *Journal of Vision*, 20(9), 18. <https://doi.org/10.1167/jov.20.9.18>
- Honda, H., & Findlay, J. (1992). Saccades to targets in three-dimensional space: Dependence of saccadic latency on target location. *Perception and Psychophysics*, 52, 167-174. <https://doi.org/10.3758/BF03206770>
- Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160(1), 106. <https://doi.org/10.1113/jphysiol.1962.sp006837>
- Hughes, A. E., Greenwood, J. A., Finlayson, N. J., & Schwarzkopf, D. S. (2019). Population receptive field estimates for motion-defined stimuli. *NeuroImage*, 199, 245-260. <https://doi.org/10.1016/j.neuroimage.2019.05.068>
- Humphreys, K., Avidan, G., & Behrmann, M. (2007). A detailed investigation of facial expression processing in congenital prosopagnosia as compared to acquired prosopagnosia. *Experimental Brain Research*, 176(2), 356-373. <https://doi.org/10.1007/s00221-006-0621-5>

- Ito, M., Tamura, H., Fujita, I., & Tanaka, K. (1995). Size and position invariance of neuronal responses in monkey inferotemporal cortex. *Journal of Neurophysiology*, 73(1), 218-226. <https://doi.org/10.1152/jn.1995.73.1.218>
- Johnson, M. H., Dziurawiec, S., Ellis, H., & Morton, J. (1991). Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition*, 40(1-2), 1-19. [https://doi.org/10.1016/0010-0277\(91\)90045-6](https://doi.org/10.1016/0010-0277(91)90045-6)
- Kalpadakis-Smith, A. V., Goffaux, V., & Greenwood, J. A. (2018). Crowding for faces is determined by visual (not holistic) similarity: Evidence from judgements of eye position. *Scientific Reports*, 8(1), 12556. <https://doi.org/10.1038/s41598-018-30900-0>
- Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, 8(5), 679-685. <https://doi.org/10.1038/nn1444>
- Kamps, F. S., Hendrix, C. L., Brennan, P. A., & Dilks, D. D. (2020). Connectivity at the origins of domain specificity in the cortical face and place networks. *Proceedings of the National Academy of Sciences*, 117(11), 6163-6169. <https://doi.org/10.1073/pnas.1911359117>
- Kanwisher, N. (2000). Domain specificity in face perception. *Nature Neuroscience*, 3(8), 759-763. <https://doi.org/10.1038/77664>
- Kanwisher, N. (2017). The Quest for the FFA and Where It Led. *The Journal of Neuroscience*, 37(5), 1056-1061. <https://doi.org/10.1523/JNEUROSCI.1706-16.2016>
- Kanwisher, N., & Dilks, D. D. (2013). The functional organization of the ventral visual pathway in humans. *The New Visual Neurosciences*, 733-748.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience*, 17(11), 4302-4311. <https://doi.org/10.1523/JNEUROSCI.17-11-04302.1997>
- Kanwisher, N., Tong, F., & Nakayama, K. (1998). The effect of face inversion on the human fusiform face area. *Cognition*, 68(1), B1-B11. [https://doi.org/10.1016/s0010-0277\(98\)00035-3](https://doi.org/10.1016/s0010-0277(98)00035-3)
- Kay, K. N., Weiner, K. S., & Grill-Spector, K. (2015). Attention reduces spatial uncertainty in human ventral temporal cortex. *Current Biology*, 25(5), 595-600. <https://doi.org/10.1016/j.cub.2014.12.050>
- Kay, K. N., Winawer, J., Mezer, A., & Wandell, B. A. (2013). Compressive spatial summation in human visual cortex. *Journal of Neurophysiology*, 110(2), 481-494. <https://doi.org/10.1152/jn.00105.2013>

- Khandhadia, A. P., Murphy, A. P., Koyano, K. W., Esch, E. M., & Leopold, D. A. (2023). Encoding of 3D physical dimensions by face-selective cortical neurons. *Proceedings of the National Academy of Sciences*, 120(9), e2214996120. <https://doi.org/10.1073/pnas.2214996120>
- Klein, B. P., Harvey, B. M., & Dumoulin, S. O. (2014). Attraction of position preference by spatial attention throughout human visual cortex. *Neuron*, 84(1), 227-237. <https://doi.org/10.1016/j.neuron.2014.08.047>
- Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3?
- Kosovicheva, A., & Whitney, D. (2017). Stable individual signatures in object localization. *Current Biology*, 27(14), R700-R701. <https://doi.org/10.1016/j.cub.2017.06.001>
- Kovacs, G., Iffland, L., Vidnyanszky, Z., & Greenlee, M. W. (2012). Stimulus repetition probability effects on repetition suppression are position invariant for faces. *NeuroImage*, 60(4), 2128-2135. <https://doi.org/10.1016/j.neuroimage.2012.02.038>
- Kovacs, P., Knakker, B., Hermann, P., Kovacs, G., & Vidnyanszky, Z. (2017). Face inversion reveals holistic processing of peripheral faces. *Cortex*, 97, 81-95. <https://doi.org/10.1016/j.cortex.2017.09.020>
- Kravitz, D. J., Kriegeskorte, N., & Baker, C. I. (2010). High-level visual object representations are constrained by position. *Cerebral Cortex*, 20(12), 2916-2925. <https://doi.org/10.1093/cercor/bhq042>
- Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S. F., & Baker, C. I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. *Nature Neuroscience*, 12(5), 535-540. <https://doi.org/10.1038/nn.2303>
- Kung, C. C., Peissig, J. J., & Tarr, M. J. (2007). Is region-of-interest overlap comparison a reliable measure of category specificity? *Journal of Cognitive Neuroscience*, 19(12), 2019-2034. <https://doi.org/10.1162/jocn.2007.19.12.2019>
- Kurzawski, J. W., Burchell, A., Thapa, D., Majaj, N. J., Winawer, J. A., & Pelli, D. G. (2021). An enhanced Bouma model fits a hundred people's visual crowding. *bioRxiv*. <https://doi.org/10.1101/2021.04.12.439570>
- Lacadie, C. M., Fulbright, R. K., Rajeevan, N., Constable, R. T., & Papademetris, X. (2008). More accurate Talairach coordinates for neuroimaging using non-linear registration. *NeuroImage*, 42(2), 717-725. <https://doi.org/10.1016/j.neuroimage.2008.04.240>
- Lage-Castellanos, A., Valente, G., Senden, M., & De Martino, F. (2020). Investigating the Reliability of Population Receptive Field Size Estimates Using fMRI [Original Research]. *Frontiers in Neuroscience*, 14. <https://doi.org/10.3389/fnins.2020.00825>

- Laguerre, R., Dormal, G., Biervoye, A., Keufner, D., & Rossion, B. (2012). Extensive visual training in adulthood significantly reduces the face inversion effect. *Journal of Vision*, 12(4). <https://doi.org/10.1167/12.10.14>
- Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D. H. J., Hawk, S. T., & van Knippenberg, A. (2010). Presentation and validation of the Radboud Faces Database. *Cognition & Emotion*, 24(8), 1377-1388. <https://doi.org/10.1080/02699930903485076>
- Le Grand, R., Mondloch, C. J., Maurer, D., & Brent, H. P. (2001). Early visual experience and face processing. *Nature*, 410(6831), 890. <https://doi.org/10.1038/35073749>
- Le, R., Witthoft, N., Ben-Shachar, M., & Wandell, B. (2017). The field of view available to the ventral occipito-temporal reading circuitry. *Journal of Vision*, 17(4), 6. <https://doi.org/10.1167/17.4.6>
- Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, 4(1), 89-94. <https://doi.org/10.1038/82947>
- Lerma-Usabiaga, G., Benson, N., Winawer, J., & Wandell, B. A. (2020). A validation framework for neuroimaging software: The case of population receptive fields. *PLoS Computational Biology*, 16(6), e1007924. <https://doi.org/10.1371/journal.pcbi.1007924>
- Levy, I., Hasson, U., Avidan, G., Hendler, T., & Malach, R. (2001). Center-periphery organization of human object areas. *Nature Neuroscience*, 4(5), 533-539. <https://doi.org/10.1038/87490>
- Lewis, M. B. (2001). The Lady's not for turning: Rotation of the Thatcher illusion. *Perception*, 30(6), 769-774. <https://doi.org/10.1068/p3174>
- Linhardt, D., Pawloff, M., Hummer, A., Woletz, M., Tik, M., Ritter, M., Schmidt-Erfurth, U., & Windischberger, C. (2021). Combining stimulus types for improved coverage in population receptive field mapping. *NeuroImage*, 238, 118240. <https://doi.org/10.1016/j.neuroimage.2021.118240>
- Liu, C. H., Young, A. W., Basra, G., Ren, N., & Chen, W. (2020). Perceptual integration and the composite face effect. *Quarterly Journal of Experimental Psychology*, 73(7), 1101-1114. <https://doi.org/10.1177/1747021819899531>
- Liu, J., Harris, A., & Kanwisher, N. (2010). Perception of Face Parts and Face Configurations: An fMRI Study. *Journal of Cognitive Neuroscience*, 22(1), 203-211. <https://doi.org/10.1162/jocn.2009.21203>
- Livingstone, M. S., Vincent, J. L., Arcaro, M. J., Srihasam, K., Schade, P. F., & Savage, T. (2017). Development of the macaque face-patch system. *Nature Communications*, 8(1), 14897. <https://doi.org/10.1038/ncomms14897>

- Loffler, G., Wilson, H. R., & Wilkinson, F. (2003). Local and global contributions to shape discrimination. *Vision Research*, 43(5), 519-530. [https://doi.org/10.1016/S0042-6989\(02\)00686-7](https://doi.org/10.1016/S0042-6989(02)00686-7)
- Logothetis, N. K., & Sheinberg, D. L. (1996). Visual object recognition. *Annual Review of Neuroscience*, 19, 577-621. <https://doi.org/10.1146/annurev.ne.19.030196.003045>
- Malach, R., Levy, I., & Hasson, U. (2002). The topography of high-order human object areas. *Trends in Cognitive Sciences*, 6(4), 176-184. [https://doi.org/10.1016/S1364-6613\(02\)01870-3](https://doi.org/10.1016/S1364-6613(02)01870-3)
- Maurer, D., Le Grand, R., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, 6(6), 255-260. [https://doi.org/10.1016/S1364-6613\(02\)01903-4](https://doi.org/10.1016/S1364-6613(02)01903-4)
- McCarthy, G., Puce, A., Gore, J. C., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. *Journal of Cognitive Neuroscience*, 9(5), 605-610. <https://doi.org/10.1162/jocn.1997.9.5.605>
- McKone, E. (2004). Isolating the special component of face recognition: peripheral identification and a Mooney face. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(1), 181-197. <https://doi.org/10.1037/0278-7393.30.1.181>
- McKone, E., Kanwisher, N., & Duchaine, B. C. (2007). Can generic expertise explain special processing for faces? *Trends in Cognitive Sciences*, 11(1), 8-15. <https://doi.org/10.1016/j.tics.2006.11.002>
- Meinhardt-Injac, B., Persike, M., & Meinhardt, G. (2010). The time course of face matching by internal and external features: Effects of context and inversion. *Vision Research*, 50(16), 1598-1611. <https://doi.org/10.1016/j.visres.2010.05.018>
- Mooney, C. M. (1957). Age in the development of closure ability in children. *Canadian Journal of Psychology*, 11(4), 219-226. <https://doi.org/10.1037/h0083717>
- Moscovitch, M., Winocur, G., & Behrmann, M. (1997). What is special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. *Journal of Cognitive Neuroscience*, 9(5), 555-604. <https://doi.org/10.1162/jocn.1997.9.5.555>
- Moutsiana, C., De Haas, B., Papageorgiou, A., Van Dijk, J. A., Balraj, A., Greenwood, J. A., & Schwarzkopf, D. S. (2016). Cortical idiosyncrasies predict the perception of object size. *Nature Communications*, 7(1), 12110. <https://doi.org/10.1038/ncomms12110>
- Nam, Y., Sato, T., Uchida, G., Malakhova, E., Ullman, S., & Tanifuji, M. (2021). View-tuned and view-invariant face encoding in IT cortex is explained by selected

- natural image fragments. *Scientific Reports*, 11(1), 1-14. <https://doi.org/10.1038/s41598-021-86842-7>
- Narumoto, J., Okada, T., Sadato, N., Fukui, K., & Yonekura, Y. (2001). Attention to emotion modulates fMRI activity in human right superior temporal sulcus. *Cognitive Brain Research*, 12(2), 225-231. [https://doi.org/10.1016/s0926-6410\(01\)00053-2](https://doi.org/10.1016/s0926-6410(01)00053-2)
- Natu, V. S., Barnett, M. A., Hartley, J., Gomez, J., Stigliani, A., & Grill-Spector, K. (2016). Development of Neural Sensitivity to Face Identity Correlates with Perceptual Discriminability. *The Journal of Neuroscience*, 36(42), 10893-10907. <https://doi.org/10.1523/JNEUROSCI.1886-16.2016>
- Nelder, J. A., & Mead, R. (1965). A Simplex-Method for Function Minimization. *Computer Journal*, 7(4), 308-313. <https://doi.org/10.1093/comjnl/7.4.308>
- Op De Beeck, H., & Vogels, R. (2000). Spatial sensitivity of macaque inferior temporal neurons. *Journal of Comparative Neurology*, 426(4), 505-518. [https://doi.org/10.1002/1096-9861\(20001030\)426:4<505::AID-CNE1>3.0.CO;2-M](https://doi.org/10.1002/1096-9861(20001030)426:4<505::AID-CNE1>3.0.CO;2-M)
- Orban de Xivry, J. J., Ramon, M., Lefevre, P., & Rossion, B. (2008). Reduced fixation on the upper area of personally familiar faces following acquired prosopagnosia. *Journal of Neuropsychology*, 2(1), 245-268. <https://doi.org/10.1348/174866407x260199>
- Pallett, P. M., & Meng, M. (2015). Inversion effects reveal dissociations in facial expression of emotion, gender, and object processing. *Frontiers in Psychology*, 6, 1029. <https://doi.org/10.3389/fpsyg.2015.01029>
- Párraga, C. A., Troscianko, T., & Tolhurst, D. J. (2000). The human visual system is optimised for processing the spatial information in natural visual images. *Current Biology*, 10(1), 35-38. [https://doi.org/10.1016/S0960-9822\(99\)00262-6](https://doi.org/10.1016/S0960-9822(99)00262-6)
- Peelen, M. V., & Downing, P. E. (2005). Within-subject reproducibility of category-specific visual activation with functional MRI. *Human Brain Mapping*, 25(4), 402-408. <https://doi.org/10.1002/hbm.20116>
- Pelli, D. G., & Vision, S. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437-442.
- Perry, V. H., & Cowey, A. (1985). The ganglion cell and cone distributions in the monkey's retina: implications for central magnification factors. *Vision Research*, 25(12), 1795-1810. [https://doi.org/10.1016/0042-6989\(85\)90004-5](https://doi.org/10.1016/0042-6989(85)90004-5)
- Petras, K., Ten Oever, S., Jacobs, C., & Goffaux, V. (2019). Coarse-to-fine information integration in human vision. *NeuroImage*, 186, 103-112. <https://doi.org/10.1016/j.neuroimage.2018.10.086>

- Piepers, D. W., & Robbins, R. A. (2012). A Review and Clarification of the Terms "holistic," "configural," and "relational" in the Face Perception Literature. *Frontiers in Psychology*, 3, 559. <https://doi.org/10.3389/fpsyg.2012.00559>
- Pinsk, M. A., Arcaro, M., Weiner, K. S., Kalkus, J. F., Inati, S. J., Gross, C. G., & Kastner, S. (2009). Neural representations of faces and body parts in macaque and human cortex: a comparative fMRI study. *Journal of Neurophysiology*, 101(5), 2581-2600. <https://doi.org/10.1152/jn.91198.2008>
- Pitcher, D. (2014). Facial expression recognition takes longer in the posterior superior temporal sulcus than in the occipital face area. *The Journal of Neuroscience*, 34(27), 9173-9177. <https://doi.org/10.1523/JNEUROSCI.5038-13.2014>
- Pitcher, D., Walsh, V., & Duchaine, B. (2011). The role of the occipital face area in the cortical face perception network. *Experimental Brain Research*, 209(4), 481-493. <https://doi.org/10.1007/s00221-011-2579-1>
- Pitcher, D., Walsh, V., Yovel, G., & Duchaine, B. (2007). TMS evidence for the involvement of the right occipital face area in early face processing. *Current Biology*, 17(18), 1568-1573. <https://doi.org/10.1016/j.cub.2007.07.063>
- Poldrack, R. A. (2007). Region of interest analysis for fMRI. *Social Cognitive and Affective Neuroscience*, 2(1), 67-70. <https://doi.org/10.1093/scan/nsm006>
- Poltoratski, S., Kay, K., Finzi, D., & Grill-Spector, K. (2021). Holistic face recognition is an emergent phenomenon of spatial processing in face-selective regions. *Nature Communications*, 12(1), 4745. <https://doi.org/10.1038/s41467-021-24806-1>
- Prkachin, G. C. (2003). The effects of orientation on detection and identification of facial expressions of emotion. *British Journal of Psychology*, 94(Pt 1), 45-62. <https://doi.org/10.1348/000712603762842093>
- Quek, G. L., & Finkbeiner, M. (2014). Face-sex categorization is better above fixation than below: Evidence from the reach-to-touch paradigm. *Cognitive, Affective, & Behavioral Neuroscience*, 14, 1407-1419. <https://doi.org/10.3758/s13415-014-0282-y>
- Quek, G. L., & Finkbeiner, M. (2016). The upper-hemifield advantage for masked face processing: Not just an attentional bias. *Attention, Perception, & Psychophysics*, 78, 52-68. <https://doi.org/10.3758/s13414-015-0965-7>
- Ramon, M., Busigny, T., & Rossion, B. (2010). Impaired holistic processing of unfamiliar individual faces in acquired prosopagnosia. *Neuropsychologia*, 48(4), 933-944. <https://doi.org/10.1016/j.neuropsychologia.2009.11.014>
- Rangarajan, V., Hermes, D., Foster, B. L., Weiner, K. S., Jacques, C., Grill-Spector, K., & Parvizi, J. (2014). Electrical stimulation of the left and right human fusiform gyrus causes different effects in conscious face perception. *The Journal of*

- Neuroscience*, 34(38), 12828-12836.
<https://doi.org/10.1523/JNEUROSCI.0527-14.2014>
- Rayner, K. (1977). Visual attention in reading: Eye movements reflect cognitive processes. *Memory and Cognition*, 5(4), 443-448.
<https://doi.org/10.3758/BF03197383>
- Rezlescu, C., Barton, J. J., Pitcher, D., & Duchaine, B. (2014). Normal acquisition of expertise with greebles in two cases of acquired prosopagnosia. *Proceedings of the National Academy of Sciences*, 111(14), 5123-5128.
<https://doi.org/10.1073/pnas.1317125111>
- Rezlescu, C., Pitcher, D., & Duchaine, B. (2012). Acquired prosopagnosia with spared within-class object recognition but impaired recognition of degraded basic-level objects. *Cognitive Neuropsychology*, 29(4), 325-347.
<https://doi.org/10.1080/02643294.2012.749223>
- Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, 2(11), 1019-1025. <https://doi.org/10.1038/14819>
- Robbins, R., & McKone, E. (2007). No face-like processing for objects-of-expertise in three behavioural tasks. *Cognition*, 103(1), 34-79.
<https://doi.org/10.1016/j.cognition.2006.02.008>
- Rosenholtz, R. (2016). Capabilities and Limitations of Peripheral Vision. *Annual Review of Vision Science*, 2, 437-457. <https://doi.org/10.1146/annurev-vision-082114-035733>
- Rossion, B. (2008). Picture-plane inversion leads to qualitative changes of face perception. *Acta Psychologica*, 128(2), 274-289.
<https://doi.org/10.1016/j.actpsy.2008.02.003>
- Rossion, B., & Gauthier, I. (2002). How does the brain process upright and inverted faces? *Behavioral and Cognitive Neuroscience Reviews*, 1(1), 63-75.
<https://doi.org/10.1177/1534582302001001004>
- Rossion, B., Gauthier, I., Tarr, M. J., Despland, P., Bruyer, R., Linotte, S., & Crommelinck, M. (2000). The N170 occipito-temporal component is delayed and enhanced to inverted faces but not to inverted objects: an electrophysiological account of face-specific processes in the human brain. *Neuroreport*, 11(1), 69-72.
- Roux-Sibilon, A., Peyrin, C., Greenwood, J. A., & Goffaux, V. (2023). Radial bias in face identification. *Proceedings of the Royal Society B: Biological Sciences*, 290(2001), 20231118. <https://doi.org/doi:10.1098/rspb.2023.1118>
- Rovamo, J., Virsu, V., Laurinen, P., & Hyvarinen, L. (1982). Resolution of gratings oriented along and across meridians in peripheral vision. *Investigative Ophthalmology and Visual Science*, 23(5), 666-670.
<https://www.ncbi.nlm.nih.gov/pubmed/7129811>

- Rubin, N., Nakayama, K., & Shapley, R. (1996). Enhanced perception of illusory contours in the lower versus upper visual hemifields. *Science*, 271(5249), 651-653. <https://doi.org/10.1126/science.271.5249.651>
- Russell, R. (2009). A sex difference in facial contrast and its exaggeration by cosmetics. *Perception*, 38(8), 1211-1219. <https://doi.org/10.1068/p6331>
- Sasaki, Y., Rajimehr, R., Kim, B. W., Ekstrom, L. B., Vanduffel, W., & Tootell, R. B. (2006). The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron*, 51(5), 661-670. <https://doi.org/10.1016/j.neuron.2006.07.021>
- Saxe, R., Brett, M., & Kanwisher, N. (2006). Divide and conquer: a defense of functional localizers. *NeuroImage*, 30(4), 1088-1096; discussion 1097-1089. <https://doi.org/10.1016/j.neuroimage.2005.12.062>
- Sayres, R., Weiner, K. S., Wandell, B., & Grill-Spector, K. (2010). Stimulus-dependent position sensitivity in human ventral temporal cortex. In.
- Schein, S., & Desimone, R. (1990). Spectral properties of V4 neurons in the macaque. *The Journal of Neuroscience*, 10(10), 3369-3389. <https://doi.org/10.1523/jneurosci.10-10-03369.1990>
- Schiltz, C., Sorger, B., Caldara, R., Ahmed, F., Mayer, E., Goebel, R., & Rossion, B. (2006). Impaired face discrimination in acquired prosopagnosia is associated with abnormal response to individual faces in the right middle fusiform gyrus. *Cerebral Cortex*, 16(4), 574-586. <https://doi.org/10.1093/cercor/bhj005>
- Schmalzl, L., Palermo, R., Harris, I. M., & Coltheart, M. (2009). Face inversion superiority in a case of prosopagnosia following congenital brain abnormalities: What can it tell us about the specificity and origin of face-processing mechanisms? *Cognitive Neuropsychology*, 26(3), 286-306. <https://doi.org/10.1080/02643290903086904>
- Schmidtman, G., Logan, A. J., Kennedy, G. J., Gordon, G. E., & Loffler, G. (2015). Distinct lower visual field preference for object shape. *Journal of Vision*, 15(5), 18. <https://doi.org/10.1167/15.5.18>
- Schuermans, J. P., Bennett, M. A., Petras, K., & Goffaux, V. (2023). Backward masking reveals coarse-to-fine dynamics in human V1. *NeuroImage*, 274, 120139. <https://doi.org/10.1016/j.neuroimage.2023.120139>
- Schwarz, L., Kreifelts, B., Wildgruber, D., Erb, M., Scheffler, K., & Ethofer, T. (2019). Properties of face localizer activations and their application in functional magnetic resonance imaging (fMRI) fingerprinting. *PloS One*, 14(4), e0214997. <https://doi.org/10.1371/journal.pone.0214997>
- Schwarzkopf, D. S. (2022). SamSrf 9.4 - Matlab toolbox for pRF analysis. <https://doi.org/https://doi.org/10.17605/OSF.IO/2RGSM>

- Schwarzkopf, D. S., Anderson, E. J., de Haas, B., White, S. J., & Rees, G. (2014). Larger extrastriate population receptive fields in autism spectrum disorders. *The Journal of Neuroscience*, 34(7), 2713-2724. <https://doi.org/10.1523/JNEUROSCI.4416-13.2014>
- Schwarzlose, R. F., Swisher, J. D., Dang, S., & Kanwisher, N. (2008). The distribution of category and location information across object-selective regions in human visual cortex. *Proceedings of the National Academy of Sciences*, 105(11), 4447-4452. <https://doi.org/10.1073/pnas.0800431105>
- Schyns, P. G., Bonnar, L., & Gosselin, F. (2002). Show me the features! Understanding recognition from the use of visual information. *Psychological Science*, 13(5), 402-409. <https://doi.org/10.1111/1467-9280.00472>
- Sekuler, A. B., Gaspar, C. M., Gold, J. M., & Bennett, P. J. (2004). Inversion leads to quantitative, not qualitative, changes in face processing. *Current Biology*, 14(5), 391-396. <https://doi.org/10.1016/j.cub.2004.02.028>
- Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., Rosen, B. R., & Tootell, R. B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science*, 268(5212), 889-893. <https://doi.org/10.1126/science.7754376>
- Sergent, J., Ohta, S., & Macdonald, B. (1992). Functional neuroanatomy of face and object processing: a positron emission tomography study. *Brain*, 115(1), 15-36. <https://doi.org/10.1093/brain/115.1.15>
- Sergent, J., & Signoret, J.-L. (1992). Varieties of functional deficits in prosopagnosia. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/2.5.375>
- Shah, P., Gaule, A., Gaigg, S. B., Bird, G., & Cook, R. (2015). Probing short-term face memory in developmental prosopagnosia. *Cortex*, 64, 115-122. <https://doi.org/10.1016/j.cortex.2014.10.006>
- Silson, E. H., Chan, A. W., Reynolds, R. C., Kravitz, D. J., & Baker, C. I. (2015). A Retinotopic Basis for the Division of High-Level Scene Processing between Lateral and Ventral Human Occipitotemporal Cortex. *The Journal of Neuroscience*, 35(34), 11921-11935. <https://doi.org/10.1523/JNEUROSCI.0137-15.2015>
- Silson, E. H., Groen, I., Kravitz, D. J., & Baker, C. I. (2016). Evaluating the correspondence between face-, scene-, and object-selectivity and retinotopic organization within lateral occipitotemporal cortex. *Journal of Vision*, 16(6), 14. <https://doi.org/10.1167/16.6.14>
- Silson, E. H., Groen, I. I. A., & Baker, C. I. (2022). Direct comparison of contralateral bias and face/scene selectivity in human occipitotemporal cortex. *Brain Structure and Function*, 227(4), 1405-1421. <https://doi.org/10.1007/s00429-021-02411-8>

- Silson, E. H., Reynolds, R. C., Kravitz, D. J., & Baker, C. I. (2018). Differential sampling of visual space in ventral and dorsal early visual cortex. *The Journal of Neuroscience*, 38(9), 2294-2303. <https://doi.org/10.1523/JNEUROSCI.2717-17.2018>
- Silva, M. F., Brascamp, J. W., Ferreira, S., Castelo-Branco, M., Dumoulin, S. O., & Harvey, B. M. (2018). Radial asymmetries in population receptive field size and cortical magnification factor in early visual cortex. *NeuroImage*, 167, 41-52. <https://doi.org/10.1016/j.neuroimage.2017.11.021>
- Silva, M. F., Harvey, B. M., Jorge, L., Canário, N., Machado, F., Soares, M., d'Almeida, O. C., & Castelo-Branco, M. (2021). Simultaneous changes in visual acuity, cortical population receptive field size, visual field map size, and retinal thickness in healthy human aging. *Brain Structure and Function*, 226, 2839-2853. <https://doi.org/10.1007/s00429-021-02338-0>
- Simion, F., Valenza, E., Cassia, V. M., Turati, C., & Umiltà, C. (2002). Newborns' preference for up-down asymmetrical configurations. *Developmental Science*, 5(4), 427-434. <https://doi.org/10.1111/1467-7687.00237>
- Stark, R., Schienle, A., Walter, B., Kirsch, P., Blecker, C., Ott, U., Schäfer, A., Sammer, G., Zimmermann, M., & Vaitl, D. (2004). Hemodynamic effects of negative emotional pictures - a test-retest analysis. *Neuropsychobiology*, 50(1), 108-118. <https://doi.org/10.1159/000077948>
- Stigliani. (2015). Temporal Processing Capacity in High-Level Visual Cortex Is Domain Specific. *The Journal of Neuroscience*, 35(36), 12412–12424. <https://doi.org/https://doi.org/10.1523/JNEUROSCI.4822-14.2015>
- Stoll, S., Infanti, E., de Haas, B., & Schwarzkopf, D. S. (2022). Pitfalls in post hoc analyses of population receptive field data. *NeuroImage*, 263, 119557. <https://doi.org/https://doi.org/10.1016/j.neuroimage.2022.119557>
- Taylor, V. K., Theodorou, M., Dahlmann-Noor, A. H., Dekker, T. M., & Greenwood, J. A. (2021). Eye movements elevate crowding in idiopathic infantile nystagmus syndrome. *Journal of Vision*, 21(13), 9-9. <https://doi.org/10.1167/jov.21.13.9>
- Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *The Quarterly Journal of Experimental Psychology*, 46(2), 225-245. <https://doi.org/10.1080/14640749308401045>
- Tanaka, J. W., & Farah, M. J. (2003). The holistic representation of faces. In *Perception of Faces, Objects, and Scenes: Analytic and Holistic Processes* (pp. 53-74).
- Tanaka, J. W., & Sengco, J. A. (1997). Features and their configuration in face recognition. *Memory and Cognition*, 25(5), 583-592. <https://doi.org/10.3758/BF03211301>

- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience*, 19, 109-139. <https://doi.org/10.1146/annurev.ne.19.030196.000545>
- Thibaut, M., Tran, T. H., Szaffarczyk, S., & Boucart, M. (2014). The contribution of central and peripheral vision in scene categorization: a study on people with central vision loss. *Vision Research*, 98, 46-53. <https://doi.org/10.1016/j.visres.2014.03.004>
- Thompson, P. (1980). Margaret Thatcher - a new illusion. *Perception*, 9(4), 483-484. <https://doi.org/10.1068/p090483>
- Tong, F., Nakayama, K., Moscovitch, M., Weinrib, O., & Kanwisher, N. (2000). Response properties of the human fusiform face area. *Cognitive Neuropsychology*, 17(1-3), 257-280. <https://doi.org/10.1080/026432900380607>
- Tootell, R. B., Hadjikhani, N. K., Vanduffel, W., Liu, A. K., Mendola, J. D., Sereno, M. I., & Dale, A. M. (1998). Functional analysis of primary visual cortex (V1) in humans. *Proceedings of the National Academy of Sciences*, 95(3), 811-817. <https://doi.org/10.1073/pnas.95.3.811>
- Towler, J., Fisher, K., & Eimer, M. (2017). The cognitive and neural basis of developmental prosopagnosia. *Quarterly Journal of Experimental Psychology*, 70(2), 316-344. <https://doi.org/10.1080/17470218.2016.1165263>
- Tsao, D. Y., Freiwald, W. A., Tootell, R. B., & Livingstone, M. S. (2006). A cortical region consisting entirely of face-selective cells. *Science*, 311(5761), 670-674. <https://doi.org/10.1126/science.1119983>
- Tulving, E., & Thomson, D. M. (1973). Encoding Specificity and Retrieval Processes in Episodic Memory. *Psychological Review*, 80(5), 352-373. <https://doi.org/10.1037/h0020071>
- Turati, C., Simion, F., Milani, I., & Umiltà, C. (2002). Newborns' preference for faces: What is crucial? *Developmental Psychology*, 38(6), 875-882. <https://doi.org/10.1037//0012-1649.38.6.875>
- Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Quarterly Journal of Experimental Psychology*, 43(2), 161-204. <https://doi.org/10.1080/14640749108400966>
- Valentine, T., & Bruce, V. (1986a). The effect of race, inversion and encoding activity upon face recognition. *Acta Psychologica*, 61(3), 259-273. [https://doi.org/10.1016/0001-6918\(86\)90085-5](https://doi.org/10.1016/0001-6918(86)90085-5)
- Valentine, T., & Bruce, V. (1986b). The effects of distinctiveness in recognising and classifying faces. *Perception*, 15(5), 525-535. <https://doi.org/10.1068/p150525>

- Valentine, T., & Bruce, V. (1986c). Recognizing familiar faces: the role of distinctiveness and familiarity. *Canadian Journal of Psychology/Revue canadienne de psychologie*, 40(3), 300. <https://doi.org/10.1037/h0080101>
- Valentine, T., Lewis, M. B., & Hills, P. J. (2016). Face-space: A unifying concept in face recognition research. *Quarterly Journal of Experimental Psychology*, 69(10), 1996-2019. <https://doi.org/10.1080/17470218.2014.990392>
- Valenza, E., Simion, F., Cassia, V. M., & Umiltà, C. (1996). Face preference at birth. *Journal of Experimental Psychology: Human Perception and Performance*, 22(4), 892-903. <https://doi.org/10.1037/0096-1523.22.4.892>
- van Dijk, J. A., de Haas, B., Moutsiana, C., & Schwarzkopf, D. S. (2016). Intersession reliability of population receptive field estimates. *NeuroImage*, 143, 293-303. <https://doi.org/10.1016/j.neuroimage.2016.09.013>
- Visconti di Oleggio Castello, M., Taylor, M., Cavanagh, P., & Gobbini, M. I. (2018). Idiosyncratic, Retinotopic Bias in Face Identification Modulated by Familiarity. *eNeuro*, 5(5). <https://doi.org/10.1523/ENEURO.0054-18.2018>
- Vul, E., Harris, C., Winkielman, P., & Pashler, H. (2009). Puzzlingly High Correlations in fMRI Studies of Emotion, Personality, and Social Cognition. *Perspectives on Psychological Science*, 4(3), 274-290. <https://doi.org/10.1111/j.1745-6924.2009.01125.x>
- Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, 56(2), 366-383. <https://doi.org/10.1016/j.neuron.2007.10.012>
- Wandell, B. A., & Winawer, J. (2011). Imaging retinotopic maps in the human brain. *Vision Research*, 51(7), 718-737. <https://doi.org/10.1016/j.visres.2010.08.004>
- Wandell, B. A., & Winawer, J. (2015). Computational neuroimaging and population receptive fields. *Trends in Cognitive Sciences*, 19(6), 349-357. <https://doi.org/10.1016/j.tics.2015.03.009>
- Wang, P., & Cottrell, G. W. (2017). Central and peripheral vision for scene recognition: A neurocomputational modeling exploration. *Journal of Vision*, 17(4), 9. <https://doi.org/10.1167/17.4.9>
- Wang, X., Zhen, Z., Song, Y., Huang, L., Kong, X., & Liu, J. (2016). The hierarchical structure of the face network revealed by its functional connectivity pattern. *The Journal of Neuroscience*, 36(3), 890-900. <https://doi.org/10.1523/JNEUROSCI.2789-15.2016>
- Weiner, K. S., & Grill-Spector, K. (2010). Sparsely-distributed organization of face and limb activations in human ventral temporal cortex. *NeuroImage*, 52(4), 1559-1573. <https://doi.org/10.1016/j.neuroimage.2010.04.262>

- Westheimer, G. (2003). The distribution of preferred orientations in the peripheral visual field. *Vision Research*, 43(1), 53-57. [https://doi.org/10.1016/s0042-6989\(02\)00398-x](https://doi.org/10.1016/s0042-6989(02)00398-x)
- Westheimer, G. (2005). Anisotropies in peripheral vernier acuity. *Spatial Vision*, 18(2), 159-167. <https://doi.org/10.1163/1568568053320611>
- Willenbockel, V., Fiset, D., Chauvin, A., Blais, C., Arguin, M., Tanaka, J. W., Bub, D. N., & Gosselin, F. (2010). Does face inversion change spatial frequency tuning? *Journal of Experimental Psychology: Human Perception and Performance*, 36(1), 122. <https://doi.org/10.1037/a0016465>
- Wilson, H. R., Loffler, G., & Wilkinson, F. (2002). Synthetic faces, face cubes, and the geometry of face space. *Vision Research*, 42(27), 2909-2923. [https://doi.org/10.1016/S0042-6989\(02\)00362-0](https://doi.org/10.1016/S0042-6989(02)00362-0)
- Winawer, J., Horiguchi, H., Sayres, R. A., Amano, K., & Wandell, B. A. (2010). Mapping hV4 and ventral occipital cortex: the venous eclipse. *Journal of Vision*, 10(5), 1. <https://doi.org/10.1167/10.5.1>
- Witthoft, N., Poltoratski, S., Nguyen, M., Golarai, G., Liberman, A., LaRocque, K. F., Smith, M. E., & Grill-Spector, K. (2016). Reduced spatial integration in the ventral visual cortex underlies face recognition deficits in developmental prosopagnosia. *BioRxiv*, 051102. <https://doi.org/10.1101/051102>
- Yamaguchi, M. K., Hirukawa, T., & Kanazawa, S. (2013). Judgment of gender through facial parts. *Perception*, 42(11), 1253-1265. <https://doi.org/10.1068/p240563n>
- Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, 81(1), 141. <https://doi.org/10.1037/h0027474>
- Young, A. W., Hellawell, D., & Hay, D. C. (1987). Configurational information in face perception. *Perception*, 16(6), 747-759. <https://doi.org/10.1068/p160747>
- Yovel, G., & Kanwisher, N. (2004). Face perception: domain specific, not process specific. *Neuron*, 44(5), 889-898. <https://doi.org/10.1016/j.neuron.2004.11.018>
- Yovel, G., & Kanwisher, N. (2005). The neural basis of the behavioral face-inversion effect. *Current Biology*, 15(24), 2256-2262. <https://doi.org/10.1016/j.cub.2005.10.072>
- Yue, X., Cassidy, B. S., Devaney, K. J., Holt, D. J., & Tootell, R. B. (2011). Lower-level stimulus features strongly influence responses in the fusiform face area. *Cerebral Cortex*, 21(1), 35-47. <https://doi.org/10.1093/cercor/bhq050>
- Zachariou, V., Nikas, C. V., Safiullah, Z. N., Gotts, S. J., & Ungerleider, L. G. (2017). Spatial Mechanisms within the Dorsal Visual Pathway Contribute to the Configural Processing of Faces. *Cerebral Cortex*, 27(8), 4124-4138. <https://doi.org/10.1093/cercor/bhw224>

Zhao, L., & Chubb, C. (2001). The size-tuning of the face-distortion after-effect. *Vision Research*, 41(23), 2979-2994. [https://doi.org/10.1016/s0042-6989\(01\)00202-4](https://doi.org/10.1016/s0042-6989(01)00202-4)