School of Optometry and Vision Science and School of Psychology
Cardiff University

Cinematic Scanning: Using Movie Watching fMRI to Investigate the Function and Organisation of the Visual Brain

Mason Taylor Wells

2022

A thesis submitted to Cardiff University for the degree of Doctor of Philosophy.

Supervisors:

Simon Rushton

Matt Dunn

Tony Redmond

Phil Clatworthy

Advisor:

Fergal Ennis

Thesis summary

Visually responsive brain regions and networks/pathways are traditionally examined using non-naturistic stimuli and time-consuming tasks. In this thesis, I investigate an alternative approach that uses a stimulus more closely resembling our typical visual input, a Hollywood movie.

First, I explore the chronotopy approach (Bartels & Zeki, 2004b, 2005b), which seeks to segment the brain into visually responsive regions using a combination of independent component analysis (ICA) and movie watching fMRI data. I find a simple division of the brain into visual sub-regions does not hold up for all visual areas. Instead, I find a high amount of variability across brains and across visual areas.

Second, through a series of timecourse correlation analyses and ICA, I investigate Strong et al.'s (2019) hypothesis that the two hemispheres have different functional roles in the processing of visual motion. At the group-level, I find only weak supporting evidence of the hypothesis. However, the examination of individual brains did allow me to identify some functional differences between hemispheres for a small number of subjects.

Third, I use ICA to look for evidence of a hypothesised third visual pathway dedicated to the processing of visual motion, the lateral motion pathway (Gilaie-Dotan, 2016). Group-level ICA provides evidence of a partitioning of the brain into distinct networks which support the three visual pathway model, which were subsequently found in over 50% of brains at the single-subject level.

Individual differences are clear with each analysis I performed, supporting a growing area of research that aims to examine brain function and organisation at the individual level.

Finally, I describe a psychophysics experiment designed to systematically investigate retained motion processing in patients with cortical vision loss.

Table of contents

List of tables

List of figures

Acknowledgements

I am grateful to have had so many great people around me over the last four years; your guidance and support have been instrumental in successfully navigating the trying circumstances of a "pandemic-PhD". First and foremost, I would like to thank my supervisory team:

Simon, thank you for your pragmatism, ingenuity in the unprecedented times the pandemic brought, and reassurance as you pushed me out of my comfort zone. It has been a privilege to have you in my corner. You have been a genuine mentor both in an academic and personal sense. Matt, thank you for providing unrivalled attention-to-detail, enthusiasm, and countless programming tutorials. Although a lot of the ideas we had never made it into the thesis, your guidance and expertise, particularly in the early development of this thesis, did not go unnoticed. Tony and Phil, I am thankful for your insights and tireless work to help facilitate access to patients.

To my advisor, Fergal, thank you for ensuring this thesis did not turn into an encyclopaedia.

I would like to offer a special thank you to Phoebe Asquith; without your generosity this thesis would be **a lot** slimmer than it is, and Mathew Glasser for your thorough advice with respect to all things pre-processing.

To my friends, Alex, Luke, Chlo, Jonny, Rory, and George. Thank you for making life outside of my PhD far more enjoyable than I could ever have hoped. For helping, sometimes forcing, me to have some form of work-life balance, from Sofa Sundays to Tuesday night football. You never failed to assure me, each in your own (very) different way, that I was good enough to do this. You are my friends, my therapists, and my family.

Mum, Dad and Grandad, your unwavering support and love made this possible. The values you instilled in me, knowingly and otherwise, helped me dig in when things got tough and gave me a real sense of pride in my work. For that, you are forgiven for reading no further than the introduction. To those who

are sadly no longer here to see this, Auntie Jean, and Grandma Beryl, I know you would be proud. To my brothers, Callum, Hayden, Ethen and Nathan, thank you for your humour and ability to regularly top up my self-confidence whilst simultaneously berating me for not having a "real job". I appreciate you all more than you know.

Chapter 1. Introduction

1.1 Progression of thesis

When I began my PhD in October 2017, the goal of the project was to investigate residual visual function in individuals with homonymous visual field defects (HVFDs), specifically hemianopia (a loss of one half of the visual field in the same relative position in each eye) following brain damage. I separated residual vision into two broad aspects: conscious and unconscious visual processing. The type of conscious vision I was interested in was macular sparing, in which a patient's central vision is preserved in the presence of damage to the contralateral occipital lobe. Unconscious visual processing refers to the phenomenon in which the behaviour of a patient can be changed by the presence of stimuli of which they may be unaware of, in their blind hemifield. To this end, the thesis had two clear research questions to be investigated:

1) Is macular sparing "true" vision or can it be explained by eye movements and/or inappropriately scaled stimuli used during perimetry?
2) What non-conscious visual abilities do individuals with HVFDs possess, and which underlying neurological pathways serve these non-conscious visual abilities?

Visual field defects, such as hemianopia, are characterised following a systematic assessment of an individual's visual field. This assessment, known as perimetry, involves the detection of the presence of a test light stimulus presented to varying locations of the visual field. To investigate the phenomenon of macular sparing and the potential role of eye movements and fixation (research question 1), a perimetric study was designed in which three techniques were utilised. The differences between the techniques were the degree to which they compensate for eye movements. 1) A standard perimetry system, the Humphrey Field Analyzer, was to be used to characterise the patient's visual field loss (no compensation). (2) Another commercially available perimetry system, the Macular Integrity Assessment (MAIA), was used to compensate for eye movements prior to stimulus presentation. 3) A gaze-contingent perimetry system

which used simultaneous non-invasive eye-tracking to present test stimuli to specific regions of the retina, regardless of the gaze position. In addition, to examine the role of stimulus size I planned to scale the perimetric stimuli in line with a published cortical magnification factor.

In addition to macular sparing (conscious visual processing), I also explored non-conscious visual processing. For some patients, unconscious stimuli can become conscious when they move. This was demonstrated by the early work of Riddoch (1917), who showed that patients with damage to primary visual cortex (V1) were able to detect motion within their blind fields. These patients were conscious of the motion but could not characterise any other attributes of the stimulus, e.g., shape or colour. Other forms of residual vision lack any conscious percept. Simply put, the patient's behaviour is changed by the blind field stimulus without the patient ever being aware of its presentation (e.g., Weiskrantz et al., 1974c)

I developed three psychophysical tasks to systematically investigate different residual visual functions, each tapping into a different level of processing of stimuli within the blind field (research question 2). I planned to measure three different types of responses: explicit, implicit, and reflexive. The background to, and the experimental design for, these tasks are covered in detail in Chapter 6 of this thesis. For reference, I summarise the three tasks and the features they were designed to assess below.

1. Explicit responses require the subject to make a direct judgement about a blind field stimulus. The temporal motion detection (TMD) task is commonly used to assess these responses, in which a stimulus, often either a Gabor patch (Ajina, Rees, et al., 2015) or an aperture of dots (Ajina & Bridge, 2018), is presented to the blind field. In one of two temporally separated intervals, the stimulus contains motion. The subject is then instructed to indicate which of the two intervals they believed contained the motion. Here, I planned to use an aperture of black and white dots that were either static or coherently moved left or right.
2. Implicit responses require the subject to make a judgement about a probe in the seeing field, whilst a second stimulus is placed in the blind field. The effects of the blind field stimulus on the seeing field stimulus are then measured. I planned to use the flow-parsing task (Rushton & Warren, 2005;

Warren & Rushton, 2009a) in which the subject is instructed to fixate centrally and judge whether a probe dot located in their seeing field moved towards or away from fixation. While the probe moves a contracting or expanding radial flow pattern is presented simultaneously to the opposite (blind) visual field. Previous work (e.g., Warren & Rushton, 2009a) has shown that in healthy brains the presence of an expanding flow field induces an inward lateral component of motion to the probe. It has been hypothesised that the flow-parsing effect relies on a non-occipital route to hMT+ (Rushton et al., 2018). If this is correct, then a flow-parsing effect may be evident when the flow field is presented to the blind field.

3. Reflexive responses do not require the subject to make any judgments about a stimulus, instead unconscious responses to blind field stimuli are measured. In the present study, I planned to measure the optokinetic nystagmus (OKN) response. When a large, continuously moving stimulus is presented, OKN minimises the retinal slip (motion on the retina) produced by the stimulus. First, the eyes follow the stimulus in the direction in which it is moving (OKN slow phase), once the maximum rotation is reached the eyes rapidly move in the opposite direction (OKN quick phase) to re-centre the image and hold it steady on the retina (Metz, 1983). To assess the OKN response, we planned to present a gaze-contingent sinusoidal grating exclusively to one visual hemifield at a time (this would be alternated between the seeing field and the blind field) while subjects were instructed to fixate the centre of the screen. Eye movements were to be measured using an eye-tracker while the grating moved upwards, downwards, left, or right.

The first question we aimed to investigate was whether there was a dissociation between performance on these tasks. For example, do some patients show a flow-parsing effect while performing at chance on the TMD task? Or are the responses all-or-nothing? e.g., if a patient is above chance at TMD, do they also show a flow-parsing effect and an OKN response?

Many pathways have been hypothesised to explain retained motion perception (see Cowey, 2010 for review), but the recent research using the TMD and diffusion-weighted MRI suggests the residual

function is reliant on a geniculo-hMT+ pathway (Ajina, Pestilli, et al., 2015; Bridge et al., 2008, 2010). Here, I had aimed to build on this work and investigate the neurological basis for each of the three potential residual functions described above (research question 2). Specifically, does a single pathway subserve all three, i.e., is the pathway itself all-or-nothing? Following Bridge et al. (2008, 2010) and Ajina et al. (2015) a combination of diffusion-weighted MRI and tractography was to be employed to assess how these residual visual functions map onto the structural pathways.

To optimise data collection, I planned to try a novel technique, chronotopy (see section 1.4.2 for more details), introduced by Bartels and Zeki (2004b, 2005b), to quickly localise visually responsive brain regions. Traditionally, the identification of functionally specialised brain regions has relied on anatomical markers and/or functional localisers. However, these approaches to divide the brain into its constituent parts are often respectively inaccurate and time-consuming.

Chronotopy applies independent component analysis (ICA) to fMRI data collected as observers watch a (~20min) movie clip and segments the brain into functional regions (similar to traditional resting state analyses but on a finer scale). Bartels and Zeki (2004b, 2005b) showed the components derived from ICA both anatomically and temporally corresponded across brains, and some coincide with known visual areas, e.g., V1, V5/MT. This holds out the possibility of identifying multiple visual areas in a single, relatively short (< 20 minutes) scanning session. The approach would not only be advantageous from a time perspective, but it is arguably more engaging than a battery of localisers, it does not require controlled fixation (subjects freely view the movie), and the naturalistic broadband visual stimuli captures the neural processing in a more ecological context.

In summary, the original aim of this thesis was to develop a robust set of tools which could then be used to investigate the prevalence, co-occurrence and neurological basis of the counterintuitive phenomenon in which patients with cortical vision loss retain some residual visual function.

When the COVID-19 pandemic hit in March 2020 and the UK entered its first lockdown, access to, and recruitment of, suitable patients ceased. This, combined with an extensive NHS ethical approval

process (more than 24 months in total), resulted in an inability to recruit patients for the project. Without the knowledge of when the pandemic was going to end, I was required to initially temporarily, later permanently, change track and switch to using archival neuroimaging data acquired from subjects during movie watching. With this deviation from the original PhD came a deviation in the questions I aimed to investigate, resulting in a new set of thesis aims.

These revised aims grew from the original topic area, and although each is addressed with the archive dataset, the tools/pipelines were developed so that should the opportunity have arisen to resume face-to-face research, I could have exploited them in patient studies. Each analysis chapter now stands on its own, but all were motivated by the original PhD topic (residual vision and hemianopia) and can be adapted for use with the clinical population I had originally aimed to investigate.

In the first instance, I utilised the archive data to validate the chronotopy approach. However, with the continued inability to acquire patient data, I used the data to examine lateralisation of function and the presence of an additional visual pathway.

As most of the work rested on the use of ICA, first, I systematically explored the sensitivity of the approach to different pre-processing pipelines, culminating in the creation of an optimal ICA pipeline (described in section 2.2). Next, using this optimal pipeline, I investigated whether movie watching, and ICA can be used to robustly identify visually responsive regions in the brain. This constituted the first revised aim, and the work is reported in Chapter 3.

The second revised aim built from work by Strong et al. (2019), who reported that left and right hemisphere hMT+ have different functions, there is a degree of lateralisation. They found that right hMT+ processes both contra- and ipsi-lateral visual fields, whereas left hMT+, in common with most other visual areas, only processes contra-lateral stimuli. Although evidence came from healthy subjects, this piqued my interest for two reasons: (1) lateralised function of hMT+ could potentially explain some residual visual function, and (2) if correct it might generate hypotheses about the relative incidence of

residual vision in left vs right hemisphere cortical damage. In Chapter 4 I explore whether lateralised function of hMT+ is evident in a healthy population.

The techniques developed could then be used to examine its prevalence in patients and more importantly assess its potential implication in residual motion perception. For example, consider a patient with left hemisphere striate cortex damage and therefore a right visual field deficit. If right hemisphere hMT+ processes both the contra- and ipsi-lateral visual field, then a visual stimulus located in the right (contra-lesional) visual field could be processed by hMT+ in the intact hemisphere and thus explain residual visual motion perception despite upstream damage to right V1.

The third revised aim, and thus fifth chapter, provides evidence compatible with a visual pathway specialised for motion processing. Originally, I aimed to examine the hypothesis that a geniculo-hMT+ projection explains residual motion vision in the blind hemifield in hemianopia. This hypothesis suggests that the primary motion processing complex has its own additional projection that bypasses occipital cortex. Gilaie-Dotan (2016) went beyond this idea to suggest that the traditional two pathway (ventral and dorsal) model of visual processing is wrong, that there is actually a third visual pathway specialised for processing visual motion, the "lateral motion pathway". I looked for evidence of the lateral motion pathway using ICA and the archive movie watching data.

The original project aims, although now more peripheral, remain interweaved in the methods developed throughout this thesis. As such, the background to, and the experimental design for, the original project aims are covered in the sixth chapter of this thesis.

As noted above, most of the work presented in this thesis relied on the use of ICA. The technique is widely understood with many excellent reviews available (e.g., Calhoun et al. 2003; see also Tharwat, 2018). In short, ICA is a data-driven technique which can be used to decompose a task-based or resting-state fMRI dataset into its constituent independent sources (Beckmann et al., 2005; Mckeown et al., 1998). A key aspect of ICA is that it does not rely on an a priori seed region or the standard general linear model. As such this model-free technique can be used to identify unexpected patterns of connectivity across the

entire brain, functional networks (S. M. Smith et al., 2009) and discrete functionally specialised regions (e.g., Bartels & Zeki, 2004a). However, the technique of applying ICA to movie watching fMRI data is not so broadly familiar. Therefore, in the following sections I provide some background to the stimulus employed, movie watching, and a brief overview of the research that has led to its inclusion in this thesis. The remaining details are found in the corresponding chapters where it is most appropriate. For completeness, I also provide an overview of the visual system, the blood-oxygen level dependent (BOLD) signal used in fMRI to infer neuronal activity, and the alternative approaches to mapping visual areas of the brain.

1.1 Background: The visual system

The transfer of sensory information throughout the brain follows defined streams, or pathways (Hudspeth & Logothetis, 2000). The geniculocortical visual pathway is the primary one and it has been suggested to have a special role in conscious perception (for detailed discussion see Leopold, 2012). Other pathways projecting from the nuclei of the thalamus to extrastriate areas, thus bypassing V1, have been implicated in non-conscious visual perception. In the following sections I briefly describe the geniculocortical pathway and the relationship between visual information processing and cortical organisation. The role of the other pathways is covered in later chapters of this thesis (see also Ajina & Bridge, 2017 for an overview of these pathways).

Binocular vision consists of two visual hemifields, left and right. The right binocular hemifield takes visual input from the temporal visual field of the right eye and the nasal visual field of the left eye. The left binocular hemifield encompasses the nasal visual field of the right eye and the temporal visual field of the left eye (Purves, Augustine, Fitzpatrick, Katz, LaMantia, McNamara, & Williams, 2001). These hemifields translate directly onto the retina, with the nasal and temporal portions of the retina in the left eye receiving input from the left visual field and right visual field respectively (vice versa for the right eye).

Phototransduction in the retinal cells converts light energy to electric activity which, via connections with bipolar cells, induces action potentials in retinal ganglion cells (RGCs). The long axons of RGCs form the optic nerve and leave the back of the eye through the optic disc. RGC axons that sit in the nasal portion of each retina cross at the optic chiasm, while the axons of RGCs from the temporal retina do not cross. The optic tracts then terminate at the left or right lateral geniculate nucleus (LGN) of the thalamus. From the LGN, the optic radiations project to the striate cortex of the occipital lobe, the seat of early visual processing. Information is then propagated along a distributed, but organised, network of specialised brain areas.

For over a century it has been known that such is the organisation of the visual system, the spatial arrangement of the retinal image is preserved to varying degrees throughout the visual cortex (Henschen, 1893; Holmes, 1918; Inouye, 1909). Spatially adjacent cells in the retina project to spatially adjacent cells in the cortex, resulting in a retinotopic organisation. For example, primary visual cortex (V1) contains a representation of the contralateral visual field, with other visual areas containing either a quadrant or other more complex visual field representation. The retinotopic organisation of the visual cortex (i.e., the spatial organisation of the retina is replicated in the cortex) was first characterised in patients with lesions (Holmes, 1918) and a map describing the relationship between the visual field and cortex was produced (I discuss this in detail below).

Horton & Hoyt (1991) correlated anatomical MRI with visual field defects arising from occipital lobe damage, finding Holmes' maps of the striate cortex to be incompatible with the visual field deficits displayed by their patients. Horton and Hoyt later modified Holmes' retinotopic map and proposed a map of their own that both extended the area of cortex dedicated to central vision and reduced the area subserving peripheral vision. This revised map was later reaffirmed by McFadzean and colleagues who studied a larger sample of patients (McFadzean et al., 1994). It is now well established that there is not a 1:1 scale between visual space and the cortical representation. The foveal region of the retina is

represented over a disproportionately large region of the visual cortex. It is thought that the central 10°–30° of the visual field is distributed over 50-60% of the posterior striate cortex (McFadzean et al., 1994).

The organisation of the visual cortex can be explored (and mapped) by investigating its representation of visual space (retinotopy) or its sensitivity to different visual features (functional localisation). Several techniques can be implemented when mapping the visual system, I discuss the four main approaches in section 1.3.

1.2 Background: What is the blood-oxygenation level dependent (BOLD) signal?

Although anatomical studies have played a significant role in developing our understanding of the organisation of the brain, the advancement of functional MRI (fMRI) has been instrumental in the progression of the field. In section 1.3 I describe the different approaches that have been developed to map the brain, from lesions studies with injured soldiers to sophisticated fMRI mapping techniques. Given that much of this literature relates to fMRI research and the analyses presented in this thesis were performed on fMRI data, below I provide a very brief overview of what MRI measures and what the MR "signal" refers to. It is important to note that fMRI, although the dominant functional imaging technique, measures neural activity indirectly.

fMRI exploits the fact that electrical impulses are the language of the human brain. When an area of the brain is "activated" the neurons subserving the process increase their firing rate resulting in an increased requirement of oxygen and glucose. This local need for an increased supply of nutrients is achieved through neurovascular coupling (Hall et al., 2016), in which blood flow and oxygenation is modulated to meet the demand. The haemodynamic response allows active neurons to receive greater levels of oxygen, via oxygenated haemoglobin (oxyhaemoglobin). Oxyhaemoglobin is diamagnetic, while deoxyhaemoglobin (deoxygenated haemoglobin) is paramagnetic. As deoxyhaemoglobin decreases, the blood-oxygenation level dependent (BOLD) signal, as detected by MRI, increases. Therefore, an increase in the BOLD signal is interpreted to reflect a local increase in blood flow and oxygenation, which in turn is

interpreted to reflect neuronal activity (see Lecrux and Hamel, 2016 for extensive review of the relationship between neuronal activity and the BOLD signal). However, BOLD signal changes do not capture purely neuronal signal, they also contain "noise".

Noise, with respect to the BOLD signal, refers to confounding non-neuronal signals elicited by sources that are not (usually) of experimental interest (e.g., heartbeat, head motion and respiration). Noise within the BOLD signal is inherent in all fMRI studies and it is impossible to remove all noise components from these signals. However, as Caballero-Gaudes and Reynolds (2017) discuss in detail there are techniques that can be applied to remove some sources of noise from the signal, i.e., the signal can be "cleaned". For example, physiological noise components can be removed by taking external cardiac signal recordings and then regressing them out the data, whereas data-driven strategies can be implemented to remove head-motion signals (e.g., ICA-AROMA as discussed in Chapter 4).

1.3 Background: Brain mapping techniques

The visual cortex can be sub-divided into regions, or areas, with specialised functions, i.e., it can be mapped. Over the decades numerous techniques have been developed to map the visual brain, the central aim of each is to "specify in as much detail as possible the localisation of function in the human brain" (Savoy, 2001, p.10). In short, the four commonly used approaches each involve mapping the brain based upon either the:

1. relationship between cortical damage and functional deficits (lesion studies)
2. difference in the microarchitecture of cell bodies of brain tissues (cytoarchitecture)
3. functional specificity of distinct areas (functional localisation)
4. visual field representation of distinct areas (retinotopy).

Each technique works from the assumption that visual areas can be defined as discrete entities based upon differential characteristics; either structural or functional. In the following section I briefly

describe the four main brain mapping techniques that are used (to varying degrees) in modern cognitive neuroscience.

Although non-invasive technologies (e.g., fMRI and PET) have played a major role in the advancement of brain mapping research, our understanding of brain organisation is not solely-dependant on it. Arguably, the most direct approach to mapping discrete areas is to study patients with focal brain lesions. Lesion studies are informative because they allow us to closely assess the relationship between the integrity of brain tissues (e.g., grey matter) and impaired function (e.g., cortical vision loss). In the absence of neuroimaging, early research investigated the visual outcomes of small survivable gunshot wounds penetrating the brain (Holmes, 1918). Holmes, like Inouye (1909) before him, estimated the trajectories of the bullets penetrating the soldier's brain and assessed their visual fields. Although limited by the technologies available at the time, Holmes generated a topographical map by correlating the observed visual field defect with the anatomical location of the damage. While crude by modern standards, Holmes' work highlighted a predictable relationship between the location of striate cortex damage and the location of the observed visual field deficit. Subsequent work has investigated the relationship between brain damage and the loss of specific visual functions. For example, Rizzo et al. (1995) reported a dissociation between the loss of shape and motion perception (akinetopsia) depending on the location of the brain damage.

With the advent of non-invasive neuroimaging techniques, high-resolution structural MRI can now be used to quantify the size and location of the lesion. From these data, researchers can infer which brain regions have likely been compromised as a result of the damage. Alongside lesion identification, behavioural/visual tasks can be administered to assess which functions are affected by the damage, allowing for the relationship between symptoms, task performance and structural damage to be measured. However, there are limited circumstances in which lesions studies can be conducted. Moreover, as with any mapping approach, lesions studies are not without limitations. I discuss these below.

Lesion-based mapping works from the premise that information about the location of the damage and task performance is sufficient for identifying which brain regions are critical for performing a given function. For example, a common approach is to superimpose individual lesions onto a standardised brain template and examine the overlap between patients with similar functional deficits (e.g., Bouvier & Engel, 2006). However, the anatomical variation between brains, combined with the often-heterogeneous cortical damage exhibited by patients, makes it difficult to establish a causal relationship between the site of the damage and the impaired function. Even if the lesions are highly localised with surrounding tissue seemingly unaffected, the function of the neighbouring areas may also be compromised due to being disconnected from the damaged region. With structural MRI alone it would be impossible to know whether the structurally intact region was also functionally intact (Rorden & Karnath, 2004). This is particularly problematic for functions that require the involvement of multiple regions organised in a distributed network. Structural scans may reveal damage to be localised to one region, with others seemingly intact. However, focal lesions can lead to widespread metabolic and blood-flow changes throughout the brain, including distal areas within the same hemisphere as well as areas in the contra-lesional hemisphere. This is known as diaschisis (see Carrera and Tononi, 2014 for a detailed review of this phenomenon). Diaschisis is characterised by abnormalities in the neuronal activity of intact brain regions distant from the lesioned area, or by changes in the strength and/or direction of connections between the regions. Although diaschisis occurs simultaneously with the insult, the effects are transitory, with the metabolic and blood flow changes recovering over time (Carrera & Tononi, 2014). This temporary neuronal dysfunction is a key limitation of using lesions to map the brain, since it makes the distinction between the recovery of post-lesion function and functional reorganisation difficult to discern.

Another anatomically principled mapping approach is cytoarchitectonic parcellation. Cytoarchitecture refers to the differences in the size, density, distribution, and shape of cell bodies between areas of the brain (Brodmann, 1909). From this parcellation of the cerebral cortex distinct areas can be identified and atlases can be defined. As with lesions studies, cytoarchitecture parcellation was described

before the advent of sophisticated imaging techniques. Following his seminal monograph at the start of the 20th century, Korbinian Brodmann is often considered the founder of anatomical brain mapping (Zilles, 2018), and the review of Zilles and Amunts (2010a), published in Nature Neuroscience, provides an excellent summary of the influence of Brodmann's illustration.

In an attempt to understand the relationship between structure and function of the brain, Broadman systematically examined the cellular structure of cortical tissues, constructing a map of the brain segregated into 43 areas. Such was the time, Brodmann's map was presented as an illustration of the lateral and medial surfaces with each area denoted by a number. Each area, now referred to as a Brodmann's Area (BA), was characterised by a specific cytoarchitecture. For example, Brodmann parcellated the occipital cortex into three areas, BA17, 18 and 19. BA17 covers the cortex of the calcarine sulcus as well as the occipital pole and corresponds to the primary visual cortex (V1), while the surrounding extrastriate area (BA18) represents secondary visual cortex (V2). The more anterior BA19 covers both the dorsal and ventral aspects of the extrastriate cortex, is much larger than the other two visual areas, and appears to encompass all visual areas outside of V1 and V2.

Early attempts to develop an atlas system to aid the localisation of functional activation in standard space adopted Brodmann's architectonic maps (e.g., Talairach and Tournoux, 1988). However, despite Brodmann's nomenclature being adopted by Talairach and Tournoux, the areas included in the atlas were not defined by cytoarchitecture, but instead Brodmann's 2D illustration of the cortical subdivisions were registered to their atlas brain.

This leads us to two, quite obvious limitations of using Brodmann's cytoarchitectonic parcellation to localise visual areas. First, the area maps lacked the sub-divisions we now believe to be present in extrastriate areas (hMT+, V3A, V4 etc.). Second, Brodmann only illustrated the left hemisphere of a single brain, which, as with lesion studies, raises an issue of inter-subject variability (Amunts et al., 2004). Moreover, as discussed in Zilles and Amunts (2010b), Talairach and Tournoux's atlas was subjected to similar criticisms since the BA maps were based on identifying corresponding sulci patterns between

Brodmann's illustrations and the atlas brain. The key criticism is that such an approach hinges on the assumption that there is a demonstrable association between cortical area boarders and sulci patterns, an assumption questioned by Brodmann himself.

More recently in vivo structural imaging has been combined with high-resolution cytoarchitectonic observations of brains post-mortem (for detailed review of these developments see Amunts and Zilles, 2015). One such development, the Julich-Brain atlas, is the culmination of over 40 projects, in which the cytoarchitecture of 24 post-mortem brains were used to generate maps for nearly 250 cortical areas and subcortical nuclei (Amunts et al., 2020). Unlike Brodmann's original map and Talairach and Tournoux's subsequent adoption, the Julich-Brain atlas is probabilistic. The map was generated by superimposing the cytoarchitectonically defined areas from different brains onto a standard brain template, thus each area map captures the inter-subject variability. Each spatial location (i.e., voxel) is assigned a probability of belonging to a given area. If the voxel was allocated to the area in all brains, it is assigned a probability of 1 (100% overlap between all brains). If the voxel was not allocated to the area in any brain, it is assigned a probability of 0. This is referred to as a full probability map (FPM) since it conveys the likelihood of each voxel belonging to each area. This approach to mapping leads to an overlap between adjacent regions, i.e., more than one area was localised to a single voxel. A complementary approach to the FPM is to generate a maximum probability map (MPM), which rather than conveying variance between brains in the location of each area, summarises the topographical organisation of the brain. Each voxel is allocated to a single area, the area with which it had the highest probability of belonging to. The MPM reduces the complexity of the map, resulting in unambiguously defined borders and therefore a more intuitive parcellation (Figure 1). Moreover, ROIs can be defined from the MPM of a specific area, from which study-related data can be extracted (e.g., cortical thickness or the fMRI timeseries). Note probabilistic mapping (both FPM and MPM) can be applied to any brain mapping technique, specifically, but not limited to, retinotopy. I discuss the retinotopic mapping technique below.

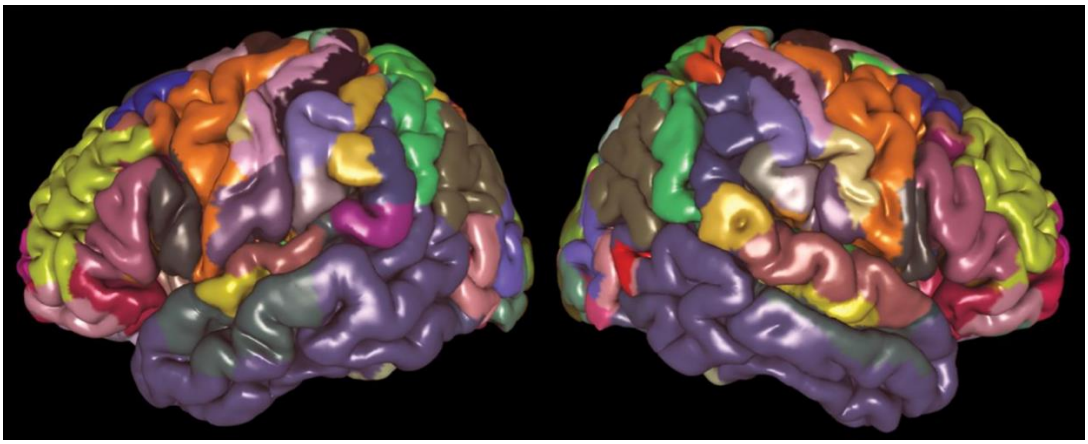

Figure 1. Adapted from Amunts et al. (2020). Cytoarchitectonic Maximum Probability Maps of the Julich-Brain atlas. Each distinct area is colour coded.

So far, I have discussed how structural MRI can be used to segment the brain, either through examining cytoarchitectural differences, or the relationship between cortical damage and subsequent visual impairment. In the following section I describe the two main fMRI approaches to mapping the brain: retinotopy and functional localisation. The former involves identifying individual visual areas based on their representations of visual space (e.g., left visual hemifield), the latter involves identifying regions that preferentially respond to a specific type of visual stimulus (e.g., motion). Both approaches allow for finer delineation of the brain into distinct areas (e.g., V1, V5/MT) compared to the more gross parcellations derived from lesion- and cytoarchitecture-based approaches.

Functional localisation refers to the broad approach in which a specialised region of the brain can be identified with stimuli designed to elicit neuronal responses in that region alone. The assumption that governs functional localisation is that the brain contains areas that serve different cognitive functions. In the lesion studies section above I reference how inferences about brain organisation can be derived from studies of cortical damage, this is effectively functional localisation in a coarse, more uncontrolled manner. Indeed, functional localisation as an approach was born from early lesion studies which demonstrated that specific lesion locations can lead to equally specific impairment. As with cytoarchitecture maps, the rapid development of non-invasive imaging technologies, such as fMRI, means that functional specialisation of the human brain continues to be a major topic in cognitive science today. It is beyond the scope of this

thesis to review the developments of the functional localisation approach, or the many discrete areas that have been defined through several years of research. However, I do provide an example of the specificity of the brain and how such sensitivities can be exploited by fMRI to parcellate the brain.

Seeing a face is a common experience for most people. Yet the face itself is a complex stimulus, made up of individual features, the perception of which gives rise to our ability to recognise who that face belongs to. Jenkins et al. (2018) estimated that on average humans can recognise around 5000 faces, yet healthy individuals can perform this cognitive task seemingly effortlessly and often rapidly. Therefore, and perhaps not surprisingly, some of the most researched specialised brain regions are the face-specific areas of the cerebral cortex.

The occipital face area (OFA; Sergent et al., 1994) has been shown to be responsible for the processing of individual features of the human face, while a separate region, the fusiform face area (FFA), has be linked with processing the face as a "whole" (Kanwisher et al., 1997). Thus, it has subsequently been implied that the FFA has facial recognition properties (Gobbini & Haxby, 2007; Haxby et al., 2000), while the OFA simply processes whether or not the combination of visual inputs represents real face parts regardless of their configuration (Liu et al., 2010). The above is just one example of how complex cognitive processes, often taken for granted, are facilitated by distinct and specialised regions. This specificity can be exploited for a range of cognitive functions, including visual processing of motion (Zeki et al., 1991) and the human body (Downing et al., 2001). However, the demonstrable specificity of the human brain makes functional localisation problematic when adopted as a brain mapping technique for two reasons: time and money.

To map the visual cortex based solely on functional localisation, would require the generation of a set of stimuli that specifically target each individual area independently. These stimuli would then need to be implemented in such a design as to allow for the elicited activation in the region of interest to be separated from a control condition, e.g., a block of faces contrasted with a block of scrambled faces. This approach is suitable if only a handful of areas are of interest but would take many hours of scan time to

map each individual visual area. This in turn drives the cost of the project up. Similarly, such an approach may not be suitable for patients who have a low threshold for fatigue, or children who may struggle to keep still in the scanner for prolonged periods of time. Moreover, a battery of localiser tasks is not only costly from a time (and cost) perspective but owing to their artificial nature they can also lack ecological validity (Sonkusare et al., 2019). This is a key limitation of current brain mapping techniques and forms the main consideration for the alternative approach described in this thesis (movie watching and ICA, see section 1.4.2).

Retinotopic mapping is often considered the 'gold standard' approach for identifying visual areas (Bridge, 2011). In contrast to functional localisation, retinotopic mapping can be performed in a single, relatively short (< 10 minutes) scan session, making it advantageous from a time perspective. Retinotopy segments the brain based on each area's eccentricity and polar angle representations (Sereno et al., 1995). Retinotopy is a measure for each voxel/vertex of the stimulus location within the visual field that generates the largest BOLD response. Retinotopic maps are defined with respect to an individual's fixation point. Traditionally, subjects maintain central fixation while a high contrast checkerboard (ring or wedge shaped) stimulus is presented at changing positions in the visual field, mapping both eccentricity (distance from fixation) and visual angle (orientation with respect to fixation) in a single scan (DeYoe et al., 1994; Engel et al., 1994, 1997).

Retinotopic mapping identifies the correspondence between visual field stimulation and the corresponding cortical activation. By combining eccentricity and visual angle measurements, the position in the visual field that elicits a maximal response for each voxel can be estimated.

Such is the organisation of the visual system, different retinotopic areas abut one another. Therefore, the boundaries of individual cortical areas are identified by the reversals in polar angle representation. The reversals represent a change in the visual field representation (e.g., from one meridian to the next). For example, the superior and inferior boundaries between V1 and V2 are defined by the reversals at the lower and upper vertical meridians, respectively. The retinotopic organisation of these early

visual areas is well-established with research continuing to shed light on the retinotopic organisation of higher visual areas (for review see Wandell & Winawer, 2011).

When comparing retinotopic mapping with functional localisation, it is important to consider the relationship between retinotopy and function. At present, it appears that multiple retinotopically defined maps overlap with category-selective regions. For example, the scene-selective region in the ventral visual cortex contains a pair of retinotopic maps, referred to as parahippocampal areas one and two (PHC1 and PHC2), with each containing a representation of the contralateral visual field (Arcaro et al., 2009). This paired-retinotopic-map characteristic is exhibited in other category-selective regions (i.e., regions that are selective to faces, scenes or objects), suggesting visuospatial coding[1] provides a "ubiquitous scaffolding" for most high-order visual functions (Groen et al., 2022). However, not all category-selective regions contain a retinotopic map. For example, the OFA shows very little overlap with any retinotopically defined maps (Silson et al., 2013). It is therefore common to define early visual areas by their representation of the visual field and higher visual areas by their sensitivity to specific visual attributes (Rosenke et al., 2021).

As with the alternative approaches described above, retinotopic mapping is not without limitations either. First, in order to define the visual areas, the subject is required to maintain fixation throughout the scan. This is problematic for studies investigating the visual function of patients, for example individuals with unstable fixation or cortical vision loss. Second, although early visual areas can be robustly identified in almost every subject, reliable identification of extrastriate areas is more challenging (Wandell & Winawer, 2011b). The lack of a map for each area, in each brain, is not necessarily evidence to reject the hypothesis that a map does indeed exist, but rather researcher seek to understand why it is absent in the other brains. The failure to identify a map could be a result of inter-subject variability (as discussed throughout this thesis) or simply a limitation of the data quality/analysis pipeline.

---

[1] Visuospatial coding includes retinotopy and visual field bias. The former is outlined above, while the latter refers to a region having a preferential visual field location, in the absence of a clear retinotopic map (see Silson et al., 2013).

In the absence of retinotopic data, it is still possible to define ROIs from a probabilistic atlas derived from the retinotopic organisation of an independent sample (e.g., Wang et al., 2015). We adopted this approach in this thesis for reasons described above, namely a lack of retinotopic data. The published topographical maps we used are described in detail in section 2.1.5. As with the Julich-Brain atlas, atlases defined through leveraging information about areas from an independent sample are built on the assumption that neuroarchitecture is highly consistent across individuals. However, there is a wealth of literature that speaks to individual differences (Gordon et al., 2017). For example, there is significant variation in the size of even the most widely studied visual areas. Benson et al., (2021) found the size of retinotopically defined V1-V3 can vary by more than three-fold across brains. This variability seen in the early visual areas is greater than the variation in the size of the cortex as a whole. This raises an important limitation of atlases; the definitions could be incorrect. The exact size and location of an area may differ between brains; thus, an atlas does not provide a 'ground truth'. This is an important limitation, and it is revisited at various points throughout this thesis.

Both the retinotopic mapping and functional localisation approaches have limitations, perhaps the most significant being they require the use of stimuli incongruent with our natural visual experience presented in highly controlled experimental designs. However, as outlined in section 1.1, there are alternative mapping techniques that use more naturalistic stimuli. I discuss these below.

1.4 Movie watching paradigms

The use of naturalistic stimuli, such as movie clips, in neuroimaging is not new. Paediatric neuroimaging studies are encouraged to utilise movies in their protocol because the engaging stimulus improves compliance and reduces head-movement in the scanner (Raschle et al., 2009).

Movies not only improve compliance, but they can serve as a "middle-ground" between traditional, highly controlled experimental designs and unconstrained resting-state paradigms. Movies have an element of structure; all subjects watch the same movie. In addition, the observers eye movements tend to be

highly correlated when watching the same movie (Davis et al., 2021; Hasson et al., 2008). Therefore, activity (at least in the early sensory cortex) is considered to be stimulus driven, time locked and thus can be compared across brains. The stimulus itself can contain a dynamic array of visual elements, from faces to complex motion. Some elements occur in isolation for an extended period (e.g., faces are the dominant feature when two characters are conversing), and then co-occur with a number of different elements (e.g., a scene in which the characters are moving through a complex environment). The inclusion of multiple visual features that come in and out of each scene creates a visually rich stimulus that more closely resembles our natural visual experience. The intersection of naturistic stimuli and analyses which are sensitive to, perhaps subtle, functional differences at the individual level allow for a more ecological investigation of a range of cognitive neuroscience research questions.

The use of movie watching in functional neuroimaging studies accelerated in 2004 following the publication of a series of landmark studies by two independent groups (Bartels & Zeki, 2004b; Hasson et al., 2004). Both groups used different approaches to demonstrate that the differential BOLD signal responses of functionally specialised regions, evoked during movie watching, were correlated across observers. I discuss both approaches below.

1.4.1 Intrersubject correlations

Hasson et al. (2004) introduced a technique for investigating the age-old philosophical question "is what you see the same as what I see?". Using movie watching fMRI data, Hasson performed intersubject correlations (ISCs) on the BOLD signal timecourses of corresponding voxels across brains. For example, the timecourse of voxel 1 in Subject A versus the timecourse of voxel 1 in Subject B. The movie stimulus evoked similar functional responses across observers, with almost 30% of the cortex showing significantly high ISCs, including almost the entire visual cortex. This highly correlated activity between corresponding voxels indicates a vast proportion of the cortex responds stereotypically across subjects during free viewing of a movie. Hasson et al. (2010) concluded that individual brains "tick together" when presented

with the same visual stimulus and these synchronised responses are time-locked to the elements of the stimulus.

The broad scope of Hasson et al.'s approach (movie watching and ISCs to assess synchronisation of activation across brains) has subsequently led to its use in investigating the neurological basis of a range of conditions. Studies into attention-deficit/hyperactivity disorder (ADHD) have benefitted from using a movie clip in which a complex social interaction is depicted. For example, Salmi et al. (2020), reported that areas in the attention networks and sensory cortex were desynchronised in ADHD subjects compared to controls. Further, the approach has also been used to identify differences in synchronicity between controls and individuals with depression (Guo et al., 2015) and psychosis (Mäntylä et al., 2018). In Chapter 4 I adopt a similar approach in which I assess the synchronicity of activity across hemispheres in the same brain. In addition, I perform ISCs between the timecourses of discrete visual areas, specifically hMT+ and V1, to investigate whether they have similar functional roles across brains.

1.4.2 Independent component analysis (ICA)

Bartels and Zeki (2004b, 2005b) went beyond Hasson et al.'s approach of examining the synchronicity of BOLD signal responses across brains on a per-voxel basis. Instead, area-specific activity timecourses (ATCs) were investigated. By performing ICA on movie watching fMRI data, Bartels and Zeki demonstrated a technique, "chronotopy", which attempts to parcellate the brain into functionally specialised regions based on their "temporal fingerprints" alone (see Figure 2).

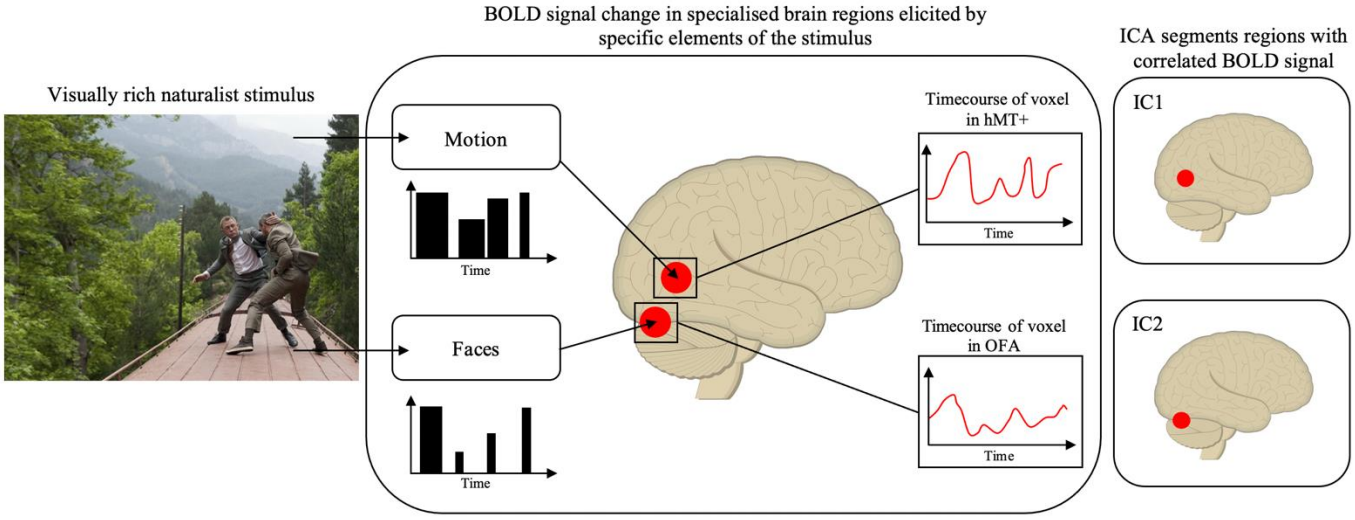

Figure 2. A schematic of the chronotopy approach. A visually rich stimulus is presented to observers while BOLD signals are measured using MRI. Different visual features appear at varying intervals throughout the movie. The magnitude of these features also vary over time (e.g. whether a face is presented face-on in a scene or whether the observer sees a side profile). Functionally specialised regions processing these different features exhibit differential BOLD signal responses, their "temporal fingerprints". ICA exploits these temporal fingerprints to segment the brain into functionally distinct regions.

The chronotopy approach conceptualises the input to natural vision as being decomposable into superimposed independent timecourses of visual features. Movies provide a rich, natural stimulus that should reveal typical processing. This is in direct contrast to traditional localiser tasks which aim to identify discrete regions in isolation which preferentially respond to an extended block of specific, often unnatural, stimuli (e.g., repeated presentation of static faces). Below I summarise the chronotopy approach.

Subjects freely viewed the opening 22 minutes of a James Bond movie (Tomorrow Never Dies) while BOLD activity was measured. ICA was then performed on these data and components were judged to anatomically correspond to functionally specialised regions based upon three spatial features: 1) The "hottest" voxels of the component were clustered to form an area which was located within the gray matter. 2) Components were bilateral, reflected by two clusters of voxels in corresponding locations in the two hemispheres. 3) Components showed spatial correspondence across brains (e.g., they were similar in shape and located in a similar position).

Figure 3 shows the 10 areas identified by Bartels and Zeki. These areas consistently satisfied the criterion detailed above, suggesting they could be reliably identified across their eight subjects during

natural viewing. Several of these regions were believed to, based on their anatomical location, correspond to the visually responsive V1, V2v/V3v, V4, V5/MT+, and LO regions.

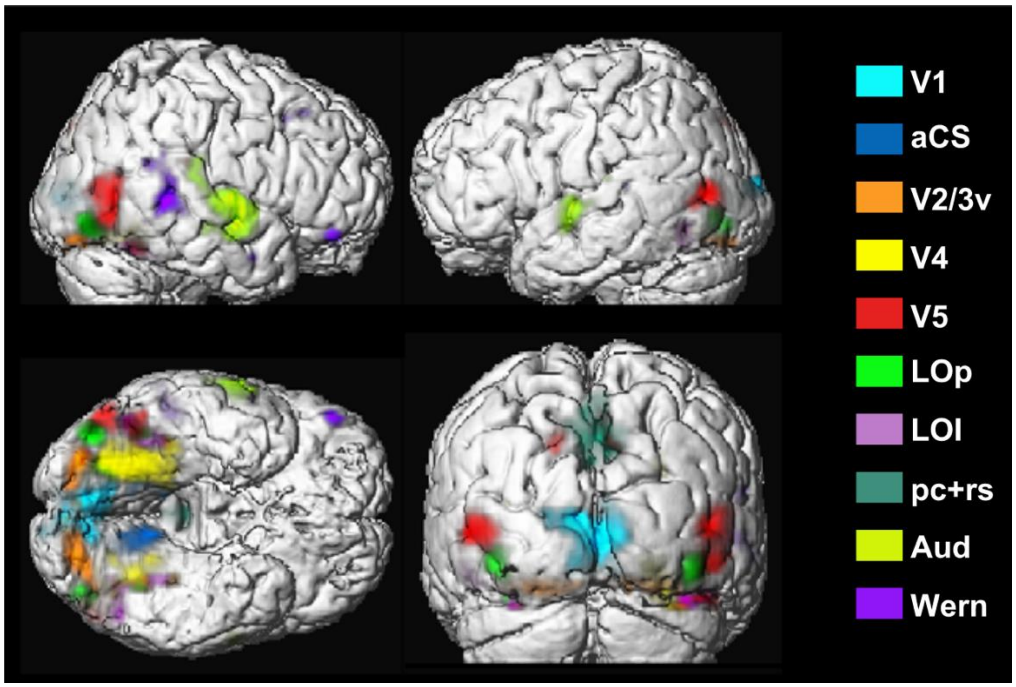

Figure 3. Adapted from Bartels and Zeki (2004). The ten chronoarchitectonically identified areas of a single example subject. Each region represents a separate component derived from performing ICA on movie watching fMRI data. These areas anatomically corresponded with the components produced by the other subjects. ACS = ventral lip of the anterior calcarine sulcus, Aud = auditory cortex, Lol = lateral part of the lateraloccipital complex, Lop = posterior part of lateraloccipital complex, pc + rs = network containing precuneus and retrosplenium, Wern = Wernicke's area.

Following this spatial correspondence screening process, the temporal correspondence of the components was then assessed. By computing cross-subject correlations, Bartels and Zeki showed the ATCs of anatomically corresponding components were significantly correlated across brains (mean $r = 0.3$). This suggests that anatomically corresponding areas across brains were similarly stimulus driven. This significant correlation was completely abolished when computing correlations between the ATCs of anatomically noncorresponding components across brains. Moreover, distinct components within the same brain showed low temporal correspondence.

Taken together, this supports the claim that different regions exhibit different responses to the same stimulus, i.e., they were functionally specialised, and ICA can be used to segment them. Following their landmark study, Bartels and Zeki have reviewed their chronoarchitectonic approach, finding six visually responsive regions can be robustly identified across brains (Bartels & Zeki, 2004a, 2004b, 2005a).

In summary, Bartels and Zeki provided preliminary evidence that during free viewing of naturalistic stimuli, the brain can be sub-divided into functionally distinct regions by their area-specific "temporal fingerprints". This characteristic temporal marker is as specific to the region as its specialisation, and crucially it is preserved across observers who are exposed to the same stimulus. This approach to mapping functionally specialised regions does not rely on a priori hypotheses about the relationship between anatomy and function, nor does it require a battery of highly controlled stimuli designed to evoke responses in isolated regions, but rather segments the brain into subdivisions in parallel with a rich, naturalistic stimulus. If the chronoarchitectonic approach can be validated, it opens the door for a briefer and more engaging way to functionally map the brain. Moreover, the approach has clear implications for clinical populations and developmental research in which long periods of fixation may be problematic and keeping scan duration to a minimum is necessary. In Chapter 3 I investigate whether performing ICA on movie watching can robustly identify multiple visual areas in parallel.

1.4.3 Naturalistic imaging

Naturalistic paradigms challenge the notion that it is appropriate to boil down complex cognitive processes, such as social interactions, to highly controlled paradigms using artificial stimuli. The pioneering work of Bartels and Zeki (2004a, 2004b, 2005b) and Hasson et al. (2004), set out methodological frameworks (chronotopy and ISCs, respectively) for using movie watching to investigate the brain, bringing naturalistic stimuli to the forefront of cognitive neuroscience research. The proceeding years have been fruitful, with research groups looking at the effectiveness of naturistic stimuli (movie clips or spoken narratives) in language, speech, auditory processing, and development. In addition, movie watching has been shown to outperform resting state for functional connectivity-based predictions of behaviour (Finn & Bandettini, 2021), further highlighting how naturalistic stimuli can be an invaluable tool in our methodological arsenal.

In the following section I outline why we, and evidently others, believe naturistic conditions provide a different window for cognitive neuroscientists to study the brain in its "natural habitat"; that is, whilst processing realistic, dynamic, and complex stimuli.

At the start of 2021, the journal NeuroImage ran a special edition collating research investigating the use of ecologically valid conditions to study brain function. From the >100 articles submitted to NeuroImage, 50 were included in the edition covering several themes within the broad cognitive neuroscience discipline, from music and language to the application of naturalistic conditions in clinical populations (e.g., autism spectrum disorder and schizophrenia; Finn et al., 2022). I urge those who are, as I am, interested in the potentially powerful technique of combining naturalistic conditions with neuroimaging to consult this special edition.

Below I provide a synopsis of a diverse selection of studies that have used movies to investigate the brain basis of social cognition, emotion processing, and visual motion processing in neurotypical observers. In the final section I discuss the use of naturalistic stimuli in studies of clinical populations. Each study has, in one way or another, built on the pioneering work of Bartels and Zeki (2004a, 2004b, 2005b) and Hasson et al. (2004).

Naturalistic conditions should, where possible, be a first 'port of call' for cognitive neuroscientist wanting to understand brain function. Nastase et al. (2020) are one of the main proponents of naturalistic paradigms and argued that "non-naturalistic experimental manipulations have occupied an overly privileged position in cognitive neuroscience" (Nastase et al., 2020) . Moreover, they caution against traditional, highly controlled experimental designs having an oversized influence in the formation of hypotheses. Instead, they believe a more ecological approach should be adopted when conceiving hypotheses, which, in turn, should lead to findings that are more generalisable to the world outside of the laboratory.

1.4.3.1 Social cognition and emotion processing

Movies allow the observer's cognitive and perceptual processes to unfold over time, in an unconstrained and uninterrupted manner (Hyon et al., 2020). Therefore, the free viewing of movies can be used to investigate the neural foundations of social cognition and emotional experiences.

Humans tend to form personal connections and wider social networks with individuals who are like themselves; our friendship groups tend to be populated with people who are similar in age, ethnicity and occupation, a principle known as homophily (McPherson et al., 2001). Parkinson et al. (2018) hypothesised that friendships are formed through similarities in the way the individuals perceive, internalise, and react to the world around them. Therefore, rather than using coarse physical traits or similarities in personality inventory scores to understand homophily, Parkinson et al. assessed neural responses to a naturalistic stimulus (an audio-visual movie clip). The aim of their study was to investigate whether or not individuals associate with those who see the world similarly to them.

First, the social networks of the subject cohort were quantified, from reciprocated social ties (i.e., both parties reported a friendship with the other), to friends of friends, to socially distant ties. Next, subjects freely viewed audio-visual clips (covering a range of topics) whilst in an MR scanner, before fMRI timecourses were extracted from anatomical ROIs in each brain. ISCs were then computed between the timecourses of corresponding ROIs for each pair of subjects (e.g., timecourse of ROI X in Subject 1 versus timecourse of ROI X in Subject 2). Parkinson et al. found that subjects who were closer to one another in the social network (e.g., friends or friends of friends) showed significantly more similar neural responses than those who were socially distant, specifically in regions associated with motivation, learning and attention. This suggests neuronal responses to movie watching reflect homophily in the social environment.

Emotive Hollywood movies have been used to investigate humans' voluntary control of empathy. Borja Jimenez et al. (2020) demonstrated that differences in brain synchronicity were dependent upon whether subjects were instructed to either empathise with the characters while rating their emotions or

make the judgements from a detached perspective. Responses in the limbic and somatomotor regions showed increased synchronicity across brains when subjects empathised compared to when they detached. In addition, functional connectivity was found to strengthen across networks involved in embodied cognition and mentalising when subjects were encouraged to empathise, suggesting intentional control of empathy changes brain activity.

Movies have also been used to study an opposing emotional response, fear. In contrast to static fear-inducing stimuli, a well-directed movie will have the power of gradually increasing the level of fear experienced by the observer before maintaining it over a period of time, as well as including scenes designed to suddenly scare (Kinreich et al., 2011). Therefore, movies can elicit both acute fear (e.g., "jump-scares") and sustained fear (e.g., anticipatory fear before the threat is presented; Hudson et al., 2020). It is this ability to capture the emergence of fear that makes movie watching paradigms particularly attractive to those aiming to study the emotion. Using Hollywood horror movies and ISCs, Hudson et al. (2020) found separate neural mechanisms involved in acute and sustained fear, with the former eliciting activation in subcortical and limbic areas, while the latter increased activation in the auditory and visual cortices.

1.4.3.2 Motion processing

A movie directors' goal is to make viewers "feel" a particular way; they contain emotional content which can be leveraged in studies seeking to understand emotion processing and social cognition. However, because the content is delivered visually, movies can also be used to investigate visual processing and perception more broadly.

In an innovative study, Bartels et al. (2008) investigated whether different aspects of visual motion, specifically global and local motion, occupy different neural pathways. In their study, image changes were extracted from each frame of the James Bond movie Tomorrow Never Dies (described below) and correlated with the subject's BOLD response measured while free viewing the movie. First, the total luminance change from one frame to the next was calculated. Next, using a motion detection algorithm, frame to frame changes in luminance were portioned into global motion flow involving the entire scene

(e.g., motion caused by the camera panning across the scene or zooming in) and local motion (e.g., incoherent, or spatially local motion). The remaining changes in luminance, not accounted for by the two motion features, were classified as residual changes (e.g., scene cuts). Correlations were then computed between the measures extracted from the movie and the BOLD responses acquired while subjects watch the same movie. Bartels et al. found early visual areas (V1-V4) only responded to non-motion-related (residual) changes, whereas the medial posterior parietal cortex (mPPC) correlated selectively with global motion (e.g., simulated observer motion) and responses in hMT+ and V3A correlated with local motion (e.g., object motion). Crucially, the two motion-selective networks showed very little overlap, suggesting a clear spatial segregation for different attributes of visual motion. More recently, Gilaie-Dotan (2016) proposed that hMT+ and V3A are involved in separate visual pathways: the lateral motion and dorsal pathway respectively. In Chapter 5 I investigate whether we find evidence of pathways specialised for visual motion when movie watching data is decomposed with ICA.

Naturalistic stimuli not only allow us to examine brain responses to complex stimuli, but also eye-movement related activity driven by the dynamic content. Research into the neural correlates of eye-movements often rely on subject's responses to synthetic stimuli (e.g., Kimmig et al., 2008). For example, subjects may be asked to follow a dot traversing a uniform background or to move their eyes to a target when it appears. However, in natural vision, humans make several different eye movements every second, even when performing well-learned or "automated" tasks, such as making a cup of tea (Land et al., 1999). Therefore, instructing subjects to make just one eye movement (or two if a smooth pursuit is followed by a saccade) does not seem appropriate when trying to understand the relationship between natural vision and eye movements. Moreover, in the "real-world", eye movements are made in the presence of complex backgrounds, often with multiple objects moving simultaneously throughout the visual field. Since movies contain different embedded visual features, including objects and complex backgrounds, they more closely resemble daily life, thus making them advantageous for investigating the neural responses associated with eye movements.

Although movies are more reflective of natural vision than static images, one concern may be that such a temporally dynamic stimulus would result in idiosyncratic gaze behaviour. However, Hollywood movies have been shown to exert consistent control over observers viewing patterns, reflected by high eye movement inter-subject correlation (Davis et al., 2021; Hasson et al., 2008).

In summary, the advantage of using movies to identify brain activity evoked by eye movements is twofold:

i. they are complex and dynamic, thus more closely resembling our everyday visual experience.
ii. they evoke a mixture of saccades, fixations and smooth pursuits that are stable across observers and can be measured with in-scanner eye-tracking systems.

One recent study examined differences in smooth pursuit and saccade related brain activation while subjects free viewed a Hollywood movie (Agtzidis et al., 2020). The authors report three significant clusters related to smooth pursuit and two related to saccades. The smooth pursuit clusters involved regions previously shown to respond during smooth pursuit and visuospatial processing, specifically hMT+ bilaterally, the middle cingulate, extending to the precuneus and the right temporo-parietal junction. However, activation in the FEF, a region consistently described in the planning and execution of smooth pursuit eye movements, was absent. Agtzidis et al. proposed that the lack of activation in the FEF may be a result of the stimulus used; traditional smooth pursuit tasks, as described above, often involve prolonged periods of fixation followed by the tracking of a synthetic stimulus. In contrast, observers free viewing a movie may be continuously engaged in the planning of eye movements. Therefore, the lack of FEF activation in the above-mentioned study may in fact be more representative of viewing behaviour in the real-world.

The studies described above highlight how naturalistic stimuli can be used to identify changes in functional synchronicity and connectivity in neurotypical brains, demonstrating their effectiveness in understanding the neural basis of emotion, social cognition, and visual motion processing. However, few studies are specifically geared towards identifying individual differences (Finn et al., 2020). In general, data

acquired through naturalistic imaging protocols are analysed in such a way as to identify similarities across brains, not differences. This approach may be inherited from Bartels, Zeki and Hasson et al. (described above) whose studies were mostly concerned with movie watching eliciting spatial and temporal corresponding activity across brains. However, movies can convey real-world complexity more ecologically than traditional paradigms using artificial stimuli, and therefore may elicit richer patterns of brain activity.

1.4.3.3 Clinical populations

Since movies tend to elicit synchronised activity across neurotypical brains (as identified through ISCs, see above), they also enable the identification of stable idiosyncratic patterns (Finn et al., 2017). Therefore, naturalistic imaging has been used to investigate psychiatric disorders, such as depression (Gruskin et al., 2020) and schizophrenia (Yang et al., 2020), as well as the neural responses of individuals in a chronic nonresponsive state (Naci et al., 2017).

There is no 'gold standard' for identifying the biological markers of psychiatric illness. One common approach used to identify consistent and generalisable deficits in brain function associated with psychiatric disorders is to average across a group of patients with similar symptomology. The average patient brain is then compared to an average brain generated from a group of healthy controls (Spaulding & Deogun, 2011). Rikandi et al. (2022) applied this approach to movie watching data and reported differences in functional connectivity between a group of patients with psychosis and a group of healthy controls. However, such an approach works from the assumption that individuals with similar clinical diagnosis also share similar neurological markers; an assumption that does not necessarily hold true. For example, Wolfers et al. (2018) showed that only 2% of individuals with schizophrenia or bipolar disorder showed comparable abnormalities in brain structure. Similarly, idiosyncratic neural responses to movies have been reported in patients with the same clinical diagnosis of autism spectrum disorder, with no such desynchronisation (i.e., lower ISCs) noted in controls (Byrge et al., 2015; Salmi et al., 2013).

An alternative to the group-versus-group method is the individualised psychiatric neuroimaging approach (Yang et al., 2020). Since healthy controls show a synchronization in brain dynamics when

watching a movie, a normative model can be generated that captures the variability in a healthy population. The activation evoked while individuals watch the same movie can then be compared to the normative spatiotemporal response template and based on deviations from the norm, individuals with schizophrenia can be detected with high accuracy (Yang et al., 2020). The approach is advantageous because it explicitly avoids averaging across patients, thus allowing for individual-specific brain deficits to be identified.

In summary, not only can movies be used to inform our understanding of how the brains of those with psychiatric illness differ to healthy controls, but as Yang et al. (2020) demonstrated they can also be adopted as part of a promising diagnostic tool.

1.4.4 Summary

Synchronised activity across healthy brains is reliably identified with different movies genres, from animated, child-friendly (e.g., Finn & Bandettini, 2021) to intense horror films (e.g., Salmi et al., 2013), demonstrating adaptability and robustness of the approach. However, the mental processes in which the observer is engaged in may differ across movie genres. For example, a comedy movie may drive different cognitive processes to a psychological thriller. Similarly, higher ISCs are dependent on the quality of the movie, with poorly directed movies showing lower ISCs compared to well directed ones (Hasson et al., 2008). Therefore, both the quality and the genre of the movie should be considered when developing a naturalistic imaging protocol, especially if hypotheses about different emotional states are being investigated, but less so if the aim is to understand broader visual perception.

To summarise, the visual richness and ability to depict complex social interactions make movies a good stimulus for simulating many aspects of everyday life. If well-directed, they can elicit synchronised activation across healthy brains in a pleather of regions. This synchronised activity can be used to generate a normative template to which individual responses can be compared and psychiatric illness subsequently detected. The studies described above highlight a growing movement of naturalistic neuroimaging towards a well-established and progressive sub-field of cognitive neuroscience. In addition, I point the reader to the

reviews of Eickhoff et al. (2020) and Sonkusare et al. (2019) who provide excellent commentaries of the benefits and future directions of naturalistic imaging.

The preceding chapter provides a description of the two primary analysis methods used in this thesis, ICA and ISCs, as well as a synopsis of the literature relating to naturalistic neuroimaging. The literature demonstrates that movies are ecologically valid stimuli that, owing to their complex nature, evoke activation throughout the brain. Importantly, this activation is synchronised across healthy individuals; neurotypical brains "tick together" when exposed to the same visually rich stimulus. The literature shows that complex cognitive processes and visual processing more broadly can be investigated through analysis of movie data with ICA or ISCs. Nevertheless, there are gaps and hypotheses in the literature that we believe can be addressed and tested, respectively, with movie watching data.

Bartels and Zeki provided preliminary evidence that during free viewing of a movie, the brain can be sub-divided into functionally distinct regions. The research that followed branched into new areas and investigated novel uses of movies, with very limited work directly assessing Bartels and Zeki's original method. Consequently, there appears to be a lack of a systematic validation of the chronotopy approach in the literature. Such a validation would be useful given how valuable the approach would be for clinical populations and research institutions. Similarly, Bartels and Zeki's provisional evidence identified six visually responsive regions: V1, V2v/V3v, V4, V5/MT+, and LO regions. Little is known about whether the approach can be extended and used to identify other visual areas of the brain, and if so, how reliably they can be identified across observers. Chapter 3 of this thesis is a direct continuation of Bartels and Zeki's work, in which we both validate and extend their chronotopy approach with the aim of assessing whether it can be adopted as an 'all-in-one' localisation tool.

In the remaining chapters I test existing hypotheses in the literature by performing ICA and ISCs on movie watching data. Therefore, synopses of the literature relating to these specific hypotheses, not covered in the sections above, are provided in the corresponding chapters where it is most appropriate. Below I provide a very brief overview of the motivations of the remaining chapters.

There is well-documented evidence in the literature of some functions being lateralised to one hemisphere, for example, visuospatial attention, language, and face processing. In Chapter 4 I adopt Hasson et al.'s methodological framework (movie watching and ISCs) to investigate whether I find evidence of lateralisation of visual motion processing, and whether this evidence is consistently found across brains.

In Chapter 5, I examined whether we find evidence in line with the literature related to visual motion processing being facilitated by a separate neural pathway. I perform ICA on movie data and investigate whether the output components correspond to well-established visual pathways (e.g., dorsal and ventral streams) or if they provide evidence to support alternative models of visual processing (e.g., an additional pathway specialised for visual motion perception).

In the final chapter I present an experiment designed to assess different forms of residual visual function in patients with cortical vision loss. As discussed in section 1.1, much of the literature relating to residual vision focuses on whether patients retain the gross ability to detect motion in their blind field. In contrast, I developed three tasks, each requiring a different response from the patient. The aim of this experiment was to investigate whether there was a dissociation between performance on these tasks.

Chapter 2. Methods development

The original aim of this thesis was to develop a set of neuroimaging and psychophysical tools to investigate and map residual visual function, and neural substrates, in patients with cortical vision loss arising from occipital damage. However, as outlined in Chapter 1, without access to patients the aims of this thesis changed; instead, I sought to investigate three distinct research questions using fMRI data acquired while subjects freely viewed a movie. This chapter details the general methods pertinent to each chapter that follows. The remaining methods are described in the relevant chapters where they are most appropriate.

In this thesis, single-subject ICA is used to investigate two of the three main aims. In Chapter 3, I explore the chronotopy approach by performing ICA on our movie watching data. The aim of Chapter 3 was to assess whether output components anatomically correspond to the visual areas included in a surface-based probabilistic atlas. Therefore, to allow for a direct comparison with the surface-based atlas, assessed through spatial correlations, I performed ICA and subsequently sampled the spatial maps to the individual surface. In Chapter 4, the ICA components identified as anatomically corresponding to hMT+ in Chapter 3 were used to assess the potential lateralised function of the complex.

Both chapters require components on the native surface and therefore, movie watching data was pre-processed, decomposed with ICA, and the resulting spatial maps were sampled to the surface. This chapter provides a detailed description of how I identified an optimal pipeline for processing these data with the aim of segmenting the brain into visually responsive regions. To be thorough, in section 2.1.5.3 I provide a detailed description of eight different pipelines along with the logic behind the pre-processing steps involved and the methods used for evaluating each pipeline.

2.1 General methods

Unless stated otherwise, the same archival dataset was used throughout this thesis[2]. The following section describes the imaging acquisitions and stimulus methods used to generate this dataset.

2.1.1 Subjects

Archival data was retrieved for seventeen subjects (seven females, nine males[3]) from Cardiff University who volunteered to take part in the original study in exchange for course credits or money. The mean age at the time of scanning was 27 years (range: 21-40)

2.1.2 Equipment

Images were generated by an LCD Canon XEED SX60 projector running at the native resolution of 1400 × 1050 pixels, at 60Hz. The images were viewed while the subjects lay on their back in the scanner. Visual stimuli were rear projected onto a screen located within the bore of the scanner. The screen was located behind the head of the subject and viewed through a mirror located above the head coil. The screen viewing distance was approximately 57cm. The viewable image was 35cm in width (~35degrees).

At the edges of the screen (±17.5 degrees laterally) the corners of the image were slightly occluded, due to the MR bore obstructing the view. Therefore, the maximum height at the centre of the stimulus was 30 degrees and at the edge was 23 degrees. The corners of the screen were slightly occluded, but this did not affect the viewing of the movie.

2.1.3 Stimulus

Subjects watched the same 20-minute clip taken from the beginning of the James Bond movie Skyfall (MGM & Columbia Pictures), with the soundtrack removed. Subjects were not instructed to fixate

---

[2] This dataset was collected by Phoebe Asquith. Sections 2.1.2 and 2.1.3 are taken directly from Asquith (2018).
[3] Due to unforeseen circumstances, information relating to the age and sex of one subject could not be retrieved from the archive.

any specific point of the stimulus but were instead instructed to "naturally" watch the movie as if they would when at home.

Bartels and Zeki (2004a, 2004b, 2005b) showed visually responsive regions can be identified using a combination of a visually rich natural stimulus and ICA. In their original studies, a James Bond movie (Tomorrow Never Dies) was selected as their stimulus. The movie was a good choice because it is visually engaging, and it contains a very wide range of visual features including faces, bodies, complex indoor and outdoor scenes, biological motion and object manipulation. In addition, the zooming and panning of the camera produced prolonged periods of global motion (optic flow) which occurred independent of what was happening in the scene. Therefore, following Bartels and Zeki's original studies (2004a, 2004b, 2005b), we also opted to use a James Bond movie as our stimulus (SkyFall).

2.1.4 Structural and functional brain imaging acquisitions

Images were acquired using a 3T General Electric (GE) Signa 3T HDxt Scanner (General Electric, Milwaukee, Wisconsin), equipped with echo-speed gradient coil and amplifier hardware using a GE 8 channel receiver RF coil.

The parameters used for the acquisition of the high-resolution anatomical scan (fast-spoiled gradient echo [FSPGR] sequence) were: 256 slices, TE = 3.29 ms, TR = 9 ms, voxel size = 1.0 × 1.0 × 1.0 mm^3, flip angle = 20°.

Activation images were acquired using echo-planar imaging (EPI) with the following parameters: TR = 2000 ms, TE = 35 ms, voxel size = 3.0 × 3.0 × 3.0 mm, flip angle = 90°. A total of 600 volumes were acquired for each of the 17 subjects. The functional scan duration was 20 minutes.

2.1.5 Visual area maps

A set of published topographical maps were used to define the visual areas throughout this thesis (Wang et al., 2015). These maps are publicly available in the standard volume (MNI space) and surface

(fsaverage space) formats. A summary of how these maps were generated by the authors is below for reference.

Polar angle and eccentricity maps generated through retinotopy were projected onto the subject's surface reconstruction and the borders between the areas were then delineated by hand. For V1, V2 and V3, dorsal and ventral representations were identified separately, resulting in the generation of 25 topographic maps ("ROIs"), comprising 22 visual areas for each hemisphere. The subject-specific ROIs were then warped from their native individual surface to a standard surface. Each subject's set of spatially normalised ROIs were then superimposed onto a single standard surface. A full probability map (FPM) and a maximum probability map (MPM) was generated for each ROI. These maps are described in detail below.

2.1.5.1 Full probability map (FPM)

The FPM for each visual area was generated by dividing the number of times a given vertex belonged to that ROI by the number of subjects who contributed to the ROI[4]. These values correspond to the likelihood that a vertex belongs to a given visual area, and therefore conveys the variance in the position of a given ROI across individuals. Higher values in the FPM represent common overlap between subjects (more brains included that voxel in the map). The peak probability values dropped from 100% for early visual areas to 40% for anterior parietal areas. Wang and colleagues concluded that this difference in the peak probability was due to an increase in the variability across brains of the alignment of ROIs within standard space for anterior regions.

An exemplary FPM of right hemisphere V1d is presented in Figure 4. The colour-coded vertices convey the likelihood of that vertex belonging to V1d across the 50 subjects who contributed to the ROI. Higher probabilities, denoted by the red colour-code, are located in the centre of the ROI.

[4] Although a total of 52 subjects were scanned, this value varied depending on the visual area (range 50 – 31). More subjects contributed to the lower visual area ROIs, whereas only ~30 subjects contributed to the IPS ROIs.

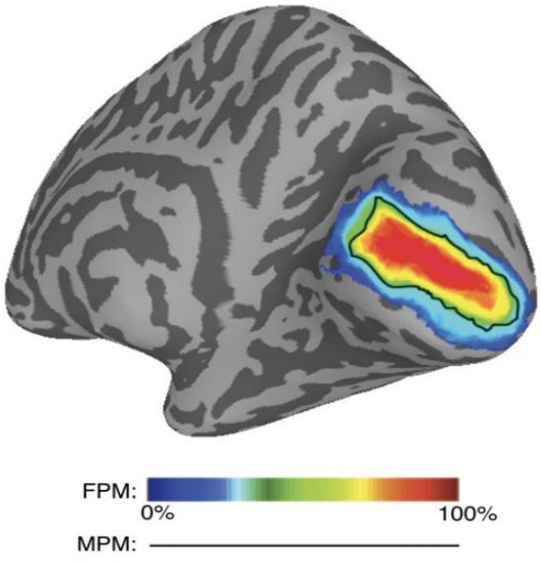

Figure 4. Adapted from Wang et al. (2015). An exemplary surface full probability map of right hemisphere V1d. A total of 50 subjects contributed to this ROI. The gradual increase in the colour code from blue to red indicates an increase in probability. The map shows the central region of the ROI has higher probability values indicating these vertices were assigned to the ROI in more subjects. The black outline denotes the V1d maximum probability map border.

2.1.5.2 Maximum probability map

The MPM (Figure 5) was generated by comparing, at each vertex, the probabilities of all the ROIs and allocating that vertex to the area with the highest probability. This approach was implemented as follows: First, the probability values for a vertex belonging to any ROI are summed to generate an "inclusion probability". For example, if the probability of vertex × belonging to region 1 or 2 was 35 and 25% respectively, the summed inclusion probability would be 60%. This is then compared to the probability of the vertex being outside of all defined regions (exclusion probability). In this example, the exclusion probability is 40% (100% - 60% inclusion probability). If the exclusion probability value was greater than the inclusion probability value, that vertex is excluded from the map. If the reverse is true, as in the example above, the probability values across ROIs were compared. The vertex is then assigned to the visual area with which it had the highest probability of belonging to. This would be region 1 in our example. By assigning each vertex to the most probable region, the MPM represents a summary of the topographic organisation across subjects.

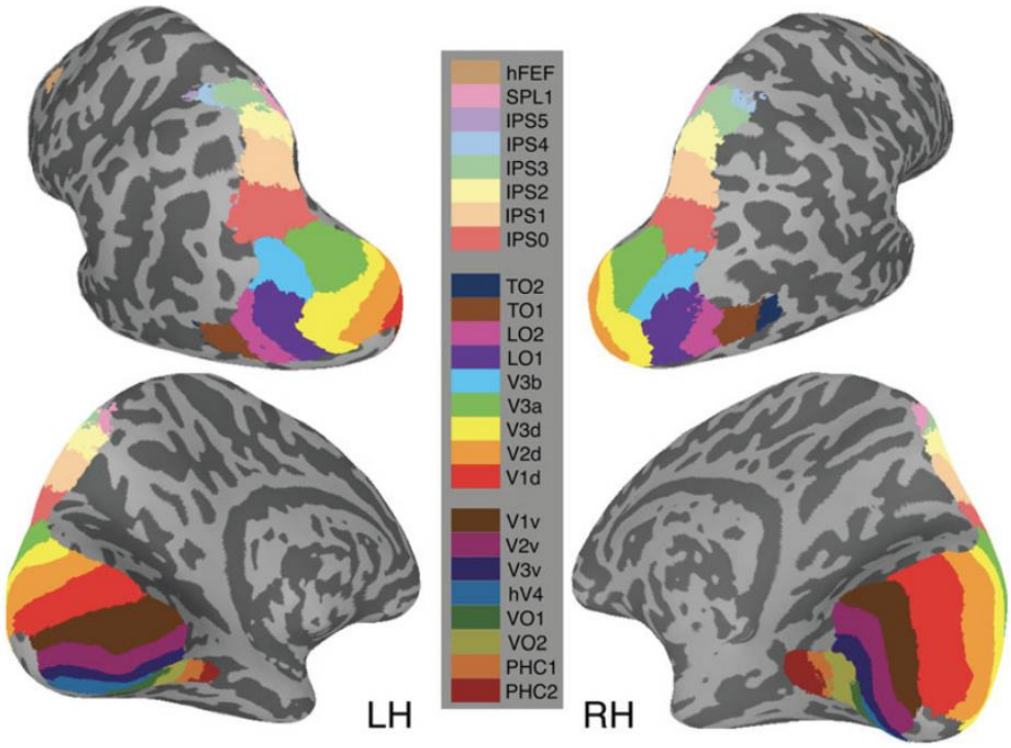

Figure 5. Adapted from Wang et al. (2015). The surface maximum probability maps. These maps represent a summary of the topographic organisation across subjects.

2.1.5.3 Limitations of probabilistic ROI-based approaches

The published topographical maps described above were used to define the 25 visual areas throughout this thesis. The use of a probabilistic atlas was beneficial because, as outlined in section 1.1, the acquisition of retinotopic data was not possible. Therefore, the atlas allowed us to define the visual areas in our subjects anatomical and functional spaces by leveraging information about each area from an independent sample of subjects. However, it is important to raise a potential limitation of these atlas-based definitions. As discussed in section 1.3, atlases are, by their nature, probabilistic and as such cannot reflect all the anatomical variation in the size or location of an area across brains. Therefore, atlases do not provide a 'ground truth' for each area.

Since we did not perform retinotopy on individual subjects or perform a battery of functional localisers, we relied on the atlas providing an approximate location of each area. Importantly, we did not

rely solely on the atlas definitions, instead we performed a range of different analyses and used the converging results to inform our conclusions.

2.1.6 Operations on Pearson's r values

Throughout this thesis, correlations are used to investigate hypotheses and to quantify results. This raises a potential issue as to how these values are used, specifically the distribution. The sampling distribution of the Pearson's r correlation coefficients is not normally distributed. Consider a situation in which the population (ground truth) correlation of two variables = .8. If we take repeated samples from the population, the correlation of each sample (the sample correlation) will not be exactly .8. As correlation values, r, range from -1 to +1, the upper tail of the sampling distribution (i.e., the distribution of r values after repeated samples) is limited by 1, resulting in a negative skew (Figure 6, bottom). On the other hand, if the population correlation is .2, the sampling distribution is closer to normal since the upper and lower tails can be more "space" symmetrical (distance from .2 to -1 and +1 is greater; Figure 6, top). This boundedness of the correlation coefficient (-1 to +1) means that the further the correlation is from zero, the greater the degree of skew.

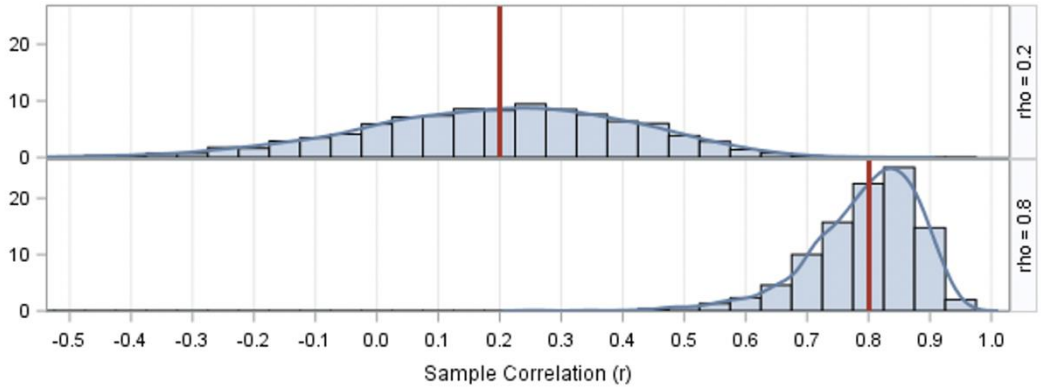

Figure 6. The sampling distribution of Pearson's r, with a population correlation of 0.2 (top) and 0.8 (bottom). Adapted from "Fisher's transformation of the correlation coefficient", by R. Wicklin, 2017. SAS. Retrieved June 1, 2022, from https://blogs.sas.com/content/iml/2017/09/20/fishers-transformation-correlation.html. Copyright 2022 by SAS Institute Inc.

There are three approaches for dealing with non-normally distributed data; (i) rely on the robustness of parametric tests (e.g., paired-samples t-tests) to deviations from normality; (ii) perform non-parametric tests (e.g., the Wilcoxon signed rank test); (iii) transform the data so that it has a normal

distribution. In this thesis we adopted approach (i) and report the results in the main body. For completeness, we also adopted approach (iii) and report the results in the appendix with any differences between raw and transformed values clearly highlighted[5]. The following section describes how the data were transformed and subsequently analysed (approach [iii]).

The statistical procedure we used to overcome the issue of non-normally distributed correlation coefficients is the Fisher's Z transform (Fisher, 1915). The Fisher's Z transform converts a Pearson's correlation coefficient (r) to a normally distributed value, known as Z'. Figure 7 shows the relationship between r and Z' values (blue line) with the identity function plotted for reference (grey line). The most significant deviations are seen when the correlation is high (e.g., > ± .6), with little change in middle values.

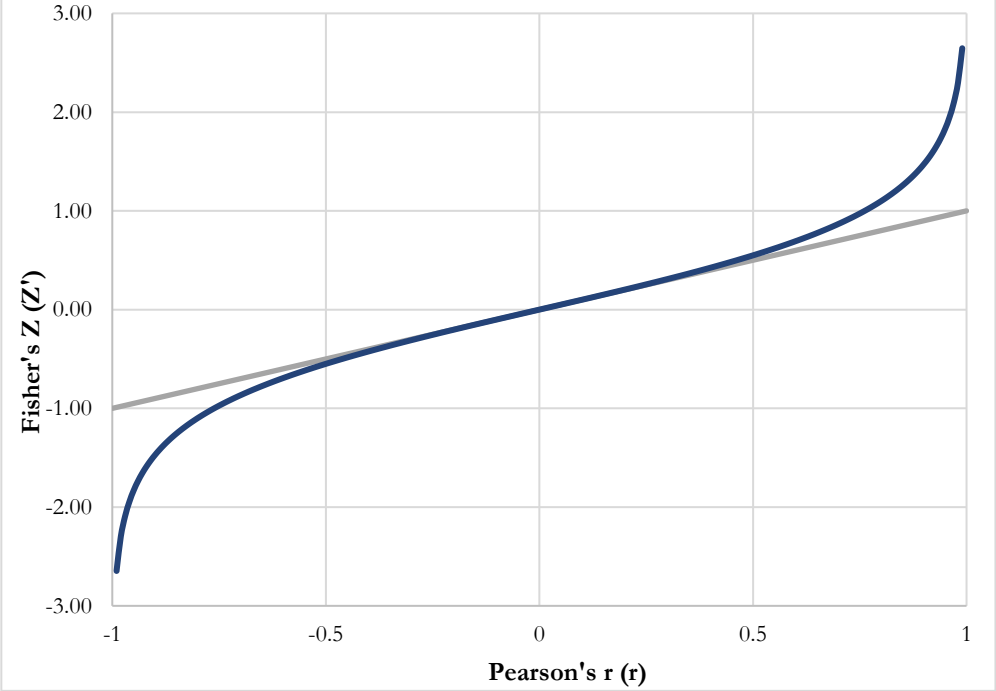

Figure 7. The relationship between Pearson's r and Fisher's Z. Note how r ≈ Z' when r is close to 0.

[5] Although these results are reported in Z', to remain consistent with the main body of the thesis I continue to refer to the results of the operations/parametric tests performed on the transformed values as "correlations" (e.g., the mean correlation was Z = 1.436 [SD = 0.176]). The reason for this is that once each correlation coefficient was transformed to a Fisher's Z value, the subsequent operations/tests were performed on the values as if they were correlation coefficients.

Let us return to the worked examples described in Figure 6. Performing the Fisher's Z transform on the correlation values effectively stretches the tails of the sampling distribution so that it is approximately normal (Figure 8). We can see that when the population correlation is .2, there is no change (Z' = .2; Figure 8, top), whereas for a population correlation of .8, Z' = 1.1 (Figure 9, bottom). Once transformed, operations such as averaging and statistical tests can then performed on these transformed values.

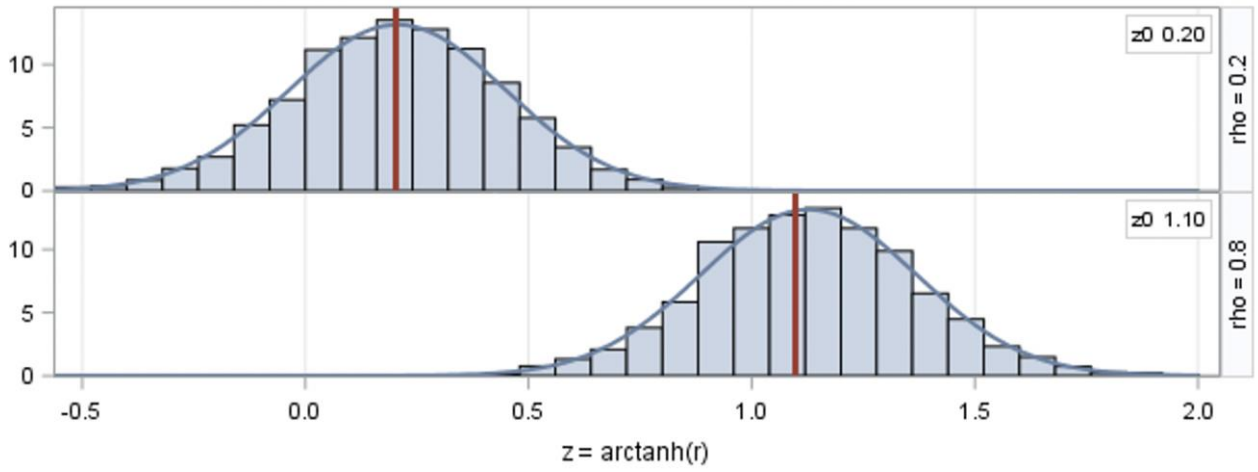

Figure 8. The sampling distribution of Fisher's Z transformed Pearson's r values with a population correlation of 0.2 (Z' = .2; top) and 0.8 (Z' = 1.1; bottom). Adapted from "Fisher's transformation of the correlation coefficient", by R. Wicklin, 2017. SAS. Retrieved June 1, 2022, from https://blogs.sas.com/content/iml/2017/09/20/fishers-transformation-correlation.html. Copyright 2022 by SAS Institute Inc.

Figure 9 shows an example of how each individual correlation coefficient was transformed to Z', before being treated identically to how they were for the equivalent statistical test performed on the raw, r values. Figure 9 also reflects how the general pattern of the results, in practice, are unaffected by the transform. This is true for both descriptive statistics (e.g., notable differences in the mean correlations of ROI timecourses are preserved when the values are transformed to Z') and parametric test results (as exhibited in Figure 9). In the case of the latter, a comparison of the p values shows only small differences between tests performed on r and Z' values. These differences were most often in the direction of a smaller p value when performed on Z' compared to r.

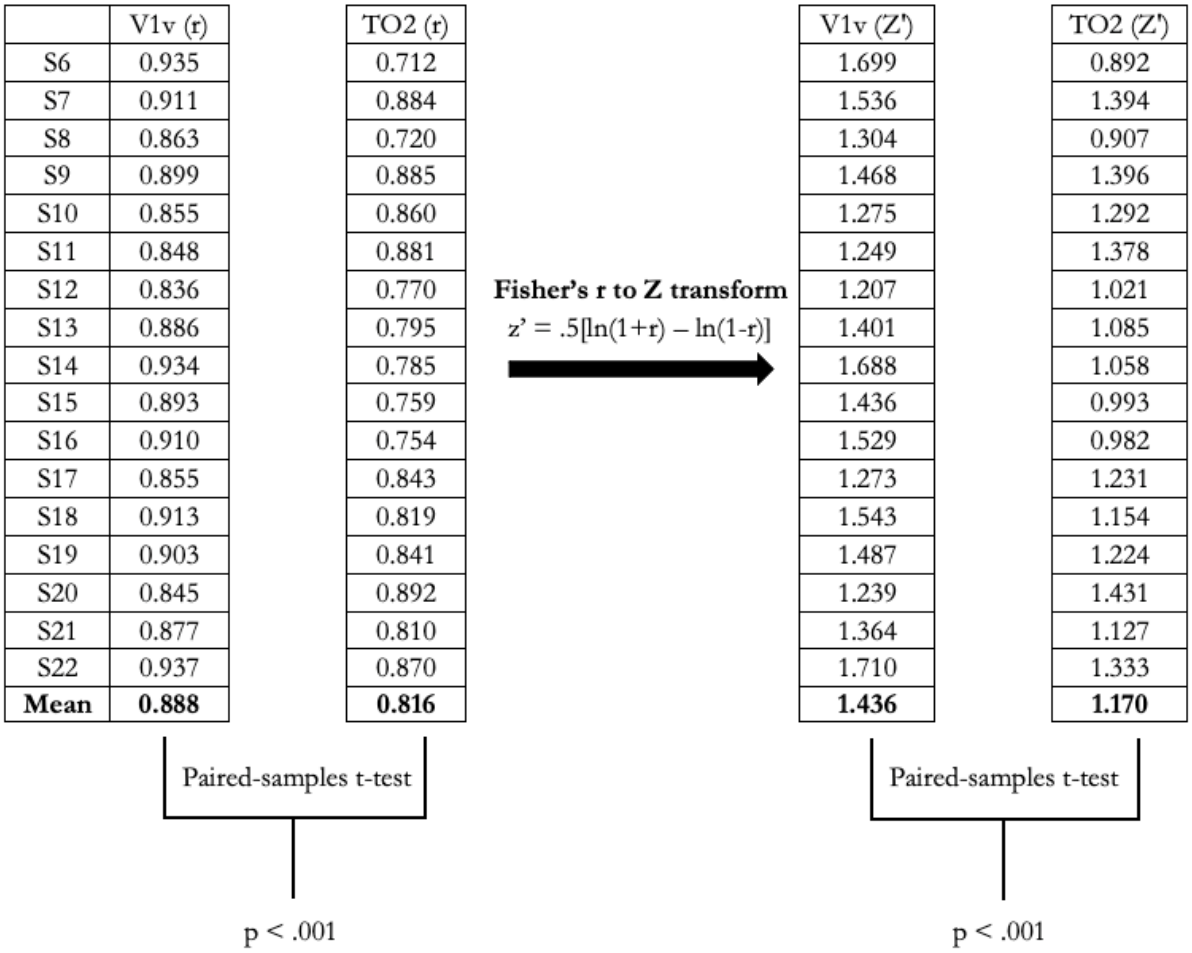

Figure 9. An example of how parametric tests were performed on both Pearson's r (r) and Fisher's Z (Z') values. To ensure the values were normally distributed, individual subject correlation coefficients were first transformed to Z'. The parametric test is then performed on both the raw and transformed values. The p value of the test is unaffected by the transform.

As outlined above, descriptive statistics, parametric test results and group average figures are presented in r in the main body of this thesis. There are three reasons why I adopted this approach. First, to aid the interpretation of our findings; r values are more familiar than Z', particularly when examining the relationship between two variables. Second, there is little-to-no difference in the results obtained from r and Z' values. Third, we followed the precedent in psychology research. For example, the extensive literature on reaction time, in which statistical analyses are performed on values following a log transformation (to remove the skew), but descriptive statistics are reported in seconds rather than log seconds (e.g., Bucsuházy and Semela, 2017).

2.1.7 Correcting for multiple comparisons

When performing multiple simultaneous null hypothesis significant tests, it is important to consider the potential impact of "multiplicity". Multiplicity refers to when multiple significance tests are performed within a single study (also known as the multiple comparison problem), and as a result there is a potential for the type I error rate (false-positive) to be inflated. It is advised that parametric tests performed in the absence of a priori hypotheses, such as t-test performed on several different variables, should be corrected for multiple comparisons (Streiner & Norman, 2011).

Consider an experiment in which data was collected for many different variables, and without planned hypotheses 20 significance tests were performed to establish whether any of the results were significant. With an alpha value for significance of 0.05 (i.e., 1 in 20), we would expect one of those 20 comparisons to be significant (i.e., $p < 0.05$) simply by chance. Therefore, without an adequate correction for multiple comparisons, the error rate remains inflated, and the significance of the finding is misleading. Correcting for multiple comparisons involves adjusting the alpha value so that the probability of a type I error remains fixed. One commonly used approach for correcting for multiple comparisons is to divide the alpha value by the number of performed comparisons (S. Lee & Lee, 2018). For example, if 20 comparisons are performed, the adjusted alpha for significance would be 0.0025. If the p value of a test is below this adjusted alpha value, it can be interpreted as significant.

On the other hand, Armstrong (2014) states that if a small number of planned comparisons are performed, no correction is advised. Similarly, post-hoc exploratory tests performed with the aim of forming hypotheses for further investigation need not be corrected for multiple comparisons (2014). The key with the latter scenario is in the interpretation of the findings. It is not appropriate to conclude there is an effect based on uncorrected and exploratory results, but rather the test should be used to guide the future exploration of an effect or to set the results of the planned comparisons in a wider context. Therefore, following Armstrong (2014) and Streiner and Norman (2011), a correction for multiple comparisons was not applied when a small number of pre-planned comparisons were performed.

Furthermore, in the case of post-hoc/exploratory comparisons we report the raw p values but explicitly label the analyses as such.

2.2 ICA pipeline validation

2.2.1 Introduction

The aim of the following section is to identify an optimal pipeline for robustly identifying visual areas using ICA and movie watching data. The objective of this pipeline was to produce, for each subject, a set of spatial maps on the individual's surface to allow for direct comparisons with the surface-based probabilistic atlas described in section 2.1.5. The optimal pipeline is then used for further analysis in subsequent chapters.

In the following section I examine eight different processing pipelines. The pipelines were generated (i) in line with Human Connectome Project (HCP) recommendations (following advice from personal communications with Mathew Glasser, July 2020), (ii) to examine the effectiveness of an emerging optimal pre-processing pipeline (fMRIPrep) compared to the default pipeline (MELODIC GUI), and (iii) to examine whether performing ICA on surface- or volume-based data results in more robust components corresponding to visual areas.

For each pipeline, standard pre-processing steps were carried out using either the MELODIC GUI or fMRIPrep. Following pre-processing I then performed ICA on the data either in the volume or directly on surface data. When ICA was performed in the volume, I explored two options for moving the spatial maps to the surface (regression of component timecourses into surface data and resampling spatial maps to the surface).

Given our aim, the effectiveness of each pipeline was evaluated through spatial correlations between the components and a probabilistic atlas of visual areas. This evaluation is based on the assumption that ICA is an appropriate tool for identifying visual areas (as demonstrated by Bartels and Zeki, 2004b; 2005b), and therefore the use of an atlas provides an appropriate measure of the effectiveness

of each pipeline. This correlation analysis gave us two metrics to use for the pipeline evaluation: (i) the overall spatial correlation between components and the atlas and (ii) the number of unique components produced by each subject. A unique component is one that is the top match (highest spatial correlation) to only one visual area. In short, more unique matches indicate the individual visual areas are segmented, and thus provides a measure of how effective the pipeline is at identifying discrete areas. A more detailed description of how the pipelines were assessed can be found in 2.2.1.4.

To pre-empt the results, I find the optimal pipeline is one in which data are pre-processed with fMRIPrep, smoothed, and then decomposed with volume-based ICA. The spatial maps of the components are then sampled to the surface following decomposition. For completeness, I provide a detailed summary of the different pipelines and the logic that led to this conclusion.

In the following section, I discuss the three key considerations when developing the processing pipelines. These were as follows:

1. Which software to use for standard pre-processing (fMRIPrep or MELODIC/FEAT)
2. Whether or not to spatially smooth the data, and if so, by how much and whether smoothing should be applied in the volume or on the surface
3. How to generate surface spatial maps of ICA components. For this, I examined three options:
    i. Sample the spatial maps to the surface following decomposition.
    ii. Perform ICA directly on surface data.
    iii. Regress component timecourses into surface data.

2.2.1.1 Pre-processing package: MELODIC vs. fMRIPrep

Pre-processing is intended to isolate and remove nuisance signals from the BOLD data to reduce false positive and false negative sources by boosting the signal to noise ratio (SNR). Pre-processing steps also allow volumes to be aligned, thus correcting for motion of the head. Multiple pre-processing steps are often incorporated into a single pre-processing pipeline designed to prepare the data for further analysis.

Well-known software packages, such as FSL, FreeSurfer, SPM, AFNI and HCP workbench allow the user to perform standard pre-processing such as head-motion correction, intra-subject registration, and surface sampling. Although these steps are considered standard in most pre-processing pipelines, the package used is often selected at the researcher's discretion.

The first decision to make when developing a processing pipeline is which software package to use for data pre-processing. Different research communities often use different software packages. However, in 2012 it was estimated that 80% of analysis pipelines implemented at least one of FSL, SPM or AFNI (Carp, 2012). An integration of two or more of these packages into a single pre-processing pipeline is often necessary. For example, one may use FSL for functional data pre-processing and FreeSurfer for surface sampling.

The MELODIC GUI (part of the FSL package) offers an all-in-one pipeline for volume-based analysis, in which typical data pre-processing is applied before ICA is performed on these data. I performed pre-processing with the pre-selected parameters; motion-correction, spatial smoothing, and high-pass filtering to remove slow drifts. I refer to this as the "default pipeline".

An alternative to using FSL, or any other combination of software packages, is to use fMRIPrep (Esteban et al., 2019, 2020). fMRIPrep is an optimised and standardised tool developed to take task-based and resting-state fMRI data from raw to fully pre-processed form. The tool allows users to optimally process their data by implementing the best tools from a range of different software packages, such as FSL and FreeSurfer, into a single easy-to-use workflow. Aside from being optimised and standardised, the ease of use of fMRIPrep makes it an enticing tool for neuroimaging researchers. Therefore, I pre-processed data either using FSL with the default parameters (using the MELODIC GUI) or with fMRIPrep, again performing standard pre-processing steps.

2.2.1.2 Spatial smoothing

Spatial smoothing is a common pre-processing step used in neuroimaging research. Carp (2012) reported that from a random sample of 300 published fMRI research papers, 88% reported applying

spatial smoothing prior to statistical analysis, typically by convolving each voxel with a 3D Gaussian kernel of a specific size (Gaussian smoothing) to increase the signal-to-noise ratio (SNR).

If smoothing is applied, the size of the filter, expressed as the width of a Gaussian FWHM kernel, must be selected by the researcher. It has been suggested that the effective smoothing kernel size is one that is twice the size of the voxels (Worsley & Friston, 1995), but in practice there is little consensus as to how to select the optimal smoothing kernel. Carp (2012) found that not one published fMRI paper in their sample detailed why they selected their chosen kernel size. In the absence of any clear justification, it appears that most studies chose the default smoothing kernel in the pre-processing package they used. For example, in a review, for studies using SPM, 64% report 8mm, the default (Carp, 2012).

Rather than relying on the default, I used empirical evidence to select the optimal smoothing level for reliably identifying visual areas with ICA. Asquith (2018) systematically investigated the sensitivity of ICA to spatial smoothing. Data were smoothed at 0, 3 and 8 mm and compared to the default 5 mm smoothing level used in MELODIC/FEAT. Results showed 5 mm smoothing to be optimal for the dataset I am using. Therefore, in my pipelines, when data is smoothed, a 5 mm FWHM kernel is used.

A primary aim of fMRI research is to precisely locate areas of cortical activation. However, the folding of the cortex can make spatial smoothing in the volume potentially problematic for achieving this aim. For example, neighbouring points on the folded 3D brain may represent locations that are distant in the cortical 2D space, e.g., two points located on opposing banks of a sulcus. One recent paper showed two points within the pre- and post-central gyrus to be 4 mm apart in the 3D volume space, but 40 mm apart when mapped to the surface (Brodoehl et al., 2020). These two points belong to functionally different systems but smoothing in the volume would lead to contamination of their functionally distinct signals. This contamination can make the separation of signals problematic and lead to issues in the localisation of functional regions.

It is important to note that the limitations of volumetric smoothing are a strong argument against its use when pre-processing data with high spatial resolution, e.g., HCP data (Glasser et al., 2016). Spatial

smoothing can negate many of the benefits of acquiring such high-resolution data, the goal of which is to allow for more accurate localisation of activation. So much so, HCP's "minimal pre-processing pipeline" does not include any overt volume-based spatial smoothing (Glasser et al., 2013). However, it has been suggested that acquiring data at a higher spatial resolution than is required and then smoothing to a lower resolution can improve SNR (Triantafyllou et al., 2006). SNR can be improved further if smoothing (from high to low resolution) is constrained to the cortical ribbon, thus minimising the contribution of CSF and white matter signals (Blazejewska et al., 2019).

Surface-based smoothing offers a partial solution to these contamination and localisation issues. Surface-based smoothing, much like the volume alternative, generates a new value for each surface point (vertex) by taking a weighted average of points within the smoothing kernel. Since the surface is a 2D mesh instead of a 3D grid as is the case with volume data, smoothing across the surface results in averaging values of neighbouring vertices, avoiding contamination across sulci and gyri. Likewise, the surface is a 2D mesh of the grey matter, smoothing across these vertices removes the issue of mixing tissue type and only activation which originates in the cortex is carried through to the analysis.

I examined several different smoothing approaches, the details of which are outlined in detail below. In short, I performed (i) standard 5 mm volumetric smoothing prior to volume ICA, (ii) surface-based smoothing prior to performing surface-based ICA and (iii) regression of the timecourses derived from performing ICA on an unsmoothed volume into smoothed surface data (see section 2.2.1.3).

2.2.1.3 Generating surface components

The goal of this pipeline was to have components on the individual's surface to allow for direct comparisons with the surface-based probabilistic atlas described in section 2.1.5. Broadly speaking, there are three ways to achieve this:

1. Perform volume-based ICA decomposition before sampling the spatial maps to the individual's surface.
2. Perform ICA decomposition directly on the surface-based data.

3. Perform volume-based ICA decomposition before regressing the component timecourses into the native mesh fMRI data (Glasser, personal communication, July 6th, 2020).[6]

Both the regression-based pipelines (ICA1, ICA2, ICA7 and ICA8) and the surface-based ICA decomposition pipelines (ICA4 and ICA5) require a surface version of the BOLD data in CIFTI (Connectivity Informatics Technology Initiative) file format to be generated. In the case of the former, the timecourse of the components are regressed into the CIFTI file, in the case of the latter MELODIC requires the input data to be in CIFTI format. The process for generating CIFTI files is described in section 2.2.3.

2.2.1.4 Evaluation of pipelines

I evaluated each of the eight pipelines with two spatial correlation measures: the mean correlation between the components and a probabilistic atlas of visual areas and the mean number of unique components produced by the subjects. Both these metrics are described in detail below. Three separate datasets, described in section 2.2.2, were used to perform the pipeline evaluations.

2.2.1.5 Estimation of dimensionality

Selecting the optimal number of output components (i.e., setting the dimensionality) is an important element of performing ICA. However, there is no clear rationale for picking one dimensionality over another. Selecting too few components may underfit the data (resulting in distinct signals being grouped into a small number of components), selecting too many components may overfit the data (resulting in individual underlying signals being split into many separate components).

By default, MELODIC attempts to solve this under/overfitting problem using a Bayesian dimensionality estimation to calculate the optimal number of output components. This automatic dimensionality estimation is utilised in ICA2, ICA3 and ICA6. For the remaining pipelines, MELODIC

[6] Variance normalisation and linear detrending was performed on data used in the pipelines developed following HCP recommendations.

cannot automatically estimate dimensionality. Instead, HCP's "*icaDim*" function is used in MATLAB, before the dimensionality estimate is manually specified in the MELODIC command line. Both *"icaDim"* and MELODIC use a similar dimensionality estimation approach (Glasser, personal communication, January 8th, 2021).

Comparing the dimensionality allows for an initial assessment of the consistency between pipelines. I calculated the mean dimensionality (the number of output components) across subjects for each pipeline and each of the three datasets. However, with respect to dimensionality, "more" does not necessarily mean "better". Many of the outputs could be noise components or could be components relating to non-grey-matter signal. Therefore, we needed a more specific measure of which pipeline produced components corresponding to the visual areas. For this, we computed spatial correlations between each set of spatial maps and the probabilistic atlas of visual areas (see below).

2.2.1.6 Spatial correlations

2.2.1.6.1 The group grand r ($\bar{r}$)

Given our goal was to employ this pipeline to identify visual areas of the brain, we use the spatial correlations between each subject set of the components derived from ICA and a probabilistic atlas of visual areas to evaluate the pipelines.

The atlas used for these correlations was taken from work by Wang et al. (2015), and is described in detail in section 2.1.5. In summary, the atlas contains 25 topographically defined regions (see Table 1) for each hemisphere in surface-based standardised space. Both the FPM and MPM are warped from standard space to each of the subject's surfaces using the Neuropathy docker (Benson & Winawer, 2018).

Table 1. The 25 topographically defined regions taken from Wang et al. (2015), split into ventral-temporal, dorso-lateral and parietal and frontal regions.

Probabilistic atlas ROIs		
Ventral–temporal	**Dorso-lateral**	**Parietal and Frontal**
V1v	V1d	IPS0
V2v	V2d	IPS1
V3v	V3d	IPS2
hV4	V3A	IPS3
VO1	V3B	IPS4
VO2	LO1	IPS5
PHC1	LO2	SPL1
PHC2	TO1	hFEF
	TO2	

Spatial correlations were performed unilaterally, i.e., correlations were computed between left hemisphere component maps and left hemisphere FPM atlas regions, using MATLAB's "*corr*" function (Pearson's r). From these correlations I identified, for each subject and each visual area, the component with the highest spatial correlation, the "top match".

An individual "grand r" was then calculated for each subject by first taking the mean of the left and right hemisphere correlation coefficient for each visual area (resulting in a bilateral visual area coefficient), and then taking the mean of the bilateral visual area coefficients for each of the visual areas. Figure 10 shows a schematic diagram of the process for calculating the individual grand r. Finally, I took the mean across all subjects individual grand r values to generate a group-level grand r (referred to by the symbol $\hat{r}$).

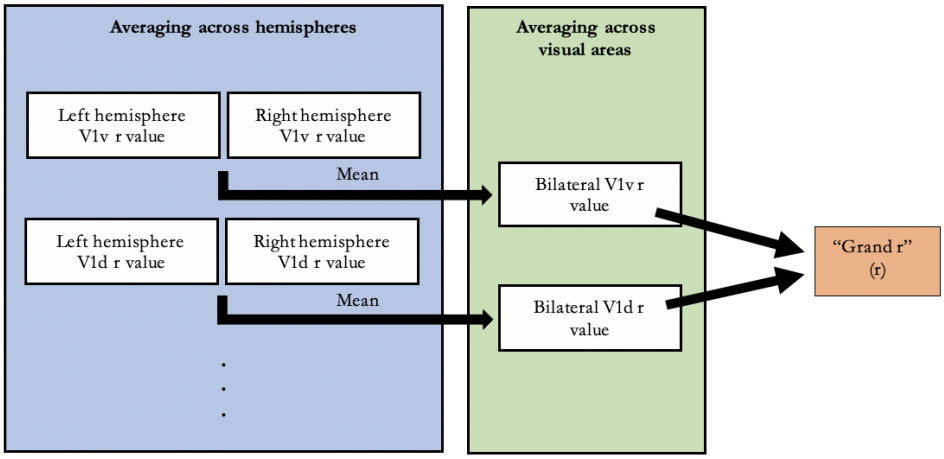

Figure 10. Schematic diagram of the individual grand r calculation.

2.2.1.6.2 Number of unique components

The number of unique components is a simple measure of the effectiveness of a pipeline. A unique component is one that was the top match to only one visual area. A higher number of unique components suggest that individual visual areas were identified, rather than neighbouring areas being grouped together into a large, single component. For each pipeline and each dataset, I calculated the mean number of unique matches across both hemispheres and all subjects.

2.2.1.6.3 Overall pipeline performance (opp)

Following the calculation of the group grand r and the number of unique components for each pipeline, we multiply one by the other to give us an overall picture of which pipelines perform best. We refer to this overall pipeline performance metric with the "*opp"* abbreviation. This method accounts for differences in the number of output components and the quality of the matches between the spatial maps and the atlas regions.

2.2.1.6.4 Data masking

In order to focus the analysis on the areas of interest a mask was applied to the component spatial maps derived from ICA prior to performing spatial correlations. The mask was used to exclude activation in vertices outside of the visual areas, for example in the frontal lobes, that may otherwise result in a poor spatial correlation between the components and the visual area maps. For example, consider a component with discrete activation centred in V1v, but with a small amount of activation located in the frontal regions. This additional frontal activation would pull the spatial correlation with the atlas region down, despite the component potentially being an accurate representation of V1v.

To generate the visual area mask, the MPM was converted to individual ROIs ("labels"), each corresponding to one of the visual areas, using FreeSurfer's "*mri_cor2label*" command. This resulted in 50 labels, 25 for each hemisphere. Each of the 25 visual area labels were then merged into a single label using

FreeSurfer's "*mri_mergelabels*" (Figure 11, blue). The combined atlas was then dilated using "*mri_label2label*" (Figure 11, yellow).

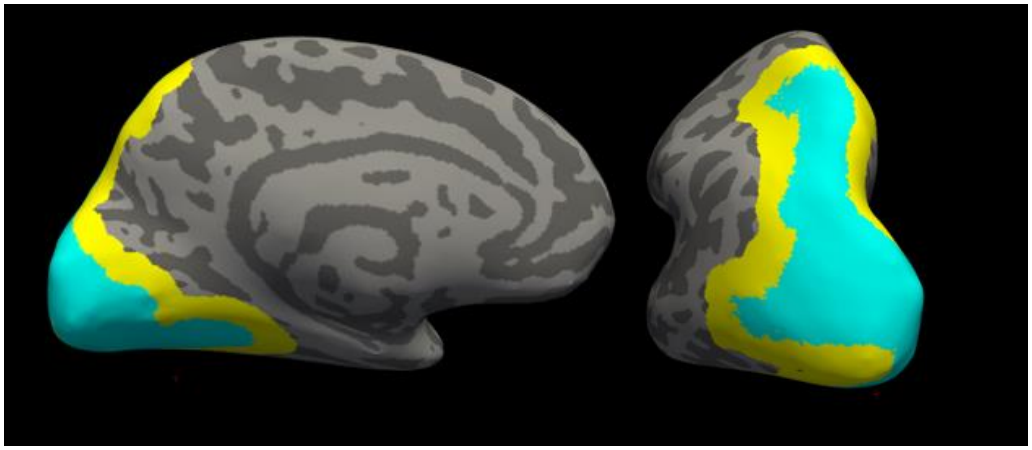

Figure 11. Surface visual area mask generated by combining all 25 visual area ROIs into a single ROI (blue). The mask was then dilated (yellow).

2.2.2 Structural and functional brain imaging acquisitions

Three different datasets were used to examine each of the eight pipelines; MTW, PMA and MEW (acronyms refer to the author of each dataset). The acquisition details of each dataset are below.

2.2.2.1 MTW dataset

Images were acquired using a 3T Siemens Magnetom® Prisma (Siemens Medical Systems, Erlangen, Germany) MRI scanner at CUBRIC fitted with a 32-channel head coil.

The parameters used for the acquisition of the high-resolution anatomical scan (3D sagittal T1 Magnetisation Prepared Rapid Gradient Echo [MPRAGE] sequence) were: 176 slices, TE = 3.24 ms, TR = 2200 ms, voxel size = 1.0 × 1.0 × 1.0 mm^3, flip angle = 9°.

Activation images were acquired using echo-planar imaging (EPI) with the following parameters: 72 slices, TE = 30 ms, TR = 1600 ms, multi-band acceleration factor = 4, voxel size = 2 × 2 × 2.5 mm^3, flip angle = 60°, phase-encoding direction = A-P. A total of 700 volumes were acquired. These EPI scans generated full brain coverage encompassing the whole of the visual cortex. A single volume was acquired with the opposite phase-encoding direction to aid distortion correction during MRI analysis. Data was collected for four subjects.

2.2.2.2 PMA dataset

Images were acquired using a 3T General Electric (GE) Signa 3T HDxt Scanner (General Electric, Milwaukee, Wisconsin), equipped with echo-speed gradient coil and amplifier hardware using a GE 8 channel receiver RF coil.

The parameters used for the acquisition of the high-resolution anatomical scan (fast-spoiled gradient echo [FSPGR] sequence) were: 256 slices, TE = 3.29 ms, TR = 9 ms, voxel size = 1.0 × 1.0 × 1.0 mm^3, flip angle = 20°.

Activation images were acquires using EPI with the following parameters: TR = 2000 ms, TE = 35 ms, voxel size = 3.0 × 3.0 × 3.0 mm^3, flip angle = 90°. A total of 600 volumes were acquired. Data was collected for five subjects.

2.2.2.3 MEW dataset

Images were acquired using a 7T Siemens Magnetom® (Siemens Medical Systems, Erlangen, Germany) MRI scanner at CUBRIC fitted with a 32-channel head coil.

The parameters used for the acquisition of the MPRAGE sequence were: TE = 2.82 ms, TR = 2200 ms, voxel size = 1.0 × 1.0 × 1.0 mm3, flip angle = 7°.

Activation images were acquires using EPI with the following parameters: TR = 2200 ms, TE = 29 ms, voxel size = 1.3 × 1.3 × 1.5 mm^3, flip angle = 60°. A total of 270 volumes were acquired. Data was collected for five subjects.

2.2.3 ICA Pipelines

fMRIPrep version 20.2.0 with default parameters was used for data pre-processing for all bar one pipeline (ICA2). The pre-processing boilerplate text generated by fMRIPrep was copied verbatim (as the developers intended) and can be found in Appendix I. For ease, I have summarised the steps below:

The anatomical image (T1w) was first bias-corrected and skulled stripped. Brain surfaces were then reconstructed, and a brain mask was refined. Segmentation of white matter, grey matter, and CSF was performed on the brain extracted T1w image.

Functional data were slice-time and motion corrected before being co-registered to the high-resolution anatomical image using brain-boundary registration (BBR). Distortion correction is also performed using AFNI commands. High-pass filtering was then performed on the BOLD time-series. These data are referred to as the "post-fMRIPrep" data onwards. Table 2 summarises the pre-processing steps used in each of the eight pipelines. A more detailed explanation of each pipeline is below.

Table 2. Pipeline summary.

	Pre-processing	**Normalisation/ detrending**	**Smoothing level of input data**	**Dimensionality estimation**	**MELODIC version**	**Resample to surface**
ICA1	fMRIPrep	Yes	None	CalcDim	Command line	Regression to unsmoothed CIFTI
ICA2	MELODIC GUI	No	5mm during MELODIC GUI	Automatic	GUI	Vol2Surf
ICA3	fMRIPrep	No	None	Automatic	Command line	Vol2Surf
ICA4	fMRIPrep	Yes	None	CalcDim	Command line	Unsmoothed CIFTI
ICA5	fMRIPrep	Yes	5mm smoothing applied to input CIFTI	CalcDim	Command line	Smoothed CIFTI
ICA6	fMRIPrep	No	5mm smoothing applied to post-fMRIPrep input data	Automatic	Command line	Vol2Surf
ICA7	fMRIPrep	Yes	None	CalcDim	Command line	Regression to Smoothed CIFTI
ICA8	fMRIPrep	Yes	Smooth post-fMRIPrep	CalcDim	Command line	Regression into unsmoothed CIFTI

2.2.3.1 ICA 1

The post-fMRIPrep BOLD data was first skull-stripped by applying the brain mask derived from fMRIPrep. The FSL's "*fslmaths*" command with the *"-mas"* flag was used to achieve this. In line with HCP

recommendations (Glasser, personal communication, June 19th, 2020), an estimation of the dimensionality of the post-fMRIPrep volume data was then generated using HCP's "*icaDim*" function.

Following this and upon advice from Glasser (personal communication, June 19th, 2020), variance normalisation was achieved by dividing the data by the unstructured noise standard deviation[7]. This data was then linearly detrended along the time dimension using MATLAB's "*detrend*" function. Finally, the NIFTI volume file was recreated from the variance normalized and detrended data. ICA was then performed on these data using the MELODIC command line, with the estimated dimensionality manually specified.

Next a CIFTI dense timeseries (a surface version of the BOLD data, see section 2.2.1.3) was created from the left and right hemisphere surface-based functional data. These surface-based data are output by fMRIPrep in GIFTI format and converted to CIFTI format with HCP's "*cifti-create-dense-timeseries*" command.

Following the completion of MELODIC, the ICA timecourses and the CIFTI dense timeseries were loaded into MATLAB. The CIFTI file was linearly detrended, before the ICA timecourses were regressed into the CIFTI timeseries. The CIFTI file, now containing the ICA timecourse was then output before it is split into left and right hemisphere GIFTI files using HCP's "*cifti-separate*" command.

2.2.3.2 ICA2

ICA2 is the "default MELODIC" pipeline, in which data were pre-processed using the standard MELODIC/FEAT parameters rather than through fMRIPrep. The three datasets were pre-processed as described below.

[7] This pre-processing step and corresponding MATLAB script was recommended by Mathew Glasser (personal communication, July 6th, 2020).

2.2.3.2.1 Distortion correction field map

In order to correct for B0 inhomogeneities, a distortion correction was applied to the EPI images within MELODIC. The process for generating the field map used for this correction varied between the MTW and MEW datasets. Both are described below. Note, no distortion correction was applied to the PMA dataset because it was not available.

2.2.3.2.2 MTW dataset distortion correction

A single volume was acquired, with identical scan parameters to those used during the movie watching task except the phase encoding (PE) direction was flipped (P-A). This acquisition followed the completion of the movie watching task. The field map was generated during pre-processing using the standard FSL approach: The final volume of the unprocessed EPI data was extracted (using *"fsl_roi"*) and subsequently merged (using *"fsl_merge"*) with the image acquired during the PE flipped acquisition. The FSL tool *"topup"* was then used to generate the field map, before it was converted from units of Hz to units of rad/s. The field map was then applied during the FEAT pre-processing steps described below.

2.2.3.2.3 MEW dataset distortion correction

A phase-difference image and a magnitude image were output by the SIEMENS scanner used for this dataset. First, the magnitude image was brain-extracted using FSL's BET tool. Next, FSL's "*fsl_prepare_fieldmap*" command was used to generate the fieldmap, using the brain-extracted magnitude image and the phase-different image.

2.2.3.2.4 Data processing and MELODIC ICA

The EPI data were pre-processed in line with the MELODIC default settings using the MELODIC graphical user interface (GUI), part of the FSL toolbox (Jenkinson et al., 2012; S. M. Smith et al., 2004; Woolrich et al., 2009). In order to reduce noise without removing valid activation, each volume of the dataset was spatially smoothed at 5 mm FWHM. Low frequency artefacts, such as low frequency temporal drifts, were removed from the dataset using a high-pass temporal filter set to 100 Hz. The EPI

was brain extracted using BET. MCFLIRT (part of the FSL package), was used to correct for motion artefacts through the application of rigid-body transformations (rotational and translational corrections only). Data was resampled to match the resolution of the input data (MTW dataset: 2 mm, PMA dataset: 3 mm, MEW dataset 1.5 mm). The distortion correction field map was applied by inputting the field map in the units of rad/s and the brain extracted magnitude image from the field map acquisition (MTW and MEW datasets only). ICA was then performed on the pre-processed fMRI data using the FSL MELODIC GUI, version 3.0 (Jenkinson et. al., 2012). The spatial maps were then co-registered to the high-resolution anatomical image using FreeSurfer's *"bbregister"* command, with the *"-feat"* flag added. This flag uses the post-MELODIC example_func image as the "moveable" volume, and outputs the anat2exf.register.dat registration file. This registration file was subsequently used when resampling the spatial maps to the individual's surface (see below).

2.2.3.2.5 Resampling to the surface

FreeSurfer's *"mri_vol2surf"* function was used to resample the spatial maps (melodic_IC file) to the individual's surface. This command assigns the values from the volume data to the surface vertices. The transformation file generated by *"bbregister"* was utilised here. In order to project to the middle of the cortical surface, the projection fraction was set to 0.5.

2.2.3.3 ICA3

ICA3 involves performing volume-based ICA on post-fMRIPrep data using the MELODIC command line. To assess the impact of smoothing, and in line with HCP recommendations, no smoothing was applied to the data in this pipeline. Unlike ICA1, dimensionality was automatically estimated by MELODIC. As with ICA2, the post-MELODIC spatial maps were co-registered to the anatomical image using *"bbregister"*. Again, this stage outputs the appropriate transformation file needed to resample the spatial maps to the subject's surface, for which *"mri_vol2surf"* was used. However, the *"bbregister"* stage in this pipeline was different to that described in ICA2. As this pipeline does not utilise any of the MELODIC pre-processing steps, including registration, the accompanying FEAT files needed to use

"bbregister" with the -feat flag are not generated. As such we use *"bbregister"* with the default FreeSurfer initialization arguments, outputting the FreeSurfer registration file in LTA format.

In order to check whether the transformation file was appropriate in the absence of the example_func image, the middle volume of the post-fMRIPrep data was extracted using *"fslroi"*. This middle volume was then viewed, with the transformation matrix applied, in tkregisterfv. This program displays the white matter surface of the anatomical image overlaid on the EPI single volume data, allowing us to ascertain whether the white matter boundaries are correctly spatially aligned. If we are satisfied that the transformation file was appropriate, we use it to resample the spatial maps to the surface using *"mri_vol2surf"*.

2.2.3.4 ICA4 and ICA5

Both ICA4 and ICA5 use CIFTI data as the input to the MELODIC command line (see ICA1 for how this file was created). The only distinction between the two pipelines was that the data used in ICA5 is smoothed before any further processing, whereas ICA4 uses unsmoothed data. As a result of this smoothing, the estimation of dimensionality (as calculated by the *"icaDim"* function) was different for the two pipelines. Surface-based smoothing of the CIFTI data was performed using HCP's "*cifti-smoothing*" command. Data were smoothed in the column (across space), i.e., 1D smoothing was applied, with a 5mm FWHM Gaussian kernel.

The following steps apply to both pipelines. Following HCP advice, the CIFTI file (smoothed or unsmoothed) was loaded into MATLAB and dimensionality was estimated. The data was then normalised and linearly detrended as described in ICA1. The normalised and detrended CIFTI data are then used as the MELODIC input file. As with ICA1, the dimensionality estimation was manually specified in the command line. Following MELODIC, the spatial maps are separated into left and right hemisphere GIFTI files using HCP's "*cifti-separate*" command.

2.2.3.5 ICA6

This pipeline is identical to ICA3, except that the post-fMRIPrep data were spatially smoothed prior to ICA. The smoothed volume data was then input to the MELODIC command line, with the dimensionality estimation set to automatic. As with ICA3, we extract the middle volume of the smoothed BOLD data using "*fslroi*" and assess whether the transformation file generated by *"bbregister"* appropriately aligned the functional data to the anatomical image. We then used this transformation file to resample the spatial maps to the subject's surface using "*mri_vol2surf*" with the same parameters as described in ICA2.

2.2.3.6 ICA7

The ICA timecourses computed in ICA1 (unsmoothed volume data) were loaded into MATLAB, linearly detrended and then regressed into the smoothed CIFTI fMRI data. The CIFTI surface data was then separated into left and right hemisphere GIFTI files using HCP's "*cifti-separate*".

2.2.3.7 ICA8

The ICA timecourses computed in ICA6 (smoothed volume data) were loaded into MATLAB, linearly detrended and then regressed into the unsmoothed CIFTI fMRI data. The CIFTI surface data was then separated into left and right hemisphere GIFTI files using HCP's "*cifti-separate*".

2.2.4 Results

2.2.4.1 Dimensionality

I report the differences in dimensionality across pipelines for each dataset in Figure 12. It is not logical to compare the dimensionality of the MEW data with the MTW/PMA datasets due to the differences in brain coverage of the EPI. We can, however, assess differences in the number of dimensions between the MTW and PMA datasets as both have near-full brain coverage and approximately the same number of volumes (MTW = 700 volumes, PMA = 600 volumes). However, one must keep in mind that surface-based pipelines will inherently output fewer components than their volume-based counterparts. This is in part due to the removal of white matter and CSF from these data. Furthermore, as the MEW

dataset has partial brain coverage, the number of output components will be reduced compared to the MTW and PMA whole brain coverage datasets.

On average, unsmoothed surface data (ICA4) yields fewer components than smoothed surface data (ICA5). This pattern is consistent across datasets and may indicate that smoothed data increased the SNR compared to unsmoothed data. Looking specifically at volume-based ICA decompositions, dimensionality is highest for ICA2 (MELODIC default pipeline, smoothed data) for the PMA dataset, whilst ICA3 (post-fMRIPrep, unsmoothed volume ICA) leads to higher dimensionality for the MTW and MEW datasets. Conversely, ICA1 (HCP recommended pipeline, unsmoothed volume ICA regressed into unsmoothed surface data) outputs the fewest components for both the MTW and PMA datasets, whereas ICA2 produces the fewest components for the MEW dataset. As ICA7 and ICA8 utilise the timecourses computed in ICA1 and ICA6 respectively, dimensionality is unchanged for these pipelines.

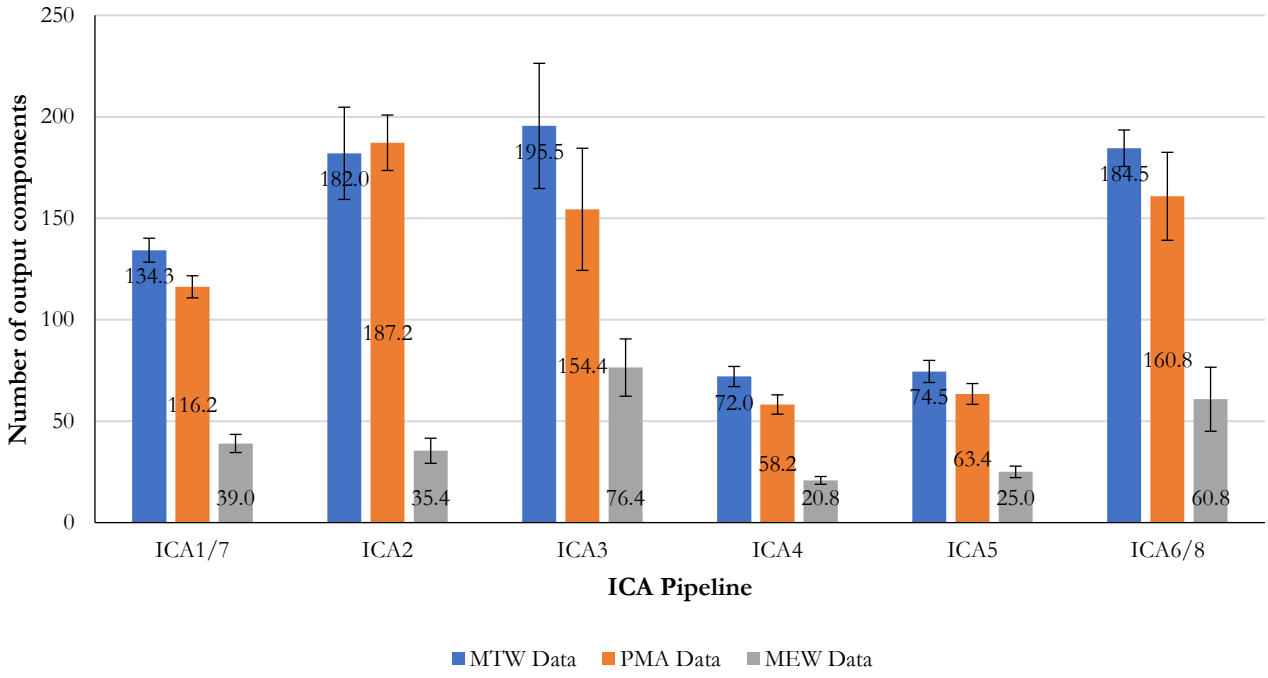

Figure 12. Mean dimensionality (number of output components) of the three datasets and eight pipelines used for pre-processing. Note, ICA7 and ICA8 utilise the timecourses computed in ICA1 and ICA6 respectively, therefore dimensionality is unchanged for these pipelines. Error bars represent standard deviation.

2.2.4.2 The group grand r ($\hat{r}$) for each pipeline

Due to the whole brain coverage of the EPI, 25 atlas regions were used when computing spatial correlations for both the MTW and PMA datasets. Whereas due to the partial FOV of the EPI in the MEW dataset only 17 atlas regions were used (IPS0-IPS5, SPL1 and FEF were not included in the correlations). This results in 50 correlation coefficients for each subject (per pipeline), for the MTW and PMA datasets, and 34 correlation coefficients for the MEW datasets. As outlined in section 2.2.1.4, a group-level grand r ($\hat{r}$) was calculated for each dataset and each pipeline. Note, $\hat{r}$ for each pipeline was calculated by taking the mean correlation across all visual areas (both the left and right hemispheres) and all brains in the dataset (Figure 13).

Overall, ICA6 produces components with the strongest mean correlation to the atlas regions for the MTW dataset ($\hat{r}$ = .448) and MEW dataset ($\hat{r}$ = .426). Whereas ICA2 produces components with the strongest mean correlation to the atlas regions ($\hat{r}$ = .475) for the PMA dataset. ICA2 and ICA6 are comparable pipelines; they both involve volume-based ICA on smoothed data. The key distinction is that for the former all pre-processing (including smoothing) is handled by the MELODIC GUI, for the latter fMRIPrep is used to first pre-process the data before it is spatially smoothed.

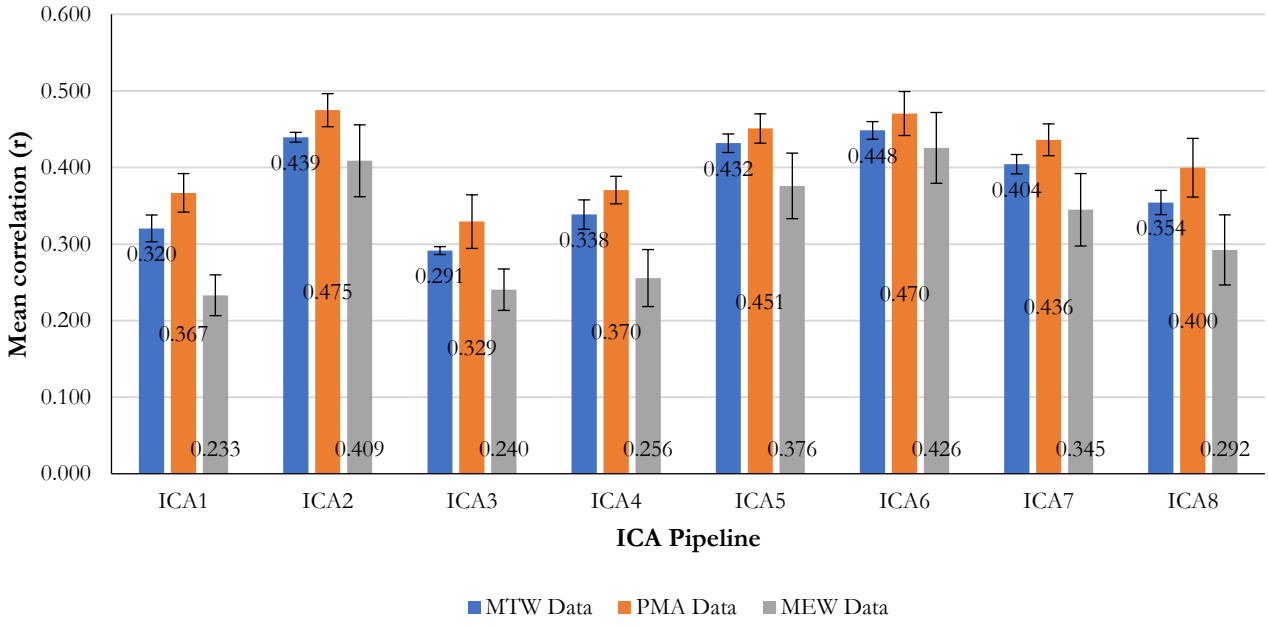

Figure 13. The group grand r of each pipeline for each dataset. The grand r conveys the mean correlation (across visual areas, hemispheres, and subjects) between components and the atlas. Error bars represent standard deviation.

Volume-based pipelines (ICA2, ICA3 and ICA6)

Of the three volume-based pipelines, ICA6 produces components which, on average, correlate best with the atlas regions for the MTW dataset ($\hat{r}$ = .439) and MEW datasets ($\hat{r}$ = .426). However, ICA2 produces components which on average correlate best with the atlas regions for the PMA dataset ($\hat{r}$ = .475).

CIFTI-based pipelines (ICA4 and ICA5)

Of the two CIFTI-based pipelines, ICA5 (smoothed CIFTI data) produces components which, on average, correlate better with atlas regions than unsmoothed data (ICA4). This is true for the MTW ($\hat{r}$ = .432), PMA ($\hat{r}$ = .451) and MEW ($\hat{r}$ = .376) datasets.

Regression-based pipelines (ICA1, ICA7 and ICA8)

Of the three regression pipelines, ICA7 (regression of unsmoothed data into a smoothed surface) produces components with the highest correlation to the atlas for the MTW ($\hat{r}$ = .404), PMA ($\hat{r}$ = .436) and MEW ($\hat{r}$ = .345) datasets.

2.2.4.3 Number of unique components in the top matches

Figure 14 shows, for each pipeline and each dataset, the mean number of unique components in the top matches across both hemispheres and all subjects. On average, ICA2 and ICA6 output the joint-highest number of unique matches (both 32.3) for the MTW dataset. For the PMA dataset, ICA7 outputs the highest number of unique components (34.8 components), with ICA6 a close second (34.4 components). For the MEW dataset, ICA7 again was the pipeline that outputs the highest number of unique components (22.2 components).

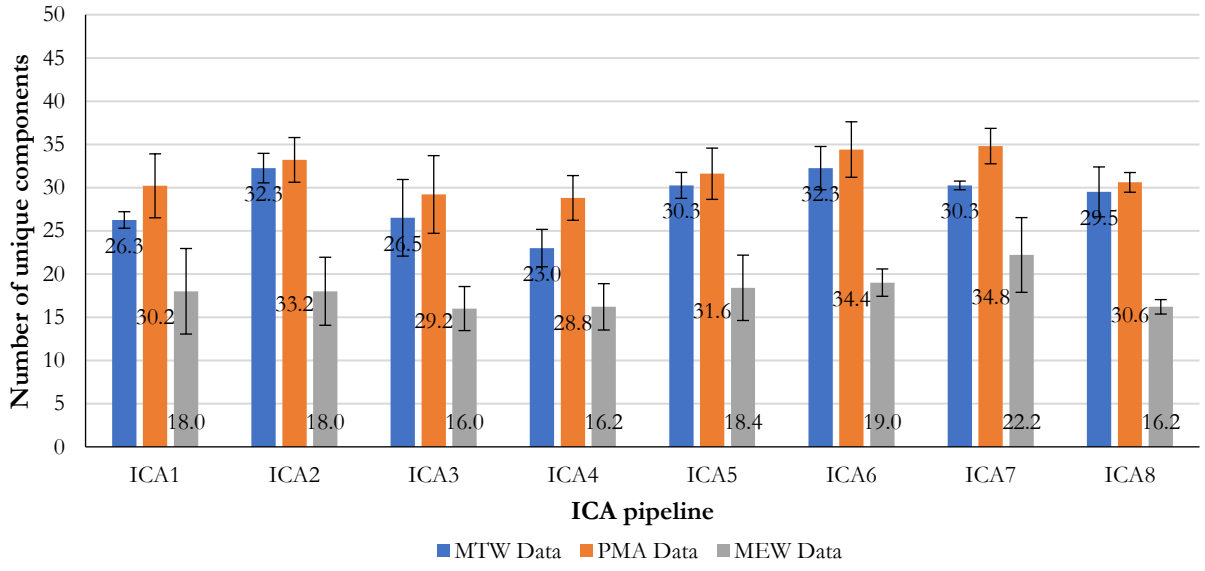

Figure 14. The mean number of unique output components within the top matches (both hemispheres combined) for each pipeline and each dataset. Error bars represent standard deviation.

2.2.4.4 Overall pipeline performance (*opp)*

The overall pipeline performance (*opp*) was measured by multiplying the number of unique components in the top matches by the group grand r (Figure 15). Using this metric, it appears that overall ICA6 is the superior pipeline for the MTW (*opp* = 14.442), PMA (*opp* = 16.147) and MEW (*opp* = 8.118) datasets. Conversely, ICA3 appears to be the least effective pipeline for the MTW (*opp* = 7.718), PMA (*opp* = 9.505) and MEW (*opp* = 3.829) datasets.

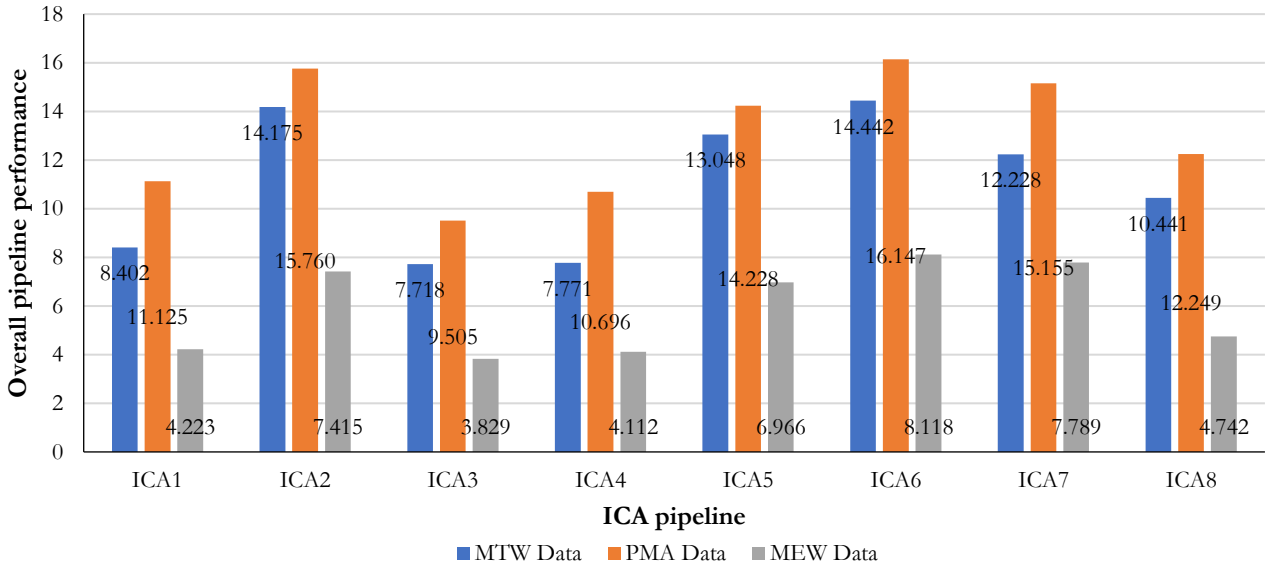

Figure 15. Overall pipeline performance for each pipeline and each dataset. Calculated by multiplying the group grand r by the number of unique components for each pipeline.

CIFTI-based pipelines (ICA4 and ICA5)

Of the two CIFTI-based pipelines, ICA5 (smoothed CIFTI data) appears to be the most effective pipeline for the MTW (*opp* = 13.048), PMA (*opp* = 14.228) and MEW (*opp* = 6.966) datasets.

Regression-based pipelines (ICA1, ICA7 and ICA8)

Of the three regression pipelines, ICA7 (regression of unsmoothed data into a smoothed surface) appears to be the most effective pipeline for the MTW (*opp* = 12.228), PMA (*opp* = 115.155) and MEW (*opp* = 7.789) datasets.

2.2.5 Conclusion

Overall, ICA6 appears to be the optimal pipeline for all datasets. This pipeline involves pre-processing data with fMRIPrep, smoothing these data following pre-processing before ICA is performed on these data. The spatial maps of these components are then sampled to the surface. Conversely, ICA3 appears to be the least optimal pipeline for all three datasets. This pipeline was identical to ICA6 with the only difference being that no smoothing was applied at any stage.

With respect to surface-based decomposition, smoothing the input CIFTI data appears to lead to better matches with the atlas regions than using unsmoothed CIFTI data as the input to MELODIC. Of the regression-based pipelines, it appears regressing the timecourses derived from unsmoothed volume ICA into smoothed CIFTI data is best for all datasets.

Table 3 shows the mean correlation (Pearson's r) for each visual area following processing using the ICA6 pipeline. This highlights which visual areas are consistently accurately identified across datasets using this pipeline. TO1 appears to be the visual area that is identified most successfully, whilst V3v is the visual area that is least so. I explore these spatial corelations in detail in subsequent chapters.

Table 3. Group mean correlation (Pearson's r) between the top matched component and each probabilistic atlas visual area.

Visual Area	Bilateral r
TO1	0.654
IPS2	0.617
TO2	0.581
IPS3	0.56
SPL1	0.535
IPS1	0.504
PHC2	0.501
V2d	0.498
IPS4	0.498
LO1	0.458
FEF	0.446
V3B	0.444
LO2	0.444
IPS0	0.438
V1d	0.438
PHC1	0.419
IPS5	0.412
V3d	0.391
V3A	0.384
VO2	0.377
V2v	0.377
V1v	0.372
hV4	0.372
VO1	0.346
V3v	0.338

Chapter 3. Use of ICA for identifying visual areas

3.1 Introduction

A primary goal of neuroscience research is to identify, and assign function to, the brain's specialised subdivisions. Traditionally this involves an extensive battery of time-consuming localisers, each designed to stimulate a specific region in isolation. Such specific activation often requires equally specific non-naturistic stimuli (e.g., long periods of faces alternated with long periods of objects), often resulting in tedious experiment paradigms. Aside from being tedious, such paradigms can have financial implications for the research group due to the costs of scan and staff time. These factors, combined with the associated analysis pipelines, makes the segmentation of the brain into constituent parts using traditional approaches a significant undertaking.

There is, however, a potential alternative. Bartels and Zeki (2004b, 2005b) demonstrated a technique they called "chronotopy" which exploits the "temporal fingerprints" of functionally distinct brain regions.

Natural visual input can be conceptualised as a series of overlaid sequences of features, a sequence of faces, a sequence of places, a sequence of global motion within the image, a sequence of contrast defined edges, etc. The timecourse of each feature in the natural visual input will be unique, for example, faces may appear every few seconds, whereas cars may seldom appear but then appear for an extended period of time. Because brain regions have differential sensitivity, e.g., MT+/V5 has a particular sensitivity to visual motion and FFA has a particular sensitivity to faces, with natural visual input, the timecourse of activation of each visual region, would be expected to be unique. The unique timecourses of features, and hence different visual regions, should therefore be dissectible with techniques for blind signal separation (the same principle that underlies traditional resting state analyses).

Bartels and Zeki used a Hollywood action film to provide naturalistic visual stimuli to observers as they were scanned in an MR scanner and then subsequently performed ICA on the data.

In their pioneering work, Bartels and Zeki showed visually responsive regions, specifically, V1, V2v, V3v, V5/MT+, and LO regions, could be reliably identified across subjects during natural viewing. The bilateral components derived from ICA showed strong anatomical correspondence across subjects, in that the significant voxels of the components were located in a similar spatial location across brains. The identification and matching of components across brains was performed visually by the authors. Figure 16 shows for V5, V1 and V4 the spatially corresponding components and corresponding ATCs for three of Bartels and Zeki's subjects. The figure shows the peak/hotspots of the components are located in approximate locations across brains, although they differ in size and shape. The components also corresponded temporally in an area-specific manner. The ATCs of spatially corresponding components across brains were significantly correlated (mean r = 0.30), as would be expected if they represented the same functionally specialised brain regions. Conversely, the was no such temporal correspondence between the ATCs of anatomically non-corresponding regions between subjects (mean r = 0.06).

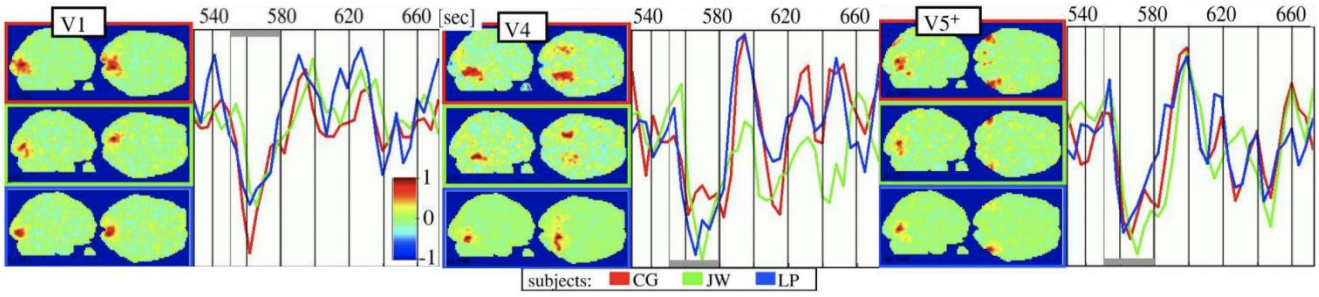

Figure 16. Adapted from Bartels and Zeki (2004b, 2005b). Anatomically corresponding areas identified with ICA in different brains (left). These components were judged to represent cortical areas V1, V4 and V5. The associated timecourses of each subject's component are superimposed, showing a high degree of temporal correspondence across brains (right).

The aim of the following chapter is to build on these findings using a larger dataset and to assess the suitability of ICA for identifying visual areas beyond those reported by Bartels and Zeki. If the approach is effective in a normal population, I had hoped to apply it to patient data, specifically from patients with hemianopia.

The aim of the work in patients was to investigate the differences between healthy brains and the brains of those with cortical vision loss. Chronotopy potentially offers a fast and accurate way of

identifying visual areas, bypassing the need to run a series of functional localisers. Further, the technique has the potential to identify physically intact regions that are functionally disconnected (not part of any bilateral signal components), and of particular interest here, regions that are functionally active, despite being severed upstream from the main (LGN-V1; geniculocortical) visual pathway. The existence of such regions would indicate the involvement of other pathways, and their identification could contribute to the understanding of some forms of spared visual function that can be present in people with hemianopia.

As we are interested in identifying visual areas across subjects, one approach could be to perform tensor ICA (TICA) on the data. TICA combines each subject's 2D time × space matrix along a third dimension (subjects) to generate a 3D time × space × subject matrix. Decomposition is then performed on this block of data. TICA identifies correlations across voxels and between subjects. As with single-subject ICA, a spatial map and timecourse are generated for each component. An additional session/subject interpretation is also output. This additional feature reflects each subject's contribution to the component.

However, in order to do this, TICA assumes the temporal pattern of activation is the same across all subjects. This assumption makes TICA applicable for task-based fMRI in which subjects perform the same time-locked task. However, to utilise this approach, one would have to assume that all subjects were attending to the same elements of the visual stimulus at the same time. This assumption may hold true as previous work has shown only relatively small individual differences in the scan paths of subjects watching movie clips (Tosi et al., 1997) In addition, free viewing of movies has been shown to elicit stable and systematic gaze behaviours across subjects (Hasson et al., 2008; Hutson et al., 2017). However, a visually rich movie stimulus could, in theory, lead to subjects attending to different elements of each scene. Furthermore, any method that is predicated upon fewer assumptions is naturally preferable. Therefore, attention could modulate activation at a given time point, meaning the pattern of activation would not necessarily be aligned across all subjects.

Given the temporal pattern of activation assumption may be violated in our paradigm and that we aimed to apply this technique to a patient population through a series of individual case studies, I opted to perform a series of single-subject ICAs. In single-subject ICA, each component represents one independent signal, and it is therefore a spatiotemporal summary of one part of the data. Each voxel within a spatial map has a value assigned to it. Voxels which contribute strongly to the component have higher values; voxels that do not contribute much to the component have lower values. Thus, when we refer to "strong activation" we are signifying the voxel (or vertex when the components are sampled to the subject's surface) strongly contributed to the component. If the visual system can be dissected into functionally distinct regions entirely by their differential ATCs, this has significant applications to patient populations. Individuals with hemianopia are known to have unstable fixation, making traditional localisers problematic (Reinhard et al., 2014). Moreover, mental fatigue during cognitive tasks is notably prevalent in stroke survivors (Colle et al., 2006). Therefore, being able to robustly identify multiple functionally distinct brain regions in a single short session, without the need for prolonged central fixation, is particularly advantageous for those with cortical vision loss.

ICA maximises the spatial independence between components so there is as little spatial similarity (overlap) as possible between components (Calhoun et al., 2013). However, a single brain region can be involved in many different networks. ICA can still capture true signal even when it overlaps, meaning in some instances a single voxel can be present and strongly contribute to more than one component (Beckmann & Smith, 2004a). This is of particular importance here as recent work investigating brain dynamics has shown that during movie watching, the brain transitions between different functional states (Meer et al., 2020).

Meer et al. (2020) examined the fluctuations in the BOLD signal of regions comprising 14 canonical brain networks. The activity in these networks reconfigured over time and was subsequently decomposed into 10 different brain states. The states, which recur over time, convey the structured patterns of activation that differentially load onto the 14 brain networks. For example, brain state seven is

characterised by high signal loading onto the high visual network, but suppression of signal in the 13 other networks. On the other hand, brain state 10 is characterised by high signal loading on both the salience and executive networks, but low signal in the primary visual network and high visual network. If ICA is sensitive to these discrete brain states, and in some subjects the visual areas are dynamically recruited into different functional states, then voxels may appear in several components representing these different functional states.

The aims of this chapter were to:

1. Investigate whether ICA can accurately identify visual areas of the brain.
2. Assess differences across subjects and across visual areas.

This chapter has two analysis elements, one spatial and one temporal. First, rather than manually identifying anatomically corresponding components, as in Bartels and Zeki (2004b, 2005b), I performed spatial correlations between each subject's set of spatial maps and each of the probabilistic atlas regions (described in section 2.1.5) to identify "matched components". In summary, our spatial correlation results show:

- Although ICA does produce components that spatially correlate with the probabilistic atlas, referred to as matches, the components are not consistently clean representations of individual visual areas.
- There is substantial variability in the quality of the matches across subjects, and across the visual areas, with a key exception being TO1/MT and TO2/MST.
- The spatial maps of TO1/MT and TO2/MST matched components showed consistency insofar as they both regularly appeared in the same component, showed the highest correlation with the atlas regions, and showed no difference across hemispheres in the number of subjects producing a matched component.

- TO1/MT and TO2/MST matched components are fundamentally poor representations of the regions compared to more traditional localisers and are not identified independently of one another, but rather the hMT+ complex is identified.

Next, I performed correlations between the ATCs of matched components across brains and compared their temporal correspondence to those of the atlas ATCs. This timecourse analysis was in-line with work by Bartels and Zeki (2004b, 2005b) who showed that components that anatomically corresponded across brains also showed temporal correspondence, as would be expected if they represented the same functionally specialised brain regions. In summary, our temporal correlation results show:

- The ATCs extracted from the atlas regions are more highly correlated than the ATCs extracted from the components. Again, TO1/MT and TO2/MST were the exception.
- TO1/MT and TO2/MST ATCs were more highly correlated across brains than atlas ATCs, with the difference in mean correlation between component and atlas ATCs being statistically significant.
- Based on our results, although it may have its uses, ICA and movie watching is not an effective tool for robustly identifying individual visual areas.

3.2 Methods

This chapter used the optimal ICA pipeline (as described in section 2.1.5.3). The EPI/functional data was first pre-processed using fMRIPrep. Following fMRIPrep pre-processing, the data was skull-stripped using a brain mask generated by fMRIPrep and before being spatially smoothed in the volume with a 5 mm smoothing kernel. The pre-processed data were then registered to the subject's high-resolution anatomical image using FreeSurfer's "*bbregister*" command. ICA is then performed on these data using the MELODIC command line program. MELODIC operates on the whole brain, both the left and right hemispheres. However, during resampling to the surface the hemispheres are split. The

transformation matrix generated by *"bbregister"* is used to resample the spatial maps to the individual surface using "*mri_vol2surf*". This results in a set of spatial maps for each hemisphere.

For each hemisphere, MATLAB's "*corr*" function is used to compute spatial correlations between each output component and each of the 25 full probability atlas regions (Wang et al., 2015). During this stage, a visual area mask is applied to the spatial maps to exclude activation in vertices outside of the visual areas, for example in the frontal lobes. Activation in these regions may otherwise result in a poor spatial correlation between the component maps and the visual areas atlas (see section 2.2.1.6.2 for details).

3.2.1 Component classification

Following ICA decomposition and the computation of spatial correlations, "matched components" are identified for each subject and each visual area (discussed in detail in section 3.3.1). Throughout this chapter, and the thesis as a whole, differences between brains in the quality and quantity of these matched components are then examined. However, it is important to note that the number of components output from the ICA is not necessarily directly proportionate to the number of underlying true signal components. Each ICA decomposition is comprised of both signal and noise components, with the former interpreted to represent a true underlying neuronal signal and the latter representing structured noise within the data. There are two main categories of noise: subject-specific effects (e.g., head movement) and acquisition related effects (e.g., MR artefacts). Given that noise can vary from subject to subject, so can the proportion of noise components derived from the decomposition. Therefore, differences between brains in the number of components must be treated with caution. There is, however, a hierarchical decision process that can be implemented to ensure that any individual differences in the number of components are genuine and not a result of one subject's dataset containing more noise. The decision process is essential for accurately identifying signal components and requires both an understanding of the likely ratio of signal to noise components, and how sources of noise present in the data. I discuss these elements below.

A high proportion of noise components is to be expected. Traditional resting state paradigms often set dimensionality to 20 or 30 with the aim of identifying the 10 canonical resting state networks (e.g., Smith et al., 2009). From this, we can see that 33-50% (10/20-10/30) of components will be identified as signal (the RSNs), with the remaining components potentially being categorised as noise. However, as Griffanti et al. (2014) notes, components should not be categorised as noise simply because they do not match a common RSN. Pamilo et al., (2012) demonstrated the ICA dimensionality impacts whether a component will encompass the entire RSN (low dimensionality) or whether the network will fragment into sub-networks (high dimensionality). In the case of the latter, these components may reflect sub-networks of the larger RSNs output at a lower dimensionality and allow researchers to estimate the "functional hierarchy" of the networks (S. M. Smith et al., 2009).

For unrestricted decompositions, or when a high dimensionality is set, the proportion of noise components typically outweigh the number of signal components (Griffanti et al., 2017). Previous work has shown that noise components can make up >70% of the total ICA output (Griffanti et al., 2014).

Griffanti et al.'s (2017) guide was used to hand classify each matched component to ensure they were BOLD-related signal components, and not a component representing structured noise. There are three main characteristics of signal components that allow researchers to differentiate them from sources of noise:

1. The spatial map of the component is localised to the grey matter (Figure 17 top, left)
2. The timeseries of the component does not contain abrupt changes (Figure 17, middle, left)
3. The power spectrum of the component shows power at low frequency (Figure 17, bottom, left)

Figure 17 shows examples of both signal (left panel) and noise (right panel) components. The signal component, likely representing the DMN, satisfies the three classification criteria; the spatial map shows voxels predominantly localised to the grey matter, the timeseries is without abrupt changes, and the power spectrum is largely low frequency. In contrast, the noise component, likely a result of head motion, shows voxels forming a brain edge ring, a sudden change in the time series (indicated by the orange circle

in Figure 17, middle, right), and a high frequency power spectrum. Here, and throughout this thesis, these "hallmarks" of signal components were used to hand classify each component match. If a component failed to exhibit all three characteristics, it was disregarded as a candidate component and removed from further analysis.

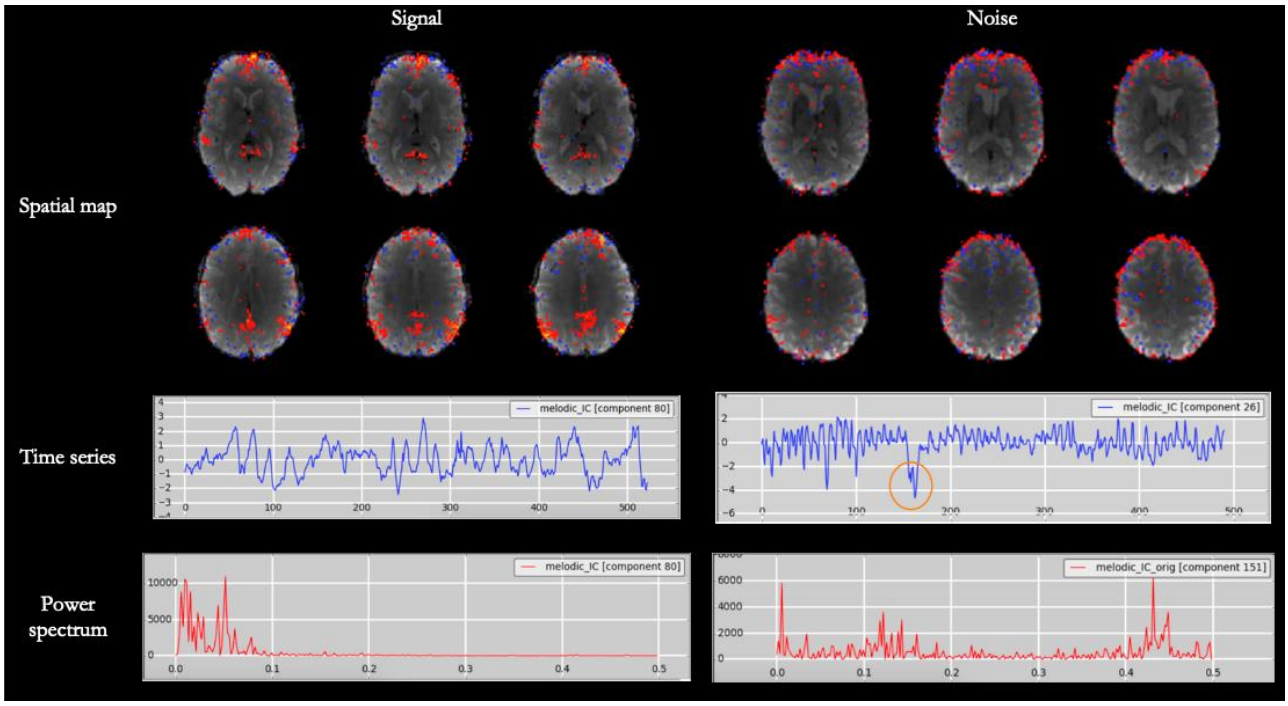

Figure 17. An example of a signal (left panel) and a noise (right) component. The signal component shows voxels localised to the gyre matter, a time series without abrupt changes and a low freeucny power spectrum. In contrast, the noise component, likely a result of head movement, shows voxels forming a ring around the edge of the brain, a time series containing abrupt changes (orange circle), and a high frequency power spectrum.

3.3 Results

3.3.1 Spatial correlations

The following sections explore whether ICA components correspond to discrete visual areas. We find a degree of inconsistency across brains and across visual areas. Broadly speaking, there are five key conclusions we draw from the spatial correlation analyses:

1) Some visual areas are more reliably identified than others.
2) Some brains are more consistent than others.
3) The two hemispheres are not the same.

4) There is not always a 1:1 mapping relationship between components and visual areas, many components match to a single visual area, and a single component can match to many visual areas.
5) The spatial maps of matched components vary across subjects.

The analyses that support these conclusions are detailed below.

3.3.1.1 Above threshold matches – visual area reliability

First, we assessed how often, for each brain, ICA produced a component match for each visual area. This analysis reveals some visual areas consistently had matched components, e.g., ≥ 80% of subjects produce a component match to the area, whilst other visual areas are more inconsistent, e.g., < 80% of subjects produce a component match. There are also some notable differences between hemispheres with respect to how consistently certain areas are identified across subjects (Table 4).

We computed correlations between each of the subject's spatial maps and each atlas region. Details of how we limited the spatial extent of both the ICA components and the probabilistic visual area maps are found in section 2.1.5.3. Smith et al. (2009) used a threshold of $r > 0.25$ as the cut-off for matching resting state components. Here we adopted a slightly stricter matching threshold of $r > 0.3$. The use of this threshold removed matches with low correspondence between the spatial maps and visual areas. For example, a noise component encompassing several visual areas may appear to correlate with each to some extent. Thresholding should reduce the occurrence of such matches. For each visual area, we counted the number of subjects which output at least one match above threshold. This metric gave us a general indication of which areas tend to have consistent component matches across subjects.

Table 4 shows, for each visual area, the number of subjects who produced at least one component match above threshold. Three areas, TO2/MST, IPS2 and SPL1 were consistently matched in all brains, in both hemispheres. Five areas, V1d, V2d, IPS1, IPS3 and IPS4 had matched components in every brain in either the right (V1d) or left (V1d, V2d, IPS1, IPS3 and IPS4) hemisphere but lacked a match for one (V1d, V2d and IPS3) or two (IPS1 and IPS4) brains in the other hemisphere.

The least consistently matched area was V3A in both left and right hemispheres, which was matched in only 12/17 (71%) of brains. V3V also showed a lower number of matched components, 13/17 (77%) of brains for the left hemisphere and 12/17 (71%) of brains for the right. Some areas showed a high number of matches in one hemisphere, but not the other. IPS5 had a matched component in every brain in the right hemisphere, but only 12/17 (71%) of brains in the left hemisphere. VO2 had a matched component in 15/17 (88%) in the right hemisphere but 10/17 (59%) of brains in the left hemisphere.

Table 4. Number of subjects who output at least one component match above the r = 0.3 threshold. Colour-coding reflects the consistency with which each area was identified across brains. Green indicates regions that were identified in ≥ 80% of brains (consistently identified regions). Red indicates regions that were identified in < 80% of brains (inconsistently identified regions). The ratio between the number of subjects who produced a match in the left hemisphere to the right hemisphere are expressed as a decimal (values >1 indicate more subjects produce a match in the left hemisphere compared to the right hemisphere. Values < 1 indicate the reverse).

	Left hemisphere		Right hemisphere		Left to right ratio (decimal form)
Visual Area	**Number of subjects**	**%**	**Number of subjects**	**%**	
V1v	15	88.2	15	88.2	1.00
V1d	16	94.1	17	100	0.94
V2v	14	82.4	15	88.2	0.93
V2d	17	100	16	94.1	1.06
V3v	13	76.5	12	70.6	1.08
V3d	16	94.1	15	88.2	1.07
hV4	14	82.4	16	94.1	0.88
VO1	13	76.5	15	88.2	0.87
VO2	10	58.8	15	88.2	0.67
PHC1	13	76.5	14	82.4	0.93
PHC2	13	76.5	15	88.2	0.87
TO2/MST	17	100	17	100	1.00
TO1/MT	16	94.1	16	94.1	1.00
LO2	14	82.4	13	76.5	1.08
LO1	15	88.2	16	94.1	0.94
V3B	14	82.4	14	82.4	1.00
V3A	12	70.6	12	70.6	1.00
IPS0	15	88.2	16	94.1	0.94
IPS1	17	100	15	88.2	1.13
IPS2	17	100	17	100	1.00
IPS3	17	100	16	94.1	1.06
IPS4	17	100	15	88.2	1.13
IPS5	12	70.6	17	100	0.71
SPL1	17	100	17	100	1.00
FEF	14	82.4	16	94.1	0.88

3.3.1.2 Mean number of matches - many:1 mapping ratio

We have seen that some brains do not show a component match to all regions above the threshold. We have also noted the inconsistencies between hemispheres, in that some areas are more regularly identified in one hemisphere than another (e.g., VO2). In the previous section we examined,

across subjects, the frequency with which each visual area receives at least one above-threshold match, but in many cases, there are multiple component matches to a visual area that are above threshold. Next, we examine how the number of above threshold matches varies across the visual pathway. We find there is not a simple 1:1 mapping between components and the atlas regions. Some areas receive more matches above threshold than others. Again, there are notable differences between the two hemispheres.

For each visual area, we calculated the mean number and range of above threshold ($r > 0.3$) matches (Table 5). The mean and range were calculated across all subjects, including those who did not output an above threshold match[8].

In the left hemisphere, VO1, VO2 and IPS5 have the fewest component matches, whilst IPS3, V2d, IPS1 and IPS2 have the most matches. PHC2 and V1d are the most variable regions with the number of matches produced by a subject ranging from 0 to 5. In the right hemisphere, V3v, LO2 and VO1 have the fewest component matches, whilst IPS3, FEF and V1d have the most component matches. FEF is the most variable region as the number of matches ranged from 0 to 6. Note, an area receiving multiple matches above threshold does not necessarily indicate that each match is a suitable representation for that visual area alone. For example, a component could be a match for multiple visual areas (which I term "component sharing" in the discussion below).

[8] An alternative approach, in which the subjects who did not output a match were excluded from these calculations was also performed. Additionally, for each subject we then calculated the percentage of the total number of matches corresponding to each visual area, rather than the mean across all subjects. This approach did not alter the overall pattern of the analysis. For completeness, the results of the alternative approach are detailed in the Appendix IV.

Table 5. The mean number (across brains) and range of components that correlated with each atlas region above the r = 0.3 threshold. To aid interpretation of differences between hemispheres, colour-coding indicates high (blue) and low (yellow) values.

Visual Area	**Left hemisphere**	Right hemisphere
V1v	1.41 (0-3)	1.59 (0-3)
V1d	1.82 (0-5)	2.24 (1-5)
V2v	1.18 (0-3)	1.24 (0-3)
V2d	2.18 (1-4)	1.88 (0-4)
V3v	0.94 (0-2)	0.82 (0-2)
V3d	1.71 (0-4)	1.29 (0-2)
hV4	1.24 (0-3)	1.71 (0-4)
VO1	0.88 (0-2)	1.12 (0-2)
VO2	0.76 (0-2)	1.47 (0-2)
PHC1	1.29 (0-4)	1.47 (0-3)
PHC2	1.59 (0-5)	1.59 (0-4)
TO2/MST	1.53 (1-2)	1.35 (1-3)
TO1/MT	1.71 (0-3)	1.35 (0-2)
LO2	1.06 (0-2)	0.94 (0-2)
LO1	1.29 (0-2)	1.29 (0-2)
V3B	1.18 (0-3)	1.47 (0-3)
V3A	1.18 (0-3)	1.18 (0-3)
IPS0	1.35 (0-3)	1.29 (0-2)
IPS1	2.12 (1-3)	1.41 (0-2)
IPS2	2.12 (1-4)	2.12 (1-4)
IPS3	2.35 (1-4)	2.41 (0-4)
IPS4	1.82 (1-3)	1.82 (0-4)
IPS5	0.88 (0-2)	2.06 (1-5)
SPL1	2.06 (1-3)	2 (1-3)
FEF	2.06 (0-4)	2.24 (0-6)

3.3.1.3 Top matches

The two analyses above have shown some visual areas are consistently identified using ICA, in that over 80% of subjects produce a component match above threshold. Likewise, some visual areas receive multiple matches above threshold suggesting either the area has sub-areas with different functions, or the area can function as part of a network of components.

Thus far, analyses have focused on the number of matches above the r = 0.3 threshold produced. In the following analysis we explore the nature of these matches to assess the spatial correspondence between a component and the atlas region.

First, we identified and explored the "top matched" component for each visual area and each subject. This component showed the highest correlation with a given visual area compared to all other output components and therefore is not constrained by the threshold we set. In theory, the top matched component should show the strongest anatomical correspondence with the atlas region compared to all other components. There are four key findings from this analysis:

1) Some atlas-to-component matches show stronger correlations than others, indicating those visual areas are more accurately identified by ICA.
2) The spatial maps of top matched components vary across subjects.
3) Components tend to be large, often encompassing adjacent regions, i.e., there is not a 1:1 mapping ratio.
4) On average, correlations are higher for TO1/MT and TO2/MST compared to all other visual areas.

We performed correlations between the atlas regions and the unthresholded (melodic_IC) spatial maps[9]. For each visual area the mean correlation across subjects of the top match was then calculated. These means are displayed in Figure 18.

[9] We also performed correlations between the atlas and the thresholded z-score spatial maps. As expected, higher correlations were found between the atlas regions and the thresholded spatial maps compared to the unthresholded maps. However, the mean correlation across all areas increased by only 0.03 when thresholded spatial maps were used (see Appendix V). For all subsequent analyses the unthresholded spatial maps were used.

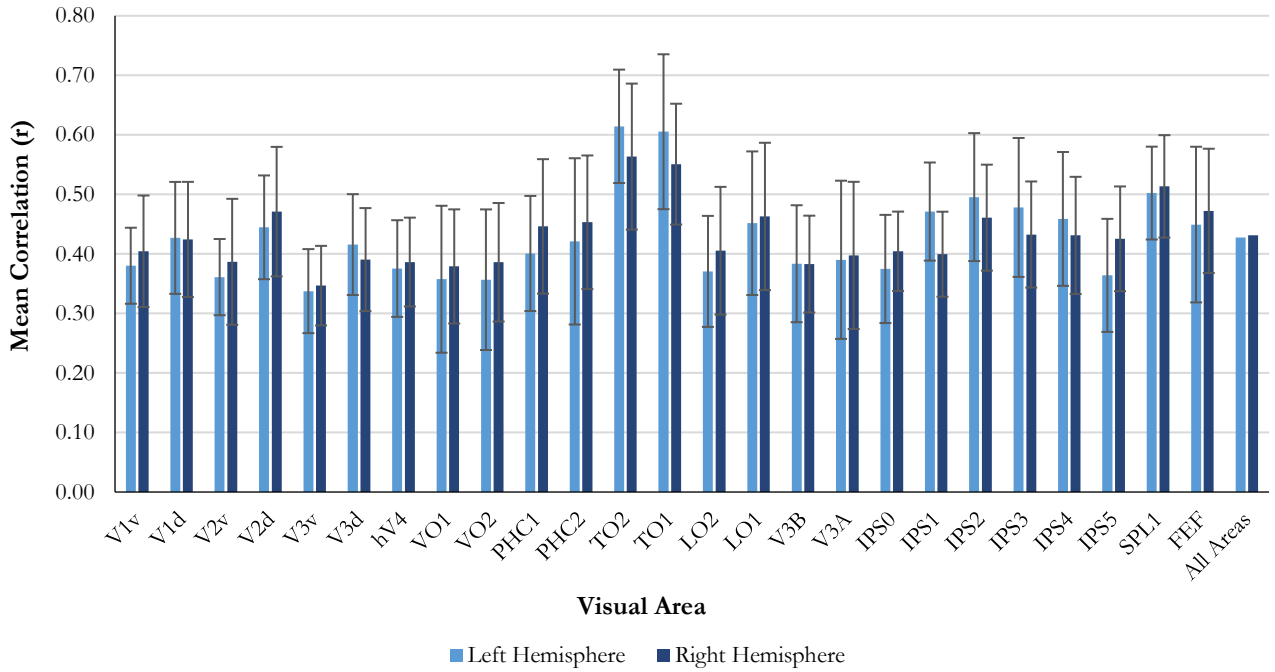

Figure 18. Mean correlation between the "top matched" component and the atlas regions. Error bars represent standard deviation. Fisher's Z equivalent figure can be found in Appendix A.

The mean across all observers for the left hemisphere correlations (r = 0.427) is comparable to the right hemisphere correlations (r = 0.431, Figure 18). TO1/MT and TO2/MST are two regions which stand out due to their high correlation. On average, the top matched component correlated with TO1/MT at r = 0.61 (SD = 0.13) in the left hemisphere and r = 0.55 (SD = 0.10) in the right hemisphere. Likewise, on average the top matched component correlated with TO2/MST at r = 0.61 (SD = 0.10) in the left hemisphere and r = 0.56 (SD = 0.12) in the right. For both regions, the left hemisphere shows a higher correlation than the right. Interestingly, on average early visual areas show lower correlations than TO1/MT and TO2/MST. This appears to be due to the fact there are inconsistencies in the size and coverage of the component maps obtained from subjects.

To illustrate the variability in component matches across brains, the three components with the highest correlations to 20 of the visual areas[10] are shown in Figure 19. The coverage of the spatial maps can be seen to vary across subjects and across visual areas. For some subjects the top matched component is large, with either activation appearing to cover several regions, or focal activation in a single region with residual activation extending into the neighbouring regions. Take V1v for example. Subject 14 produced a component that correlated with the region at $r = 0.51$. The spatial map shows activation encompassing the entire region, which extends dorsally and ventrally into V1d and V2v respectively. For some subjects, the spatial maps are more confined, with very little residual activation extending into neighbouring regions. However, in these instances, the coverage of the visual area tends to be incomplete. For example, Subject 6 produced a component that correlated with V1v at $r = 0.46$. The spatial map of this component shows activation in the anterior region of V1v with very little activation in the posterior of V1v. Unlike Subject 14, the activation does not extend ventrally into V2v. This illustrates the differences between subjects in the coverage of the top matched components.

The characteristic output of large components has a two-fold effect on our spatial correlation analysis. First, large components may drive the correlation with the atlas region down because they encompass vertices outside of the area with which they correlate. Second, large components are likely to match with multiple regions. This lack of uniqueness amongst the top matches of visual areas is discussed further below.

[10] We excluded the following visual areas due to the inherent variability of the atlas regions leading to potentially spurious results from the spatial correlations: IPS3, IPS4, IPS5, SPL1 and FEF. See section 3.3.1.5 for further details as to why these regions were not explored further.

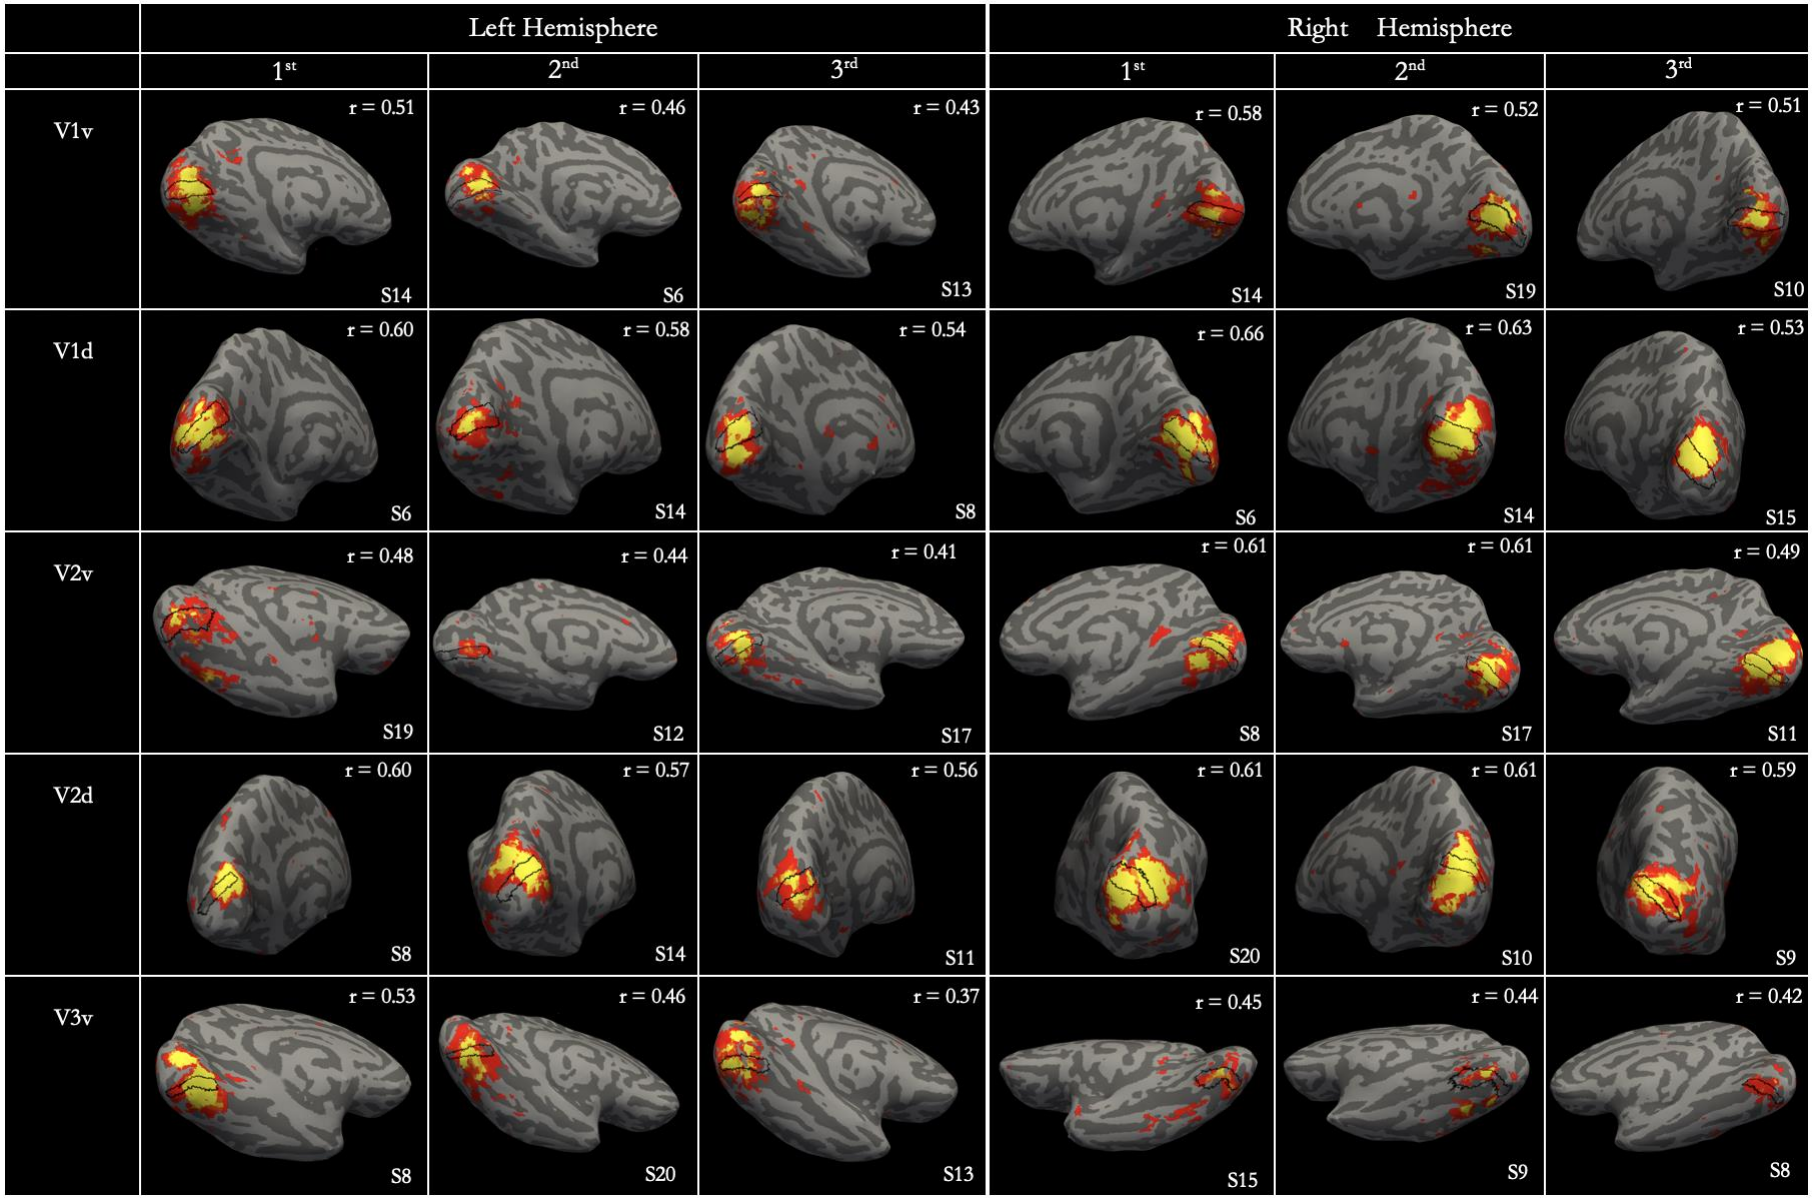
Left Hemisphere
Right Hemisphere
1st
2nd
3rd
1st
2nd
3rd
V1v
r = 0.51
S14
r = 0.46
S6
r = 0.43
S13
r = 0.58
S14
r = 0.52
S19
r = 0.51
S10
V1d
r = 0.60
S6
r = 0.58
S14
r = 0.54
S8
r = 0.66
S6
r = 0.63
S14
r = 0.53
S15
V2v
r = 0.48
S19
r = 0.44
S12
r = 0.41
S17
r = 0.61
S8
r = 0.61
S17
r = 0.49
S11
V2d
r = 0.60
S8
r = 0.57
S14
r = 0.56
S11
r = 0.61
S20
r = 0.61
S10
r = 0.59
S9
V3v
r = 0.53
S8
r = 0.46
S20
r = 0.37
S13
r = 0.45
S15
r = 0.44
S9
r = 0.42
S8

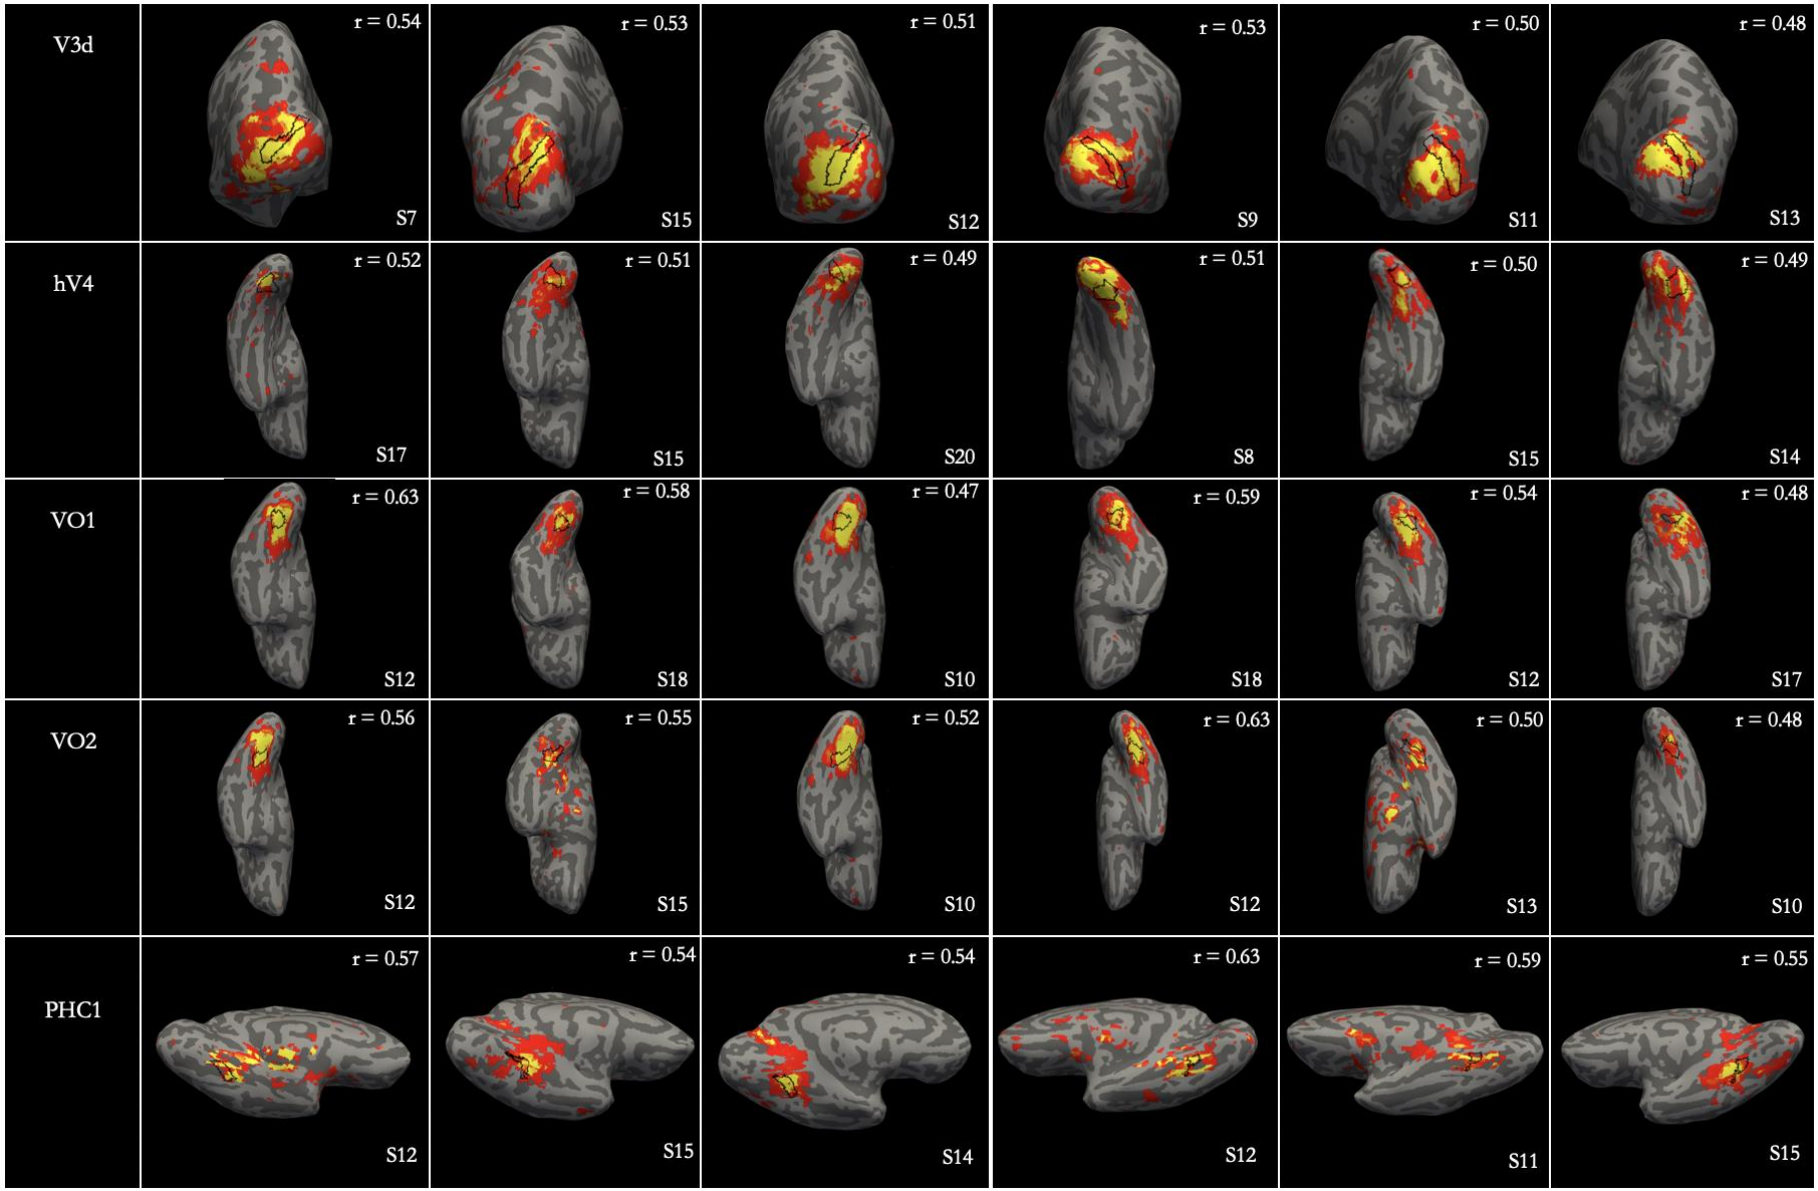
V3d
r = 0.54
S7
r = 0.53
S15
r = 0.51
S12
r = 0.53
S9
r = 0.50
S11
r = 0.48
S13
hV4
r = 0.52
S17
r = 0.51
S15
r = 0.49
S20
r = 0.51
S8
r = 0.50
S15
r = 0.49
S14
VO1
r = 0.63
S12
r = 0.58
S18
r = 0.47
S10
r = 0.59
S18
r = 0.54
S12
r = 0.48
S17
VO2
r = 0.56
S12
r = 0.55
S15
r = 0.52
S10
r = 0.63
S12
r = 0.50
S13
r = 0.48
S10
PHC1
r = 0.57
S12
r = 0.54
S15
r = 0.54
S14
r = 0.63
S12
r = 0.59
S11
r = 0.55
S15

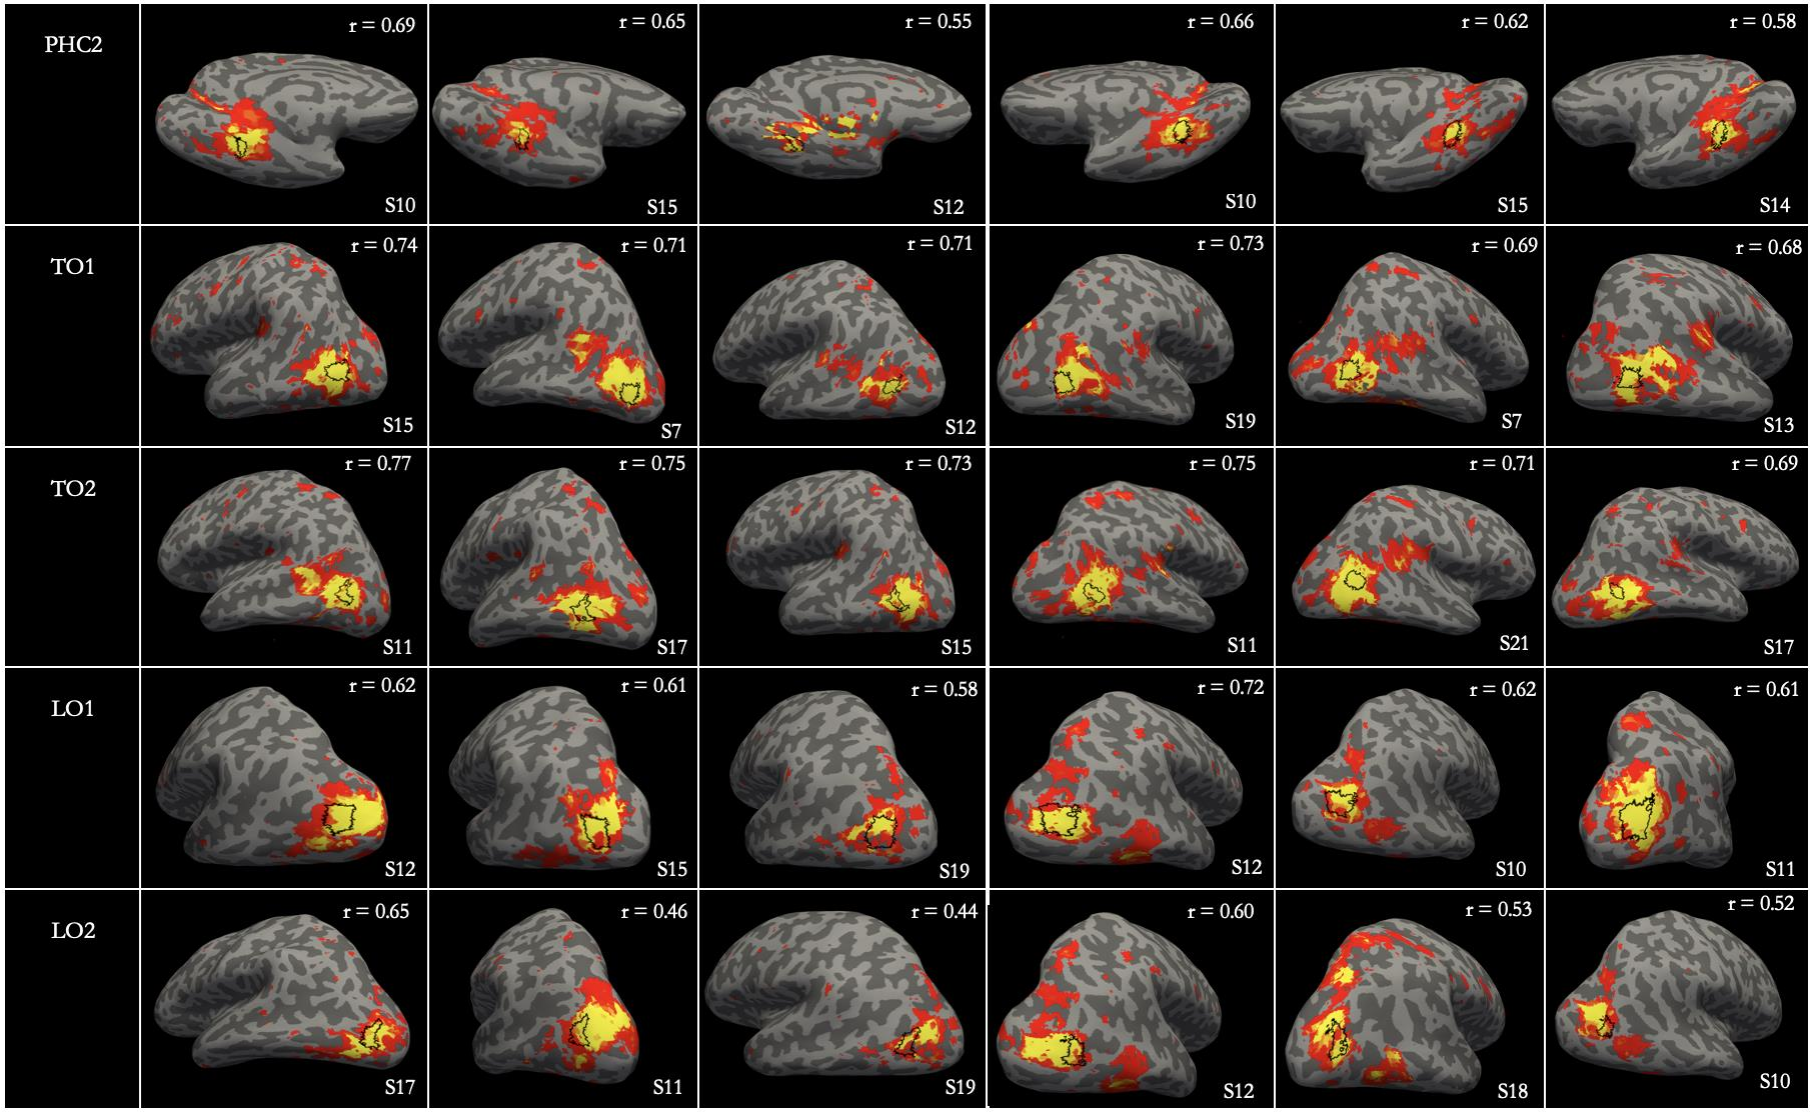
PHC2
r = 0.69
S10
r = 0.65
S15
r = 0.55
S12
r = 0.66
S10
r = 0.62
S15
r = 0.58
S14
TO1
r = 0.74
S15
r = 0.71
S7
r = 0.71
S12
r = 0.73
S19
r = 0.69
S7
r = 0.68
S13
TO2
r = 0.77
S11
r = 0.75
S17
r = 0.73
S15
r = 0.75
S11
r = 0.71
S21
r = 0.69
S17
LO1
r = 0.62
S12
r = 0.61
S15
r = 0.58
S19
r = 0.72
S12
r = 0.62
S10
r = 0.61
S11
LO2
r = 0.65
S17
r = 0.46
S11
r = 0.44
S19
r = 0.60
S12
r = 0.53
S18
r = 0.52
S10

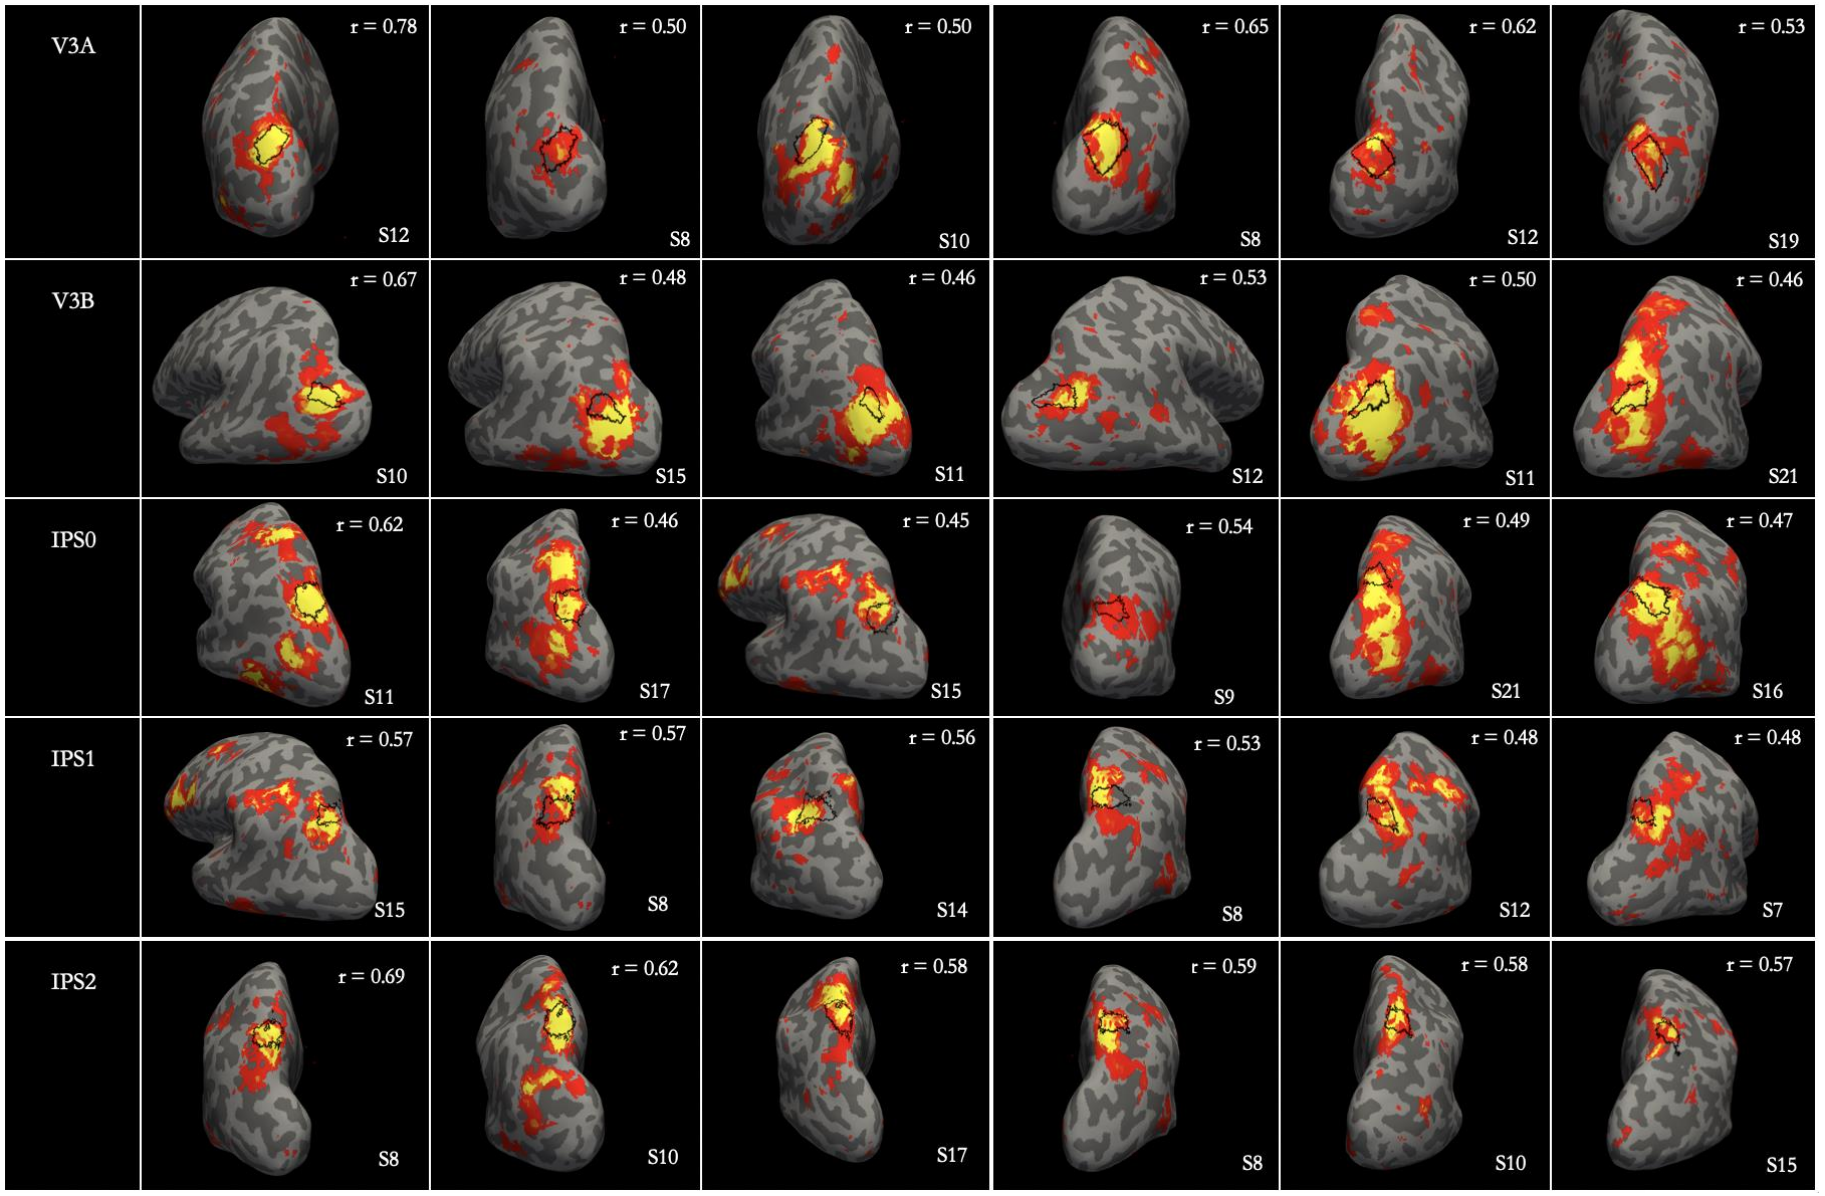

Figure 19. The three components with the highest correlations to 20 of the visual areas (Note, IPS3, IPS4, IPS5, SPL1 and FEF are excluded due to the inherent variability of their probabilistic maps). The subject ID (S) and spatial correlation (r) between the component and the atlas region (indicated by the black borders overlaid on the spatial maps) are also displayed.

3.3.1.4 Uniqueness of components – 1:many mapping ratio

For each subject the number of unique components within the top matches was calculated: A unique component is one which was the top match for only one visual area. The total number of unique matches is consistent across subjects and hemispheres. On average, in each brain, in each hemisphere, 15 components within the top matches are unique, with 10 components matching with more than one visual area (Figure 20). However, we find there is not a simple 1:1 mapping between the visual areas and the top matched components and there are notable differences between the two hemispheres. These conclusions are based upon four key findings:

1) The likelihood of a unique match varies across visual areas and hemispheres.
2) When the top matched component is not limited to a single visual area, it tends to include adjacent regions. We refer to this as "adjacent component sharing".
3) The visual areas included in a 1:many component are not consistent across brains.
4) When a component is unique, it is not always a clean representation of the region alone.

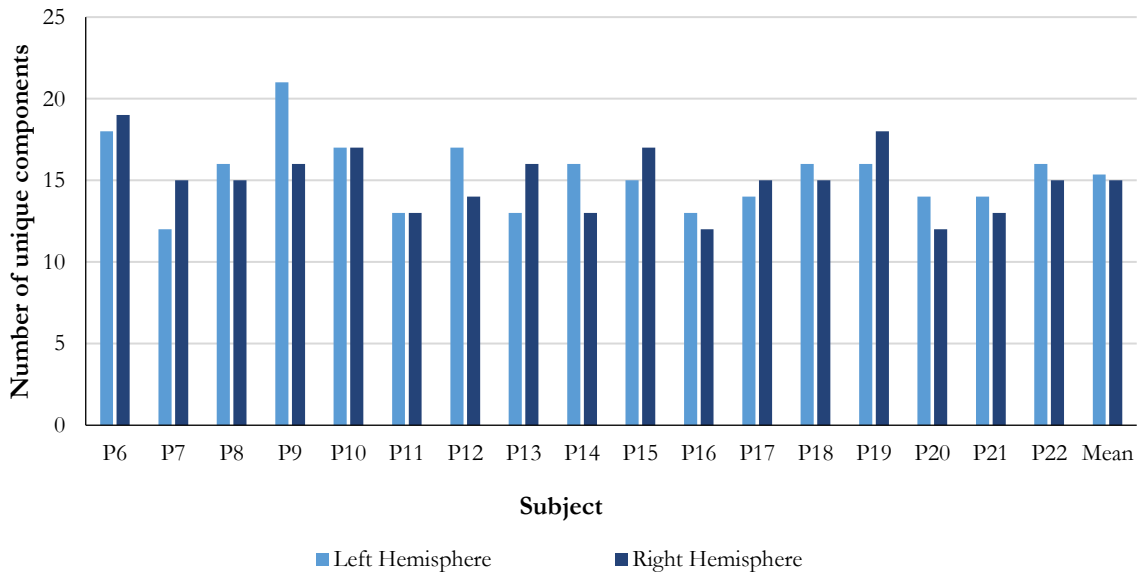

Figure 20. Number of unique components in the top matches for each subjet.

3.3.1.4.1 Uniqueness across visual areas

The likelihood of a component being unique varies across visual areas and hemispheres (see Figure 21). Some areas typically have a unique component, FEF and SPL have the highest proportion of unique matches. Some areas seldom have a unique component, TO1/MT and TO2/MST have the lowest proportion. In the left hemisphere, uniqueness of top matches varies from 94.11% for the FEF to 0% for IPS4 (mean = 35.76%, SD = 23.77). In the right hemisphere, uniqueness of top matches varies from 94.11% for the FEF to 0% for TO1/MT (mean = 31.76%, SD = 21.55).

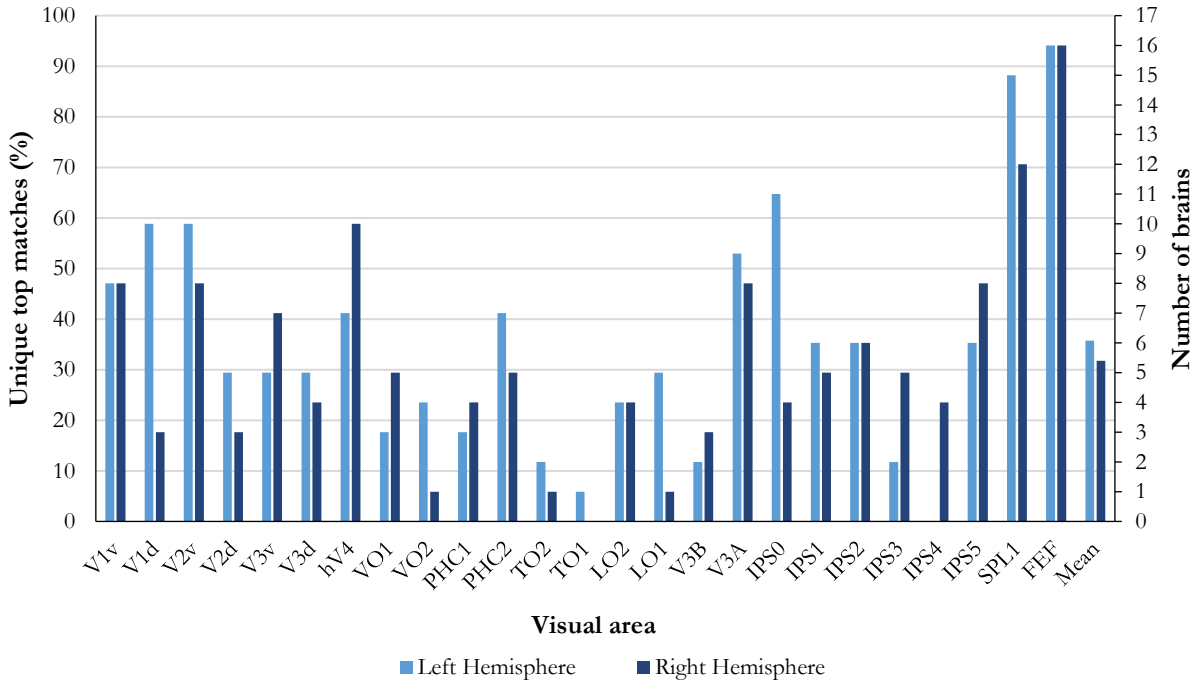

Figure 21. The number of brains that produced a unique match for a given visual area, expressed as the raw number of brains (right axis) and a per cent of the total number of brains (left axis).

The uniqueness of the matched components was consistent across hemispheres for some areas. V1v, LO2, IPS2 and FEF were equally unique in both hemispheres. For all other areas there is an asymmetry, indicating that uniqueness varies across hemispheres. The most marked asymmetries are V1d (10/17, 59% of subjects output a unique match in the left hemisphere, compared to 3/17, 18% in the right hemisphere) and IPS4 (no unique matches in the left hemisphere, compared to 4/17, 24% in the right).

Next, for each hemisphere we investigated the differences in uniqueness across the visual areas and the consistent patterns of component sharing across subjects.

3.3.1.4.2 Adjacent component sharing

When the top matched component is not limited to a single visual area, it tends to include adjacent regions. We refer to this as "adjacent component sharing". There is little consistency across brains, different regions were included in the same component for different subjects (Figure 22), and therefore the visual areas included in a 1:many component are not consistent across brains. There were some notable exceptions which are discussed below.

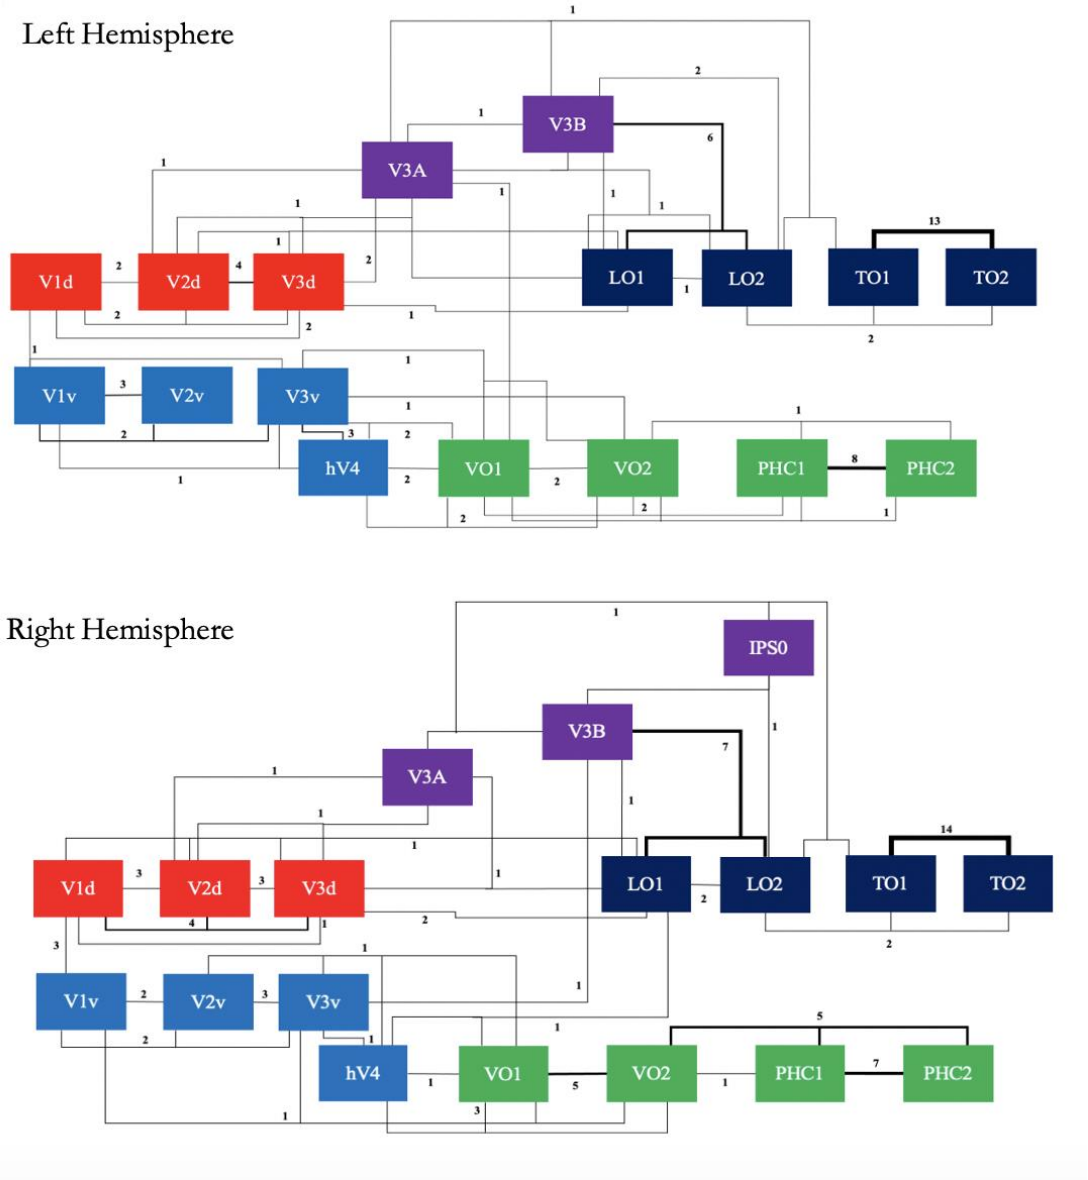

Figure 22. Pattern of component sharing (the instances in which a single component was the top match to more than one visual area) across brains. Lines connecting visual areas indicate that the regions were included in a single component. Thicker lines indicate that more brains showed that pattern of component sharing. Note, the number above each line shows the number of brains that displayed this pattern of component sharing.

There are four main component sharing patterns in the left hemisphere:

1) 13/17 (76%) of subjects produced a component that was the top match to TO1/MT and TO2/MST.
2) 8/17 (47%) of subjects produced a component that was the top match to PHC1 and PHC2.
3) 6/17 (35%) of subjects produced a component that was the top match to V3B, LO1 and LO2.
4) 4/17 (24%) of subjects produced a component that was the top match to V2d and V3d.

Component sharing is more consistent in the right hemisphere than in the left hemisphere in that when a component is not unique, more subjects tend to show the same pattern of component sharing. There are six main component sharing patterns in the right hemisphere, of which three are evident in the left hemisphere (TO1/MST- TO2/MST, PHC1-PHC2, V3B-LO1-LO2).

1) 14/17 (82%) of subjects produced a component that was the top match to TO1/MT and TO2/MST.
2) 7/17 (41%) of subjects produced a component that was the top match to V3B, LO1 and LO2.
3) 7/17 (41%) of subjects produced a component that was the top match to PHC1 and PHC2.
4) 5/17 (29%) of subjects produced a component that was the top match to VO2, PHC1 and PHC2.
5) 5/17 (29%) of subjects produced a component that was the top match to VO1 and VO2.
6) 4/17 (24%) of subjects produced a component that was the top match to V1d, V2d and V3d.

We then explored the spatial maps of components matching to multiple visual areas and found three main characteristics of these components.

1) When multiple regions share a component, there are differences in the coverage of the spatial maps between the subjects.
2) When large, network-like components are output they appear to include multiple visual areas despite not being the top matched component to each individually.
3) There is evidence of a misalignment of the probabilistic atlas for some visual areas.

When multiple regions share a component, there are differences in the coverage of the spatial maps between the subjects. These differences are evident across different visual areas, three examples of which are below.

Subject 14 produced a component that was the top match for V2d (r = 0.57) and V3d (r = 0.44). The spatial map shows activation in the anterior region of V2d and not the posterior (Figure 23, top, left). On the other hand, Subject 20 had similar matches, V2d (r = 0.5) and V3d (r = 0.46), but the spatial map shows activation in the posterior of V2d, and not in the far anterior of the region (Figure 23, top, right).

Likewise, a single component was the top match to V3b (r = 0.67), LO1 (r = 0.45) and LO2 (r = 0.31) for Subject 10 (Figure 23, middle, left). The spatial map shows strong activation centred around V3B with activation in the superior portion of the LO regions. On the other hand, a single component was the top match to V3B (r = 0.48), LO1 (r = 0.61) and LO2 (r = 0.38) for Subject 15 (Figure 23, middle, right). This subject's spatial map shows stronger activation in LO1 compared to V3B and LO2.

Subject 10 produced a component that correlated with V1d (r = 0.48), V2d (r = 0.61) and V3d (r = 0.42). The spatial map shows good coverage across all three regions (Figure 23, bottom, left). Subject 19 produced a component that correlated also with V1d (r = 0.43), V2d (r = 0.36) and V3d (r = 0.27). However, the spatial map shows activation focused on the centre of the areas which reduces towards the posterior (Figure 23, bottom, right).

Therefore, components correlating with multiple visual areas often include different portions of the areas for different subjects.

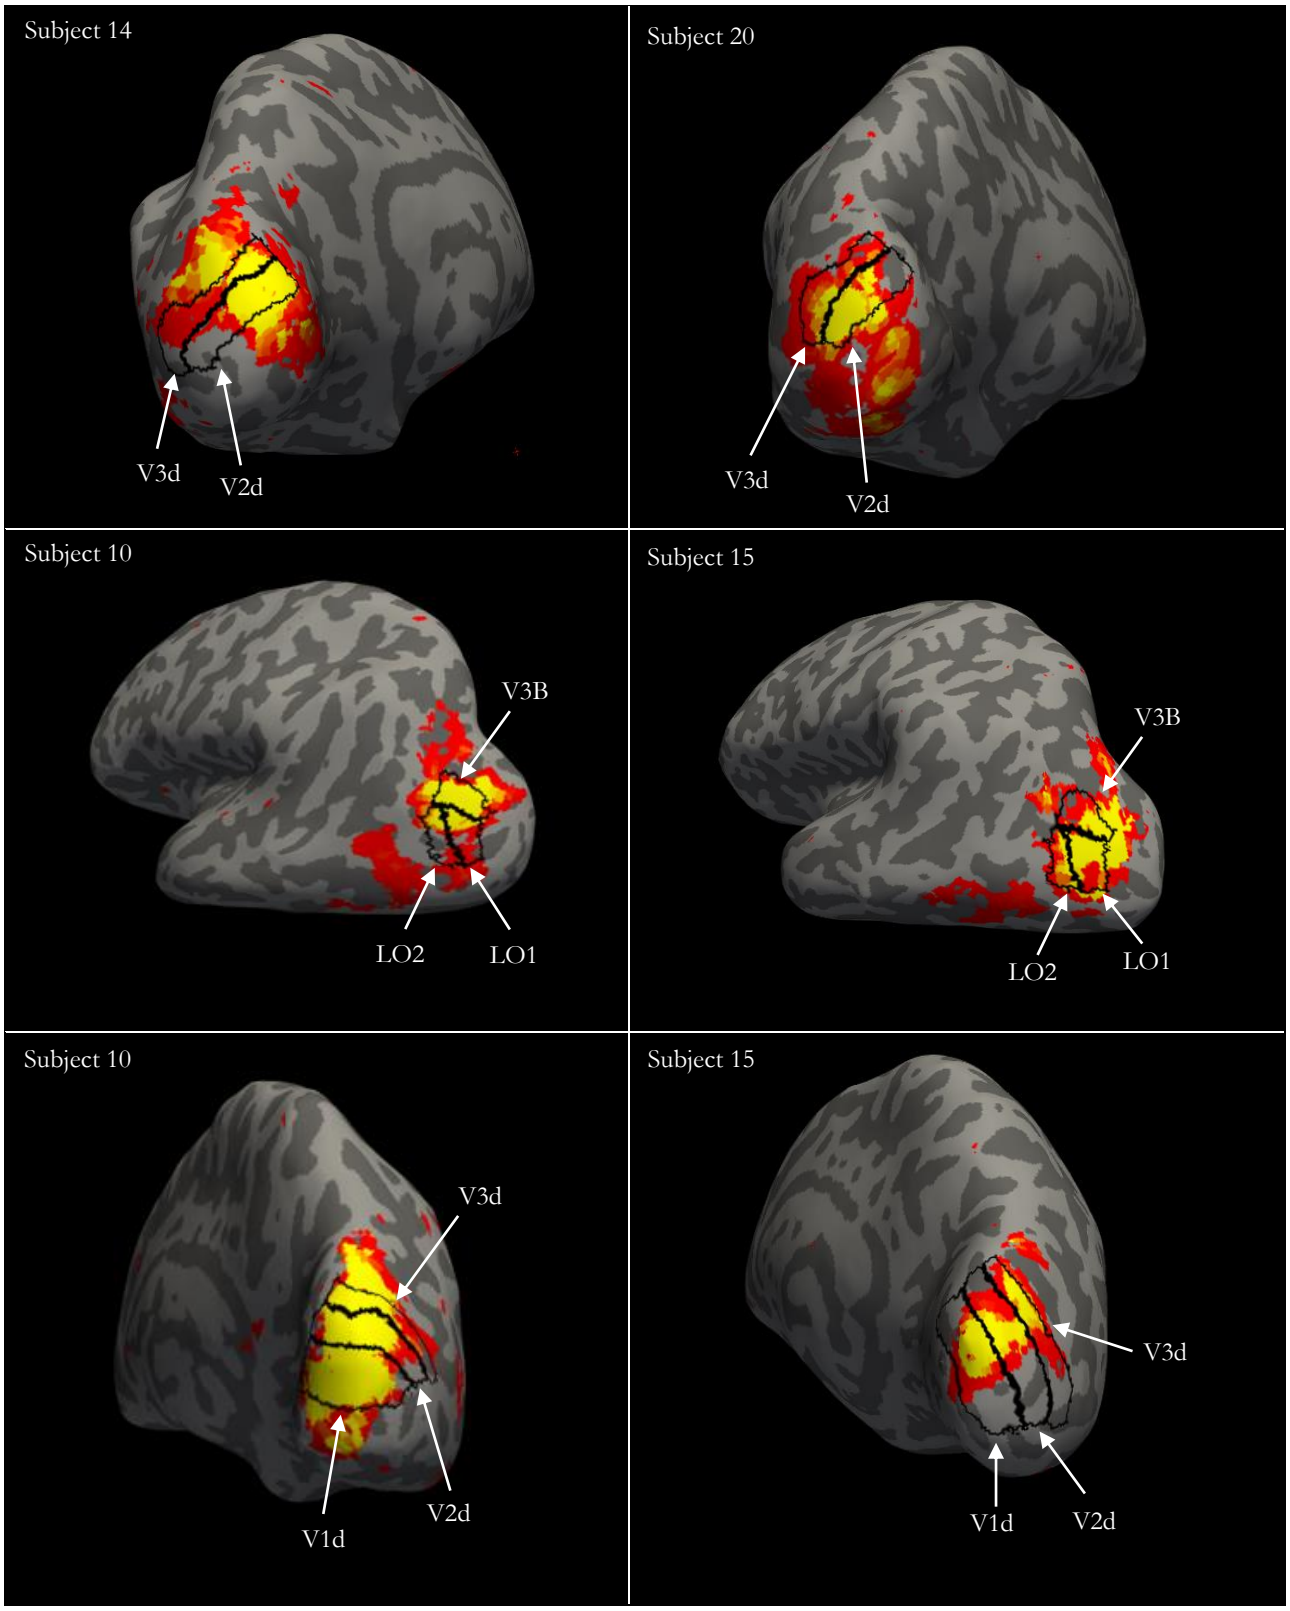

Figure 23. Six example components (from five subjects) showing differing coverage of the visual areas to which they were the top match, highlighting components that match with multiple visual areas often include different portions of the areas for different subjects. Visual area ROIs are indicated by the black outlines.

TO1/MT and TO2/MST do not appear to be segmented into separate components, but rather the whole motion complex ("hMT+") is segmented out. 76% (13/17) and 82% (14/17) of subjects

produced a component that was the top match to both TO1/MT and TO2/MST in the left and right hemisphere respectively. This was the most common component sharing pattern in both hemispheres. TO1/MT in the left hemisphere is represented by a unique component by only one subject, whereas TO1/MT in the right hemisphere is never represented by a unique component. A unique match to TO2/MST in the left hemisphere was produced by two subjects (12%), whereas only one subject produced a unique match TO2/MST in the right hemisphere.

Further evidence of the assertion that there are differences in the coverage of the spatial maps between the subjects comes from the relative correlation with the two atlas regions. Some subjects produce a component that correlates with both atlas regions equally. For example, Subject 15 produced a component that correlated with TO1/MT at $r = 0.74$ and TO2/MST at $r = 0.73$ (Figure 24, left). Other subjects output a component that correlates more highly with one region than the other. For example, Subject 11 produced a component that was the top match to both TO1/MT ($r = 0.43$) and TO2/MST ($r = 0.76$). The spatial map of this component shows activation centred around TO2/MST which extends towards TO1/MT (Figure 24, right). There is also residual activity in the posterior end of the superior temporal sulcus (STS). We also find this pattern of activation in other brains. This is interesting as previous work has suggested there is a "social perception pathway" located on the lateral surface (Pitcher & Ungerleider, 2021). This potential pathway is discussed in Chapter 5.

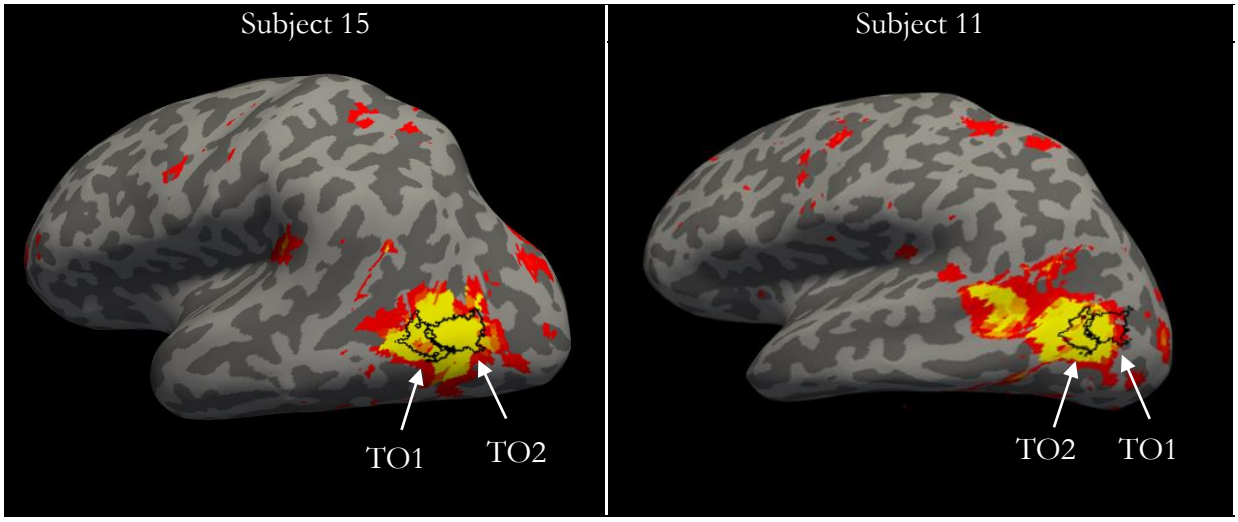

Figure 24. Two example subjects (Subject 15 and 11) who produced a single component that was the top match to both TO1/MT and TO2/MST. For both subjects, the spatial maps show strong activation within the TO1/TO2 atlas ROIs (note, the activation extends beyond the ROI borders). This pattern of component sharing and activation extending beyond the ROI is consistent across brains, indicating the whole motion complex ("hMT+") is segmented out. Visual area ROIs are indicated by the black outlines.

Another key finding from this analysis is that when a component is a unique match to a visual area, it does not always indicate the component is a clean representation of that visual area alone. For example, Subject 10 produced a unique match to TO1/MT ($r = 0.45$) in the left hemisphere (Figure 25). Despite this component being a match to TO1/MT alone, the spatial map shows strong activation in TO2/MST. These "pseudo-unique" matches are consistently found across brains and across visual areas, providing further support for the general conclusion that there is not always a simple 1:1 mapping between the visual areas and components.

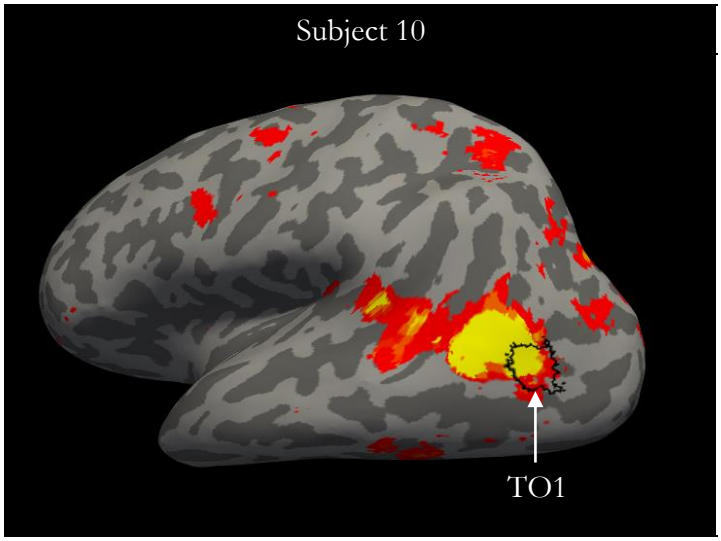

Figure 25. An example subject (Subject 10) who produced a unique match to TO1 in the left hemisphere. The spatial map (displayed on the inflated surface) shows activation extending anteriorly into TO2 and beyond. We refer to these matches as "pseudo-unique" as although they are the top match to only one visual area, activation extends into adjacent regions. Visual area ROI is indicated by the black outline.

3.3.1.4.3 Network or activation bleeding?

So far, we have seen that some subjects produce large, network-like components. These components appear to include multiple visual areas despite not being the top matched component to each individually. Instead, we see separate components produced for some of the sub-regions that correlate greater with the atlas region. However, "activation bleeding" needs to be considered when interpreting the spatial maps and judging whether an area is included in both a network component and a single, discrete component.

Activation bleeding refers to activation spreading from one region to another. This can be an inherent issue arising from the low spatial resolution of the EPI data or can be the result of volumetric smoothing. To investigate this potential issue, we also examine the components on the non-inflated midthickness surface[11]. Below I present an example of how the midthickness surface is used when inspecting the spatial maps of network-like components.

Figure 26 presents an example of a subject who produced a network-like component that was the top match to both V1v ($r = 0.57$) and V2v ($r = 0.44$). The component also has activation around the anterior portion of V1d despite it not being the top match to the region (Figure 26, top, left). The midthickness surface can be used to judge whether the network component is a true representation of all three regions, or whether it shows the hallmarks of activation bleeding. From inspecting the spatial map on the midthickness surface, it is clear that the activation extending from V1v into V1d is not a result of activation bleeding, suggesting the regions may be functioning as part of a network. Moreover, it supports our conclusion that a top matched component does not necessarily always indicate a clean representation of the area(s) alone.

The midthickness surface can also be used to inspect components that appear to provide evidence of segmentation within visual areas. We note above that the anterior portion of V1d is represented in the network component, but that component is not the top match to V1d. The subject

[11] The midthickness surface is generated by averaging the white and pial surface coordinates and allows us to see where the activation is located relative to the sulci and gyri.

instead produced a unique component that correlated with V1d at r = 0.58 (Figure 26, top, right). The spatial map shows activation mainly in the posterior portion of the region. Inspecting the component on the midthickness surface shows the two components cover opposing ends of the cortical fold (Figure 26, green circles), suggesting that different portions of a visual area can either function as part of a network or function independently.

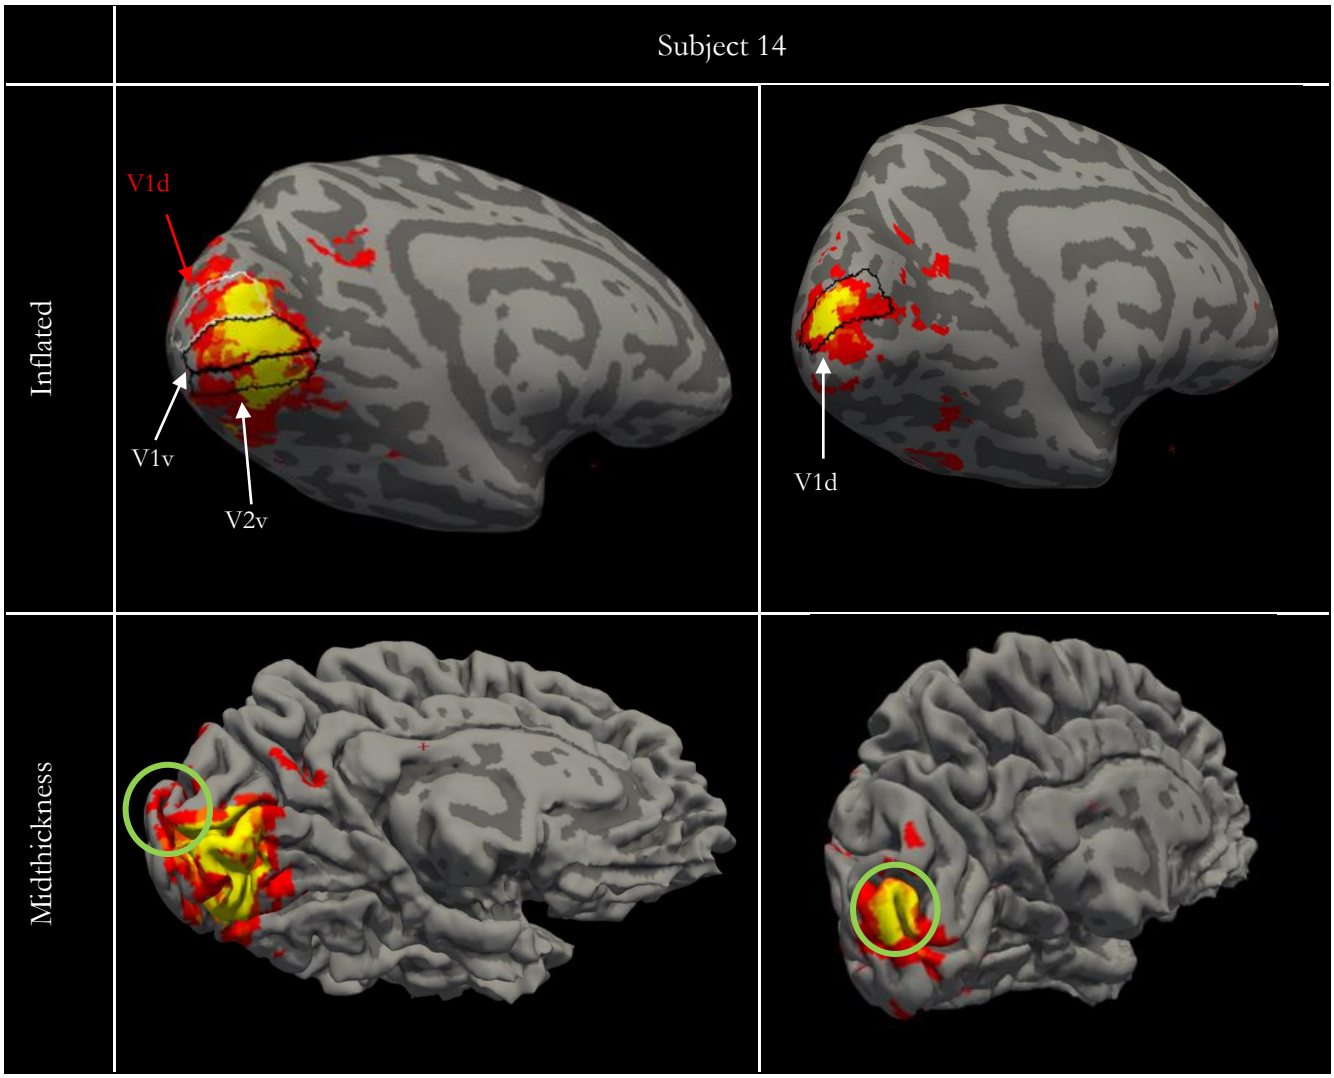

Figure 26. An example subject (Subject 14) who produced a network-like component that was the top match to V1v and V2v in the left hemispheres. When viewed on the inflated surface (top, left), activation is also present in the anterior regions of V1d despite it not being the top match to this region. Displaying the component on the midthickness surface shows the activation in V1d is not a result of activation bleeding (bottom, left). The subject also produced a separate component that was the top match to V1d only (right). Comparing the spatial maps of the two components on the midthickness surface shows that activation in V1d covers opposing ends of the cortical fold in each component (green circles), suggesting that different portions of a visual area can either function as part of a network or function independently.

3.3.1.4.4 Misalignment of the atlas

A key finding from our analysis is that the spatial maps of matched components vary across subjects. This could be indicative of individual differences in the function of certain visual areas. It

could, however, also be due to individual differences in anatomy leading to the potential misalignment of the probabilistic atlas for some visual areas in some subjects.

The probabilistic atlas cannot account for all the anatomical variation of certain visual areas and therefore could be slightly misplaced in some brains. Given the probabilistic atlas is susceptible to misalignment due to this anatomical variability, the components derived from ICA could allow for more accurate identification of some visual areas.

Evidence of this potential atlas misalignment can be seen in TO1/TO2. For example, Subject 16 produced a unique top match to TO2/MST in the left hemisphere (r = 0.66). The spatial map shows strong activation in TO2/MST that extends anteriorly along the lateral surface towards the STS (Figure 27, left). Despite this component being a unique match to TO2/MST, the spatial map also shows residual activation in the anterior of TO1/MT with no activation in the posterior. The atlas suggests that relative to the component the hMT+ complex is located more posteriorly and therefore the component is not a good match to TO1/MT. The top match component to TO1/MT (r = 0.26) was also the top match to LO2 (r = 0.36), V3A (r = 0.31) and V3B (r = 0.36). The spatial map shows widespread activation that encompasses only the posterior of TO1/MT. The two components cover opposing halves of TO1/MT, rather than a complete representation of the region. However, as above, the component uniquely matching TO2/MST (Figure 27, left) could itself be a more accurate representation of hMT+ than the atlas region.

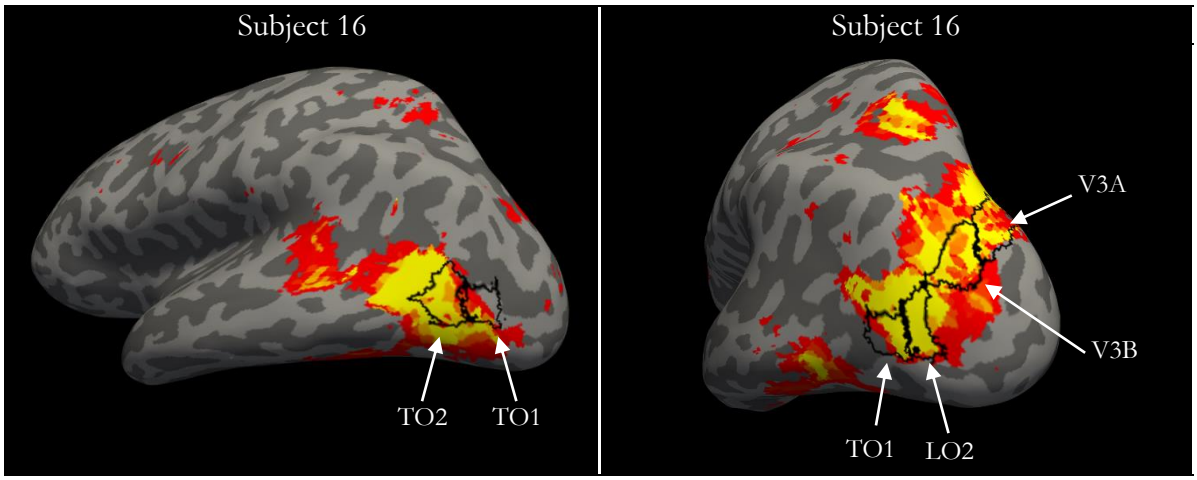

Figure 27. An example subject (Subject 16) showing potential misalignment of the TO1/TO2 atlas ROIs. The subject produced a unique match to TO2/MST (left). The spatial map shows strong activation in TO2/MST which extends anteriorly along the lateral surface; only residual activation is present in the anterior of TO1/MT. The top matched

component to TO1/MT was also the top match to LO2, V3B and V3A (right). The spatial maps suggest hMT+ may be located more anteriorly than the atlas indicates. Therefore, ICA may be a more accurate way of identifying the complex.

Additional evidence of this misalignment can be seen in PHC regions. For example, Subject 20 produced a component that was the top match to VO2 (r = 0.35), PHC1 (r = 0.38) and PHC2 (r = 0.39). The correlation values suggest the component is a moderate and equal match to all three regions. However, the spatial map suggests this might not be accurate. The component shows weak activation inside the atlas regions, with stronger activation displaced laterally on the ventral surface (Figure 28, left). The atlas suggests that relative to the component, VO2/PHC1/PHC2 are located more medially. Similarly, Subject 14 produced a component that was the top match to VO2 (r = 0.31), PHC1 (r = 0.54) and PHC2 (r = 0.58). Activation is focused around PHC2, reflected by a higher correlation to this atlas region compared to PHC1 and VO2, with activation present in the parieto-occipital sulcus (Figure 28, right). Again, it appears the regions are located more posteriorly relative to the component. It could, however, be that these components are an accurate representation of the regions and in fact the atlas is slightly misplaced.

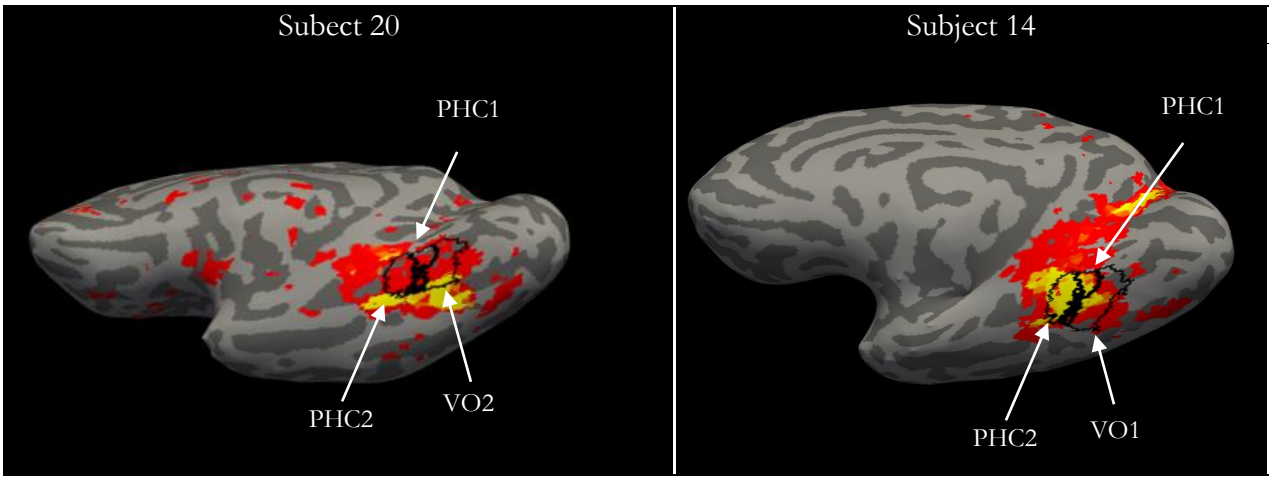

Figure 28. Two example subjects (Subject 20 and 14) showing potential misalignment of the PHC atlas ROIs. Both components show stronger activation outside of the ROIs than inside, indicating the component may be a more accurate representation of the areas than the atlas.

3.3.1.5 Problematic regions

3.3.1.5.1 Frontal eye fields (FEF)

The FEF region is almost always represented by a unique component. This is not surprising as the FEF is located more anteriorly than any other regions and has no neighbouring regions to share a

component with. However, this makes the region more susceptible to an inaccurate match. If a component is located in the anterior region of the cortex, it is likely to correlate with the FEF only. Furthermore, since spatial maps are masked prior to the correlations being computed, some of the activation outside the FEF is removed (Figure 29). This masking of components may result in a misleading matching of components by only leaving activation in FEF. For example, Figure 29 shows a component that correlated with the FEF at r = 0.75. The black outline represents the visual area mask used prior to correlations being computed. It is clear that the component is large and covers more of the cortex than just the FEF. However, when the mask is applied the component is a much cleaner representation. Therefore, we must interpret any spatial correlations with regards to FEF with caution.

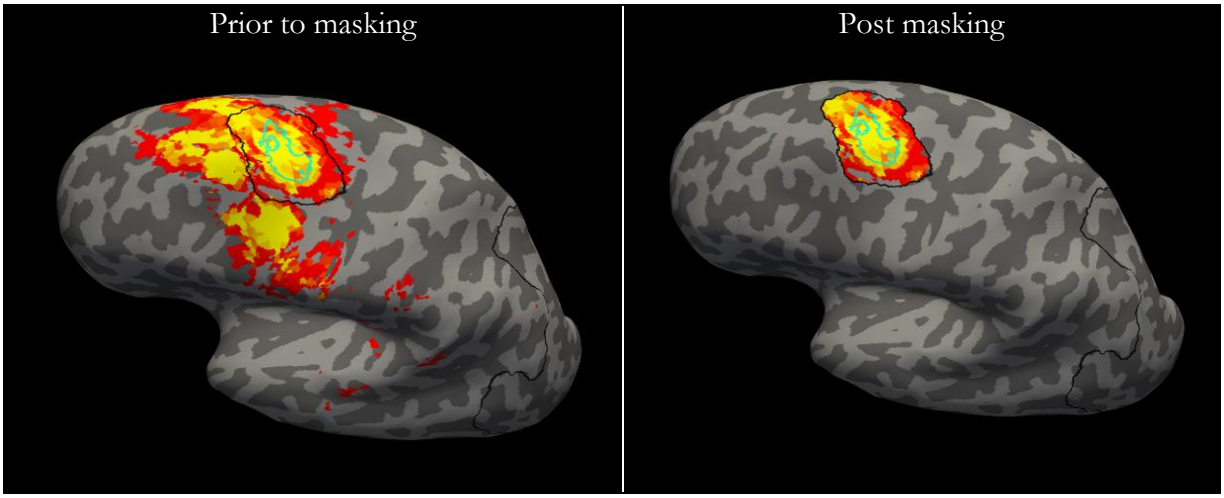

Figure 29. An example component, displayed on the inflated surface, matching with FEF before (left) and after (right) the visual area mask is applied. The visual area mask is indicated by the black outline, the FEF ROI is indicated by the cyan outline. Prior to masking, the spatial map shows widespread activation in the frontal cortex, beyond the visual area mask. When the visual area mask is applied, a substantial amount of this activation is removed, leading to the component appearing to be a good match to FEF.

3.3.1.5.2 IPS regions

The probabilistic maps of IPS3-IPS5 and SPL1 are relatively poor representations of the regions. The authors of the atlas note that this was most likely due to the high anatomical variance across subjects in these regions (Wang et al., 2015). Therefore, we did not explore component matches in the IPS3-IPS5 or SPL1 regions. We did, however, explore component matches to IPS0, IPS1 and IPS2. Overall, we find components matches to these regions are much larger than other regions, often including non-adjacent areas. The spatial maps of the components with the highest correlations with IPS0-IPS2 show widespread activation along the intraparietal sulcus which extends into the superior

and inferior parietal lobules, as well as activation in the frontal lobes for some subjects. Furthermore, the component sharing patterns observed for these regions are far more inconsistent than the other visual areas (see Appendix II) Therefore, we conclude that ICA is not a suitable tool for identifying IPS regions or SPL alone.

3.3.1.6 Bilaterality of components

As the uniqueness analysis is performed on individual hemispheres, we do not get a measure of how consistent component matches are across hemispheres. Next, we investigated how often the corresponding visual areas in the left and right hemispheres were matched to the same components, i.e., how often they were bilateral. If the top matched component to a region is consistently unilateral, it might indicate there are functional differences between hemispheres. We find that the frequency at which top matches are bilateral varies across visual areas and across brains. This is consistent with the differences between hemispheres and across subjects reported in the analyses above. In some instances, genuine unilateral or bilateral components are produced, but this is not always the case. In addition, we report three key findings from this bilaterality measure:

1) When a bilateral component is unique to a region in one hemisphere, it is not always unique to that region in the other hemisphere.
2) When a component is both bilateral *and* a unique match to the visual area, the spatial maps vary across hemispheres.
3) A unilateral top match does not always indicate a unilateral component.

Table 6 displays, for each visual area and each hemisphere, the mean correlation of the top matched components (section 3.3.1.3), the percentage of unique matches (section 3.3.1.4) and the percentage of bilateral matches. Assessing these three metrics alongside Figure 19 offers a more holistic overview of whether performing ICA on movie watching data is an effective tool for identifying visual areas. There appears to be an independent relationship between the three metrics in that an increase in one does not always predict an increase in another. It is, however, clear that there are differences both between hemispheres and across visual areas.

Table 6. Summary of spatial correlations. For each visual area and each hemisphere, I calculated the mean correlation across brains (Pearson's r) between the top matched component and atlas region (mean correlation). The per cent of unique matches (how often the component was the top match to only one visual area) is also presented for each visual area and each hemisphere. Top match bilateral (percent of all brains) reflects the frequency with which the top match in one hemisphere was also the top match to the same region in the opposite hemisphere. Blue cells indicate a higher value, e.g., higher correlation between the component and atlas, higher percentage of unique matches across subjects and higher percentage of bilateral matches across subjects. Yellow cells indicate the reverse. Fisher's Z equivalent table can be found in Appendix B.

	Left Hemisphere		Right Hemisphere		Top match bilateral (%)
	Top match unique (%)	Mean Correlation (r)	Top match unique (%)	Mean Correlation (r)	
V1v	47	0.38	47	0.404	71
V1d	59	0.427	18	0.424	47
V2v	59	0.361	47	0.387	35
V2d	29	0.445	18	0.471	65
V3v	29	0.337	41	0.347	24
V3d	29	0.416	24	0.39	35
hV4	41	0.375	59	0.386	35
VO1	18	0.357	29	0.379	53
VO2	24	0.357	6	0.386	53
PHC1	18	0.401	24	0.446	35
PHC2	41	0.421	29	0.453	53
TO2	12	0.614	6	0.563	76
TO1	6	0.605	0	0.551	71
LO2	24	0.371	24	0.405	65
LO1	29	0.452	6	0.463	53
V3B	12	0.384	18	0.383	71
V3A	53	0.39	47	0.397	18
IPS0	65	0.375	24	0.404	41
IPS1	35	0.471	29	0.399	53
IPS2	35	0.495	35	0.461	41
IPS3	12	0.478	29	0.433	24
IPS4	0	0.459	24	0.431	29
IPS5	35	0.364	47	0.425	24
SPL1	88	0.502	71	0.513	47
FEF	94	0.449	94	0.472	35

The frequency at which top matches are bilateral/unilateral varies across visual areas. Over 70% of matches to V3v and V3A were unilateral, some of which were genuine unilateral components. For example, Subject 20 produced a component that was the top match to V3v (r = 0.46), hV4

(r = 0.49) and VO1 (r = 0.36) in the left hemisphere. The spatial map shows this to be a clear unilateral component that encompasses the three regions in the left hemisphere, with no activation in the right hemisphere (Figure 30A). A separate component was produced that was the top match to V3v (r = 0.40) and hV4 (r = 0.31) in the right hemisphere (Figure 30B). Again, this appears to be a clear unilateral component. Taken together, it could be suggested this subject produced a unilateral ventral visual network component for each hemisphere.

However, unilateral components are not always produced for each hemisphere. We find that despite the analysis indicating the top match was unilateral, the components for some subjects show activation in the opposite hemisphere, sometimes outside of the region with which the component was the top match. For example, Subject 6 produced a component that was a unique match to V3A in the left hemisphere (r = 0.47). The spatial map reveals good correspondence between the component and the V3A atlas region with activation focused on the centre of the region but extending towards the medial surface (Figure 30C). The component shows no activation in the right hemisphere, suggesting a unilateral and unique left hemisphere V3A component was produced. A separate unique match to V3A in the right hemisphere (r = 0.51) was also produced. However, the spatial map shows activation within V3A that extends into the anterior portions of V2d and V3d (Figure 30D). Moreover, the component also showed activation in the opposite hemisphere outside of the V3A atlas region, suggesting the right hemisphere top match was not a unilateral component. These examples highlight the differences between brains and visual areas, whilst also underlining our assertion that a unilateral top match does not always indicate a unilateral component.

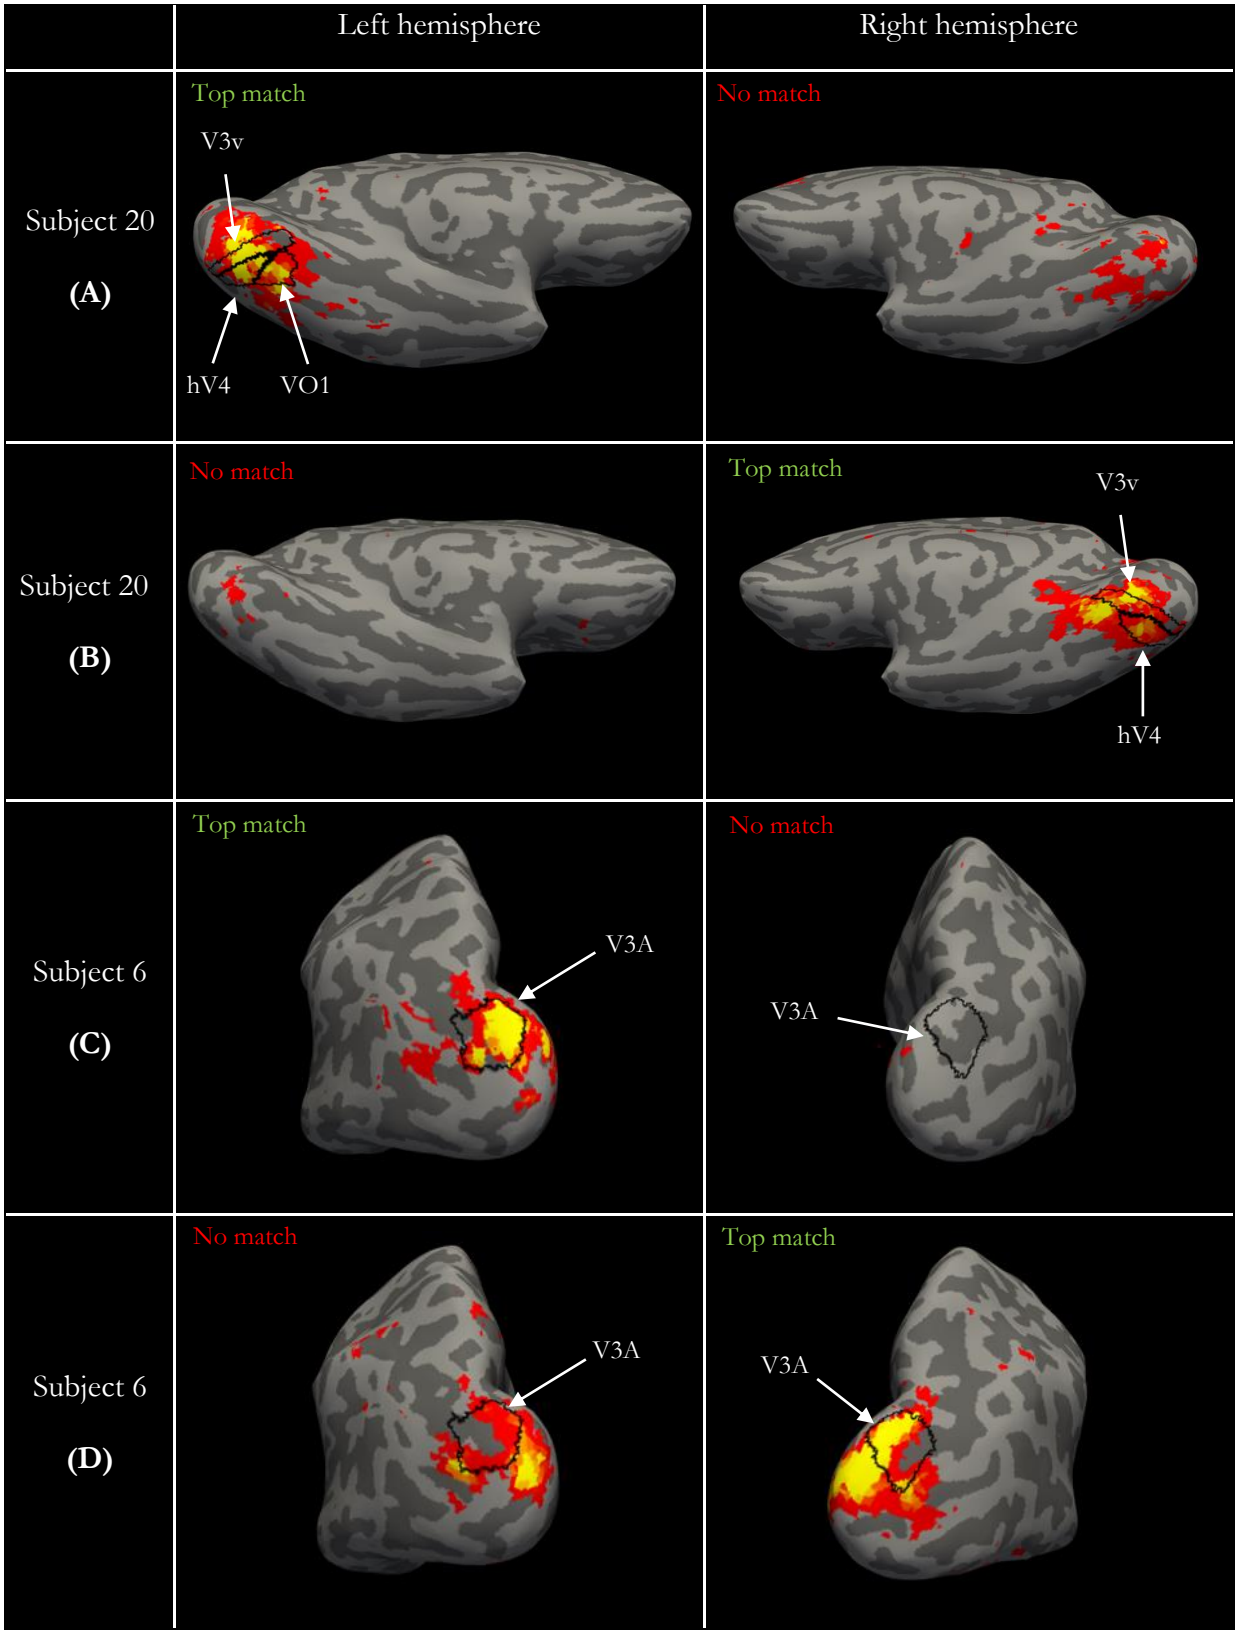

Figure 30. Four example spatial maps of unilateral top matched components from two subjects (Subject 20 and 6). (A) and (B) show a unilateral component was produced for each hemisphere, with little-to-no activation in the opposite hemisphere. (C) shows the top match to left hemisphere V3A is a unilateral component with no activation in the right hemisphere. (D) shows the top match to V3A in the right hemisphere, but the component also shows activation in the left hemisphere, highlighting that a unilateral top match does not always indicate a unilateral component.

A component can be classified as unilateral, but activation can still be present in the regions in both hemispheres. For example, Subject 20 produced a component that was a unilateral and unique match to V1v in the left hemisphere (r = 0.39. Figure 31, top, left). Despite the component only being a top match to left hemisphere V1v, the spatial map of the component shows strong activation throughout the posterior of V1v which extends both dorsally and ventrally in the right hemisphere (Figure 31, top, right). The subject produced a separate component that was the top match to V1v (r = 0.42) and V2v (r = 0.38) in the right hemisphere (Figure 31, bottom, right). However, again, the spatial map of the component shows activation straddling the centre of the two regions in the left hemisphere, albeit at a lower magnitude compared to the right (Figure 31 bottom, left). Therefore, a unilateral top match does not always indicate a unilateral component.

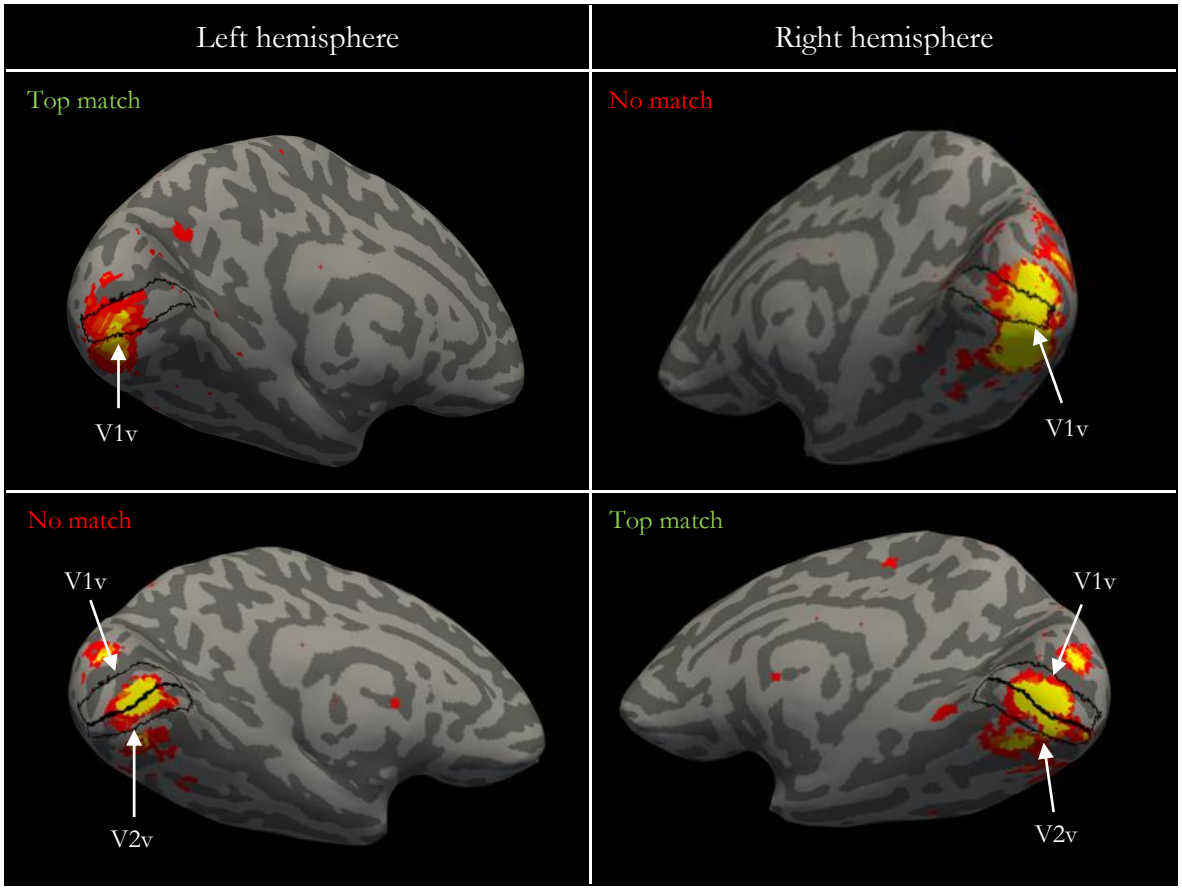

Figure 31. An example subject showing that components can be classified as unilateral, but activation can still be present in the regions in both hemispheres. The subject produced a unilateral top match to V1v in the left hemisphere (top, left) and a unilateral top match to V1v (shared with V2v) in the right hemisphere (bottom, right). The spatial maps of both top matched components show activation in the opposite hemisphere, indicating unilateral top matches can be bilateral components.

Although over 70% of matches to V3A were unilateral, some subjects produce genuine bilateral V3A components. For example, Subject 8 produced a component that was the top match to V3A in the left (r = 0.51) and right hemisphere (r = 0.65) only, suggesting a unique and bilateral component was output. As the correlation values suggest, the spatial map of the component shows strong activation in right hemisphere V3A and activation at a weaker magnitude covering V3A in the left hemisphere (Figure 32). This suggests that visual areas with a high proportion of unilateral matches can still be represented by bilateral components for some subjects. Again, this highlights the inconsistency across brains.

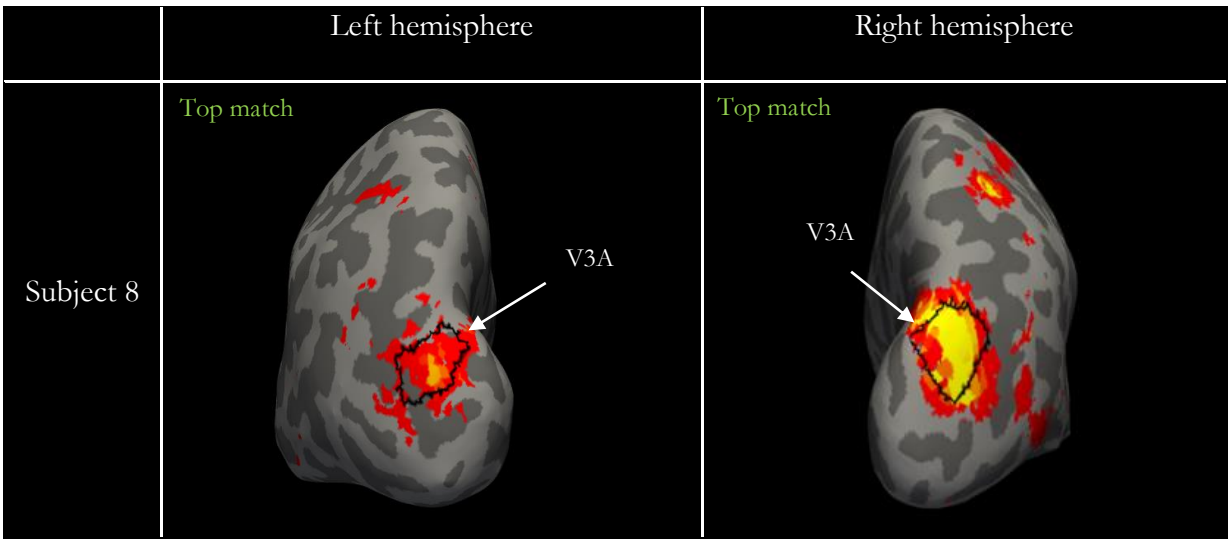

Figure 32. An example of a genuine bilateral V3A component produced by Subject 8. The spatial map shows stronger activation in the right hemisphere (right) compared to the left hemisphere (left).

We find evidence of visual areas functioning bilaterally in one component and the same region functioning unilaterally in another. For example, Subject 12 produced a unique component match to V3A in the left hemisphere (r = 0.78), yet the spatial maps reveal coherent activation around V3A in both hemispheres (Figure 33, top). This suggests the component is in fact a bilateral representation of V3A, despite the right hemisphere component not being the top match to the region. A separate component was the top match to V3A in the right hemisphere (r = 0.62. Figure 33, bottom). The spatial map of this component shows activation in the right hemisphere only, suggesting a genuine unilateral component was produced. Here we find a unilateral match, but a bilateral component and what appears to be a genuine unilateral component match. This suggests V3A sometimes functions bilaterally, and other times functions unilaterally.

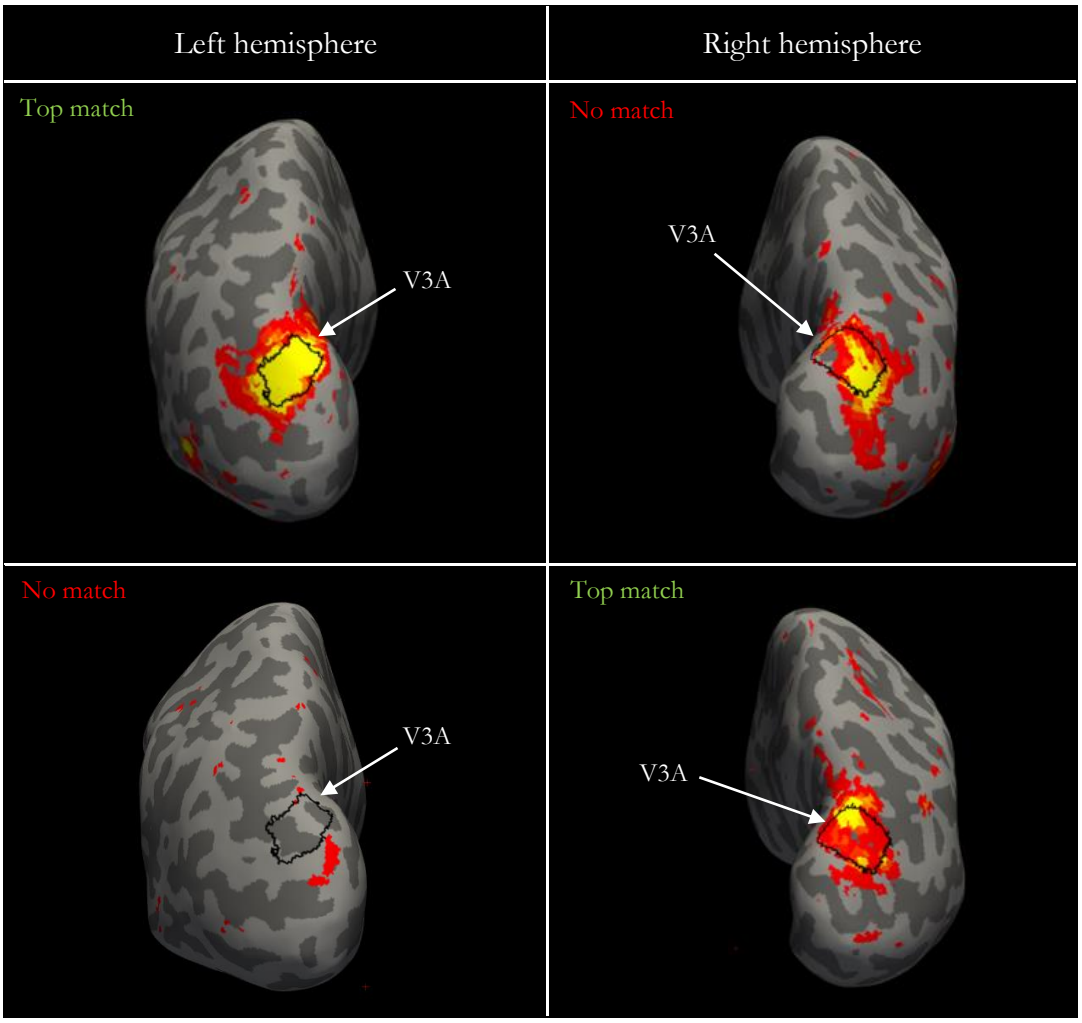

Figure 33. Top: an example of a unilateral top match to V3A that shows activation in both hemispheres, indicating the component is bilateral. Bottom: a genuine unilateral component match to right hemisphere V3A. The component does not show activation in the opposite hemisphere.

When a component is both bilateral *and* a unique match to the visual area, the spatial maps vary across hemispheres. For example, Subject 10 produced a component that was a unique match to V1v in the left ($r = 0.39$) and right ($r = 0.51$) hemisphere. The spatial map of the left hemisphere component shows strong activation throughout V1v which extends both dorsally and ventrally into adjacent regions (Figure 34, top, left). On the other hand, the right hemisphere spatial map shows activation focused on the centre of V1v with only residual activity in adjacent regions (Figure 34, top, right).

Subject 6 produced a component that was a unique match to V1v in the left ($r = 0.46$) and right ($r = 0.33$). The spatial map of the component shows strong activation in the anterior region of V1v with residual activity extending dorsally (Figure 34, bottom, left). The spatial map of the right

hemisphere component again shows activity located in the anterior of the region but at a much weaker magnitude than that seen in the left hemisphere (Figure 34, bottom, right). This suggest that multiple visual areas may function as part of a network in one hemisphere, while individual areas function independently in the opposite hemisphere. However, there is no consistent pattern as to which regions function in this manner or which hemisphere exhibits the larger component.

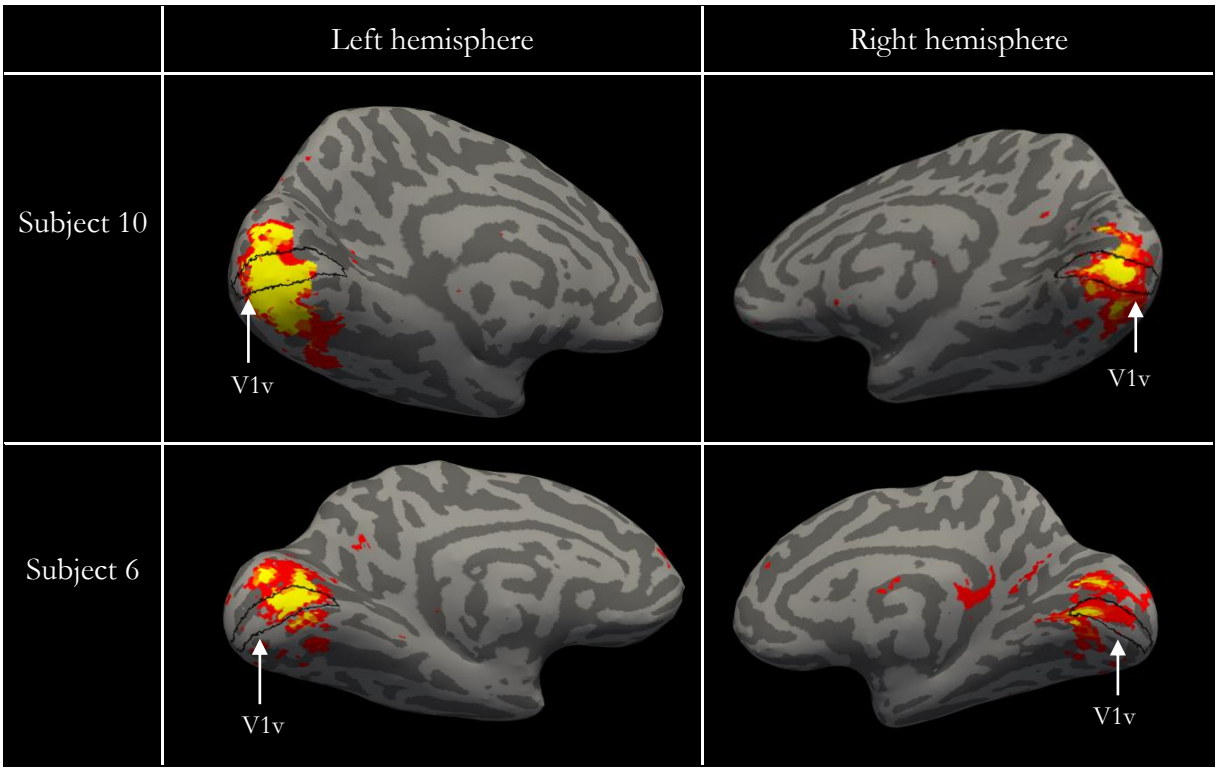

Figure 34. Two example subjects (Subject 10 and 6) who produced bilateral and unique matches to V1v. In both examples, there are clear differences between hemispheres in the coverage of the spatial maps.

A bilateral component can be a unique match to a region in one hemisphere but can also be a top match to multiple regions in the other hemisphere. For example, Subject 14 produced a component that was the top match to V1v in the left ($r = 0.51$) and right ($r = 0.58$) hemisphere (Figure 35, top). The component was unique to V1v in the right hemisphere but was also the top match to V2v in the left hemisphere ($r = 0.39$). Subject 8 produced a similar component that was the top match to V3v in the left ($r = 0.53$) and right ($r = 0.24$) hemisphere. It was also the top match to left hemisphere hV4 ($r = 0.32$). The spatial map reveals strong activation in the left hemisphere V3v and hV4 which extends ventrally into V2v and V3v, as well as dorsally into VO1 (Figure 35, bottom) It is, however, not the top match to these regions. The component also shows good correspondence to V3v in the right hemisphere, but at a weaker magnitude than is seen in the left hemisphere, highlighting the

differences between hemispheres. Again, this suggests a lateralisation in function with multiple regions involved in a network in one hemisphere only.

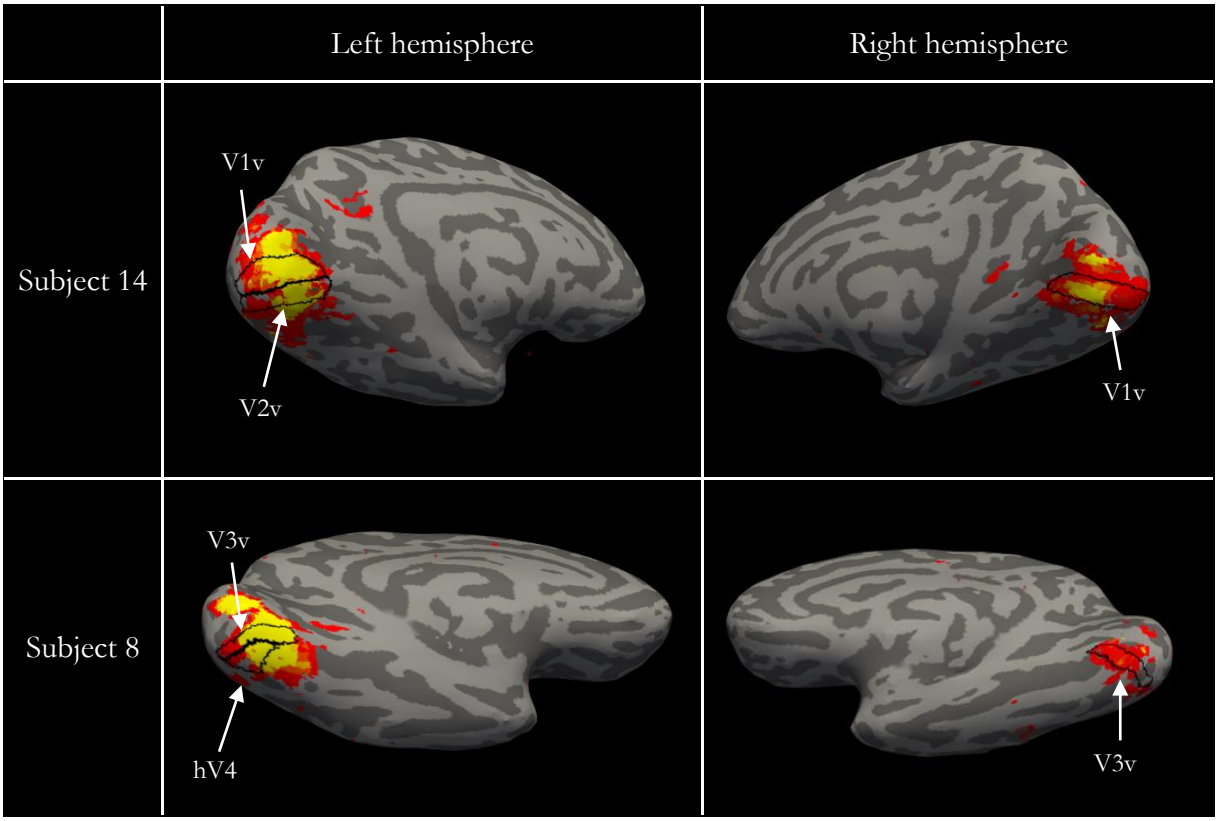

Figure 35. Two examples showing a bilateral component can be a unique match to a region in one hemisphere but can also be a top match to multiple regions in the other. Subject 14 produced a component that was the top match to both V1v and V2v in the left hemisphere, but only V1v in the right hemisphere (top). Subject 8 produced a component that was the top match to both V3v and hV4 in the left hemisphere, but only V3v in the right hemisphere (bottom).

3.3.1.6.1 TO1 and TO2

TO1/MT and TO2/MST are the only regions to show consistency and are therefore discussed in detail here. The two regions showed the highest correlations between the component and atlas, with over 70% of component matches being bilateral. However, as discussed above these regions almost always share the top match component and therefore are rarely unique matches to one of the two regions. As such, we refer to these as hMT+ components.

For the most part, hMT+ bilateral top matches are in fact bilateral components. For example, Subject 15 produced a single component that was the top match to TO1 ($r = 0.74$) and TO2 ($r = 0.73$) in the left hemisphere and TO1 ($r = 0.53$) and TO2 ($r = 0.55$) in the right hemisphere (Figure 36, top). As the correlations suggests, the spatial maps show stronger activation in the left hemisphere compared

to the right. Subject 18 produced a single component that was the top match to TO1/MT (r = 0.69) and TO2/MST (r = 0.64) in the left hemisphere and TO1/MT (r = 0.57) and TO2/MST (r = 0.65) in the right hemisphere (Figure 36, bottom). The spatial maps of these components indicate a bilateral hMT+ component was produced. However, these components are large with activation extending along the lateral surface. Therefore, it is difficult to conclude that hMT+ is accurately and individually identified using ICA.

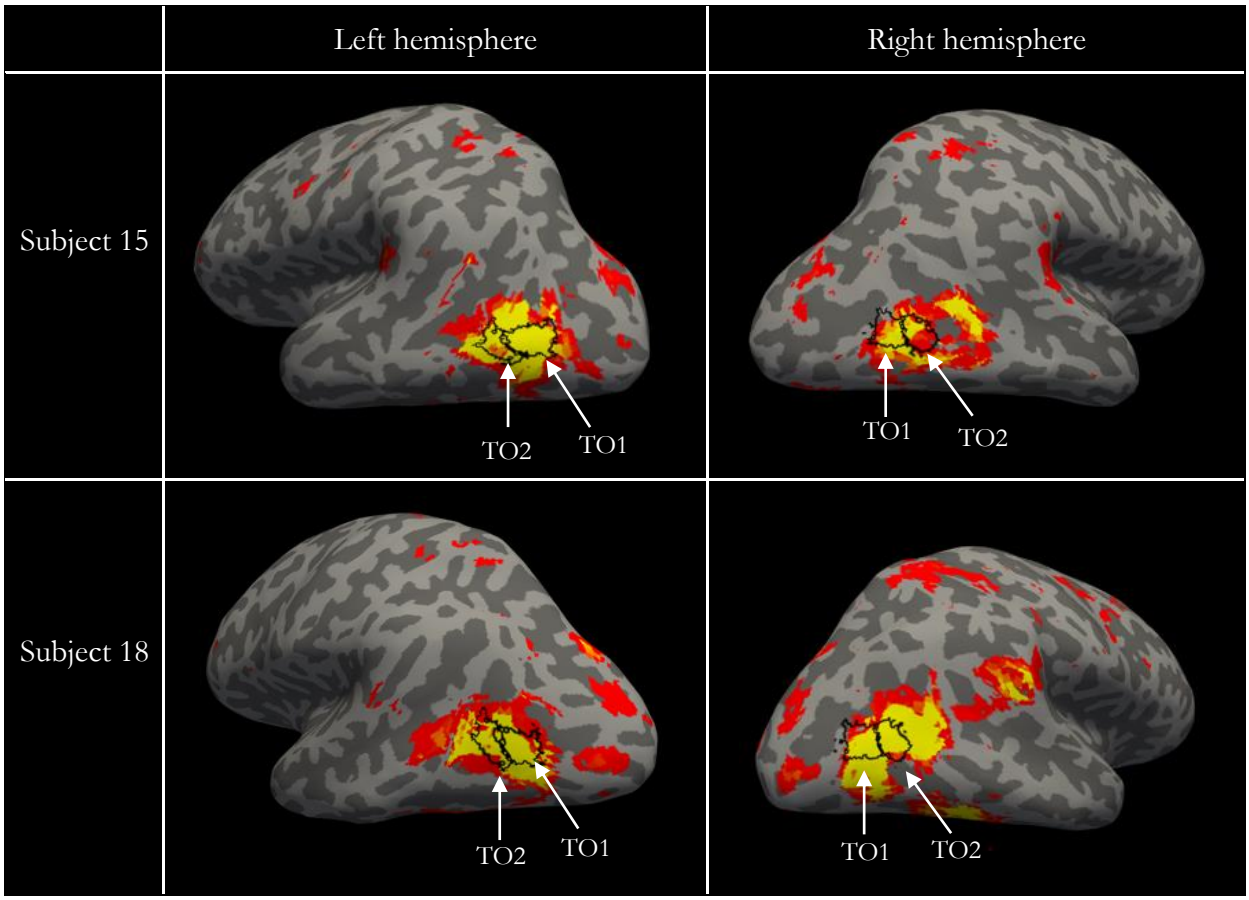

Figure 36. Two examples of bilateral TO1/MT and TO2/MST components. The spatial maps show activation extending beyond the TO1/TO2 atlas regions (black outlines), which is consistently found across brains.

As with other visual areas we find potential evidence of lateralisation in hMT+ when examining unilateral matches. Again, we also find that a unilateral top match does not always indicate a unilateral component.

Subject 8 produced a component that was the top match to TO1/MT (r = 0.55) and TO2/MST (r = 0.60) in the left hemisphere (Figure 37, top) and a separate component that was the top match to TO1/MT (r = 0.49) and TO2/MST (r = 0.35) in the right hemisphere (Figure 37, bottom). This suggests two unilateral hMT+ components had been output. The spatial map of the left

hemisphere top match shows activation in TO1/MT and TO2/MST that extends inferiorly past the atlas regions in the left hemisphere. The component displays only weak/residual activity around TO1/MT and TO2/MST in the right hemisphere. The right hemisphere top match shows activation centred at TO1/MT and TO2/MST in the right hemisphere which, again, extends inferiorly. However, the component also displays activation in these regions, particularly TO1/MT, in the left hemisphere despite it not being the top matched component for the left hemisphere. This suggests a left hemisphere unilateral and a bilateral hMT+ component, with a preference for the right hemisphere, have been output.

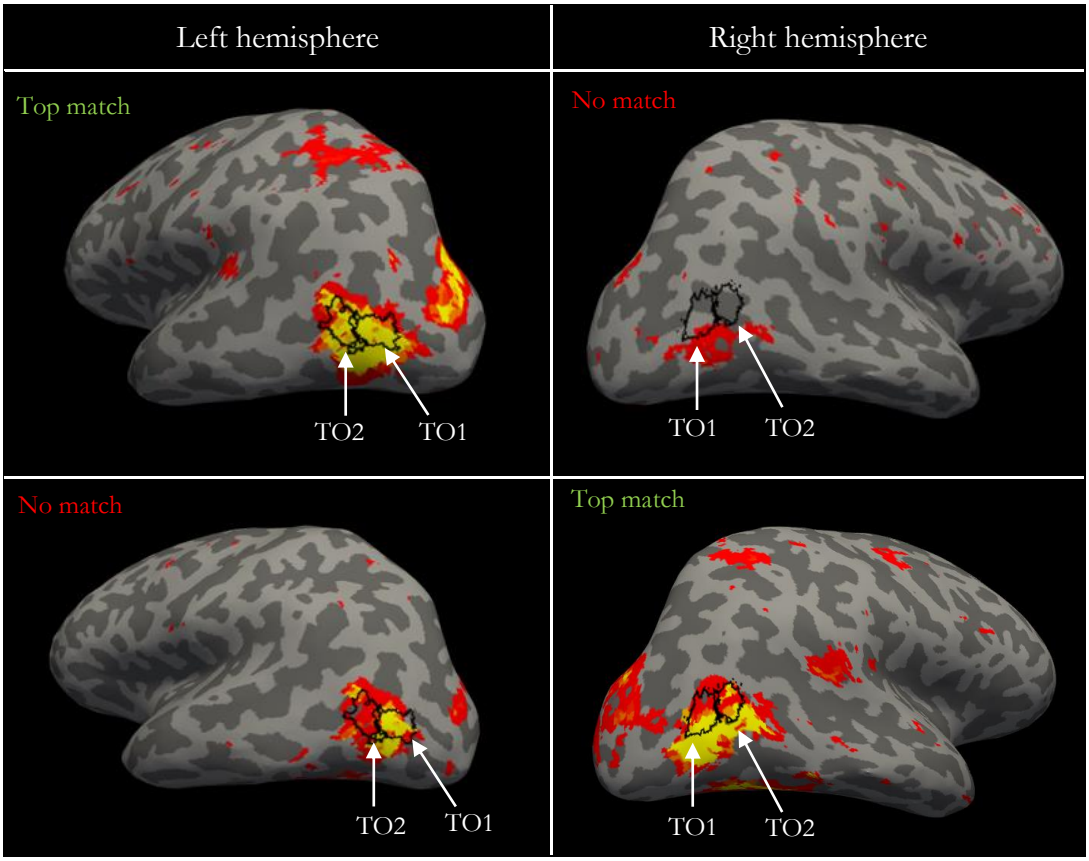

Figure 37. Top: a unilateral top match to TO1/MT and TO2/MST in the left hemisphere. The spatial map shows strong activation in the left hemisphere and only residual activation in the right hemisphere, suggesting the component is indeed unilateral. Bottom: a unilateral top match to TO1/MT and TO2/MST in the right hemisphere. The spatial map shows strong activation in both hemispheres, suggesting the component is in fact bilateral.

Interestingly, the reverse is observed for Subject 12. Again, this subject produced a component that was the top match to TO1/MT ($r = 0.71$) and TO2/MST ($r = 0.57$) in the left hemisphere (Figure 38, top) and a separate component that was the top match to TO1/MT ($r = 0.51$) and TO2/MST ($r = 0.43$) in the right hemisphere (Figure 38, bottom). However, the top match to left hemisphere

TO1/TO2 also shows strong activation in these regions in the right hemisphere which extends inferiorly. Whereas the top match to right hemisphere TO1/TO2 shows no activation in these regions in the left hemisphere. This suggests a right hemisphere unilateral and a bilateral hMT+ component, with a preference for the left hemisphere, have been output.

These findings are broadly in line with Strong et al. (2019) who reported that right hemisphere hMT+ has an enhanced role for processing both the contra- and ipsi-lateral visual fields, i.e., right hMT+ exhibits lateralised function. However, in our data, it is not always right hemisphere hMT+ that appears to show lateralised function (i.e., the production of a bilateral and an additional unilateral component). The potential lateralised function of hMT+ is examined in detail in Chapter 4.

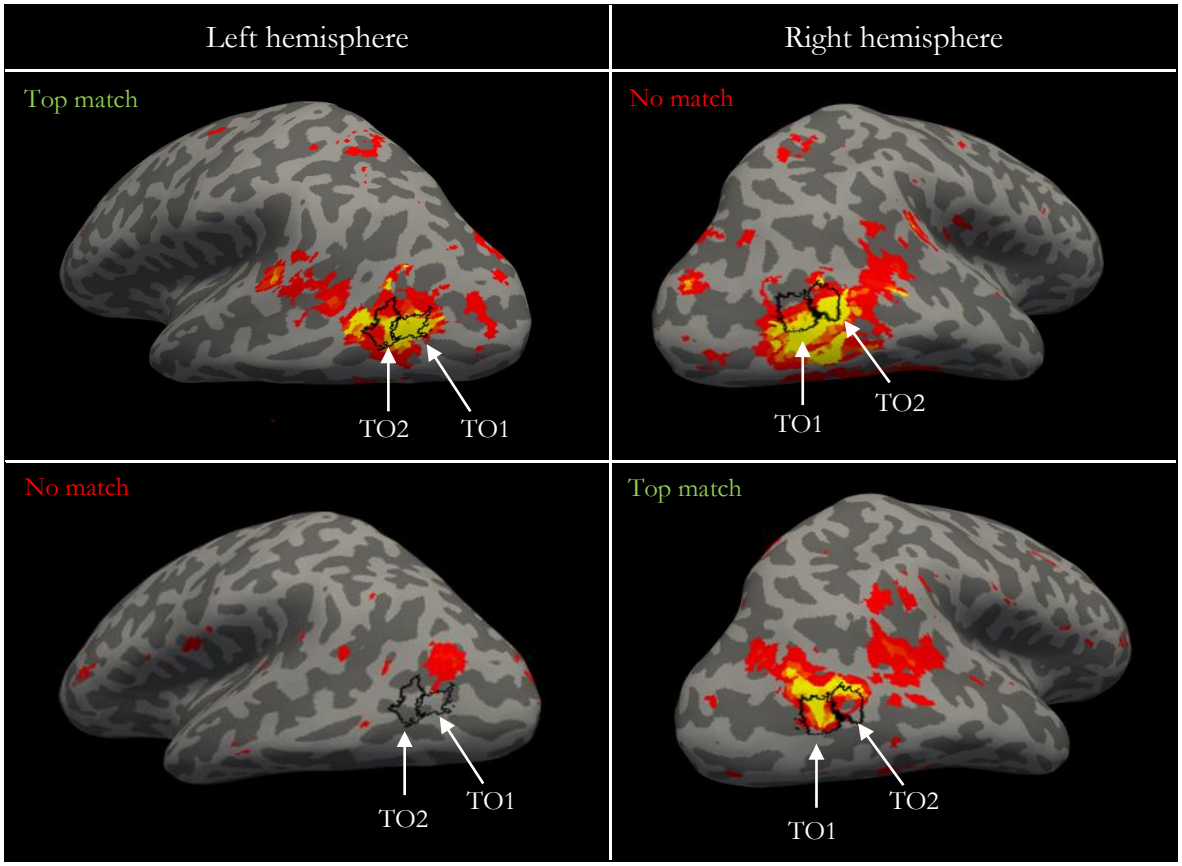

Figure 38. Top: a unilateral top match to TO1/MT and TO2/MST in the left hemisphere. The spatial map shows strong activation in both hemispheres, suggesting the component is in fact bilateral. Bottom: a unilateral top match to TO1/MT and TO2/MST in the right hemisphere. The spatial map shows strong activation in the right hemisphere and only residual activation in the left hemisphere, suggesting the component is indeed unilateral. Note, this pattern is the reverse of that described in Figure 37.

3.3.1.7 Summary: spatial correlations

Bartels and Zeki (2004b, 2005b) found a number of functional regions could be identified when ICA was performed on movie watching data. V1, V2v, V3v, V4 and V5/hMT+ were among the

regions identified in their eight subjects. These findings suggest visually responsive regions of the brain can be segmented in a data-driven manner, without the need for specific, and often time-consuming, functional localisers.

Bartels and Zeki reported that components showed spatial and temporal correspondence across brains. The analyses thus far have explored the former. We performed single-subject ICA on movie watching data to investigate whether visual areas can be segmented in healthy controls with the aim of applying the technique to individuals with cortical vision loss. The results suggest that when viewing broadband naturistic stimuli, a simple division of the brain into visual sub-regions does not hold up for all visual areas.

A combination of an atlas and ICA appears to be a more optimal approach than using one or the other in isolation. Discerning meaningful ICs from noise can be time consuming as many ICs are output. On average, 145 ICs were output from our data, with a range of 55-214 (see Appendix III for individual subject breakdown). Therefore, computing correlations between the ICs and a probabilistic atlas serves as a filter, quickly indicating to the researcher which region(s) a component is likely to correspond to. However, we must note that a higher correlation between the components and the atlas does not always translate to a better match. It is, therefore, imperative that the spatial maps of the components are visually inspected and their location relative to the atlas is examined. For example, the spatial maps of the components showing the highest correlation with the PHC atlas regions shows it to be a worse representation of the region than the 2nd and 3rd ranked component matches (see Figure 19). This is one example of the inconsistencies in the size and coverage of the component maps across subjects.

Overall, there is not always a 1:1 mapping relationship between components and visual areas. For some subjects, components show good spatial correspondence to the atlas regions, but this is not a consistent finding across brains or across the visual areas. Some subjects output unique matches despite the spatial maps showing activation in adjacent regions. Others output large, network components that correlate with multiple visual areas. Even when components match to multiple visual areas, the regions sharing the component are not consistent across subjects (except for TO1-TO2). In

some instances, genuine unique matches are produced but there is no consistency across hemispheres. This suggests some regions are recruited into a larger network and sometimes they function independently.

The bilaterality analysis underlines these differences between hemispheres and indicates potential lateralisation of some areas, specifically TO1/TO2. We find for subjects a single component with activation in both hemispheres, suggesting bilateral function. In the same brain we also find activation in only one hemisphere. This suggests that one hemisphere functioned independently of the other, resulting in a unilateral component, but also functioned bilaterally during free viewing. However, we find no consistent preference as to which hemisphere shows lateralised function.

Similarly, in some subjects, network components are evident in the left hemisphere whilst more constraint activation is present in the right. For other subjects, we find the reverse. This indicates some areas function as part of a network in one hemisphere but independently in the other. Again, we find this varies across subjects and across visual areas.

These findings suggests that ICA is not an accurate tool for consistently identifying single, discrete visual areas. However, ICA may be utilised for identifying differences in function across hemispheres at the individual level.

There was, however, one consistent finding: the production of a hMT+ component. TO1 and TO2 (two sub-regions of V5/hMT+) appear to be consistently identified across subjects, indicated by the highest mean correlation between components and the atlas regions. However, given the top match component for one tends to be the top match to both, it appears the hMT+ complex is identified rather than the subdivisions. This is in line with Bartels and Zeki's findings.

For some subjects, there is evidence of a potential misalignment of some atlas regions, specifically TO1/TO2. As the atlas is taken from a standard surface and resampled to the individual subject surface, some visual areas may be slightly misaligned due to individual differences of the subject's surface. The misalignment could also be due to the anatomical variations in the locations of certain visual areas. As the components derived from ICA are functional, data-driven outputs this may

overcome the issues of misalignment, and thus ICA may be a more optimal approach for identifying the complex.

Comparing the ATCs of matched components across brains, e.g., the top match TO1 component in one brain versus the top match from another brain, allows us to assess whether components which correlate with the atlas genuinely correspond to the same visual areas across brains. The following section examines the temporal correspondence of components across brains. If the ICs represent the regions better than the atlas, we expect their ATCs to be more consistent across brains than the ATCs extracted from the atlas.

3.3.2 Temporal correlations

The preceding section explored the spatial correspondence of ICA components by computing correlations between ICA components and a well-established probabilistic atlas of the visual areas (Wang et al., 2015). We found inconsistencies across visual areas, subjects, and hemispheres. The following section explores the temporal correspondence of the ICA components by computing between- and within-subject ATC correlations.

Bartels and Zeki (2004b, 2005b) proposed that the ATCs of functionally distinct cortical sub-regions should be unique, and therefore can be used to segment the brain into functionally specific regions. They refer to this as a "temporal fingerprint". ICA takes advantage of differential ATCs to cluster together voxels with correlated BOLD activity into components. The idea is that the ATCs of voxels (vertices in our case as components are sampled to the cortical surface) belonging to the same functional region will be correlated and therefore can be dissected based upon their temporal fingerprint alone. Given all subjects watched the same movie clip, Bartels and Zeki hypothesised that if a region was accurately identified, the ATCs of the components corresponding to that region should be highly correlated across brains. Indeed, they found significant correlations between the ATCs of components corresponding to V1, V2v, V3v, V5/MT+, V4 and LO regions, indicating the same region was identified across their eight subjects.

We too take advantage of this approach to assess the temporal correspondence of matched components across brains for all visual areas included in the atlas. First, we extracted the average timecourse of a sub-set of vertices within the atlas region to generate an average ATC. We then generated an average ATC for the component showing the highest correlation to the atlas region. Correlations were then computed between the ATCs of paired components, e.g., the ATC of the top match component to V1v in brain 1 versus the ATC of the top match component to V1v in each of the other brains.

If the same functional regions are represented by spatially corresponding components, their ATCs should be area-specific and thus should show temporal correspondence across brains.

3.3.2.1 Average timecourse and label generation

For each subject and each visual area, an ROI ("label") was generated from which the average timecourse of vertices within it was calculated. First, we warped the Wang et al. (2015) full probability maps to each brain and then for each visual area we identified the 10% of vertices that were most likely to belong to a given visual area. For example, Subject 6's left hemisphere V1v atlas region contains 4230 vertices. We then take the top 10% of vertices from the region (423 vertices)[12].

Next, we create a set of labels for the components which showed the highest correlation with each visual area. These component labels contained the same number of vertices as was included in the atlas region label with which it was the top match. In some instances, the top match component for a visual area is not unique to that region (see section 3.3.1.4). Therefore, if a component correlates with more than one region, we create a component label containing the same number of vertices as the region it matched to. For example, if component x correlates with region y which contains 1000 vertices and region z which contains 1500 vertices, we generate two component labels, one containing the top 100 vertices, one containing the top 150 vertices. Labels are generated for the left and right

[12] This value varies both between areas and between subjects. This is due to differences in the size of each area relative to another and differences in the size of the subject's brain. Larger brains contain more vertices, and therefore when the atlas is mapped to the subject's surface, more vertices are used to represent the area.

hemispheres. The average timecourse is then extracted from each atlas region and the corresponding top matched component labels. We then removed the best straight-fit line from the data using MATLAB's "detrend" command. Within-subject correlations and between subject correlations were then computed.

3.3.2.2 Within-subject correlations

Next, for each subject we explored the temporal correspondence between the components and the atlas regions. For each subject and each hemisphere, the correlation between the mean timecourses of each atlas region and the component showing the highest correlation to that region were computed. We then took the mean across all subjects for each of these pairwise comparisons.

As expected, correlations between the timecourses of the atlas regions and matched components are high. This is due to a) the correlations being computed within-subjects and b) the nature of the component-atlas matching. Components are identified by their spatial correspondence with the atlas region, and therefore the ATCs are extracted from spatially similar locations. Consequently, the ATCs extracted from an atlas region correlated highly with the ATCs extracted from the matched component from the same brains. Figure 39 shows the mean correlations between the ATC for each atlas region and the top matched components. The mean across all observers for left hemisphere correlations (r = 0.78) is comparable to the right hemisphere correlations (r = 0.789).

In the left hemisphere, correlations are highest for TO1 (mean r = 0.926, SD = 0.053), TO2 (mean r = 0.9, SD = 0.047) and V2d (mean r = 0.884, SD = 0.051). The atlas and component timecourses of VO2 (mean r = 0.626, SD = 0.34), PHC2 (mean r = 0.629, SD = 0.280), and PHC1 (mean r = 0.639, SD = 0.319) show the lowest correspondence. In the right hemisphere, correlations are highest for TO1 (mean r = 0.894, SD = 0.067), V1v (mean r = 0.881, SD = 0.112) and V1d (mean r = 0.863, SD = 0.123). The atlas and component timecourses of PHC1 (mean r = 0.642, SD = 0.326), PHC2 (mean r = 0.655, SD = 0.317) and FEF (mean r = 0.689, SD = 0.195) show the lowest correspondence. This initial analysis indicates the peaks of the components, in this case the top 10% of vertices, show good spatial correspondence with the peaks of the probabilistic atlas regions.

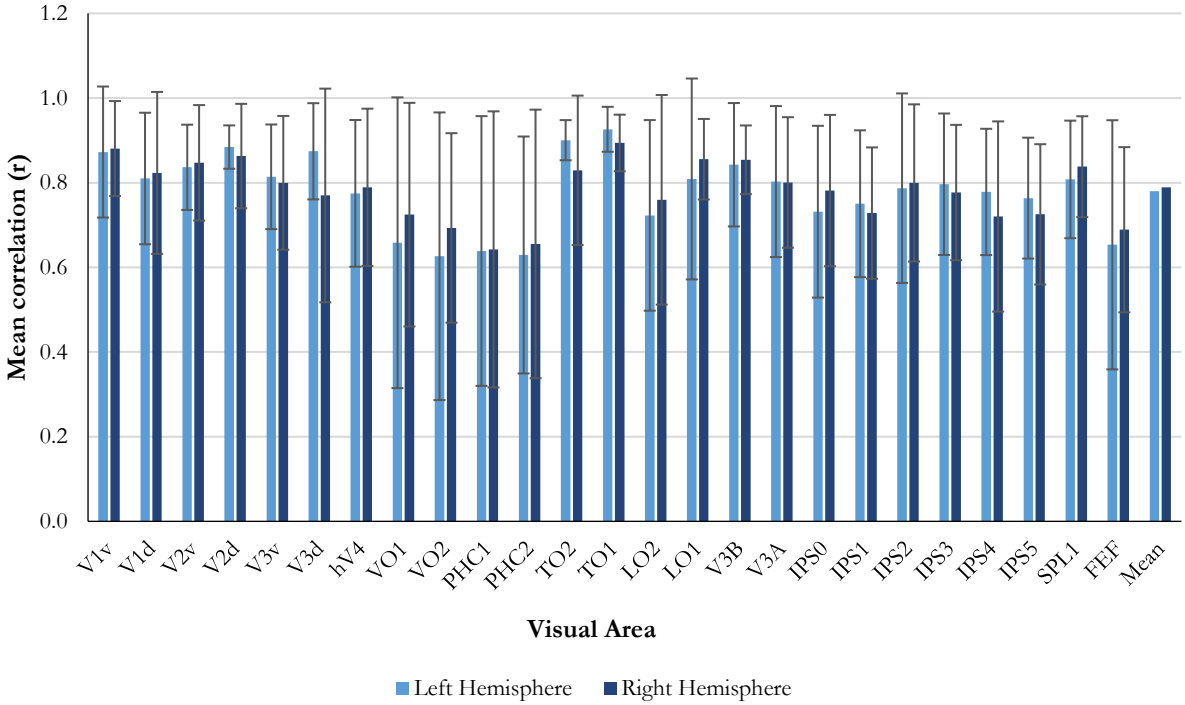

Figure 39. Group mean correlation between the timecourses of each atlas region and the component showing the highest correlation to that region. Correlations were computed for each hemisphere separately. Overall, correlations between atlas and component timecourses are high, indicating the peaks of the components show good spatial correspondence with the peaks of the probabilistic atlas regions. Error bars represent standard deviation. Fisher's Z equivalent figure can be found in Appendix C.

3.3.2.3 Between-subject correlations

Next, we compared the temporal correspondence of components and atlas regions across brains. We found that for TO1/MT and TO2/MST the ATCs extracted from the components were significantly more consistent across brains than the ATCs extracted from the atlas regions. This suggest that for these visual areas at least, components derived from ICA are a better representation of the areas than an atlas. For the remaining visual areas, atlas ATCs were more consistent across brains.

For each subject and each visual area, we computed correlations between the ATC extracted from the component and the ATCs extracted from all other subject's components (pairwise correlations). We then repeated this process for ATCs extracted from the atlas regions and compared the results. A higher correlation indicates that the ATCs were more consistent across brains, and therefore are more likely to belong to the same functionally related regions. If the correlation between the component ATCs is lower, it suggests the activation pattern differed across subjects, and therefore the components may represent functionally distinct regions.

Taking the mean across all visual areas, the ATCs from both the components and the atlas are more consistent in the right hemisphere than the left (Table 7). Averaging across visual areas and hemispheres (grand mean), the ATCs extracted from the atlas regions (r = 0.194) are more consistent than when extracted from the components (r = 0.163), at a difference of 0.031.

Table 7. Atlas and component timecourse consistency, inferred from the group mean pairwise timecourse correlations (Pearson's r). For each subject we computed the correlation between the timecourse extracted from the component and the timecourse extracted from all other subject's corresponding components. These correlations were computed for each visual area and hemisphere separately. We then took the mean across all areas and all brains for each hemisphere (mean timecourse correlation). Next, we took the mean across both hemispheres (grand mean correlation). We then repeated this process for atlas timecourses. Results show timecourses from both the components and the atlas are more consistent in the right hemisphere than the left, and that atlas timecourses are more consistent than component timecourses. Fisher's Z equivalent table can be found in Appendix D.

ROI	Hemisphere	Mean Timecourse Correlation (r)	Grand Mean Correlation (r)
Component	Left	0.160	0.163
	Right	0.166	
Atlas	Left	0.192	0.194
	Right	0.196	

Figure 40 shows the bilateral means (the mean across left and right hemisphere correlations) for each visual area. The highest correlations are observed for TO1 and TO2 regardless of whether the ATCs were extracted from the atlas or components. For TO1 and TO2, the ATCs extracted from the components were more consistent across brains (r = 0.392 and r = 0.413 respectively) than when they were extracted from the atlas (TO1 r = 0.346, TO2 r = 0.322). Conversely, the lowest correlations are observed in FEF for both component (r = 0.024) and atlas (r = 0.058) ATCs.

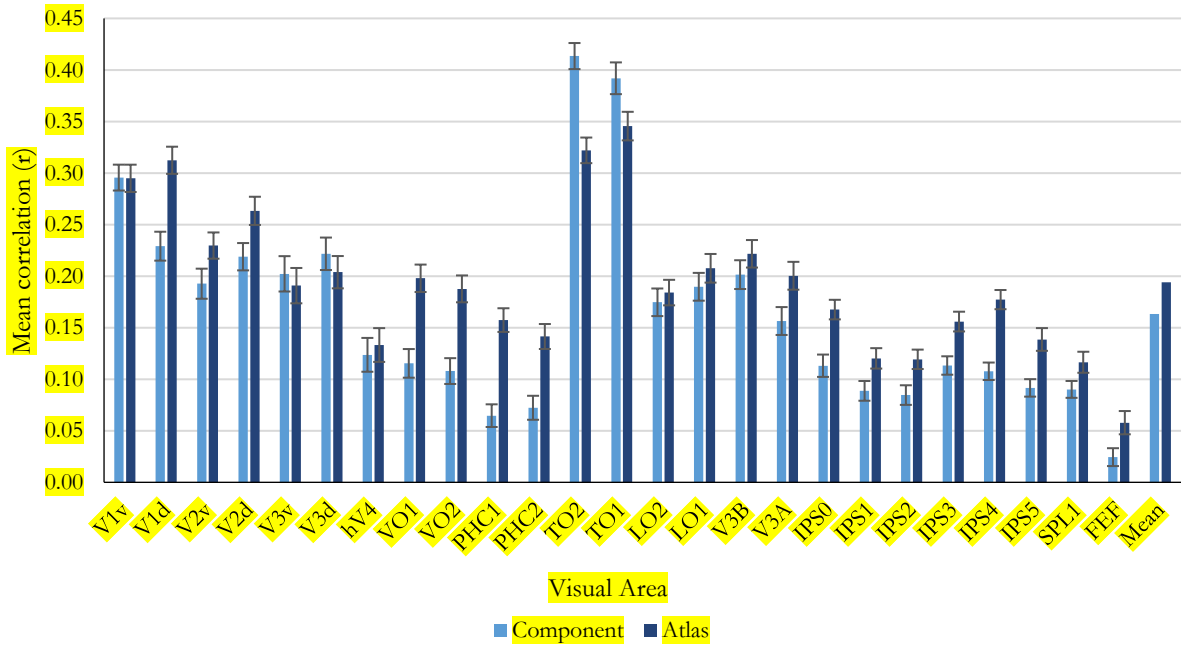

Figure 40. Group mean bilateral timecourse correlations for each visual area. For each subject and each visual area, we computed the correlation between the timecourse extracted from the component and the timecourse extracted from all other subject's components. These correlations were computed for each hemisphere separately before the mean across left and right hemisphere correlations was taken (bilateral mean). We then repeated this process for atlas timecourses. Results show the highest correlations are observed for TO1/MT and TO2/MST regardless of whether the timecourses were extracted from the atlas or components. Fisher's Z equivalent figure can be found in Appendix E.

Table 8 shows, for each visual area, the ratios of the mean pairwise correlation. The ratios were calculated by dividing the mean of the pairwise correlations for the ICA components by the mean of the pairwise correlations for atlas regions. Values = 1 (white) indicate the component and atlas timecourses were equally consistent across brains. Values >1 (blue) indicate the component timecourses were more consistent across brains, whereas values <1 (yellow) indicate the reverse. For 5/25 visual areas (V1v, V3v, V3d, TO1, TO2) the ATCs extracted from the components were more consistent across brains than when the ATCs were extracted from the atlas.

Next, for each of the five visual areas, we assessed whether ATCs extracted from the components were significantly different to those extracted from the atlas. The pairwise correlation matrices were reshaped into vectors (see Figure 41 for schematic example) before paired-samples t-tests were performed between the atlas and component pairwise correlation vectors. The analysis revealed there to be no statistically significant difference between atlas and component ATCs for V1v, V3v or V3d.

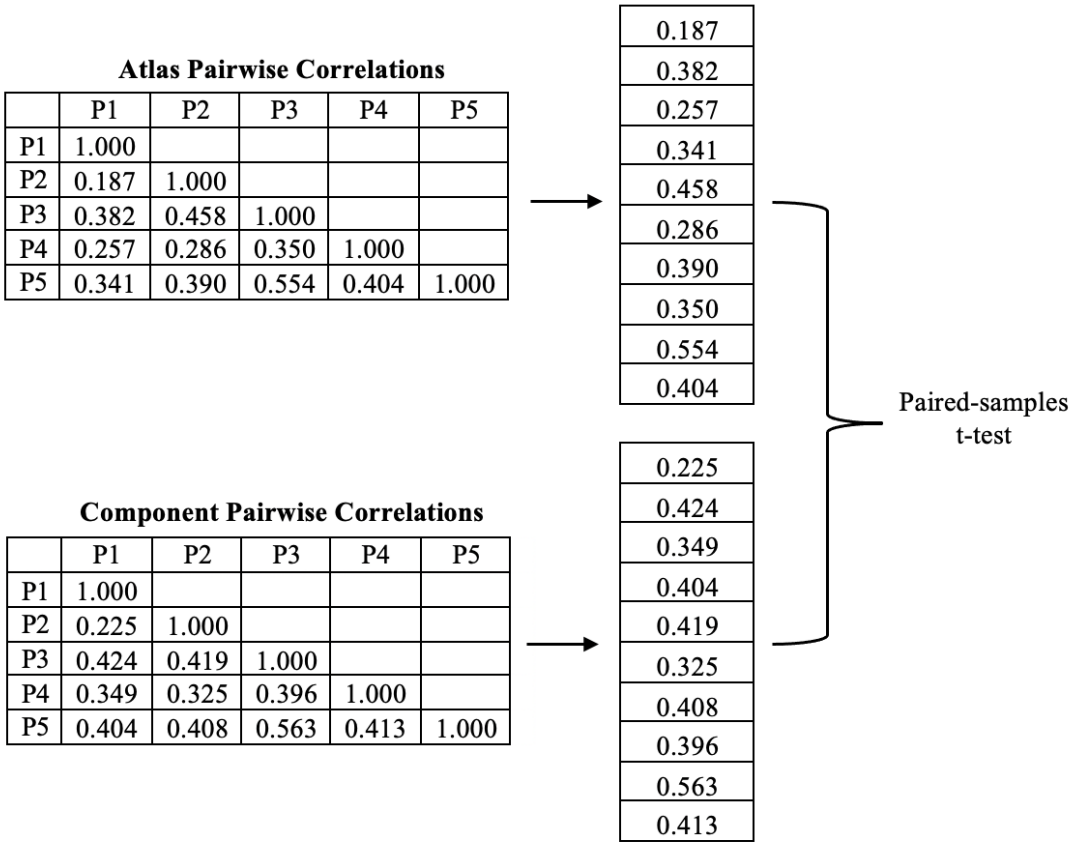

Figure 41. A schematic example of how component and atlas subject-to-subject timecourse correlation values (Pearson's r) were transformed from a correlation matrix (left) to a vector (right). Paired-samples t-tests were then performed between the component and atlas timecourse correlation vectors. Values in the matrices/vectors represent subject-to-subject pairwise correlation values.

The analysis did reveal a significant difference between the component and atlas ATCs extracted from TO1 and TO2. This supports the spatial correlation findings in that it appears the components derived from ICA are a better representation of the areas. As noted above, the two regions almost exclusively appear in the same component, and therefore we caveat this conclusion by stating that ICA consistently identifies the hMT+ complex, rather than its subdivisions.

The ATCs of the remaining regions were more consistent when extracted from the atlas, suggesting that overall, the probabilistic atlas is a more accurate way of identifying visual areas across subjects. The most pronounced differences in the consistencies of the ATCs are observed in V1d, PHC1/2 and VO1/2. We discuss these regions below.

It is important to note that the reasons for the component ATCs being less consistent are not exclusive to the regions discussed below. A lack of uniqueness, below threshold matches, and large

spatial maps covering adjacent regions are characteristics reported for all visual areas (see section 3.3.1). It does, however, appear that these characteristics are more common in PHC1/2, VO1/2 and V1d compared to the other visual areas[13].

Table 8. Ratios of the mean pairwise timecourse correlations. The ratios were calculated by dividing the mean of the pairwise correlations for the ICA component timecourses by the mean of the pairwise correlations for atlas region timecourses. Values = 1 (white) indicate the component and atlas timecourses were equally consistent across brains. Values >1 (blue) indicate the component timecourses were more consistent across brains, whereas values <1 (yellow) indicate the reverse. For 5/25 visual areas (V1v, V3v, V3d, TO1, TO2) the ATCs extracted from the components are more consistent across brains than when the ATCs are extracted from the atlas (indicated by red borders).

Visual Area	Ratio of the component grand mean to atlas grand mean
V1v	1.002
V1d	0.733
V2v	0.839
V2d	0.831
V3v	1.060
V3d	1.088
hV4	0.928
VO1	0.583
VO2	0.575
PHC1	0.411
PHC2	0.511
TO2	1.284
TO1	1.134
LO2	0.949
LO1	0.914
V3B	0.909
V3A	0.781
IPS0	0.675
IPS1	0.738
IPS2	0.709
IPS3	0.727
IPS4	0.608
IPS5	0.661
SPL1	0.774
FEF	0.422

[13] The following section details the results of paired samples t-tests performed assess whether ATCs extracted from the components were significantly different to those extracted from the atlas. We performed identical tests on values following the transformation to Fisher's Z, the results of which can be found in Appendix F. Note, the transformation did not alter the results presented here.

3.3.2.3.1 PHC1/2

On average, the ATCs extracted from the atlas correlated higher across brains (M = 0.157, SD = 0.134) than ATCs extracted from PHC1 components (M = 0.064, SD = 0.129). This difference, 0.093, was statistically significant, t(271) = -10.796, p < .001 (uncorrected). Similarly, for PHC2, the ATCs extracted from the atlas correlated higher (M = 0.141, SD = 0.141) than the ATCs extracted from the components (M = 0.072, SD = 0.136). Again, this difference, 0.069, was statistically significant, t(271) = -7.891, p < .001 (uncorrected). There was no significant difference between hemispheres for both the atlas and the component ATCs.

We previously noted that PHC regions are problematic because despite showing strong correlations with the atlas, the spatial maps of the components are either incomplete or widespread. Therefore, it is not surprising that ATCs from the atlas are more consistent across brains than the component ATCs since the components themselves are inconsistent.

3.3.2.3.2 VO1/2

On average, ATCs extracted from the VO1 atlas correlated higher across brains (M = 0.198, SD = 0.155) than the ATCs extracted from the components (M = 0.166 SD = 0.162). This difference, 0.082, was statistically significant, t(271) = -7.915, p < .001 (uncorrected). Similarly, for VO2, the ATCs extracted from the atlas correlated higher (M = 0.188, SD = 0.152) than the ATCs extracted from the components (M = 0.108, SD = 0.146). Again, this difference, 0.08, was statistically significant, t(271) = -7.257, p < .001 (uncorrected). There was no significant difference between hemispheres for both the atlas and the component ATCs.

A combination of subjects producing fewer unique matches and fewer matches above the r = 0.3 threshold leads to the ATCs of VO components being less comparable across brains than ATCs from the atlas. As noted above VO regions showed clear hemispheric differences. More subjects produced VO1 and VO2 matched components above threshold in the right hemisphere (both 88%) than in the left hemisphere (76.5% and 59% respectively). Although we have noted that a stronger correlation between a component and the atlas region does not always translate to a better

representation of the visual area (e.g., PHC matches), a below threshold match does almost always indicate a poor representation of the visual area. Given several subjects produced below threshold VO top matches, the ATCs extracted from these components are likely to be dissimilar.

Even when a component match is above threshold, the uniqueness of matches impacts the correlations. For VO1, 18% of subjects produced a unique match in the left hemisphere, compared to 30% in the right. For VO2, 24% of subjects produced a unique match in the left hemisphere, compared to 6% in the right. For both regions, there was no consistent component sharing pattern in the left hemisphere, in that different regions were included in the VO1/VO2 components (see Figure 22). In the right hemisphere, 29% of subjects produced a component that was the top match to VO1 and VO2, and a further 29% of subjects produced a component that was the top match to VO2, PHC1 and PHC2. It is clear that VO1 and VO2 matches are seldom unique, and when they are shared the regions included in the component varies across brains. This indicates a high amount of variability in the spatial maps across subjects.

The variability of the spatial maps of these components results in the ATCs being extracted from vertices outside of the regions (when compared to the atlas). Figure 42 shows two examples of components which were the top match to both VO1 and VO2. Overlaid on the components are the atlas regions (black) along with the label encompassing the top 10% of vertices from the component (white).

Subject 6's component correlated with VO1 and VO2 at $r = 0.359$ and $r = 0.316$ respectively (Figure 42, left). Subject 17's component correlated with VO1 and VO2 at $r = 0.324$ and $r = 0.354$ respectively (Figure 42, right). Both components correlated with the atlas regions above threshold, but the spatial maps show only residual activation present within the VO regions. The strongest activation, and thus the top 10% of vertices, is located outside of the atlas labels. Naturally, the ATCs extracted from these components show low correlations. The ATCs of the two subjects VO1 and VO2 components correlated with each other at $r = -0.077$ and $r = -0.076$ respectively. The nature of VO1 and VO2 matches leads to the ATCs being extracted from spatially different locations across brains, which in turn leads to lower correlations between the ATCs.

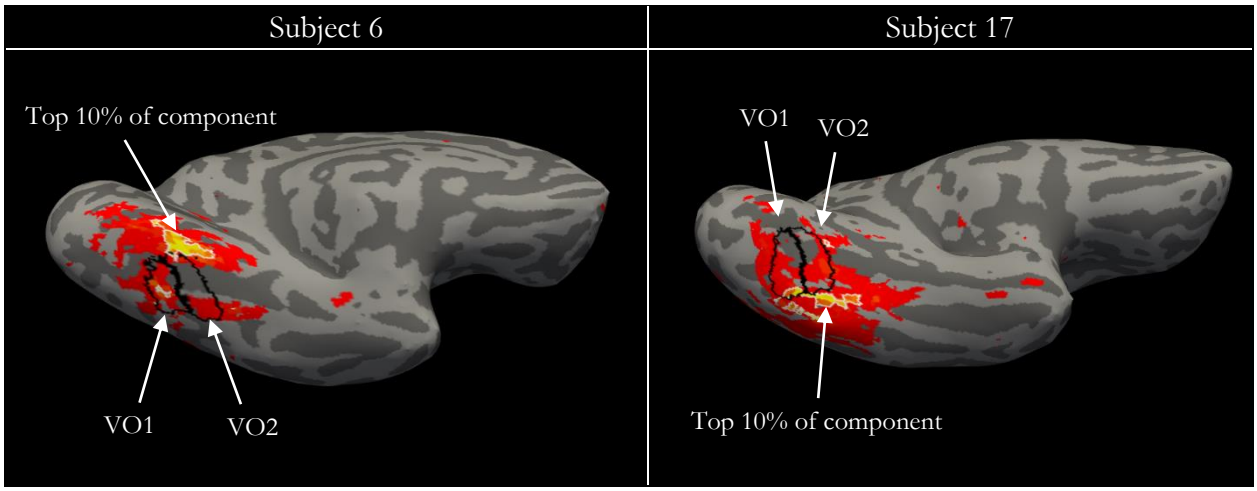

Figure 42. Two examples showing how the variability across brains in the spatial maps of components results in ATCs being extracted from vertices outside of the regions. Both components correlated with the VO1 and VO2 atlas regions above threshold, but the spatial maps show only residual activation present within the VO regions (black). The strongest activation, and thus the top 10% of vertices (white), is located outside of the atlas labels. The ATCs extracted from these two components show low correlations, likely due to the fact they are extracted from anatomically non-corresponding locations.

3.3.2.3.3 V1d

In V1d, the ATCs are more consistent across brains when extracted from the atlas because of the inconsistencies in the coverage of the component matches. Only one subject failed to produce an above threshold match to V1d (in the left hemisphere only), meaning the low correlations are not a result of below threshold matches being output. As with the VO regions, the lower correlations seen for V1d component ATCs are being driven by a high degree of variability in the spatial maps of the components.

Uniqueness of V1d matches was shown to vary across hemispheres, with 59% of subjects producing a unique match in the left hemisphere, compared to 18% in the right hemisphere. On average, the top match to V1d in the left hemisphere was the top match to 1.71 visual areas, compared to 2.29 in the right. If V1d was accurately identified using ICA, then more unique matches in the left hemisphere should be reflected by more consistent ATCs compared to the right hemisphere. However, on average the ATCs from the left hemisphere components are less consistent ($M = 0.197$, $SD = 0.167$), compared to ATCs from the right ($M = 0.261$, $SD = 0.160$). This difference, 0.064, was statistically significant, $t(135) = -5.368$, $p < .001$ (uncorrected). On the other hand, left hemisphere atlas ATCs ($M = 0.307$, $SD = 0.150$) are comparable to the right hemisphere ATCs ($M = 0.318$,

SD = 0.157). This difference, 0.011, was not statistically significant, t(135) = -1.340, p = .183 (uncorrected).

The reason for the large discrepancy between hemispheres appears to be due in part to a particularly large component match produced by a single subject. Subject 16 produced a unique match to left hemisphere V1d which correlated with the atlas at r = 0.284. The spatial map shows widespread activation on the medial surface of the occipital lobe (Figure 43). On average, the ATC extracted from this component correlated with the ATCs of the remaining subjects at r = -0.040. Excluding this subject from the analysis results in an increase in the component mean correlations for both the left (M = 0.228, SD = 0.150) and right hemispheres (M = 0.265, SD = 0.163). Nonetheless, the difference between hemispheres, 0.037, remains statistically significant[14], t(119) = -2.725, p = .007 (uncorrected).

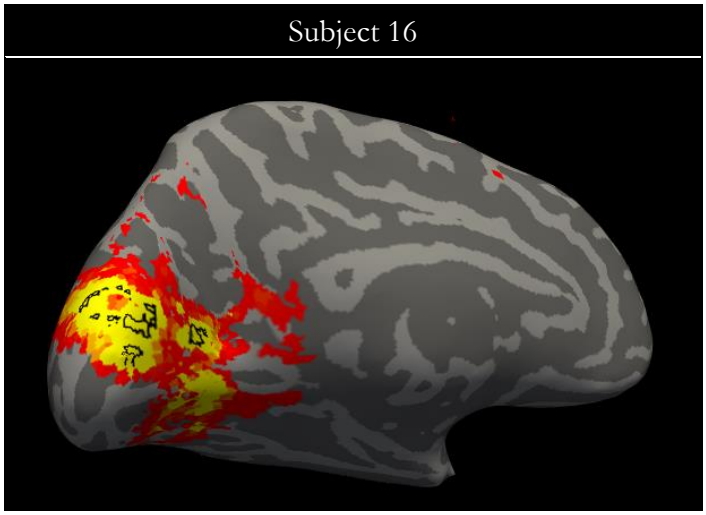

Figure 43. Subject 16's unique match to left hemisphere V1d. The spatial map shows widespread activation on the medial surface of the occipital lobe. The top 10% of vertices, from which the mean timecourse is extracted from (indicated by the black ROI), are dispersed throughout the component. This results in low correlations with the timecourses of other subjects V1d components as the timecourses are not extracted from anatomically corresponding locations across brains.

Given there are significant differences between hemispheres, we examined the correlations within the same hemisphere (e.g., left hemisphere atlas ATCs against left hemisphere component ATCs). In the left hemisphere, atlas ATCs were more highly correlated (M = 0.314, SD = 0 .145) than component ATCs (M = 0.228, SD = 0.150). This difference was statistically significant, t(119) = -

[14] When Subject 16 was excluded from the analysis there was no statistical difference between left hemisphere (M = 0.314, SD = 0.145) and right hemisphere (M = 0.333, SD = 0.157) ATCs extract from the atlas, t(119) = -1.880, p = 0.063 (uncorrected).

7.124, $p < .001$ (uncorrected). Likewise, right hemisphere atlas ATCs were more highly correlated ($M = 0.333$, $SD = 0.157$) than component ATCs ($M = 0.265$, $SD = 0.163$). Again, this difference was statistically significant, $t(119) = -8.335$, $p < .001$ (uncorrected). Therefore, ATCs extracted from the components are less comparable across brains than atlas ATCs, with component ATCs showing hemispheric differences. Below I discuss the possible reasons for these differences.

When the atlas is used, the top 10% of vertices are clustered around the centre of the area (Figure 44) and therefore, the average ATC is extracted from neighbouring vertices. Providing the atlas is not misplaced, this results in the ATCs being correlated across brains ($r = 0.345$ for the example in Figure 44). This is in line with Bartels and Zeki (2004b, 2005b) who showed that ATCs extracted from anatomically corresponding region show strong temporal correspondence

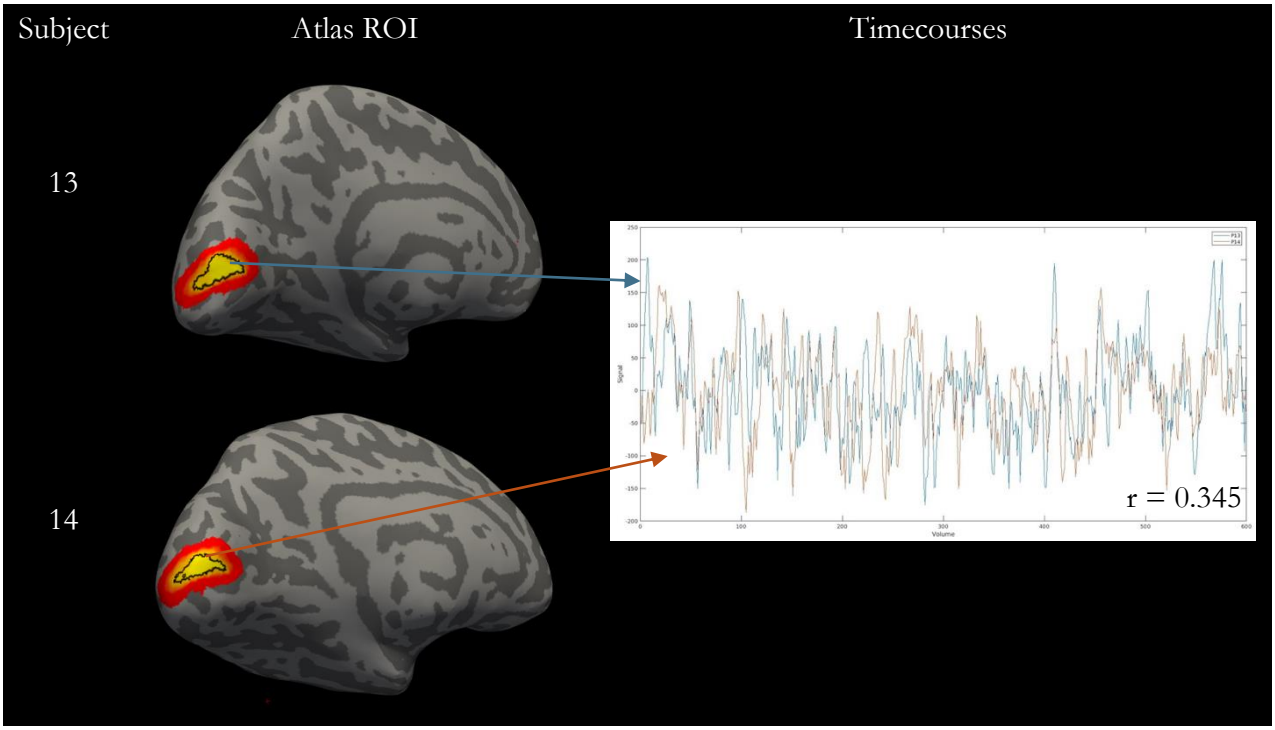

Figure 44. An example showing why timecourses extracted from the probabilistic atlas ROIs are more highly correlated across brains than timecourses extracted from components. When an atlas is used, the timecourses are extracted from the top 10% of vertices (i.e., the vertices with the highest probability of belonging to that visual area) which are often clustered around the centre of the area. This results in timecourses being extracted from anatomically corresponding locations across brains. Here, I show the V1d full probability map (red-yellow) and the top 10% of vertices ROI (black) from which ATCs are extracted. Timecourses extracted from the V1d ROI are more highly correlated than when the component is used for ATC extraction.

However, the variation in the coverage of the component spatial maps leads to the ATCs being extracted from differing vertex locations for each subject. This is a result of large components

correlating with multiple visual areas, or a result of what I refer to as pseudo-unique matches[15]. ATCs are therefore extracted from vertices distributed throughout a component which in turn leads to lower correlations between the ATCs. In short, ATCs are extracted from non-corresponding spatial locations for some subjects.

To recap, there are three typical classes of component matches and it is the proportions of these produced by the subjects which impact the temporal correspondence of matched components for a given visual area:

1. A pseudo-unique match (Figure 45A and B).
2. A genuine unique component (Figure 45C).
3. A shared component (Figure 45D and E).

In the case of V1d, there are more pseudo-unique matches in the left hemisphere and more shared components in the right hemisphere. Subject 14 produced a genuine unique match which correlated with the atlas region at $r = 0.58$, suggesting good spatial correspondence between the two (Figure 45C). However, as such components are seldom produced, we are fundamentally computing correlations between the ATC of a good unique match, and the ATCs of the more often produced pseudo-unique and shared components. On average, the ATC of the genuine unique match produced by Subject 14 correlated with all other V1d ATCs at $r = 0.199$.

Now, as these matches are the exception, rather than the rule, the correlations between the ATCs of genuine unique matches, which appear to be good representations of a single visual area, and pseudo-unique or shared component matches are going to be low.

Interestingly, when components spatially correspond across brains, their ATCs are more highly correlated. Take Subject 20 and 15 for example (Figure 45A and B), both of which produced pseudo-unique matches. The spatial maps of these components show activation mostly within V1d but also in

[15] Pseudo-unique matches are components which are the top match to only one visual area despite activation being present in adjacent regions. These types of matches are frequently produced by subjects and led to the conclusion that a unique component match does not always indicate a clean representation of the region alone.

adjacent regions. The top 10% of vertices also follow a similar pattern in that most are clustered in the centre of the component. The ATCs of these two components correlate at r = 0.513. Whereas the ATCs correlated with the unique match ATC at r = 0.329 and r = 0.377 respectively. This suggests the ATCs of a genuine and pseudo-unique component are less correlated than the ATCs of spatially comparable, but not genuinely unique, components.

Much like pseudo-unique matches, shared components show activation spread across the adjacent regions, and therefore the top 10% of vertices are dispersed throughout the component, rather than a focal cluster of vertices as in the case with the atlas and genuine unique components. When there is consistency with respect to which regions are included in a component, the ATCs show good correspondence. For example, Subjects 17 and 18 produced components that were the top match to V1d and V2d (Figure 45D and E). The ATCs of these components correlated at r = 0.403. However, as noted above, in the left hemisphere 41% of matches were shared, with no more than two subjects showing the same pattern of component sharing. This highlights the high level of inconsistency with respect to which regions are recruited into larger network components and this drives the correlations down. Moreover, this inconsistency offers an explanation for the significant hemispheric differences in the ATCs.

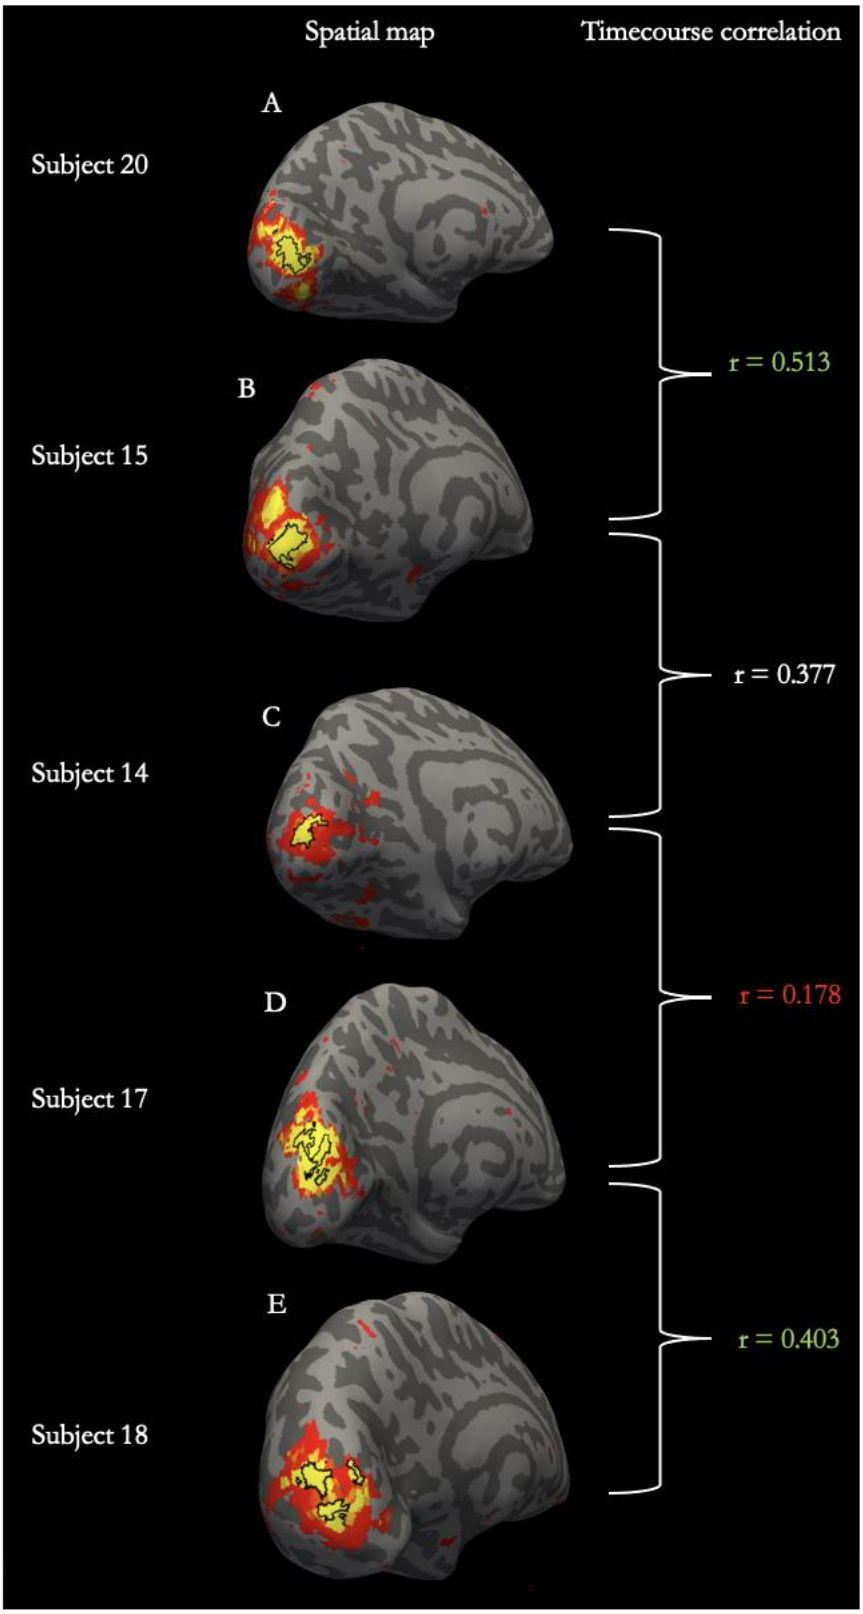

Figure 45. The spatial maps of five example subjects (middle) and the pairwise correlations computed between timecourses extracted from the component ROIs (right). Black outlines indicate the 10% of vertices from which the timecourses were extracted. Timecourses extracted from components that anatomically correspond across brains show high correlation (e.g., component A versus B), even when the components match more than one visual area (e.g., component D versus E). Timecourse extracted from components that are not anatomically aligned across brains (e.g., component B versus C), result in lower correlations. Likewise, correlations computed between the timecourse of a component that matches with more than one visual area (resulting in the 10% of vertices being dispersed throughout the component), and a discrete component (resulting in the top 10% of vertices being more focal cluster) show lower correlations (e.g., component C versus D).

3.3.2.3.4 TO1/TO2

The spatial maps of TO1/TO2 components indicated that for some subjects the atlas was slightly misaligned. If this was the case, and the components derived from ICA represent the regions better than the atlas, the ATCs of components should be more consistent than the atlas. Indeed, we find this to be the case.

For TO1/MT, the ATCs extracted from the components were more consistent (M = 0.392, SD = .180) than when extracted from the atlas (M = 0.346, SD = 0.162). This difference, 0.46, was statistically significant, $t(271) = 9.670$, $p < .001$ (uncorrected). Likewise, the ATCs extracted from TO2/MST components were more consistent (M = 0.413, SD = 0.148) than when extracted from the atlas (M = 0.322, SD = 0.147). Once again, this difference, 0.091, was statistically significant, $t(271) = 16.081$, $p < .001$ (uncorrected).

For the overwhelming majority of subjects, TO1/MT and TO2/MST share a top match component. As expected, the ATCs from these components are comparable (difference of 0.021). It should be noted that for these subjects the ATCs are not being extracted from TO1/MT or TO2/MST individually, but rather from the hMT+ complex (for example, Figure 46, top). This is, however, not the case for all subjects.

Large, non-unique components reduce the correlations across subjects. We noted in the spatial correlation section that Subject 16 produced a component that was the top match to TO1/MT, LO2, V3A, and V3B (Figure 46, bottom, left). The top 10% of vertices are located throughout the component, rather than clustered locally. Not surprisingly, on average the ATC extracted from the TO1/MT component correlated with all other subjects TO1/MT ATCs at $r = 0.022$, suggesting the component is not an accurate representation of TO1/MT. Excluding Subject 16 from the TO1/MST mean correlation calculation leads to an increase of 0.059, to $r = 0.445$.

Subject 16 also produced a unique TO2/MST component, the ATC of which correlated with the remaining subjects TO2/MST ATCs at $r = 0.148$. The spatial map suggests this is in fact a hMT+ component (Figure 46, bottom, right), and the atlas is slightly misplaced due to anatomical variation. Interestingly, Subject 16's TO2/MST component ATC correlated on average with all other subjects

TO1/MT component ATCs at r = 0.171. Given the probabilistic atlas appears to be susceptible to misalignment due to anatomical variability of hMT+, the components derived from ICA may be a better representation of the complex.

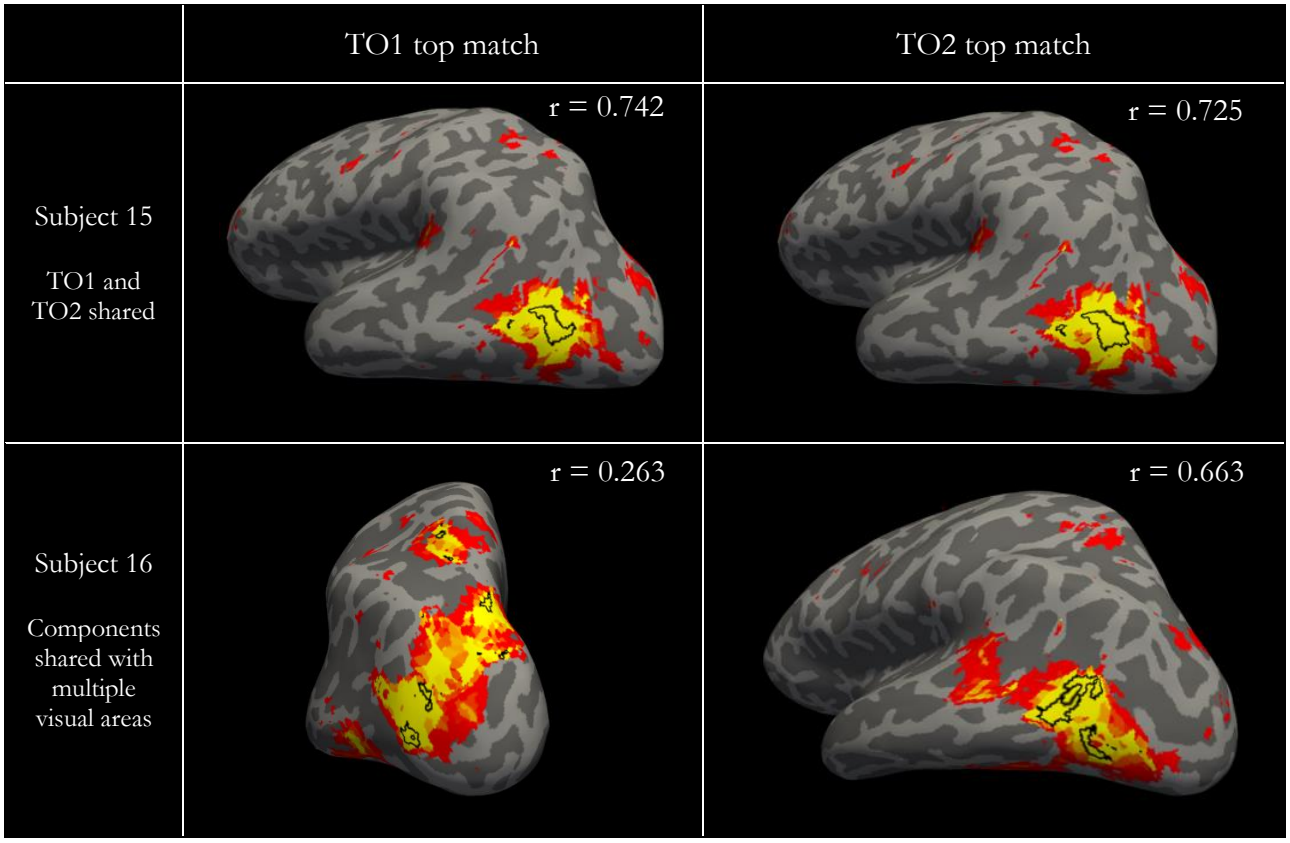

Figure 46. Top: an example of a hMT+ component (produced by Subject 15). The component was the top match to both TO1/MT and TO2/MST. The peak of the component, indicated by the 10% of vertices (black border overlaid on the component), is located in the centre of the component. Bottom: Subject 16 produced a separate component match for TO1/MT (left) and TO2/MST (right). The TO1/MT component was also the top match to LO2, V3A, and V3B, whereas the TO2/MST component was unique. The spatial map suggests the TO2/MST component (bottom, right) may be a hMT+ component and the atlas used for spatial correlations may be misplaced in this subject due to anatomical variability. For reference, the spatial correlation, r, between the component and atlas regions are also displayed.

Subject 10 produced a unique match to both TO1/MT (Figure 47, top) and TO2/MST (Figure 47, middle) in the left hemisphere. The spatial map of the TO1/MT component closely resembles the hMT+ components produced by the other subjects, albeit slightly more confined, suggesting the component is not a genuine unique TO1/MT component. This assertion is supported by the ATC of the component correlating on average with the remaining subject's TO1/MT component ATCs at r = 0.511. The TO2/MST component is much more confined, with activation present within both the TO1/MT and TO2/MST atlas regions, again indicating the component is not unique to TO2/MST. This component ATC correlated with the remaining subjects ATCs at r = 0.532. The peak of both

components, indicated by the top 10% of vertices, show substantial overlap (Figure 47, bottom). This suggests the same area of cortex strongly modulated both components and could be indicative of the brain dynamically transitioning between brain states, one in which a large area of the cortex around hMT+ is co-active, another in which a more confined region is processing the stimulus.

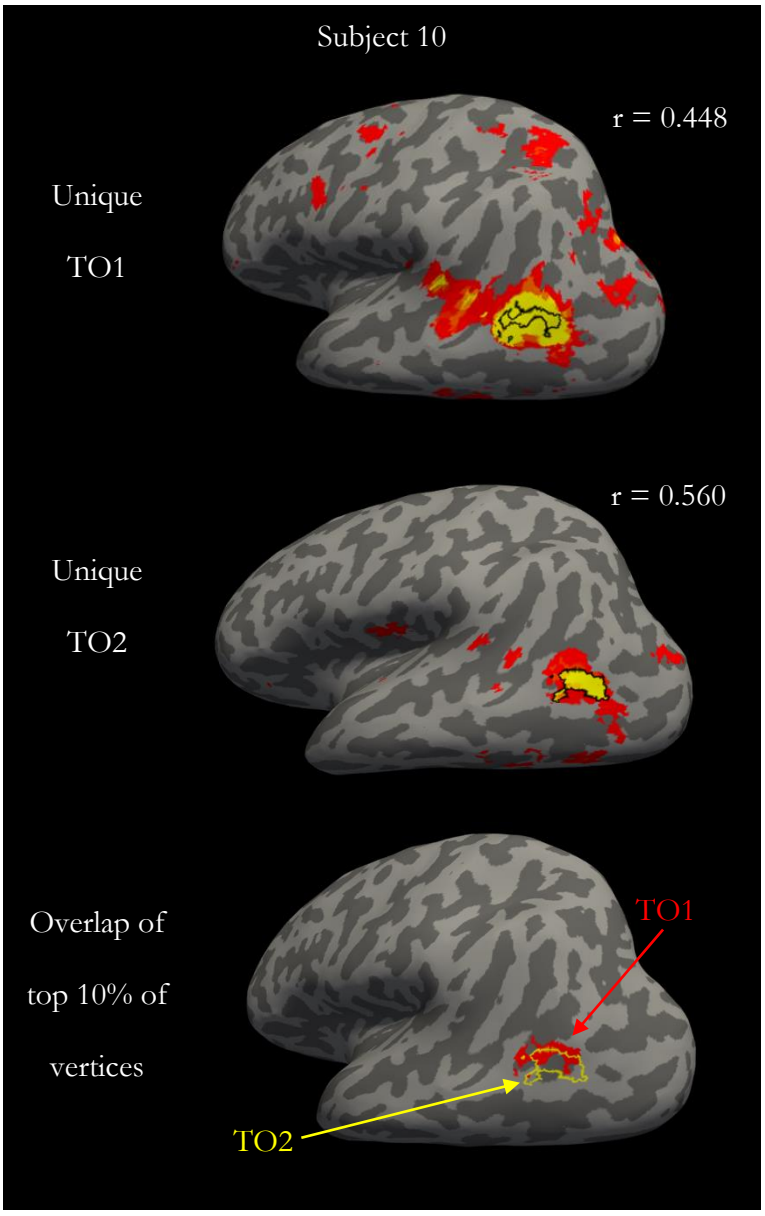

Figure 47. The unique component matches to TO1/MT (top) and TO2/MST (middle) produced by Subject 10. The peaks of both components (the top 10% of vertices) are indicated by the black ROIs overlaid on the spatial maps. The spatial correlation, r, between component and the atlas are also displayed. The peaks of both components show substantial overlap (bottom), suggesting the same area of cortex strongly modulated both components. This could be indicative of the brain dynamically transitioning between brain states.

3.4 Chapter summary

In traditional visual mapping paradigms, subjects are often presented with a simplistic stimulus and instructed to maintain fixation for a prolonged period of time. These tasks are often attentionally

demanding, highly controlled, and therefore incongruent with our natural visual experience. In contrast, the chronotopy approach allows subjects to freely move their eyes around a complex scene embedded with multiple visual stimuli and complex motion. Therefore, free viewing of a movie offers a more ecological way to investigate the brain's functional architecture.

However, the results from our spatial analyses show that during natural viewing, a simple division of the brain into visual sub-regions does not hold up for all visual areas. Instead, we find a high amount of variability across brains and across visual areas for each of our analyses. The inconsistencies between subjects reported for our top match, uniqueness and bilaterality analyses can be summarised by our key conclusion - we do not always find a 1:1 mapping relationship between components and visual areas, a single component can match to many visual areas and many components can match to a single visual area.

The network-like components (i.e., components that span several visual areas) derived from our data are not akin to the large resting-state networks that are reliably identified across subjects when performing ICA on resting state data (S. M. Smith et al., 2009). In this thesis, I perform ICA at a higher dimensionality than Smith et al., therefore, as expected, larger networks fragment into smaller, more local networks. However, there are inconsistencies between subjects with respect to which regions are included in these network-like components. Moreover, some subjects also produce smaller, more discrete component representations of the regions recruited in the network. Taken together, our results could (i) be indicative of the brain dynamically switching between different states during movie watching, as reported by Meer et al. (2020), or (ii) reflect the correlations in the visual input. I discuss these below.

Meer and colleagues (2020) investigated the dynamic fluctuations in the BOLD signal of 14 canonical brain networks (BNs) during movie watching. It was found that the fluctuations of the BNs could be decomposed into 10 discrete brain states, each of which recur over time. Each brain state was characterised by the temporal dynamics of the BNs which systematically fluctuate over time. These brain states can, and do, periodically alternate during movie watching. Of interest to us are brain states four and seven as they involve the primary and high visual networks. Brain state four is characterised

by the fMRI signal preferentially loading on the high visual network only. Brain state seven displays high fMRI signal in the primary visual network, while all other networks, including high visual, show low signal loading. As there is clear variability across brains as to which visual areas are included in the larger network-like components, we could infer the components reflect the brain transitioning between different functional states.

The inclusion of several areas in a single component could in fact reflect correlations of visual features in the scene. Consider a segment of the movie which contains a minute of dialogue between characters with close-up shots of the faces. Activity in regions specialised for face perception (e.g., FFA and OFA) would have correlated activity, and activity in other regions are likely to be less stimulus driven. ICA is therefore likely to output a component corresponding to OFA and FFA, and if the pattern of activation is sufficiently different between the two may output a discreet component representing each area. Now, consider a scene in which two people are walking through a complex background. The stimulus is likely to simultaneously drive responses in many different visual areas, and therefore the presence of each visual feature, e.g., faces, bodies, edges etc, is likely to be correlated with each other. This synchronised activity may result in visual areas being "grouped" into a single component.

Rather than hand classifying components as anatomically corresponding across brains (as Bartels and Zeki did), we used an atlas to identify potential candidate components. Although components derived from ICA did spatially correspond with the atlas, the components are substantially larger than those obtained from traditional localisers. However, it is important to note that the absolute size of the component is specified by the threshold applied to the spatial map. When inspecting, and presenting figures of, the components it is at the researcher's discretion as to the level of thresholding to apply to the spatial maps (z-maps). Setting a higher threshold will remove more residual activation, making the peak of the component easier to identify. This can be at the expense of potentially misleading the reader as to the size/coverage of the component. Setting the threshold too low can be equally misleading in that it can make the spatial extent of the component difficult to interpret. The default threshold is usually z = 2-3 (Griffanti et al., 2017), therefore throughout this chapter, I

employed a fixed and consistent threshold of z = 2.75. However, even with a threshold at the higher end of the default, component matches in our data are not entirely analogous to previous regions identified with traditional approaches.

For example, Figure 48A shows TO1/MT and TO2/MST as identified by Huk et al. (2002) using a traditional motion localiser task. TO1/MT and TO2/MST are represented by small patches of cortex located in the posterior and anterior banks of the inferior temporal sulcus (ITS) respectively. In contrast, Figure 48B shows a component derived from our data that correlated with both the TO1/MT and TO2/MST atlas regions at r > 0.5. The spatial map shows good spatial correspondence between the peak of the component and Huk et al.'s example, but with much more widespread activation along the lateral surface of the brain.

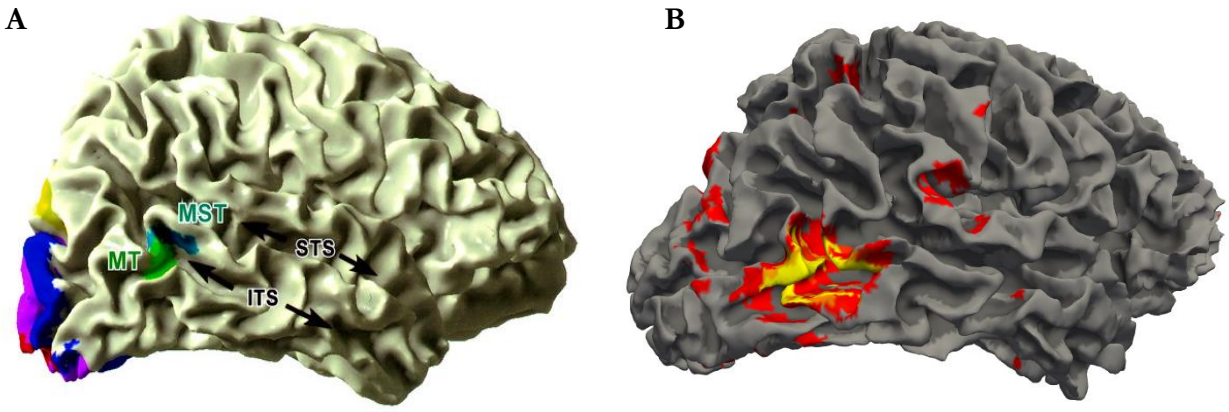

Figure 48. (A) Adapted from Huk et al. (2002). The right hemisphere surface showing the location of MT (green) and MST (blue) as identified with a traditional motion localiser task. Both regions occupy a small patch of the cortex. (B) a component derived from our movie watching data using ICA. The peak of the component (yellow vertices) shows good spatial correspondence with the location of MT and MST in (B), but with much more widespread activation along the lateral surface of the brain.

Thus far, I have discussed the results of analyses focused on the spatial domain, finding inconsistencies between brains and visual areas. The results from our temporal analysis support the conclusions of its spatial counterpart. We find that on average the ATCs of anatomically corresponding components, e.g., the top matched components for a given visual area, are less correlated across brains than the ATCs of the atlas regions.

An explanation for the lower correlations seen for component ATCs comes from the observed variability between subjects in the spatial correlation analysis. For each visual area, subjects show an inconsistent combination of non-unique matches, below threshold matches, and large spatial maps

covering adjacent regions. This in turn leads to equally inconsistent peaks of the top matched component (i.e., the vertices that contributed most to the component) across brains. Given we extract the mean ATC from a predetermined number of these peak vertices (we matched this number to the 10% of vertices from the probabilistic atlas), the ATCs are sometimes extracted from anatomically non-corresponding vertices. This variability in the location of the peak vertices could explain why the ATCs of components are less correlated across brains than atlas ATCs. We did, however, find the ATCs of hMT+ components to be significantly more consistent across brains when extracted from the components than from the atlas.

A potential explanation as to why components for TO1/TO2 are consistently found in our data, and for TO1/MT and TO2/MST to regularly map to a single component is that the visual motion cues are decorrelated from other cues and correlated with each other. For example, consider edges: edges are identified by differences in luminance, colour, specularity, texture, disparity, and feature boundaries (e.g., where a face feature ends). Therefore, as objects appear and disappear in the scene there will be correlated activity across a series of visual areas. The one cue that is not correlated with others is visual motion. The movement of the camera is independent of what is happening in the scene. But as the camera moves, various forms of motion cues are created in synchrony – from large field optic flow to motion defined boundaries.

In summary, the results of both our spatial and temporal analyses suggest performing ICA on movie watching data does not produce clean and consistent representations of the discrete regions of the visual system across brains. However, the approach could still have applications for case studies of patients with cortical vision loss. In particular, the bilateral nature of most components provides opportunities to identify unilateral components when tissue is spared (so indicating a functional disconnection), and bilateral components when areas upstream of an area are damaged (so indicating the involvement of non-V1 pathways and potential substrates for residual visual function).

Chapter 4. Convergent evidence for lateralisation of function in hMT+ in some individuals

4.1 Introduction

Strong et al. (2019) reported an "enhanced" role for right hemisphere hMT+ (human motion complex) in the processing of low-level motion stimuli in both the contra- and ipsi-lateral visual hemifields, indicating the region exhibits functional lateralisation. In the following chapter I investigate whether lateralised function of two hMT+ sub-regions, TO1/MT and TO2/MST, is evident when a visually rich broadband stimulus is used, here a movie. Such a stimulus more closely resembles our natural visual input than an aperture of moving dots and as such offers a more ecological approach to addressing this question. In the course of this work, I develop a series of tools that can be used to investigate lateralisation in patients with cortical vision loss.

In the following sections I present some background to the phenomenon of functional lateralisation and provide a brief overview of what is known about visual motion processing. Finally, I discuss evidence to suggest that visual motion perception is lateralised.

4.1.1 Functional lateralisation

Functional lateralisation refers to the hemispheric asymmetries in the function of specialised brain regions. There are three well known examples of functional lateralisation: language processing, face perception and visuospatial attention.

Language processing has been shown to be lateralised to the left hemisphere in 94% of normal subjects, with 6% showing bilateral language representation (Springer et al., 1999). During a semantic monitoring task in which subjects respond to spoken words, strong activation has been observed in the left hemisphere but not the right, indicating a hemispheric asymmetry (Frost et al., 1999).

Face perception is somewhat lateralised to the right hemisphere also. McCarthy and colleagues (1997) reported a small region in the right lateral fusiform gyrus which selectively responded to faces. Similarly, faces have been shown to elicit greater activation in the right hemisphere fusiform than the left (Kanwisher & Yovel, 2006).

Another particularly interesting area with respect to lateralisation is the posterior parietal cortex (PPC). The PPC, located between the somatosensory and visual cortices, is considered an important neural substrate of visuospatial function (Kravitz et al., 2011). Evidence from patients with damage to the PPC points to a role in the encoding of visual space, both for perception and action. Hemispatial neglect and extinction are two common disorders following damage to the PPC. Both disorders provide evidence of the right hemisphere playing a dominant role in visuospatial attention. I discuss this below.

Hemispatial neglect refers to an inability to perceive and respond to stimuli in the contra-lesional side of space (Heilman et al., 2000). Patients display a bias for the ipsi-lesional side of space at the expense of the contra-lesional side. Hemispatial neglect is often more severe and prolonged in patients with right hemisphere lesions (Ringman et al., 2004).

Extinction is a visual attention disorder in which patients fail to detect contra-lesional stimuli, but only when a similar stimulus is presented simultaneously to the ipsi-lesional side of space (Karnath & Rorden, 2012). Although previously considered a "milder" form of neglect due to their similar symptomology, Vossel et al. (2011) provided evidence of the two deficits being dissociable.

Both spatial neglect and extinction are more prevalent and severe following right hemisphere damage, and therefore, it is believed that the right hemisphere displays attentional "dominance" (Becker & Karnath, 2007).

Not only is there evidence that the PPC shows a lateralisation of function but, of particular interest here, the hemispatial theory (Heilman & van den Abell, 1980) explains the left/right asymmetry by proposing that right hemisphere PPC is involved in directing attention to both the contra- and ipsi-lateral hemispaces, whereas left PPC only directs attention to the contra-lateral (right) hemispace (Heilman & van den Abell, 1980; Mesulam, 1981). Therefore, in the case of damage to left PPC, the right hemisphere is still able to direct attention to its ipsi-lateral hemispace, and thus neglect is absent or less severe in these patients. In contrast, damage to right PPC results in inattention to the left visual space, but crucially the left hemisphere is unable to compensate for this since it only directs attention to the contra-lateral visual space. Consequently, the prevalence and severity of neglect is greater in

patients with right PPC damage. The proposal that right PPC covers both hemifields whereas left covers only the contralateral field foreshadows the recent proposal of lateralisation of function in hMT+.

Research into the functional lateralisation of motion perception is limited. It has been reported that trajectory discrimination is both lateralised to the right hemisphere and independent of hand preference (Boulinguez et al., 2003): participants were instructed to respond, as fast as possible, with a button press if they believed the trajectory of an on-screen dot, presented in either the left or right visual hemifield, was to end in a target location in the centre of the screen. Lateralisation was inferred from shorter reaction times to stimuli presented to the left compared to the right visual field. This raises a question of whether there are asymmetries between hemispheres for other forms of motion perception. In section 4.1.5 I discuss evidence that indicates the human motion complex (hMT+) exhibits lateralised function.

4.1.2 Motion processing – non-human primates

Much of our understanding of the neural substrates underpinning human motion processing came from research into the anatomy of the monkey brain. Similar distinctions (e.g., retinotopic organisation, motion sensitivity and direction selectivity) were then identified in the human brain. Therefore, before discussing evidence to suggest the hMT+ exhibits functional lateralisation, it is important to first understand the region's functional specialisation and why it is believed to be homologous to the motion areas, MT and MST, of the monkey brain.

The motion sensitive area V5, also referred to as middle temporal area (MT), was first identified in the macaque brain (Dubner & Zeki, 1971). The first visual region outside of V1 to have a specialised function assigned to it by researchers, V5/MT is characterised by a high proportion of directionally selective (DS) neurons: neurons that respond selectively to stimulus motion in a particular direction (Dubner & Zeki, 1971; Zeki, 1974). These DS neurons are reported to have larger receptive fields compared to those located within V1 (Albright, 1984), and have also been shown to be tuned for the speed of visual motion (Perrone & Thiele, 2001; Priebe et al., 2003). Some functional properties of

hMT+ (speed and direction selectivity) are said to be inherited from V1 (Born & Bradley, 2005). V1 neurons projecting to V5/MT have also been shown to be directionally selective (Movshon & Newsome, 1996) and tuned for speed (Priebe et al., 2006).

Cortical regions, V1 in particular, dominate the input to V5/MT (Felleman & van Essen, 1991). However, some V5/MT neurons have been shown to remain visually responsive following inactivation of V1 (Rodman et al., 1989, 1990). Sincich el al. (2004) demonstrated a pathway in the macaque brain in which neurons project directly from the lateral geniculate nucleus (LGN) to V5/MT. Such a pathway was hypothesised to facilitate the residual activation in V5/MT in the absence of V1 input. The pathway is a fraction of the size of the primary visual pathway. The number of neurons involved in the direct LGN-V5/MT pathway is estimated to be 10% of the number of neurons projecting from V1 to MT, and approximately 1% of the total number of LGN relay neurons (Sincich et al., 2004).

Neurons in V5/MT directly project to an adjacent area, the medial superior temporal (MST) area, which itself contains a large proportion of DS neurons (Lagae et al., 1994). The DS properties of MST neurons are more complex than MT, in that they also respond to expansion/contraction and rotation. Compared to MT, MST neurons have larger receptive fields which extend by up to 40° into the ipsi-lateral visual field (Raiguel et al., 1997).

4.1.3 Motion processing – humans

In the human brain, the first motion area to be identified was area V5 (Zeki et al., 1991). Using PET, V5 was located in the temporo-parieto-occipital junction and showed stronger activation in response to moving than stationary stimuli. The area was believed to be homologous to the macaque V5/MT and MST regions and thus is often referred to as the human middle temporal complex (hMT+). FMRI studies have confirmed the earlier PET findings, often reporting bilateral activation in hMT+ following visual motion stimulation (Tootell et al., 1995; Watson et al., 1993).

The hMT+ complex responds to a broad range of visual motion stimuli such as incoherent flicker and moving gratings or dots (Tootell et al., 1995). As the name suggests, hMT+ is a complex and fMRI research has focused on parcellating its subregions (in line with the monkey

neurophysiology), of which there are believed to be at least two (Dukelow et al., 2001; Huk et al., 2002; Morrone et al., 2000). Huk and colleagues (2002) divided hMT+ into two subregions, one located in the posterior of the complex which responded to standard wedge retinotopic stimuli, (homologous with the macaque area MT), and one in the anterior which was responsive to motion stimuli presented to ipsi-lateral visual field, (homologous with macaque area MST). Note, in the macaque brain, MST is composed of two subregions, MSTd and MSTl (Eifuku & Wurtz, 1998). However, Huk et al. (2002) reported that with their approach hMT+ was divided into only two regions. Morrone et al. (2000) reported functional distinctions between hMT+ subregions based on motion-type specificity, finding an area responsive to radial and circular flow motion (only when the direction of motion periodically changed), and another responsive to translational motion. Subsequent work has found a similar division based upon sensitivity to different types of optic flow. MT responds equally to random dot motion and coherent global flow patterns, whilst MST responds strongest to complex flow patterns containing rotation, expansion, and contraction components (A. T. Smith et al., 2006). Activation in MST is reduced when the flow pattern contains only one of these components and is reduced further with 'low-level' translation motion stimuli (A. T. Smith et al., 2006).

The retinotopic organisation of, and population receptive field (pRF) sizes in, hMT+ can be used to distinguish its subregions (Amano et al., 2009). Contra-lateral stimulation elicits activation in hMT+ as a whole. However, ipsi-lateral stimuli elicit strong activation in MST, but only weak activation in MT. As is the case with non-human primates, it is suggested that responses to ipsi-lateral stimulation are due to the receptive fields of MST neurons extending into the ipsi-lateral visual field (Amano et al., 2009). The weak response to ipsi-lateral stimulation in MT can be explained by its neurons having relatively smaller receptive fields which do not encroach into the ipsi-lateral visual field to the same extent MST neurons do (Amano et al., 2009).

It is clear that the two subregions of hMT+: MT and MST, demonstrate differences in their functional properties and receptive field coverage.

The nomenclature used to refer to the subdivisions of hMT+ varies across studies: some adhere to MT and MST as used in the early monkey research (e.g., Huk et al. 2002), others make

distinctions between the sub-regions based on the differing retinotopic and pRF properties, referring to MT and MST (owing to their location in the temporal-occipital cortex) as TO1 and TO2 respectively (e.g., Amano et al. 2009). Here, we adopt the hybrid term TO1/MT and TO2/MST used by Strong et al. (2017).

4.1.4 Anatomical variability of hMT+

In humans, hMT+ shows high anatomical variability, both between subjects and between hemispheres of the same brain (Huang et al., 2019). Wilms et al. (2005) found that functionally defined hMT+ (fMRI paradigm contrasting moving and stationary dots) in the left hemisphere was around 65% of the size of right hemisphere hMT+.

Not only does hMT+ differ in size between hemispheres, but evidence from retinotopic mapping (Wang et al., 2015) and cytoarchitectural studies of post-mortem brains (Malikovic et al., 2007) indicate there is considerable interindividual variability in the anatomical location of hMT+.

The inter-subject variability of the anatomical location of hMT+ is closely related to the variability in individual brain topography. fMRI studies locate hMT+ on the lateral surface of the anterior occipital lobe (e.g., Huk et al., 2002), an area of the cortex with considerable inter-subject variation in sulcal pattern (Thompson et al., 1996). Using a functional localiser, Dumoulin et al. (2000) found that for the majority of brains (~95%), hMT+ lies within a sulcus. However, the exact sulcus in which it was located varied across brains.

The anatomical location of hMT+ is frequently correlated with the junction between the inferior temporal sulcus (ITS) and the anterior occipital sulcus (AOS; also referred to as the ascending limb of the inferior temporal sulcus; Dumoulin et al., 2000). However, substantial variability has been reported in both the exact anatomical location of hMT+ relative to these two sulci and the morphology of the sulci themselves (Dumoulin et al., 2000; Malikovic et al., 2007).

The anatomical location of hMT+ can also vary between hemispheres. Anyone who has visually inspected the surface reconstruction of a subject's anatomical image will have seen how the folding pattern of the cortex varies across hemispheres. The AOS, which as discussed above is

associated with the location of hMT+, is one such sulcus that has been shown to markedly differ between hemispheres, whereas the calcarine sulcus, for example, is more consistent (van Essen & Drury, 1997). This variation in the location and extent of sulci correlated with hMT+, means the anatomical location of the complex can vary across hemispheres of the same brain.

As discussed in section 1.3, there are several ways to map the visual brain (e.g., cytoarchitecture, functional localisation, retinotopy etc.), and different techniques use different "markers" to define visual areas (e.g., structure or function). The anatomical location of hMT+ has been shown to vary depending on which marker is used to define it. For example, Figure 49 shows the anatomical location of probabilistic hMT+ as defined by retinotopy (TO1/TO2 sub-divisions) and cytoarchitecture (also referred to as hOc5; Malikovic et al., 2007), relative to the AOC and ITS. When defined with retinotopy, hMT+ occupies an area of the cortex superior to the AOS-v and inferior to the AOS-d , whereas when defined with cytoarchitecture, the complex is mostly confined to the AOS-v.

Indeed, Large et al. (2016) used high-resolution in vivo imaging to examine the relationship between structural (myelination) and functional (fMRI responses to visual motion) markers of hMT+ in a single group of subjects. When myelination maps were averaged across brains, there was clear spatial association with the functional hMT+ map. However, at the individual level there was greater variation in the alignment of the functional and structural maps (i.e., the two markers vary in their spatial congruence), suggesting group maps may not be a good prediction of the individual brain.

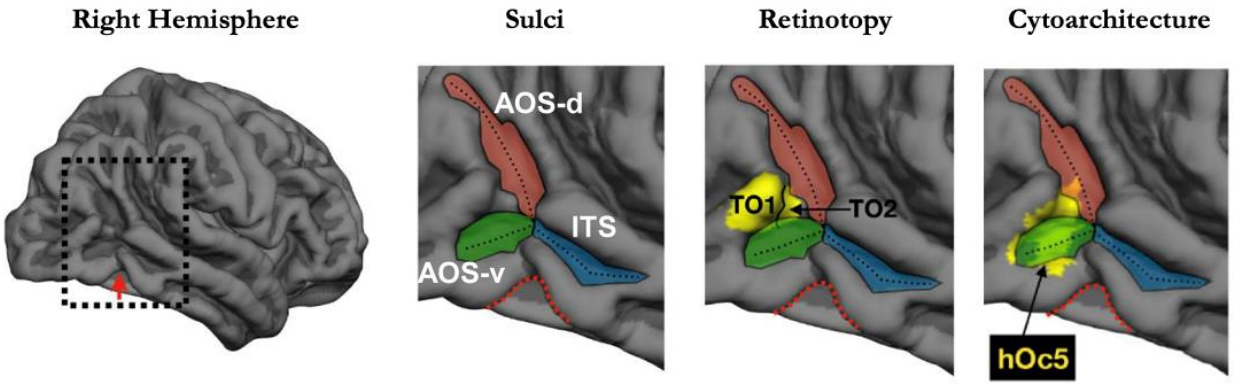

Figure 49. Adapted from Huang et al., (2019). Left: the pial surface of the right hemisphere fsaverage brain (left). Red arrow indicates the anatomical location of TO1/MT. Middle left: zoomed portion of the lateral occipital temporal cortex (as indicated by the dotted square in the left image. The three sulci of interest are defined by different colours: red = dorsal portion of the anterior occipital sulcus (AOS-d), green = ventral portion of the AOS (AOS-v), blue = inferotemporal sulcus (ITS). Middle right: the anatomical location of the Wang et al. (2015) retinotopically defined TO1/MT and TO2/MST ROIs relative to the three sulci. Right: the anatomical location of hMT+ when defined with cytoarchitecture (hOc5; Malikovic et al., 2007).

Individual differences in brain architecture clearly impact the anatomical location of hMT+, both within the same brain (hemispheric differences in gyrification) and across subjects. Similarly, the criteria used to identify hMT+ (e.g., structural markers or functional localisation) can also influence the area definition. However, Huang et al. (2019) examined the spatial consistency of functionally defined hMT+ and produced a probabilistic map of the complex based on fMRI data from over 450 brains. Following this, they aimed to investigate the correspondence between their functionally defined hMT+ map and maps defined by other techniques. Of the maps that were considered, the Wang et al. (2015) atlas was used to represent retinotopically defined hMT+ (TO1/MT and TO2/MST maps were combined). This comparison is particularly relevant here since we adopted the Wang et al. (2015) atlas to define the visual areas in this thesis. Huang et al. found the size and the peak location of the functionally defined hMT+ atlas varied across brains, but converged (i.e., highest probability) at the ventral AOS (ASO-v). Most importantly, the functionally and retinotopically defined maps showed good spatial agreement. This finding not only indicates good anatomical correspondence between approaches but enabled us to be satisfied that the area definitions, in particular that of hMT+, used in this thesis were appropriate, if not a perfect 'ground truth'.

4.1.5 Lateralised function of hMT+

Thus far, I have provided a brief overview of the evidence of functional lateralisation for some capacities, including motion perception in the form of trajectory discrimination, and the motion complex (hMT+) containing (at least) two subregions each with demonstrated differences in their functional properties and receptive field coverage. In the following section, I discuss the findings of two groups that, through the application of TMS, sought to investigate the causal relationship between hMT+ and its function (Strong et al., 2019; Thakral & Slotnick, 2011). Both groups found functional asymmetries in hMT+ across hemispheres, indicating the region exhibits lateralised function.

Transient perceptual deficits can be induced in individuals using TMS, with the principal aim of temporarily disrupting a region believed to be necessary for a certain cognitive process. If disruption leads to an impairment of the corresponding behaviour, one may conclude this region is necessary to perform the process.

Thakral and Slotnick (2011) sought to specifically target TO1/MT with TMS while presenting stimuli to both the ipsi- and contra-lateral visual fields. First, hMT+ was identified using an fMRI motion localiser in which subjects viewed a full field dot stimulus that alternated between moving and stationary. TO1/MT was identified as the activation elicited by the stimulus that was constrained to the posterior bank of the ascending limb of the inferior temporal sulcus. Next, TMS was applied to TO1/MT during a spatial attention task. Importantly, the TMS coil was positioned in such a way as to specifically avoid stimulation of the anterior bank of the ascending limb of the inferior temporal sulcus, the known location of TO2/MST.

The spatial attention task was composed of two squares of dots which moved with 70% coherence, one presented in each hemifield. Participants shifted their attention to the opposite hemifield stimulus if dots within the currently attended hemifield increased in speed and responded with a 'shift' button press. If dots in the attended hemifield decreased in speed, participants responded with a 'detect' button press but did not shift their attention to the opposite hemifield.

Thakral and Slotnick found that TMS applied to left TO1/MT resulted in significantly lower target detection of contra-lateral stimuli compared to ipsi-lateral stimuli. Furthermore, detection of

both contra- and ipsi-lateral stimuli was significantly lower when TMS was applied compared to when it was not. Suggesting disruption of left TO1/MT results in a deficit in motion perception for both ipsi- and contra-lateral stimuli which is more pronounced for contra-lateral stimuli.

On the other hand, when TMS was applied to right TO1/MT there was no significant difference between the detection of ipsi- and contra-lateral targets, suggesting that disruption of right TO1/MT produced equal impairments in the detection of both contra- and ipsi-lateral stimuli. These findings not only highlight the importance of TO1/MT in motion processing of contra-lateral stimuli but also indicate a potential functional difference between left and right TO1/MT in the processing of the ipsi-lateral visual field. This difference is demonstrated by the fact right TO1/MT showed comparable deficits in the processing of ipsi- and contra-lateral stimuli, whereas left TO1/MT showed greater deficits for contra-lateral than ipsi-lateral stimuli. Therefore, these findings may be interpreted as evidence of lateralisation of TO1/MT, specifically implicating the right hemisphere in the processing of motion in both contra- and ipsi-lateral visual fields.

Strong at al. (2019) rebutted Thakral and Slotnick's conclusions, suggesting that as the hMT+ complex was identified as a whole, and TMS was subsequently applied to TO1/MT based on the anatomical location of the fMRI activation, TO2/MST may have in fact been stimulated unequally across hemisphere. Owing to TO2/MST's larger receptive fields, such stimulation could explain the ipsi-lateral deficits. Following this, Strong et al. aimed to investigate the functional asymmetries of TO1/MT and TO2/MST across hemispheres, specifically targeting the sub-regions independently.

Strong et al. (2019) used a functional localiser to identify TO1/MT and TO2/MST in their subjects. TMS was then applied independently to the subdivisions whilst subjects performed a motion discrimination task. A circular aperture of moving dots was presented 15° to the left or right of a central fixation target. A predetermined number of the dots moved coherently up or down, while the remaining dots moved randomly. Deficits in the motion perception of stimuli located in both the ipsi- and contra-lateral hemifield were induced when TMS was applied to right hemisphere TO1/MT and TO2/MST. In contrast, when TMS was applied to these regions in the left hemisphere, deficits in motion perception were only evident for contra-lateral hemifield stimuli. Consistent with early work by

Thakral and Slotnick (2011), these findings suggest a degree of right hemisphere lateralisation reflected by an increased responsibility for processing motion in both the contra- and ipsi-lateral visual fields.

Strong et al.'s right hemisphere lateralisation hypothesis complements the well-known literature describing a right hemisphere dominance in visuospatial attention (see section 4.1.1). In both cases right hemisphere areas (hMT+ and PPC) have been implicated in the processing of both sides of visual space, whereas the left hemisphere counterparts are concerned with the contra-lateral space only.

4.1.6 Implications of functional lateralisation

Investigating whether hMT+ is lateralised not only develops our understanding of the brain's functional organisation, but it also has clinical relevance. Consider the primary visual pathway in which V1 is compromised. This in turn means the primary inputs to hMT+ are also extinguished. Therefore, V1 damage in the left hemisphere leads to a loss of conscious vision in the right visual field. For all intents and purposes, the patient is perceptually blind in the right visual field. However, some patients show an intact sensitivity to motion stimuli presented to the blind hemifield

Recently, it has been suggested that intact motion perception could be facilitated by a small number of neurons projecting directly from the LGN to hMT+ (Ajina, Pestilli, et al., 2015; Ajina & Bridge, 2018). If remained intact following post-LGN damage, this non-occipital route could be sufficient for preserving motion perception. On the other hand, the studies discussed above raise the question of whether the enhanced role of right hemisphere hMT+ could facilitate intact motion processing in patients with left hemisphere post-LGN damage. Specifically, could the bilateral visual field coverage of the right hemisphere (contra-lesional) hMT+ facilitate the processing of the ipsi-lateral visual field?

In this chapter I investigate functional lateralisation of TO1/MT and TO2/MST by performing (i) cross-hemisphere timecourse correlations, (ii) between-brains timecourse correlations, and (iii) ICA on movie watching data.

For (i) and (ii), I treat V1 as a baseline measure. V1 is as close as one can get to examining the correlation in the visual input, and therefore, any differences between hemispheres would likely be

driven by differences in the visual input across hemifields, rather than differences in function across hemispheres.

Results from the cross-hemispheres analyses show that correlations between left/right TO2/MST are lower than between left/right TO1/MT and left/right V1 (TO1/MT and V1 correlations being comparable). The lower TO2/MST correlations can be taken as weak support for the lateralisation hypothesis.

Between-brains correlations point to individual differences. The between-brains correlations are lower for the TO1/MT and TO2/MST timecourses in the right hemisphere than the left hemisphere, with the difference between hemispheres largest for TO2/MST. This finding is compatible with individual differences in the function of TO2/MST, lateralisation of function in some brains and not others.

ICA supports the same conclusion, with evidence of lateralisation, an unequal number of components including the left and right hemispheres, in ~25% of brains.

Taken together, the analyses do not support the Strong et al. (2019) version of the lateralisation hypothesis. They do, however, provide support for a weaker version, lateralisation of TO2/MST function in a subset of brains, suggesting that it might be fruitful to look for individual differences in the tasks performed by Strong et al. (2019) and Thakral and Slotnick (2011).

4.2 General methods

For all analyses in this thesis the same dataset was used, and therefore the acquisition and stimulus methods were identical to those reported before (see section 2.1). All 17 subjects watched the same 20-minute clip taken from the beginning of a Hollywood action film (James Bond, Skyfall) whilst in an MR scanner.

4.2.1 Data analysis and pre-processing

Pre-processing of functional data was carried out in FEAT (fMRI Expert Analysis Tool) Version 6.0, part of the FSL toolbox (Jenkinson et al., 2012; S. M. Smith et al., 2004; Woolrich et al., 2009) with the default parameters described below.

Low frequency artefacts, such as low frequency temporal drifts, were removed from the dataset using a high-pass temporal filter set to 100 Hz. The EPI data was brain extracted using BET. MCFLIRT, part of the FSL package, was used to realign the data to compensate for changes in head position and/or orientation through the application of rigid-body transformations (rotational and translational corrections only). No spatial smoothing was applied at this stage.

Structural data were pre-processed using FreeSurfer version 7.1.1. The *"recon-all"* command in FreeSurfer was used for its auto-segmentation of subcortical white matter and deep grey matter structures (Fischl et al., 2004) from which cortical surface reconstructions can be generated. This surface-based stream is described in detail elsewhere (see; Dale et al., 1999; Fischl et al., 1999)

The processed EPI data were then registered to the subject's high-resolution anatomical image using FreeSurfer's *"bbregister"* command, with the *"–feat"* flag added. Functional data was then resampled to the fsaverage_sym surface, via the fsaverage surface (see section 4.2.2).

4.2.2 Surface generation and resampling

In order to compute the cross-hemisphere correlations, we needed a surface in which every vertex in one hemisphere has a match in the other hemisphere. The folding pattern of the fsaverage[16] surface is not symmetrical across hemispheres, and therefore we used fsaverage_sym. fsaverage_sym is a left-right symmetrical pseudo-hemisphere. The right hemisphere is a mirror-flipped version of the left hemisphere, with both hemisphere surfaces using a corresponding vertex space. Data was transformed between surfaces/hemispheres using FreeSurfer's *"surfreg/xhemi"* commands.

The sequence of transformation and processing was as follows:

1. FreeSurfer's *"mri_vol2surf"* function was used to resample the pre-processed EPI data to the fsaverage surface with the projection fraction set to 0.5 (middle of the cortical surface).
2. Data are resampled from the fsaverage surface to the fsaverage_sym surface using FreeSurfer's *"mri_surf2surf"* command.

[16] The fsaverage surface is an average of 40 brains. fsaverage is the default average surface used in FreeSurfer.

3. To improve the signal to noise ratio, data were spatially smoothed with a conventional 5 mm FWHM kernel (Blazejewska et al., 2019) during sampling to fsaverage_sym[17].

4.2.3 Data analysis summary

In this chapter I perform correlations between timecourses (both within and across hemispheres) and use the relative correlation values to infer whether an area appears to exhibit lateralised function.

I first calculated correlations between the mean timecourses of corresponding visual areas (ROIs derived from an atlas, details in section 4.2.4 below) in opposite hemispheres. The ROI-based mean timecourse analysis could confound the size of the area with the degree of the cross-hemisphere correlation; visual areas are of different sizes meaning the mean timecourse is calculated over a different number of vertices, therefore effectively smoothing to differing degrees across areas. As smoothing can increase the signal to noise ratio (SNR), this approach could alter the correlations. Reciprocally, if an area consists of functionally distinct sub-regions, e.g., MST contains MSTd and MSTl, then taking the mean timecourse over the entire ROI could reduce the cross-hemisphere correlation. Therefore, to control for this, I performed a second analysis in which I calculated correlations between corresponding vertices in the two hemispheres. The key difference between these two analyses is that in the ROI-based analysis I perform correlations between the mean timecourses of two ROIs (correlation of the mean). Whereas in the vertex-based analysis I perform timecourse correlations on a per-vertex basis and then take the mean of these correlations for each ROI (mean of the correlation). To allow direct comparisons between the ROI- and vertex-based analyses, both were performed on data resampled to the symmetrical average surface[18].

---

[17] I investigated how sensitive our results were to the choice of smoothing pipeline, the details of which are described in Appendix XII.

[18] An alternative to this approach would be to perform these analyses on the subject's native surface. However, such an approach would not allow direct vertex-to-vertex correlations to be performed, as the two hemispheres of the native surface are not symmetrical. To confirm the conclusions drawn were robust I also performed the mean timecourse analyses on data resampled to the native surface, the details of which can be found in Appendix XI.

For ICA, data were pre-processed using the optimal pipeline we described in section 2.1.5.3 (in brief, data were pre-processed using fMRIPrep before being spatially smoothed in the volume). Next, I performed single-subject ICA on the native volume data. Component spatial maps were then resampled from the volume to the individual subject's surface. TO1/2 components were identified by finding the components that showed a spatial overlap, a $r > 0.3$ correlation, with the full probability maps from the probabilistic atlas of visual areas (see section 2.1.5 for details). I then examined the subset of components that had a correlation with TO1/MT and TO2/MST above a $r = 0.3$ threshold.

4.2.4 Visual area ROIs

To investigate the differences across hemispheres in the function of TO1/MT and TO2/MST, I use a set of visual area labels ("ROI's") to extract the mean timecourses and the mean cross-hemisphere correlation values. However, the use of the maximum probability map (MPM, see section 2.1.5) ROIs raised a potential confound to these analyses. The left hemisphere MPM of TO2/MST was substantially smaller than any other region and over 5× smaller than the right hemisphere equivalent. I sought to control for this by creating an ROI that was a more plausible representation of the region.

A MPM and FPM is produced for each hemisphere separately and the ROIs differ in size and shape between hemispheres. However, as our analyses are performed on a symmetrical surface, opting to use one hemisphere's set of ROIs over the other does not seem appropriate.

I created an ROI for each visual area that included only vertices that were common across the two hemispheres *and* were not included in any other visual area (I refer to these as "exclusive ROIs"). These ROIs were adapted from the topographical maps of the visual areas described in section 2.1.5 (Wang et al., 2015) and underwent a number of revisions which I discuss in detail below. It is important to note that each set of ROIs (MPM, FPM, intersection and exclusive) have their own limitations, but ultimately, I use the converging results of each stage to generate a set of ROIs that were plausible.

To ensure the results obtained from the analyses performed in this chapter were not dependent on the use of these exclusive ROIs, I checked the same conclusions were reached when I used alternative ROIs. The results are robust in that the pattern does not change irrespective of which ROIs were used.

4.2.4.1 Max probability maps (MPM)

First, I warped the MPM from the fsaverage surface to the symmetrical surface (Figure 50). Table 9 shows the size, in terms of the number of included vertices, of the MPM. As already noted, there are noticeable differences in the size of left and right hemisphere visual areas. Furthermore, left hemisphere TO2/MST (28 vertices) is over 5× smaller than its right hemisphere counterpart (148 vertices). Whereas the size TO1/MT is more consistent across hemispheres. The authors of the paper noted that the size of the TO2/MST MPM, as well as PHC1/2, IPS4, IPS5, SPL1, and FEF, showed a considerable deviation from the subject level mean area. They concluded that this deviation was due to the fact the MPM represents the maximal overlap between subjects and therefore if a region has a high level of variability in its anatomical location between subjects, the size of the MPM of this region will be reduced.

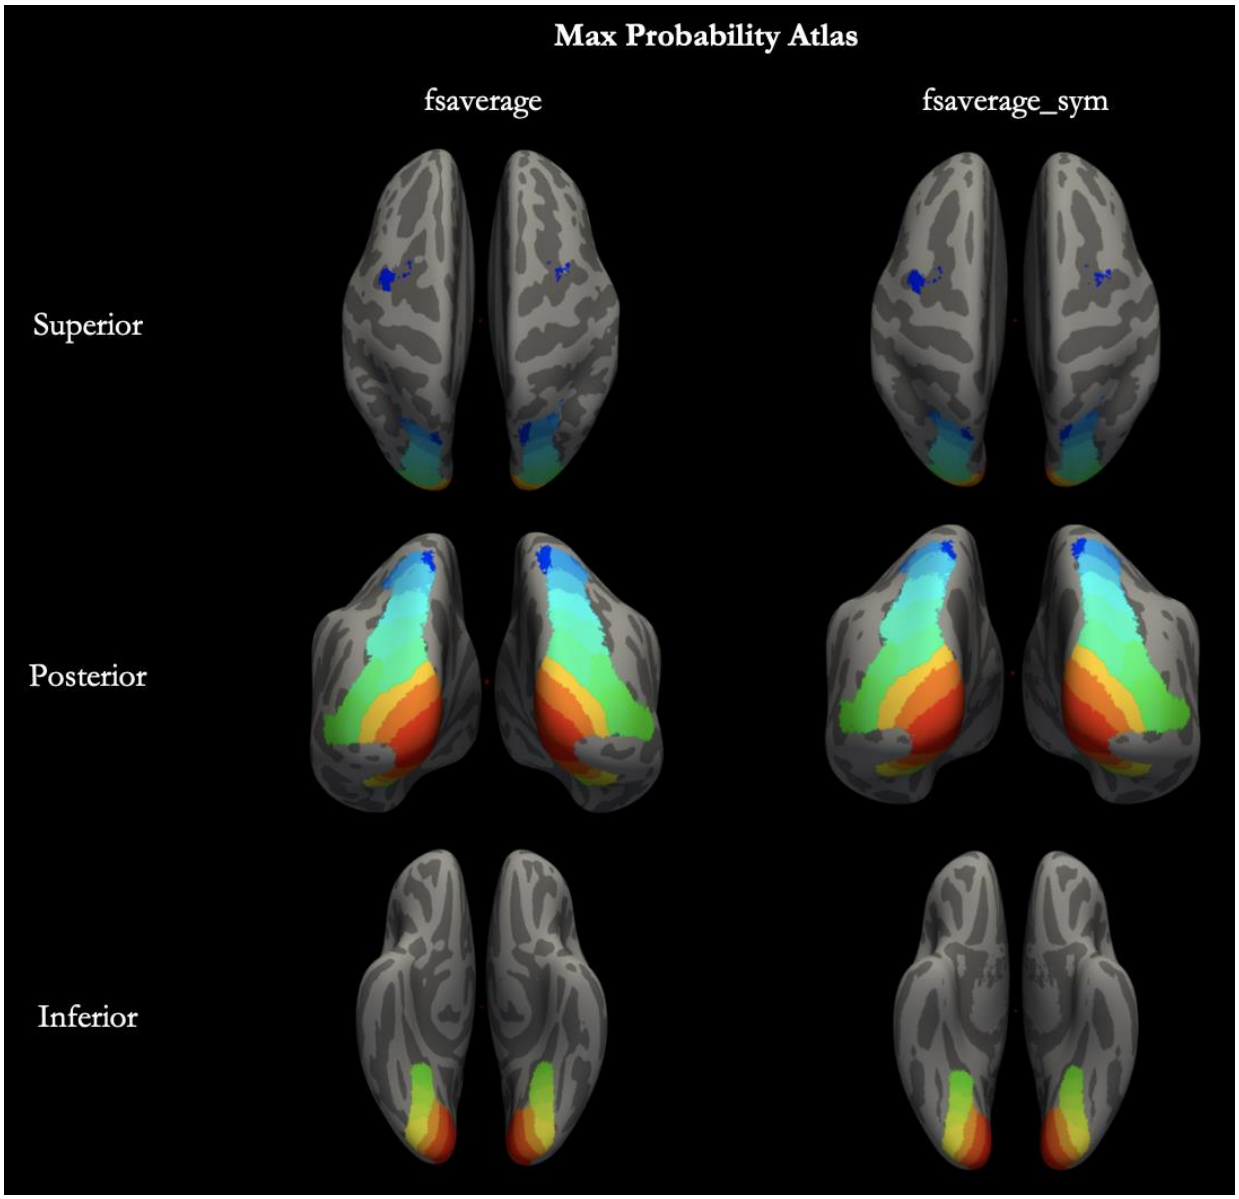

Figure 50. Max probability map taken from Wang et al. (2015) overlaid on the fsaverage (left) and fsaverage_sym (right) surfaces. Each visual area is represented by a different coloured ROI.

Table 9. Size (in terms of the number of included vertices) of the maximum and full probability maps taken from Wang et al. (2015) after sampling to the fsaverage_sym surface. Green cells indicate which hemisphere contained a higher number of vertices.

	Number of vertices in ROI			
	fsaverage_sym max probability atlas		fsaverage_sym full probability atlas (0.2 threshold)	
Visual Area	**Left Hemisphere**	**Right Hemisphere**	**Left Hemisphere**	**Right Hemisphere**
V1v	1052	993	1418	1339
V1d	1125	1101	1548	1549
V2v	926	868	1312	1205
V2d	743	861	1242	1210
V3v	593	609	955	1007
V3d	743	765	1202	1311
hV4	444	505	871	898
VO1	280	212	514	443
VO2	317	381	630	700
PHC1	342	139	777	635
PHC2	249	382	791	905
TO2/MST	28	148	683	604
TO1/MT	404	375	737	737
LO2	330	331	638	719
LO1	454	651	862	1143
V3B	426	510	769	918
V3A	688	841	1130	1272
IPS0	1185	935	2035	1991
IPS1	981	736	2008	1787
IPS2	965	621	1841	1396
IPS3	510	760	1154	1449
IPS4	123	155	1042	1097
IPS5	49	18	850	648
SPL1	255	304	1332	1536
FEF	349	108	2637	2868
All areas	13561	13309	28978	29367

4.2.4.2 Thresholded full probability maps (tFPM)

As outlined above, the MPM of left hemisphere TO2/MST includes just 28 vertices, while the right hemisphere equivalent includes 148. However, we wanted to utilise the full probability maps (FPM) to ensure the left hemisphere ROI of TO2/MST was more plausibly sized. Likewise, as is the nature of the MPM, some vertices which may be almost equally probable of belonging to two neighbouring regions will be allocated to the region with the highest probability value. This results in some visual area maps being unnecessary reduced in size. On the other hand, as is the nature of the FPM, a given visual area will include a vast number of vertices with varying probabilities of belonging to multiple visual areas. This results in visual area maps which are unrealistically large if left unthresholded. Therefore, we sought to generate a set of consistent maps which made use of the data in both the MPM and FPM.

The FPM was converted into individual surface overlays in MATLAB. We expected the MPM of V1v to be an accurate representation of the area, due to the low level of variability across brains reported by the authors of the maps (reflected by a peak probability of 100%). Probability values ranged from 0 (vertex was not assigned to the ROI in any brain) to 1 (vertex was assigned to the area in all brains that contributed to the ROI). Therefore, I used V1v as a starting point.

I found that setting the threshold of the FPM to 0.2 (indicating the vertex was present for at least 20% of the brains) gave an overlay that closely resembled the V1v MPM (Figure 51). I refer to these maps as thresholded full probability maps (tFPM). I then examined whether thresholding the FPM overlay for each region by the same level (0.2) resulted in plausible representations of the regions. This approach is effectively an atheoretical, empirical solution to generating visual maps; we take a known good map (e.g., V1v) to estimate the appropriate parameters and then apply them elsewhere.

Next, I inspected maps for other early visual areas that were also consistent across brains in the Wang et al. (2015) atlas. It again appeared that thresholding the FPM to 0.2 resulted in plausibly sized overlays which encompassed the region of interest (relative to the MPM represented by the black outlines in Figure 52), whilst allowing a small amount of dilation around the edges of the region.

Figure 51. V1v max probability map (black outline) and V1v full probability map with 0.2 threshold (red-yellow).

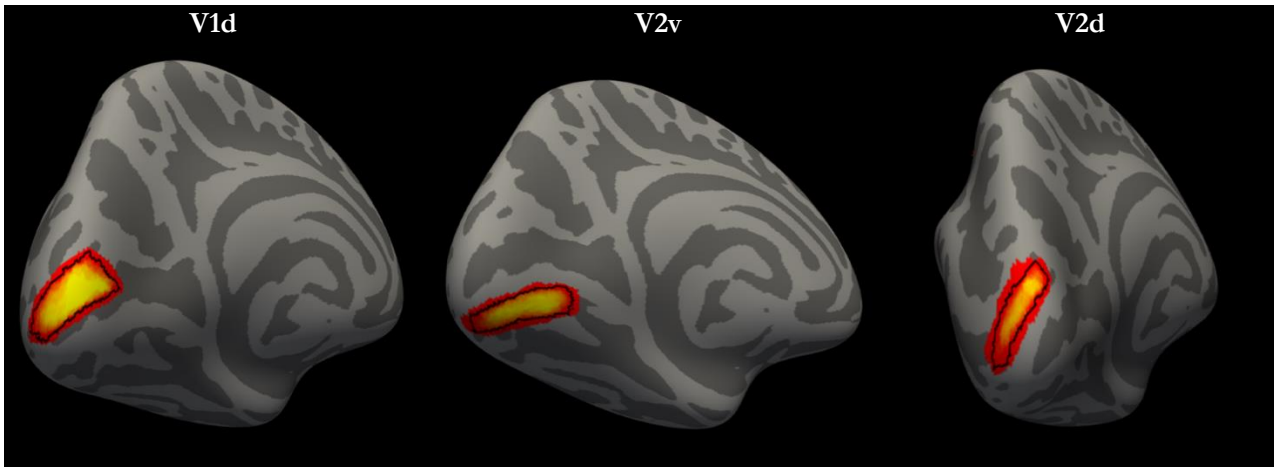

Figure 52. Full probability map overlays (red-yellow, threshold = 0.2) for V1d, V2v and V2d with the max probability map ROIs for the corresponding area (black).

Importantly, using the thresholded FPM approach produces a more plausible representation of TO2/MST (Figure 53, right) than the original MPM ROI (Figure 53, left).

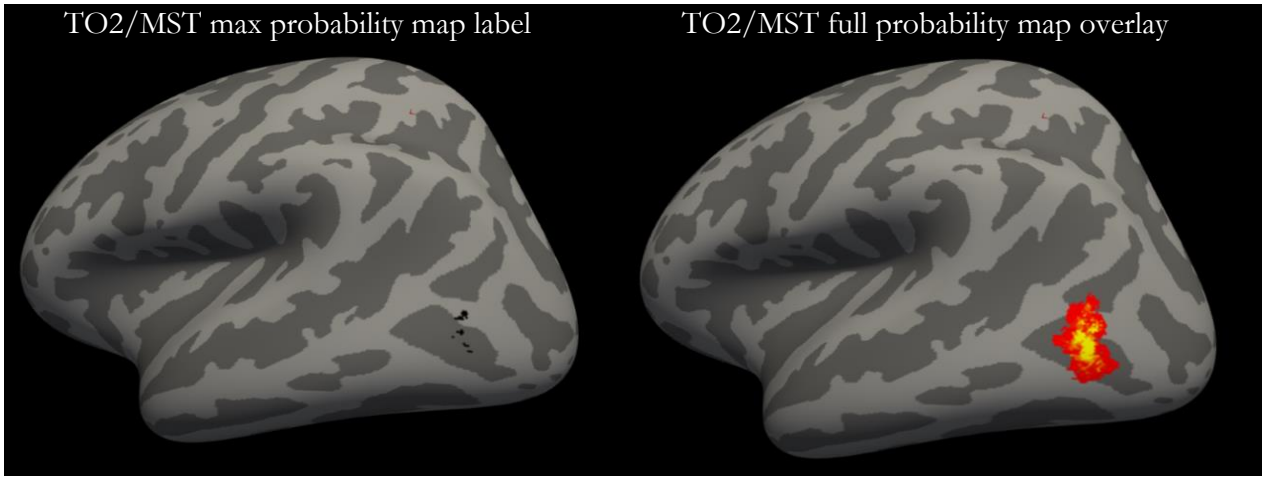

Figure 53. TO2/MST max probability map label (left) and full probability map overlay thresholded to 0.2 (right).

4.2.4.3 Left-right intersection ROIs

The above approach generated a separate set of ROIs for each hemisphere. These ROIs differ in size and shape across hemispheres. However, our analyses are performed on a symmetrical surface and therefore, opting to use one hemisphere's set of ROIs over the other does not seem appropriate. Therefore, we generated a set of intersection ROIs.

The intersection ROIs were generated by taking the left and right hemisphere maps for a given visual area and including only vertices which are common to both the left and right hemispheres. However, the intersection ROIs have a significant amount of overlap between neighbouring visual areas which may confound the subsequent mean correlation coefficient extraction (Figure 54).

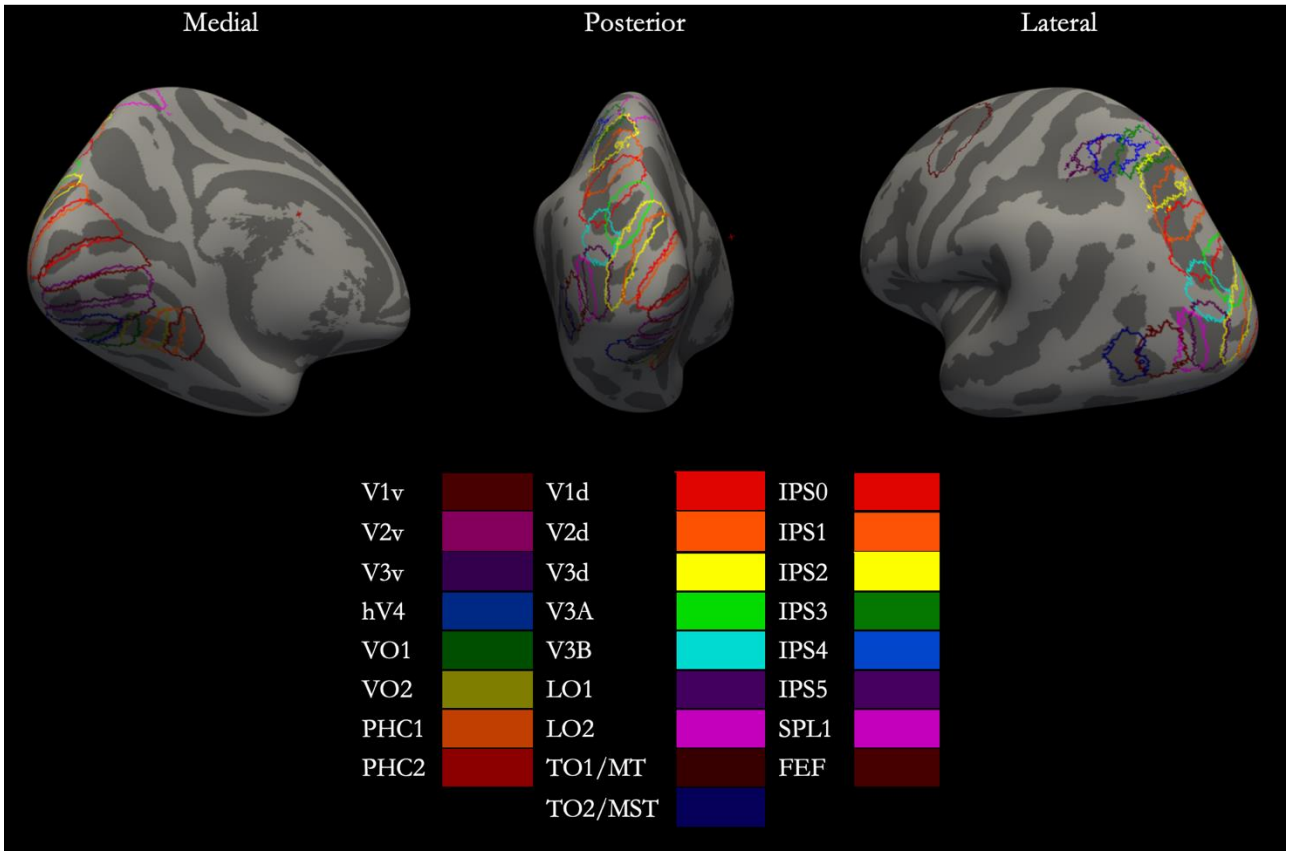

Figure 54. The left-right intersection visual area ROIs overlaid on the left symmetrical surface. ROIs were generated by taking, for each visual area, vertices that were included in the area for both left and right hemispheres.

4.2.4.4 Exclusive ROIs

One way to control for this level of overlap is to create an ROI for each visual area that does not include vertices which are common to other visual areas. We refer to these as "exclusive" visual area ROIs (Figure 56). Exclusive ROIs were generated by taking the intersection ROIs and excluding

vertices which fall within a given visual area and the neighbouring visual area(s). Figure 55 shows a schematical example of how these exclusive ROIs were generated. The exclusive ROIs were used throughout this chapter[19].

Vertex	Visual Area					Visual Area				Outcome
	V1v	V1d	V2v	V2d		V1v	V1d	V2v	V2d	
1	**1**	**1**	0	0		0	0	0	0	Excluded
2	**1**	0	0	0	->	**1**	0	0	0	Included
3	0	0	**1**	**1**		0	0	0	0	Excluded
4	0	0	0	**1**		0	0	0	**1**	Included
5	0	0	**1**	**1**		0	0	0	0	Excluded

Figure 55. A schematic representation of how the exclusive ROIs were generated. Vertices which were included in more than one visual area were excluded. Vertices that were assigned to only one visual area were included in the ROI for that area.

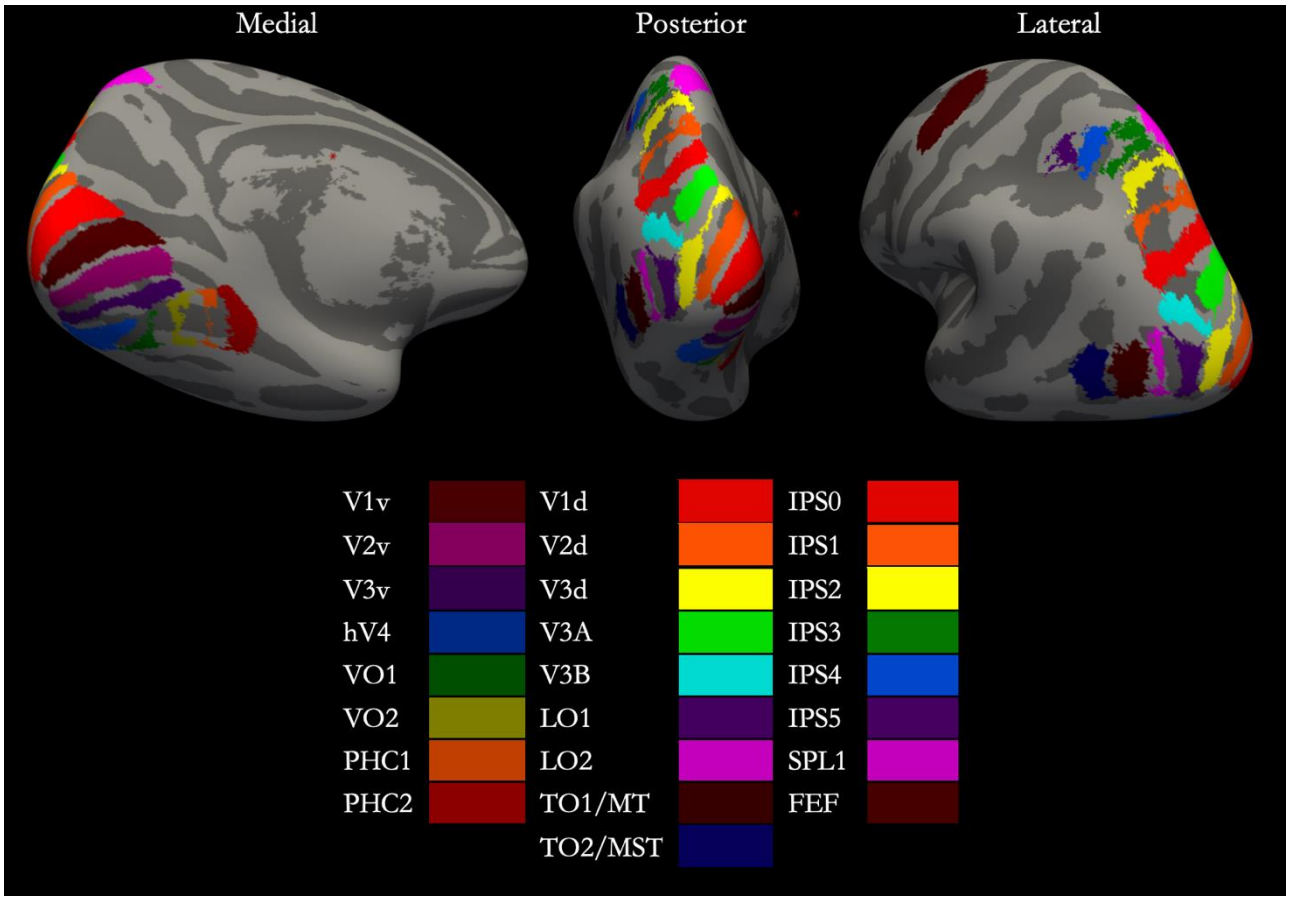

Figure 56. The exclusive visual area ROIs overlaid on the left symmetrical surface. An ROI for a given visual area was generated by taking the left-right intersection ROIs and excluding any vertices that were assigned to more than one visual area.

[19] For completeness I compared the results obtained from the use the MPM, left-right intersection and exclusive ROIs, finding the results to be robust. The details of these additional analyses are found in Appendix VIII.

Figure 57 shows TO1/MT (yellow) and TO2/MST (red) ROIs generated through the four approaches. In summary, the MPM generates ROIs which are inconsistent in size across hemispheres, particularly TO2/MST. The tFPM generates ROIs which are large with a significant amount of overlap between neighbouring areas. Intersection ROIs reduce this overlap to some degree but do not remove it. Exclusive ROIs appear to be the most appropriate approach as they result in plausibly size ROIs with no overlap between a region and its neighbours.

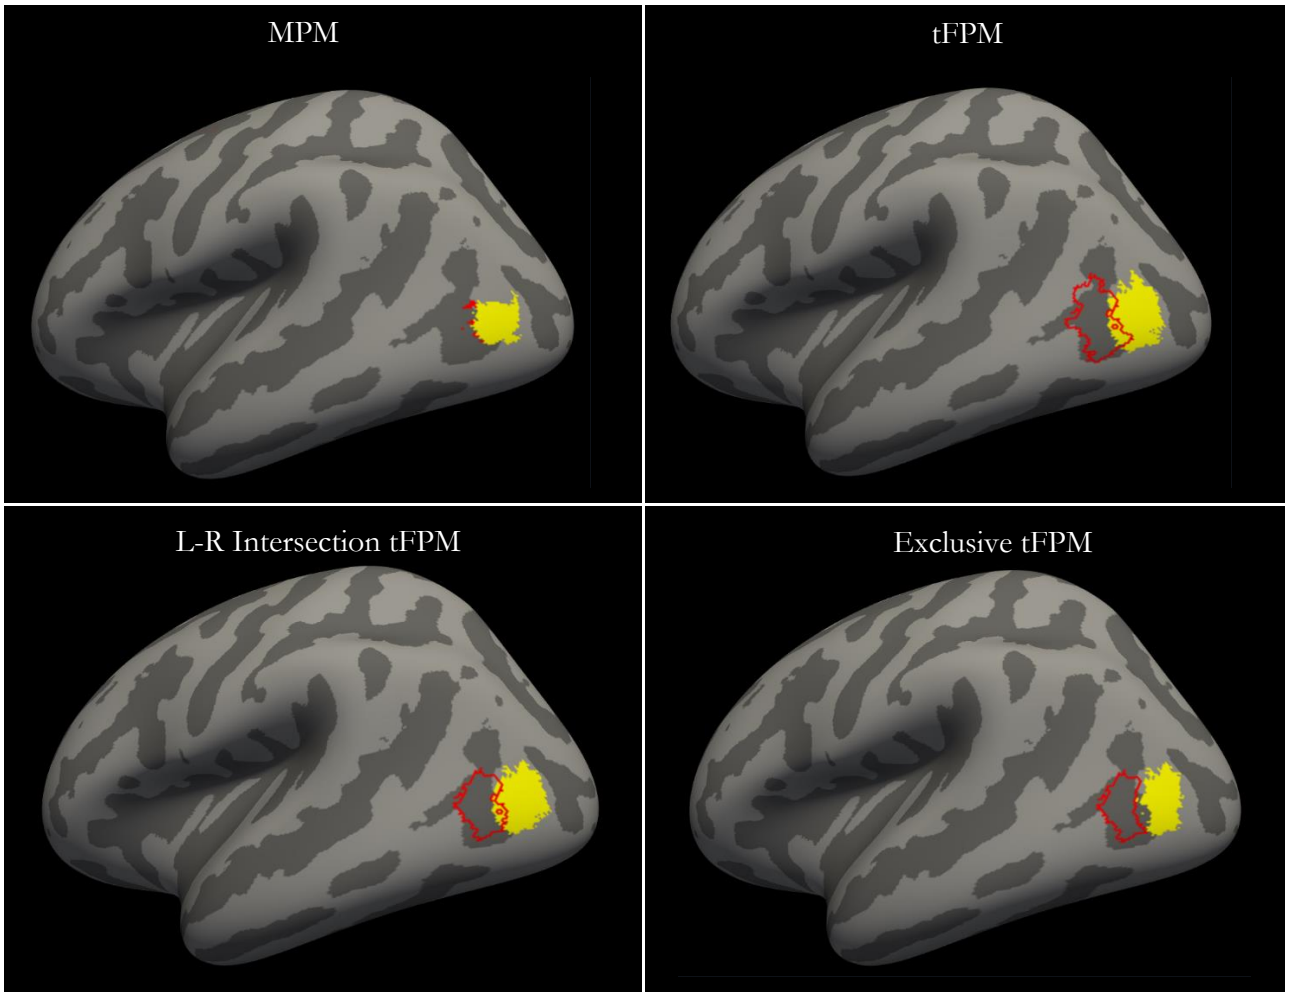

Figure 57. TO1/MT (yellow) and TO2/MST (red) ROIs generated through four different approaches; max probability (top left), thresholded full probability (top right), left-right intersection (bottom left) and exclusive (bottom right).

4.2.5 Interim summary

Thus far, I have described the general methods used in this chapter relating to data pre-processing, symmetrical surface generation and ROI creation. The following sections detail the ROI-to-ROI mean timecourse correlation analysis, the vertex-to-vertex cross-hemisphere timecourse

correlation analyses and the between-brains timecourse consistency analysis. The corresponding methods used for, and results obtained from, each analysis are presented in the relevant sections below.

4.3 ROI-to-ROI mean timecourse correlation analysis

The aim of the ROI-to-ROI cross-hemisphere timecourse correlation analysis was to assess the consistency of the timecourse between hemispheres, resulting in an index of bilaterality. If the ROI in the left hemisphere has a different functional role to the right hemisphere, the correlation between the two timecourses should be lower than if the two ROIs have similar functional roles. We use the cross-hemisphere correlations between left and right V1 as a benchmark for comparison with TO1/MT and TO2/MST. I use V1 because it is the earliest stage of visual processing and there is minimal opportunity for the two hemispheres to be specialised for different tasks. Correlations between left and right V1 would primarily be driven by correlations within the parts of the visual image to the left and right of fixation.

To pre-empt the results, I find that within the same brain, the timecourses of corresponding areas are highly correlated across hemispheres, with no clear evidence of lateralisation of TO1/MT. However, whilst we find cross-hemisphere correlations of TO2/MST timecourses to be high, the correlations are lower than in TO1/MT, V1v and V1d, potentially indicating a degree of functional lateralisation.

4.3.1 Methods

For each subject, the mean timecourse of vertices within a given exclusive label (referred to as "ROIs" hereafter) is extracted from the surface BOLD data using FreeSurfer's "*mri_segstats*". Details of how these ROIs were generated are described in section 4.2.4. For each visual area, pairwise correlations were computed between the individual subject's average timecourse in one hemisphere and the average timecourse from the opposite hemisphere using MATLAB's "*corr"* function (Pearson's r). Once individual subject correlations were computed, a group mean across all subjects was calculated for each visual area (Figure 58). All visual areas within the Wang et al. (2015) atlas are included but V1 is the primary area for comparison.

The Wang atlas splits V1 into dorsal (V1d) and ventral (V1v) ROIs, thus each contains a quadrant representation of the contralateral visual field. However, we are interested in the potential lateralisation of TO1/MT and TO2/MST, each of which contains a hemifield representation, relative to V1. To guard against apparent area differences being driven by differences in image content in the upper and lower visual field, V1v and V1d correlations were averaged to create a single measure for comparison to TO-1 and TO-2.

4.3.2 Results

Overall, the timecourses of the visual areas were highly correlated across hemispheres (mean r = 0.806; Figure 58). As expected, the average timecourses of left and right hemisphere V1 timecourses were most highly correlated (mean r = 0.872, SD = 0.038) compared to the other visual areas. Given we treat V1 as our baseline, from which the relative correlations of other areas are compared, we examine the correlations of TO1/MT and TO2/MST.

Interestingly, the left and right hemisphere mean timecourses of TO1/MT (mean r = 0.864, SD = 0.036) and TO2/MST (mean r = 0.816, SD = 0.059), were also particularly highly correlated. However, they were lower than those seen for V1. Following the logic that differences in function in the two hemispheres will depress between hemisphere correlations, these data could indicate some functional differences across hemispheres in the TO regions. We next explored this with inferential statistics.

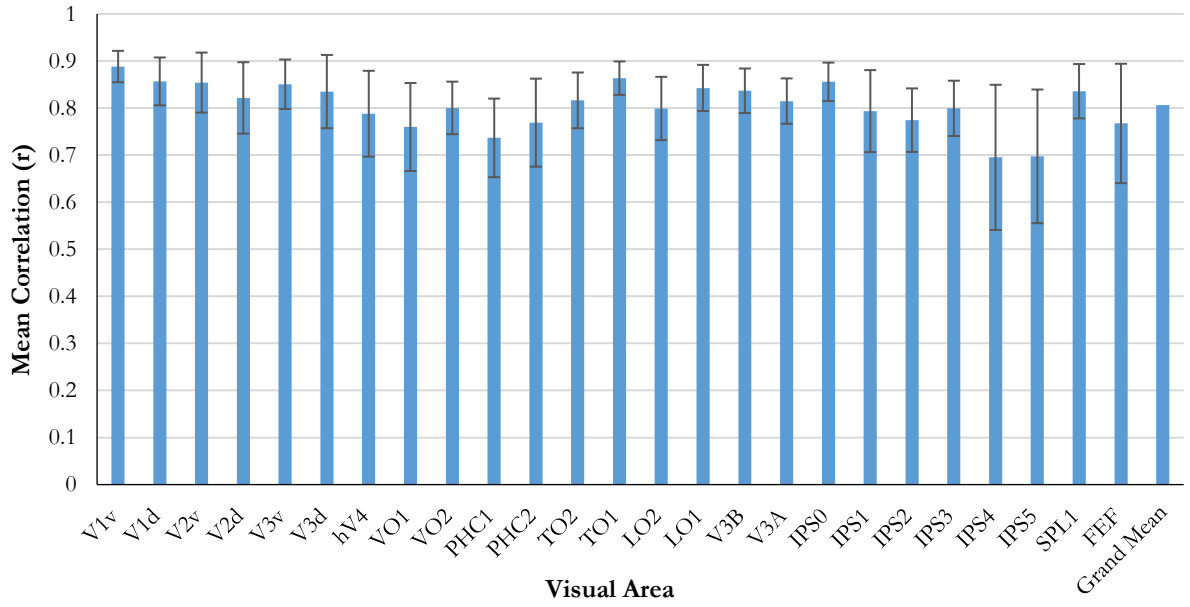

Figure 58. Group mean ROI-to-ROI timecourse correlation coefficients for each of the 25 visual areas. Error bars represent standard deviation. Results show the timecourses of the visual areas are highly correlated across hemispheres. TO2/MST is less correlated than TO1/MT, V1v and V1d, suggesting there may be differences in function across hemispheres for TO2/MST. Fisher's Z equivalent figure can be found in Appendix G.

Paired samples t-tests reveal that the mean cross-hemisphere timecourse correlation for TO1/MT was not significantly different to the mean correlation for V1, $t(16) = .997$, $p = .333$ (uncorrected). In contrast, the mean cross-hemisphere timecourse correlation for TO2/MST was significantly different to that seen for V1, $t(16) = 2.942$, $p = .010$ (uncorrected)[20].

As the TO1/MT and TO2/MST cross-hemisphere timecourse correlations are still relatively high (compared to V1), and only the difference between V1 and TO2/MST was significant, this provides only weak evidence that TO1/MT and TO2/MST are lateralised in each brain. However, we do note that TO1/MT and TO2/MST show lower mean correlations than V1, and therefore we explored whether this was consistent across brains.

Figure 59 shows the individual subject cross-hemisphere timecourse correlations plotted in paired format for V1 and TO1/TO2 (group means represented by bars). The figure shows a depression in the mean correlation for TO regions compared to V1, which is largest for TO2/MST.

[20] The V1 to TO1/TO2 comparisons were also performed on Fisher's Z transformed values and can be found in Appendix H. Note, the transform did not alter the results presented here.

Looking specifically at the individual subject data, there is clear variability between brains. For some subjects V1 and TO correlations are comparable, others show substantial differences between the regions.

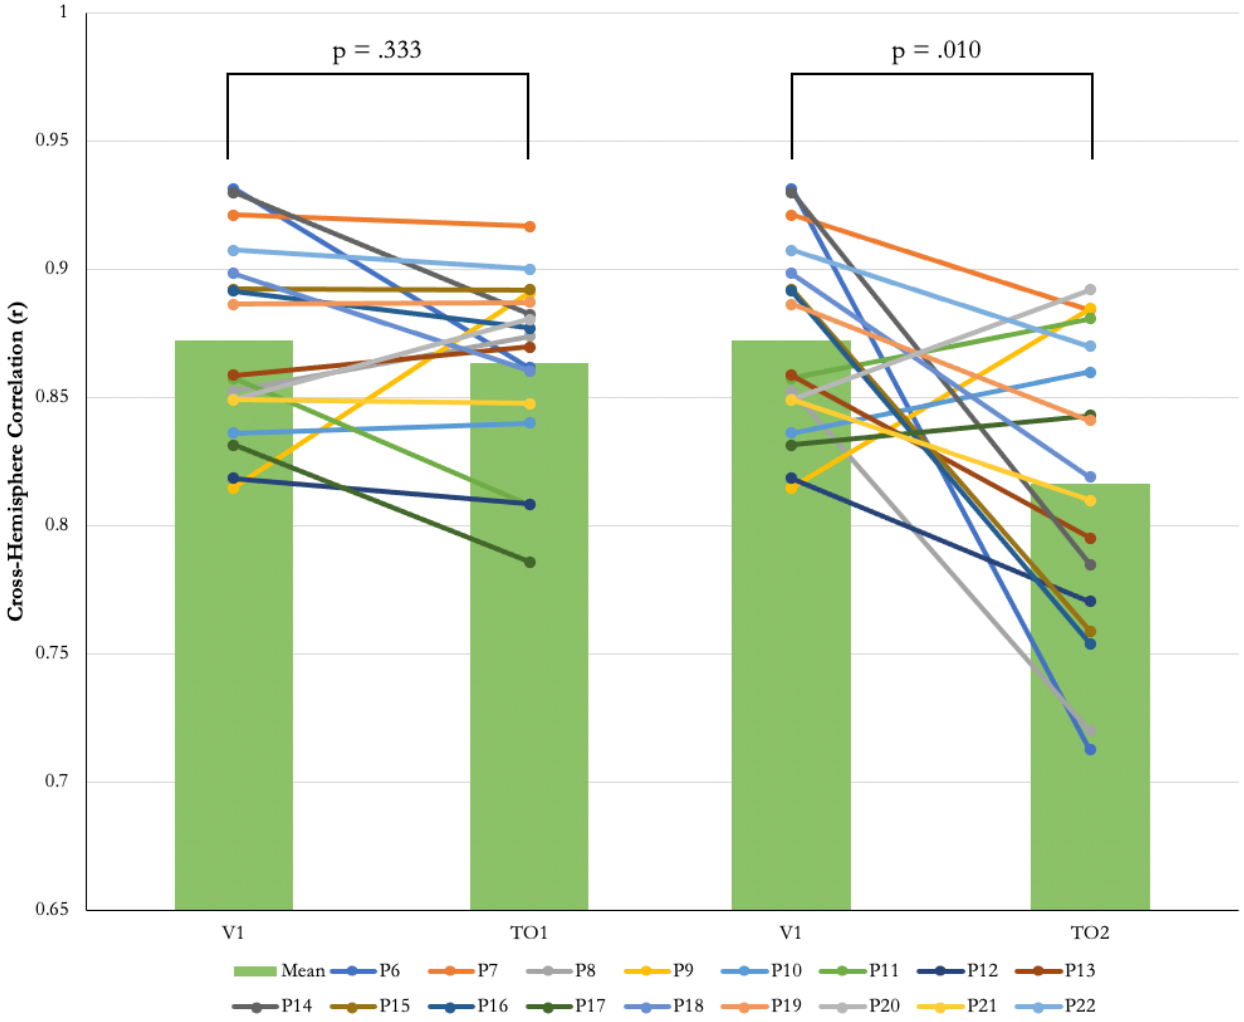

Figure 59. Individual subject cross-hemisphere timecourse correlations presented in paired format for V1 and TO1/TO2. Group means are represented by bars. Data shows for both TO1/MT and TO2/MST there is a depression in the mean correlation relative to V1, with clear inter-subject variability (i.e., for some brains the timecourses of V1 are more highly correlated across hemispheres than TO1/TO2 timecourses, for others we find the reverse). Fisher's Z equivalent figure can be found in Appendix I.

4.3.3 ROI-to-ROI correlation validation

The results above cannot be construed as providing strong support for the lateralisation hypothesis. One concern with these results however is that correlations might be biased by global signal changes and more localised effects of head movements. Global signal is defined as the average timecourse of all voxels within the brain (Zarahn et al., 1997) and contains a variety of noise signal components, e.g., motion-related, and physiological signals. A particular concern here is head

movements that create false activation around the edge of the brain. Head movement could have differential impacts on measured BOLD activity across the brain and so elevate some between hemisphere correlations. To investigate the impact of global signal changes on our results, I performed two validations with a subset of subjects (n = 10).

First, I examined the contribution of global signal to the correlations. To estimate this, I computed cross-hemisphere correlations between the timecourses of distant brain regions that do not have obvious functional connections.

Next, I factored out the contribution of head movements (given the nature of the stimulus and the scan duration [20 minutes] some subject head movement is expected during the acquisition). I removed head movement noise signals from the data (denoised) through the use of ICA-AROMA (Pruim et al., 2015) and examined its impact on the results.

We estimate a substantial portion of the correlations between areas are due to global signal changes across the brain. This level of correlation appears to be consistent across the brain, and analyses performed on denoised data indicate head movements do not have a disproportionate impact on TO areas.

4.3.3.1 Distant brain region correlations

To provide some insight into the role of global signal changes we extracted the mean timecourses from five ROIs (Figure 60):

1. V1v
2. Broca's area
3. An auditory region located in the middle of the Heschl's gyrus (HG)
4. Combined visual
5. Non-visual

The large combined visual area ROI was created by combining the full probability ROIs for each visual area into a single ROI. The large non-visual ROI was created by taking all vertices outside of the combined visual area ROI.

We then computed correlations between ROIs located in opposite hemispheres to estimate what proportion of the correlation between areas is due to global signal. Our results suggest that global signal changes across the brain accounts for a correlation, r, of ~0.35.

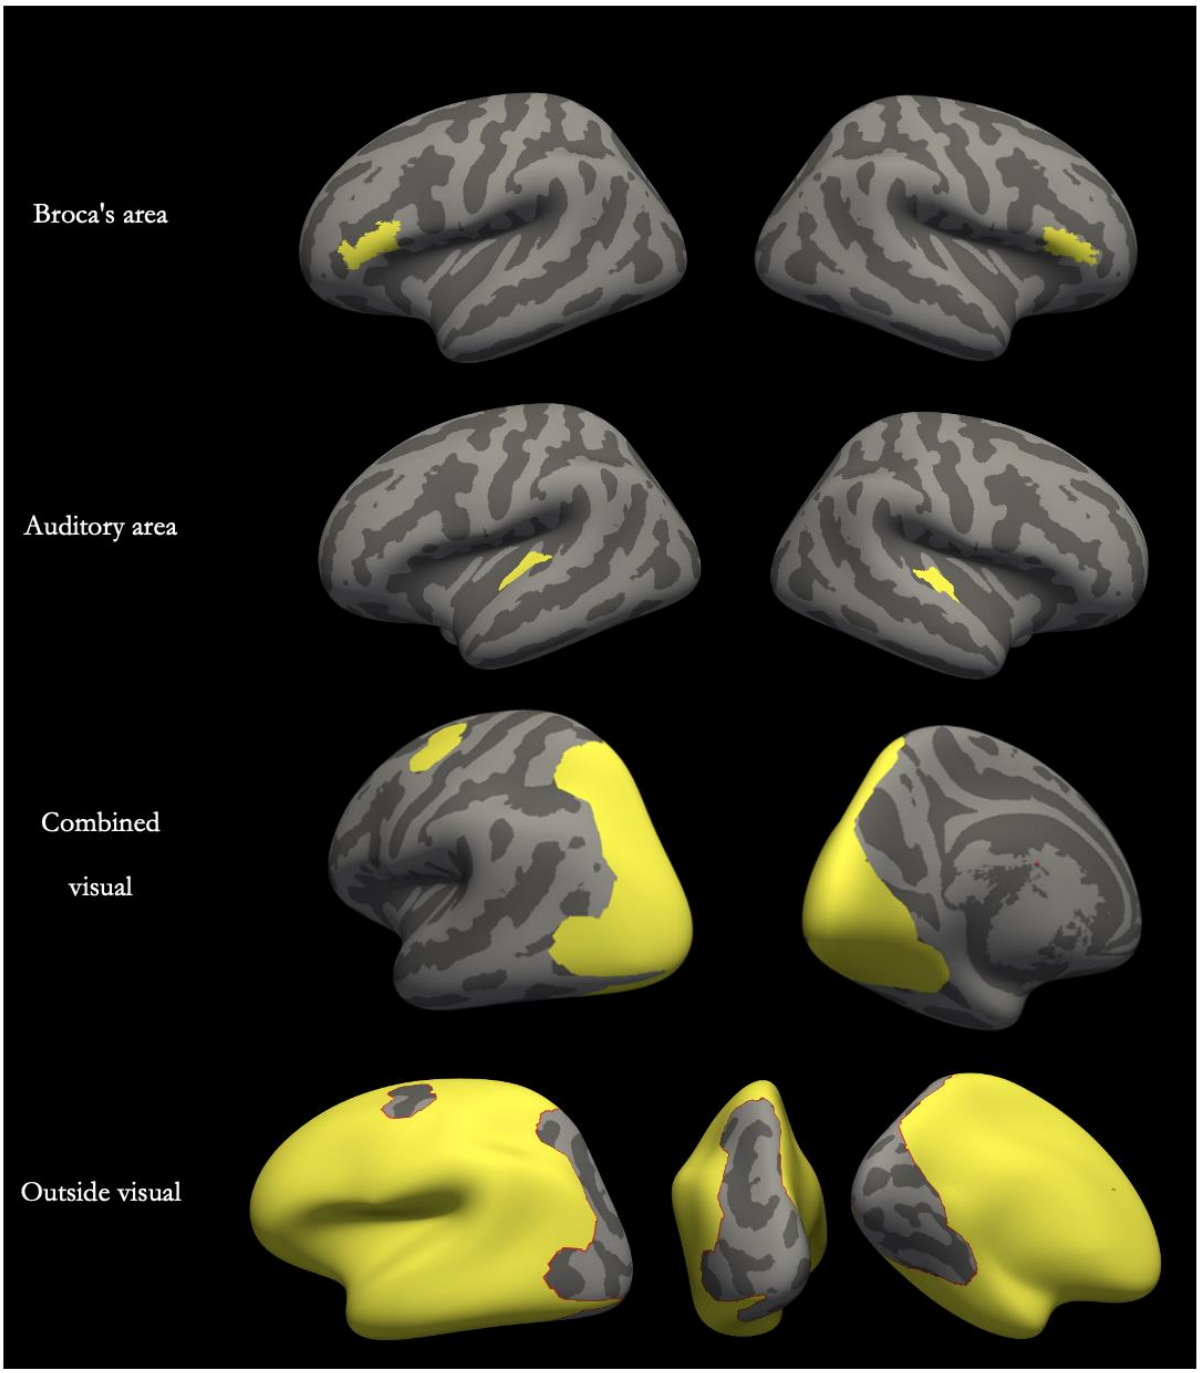

Figure 60. The Broca's area, auditory region, combined visual and non-visual ROIs used for the ROI-to-ROI correlation validation. The V1v ROI is presented in Figure 51. The average timecourse of vertices within each ROI (yellow) was extracted. Correlations were then computed between the timecourses of ROIs located in opposite hemispheres.

For each subject, we first computed correlations between the timecourses of ROIs located in opposite hemispheres, e.g., V1v in the left hemisphere and BA in the right. We took the mean across subjects for each pairwise correlation to generate a mean left versus right and right versus left

coefficient for each comparison. Finally, we then took the mean across the cross-hemisphere correlations to generate a grand mean coefficient for each comparison (Table 10)[21].

Table 10. The group mean cross-hemisphere correlations computed between ROIs in opposite hemispheres. Correlations were computed within the visual system (V1v vs combined visual), brain wide (V1v vs outside visual) and between the visual areas and two anatomically distant and functionally distinct control regions, Broca's area (BA) and auditory cortex.

Comparison	Grand Mean	Std Dev
V1v vs. V1v	0.886	0.035
V1v vs. combined visual	0.805	0.057
V1v vs. outside visual	0.532	0.110
V1v vs. BA	0.314	0.129
V1v vs. auditory	0.327	0.129
Combined visual vs. BA	0.394	0.136
Combined visual vs. auditory	0.343	0.145
BA vs. auditory	0.353	0.175

As already reported above, the mean timecourses of left and right V1v were strongly correlated (mean r = 0.886, SD = 0.035). On average, V1v correlated with regions within the visual pathway in the opposite hemisphere at r = 0.805 and regions outside of the visual pathway at r = 0.532. This V1v-to-non-visual correlation is an initial indication of what proportion of the correlations are due to global signal changes. However, the non-visual ROI includes many sites throughout the cortex and as such some are likely to be functionally connected. Therefore, to get a better estimate of global signal change, we computed correlations between V1v and the two control regions located in the opposite hemisphere which are not expected to be functionally connected (BA and auditory). On average, correlations reveal a moderate-to-low positive relationship between the timecourses of V1v and the two control regions, suggesting V1v correlates with non-visual regions at ~0.32.

Finally, we computed correlations between the combined visual ROI and both control regions. These results suggest non-visual areas correlate with the combined visual system in the opposite hemisphere at ~0.37. The two control regions differ both in function and location (BA is more

[21] Separate left versus right and right versus left correlations can be found in Appendix VI.

anterior than the auditory region, and therefore more distant from V1v) yet correlate with V1v and the visual system as a whole at ~0.35 on average. Taking these comparisons together, we estimate that global signal changes across the brain accounts for a correlation of ~0.35.

Given the timecourses of functionally distinct regions positively correlate to some extent, a "confidence check" was performed to ensure the method used was appropriate. The V1v timecourse of an individual subject (Subject 6) was correlated with each of the other subject's auditory area timecourse. As expected, Subject 6's V1v timecourse did not positively correlate with the other subject's auditory timecourses (Table 11 and Table 12).

This confidence check rules out induced cross-modal effects. Given all subjects watched the same movie and we find synchronised activity across brains in the visual cortex, if a visual scene primed activity of an associated sound in the auditory cortex, we would expect the correlation to be non-zero.

Table 11. Correlation coefficients computed between Subject 6's left V1v timecourse and the right auditory (MGH) timecourses of each subject.

	Right MGH										
	S7	S8	S9	S10	S11	S12	S13	S14	S15	Mean	Std Dev
S6 Left V1v	-0.072	0.01	0.058	0.019	0.018	0.003	0.018	-0.036	-0.05	-0.004	0.041

Table 12. Correlation coefficients computed between Subject 6's right V1v timecourse and the left auditory (MGH) timecourses of each subject.

	Left MGH										
	S7	S8	S9	S10	S11	S12	S13	S14	S15	Mean	Std Dev
S6 right V1v	-0.036	0.021	0.085	0.006	0.046	0.026	-0.018	-0.024	-0.097	0.001	0.053

4.3.3.2 ICA-AROMA

The analysis above suggests global signal change across the brain accounts for a correlation, r, of ~0.35. Next, I examined the impact of signal changes due to head movements on the correlations.

To compensate for the spatial consequences of head movements, our data were aligned to the middle volume of the timeseries using MCFLIRT (Jenkinson et al., 2002). However, the algorithm does not remove noise signals due to head movements from the data. For this, additional fMRI denoising strategies can be implemented.

I used an ICA-based strategy for Automatic Removal of Motion Artifacts (ICA-AROMA), which provides a simple way to "factor out" the effect of head movements and has been validated for use with resting state and task-based fMRI data (Pruim et al., 2015). ICA-AROMA first decomposes the data into independent components using ICA, before automatically identifying components relating to head motion based on a combination of four features. These four features, two temporal and two spatial, are summarised below:

1. High-frequency content (HFC). The time series of signal components are dominated by low frequencies, whereas components relating to motion will show power at higher frequency bands.
2. Maximum correlation with realignment parameters. The timeseries of components relating to motion artifacts should correlate to some extent with the realignment parameter timeseries derived from MCFLIRT (or other volume-realignment algorithm). In contrast, the time series of signal components should not correlate with the realignment parameters.
3. Edge fraction. Head movements result in a shift in the brain location relative to the voxel location. As a result, there is significant variation in the signal for voxels located near to the edge of the brain, due to voxels not representing identical brain regions over the duration of the scan. ICA-AROMA creates a mask of the edge of the brain and calculates the edge fraction of each component by taking the sum of absolute z-values of voxels overlapping the mask, divided by the sum of absolute z-values of all voxels.
4. CSF fraction. Much like the edge fraction, a CSF mask is defined from a CSF segmentation. The CSF fraction for each component is calculated by taking the sum of the absolute z-values of voxels overlapping the mask and dividing by the sum of absolute z-values of all voxels.

The timecourses of components classified as being motion-related are removed from the BOLD data through linear regression, resulting in "denoised" or "clean" data.

I extracted the mean timecourses, this time from the denoised BOLD data (post ICA-AROMA), and computed correlations between these timecourses. If a high proportion of the correlation between regions is being driven by signal changes due to head movements, we would expect to see lower correlations in these regions for the denoised data compared to data pre-processed using standard techniques. If head movements have a differential impact on correlations across the brain, then the relative magnitude of correlations may change.

Table 13 shows the mean correlation for each pairwise comparison to be lower for denoised data, with the mean difference between data types being 0.066. However, the pattern of the results does not change (left versus right and right versus left comparisons can be found in Appendix VII).

Table 13. The group mean cross-hemisphere correlations computed between ROIs in opposite hemispheres. The mean of the left versus right and right versus left comparisons are displayed here. Results are shown for data pre-processed using the standard approach (raw) and data denoised with ICA-AROMA (AROMA). The overall pattern of the results does not change irrespective of the approach adopted. Fisher's Z equivalent table can be found in Appendix J.

	Raw		AROMA		
	Mean	Std Dev	Mean	Std Dev	Difference
V1v vs. V1v	0.886	0.035	0.826	0.067	0.060
V1v vs. combined visual	0.805	0.057	0.745	0.086	0.060
V1v vs. outside visual	0.532	0.110	0.471	0.111	0.062
V1v vs. BA	0.314	0.129	0.254	0.130	0.060
V1v vs. auditory	0.327	0.129	0.268	0.120	0.059
Combined visual vs. BA	0.394	0.136	0.323	0.130	0.071
Combined visual vs. auditory	0.343	0.145	0.274	0.117	0.069
BA vs. auditory	0.353	0.175	0.268	0.120	0.084
				Mean	0.066

Next, we performed the visual area ROI-to-ROI correlation analysis on the denoised data for the same sub-set of subjects (Figure 61). We find the correlations are lower for denoised data compared to the raw data, but the overall pattern of our result did not change. The grand mean of the

AROMA data (r = 0.718) remains high, and on average the timecourses of TO1/MT (r = 0.796)[22] are similarly correlated across hemispheres as V1v (r = 0.826) and V1d (r = 0.755) timecourses. V1v and TO1/MT are the highest and second highest correlated regions respectively, irrespective of whether denoised or raw data were used. Again, we find the timecourses of TO2/MST (r = 0.753) are less correlated across brains than V1 and TO1/MT.

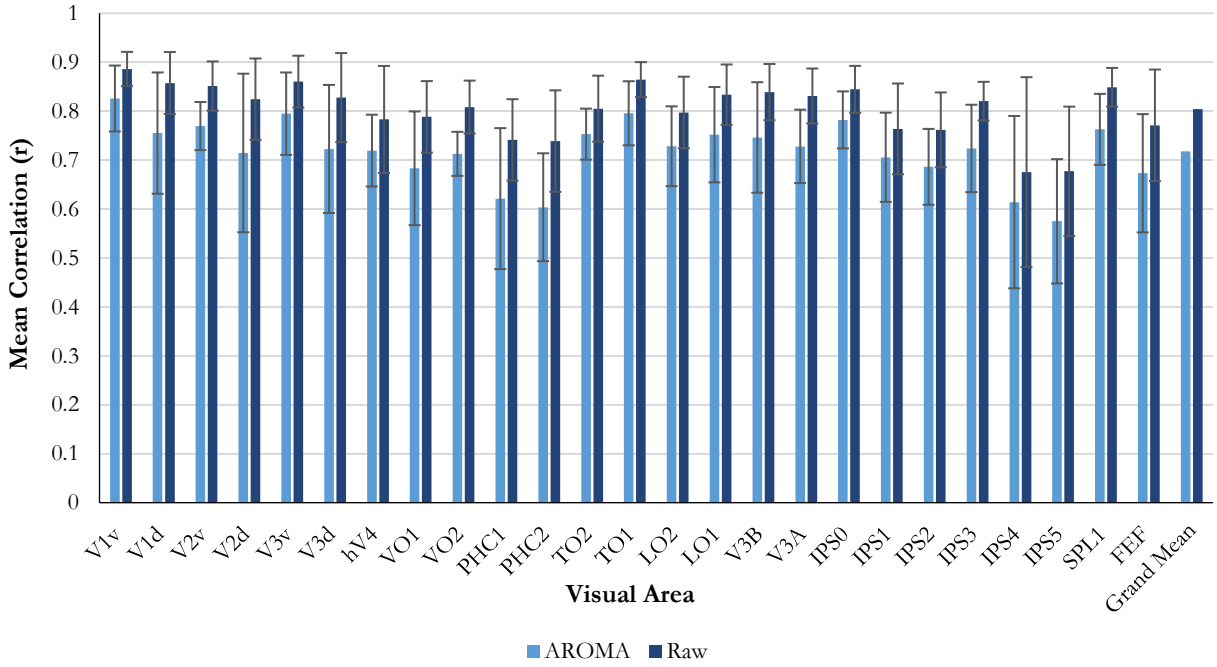

Figure 61. Group mean ROI-to-ROI timecourse correlation coefficients for each of the 25 visual areas. Error bars represent standard deviation. Results are shown for data pre-processed using the standard approach (raw) and data denoised with ICA-AROMA (AROMA). Correlations are lower for denoised data than raw data, but crucially the pattern of the results did not change irrespective of the approach used. Fisher's Z equivalent figure can be found in Appendix K.

In summary, we draw two conclusions from these analyses: 1) within the same brain, the timecourses of corresponding areas are highly correlated across hemispheres, 2) the timecourses of TO2/MST are, on average, less correlated across hemispheres than V1 and TO1. These conclusions hold true for denoised data, suggesting head movements do not have a disproportionate impact on TO areas.

[22] Correlations computed with AROMA data were lower by 0.069 and 0.052 compared to raw data for TO1/MT and TO2/MST respectively. Correlations with AROMA data were lower by 0.06 and 0.1 compared to raw data for V1v and V1d respectively.

So far, our findings do not provide strong supporting evidence of the lateralisation hypothesis proposed by Strong et al. (2019). However, there are notable differences between hemispheres observed in TO2 which could speak to differences in function across the two hemispheres.

4.4 Vertex-to-vertex cross-hemisphere correlations

The results from the analysis above may be confounded by the size of the ROIs, which vary across visual areas. With a large ROI, the timecourse is generated by averaging over more vertices than if a smaller ROI is used. This in turn will "smooth" the data and increase the signal-to-noise ratio for larger ROIs. Given V1 contains substantially more vertices than the TO regions, the significant difference in the cross-hemisphere correlation between V1 and TO2/MST could be driven by the relative size of the ROIs. For example, as a function of ROI size, the V1 timecourse may be smoothed to a greater extent compared to TO1/2 timecourses. Consequently, the larger ROI (V1), thus "smoother" timecourse has less noise and so is more highly correlated across hemispheres, than the timecourse of the smaller (TO1/2) ROIs. Therefore, I conducted a second analysis to control for this potential bias.

I performed a vertex-to-vertex analysis, in which correlations between the timecourses of corresponding vertices in opposite hemispheres were computed first. I then took the mean of the correlation coefficients of the vertices within a given ROI[23] and compared across the visual areas. To reiterate the above, the ROI-to-ROI analysis performs correlations between the mean timecourses of vertices within a ROI, here I perform correlations on a per-vertex basis before taking the mean over the correlations within a given ROI (see section 4.4.1 below).

I find visual areas remain highly correlated across hemispheres when the mean correlation is extracted from vertices within the ROIs. The results reflect those obtained from the previous analysis:

[23] Exclusive labels were again used as the ROIs for this analysis. We performed the same analysis, with both the intersection and max probability labels, finding the same result as is presented here. See Appendix VIII for details on the results of these additional analyses.

1. Relative to the V1 baseline, timecourses of vertices within TO1/MT are highly correlated. Based on the relative magnitude of correlations, we infer the function of the TO1/MT is no more lateralised than any other region.
2. Timecourses of vertices within TO2/MST are less correlated across hemispheres than V1v, V1d and TO1/MT.

4.4.1 Methods

Once the subject's individual data had been resampled to the fsaverage_sym surface (see section 4.2.2 for details), pairwise cross-correlations were computed between the timecourses of corresponding vertices in the two hemispheres using MATLAB's "*corr*" function (Pearson's r). For each vertex, a single correlation coefficient is generated. This subject-specific correlation coefficient map can then be displayed on the fsaverage_sym surface (see Figure 62 for examples). I then extracted the correlation coefficients for vertices within a given ROI and took the mean of those values for each ROI and each subject. This results in a visual area specific correlation coefficient that represents the similarity of the area's timecourse across the two hemispheres for each subject.

As with the ROI-to-ROI analysis, if two areas functionally correspond, i.e., they are processing the same elements of the stimulus, their timecourses should be more highly correlated than if they were functionally distinct. We expected the correlations from the vertex-to-vertex analysis to be lower than those reported in the mean timecourse analysis due to the fact this analysis computes correlations before averaging, i.e., there was no additional smoothing. It was, however, important to compare the pattern of the results between approaches.

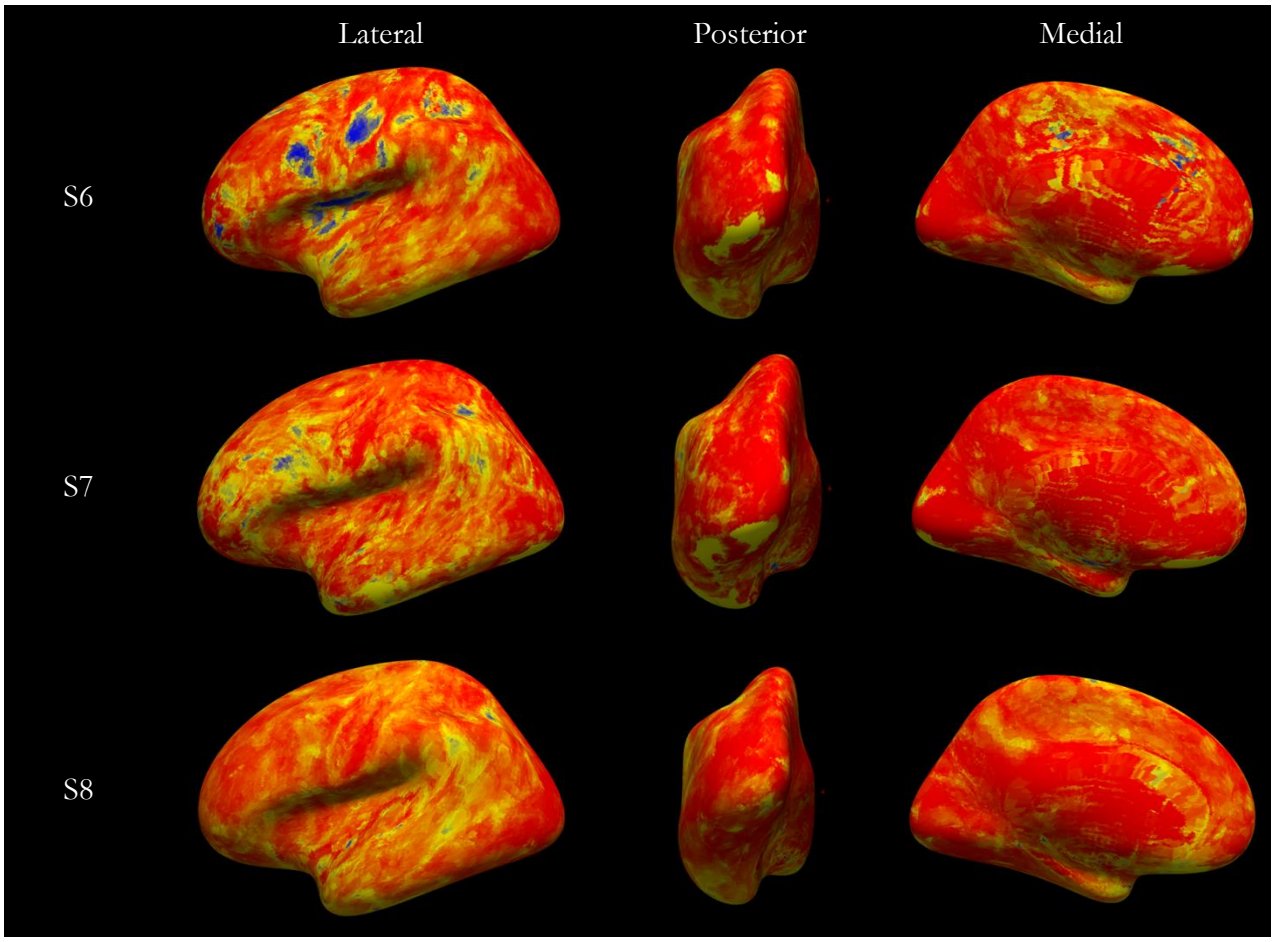

Figure 62. Three examples subject's vertex-to-vertex cross-hemisphere correlations overlaid on the inflated surface. "Hotter" coloured vertices indicate higher correlations between hemispheres.

4.4.2 Results

Figure 63 shows the mean correlations generated for each visual area for both the ROI-to-ROI and vertex-to-vertex analyses. As with the ROI-to-ROI analysis, the vertex-to-vertex analysis was performed on data following surface-based smoothing[24]. As expected, the grand mean, calculated by taking the mean r across all the visual areas, is lower for the mean vertex analysis (r = 0.599) than it is for the ROI analysis (r = 0.806).

[24] We also performed the vertex-to-vertex analysis on volume-smoothed data. Both smoothing approaches yield similar results. Details of the volume-smoothed results are in Appendix IX.

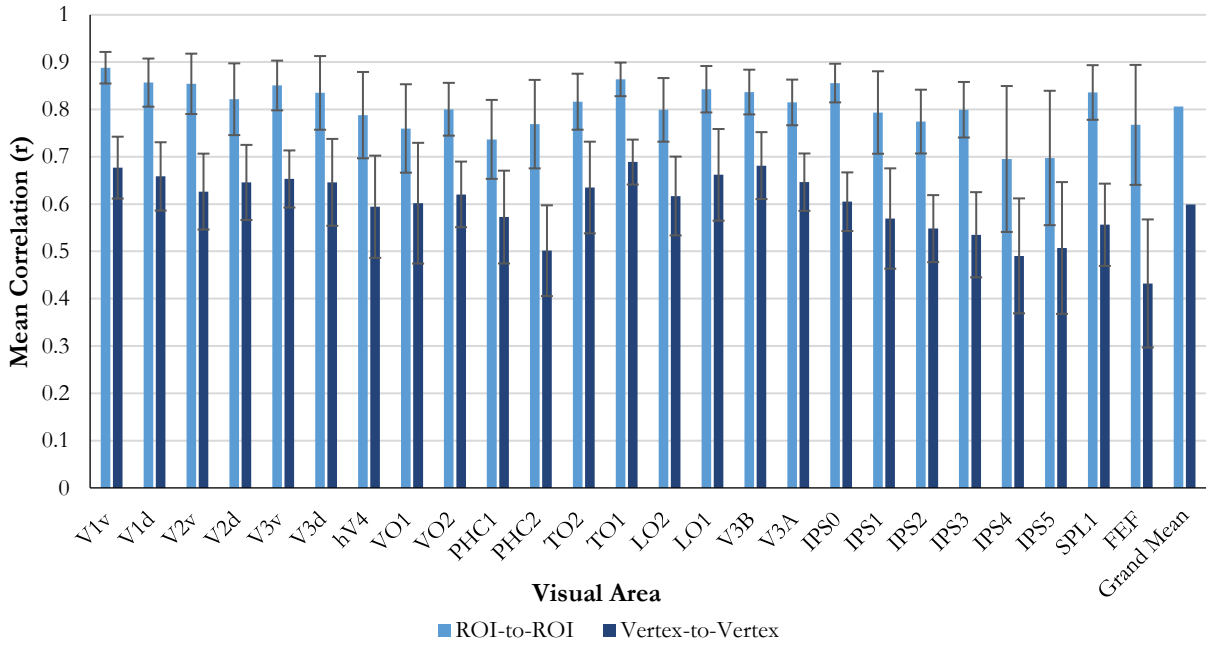

Figure 63. Group mean cross-hemisphere correlations for each visual area. Results are displayed for both the ROI-to-ROI and vertex-to-vertex analyses. For the ROI analysis the mean timecourse of vertices within a given ROI in each hemisphere is extracted. Correlations are then computed between the timecourses from opposite hemispheres. For the vertex analysis, correlations are computed between corresponding vertices in each hemisphere. The mean coefficient is then extracted from vertices within each ROI. Results show the timecourses of TO2/MST are less correlated across hemispheres than V1v, V1d and TO1/MT. Error bars represent standard deviation. Fisher's Z equivalent figure can be found Appendix L.

The pattern across visual areas is comparable between analyses. In the ROI analysis, TO1/MT and TO2/MST are the 2nd and 12th most highly correlated regions, with V1v and V1d ranking 1st and 3rd respectively. Similarly, for the vertex analysis TO1/MT and TO2/MST are ranked 1st and 10th, with V1v and V1d ranking 3rd and 5th respectively. These findings are consistent with the example surface plots presented in Figure 62, in that there is not an obvious patch of lower correlation in the areas where TO1/MT and TO2/MST are found.

As stated above, we expected the left and right timecourses of V1 to be highly correlated as there is less opportunity for the two hemispheres to be specialised at such an early stage of processing and correlated activity will reflect correlations in the retinal image. The analysis appears to support that assertion, with the left and right hemisphere timecourses of V1 being highly correlated (r = 0.667, SD = 0.056).

The timecourses of TO1/MT correlated across hemispheres at r = 0.689 (SD = 0.047). Relative to V1, we expected to see lower correlations if left and right hemisphere TO1/MT had

different functional roles. However, given the timecourses are in fact more highly correlated across hemispheres than V1, we interpret this finding as not providing supporting evidence for the lateralisation hypothesis. Paired-samples t-tests performed on the correlations support this interpretation. There was no statistically significant difference between V1 and TO1/MT, $t(16) = -1.379$, $p = .187$ (uncorrected).

As with the ROI analysis, we find the timecourses of TO2/MST are less correlated across hemispheres ($r = 0.635$, $SD = 0.097$) than V1 and TO1/MT. This may indicate a degree of functional difference between hemispheres. However, paired-samples t-tests revealed, on average, cross-hemisphere timecourse correlations in TO2/MST were not significantly different to V1, $t(16) = 1.086$, $p = .294$ (uncorrected)[25].

Although the inferential statistics reveal the mean difference between V1 and TO regions to be non-significant, we find, as with the ROI analysis (Figure 59), a depression in the mean correlation for TO2/MST relative to V1 and considerable inter-subject variability in the cross-hemisphere correlations (Figure 64, right); some subjects show substantially lower correlations in TO2/MST compared to V1, while others show the reverse. In contrast, V1 and TO1/MT correlations are more comparable, as reflected by a mean difference of 0.012 (Figure 64, left). However, there are still notable individual differences in the correlations between the areas.

[25] The V1 to TO1/TO2 comparisons were also performed on Fisher's Z transformed values and can be found in Appendix M. Note, the transform did not alter the results presented here.

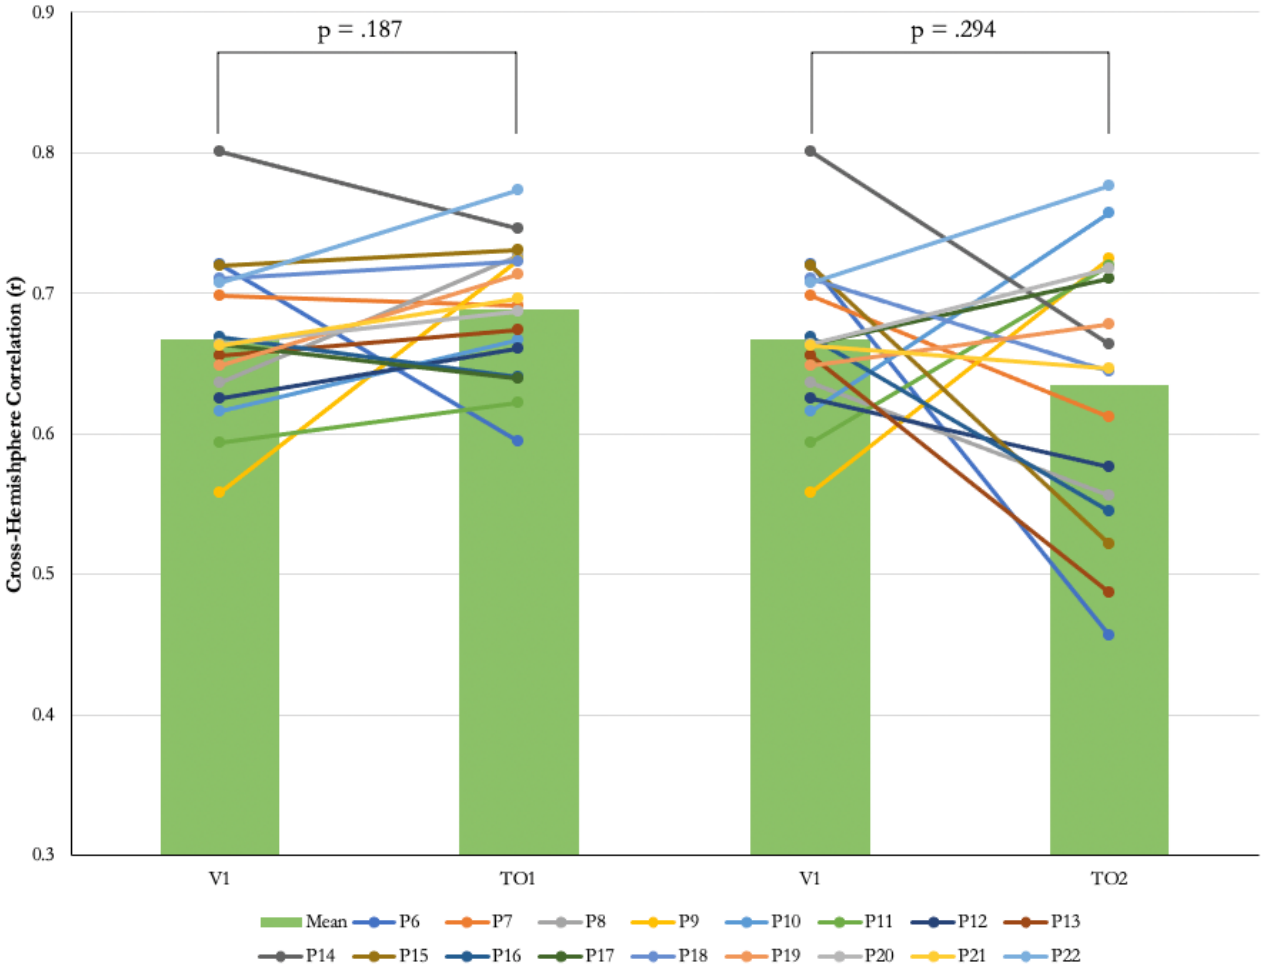

Figure 64. Individual subject cross-hemisphere timecourse correlations presented in paired format for V1 and TO1/TO2. Group means are represented by bars. Data shows, on average, TO1/MT timecourses are more highly correlated across hemispheres than V1 timecourses. For TO2/MST, we find a depression in the mean correlation relative to V1. Both V1 to TO1/MT and V1 to TO2/MST comparisons show clear inter-subject variability (i.e., for some brains the timecourses of V1 are more highly correlated across hemispheres than TO1/TO2 timecourses, for others we find the reverse). Fisher's Z equivalent figure can be found in Appendix N.

To summarise the findings so far, when the correlation between the mean timecourse of a ROI in each hemisphere is calculated, the difference in the correlation between V1 and TO2/MST is statistically significant. However, when we compute vertex-to-vertex correlations and then extract the mean coefficient from each ROI, the differences in the correlations between the areas of interest do not reach significance. As discussed in section 4.4, one possible explanation for this contrasting statistical result is the relative size of the ROIs from which the mean timecourse/correlation is extracted. When we control for the differences in the size of the ROIs, significant differences between areas are abolished. Therefore, the combined results from both the ROI- and vertex-based analyses suggest that on average, neither TO1/MT or TO2/MST show *clear evidence* of being more functionally

lateralised than V1 (i.e., on average, there appears to be no functional difference between hemispheres in these areas). We do note that on average the timecourses of TO2/MST are less correlated across hemispheres than V1 and TO1/MT.

Examination of the individual subject correlations led us to consider an alternative conclusion: it is possible that there is lateralised function in hMT+, but not in all brains and therefore, examining group data could obscure this. Consequently, it would be remiss of us to ignore these clear individual differences and therefore we explore this hypothesis in the following sections.

4.5 Between-brains timecourse consistency

An explanation as to why we do not find evidence to support Strong et al.'s (2019) conclusion that TO1/MT and TO2/MST have different functional roles across hemispheres might be that in their study, data was collected from seven subjects and only the group average data are presented. It is possible that their finding is not reflective of all brains, that the lateralisation of hMT+ is not a universal phenomenon. Therefore, we wanted to explore whether there are individual differences in the function of TO1/MT and TO2/MST across hemispheres.

Consider TO2/MST in the right hemisphere. If the region has the same functional role in all brains, we would expect the timecourses to be more highly correlated across brains than if there are individual differences in the function of the region. In the case of the latter, we would expect the correlations to be lower since we are not correlating like-for-like across all brains. Therefore, we compare the group average timecourse to each individual subjects timecourse with the aim of identifying possible individual differences with respect to lateralisation.

We performed, for each hemisphere individually, a between-brains correlation analysis, in which the timecourse of an individual was correlated with the group average timecourse. This analysis was performed on surface-smoothed data from each hemisphere individually[26], and gives us an index of the consistency of the timecourses relative to the group for each hemisphere. We find the

[26] We also performed the analyses on volume-smoothed data and found the results to be comparable to those obtained from surface-smoothed data. See Appendix X for results with volume-smoothed data.

timecourses of V1 are similarly correlated across brains in both hemispheres. On the other hand, we find lower correlations for TO1/MT and particularly TO2/MST. Based on the relative correlations, the results suggests that V1 functions similarly across brains and across hemispheres, whereas the lower correlations in TO1/MT and TO2/MST are compatible with differences between brains. We discuss these below.

4.5.1 Methods

First the average timecourse of vertices within the visual area ROIs for each subject was extracted. Next, correlations between the individual subject's timecourse, and the group timecourse were computed. The group timecourse was calculated by taking the mean across all subjects, excluding the subject with which we are correlating. For example, we compute the correlation between Subject 1's V1v timecourse and the group mean V1v timecourse, excluding subject 1's data from the mean. We then took the mean across these individual-to-group pairwise correlations for each visual area. A higher correlation indicates the timecourse was more comparable across brains.

4.5.2 Results

For 14/25 visual areas, timecourses extracted from the left hemisphere are more consistent across brains than the right hemisphere (Figure 65). The grand mean, calculated by taking the mean of all the individual visual area coefficients, was identical for both hemispheres ($r = 0.399$), suggesting that overall, the visual areas are consistent across subjects and across hemispheres.

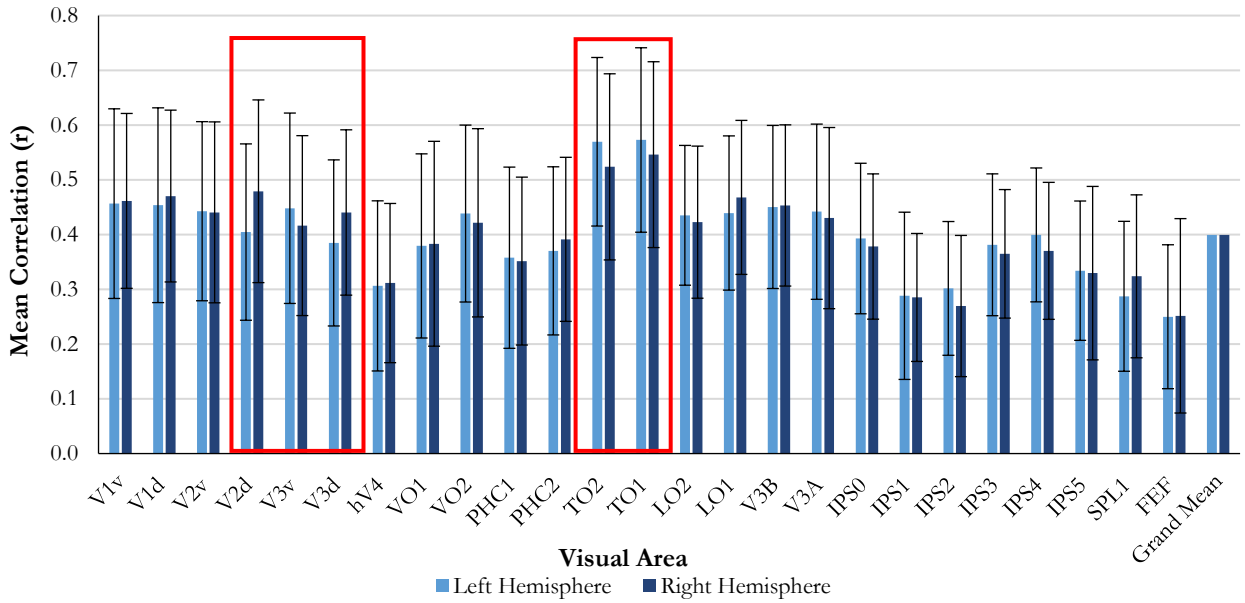

Figure 65. Group mean between-brains correlations for both the left and right hemispheres. Correlations were computed between each individual's visual area timecourse and the group average timecourse for that visual area (the group average timecourse did not include the brain with which the correlation was being computed). The mean of each of these subject-to-group correlations was then calculated. This analysis gives a measure of the consistency of the timecourses for each visual area. Results show that for TO1/MT and TO2/MST, left hemisphere timecourses are more highly correlated across brains than right hemisphere timecourses. Conversely, V1v and V1d timecourses are similarly correlated across hemispheres. Red boxes indicate significant differences between hemispheres in the consistency of the timecourses (V2d, V3v, V3d, TO1/MT and TO2/MST). Error bars represent standard deviation. Fisher's Z equivalent figure can be found in Appendix O.

As we are interested in the potential lateralisation of TO1/MT and TO2/MST, relative to the other areas, in particular V1, we first examined the relationship between hemispheres in the consistency (across brains) of the timecourses in these areas.

For each subject, we plotted the correlation between the individual timecourse, and the group mean timecourse for left and right hemispheres (Figure 66). We then examined how the individual subject data points compare between V1 and TO regions.

For V1, we find the datapoints are mostly clustered around the line of equality (Figure 66, top). The location of the datapoints indicate there is little hemispheric variability in the timecourse (i.e., they are similarly correlated across hemispheres). In contrast, for TO1/MT and TO2/MST, many data points are shifted away from the line of equality (Figure 66, bottom). This is more pronounced for TO2/MST relative to TO1/MT. This suggests that for some subjects the timecourses are more consistent in one hemisphere than the other. These exploratory findings were further investigated with statistical analyses.

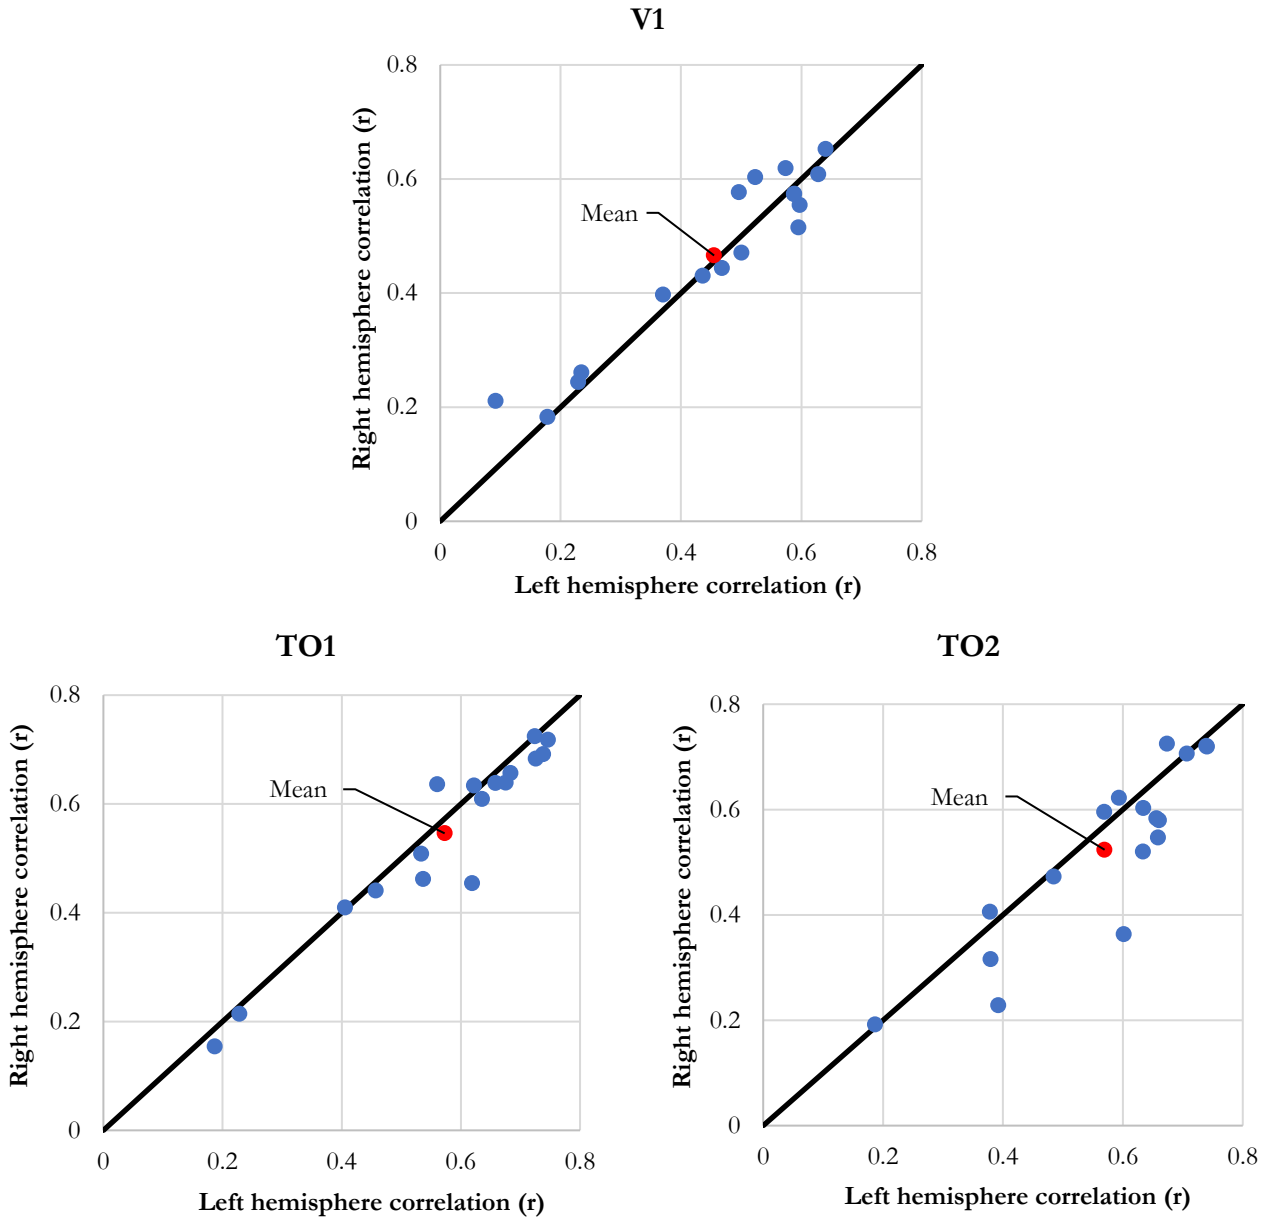

Figure 66 Individual subject timecourse consistency of the left (x-axis) and right (y-axis) hemisphere. Correlations were computed between each individual's visual area timecourse and the group average timecourse for that visual area (the group average timecourse did not include the brain with which the correlation was being computed). This analysis gives a measure of the consistency of the timecourses for each visual area. Group means represented by red markers. Black line represents the line of equality. Results show that for TO regions, in particular TO2/MST, left hemisphere timecourses are more highly correlated across brains than right hemisphere timecourses (more data points below the line of equality). Conversely, V1 timecourses are similarly correlated across hemispheres (data points are distributed more symmetrically relative to the line of equality). Fisher's Z equivalent figure can be found in Appendix P.

Next, we examined whether there was a significant difference in the consistency of the timecourses between hemispheres for these regions of interest[27]. As expected, timecourses extracted from the left hemisphere V1 ROIs were correlated (mean r = 0.455, SD = 0.173) across brains to a

[27] The same analyses were also performed on Fisher's Z transformed values and can be found in Appendix Q. Note, the transform did not alter the results presented here.

comparable degree to right hemisphere V1 timecourses (mean r = 0.466, SD = 0.156). This mean difference between hemispheres, 0.011, was not statistically significant, t(16) = -0.904, p = 0.379 (uncorrected).

On the other hand, the timecourses of TO1/MT in the left hemisphere were more consistent across brains (mean r = 0.573, SD = 0.168) than right hemisphere TO1/MT (mean r = 0.546, SD = 0.170). This difference, 0.027, was statistically significant, t(16) = 2.322, p = 0.034 (uncorrected). Likewise, the timecourses of left TO2/MST (mean r = 0.569, SD = 0.154) were more consistent than right hemisphere TO2/MST timecourses (r = 0.524, SD = 0.170). Again, this difference, 0.045, was statistically significant, t(16) = 2.454, p = 0.026 (uncorrected). These comparisons show that, on average, the timecourses extracted from the left hemisphere were more highly correlated across brains than timecourses from the right hemisphere.

This analysis gives us an index of how, on average, the individual timecourses compared to the group mean timecourse. To this end, the analysis suggests the timecourses of left hemisphere TO1/MT and TO2/MST are more comparable across brains, suggesting little individual differences in their functional role across brains. On the other hand, in the right hemisphere the timecourses are less correlated across brains on average. This lower correlation observed for right hemisphere timecourses suggests there are variations in the timecourses across brains, and this variation is higher relative to the left hemisphere.

Following the planned TO/V1 comparisons, we then performed a post-hoc and exploratory analysis in which we examined whether we find such differences anywhere else in the visual system. For each of the 25 visual area paired samples t-test were performed on the correlations of each hemisphere. In addition to TO1/MT and TO2/MST, we find significant differences between hemispheres for V2d ($p < .001$), V3v ($p = 0.027$), V3d ($p < .001$)[28, 29].

---

[28] V2d, V3v and V3d p values are uncorrected.
[29] The same a post-hoc and exploratory analysis was performed on Fisher's Z transformed values and can be found in Appendix R. Note, the transform did not alter the results presented here.

A key question was whether the differences in the consistency of the timecourses between hemispheres observed for TO1/MT and TO2/MST were significantly larger than those exhibited by V1. Although we find a significant difference between hemispheres for TO1/MT and TO2/MST ($p < .05$) and a nonsignificant difference between hemispheres for V1 ($p > .05$), it would be remiss of us to conclude that the discrepancy between hemispheres in the consistency of TO1/TO2 timecourses were significantly different to V1. Therefore, as conveyed in Nieuwenhuis et al. (2011), to draw appropriate comparisons between two regions (e.g., V1 vs. TO1/MT), we investigated the statistical significance of their differences, not just the differences between significance levels (e.g., $p = 0.379$ vs. $p = 0.034$). This analysis indicates whether the differences in the consistency of the timecourses across hemispheres were specific to those regions.

For V1, TO1/MT and TO2/MST, we calculated the difference between the two hemispheres for each subject-to-group timecourse comparison by subtracting the correlation in the left hemisphere from the correlation in the right hemisphere. For example, Subject 6's TO1/MT timecourse correlated with the group timecourse at $r = 0.619$ in the left and $r = 0.455$ in the right, resulting in a difference between hemispheres of 0.164 (note a positive difference indicates correlations were higher in the left hemisphere than the right). We then performed paired samples t-tests on these cross-hemisphere differences for each of the visual areas (Table 14). This analysis revealed the mean difference between the left and right hemispheres seen for TO1/MT ($M = 0.027$, $SD = 0.048$) and TO2/MST ($M = 0.045$, $SD = 0.077$) was significantly larger than the mean differences observed for V1 ($M = -0.011$, $SD = 0.05$) [30].

[30] The same analysis was performed on Fisher's Z transformed values and can be found in Appendix S. Note, the transform did not alter the results presented here.

Table 14. Results of paired-samples t-test performed on the differences between hemispheres in the timecourse consistency for V1, TO1 and TO2. For each subject we calculated the difference between hemispheres for the subject-to-group timecourse correlations. We then performed paired-samples t tests on these differences to ascertain whether the level of timecourse consistency in TO1 and TO2 differed significantly from the level of consistency in V1. p values are uncorrected.

Comparison	Mean	Std. Deviation	t	df	p (2-tailed)
V1 - TO1	-0.038	0.06	-2.599	16	.019
V1 - TO2	-0.057	0.096	-2.431	16	.027

Our results are in line with variations across brains in the function of right hemisphere hMT+, specifically TO2/MST. Timecourses are less consistent across brains in the right hemisphere compared to the left, and the differences between hemispheres is statistically significant. These lower correlations could reflect comparisons being drawn between two different functional processes, one in which both the ipsi- and contra-lateral visual field is processed ("enhanced" role hMT+) and one in which only contra-lateral visual field is being processed. Whereas in the left hemisphere, timecourses are more correlated across brains suggesting the region functions comparably across subjects.

In summary, when we examine each hemisphere individually, we find the timecourses of left TO1/MT and TO2/MST are significantly more correlated across brains than right hemisphere timecourses. We find no such discrepancy between hemispheres for V1. In the right hemisphere there are individual differences in the timecourses, while in the left the timecourses are more comparable across brains, suggesting little functional difference. Individual differences that impact one hemisphere and not the other are compatible with the area having a different functional role across hemispheres. Therefore, our results indicate that TO2/MST, and potentially TO1/MT, exhibit lateralised function for some subjects, and not for others.

4.5.3 Differences across brains

The analysis above suggests that in some brains TO2/MST, and potentially TO1/MT, have different functional roles across hemispheres. Given that, a logical progression is to identify which brains are different, and what proportion of our sample they make up.

To assess individual differences, we examined cross-hemisphere timecourse correlations in each brain individually. We find evidence that some subjects exhibit different activity patterns to the other subjects. These individual differences are compatible with lateralised function in some brains and not others. From this, we infer that ~30% of brains show evidence of lateralised function of TO2/MST.

For each subject we divided the cross-hemisphere mean timecourse (ROI-based) correlation for TO1/MT and TO2/MST (separately) by the correlation for V1 (mean of V1v and V1d). This generated a ratio for each subject. A ratio > 1 indicates the TO1/TO2 timecourse was more highly correlated across hemispheres than the V1 timecourse. A ratio < 1 indicates the reverse (TO1/TO2 timecourse was less correlated across hemispheres than V1). For example, Subject 6's V1 timecourses correlated across hemispheres at $r = 0.932$, whereas the timecourses of TO2/MST correlated across hemispheres at $r = 0.712$. Dividing the TO2/MST correlation by the V1 correlation results in a ratio of 0.765, indicating the V1 timecourse was more highly correlated across hemispheres than the TO2/MST timecourses. We categorised the V1 and TO1/TO2 cross-hemisphere timecourses correlations as being notably different from one another based on a ratio threshold of 1 ± 0.1. Table 15 displays these individual subject ratios.

Table 15. The individual subject ratios of the cross-hemisphere timecourse correlations of TO1 and TO2 to V1. For each subject, the ratios were calculated by dividing the cross-hemispheric correlation coefficient for TO1 and TO2 (separately) by the cross-hemisphere correlation coefficient for V1 (mean of V1v and V1d). Values > 1 indicate the timecourses of TO1/TO2 were more highly correlated across hemispheres than V1, values < 1 indicate the reverse. A threshold of 1 ±0.1 was employed to indicate notable differences in the cross-hemisphere timecourse correlations of V1 and TO1/TO2. Based on this threshold, our results show that for 5 brains the timecourses of TO2 were less correlated across hemispheres than the timecourses of V1 (indicated by the red borders).

Subject	TO1 vs. V1	TO2 vs. V1
6	0.925	0.765
7	0.995	0.960
8	1.025	0.844
9	1.088	1.080
10	1.004	1.028
11	0.942	1.027
12	0.988	0.941
13	1.012	0.926
14	0.949	0.844
15	1.000	0.850
16	0.983	0.845
17	0.945	1.013
18	0.958	0.912
19	1.001	0.949
20	1.037	1.050
21	0.998	0.953
22	0.991	0.958

For all subjects, TO1/MT timecourses were as similarly correlated across hemispheres as V1 timecourses, indicated by ratios close to 1. This suggests, on the whole, TO1/MT has a similar functional role across hemispheres. More specifically, TO1/MT is no more functionally lateralised than V1. However, for 5/17 subjects we find TO2/MST timecourses are less correlated across hemispheres than V1 timecourses (ratios above the 1 ±0.1 threshold). These differences in the cross-hemisphere timecourse correlations are compatible with TO2/MST having a different functional role across hemispheres for some individuals, in our case ~30% of brains.

4.6 Correlations relative to V1

A potential limitation of all the analyses presented above is that they involve taking an average over the entire ROI, whether that be averaging the signal of an area to generate a mean timecourse

(ROI-to-ROI) or generating a mean correlation from vertices within an area (mean vertex-to-vertex). Both approaches could, in theory, result in differences between hemispheres in the sub-regions of an area being lost, as signal/correlations are averaged over large areas. For example, in the primate brain, MST is composed of two functionally distinct regions. The dorsal medial (MSTd) region is responsive to complex optic flow patterns, suggesting it plays an important role in heading perception (Tanaka et al., 1986). In contrast, the lateral ventral (MSTl) region is more responsive to smaller stimuli, suggesting it is involved in the processing of object motion (Eifuku & Wurtz, 1998). It is possible one of these sub-regions are lateralised in the human brain, and one is not. However, if an average is taken across the entire TO2/MST ROI, any differences will be absorbed by the mean. One way to explore this issue is to examine correlations on a per-vertex basis. I detail this approach below.

4.6.1 Group-mean vertex correlation map

A group-mean vertex correlation map was generated by calculating, for each vertex, the group-mean cross-hemisphere correlation. To reveal highly correlated areas, I used the V1 correlation as a baseline: the group mean r value of V1 (dorsal and ventral regions combined, $r = 0.667$) was subtracted from each vertex group-mean correlation coefficient. These values can then be displayed on the surface and are therefore not confined by the ROIs from which we have previously extracted values from (Figure 67, left). Areas with $r > 0$ are represented by red regions, areas with $r = \sim 0$ are represented by grey regions, and areas with $r < 0$ are represented by blue regions.

Figure 67 reveals, on average, vertices within the anterior region of TO1/MT are more highly correlated than posterior vertices, suggesting that taking the mean across the entire ROI may mask important differences between hemispheres in the sub-regions. On average the vertices within TO2/MST are less correlated across hemispheres than TO1/MT, but it does not appear that the ROI covers multiple distinct regions as vertices within the ROI are more uniformly correlated. However, the map represents the average brain, and therefore it is possible that individual differences in the function of the regions, and/or sub-regions, are absorbed when taking the mean across brains at each vertex. In addition, the standard deviation map (Figure 67, right), shows the vertices within the

TO1/MT ROI that are, on average, more highly correlated also show a larger standard deviation. Consequently, the higher correlations seen in anterior TO1/MT could be a result of a small number of brains driving the mean correlation up, therefore a reasonable progression from this analysis was to explore individual-subject data.

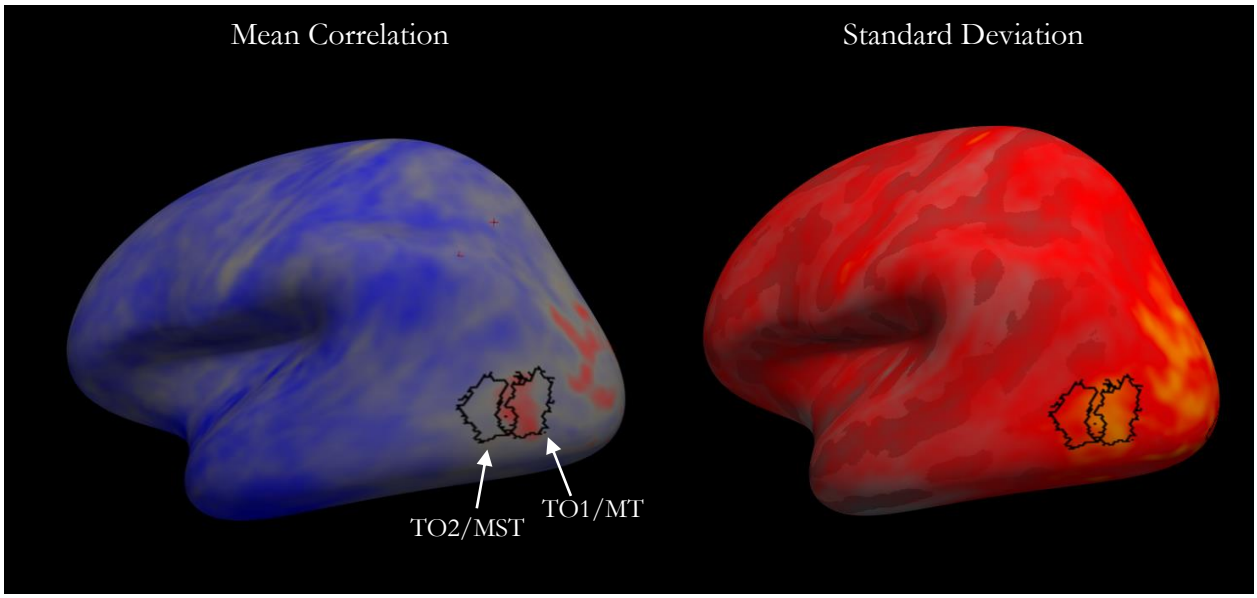

Figure 67. Group mean vertex correlation map displayed on the left hemisphere (left). The map was generated by taking the mean correlation across brains at each vertex. The group mean r value of V1 was subtracted from each vertex correlation coefficient to reveal areas more/less correlated than V1. "Hotter" colours represent higher r values, indicating the vertex was more highly correlated across hemispheres. The standard deviation map (right) was generated by replacing each vertex mean correlation value with the standard deviation of the mean for that vertex. Lower standard deviations are represented by transparent-red vertices, higher standard deviations are represented by to opaque-yellow vertices. TO1/MT and TO2/MST exclusive ROIs are also displayed.

4.6.2 Individual-subject vertex correlation map

At the group-level we find evidence to suggest there are differences between hemispheres in the sub-regions of the TO1/MT ROI, whereas vertices within the TO2/MST ROI are uniformly less correlated across hemispheres than TO1/MT and V1. However, given the results of our analyses thus far point towards differences between brains in the function of TO1/MT and TO2/MST, we next explored the individual data. To do this, individual-subject vertex correlation maps were generated, this time subtracting the corresponding subject's V1 correlation coefficient (rather than the group mean) from each of their individual vertex correlation values.

The individual-subject maps served two purposes. First, they allowed me to assess whether the individual data reflects the group correlation map, i.e., do we find differences between hemispheres in sub-regions of the ROIs? Second, they allowed me to assess whether there was evidence of a

misalignment of the TO1/MT and TO2/MST ROIs. As mentioned in previous chapters, the use of an atlas has inherent limitations, especially when individual data is warped to a standard surface. An atlas cannot reflect all the anatomical variation in the location of the areas across brains; it is by its nature, probabilistic. Therefore, when signal or correlations are averaged across a ROI, there is always the possibility the ROI is misaligned relative to the true location of the area. The individual-subject maps helped us to identify potential variability in anatomical location of the regions resulting in the atlas being misplaced. For example, do we find a region of vertices with low correlation just outside the ROI?

Once the individual maps were generated for each subject, they were compared to the results obtained from the ROI-to-ROI mean timecourse analysis (specifically the ratios reported in Table 15). In general, the vertex correlation maps correspond with the mean timecourse results, indicating converging results from the analyses. For example, when a subject's mean timecourse of TO2/MST was more highly correlated across hemispheres than the V1 timecourse, vertices within the region also show high cross-hemisphere correlations.

The individual-subject maps show a large amount of variability across brains. They do, however, reflect the findings at the group-level to some extent: there are differences between hemispheres in the sub-regions of the ROIs, i.e., the entire ROI is not uniformly correlated across hemispheres. However, differences within TO1/MT are not evident for all brains. Instead, we find there are differences between brains as to which regions, and which sub-regions, are more/less correlated across hemispheres than V1 (I provide some example vertex correlation maps in Figure 68). Indeed, we find for some subjects the timecourses of vertices within the TO2/MT ROI are more highly correlated than V1. This pattern was evident in the group-level map.

Figure 68 shows four examples of the individual-subject maps. Each example subject was selected based on their differing results from the ROI-to-ROI mean timecourse analysis (see Table 15 in section 4.5.3):

1) Subject 6 – Mean timecourse results showed TO1/MT and TO2/MST were less correlated across hemispheres than V1 (TO2/MST less correlated than TO1/MT).

2) Subject 9 – The mean timecourses of both TO1/MT and TO2/MST were similarly correlated across hemispheres as V1 timecourses.
3) Subject 11 – Mean timecourse of TO1/MT was less correlated across hemispheres than V1, whilst mean timecourse of TO2/MST was slightly more correlated across hemispheres that V1.
4) Subject 15 – TO1/MT timecourses were equally correlated across hemispheres as V1, TO2/MST timecourses were less correlated across hemispheres than V1.

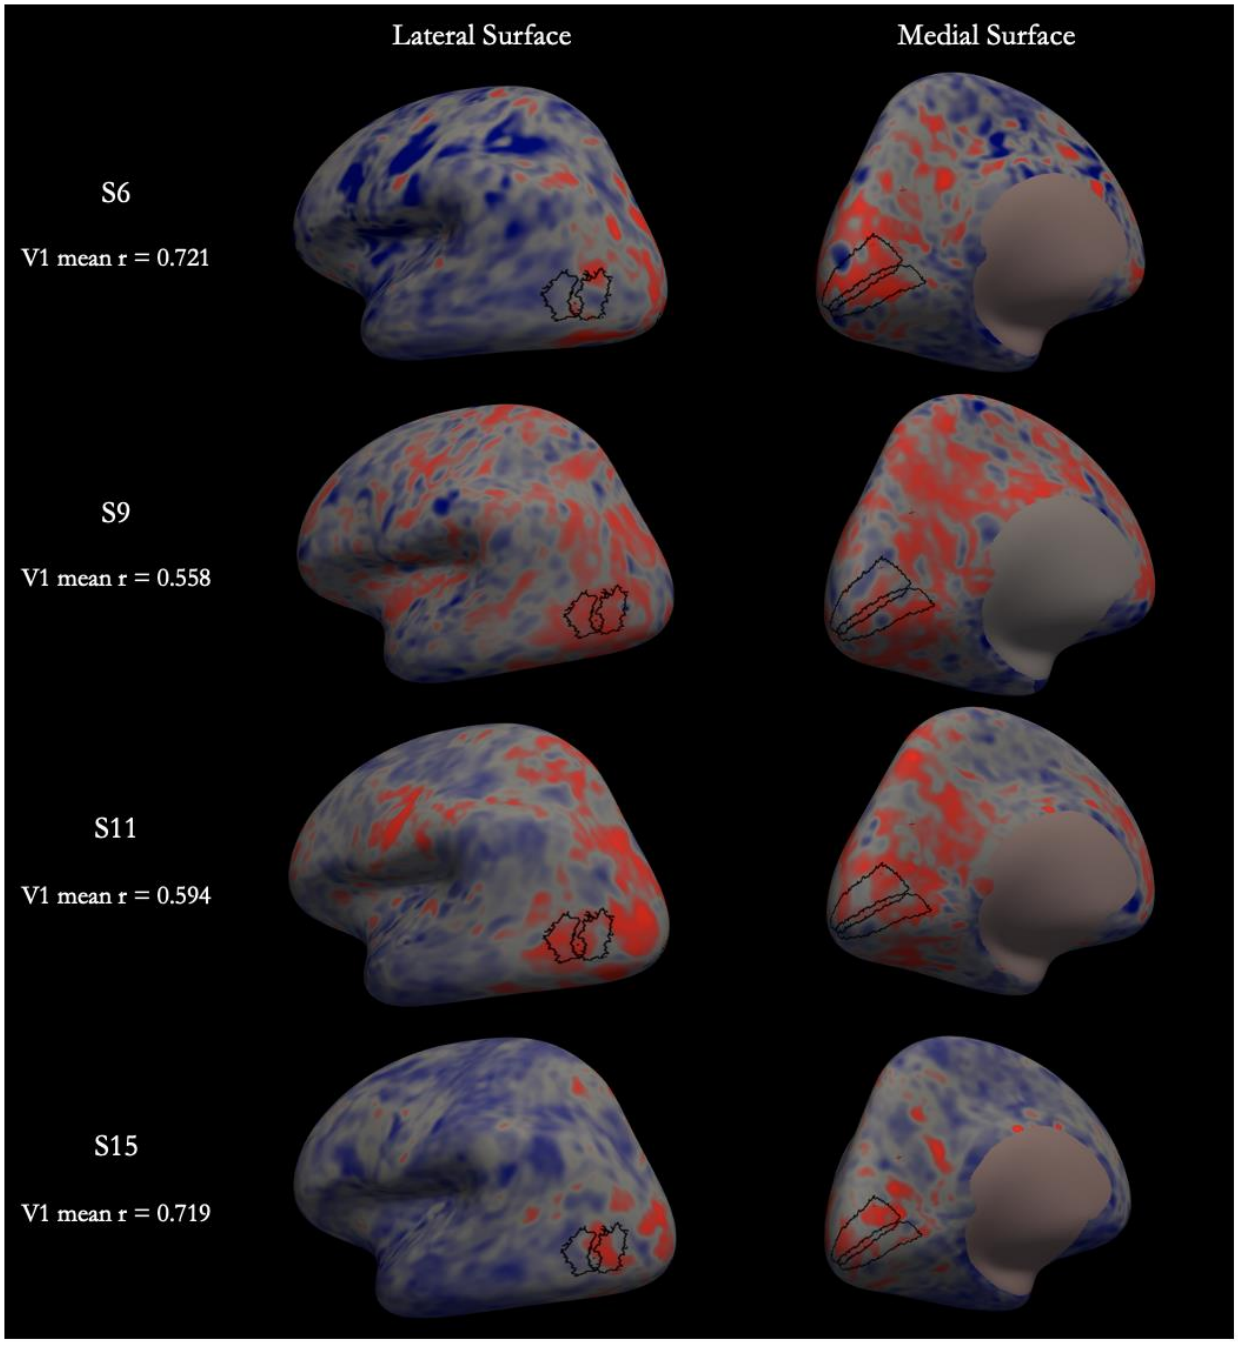

Figure 68. Four examples of individual-subject vertex correlation maps, displayed on the left hemisphere. Each subject's mean r value of V1 was subtracted from each vertex correlation coefficient. Areas with $r > 0$ are represented by red regions (more highly correlated than average V1), areas with $r = \sim 0$ are represented by grey regions, and areas with $r < 0$ are represented by blue regions. The V1v, V1d, TO1 and TO2 exclusive ROIs are also displayed.

Subject 15's map closely resembles the group-level map in that the anterior region of TO1/MT is more highly correlated than the posterior. However, this pattern is not consistent across brains. For example, Subject 11's map differs substantially from the group-level map, TO2/MST is more highly correlated than TO1/MT. The highly correlated vertices extend anteriorly beyond the ROI, which could indicate either TO2/MST or the hMT+ complex is located more anteriorly than the ROIs suggests.

Judging ROIs as being misplaced is potentially problematic since we do not have access to a "ground truth" of the anatomical location of the area. However, when visually inspecting the vertex correlation maps, we do find evidence in line with misalignment in some brains. For example, Subject 16's map (Figure 69) shows a patch of highly correlated vertices within TO1/MT and a second patch of vertices located on the anterior border of TO2. This could indicate the true location of TO2/MST is more anterior than the atlas ROI. However, we acknowledge this approach is subjective and therefore future research may opt to identify the regions using traditional localisers before performing cross-hemispheres timecourse correlations of movie watching data.

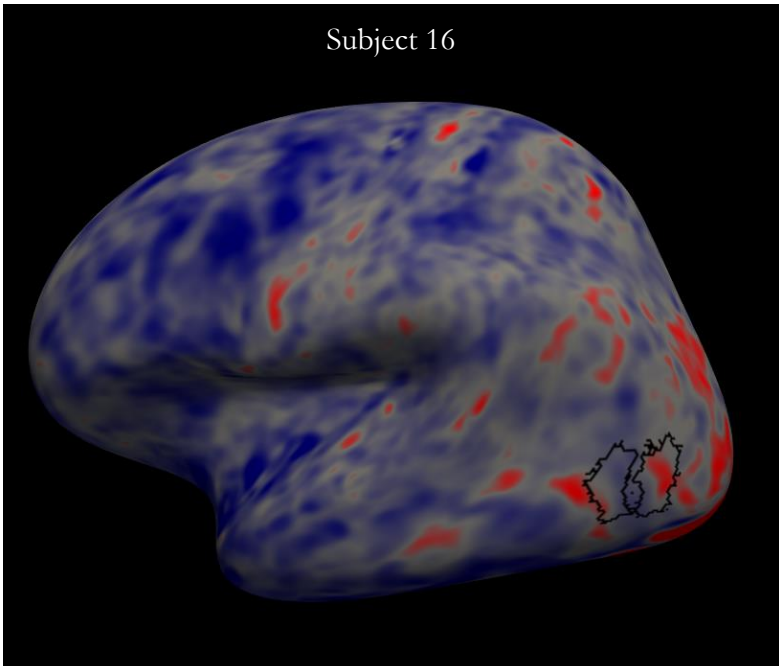

Figure 69. An example subject's vertex correlation map showing two areas of highly correlated vertices: one area is located within the TO1/MT ROI, another is located on the anterior border of the TO2/MST ROI. The latter may indicate the true anatomical location of TO2/MST is more anterior than the ROI suggests.

In summary, the example maps presented in Figure 68 highlight that the results obtained from the ROI-based mean timecourse and the vertex-based approaches are robust, e.g., if a subject shows

lower timecourse correlation in one region than the other, the pattern is evident across analyses. Moreover, there is, again, evidence of individual differences in the function of TO1/MT and TO2/MST across hemispheres: the individual-subject vertex correlation maps vary across brains.

4.7 Interim summary

The results of the timecourse correlation analyses presented above suggest that TO1/MT is no more lateralised than V1, and in fact appears to be more bilateral in its function compared to other visual areas, e.g., TO1 was the 2nd (ROI-based analysis) and 1st (vertex-based analysis) most highly correlated region. This may simply reflect the image statistics, i.e., motion is more highly correlated across hemifields than other features. I discuss this point later in this chapter. Conversely, we find evidence that TO2/MST may be lateralised in some brains, offering partial support to the lateralisation hypothesis.

First, I performed an ROI-to-ROI correlation analysis in which the mean timecourse of vertices within a given ROI were extracted from each hemisphere. We then computed correlations between these corresponding timecourses. We found TO1/MT was comparably correlated across hemispheres to V1, whilst TO2/MST was less correlated across hemispheres than V1 and TO1/MT. This may indicate a difference in TO2/MST function across hemispheres.

I ruled out the possibility that differences in the sizes of the visual areas were obscuring underlying differences between hemispheres. Rather than extracting the mean timecourse of all vertices within a given region, I performed correlations at the vertex level and took the means. Again, our results reflect those obtained by the mean timecourse analysis. TO1/MT was more highly correlated than V1, TO2/MST was less correlated than V1 and TO1/MT.

The results from both our cross-hemisphere (ROI and vertex) analyses are robust. The overall pattern of the results did not change irrespective of the analysis pipeline used (e.g., volume-smoothed or surface-smoothed; see Appendix IX) or the type of ROI from which values were extracted (intersection, exclusive or max probability; see Appendix VIII).

I then assessed the potential contribution of global signals to the cross-hemisphere correlations and to rule out any confounding role. First, I examined global signal changes from a broad perspective by computing correlations between the timecourses of distant brain regions that do not have obvious functional connections. I estimate that global signal change across the brain account for a correlation, r, of ~0.35. Importantly, this level of correlation is consistent across the brain. Next, I looked specifically at the contribution of global signal changes due to head movements, which could have an un-equal impact across the brain, by denoising our data with ICA-AROMA. Our results remained essentially the same for both denoised and raw data, signifying that head movements do not have a disproportionate impact on TO areas.

We then examined how, on average, the individual subject timecourses compared to the group mean timecourse. The analysis showed the timecourses of left hemisphere TO1/MT and TO2/MST were more comparable across brains than right hemisphere timecourses. Taken together, it appeared there were differences across brains in the timecourses of right hemisphere TO2/MST that were not evident in the left hemisphere. This raised a question - is TO2/MST lateralised in some brains and not others?

We further explored whether individual differences in the function of TO2/MST were underpinning the lower correlations relative to TO1/MT and V1. We took the ratio of V1 cross-hemisphere correlations to TO1/MT and TO2/MST correlations. We found for ~30% of brains TO2/MST timecourses were less correlated across hemispheres than V1 timecourses, suggesting in these brains TO2/MST was lateralised. On the other hand, the ratios of V1 timecourses to TO1/MT timecourses were not as pronounced.

To ensure I was not being misled by the accuracy of the boundaries of the ROIs and that the results did not obscure sub-region differences, I examined the mean correlations on a per-vertex basis by displaying the values on the surface. Since I was interested in lateralisation of TO1/MST and TO2/MST relative to V1, I subtracted the mean cross-hemisphere correlation of V1 from each vertex correlation value to reveal areas of high/low correlation relative to the baseline. This analysis allowed us to also visually inspect the correlations without the need to extract the values from the visual area

ROIs. If there was some misplacement of the ROIs, but regions were functioning unilaterally the surface overlay would reveal this.

The group-mean vertex correlation map (discussed in section 4.6.1) reveals some differences within the TO regions. The anterior region of TO1/MT is more highly correlated than the posterior region, therefore taking the mean across the entire ROI may mask important differences between hemispheres in the sub-regions. On the other hand, TO2/MST was uniformly less correlated across hemispheres than TO1/MT and V1, adding weight to our assertion that TO2/MST may have a different role across hemispheres. However, the group-level map is essentially a summary of the vertex-to-vertex correlation results. Given our analyses up to that point suggested that there are individual differences in the function of the region across hemispheres, individual-subject maps were then generated.

The individual-subject maps show a large amount of variability across brains, but importantly they reflect the corresponding subject's results obtained in both the ROI- and vertex-based analyses. Based on the visual inspection of the correlation maps, we find evidence to suggest that in some brains the atlas ROIs may be slightly misplaced relative to the true location of TO1/TO2, highlighting an important potential limitation of using probabilistic atlases.

In summary, the converging results of each analysis lead to the following conclusions:

1. Timecourses of TO1/MT are comparably correlated across hemispheres to V1, which I treat as our baseline for bilateral function.
2. While still highly correlated, the timecourses of TO2/MST are less correlated across hemispheres than TO1/MT and V1.
3. There are variations across brains in the timecourses of right hemisphere hMT+, specifically TO2/MST, that are not evident in the left hemisphere.

These findings are compatible with individual differences in the function of hMT+ across hemispheres. In the proceeding section I performed ICA on our movie watching data to further investigate the potential lateralised function of TO1/MT and TO2/MST.

4.8 ICA

So far in this chapter I have investigated the potential functional lateralisation of TO1/MT and TO2/MST by performing cross- and within-hemisphere correlations, arriving at a tentative conclusion that there are individual differences in the function of hMT+ across hemispheres. Next, I sought to address the question of lateralised function of hMT+ from an entirely different direction and assess whether I arrive at the same conclusion. An alternative approach is to perform ICA on our movie watching data and examine differences between hemispheres in the output components. ICA is advantageous in that it does not make any assumptions about which areas of the brain should have correlated activity; it segments the brain in a more organic manner. We then identify components that show anatomical correspondence with the atlas regions (through spatial correlations), rather than manually identifying components based on where we believe the true location of TO1/TO2 to be. Therefore, the issues I have flagged previously about the differences between brains in the exact location of TO1/TO2 are less of an issue with this analysis. The following sections detail this approach.

4.8.1 Methods

Details of how data were processed and how component matches were identified are described in Chapter 3. For ease, I summarise the approach below.

We performed single-subject ICA and computed correlations between each subject's set of spatial maps and a probabilistic atlas of visual areas. Components which correlated with the atlas regions above the $r > 0.3$ threshold were explored.

First, for each hemisphere, I examined the mean number of above threshold matches. More matches indicate that the region is represented in multiple components. This is particularly important for investigating lateralised function, as if there are differences between hemispheres in the visual field coverage of hMT+, we would expect an unequal number of components that include left and right hemispheres. The logic behind this assumption is discussed in detail below.

Next, we identified the component showing the highest correlation to that area (the "top match"). We then took the mean correlation of the top matches across all subjects. This reveals the spatial correspondence between a component and an atlas region. If a region is represented by multiple fragmented components, correlations between the components and the atlas may be lower than if a single component represents the entire area (see Figure 70 for schematic example of this logic).

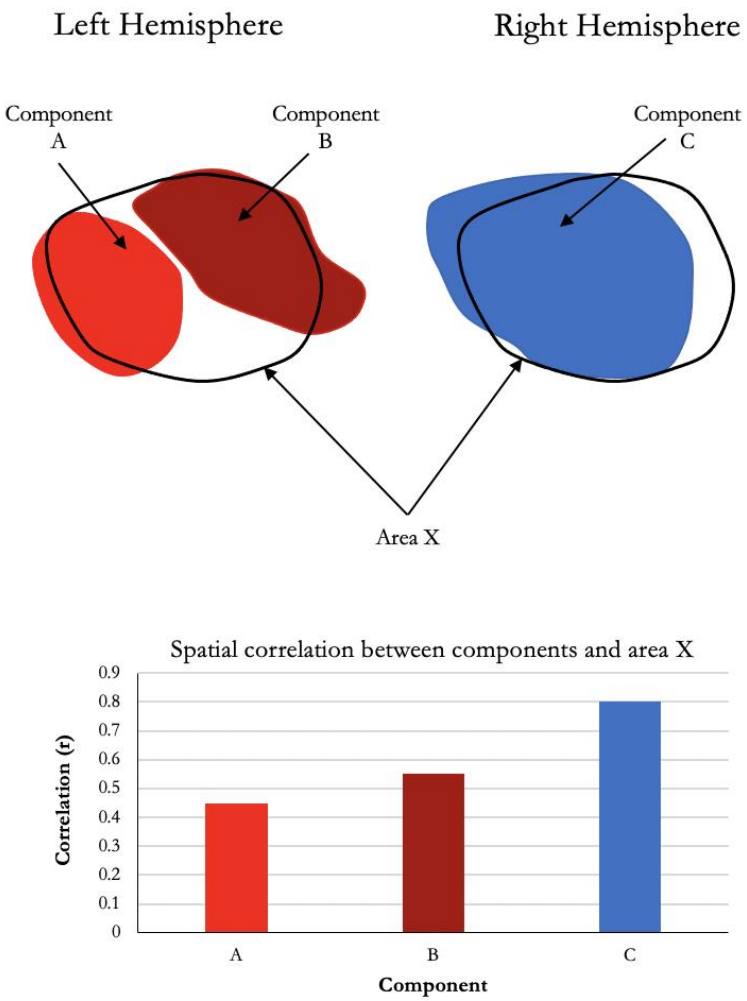

Figure 70. A schematic example of how a region represented by multiple fragmented components (components A and B) may result in lower spatial correlations with the corresponding atlas region (area X) than if a single, more complete component was produced (component C).

Finally, we calculated how often the top matched component was the top match to only one visual area i.e., how often the top match was unique. A unique match to TO1/MT or TO2/MST would indicate the regions segment out from one another, and differences across hemispheres in the uniqueness of matches could indicate differences in the functional role of the region.

4.8.2 Results

It is difficult to draw any clear distinctions between TO1/MT and TO2/MST because neither region is identified independently in our ICA analysis (see Chapter 3). This suggests it may be more appropriate to investigate differences between brains in the whole motion complex instead. The group average results of the three analyses (presented in Table 16) confirm that the TO1/TO2 division is not suitable with ICA. We find values are slightly lower in the left than the right hemisphere for both TO regions. However, the results suggest that at the group level there is little difference between hemispheres.

Table 16. The group mean values for three measures used to assess differences between hemispheres in the output of components derived from ICA. 1) The mean number of components correlating with the TO1/MT and TO2/MT atlas regions above the r = 0.3 threshold. 2) The mean across subjects of correlation between the top matched components and the atlas regions. 3) The percentage of components that were the top match to only one visual area (i.e., how often the top match was unique). For all three measures, we find values are slightly lower in the left than the right hemisphere for both TO regions. However, the results suggest that at the group level there is little difference between hemispheres.

	Mean above threshold matches				Mean correlation between top matched component and atlas				Top matches unique	
	Left		Right		Left		Right		Left	Right
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	%	%
TO2	1.529	0.514	1.353	0.606	0.614	0.095	0.563	0.123	11.765	5.882
TO1	1.706	0.985	1.353	0.606	0.605	0.130	0.551	0.101	5.882	0

It is appropriate that we qualify the question we are trying to address with this ICA approach. This qualification involves two steps.

First, we are no longer looking for differences between hemispheres in TO1/MT or TO2/MST individually, but rather the whole motion complex (hMT+), which we find is consistently identified across brains in our pipeline (see Chapter 3).

Second, to investigate if hMT+ is lateralised in some brains and not others, I inspect each subject's set of spatial maps. If hMT+ has a different functional role across hemispheres in some brains and not others, we would expect to see differences between these brains in the output of

components spatially overlapping with hMT+ (see below). It is therefore the differences in the spatial maps between subjects that is the important distinction to make for this analysis.

If hMT+ in each hemisphere processes the corresponding contra-lateral visual field only (e.g., there is no difference in role across hemispheres), we would expect an equal number of components that include left and right hemispheres. For example, we might expect ICA to output:

i. a bilateral component that represents the commonality of the visual fields, i.e., elements within the scene are correlated across the two visual hemifields. For example, consider the full field motion stimulation when looking out of a window on a train. hMT+ in each hemisphere is therefore exposed to the same/similar stimulus and responses are similar across hemispheres,

and/or

ii. a unilateral component in each hemisphere which represents the differences in the scene between the two visual hemifields. For example, consider a scene in which two characters are interacting in the right visual field, whilst stationary objects are present in the left hemifield.

If right hMT+ processes ipsi- and contra-lateral visual fields and left hMT+ processes only the contra-lateral visual field, then we would expect an unequal number of components that include left and right hemispheres, with the additional component(s) sitting in the hemisphere that has coverage of both visual fields. For example, we might expect ICA to output:

i. a bilateral component that represents the right hMT+ processing the ipsi-lateral (right) visual field and left hMT+ processing the contra-lateral (right) visual field,

and

ii. an additional unilateral component that represents right hMT+ processing the contra-lateral (left) visual field.

With these characteristics in mind, we examined the spatial maps of each subject's above threshold matches, finding evidence of lateralised function in 4/17 (24%) brains. Note, of these four brains, three showed higher cross-hemisphere correlations for V1 than TO2/MST when the ratio was

calculated in the earlier analysis (see Table 15 for reference), suggesting some convergence between approaches. However, our findings are not consistent with Strong et al. (2019) in that it is not always right hemisphere hMT+ that shows lateralised function in our data.

Most subjects (11/17 – 65%) produced a single bilateral component with activation of equal magnitude present in both hemispheres (for example, Figure 71), suggesting function of hMT+ was not lateralised in these brains. Note the component extends beyond the expected boundaries of the TO1/TO2 ROIs; this is common to all bilateral hMT+ components.

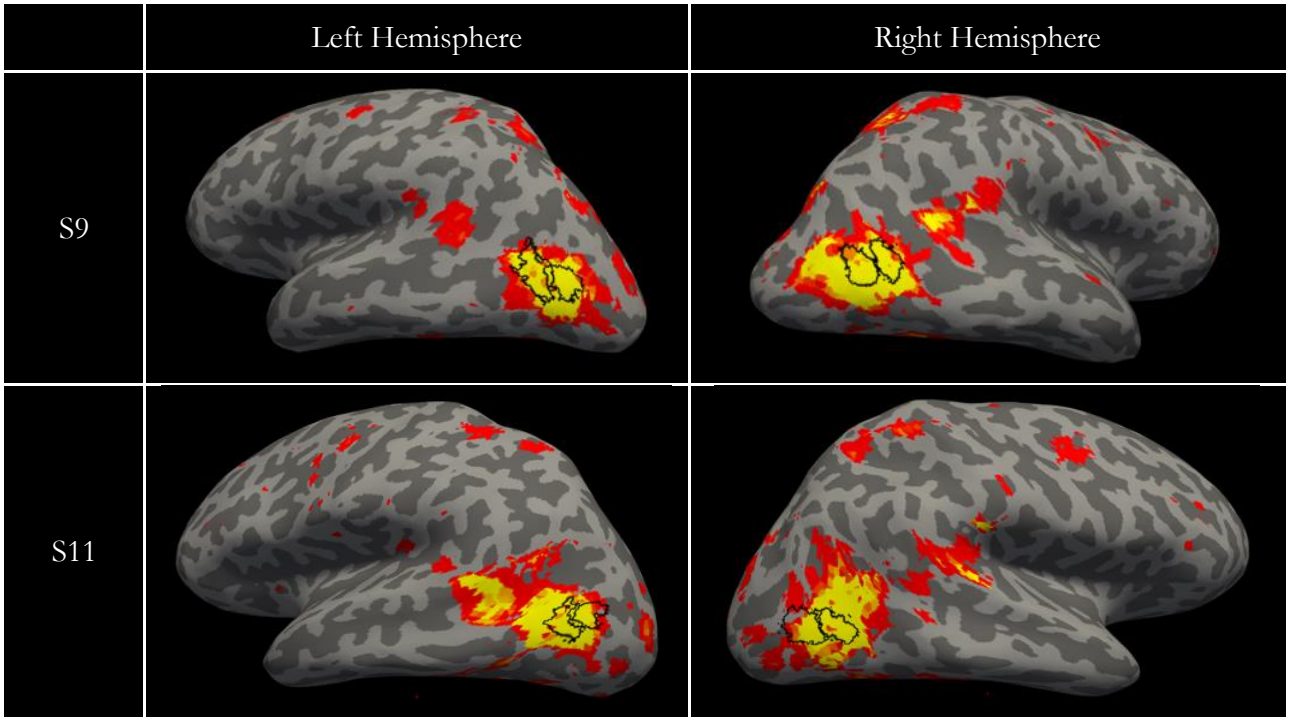

Figure 71. Two examples of TO1/TO2 matched components. The spatial maps of these components show strong bilateral activation in hMT+. The probabilistic atlas ROIs of TO1/MT and TO2/MST are overlaid (black outlines). The components indicate that hMT+ was not lateralised in these brains.

For two subjects we find a unilateral component is produced for each hemisphere with no clear bilateral components being produced (Figure 72). This suggests TO1/TO2 are functioning independently across the two hemispheres. These unilateral components show activation predominantly in one hemisphere. For both subjects, the right hemisphere top matched component has stronger activation in the opposite hemisphere compared to the left hemisphere top matched component. However, in both instances the magnitude of the activation in the non-matched hemispheres is much lower than seen for the bilateral components discussed above.

Given in both instances above (a single bilateral or two unilateral components are produced), we find an equal number of components in the two hemispheres, we infer these subjects as not displaying lateralised function of hMT+.

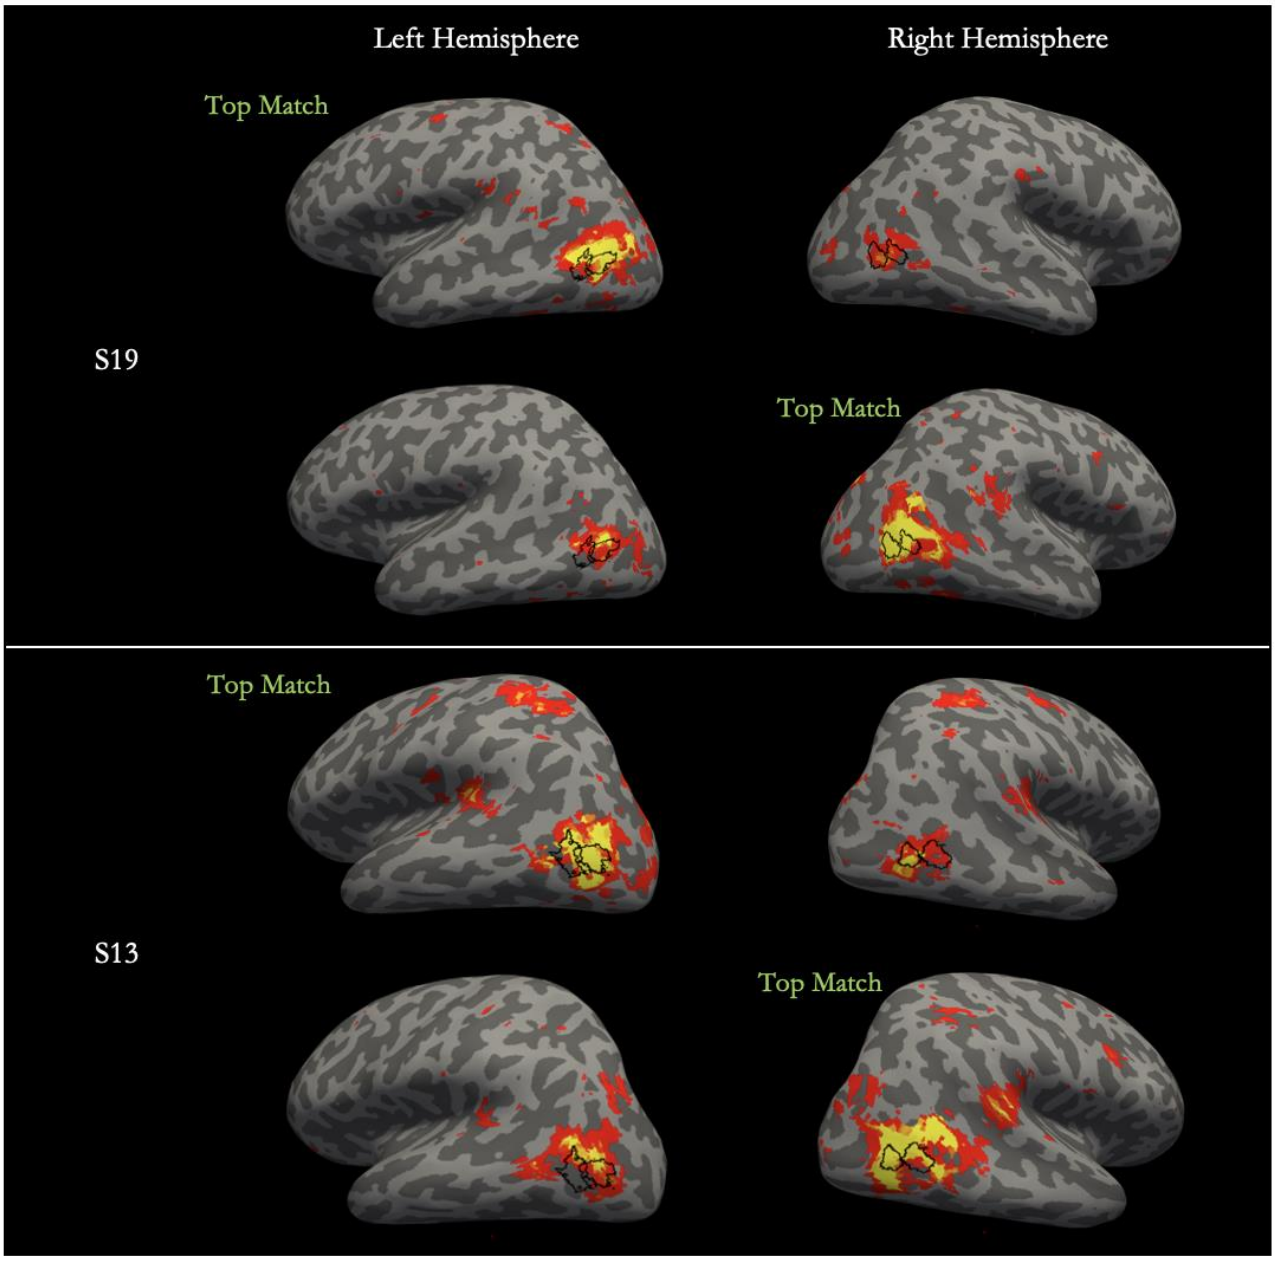

Figure 72. Two subjects who produced a unilateral top match component to TO1/TO2 for each hemisphere and did not produce a clear bilateral component. Unlike the bilateral components which show activation of equal magnitude in both hemispheres, the spatial maps of these unilateral components show activation predominantly in one hemisphere. We do note some residual activation in the opposite hemisphere of these unilateral components, but it is at a much weaker magnitude than is seen in the bilateral components. TO1/MT and TO2/MST ROIs are also displayed (black outlines).

Four subjects produced a bilateral hMT+ component and an additional unilateral component (Figure 73). These are the only subjects to produce an unequal number of components that included the left and right hemispheres. Based on the principles of ICA outlined above, the production of a

bilateral component *and* a unilateral component is compatible with hMT+ having a different functional role across hemispheres. We infer the hemisphere containing the additional hMT+ component as having coverage of both the contra- and ipsi-lateral visual fields. For half of the brains (Subjects 8 and 10), the additional component shows good and equal coverage of both TO1/MT and TO2/MST. For the other brains, activation is stronger in TO1/MT than TO2/MST (Subjects 12 and 14). For 3/4 brains, the unilateral component was located in the left hemisphere, indicating that for the majority of subjects it is left hemisphere hMT+ that has coverage of both visual hemifields. This is inconsistent with Strong et al.'s (2019) conclusion that right hemisphere hMT+ processes both the contra- and ipsi-lateral visual fields, while left hemisphere hMT+ processes the contra-lateral visual field only. However, of the four subjects that showed evidence of lateralised function with ICA, three also showed lower cross-hemisphere timecourse correlations for TO2 than V1. Further, two of the subjects were above the 1 ±0.1 threshold we employed to indicate notable differences in the timecourse correlations (see Table 15).

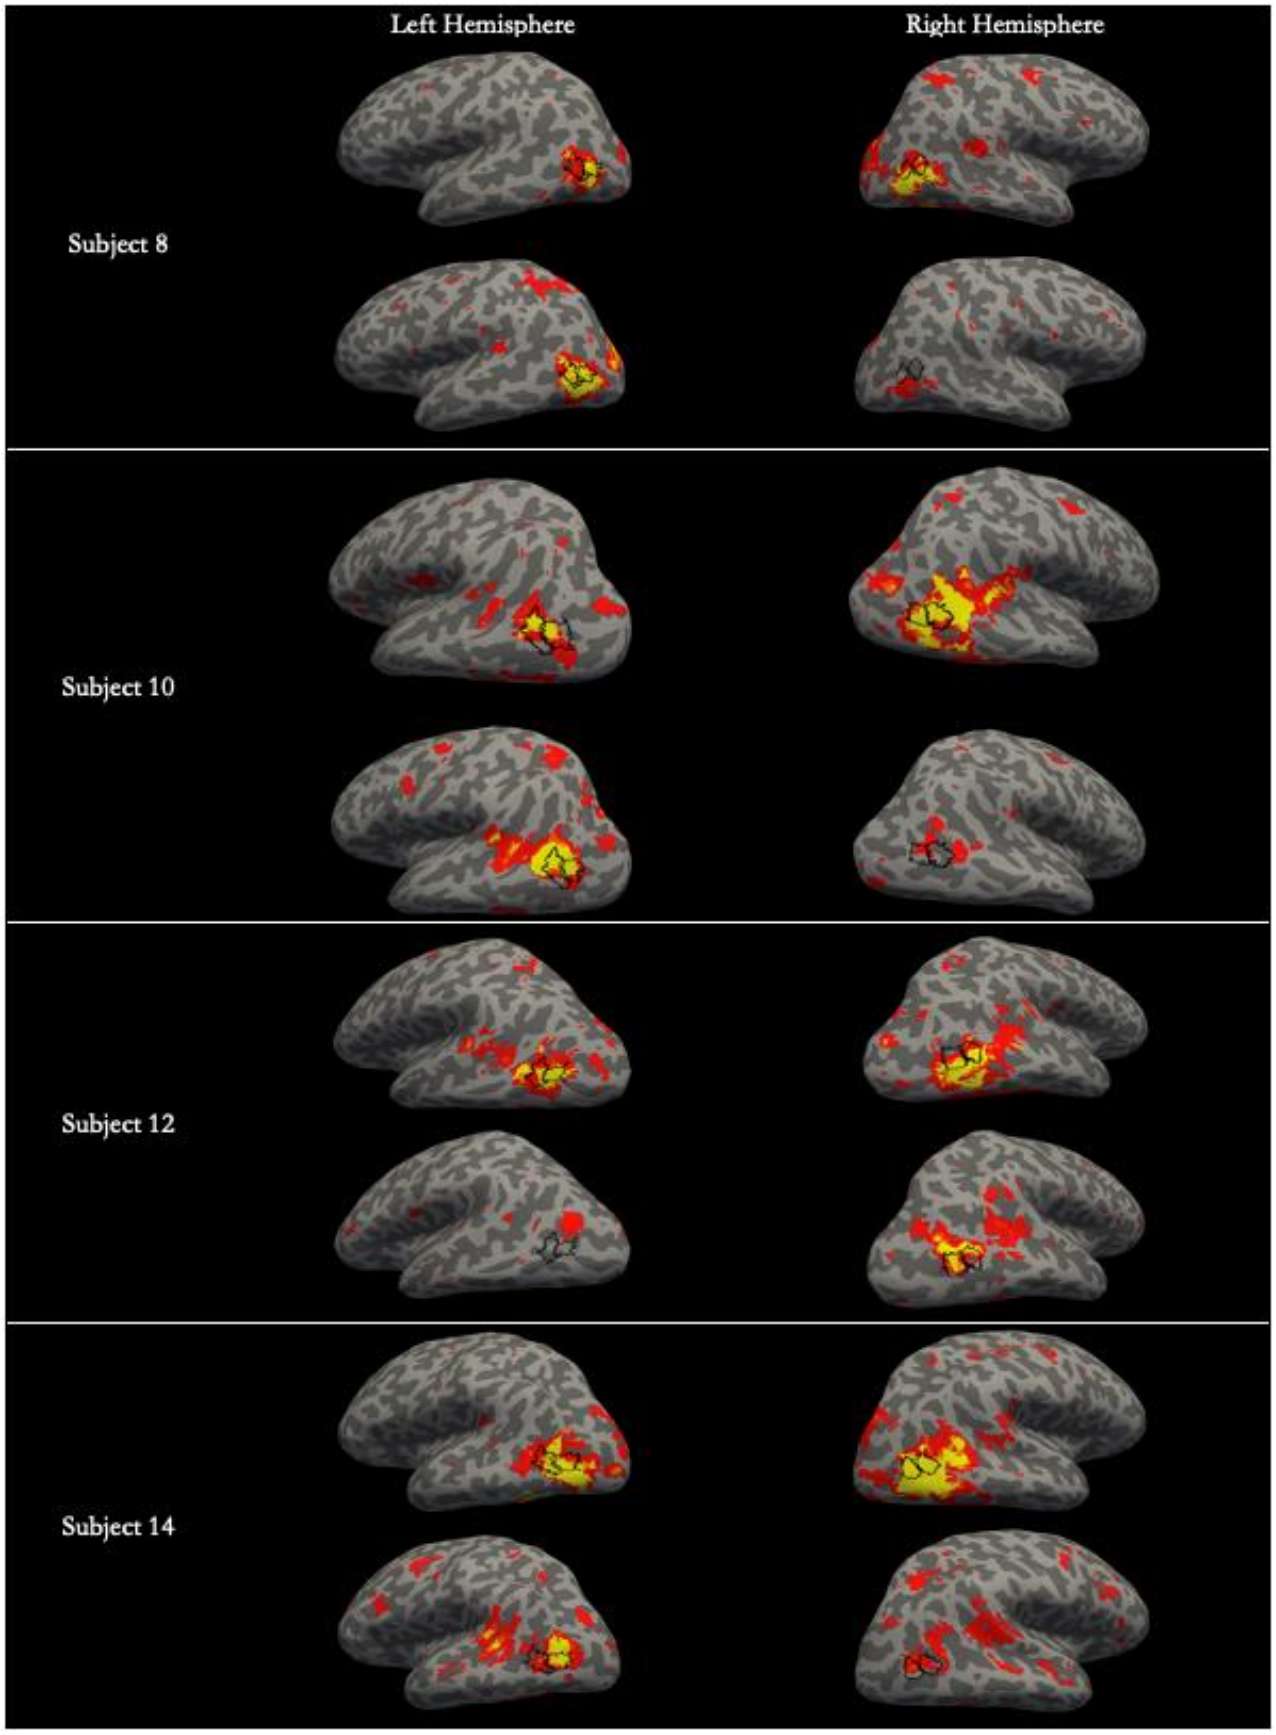

Figure 73. Four subjects showing lateralised function of hMT+. These subjects produced a bilateral hMT+ component with strong activation in TO1/TO2 that was comparable across hemispheres. The subjects also produced an additional unilateral component that showed strong activation in one hemisphere, with much weaker activation in the opposite hemisphere. For 3/4 brains, the additional unilateral component is located in the left hemisphere. In contrast to Strong et al.'s (2019) findings which stated right hemisphere hMT+ has an enhanced role, our results suggest left hemisphere hMT+ had an enhanced coverage of the visual field in 3/4 brains.

4.9 Chapter summary

Strong et al. (2019), in line with earlier investigators (Thakral & Slotnick, 2011) found differences in the function of hMT+ in the left and right hemispheres. TMS applied to either TO1/MT or TO2/MST in the left hemisphere induced deficits in motion perception of stimuli presented in the contra-lateral visual field only. Conversely, TMS applied to the regions in the right hemisphere induced deficits for both contra- and ipsi-lateral visual field stimuli, suggesting an enhanced role for right hemisphere hMT+. This is an interesting and novel hypothesis and using naturalistic stimuli and different analysis techniques, I sought further evidence to support it .

The results from each analysis performed in this chapter point to the same conclusion; there is some lateralisation of function of hMT+ but it is not a universal phenomenon. When group-level analyses were performed, TO1/MT and TO2/MST do not stand out from any other region with respect to clear hemispheric differences, although we note correlations are lower in TO2/MST than TO1/MT and V1 (see sections 4.3 and 4.4). However, when we systematically compared the timecourses of each individual to the group average timecourse, it became apparent there were differences between brains and between hemispheres in the TO regions.

Some brains appear to provide supporting evidence of the lateralisation hypothesis, others show no clear hemispheric differences. It is evident that there are individual differences in the brains organisation, particularly with respect to the function of hMT+. When examined, the components derived from performing ICA on our movie watching data suggest that ~25% of brains show a form of lateralisation, indicated by a bilateral component with an additional unilateral component in one hemisphere (see section 4.8). The remaining ~75% of subjects produced either a clear bilateral hMT+ component or a unilateral component in each hemisphere, which we infer as not providing evidence to support the hypothesis that there is lateralised function in all brains.

The results of correlational and ICA analyses converge – 75% of the subjects that showed evidence of lateralised function with the ICA approach also showed lower cross-hemisphere timecourse correlations for TO2/MST than V1.

To summarise, our results do not show a consistent pattern of lateralised function in hMT+ across all brains, nor is our evidence consistent with Strong et al.'s findings that right hemisphere hMT+ has an enhanced role compared to left hemisphere hMT+ (e.g., the additional unilateral component does not always sit in the right hemisphere). However, the examination of individual brains did allow us to identify some functional differences between hemispheres in TO2/MST for a small number of subjects (see section 4.5). It is the convergent evidence of the different analyses presented in this chapter that informed our conclusion that the lateralised function of hMT+ is not a universal phenomenon. This highlights the importance of conducting individual subject analyses. A common focus of human neuroimaging research has been on averaging data across groups. These group-average analyses involve the aggregation of small amounts of data from multiple individuals and can, often unintentionally, limit the detail and specificity of the results. Although an invaluable tool for establishing generalised basic principles of function, analyses on the group-average brain can obscure subject-specific features. "Precision fMRI" offers an alternative approach, the central focus of which is the characterisation of brain function and organisation at the individual level (Gordon et al., 2017).

We adopted this approach in this Chapter, and throughout this thesis as a whole. Our results provide evidence that a phenomenon inferred to be consistent across brains, the enhanced role of right hMT+ in motion processing, may in fact be more individualised than first thought.

It is important to note the potential alternative explanations of our findings. As discussed in section 4.1.4, the anatomical location of hMT+ can vary across brains. This is an important limitation of the atlas-based approach adopted in this chapter. Atlases are the direct result of the aggregation of data from different brains, and although they are useful tools for providing spatial approximations of the areas, they cannot account for all the anatomical variability between brains. This is potentially problematic when, as they were here, ROIs are defined from the atlas. ROIs defined from group data may not reflect the individual brain. Therefore, we must continue to remind the reader that the atlas definitions of the visual areas, in particular hMT+, are susceptible to a potential misplacement of the area in individual brains. The ROIs derived from the Wang et al. (2015) atlas may provide an approximation of the area, but in the absence of a 'ground truth' it is possible that the

timecourses/mean correlations extracted from the ROI may not encompass the entire hMT+ complex in each brain. However, as Huang et al. (2019) reported, there is good spatial correspondence between functionally defined hMT+ and Wang et al.'s TO1/TO2 atlas maps.

Hemispheric differences in brain topography must also be considered when interpreting our results since the folding pattern of the cortex can vary across hemispheres within the same brain. As discussed in section 4.1.4, the AOS, a sulcus that is associated with the anatomical location of hMT+, has been shown to vary across hemispheres. Given our fMRI data were resampled to a symmetrical surface (fsaverage_sym) that is a mirror flip of the left hemisphere, one may suggest our results are susceptible to misplacement of activation relative to sulci/gyri. However, the ROI-based analysis was also performed in native surface space, therefore circumventing the issues potentially arising from registering the data to a surface with an average folding pattern. We find the results are almost identical to those derived from performing the analysis on the average surface (see Appendix XI). In addition, ICA was also performed in native volume space before the spatial maps were resampled to the native surface, which, again, alleviates concerns about the differences between brains in the exact location of TO1/TO2 influencing the results. Although we acknowledge that none of our analyses are conclusive in isolation, they do all point in the same direction; they converge.

Finally, it is possible that differences in fixation could have led to the unequal number of components in the left and right hemispheres. For example, subjects who seem to have lateralised motion processing, may have been looking at the edge of the screen when the most significant motion in the movie occurred. Therefore, correlated activity between hemispheres could simply reflect the image statistics. Future work could explore this through the implementation of eye-tracking whilst subjects are in the scanner.

Chapter 5. Three visual pathways?

5.1 Introduction

This chapter will investigate whether performing ICA on fMRI data collected during natural viewing of a movie can be used to provide evidence of a hypothesised third visual pathway, the lateral motion pathway (Gilaie-Dotan, 2016). Separate from the traditional dorsal and ventral pathway model, the lateral motion pathway is believed to be dedicated to the processing of visual motion.

The traditional dual stream hypothesis suggests visual processing is mediated by two functionally distinct cortical pathways: the ventral pathway involving occipito-temporal regions and the dorsal pathway involving occipito-parietal regions (Goodale & Milner, 1992; Ungerleider & Mishkin, 1982). This dual stream model was first proposed by Ungerleider and Mishkin (1982) who suggested the striate cortex of the rhesus monkey was segregated into two corticocortical pathways, evidence of which came from lesions studies (Mishkin et al., 1983). One such study showed lesions to the inferior-temporal cortex induced significant impairments in visual pattern discrimination and recognition. On the other hand, there appeared to be little impact on the monkey's ability to accurately perform tasks in which visual cues primed the monkey to which of two locations would be rewarded (landmark task).

The model suggests that the two pathways process different visual attributes, and therefore the nature of the visual inputs themselves dictate which stream is charged with processing the information. It is this distinction between the inputs of these two streams that directs the division of labour. The dorsal pathway was believed to mediate spatial perception relating to an object's spatial location, while the ventral pathway was believed to mediate object identification, including recognition and categorisation, relating to the shape and identity of an object.

Later, Goodale and Millner proposed a revised and somewhat contrasting "perception-action" model in which the functional differences between the two pathways centred around the output systems they serve (Goodale & Milner, 1992). Both models involve anatomically segregated pathways which are hierarchical in nature, in that information processing becomes increasingly complex the further along the respective pathway it travels. However, unlike Ungerleider and Mishkin, the revised

model suggests the processing goal dictates which pathway is utilised. The key distinction between the two models is that Goodale and Milner proposed the ventral pathway constructs a perceptual representation of the visual world and objects within it, while the dorsal pathway supports the visual control of actions directed towards those objects, such as reaching or grasping. In their model, Goodale and Milner propose the output, rather than the input, characteristics of the two pathways are responsible for the division of labour. The dorsal and ventral pathways are often referred to as the "where"/"action" and "what"/"perception" pathways respectively.

Compelling evidence for the perception-action model was initially provided by the double-dissociations seen in individuals with lesions either to the lateral-occipital regions (ventral stream) or the intraparietal sulcus (IPS) and superior occipito-parietal regions (dorsal stream). Damage to the dorsal stream can give rise to optic ataxia, a neuropsychological disorder in which an individual's ability to aim movements towards visual targets is compromised (Pisella et al., 2000). Patients with optic ataxia often fail to accurately direct their hand towards an object they remain able to describe the relative position of (Perenin & Vighetto, 1983). On the other hand, damage to the ventral pathway can give rise to visual agnosia (Lissauer, 1890), a deficit in the ability to recognise visual objects.

Patient DF suffered bilateral damage to the ventro-lateral regions of her occipital lobes as the result of carbon monoxide poisoning (Milner et al., 1991). When DF was presented with a pair of shapes and asked to determine whether the two were the same or different, her accuracy was around chance (Goodale et al., 1994). Essentially, DF was unable to perceive whether two objects were the same. Likewise, DF could only identify 10% of the objects conveyed in several line drawings. Conversely when asked to pick an object up her grasp points were similar to those selected by a control subject. In addition to her object shape recognition deficits DF has been shown to be similarly impaired at object orientation recognition, even judging horizontal lines to be vertical (Goodale et al., 1991).

In contrast, patient RV suffered bilateral damage to the occipito-parietal cortex as a result of stroke. RV showed no evidence of visual form agnosia, in that she was able to identify all the objects conveyed in the same line drawings that were presented to DF. When asked to determine whether

shape pairs were the same or different, RV scored 90%. However, unlike DF, RV was unable to use the visual information about the objects to control her actions, in that her grasp points were often unstable (Goodale et al., 1994). Here there are two individuals, one is able to perceive the shape of an object but unable to modify their grasping action based upon this information, the other is able to guide the placement of their grasp but unable to distinguish between two objects. But where does motion perception fit within this model?

Visual motion perception has long believed to fall under the purview of the dorsal stream. However, there is emerging evidence that speaks against the traditional dual stream model, with particular attention being paid to the notion of a third, lateral pathway dedicated to the processing of visual motion (Gilaie-Dotan, 2016). This three-stream model suggests the ventral pathway is charged with vision for perception, the dorsal pathway is charged with vision for action, and a lateral pathway including V5/MT+ is charged with visual motion perception. This "lateral visual motion pathway" proposes motion signals are propagated to TO1/MT via three non-hierarchical routes. Signals are then relayed to TO2/MST and multiple other brain regions depending on task needs (Figure 74).

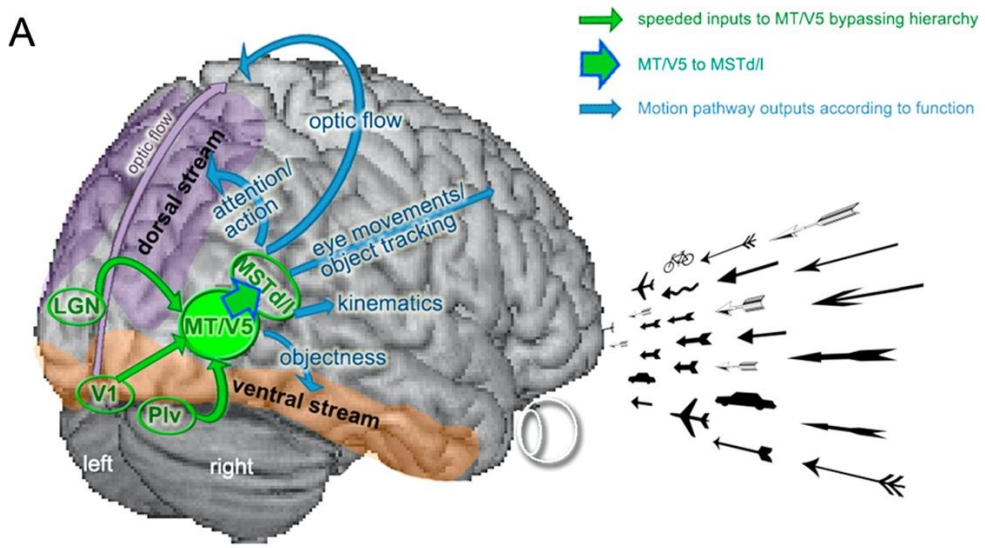

Figure 74. A visualisation of the lateral motion pathway adapted from Gilaie-Dotan (2016). Motion information is swiftly propagated directly to V5/MT+ via V1, lateral geniculate nucleus (LGN) and pulvinar (Plv) in parallel. Information flow is therefore non-hierarchical. Information is then transferred to a multiplicity of brain regions depending on task needs.

Gilaie-Dotan proposed a clear distinction between the dorsal and motion pathways based upon several factors. First, the neuroanatomy of the visual motion pathway which is split into three stages. V1, the pulvinar and the LGN (first stage) each swiftly propagate visual motion information in parallel

to V5/MT+ (second stage). Information is then propagated from V5/MT+ to the subdivisions of MST (third stage).

The differences in sensitivity to spatial resolution between the dorsal and visual motion pathways is another key distinction. The visual motion pathway is hypothesised to mediate the swift detection of any kind of motion at any visual field location, and therefore is sensitive to medium-to-high spatial information. On the other hand, since the dorsal pathway mediates one's ability to perform actions directed towards objects, such as reaching and grasping, it does not require high spatial information.

Another distinction between the two perspective pathways is the sensitivity to the spatial proximity of visual stimuli. Again, since the visual motion pathway is sensitive to all visual motion across the entire visual field, it is hypothesised the pathway will not show a preferential sensitivity to extrapersonal (far) or peripersonal (near) space. On the other hand, dorsal visuomotor regions show greater activation when subjects are required to perform tasks in peripersonal space, compared to extrapersonal space (Weiss et al., 2000).

According to Gilaie-Dotan, not all motion processing is performed in the separate lateral motion pathway. Despite V3A exhibiting similar motion response properties to those reported for V5/MT+ (Tootell et al., 1997), the three-stream model states the region falls under the purview of the dorsal pathway, and not the lateral motion pathway.

To investigate the lateral motion pathway hypothesis, I performed ICA on data collected with a broadband naturalistic stimulus (a movie). ICA is a data-driven technique and so provides an unbiased way of assessing this hypothesis.

ICA has previously been used with movie data to reliably identify discrete visual areas (discussed in Chapter 3) and functional networks, composed of regions with coherent fluctuations which are temporally correlated at rest. These "resting-state networks" (RSNs) are consistent across observers and contain regions known to be involved in specific functions, such as visual processing (Beckmann et al., 2005; Damoiseaux et al., 2006; Smith et al., 2009).

Smith et al. (2009) performed ICA on resting-state fMRI (rs-fMRI) data and compared the spatial maps to a set of task-based networks. Ten functional networks were identified which spatially correlated with regions known to show co-activation during specific tasks. Smith at al.'s maps include the following networks: 1-3) visual, 4) the default mode network (DMN), 5) cerebellum, 6) sensory-motor, 7) auditory, 8) executive control, 9-10) frontal-parietal. Visual network 1 represents the primary visual network, while visual networks 2 and 3 appear to represent the traditional ventral and dorsal streams of processing respectively. Comparable networks have been reported elsewhere (Beckmann et al., 2005; Damoiseaux et al., 2006)

Here, visual network 3 was of particular interest as it includes both V3A and hMT+, two regions believed to belong to different pathways under Gilaie-Dotan's model. We examined whether the two regions are fragmented when movie watching data is used. An advantage of using a broadband naturalistic stimulus to segment the brain is that neural activity is driven by both internal and external stimulation, whilst resting-state data relies solely on internally driven activity. Therefore, if Gilaie-Dotan's model is correct and hMT+ is part of a pathway separate from the dorsal stream, we might expect activity driven by the movie in the regions involved, namely V5/hMT+, to differ to those involved in the separate dorsal pathway, such as V3A. The technique of performing ICA with a fixed dimensionality on our movie watching data should identify brain regions within a network that have correlated activity and segment them from other networks with which they are uncorrelated.

First, we assessed whether the 10 well-established RSNs can be accurately identified with ICA and movie watching data, paying particular attention to visual network 3 which represents the dorsal pathway. We find each network was well represented by our data, indicated by high spatial correlations between the components and the network maps. Visual network 3 was a notable deviation as the correlation between the component and the network was substantially lower compared to the other two visual networks. Second, we examined whether the lower correlations seen for visual network 3 were symptomatic of the network splitting into a dorsal and lateral motion pathway, i.e., was this evidence of a partitioning of the brain into distinct networks which support Gilaie-Dotan's three visual pathway model? At the group level, we find evidence of distinct dorsal and lateral motion pathways,

which were subsequently found in the majority of brains (>50%) when single-subject ICA was performed.

5.2 Methods

For all analyses the same dataset was used, and therefore the acquisition and stimulus methods were identical (see section 2.1). All 17 subjects watched the same 20-minute clip taken from the beginning of a Hollywood action film (James Bond, Skyfall) whilst in an MR scanner.

Individual subject data were pre-processed using the Feat GUI with default parameters; motion correction (MCFLIRT), brain extraction (BET) and high pass temporal filtering (100 Hz). As with Smith et al. (2009), data were spatially smoothed with the FEAT default 5 mm FWHM kernel. Pre-processed data were resampled to a 3 mm MNI template using FSL's *"applywarp"* command, before temporally concatenated ICA was performed on the data. A grey matter mask derived from the MNI template was also applied to the data. Resting state networks were identified by Smith et al. with dimensionality set to 20. Therefore, to allow for direct comparison between components derived from movie watching data and networks reported by Smith et al., dimensionality was also set to 20 in our analysis. Smith's networks were downsampled from 2 mm MNI to 3 mm MNI space to match the native resolution of the individual subject and therefore the group data. Spatial correlations were computed between the unthresholded spatial maps (melodic_IC) and Smith et al.'s networks using FSL's *"fslcc"* command. To identify a candidate component for each of Smith et al.'s networks, a minimum correlation threshold of $r = 0.25$ was implemented. This threshold was also employed by Smith et al. to compute correlations between components derived from resting-state and task-based (BrainMap) data. Both the spatial maps and Smith's networks were resampled to 1 mm MNI space for visualisation purposes.

5.3 Results

Table 17 presents the spatial correlations between Smith et al.'s networks and the best matched components derived from the group-level ICA. All networks are represented by an ICA component above the minimum threshold, with visual network 1, DMN and the right hemisphere frontoparietal

networks being represented best by the components (indicated by highest correlations). The executive control network, auditory network, and visual network 3 have the lowest correlations, but the best matched components still correlate above the r = 0.25 threshold. The top matched component for each network is unique to that network, in that it was not a top match for another network. This suggests each network had been identified separately. The top matched component for each network is displayed in Figure 76, Figure 77 and Figure 78.

Table 17. Spatial correlations between the top matched components derived from performing group ICA on movie watching data and the resting state networks published in Smith et al. (2009).

Network	Movie vs. Smith et al.'s network r	Rank
Visual 1	0.81	1
Visual 2	0.74	4
Visual 3	0.53	7
DMN	0.8	2
Sensorimotor	0.74	4
Auditory	0.56	6
Executive control	0.41	8
LH frontoparietal	0.7	5
RH frontoparietal	0.76	3
Mean	0.67	
SD	0.14	

5.3.1 Visual network 1 – "primary visual"

Smith's visual network 1 encompasses the medial-occipital regions, including the primary visual cortex, extending dorsally (Figure 76, left, top row). This network, also referred to as the primary visual network, is believed to be responsible for processing static and moving objects (Shen et al., 2019). The top matched component represents the network well, with a correlation of r = 0.81 (Figure 76, left,

bottom row). The peak voxel location of the component ([26, 16, 25], MNI coordinates [12, -78, 3]) lies within the right hemisphere, near to the centre of calcarine sulcus.

5.3.2 Visual network 2 – "ventral pathway"

Smith's visual network 2 encompasses the ventro-lateral regions of the occipital lobe (Figure 76, middle, top). Visual networks 1 and 2 overlap to a small degree in posterior regions of the occipital lobe (Figure 75) before the two networks extend dorsally and laterally respectively. We may consider visual network 2, also referred to as the ventral visual network (Shen et al., 2019), to be part of the ventral stream of visual processing. Visual network 2 is represented well by the top matched component, with a correlation of r = 0.74 (Figure 76, middle, bottom). The peak voxel location of the component ([34, 10, 23], MNI coordinates [-12, -96, -3]) lies within the left hemisphere at the posterior end of the calcarine sulcus.

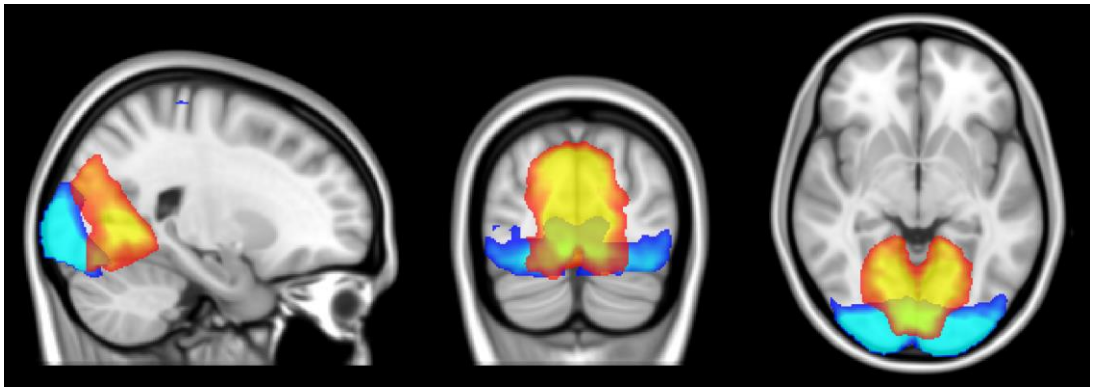

Figure 75. Smith's visual network 1 (yellow) and visual network 2 (blue) overlaid on the MNI template. The two networks show a small degree of overlap at the posterior regions of the occipital lobe, before they extend dorsally and laterally respectively.

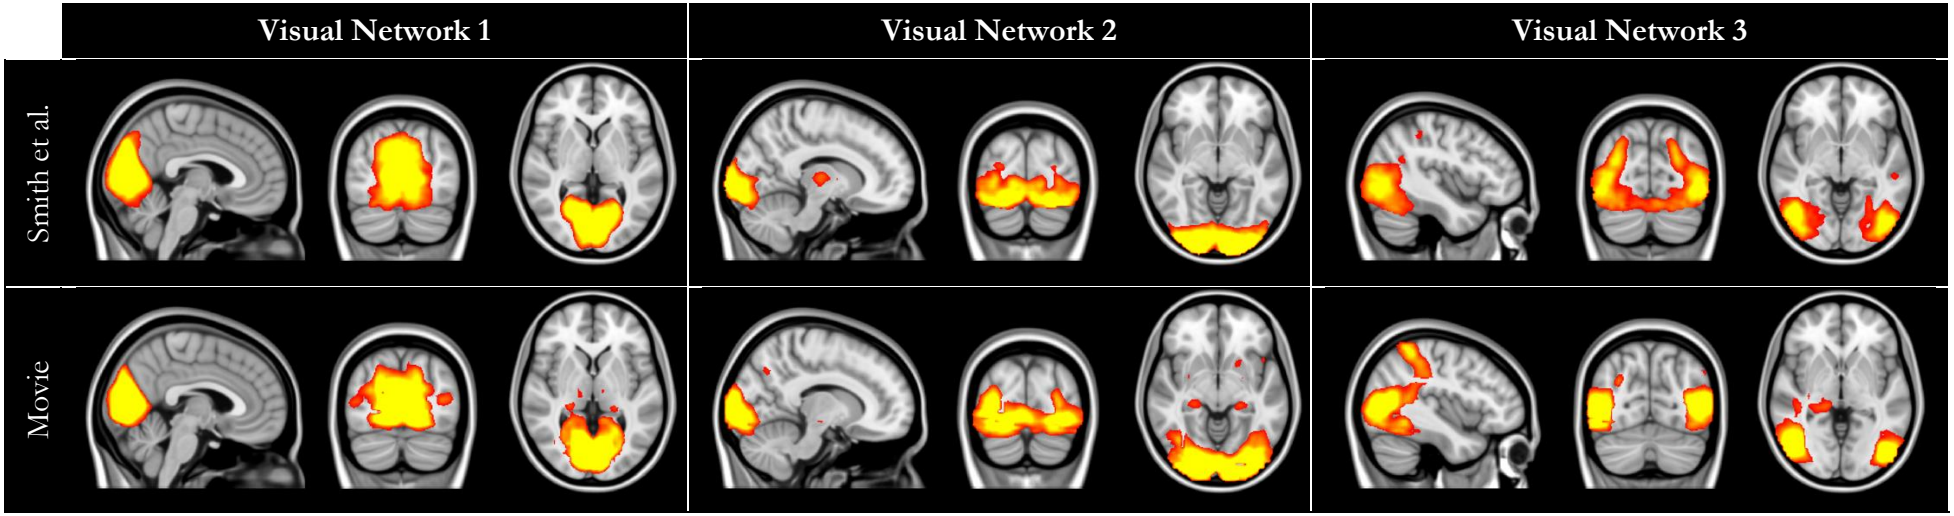

Figure 76. Smith et al.'s three visual networks (top) and the components showing the highest spatial correlation with each network (bottom). Components were derived from movie watching data.

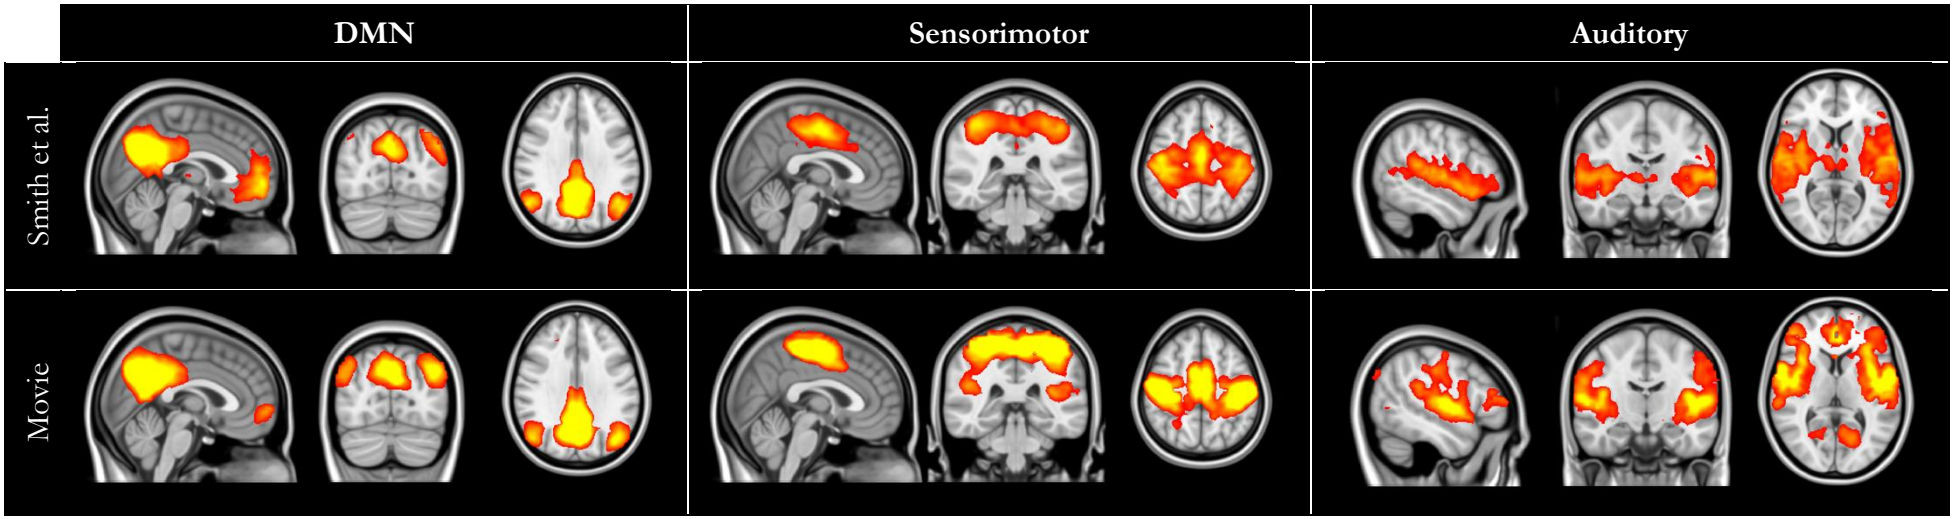

Figure 77. Smith et al.'s DMN, sensorimotor and auditory networks (top) and the components showing the highest spatial correlation with each network (bottom). Components were derived from movie watching data.

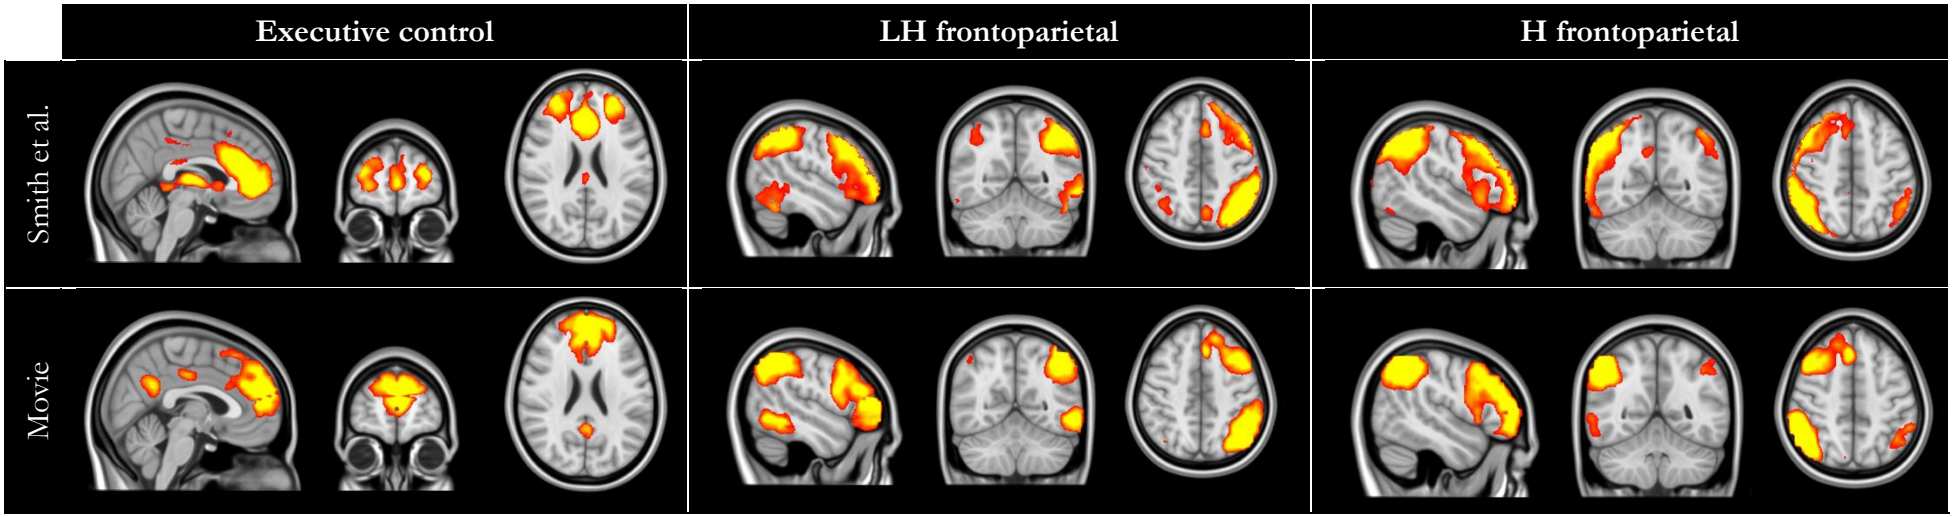

Figure 78. Smith et al.'s executive control, left and right hemisphere frontoparietal networks (top) and the components showing the highest spatial correlation with each network (bottom). Components were derived from movie watching data.

5.3.3 Visual network 3 – "dorsal pathway"

Visual network 3, also referred to as the dorsal visual network (Shen et al., 2019), covers the lateral occipital regions, including V5/hMT+ and extending dorsally past region V3A (Figure 76, right, top). The component showing the highest correlation (top matched component) with Smith's visual network 3 (r = 0.53) includes bilateral V5/hMT+ with activity in the post-central sulcus (Figure 76, right, bottom). The peak voxel location lies within the left lateral-occipital region ([46, 18, 27], MNI coordinates [-48, -72, 9]).

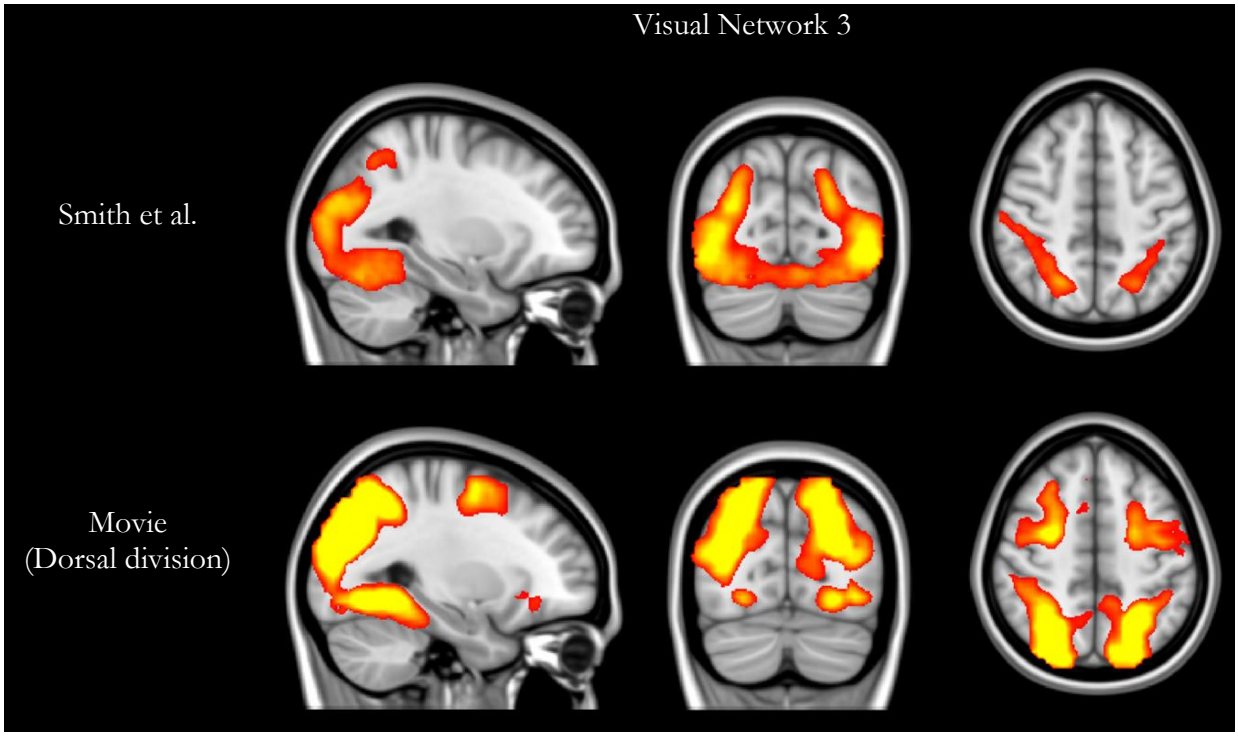

Figure 79. Smith et al.'s visual network 3 (top) and an additional component derived from performing ICA on our movie watching data (bottom). The additional component correlates with Smith's visual network 3 at r = 0.4 and covers the dorsal occipital regions but does not include hMT+. We refer to this as the "dorsal division" component.

The top matched component with visual network 3 does not include a dorsal representation. Instead, the dorsal regions are present in a separate component which correlates with Smith et al.'s visual network 3 at r = 0.4 (Figure 79). The additional component shows bilateral activation in dorsal regions of the occipital lobe including V3A, with very little activation in the lateral-occipital regions. The peak voxel of the additional component lies dorsally in the left hemisphere ([39, 14, 32], MNI coordinates [-27, -84, 24]). To assess whether the two components represent subdivisions of the network, the two components were combined. Correlations between the combined component and

Smith's visual network 3 increased to r = 0.61 (Figure 80, bottom), indicating visual network 3 had split into a "dorsal" and "lateral" component.

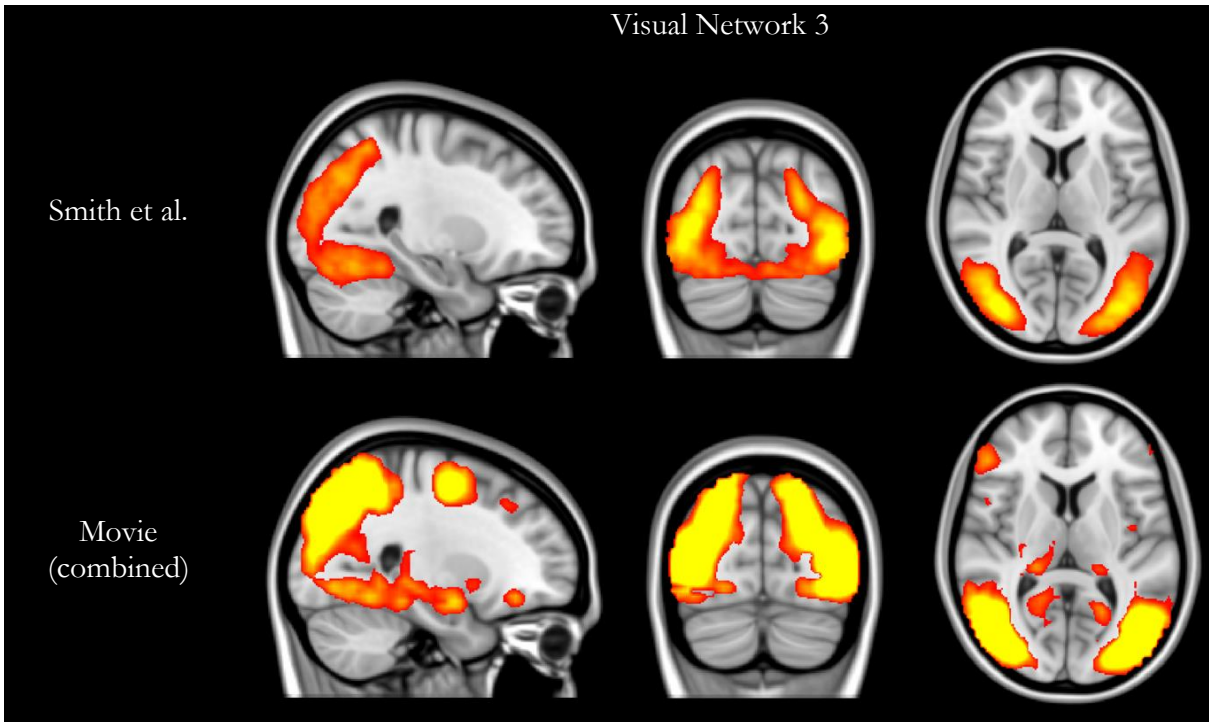

Figure 80. Smith et al.'s visual network 3 (top) and the spatial map generated by combining the top matched component for this network (dorsal division) with a sub-component covering the lateral-occipital regions (lateral division). The combined component correlates with Smith's network at r = 0.61, suggesting the two components combined represent visual network 3 better than either does on its own. From this, we infer the network has split into two distinct components, one centred on hMT+ (lateral division) and another centred on V3A (dorsal division).

To assess which areas were included in the dorsal and lateral divisions, correlations were computed between the components and the probabilistic atlas of visual areas used in previous chapters (Wang et al., 2015). The lateral division (Figure 81) correlates with TO1/MT (r = 0.41), TO2/MST (r = 0.42) and IPS5 (r = 0.28). The dorsal division (Figure 82) correlates with V3b (r = 0.28), and regions located in the intraparietal sulcus (IPS0 – IPS4 mean r = 0.37). Although the component does not correlate with V3A above the r = 0.25 threshold, there is clearly an involvement of the region in the component (Figure 82). Furthermore, we calculated the percent overlap between the component and the probabilistic atlas region of V3A[31]. This analysis reveals that 50% of the total volume of V3A is

[31] First, the spatial maps were thresholded at z score = 3. We then calculated the intersection between the component and the atlas region by multiplying the binarized component by the binarized atlas region. Next, we calculated the size in terms of non-zero voxels of the intersection (630) and the atlas region (1260). To calculate how big the intersection is relative to the atlas region, we divided the intersection voxels by the total size of the atlas region.

represented by the component. The lateral portions of V3A are well represented by the component, while the posterior/medial regions are not.

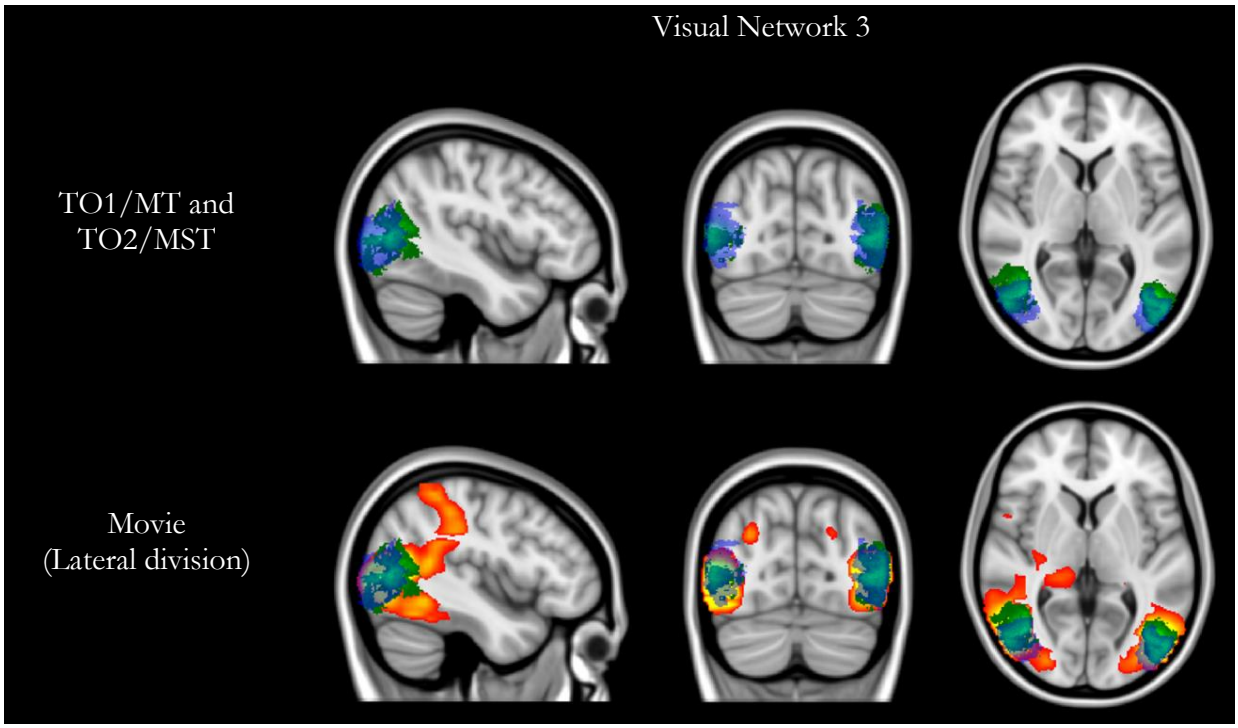

Figure 81. Wang et al.'s probabilistic TO1/MT (blue) and TO2/MST (green) atlas regions (top). When these atlas regions are overlaid on top of the component with the highest spatial correlation with Smith's visual network 3 (lateral division), there is a clear overlap.

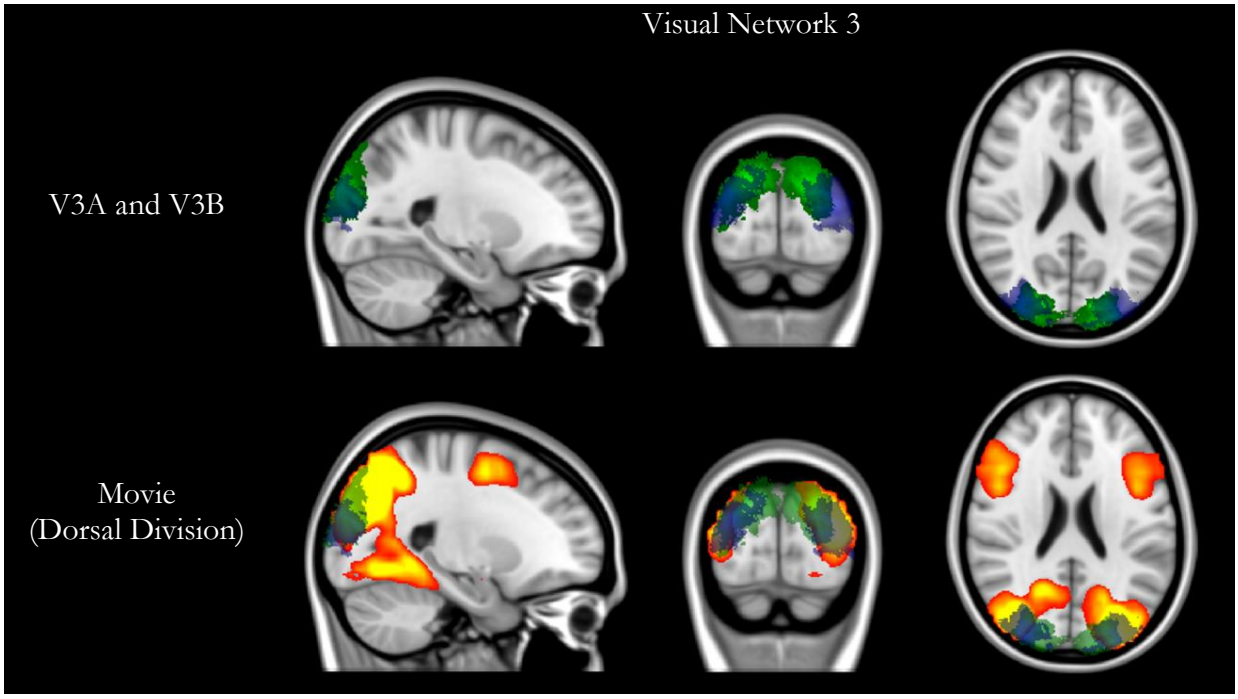

Figure 82. Wang et al.'s probabilistic V3A (green) and V3B (blue) atlas regions (top). When overlaid on top of the sub-component covering the dorsal-occipital regions (dorsal division), there is a clear overlap between the regions and the component.

5.3.4 Comparison to resting state fMRI

In addition to movie watching, five minutes of resting state (RS) fMRI data was acquired for 14 subjects[32]. These data were used to explore whether the splitting of visual network 3 into dorsal and lateral sub-components was also evident when the subjects were at rest.

5.3.4.1 Methods

Resting state fMRI data were acquired using the same scan parameters described in section 2.1, with the only difference being that 150 volumes were acquired over the five-minute scan duration. Subjects were instructed to relax, think about nothing in particular and fixate a white cross presented in the centre of a grey background. Data pre-processing, as well as the approach adopted for identifying a candidate component for each network, was identical to that described in section 5.2. To recap, data were pre-processed using the FEAT default parameters before being resampled to a 3 mm MNI template. Temporally concatenated ICA was then performed on the spatially normalised data with dimensionality set to 20. To identify a candidate component for each network, spatial correlations were computed between the unthresholded spatial maps and each of Smith et al.'s networks.

5.3.4.2 Results

Table 18 presents the spatial correlations between Smith et al.'s networks and the best matched components derived from the group-level ICA performed on RS and movie watching data. With both data sets, each network is represented by a unique ICA component, with a spatial correlation above the $r = 0.3$ threshold. Again, this suggests each network had been identified separately. The top matched component for each network is displayed in Figure 83, Figure 84 and Figure 85. There is, however, a clear difference between the two data sets: visual network 3. With movie data, visual network 3 shows the second lowest spatial correlation. As discussed above, this low correlation is likely being driven by the network sub-dividing into two distinct components. However, when the subjects are at rest, visual network 3 shows the second highest spatial correlation. In addition, with RS data, visual network 3 receives only one component match above threshold, indicating there is a single candidate component

[32] This dataset was collected by Phoebe Asquith.

for the entire network. This is in contrast to the movie watching data in which the network received two matches above threshold, indicating the network had fragmented. Visual inspection of the spatial map derived from RS data shows the component covers the lateral occipital regions, including V5/hMT+ and extends dorsally past region V3A. Again, this is in direct contrast to the component derived from movie watching, in which the dorsal extension is absent and is found in a separate component (see section 5.3.3). The inclusion of dorsal V3A in the RS component was then confirmed through spatial correlations with the probabilistic atlas of visual areas (Wang et al., 2015). The component matching with visual network 3 correlates with the V3A atlas region at $r = 0.36$. The result of this RS analysis indicates that visual network 3 does not split into dorsal and lateral sub-components when subjects are at rest.

Table 18. Spatial correlations between the top matched components derived from group ICA and the resting state networks published in Smith et al. (2009). Correlations are shown for resting state (left) and movie watching fMRI data (right).

Network	Rest Vs. Smith et al.'s network r	Rank	Movie Vs. Smith et al.'s network r	Rank
Visual 1	0.8	1	0.81	1
Visual 2	0.58	7	0.74	4
Visual 3	0.7	2	0.53	7
DMN	0.61	6	0.8	2
Sensorimotor	0.7	2	0.74	4
Auditory	0.68	3	0.56	6
Executive control	0.49	8	0.41	8
LH frontoparietal	0.64	5	0.7	5
RH frontoparietal	0.66	4	0.76	3
Mean	0.65		0.67	
SD	0.09		0.14	

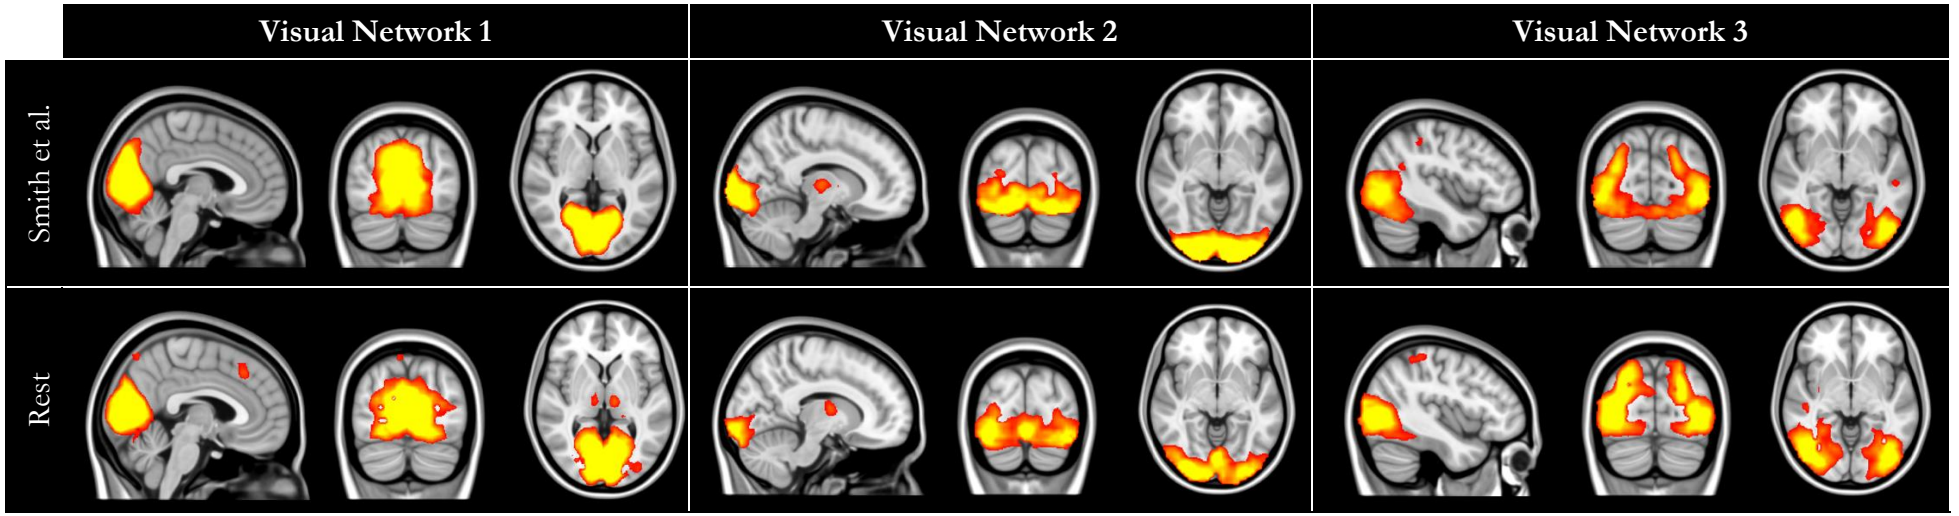

Figure 83. Smith et al.'s three visual networks (top) and the components showing the highest spatial correlation with each network (bottom). Components were derived from resting state data.

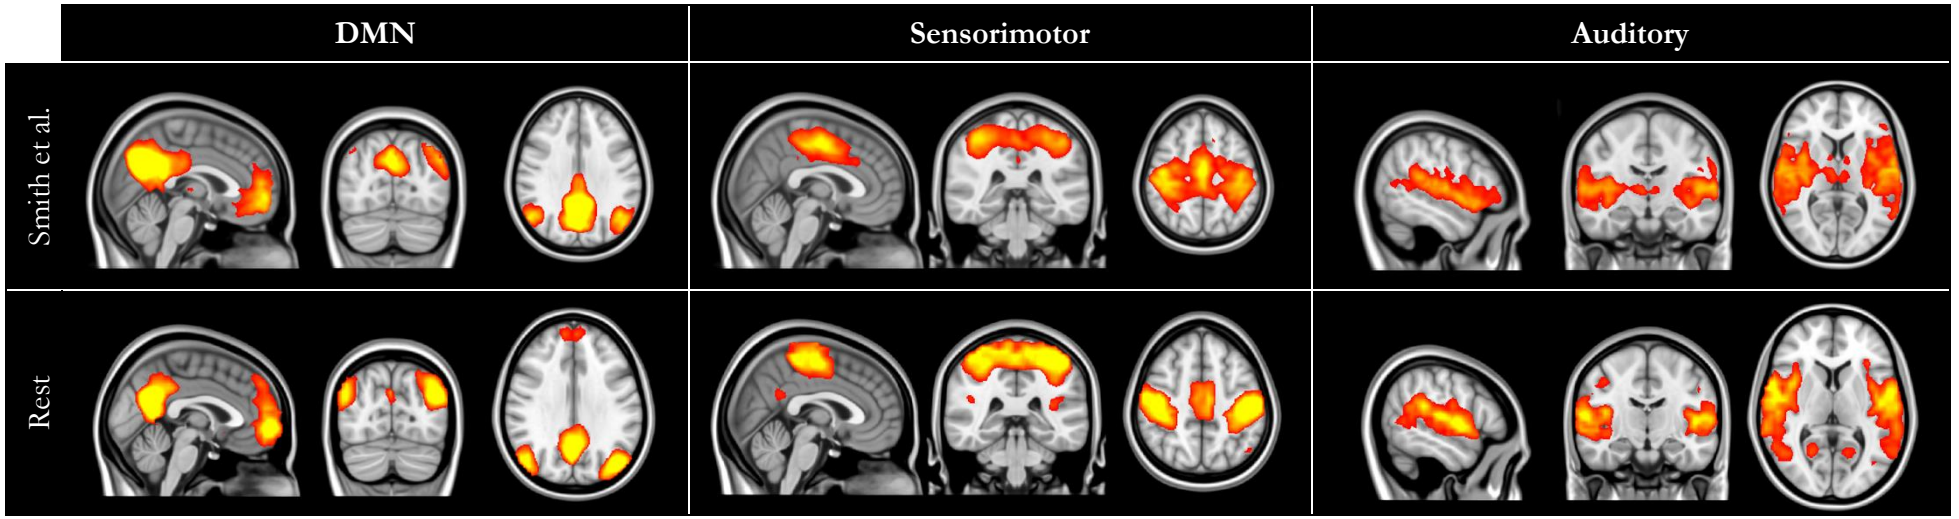

Figure 84. Smith et al.'s DMN, sensorimotor and auditory networks (top) and the components showing the highest spatial correlation with each network (bottom). Components were derived from resting state data.

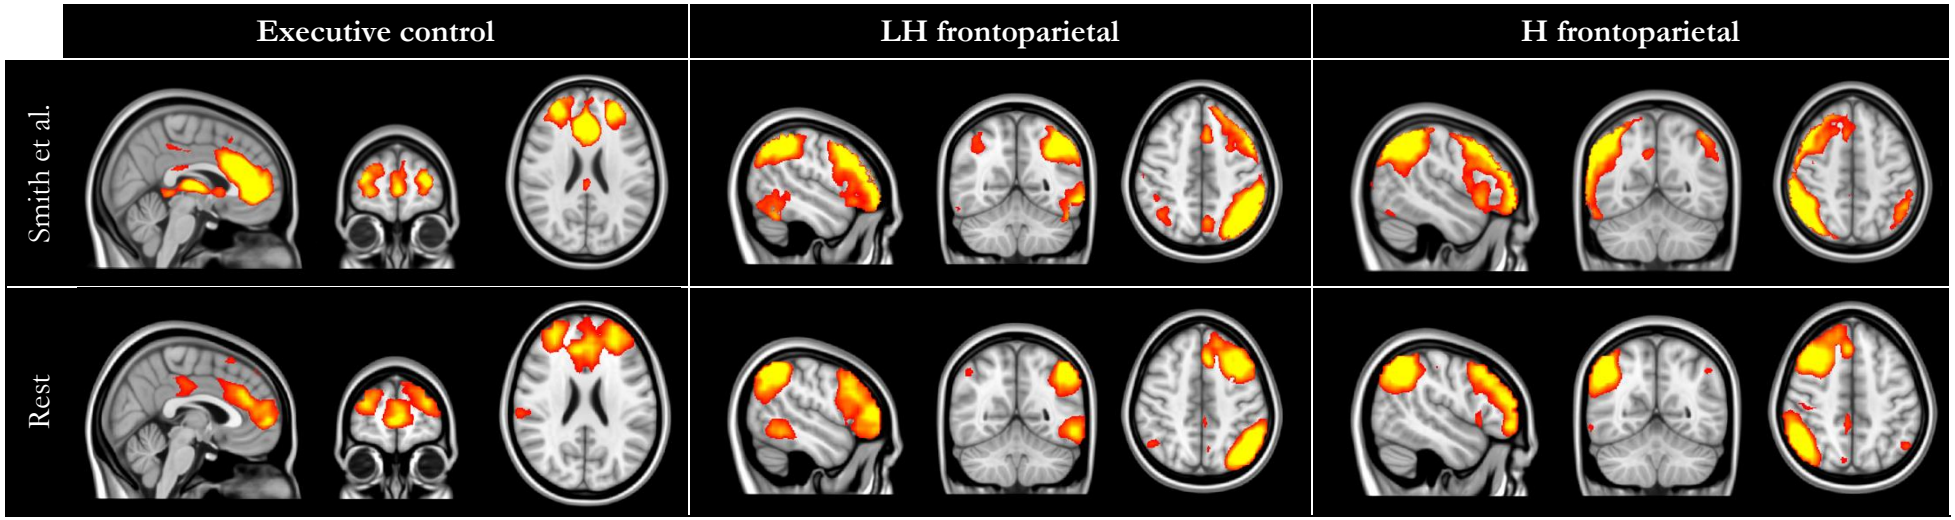

Figure 85. Smith et al.'s executive control, left and right hemisphere frontoparietal networks (top) and the components showing the highest spatial correlation with each network (bottom). Components were derived from resting state data.

5.3.5 Interim discussion

Thus far, performing ICA on movie watching fMRI data has output components which correlate with well-established resting state networks (S. M. Smith et al., 2009). Visual network 1 corresponds to the primary visual network, while visual network 2 appears to represent the ventral visual pathway. However, visual network 3, considered to represent the dorsal pathway, is represented by two distinct components; one shows activation in the dorsal V3A region which extends to the IPS (dorsal division), while the other shows activation in V5/MT+ and the post-central sulcus (lateral division). In contrast, when subjects are at rest visual network 3 does not fragment, but rather the whole network is represented by a single component covering both V5/MT+ and V3A. Our finding from the movie data provide supporting evidence of Gilaie-Dotan's (2016) alternative model of visual processing which postulates hMT+ is part of a visual motion pathway, separate from the dorsal pathway which includes V3A. Figure 86 presents a summary of the spatial correspondence between the components derived from our movie watching data and Smith's three visual networks, along with the network or pathway we infer our components to represent.

The order in which the subject data is input into the group-level ICA has been shown to affect the reliability of the output components. This is known as the brain order effect (Asquith, 2018; Poppe et al., 2013). To examine this, I systematically swapped the first brain input to MELODIC and found a negligible effect on the outputs. The mean correlation between the components and Smith's networks decreased on average by 0.01. Likewise, the spatial maps of the top matched components were comparable across swapping permutations and the splitting of dorsal/lateral components remained.

In summary, at the group-level our results demonstrate two things. First, well-established resting state networks do not reorganise when observers free view a movie, with the exception of visual network 3. Second, our data provides evidence to support Gilaie-Dotan's (2016) hypothesis of a third visual pathway specialised for processing visual motion, the lateral motion pathway. Following this, a logical question to ask is – does the group data reflect individual brains? i.e., do all observers

show evidence of the lateral motion pathway? To address this question, I next performed single-subject ICA.

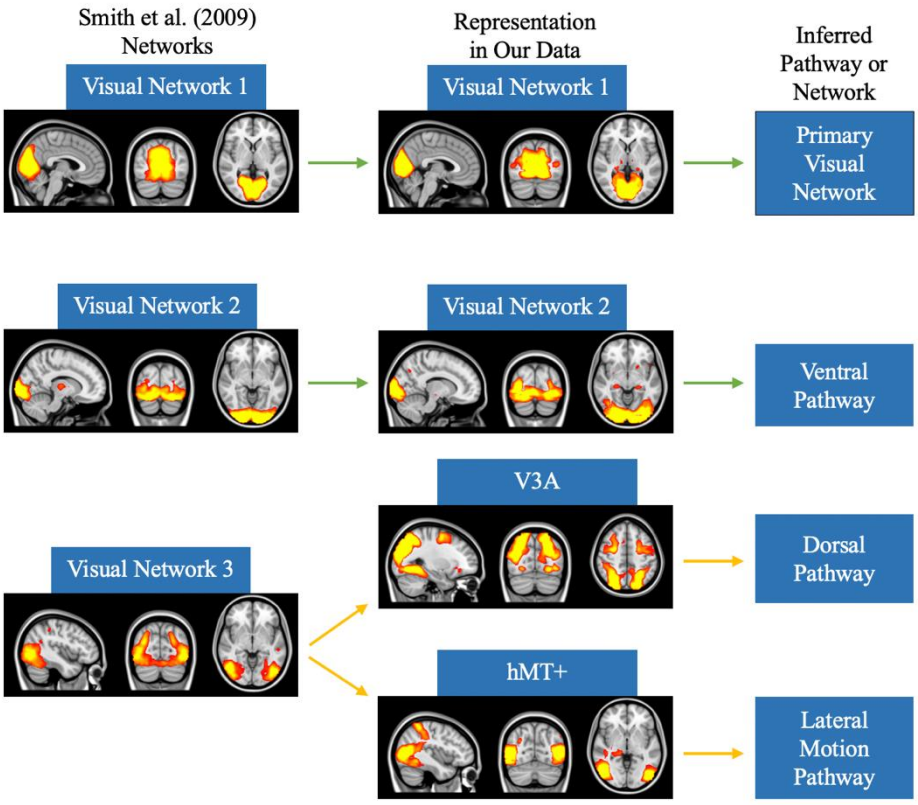

Figure 86. Summary of the spatial correspondence between Smith et al.'s three visual networks and the components derived from performing ICA on movie watching fMRI data. Visual networks 1 and 2 are well represented by our data (indicated by the green arrows), which we infer as corresponding to the primary visual network and the ventral pathway of visual processing respectively. Visual network 3 is represented by two components in our analysis, one with activation centred at V3A (dorsal division), another with activation centred at hMT+ (lateral division). The splitting of V3A and hMT+ into two distinct components (represented by the yellow arrows) provides supporting evidence of Gilaie-Dotan's (2016) lateral motion pathway which is believed to include hMT+ and not V3A.

5.3.6 Single-subject ICA

In order to assess whether dorsal and lateral representations are separated at the individual level, I performed a series of single brain analyses. The data used in the group-level analysis, prior to resampling to MNI space, was input to MELODIC. Again, the output spatial maps were resampled to 3 mm MNI space to allow for correlations to be computed with Smith's networks. As with previous single-subject analyses performed in this thesis, we adopted a slightly stricter matching threshold of $r > 0.3$. To aid in the identification of components, correlations were also computed between the subject's spatial maps and the group-based ICA components (dorsal and lateral subdivisions and the two combined). The spatial maps were resampled to 1 mm MNI space to aid visualisation.

5.3.6.1 Dimensionality 20

As with the group analysis above, dimensionality was first set to 20. Spatial correlations were computed between the individual subject's spatial maps and Smith at al.'s networks (Figure 87). The mean across all networks and all subjects was r = 0.42. The mean across the three visual networks was r = 0.46. Looking exclusively at the top matched components for Smith's three visual networks:

- 94% (16/17) of subjects produced a component that correlated with visual network 1 at r > 0.3.
- 82% (14/17) of subjects produced a component that correlated with visual network 2 at r > 0.3.
- 100% (17/17) of subjects produced a component that correlated with visual network 3 at r > 0.3.

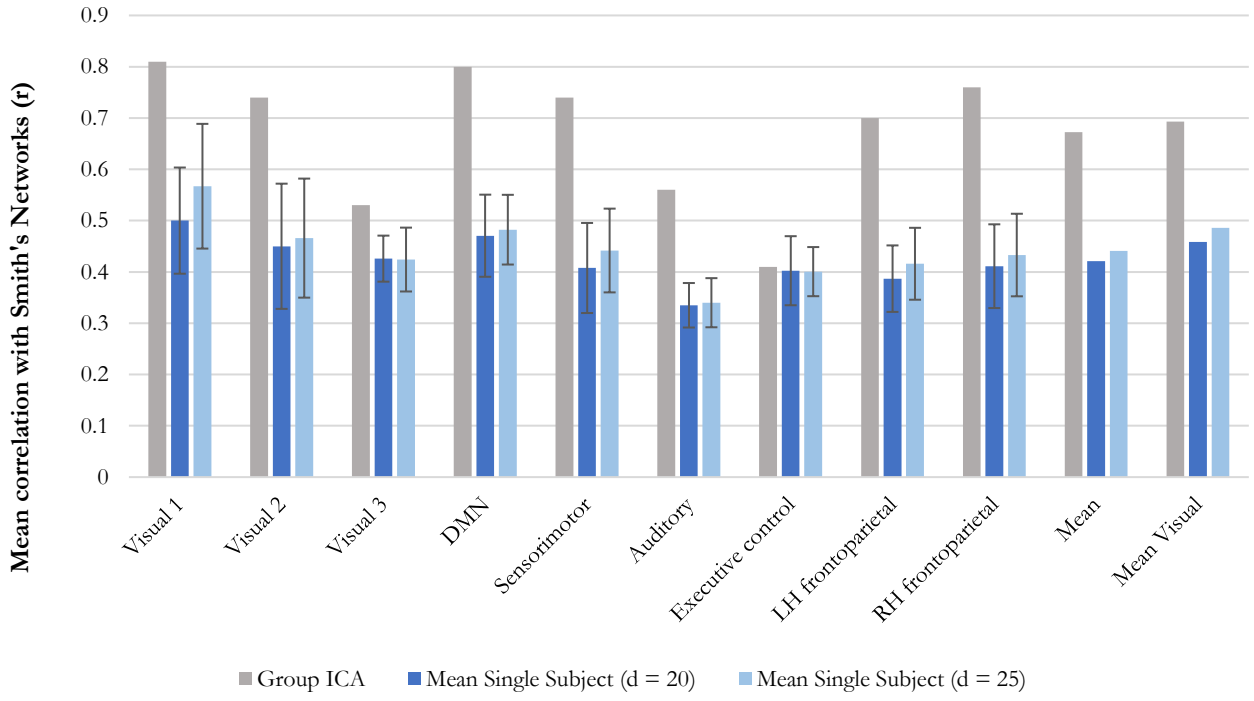

Figure 87. Correlations between each of Smith et al.'s networks and the top matched component derived from performing ICA on movie watching data. The group mean correlation between each of Smith et al.'s networks and the top matched single-subject components, with dimensionality (d) set to 20 or 25, are presented. Correlations between the top matched components derived from group-based ICA with dimensionality set to 20 are also presented. Fisher's Z equivalent figure can be found in Appendix T.

However, in some cases the top matched component (component with the highest spatial correlation) for one visual network was also the top match for another visual network. Only 35%

(6/17) of subjects had a unique component for each of the three visual networks. Likewise, only 35% (6/17) of subjects produced more than one component which correlated with Smith's visual network 3 above the threshold of r = 0.3. Of these 6 subjects, only 50% (3/6) produced additional components which correlated with visual network 3 alone, the three remaining subjects produced more than one component, but the additional components also correlated with another network. This suggests that at a dimensionality of 20, not only does visual network 3 not split into sub-regions, but the three networks themselves are not separated.

5.3.6.2 Dimensionality 25

At the group level, dimensionality of 20 is sufficient to allow visual network 3 (dorsal pathway) to split into two separate components (dorsal V3A and lateral hMT+). With single-subject ICA, 20 dimensions appears to underfit the data, resulting in the network not splitting into two components. One possible explanation for this is that when performing single-subject ICA, there is a reduction in data points compared to when concatenated group-ICA is performed. It is unlikely that noise will be consistent across brains and therefore the components derived from group-ICA may be less prone to the inclusion of multiple noise components within the output. This makes the spatial maps less noisy and the RSNs more easily identifiable (Griffanti et al., 2017). The lack of consistent noise across brains also means at a dimensionality of 20, ICA has more free "slots" meaning the network can split off into two components.

On the other hand, when ICA is forced to summarise the single-subject data into a predefined number of components (for example, 20), noise components may be included at the expense of signal components. These additional noise components do not get "washed out" as they might in group-ICA, potentially resulting in ICA being unable to split the network at a low dimensionality: the noise components are included in the 20 output components. In RSN research, 20-30 dimensions is considered "conventional" (Griffanti et al., 2014) and therefore, I investigated whether increasing dimensionality to 25 results in the output of components that better match Smith's networks, and whether the increase results in a splitting of visual network 3 into separate dorsal and lateral

components. Again, spatial correlations were computed between the individual subject's spatial maps and Smith's networks (Figure 87).

When dimensionality was set to 25, there was only a modest increase in the correlations between the networks and the top matched components compared to when dimensionality was set to 20. The mean correlation across all networks and all observers was $r = 0.44$, an increase of 0.02. The mean across the three visual networks was $r = 0.49$, an increase of 0.03. However, an increase in dimensionality does produce notable differences in the output of components relative to when dimensionality was set to 20. Again, looking exclusively at the top matched components for Smith's three visual networks:

- All subjects produced a component that correlated with visual network 1 at $r > 0.3$.
- 94% (16/17) subjects produced a component that correlated with visual network 2 at $r > 0.3$.
- All subjects produced a component that correlated with visual network 3 at $r > 0.3$.

Seventy per cent (12/17) of subjects produced a unique component for each of the three visual networks. 82% (14/17) of subjects produced more than one component that correlated with visual network 3, 64% (9/14) of which produced at least two components that correlated with the network alone e.g., the additional component(s) were not a top match to any other network.

To assess whether visual network 3 splits into dorsal and lateral representations as per the group analysis, we first inspected all matches to the network above the $r = 0.3$ threshold. If a component correlated with visual network 3 above the threshold but it was also a top match for a different network, we disregarded this component as a candidate for dorsal/lateral splitting.

Components surviving this initial screening process (above threshold and a unique match to visual network 3) were then assigned to one of four categories based upon their splitting characteristics and spatial correspondence with visual network 3 (Figure 88):

1. Dorsal and lateral splitting. Separate components centred on V3A (dorsal division) and hMT+ (lateral division) were produced.

2. Visual network 3 only. A component corresponding to Smith's visual network 3 was produced, i.e., the dorsal and lateral divisions did not split into separate components.
3. Different splitting characteristics. These are discussed in section 5.3.6.5.
4. No network or splitting. A component corresponding to visual network 3 was not produced, nor was a dorsal and/or lateral component produced.

In most cases it is clear from visually inspecting the spatial maps whether the components reflect a dorsal/lateral splitting or whether they correspond to visual network 3. However, to be thorough, the classifications were validated through spatial correlations between the individual subject's spatial maps and three spatial maps generated from the group-level analysis. The group-based lateral (Gl, Figure 76, right, bottom) and dorsal (Gd, Figure 79) components were taken directly from the group-level ICA decomposition. The group-based combined component (Gc) was generated by combining the Gl and Gd components together (Figure 80, bottom). The proportions of each of the four categories, along with example subjects for each, are discussed below.

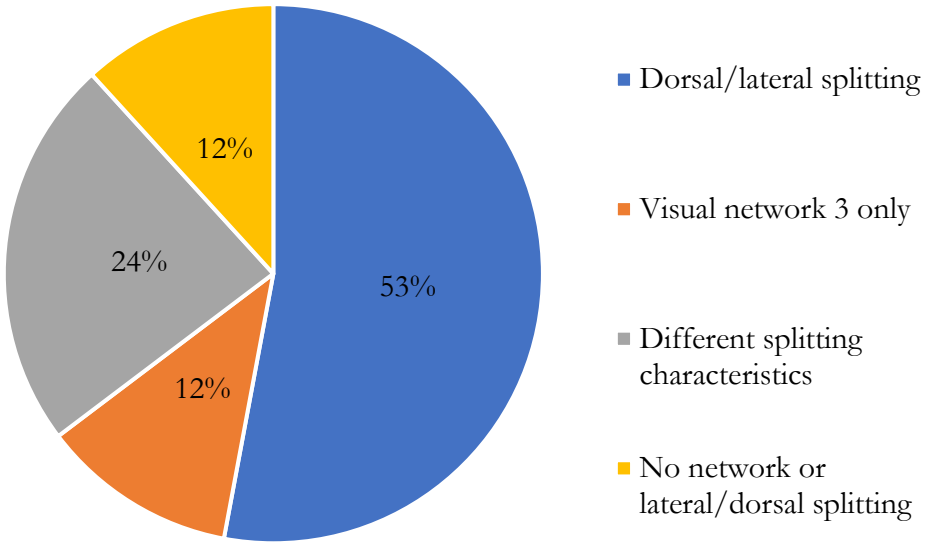

Figure 88. The proportion of splitting characteristics of components resulting from single-subject ICA.

5.3.6.3 Dorsal and lateral splitting

Fifty-three per cent (9/17) of subjects showed splitting of dorsal and lateral components, with the former showing activation in V3A and the latter showing activation in hMT+. These components did not correlate with any other visual network but did correlate with the corresponding group-based dorsal and lateral components. An example subject showing splitting of dorsal and lateral components is shown in Figure 89. One component correlated with the Gl component at r = 0.53 but did not correlate with the Gd component above r = 0.25 (Figure 89, top). A second component correlated with Gd component at r = 0.66 but did not correlate with Gl above r = 0.25. (Figure 89, bottom).

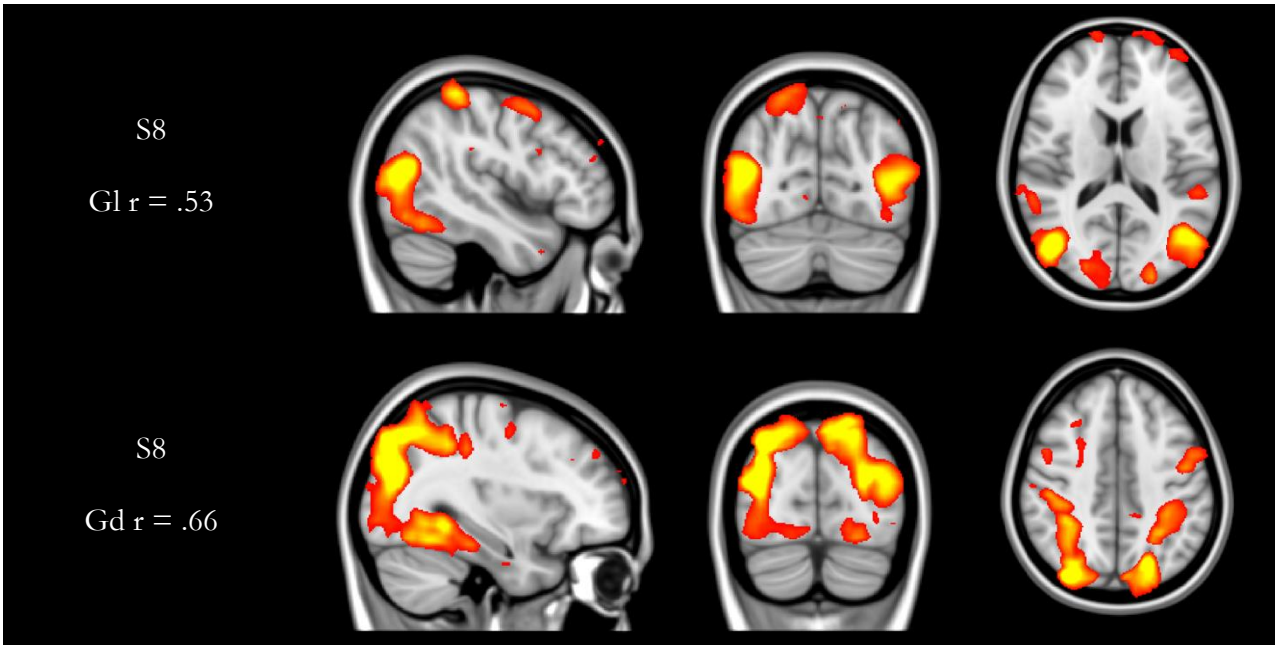

Figure 89. An example of dorsal and lateral splitting of components derived from single-subject ICA (Subject 8). The lateral component (top) correlated with the group-based lateral (Gl) component at r = 0.53. The dorsal component (bottom) correlated with the group-based dorsal (Gd) component at r = 0.66.

5.3.6.4 Visual network 3 only

Eleven per cent (2/17) of subjects did not show a clear splitting of dorsal and lateral components but did output a component that represented Smith's visual network 3 as a whole (Figure 90). These components correlated highly with Smith's visual network 3 and the Gc component.

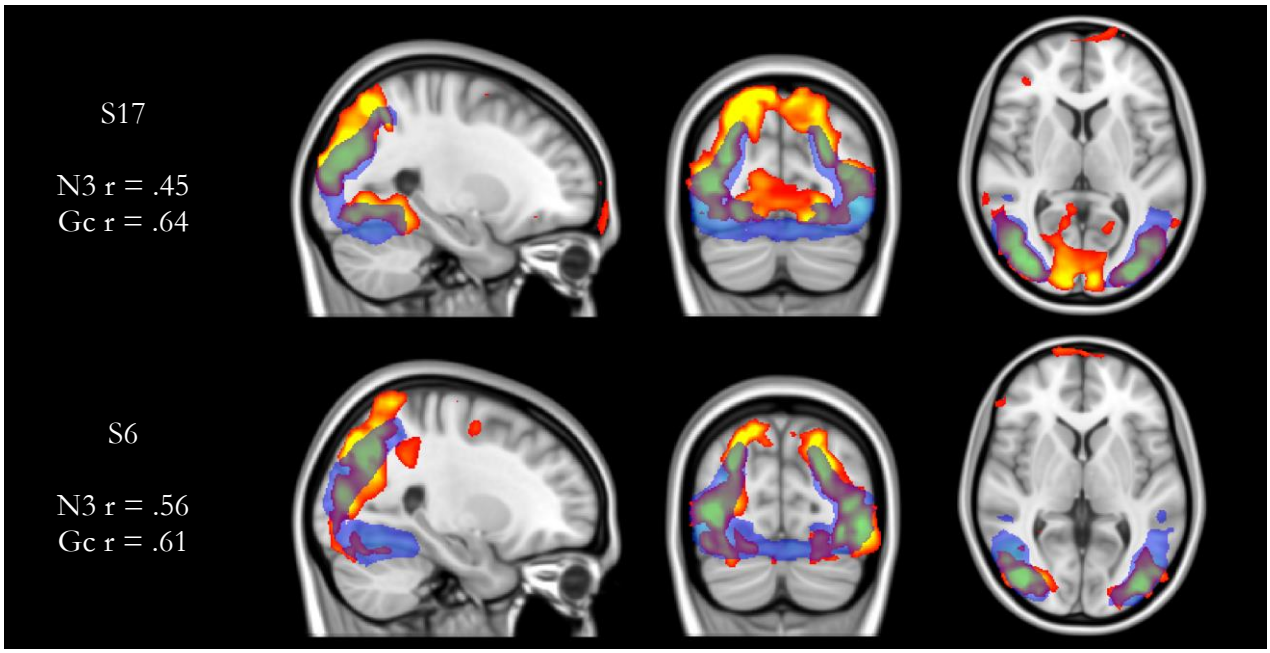

Figure 90. Two examples of components which represent Smith et al.'s visual network 3. The component derived from subject 17's data (top) correlated with Smith's visual network 3 (N3) at r = 0.45 and the group-based lateral and dorsal combined component (Gc) at r = 0.64. The component derived from Subject 5's data (bottom) correlated with N3 at r = 0.56 and the Gc component at r = 0.61.

5.3.6.5 Different splitting characteristics

Twenty-four percent (4/17) of subjects produced components that did not follow any specific trend. These are discussed below.

Subject 9 produced a component which upon visual inspection appeared to represent Smith's visual network 3 as a whole (Figure 91, top). This was reflected by a correlation with the network of r = 0.44. A separate dorsal component was also produced which did not correlate with Smith et al.'s visual network 3 above threshold but did correlate with the Gd component at r = 0.48 (Figure 91, bottom). This component appears to involve IPS regions, rather than V3A.

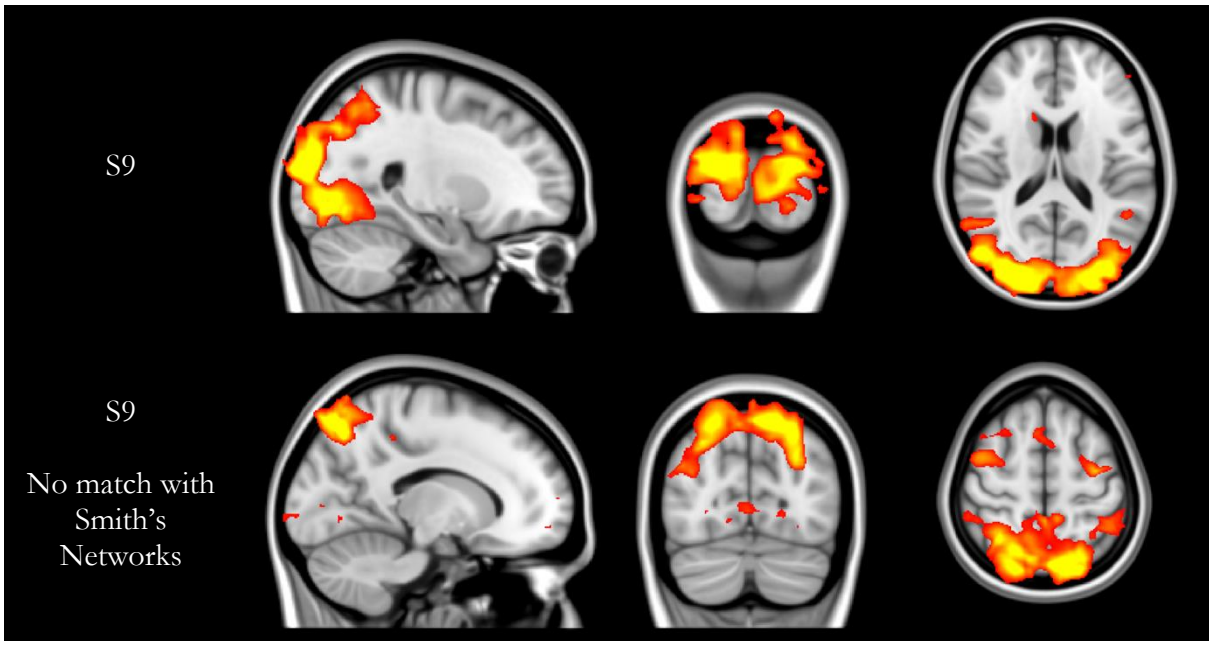

Figure 91. Two components derived from single-subject ICA (Subject 9). One component appears to represent Smith et al.'s visual network 3 (top) and correlated with the network at r = 0.44. A second component was produced which appears to include the IPS regions with no clear activation in the lateral-occipital regions or V3A.

When dimensionality was set to 25, Subject 12 produced a single match with visual network 3 (r = 0.36). The component is not a clear representation of Smith's network, but rather a more medial-occipital component which extends dorsally past V3A and does not appear to include the lateral-occipital regions (Figure 92, top). In fact, this component shows greater correlation with Smith's visual network 1 (r = 0.42) than it does with visual network 3.

To be thorough and following the logic that with single-subject ICA a greater dimensionality is necessary for visual network 3 to split into dorsal and lateral divisions, I performed ICA with dimensionality set to 30. At the higher dimensionality, the only component to correlate with visual network 3 above threshold (r = 0.36) correlated greater with visual network 2 (r = 0.41). However, correlations between the spatial maps and the group-based components reveal a clean lateral component which correlates with the Gl component at r = 0.47 (Figure 92, bottom).

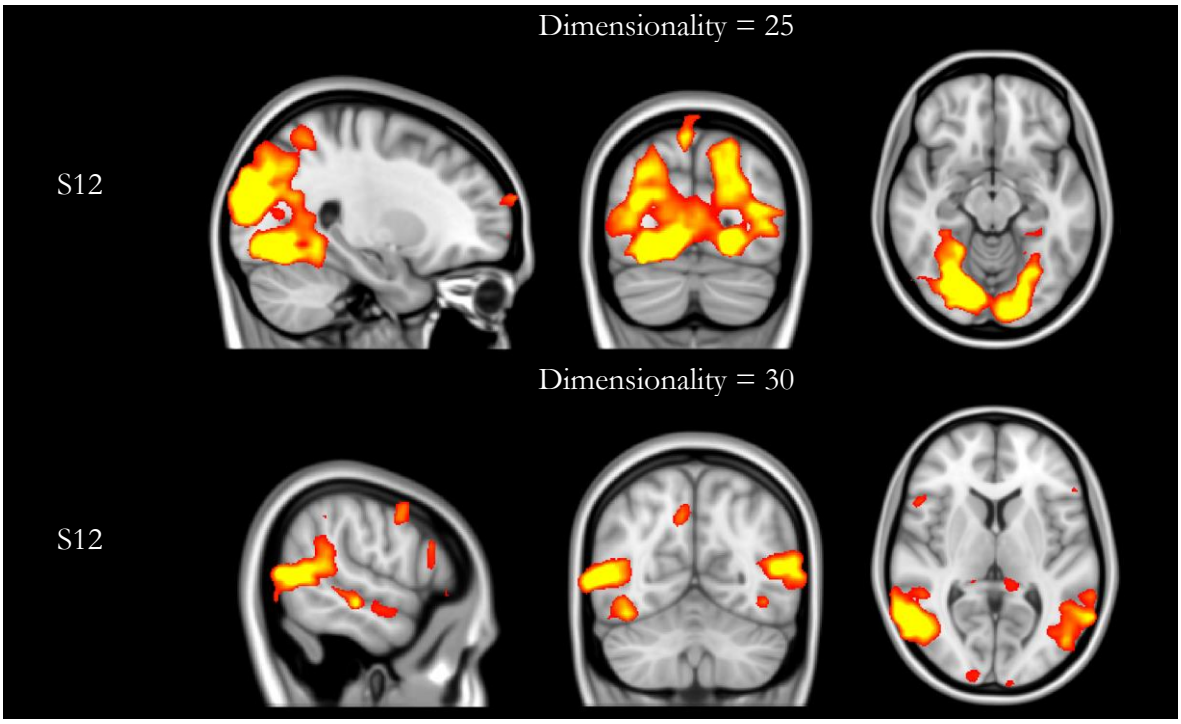

Figure 92. Two components derived from single-subject ICA (Subject 12) with dimensionality set to 25 or 30. At the lower dimensionality (20) a medial-occipital component which extends dorsally past V3A is produced (top). This component does not appear to include the lateral-occipital regions. When dimensionality was increased to 30 a clean lateral component with no dorsal extension was produced (bottom).

Subject 15 produced a clear lateral component, encompassing hMT+, with no clear dorsal extensions (Figure 93, top). The component correlated with the Gl component at r = 0.44. A second component (Figure 93, bottom), which correlated with Smith's visual network 3 (r = 0.35) and the Gd component (r = 0.31), does not represent the dorsal regions as cleanly as seen for the other subjects. Therefore, although there is splitting of a lateral component, we did not classify this as evidence of dorsal/lateral splitting.

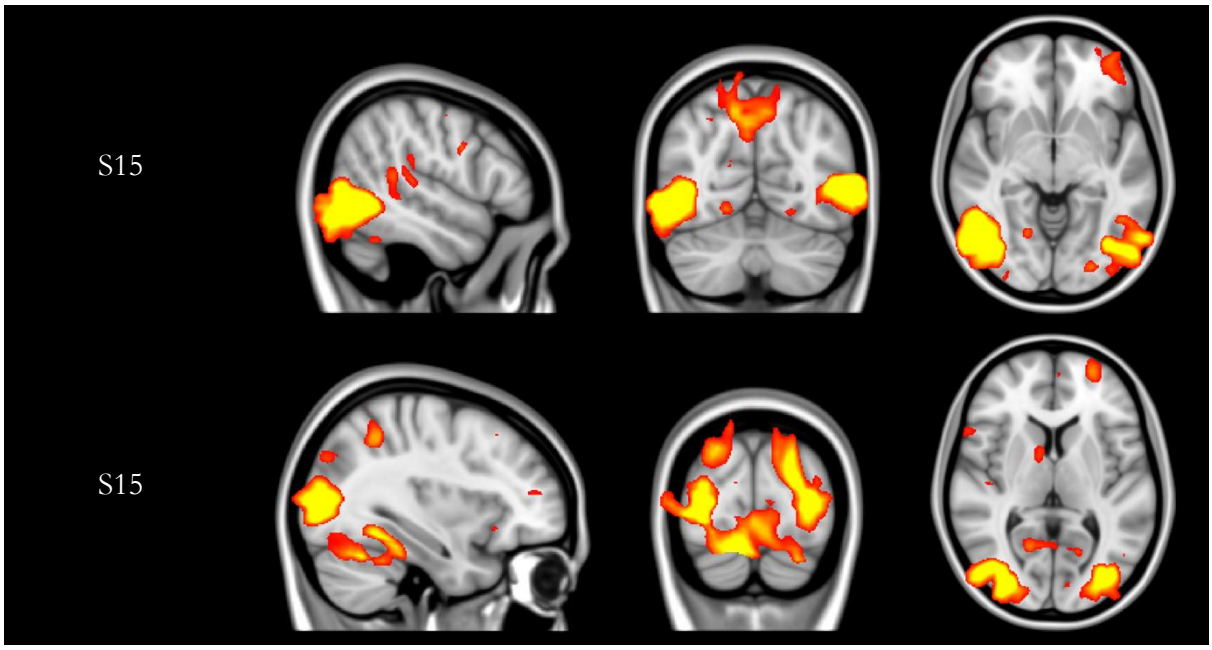

Figure 93. Two components derived from single-subject ICA (Subject 15). This subject produced a clear lateral-occipital component (top) which correlated with the group-based lateral component at r = 0.44 and an incomplete dorsal component (bottom). The dorsal component does not represent the dorsal regions and clearly as the other subjects.

When dimensionality was set to 25, Subject 22 did not output a unique match to visual network 3, the component was also the top match to visual network 2. When dimensionality was set to 30, a dorsal network component was output which correlated with Smith's visual network 3 at r = 0.36 (Figure 94). The component does not include lateral regions but does correlate with the Gd component at r = 0.51. As was the case with Subject 9, this component appears to involve IPS regions, rather than V3A.

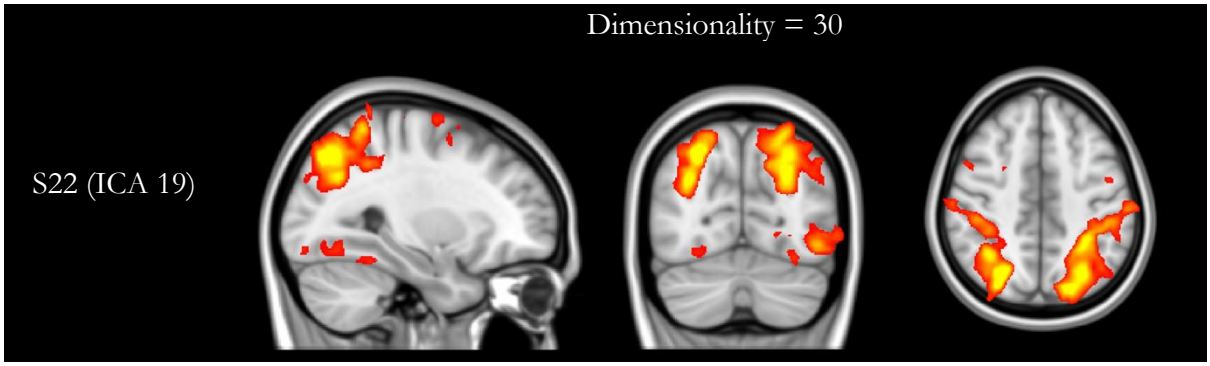

Figure 94. A component derived from single-subject ICA (Subject 22) with dimensionality set to 30. This component appears to include the IPS regions with little-to-no activation seen in the lateral-occipital regions or V3A.

5.3.6.6 No network or dorsal/lateral splitting

Finally, two subjects did not produce a clear match with Smith's visual network 3 or the splitting of dorsal and lateral sub-regions. A unique component for visual network 3 was not output by either of these subjects at any dimensionality.

5.4 Chapter summary

The results presented in this chapter provide supporting evidence of a third visual pathway which is distinct from the dorsal and ventral streams (Gilaie-Dotan, 2016). The third, lateral visual motion, pathway is hypothesised to mediate rapid visual motion perception across the entire visual field, which is facilitated by non-hierarchical connections that swiftly propagate motion information in parallel to TO1/MT. This information is then relayed to several other regions, including TO2/MST. Gilaie-Dotan (2016) hypothesised that V3A and other dorsal stream regions are not directly involved in the lateral visual motion pathway.

I performed group-ICA on fMRI data acquired while subjects free-viewed a movie and compared, through spatial correlations, the components to a set of well-established resting state networks (Smith et al., 2009). Two of these resting-state networks, visual networks 2 and 3, appear to represent the traditional ventral and dorsal visual pathways, which have long been considered to mediate vision for perception and action respectively (Goodale & Milner, 1992; Ungerleider & Mishkin, 1982).

Each resting-state network is represented by a single component in our data, except for visual network 3 for which we find two candidate components. According to Smith et al. (2009), visual network 3 covers the lateral-occipital regions, including the motion-sensitive V5/hMT+, and a dorsal extension past V3A, another region implicated in visual motion perception. In our movie watching data, we find these two motion-related regions split into separate components, and when combined correlate greater with visual network 3 than either does on its own. These areas are not separated when ICA is performed on rs-fMRI data (Smith et al., 2009).

One explanation as to why hMT+ (lateral motion pathway) and V3A (dorsal pathway) are segmented into two separate components in our data but not in resting state could be that the lateral motion pathway needs to be "active" in order to fragment from the dorsal pathway. Resting-state data is acquired in the absence of a visual stimulus, specifically the observers are not exposed to visual motion. This lack of a motion stimulus could mean the spontaneous activity in hMT+ and V3A at rest is not sufficiently different and therefore ICA does not segment visual network 3 into separate components. On the other hand, the movie contains a range of motion features which may drive different responses in the two motion areas which have been shown to have different roles in processing global and local motion (Cai et al., 2014). ICA is sensitive to these different responses and therefore segments the two areas.

Furthermore, correlations between the components and a probabilistic atlas of the visual areas (Wang et al., 2015) suggests the dorsal component includes IPS0 (also referred to as V7) through to IPS4, but not IPS5, which is included in the lateral component. The residual activity around the post-central sulcus present in the lateral component may represent IPS5 which is located in the anterior-lateral branch of the IPS and extends toward the post-central sulcus. However, we acknowledge that drawing any firm conclusions about the inclusion of specific IPS regions comes with the caveat that there is a high level of anatomical variance in these regions between subjects, and therefore it is difficult to say with certainty that IPS5 was not represented by the dorsal component. This variability is highlighted by the authors of the probabilistic atlas used in the present analysis. However, the inclusion of IPS regions in the dorsal component suggests the activity in these regions are more closely related with V3A than to hMT+. This could be indicative of the hierarchical dorsal pathway in which information is propagated along the areas from V3A anteriorly along the IPS (Cai et al., 2014).

5.4.1 Social perception pathway

More recently, the concept of a third pathway has been further explored. Pitcher and Ungerleider (2021) presented evidence of a third cortical pathway on the lateral surface that mediates the dynamic aspects of social perception, referred to as the "social perception" pathway. Although not

acknowledged in the paper, the social perception pathway has similar characteristics to the lateral motion pathway hypothesised by Gilaie-Dotan in 2016. While Gilaie-Dotan (2016) placed emphasis on faster response latencies and anatomical connections to distinguish between the dorsal and visual motion pathways, the social perception pathway places particular emphasis on the regions involved in the processing of biological motion. The pathway is proposed to be charged with propagating information about moving bodies and faces to the STS, via V5/MT+.

This pathway is of particular relevance here as the movie stimulus used in the present study contains elements of biological motion, such as moving faces and bodies, and object motion. However, the lateral component which centres around hMT+ does not appear to include activation in the STS. Therefore, the components output from movie watching data appear to resemble the lateral visual motion pathway more than the social perception pathway.

5.4.2 Single-subject ICA

At the group level, setting dimensionality to 20 is sufficient to enable components representing the visual networks to split into dorsal and lateral divisions. This splitting was then examined at the subject level by performing single-subject ICA. We found dimensionality of 25 to be more optimal relative to a dimensionality of 20 at the subject-level. Increasing the dimensionality from 20 to 25 did not alter the overall network representativeness of the components, as reflected by a modest increase in the mean correlation across networks. It did, however, allow visual network 3 to split into dorsal and lateral sub-components. For over half of the subjects (52%), clear splitting of dorsal and lateral regions was evident. The majority of the remaining subjects were shown to either output a network component representing Smith's visual network 3 as a whole or showed some other form of dorsal/lateral splitting, e.g., a clear lateral component but an incomplete or disordered dorsal component.

For subjects who did not show clear evidence of splitting at dimensionality of 25, we examined whether increasing dimensionality further results in a separation of dorsal and lateral components. We found increasing dimensionality to 30 did not lead to a clear splitting of the regions. However, there was an interesting difference noted for two subjects. At either dimensionality, neither Subject 12 nor

Subject 22 produced a component that corresponded with Smith's visual network 3. However, at the higher dimensionality an additional lateral component was output for Subject 12 which did not appear at the lower dimensionality. On the other hand, when dimensionality was increased, a dorsal network component excluding the lateral regions was output from Subject 22's data.

Finally, given that movie watching yields comparable results to those obtained from resting-state data, it offers a suitable and more engaging substitution to asking subjects to fixate a central target or lay with their eyes closed for the duration of the scan.

In summary, we find when subjects watch a movie, visual network 3, believed to represent the dorsal visual pathway, splits into two components: one covering V3A, one covering hMT+. This finding is broadly in line with Gilaie-Dotan's (2016) hypothesis which opposes the traditional two pathways model of visual processing, instead suggesting a third pathway (the lateral motion pathway) is responsible for the processing of visual motion. The remaining resting-state networks can be accurately and reliably identified using movie watching fMRI data, suggesting on the whole the networks do not reorganise when observers are exposed to a rich naturistic stimulus.

Chapter 6. Motion psychophysics

6.1 Introduction

In Chapter 1 I outlined the original research questions to be investigated in this thesis: (i) is macular sparing "true" vision or can it be explained as an artefact of visual field testing? And (ii) what non-conscious visual abilities do individuals with cortical vision loss possess, and which underlying neurological pathways serve these non-conscious visual abilities? Note, although I refer to these residual functions as non-conscious, patients are often more "aware" of moving stimuli than they are of static ones (Azzopardi & Cowey, 1998). However, as discussed in Chapter 1, due to the COVID-19 pandemic patient data collection was not possible. I changed track and made use of an archival dataset collected with healthy adults. The chapters that followed discussed the neuroimaging methods I developed and/or employed to instead address the revised PhD aims using this archival data. The revised aims grew from, and were motivated by, the original topic area and can be adapted for use with the clinical population I had originally aimed to investigate.

In the following chapter I describe the experiment we had previously designed to systematically investigate different forms of residual visual function in individuals with homonymous hemianopia. Each task was designed to tap into a different level of processing of stimuli within the blind field.

6.1.1 Homonymous hemianopia

Vision loss following brain damage is not uncommon. In a sample of 220 patients with acquired brain injury, 46% displayed some form of visual field loss (Suchoff et al., 2008). Homonymous hemianopia (HH) is a common visual field defect arising from unilateral damage to the visual cortex[33]. In the case of HH, patients report a loss of one half of their visual field in the same relative position in each eye (a homonymous defect). Although there are mixed reports as to the exact prevalence of HH, it is widely reported that 30% of stroke admission patients display some form of

[33] Note, lesions anywhere between the optic tract and the primary visual cortex (retrochiamisal pathway) can also give rise to a range of cortical visual impairments, including HH (Zhang et al., 2006a).

homonymous visual field defect (HVFD; Pambakian et al., 2004), of which the dominate aetiology is ischemic or haemorrhagic stroke (Townend et al., 2007).

In clinical practice, perimetry is used to assess and characterise HH; perimetry is a systematic measurement of an individual's visual field, involving the detection of the presence of a test light stimulus presented at varying locations in the visual field, from close to fixation to the far periphery. HH is defined by a deficit in the ability to respond to stimuli presented in the contralesional visual field.

6.1.2 Macular sparing

Patients with HH sometimes retain the ability to detect stimuli presented close to fixation in their otherwise "blind" hemifield when tested with perimetry. This phenomenon is known as macular sparing. It is estimated that macular sparing is found in 7% of patients with HH (Zhang et al., 2006b), and although it has been extensively debated for over a century (Holmes, 1918), an accurate interpretation of the phenomenon remains unusually elusive (for historical reviews see Leff, 2004; Horton et al., 2021).

One hypothesis is that macular sparing is not "true" spared vision and is in fact an artefact of perimetry resulting from fixation instability. As outlined in Chapter 1, prior to the pandemic, we had aimed to systematically investigate this hypothesis using three perimetry systems (two of which are commercially available), each with different levels of gaze stabilisation. Though the change of direction resulting from the pandemic meant this experiment is omitted from this thesis, I do, for completeness, summarise our proposed approach below.

The Humphrey Field Analyzer 3 (HFA3; Carl Zeiss Meditec, Dublin, CA, USA) is an automated perimetry system that is widely used for assessing visual field defects, such as HH, both in a clinical and research setting. However, the HFA3 has no means of accounting for eye movements during the assessment making it particularly vulnerable to eye movement artefacts.

Microperimetry offers a partial solution. Microperimetry, is a technique that maps contrast sensitivity very much the same way standard perimetry does, but with specific focus on the central 10°

of the visual field and has a crude gaze stabilisation feature. The Macula Integrity Assessment (MAIA, CenterVue, Padova, Italy) is a microperimetry system that updates the stimulus presentation location, with respect to the location on the retina, prior to the presentation. This can account for eye movements made prior to the stimulus presentation, but the stimulus remains fixed during the 200 ms it is presented.

The third, gaze-contingent system, uses simultaneous non-invasive eye tracking to present stimuli to specific regions of the retina, regardless of the gaze direction (Thomas et al., 2019). This custom perimetry system tracks the eye position and, unlike the microperimetry system, updates the stimulus position accordingly, even during presentation.

We had aimed to use the three perimetry systems outlined above to assess whether or not eye movements play a critical role in the expression of macular sparing. This could be achieved by comparing, across systems, the sensitivity values obtained for stimuli presented to the blind hemi field. For example, if macular sparing is an artefact of eye movements, we would expect to see sparing when using the HFA3, and potentially with the MAIA. The latter would depend on when the eye movements occur relative to stimulus presentation. However, if sparing is simply a consequence of eye movements during the visual field assessment, we would not expect to see spared vision when tested with the gaze-contingent system.

Although the perimetric experiment outlined above is omitted from this thesis, both the HFA3 and MAIA are used to assess the visual field of patients prior to our motion psychophysics experiment (described in detail below).

6.1.3 Residual visual function

This chapter focuses on the counter-intuitive phenomenon in which, despite being apparently blind, the behaviour of some patients with cortical vision loss can be changed by stimuli of which they may be unaware. This is referred to as residual visual function. Some patients with cortical vision loss are able to look (Pöppel et al., 1973) or point (Weiskrantz et al., 1974) towards the location of objects in their blind field, while at the same time denying that they can 'see' them in any conventional sense.

In some cases, individuals report an awareness of moving objects in their blind field (Riddoch, 1917; Zeki & Ffytche, 1998), and when forced to make a judgment about the object perform significantly above chance.

The first report in the literature of residual visual function traces back to 1917 when Riddoch conducted a series of perimetric studies on soldiers following gunshot wounds that primarily damaged the calcarine cortex (V1). These patients showed a limited but conscious percept of moving stimuli within their blind fields; limited because they could not characterise any other attributes of the stimulus, e.g., shape or colour, but conscious insofar as they knew something had moved (Riddoch, 1917). Remarkably, research into the phenomenon stagnated for over half a century.

Research into residual visual function accelerated in the early 1970's, when studies using two-alternative force choice (2AFC) and two-alternative forced-response (2AFR) paradigms found patients with cortical vision loss could accurately detect, discriminate (e.g., horizontal vs. vertical lines) and perform visually guided behaviours (e.g., pointing or saccades) towards blind field stimuli, despite denying they ever saw them (Pöppel et al., 1973; Weiskrantz et al., 1974). Following this early work, a range of intact visual capabilities have been documented (for review see Stoerig, 2006). For example, Solcà et al. (2015) described a patient with cortical blindness who could accurately discriminate between jumbled and normal faces, as well as between known and unknown faces, indicating intact facial processing in the absence of awareness.

Of the multiplicity of residual abilities, motion processing may be considered the most robust, with moving stimuli typically more detectable than static stimuli (Bridge, 2020). Again, through the use of 2AFC paradigms, patients with cortical vision loss have been shown to be consistently able to accurately detect the presence of and discriminate the direction of a range of motion stimuli presented in the blind field (Azzopardi & Cowey, 1998; Weiskrantz et al., 1995; Zeki & Ffytche, 1998)

Retained motion processing in the blind field is likely to be facilitated by an intact, direct pathway from LGN to hMT+ (Bridge et al., 2008). The preservation of the LGN-hMT+ pathway can predict the presence of intact motion processing in patients with cortical vision loss (Ajina, Pestilli, et al., 2015). Despite the destruction of V1, the human motion area (hMT+) in the damaged hemisphere

is activated by moving stimuli presented to the blind hemifield (Barbur et al., 1993; Zeki & Ffytche, 1998). It is important to note that it is unlikely that patients with cortical damage develop this "alternative pathway" which avoids V1, rather it is present in all brains but is dominated by the geniculo-striate pathway (Bridge, 2020).

Much of the previous research investigating residual motion processing has focused on whether or not an individual with cortical vision loss can detect the presence of moving stimuli in their blind field. However, it is important to note that other forms of residual motion processing may remain intact. Here, we aimed to investigate the relationship between different forms of motion processing that are retained[34]. To investigate this, I adopted three psychophysical tasks sensitive to different forms of motion processing. Each task measured three different types of responses to blind field stimuli:

1. Explicit responses require the subject to make a direct judgement about a blind field stimulus, much like the tasks described above. Here I use the temporal motion detection (TMD) in which an aperture of black and white dots is presented to the blind field. In one of two temporally separated intervals the dots coherently moved left or right. The subject is instructed to indicate which of the two intervals they believed contained the motion.

2. Implicit responses require the subject to make a judgement about a seeing field stimulus, whilst a second stimulus is placed in the blind field. The effects of the blind field stimulus on the seeing field stimulus are then measured. Here, I use the flow-parsing task (Rushton & Warren, 2005) in which the subject fixates a central target and judges whether a probe dot located in their seeing field moved towards or away from fixation. While the probe moves a contracting or expanding radial flow pattern is presented simultaneously to the opposite (blind) visual field. Previous work (e.g., Warren & Rushton, 2009) has shown that in healthy brains the presence of an expanding flow field induces an inward motion component to the probe.

[34] Prior to the pandemic, we had also aimed to assess the relationship between task performance and the location of the cerebral damage. Due to the inability to collect patient data, the neuroimaging methods are omitted here.

3. Reflexive responses do not require the subject to make any judgments about a stimulus, instead unconscious responses to blind field stimuli are measured. In the present study, I measured the subjects eye movements, specifically optokinetic nystagmus (OKN), in response to a blind field sinusoidal grating.

The present study aimed to assess whether there was a dissociation between performance on each of the three tasks. For example, do some patients show a flow-parsing effect while performing at chance on the TMD task? Or are the responses all-or-nothing? e.g., if a patient is above chance at TMD, do they also show a flow-parsing effect and an OKN response?

The TMD is considered to be the standard approach for detecting spared motion processing in individuals with cortical vision loss (e.g., Ajina, Rees, et al., 2015; Ajina & Bridge, 2018), hence its inclusion in the present study, while the use of the flow-parsing and OKN tasks to investigate residual visual function is more novel. In the following sections, I provide some background to the visual processes at the centre of these two novel tasks.

6.1.4 Eye movements

Humans possess a repertoire of dynamic eye movements, of which there are seven basic classes: saccades, smooth pursuit movements, vergence movements and nystagmus quick phases are responsible for gaze-shifting, while vestibulo-ocular reflex (VOR), optokinetic nystagmus (OKN) and fixation maintenance movements are responsible for gaze-stabilisation (Metz, 1983). The eye movements pertinent to this chapter are described in detail below.

6.1.4.1 Saccades

Saccades are ballistic eye movements that allow us to rapidly shift our gaze from one fixation point to the next. This mechanism is fundamental in placing the fovea, the portion of the retina responsible for high-acuity visual processing, on objects of interest in the visual field. Saccades are ballistic in nature as they have a pre-determined destination, that is to say that once a saccade has been generated the eyes cannot respond to changes in the position of the target during the movement (Purves, Augustine, Fitzpatrick, Katz, LaMantia, McNamara, & Williams., 2001). Saccades can range in

amplitude, from small movements whilst reading, to large movements when scanning a room of people.

6.1.4.2 Optokinetic nystagmus (OKN)

OKN is an eye movement which allows us to stabilise the moving world on the retina while the head is still. OKN is comprised of two phases: the OKN quick phase and the OKN slow phase. When a large, continuously moving stimulus is presented, OKN minimises the retinal slip provoked by the stimulus by cycling these two phases. First, the eyes follow the stimulus in the direction in which it is moving (OKN slow phase), once the maximum rotation is reached the eyes rapidly move in the opposite direction (OKN quick phase) to re-centre the image and hold it steady on the retina (Leigh & Zee, 2015). A frequently used example of the OKN response is that of watching from the platform as a train passes through the station. The observer will track the train as it passes, this is OKN slow phase, once maximum rotation is reached the eyes make a saccade in the opposite direction, bringing gaze back to the centre of the scene, this is OKN quick phase.

6.1.5 Flow-parsing

Retinal motion can be caused by self-movement, movement of objects within the scene, or more often than not, a combination of the two. During self-movement the image of a stationary object within the scene is moving on the retina, yet we experience this object to be stationary. Likewise, objects moving within the scene during self-motion may be stationary on the retina if we fixate them, yet we are able to detect object movement. In order to successfully interact with the world, the brain needs to resolve this ambiguity and separate retinal motion into constituent parts, retinal motion due to movement of objects in the scene and retinal motion due to movement of the observer.

The flow-parsing hypothesis (Rushton & Warren, 2005; P. A. Warren & Rushton, 2007) offers a purely visual solution to the problem. The hypothesis is that the brain utilises sensitivity to optic flow information to parse retinal motion into self-movement and object-movement components.

Optic flow is the pattern of motion in the optic array (light that hits the back of the eye) caused by relative movement between the scene and the observer (Gibson, 1950). The optic flow we

experience is a valuable source of information relating to self-movement and has a role in the guidance of movement. For example, the processing of optic flow is implicated in control of posture, so much so that infants are caused to sway and, in some cases, fall forward or backward in response to discordant optic flow stimuli (D. N. Lee & Aronson, 1974).

It is well documented that the human brain has a particular sensitivity to optic flow (W. H. Warren & Hannon, 1988). The flow parsing hypothesis states that if the brain knows how the observer has moved, it should be able to isolate retinal motion due to object movement by 'subtracting' optic flow associated with self-movement from the retinal flow containing retinal motion due to both self- and object-motion. Following this subtraction, the remaining motion is attributed to object motion (P. A. Warren & Rushton, 2009b). This process is illustrated in Figure 95.

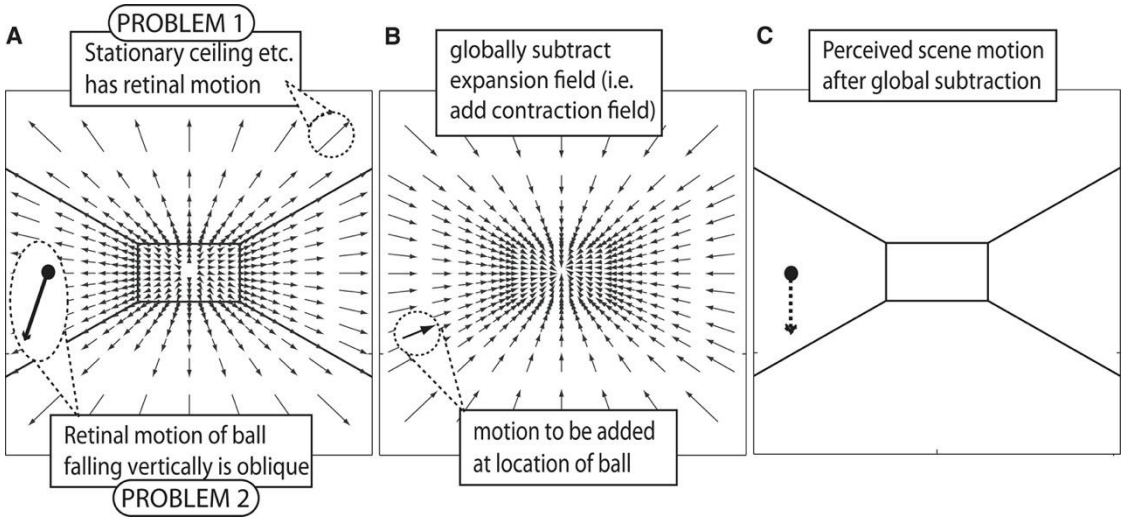

Figure 95. An illustration of the flow parsing hypothesis depicting an observer moving down a corridor while a ball falls vertically downward. Taken from Warren and Rushton (2009a). (A) As the observer moves down the corridor, the retinal image of the stationary elements of the scene (floor, walls, and ceiling) contain motion associated with forward translation. The retinal image of a vertically falling ball contains the true object motion (vertical) plus a radial component of motion caused by forward translation. This results in the retinal motion of the ball being oblique. (B) The brain detects and 'subtracts' the flow field associated with forward translation of the observer from the retinal image by adding to it a flow field in the opposite direction. (C) The perceived scene resulting from adding flow field (B) to flow field (A). The stationary elements of the scene are perceived to be stationary (perceptual stability), and the physical trajectory of the ball (vertical) has been recovered.

Consider an example in which an observer walks through a static scene. As the observer moves forward, the images of the stationary elements of the scene (e.g., the walls, ceiling, and floor) move across the retina and create an expanding radial pattern of motion. If the brain can identify and

subtract ("parse out") expansion flow field associated with forward translation the result is perceptual stability, stationary elements are indeed perceived to be stationary.

Consider an observer moving through a scene that contains a moving object. Figure 95A shows a ball that is, due to gravity, falling vertically. However, as the observer moves forward, the retinal motion of the ball is oblique, i.e., the retinal image of the ball contains the true object motion (vertical) plus a radial component of motion caused by forward translation. Because the brain is sensitive to optic flow patterns associated with self-movement, it is able to detect and subtract the flow field from the retinal image by adding to it a flow field in the opposite direction (Figure 95B). Flow-parsing (i) solves the perceptual stability problem, resulting in the static elements of the scene (e.g., the walls) being perceived to be stationary, and (ii) recovers the physical trajectory of objects (e.g., the ball is perceived to be moving vertically). The perceived scene is represented in Figure 95C.

One way to examine the flow-parsing effect is to use a radial flow field and a probe dot. The presence of an expanding radial flow field creates an optic flow pattern consistent with forward translation. Under the flow-parsing hypothesis, due to the global subtraction of the flow field, inward lateral movement of a stationary probe is induced. One way to quantify the induced component of motion is to null it out with physical (on-screen) movement, resulting in the probe being perceived to be stationary. Note, the perceived speed of the probe reflects the physical speed plus the induced speed. Therefore, if the perceived speed is zero then the induced speed is equal to the physical speed in the opposite direction. For example, in order to null the induced movement of the probe in the presence of an expanding flow field, the on-screen movement of the probe needs to be outward. Here, we use the speed of the probe at which the subject is equally likely to judge it as moving inward and outward as the "nulling speed", from this we can infer the induced movement caused by the parsing out of the flow field, this is known as the flow-parsing effect (Rushton et al., 2018).

If the flow field is not perceived, there should be no induced movement of the probe, and therefore the nulling speed of the probe should be zero. If, however, a patient shows a flow-parsing effect, i.e., an induced lateral movement of the probe, we can infer the brain is able to perceive and

globally subtract optic flow fields in the absence of V1, indicating the process may rely on a non-occipital route to hMT+ (Rushton et al., 2018).

6.2 Methods

6.3 The behavioural inattention task (BIT)

Unilateral visual neglect (UVN) is an attention deficit arising from cerebral damage, in which patients often fail to respond, report, or orient toward stimuli presented to the affected side (Menon & Korner-Bitensky, 2004). As research has shown UVN can masquerade as HH (Kooistra & Heilman, 1989), it is essential that we ensure the patient has a visual field defect, not an attentional defect. The BIT is a standardised battery of attentional tasks that can be administered to assess the presence of UVN. Here, we used the line bisection, star cancellation and line crossing tasks. Only subjects with scores in the normal range were invited to continue in the study.

6.3.1 Visual field testing

Each subject's visual field is assessed with two forms of perimetry. First, the Humphrey Field Analyzer 3 (HFA3) is utilised to identify and characterise HH in the subjects with visual field loss. For the flow-parsing and OKN tasks it is essential that the flow field and sinusoidal grating are presented exclusively to the blind field. Therefore, if a patient displays macular sparing, or islands of spared vision in their scotoma, they must be masked with an overlay matching the background of the seeing field stimulus, thus ensuring the potentially spared visual field does not receive unwanted stimulation.

If patients show macular sparing with the HFA, microperimetry is then performed using the MAIA system to further examine the spared vision. The HFA3 and MAIA are used in conjunction to indicate which regions of the visual field need to be masked. The HFA offers a way to assess for islands of spared vision in the central 60° of the visual field, but only starting 3° from the meridians, while the MAIA allows us to assess the central 2.25°, identifying locations that may be missed by the HFA[35].

[35] Pilot data collected from two patients with cortical vision loss using the HFA3 and MAIA can be found in section 6.4.

6.3.1.1 Humphrey field analyser 3 (HFA3)

The HFA3 projects a series of white light stimuli of varying intensities throughout a uniformly illuminated bowl. Subjects are instructed to fixate a yellow light in the centre of the bowl throughout the test and to respond with a button press if they believe a stimulus presentation has been made. Prior to the main assessment, subjects undergo a short training session in order to familiarise themselves with the protocol, this data is not analysed. The main assessment is comprised of 76 test points over the central 60° using a 30-2 grid (Figure 96). Stimuli are presented 3° from the x and y axes and at every 6° thereafter. Each stimulus was white in colour and the size was comparable to a Goldmann III (0.43° radius). The stimuli are presented, one at a time, for 200ms onto a background with the luminance of 31.5 apostilbs (asb; 10 cd/m^2). The Swedish Interactive Threshold Algorithm (SITA) standard threshold strategy was employed. Visual field testing is conducted monocularly, first on the right eye followed by the left eye.

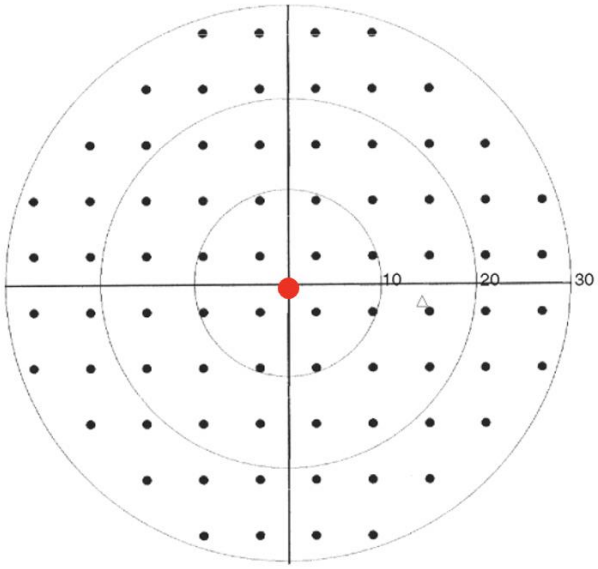

Figure 96: The test grid used by the Humphrey Field Analyser 3 (HFA3, Carl Zeiss Meditec, Dublin, CA, USA). Red circle indicates the fixation target. Black dots indicate the locations to which visual stimuli were presented during the visual field assessment.

6.3.1.2 Macular integrity assessment (MAIA) microperimeter

The Macular Integrity Assessment (MAIA; CenterVue, Padova, Italy) system is used to conduct microperimetry. The patient's gaze is directed to a central fixation cross. Identically to the HFA3 assessment, subjects are instructed to respond to a stimulus presentation with a button press and a

short training session precedes the main assessment. This training session uses a different grid (i.e., the visual field locations to which stimuli are presented) to that employed in the main assessment.

For the main assessment, a custom grid was designed consisting of 52 test points with 36 points sitting within the central 2.25° (Figure 97, right). The grid also has a further 12 points that extend to 9°. No test points are projected directly onto the vertical or horizontal meridians. The MAIA grid included 16 control test points that were shared with HFA3 (indicated by the black dots in Figure 97, right). These control test points allowed us to assess whether the results were consistent across systems.

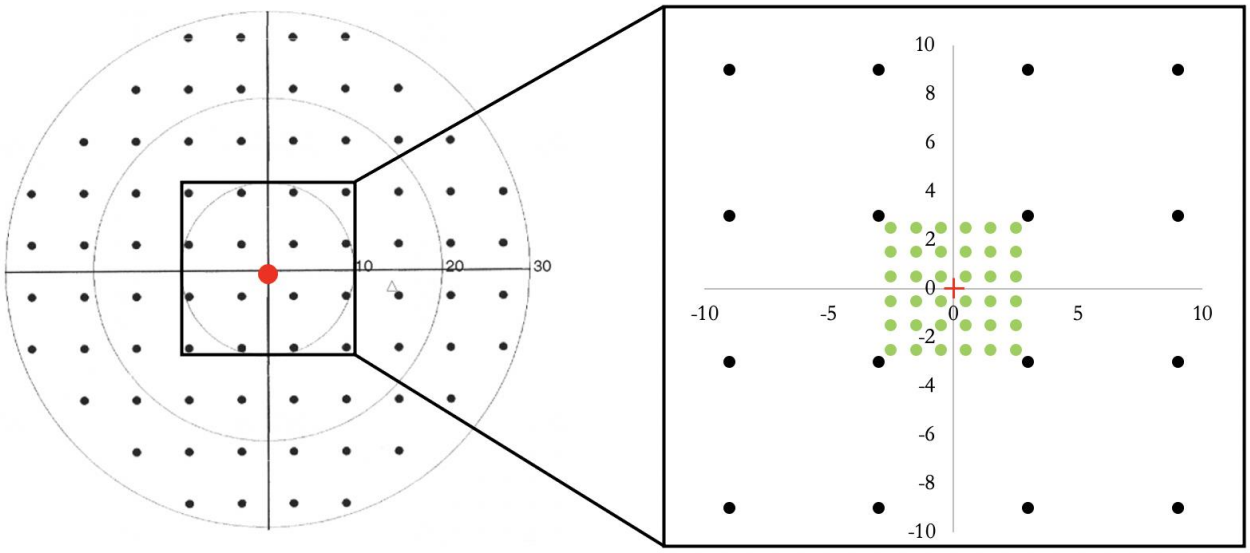

Figure 97. The testing grid used for the presentation of stimuli during standard perimetry (HFA3, left) and the custom grid used for microperimetry (MAIA, right). Red circle/cross represents the fixation target location. Black points on the MAIA grid denote shared presentation locations with the HFA, green points denote the additional presentation locations used in microperimetry.

Following the perimetric standard, the full threshold 4-2 strategy is administered which changes the light intensity in 4 decibels (dB) steps until there is a change from not seen to seen (or from seen to not seen), the intensity then changes in 2 dB steps until the stimulus is not seen again. Additional parameters for the MAIA were: stimulus size = Goldmann III (0.43° radius), background luminance = 4 asb (1.27 cd/m^2), maximum target luminance = 1000 asb (318.47 cd/^{2}) with a dynamic range of 0 to 36 dB, stimulus duration = 200ms. The field of view of the system was 20° × 20°. As

with the HFA3, testing is conducted monocularly, first on the right eye followed by the left eye. Prior to starting the experiment, the subject is dark adapted for 4 minutes.

6.3.2 Motion psychophysics

Following visual field testing, each subject completes three psychophysical computer-based tasks designed to target a specific visual function. Each of these tasks are outlined in detail below.

6.3.2.1 Display equipment

Stimuli are rendered using a Windows 7 PC, with an Intel core i5 processor. Stimuli are presented using an iiyama Vision Master Pro 513 CRT monitor. This monitor has a resolution of 1280 × 960 pixels and a refresh rate of 100 Hz. The viewing distance from the subject to the screen is 65 cm. The stimuli used in each task were generated in MATLAB using the Psychophysics Toolbox 3 (Psychtoolbox).

6.3.2.2 Eye tracking equipment

Eye movements are measured with a Tobii Spectrum Pro eye tracker (Figure 98). This eye tracker has a sample rate of 1200 Hz. The tracker is situated under the monitor with an optical path to the eyes of 65 cm.

Figure 98: Tobii Spectrum Pro eye tracker.

6.3.2.3 Study procedure

Tasks are performed in a descending order with regards to cognitive demand. Subjects first complete the optic flow task (considered to be the most cognitively demanding task), followed by the temporal motion detection task and finally the OKN task (considered to be the least cognitively demanding task). Subjects are seated with the monitor located at their midsagittal plane. An eye tracker calibration is conducted prior to the start of each individual task to ensure accurate tracking throughout the experiment. If the experimenter notes a series of fixation losses throughout a task, a recalibration is conducted. All tasks requiring a fixation target use the same bullseye/cross hair combination target. It has been documented that the use of this target (see Figure 99) results in a low dispersion and microsaccade rate (Thaler et al., 2013).

Figure 99: Target used in all tasks requiring a fixation target, comprised of a combination of "bullseye" and "cross hair" shapes.

6.3.2.4 Task 1: 2AFC flow-parsing

6.3.2.4.1 Visual stimulus

The stimulus for the flow-parsing task consists of a 1° red central fixation target, a hemifield radial flow field, and a dot probe, all of which were presented to a black background (Figure 100). This is a standard display used in optic flow processing research, as it is compatible with forward translation of the observer. The flow field consists of 8000 limited lifetime dots ranging in size from 0.5mm to 9mm. The dots move in virtual space in the x-axis only at 1500 mm/s. The flow field is rendered in 3D to create a "cloud" of dots, which is then compressed to the screen. The dots on the screen move radially, and at different speeds (dot speed is inversely proportional to the distance away from the centre of the screen). This stimulus creates a flow pattern similar to what would be experienced during forward/backward translation. The probe is 0.2° in diameter and is located in the opposite visual hemifield to the optic flow field, at 2° or 4° from fixation. To decrease the perception of a hard-edge as

the flow moves past the edge of the monitor, a mask starting 5° from the edge of the screen is applied to linearly fade the optic flow field.

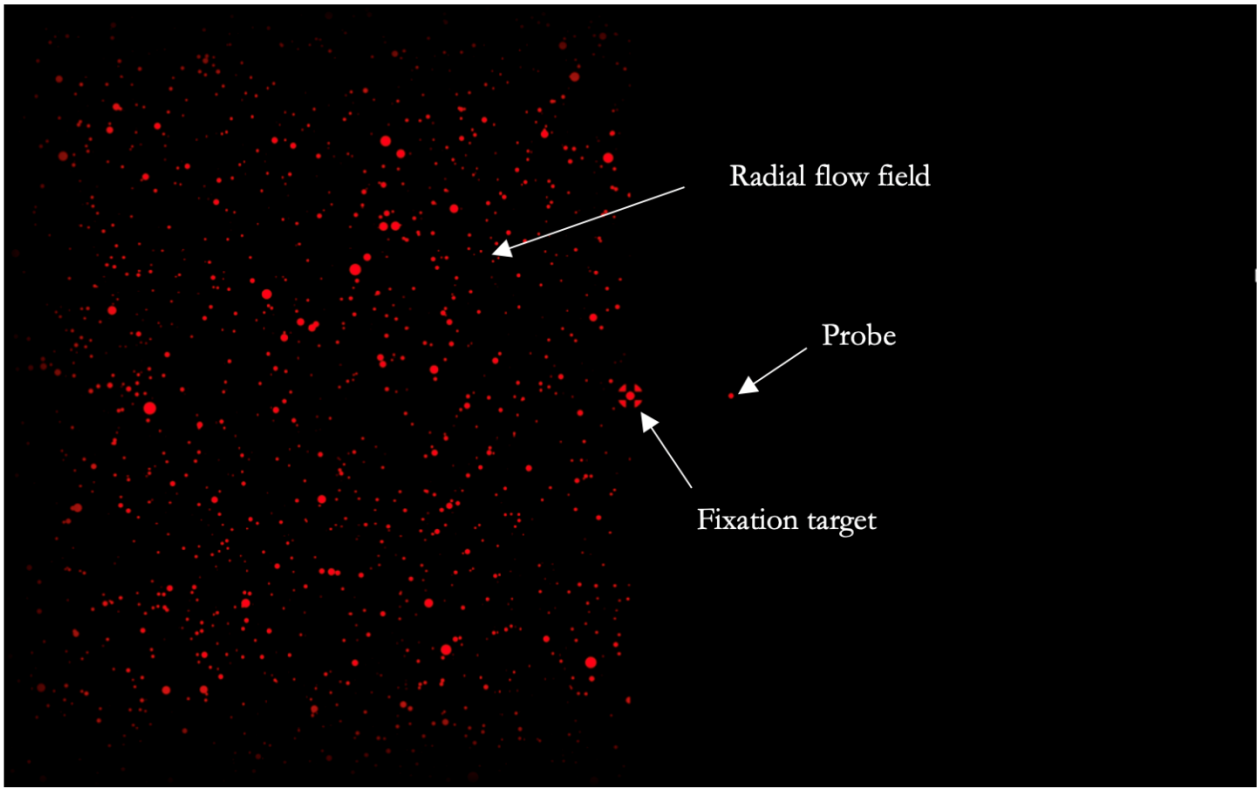

Figure 100: Stimulus used during flow-parsing task, consisting of a radial flow field, fixation target and probe.

6.3.2.4.2 Task procedure

Subjects perform a 2AFC task in which they are instructed to fixate the central target and indicate whether the dot probe moved inward or outward. Whilst the observer makes this direction judgment, the flow field moves in either an expanding radial pattern, consistent with forward translation of the subject, or in a contracting radial pattern, consistent with backward translation of the subject. A single trial begins with a 1000 ms period in which the flow field does not expand or contract but appears to "twinkle" and the probe is not present. Following this, the probe appears at one of the two locations depending on the condition (2 or 4°) and moves linearly either towards or away from fixation. The expansion/contraction of the flow field and the movement of the probe are simultaneous over a 500 ms period. Eye position is monitored during each trial. An annulus, 1° in diameter, is used as an eye tracking tolerance zone. In the event that the subject's gaze deviates outside this annulus, the

fixation target increases in size, an auditory beep is presented to instruct the observer to bring their gaze back to the centre of the screen, and the trial is repeated.

The QUEST algorithm identifies the subject's point of subjective equality (PSE) using a staircase procedure that controls the speed and direction of the probe. The PSE represents the physical movement of the probe that is equally likely to be reported as inward or outward, i.e., the physical movement nulls the induced movement, resulting in the observer perceiving the probe to be stationary. Therefore, the nulling speed provides a measure of the induced speed[36].

One staircase was completed for each of the four conditions, with each condition having a minimum of 20 trials completed:

1. 2° probe with flow expanding,
2. 2° probe with flow contracting,
3. 4° probe with flow expanding,
4. 4° probe with flow contracting.

6.3.2.5 Task 2: 2AFC temporal motion detection

6.3.2.5.1 Visual stimulus

The TMD stimulus consists of an aperture, 8° in diameter, containing 400 limited-lifetime dots, each 0.075° in diameter. 50% of the dots are white, 50% are black. The dots within the aperture are either static or coherently moving left or right. A 1° black fixation target is present throughout the task. The aperture is presented in the blind hemifield (ensuring it is not close to any previously identified islands of spared vision) on a mean luminance grey background (see Figure 101).

[36] Pilot data collected from four control subjects can be found in Appendix XIV.

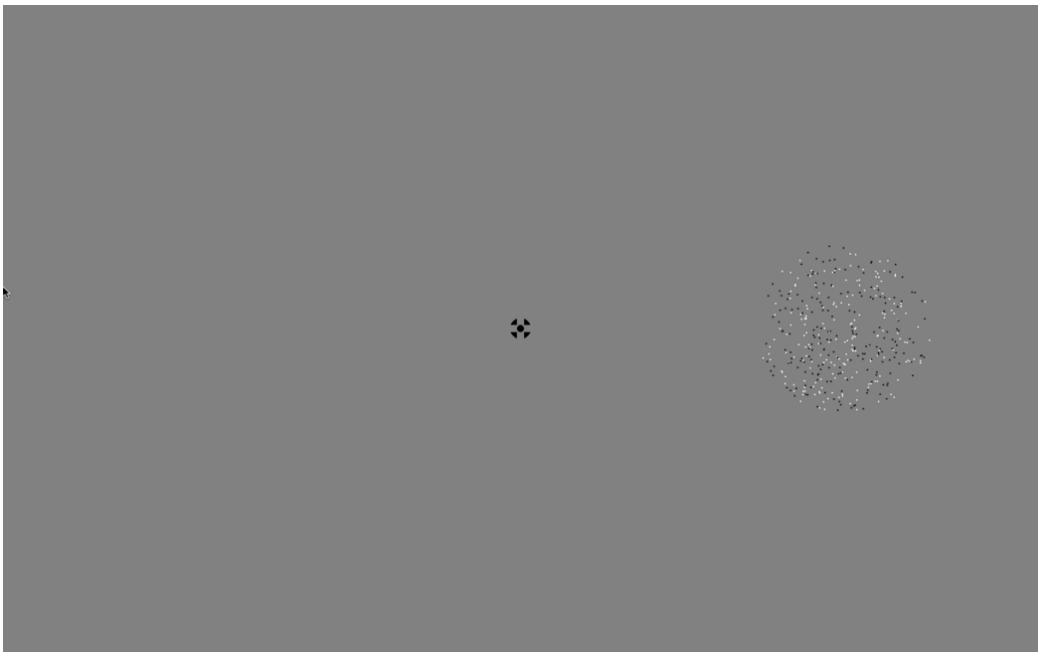

Figure 101: Stimulus used during the temporal motion detection task, consisting of a fixation target and aperture containing either static or coherently moving dots.

6.3.2.5.2 Task procedure

Using a 2AFC paradigm, subjects are instructed to indicate whether the stimulus appeared in the first or second interval (Figure 102).

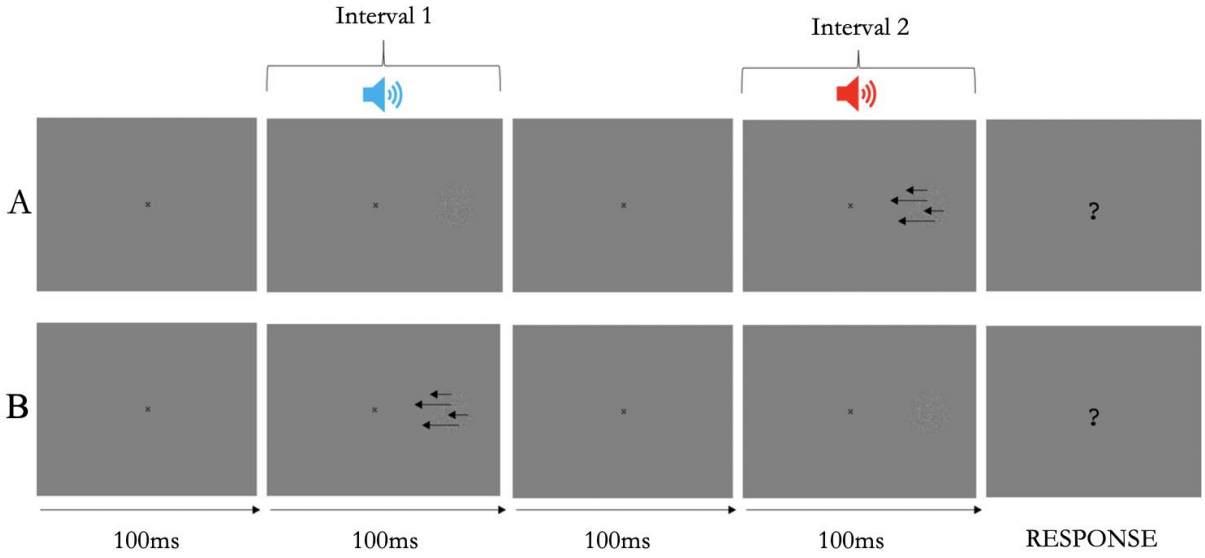

Figure 102: Temporal motion detection paradigm. Condition A: coherent motion of dots within the aperture is present in interval 2. Condition B: coherent motion of dots within the aperture is present in interval 1.

The start of each interval is indicated by a 500 ms auditory tone. The onset of the tone is synchronised with the onset of the stimulus. A 300 Hz auditory tone marks the start of interval 1, while a 600 Hz auditory tone marks the start of interval 2. Each interval lasts 1000 ms and is separated by a

1000 ms inter stimulus interval (ISI) consisting of a blank grey screen on which the fixation target remains. To indicate to the subject when to respond, a black question mark replaces the fixation target 1000 ms after the offset of the second interval. Responses are made with a keyboard press (1 for first interval, 2 for second interval). If the subject states they saw nothing, they are instructed to guess. If a correct response is made, the question mark turns green, if an incorrect response is made, the question mark turns red. Feedback is included in the task because the literature documents patients having a "sense" of motion in the blind field, rather than a traditional percept. Positive feedback gives the subjects an indication of what they are looking for in terms of a correct response, making the distinction between no motion and a "sense" of motion easier to make. Following this, the fixation target returns in place of the question mark and the subject is instructed to press spacebar to move onto the next trial.

Subjects first complete 10 practice trials with the stimulus presented to the seeing field, followed by 40 test trials with the stimulus presented to the blind field. In both the practice and the test trials, the same paradigm is used (see Figure 102). Task performance is measured by calculating the per cent correct for each observer. Practice data are not analysed.

6.3.2.6 Task 3: Optokinetic nystagmus (OKN)

6.3.2.6.1 Visual stimulus

A gaze-contingent sinusoidal grating is presented to each subject in order to elicit an OKN response (Figure 103). The grating is confined to one of the subject's hemifield at a time using an eye-tracker. The visual field to which the grating was presented alternates between the seeing and blind hemifield. If the subject's gaze deviates by 1 degree or more, the stimulus position is updated to this new gaze position. The grating has a spatial frequency of 0.5 cycles per degree and a temporal frequency of 5 cycles per second. The opposite hemifield consists of a blank, grey background.

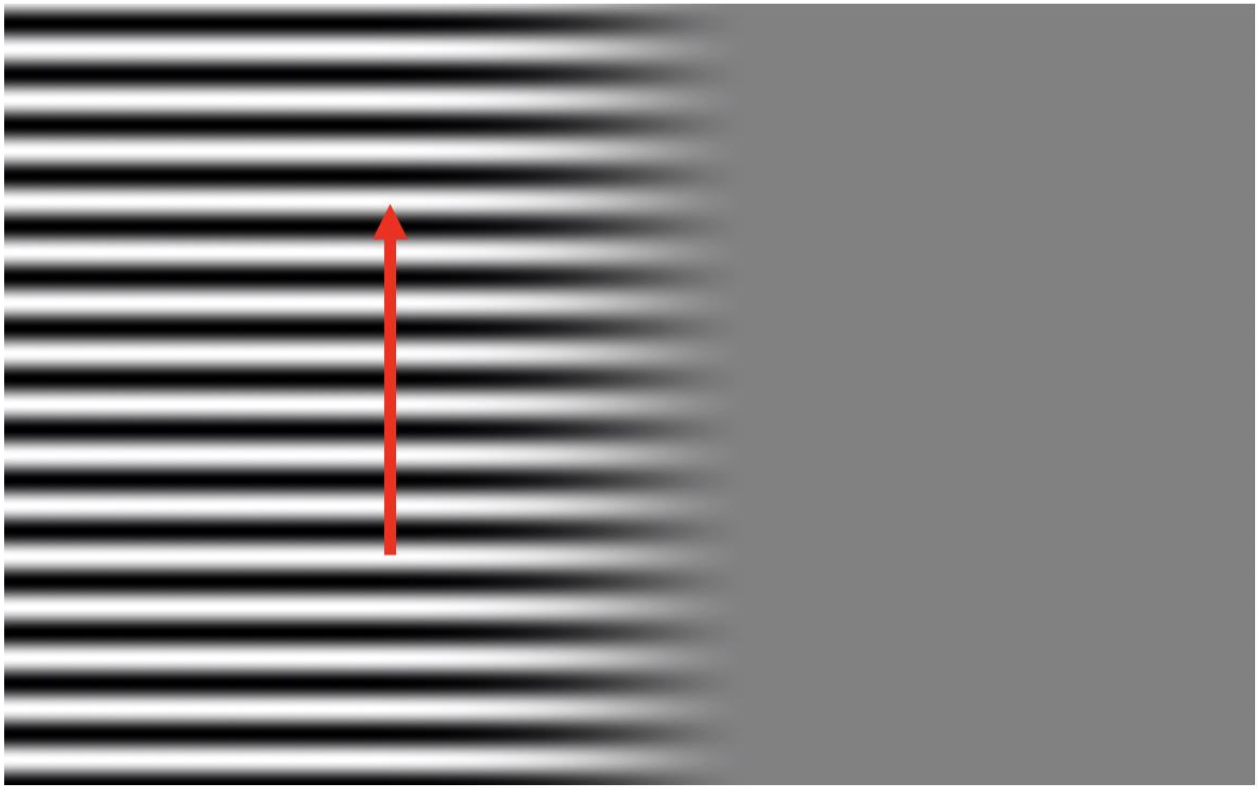

Figure 103: Stimulus used during OKN task, consisting of a hemifield sinusoidal grating.

6.3.2.6.2 Task procedure

Four stimulus direction (upwards, downwards, left and right) and two visual field location (blind or seeing) conditions are investigated in this task, resulting in eight task conditions. The order in which these conditions are presented is randomised for each subject. As this stimulus is novel to most observers, subjects are shown the stimulus moving in one of the four directions (selected at random) before the main experiment. Subjects are instructed that their task is to keep their eyes in the centre of the screen. Prior to the start of the task, a target is presented to indicate the centre of the screen. When the task begins, the target is removed. Subject's eye position is monitored throughout the task.

6.3.2.6.3 Data analysis

An automated segmentation algorithm (see Dunn et al., 2019 for detailed description) is utilised to quantify the waveforms extracted from the eye trace data collected during this task. In short, the eye trace is cleaned, with data associated with blinks and tracking artefacts removed. The waveform is then split into cycles, identifying peaks, slow phases, and quick phases (saccades). The median slow phase gain was then calculated for each condition. Slow phase gain is the ratio of the stimulus velocity to the

observed slow phase eye movement velocity: a gain of 1 equates to perfect synchronisation between the eyes and the stimulus, while a gain of 0 equates to no eye movements in the presence of the stimulus.

6.4 Pilot perimetry results

We acquired pilot visual field data from two male subjects (S1 and S2) previously diagnosed with homonymous hemianopia. Both subjects were recruited from the Cardiff University Neurological Vision Loss (NVL) Database. At the time of testing, S1 was 37 years old, S2 was 29 years old. The relevant information and results from the BIT are displayed in Table 19.

Table 19. Overview of subject information.

	Age	Diagnosed visual field defect	BIT Neglect Score
Subject 1 (S1)	37	L. homonymous hemianopia	Within normal limits
Subject 2 (S2)	29	L. homonymous hemianopia	Within normal limits

Subject 1 (S1)

The HFA characterised Subject 1 with a left homonymous hemianopia with a clearly defined boarder of vision down the vertical meridian (Figure 104). We note there was some light sensitivity measured in the left superior quadrant on the left eye and potentially in the inferior quadrant close to fixation. However, this was not clear macular sparing and potentially a result of eye movements into the blind field during the assessment.

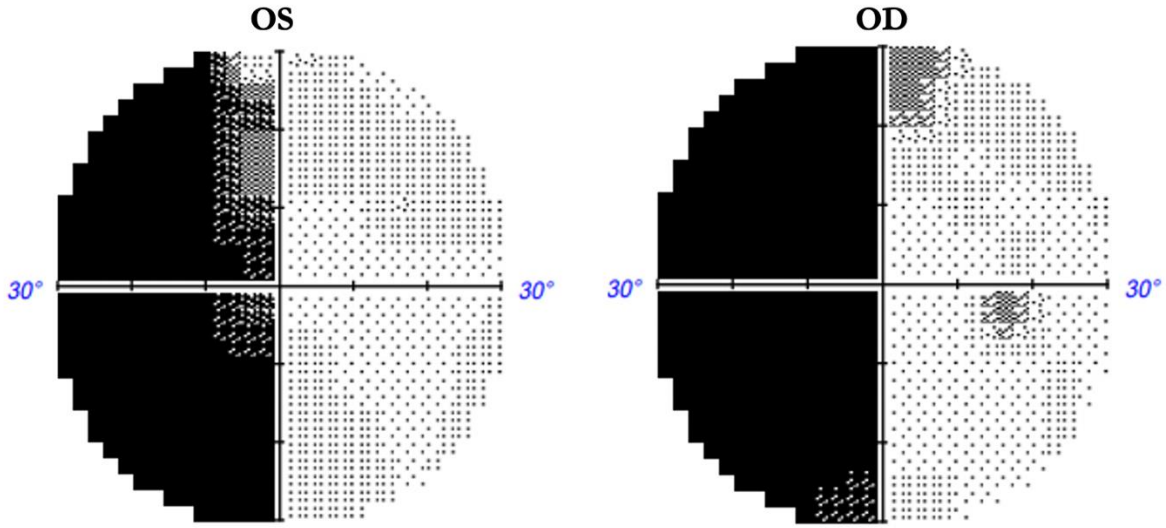

Figure 104: Humphrey visual fields for Subject 1. We note a left homonymous hemianopia without clear macular sparing. OS = left eye, OD = right eye.

Figure 105 shows the microperimetry data for the right eye of Subject 1. Note, the MAIA results are flipped vertically, thus the top of the figure represents the inferior visual field. Visually, the HFA and MAIA data are comparable showing a clear blind left hemifield. The small cluster of turquoise dots in the centre show Figure 105 (right) the fixation instability of the subject, although the MAIA attempts to counteract this instability, the system has a fairly low (25Hz) frame rate.

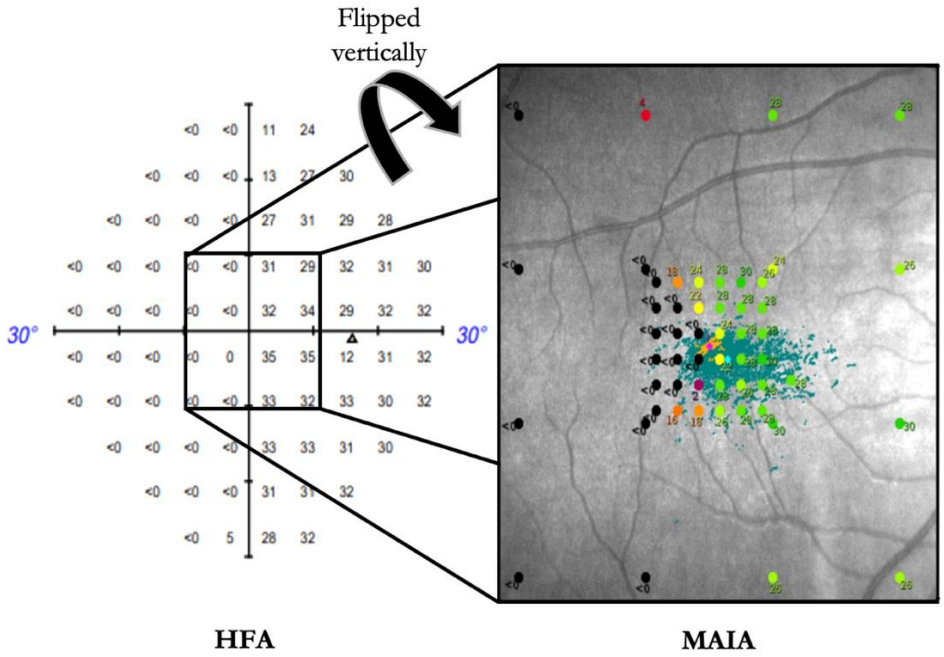

Figure 105: Comparison of thresholds acquired for Subject 1 using standard perimetry (HFA, left) and microperimetry (MAIA, right). The MAIA plot shows sensitivity thresholds superimposed onto the fundus image. Brighter colours show high threshold values, black plot points show a failure to respond to stimuli. The turquoise dots in the centre of the MAIA plot show the subject's fixation instability. Note, the MAIA plot is flipped vertically compared to the HFA plot.

Subject 2 (S2)

The results from the HFA for Subject 2 were somewhat unexpected. This subject had a long-standing diagnosis of left-sided homonymous hemianopia. However, when characterising the visual field defect with the HFA, the data appears to reveal a homonymous left-sided inferior quadrantanopia (Figure 106). This visual defect cleanly follows the vertical meridian in the inferior quadrant. However, the boundary between quadrants is not as clearly defined. Subject 2 shows good sensitivity in the most upper section of the quadrant. An interning note on Subject 2 was that he believed for many years he had hemianopia. However, on the day of testing he did explicitly state that on "good days" his vision in the upper left quadrant can be better than that in his seeing hemifield (right visual field). Subject 2 also stated that the vision in this left superior quadrant was not "normal" vision but was more of a visual "sense". Subject 2 further claimed that he could accurately sense the colour of a car that were situated in his superior left quadrant but could not "see" the car itself. To my knowledge, the day-to-day fluctuations in visual ability described by Subject 2 do not appear in existing literature

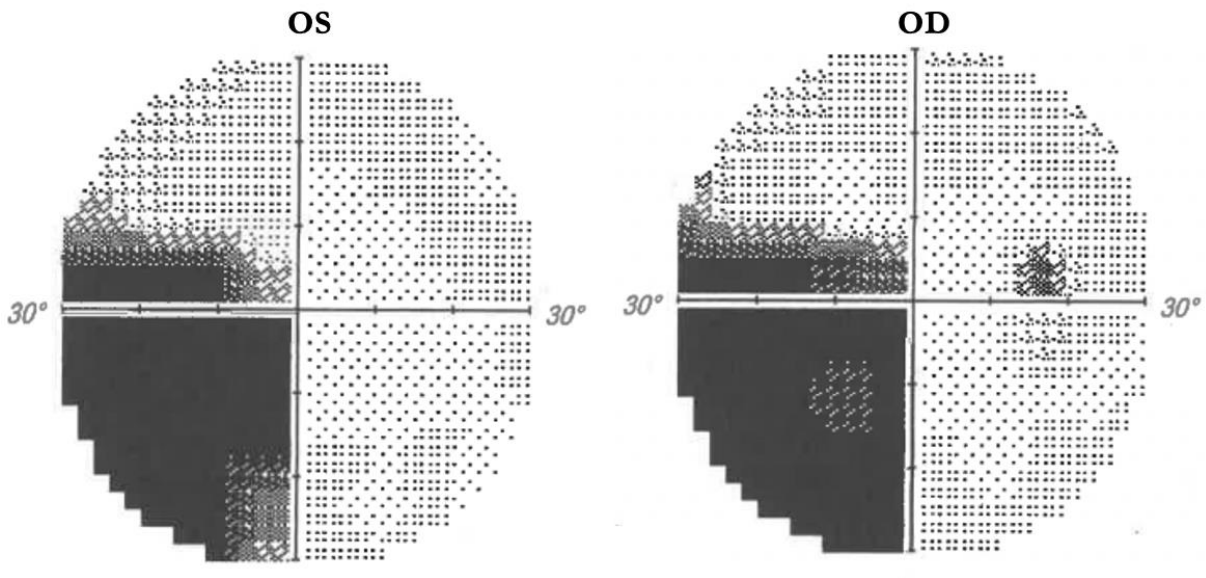

Figure 106: Humphrey visual fields for Subject 2. We note a left homonymous inferior quadrantanopia. OS = left eye, OD = right eye.

The HFA findings are somewhat replicated in the MAIA data, even though the majority of stimulus presentations were made to the central 2.25°. The microperimetry assessment (Figure 107) shows spared vision in the upper quadrant of the otherwise blind hemifield and fixation appears

relatively stable throughout the assessment (Figure 107). Although Subject 2 states he cannot necessarily "see" stimuli in his blind hemifield, his ability to respond to stimuli indicates a level of residual visual function. The extent of this spared vision is unknown.

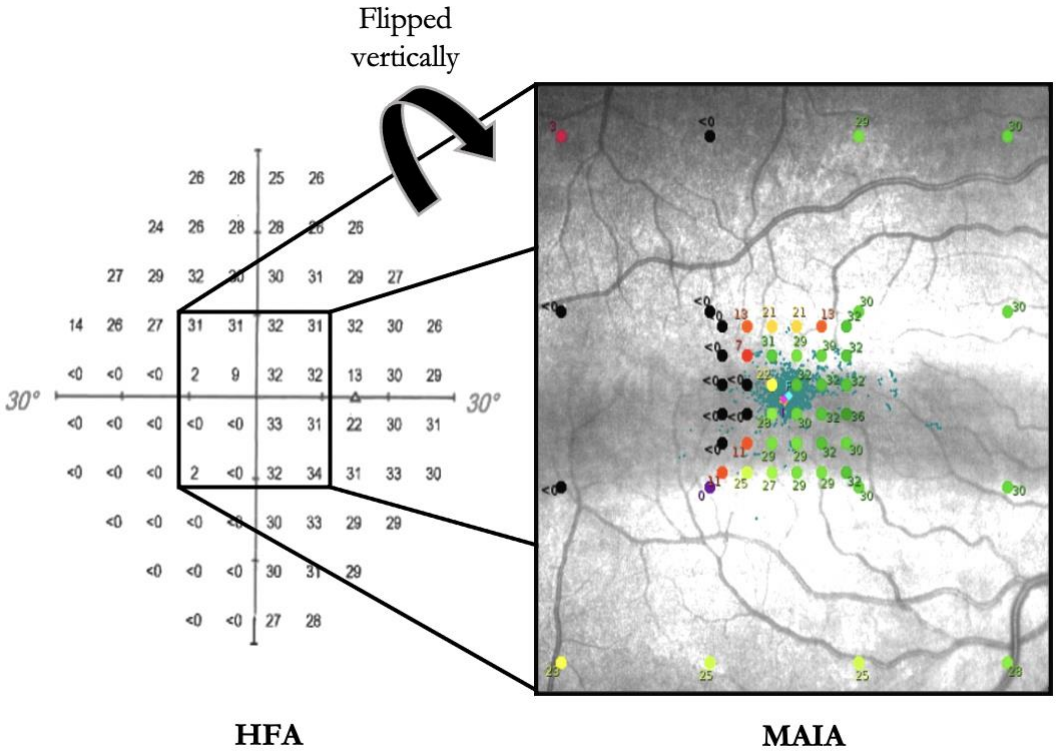

Figure 107: Microperimetry plot showing Subject 2's sensitivity map (A) and retina sensitivity thresholds superimposed onto the fundus image (B). Brighter colours show high threshold values, black plot points show a failure to respond to stimuli. Fixation stability is also displayed (C).

Chapter 7. Thesis summary and general discussion

This revised thesis sought to use rich, naturalistic stimuli (movie watching) to investigate:

1) whether ICA can be used to accurately identify visual areas of the brain (Chapter 3)
2) the potential lateralised function of TO1/MT and TO2/MST (Chapter 4)
3) the potential existence of a third visual pathway, distinct from the dorsal and ventral streams, that is specialised for processing visual motion (Chapter 5)

In addition, I outline the methods developed to address each of the PhD aims, including an optimal ICA pipeline for identifying visual areas (Chapter 2). In the final chapter I present a detailed psychophysical experiment designed to assess three different forms of residual visual function (Chapter 6).

7.1 Summary and interpretations

7.1.1 Chapter 3

Chapter 3 builds on the pioneering work of Bartels and Zeki who demonstrated a technique they called chronotopy. At the centre of their approach is the notion that different specialised regions have characteristic ATCs, i.e., they have a temporal fingerprint, which are preserved across observers who are exposed to the same visually rich stimulus. By performing ICA on fMRI data acquired while subjects freely viewed a movie, Bartels and Zeki found a small number of visually responsive regions can be identified.

In this thesis I adopted the chronotopy approach and explored whether the components derived from ICA correspond to visual areas. I performed spatial correlations between each subject's set of spatial maps and a probabilistic atlas of 25 visual areas (Wang et al., 2015).

First, I report there is not always a 1:1 mapping relationship between components and visual areas, a single component can match to many visual areas and many components can match to a single visual area. Next, I identified, for each brain, the top matched component (highest spatial correlation) for each of the 25 visual areas included in the atlas. I report that during natural viewing, a simple division of the brain into visual sub-regions does not hold up for all visual areas. Instead, I find a high

amount of variability across brains and across visual areas for each of the spatial analyses, with a key exception being TO1/MT and TO2/MST. In particular, I find the components derived from our ICA are often large, frequently encompassing adjacent regions. The results presented in this thesis may represent individual differences in functionality across the visual system and may add to a particularly exciting body of work that suggests the brain dynamically and temporally switches between different states (Meer et al., 2020).

Correlations within the visual input offer another possible reason for components straddling multiple visual areas. For example, consider when a face appears in the movie, not only is there a face but there are also edges that are redundantly coded by differences in luminance, colour, texture, depth, and motion. The co-occurrence of these features within the scene may elicit temporally synchronised activity, resulting in the areas being "grouped" into a single, large component.

Next, I performed a between-brains correlation analysis (analogous with Bartels and Zeki) to assess whether the ATCs of candidate components showed correspondence across subjects. This temporal analysis was central to Bartels and Zeki's approach as they claimed that functionally distinct regions have equally distinct ATCs. Therefore, if the same functional region is identified across brains, their timecourses should be correlated. We extracted timecourses from the components derived from ICA and from the atlas with which the component spatially matched and compared the results. We found, on average, the ATCs extracted from the atlas regions were more correlated across brains than the ATCs extracted from the components. We did, however, find the ATCs of hMT+ components to be significantly more consistent across brains than ATCs extracted from the atlas. TO1/MT and TO2/MST are not segmented into independent components, but rather a component corresponding to the entire motion complex is produced. Future research may utilise higher resolution acquisitions which might allow for distinct signals within the two regions to be segmented.

An important limitation of the temporal analysis is the inherent spatial variability of the top matched components across brains. Take V1v as an example. In some brains the top matched component is a cluster of activation that covers both V1v and V1d, in other brains, the component covers V1v and V2v, and in some instances the top matched component is discrete and does not

extend into any adjacent regions. This variability between brains means the location from which the mean ATC is extracted from can differ substantially across brains. Therefore, the correlations computed between these ATCs are not like for like.

Another key difference between the analysis performed here and the work of Bartels and Zeki is that the latter only examined bilateral components that showed anatomical correspondence across observers, at the cost of ignoring components that were unilateral or appeared to represent sub-divisions of the areas that were present in only a few brains. In contrast, we did not confine our spatial analysis to bilateral components, but instead used an atlas to identify candidate components in each hemisphere. This approach meant that we could identify both unilateral and bilateral components. Since the components are split into left and right hemisphere maps, and correlations are then computed within the same hemisphere, our analysis is sensitive to the potential differences across hemispheres and across brains. This has potential applications for case studies of patients with cortical vision loss. In particular the approach provides an opportunity to identify unilateral components in spared regions of the damaged hemisphere and bilateral components in areas upstream of the damage, for example in hMT+, which may indicate the involvement of an intact non-occipital pathway. These differences may be missed if the analysis only sought to identify bilateral components as Bartels and Zeki did.

The goal of chronotopy is to split the brain into specialised regions in an organic manner, in the hope that the processes serving our normal visual experience are replicated when free viewing naturalistic stimuli. The fact we do not find a single discrete component for each visual area, in each brain, is not necessarily indicative that the approach is ineffective. Rather, chronotopy may offer a more accurate way to segment the visual brain; although there are distinct areas, during natural vision they may act in concert.

7.1.2 Chapter 4

In Chapter 4, I investigated Strong et al.'s (2019) hypothesis that right hemisphere hMT+ has an "enhanced" role for processing visual motion in the contra- and ipsi-lateral visual fields, compared

to left hemisphere hMT+ which is responsible for processing motion in the contra-lateral visual field only. I explored this hypothesis by first performing a series of cross-hemisphere timecourse correlations.

The timecourse analyses were motivated by the premise that if a region in the left hemisphere has a different functional role to the right hemisphere, the correlation between the two timecourses should be lower than if a region has a similar functional role in both hemispheres. First, I performed an ROI-based analysis in which the mean timecourse of vertices within each ROI (defined from a probabilistic atlas) was extracted. Correlations were then performed between the mean timecourses of corresponding ROIs in the two hemispheres. Next, I performed a vertex-based analysis in which the correlations between the timecourses of corresponding vertices in the two hemispheres were calculated before taking the mean correlation across vertices within each probabilistic-atlas-defined ROI. The two analyses are conceptually different; the ROI-based analysis examines the correlation of the mean timecourses, the vertex-based analysis examines the mean of the timecourse correlation of vertices within each ROI.

I find, for both analyses, the timecourses of TO1/MT are as highly correlated across hemispheres as V1 (which I treat as a baseline of bilateral function). From this I infer the function of TO1/MT is no more lateralised than any other region. However, TO2/MST does not follow this trend. I find, on average, TO2/MST timecourses are less correlated across hemispheres than V1v, V1d and TO1/MT. However, as the correlations between hemispheres were still relatively high, this provides only weak evidence that TO2/MST is lateralised in each brain.

Although cross-hemisphere correlations of TO2/MST timecourses were still high, one possibility is that they were being driven down by individual differences. Therefore, to assess individual differences in the function of TO2/MST, I performed a between-brains timecourse analysis for each hemisphere separately. This analysis assessed the consistency of the timecourses across brains by computing correlations between each individual and the group mean timecourse.

First, I examined V1 and found timecourses were equally correlated across brains in both the left and right hemispheres, indicating there is little functional difference between hemispheres or

brains. Next, I examined hMT+. If right hemisphere hMT+ has the same functional role in each brain, e.g., processes both contra- and ipsi-lateral visual fields, we would have expected the timecourses to be more highly correlated than if there were individual differences in function. Moreover, if hMT+ had the same functional role across hemispheres we would not expect to find a significant difference between the correlations in each hemisphere. I find timecourses of TO1/MT and TO2/MST in the right hemisphere are significantly less correlated across brains than left hemisphere timecourses. This finding suggests right hemisphere hMT+, in particular TO2/MST, has a different functional role across brains. These timecourse correlations led to the tentative conclusion that there are individual differences in the function of TO2/MST across hemispheres.

A possible limitation of the timecourse analyses is that they all involved taking a mean across the entire ROI, potentially resulting in hemispheric differences in the sub regions of a visual area (e.g., MSTd and MSTl) being lost. To address this, I explored cross-hemisphere timecourse correlations on a per-vertex basis. I displayed each subject's cross-hemisphere correlation coefficients calculated in the vertex-based analysis on the surface of the brain. The subjects cross-hemisphere correlation of the baseline region, V1, was subtracted from each vertex correlation value to reveal areas of the brain that were more/less correlated than the baseline. These individual-subject vertex correlation maps not only allowed me to assess differences across hemispheres in the sub-divisions of MST, but they also allowed me to assess whether the probabilistic atlas ROIs were misplaced. For example, whether we find an area of highly correlated vertices slightly outside of the ROI.

The correlation maps indeed support the conclusions derived from previous analyses. I find individual differences in the cross-hemisphere correlations, of which I provide four examples in the chapter. In some brains the timecourses of both TO1/MT and TO2/MST were less corelated across hemispheres than V1, indicating the regions have different functional roles across hemispheres. For other brains I find the reverse. In some cases, one of TO1/TO2 was more highly correlated than the other. These findings added further weight to our initial conclusion that there are individual differences in the function of TO1/MT and TO2/MST.

Given the timecourse correlation analysis pointed to individual differences, I examined the functional lateralisation of hMT+ with a different approach. I performed single-subject ICA and inspected components that anatomically corresponded to TO1/MT and TO2/MST in each brain. We classify a subject who produced an unequal number of components including the left and right hemispheres as providing evidence of lateralised function, e.g., a bilateral component and an additional unilateral component. The logic for this clarification was as follows. Consider a case in which right hMT+ processes both the ipsi- and contra-lateral visual field. The bilateral component represents the commonality between hemispheres, i.e., left hMT+ processing the contra-lateral (right) visual hemifield and right hMT+ processing the ipsi-lateral (right) visual hemifield. In addition, we know right hMT+, by default, should have sensitivity to the contra-lateral (left) visual hemifield, therefore ICA will output a unilateral component in the right hemisphere that reflects right hMT+ processing the contra-lateral visual hemifield.

Based on this classification, we find evidence of lateralised function of hMT+ in ~25% of brains. The remaining subjects either produced a single bilateral component only (~65%) or a unilateral component in each hemisphere (~10%), which I infer as not providing evidence of lateralised function of hMT+.

In summary, the results from our three different analyses suggest that if Strong et al.'s hypothesis of the lateralised function of hMT+ is correct, it may not be a universal phenomenon. With each analysis, I arrived at the same conclusion – there are individual differences in the function of hMT+, in particular TO2/MST.

7.1.3 Chapter 5

In Chapter 5, I looked for evidence to support Gilaie-Dotan's (2016) hypothesis of a functionally distinct lateral motion pathway. Initially, I determined whether well-established resting state networks can be identified by performing group ICA on movie watching fMRI data. This analysis served two purposes. First, it is useful for us to know whether these networks reorganise when subjects watch a movie. I find a spatially corresponding component for each of the resting-state networks is

produced, with the exception of visual network 3. This suggests, for the most part, the networks do not reorganise.

Second, as three of the networks encompass different portions of the visual cortex, the analysis allowed me to explore their potential representation of the three visual pathways: ventral, dorsal, and lateral motion. At the group-level, components derived from ICA spatially correspond to visual network 1 (primary visual network) and visual network 2 (ventral pathway). Visual network 3 covers the lateral occipital regions, including V5/hMT+ and extending dorsally past V3A, and is considered to represent the dorsal pathway of visual processing. In our movie data, we find this network splits off into two separate components, one with activation centred on V5/hMT+ and another with activation centred on V3A. This splitting is in line with what would be expected if Gilaie-Dotan's (2016) hypothesis is correct, that the lateral motion pathway includes V5/hMT+ but not V3A. In contrast, we find no evidence of the network splitting when the analysis is performed on our resting state dataset.

An explanation as to why the network fractionates with movie watching but not rs-fMRI data (Smith et al., 2009) could indeed be that the components represent two distinct visual pathways (dorsal and lateral motion). However, the splitting may reflect motion content in the movie being dissociated with other temporal changes. As outlined previously, motion features are decorrelated with the content of the movie, i.e., because of the camera movement, motion occurs independently of what is happening in the scene. Therefore, the segmentation of V3A and hMT+ from the other visual areas represented in visual networks 1 and 2 may reflect this decorrelated activity.

Finally, I examined whether the group-level characteristic splitting of visual network 3 (dorsal pathway) was evident in individual subject data. I find over half of the subjects (52%) showed clear evidence of splitting of visual network 3 into dorsal and lateral regions.

7.2 Implications and future directions

Chapter 3 is a direct continuation of Bartels and Zeki's work, in which we validate and extend their chronotopy approach. Such a validation is useful given how valuable the approach would be for mapping the visual brain. Traditional functional brain mapping techniques can be problematic for individuals with vision loss because they almost always require the subject maintains fixation throughout. Chronotopy has no fixation requirements, the subject simply watches the movie as they would at home. Similarly, children or patients with cognitive deficits that impact their ability to understand or follow instructions (e.g., aphasia), may benefit from a paradigm that is engaging and familiar to them, such as freely viewing a movie. Consequently, compliance may be improved in these subjects.

In their seminal publications, Bartels and Zeki report a small number (10) of bilateral ICs that showed spatial, and temporal correspondence across brains. Based on their anatomical location, six ICs were believed to correspond to the visually responsive V1, V2v/V3v, V4, V5/MT+, LOp and LOi regions. Our aims were to (i) validate their approach by investigating whether components derived from our data correspond to Bartels and Zeki's ICs, or, for example, do we find V2v and V3v split into separate components? and (ii) extend it by assessing whether more than six visually responsive regions can be identified.

Unlike Bartels and Zeki who identified ICs manually, we computed spatial correlations between each output component and the full probability maps of 25 visual areas (Wang et al., 2015). We found that rather than components representing discrete visual areas, the brain generally divides into 'early', 'dorsal', 'ventral' and 'lateral' areas. However, as noted throughout this thesis, an atlas cannot account for all the anatomical variability across brains, and therefore conclusions drawn from our analysis are limited by the fact that we do not provide a ground truth for each area in each brain. The spatial correlation measure conveys whether the peak of a component spatially corresponds with the vertices that show a high probability of belonging to a given area. Therefore, in order to truly assess whether components derived from ICA correspond to discrete visual areas, future research might assess their spatial extent relative to retinotopically defined boundaries or ROIs defined by functional localisers.

The use of these traditional approaches means a "ground truth" could be established and the components derived from ICA could then be compared to this.

Future work should use the chronotopy approach to examine how the brains of patients with cortical vision loss segment when watching a movie. The approach provides an opportunity to identify individual differences in the frequency of unilateral and bilateral components. For example, do we find unilateral 'early' components and bilateral 'lateral' components in the patients? Such a finding could indicate the involvement of non-V1 pathways providing input to extrastriate areas. This in turn could inform our understanding of why some patients exhibit retained visual processing despite damage to geniculocortical pathway.

Chapter 4 explores whether we find evidence to support the hypothesis, most recently championed by Strong et al.'s (see also; Thakral & Slotnick, 2011), that hMT+ has a different functional role across hemispheres. Investigating whether hMT+ is lateralised not only develops our understanding of the brain's functional organisation in healthy subjects, but the techniques developed could also be used to examine the role lateralised function plays in facilitating residual vision in patients with cortical vision loss. Of particular relevance to us is that lateralised function of hMT+ could potentially explain why some patients exhibit intact motion processing despite damage to V1. If consistently identified across healthy subjects, it may inform clinicians about the relative incidence of residual vision in left vs right hemisphere cortical damage patients.

As mentioned above, we used a probabilistic atlas to define two hMT+ sub-regions; TO1/MT and TO2/MST, from which timecourses and mean correlations were extracted. Given the size and anatomical location of hMT+ shows variability both between subjects and between hemispheres of the same brain (Huang et al., 2019; Large et al., 2016; Wilms et al., 2005), individualised ROIs would provide a more accurate definition than an atlas. Therefore, future research could use a traditional TO1/MT and TO2/MST localiser (e.g., Huk et al. 2002) to identify the regions more accurately. The analyses presented in Chapter 4 could then be performed on the timecourses extracted from the localiser ROIs, rather than from an approximation as is the case with a probabilistic atlas. Alternatively, if it is not possible to perform a functional localiser, probabilistic ROIs could be restricted so they only

include vertices that contain the highest probability of belonging to that area. For example, Wang et al. (2009) reported the peak probability (i.e., the highest probability value in each area) ranged from 1 (100% of brains included this vertex) in early visual areas to ~0.4 in anterior parietal areas. Therefore, the peak probability value of each area could be identified, and ROIs then defined to only include vertices that contain that probability value.

In a separate analysis, ICA was performed on our movie data and components spatially corresponding to hMT+ were inspected. We interpreted subjects who produced a bilateral component, or a unilateral component in each hemisphere, as not providing evidence of lateralisation. The fact some subjects produce a single bilateral component while others produce two unilateral components may reflect differences in the input between the two groups of subjects. A unilateral component in each hemisphere could reflect differences in fixation/attention during the movie resulting in the two visual hemifields having uncorrelated motion stimulation. This could be a fixation to one visual hemifield resulting in inhomogeneous visual motion stimulation across the visual field. Whereas subjects who maintain fairly stable central fixation throughout the movie are more likely to have homogeneous motion stimulation across the entire visual field, resulting in a bilateral component. Given subjects free view the movie we have no control over which elements of the stimulus the subject attends to. However, when watching a movie, the observers eye movements do tend to be highly correlated (Davis et al., 2021; Hasson et al., 2008). Therefore, the suggestion that differences in attention and/or fixation play a role in the differences in the output of components are no more than speculation in this thesis.

Future research could explore the impact of fixation on the number of components in each hemisphere with a more controlled, albeit less naturalistic, stimulus. For example, an experiment in which a full field of moving dots is presented to the observer while they fixate a central target. The experiment could contain conditions in which motion components known to elicit activation in TO1/MT and TO2/MST are interleaved. For example, random motion, pure translation, expansion/contraction, and a complex flow pattern containing rotation, expansion, and contraction elements (Smith et al., 2006). A condition in which static dots are presented may also be included. If

hMT+ has a similar functional role across hemispheres, we would expect ICA to produce a bilateral component. If hMT+ has a different role across hemispheres, we would expect a bilateral component and an additional unilateral component in the right hemisphere. To test this further, the stimulus could be presented to each hemifield separately, and/or be confined to an aperture which is presented at varying eccentricities in each hemifield. If right hMT+ has enhanced coverage, then we would expect left hemifield stimulation to produce a unilateral component and right visual hemifield stimulation to produce a bilateral component. Furthermore, since this experiment would be controlled, stimulus timings would be available and therefore data could also be analysed with a traditional general linear model for comparison.

Chapter 5 focused on whether we find evidence to extend the traditional dual stream model of visual processing (Goodale & Milner, 1992; Ungerleider & Mishkin, 1982). There is evidence in the literature of anatomically discrete pathways involved in local and global motion processing (Bartels et al., 2008), as well as for social perception (Pitcher & Ungerleider, 2021). However, the pathway of primary interest in this thesis was the lateral motion pathway proposed by Gilaie-Dotan (2016).

Under the dual stream model of visual processing, the dorsal pathway is proposed to facilitate the visual control of actions and includes both V3A and hMT+ (Goodale & Milner, 1992). In contrast, Gilaie-Dotan (2016) hypothesised that V3A falls under the purview of the dorsal stream, while hMT+ is the central hub of the lateral motion pathway. We explored whether functional network components derived from performing ICA on movie watching support Gilaie-Dotan's hypothesis.

Following the precedent in resting state literature, fixed, low dimensionality ICA was performed. Next, components corresponding to well-established functional networks were identified. At the group level, our results demonstrate the functional networks can be identified with movies, providing supporting evidence for an alternative approach to traditional resting state paradigms in which subjects lay with their eyes closed or fixate a target for the duration of the scan. This alternative approach may in fact be advantageous from a compliance perspective; movies are an engaging stimulus and therefore may improve data quality (observers can stay still for longer if engaged; Raschle et al.,

2009) and data quantity (subjects may be more comfortable staying in the scanner for longer if watching a movie).

Although our movie-derived components correspond well with the functional networks, there was one important exception – the network representing the dorsal visual pathway (visual network 3). Our data is broadly in line with Gilaie-Dotan's predictions; we find hMT+ and V3A are found in separate components when subjects watch a movie. No such splitting is evident in our resting state data, but rather both areas are included in the component corresponding to the dorsal pathway.

The existence of the lateral motion pathway has particular relevance for informing our understanding of residual vision in patients with cortical vision loss. If, as Gilaie-Dotan predicts, V1, the pulvinar and the LGN swiftly propagate motion information in parallel to hMT+, then despite damage to V1, the spared subcortical relay structures could still deliver visual signals to hMT+, thus facilitating some form of visual perception. Given the approach reveals bilateral activation in healthy brains, it could be used to examine whether we find a dissociation between unilateral brain damage patients who display residual vision and those who do not. For example, Ajina and Bridge (2018) demonstrated that when presented with a traditional motion stimulus (aperture of static or moving dots), patients with preserved motion processing showed significant activity and intact sensitivity to speed in hMT+, while those without did not. Future research should investigate whether patients show evidence of the lateral motion pathway by combining movie watching and ICA.

At the individual level, we found over half of the subjects showed clear evidence of visual network 3 splitting into dorsal and lateral regions. The lack of splitting in all brains could be indicative of individual differences with respect to functionality, attention, or fixation. For example, some observers may focus primarily on the faces or objects of the scene, while others may visually explore the scene more freely. In the former, activation in motion areas may be reduced. Future research could benefit from the use of eye-tracking in the scanner to explore the potential impact of gaze/fixation on the splitting of visual network 3.

Our analysis shows that V3A and hMT+ split into separate components when subjects watch a movie, but not when they are at rest. Another avenue to explore with future research is whether the

splitting is evident with other types of visual stimuli. For example, one could use clips with different spatial processing elements compared to a Hollywood movie; a scene only clip that does not include human bodies or faces (e.g., a nature documentary) and the Inscapes stimulus (Vanderwal et al., 2015).

The scene only clip would not contain biological motion, and therefore could probe whether the organisation of the output components support Gilaie-Dotan's model in which hMT+ processes all visual motion, regardless of what/where the motion stimulus is (e.g., do V3A and hMT+ still split?), or whether they support Pitcher & Ungerleider's (2021) social perception pathway that is specifically concerned with propagating information about moving bodies and faces to the STS, via hMT+ (e.g., does the hMT+ component also include STS activation?).

Inscapes is a movie-like stimulus designed to be a middle ground between rest and movie watching. Inscapes are abstract shapes that gradually evolve from one scene to the next, meaning there are no cut-away transitions. Vanderwal et al. (2015) reported that relative to rest, Inscapes reduce head motion and increase wakefulness. However, movies outperformed both Inscapes and rest in improving these compliance measures.

In Chapter 5, we suggest that the splitting of V3A and hMT+ from the other visual areas may be a result of motion being decorrelated from other features; because of the camera movement, motion occurs independent of what is happening in the scene. Inscapes offer a way to investigate this because a fixed-camera based perspective was adopted, and therefore the Inscapes movies do not contain global flow or cut-scenes.

Chapter 6 describes an extensive psychophysical experiment designed to test residual visual function in patients with cortical vision loss. Due to unforeseen circumstances, it was not possible to test the psychophysical tasks with patients (see section 1.1). A future study could implement this experiment, in its entirety, to investigate the prevalence and co-occurrence of different forms of residual visual function in patients with cortical vision loss.

7.3 Final remarks

In summary, this thesis used movie watching fMRI data to investigate the function and organisation of the visual brain. Movies offer a more engaging and ecological method for conducting such investigations. The approach does not require prolonged fixation on a static target and can make the scanning experience more pleasant for subjects, making it advantageous from a developmental and clinical perspective. A theme that runs through all analyses in this thesis is that all brains are not the same. Individual differences are clear with each analysis I performed. These individual differences need to be recognised. My results complement a growing movement within the neuroscientific research community, known as precision fMRI (Gordon et al., 2017). Precision fMRI focuses on the examination of the idiosyncrasies in the functioning of the brain. Precision fMRI is beginning to characterise individual variations in network organisation and generate individualised maps using highly-sampled data (e.g., each subject provides 300 minutes of RS data; see Gordon et al., 2017), such as the publicly available Midnight Scan Club dataset. The in-depth study of brain function and organisation at the individual level is an important avenue to explore in future research endeavours.

It is important to note that if my PhD was not disrupted by the COVID-19 pandemic, much of the work presented in this thesis would have concentrated on patients with hemianopia. The methodological and analysis techniques described throughout this thesis were developed to compliment the original PhD aims in the hope that should access to patients become possible, each analysis could be performed on these patients. As a result, this thesis now contains a thorough analysis framework that could, in future, be applied to data acquired from patients with cortical vision loss.

References

Agtzidis, I., Meyhöfer, I., Dorr, M., & Lencer, R. (2020). Following Forrest Gump: Smooth pursuit related brain activation during free movie viewing. *NeuroImage*, *216*, 116491. https://doi.org/10.1016/j.neuroimage.2019.116491

Ajina, S., & Bridge, H. (2017). Blindsight and Unconscious Vision: What They Teach Us about the Human Visual System. In *Neuroscientist* (Vol. 23, Issue 5, pp. 529–541). https://doi.org/10.1177/1073858416673817

Ajina, S., & Bridge, H. (2018). Blindsight relies on a functional connection between hMT+ and the lateral geniculate nucleus, not the pulvinar. *PLoS Biology*, *16*(7), e2005769. https://doi.org/10.1371/journal.pbio.2005769

Ajina, S., Pestilli, F., Rokem, A., Kennard, C., & Bridge, H. (2015). Human blindsight is mediated by an intact geniculo-extrastriate pathway. *ELife*, *4*(OCTOBER2015). https://doi.org/10.7554/ELIFE.08935

Ajina, S., Rees, G., Kennard, C., & Bridge, H. (2015). Abnormal contrast responses in the extrastriate cortex of blindsight patients. *Journal of Neuroscience*, *35*(21), 8201–8213. https://doi.org/10.1523/JNEUROSCI.3075-14.2015

Albright, T. D. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. *Journal of Neurophysiology*. https://doi.org/10.1152/jn.1984.52.6.1106

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

Amunts, K., Mohlberg, H., Bludau, S., & Zilles, K. (2020). Julich-Brain: A 3D probabilistic atlas of the human brain's cytoarchitecture. *Science*, *369*(6506), 988–992. https://doi.org/10.1126/science.abb4588

Amunts, K., Weiss, P. H., Mohlberg, H., Pieperhoff, P., Eickhoff, S., Gurd, J. M., Marshall, J. C., Shah, N. J., Fink, G. R., & Zilles, K. (2004). Analysis of neural mechanisms underlying verbal fluency in cytoarchitectonically defined stereotaxic space - The roles of Brodmann areas 44 and 45. *NeuroImage*, *22*(1), 42–56. https://doi.org/10.1016/j.neuroimage.2003.12.031

Amunts, K., & Zilles, K. (2015). Architectonic Mapping of the Human Brain beyond Brodmann. In *Neuron* (Vol. 88, Issue 6, pp. 1086–1107). Cell Press. https://doi.org/10.1016/j.neuron.2015.12.001

Arcaro, M. J., McMains, S. A., Singer, B. D., & Kastner, S. (2009). Retinotopic organization of human ventral visual cortex. *Journal of Neuroscience*, *29*(34), 10638–10652. https://doi.org/10.1523/JNEUROSCI.2807-09.2009

Armstrong, R. A. (2014). When to use the Bonferroni correction. In *Ophthalmic & physiological optics : the journal of the British College of Ophthalmic Opticians (Optometrists)* (Vol. 34, Issue 5, pp. 502–508). John Wiley & Sons, Ltd. https://doi.org/10.1111/opo.12131

Asquith, P. (2018). *Using independent components analysis to identify visually driven regions and networks in the human brain, using data collected during movie watching*. http://orca.cf.ac.uk/id/eprint/120197

Azzopardi, P., & Cowey, A. (1998). Blindsight and Visual Awareness. *Consciousness and Cognition*, *7*(3), 292–311. https://doi.org/10.1006/ccog.1998.0358

Barbur, J. L., Watson, J. D. G., Frackowiak, R. S. J., & Zeki, S. (1993). Conscious visual perception without VI. *Brain*, *116*(6), 1293–1302. https://doi.org/10.1093/brain/116.6.1293

Bartels, A., & Zeki, S. (2004a). Functional Brain Mapping during Free Viewing of Natural Scenes. *Human Brain Mapping*, *21*(2), 75–85. https://doi.org/10.1002/hbm.10153

Bartels, A., & Zeki, S. (2004b). The chronoarchitecture of the human brain - Natural viewing conditions reveal a time-based anatomy of the brain. *NeuroImage*, *22*(1), 419–433. https://doi.org/10.1016/j.neuroimage.2004.01.007

Bartels, A., & Zeki, S. (2005a). Brain dynamics during natural viewing conditions - A new guide for mapping connectivity in vivo. *NeuroImage*, *24*(2), 339–349. https://doi.org/10.1016/j.neuroimage.2004.08.044

Bartels, A., & Zeki, S. (2005b). The chronoarchitecture of the cerebral cortex. In *Philosophical Transactions of the Royal Society B: Biological Sciences* (Vol. 360, Issue 1456, pp. 733–750). The Royal Society. https://doi.org/10.1098/rstb.2005.1627

Bartels, A., Zeki, S., & Logothetis, N. K. (2008). Natural vision reveals regional specialization to local motion and to contrast-invariant, global flow in the human brain. *Cerebral Cortex*, *18*(3), 705–717. https://doi.org/10.1093/cercor/bhm107

Becker, E., & Karnath, H. O. (2007). Incidence of visual extinction after left versus right hemisphere stroke. *Stroke*. https://doi.org/10.1161/STROKEAHA.107.489096

Beckmann, C. F., DeLuca, M., Devlin, J. T., & Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *360*(1457), 1001–1013. https://doi.org/10.1098/rstb.2005.1634

Beckmann, C. F., & Smith, S. M. (2004). Probabilistic Independent Component Analysis for Functional Magnetic Resonance Imaging. *IEEE Transactions on Medical Imaging*, *23*(2), 137–152. https://doi.org/10.1109/TMI.2003.822821

Benson, N. C., & Winawer, J. (2018). Bayesian analysis of retinotopic maps. *ELife*. https://doi.org/10.7554/eLife.40224

Benson, N. C., Yoon, J. M. D., Forenzo, D., Kay, K. N., Engel, S. A., & Winawer, J. (2021). Variability of the surface area of the V1, V2, and V3 maps in a large sample of human observers. In *bioRxiv*

(p. 2020.12.30.424856). Cold Spring Harbor Laboratory. https://doi.org/10.1101/2020.12.30.424856

Blazejewska, A. I., Fischl, B., Wald, L. L., & Polimeni, J. R. (2019). Intracortical smoothing of small-voxel fMRI data can provide increased detection power without spatial resolution losses compared to conventional large-voxel fMRI data. *NeuroImage*, *189*, 601–614. https://doi.org/10.1016/j.neuroimage.2019.01.054

Borja Jimenez, K. C., Abdelgabar, A. R., de Angelis, L., McKay, L. S., Keysers, C., & Gazzola, V. (2020). Changes in brain activity following the voluntary control of empathy. *NeuroImage*, *216*. https://doi.org/10.1016/j.neuroimage.2020.116529

Born, R. T., & Bradley, D. C. (2005). Structure and function of visual area MT. In *Annual Review of Neuroscience*. https://doi.org/10.1146/annurev.neuro.26.041002.131052

Boulinguez, P., Ferrois, M., & Graumer, G. (2003). Hemispheric asymmetry for trajectory perception. *Cognitive Brain Research*. https://doi.org/10.1016/S0926-6410(02)00276-8

Bouvier, S. E., & Engel, S. A. (2006). Behavioral deficits and cortical damage loci in cerebral achromatopsia. *Cerebral Cortex*, *16*(2), 183–191. https://doi.org/10.1093/cercor/bhi096

Bridge, H. (2011). Mapping the visual brain: How and why. *Eye*, *25*(3), 291–296. https://doi.org/10.1038/eye.2010.166

Bridge, H. (2020). Brain Basis of Blindsight. In *Oxford Research Encyclopedia of Psychology*. Oxford University Press. https://doi.org/10.1093/acrefore/9780190236557.013.733

Bridge, H., Hicks, S. L., Xie, J., Okell, T. W., Mannan, S., Alexander, I., Cowey, A., & Kennard, C. (2010). Visual activation of extra-striate cortex in the absence of V1 activation. *Neuropsychologia*. https://doi.org/10.1016/j.neuropsychologia.2010.10.022

Bridge, H., Thomas, O., Jbabdi, S., & Cowey, A. (2008). Changes in connectivity after visual cortical brain damage underlie altered visual function. *Brain*. https://doi.org/10.1093/brain/awn063

Brodmann, K. (1909). Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestelle auf Grund des Zellenbaues. In *The Journal of Nervous and Mental Disease* (Vol. 44, Issue 0). https://archive.org/details/b28062449/page/n4

Brodoehl, S., Gaser, C., Dahnke, R., Witte, O. W., & Klingner, C. M. (2020). Surface-based analysis increases the specificity of cortical activation patterns and connectivity results. *Scientific Reports*. https://doi.org/10.1038/s41598-020-62832-z

Bucsuházy, K., & Semela, M. (2017). Case Study: Reaction Time of Children According to Age. *Procedia Engineering*, *187*, 408–413. https://doi.org/10.1016/j.proeng.2017.04.393

Byrge, L., Dubois, J., Tyszka, J. M., Adolphs, R., & Kennedy, D. P. (2015). Idiosyncratic brain activation patterns are associated with poor social comprehension in autism. *Journal of Neuroscience*, *35*(14), 5837–5850. https://doi.org/10.1523/JNEUROSCI.5182-14.2015

Caballero-Gaudes, C., & Reynolds, R. C. (2017). Methods for cleaning the BOLD fMRI signal. *NeuroImage*, *154*, 128–149. https://doi.org/10.1016/j.neuroimage.2016.12.018

Cai, P., Chen, N., Zhou, T., Thompson, B., & Fang, F. (2014). Global versus local: double dissociation between MT+ and V3A in motion processing revealed using continuous theta burst transcranial magnetic stimulation. *Experimental Brain Research*, *232*(12), 4035–4041. https://doi.org/10.1007/s00221-014-4084-9

Calhoun, V. D., Adali, T., & Hansen, L. K. (2003). ICA of functional MRI data: an overview. *In Proceedings of the …*, *April*, 281–288. http://www.fmrib.ox.ac.uk/

Calhoun, V. D., Potluru, V. K., Phlypo, R., Silva, R. F., Pearlmutter, B. A., Caprihan, A., Plis, S. M., & Adali, T. (2013). Independent Component Analysis for Brain fMRI Does Indeed Select for Maximal Independence. *PLoS ONE*. https://doi.org/10.1371/journal.pone.0073309

Carp, J. (2012). The secret lives of experiments: Methods reporting in the fMRI literature. *NeuroImage*, *63*(1), 289–300. https://doi.org/10.1016/j.neuroimage.2012.07.004

Carrera, E., & Tononi, G. (2014). Diaschisis: Past, present, future. In *Brain* (Vol. 137, Issue 9, pp. 2408–2422). Oxford Academic. https://doi.org/10.1093/brain/awu101

Colle, F., Bonan, I., Gellez Leman, M. C., Bradai, N., & Yelnik, A. (2006). Fatigue after stroke. *Annales de Readaptation et de Medecine Physique*, *49*(6), 361–364. https://doi.org/10.1016/j.annrmp.2006.04.010

Cowey, A. (2010). The blindsight saga. In *Experimental Brain Research*. https://doi.org/10.1007/s00221-009-1914-2

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

Damoiseaux, J. S., Rombouts, S. A. R. B., Barkhof, F., Scheltens, P., Stam, C. J., Smith, S. M., & Beckmann, C. F. (2006). Consistent resting-state networks across healthy subjects. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(37), 13848–13853. https://doi.org/10.1073/pnas.0601417103

Davis, E. E., Chemnitz, E., Collins, T. K., Geerligs, L., & Campbell, K. L. (2021). Looking the same, but remembering differently: Preserved eye-movement synchrony with age during movie watching. *Psychology and Aging*, *36*(5), 604–615. https://doi.org/10.1037/pag0000615

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

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/10.1126/science.1063414

Dubner, R., & 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*, *35*(2), 528–532. https://doi.org/10.1016/0006-8993(71)90494-X

Dukelow, S. P., DeSouza, J. F. X., Culham, J. C., van den Berg, A. v., Menon, R. S., & Vilis, T. (2001). Distinguishing subregions of the human MT+ complex using visual fields and pursuit eye movements. *Journal of Neurophysiology*, *86*(4), 1991–2000. https://doi.org/10.1101/sqb.1990.055.01.069

Dumoulin, S. O., Bittar, R. G., Kabani, N. J., Baker, C. L., le Goualher, G., Pike, G. B., & Evans, A. C. (2000). A new anatomical landmark for reliable identification of human area V5/MT: A quantitative analysis of sulcal patterning. *Cerebral Cortex*, *10*(5), 454–463. https://doi.org/10.1093/cercor/10.5.454

Dunn, M. J., Harris, C. M., Ennis, F. A., Margrain, T. H., Woodhouse, J. M., McIlreavy, L., & Erichsen, J. T. (2019). An automated segmentation approach to calibrating infantile nystagmus waveforms. *Behavior Research Methods*. https://doi.org/10.3758/s13428-018-1178-5

Eickhoff, S. B., Milham, M., & Vanderwal, T. (2020). Towards clinical applications of movie fMRI. *NeuroImage*, *217*, 116860. https://doi.org/10.1016/j.neuroimage.2020.116860

Eifuku, S., & Wurtz, R. H. (1998). Response to motion in extrastriate area MSTI: Center-surround interactions. *Journal of Neurophysiology*, *80*(1), 282–296. https://doi.org/10.1152/jn.1998.80.1.282

Engel, S. A., Glover, G. H., & Wandell, B. A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cerebral Cortex*, *7*(2), 181–192. https://doi.org/10.1093/cercor/7.2.181

Engel, S. A., Rumelhart, D. E., Wandell, B. A., Lee, A. T., Glover, G. H., Chichilnisky, E. J., & Shadlen, M. N. (1994). fMRI of human visual cortex. *Nature*, *369*(6481), 525. https://doi.org/10.1038/369525A0

Esteban, O., Ciric, R., Finc, K., Blair, R. W., Markiewicz, C. J., Moodie, C. A., Kent, J. D., Goncalves, M., DuPre, E., Gomez, D. E. P., Ye, Z., Salo, T., Valabregue, R., Amlien, I. K., Liem, F., Jacoby, N., Stojić, H., Cieslak, M., Urchs, S., … Gorgolewski, K. J. (2020). Analysis of task-based functional MRI data preprocessed with fMRIPrep. *Nature Protocols*, *15*(7), 2186–2202. https://doi.org/10.1038/s41596-020-0327-3

Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., Kent, J. D., Goncalves, M., DuPre, E., Snyder, M., Oya, H., Ghosh, S. S., Wright, J., Durnez, J., Poldrack, R. A., & Gorgolewski, K. J. (2019). fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nature Methods*, *16*(1), 111–116. https://doi.org/10.1038/s41592-018-0235-4

Felleman, D. J., & van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*. https://doi.org/10.1093/cercor/1.1.1-a

Finn, E. S., & Bandettini, P. A. (2021). Movie-watching outperforms rest for functional connectivity-based prediction of behavior. *NeuroImage*, *235*, 117963. https://doi.org/10.1016/j.neuroimage.2021.117963

Finn, E. S., Glerean, E., Hasson, U., & Vanderwal, T. (2022). Naturalistic imaging: The use of ecologically valid conditions to study brain function. In *NeuroImage* (Vol. 247). Academic Press Inc. https://doi.org/10.1016/j.neuroimage.2021.118776

Finn, E. S., Glerean, E., Khojandi, A. Y., Nielson, D., Molfese, P. J., Handwerker, D. A., & Bandettini, P. A. (2020). Idiosynchrony: From shared responses to individual differences during naturalistic neuroimaging. *NeuroImage*, *215*, 116828. https://doi.org/10.1016/j.neuroimage.2020.116828

Finn, E. S., Scheinost, D., Finn, D. M., Shen, X., Papademetris, X., & Constable, R. T. (2017). Can brain state be manipulated to emphasize individual differences in functional connectivity? In *NeuroImage* (Vol. 160, pp. 140–151). Academic Press. https://doi.org/10.1016/j.neuroimage.2017.03.064

Fischl, B., Sereno, M. I., & Dale, A. M. (1999). Cortical Surface-Based Analysis. *NeuroImage*, *9*(2), 195–207. https://doi.org/10.1006/nimg.1998.0396

Fischl, B., van der Kouwe, A., Destrieux, C., Halgren, E., Ségonne, F., Salat, D. H., Busa, E., Seidman, L. J., Goldstein, J., Kennedy, D., Caviness, V., Makris, N., Rosen, B., & Dale, A. M. (2004). Automatically Parcellating the Human Cerebral Cortex. *Cerebral Cortex*, *14*(1), 11–22. https://doi.org/10.1093/cercor/bhg087

Fisher, R. A. (1915). Frequency Distribution of the Values of the Correlation Coefficient in Samples from an Indefinitely Large Population. *Biometrika*, *10*(4), 507. https://doi.org/10.2307/2331838

Frost, J. A., Binder, J. R., Springer, J. A., Hammeke, T. A., Bellgowan, P. S. F., Rao, S. M., & Cox, R. W. (1999). Language processing is strongly left lateralized in both sexes. Evidence from functional MRI. *Brain*, *122*(2), 199–208. https://doi.org/10.1093/brain/122.2.199

Gibson, J. J. (1950). Perception of the visual world. In *Boston: Houghton-Mifflin*. https://doi.org/10.1007/s10803-012-1669-7

Gilaie-Dotan, S. (2016). Visual motion serves but is not under the purview of the dorsal pathway. In *Neuropsychologia*. https://doi.org/10.1016/j.neuropsychologia.2016.07.018

Glasser, M. F., Smith, S. M., Marcus, D. S., Andersson, J. L. R., Auerbach, E. J., Behrens, T. E. J., Coalson, T. S., Harms, M. P., Jenkinson, M., Moeller, S., Robinson, E. C., Sotiropoulos, S. N., Xu, J., Yacoub, E., Ugurbil, K., & van Essen, D. C. (2016). The Human Connectome Project's neuroimaging approach. In *Nature Neuroscience* (Vol. 19, Issue 9, pp. 1175–1187). NIH Public Access. https://doi.org/10.1038/nn.4361

Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., van Essen, D. C., & Jenkinson, M. (2013). The minimal

preprocessing pipelines for the Human Connectome Project. *NeuroImage*, *80*, 105–124. https://doi.org/10.1016/j.neuroimage.2013.04.127

Gobbini, M. I., & Haxby, J. v. (2007). Neural systems for recognition of familiar faces. *Neuropsychologia*, *45*(1), 32–41. https://doi.org/10.1016/j.neuropsychologia.2006.04.015

Goodale, M. A., Meenan, J. P., Bülthoff, H. H., Nicolle, D. A., Murphy, K. J., & Racicot, C. I. (1994). Separate neural pathways for the visual analysis of object shape in perception and prehension. *Current Biology*. https://doi.org/10.1016/S0960-9822(00)00132-9

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*. https://doi.org/10.1016/0166-2236(92)90344-8

Goodale, M. A., Milner, A. D., Jakobson, L. S., & Carey, D. P. (1991). A neurological dissociation between perceiving objects and grasping them. *Nature*. https://doi.org/10.1038/349154a0

Gordon, E. M., Laumann, T. O., Gilmore, A. W., Newbold, D. J., Greene, D. J., Berg, J. J., Ortega, M., Hoyt-Drazen, C., Gratton, C., Sun, H., Hampton, J. M., Coalson, R. S., Nguyen, A. L., McDermott, K. B., Shimony, J. S., Snyder, A. Z., Schlaggar, B. L., Petersen, S. E., Nelson, S. M., & Dosenbach, N. U. (2017). Precision Functional Mapping of Individual Human Brains. *Neuron*, *95*(4), 791-807.e7. https://doi.org/10.1016/j.neuron.2017.07.011

Griffanti, L., Douaud, G., Bijsterbosch, J., Evangelisti, S., Alfaro-Almagro, F., Glasser, M. F., Duff, E. P., Fitzgibbon, S., Westphal, R., Carone, D., Beckmann, C. F., & Smith, S. M. (2017). Hand classification of fMRI ICA noise components. *NeuroImage*, *154*, 188–205. https://doi.org/10.1016/j.neuroimage.2016.12.036

Griffanti, L., Salimi-Khorshidi, G., Beckmann, C. F., Auerbach, E. J., Douaud, G., Sexton, C. E., Zsoldos, E., Ebmeier, K. P., Filippini, N., Mackay, C. E., Moeller, S., Xu, J., Yacoub, E., Baselli, G., Ugurbil, K., Miller, K. L., & Smith, S. M. (2014). ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *NeuroImage*, *95*, 232–247. https://doi.org/10.1016/j.neuroimage.2014.03.034

Groen, I. I. A., Dekker, T. M., Knapen, T., & Silson, E. H. (2022). Visuospatial coding as ubiquitous scaffolding for human cognition. In *Trends in Cognitive Sciences* (Vol. 26, Issue 1, pp. 81–96). Elsevier Ltd. https://doi.org/10.1016/j.tics.2021.10.011

Gruskin, D. C., Rosenberg, M. D., & Holmes, A. J. (2020). Relationships between depressive symptoms and brain responses during emotional movie viewing emerge in adolescence. *NeuroImage*, *216*, 116217. https://doi.org/10.1016/j.neuroimage.2019.116217

Guo, C. C., Nguyen, V. T., Hyett, M. P., Parker, G. B., & Breakspear, M. J. (2015). Out-of-sync: Disrupted neural activity in emotional circuitry during film viewing in melancholic depression. *Scientific Reports*, *5*(1), 1–12. https://doi.org/10.1038/srep11605

Hall, C. N., Howarth, C., Kurth-Nelson, Z., & Mishra, A. (2016). Interpreting BOLD: Towards a dialogue between cognitive and cellular neuroscience. In *Philosophical Transactions of the Royal Society B: Biological Sciences* (Vol. 371, Issue 1705). The Royal Society. https://doi.org/10.1098/rstb.2015.0348

Hasson, U., Landesman, O., Knappmeyer, B., Vallines, I., Rubin, N., & Heeger, D. J. (2008). Neurocinematics: The Neuroscience of Film. *Projections*, *2*(1), 1–26. https://doi.org/10.3167/proj.2008.020102

Hasson, U., Malach, R., & Heeger, D. J. (2010). Reliability of cortical activity during natural stimulation. In *Trends in Cognitive Sciences* (Vol. 14, Issue 1, pp. 40–48). Elsevier Current Trends. https://doi.org/10.1016/j.tics.2009.10.011

Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., & Malach, R. (2004). Intersubject Synchronization of Cortical Activity during Natural Vision. *Science*, *303*(5664), 1634–1640. https://doi.org/10.1126/science.1089506

Haxby, J. v., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. In *Trends in Cognitive Sciences* (Vol. 4, Issue 6, pp. 223–233). https://doi.org/10.1016/S1364-6613(00)01482-0

Heilman, K. M., Valenstein, E., & Watson, R. T. (2000). Neglect and related disorders. In *Seminars in Neurology* (Vol. 20, Issue 4, pp. 463–470). Semin Neurol. https://doi.org/10.1055/s-2000-13179

Heilman, K. M., & van den Abell, T. (1980). Right hemisphere dominance for attention: The mechanism underlying hemispheric asymmetries of inattention (neglect). *Neurology*. https://doi.org/10.1212/wnl.30.3.327

Henschen, S. E. (1893). On the visual path and centre. *Brain*, *16*(1–2), 170–180. https://doi.org/10.1093/brain/16.1-2.170

Holmes, G. (1918). Disturbances of vision by cerebral lesions. *The British Journal of Ophthalmology*, *2*(7), 353–384. https://doi.org/10.1136/bjo.2.7.353

Horton, J. C., Economides, J. R., & Adams, D. L. (2021). The Mechanism of Macular Sparing. In *Annual Review of Vision Science* (Vol. 7, pp. 155–179). https://doi.org/10.1146/annurev-vision-100119-125406

Horton, J. C., & Hoyt, W. F. (1991). The representation of the visual field in human striate cortex. A revision of the classic Holmes map. *Archives of Ophthalmology*, *109*(6), 816–824. https://doi.org/10.1001/archopht.1991.01080060080030

Huang, T., Chen, X., Jiang, J., Zhen, Z., & Liu, J. (2019). A probabilistic atlas of the human motion complex built from large-scale functional localizer data. *Human Brain Mapping*, *40*(12), 3475–3487. https://doi.org/10.1002/hbm.24610

Hudson, M., Seppälä, K., Putkinen, V., Sun, L., Glerean, E., Karjalainen, T., Karlsson, H. K., Hirvonen, J., & Nummenmaa, L. (2020). Dissociable neural systems for unconditioned acute and sustained fear. *NeuroImage*, *216*, 116522. https://doi.org/10.1016/j.neuroimage.2020.116522

Hudspeth, A., & Logothetis, N. K. (2000). Sensory systems. In *Current Opinion in Neurobiology* (Vol. 10, Issue 5, pp. 631–641). Elsevier Current Trends. https://doi.org/10.1016/S0959-4388(00)00133-1

Huk, A. C., Dougherty, R. F., & Heeger, D. J. (2002). Retinotopy and functional subdivision of human areas MT and MST. *Journal of Neuroscience*, *22*(16), 7195–7205. https://doi.org/10.1523/jneurosci.22-16-07195.2002

Hutson, J. P., Smith, T. J., Magliano, J. P., & Loschky, L. C. (2017). What is the role of the film viewer? The effects of narrative comprehension and viewing task on gaze control in film. *Cognitive Research: Principles and Implications*, *2*(1). https://doi.org/10.1186/s41235-017-0080-5

Hyon, R., Kleinbaum, A. M., & Parkinson, C. (2020). Social network proximity predicts similar trajectories of psychological states: Evidence from multi-voxel spatiotemporal dynamics. *NeuroImage*, *216*, 116492. https://doi.org/10.1016/j.neuroimage.2019.116492

Inouye, T. (1909). Die Sehstörungen bei Schussverletzungen der kortikalen Sehsphäre: nach Beobachtungen an Verwundeten der letzten japanischen Kriege. In *Engelmann*. W. Engelmann.

Jenkins, R., Dowsett, A. J., & Burton, A. M. (2018). How many faces do people know? *Proceedings of the Royal Society B: Biological Sciences*, *285*(1888). https://doi.org/10.1098/rspb.2018.1319

Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*. https://doi.org/10.1016/S1053-8119(02)91132-8

Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., & Smith, S. M. (2012). FSL. *NeuroImage*, *62*(2), 782–790. https://doi.org/10.1016/J.NEUROIMAGE.2011.09.015

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, *17*(11), 4302–4311. https://doi.org/10.1523/jneurosci.17-11-04302.1997

Kanwisher, N., & Yovel, G. (2006). The fusiform face area: A cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *361*(1476), 2109–2128. https://doi.org/10.1098/rstb.2006.1934

Karnath, H. O., & Rorden, C. (2012). The anatomy of spatial neglect. In *Neuropsychologia* (Vol. 50, Issue 6, pp. 1010–1017). NIH Public Access. https://doi.org/10.1016/j.neuropsychologia.2011.06.027

Kimmig, H., Ohlendorf, S., Speck, O., Sprenger, A., Rutschmann, R. M., Haller, S., & Greenlee, M. W. (2008). fMRI evidence for sensorimotor transformations in human cortex during smooth pursuit eye movements. *Neuropsychologia*, *46*(8), 2203–2213. https://doi.org/10.1016/j.neuropsychologia.2008.02.021

Kinreich, S., Intrator, N., & Hendler, T. (2011). Functional Cliques in the Amygdala and Related Brain Networks Driven by Fear Assessment Acquired During Movie Viewing. *Brain Connectivity*, *1*(6), 484–495. https://doi.org/10.1089/brain.2011.0061

Kooistra, C. A., & Heilman, K. M. (1989). Hemispatial visual inattention masquerading as hemianopia. *Neurology*, *39*(8), 1125–1127. https://doi.org/10.1212/wnl.39.8.1125

Kravitz, D. J., Saleem, K. S., Baker, C. I., & Mishkin, M. (2011). A new neural framework for visuospatial processing. In *Nature Reviews Neuroscience* (Vol. 12, Issue 4, pp. 217–230). Nature Publishing Group. https://doi.org/10.1038/nrn3008

Lagae, L., Maes, H., Raiguel, S., Xiao, D. K., & Orban, G. A. (1994). Responses of macaque STS neurons to optic flow components: A comparison of areas MT and MST. *Journal of Neurophysiology*. https://doi.org/10.1152/jn.1994.71.5.1597

Land, M., Mennie, N., & Rusted, J. (1999). The roles of vision and eye movements in the control of activities of daily living. *Perception*, *28*(11), 1311–1328. https://doi.org/10.1068/p2935

Large, I., Bridge, H., Ahmed, B., Clare, S., Kolasinski, J., Lam, W. W., Miller, K. L., Dyrby, T. B., Parker, A. J., Smith, J. E. T., Daubney, G., Sallet, J., Bell, A. H., & Krug, K. (2016). Individual Differences in the Alignment of Structural and Functional Markers of the V5/MT Complex in Primates. *Cerebral Cortex*, *26*(10), 3928–3944. https://doi.org/10.1093/cercor/bhw180

Lecrux, C., & Hamel, E. (2016). Neuronal networks and mediators of cortical neurovascular coupling responses in normal and altered brain states. In *Philosophical Transactions of the Royal Society B: Biological Sciences* (Vol. 371, Issue 1705). The Royal Society. https://doi.org/10.1098/rstb.2015.0350

Lee, D. N., & Aronson, E. (1974). Visual proprioceptive control of standing in human infants. *Perception & Psychophysics*, *15*(3), 529–532. https://doi.org/10.3758/BF03199297

Lee, S., & Lee, D. K. (2018). What is the proper way to apply the multiple comparison test? *Korean Journal of Anesthesiology*, *71*(5), 353–360. https://doi.org/10.4097/kja.d.18.00242

Leff, A. (2004). A historical review of the representation of the visual field in primary visual cortex with special reference to the neural mechanisms underlying macular sparing. *Brain and Language*, *88*(3), 268–278. https://doi.org/10.1016/S0093-934X(03)00161-5

Leigh, R. J., & Zee, D. S. (2015). *The neurology of eye movements*. Contemporary Neurology.

Leopold, D. A. (2012). Primary visual cortex: Awareness and blindsight. In *Annual Review of Neuroscience* (Vol. 35, pp. 91–109). NIH Public Access. https://doi.org/10.1146/annurev-neuro-062111-150356

Lissauer, H. (1890). Ein Fall von Seelenblindheit nebst einem Beitrage zur Theorie derselben. *Archiv Für Psychiatrie Und Nervenkrankheiten*. https://doi.org/10.1007/BF02226765

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

Malikovic, A., Amunts, K., Schleicher, A., Mohlberg, H., Eickhoff, S. B., Wilms, M., Palomero-Gallagher, N., Armstrong, E., & Zilles, K. (2007). Cytoarchitectonic analysis of the human extrastriate cortex in the region of V5/MT+: A probabilistic, stereotaxic map of area hOc5. *Cerebral Cortex*, *17*(3), 562–574. https://doi.org/10.1093/cercor/bhj181

Mäntylä, T., Nummenmaa, L., Rikandi, E., Lindgren, M., Kieseppä, T., Hari, R., Suvisaari, J., & Raij, T. T. (2018). Aberrant Cortical Integration in First-Episode Psychosis During Natural Audiovisual Processing. *Biological Psychiatry*, *84*(9), 655–664. https://doi.org/10.1016/j.biopsych.2018.04.014

McCarthy, G., Puce, A., Gore, J. C., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. *Journal of Cognitive Neuroscience*. https://doi.org/10.1162/jocn.1997.9.5.605

McFadzean, R., Brosnahan, D., Hadley, D., & Mutlukan, E. (1994). Representation of the visual-field in the occipital striate cortex. *British Journal of Ophthalmology*, *78*(3), 185–190. https://doi.org/10.1136/bjo.78.3.185

Mckeown, M. J., Makeig, S., Brown, G. G., Jung, T.-P., Kindermann, S. S., Bell, A. J., Sejnowski, T. J., & Mckeown, M. J. (1998). Analysis of fMRI Data by Blind Separation Into Independent Spatial Components. *Hum. Brain Mapping*, *6*, 160–188. https://doi.org/10.1002/(SICI)1097-0193(1998)6:3

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, *27*, 415–444. https://doi.org/10.1146/annurev.soc.27.1.415

Meer, J. N. van der, Breakspear, M., Chang, L. J., Sonkusare, S., & Cocchi, L. (2020). Movie viewing elicits rich and reliable brain state dynamics. *Nature Communications*, *11*(1), 1–14. https://doi.org/10.1038/s41467-020-18717-w

Menon, A., & Korner-Bitensky, N. (2004). Evaluating unilateral spatial neglect post stroke: Working your way through the maze of assessment choices. *Topics in Stroke Rehabilitation*, *11*(3), 41–66. https://doi.org/10.1310/KQWL-3HQL-4KNM-5F4U

Mesulam, M. -Marchsel. (1981). A cortical network for directed attention and unilateral neglect. *Annals of Neurology*, *10*(4), 309–325. https://doi.org/10.1002/ana.410100402

Metz, H. S. (1983). The Neurology of Eye Movements. In *Journal of Pediatric Ophthalmology & Strabismus* (Fourth Edi, Vol. 20, Issue 4). New York: Oxford University Press. https://doi.org/10.3928/0191-3913-19830701-12

Milner, A. D., Perrett, D. I., Johnston, R. S., Benson, P. J., Jordan, T. R., Heeley, D. W., Bettucci, D., Mortara, F., Mutani, R., Terazzi, E., & Davidson, D. L. W. (1991). Perception and action in "visual form agnosia." *Brain*. https://doi.org/10.1093/brain/114.1.405

Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: two cortical pathways. In *Trends in Neurosciences*. https://doi.org/10.1016/0166-2236(83)90190-X

Morrone, M. C., Tosetti, M., Montanaro, D., Fiorentini, A., Cioni, G., & Burr, D. C. (2000). A cortical area that responds specifically to optic flow, revealed by fMRI. *Nature Neuroscience*, *3*(12), 1322–1328. https://doi.org/10.1038/81860

Movshon, J. A., & Newsome, W. T. (1996). Visual response properties of striate cortical neurons projecting to area MT in macaque monkeys. *Journal of Neuroscience*. https://doi.org/10.1523/jneurosci.16-23-07733.1996

Naci, L., Sinai, L., & Owen, A. M. (2017). Detecting and interpreting conscious experiences in behaviorally non-responsive patients. *NeuroImage*, *145*, 304–313. https://doi.org/10.1016/j.neuroimage.2015.11.059

Nastase, S. A., Goldstein, A., & Hasson, U. (2020). Keep it real: rethinking the primacy of experimental control in cognitive neuroscience. *NeuroImage*, *222*, 117254. https://doi.org/10.1016/j.neuroimage.2020.117254

Nieuwenhuis, S., Forstmann, B. U., & Wagenmakers, E. J. (2011). Erroneous analyses of interactions in neuroscience: A problem of significance. In *Nature Neuroscience* (Vol. 14, Issue 9, pp. 1105–1107). Nature Publishing Group. https://doi.org/10.1038/nn.2886

Pambakian, A. L. M., Mannan, S. K., Hodgson, T. L., & Kennard, C. (2004). Saccadic visual search training: A treatment for patients with homonymous hemianopia. *Journal of Neurology, Neurosurgery and Psychiatry*, *75*(10), 1443–1448. https://doi.org/10.1136/jnnp.2003.025957

Pamilo, S., Malinen, S., Hlushchuk, Y., Seppä, M., Tikka, P., & Hari, R. (2012). Functional subdivision of group-ICA results of fMRI data collected during Cinema viewing. *PLoS ONE*, *7*(7), e42000. https://doi.org/10.1371/journal.pone.0042000

Parkinson, C., Kleinbaum, A. M., & Wheatley, T. (2018). Similar neural responses predict friendship. *Nature Communications*, *9*(1), 1–14. https://doi.org/10.1038/s41467-017-02722-7

Perenin, M. T., & Vighetto, A. (1983). Optic Ataxia: A Specific Disorder in Visuomotor Coordination. In *Spatially Oriented Behavior*. https://doi.org/10.1007/978-1-4612-5488-1_17

Perrone, J. A., & Thiele, A. (2001). Speed skills: Measuring the visual speed analyzing properties of primate MT neurons. *Nature Neuroscience*. https://doi.org/10.1038/87480

Pisella, L., Gréa, H., Tilikete, C., Vighetto, A., Desmurget, M., Rode, G., Boisson, D., & Rossetti, Y. (2000). An "automatic pilot" for the hand in human posterior parietal cortex: Toward reinterpreting optic ataxia. *Nature Neuroscience*. https://doi.org/10.1038/76694

Pitcher, D., & Ungerleider, L. G. (2021). Evidence for a Third Visual Pathway Specialized for Social Perception. In *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2020.11.006

Poppe, A. B., Wisner, K., Atluri, G., Lim, K. O., Kumar, V., & MacDonald, A. W. (2013). Toward a neurometric foundation for probabilistic independent component analysis of fMRI data. *Cognitive, Affective and Behavioral Neuroscience*, *13*(3), 641–659. https://doi.org/10.3758/s13415-013-0180-8

Pöppel, E., Held, R., & Frost, D. (1973). Residual visual function after brain wounds involving the central visual pathways in man. *Nature*, *243*(5405), 295–296. https://doi.org/10.1038/243295a0

Priebe, N. J., Cassanello, C. R., & Lisberger, S. G. (2003). The neural representation of speed in macaque area MT/V5. *Journal of Neuroscience*. https://doi.org/10.1523/jneurosci.23-13-05650.2003

Priebe, N. J., Lisberger, S. G., & Movshon, J. A. (2006). Tuning for spatiotemporal frequency and speed in directionally selective neurons of macaque striate cortex. *Journal of Neuroscience*. https://doi.org/10.1523/JNEUROSCI.3936-05.2006

Pruim, R. H. R., Mennes, M., van Rooij, D., Llera, A., Buitelaar, J. K., & Beckmann, C. F. (2015). ICA-AROMA: A robust ICA-based strategy for removing motion artifacts from fMRI data. *NeuroImage*, *112*, 267–277. https://doi.org/10.1016/j.neuroimage.2015.02.064

Purves, D., Augustine, G., Fitzpatrick, D., Katz, L., LaMantia, A.-S., McNamara, J., & Williams, M. (2001). *Neuroscience. 2nd edition*. Sunderland (MA): Sinauer Associates; 2001. https://doi.org/978-0878937257

Purves, D., Augustine, G. J., Fitzpatrick, D., Katz, L. C., LaMantia, A.-S., McNamara, J. O., & Williams., S. M. (2001). Neuroscience, 2nd Edition. In *Sinauer Associates, Inc.* https://doi.org/978-0878937257

Raiguel, S., van Hulle, M. M., Xiao, D. K., Marcar, V. L., Lagae, L., & Orban, G. A. (1997). Size and shape of receptive fields in the medial superior temporal area (MST) of the macaque. *NeuroReport*. https://doi.org/10.1097/00001756-199708180-00030

Raschle, N. M., Lee, M., Buechler, R., Christodoulou, J. A., Chang, M., Vakil, M., Stering, P. L., & Gaab, N. (2009). Making MR imaging child's play - Pediatric neuroimaging protocol, guidelines and procedure. *Journal of Visualized Experiments*, *29*, 29. https://doi.org/10.3791/1309

Reinhard, J. I., Damm, I., Ivanov, I. v., & Trauzettel-Klosinski, S. (2014). Eye movements during saccadic and fixation tasks in patients with homonymous hemianopia. *Journal of Neuro-Ophthalmology*, *34*(4), 354–361. https://doi.org/10.1097/WNO.0000000000000146

Riddoch, G. (1917). Dissociation of visual perceptions due to occipital injuries, with especial reference to appreciation of movement. *Brain*, *40*(1), 15–57. https://doi.org/10.1093/brain/40.1.15

Rikandi, E., Mäntylä, T., Lindgren, M., Kieseppä, T., Suvisaari, J., & Raij, T. T. (2022). Functional network connectivity and topology during naturalistic stimulus is altered in first-episode psychosis. *Schizophrenia Research*, *241*, 83–91. https://doi.org/10.1016/j.schres.2022.01.006

Ringman, J. M., Saver, J. L., Woolson, R. F., Clarke, W. R., & Adams, H. P. (2004). Frequency risk factors, anatomy, and course of unilateral neglect in an acute stroke cohort. *Neurology*. https://doi.org/10.1212/01.WNL.0000133011.10689.CE

Rizzo, M., Nawrot, M., & Zihl, J. (1995). Motion and shape perception in cerebral akinetopsia. *Brain*, *118*(5), 1105–1127. https://doi.org/10.1093/brain/118.5.1105

Rodman, H. R., Gross, C. G., & Albrigh, T. D. (1990). Afferent basis of visual response properties in area MT of the macaque. II. Effects of superior colliculus removal. *Journal of Neuroscience*, *10*(4), 1154–1164. https://doi.org/10.1523/jneurosci.10-04-01154.1990

Rodman, H. R., Gross, C. G., & Albright, T. D. (1989). Afferent basis of visual response properties in area MT of the macaque. I. Effects of striate cortex removal. *Journal of Neuroscience*, *9*(6), 2033–2050. https://doi.org/10.1523/jneurosci.09-06-02033.1989

Rorden, C., & Karnath, H. O. (2004). Using human brain lesions to infer function: A relic from a past era in the fMRI age? *Nature Reviews Neuroscience*, *5*(10), 812–819. https://doi.org/10.1038/nrn1521

Rosenke, M., van Hoof, R., van den Hurk, J., Grill-Spector, K., & Goebel, R. (2021). A Probabilistic Functional Atlas of Human Occipito-Temporal Visual Cortex. *Cerebral Cortex*, *31*(1), 603–619. https://doi.org/10.1093/cercor/bhaa246

Rushton, S. K., Niehorster, D. C., Warren, P. A., & Li, L. (2018). The primary role of flow processing in the identification of scene-relative object movement. *Journal of Neuroscience*, *38*(7), 1737–1743. https://doi.org/10.1523/JNEUROSCI.3530-16.2017

Rushton, S. K., & Warren, P. A. (2005). Moving observers, relative retinal motion and the detection of object movement [2]. In *Current Biology* (Vol. 15, Issue 14, pp. R542–R543). Cell Press. https://doi.org/10.1016/j.cub.2005.07.020

Salmi, J., Metwaly, M., Tohka, J., Alho, K., Leppämäki, S., Tani, P., Koski, A., Vanderwal, T., & Laine, M. (2020). ADHD desynchronizes brain activity during watching a distracted multi-talker conversation. *NeuroImage*, *216*, 116352. https://doi.org/10.1016/j.neuroimage.2019.116352

Salmi, J., Roine, U., Glerean, E., Lahnakoski, J., Nieminen-Von Wendt, T., Tani, P., Leppämäki, S., Nummenmaa, L., Jääskeläinen, I. P., Carlson, S., Rintahaka, P., & Sams, M. (2013). The brains of high functioning autistic individuals do not synchronize with those of others. *NeuroImage: Clinical*, *3*, 489–497. https://doi.org/10.1016/j.nicl.2013.10.011

Savoy, R. L. (2001). History and future directions of human brain mapping and functional neuroimaging. *Acta Psychologica*, *107*(1–3), 9–42. https://doi.org/10.1016/S0001-6918(01)00018-X

Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., Rosen, B. R., & Tootell, R. B. H. (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., & Zuck, E. (1994). Segregated Processing of Facial Identity and Emotion in the Human Brain: A PET Study. *Visual Cognition*, *1*(2–3), 349–369. https://doi.org/10.1080/13506289408402305

Shen, W., Tu, Y., Gollub, R. L., Ortiz, A., Napadow, V., Yu, S., Wilson, G., Park, J., Lang, C., Jung, M., Gerber, J., Mawla, I., Chan, S. T., Wasan, A. D., Edwards, R. R., Kaptchuk, T., Li, S., Rosen, B., & Kong, J. (2019). Visual network alterations in brain functional connectivity in chronic low back pain: A resting state functional connectivity and machine learning study. *NeuroImage: Clinical*, *22*, 101775. https://doi.org/10.1016/j.nicl.2019.101775

Silson, E. H., McKeefry, D. J., Rodgers, J., Gouws, A. D., Hymers, M., & Morland, A. B. (2013). Specialized and independent processing of orientation and shape in visual field maps LO1 and LO2. *Nature Neuroscience*, *16*(3), 267–269. https://doi.org/10.1038/nn.3327

Sincich, L. C., Park, K. F., Wohlgemuth, M. J., & Horton, J. C. (2004). Bypassing V1: A direct geniculate input to area MT. *Nature Neuroscience*, *7*(10), 1123–1128. https://doi.org/10.1038/nn1318

Smith, A. T., Wall, M. B., Williams, A. L., & Singh, K. D. (2006). Sensitivity to optic flow in human cortical areas MT and MST. *European Journal of Neuroscience*, *23*(2), 561–569. https://doi.org/10.1111/j.1460-9568.2005.04526.x

Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R., & Beckmann, C. F. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(31), 13040–13045. https://doi.org/10.1073/pnas.0905267106

Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E. J., Johansen-Berg, H., Bannister, P. R., de Luca, M., Drobnjak, I., Flitney, D. E., Niazy, R. K., Saunders, J., Vickers, J., Zhang, Y., de Stefano, N., Brady, J. M., & Matthews, P. M. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, *23*(SUPPL. 1). https://doi.org/10.1016/j.neuroimage.2004.07.051

Solcà, M., Guggisberg, A. G., Schnider, A., & Leemann, B. (2015). Facial blindsight. *Frontiers in Human Neuroscience*, *9*(SEPTEMBER), 522. https://doi.org/10.3389/fnhum.2015.00522

Sonkusare, S., Breakspear, M., & Guo, C. (2019). Naturalistic Stimuli in Neuroscience: Critically Acclaimed. *Trends in Cognitive Sciences*, *23*(8), 699–714. https://doi.org/10.1016/J.TICS.2019.05.004

Spaulding, W., & Deogun, J. (2011). A pathway to personalization of integrated treatment: Informatics and decision science in psychiatric rehabilitation. *Schizophrenia Bulletin*, *37*(SUPPL. 2), S129–S137. https://doi.org/10.1093/schbul/sbr080

Springer, J. A., Binder, J. R., Hammeke, T. A., Swanson, S. J., Frost, J. A., Bellgowan, P. S. F., Brewer, C. C., Perry, H. M., Morris, G. L., & Mueller, W. M. (1999). Language dominance in neurologically normal and epilepsy subjects. A functional MRI study. *Brain*. https://doi.org/10.1093/brain/122.11.2033

Stoerig, P. (2006). Chapter 12 Blindsight, conscious vision, and the role of primary visual cortex. In *Progress in Brain Research: Vol. 155 B* (pp. 217–234). https://doi.org/10.1016/S0079-6123(06)55012-5

Streiner, D. L., & Norman, G. R. (2011). Correction for multiple testing: Is there a resolution? In *Chest* (Vol. 140, Issue 1, pp. 16–18). https://doi.org/10.1378/chest.11-0523

Strong, S. L., Silson, E. H., Gouws, A. D., Morland, A. B., & McKeefry, D. J. (2017). A direct demonstration of functional differences between subdivisions of human V5/MT+. *Cerebral Cortex*, *27*(1), 1–10. https://doi.org/10.1093/cercor/bhw362

Strong, S. L., Silson, E. H., Gouws, A. D., Morland, A. B., & McKeefry, D. J. (2019). An enhanced role for right hV5/MT+ in the analysis of motion in the contra- and ipsi-lateral visual hemi-fields. *Behavioural Brain Research*. https://doi.org/10.1016/j.bbr.2019.112060

Suchoff, I. B., Kapoor, N., Ciuffreda, K. J., Rutner, D., Han, E., & Craig, S. (2008). The frequency of occurrence, types, and characteristics of visual field defects in acquired brain injury: A retrospective analysis. *Optometry - Journal of the American Optometric Association*, *79*(5), 259–265. https://doi.org/10.1016/j.optm.2007.10.012

Talairach, J., & Tournoux, P. (1988). Co-planar stereotaxic atlas of the human brain: 3-dimensional proportional .... In *Thieme New York* (Vol. 270, Issue 99).

Tanaka, K., Hikosaka, K., Saito, H. A., Yukie, M., Fukada, Y., & Iwai, E. (1986). Analysis of local and wide-field movements in the superior temporal visual areas of the Macaque monkey. *Journal of Neuroscience*, *6*(1), 134–144. https://doi.org/10.1523/jneurosci.06-01-00134.1986

Thakral, P. P., & Slotnick, S. D. (2011). Disruption of MT impairs motion processing. *Neuroscience Letters*. https://doi.org/10.1016/j.neulet.2010.12.057

Thaler, L., Schütz, A. C., Goodale, M. A., & Gegenfurtner, K. R. (2013). What is the best fixation target? The effect of target shape on stability of fixational eye movements. *Vision Research*. https://doi.org/10.1016/j.visres.2012.10.012

Tharwat, A. (2018). Independent component analysis: An introduction. In *Applied Computing and Informatics* (Vol. 17, Issue 2, pp. 222–249). https://doi.org/10.1016/j.aci.2018.08.006

Thomas, N., Acton, J. H., Erichsen, J. T., & Dunn, M. J. (2019). Validation of a novel gaze-contingent perimeter with high-speed eye tracking. *Investigative Ophthalmology and Visual Science*, *60*(9), 522–522. https://www.embase.com/search/results?subaction=viewrecord&id=L629939115&from=export

Thompson, P. M., Schwartz, C., Lin, R. T., Khan, A. A., & Toga, A. W. (1996). Three-dimensional statistical analysis of sulcal variability in the human brain. *Journal of Neuroscience*, *16*(13), 4261–4274. https://doi.org/10.1523/jneurosci.16-13-04261.1996

Tootell, R. B. H., Mendola, J. D., Hadjikhani, N. K., Ledden, P. J., Liu, A. K., Reppas, J. B., Sereno, M. I., & Dale, A. M. (1997). Functional analysis of V3A and related areas in human visual cortex. *Journal of Neuroscience*, *17*(18), 7060–7078. https://doi.org/10.1523/jneurosci.17-18-07060.1997

Tootell, R. B. H., Reppas, J. B., Kwong, K. K., Malach, R., Born, R. T., Brady, T. J., Rosen, B. R., & Belliveau, J. W. (1995). Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging. *Journal of Neuroscience*. https://doi.org/10.1523/jneurosci.15-04-03215.1995

Tosi, V., Mecacci, L., & Pasquali, E. (1997). Scanning eye movements made when viewing film: Preliminary observations. *International Journal of Neuroscience*, *92*(1–2), 47–52. https://doi.org/10.3109/00207459708986388

Townend, B. S., Sturm, J. W., Petsoglou, C., O'Leary, B., Whyte, S., & Crimmins, D. (2007). Perimetric homonymous visual field loss post-stroke. *Journal of Clinical Neuroscience*, *14*(8), 754–756. https://doi.org/10.1016/j.jocn.2006.02.022

Triantafyllou, C., Hoge, R. D., & Wald, L. L. (2006). Effect of spatial smoothing on physiological noise in high-resolution fMRI. *NeuroImage*, *32*(2), 551–557. https://doi.org/10.1016/j.neuroimage.2006.04.182

Ungerleider, & Mishkin. (1982). Two Cortical Visual Systems. In *Analysis of Visual Behavior*.

van Essen, D. C., & Drury, H. A. (1997). Structural and functional analyses of human cerebral cortex using a surface-based atlas. *Journal of Neuroscience*, *17*(18), 7079–7102. https://doi.org/10.1523/jneurosci.17-18-07079.1997

Vanderwal, T., Kelly, C., Eilbott, J., Mayes, L. C., & Castellanos, F. X. (2015). Inscapes: A movie paradigm to improve compliance in functional magnetic resonance imaging. *NeuroImage*, *122*, 222–232. https://doi.org/10.1016/j.neuroimage.2015.07.069

Vossel, S., Eschenbeck, P., Weiss, P. H., Weidner, R., Saliger, J., Karbe, H., & Fink, G. R. (2011). Visual extinction in relation to visuospatial neglect after right-hemispheric stroke: Quantitative assessment and statistical lesion-symptom mapping. *Journal of Neurology, Neurosurgery and Psychiatry*, *82*(8), 862–868. https://doi.org/10.1136/jnnp.2010.224261

Wandell, B. A., & Winawer, J. (2011a). Imaging retinotopic maps in the human brain. In *Vision Research* (Vol. 51, Issue 7, pp. 718–737). Pergamon. https://doi.org/10.1016/j.visres.2010.08.004

Wandell, B. A., & Winawer, J. (2011b). Imaging retinotopic maps in the human brain. In *Vision Research* (Vol. 51, Issue 7, pp. 718–737). Pergamon. https://doi.org/10.1016/j.visres.2010.08.004

Wang, L., Mruczek, R. E. B., Arcaro, M. J., & Kastner, S. (2015). Probabilistic maps of visual topography in human cortex. *Cerebral Cortex*. https://doi.org/10.1093/cercor/bhu277

Warren, P. A., & Rushton, S. K. (2007). Perception of object trajectory: Parsing retinal motion into self and object movement components. *Journal of Vision*, *7*(11). https://doi.org/10.1167/7.11.2

Warren, P. A., & Rushton, S. K. (2009a). Optic Flow Processing for the Assessment of Object Movement during Ego Movement. *Current Biology*, *19*(18), 1555–1560. https://doi.org/10.1016/j.cub.2009.07.057

Warren, P. A., & Rushton, S. K. (2009b). Perception of scene-relative object movement: Optic flow parsing and the contribution of monocular depth cues. *Vision Research*, *49*(11), 1406–1419. https://doi.org/10.1016/j.visres.2009.01.016

Warren, W. H., & Hannon, D. J. (1988). Direction of self-motion is perceived from optical flow. *Nature*, *336*(6195), 162–163. https://doi.org/10.1038/336162a0

Watson, J. D. G., Shipp, S., Zeki, S., Watson, J. D. G., Myers, R., Frackowiak, R. S. J., Hajnal, J. v., Woods, R. P., & Mazziotta, J. C. (1993). Area v5 of the human brain: Evidence from a combined study using positron emission tomography and magnetic resonance imaging. *Cerebral Cortex*. https://doi.org/10.1093/cercor/3.2.79

Weiskrantz, L., Barburt, J. L., & Sahraie, A. (1995). Parameters affecting conscious versus unconscious visual discrimination with damage to the visual cortex (VI). In *Psychology* (Vol. 92, Issue 13, pp. 6122–6126). https://doi.org/10.1073/pnas.92.13.6122

Weiskrantz, L., Warrington, E. K., Sanders, M. D., & Marshall, J. (1974). Visual capacity in the hemianopic field following a restricted occipital ablation. *Brain*, *97*(4), 709–728. https://doi.org/10.1093/brain/97.4.709

Weiss, P. H., Marshall, J. C., Wunderlich, G., Tellmann, L., Halligan, P. W., Freund, H. J., Zilles, K., & Fink, G. R. (2000). Neural consequences of acting in near versus far space: A physiological basis for clinical dissociations. *Brain*, *123*(12), 2531–2541. https://doi.org/10.1093/brain/123.12.2531

Wilms, M., Eickhoff, S. B., Specht, K., Amunts, K., Shah, N. J., Malikovic, A., & Fink, G. R. (2005). Human V5/MT+: Comparison of functional and cytoarchitectonic data. *Anatomy and Embryology*, *210*(5–6), 485–495. https://doi.org/10.1007/s00429-005-0064-y

Wolfers, T., Doan, N. T., Kaufmann, T., Alnæs, D., Moberget, T., Agartz, I., Buitelaar, J. K., Ueland, T., Melle, I., Franke, B., Andreassen, O. A., Beckmann, C. F., Westlye, L. T., & Marquand, A. F. (2018). Mapping the Heterogeneous Phenotype of Schizophrenia and Bipolar Disorder Using Normative Models. *JAMA Psychiatry*, *75*(11), 1146–1155. https://doi.org/10.1001/jamapsychiatry.2018.2467

Woolrich, M. W., Jbabdi, S., Patenaude, B., Chappell, M., Makni, S., Behrens, T., Beckmann, C., Jenkinson, M., & Smith, S. M. (2009). Bayesian analysis of neuroimaging data in FSL. *NeuroImage*, *45*(1 Suppl). https://doi.org/10.1016/j.neuroimage.2008.10.055

Worsley, K. J., & Friston, K. J. (1995). Analysis of fMRI time-series revisited — Again. *NeuroImage*, *2*(3), 173–181. https://doi.org/10.1006/nimg.1995.1023

Yang, Z., Wu, J., Xu, L., Deng, Z., Tang, Y., Gao, J., Hu, Y., Zhang, Y., Qin, S., Li, C., & Wang, J. (2020). Individualized psychiatric imaging based on inter-subject neural synchronization in movie watching. *NeuroImage*, *216*, 116227. https://doi.org/10.1016/j.neuroimage.2019.116227

Zarahn, E., Aguirre, G. K., & D'Esposito, M. (1997). Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *NeuroImage*, *5*(3), 179–197. https://doi.org/10.1006/nimg.1997.0263

Zeki, S. (1974). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *The Journal of Physiology*, *236*(3), 549–573. https://doi.org/10.1113/jphysiol.1974.sp010452

Zeki, S., & Ffytche, D. H. (1998). The Riddoch syndrome: Insights into the neurobiology of conscious vision. *Brain*, *121*(1), 25–45. https://doi.org/10.1093/brain/121.1.25

Zeki, S., Watson, J. D. G., Lueck, C. J., Friston, K. J., Kennard, C., & Frackowiak, R. S. J. (1991). A direct demonstration of functional specialization in human visual cortex. *Journal of Neuroscience*, *11*(3), 641–649. https://doi.org/10.1523/jneurosci.11-03-00641.1991

Zhang, X., Kedar, S., Lynn, M. J., Newman, N. J., & Biousse, V. (2006a). Homonymous hemianopia in stroke. *American Journal of Ophthalmology*, *142*(6), 1097. https://doi.org/10.1016/j.ajo.2006.10.017

Zhang, X., Kedar, S., Lynn, M. J., Newman, N. J., & Biousse, V. (2006b). Homonymous hemianopias: Clinical-anatomic correlations in 904 cases. In *Neurology* (Vol. 66, Issue 6, pp. 906–910). https://doi.org/10.1212/01.wnl.0000203913.12088.93

Zilles, K. (2018). Brodmann: A pioneer of human brain mapping-His impact on concepts of cortical organization. In *Brain* (Vol. 141, Issue 11, pp. 3262–3278). Oxford University Press. https://doi.org/10.1093/brain/awy273

Zilles, K., & Amunts, K. (2010a). Centenary of Brodmann's map — conception and fate. In *Nature Reviews Neuroscience 2010 11:2* (Vol. 11, Issue 2, pp. 139–145). Nature Publishing Group. https://doi.org/10.1038/nrn2776

Zilles, K., & Amunts, K. (2010b). Centenary of Brodmann's map conception and fate. In *Nature Reviews Neuroscience* (Vol. 11, Issue 2, pp. 139–145). Nature Publishing Group. https://doi.org/10.1038/nrn2776

Main Appendices

Appendix I **fMRIPrep boilderplate**

The following section is taken directly from the report generated following fMRIPrep pre-processing.

FMRIPrep (MTW dataset)

Results included in this manuscript come from pre-processing performed using fMRIPrep 20.2.0 (Esteban, Markiewicz, et al. (2018); Esteban, Blair, et al. (2018); RRID:SCR_016216), which is based on Nipype 1.5.1 (Gorgolewski et al. (2011); Gorgolewski et al. (2018); RRID:SCR_002502).

Anatomical data preprocessing

A total of 1 T1-weighted (T1w) images were found within the input BIDS dataset. The T1-weighted (T1w) image was corrected for intensity non-uniformity (INU) with N4BiasFieldCorrection (Tustison et al. 2010), distributed with ANTs 2.3.3 (Avants et al. 2008, RRID:SCR_004757), and used as T1w-reference throughout the workflow. The T1w-reference was then skull-stripped with a Nipype implementation of the antsBrainExtraction.sh workflow (from ANTs), using OASIS30ANTs as target template. Brain tissue segmentation of cerebrospinal fluid (CSF), white-matter (WM) and gray-matter (GM) was performed on the brain-extracted T1w using fast (FSL 5.0.9, RRID:SCR_002823, Zhang, Brady, and Smith 2001). Brain surfaces were reconstructed using recon-all (FreeSurfer 6.0.1, RRID:SCR_001847, Dale, Fischl, and Sereno 1999), and the brain mask estimated previously was refined with a custom variation of the method to reconcile ANTs-derived and FreeSurfer-derived segmentations of the cortical gray-matter of Mindboggle (RRID:SCR_002438, Klein et al. 2017). Volume-based spatial normalisation to one standard space (MNI152nLin2009cAsym) was performed through nonlinear registration with antsRegistration (ANTs 2.3.3), using brain-extracted versions of both T1w reference and the T1w template. The following template was selected for spatial normalisation: ICBM 152 Nonlinear Asymmetrical template version 2009c [Fonov et al. (2009), RRID:SCR_008796; TemplateFlow ID: MNI152nLin2009cAsym],

Functional data preprocessing

For each of the 1 BOLD runs found per subject (across all tasks and sessions), the following preprocessing was performed. First, a reference volume and its skull-stripped version were generated using a custom methodology of fMRIPrep. A B0-nonuniformity map (or fieldmap) was estimated based on two (or more) echo-planar imaging (EPI) references with opposing phase-encoding directions, with 3dQwarp Cox and Hyde (1997) (AFNI 20160207). Based on the estimated susceptibility distortion, a corrected EPI (echo-planar imaging) reference was calculated for a more accurate co-registration with the anatomical reference. The BOLD reference was then co-registered to the T1w reference using bbregister (FreeSurfer) which implements boundary-based registration (Greve and Fischl 2009). Co-registration was configured with six degrees of freedom. Head-motion parameters with respect to the BOLD reference (transformation matrices, and six corresponding rotation and translation parameters) are estimated before any spatiotemporal filtering using mcflirt (FSL 5.0.9, Jenkinson et al. 2002). BOLD runs were slice-time corrected using 3dTshift from AFNI 20160207 (Cox and Hyde 1997, RRID:SCR_005927). The BOLD time-series were resampled onto the following surfaces (FreeSurfer reconstruction nomenclature): fsnative. The BOLD time-series (including slice-timing correction when applied) were resampled onto their original, native space by applying a single, composite transform to correct for head-motion and susceptibility distortions. These resampled BOLD time-series will be referred to as preprocessed BOLD in original space, or just preprocessed BOLD. Several confounding time-series were calculated based on the preprocessed BOLD: framewise displacement (FD), DVARS and three region-wise global signals. FD was computed using two formulations following Power (absolute sum of relative motions, Power et al. (2014)) and Jenkinson (relative root mean square displacement between affines, Jenkinson et al. (2002)). FD and DVARS are calculated for each functional run, both using their implementations in Nipype (following the definitions by Power et al. 2014). The three global signals are extracted within the CSF, the WM, and the whole-brain masks. Additionally, a set of physiological regressors were extracted to allow for component-based noise correction (CompCor, Behzadi et al. 2007). Principal components are estimated after high-pass filtering the preprocessed BOLD time-series (using a discrete cosine filter with 128s cut-off) for the two CompCor variants: temporal (tCompCor) and anatomical

(aCompCor). tCompCor components are then calculated from the top 2% variable voxels within the brain mask. For aCompCor, three probabilistic masks (CSF, WM and combined CSF+WM) are generated in anatomical space. The implementation differs from that of Behzadi et al. in that instead of eroding the masks by 2 pixels on BOLD space, the aCompCor masks are subtracted a mask of pixels that likely contain a volume fraction of GM. This mask is obtained by dilating a GM mask extracted from the FreeSurfer's aseg segmentation, and it ensures components are not extracted from voxels containing a minimal fraction of GM. Finally, these masks are resampled into BOLD space and binarized by thresholding at 0.99 (as in the original implementation). Components are also calculated separately within the WM and CSF masks. For each CompCor decomposition, the k components with the largest singular values are retained, such that the retained components' time series are sufficient to explain 50 percent of variance across the nuisance mask (CSF, WM, combined, or temporal). The remaining components are dropped from consideration. The head-motion estimates calculated in the correction step were also placed within the corresponding confounds file. The confound time series derived from head motion estimates and global signals were expanded with the inclusion of temporal derivatives and quadratic terms for each (Satterthwaite et al. 2013). Frames that exceeded a threshold of 0.5 mm FD or 1.5 standardised DVARS were annotated as motion outliers. All resamplings can be performed with a single interpolation step by composing all the pertinent transformations (i.e. head-motion transform matrices, susceptibility distortion correction when available, and co-registrations to anatomical and output spaces). Gridded (volumetric) resamplings were performed using antsApplyTransforms (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of other kernels (Lanczos 1964). Non-gridded (surface) resamplings were performed using mri_vol2surf (FreeSurfer).

Appendix II **Component sharing pattern of IPS regions**

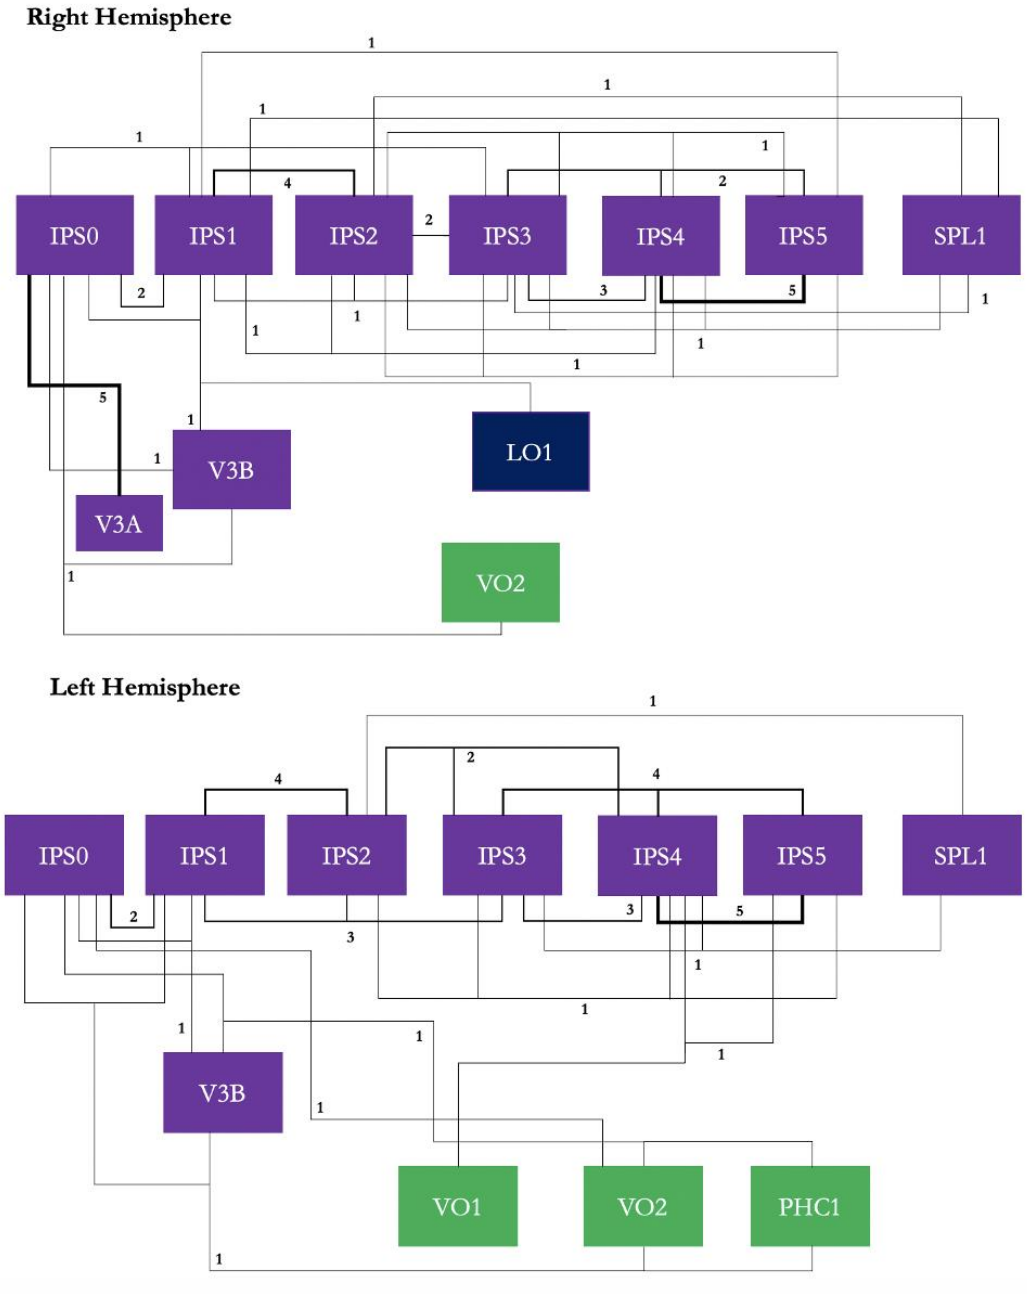

Figure 108. Pattern of component sharing (the instances in which a single component was the top match to more than one visual area) across brains in the IPS regions. Lines connecting visual areas indicate that the regions were included in a single component. Thicker lines indicate that more brains showed that pattern of component sharing. Note, the number above each line shows the number of brains that displayed this pattern of component sharing.

Appendix III **Number of components**

Table 20. The number of output components produced by each subject following single-subject ICA.

Subject	Number of ICs output
1	145
2	155
3	173
4	139
5	192
6	124
7	179
8	125
9	136
10	149
11	81
12	214
13	155
14	164
15	170
16	55
17	111
Mean	145.12
Range	55 – 214

Appendix IV **Mean number of matches**

For each subject we calculated the percentage of the total number of matches corresponding to each visual area (Table 21). For example, Subject 6 produced a total of 28 matches above threshold. Of which three, and therefore 7.89%, of the total number of matches corresponded to V1v. We then took the mean of these percentages across subjects to assess whether on average areas tend have multiple component matches, or whether a single component match is typically output. The subjects who did not output a component match above threshold were removed from the mean calculation for that visual area. For example, 15 subjects produced at least one component which correlated with the left hemisphere V1v atlas region, and therefore the mean number of V1v matches was calculated using the values of those 15 subjects.

Table 21. The mean number, across brains, of above threshold matches for each visual area and the mean per cent of the total number of matches produced for each area. The brains that did not produce an above threshold match were excluded from these mean calculations.

	Non match brains excluded			
	Left Hemisphere		Right Hemisphere	
	Mean number of matches	Mean % of total matches	Mean number of matches	Mean % of total matches
V1v	1.60	4.34	1.80	4.39
V1d	1.94	4.95	2.24	5.89
V2v	1.43	3.81	1.40	3.51
V2d	2.18	5.64	2.00	4.97
V3v	1.23	3.09	1.17	2.99
V3d	1.81	4.69	1.47	3.69
hV4	1.50	4.05	1.81	4.49
VO1	1.15	2.86	1.27	3.21
VO2	1.30	3.43	1.67	4.07
PHC1	1.69	4.16	1.79	4.30
PHC2	2.08	5.01	1.80	4.38
TO2	1.53	4.09	1.35	3.46
TO1	1.81	4.72	1.44	3.57
LO2	1.29	3.52	1.23	2.99
LO1	1.47	4.07	1.38	3.51
V3B	1.43	3.71	1.79	4.60
V3A	1.67	4.48	1.67	4.11
IPS0	1.53	4.27	1.38	3.58
IPS1	2.12	6.04	1.60	4.02
IPS2	2.12	5.77	2.12	5.71
IPS3	2.35	6.32	2.56	6.82
IPS4	1.82	4.76	2.07	5.34
IPS5	1.25	3.26	2.06	5.18
SPL1	2.06	5.81	2.00	5.20
FEF	2.50	6.57	2.38	6.16

Appendix V **Mean correlation between top matched component and atlas regions**

For each subject, we performed spatial correlations between the atlas and the thresholded z-score spatial maps. We then took the mean correlation across brains for each visual area (Figure 109) As expected, higher correlations were found between the atlas regions and the thresholded spatial maps compared to the unthresholded maps. However, the mean correlation across all areas increased by only 0.03 when thresholded spatial maps were used.

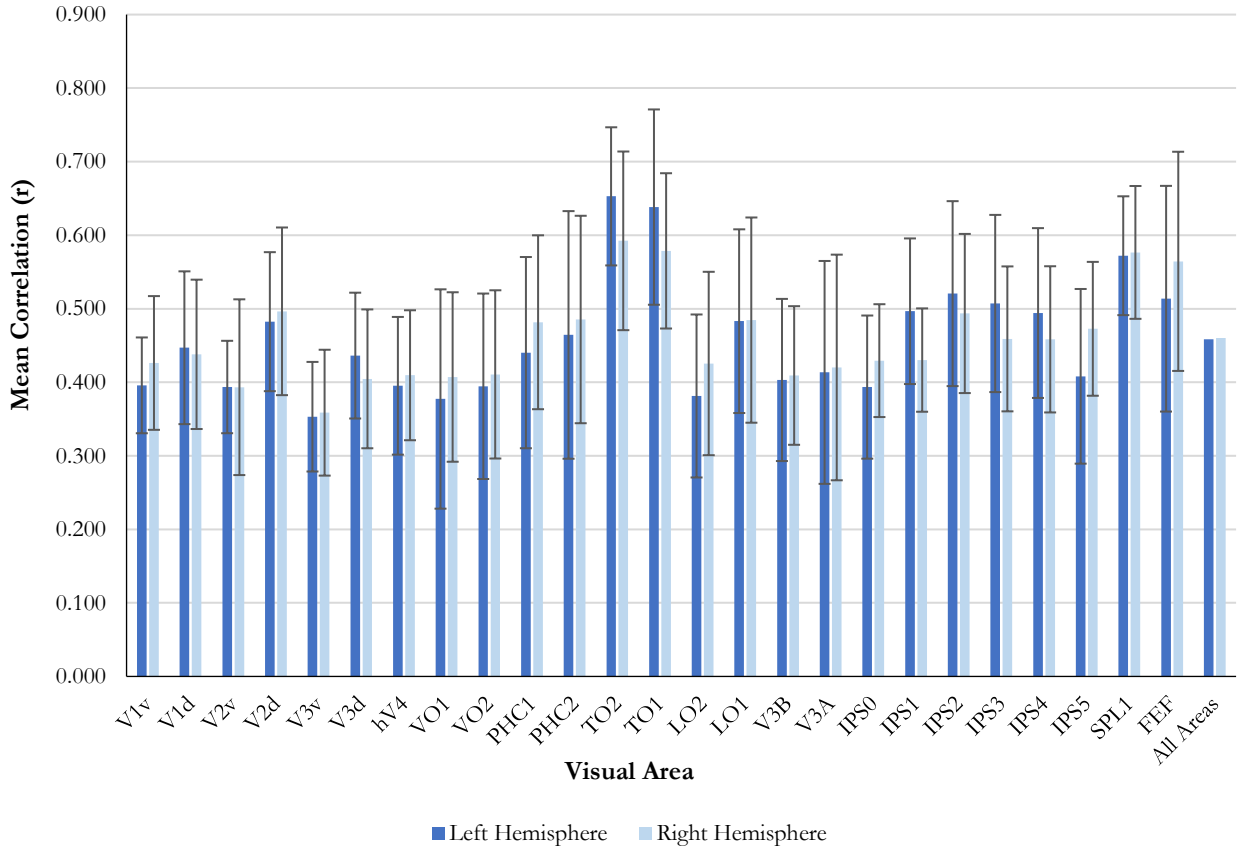

Figure 109. Mean correlation between "top matched" components and atlas regions using thresholded z stat spatial maps.

Appendix VI **ROI-to-ROI validation in detail**

Cross hemisphere correlations between the timecourses extracted from visual and non-visual are presented in Table 22.

Table 22. Group mean correlation coefficients computed between the mean timecourses extracted from V1v and the combined visual labels in the opposite hemisphere.

	Mean	Std Dev
Left V1v vs. Right combined visual	0.816	0.055
Right V1v vs. Left combined visual	0.794	0.059
Left V1v vs. Right outside visual	0.546	0.098
Right V1v vs. Left outside visual	0.519	0.121
Left V1v vs. Right BA	0.343	0.112
Right V1v vs. Left BA	0.284	0.146
Left V1v vs. Right auditory	0.328	0.132
Right V1v vs. Left auditory	0.327	0.127
Right combined visual vs. Left BA	0.331	0.160
Left combined visual vs. Right BA	0.458	0.111
Left combined visual vs. Right auditory	0.356	0.147
Right combined visual vs. Left auditory	0.329	0.143
Left BA vs. Right auditory	0.372	0.158
Right BA vs. Left auditory	0.334	0.193

Appendix VII **AROMA data versus standard pre-processing**

AROMA data versus standard pre-processing

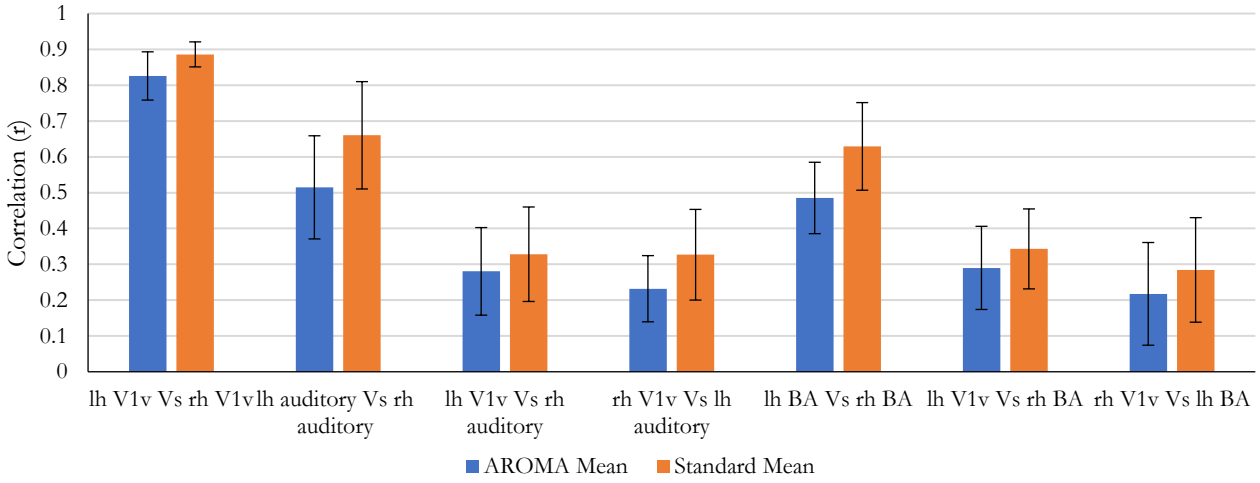

Figure 110. V1v versus Broca's area and auditory cortex. Group mean correlations computed between the timecourses of V1v and two non-visual regions located in the opposite hemisphere; Broca's area and an auditory region located in the middle of the Heschl's gyrus (HG). Correlations were computed with denoised data (AROMA mean) and data following standard pre-processing (standard mean).

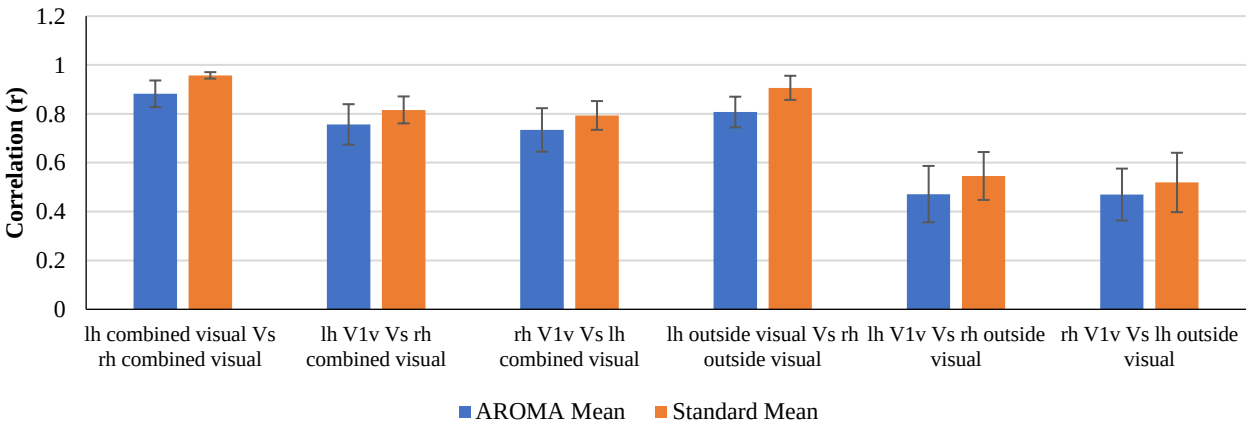

Figure 111. V1v versus combined visual and outside visual. Group mean correlations computed between the timecourses of V1v and two large regions located in the opposite hemisphere: one visual and one non-visual. Correlations were computed with denoised data (AROMA mean) and data following standard pre-processing (standard mean).

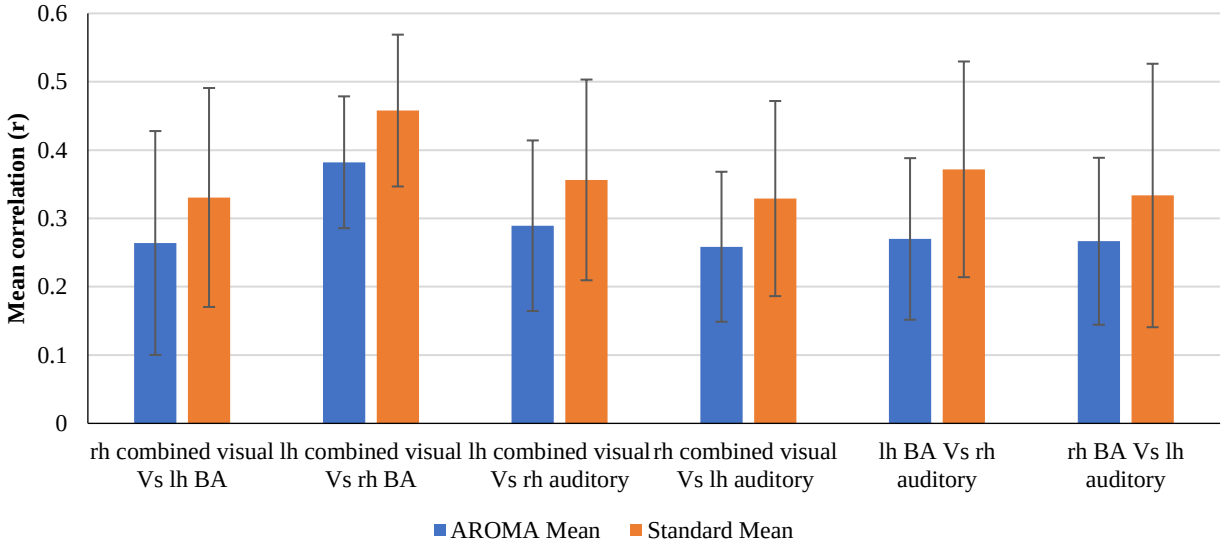

Figure 112. Combined visual versus Broca's area and auditory cortex. Group mean correlations computed between the timecourses of the combined visual region and two non-visual regions located in the opposite hemisphere; Broca's area and an auditory region located in the middle of the Heschl's gyrus (HG). Correlations were computed with denoised data (AROMA mean) and data following standard pre-processing (standard mean).

Appendix VIII **Vertex-to-vertex analysis (MPM, intersection and exclusive ROIs)**

For completeness we ran the vertex-to-vertex analysis, this time extracting mean correlation values from vertices within the set of intersection labels as well as the maximum probability maps (MPM) ROIs. The results hold true regardless of whether the correlations are computed using exclusive or intersection labels (Figure 113). The pattern of the results holds true when extracted mean correlations from MPM (Figure 114).

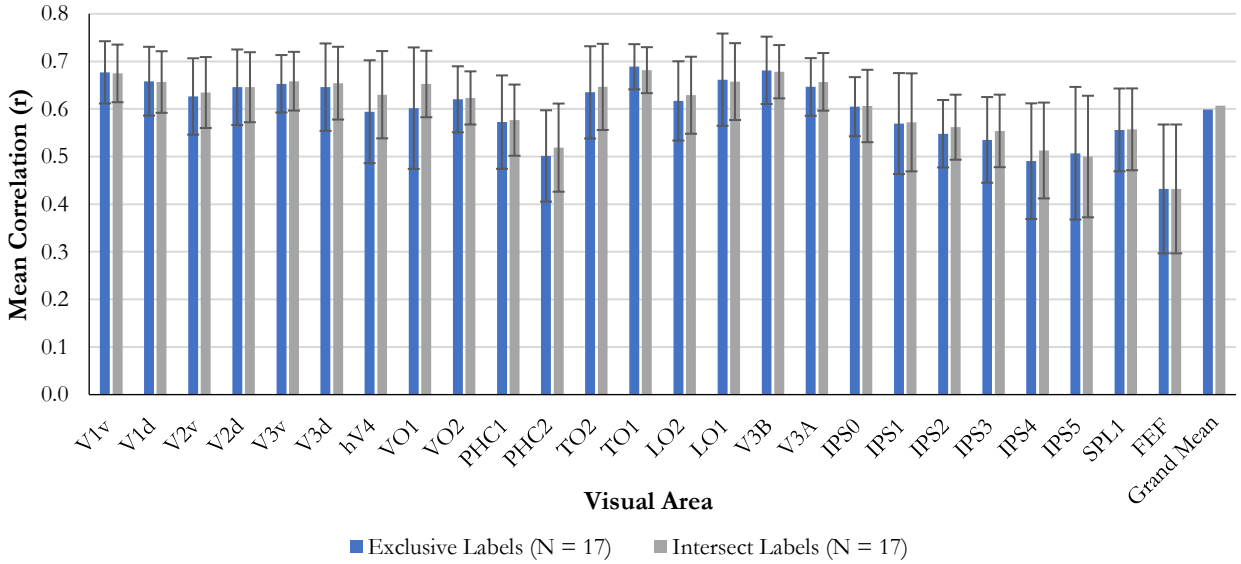

Figure 113. Vertex-to-vertex cross hemisphere timecourse correlation analysis results. Mean correlations were extracted from vertices within the intersection and exclusive ROIs.

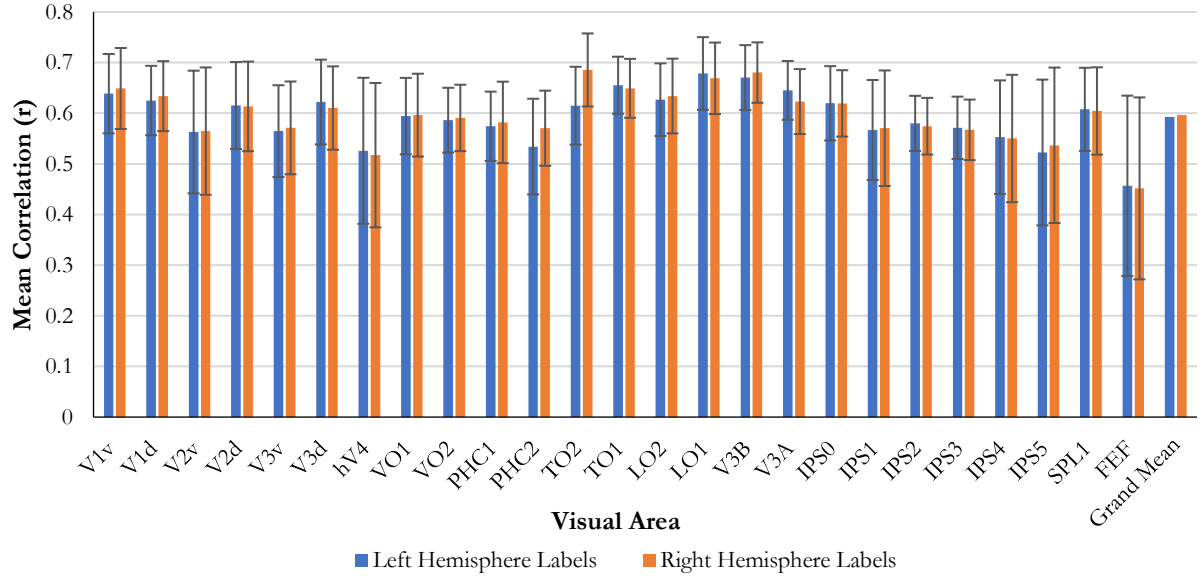

Figure 114. Vertex-to-vertex cross hemisphere timecourse correlation analysis results. Mean correlations were extracted from vertices within the left and right hemisphere MPM ROIs. Results are consistent with those obtained from using other ROIs.

Appendix IX **Vertex-to-vertex correlation analysis with volume-smoothed data**

The grand means, calculated by taking the mean r across all the visual areas, for volumetric smoothing (r = 0.582) and surface-based smoothing (r = 0.599) is comparable. The results of the analysis with volume-smoothed data reflects those obtained from surface-smoothed approach. There is two apparent discrepancies between the volume- and surface-based smoothing results, V4 and V2v. Following volume-based smoothing, results show that the mean correlation for hV4 (mean r = 0.465, SD = 0.148) is the second lowest. Likewise, V2v (r = 0.562, SD = 0.115) is lower than neighbouring visual areas. No such trend appears in surface-smoothed data. However, the low correlations seen for hV4 and V2v could be explained by a lack of EPI coverage around these regions. From the raw data sampled to the fsaverage surface, we identified seven brains without full coverage on the inferior surface (Figure 115). The surface-based smoothing approach appears to "fill in" these gaps in data.

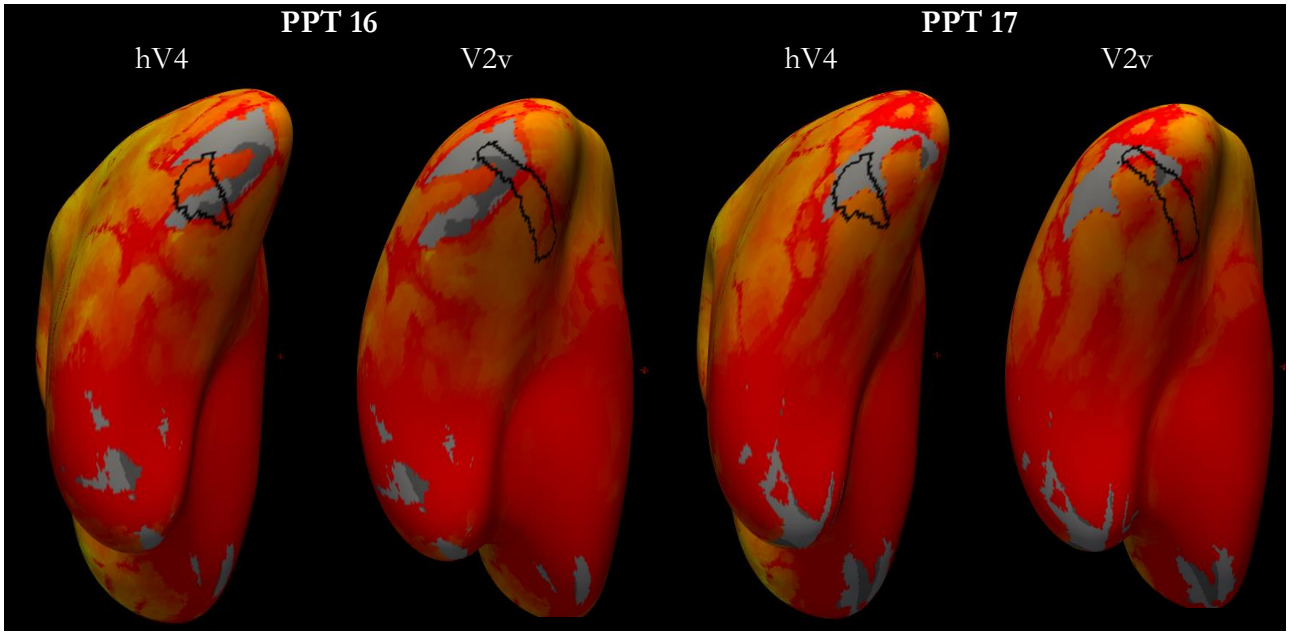

Figure 115. Inferior surface of two example subjects (Subject 16 and 17) showing fMRI data resampled to the fsaverage surface with hV4 and V2v labels overlaid (black).

To investigate whether the low correlations in hV4 and V2v are a result of a lack of coverage, we recalculated the group mean correlations for each visual area, this time omitting the seven brains (Figure 116). As expected, following the omissions of the brains with a lack of coverage, the mean correlation for hV4 increased for volume-smoothed (mean r = 0.569, SD = 0.093) and surface-smoothed data (mean r = 0.626, SD = 0.078). Likewise, the mean correlation of V2v increased for

both the volume-smoothed (mean r = 0.625, SD = 0.064) and surface-smoothed data (mean r = 0.667, SD = 0.059). With the smaller sample of whole-coverage brains (n = 10), the grand mean for the volume-smoothed data increased by 0.011 (r = 0.593), whereas the increase seen for the surface-smoothed data was 0.004 (r = 0.603). When we control for the lack of data issue in the seven brains by omitting them from the group mean, surface-based and volumetric smoothing correlation results are comparable across visual areas.

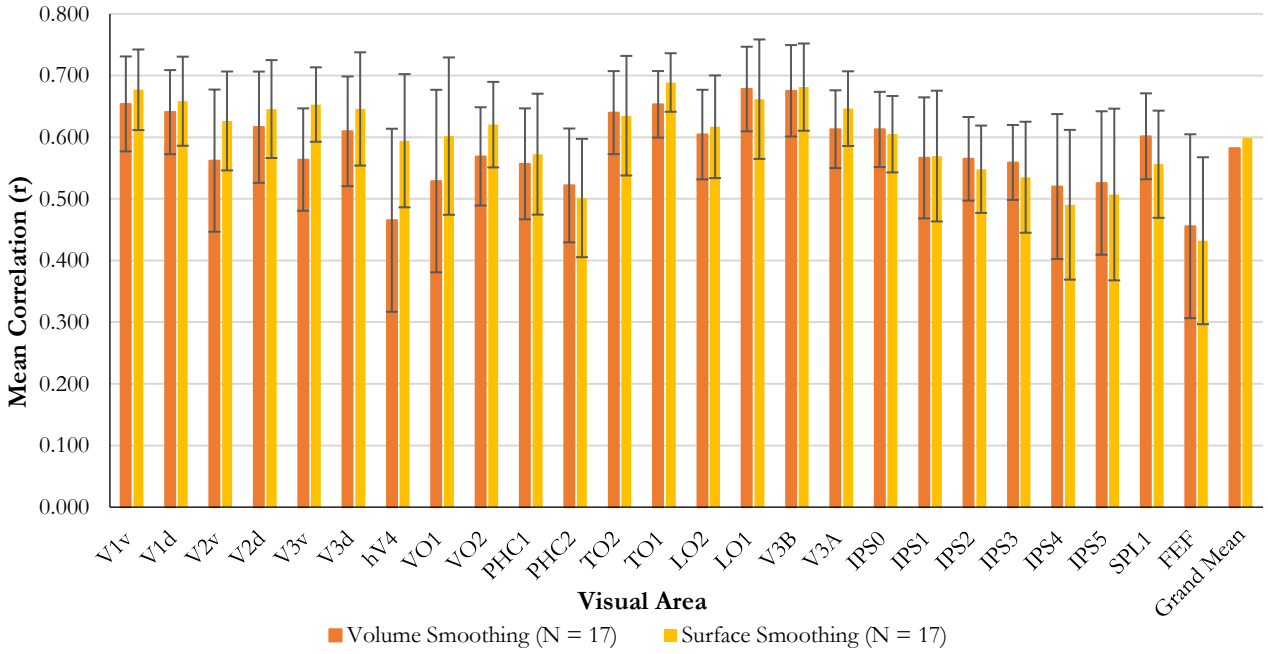

Figure 116. Mean interhemispheric correlation for each visual area, following volume- or surface-based smoothing. Mean timecourses were extracted from exclusive surface labels. Error bars represent standard deviation.

Appendix X **Between brains timecourse consistency – volume smoothed data**

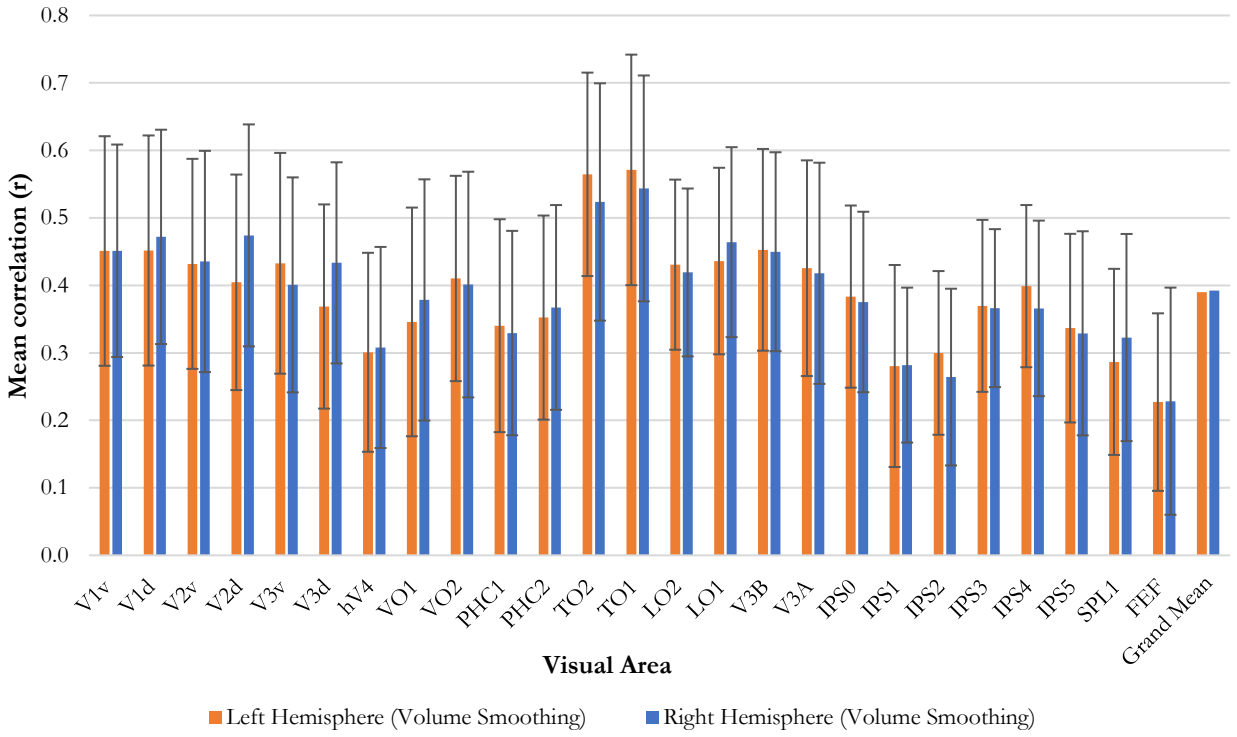

Figure 117. Group mean between-brains correlations for both the left and right hemispheres, using volume-smoothed data. Correlations were computed between each individual's visual area timecourse and the group average timecourse for that visual area (the group average timecourse did not include the brain with which the correlation was being computed). The mean of each of these subject-to-group correlations was then calculated. This analysis gives a measure of the consistency of the timecourses for each visual area. Results show that for TO1 and TO2, left hemisphere timecourses are more highly correlated across brains than right hemisphere timecourses. Error bars represent standard deviation.

Appendix XI **ROI-to-ROI cross hemisphere correlations - native surface validation**

Our results seem to suggest there are individual differences in the function of TO1/MT and TO2/MST across hemispheres. A sensible question to ask is: do we find these differences when analyses are performed on the subject's individual surface? I address this question in the following section by projection each subject's BOLD data to their individual surface and performing analyses on these data.

Methods

There were two differences between the average surface and native surface analyses. 1) pre-processed data were resampled to the native surface, rather than the fsaverage surface (using mri_vol2surf). Again, data were spatially smoothed on the surface with a 5mm FWHM kernel (using mri_surf2surf). 2) The mean timecourse of each visual area was extracted from the max probability maps (MPM), rather than the exclusive ROIs. The ROI-to-ROI cross-hemisphere correlations and the within-hemisphere timecourse consistency analyses were then performed on these native-space timecourses. Results from these analyses are consistent with analyses performed on the average surface.

Results

Results are consistent with those derived from the average surface analysis (Figure 118) TO1/MT does not appear to be no more lateralised than any other visual area. When extracted from the native surface, TO1/MT timecourses were similarly correlated across hemispheres ($r = 0.844$, $SD = 0.044$) as V1v ($r = 0.883$, $SD = 0.035$) and V1d ($r = 0.849$, $SD = 0.052$). As with the average brain analysis, TO2/MST timecourses were less correlated across hemispheres ($r = 0.780$, $SD = 0.072$) than TO1/MT, V1v and V1d.

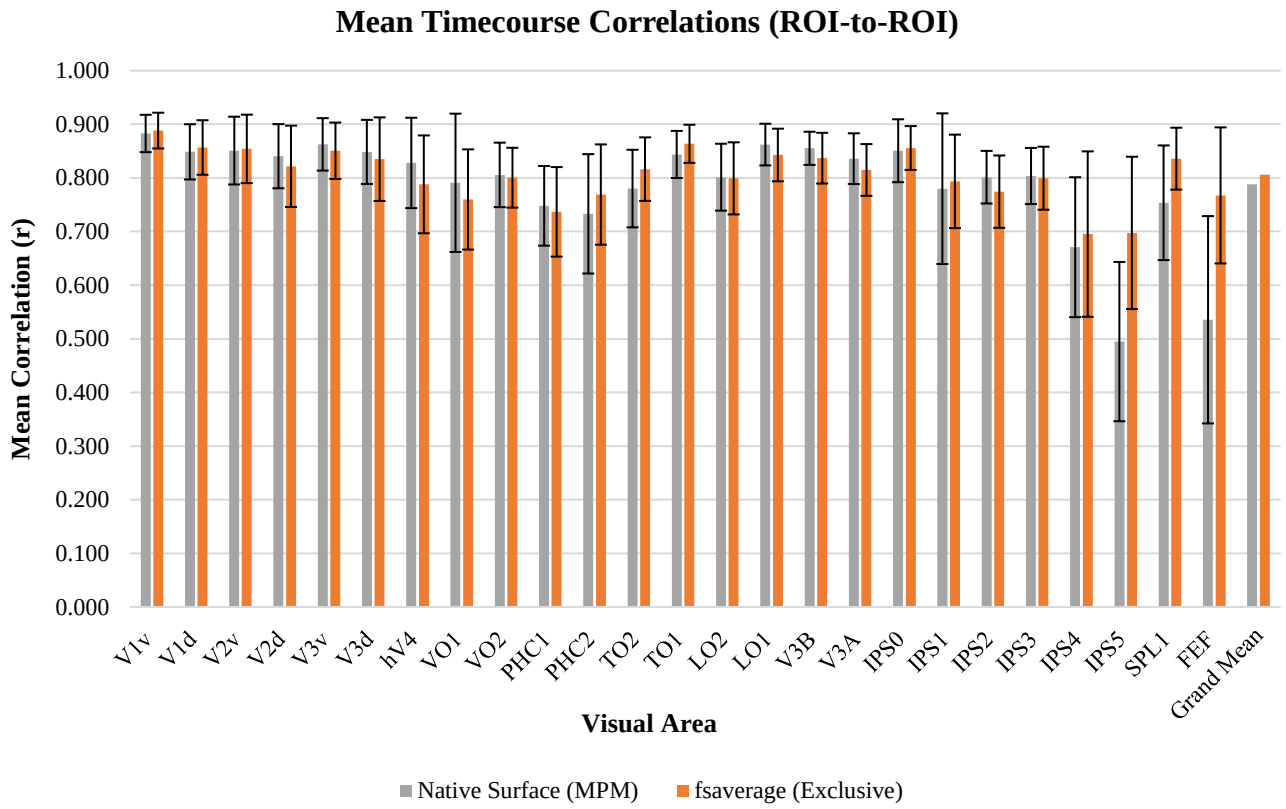

Figure 118. Group mean cross-hemisphere timecourse correlation coefficients for each of the 25 visual areas. Error bars represent standard deviation. Results from both the average and native surface analyses are presented. The timecourses of the visual areas are highly correlated across hemispheres. TO2 is less correlated than TO1, V1v and V1d, suggesting there may be differences in function across hemispheres for TO2.

Ratio of cross-hemisphere V1 correlations to TO1/TO2 correlations

Again, for each subject we divided the cross-hemispheric correlation for TO1/MT and TO2/MST (separately) by the correlation for V1 to generate a timecourse ratio of V1 to TO1/TO2 (Table 23). For TO1, the timecourses are similarly correlated across hemispheres as V1 for all but one subject, indicated by ratios close to 1. This suggests on the whole TO1/MT is no more lateralised than V1. However, for 7/17 subjects we find V1 timecourses are more highly correlated across hemispheres than TO2/MST timecourses. When this analysis was performed on the average surface, 5/17 brains showed this pattern. Of the seven brains showing lower correlations in TO2/MST compared to V1 in native surface space, four also showed the same pattern in average surface space. The additional three brains did show differences between the timecourses of TO2/MST and V1, but they are more pronounced in the native space analysis.

Table 23. The individual subject ratios of the cross-hemisphere timecourse correlations of TO1 and TO2 to V1. For each subject, the ratios were calculated by dividing the cross-hemispheric correlation coefficient for TO1 and TO2 (separately) by the cross-hemisphere correlation coefficient for V1 (mean of V1v and V1d). Values > 1 indicate the timecourses of TO1/TO2 were less correlated across hemispheres than V1. Results from the native and average surface analysis are shown. Analysis on the native surface suggests that for 7 brains V1 timecourses were more highly correlated across hemispheres than TO2 timecourses. For average surface data, 5 brains show this pattern.

	Native Surface		Average Surface	
	TO1 versus V1	TO2 versus V1	TO1 versus V1	TO2 versus V1
6	0.884	0.778	0.925	0.765
7	1.011	1.007	0.995	0.96
8	1.001	0.795	1.025	0.844
9	1.06	1.002	1.088	1.08
10	0.946	0.98	1.004	1.028
11	0.93	0.949	0.942	1.027
12	0.999	0.87	0.988	0.941
13	1.006	0.878	1.012	0.926
14	0.95	0.876	0.949	0.844
15	0.986	0.944	1	0.85
16	0.983	0.745	0.983	0.845
17	0.93	0.908	0.945	1.013
18	0.957	0.891	0.958	0.912
19	1.012	0.922	1.001	0.949
20	0.983	0.961	1.037	1.05
21	0.967	0.904	0.998	0.953
22	0.972	0.91	0.991	0.958

Between-brains correlations

The results from between-brains analyses performed on the native surface reflect those obtained from average surface analyses. We find left hemisphere TO1 and TO2 timecourses were less correlated across brains than the right hemispheres timecourses. This differences in the consistency of timecourses is not observed in V1v or V1d. This is consistent with analyses performed on the average surface.

Timecourses extracted from left V1v ($r = 0.471$, $SD = 0.176$) and right V1v ($r = 0.493$, $SD = 0.165$) were similarly correlated across brains. This difference, 0.022, was not statistically significant, $t(16) = -1.48$, $p = .158$. Timecourses from left V1d ($r = 0.477$, $SD = 0.176$) and right V1d ($r = 0.494$, $SD = 0.159$) were similarly correlated across brains. Again, this difference, 0.017, was not statistically significant $t(16) = -1.202$, $p = .247$ (uncorrected).

Timecourses extracted from left hemisphere TO1 ($r = 0.580$, $SD = 0.162$) were more highly correlated across brains than right hemisphere timecourses ($r = 0.548$, $SD = 0.171$). This difference, 0.032, was statistically significant, $t(16) = 2.410$, $p = .028$ (uncorrected). Likewise, timecourses extracted from left hemisphere TO2 ($r = 0.579$, $SD = 0.167$) were more highly correlated across brains than right hemisphere timecourses ($r = 0.516$, $SD = 0.191$). Again, this difference, 0.063, was statistically significant, $t(16) = 3$, $p = .008$ (uncorrected).

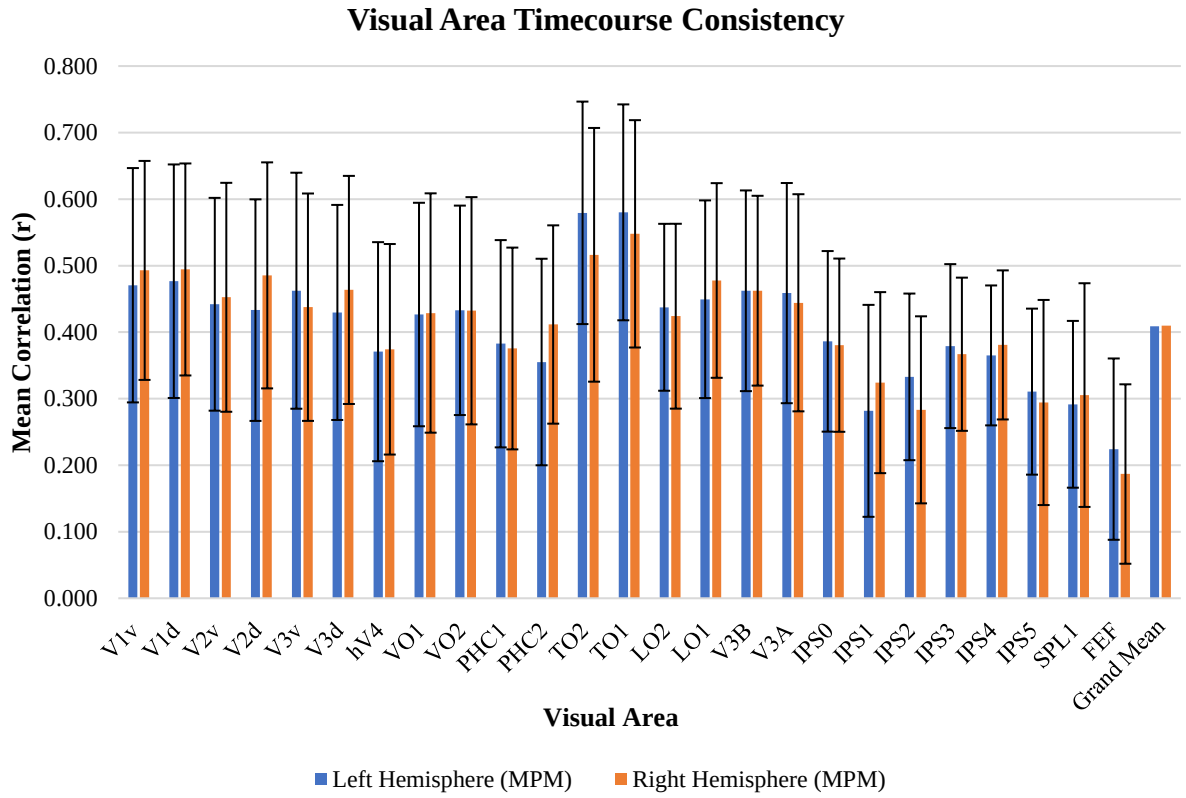

Figure 119. Group mean within-hemisphere correlations for both the left and right hemispheres. Timecourses were extracted from the max probability maps (MPM). Correlations were computed between each individual's visual area timecourse and the group average timecourse for that visual area (the group average timecourse did not included the brain with which the correlation was being computed). The mean of each of these subject-to-group correlations was then calculated. Results show that for TO1 and TO2, left hemisphere timecourses are more highly correlated across brains than right hemisphere timecourses. Conversely, V1v and V1d timecourses are similarly correlated across hemispheres. Error bars represent standard deviation.

As with the average surface analysis, we next examined whether the discrepancy between hemispheres in the timecourses consistency of TO1 and TO2 were significantly different to those exhibited by V1v and V1d. First, we calculated the difference between hemispheres for each of the subject-to-group timecourse correlation. We then performed paired-samples t tests on these data.

Results show that the discrepancy between hemispheres in both TO1 and TO2 were significantly different to the discrepancy observed for V1v and V1d Table 24. When this analysis was performed on the average surface, the V1v – TO1 comparison did not reach significance ($p = 0.061$, uncorrected). However, when performed on the native surface, the comparison was statistically significant ($p = 0.005$, uncorrected).

Table 24. Results of paired-samples t-test performed on the differences between hemispheres in the timecourse consistency for V1v V1d, TO1 and TO2. Analysis was performed on data on native surface. For each subject we calculated the difference between hemispheres for the subject-to-group timecourse correlations. We then performed paired-samples t tests on these differences to ascertain whether the level of timecourse consistency in TO1 and TO2 differed significantly from the level of consistency in V1. p values are uncorrected.

Comparison	Mean	Std. Deviation	t	df	Sig. (2-tailed)
V1v – TO1	-0.085	0.108	-3.265	16	.005
V1v – TO2	-0.055	0.078	-2.89	16	.011
V1d – TO1	-0.081	0.118	-2.818	16	.012
V1d – TO2	-0.05	0.082	-2.5	16	.024

Appendix XII **Sensitivity to choice of smoothing pipeline**

Surface-based smoothing appears to be the most optimal approach to take with respect to reducing the inherent issues of spatial smoothing. However, I sought to investigate how sensitive our results were to the choice of smoothing pipeline. Specifically, whether the stage at which the functional data was smoothed influenced the quality of the data and the subsequent interhemispheric correlation overlays. Data were smoothed with a 5mm Gaussian kernel at one of 5 different stages. As per the standard approach, during surface-based smoothing a cortex mask was applied to the data (Figure 120). This mask ensured only vertices within the cortex were spatially smoothed by excluding the medial wall (mainly the corpus callosum) from the smoothing operation.

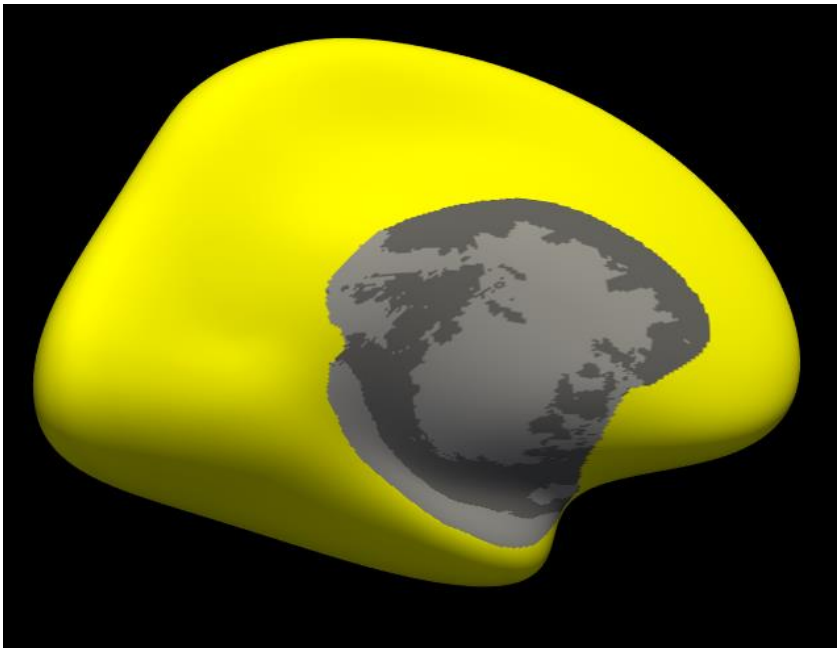

Figure 120. The cortex mask (yellow) applied to functional data during resampling from fsaverage to fsaverage_sym. This mask removes vertices outside of the cortex and smooths the signals of vertices within the cortex ROI.

For approaches 1-4, spatial smoothing within the FEAT GUI was set to 0. For these approaches, data were smoothed with a 5mm full width at half-maximum (FWHM) kernel during resampling either from the volume to the surface or from one surface to another.

1) No smoothing was applied at any stage of pre-processing.
2) Unsmoothed volume data were resampled to the fsaverage surface (mri_vol2surf) and then smoothed during resampling from fsaverage to fsaverage_sym (mri_surf2surf).

3) Unsmoothed volume data were resampled to fsaverage_sym (via fsaverage). Data was then smoothed after resampling to fsaverage_sym (mri_surf2surf).
4) Unsmoothed volume data were smoothed during resampling to fsaverage (mri_vol2surf). These data were then resampled to fsaverage_sym.
5) Raw data were spatially smoothed in the volume during FEAT pre-processing with a 5mm FWHM kernel.

Figure 121 shows the BOLD data resampled to the fsaverage_sym surface following each of the smoothing approaches. I also present the resulting overlays following the cross-hemisphere correlation analysis (described in detail in Chapter 4). Approaches 4 and 5 yield almost identical results. This is to be expected as both smooth the source volume data either before (approach 5) or during (approach 4) resampling to the fsaverage surface. Given the issues raised with volume smoothing, a surface-based smoothing approach is preferable. The functional overlays resulting from approaches 2 and 3 again appear similar. However, approach 2 requires fewer resampling stages and for this reason we believe it is more advantageous than approach 3. For this reason and to minimise the issues raised with regards to smoothing in the volume, approach 2 in which functional data is spatially smoothed during sampling to the average surface is adopted for the analyses presented in Chapter 4.

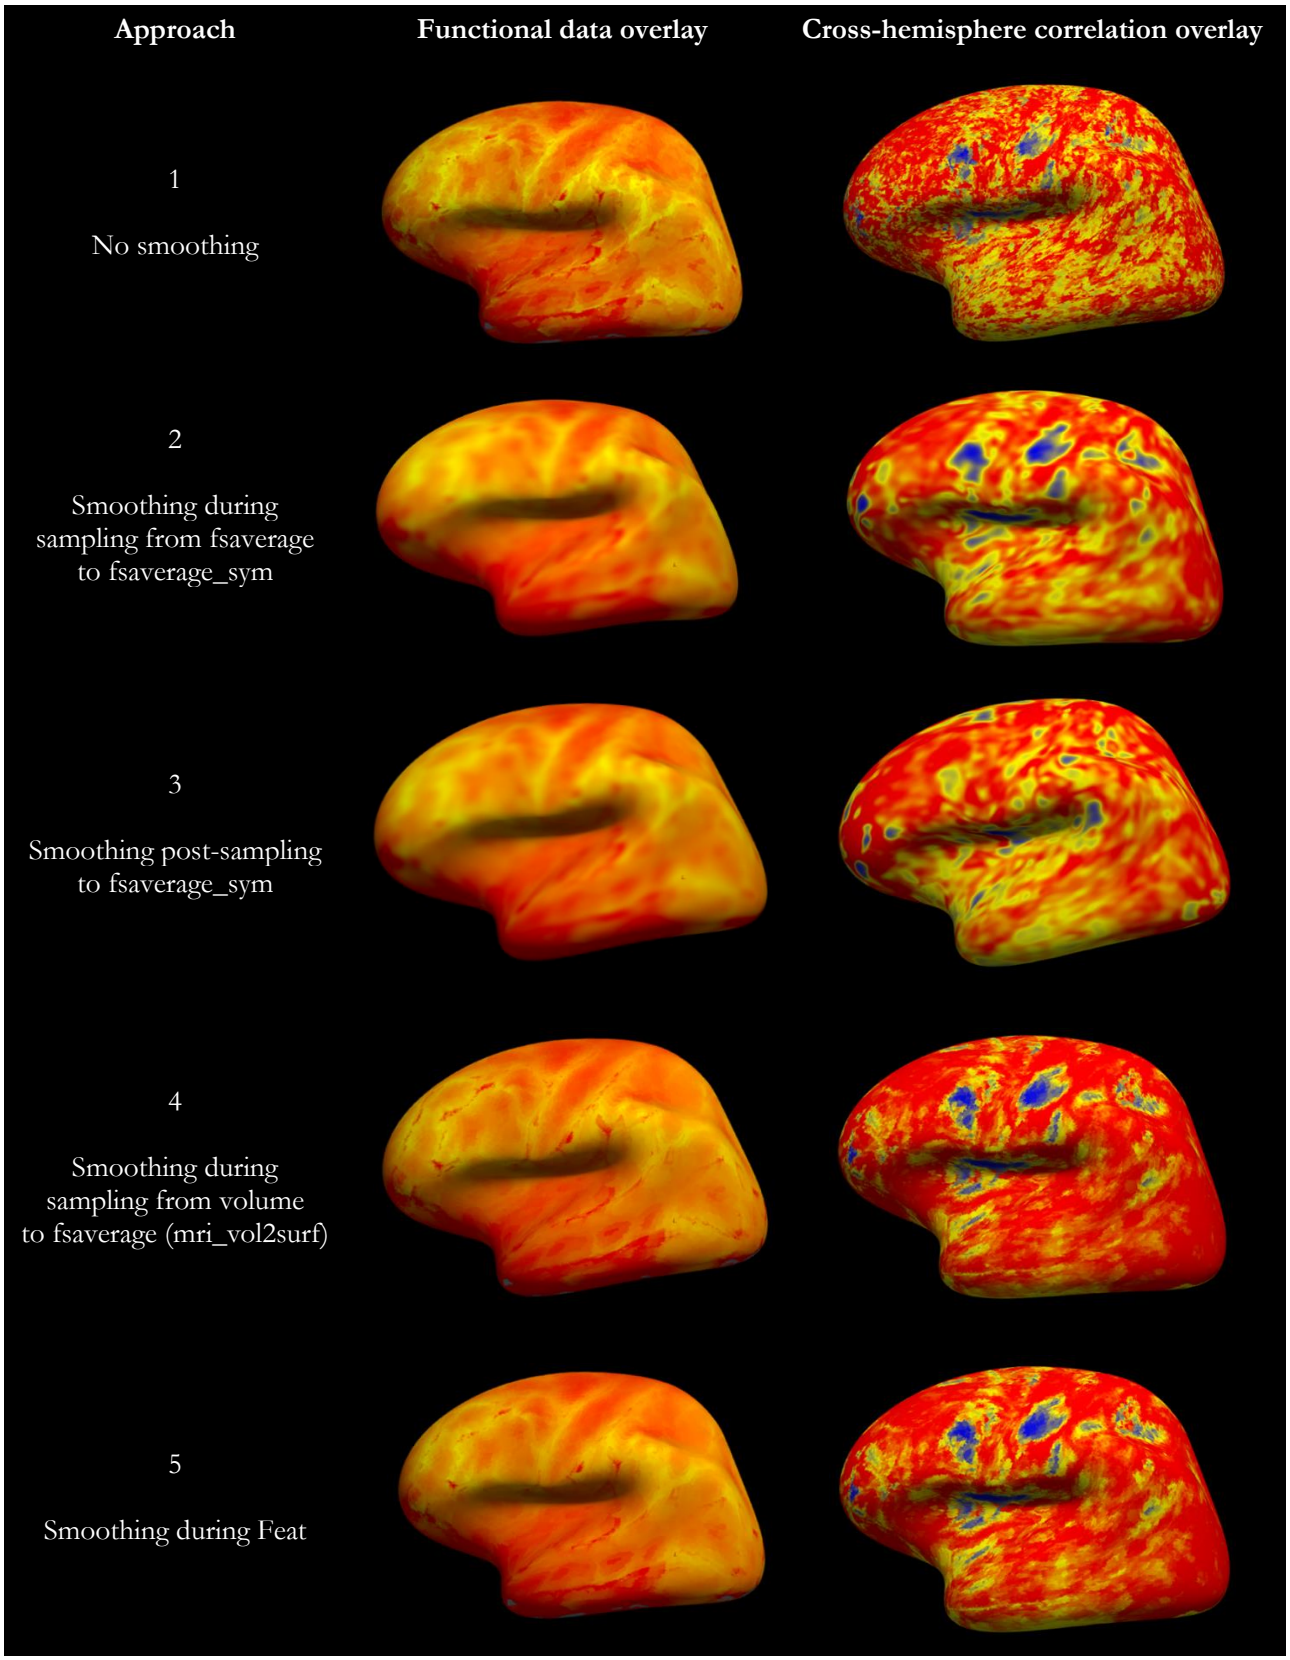

Figure 121. Surface data following resampling and smoothing (left) and subsequent interhemispheric correlation overlays for an example subject.

Appendix XIII **Dorsal and lateral components derived from single-subject ICA**

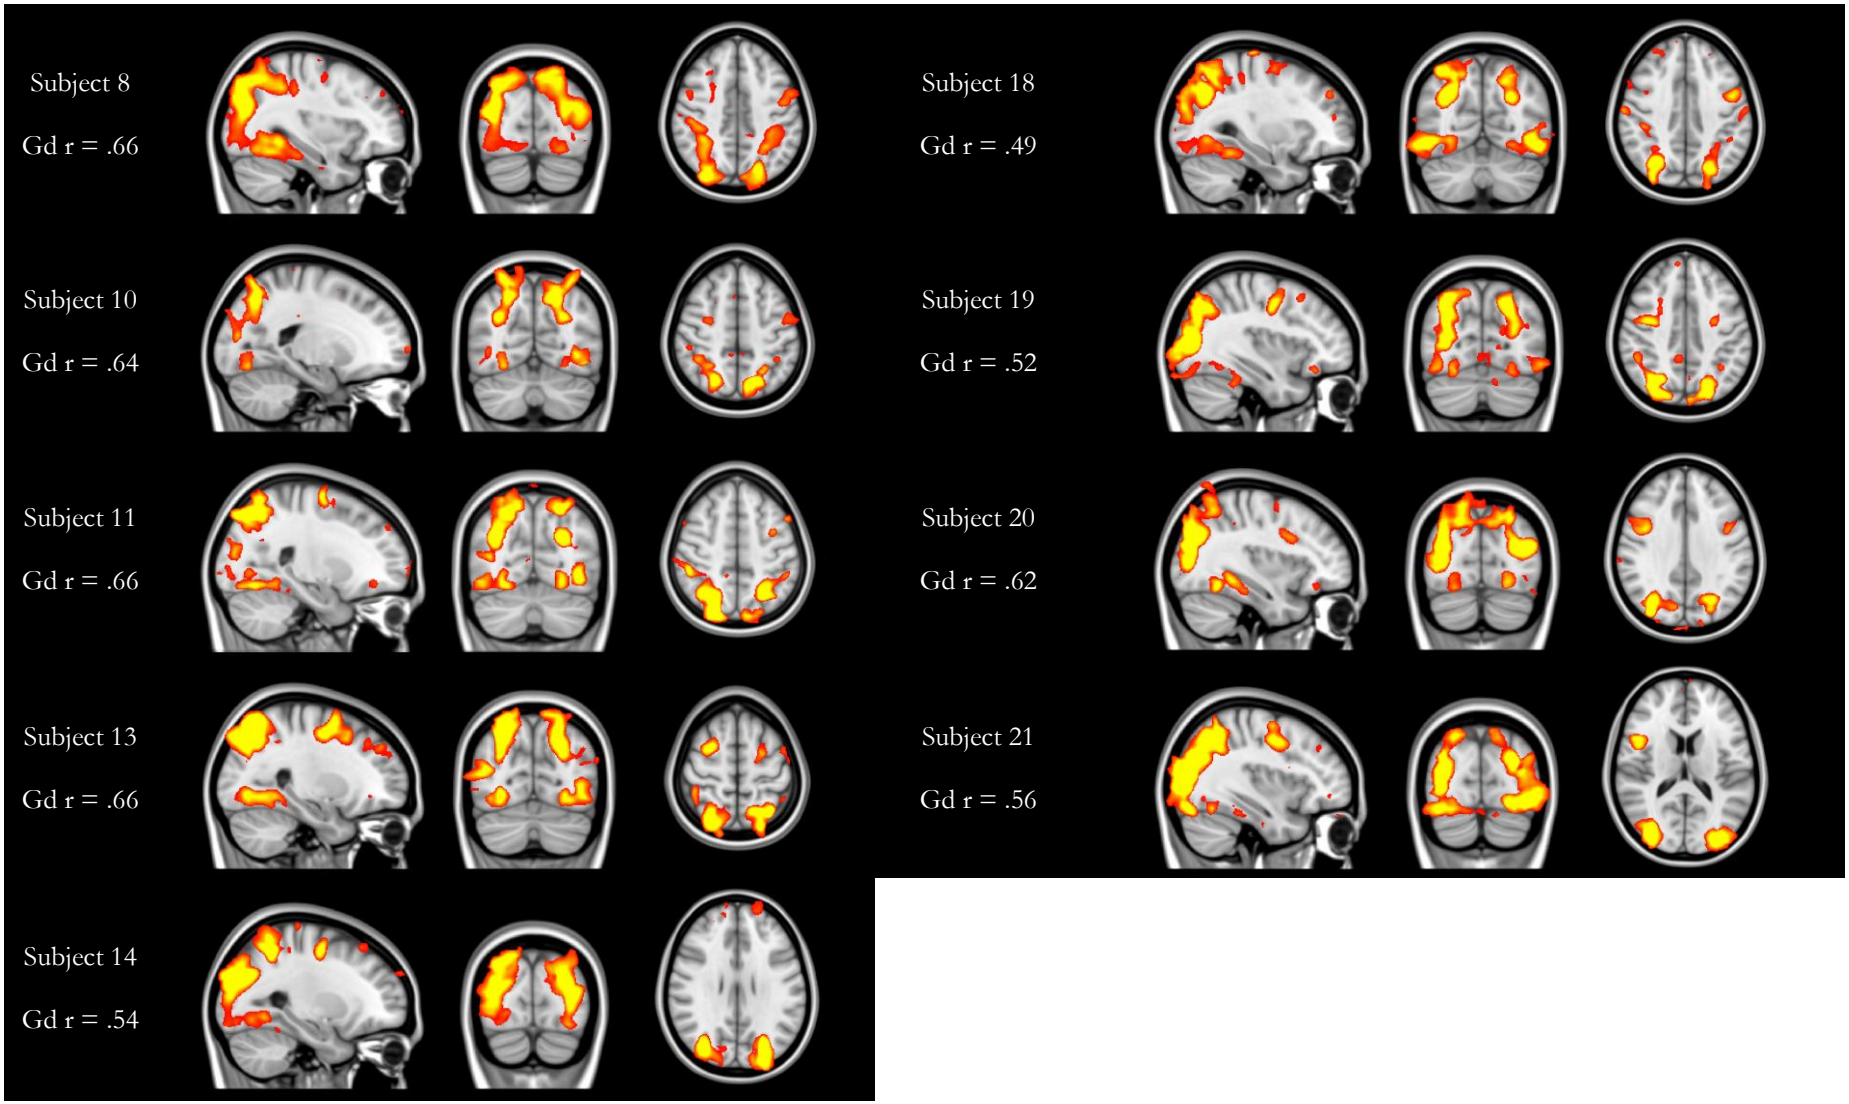

Figure 122. Dorsal components derived from single-subject ICA and the corresponding correlation coefficient (r) with the group-based dorsal component.

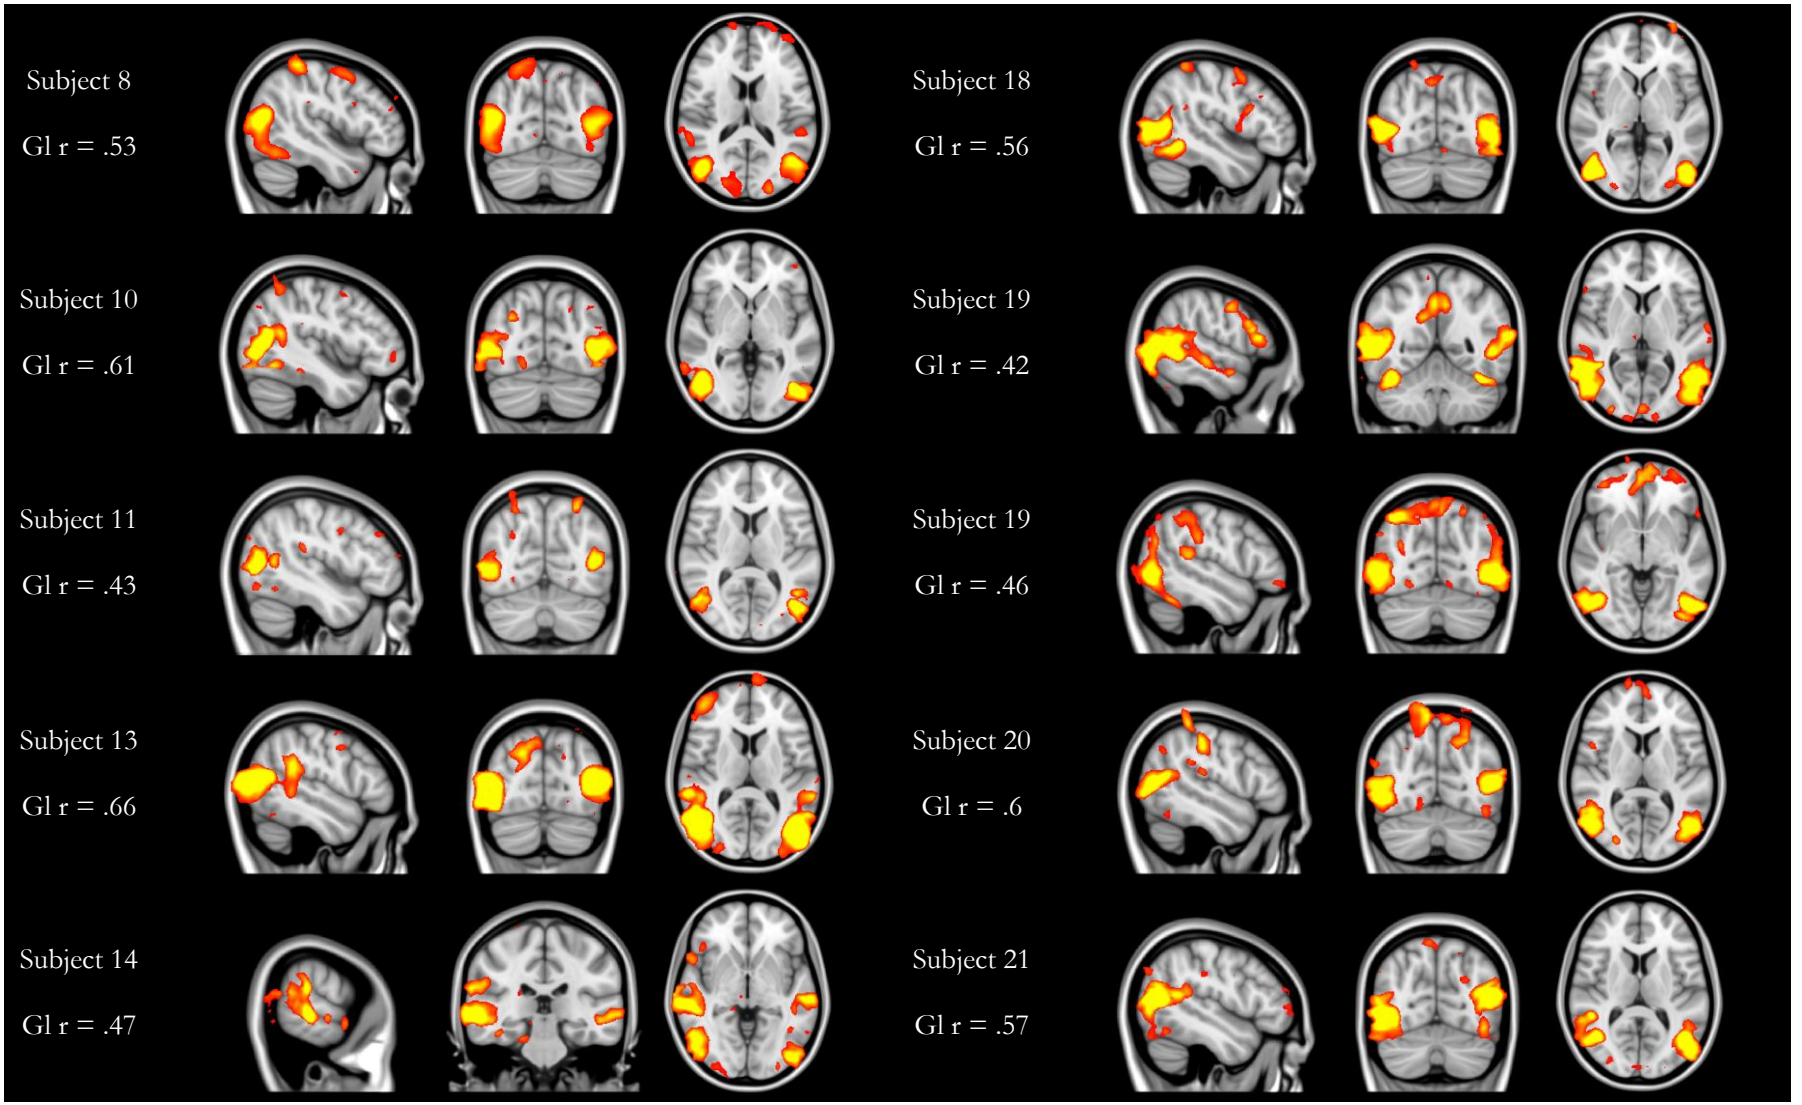

Figure 123. Lateral components derived from single-subject ICA and the corresponding correlation coefficient (r) with the group-based lateral component.

Appendix XIV **Pilot motion psychophysics results**

Subjects

Four healthy-control subjects (one female, three males) with normal/corrected to normal vision were recruited from the Cardiff University Neurological Vision Loss Panel. The average age at the time of participation was 23.75 years (range 20 – 31).

Flow-parsing

The data presented in Table 25 show an effect of optic flow direction (expanding or contracting) and an effect of probe location (2 degrees or 4 degrees). A positive PSE value indicates the probe moved outward the given speed, whereas a negative value indicates the probe moving inward at the given speed.

Table 25. Point of subjective equality values for all subjects and mean values for each condition.

	CONDITION							
	2° expanding		**2° contracting**		**4° expanding**		**4° contracting**	
SUBJECT	Mean probe speed (°/s)	SD (+/-)	Mean probe speed (°/s)	SD (+/-)	Mean probe speed (°/s)	SD (+/-)	Mean probe speed (°/s)	SD (+/-)
S1	0.20	0.04	-0.12	0.04	0.24	0.04	-0.45	0.05
S2	0.20	0.05	0	0.04	0.21	0.04	-0.06	0.04
S3	0.16	0.04	0.01	0.04	0.18	0.04	-0.09	0.04
S4	0.17	0.05	-0.16	0.04	0.21	0.04	-0.13	0.05
MEAN	0.18		-0.07		0.21		-0.18	
SD	0.02		0.09		0.03		0.18	

Expanding optic flow field

The predicted effect of expanding optic flow is observed at both probe locations, in that the average threshold for this condition, is a positive value. These data suggest that the probe needs to be moving in outward in order to be perceived to be stationary. Furthermore, the effect of the optic flow

field is larger when the probe is located at 4 degrees from fixation (M = 0.21, SD = 0.03) compared to 2 degrees (M = 0.18, SD = 0.02).

Contracting optic flow field

The predicted effect of contracting optic flow is observed, again at both probe locations, in that the average threshold for this condition is a negative value. These data suggest that the probe needs to be moving inward in order to be perceived to be stationary. As with the expanding optic flow conditions, the effect of the flow field is larger when the probe is located at 4 degrees (M = -0.18, SD = 0.18) compared to at 2 degrees (M = -0.07, SD = 0.09).

Magnitude of flow parsing effect

The key measure when analysing optic flow-parsing is the magnitude of the effect. As an expanding flow field induces an inward component of motion to the probe and a contracting flow field induces an outward component of motion the probe, the magnitude of the effect can be calculated by taking the difference between the thresholds computed for each condition (Table 26). This not only allows for a useful assessment of the effect, but also avoids any response or perceptual biases present in the observer. The data presented in Table 26 replicate previous findings and, importantly, the flow-parsing effect is also found at the individual level making the task appropriate for patient research.

Table 26. Magnitude of flow parsing effect, calculated by the taking difference in threshold values between expanding and contracting conditions.

	Condition	
	2°	**4°**
	Magnitude of flow parsing (°/s)	Magnitude of flow parsing (°/s)
C1	0.32	0.69
C2	0.20	0.27
C3	0.15	0.17
C4	0.33	0.34
Mean	0.25	0.38
SD	0.09	0.28

Fisher's Z transformations appendices

The following appendices (Appendix A – S) present the Fisher's Z transformed data. See section 2.1.6 for details as to why this transform was implemented. Although values are reported in Z' throughout the following section, to remain consistent with the main body of the thesis I refer to correlations. This is because each correlation coefficient (r) was transformed to a Fisher's Z value (Z') and then operations were performed on these values as if it were correlation coefficients.

Appendix A **Chapter 3: Fisher's Z transform of Figure 18**

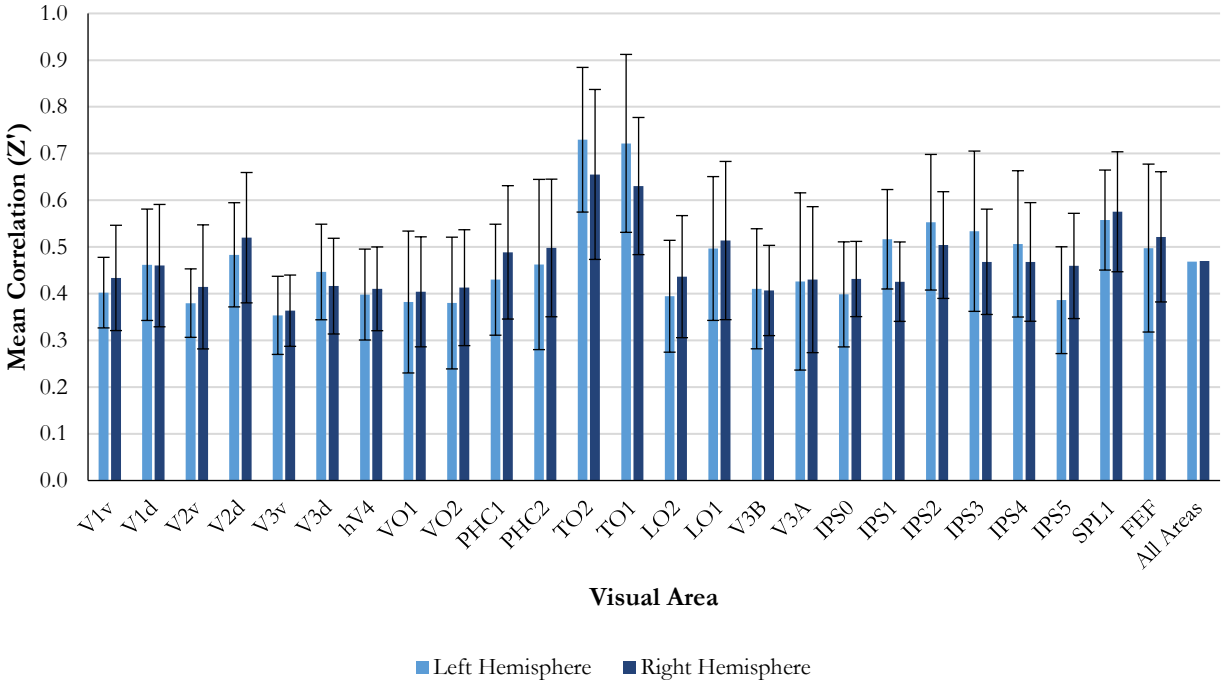

Figure 124. Mean correlation between the "top matched" component and the atlas regions. Values are presented following transformation from Pearson's r to Fisher's Z. Error bars represent standard deviations.

Appendix B **Chapter 3: Fisher's Z transform of Table 6**

Table 27. Summary of spatial correlations. For each visual area and each hemisphere, I calculated the mean correlation across brains (Transformed from Pearson's r to Fisher's Z) between the top matched component and atlas region (mean correlation). The per cent of unique matches (how often the component was the top match to only one visual area) is also presented for each visual area and each hemisphere. Top match bilateral (percent of all brains) reflects the frequency with which the top match in one hemisphere was also the top match to the same region in the opposite hemisphere. Blue cells indicate a higher value, e.g., higher correlation between the component and atlas, higher percentage of unique matches across subjects and higher percentage of bilateral matches across subjects. Yellow cells indicate the reverse.

	Left Hemisphere		Right Hemisphere		Top matches bilateral (%)
	Top match unique (%)	Mean Correlation (Z')	Right top match unique (%)	Mean Correlation (Z')	
V1v	47	0.402	47	0.434	71
V1d	59	0.462	18	0.460	47
V2v	59	0.380	47	0.415	35
V2d	29	0.483	18	0.520	65
V3v	29	0.354	41	0.364	24
V3d	29	0.447	24	0.416	35
hV4	41	0.398	59	0.410	35
VO1	18	0.382	29	0.404	53
VO2	24	0.380	6	0.413	53
PHC1	18	0.430	24	0.488	35
PHC2	41	0.462	29	0.498	53
TO2	12	0.730	6	0.655	76
TO1	6	0.722	0	0.630	71
LO2	24	0.395	24	0.437	65
LO1	29	0.497	6	0.514	53
V3B	12	0.411	18	0.407	71
V3A	53	0.426	47	0.430	18
IPS0	65	0.399	24	0.431	41
IPS1	35	0.517	29	0.426	53
IPS2	35	0.553	35	0.504	41
IPS3	12	0.534	29	0.468	24
IPS4	0	0.507	24	0.468	29
IPS5	35	0.386	47	0.459	24
SPL1	88	0.558	71	0.575	47
FEF	94	0.498	94	0.522	35

Appendix C **Chapter 3: Fisher's Z transform on Figure 39**

The mean across all observers for left hemisphere correlations (Z' = 1.246) is comparable to the right hemisphere correlations (Z' = 1.235). In the left hemisphere, correlations are highest for TO1 (mean Z' = 1.759, SD = 0.398), TO2 (mean Z' = 1.543, SD = 0.303) and V1v (mean Z' = 1.559, SD = 0.494). The atlas and component timecourses of VO2 (mean Z' = 0.927, SD = 0.649), FEF (mean Z' = 0.925, SD = 0.506), and PHC1 (mean Z' = 0.895, SD = 0.504) show the lowest correspondence. In the right hemisphere, correlations are highest for TO1 (mean Z' = 1.543, SD = 0.365), V1v (mean Z' = 1.553, SD = 0.462) and V2d (mean Z' = 1.481, SD = 0.465). The atlas and component timecourses of PHC1 (mean Z' = 0.898, SD = 0.516), VO2 (mean Z' = 0.995, SD = 0.524) and FEF (mean r = 0.964, SD = 0.453) show the lowest correspondence.

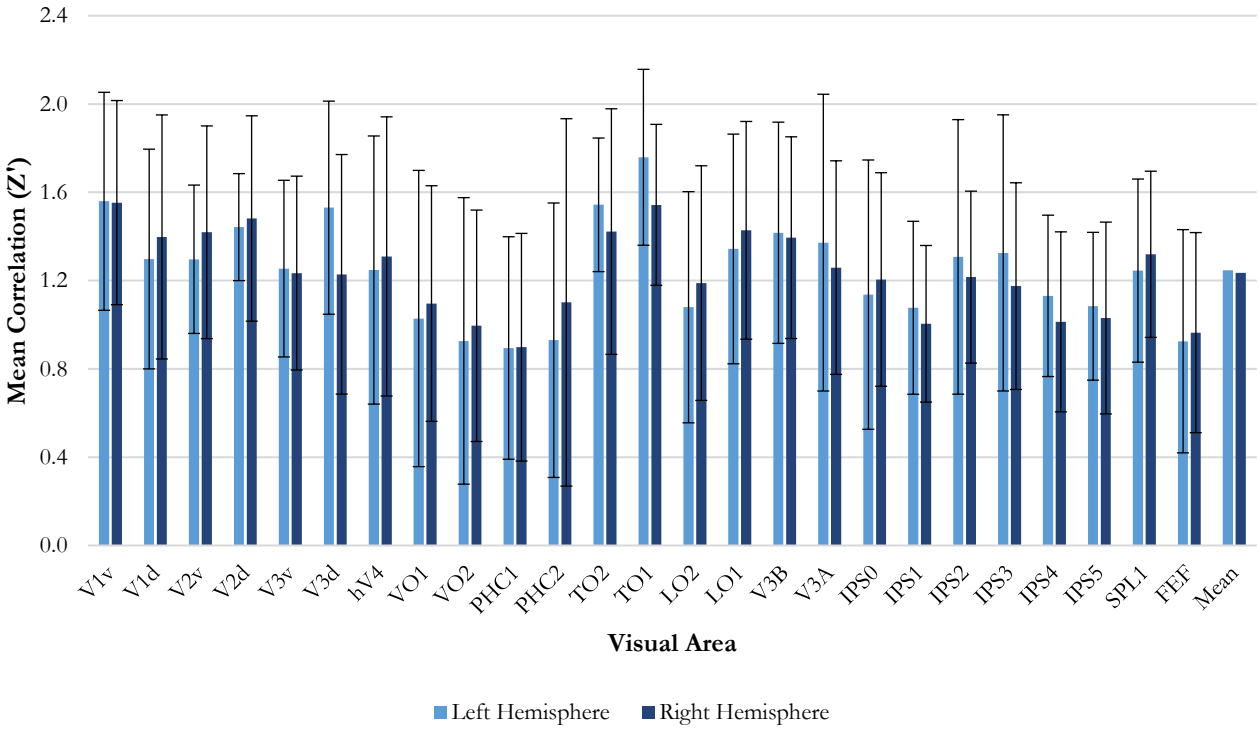

Figure 125. Group mean correlation between the timecourses of each atlas region and the component showing the highest correlation to that region. Correlations were computed for each hemisphere separately and then converted from Pearson's r to Fisher's Z. Overall, correlations between atlas and component timecourses are high, indicating the peaks of the components show good spatial correspondence with the peaks of the probabilistic atlas regions. Error bars represent standard deviation.

Appendix D **Chapter 3: Fisher's Z transform of Table 7**

Table 28. Atlas and component timecourse consistency, inferred from the group mean pairwise timecourse correlations. For each subject we computed the correlation between the timecourse extracted from the component and the timecourse extracted from all other subject's corresponding components. These correlations were computed for each visual area and hemisphere separately, and then transformed from Pearson's r to Fisher's Z. We then took the mean across all areas and all brains for each hemisphere (mean timecourse correlation). Next, we took the mean across both hemispheres (grand mean correlation). We then repeated this process for atlas timecourses. Results show timecourses from both the components and the atlas are more consistent in the right hemisphere than the left, and that atlas timecourses are more consistent than component timecourses

ROI	Hemisphere	Mean Timecourse Correlation (Z)	Grand Mean Correlation (Z)
Component	Left	0.169	0.172
	Right	0.175	
Atlas	Left	0.2	0.202
	Right	0.205	

Appendix E **Chapter 3: Fisher's Z transform of Figure 40**

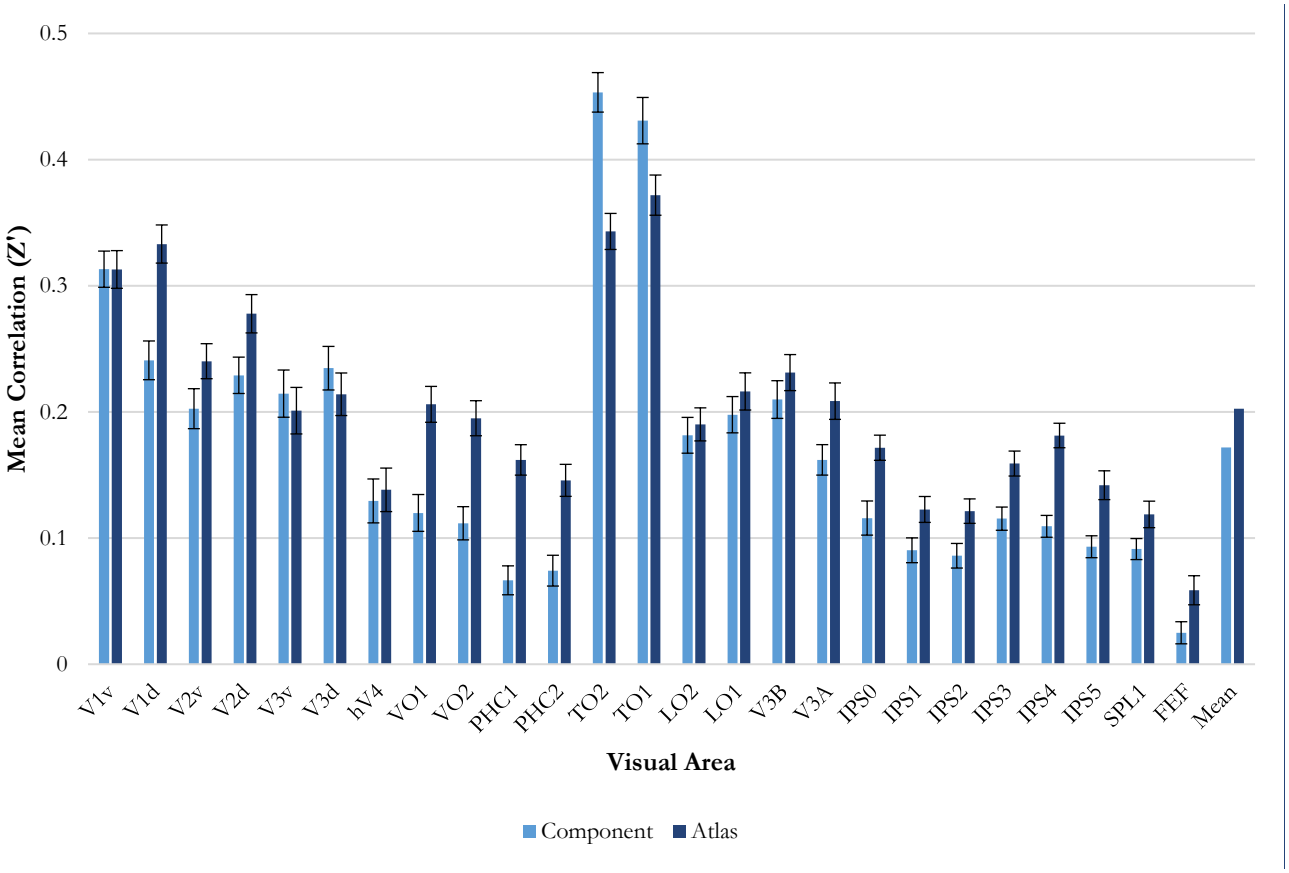

Figure 126. Group mean bilateral timecourse correlations for each visual area. For each subject and each visual area, we computed the correlation between the timecourse extracted from the component and the timecourse extracted from all other subject's components. These correlations were computed for each hemisphere separately and transformed from Pearson's r to Fisher's Z before the mean across left and right hemisphere correlations was taken (bilateral mean). We then repeated this process for atlas timecourses. Results show the highest correlations are observed for TO1/MT and TO2/MST regardless of whether the timecourses were extracted from the atlas or components.

Appendix F **Chapter 3: Fisher's Z transform of component vs. atlas timecourse correlation analysis (section 3.3.2.3)**

Table 29. Atlas and component timecourse consistency, inferred from the group mean pairwise timecourse correlations. For each subject we computed the correlation between the timecourse extracted from the component and the timecourse extracted from all other subject's corresponding components. These correlations were computed for each visual area and hemisphere separately, and then transformed from Pearson's r to Fisher's Z. For each area, we then took the mean across brains and hemispheres (mean timecourse correlation). We then repeated this process for atlas timecourses.

ROI	Mean Correlation (Z')	Std. Deviation
PHC1 component	.067	.134
PHC1 atlas	.162	.141
PHC2 component	.074	.142
PHC2 atlas	.146	.148
VO1 component	.120	.170
VO1 atlas	.206	.166
VO2 component	.112	.154
VO2 atlas	.195	.163
V1d component	.241	.182
V1d atlas	.333	.176
TO1 component	.431	.214
TO1 atlas	.372	.186
TO2 component	.453	.183
TO2 atlas	.343	.169

Table 30. Results of paired-samples t-test performed on the atlas and component timecourse consistency measures. Results show the consistency of atlas ATCs is significantly different to the consistency of component ATCs.

Component Vs. Atlas	Mean	Std. Deviation	t	df	p (2-tailed)
PHC1	-.095	.148	-10.619	271	<.001
PHC2	-.072	.151	-7.798	271	<.001
VO1	-.086	.183	-7.764	271	<.001
VO2	-.083	.193	-7.125	271	<.001
V1d	-.092	.113	-13.458	271	<.001
TO1	.059	.095	10.212	271	<.001
TO2	.110	.110	16.537	271	<.001

Appendix G **Chapter 4: Fisher's Z transform of Figure 58**

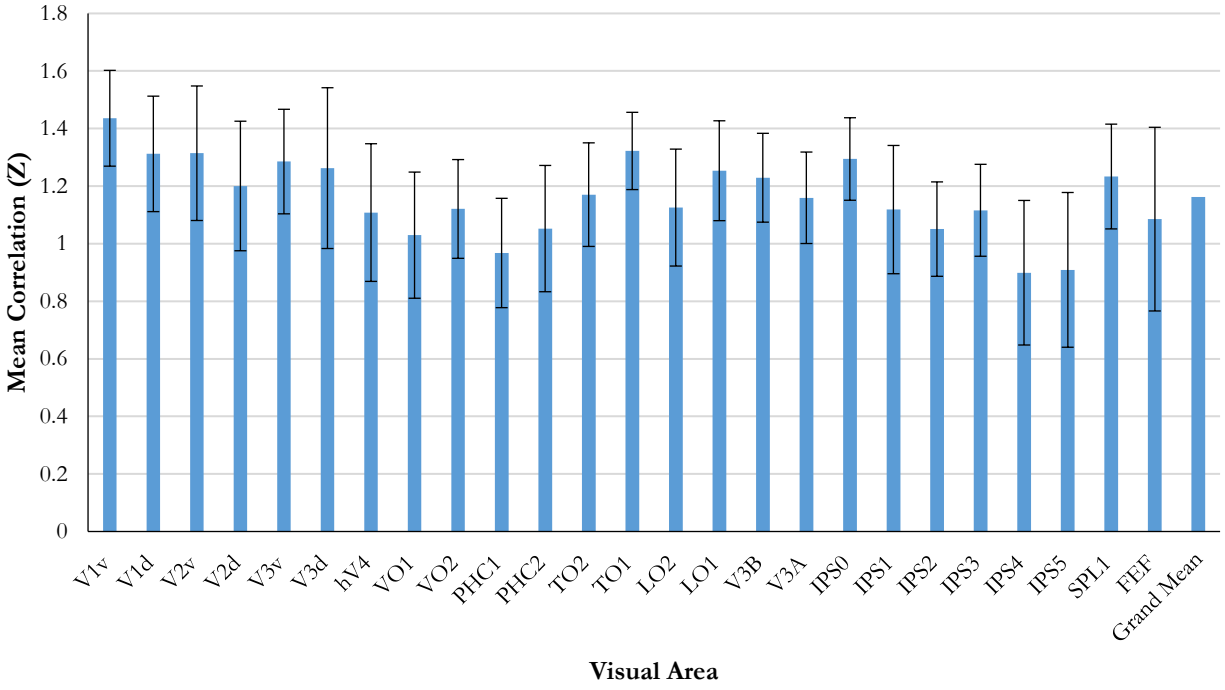

Figure 127. Group mean ROI-to-ROI timecourse correlation coefficients for each of the 25 visual areas. Values are presented following transformation from Pearson's r to Fisher's Z. Error bars represent standard deviation. Results show the timecourses of the visual areas are highly correlated across hemispheres. TO2/MST is less correlated than TO1/MT, V1v and V1d, suggesting there may be differences in function across hemispheres for TO2/MST.

Appendix H **Chapter 4: Fisher's Z transform of section 4.3.2 results (ROI-to-ROI)**

Table 31. Mean cross-hemisphere timecourse correlations of V1, TO1/MT and TO2/MST. For each subject, corelations were performed between the mean timecourse of an ROI in one hemisphere and the mean timecourse of the same ROI in the opposite hemisphere. Prior to averaging across brains, individual correlation coefficients were transformed from Pearson's r to Fisher's Z.

ROI	Mean Correlation (Z)	Std. Deviation
V1	1.374	.168
TO1/MT	1.322	.134
TO2/MST	1.170	.180

Table 32. Results of paired-samples t-test performed on the cross-hemisphere timecourse correlations (ROI-analysis) of V1 and TO2/TO2 (all individual correlation coefficients were first transformed from Pearson's r to Fisher's Z).

Comparison	Mean	Std. Deviation	t	df	p (2-tailed)
V1 - TO1	0.052	0.143	1.493	16	.155
V1 - TO2	0.204	0.096	3.092	16	.007

Appendix I **Chapter 4: Fisher's Z transform of Figure 59**

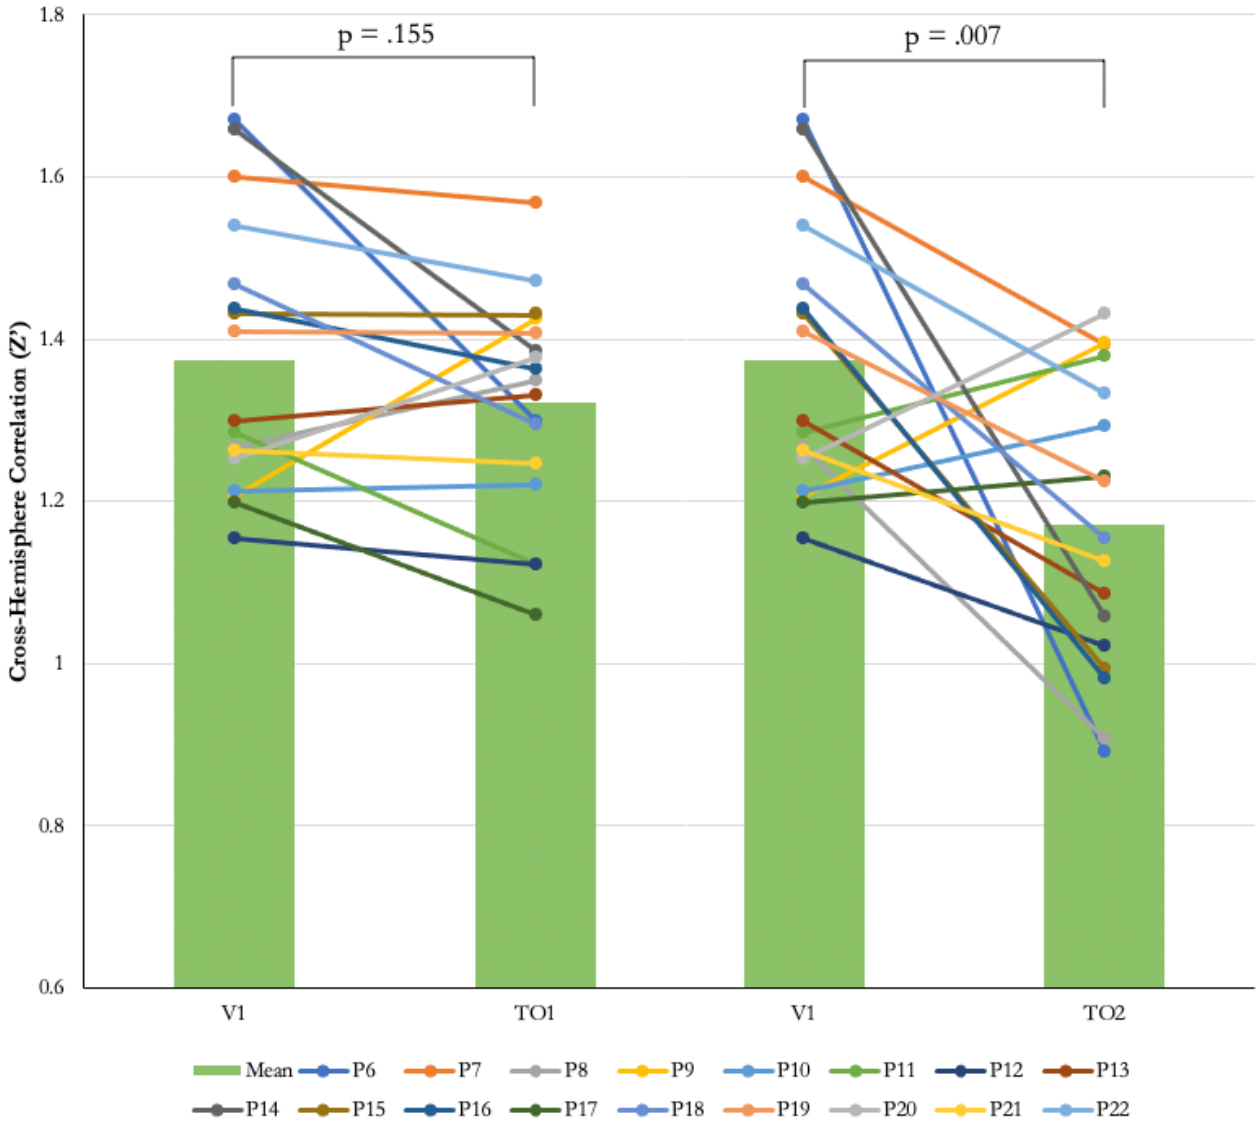

Figure 128. Individual subject cross-hemisphere timecourse correlations presented in paired format for V1 and TO1/TO2. Individual subject data were transformed from Pearson's r to Fisher's Z (Z'). Group means are represented by bars. Data shows for both TO1/MT and TO2/MST there is a depression in the mean correlation relative to V1, with clear inter-subject variability (i.e., for some brains the timecourses of V1 are more highly correlated across hemispheres than TO1/TO2 timecourses, for others we find the reverse).

Appendix J **Chapter 4: Fisher's Z transform of Table 13**

Table 33. The group mean cross-hemisphere correlations computed between ROIs in opposite hemispheres. Values were transformed from Pearson's r to Fisher's Z. The mean of the left versus right and right versus left comparisons are displayed here. Results are shown for data pre-processed using the standard approach (raw) and data denoised with ICA-AROMA (AROMA). The overall pattern of the results does not change irrespective of the approach adopted.

	Raw		AROMA	
	Mean	Std Dev	Mean	Std Dev
V1v vs. V1v	1.426	0.174	1.218	0.254
V1v vs. combined visual	1.135	0.179	0.996	0.231
V1v vs. outside visual	0.603	0.145	0.519	0.142
V1v vs. BA	0.330	0.143	0.264	0.141
V1v vs. auditory	0.345	0.141	0.265	0.115
Combined visual vs. BA	0.428	0.164	0.344	0.154
Combined visual vs. auditory	0.365	0.165	0.285	0.129
BA vs. auditory	0.381	0.200	0.279	0.132

Appendix K **Chapter 4: Fisher's Z transform of Figure 61**

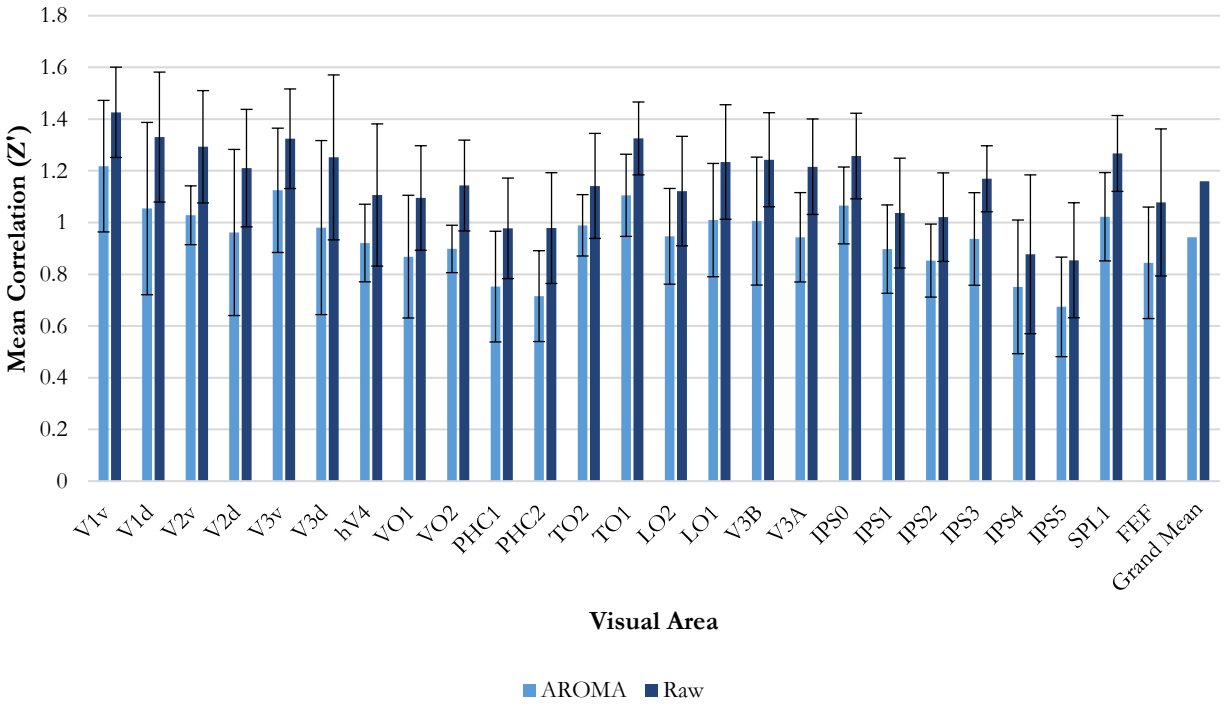

Figure 129. Group mean ROI-to-ROI timecourse correlation coefficients for each of the 25 visual areas. Values are presented following transformation from Pearson's r to Fisher's Z. Error bars represent standard deviation. Results are shown for data pre-processed using the standard approach (raw) and data denoised with ICA-AROMA (AROMA). Correlations are lower for denoised data than raw data, but crucially the pattern of the results did not change irrespective of the approach used.

Appendix L **Chapter 4: Fisher's Z transform of Figure 63**

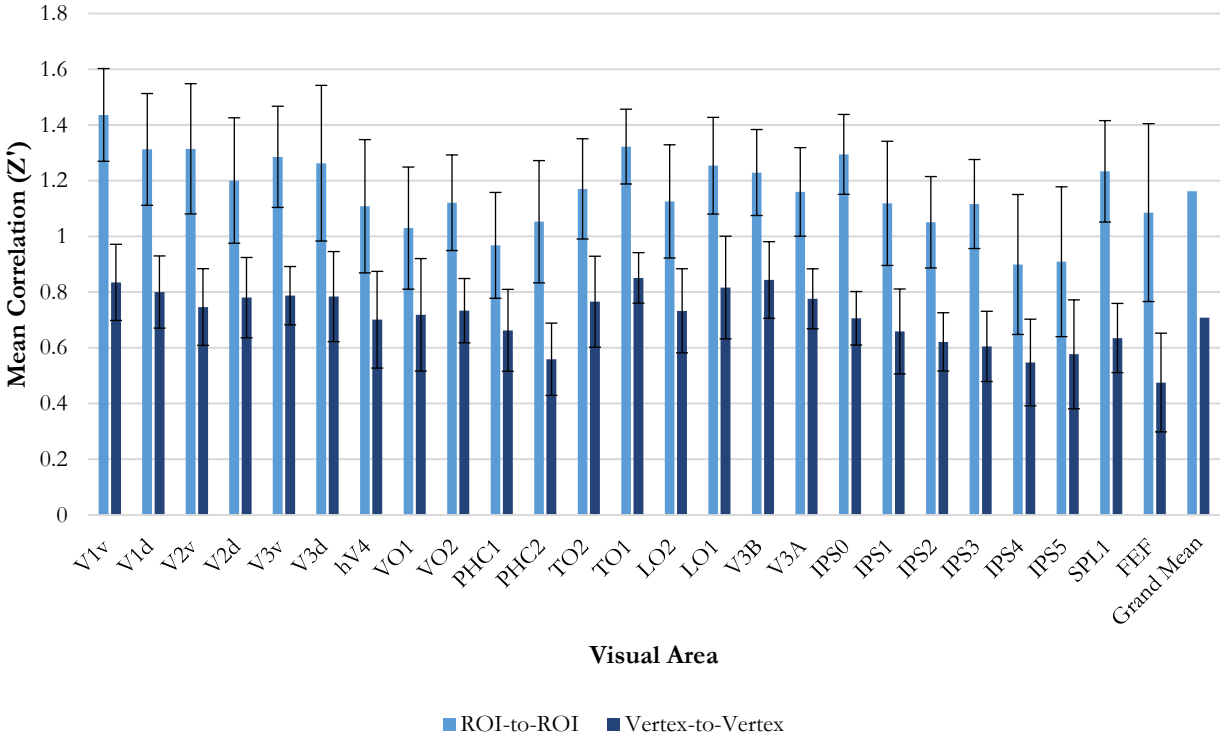

Figure 130. Group mean cross-hemisphere correlations for each visual area. Results are displayed for both the ROI-to-ROI and vertex-to-vertex analyses. For the ROI analysis the mean timecourse of vertices within a given ROI in each hemisphere is extracted. Correlations are then computed between the timecourses from opposite hemispheres. For the vertex analysis, correlations are computed between corresponding vertices in each hemisphere. The mean coefficient is then extracted from vertices within each ROI. For both approaches, correlation coefficients were transformed from Pearson's r to Fisher's Z before the mean across brains was computed. Results show the timecourses of TO2/MST are less correlated across hemispheres than V1v, V1d and TO1/MT. Error bars represent standard deviation.

Appendix M **Chapter 4: Fisher's Z transform of section 4.4.2 results (Vertex-to-Vertex)**

Table 34. Mean cross-hemisphere timecourse correlations of V1, TO1/MT and TO2/MST. Timecourse correlations were performed on a per-vertex basis (e.g., vertex 1 in the left hemisphere vs. vertex 1 in the right hemisphere), before the mean correlation coefficient was extracted from each ROI. Prior to averaging across brains, individual correlation coefficients were transformed from Pearson's r to Fisher's Z.

ROI	Mean Correlation (Z')	Std. Deviation
V1	.817	.109
TO1/MT	.851	.091
TO2/MST	.765	.163

Table 35. Results of paired-samples t-test performed on the cross-hemisphere timecourse correlations (vertex-analysis) of V1 and TO1/TO2 (all individual correlation coefficients were first transformed from Pearson's r to Fisher's Z).

Comparison	Mean	Std. Deviation	t	df	p (2-tailed)
V1 - TO1	-0.033	0.114	-1.205	16	.246
V1 - TO2	0.052	0.214	1.003	16	.331

Appendix N **Chapter 4: Fisher's Z transform of Figure 64**

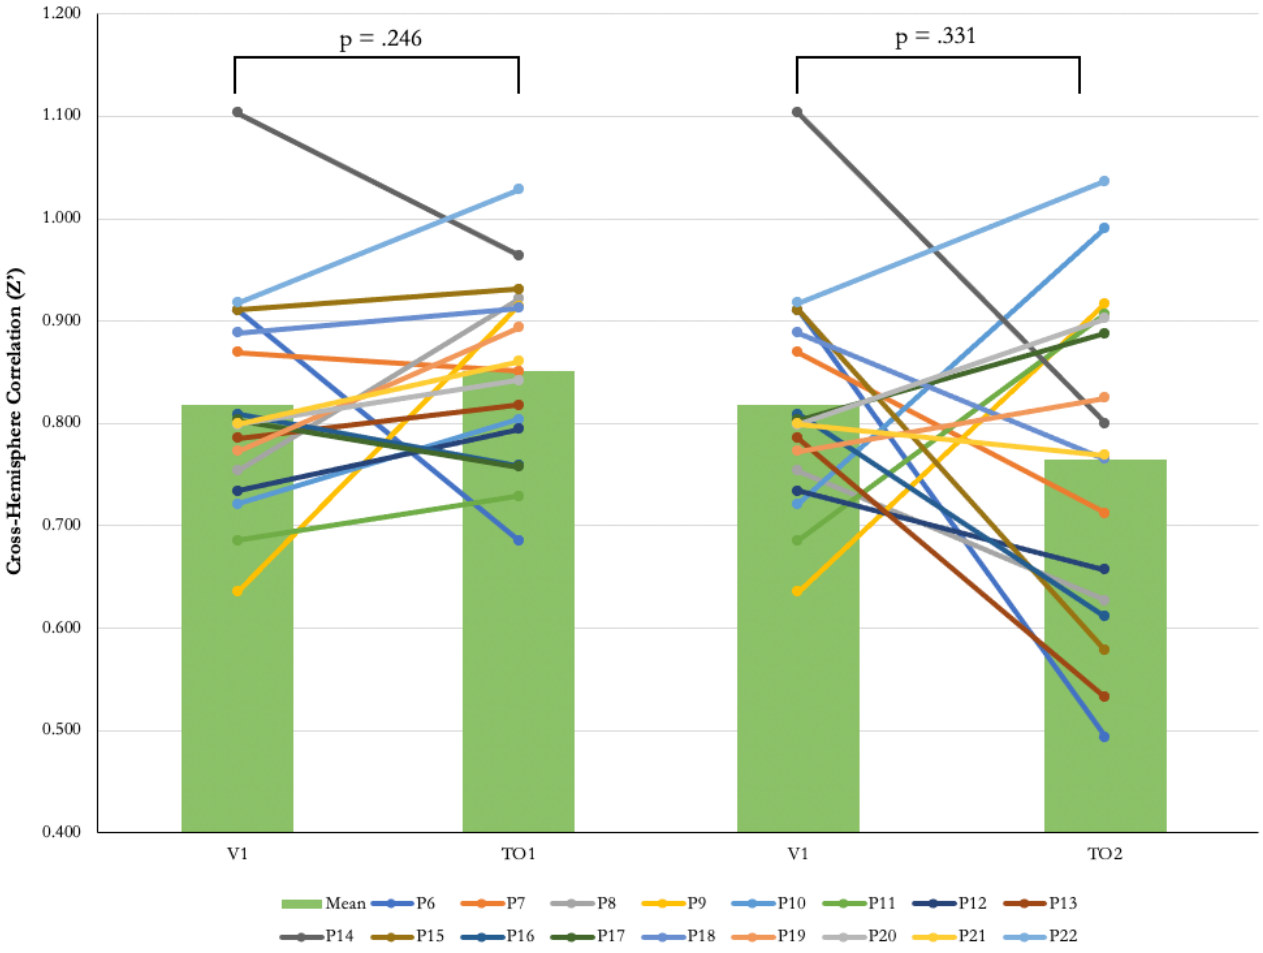

Figure 131. Individual subject cross-hemisphere timecourse correlations presented in paired format for V1 and TO1/TO2. Individual subject data were transformed from Pearson's r to Fisher's Z (Z'). Group means are represented by bars. Data shows, on average, TO1/MT timecourses are more highly correlated across hemispheres than V1 timecourses. For TO2/MST, we find a depression in the mean correlation relative to V1. Both V1 to TO1/MT and V1 to TO2/MST comparisons show clear inter-subject variability (i.e., for some brains the timecourses of V1 are more highly correlated across hemispheres than TO1/TO2 timecourses, for others we find the reverse).

Appendix O **Chapter 4: Fisher's Z transform of Figure 65**

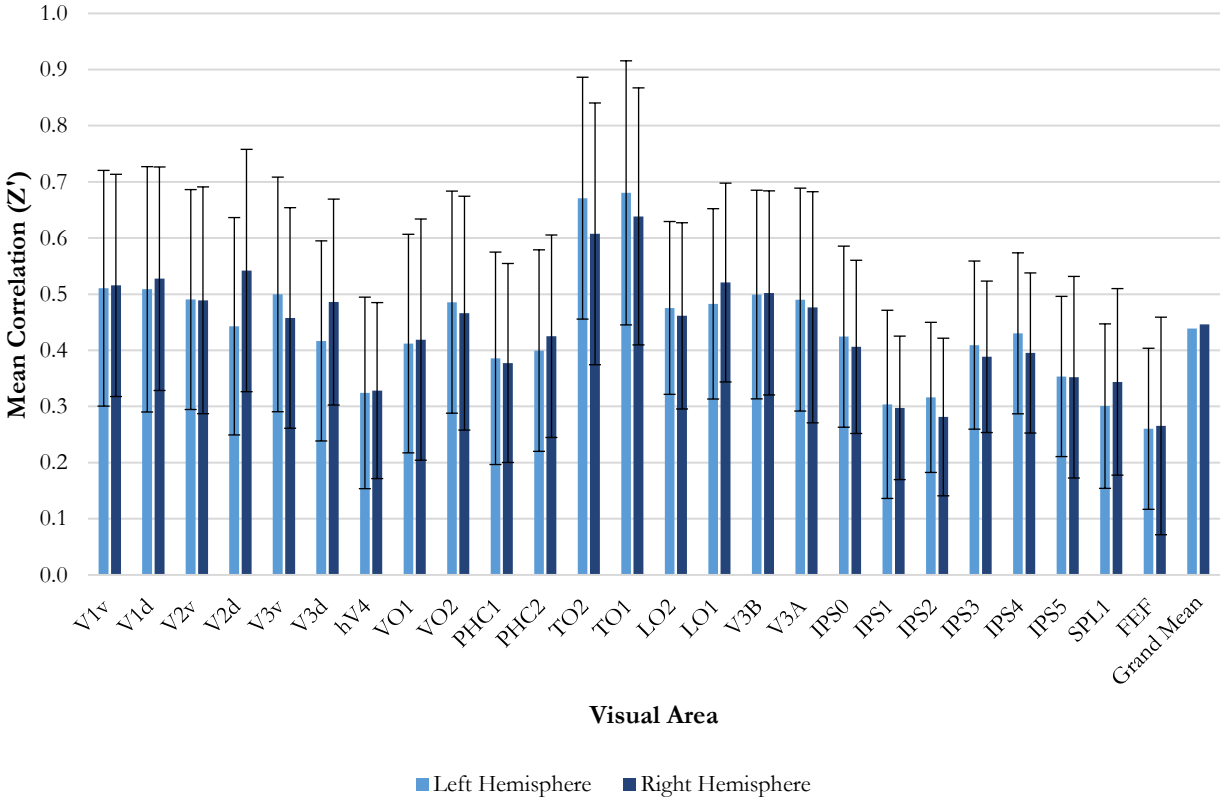

Figure 132. Group mean between-brains correlations for both the left and right hemispheres. Correlations were computed between each individual's visual area timecourse and the group average timecourse for that visual area (the group average timecourse did not include the brain with which the correlation was being computed). Values were then transformed from Pearson's r to Fisher's Z, before the mean of each of these subject- to-group correlations was calculated. This analysis gives a measure of the consistency of the timecourses for each visual area. Results show that for TO1/MT and TO2/MST, left hemisphere timecourses are more highly correlated across brains than right hemisphere timecourses. Conversely, V1v and V1d timecourses are similarly correlated across hemispheres. Error bars represent standard deviation.

Appendix P **Chapter 4: Fisher's Z transform of Figure 66**

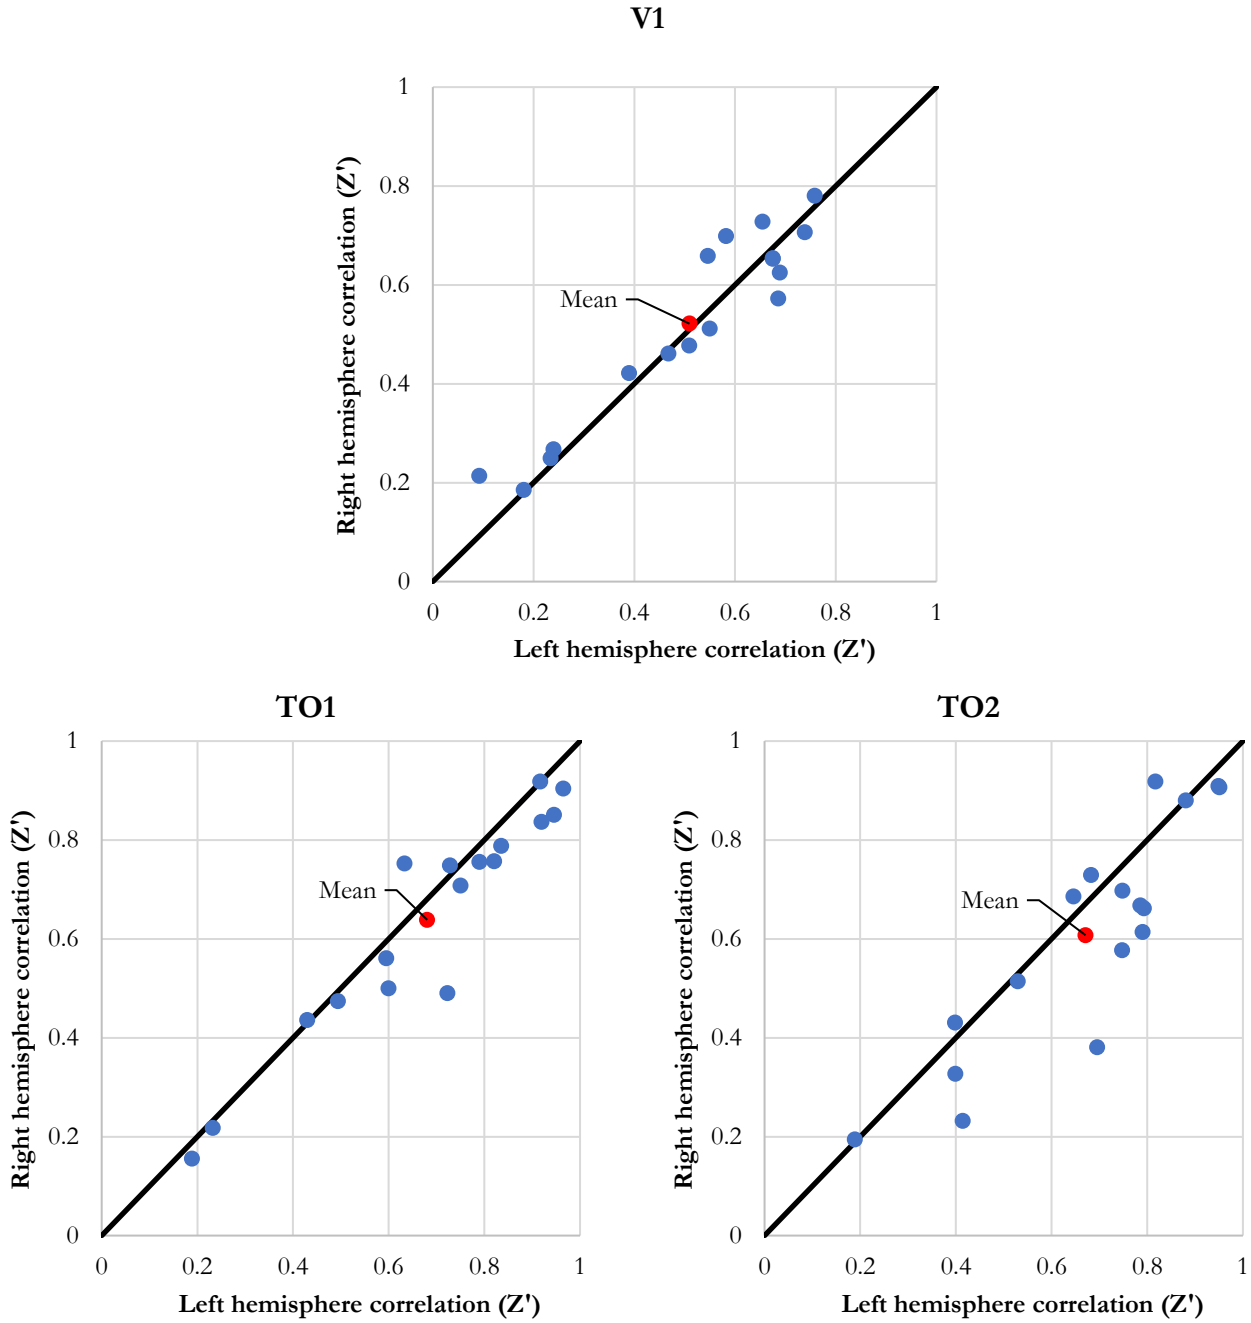

Figure 133. Individual subject timecourse consistency of the left (x-axis) and right (y-axis) hemisphere. Correlations were computed between each individual's visual area timecourse and the group average timecourse for that visual area (the group average timecourse did not include the brain with which the correlation was being computed). This analysis gives a measure of the consistency of the timecourses for each visual area. Individual correlations were then transformed from Pearson's r to Fisher's Z (Z'). Group means represented by red markers. Black line represents the line of equality. Results show that for TO regions, in particular TO2/MST, left hemisphere timecourses are more highly correlated across brains than right hemisphere timecourses (more data points below the line of equality). Conversely, V1 timecourses are similarly correlated across hemispheres (data points are distributed more symmetrically relative to the line of equality). Fisher's Z equivalent figure can be found in Appendix P.

Appendix Q **Chapter 4: Fisher's Z transform of section 4.5.2 results (timecourse consistency)**

Table 36. Group mean between-brains correlations for both the left and right hemispheres. Correlations were computed between each individual's visual area timecourse and the group average timecourse for that visual area (the group average timecourse did not include the brain with which the correlation was being computed). Values were then transformed from Pearson's r to Fisher's Z, before the mean of each of these subject- to-group correlations was calculated for each ROI and each hemisphere.

ROI	Mean Correlation (Z')	Std. Deviation
V1 left hemisphere	.510	.211
V1 right hemisphere	.522	.194
TO1 left hemisphere	.680	.235
TO1 right hemisphere	.639	.229
TO2 left hemisphere	.671	.215
TO2 right hemisphere	.607	.233

Table 37. Results of paired-samples t-test performed on the differences between hemispheres in the timecourse consistency for V1, TO1 and TO2 (all individual correlation coefficients were first transformed from Pearson's r to Fisher's Z). There was no significant difference between hemispheres for V1. However, for TO1 and TO2, the difference between hemispheres was statistically significant.

Left hemisphere Vs. Right Hemisphere	Mean	Std. Deviation	t	df	p (2-tailed)
V1	-.012	.065	-.754	16	.462
TO1	.042	.071	2.437	16	.027
TO2	.064	.106	2.471	16	.025

Appendix R **Chapter 4: Fisher's Z transform of post-hoc exploratory comparisons (section 4.5.2). All visual areas left versus right**

Table 38. Group mean between-brains correlations for both the left and right hemispheres. Correlations were computed between each individual's visual area timecourse and the group average timecourse for that visual area (the group average timecourse did not include the brain with which the correlation was being computed). Values were then transformed from Pearson's r to Fisher's Z, before the mean of each of these subject- to-group correlations was calculated for each ROI and each hemisphere.

ROI	Mean Correlation (Z')	Std. Deviation
V2d left hemisphere	.443	.194
V2d right hemisphere	.542	.216
V3v left hemisphere	.5	.209
V3v right hemisphere	.458	.196
V3d left hemisphere	.417	.178
V3d right hemisphere	.486	.183

Table 39. Results of post-hoc exploratory analysis, in which paired-samples t-test were performed on the differences between hemispheres in the timecourse consistency (all individual correlation coefficients were first transformed from Pearson's r to Fisher's Z). As with the equivalent analysis performed on Pearson's r values, there is a significant difference between hemispheres for V2d, V3v and V3d.

Left hemisphere Vs. Right Hemisphere	Mean	Std. Deviation	t	df	p (2-tailed)
V2d	-.099	.091	-4.520	16	<.001
V3v	.042	.068	2.529	16	.022
V3d	-.069	.054	-5.242	16	<.001

Appendix S **Chapter 4: Fisher's Z transform of Table 14**

Table 40. Left versus right correlation differences. Correlations were computed between each individual's visual area timecourse and the group average timecourse for that visual area (the group average timecourse did not include the brain with which the correlation was being computed). Values were then transformed from Pearson's r to Fisher's Z. For V1, TO1/MT and TO2/MST, we then calculated the difference between the two hemispheres for each subject-to-group timecourse comparison by subtracting the correlation (Z' value) in the left hemisphere from the correlation in the right hemisphere (Z' value). For example, Subject 6's TO1/MT timecourse correlated with the group timecourse at Z' = 0.619 in the left and Z' = 0.455 in the right, resulting in a difference between hemispheres of 0.164 (note a positive difference indicates correlations were higher in the left hemisphere than the right). Finally, we took the mean across brains.

Left Vs. right difference	Mean Correlation (Z')	Std. Deviation
V1	-.012	.065
TO1	.042	.071
TO2	.064	.106

Table 41. Results of paired-samples t-test performed on the differences between hemispheres in the timecourse consistency for V1, TO1 and TO2 (all individual correlation coefficients were first transformed from Pearson's r to Fisher's Z). For each subject we calculated the difference between hemispheres for the subject-to-group timecourse correlations. We then performed paired-samples t tests on these differences to ascertain whether the level of timecourse consistency in TO1 and TO2 differed significantly from the level of consistency in V1.

Comparison	Mean	Std. Deviation	t	df	p (2-tailed)
V1 – TO1	-.054	.081	-2.747	16	.014
V1 – TO2	-.076	.129	-2.419	16	.028

Appendix T **Chapter 5: Fisher's Z transform of Figure 87**

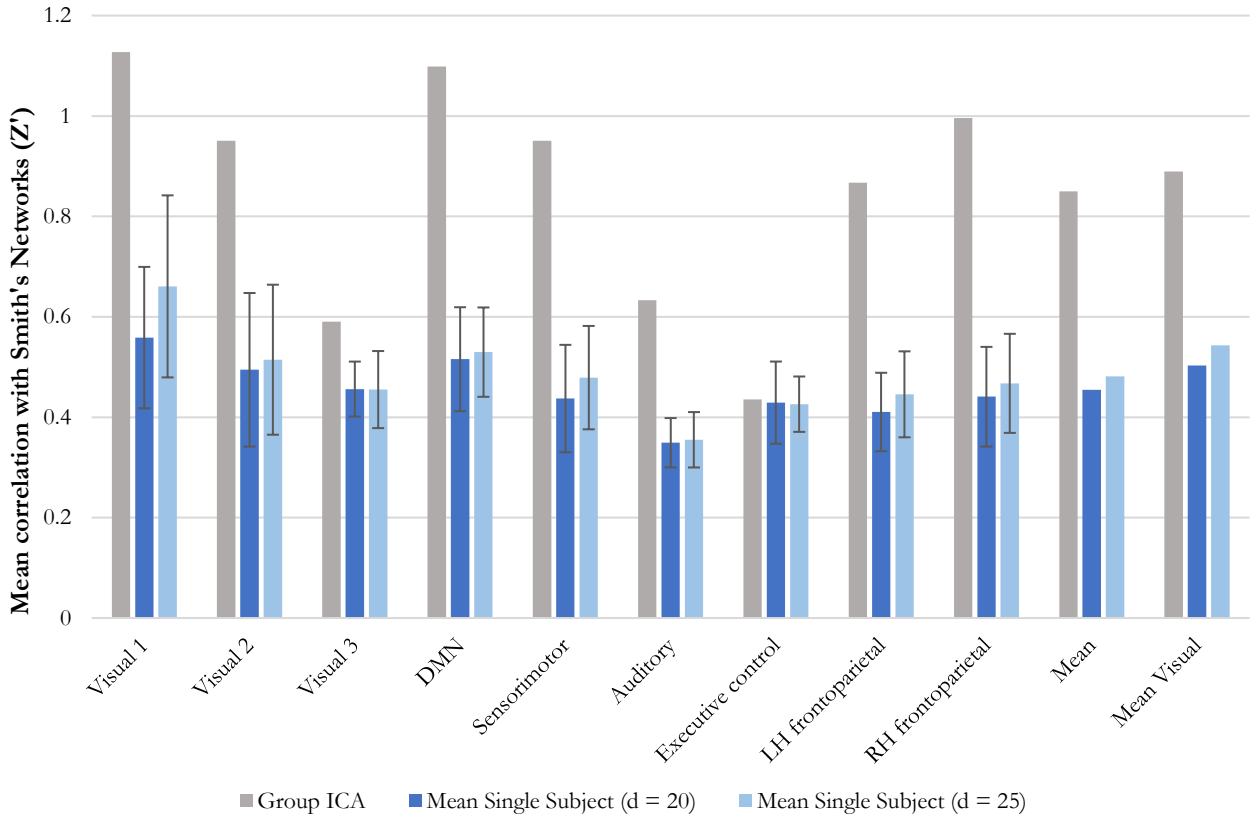

Figure 134. Correlations between each of Smith et al.'s networks and the top matched component derived from performing ICA on movie watching data (following transformation from Pearson's r to Fisher's Z). The group mean correlation between each of Smith et al.'s networks and the top matched single-subject components, with dimensionality (d) set to 20 or 25, are presented. Correlations between the top matched components derived from group-based ICA with dimensionality set to 20 are also presented.