Feature selectivity within and across areas of the primate visual cortex

Declan Rowley

Bachelor of Science (Honours)

A thesis submitted for the degree of Doctor of Philosophy at

Monash University in 2019

Department of Physiology, Faculty of Medicine, Nursing and Health Sciences

Copyright notice

Abstract

A fundamental question in visual neuroscience is how an image on the retina is represented in the activity of neurons in the visual cortex. This is an important question because how information is encoded in the visual cortex directly effects our perception of the world around us. Additionally, the processes used by visual areas to encode and transform information may be generalizable to the cortex as a whole and inform the development of future artificial processing techniques mimicking natural brain function. In the visual cortex, huge amounts of visual information need to be distilled into the perception of a small number of perceptually relevant objects while irrelevant information needs to be appropriately discarded. In the process of doing this, the cortex represents visual information in terms of particular structures (or "features") at various levels of complexity. Features can be as simple as orientated contours in the visual field or as complex as specific objects or faces. Individually, any particular neuron has a representation that encodes for visual information in just a small part of a larger feature space, with precise selectivity to multiple features and invariance to others. As a population, neurons in a visual area project information into this abstract feature space, representing the visual information in terms of a set of features. Because this set tends to be consistent within each visual area, by looking at how information is represented by a sample of cells in an area, we can gain insight into the potential functional role it plays in terms of perception. Each visual area is made up of a dense network of local circuits and organised such that neurons located nearby in the cortex tend to have similar representations. This results in complex maps of activity as multiple dimensions of visual information are encoded simultaneously within the same tissue. This thesis uses the visual cortex of the marmoset monkey (*Callithrix jacchus*) as a model system to look at feature selectivity of single cells as well the organisation of the functional architecture in which they are embedded. In Chapter 3, I describe the feature selectivity of neurons in two areas in the visual cortex, the primary visual cortex (V1) and the dorsomedial area (DM), using electrophysiology and white noise analysis. While the behaviour of neurons in V1 has been well documented, mid-tier areas such as DM have resisted characterisation. I found that cells in DM formed more complex representations than in V1 in terms of their feature selectivity, and likely also their invariance. In Chapter 4, I develop and validate state of the art calcium imaging tools to probe the structure of feature selectivity of single cells and tuning to orientation and spatial frequency over the surface of V1. In Chapter 5, I use wide field imaging to describe the relationships between maps of selectivity to different parameters in V1. Together, these results provide new insight into feature-based processing in visual cortex in terms of both feature selectivity and how it is mapped across the cortical surface.

Declaration

This thesis is an original work of my research and contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

Declan Rowley

Date: 15/10/2019

Acknowledgements

I would like to thank my supervisors, Dr. Elizabeth Zavitz, Dr. Hsin-Hao Yu, & Prof. Marcello Rosa for their tireless guidance and unwavering support over the last few years and their help with every aspect of the work presented in this thesis. I'd particularly like to thank Hsin-Hao for finding the time to continue to help and guide me despite leaving academia and being very busy with both with work and taking care of Zoe. I'd also like to thank Liz for agreeing to supervise me when Hsin-Hao left Monash particular for your help in developing my writing skills and providing constant constructive feedback with the drafting of the chapters of this thesis.

The experiments at Monash would not have possible without the assistance of experienced lab members like Katrina Worthy, who along with Hsin-Hao Yu, performed much of the histological staining and sectioning of marmoset visual cortex. I would also like to thank Nicholas Price, Adam Morris, Yan Wong, Shaun Cloherty, Maureen Hagan, Tristan Chaplin and all the members of the lab for always being happy to answer questions or give advice. I'd also like to acknowledge Saman Haghgooie who set up the first white noise stimulus and recordings in the lab.

Part of the experimental work presented in this thesis includes results from data collected in the Yamamori lab at RIKEN centre for brain science Japan in collaboration with Osamu Sadakane and under the supervision of Tetsuo Yamamori. I would like to thank the entire Yamamori lab for making me feel welcome in a very different culture. I would also like to thank Akiya Watakabe for performing the virus injections for these experiments and performing the histology, Masafumi Takaji for genotyping the marmosets as well as Toshiki Tani and Abe Hiroshi for their assistance with the surgeries and terminal experiments.

Additionally I would like to thank Paul Martin and Calvin Eiber at the University of Sydney for their invaluable assistance with setting up colour stimuli and calibrating our monitor for cone isolating stimuli.

Finally, I'd like to thank the marmosets. I hope I can make it worth your sacrifice.

My doctoral studies were supported by MBio Postgraduate Discovery Scholarship (MPDS)

Contents

1 Introduction

A fundamental question in visual neuroscience is how the information of an image on the retina is represented in the activity of neurons in the visual cortex. Humans and non-human primates are highly visual animals, who rely on this sense for navigation, socialising and communication among many other functions central to evolutionary success. This range of function is supported biologically by the visual cortex, the collection of cortical areas that are responsible for the processing of visual information and the emergence of visual perception and cognition. Studying the visual cortex allows us to look at how our sensory experience is encoded in measurable, physiological signals recorded from groups of neurons. Moreover, given that the cerebral cortex as a whole is formed by cells organised in similar circuits, it is likely that some of the processes used by the visual cortex to encode and transform information may be generalizable to the cortex as a whole. As a model system for investigating functional motifs, the visual cortex is very useful in an experimental setting because inputs to the system, visual stimuli, can be tightly controlled. By understanding the functional circuitry and underlying mechanisms behind information processing here we can learn more about our own sensory experience.

In addition, research into information processing in the visual cortex can inform the development of future artificial processing techniques mimicking natural brain function. In recent years there has been a boom in artificial neural network technology designed loosely on the architecture of the brain. This has allowed computers to accomplish tasks that were previously difficult for a computer to do such as voice and face recognition (Alexnet; Krizhevsky et al., 2012) and beating top-rated human players in games like chess and go (AlphaGo; Silver et al., 2016). As these technologies continue to develop and hardware approaches the processing power of animal brains, understanding the information processing strategies refined over millions of years of biological evolution will only become more relevant.

The visual cortex can be split into many discrete areas defined by their coherent and continuous anatomical structure and connections (Felleman & Van Essen, 1991). Although there is a long history of using anatomical markers to define boundaries between areas, these can also be defined functionally as regions of cortex with a full representation of the visual field. In early sensory areas anatomical and functional definitions of areas are generally aligned. However, uncertainty in areal designations increases in higher areas where complex function makes clear distinctions difficult. These areas are often considered discrete modules with distinct functions in terms of information processing and perception.

This thesis aims to investigate representation of visual information at the level of single cells and also in how representations are organised in cortex within a cortical area by using the visual cortex of the marmoset monkey (*Callithrix jacchus*) as a model system.

1.1 Feature selectivity

In the visual cortex, huge amounts of visual information that reach the retinas need to be distilled into the perception of a small number of perceptually relevant objects, while irrelevant information needs to be appropriately discarded (Attneave, 1954). In the process of doing this the information is represented in the cortex in terms of particular structures (or "features") of the image at various levels of complexity. Features can be as simple as regions of different contrast or oriented contours in the visual field, or as complex as entire objects comprised of multiple parts, or faces. Individually, any particular neuron is capable of encoding visual information in just part of a larger feature space, with precise selectivity to multiple features and invariance to others. For robust perception, we need the ability to generalise visual information in order to recognise an object independent of the size, shape or context in which it appears (Rust & Dicarlo, 2010). This implies we have a population of neurons that is sensitive to the abstract idea of the object, while not being affected by a set of non-relevant transformations that would drastically change the image of the object on the retina and its initial encoding in the early visual system. The presence of feature selectivity and invariance at all levels of sensory processing, indicates that this strategy is key to our encoding of sensory information.

Although the contribution of any individual cell to perception and behaviour is minimal, what information is encoded by cells in an area can tell us something about what functional role the area is playing. The activity of a visual neuron is largely driven by the visual information within a small area of the visual field, the neuron's "receptive field" (RF). Visual neurons in different areas tend to respond to different types of visual stimuli with selectivity to different sets of features. By looking at what visual properties neuron in an area are sensitive to we can infer what information is encoded in their activity. This can inform our models of what role the area has in terms of information processing.

1.1.1 Functional specialisation in parallel and hierarchical pathways

The organisation of the visual cortex is often conceptualised as a hierarchy of areas. Visual information first reaches the cortex from thalamic afferents connected to the primary visual cortex, V1, which forms the foundation of the visual system in the cortex. Even at the stage where the signal first reaches V1 a lot of information processing has already been done by retinal cells and neurons in the lateral geniculate nucleus, which have their own receptive fields and types of feature selectivity (Reid & Alonso, 1995). Axonal projections from V1 send signals to the next levels in the visual system hierarchy that includes areas such as the secondary visual cortex, V2 and middle temporal area, MT. As information passes through the hierarchy of areas it is transformed to create representations in terms of increasingly complex features. However, the visual system is not just a one-way processing stream. Feedback signals from higher areas modulate the activity of lower regions, and may provide contextual information that is relevant to the incoming signal. This allows for neurons in the lower level areas to adjust, for instance increasing the signal for an unexpected signal (Rao & Ballard, 1999) or to boost sensitivity and precision where attention is directed (Carrasco, Ling, Gobel, Fuller, & Read, 2004).

Within areas there can be further specialisation and organisation. Parallel pathways of information processing exist in the visual system which are responsible for different types of information. In the early visual system in primates, visual information is segregated into the parvocellular (P), koniocellular (K) and magnocellular (M) pathways originate in the retina as separate populations of retinal ganglion cells and project in parallel to V1 via different compartments of the LGN (Callaway, 2005; Casagrande, 1994; Livingstone & Hubel, 1984a). These pathways are thought to be responsible for different types of visual information. For example, the parvocellular pathway carrying more wavelength-specific and higher spatial resolution information than the magnocellular pathway,

which in turn carries information with the high temporal resolution necessary for motion perception (Atkinson, 1992). In V1 the P and M cells project to layer 4 whereas the K cells project to both to layer 1 and to the cytochrome oxidase (CO) blobs in layer 3 (Casagrande, Yazar, Jones, & Ding, 2007; Ding & Casagrande, 1997; Hendry & Yoshioka, 1994). In V1, the signals from all pathways are mixed together, although there is still some segregation into blob vs interblob regions.

Blob-related segregation of V1 by geniculate inputs supports early findings of Livingstone and Hubel (1984a) that suggested that CO blobs are colour processing domains. This has also been supported by 2-deoxy-D-glucose autoradiography in macaques where CO blobs were colocalised with regions of increased activity following the presentation of coloured stimuli (Tootell, Silverman, Hamilton, De Valois, & Switkes, 1988). Also in macaques, optical imaging studies (Landisman & Ts'o, 2002b; Lu & Roe, 2008; Xiao & Felleman, 2004) have revealed colour sensitive regions aligned to CO blobs and even more recently single photon wide field imaging studies have supported the link between CO blobs and colour processing (Garg, Li, Rashid, & Callaway, 2019).

Most of our understanding of colour domains in primate V1 comes from work done in the macaque. But in marmoset, intrinsic imaging studies have had mixed success in recovering colour domains. Roe, Fritsches, and Pettigrew (2005) did not find colour domains in V1, nor did Buzas, Szmajda, Hashemi-Nezhad, Dreher, and Martin (2008). Valverde Salzmann et al. (2012) did manage to find colour sensitive regions in V1 aligned to cytochrome oxidase blobs that are especially sensitive to red/green changes. However, so far the evidence for colour domains from single cell electrophysiology in marmoset has been inconclusive (P. R. Martin, 2004).

Functional segregation is also evident in V2, where cytochrome oxidase staining reveals stripes perpendicular to the border of V1/V2 (Tootell & Hamilton, 1989). These stripes form a repeating pattern in V2 and contain separate representations of the same visual space with sensitivity for different types of information and distinct cortical and subcortical connections (Livingstone & Hubel, 1984b; Shipp & Zeki, 1985). Thin stripes have been suggested to be responsible for colour processing (Xiao, Wang, & Felleman, 2003), thick for disparity tuning (Chen, Lu, & Roe, 2008) and pale stripes for orientation (Roe & Ts'o, 1995). Additionally, pale stripes can be further separated into two different types (Federer, Williams, Ichida, Merlin, & Angelucci, 2013; Felleman et al., 2015). Livingstone and Hubel's early work (Livingstone & Hubel, 1983, 1987) suggested that layer 4B of V1 projected to thick stripes, blobs to thin stripes and interblobs to pale stripes. More recent studies have disputed the idea of clean correspondence between compartments in V1 and V2 (Sincich & Horton, 2002; Xiao & Felleman, 2004), but there is also new evidence that supports the idea (Yarch, Federer, & Angelucci, 2017). Stripes in macaque V2 are well established and have been mapped at high resolution using fMRI techniques (X. Li, Zhu, Janssens, Arsenault, & Vanduffel, 2019) and there is now evidence of stripe-like compartments in both V2 and V3 of humans (Dumoulin et al., 2017; Nasr, Polimeni, & Tootell, 2016).

Although the scheme is possibly an oversimplification (K. A. Martin, 1988) the idea of parallel pathways for motion, colour and form prevails in the literature for the basic organisation of the early visual system (Nassi & Callaway, 2009). The two-streams hypothesis postulates that there are two separate visual pathways, a ventral "what" stream processing for object identification and a dorsal "where" stream processing for spatial location (Goodale & Milner, 1992; Jeffs, Federer, & Angelucci, 2015). The thick stipes in V2 project to the middle-temporal area (Shipp & Zeki, 1989), "MT", an area in the dorsal stream generally associated with motion (Newsome & Paré, 1988), whereas the pale and thin stripes of V2 project to visual area 4 (Shipp & Zeki, 1985; Xiao, Zych, & Felleman, 1999),

"V4", an area in the ventral stream specialised for processing form (Pasupathy & Connor, 1999) and colour (Conway, Moeller, & Tsao, 2007). The dorsal and ventral streams are thought to make up two relatively separate, but interacting streams. This interaction is evident in the connections between areas V4 and MT (Maunsell & van Essen, 1983).

1.1.2 Hierarchy of feature selectivity

The visual system is a complex network of interconnected areas where information flows both up and down a rough hierarchy with information shared and transformed between them. As information is sequentially transformed as it moves up the hierarchy, the way it is represented becomes increasingly complex. Moving up through the hierarchy the features become more complex as many inputs from areas lower in the hierarchy converge and create larger receptive fields. Intuitively, this makes sense because just as you can draw simple objects out of arrangements of line segments and curves, signals encoding for contrast or presence of a contour within small regions of the visual field containing could be combined to create sensitivity to those simple objects. Combining the signals from a set of set of simple objects could then be used to recognise a full scene. In higher levels of the visual cortex, selectivity for individual objects and faces have been reported (Kanwisher, McDermott, & Chun, 1997; Moeller, Crapse, Chang, & Tsao, 2017). However, this is an area of some controversy since it is very difficult to exclusively identify the correct set of relevant features in high order areas.

1.1.3 Selectivity vs invariance in V1

Most of the work for feature selectivity has been done on perhaps the most well-studied area in the visual cortex: the primary visual cortex, V1. This area is critical to visual perception as it forms the foundation of the visual system hierarchy where visual information first reaches the cortex. Neurons in this low level area tend to have very small receptive fields with widths between about 0.25 to 10

degrees of the visual field in primates (Rosa, Fritsches, & Elston, 1997). Many neurons in V1 are tuned for specific orientations and widths of luminance defined boundaries within their RFs (De Valois, Albrecht, & Thorell, 1982; Hubel & Wiesel, 1968). V1 neurons are also selective to features in the time domain, such as the direction of a moving stimulus (DeAngelis, Ohzawa, & Freeman, 1993) and can be described in terms of spatial frequency filters (De Valois, Albrecht, & Thorell, 1978). Even in V1, selectivity is diverse and two different types of orientation tuned cells, simple and complex cells, have been reported. This dichotomy is really the two extremes of a continuum of neuronal behaviour (Chance, Nelson, & Abbott, 1999; Mechler & Ringach, 2002; Priebe, Mechler, Carandini, & Ferster, 2004), but the distinction is useful for illustrating the concept of invariance. Complex cells are a classic example of the tension between selectivity and invariance in tuning. They combine the signals of multiple simple cells encoding very precise spatial information (location, phase and polarity) to create a response which sacrifices some of this precision to encode more general information about the visual stimulus.

Simple cells in V1 are very sensitive to the exact location, spatial phase and size of stimuli presented. These cells are activated by a contour of light/dark contrast in a specific location. The responses of these can be characterised in terms of responding to a single pattern within their receptive field. If the image in the receptive field somewhat matches the pattern the cell is selective for then the cell will be much more likely to fire an action potential. Conversely, if the image in the receptive field is the inverse of the preferred pattern the activity of the neuron will be suppressed. These patterns often resemble and can be effectively modelled using a 2D sinusoid and a Gaussian window: a Gabor filter (Daugman, 1985). The fact that V1 receptive fields can be conveniently characterised by Gabor filters fits very well with our interpretation of the function of V1 neurons that they encode fundamental features within small regions of the visual field. The family of mathematical functions called two-dimensional Gabor filters represent the theoretical lower limit for the joint uncertainty

for orientation, spatial frequency and spatial position. These filters therefore are the most efficient possible linear filter for sampling these variables simultaneously. The tendency of these patterns to represent Gabor filters is consistent with presence of a near optimal representation of space and frequency in V1.

Similarly to simple cells, complex cells have tuned responses which are selective for orientation and spatial frequency. However, the responses of these cells are invariant to changes to the spatial phase of the stimuli. This was first observed using bar stimuli, a bar presented anywhere in the receptive field would evoke a response, without the patterns seen for simple cells. This behaviour can still be described in terms of matching the image in a receptive field to patterns in the receptive field. However multiple patterns, in this case two, are needed to describe this behaviour. The receptive fields of complex cells can be modelled using a pair of Gabor filters that are tuned for the same spatial frequency and orientation, but with sinusoidal components that are 90° out of phase (Adelson & Bergen, 1985; Movshon, Thompson, & Tolhurst, 1978a). Nonlinear combination of the responses to the pair allow for a consistent response across spatial phase. In this way complex cells lose selectivity for spatial phase, their responses no longer carries precise spatial information, but now encodes a more general selectivity for contours.

This model is one of a larger group of models that use a set of patterns, or linear filters, to model the receptive field properties of visual neurons which can be extended to include any number of linear filters (Rust, Schwartz, Movshon, & Simoncelli, 2005). Beyond V1, in higher tier areas in the visual system we expect neurons to be sensitive to even more complicated spatial information and more complicated invariances. Add one more filter to the model of a simple cell and the invariant behaviour and more general selectivity of complex cells can be described. To characterise the

selective behaviour of cells in higher tier areas we would expect that even more filters would be needed.

1.1.4 Functional selectivity in mid-tier areas

The presence of feature selectivity and invariance at all levels of sensory processing, indicates that this strategy is key to our encoding of sensory information. But outside of V1, characterising the features and modelling responses is much more difficult. Much less is known about the feature selectivity of mid-tier areas. The hope is that by studying feature selectivity of cells within the area, we can learn something about the functional role that mid-tier areas might play in perception. In order to study feature selectively in the visual cortex, we have chosen two areas, V1 and DM, of the marmoset monkey to study in detail. These areas are a useful model for hierarchical processing because, along with V2 and middle-temporal area (MT), DM is one of the main destinations of V1 projections (Allman & Kaas, 1975; Krubitzer & Kaas, 1993). It also receives a large portion of its inputs from V2, and as such is considered a "third-tier" area in the hierarchy.

Studying DM in the marmoset brain is advantageous because DM is readily accessible by electrode array in contrast to, for example, macaques where the rostral border of dorsal surface of V2 is hidden in the lunate sulcus (Sincich, Adams, & Horton, 2003). Because of the difficulty in accessing this location in the macaque visual cortex for experiments, we know much less about the third tier areas than for many other areas. Even in marmoset, there have only been a few published studies of feature selectivity and tuning properties in DM. In marmoset DM Lui, Bourne, and Rosa (2006) found that 82.4% of neurons are orientation selective, with a further 6.9% being both orientation and direction selective. In owl monkeys, (*Aotus trivirgatus*), Baker, Petersen, Newsome, and Allman

(1981) found that DM neurons were the more sharply tuned to orientation than those from MT, dorsolateral (DL), and medial (M) areas, but were less likely to be selective for direction.

Although there has been minimal research into feature selectivity in DM, findings from these studies are consistent in the idea that DM is highly selective to orientation, and to a lesser extent, direction. Beyond this, Lui et al. (2006) observed that the response of DM neurons to an optimally oriented bar within its receptive field could often be facilitated by extending the stimuli significantly beyond the hand-mapped receptive field. This suggests that neurons in DM may tend exhibit some higher order context effect that may allow them to be sensitive to contour completion.

DM is part of the dorsal "where" stream of the two-stream hypothesis which begins in V1, continues through V2 to both DM and MT, and terminates in the posterior parietal region. As well as location information, the dorsal stream is associated with motion and hand eye coordination (Ungerleider & Haxby, 1994). As DM is an intermediate-level area within this stream, it could be expected from this hypothesis that neurons in DM may be sensitive to motion or fine edge detection. The proportion of orientation tuned cells in the studies above suggest that edge detection may be an important part of the function of DM. Another possibility is that DM contributes to both streams as it has almost balanced connections with medial and ventral stream areas, including the homologues of area V4; the ventrolateral anterior area (Rosa et al., 2009).

A key motivation behind studying the brain in animal models is that the cortical area may be homologous to an area in the human brain. DM in marmosets has been proposed as a homolog to human visual area 6 (V6) (Rosa & Tweedale, 2001). fMRI studies of human V6 show a strong response to optic flow fields: coherent motion patterns that simulate the motion signals created in

the visual field when the viewer moves in space (Pitzalis et al., 2010). However, whether these properties are also present in marmoset DM remains to be determined. This interpretation is complicated by the fact that macaque V6 has been defined as an area that only overlaps partially with the putative DM. This traditional designation of V6 corresponds to the peripheral representation of DM, where visual acuity is expected to be low and functional specialisation for motion is expected (Hadjidimitrakis et al., 2019).

The organisation of the third-tier visual areas of the primate has been the source of controversy due to the difficulty in reconciling areas defined by their histology, connectivity and single-unit electrophysiology with the retinotopic organisation seen in these areas. This is especially true for marmoset DM where there has been longstanding debate over its areal designation and whether this area should be considered a single area or the union of multiple areas with more conventional retinotopy (Lyon & Connolly, 2012). In addition, the homologue of DM for macaques has been difficult to identify for the widely accepted models. However, a recent submillimeter fMRI study which mapped the macaque visual cortex at unprecedented high resolution has revealed a retinotopic organisation that is remarkably similar to the layout in marmoset monkeys (Zhu & Vanduffel, 2018) . This is supported by work done by members of our lab by Hadjidimitrakis et al. (2019) who found similar retinotopy using electrophysiology. They concluded that V6, the tentative DM homologue, is more extensive than its typical description, and includes parts of what is typically considered V3. In this view, the extent of orientation selectivity found in marmoset DM is consistent as it matches what can be found in macaque V3 (Felleman & Van Essen, 1987; Gegenfurtner, Kiper, & Levitt, 1997).

The difficulty in reconciling the retinotopic organisation with the other methods of designating an area arises from the idea that each area should have a smooth, continuous representation of the

visual field. This is something that is very clear to see in V1 and V2 where the retinotopic organisation progresses smoothly over the cortical surface and neatly reverses direction at the V1/V2 border. However, a patchwork of third tier areas border the rostral edge of V2 which need to form continuous retinotopic representations at the border. Based on the hypothesis that areas will continue to form mirrored representations across the borders of areas, one strategy for separating these areas is based on the progression of retinotopy over cortex (M. I. Sereno, McDonald, & Allman, 1994). In this view, DM has both an internal change in retinotopic progression and would be split into multiple areas (Lyon & Connolly, 2012).

In V1 and V2, where the upper and the lower quadrants of the visual field are represented separately in ventral and dorsal cortex (Allman & Kaas, 1975). However, third tier areas introduce a new level of complexity into retinotopy, because some third tier areas, including DM, contain representations of both the upper and the lower quadrants adjacent to the lower field representation of V2 at its rostral border (Angelucci, Roe, & Sereno, 2015; Angelucci & Rosa, 2015; Rosa & Schmid, 1995; Martin I. Sereno, McDonald, & Allman, 2015). This requires reversing the progression of receptive fields within the area in order to accommodate a smooth representation across the border of V2. In a recent preprint on BioxRiv (Yu, Rowley, Zavitz, Price, & Rosa, 2019) we demonstrate the unconventional receptive field progression of DM is consistent with theoretical mechanisms that maintains topographic continuity and the idea that DM is a single functional area.

1.1.5 Challenges in mapping complex selectivity

While the behaviour of V1 neurons has been described since the work of Hubel and Wiesel in the 1950s and 60s, mid-tier areas such as DM have resisted characterisation. Early electrophysiology work in V3 characterised responses in terms of fundamental features such as orientation, motion direction and colour selectivity (Felleman, 1987; Gegenfurtner, 1997), but did not attempt more

complex visual stimuli. This is partly due to the difficulty in choosing and designing a relevant stimulus set for an area where the behaviour is unknown. We expect the function of cells in DM to be more complex than in V1, because they can combine the signals of many V1 units along with signals from other early areas to build a more complicated representation. While this arguably makes it more interesting, it makes the process of designing relevant stimulus more difficult as the set of possible behaviours is huge. In V1 highly parameterised stimuli such as drifting gratings have been used with great success, but this type of analysis is restricted by the assumption that the preferred stimuli of a neuron exists in this set. But for an area without a guiding hypothesis of behaviour how can you start to work out what it is doing?

One strategy is to expand the set of stimuli to include more complex but still parameterisable shapes. For example shapes such as arcs, sharp angles, circles, intersecting lines, hyperbolic gratings and polar gratings (Hegde & Van Essen, 2000) have been tested with some success in the secondary visual cortex V2. The average diameter of RFs of V2 neurons is roughly twice that of neurons in V1 (Gattass, Gross, & Sandell, 1981). V2 receives the majority of its input connections from V1 (Cowey, 1964; Van Essen et al., 1986; Boynton & Hedge, 2004) and neurons in this area are generally selective to more complicated visual features, but pinning down which features are relevant has been difficult and controversial. J. Freeman, C. M. Ziemba, D. J. Heeger, E. P. Simoncelli, and J. A. Movshon (2013) found that cells in macaque V2 were sensitive to the high-order statistics of naturalistic textures. This is a good clue for what type of information V2 is sensitive to, but we are still lacking a clear specification of feature selectivity in V2. This is only the second level in the visual hierarchy and already the complexity of information representation makes classification difficult.

Beyond V2, more complex selectivities have been found using diverse stimuli sets. In V4, selectivity to a variety of different shapes is well-established (Desimone & Schein, 1987; Gallant, Connor,

Rakshit, Lewis, & Van Essen, 1996; Kobatake & Tanaka, 1994). In V4 there is also evidence of tuning for colour (Kotake, Morimoto, Okazaki, Fujita, & Tamura, 2009; Zeki, 1980), textures (Hanazawa & Komatsu, 2001) and disparity (Hinkle & Connor, 2001, 2002; Watanabe, Tanaka, Uka, & Fujita, 2002). Continuing along the ventral stream of visual processing through to the inferotemporal cortex neurons are selective to complex features such as faces (Kanwisher et al., 1997; Sergent, Ohta, & MacDonald, 1992) hands and seemingly random objects (Desimone, Albright, Gross, & Bruce, 1984; Logothetis & Sheinberg, 1996; K. Tanaka, 1996). The problem with choosing a stimulus set is there is no way to know if it will contain all of the relevant features or just a subset that happens to promote activity. Also it requires that you have some hypothesis guiding stimulus design.

One alternative approach is to use white noise analysis to probe the feature selectivity of neurons. The basic idea of white noise analysis is to present a large ensemble of randomly-generated, uncorrelated visual pattern to a neuron and use its responses to infer the features that the neuron is selective to. Occasionally and purely by chance, the noise on screen will display some pattern that excites or suppresses the response of the neurons. By correlating the onscreen stimuli with neural responses we can build up a statistical understanding of what visual information is relevant to the activity of a neuron. Because the full set of stimuli is random, any statistical structure found in the subset of frames that the neuron responds to is due to the selectivity of the neuron. The statistical structure of the subset exists in the high dimensional space of all possible visual patterns, but we can use dimension reduction techniques under the assumption that only a low-dimensional subspace contains the features that drive the neuron (Attneave, 1954) to try to find what determines the neuron's behaviour. Compared to the traditional methods where stimuli are highly restricted white noise analysis is unbiased as it randomly samples from the space of all possible images.

By looking at how information is represented by a sample of cells in an area, we can gain insight into the potential functional role it plays in terms of perception. However, individual cells do not form or act on their feature selectivity in a vacuum. They are part of a larger network of cells and embedded in functional architecture that determines their response properties. Understanding a visual area requires not only understanding the responses of individual cells, but an understanding of the organisation that gives rise to their behaviour.

1.2 Feature Maps

A well-documented property of the visual cortex is that neurons that are organised into columns of cortex which have consistent tuning through depth. Assuming columnar organisation means that tuning can be described in terms of two dimensions of cortical surface, disregarding depth. Adjacent columns tend to have very similar representations such that over horizontal distance, receptive field properties change systematically and often smoothly creating topographic maps. Topographic maps are not restricted to the visual cortex, but can be found in sensory areas for touch (Friedman, Chen, & Roe, 2004), hearing (Reale & Imig, 1980) as well as in motor areas (Laskowski & Sanes, 1987) suggesting that in general information is represented in maps throughout cortex.

It has been suggested that topographic maps are formed by constraints imposed during development that determine that neurons near to each other in cortex form very similar receptive field properties (Durbin & Mitchison, 1990). The clearest example of topology in the visual cortex is retinotopy: the smooth progression of receptive field centres as you move from one side of a visual area to the other. Receptive fields of neurons are organised to maximise coverage of the visual field while maintaining smooth overlapping representations (Swindale, Shoham, Grinvald, Bonhoeffer, & Hubener, 2000). Retinotopy is ubiquitous in the visual system and areas are often defined in terms of their retinotopy. Adjacent areas have congruent retinotopy, with maps being continuous across the borders typically resulting in mirror images of each other (Rosa, 2002). As a consequence, borders between visual areas can be identified by detecting reversals in the progression of retinotopy (i.e., “field sign” introduced by Sereno et al., 1994). Although experience plays a role in refining retinotopic maps, these maps can be formed completely independently of visual experience (Tiriac & Feller, 2019).

1.2.1 Map formation

Molecular signals that guide the initial formation of the retinotopic map in V1 by guiding the projection of thalamic afferents (Espinosa & Stryker, 2012; Sperry, 1963). This is at least partly mediated by EphA-ephrin signalling (Frisén et al., 1998; Nakamoto et al., 1996). The EpA family of receptors are expressed on the growth cones of thalamic cells that project to V1 and interact with ephrin-A ligands bound to the surface of V1 cells. The ephrin-A ligands are expressed in a gradient across V1, which creates a smooth projection of afferents across the surface. EphA-ephrin signalling is critical to the formation of retinotopy in V1 and may be just one of multiple molecular signals that guide the initial formation of maps (McLaughlin & O'Leary, 2005).

Before eye opening in mammals, the initial retinotopy in V1 is refined by waves of spontaneous activity in retinal ganglion cells: retinal waves (Firth, Wang, & Feller, 2005). Retinal waves are critical for the refinement of visual maps in V1 and also for the retinotopic organisation of the LGN. Retinal waves are caused by the propagation of excitation through several mechanisms (Zhou & Zhao, 2000). This is evident in studies where these mechanisms are blocked and retinotopy in the adult animal is disturbed (Cang et al., 2005). Columnar architecture is likely to be created through local inhibitory connections (Hensch et al., 1998; Hensch & Stryker, 2004). These shape the geometry of the thalamic projections by competitive plasticity: if the same inputs activate neighbouring regions then responses are suppressed and the inputs are pruned promoting differentiation. Retinal waves are also suggested to be the driving force behind the organisation of orientation in V1, which has been shown to form around the time of eye opening without relying on visual experience (cats: Crair et al., 1998; mouse: Hagihara et al., 2015, ferret: Tiriac et al., 2018)

1.2.2 Mapping multiple features in cortical tissue

Topological organisation applies not just to retinotopy and orientation tuning, but also selectivity to other visual information. This means that selectivity to multiple features need to be organised to maximise **coverage** of all possible combinations of values for the encoded features while also maximising **smoothness** (N. V. Swindale, 2004). This leads to complex interactions between many maps as selectivity to many feature dimensions need to be smoothly projected onto the two-dimensional cortical surface. This complex structure can be created with simple biologically plausible mechanisms that enforce the simple rule: neurons that are located near to each other tend to develop similar representations (Durbin & Mitchison, 1990; Goodhill, 2007; Goodhill, Bates, & Montague, 1997).

Maps for multiple parameters were first famously described in Hubel and Weisel's (1977) "ice cube" model which showed a schematic of V1 made up of many columns tuned for both orientation and ocular dominance. In this model the progression of orientation tuning and ocular dominance progressed in perpendicular directions so as to cover all combinations of both parameters. This is an early model that did not account for how more than two parameters might be encoded in the same surface. Later work, mapping orientation tuning across two dimensional cortical surface revealed that adjacent to the orientation columns there were singularities in the orientation tuning map where the preference rapidly changed across cortex: pinwheels.

The organisation of V1 has a precise functional architecture that includes maps for tuning to many more parameters than just orientation and ocular dominance. Functional maps in V1 have been studied in multiple mammalian species (monkeys, cats, tree shrews, ferrets, humans) using

electrophysiology (Hubel & Wiesel, 1974), autoradiography (Tootell, Switkes, Silverman, & Hamilton, 1988), voltage-sensitive dyes (Blasdel & Salama, 1986), intrinsic optical imaging (Grinvald, Lieke, Frostig, Gilbert, & Wiesel, 1986) and functional magnetic resonance imaging (Yacoub, Harel, & Ugurbil, 2008). These studies have revealed columnar structures and topological maps for orientation selectivity, direction selectivity and spatial frequency selectivity. Maps for tuning have also been reported in V2 (Felleman et al., 2015), V4 (Li et al., 2013) and MT (Albright & Desimone, 1987; Kaskan, Dillenburger, Lu, Roe, & Kaas, 2010).

For marmoset the bulk of the work done on V1 topology has been done using intrinsic optical imaging (Buzas et al., 2008; Roe et al., 2005; Schiessl & McLoughlin, 2003). However, since the signal is driven by blood flow, the signals have poor spatial resolution (~50µm) and temporal resolution (a couple of seconds), and are only indirectly related to the underlying electrophysiological activity. Calcium imaging is an alternative method for monitoring the activities of neurons from large sections of the cortex. Since the fluorescent signal corresponds to the intracellular calcium concentration which is driven by action potentials, the recordings are more directly related to the signals recorded with electrophysiology. In macaque this has been successfully used to more precisely map tuning in V1 (Garg et al., 2019; Nauhaus, Nielsen, & Callaway, 2016). Because many of the features that V1 neurons are selective for are already known, V1 present a unique experimental opportunity for furthering our understanding of how topographic maps are formed and organised. Maps in V1 could be especially useful in understanding about how the maps for different parameters in the same cortical surface interact with each other.

1.2.3 Overlapping maps create feature space bottlenecks

Spatial receptive fields in V1 are the smallest in the visual cortices. In primate V1 receptive fields sizes go from 0.25 degrees in diameter at the fovea up to 10 degrees in the periphery (Rosa et al.,

1997). As V1's retinotopy samples the visual field at the smallest resolution with many tiny overlapping receptive fields. This in turn means that the amount of cortical surface dedicated to the same area in visual space is the smallest in the visual cortex. The area of cortical surface dedicated to part of the visual field is termed the "cortical point image" (Capuano & McIlwain, 1981; Chaplin, Yu, & Rosa, 2013; Van Essen, Newsome, & Maunsell, 1984). The size of the cortical point image limits the numbers of feature combinations that be encoded for any part of the visual space. Given that the amount of cortical tissue needed to represent each parameter combination is fixed (Buxhoeveden & Casanova, 2002; Mountcastle, 1997), the number of features that can be represented per cortical point image in V1 has an upper limit at around 8 to 10 (Swindale, 2000).

Because the number of available features is limited by cortical magnification, V1 has the smallest number of features to work with in the visual cortex. If V1 is restricted to a few features due to tight constraints on its retinotopy compared to the cortical resources available, then this could go some way to explain why its size is so well preserved across primates: any further reduction of cortical point image would require discarding a representation of a fundamental visual feature. It also stresses the importance of V1 representations, as all visual information must first pass through V1 and be filtered through the restricted feature space as dictated by the V1's physiology. Given the competing constraints on the topographical organisation of V1 that stem from balancing a high-acuity representation with a range of fundamental visual features, we expect V1 to have highly optimised topographical maps with structured relationships between maps of tuning to different features.

In V2 where the cortical point image is much larger (Rosa et al., 1997), 14 features could be represented and given that receptive field diameters in DM are approximately ten times larger than V1 (Lui et al., 2006), and much larger than those of V2, even more features could be represented.

Moving up the visual hierarchy the spacing of columns (Fujita & Fujita, 1996) and sizes of pyramidal cells (Elston & Rosa, 2000) increase, and the cortical point image increases in size (Gattass et al., 2005). This increases the maximum number of features that can be simultaneously represented by the same set of neurons increases and the complexity of receptive fields can expand into an increasing number of abstract dimensions.

Increasing the number of dimensions may have profound effects on how specific the selectivities of neurons in higher order areas can be. If the number of dimensions is no longer a limiting factor for representation then cells are free to form representations combining features encoding wide classes of objects and still be highly selective. This may be used beyond visual areas to combine information across modalities and allow for associations of very different stimuli, percepts and ideas. The role of maps in the co- representation of multiple features is still not well understood. V1, where the features encoded may be highly parameterised and controlled for, serves as a perfect test area for probing the organisation of maps.

1.3 Conclusions

We study the visual cortex to learn about how the brain encodes and processes information. Feature selectivity and invariance is a theme throughout the visual cortex suggesting that it an important strategy in processing visual information. Despite this, we know very little about the features that are relevant to visual areas beyond V1. By looking at how information is represented by a sample of cells in an area, we can gain insight into the potential functional role it plays in terms of perception. Likewise, while topological maps are ubiquitous in sensory areas, our understanding of how multiple maps in the same tissue are organised is superficial.

This thesis will look at the feature selectivity of single cells as well the organisation of the functional architecture in which they are embedded. In chapter 3, I describe the feature selectivity of neurons in two areas in the visual cortex, the primary visual cortex (V1) and the dorsomedial area (DM), using electrophysiology and white noise analysis. In chapter 4, I develop and validate state of the art calcium imaging tools to probe the structure of feature selectivity of single cells and tuning to orientation and spatial frequency over the surface of V1. In chapter 5, I use wide field imaging to describe the relationships between maps of selectivity to different parameters in V1. These results are some of the most complete characterisations of V1 map structure for multiple parameters in the same animals.

2 Methods

2.1 Overview

The results presented in this thesis were generated from datasets collected from adult marmoset monkeys (*Callithrix Jacchus*) in two different labs and using two different methodologies. At Monash University, a series of electrophysiological recordings were obtained from anaesthetised marmosets. At the Yamamori lab at RIKEN Centre for Brain Science, in Saitama, Japan, calcium imaging recordings were obtained in both awake and anesthetised preparations.

All procedures done at Monash University were approved by the Monash University Animal Ethics Experimentation Committee, and the experiments were conducted in accordance with the Australian Code of Practice for the Care and Use of Animals for Scientific Purposes. All experiments done at the RIKEN adhered to guidelines set out by the Institutional Animal Research Committee at RIKEN.

This chapter will outline animal preparation and recording methods that are common to multiple results sections. More chapter specific methods, especially for data analysis, will accompany the results in the relevant chapter.

2.2 Electrophysiology preparation

2.2.1 Surgery and anaesthesia

Electrophysiological data were collected from individual recordings sessions that took place in the Department of Physiology at Monash over a number of years. In each session, data were collected from a single anaesthetised adult marmoset monkey (*Callithrix jacchus*) over the course of 1 to 4 days. These experiments were non-recovery, and involved several minor surgical procedures, including a craniotomy to allow the insertion of multi-electrode arrays into areas V1 and DM of the visual cortex. An outline of the general procedure can be found in Bourne and Rosa's 2003 paper, and an update with respect to the drug protocol in Yu and Rosa 2010.

The animal was first given a pre-anaesthesia medication consisting of a mild long-acting sedative, diazepam (2mg/kg), and atropine (0.2mg/kg), to decrease salivation and bronchial secretions as well as to minimise disturbances in the cardiac rhythm. After 30 minutes, surgical anaesthesia was induced by intramuscular injection of alfaxalone (Alfaxan, 8mg/kg). After anaesthesia, the animal was injected with an antibiotic (Norocillin, 25mg/kg) and dexamethasone (Dexason, 0.3mg/kg) to prevent cerebral oedema during the experiment.

The surgical procedures included a tracheotomy, cannulation of the femoral artery, craniotomy, and array implantation. The tracheotomy was required in order to insert a tube for artificial ventilation of a nitrous oxide and oxygen mixture (7:3), which provided a background level of analgesia and sedation. The cannulation of the femoral artery allowed for a maintenance solution to be continuously infused during the experiment.

The animal was then secured to a stereotaxic frame to allow for stability during thethe craniotomy and array insertion. Before the craniotomy a local anaesthetic, Marcain 0.5% (bupivacaine), was injected at areas of surgical incision. The skin and muscles over the region of interest were retracted

and the craniotomy was performed with a fine dental drill bathed in cooled saline to create a window exposing the areas of interest. The dura matter was then resected to allow for electrode array insertion.

After surgery, a maintenance solution was continuously infused through the cannula in the femoral artery initially at a rate of 5ml/hr for the first 30 minutes, then at 2ml/hr. The rate was then adjusted over the course of the recording session to keep a steady heart-rate. This solution comprised the opioid anaesthetic sufentanil citrate (Sufenta Forte, 50µg/ml), the neuromuscular blocker pancuronium bromide (2mg/ml), dexamethasone (Dexapent, 5mg/ml), xylazine (Xylazil-20, 20mg/ml) and salts and nutrients (.18% NaCl/4% glucose solution, Synthamine-13 and Hartmann's solution). The use of pancuronium bromide was required by the need to have a stable eye position for the purposes of the experiments. However, other experiments, conducted with a similar solution but without pancuronium bromide, have confirmed that this regime results in a stable level of anaesthesia (S. Zhu et al., 2019).

Throughout the duration of the recording session, the animal's body temperature was kept at a steady 38 degrees, measured by rectal thermometer connected to a homeostatic electric blanket unit. The heart rate and blood oxygenation levels were monitored constantly using a digital vital signs monitor (Surgivet) equipped with a cutaneous contact probe, which was opposed to the animal's inner thigh. Adjusting the rate of infusion of the maintenance solution along with constant ventilation of the gas mixture allowed the animal to be kept in a state of anaesthesia for the entire period of time required for the experiment.

In order to present visual stimuli to the anaesthetised animal, one eye needed to be held open, and focused at the correct distance corresponding to the surface of a monitor, while preserving the cornea in a healthy state for a prolonged period of time. The eye contralateral to array implantation

was held open with custom-made microretractors, and atropine (Atropt, 1%) was used to dilate the pupil (mydriasis). Phenylephrine was used to paralyse the ciliary muscle, which brought the focal distance of the eye to infinity (cycloplegia). A contact lens of appropriate curvature was then placed on the frontal surface of the cornea to focus the eye onto the stimulus monitor. In order to prevent the cornea from drying a lachrymal film was created between inner surface of the lens and the cornea with carmellose sodium (Celluvisc). The other eye was covered with chloramphenicol eye ointment to prevent epithelial dehydration and held closed with tape.

2.2.2 Data Collection

Two different types of multi-electrode arrays were used in the electrophysiology experiments. A 10x10 “Utah” array (Blackrock Microsystems, Salt Lake City, USA) with 96 active channels. Electrodes were 1.5mm long and implanted in cortex with a pneumatic insertion tool covering a small area of the cortical surface at intervals of 400µm. Raw voltage signals were collected at 30kHz and high-pass filtered at 750Hz using a Cerebus system (Blackrock Microsystems, Salt Lake City, USA). This filtering choice is standard for our lab using this setup (Zavitz, Yu et al. 2016, Zavitz, Yu et al. 2019).

We also used a single shank linear “NeuroNexus” probe (NeuroNexus, Michigan, USA) with 32 channels spanning the depth of the cortex at intervals of (50µm). The signals from the “NeuroNexus” were collected with the same parameters using a Cereplex system (Blackrock Microsystems, Salt Lake City, USA). In all cases, spikes were detected on each channel using an automatic threshold of the local signal. After recording, manual spike sorting was performed offline using Plexon Offline Sorter (Plexon Inc., Dallas, USA).

The dorsomedial area (DM) was located using gross physiological landmarks and then confirmed using progression of receptive fields with a single electrode. Once the caudal border between V2 and

DM was localised, arrays were implanted into the region rostral to this border. The “Utah” array was implanted in area DM using a pneumatic insertion tool whereas the “NeuroNexus” probes were inserted into the cortex at an approximately perpendicular angle to the cortical surface and advanced such that spiking activity could be seen across as many as channels along the 1.6mm probe as possible. The position of the electrode arrays was confirmed using retinotopy and histology.

2.2.3 Visual Stimulus

Visual stimuli were generated in Psychophysics toolbox (Brainard, 1997; Pelli, 1997; Kleiner et al, 2007) on MATLAB (Mathworks Inc.) and displayed on either a Display++ monitor (Cambridge Research Systems Ltd, Kent, UK) or a VIEWPixx monitor (VPixx Technologies Inc., Saint-Bruno, Canada) with a refresh rate of 120Hz. These monitors have been fully characterised in terms of their luminance ranges (Display++: 0.40 to 234.0 cd/m^2 and VIEWPixx: 0.25 to 106.6 cd/m^2) and were configured to present luminance values on a linear scale (Ghodrati et al., 2015). The stimuli for this experiment was a dense tertiary white noise, with each stimulus frame generated at random. Each frame consisted of a square grid of equally sized elements which could be either white, mid-grey or black. The stimulus was presented to the animal at a high rate (60 stimulus frames per second) over 2-4 hours.

2.2.4 Staining and histology

After the completion of data collection, the animal was given a lethal overdose of sodium pentobarbitone (100mg/kg). Histology was performed for the experiments that used the “Utah” array to ascertain which electrodes were implanted within the boundaries of DM. The array was removed and the animal transferred to a fume hood where it was perfused with buffered saline. The unfixed brain was immediately extracted and the two hemispheres were separated and physically flatmounted (Tootell & Silverman, 1985).

Flatmounting was performed by gently dissecting away the white matter of the cortex with dry cotton swabs, with the cortex supported on a piece of moist filter paper (pial surface down). Relaxation cuts were made in the fundus of the calcarine sulcus, and at the anterior end of the sylvian sulcus to allow the cortex to lie flat. The cortex was held in fixative (4% paraformaldehyde) between two large glass slides under a small weight overnight, and then was soaked in sucrose solution in increasing concentrations (10%, 20% and 30%). The flat-mounted hemisphere was then cut in a cryostat to a thickness of 40µm. Alternate sections were stained for myelin and cytochrome oxidase. Cytochrome oxidase histochemical reaction followed the procedure modified by Silverman and Tootell (1987) from the protocol of Wong-Riley (1979).

2.3 Imaging preparation

2.3.1 Virus vectors and calcium indicator

The experiments conducted in the Yamamori lab in Japan used GCaMP6s indicator carried by an adeno-associated virus (AAV) in concert with a tetracycline-inducible (TET) gene expression system to amplify the transgene expression (Sadakane et al., 2015ab). This combination of indicator and virus was chosen to produce the strongest and most reliable fluorescence response. The GCaMP6s is a variant of the GCaMP6 family of calcium indicators, characterised by its slow response time ("s" in GCaMP6s stands for "slow"). The faster version of GCaMP6 was rejected after it failed to produce strong signals in previous experiments. The protocol for anaesthesia was similar to the one described above

2.3.2 Virus injection

AAV viral vectors (rAAV2 with Thy1S promoter; Ako et al, 2011) carrying the GCaMP6s calcium indicator (amplified with the TET-Off system) were injected into the primary visual cortex (V1) of four adult marmoset monkeys. All injections were performed under anesthesia induced by an intramuscular injection of ketamine hydrochloride (Ketalar, 25 mg/kg) and under aseptic conditions. The marmoset was then placed in a stereotaxic apparatus with anaesthesia maintained using inhaled isoflurane (1.5%-2.5% in oxygen). Pulse oxygen (SpO2), heart rate, and rectal temperature were continuously monitored to judge the animal's condition (Sadakane et al., 2015b). The virus was injected by first drilling small holes in the skull over V1, which was identified by stereotaxic coordinates. Four injections of 0.5µL aliquots were made at 0.1 µL/min for 5 min an approximate depth of 500µm below the cortical surface using a pulled glass pipette (70 µm outer diameter). The injection sites were spread out to maximise the coverage of the indicator expressed over the region to be imaged.

2.3.3 Window Construction

Surgery for implantation of an imaging chamber was performed a month after the virus injection. The craniotomy and durotomy were performed under anaesthesia as above to make a circular window (5mm diameter) in the skull above the injection site. The dura was cut and peeled back onto the exposed skull and a saline solution was applied to the exposed cortex. We used a custom-made glass window consisting of an outer ring of glass (7mm in diameter) around a flat indented circle (3mm in diameter). The inner surface was pressed against the cortex while the elevated ring sits above the skull. To fix the window into position the outer ring was secured with dental cement to the skull while the inner window was pressed firmly against the cortex. A head post was also attached to the skull with dental cement so the marmosets head could be fixed in place for awake recording. The skin was then stitched back in place and the marmoset recovered for 5-7 days. The quality of the window and virus expression were checked by eye immediately after surgery with the animal still under anaesthetic, using an ultraviolet LED and filtered glasses.

2.3.4 Data collection

The marmoset was head-fixed in a stereotaxic frame at a distance of 40 cm from the screen used to display visual stimuli. For awake experiments, an IR eye tracker was used to confirm that the marmoset was looking at the (full-screen) stimuli. The rotating head stage of the microscope was positioned carefully so that the lens of the microscope approached the imaging window from a perpendicular angle. The histogram of the intensity of the recorded signal over every pixel was monitored during the course of every recording session to ensure the laser power was not high enough to cause photo-bleaching.

For anaesthetised preparations, anaesthesia was induced by intramuscular injection of Alfaxan (alfaxalone 10 mg/kg), ampicillin (40 mg/kg) as an antibiotic, and carprofen (5 mg/kg) as an anti-inflammatory agent. Either a tracheotomy (n=2) was then performed and a breathing tube was inserted into the trachea or a ventilation mask (n=2) was fixed for artificial ventilation. The marmoset was then placed on a heated mat under the microscope with its head post fixed to a stereotaxic apparatus. During the experiment, an intravenous infusion of a mixture of pancuronium bromide (0.1 mg/kg/h), sufentanil (6 – 8 g/kg/h), and dexamethasone (0.4 mg/kg/h), in a saline-glucose solution was continuously administered. The electrocardiogram and SpO2 level were continuously monitored to judge the marmoset's condition during recording and the rate of infusion was adjusted to maintain a stable heart rate. Before recording data from the animal, the contralateral eye was held closed with a piece of tape and the ipsilateral eye was held open with micro retractors, and a contact lens and protective eye drops were inserted. The appropriate strength of the lens was found using a retinoscope and confirmed by displaying drifting gratings and looking for changes in fluorescence in real time using the single photon camera.

2.3.5 Visual stimuli

Visual Stimuli were generated in Psychophysics toolbox (Brainard, 1997; Pelli, 1997; Kleiner et al, 2007) on MATLAB (Mathworks Inc.). The visual stimuli for were presented on a 4K LG LCD screen (24UD58-B) with a refresh rate of 60 Hz. These experiments used a variety of stimuli presented at a distance of 40cm from the subject's eye using a Linux machine running Matlab version 2014a.

Full screen drifting sinusoidal gratings were used to measure orientation and spatial frequency tuning. Gratings with a temporal frequency of 1Hz, in 12 different directions (30° intervals) and at 6 different spatial frequencies: 0.25, 0.5, 1, 2, 4 and 8 cycles per degree were presented. Trials were randomly interleaved and presented for 1.5 s following 1 s of blank grey screen. For awake experiments, slight adjustments to the stimuli were made in order to keep the marmoset's attention on the screen. Presentation of the full screen was reduced to 0.75 seconds and blanks were replaced with half a second of greyscale natural images. An infrared eye tracker was used to confirm that the marmoset was looking at the stimuli.

Orientated black and white bars of fixed width were used to map receptive fields. These bars had 6 different orientations for which there were 7 positions along an axis orthogonal to the orientation at which the bars were displayed. The order of trials was randomised and an equal number of black and white bars were presented for every combination of orientation and position (Nauhaus et al., 2016). As this stimulus required a rough estimate of the relevant area of the screen this was only presented in anaesthetised conditions where fixation was stationary.

Isoluminant (as determined by a handheld photometer) colour stimuli were used to reveal functional maps related to colour processing, similar to those used in Valverde Salzmann et al. (2012). These were colour opponent pairs of blue-yellow, red-green as well as an isoluminant grey paired with every colour used. Trials would include both changing back and forth between each of the opponent

coloured pairs (at 0.2 to 0.5s intervals) as well as unipolar excursions to and from grey. Rather than matching every stimuli in the trial to the lowest luminance colour (blue), pairs were luminance matched to each other and proceeded by a blank period of 2s with a luminance matched grey to avoid interference between trials.

Additionally, cone isolating stimuli were employed for some animals (animals 2 and 3) where prior knowledge of the genotype of the animal and emission spectra of our screen allowed us to construct stimuli that specifically modulate single cones or combinations while not driving others. We measured the emission spectra of an LG 4K LCD screen in absolute units using a spectrophotometer at the Martin lab at University of Sydney. This was a different, but identical model to the screen used in the experiment. For confirmation, both screens were measured using an iPro (X-Rite inc., Grand Rapids, Michigan, United States) portable spectrometer with relative units and found to be a close match.

2.3.6 Multi-scale imaging

Before imaging data were digitally acquired, the fluorescence levels of the signals were visually inspected using an ultraviolet LED light source and filtered glasses. When the LED was aimed at the window we could see a powerful glow confirming the strength of the virus expression. We imaged the fluorescent signals at 3 different spatial scales: We used one-photon microscopy to perform wide-field imaging over a 3x3mm window with 512x512 pixels (5.77µm/pixel) at 23.3Hz (left box in Figure 2.1). We then used two-photon microscopy with a x10 optical lens. Using different digital zoom, we imaged at two spatiotemporal scales: over a 1200x1200µm area at 1024x1024 pixels (1.17µm/pixel), 15Hz (the blue box in Figure 2.1), and over a 600x600µm area at 512x512 pixels (1.17µm/pixel), 30.5Hz (the red box in Figure 2.1).

2.3.7 Single photon imaging

Wide field single-photon microscopy was performed over the full 3x3mm window with a CFI Plan Apo λ x4 optical lens (Nikon) at 512x512 (5.77µm/pixel) at 23.3Hz with a Bergamo II microscope (ThorLabs, Newton, NJ, USA). The One-photon imaging was performed using a 1500MGE Monochrome gig-E CCD scientific camera (ThorLabs, Newton, NJ, USA) or a CMOS camera CS2100M-USB, scientific camera (ThorLabs, Newton, NJ, USA) using an Epi-Illuminator with M470L3 470 nm, 650 mW LED (ThorLabs).

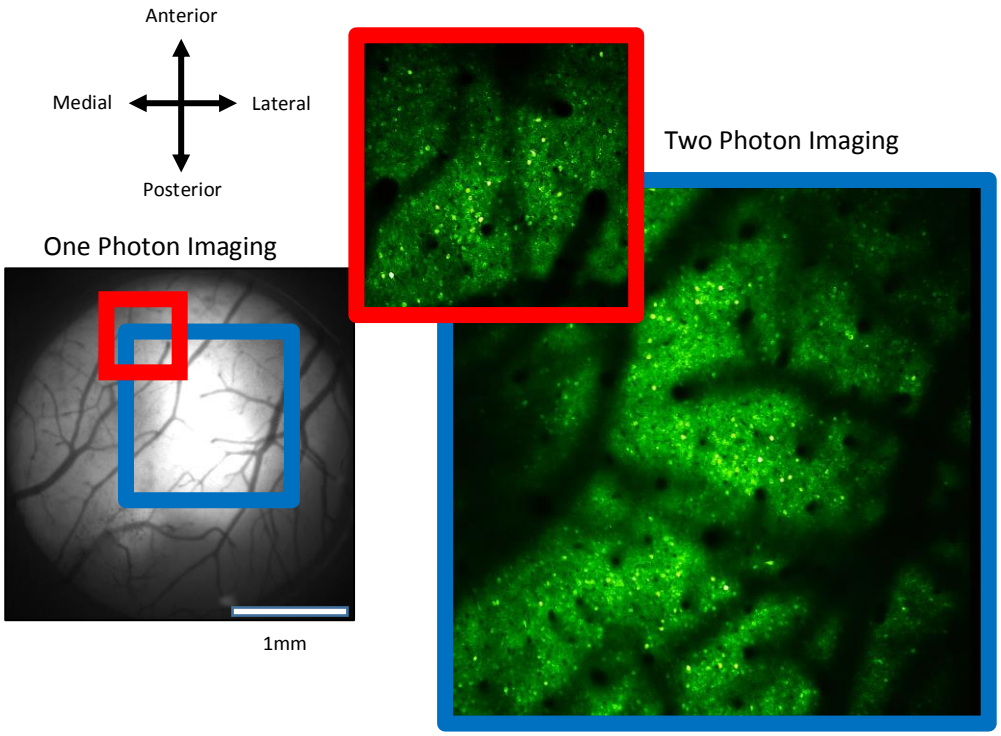

Figure 2.1: Calcium imaging from the primary visual cortex (V1) in the right hemisphere of a common marmoset. The strength of signal achieved allowed us to record from 3 different scales: over the entire 3mm window, over a large population (Blue, 1024x1024p, 1200µm 15Hz) and at a very fine spatial and temporal scale (Red, 512x512p, 600µm, 30.5Hz). The data presented here was recorded from a single marmoset monkey infected with an adeno-associated virus expressing GCaMP6s in cortical neurons within the superficial layers of V1.

Two-photon imaging was performed using a Bergamo II (Thorlabs) Multiphoton Laser Scanning Microscope with and MaiTai HP (Spectra Physics) broadly tunable ultra-fast laser set to a wavelength of 940nm. Regions of interest were chosen by finding a bright area with as few blood vessels as possible, depth was recorded from the surface of the cortex at 150µm.

2.3.8 Stabilisation and image registration

For single photon data, registration was done in imageJ (Schindelin et al., 2012; Schneider, Rasband, & Eliceiri, 2012) using a custom script to run Turboreg (Thevenaz, Ruttimann, & Unser, 1998) first over averaged sets of between 20-50 frames against a reference image and then within each subset

to the mean. This provides a steady image even when there are large slow movements and also very small quick movements which are the main problems we faced due to animal moving in awake experiments. The same registration was done in both awake and anaesthetised experiments.

2.3.9 Image segmentation

For images acquired with two-photon microscopy, a matrix factorization algorithm (Pnevmatikakis et al., 2016) was used to isolate single neurons. This method automatically segments overlapping cell bodies (i.e., demixing) using the spatiotemporal correlations of the signals (Maruyama et al., 2014). After neurons were identified by the algorithm, signals from all pixels inside the identified neurons were averaged for the subsequent analyses.

2.3.10 Staining and histology

For the imaging experiments, after the terminal experiments were performed the brains were perfused and stored. Later, the schedule varied across animals, the brains were stained for cytochrome oxidase (Rosa et al., 1997; Wong-Riley, 1979) and sectioned and scanned using a high speed multiphoton microscope with integrated vibratome sectioning (TissueCyte, TissueVision, Inc.). This microscope shares the same laser with the two-photon microscope used in our in-vivo experiments. Using TissueCyte we recreated a 3D image of the brain and registered it to the imaging window by matching the location of blood vessels. Cytochrome oxidase histochemical reaction followed the procedure modified by Silverman and Tootell (1987) from the protocol of Wong-Riley (1979).

Unlike the registration process between the one and two photon images, for the sliced sections the vertical tracts of the blood vessels through the cortical layers was used rather than the horizontal pattern of branching across the surface of the cortex. Tears in sliced sections were first removed in Photoshop (Adobe, Inc.) before slices were overlaid with partial transparency until darker regions of cortex (blobs) became apparent.

3 Feature selectivity of neurons in the primate visual cortex

3.1 Introduction

A fundamental question in visual neuroscience is how information about the image on the retina is represented in the activity of neurons in the visual cortex. As described in chapter 1, the visual cortex is made up of a number of hierarchically-arranged cortical areas. The primary visual cortex (V1) processes inputs from the lateral geniculate nucleus, and feeds this processed information forward to subsequent areas, such as second-order area V2, and mid-tier areas like the middle temporal area (MT) and the dorsomedial area (DM) that receive input from both V1 and V2. Visual neurons in different areas tend to respond to different types of visual stimuli with selectivity and invariance to different sets of features. By looking at what visual properties neuron in an area are sensitive to we can infer what information is encoded in their activity. This can inform our models of what role the area has in terms of information processing.

In order to study feature selectively in the visual cortex, we have chosen two areas, V1 and DM, of the marmoset monkey (*Callithrix jacchus)* to study in detail. These areas are a useful model for hierarchical processing because, along with V2 and MT, DM is one of the main recipients of V1 projections (Allman & Kaas, 1975; Krubitzer & Kaas, 1993). Conversely, approximately 40% of all neurons providing cortical afferents to DM are located in V1 (Rosa et al., 2009), with V2 cells providing the next most numerous source of projections (~15%). Although there has been minimal research into feature selectivity in DM, findings from these studies are consistent in the idea that DM is highly selective to orientation direction (Baker et al 1981; Lui et al 2006). Part of the difficulty in studying selectivity in DM comes from the fact that in macaques, a very common animal model for visual neuroscience, the rostral border of dorsal surface of V2 is hidden deep in the lunate sulcus (Sincich et al., 2003). In contrast, the dorsal surface of the visual cortex in the marmoset brain is flat and DM is readily accessible by electrode array.

3.1.1 Reverse correlation analysis

Reverse correlation may be used to determine the selectivity of a neuron from white noise stimulation. Briefly, it involves finding the spike times of every spike during the recording and checking what image was onscreen shortly beforehand. The patterns that are displayed at a fixed time interval, to account for delay, prior to each spike are collected (De Boer & Kuyper, 1968). The collection is referred to as the Spike Triggered Ensemble (STE). Reverse correlation analysis quantifies how these stimulus frames that correlate with spikes are different to the full set of stimuli presented. In white noise analysis, because the full set of stimuli is just noise, statistics of stimuli over or underrepresented in the STE reflect the behaviour of the neuron.

The simplest useful statistic of the STE is the Spike-Triggered Average (STA), which is simply the spatial average of all the patterns in the STE (Chichilnisky, 2001). For visual neurons that act like an approximately linear filter (such as neurons in the retina, the LGN, and "simple cells" in the primary visual cortex V1), it can reveal the receptive field structure (J. P. Jones & Palmer, 1987). However, the Spike-Triggered Average fails when it is used to analyse a neuron whose receptive cannot be approximated by a single linear filter (Schwartz, Pillow, Rust, & Simoncelli, 2006). For example, "complex cells" in V1 are neurons that are selective to the orientation and the spatial frequency of a contrast defined pattern, but at least partially invariant to the phase (Dean & Tolhurst, 1983; Mechler & Ringach, 2002). This invariance to the spatial phase of the stimulus means that patterns in the STE (that are by definition defined over space) cancel each other out.

For neurons like complex cells, whose behaviour relative to a stimulus cannot by described in terms of how similar the stimulus is to a single linear filter, may instead be described in terms of the pattern of correlations between pixels in the stimulus (Simoncelli, Pillow, Paninski, & Schwartz,

2004). Simply taking the average disregards the correlation information present in the STE, however other reverse correlation techniques can use this information to characterise nonlinear selectivity. Spike-Triggered Covariance (Rust et al., 2005; J. Touryan, Lau, & Dan, 2002) is the most commonly used method for finding correlation structure of the STE. Spike-triggered covariance (STC) is performed by calculating the covariance matrix of the STE, and then extracting the eigenvectors corresponding to the larger eigenvalues (Brenner, Bialek, & de Ruyter van Steveninck, 2000; De Ruyter van Steveninck & Bialek, 1988). It is an application of the Principal Components Analysis in multivariate statistics. Because we expect that neurons in DM have more complex selectivity than simple cells spike-triggered averaging is likely inappropriate, but spike-triggered covariance may be a fruitful approach.

3.1.2 The Linear-Nonlinear-Poisson model

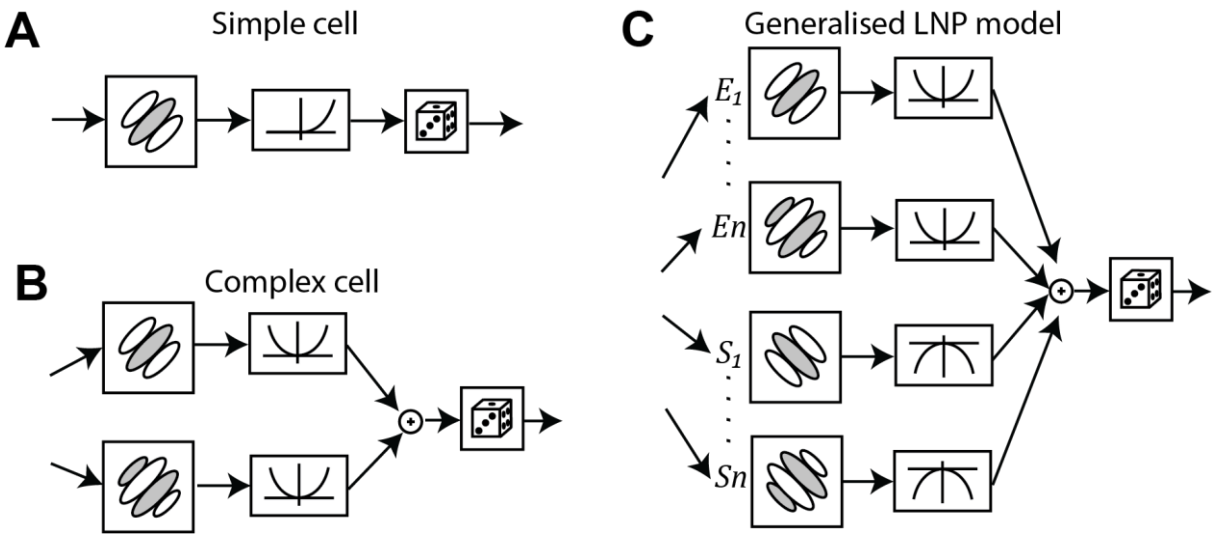

Figure 3.1. Functional models for V1 neurons. A: A model for simple cell function including a receptive field that is represented by a single orientated linear filter. In this model of simple cell function the output of this filter is half-squared (Heeger, 1992) and a Poisson process is used to covert the continuous firing rate into a series of discrete spike times. B: The "energy model" of complex cell involves summing the squared outputs of two Gabor filters that are tuned for the same spatial frequency and orientation, but with sinusoidal components that are 90° out of phase (Adelson & Bergen, 1985). The firing rate produced by this process is independent to the spatial location of contours within the neuron's receptive field. A Poisson process can be used to convert the firing rate into a series of discrete spike times. C: The generalized Linear-Nonlinear Poisson (LNP) response model extends the idea of modelling the firing rate by nonlinearly combining the outputs of linear filters to any number of both excitatory (E_1 to E_n) and suppressive (S_1 to S_n) filters.

In additional to being a form of dimensionality reduction, the STA/STC analysis (as well as iSTAC discussed below) can also understood as a method for explaining the responses of a neuron in the form of a computational model called the generalised Linear-Nonlinear Poisson (LNP) model. This model relates the receptive field structure to a modelled spike rate by the combination of three distinct stages. For example for an idealised simple cell (Figure 3.1a), the three stages are as follows. First, the image in the neuron's receptive field is convolved with a single linear filter, represented by

pattern of ovals representing alternating orientated on and off regions. The output of this stage, a measure of how similar the image on the retina is to the linear filter, is rectified by half-squaring to give the cell's instantaneous firing rate (Heeger, 1992). Finally, a Poisson process is used to covert the continuous firing rate into a series of discrete spike times. The "energy model" of complex cell (Figure 3.1b) sums the squared outputs of two Gabor filters that are tuned for the same spatial frequency and orientation, but with sinusoidal components that are 90° out of phase (Adelson & Bergen, 1985). The firing rate produced by this process is insensitive to the spatial location of contours within the neuron's receptive field, and is thus phase invariant.

The LNP model extends the idea of nonlinearly combining the outputs linear filters to any number of both excitatory and suppressive filters (Figure 3.1c; Rust et al., 2005) . In the LNP model the response of a neuron is driven by the convolution of the image in the neuron's RF with a set of statistically independent linear filters. The output of each filter is a measure of how similar the image on the retina is to the linear filter. By combining the responses to multiple filters non-linearly, selectivity for combinations of patterns can be modelled. In the LNP model, possible non-linear functions applied to the output of a filter can be monotonic excitatory functions like that of a simple cell (Figure 3.1a), quadratic excitatory functions like the non-linearity applied to the quadrature pair in Adelson and Bergen's (1985) "energy model" ('U'-shaped Figure 3.1b) or quadratic suppressive functions (inverted 'U'-shaped in Figure 3.1c).

Beyond V1, in areas where we expect a neuron to be sensitive to even more complicated spatial information, we expect that even more filters may be needed to capture the selectivity of the cells. If a neuron's response can be characterised by the nonlinear combination of a set of linear filters it can be fit using the LNP model and the eigenvectors returned by STC analysis will be unbiased estimations of these filters (Schwartz et al., 2006). STC analysis estimates the linear component of a

LNP model neuron, but for a full functional model further analysis is needed to estimate the static non-linearities associated with each linear filter. This will be achieved using a variant of the spike-triggered analysis called ‘‘information-theoretic spike-triggered average and covariance’’ analysis (iSTAC; Pillow and Simoncelli (2006)) to produce LNP models for each measured neuron.

Similar to STA/STC, iSTAC is a dimension reduction technique that looks for a linear subspace (that is, a set of orthogonal bases) in a high-dimensional space that span the relevant features. The primary difference between STA/STC and iSTAC is that iSTAC is derived from information theory (Paninski, 2004; Shannon, 1948). The method uses an iterative gradient-descent method to look for dimensions where the statistical distributions of STE are maximally different from the distribution of the raw stimulus, measured by an information theoretical quantity called the Kullback–Leibler (KL) divergence (Thomas & Cover, 1991).

iSTAC belongs to a family of information-theory-based methods for finding “maximally informative dimensions” (Paninski, 2004; I. Park & Pillow, 2011; Sharpee, Rust, & Bialek, 2004). Whereas previous formulations of the concept tend to be computationally demanding or even intractable in practice, iSTAC reduces the computational complexity by imposing a Gaussian assumption on the distributions. The algorithm therefore operates on mean vectors and covariance matrices, in a way similar to STA/STC.

Compared to STA/STC, iSTAC offers several advantages:

1. The STA and the STC have different geometrical and statistical meanings, making it difficult to compare the STA and STC as feature filters. In iSTAC, the two are unified in a single framework.

2. The “magnitude” of the STC eigenvectors is reflected by eigenvalues, whose meaning can be difficult to interpret. In iSTAC, a metrical measure based on information theory, the KL divergence, is used. The KL divergence is measured in an absolute unit (bits). It is therefore straightforward to compare the performance of iSTAC across different cases.
3. iSTAC offers a straightforward procedure for generating a parametric model (called the ratio of Gaussians, or ROG) for predicting the responses of a neuron to an arbitrary stimulus. The parametric nature of the model makes it possible to apply a technique for finding direction of invariance to the results of iSTAC. In STA/STC, a more ad-hoc, non- parametric procedure is usually needed (Schwartz, Chichilnisky, & Simoncelli, 2001).

3.1.3 Invariance of selective responses

Linear nonlinear Poisson models predict how a neuron responds to stimuli based on their similarity to a set of filters. However, these models do not explicitly make clear the relationships between the patterns in the filters and the selectivity and invariances of the responses. This is because a feature that a neuron is selective to is rarely a specific combination of “pixels” that make up the visual input. It can be more abstract. For example, imagine a visual neuron that is selective for the lowercase letter “a” in its RF. This cell needs to respond equally to the letter “a” independent of its exact location, size, colour, font or style and the properties of the background context. To be selective to the essential feature of the letter a in a way that is not influenced by incidental features (such as location, size, colour, font… etc.), neurons in the visual cortex need to disregard irrelevant information. In other words, the responses of the neurons need to be invariant to irrelevant parameters. Feature invariance is traditionally studied by analysing the responses of a neuron to stimuli under a set of geometrical transformations. Size invariance, for example, can be examined by studying of the responses of a neuron to scaled-up or scaled-down versions of the same pattern. If the responses remain relatively constant, the neuron is said to be size invariant (Cadieu et al., 2007).

Neurons in higher-order visual areas are believed to be invariant to more complex transformations. Early in the ventral visual stream, simple cells in V1 are very sensitive to the exact location, spatial phase and size of stimuli presented. In complex cells, the tuned responses remain selective for orientation but can tolerate changes to the spatial phase of the stimuli. Beyond complex cells in V1 there is less experimental literature with a focus on the invariance, as opposed to selectivity, of tuned responses (Riesenhuber & Poggio, 1999; Serre, Wolf, Bileschi, Riesenhuber, & Poggio, 2007). This may be due to the increased complexity of invariance in higher-order areas. Rust and Dicarlo (2010) compared the both the population activity and single unit responses in areas V4 and IT in macaques in response to the same set of controlled image stimuli. Ten objects were presented under six different transformed conditions defined by changes in scale, position and contrast. As expected the responses of the V4 neurons were much more sensitive to these changes, especially changes in position and context.

For investigating invariances in the visual system, choosing a small number of geometrical transformations to test has been the most common strategy (Bruce, Desimone, & Gross, 1981). However, this method requires a good hypothesis about the transformations that the neuron is invariant to. Although invariances to 2D and 3D rigid affine transformations are believed to be important for the efficient encoding of visual information, this remains a convenient assumption. It is entirely conceivable that visual neurons are invariant to other types of transformations that are not as easily parameterised as geometrical operations. For example, a neuron might be invariant to nonlinear morphing.

The white noise method characterises the responses of a neuron by systematic exploration of the stimulus space, and therefore offers the possibility of studying invariance without making strong prior assumptions about what transformation a neuron might be invariant to. In some conditions, the results of white noise analysis can be interpreted as characterisation of response invariance. For example, for complex cells in V1 (J. Touryan et al., 2002), the basis of the low-dimensional subspace typically forms quadrature pairs of Gabor filters. Assuming squaring nonlinearity, the responses can be interpreted as invariant to the spatial phase of Fourier components (because the squares of sine and cosine sum to 1).

Here I will describe the response properties of single neurons in marmoset DM by using iSTAC analysis with novel modifications to show that neurons in DM are primarily driven by long contours in contrast with a specific orientation and spatial frequency, but invariant to the exact location. This is similar behaviour to V1 complex cells, but over much larger areas in the visual field.

3.2 Methods

Data presented here are from electrophysiological recordings from 8 anaesthetised adult marmoset monkeys (*Callithrix jacchus*) with recordings made from either a Utah array inserted into DM, or with a linear probe into DM, V1 or both (Table 1). Recording details and the visual stimulation are described in the general methods (Chapter 2). Briefly, spiking activity was recorded from neurons in DM while the marmoset was presented with a spatiotemporally white noise visual stimulus.

	DM		V1	
Animal	Utah	Linear	Utah	Linear
CJ138	X			
CJ139	X			
CJ140	X			
CJ158	X			
CJ165		X		
CJ166				X
CJ169		X		X
CJ178		X		X

.

3.2.1 Masked estimator

A modified version of iSTAC (Pillow and Simoncelli (2006)) was used to determine the locations of the receptive fields, and to characterise their feature selectivity. The true covariance matrix is typically estimated by calculating the sample covariance matrix:

$$\Lambda_S = \frac{1}{N-1}\sum_{n=1}^{N}(\vec{s}(t_n) - \mu_S)(\vec{s}(t_n) - \mu_S)^T \quad (1)$$

where t_n is the time of the nth spike, $\vec{s}(t_n)$ is a vector of pixels in the stimuli presented a short, fixed time (τ) before the spike, N is the total number of spikes, and μ_S is the sample mean of $\vec{s}(t_n)$ associated with all spikes. A more efficient estimation can be achieved by giving weights to entries of Λ_S, according to the prior knowledge about the structure of the covariance matrix (R. Y. Chen, Gittens, & Tropp, 2012). Using the Hadamard product (symbolised by ∘) to denote entry-wise multiplication, the masked estimator is

$$\Lambda_M = M \circ \Lambda_S \quad (2)$$

where M, the mask, has the same dimension as Λ_S. Assuming that correlation between two pixels decreases with their on-screen distance, entries in M are determined by a Gaussian distribution

$$m_{i,j} = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{{d_{i,j}}^2}{2\sigma}} \quad (3)$$

where $d_{i,j}$ is the distance (in pixels) between the on screen locations of elements in corresponding (i,j)-entry in Λ_S. The standard deviation σ was set to a value that translates to 5° in the visual space. The value was chosen to reflect the prior knowledge about the radius of the receptive fields.

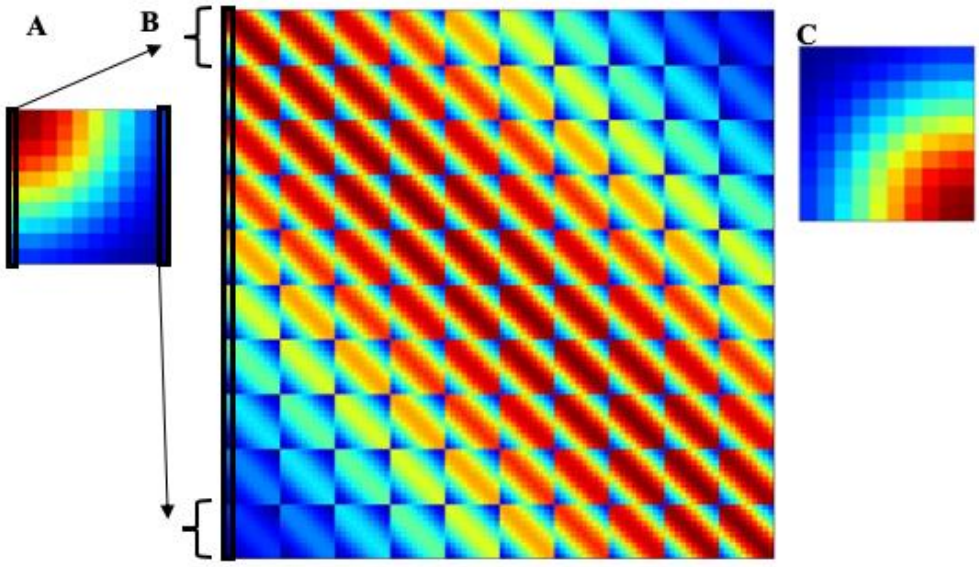

Figure 3.2. A 2D Gaussian function of the distance from the first pixel in the top-left corner (A) forms the first column of M (B). Reshaping the 2D image in this way causes discontinuities where the ends of columns are stitched together. This gives the resulting mask estimator M a blocky appearance. Values within each column (i) of M are a function of the distance (i) of each other pixel, progressing in column major order until the bottom-rightmost pixel (C).

Figure 3.2 illustrates the values in M. Because $\vec{s}(t_n)$ has been reshaped into a single column vector from a 2D image, there are discontinuities in the distance as a function of pixel index. The elements in the reshaped column vector corresponded to the pixels on the screen in column major order.

$$d_{i,j} = \sqrt{(x_i - x_j)^2 + (y_i + y_j)^2} \tag{4}$$

M is a symmetric matrix, with a structure similar to the "banded" mask-estimator in the statistics literature (Cai, Ren, & Zhao, 2016). I used the masked estimator to estimate both the covariance matrix of the raw stimulus and the covariance matrix of the STE (see below).

3.2.2 iSTAC Procedure

For two Gaussian distributions P and Q, where P is assumed to have zero mean and identity covariance, and Q is assumed to have the mean vector μ and the covariance matrix Λ, the Kullback-Leibler (KL) divergence between P and Q is

$$D(Q,P) = \frac{1}{2}(Tr(\Lambda) - \log|\Lambda| + \mu\mu^T - n) \quad (5)$$

where $Tr(\cdot)$and $|\cdot|$ are the trace and the determinant of a matrix, and *n* is the dimensionality of data. iSTAC searches for a subspace where the KL divergence between the raw stimulus (P) and the STE (Q) within the subspace is maximal. The subspace is represented by the matrix *B*, whose *m* columns form an orthonormal basis for the subspace. In other words, the objective function is:

$$D_B(Q,P) = \frac{1}{2}(Tr[B^T(\Lambda + \mu\mu^T)B] - \log|B^T(\Lambda)B| - m) \quad (6)$$

The gradient (used in iterative optimisation, see below) of this objective function is:

$$\frac{dD(Q,P)}{dB} = 2(\Lambda + \mu\mu^T - \Lambda B B^T \Lambda^{-1})B \quad (7)$$

In the case of a stimulus that has a non-Gaussian distribution, such as our tertiary white noise, a "whitening" transformation is needed to transform the distribution of P into a zero-mean, unit covariance distribution. Following whitening, the distribution of the STE (Q) is parameterised by:

$$\Lambda = \Lambda_0^{-\frac{1}{2}} \Lambda_1 \Lambda_0^{-\frac{1}{2}T} \quad (8)$$

and

$$\mu = \Lambda_0^{-\frac{1}{2}}(\mu_1 - \mu_0) \quad (9)$$

where (μ_0, Λ_0) and (μ_1, Λ_1) are the parameters for the raw stimulus (P) and the STE (Q) prior to whitening.

During optimization, basis vectors are searched iteratively starting with one maximally informative dimension, and by subsequently adding more basis vectors as columns in B such that the KL divergence is maximized. To increase the efficiency of the search, the newly added basis vector in B is initialised to multiple starting points, which include the STA and the largest and the smallest eigenvectors of the STC.

Since the objective function is evaluated thousands of times in the iterative optimisation process, a computationally more efficient form of the objective function is used:

$$\log|B^T \Lambda B| - Tr[B^T \Lambda B] - (B^T \mu)^2 \qquad (10)$$

In this formulation, the evaluation of the objective function can recycle the values used in the previous step (before a new basis in B is added). This form also useful for applying the masked estimator (see below).

Optimization was performed using MATLAB's constrained nonlinear optimisation function fmincon() function, which searches for a new basis vector under the constraint that the new vector is orthonormal to all previously found vectors. In each iteration, the maximally informative KL was evaluated using the new and previously found dimensions. Through this, the iSTAC algorithm takes into account correlations between the projection of the spike-triggered stimuli onto the new dimension and the previously found dimensions.

The set of orthonormal basis vectors produced by the iSTAC method are vectors in the same space as the distributions P and Q and not the 'stimulus space' of the STE. As such, the whitening step needs to be reversed according to:

$$B_1 = \Lambda_0^{-\frac{1}{2}T} B \tag{11}$$

Which gives B_1, a basis for the feature space found by iSTAC where rows correspond to location of pixels on screen.

3.2.3 iSTAC with masked covariance estimator

The covariance matrix of the STE, after the application of the mask and whitening, is:

$$\Lambda = \ (M \circ \Lambda_0)^{-\frac{1}{2}} (M \circ \Lambda_1) \ (M \circ \Lambda_0)^{-\frac{1}{2}T} \tag{12}$$

The masked estimator is applied to covariance matrices. It would therefore appear that the mean of the STE (μ) is not affected by the mask. However, it should be noted that the objective function does not refer to μ, but to $\mu\mu^T$. Since it is an estimate of the covariance of μ, the mask also needs to be used. Without this correction, the importance of the STA is artificially inflated by iSTAC, because entries in other covariance matrices are reduced by the mask. After whitening, the masked estimator of $\mu\mu^T$ is:

$$\mu\mu^T = \ (M \circ \Lambda_0)^{-\frac{1}{2}} [\, M \circ (\mu_1 - \mu_0)(\mu_1 - \mu_0)^T] \ (M \circ \ \Lambda_0)^{-\frac{1}{2}T} \tag{13}$$

3.2.4 Interpreting KL divergence

One key benefit of using iSTAC is that KL divergence is measured in the unit of bits, which can be interpreted in a straightforward way in the framework of information theory. More precisely, the goal of iSTAC is to maximise $D(Q, P)$, the KL divergence between the raw stimulus (P) and the STE(Q). The measure can be more intuitively interpreted in light of the following equation:

$$H(Q, P) = H(Q) + D(Q, P) \tag{14}$$

In this equation, H(Q,P) is the cross entropy between Q and P, which is the average number of bits that is needed to encode a pattern in the STE (Q), if the distribution of the raw stimulus (P) is used. This value is the sum of the entropy of the STE (the average number of bits needed to encode the STE using its true distribution), and the KL divergence between the two distributions.

The optimal average code length needed to encode a pattern in the white noise stimulus that causes a neuron to fire is the entropy of the STE, that is is $H(Q)$. In a situation where the distribution of the raw stimulus, instead of the true distribution, is used, the code length must increase from $H(Q)$ because the code is suboptimal. The penalty associated is the difference between $H(Q,P)$, and $H(Q)$, which is the KL divergence. iSTAC seeks to maximise this penalty, because its goal is to look for dimensions where the difference between the raw stimulus and the STE are as large as possible.

When iSTAC is applied to my data, the maximal KL divergence returned is usually a very small number, which might appear to be unexpected for the high entropy white noise stimulus. This is, however, expected because the KL divergence represents the difference between STE and the raw stimulus. Because the results of iSTAC indicate that the features encoded by neurons are typically captured by a low dimensional subspace, the distributions are identical in almost all of the dimensions, which therefore produce zero KL divergence. The value returned from iSTAC is therefore the KL divergence of a small number of Gaussian distributions.

3.2.5 Temporal dynamics

The current analysis concentrates of the spatial aspect of feature selectivity and ignores the temporal dynamics. The analysis therefore looks for the spatial features that cause a neuron to fire at a certain time delay τ after the onset of the stimulus. The correlation time τ is determined for individual neuron by searching for τ that produces the largest KL-divergence for the first dimension. For V1 the delay search was initialised at a multiple shorter time points (30ms, 40ms and 50ms) than DM (50ms) and limited to a range between 0 and 120ms at 5ms intervals.

This range was chosen based on preliminary results where some V1 cells were found to have a delay of up to 220ms when unrestricted which is out of line with what we would expect in the very early area. These responses typically produced filters which were spatially unstructured despite having a large KL-divergence. This effect could have been due to a greater chance of finding the highest KL-divergence due to noise with a larger number of tested time points or possibly due to the accumulation of information in V1 from feedback information over time.

In this case, iSTAC analysis is unlikely to be a useful method as feedback information can encode spatial information that cannot be easily described in terms of point to point covariance such as surround modulation. The feedback issue especially important for or comparisons between DM and V1 because we would expect V1 to receive much more feedback information than DM under anaesthetised conditions. The limits on the time delay allows for comparisons between the filters that can be expected to be produced predominantly from feedforward information in both V1 and DM.

3.2.6 Cropped estimation

After the masked estimator has found the receptive field at optimal τ, a Gabor function is fitted to the receptive field to estimate its location and geometry. A small region encompassing the entire receptive field is selected. iSTAC is then used again on this cropped stimulus space (without the mask).

3.2.7 Ratio of Gaussians model

In addition to being a method of dimension reduction, iSTAC also produces a model that predicts the response of a neuron to arbitrary stimulus x, taking the form of the ratio of two gaussians (ROG):

$$P(spike|x) = \alpha \frac{\hat{Q}(x)}{\hat{P}(x)} \tag{15}$$

where $\hat{Q}(x)$ is the estimated Gaussian distribution for the STE, evaluated at x, and $\hat{P}(x)$ is the estimated Gaussian distribution of the raw stimulus, evaluated at x, and α is a constant. It can be written as an exponential function applied to a quadratic form

$$P(spike|x) = ae^{x^t Mx + b^t x} \tag{16}$$

where $M = \frac{1}{2}(I - \widehat{\Lambda}^{-1})$, and $b = \widehat{\Lambda}^{-1}\hat{\mu}$, $\widehat{\Lambda} = B^T \Lambda B$ are the parameters of $\hat{Q}$.

The static nonlinearity for each dimension can be calculated as the marginal distribution of $P(spike|x)$ along that dimension.

3.2.8 Bootstrapping

A novel bootstrapping procedure for testing the statistical significance of the basis vectors found by iSTAC was developed. Computational simulations of the responses of white noise of a LNP model showed that the procedure was able to correctly recover the dimensionality of the subspace.

Assume that iSTAC has found a maximally informative subspace of dimension k, the bootstrapping procedure aims to establish the variations in the amount of information gained with the addition of the (k+1)th dimension, if the addition of the (k+1)th dimension is simply due to sampling error (that is, uninformed by the activities of the neuron).

To achieve this goal, a new distribution parameterised by $\overline{\mu_B}$ and $\overline{\Lambda_B}$ is formed. The new distribution has the same distribution as the STE on the k-dimensional subspace found by iSTAC, but on other dimensions, its distribution is estimated by a "shuffled" version of the STE, where the spike times are randomly shifted to remove their relationship to the stimulus.

The procedure first extends the k orthonormal basis found by iSTAC to a full set (dimension=n) of orthonormal basis B, using the Gram–Schmidt process. Two sets of parameters are then estimated (using the whitening and masked estimator described above): μ and Λ estimated from the STE, as well as $\tilde{\mu}$ and $\tilde{\Lambda}$ from the shuffled STE. The covariance matrices are then transformed to the coordinates of B:

$$\Lambda_B = B^T \Lambda B \tag{17}$$

$$\widetilde{\Lambda_B} = B^T \tilde{\Lambda} B \tag{18}$$

The distribution used for bootstraping is formed by combining blocks from Λ_B and $\widetilde{\Lambda_B}$. To describe this new distribution, first express Λ_B and $\widetilde{\Lambda_B}$ as blocks:

$$\Lambda_B = \begin{bmatrix} \Lambda_1 & \Lambda_3 \\ \Lambda_3 & \Lambda_2 \end{bmatrix} \tag{19}$$

$$\widetilde{\Lambda_B} = \begin{bmatrix} \widetilde{\Lambda_1} & \widetilde{\Lambda_3} \\ \widetilde{\Lambda_3} & \widetilde{\Lambda_2} \end{bmatrix} \tag{20}$$

where $\Lambda_1$and $\widetilde{\Lambda_1}$ are k-by-k matrices, representing the covariance of the STE and the shuffled STE on the subspace spanned by the basis found by iSTAC, and $\Lambda_2$and $\widetilde{\Lambda_2}$ are (n-k)-by-(n-k) matrices, representing the covariance on the rest of the dimensions.

The covariance matrix of the new distribution is set to:

$$\overline{\Lambda_B} = \begin{bmatrix} \Lambda_1 & 0 \\ 0 & \widetilde{\Lambda_2} \end{bmatrix} \tag{21}$$

The mean vector $\overline{\mu_B}$ of the bootstrapping distribution is calculated using a similar logic: the mean vector of the STE (μ_B) and the shuffled STE ($\widetilde{\mu_B}$) are first calculated. μ_B is then projected to the subspace spanned by the vectors found by iSTAC, and $\widetilde{\mu_B}$ is projected to the other dimensions. Finally, the two (orthogonal) projected vectors are added. The method was tested on simulated neurons with different numbers of significant dimensions.

3.2.9 Directions of Invariance

Invariances can be deduced based on the filters and nonlinearities recovered by an iSTAC analysis, but they can also be recovered explicitly. For each receptive field, iSTAC produces a Ratio of Gaussian (ROG) model, which predicts the responses of the neuron to an arbitrary stimulus. That model is used as the basis of the invariance analysis. Berkes & Wiskott (Berkes & Wiskott, 2006, 2007) introduced a general framework for characterising invariance. In their formulation, a direction of invariance *w* around a pattern *s* is a direction in the stimulus space, such that if the pattern is transformed in direction of *w*, the change in the responses is minimised. They showed that if the responses of a neuron can be modelled by a quadratic form (figure 3.3), the direction of invariance can be found by solving an eigenvalue decomposition problem by calculating the Hessian matrix (the second derivative). In fact, under this framework, a neuron can have multiple directions of invariance, ranked by how invariant it is to the transformations.

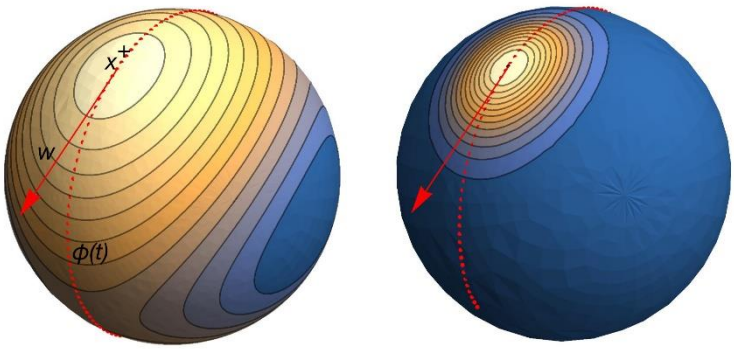

Figure 3.3. The direction of invariance, in a three dimensional case. The figure on the left shows the algorithm applied to neuronal responses modelled as quadratic form (the colour scale represents firing rate). The red vector is the first direction of invariance. The figure on the right illustrates the responses of the Ratio of Gaussian model produced by iSTAC. In both the quadratic form and Ratio of Gaussian models the same algorithm can be used to find the direction of invariance.

The method was introduced to analyse the behaviour of model neurons in artificial neural networks (Le et al., 2010) but is here used on the statistics of the STE. To quantify response invariance for units

in an artificial neural network, under the assumption that the responses of unit are characterized by quadratic forms:

$$q(x) = w^T Hw + w^T \cdot f + c \tag{22}$$

where x is a n-dimensional stimulus, H is a n-by-n symmetric matrix, and f is a n-dimensional vector. However, quadratic forms tend to produce units with very broad stimulus selectivity, which is unrealistic for cortical neurons. A more realistic descriptor is the ratio of Gaussian (ROG) model introduced by iSTAC, where the quadratic form is followed by an exponential nonlinearity:

$$g(x) = e^{q(x)} \tag{23}$$

Here, I demonstrate that the method described in Berkes & Wiskott (2006) can be used without modification for ROG.

The method of Berkes & Wiskott (2006) starts by locating the optimal stimulus x^+ for $g(x)$, under the constraint that the energy of the stimulus is constant (that is, x is on the sphere S: $|x| = r$). Since x^+ is also optimal for $g(x)$, the trust-region method described in Berkes & Wiskott (2006) can be used without modification.

Let $\tilde{g}(x)$ be $g(x)$ restricted to S, and let $\varphi(t)$ be a geodesic on S:

$$\varphi(t) = \cos\left(\frac{1}{r}\right)x^+ + \sin\left(\frac{1}{r}\right)r * w \tag{24}$$

where $\varphi(0) = x^+$, and w is a vector tangent to S at x^+. In Berkes & Wiskott (2006), a direction of invariance is defined as w that minimizes the second derivative of $\tilde{g} \circ \varphi(t)$.

Let $k = \tilde{g}(x^+)$. Using the chain rule, the second derivative can be derived:

$$\frac{d^2}{dt^2}(\tilde{g} \circ \varphi)\Big|_{t=0} = k(w^T Hw + \frac{1}{k^2}(\nabla g(x^+) \cdot w)^2 - \frac{1}{r^2} f \cdot x^+ - \frac{1}{r^2} x^{+T} H x^+) \tag{25}$$

To minimize this equation, note that the equation can be simplified by the constrained imposed by the first derivative:

$$\frac{d}{dt}(\tilde{g} \circ \varphi)\Big|_{t=0} = k(\nabla g(x^+) \cdot w) \tag{26}$$

Since $\tilde{g}$ is maximal at x^+, $k(\nabla g(x^+) \cdot w) = 0$. It follows $\nabla g(x^+) \cdot w = 0$. In other words, the gradient of g at x^+ is always perpendicular to the tangent plane.

Consequentially,

$$\frac{d^2}{dt^2}(\tilde{g} \circ \varphi)\Big|_{t=0} = k(w^T H w - \frac{1}{r^2} f \cdot x^+ - \frac{1}{r^2} x^{+T} H x^+) \tag{27}$$

This is exactly the same equation derived from quadratic models in Berkes & Wiskott (2006). The equation can therefore be minimized by the same eigenvalue decomposition procedure described in Berkes & Wiskott (2006).

3.3 Results

The initial dataset consisted of data recorded from area DM of four animals using a 96 channel "Utah" array. The data collected on "Utah" arrays were useful for comparing the responses of neurons across the cortical surface, because the electrodes are implanted at approximately the same depth in a 2D grid. Recordings were also made using the linear arrays in area DM and V1. This type of array records the neuronal activities across the depth of the cortex, allowing the study of receptive field properties in different layers of the cortex. Unless otherwise stated, examples from the following results section show data from single penetrations using a linear probe in both V1 and DM for a single case: CJ178.

In the results presented below, the masked estimator was used as a preliminary step to locate the coordinates of the receptive field. The first informative vector of the masked estimator was fitted to a 2D Gabor filter, which was then used to crop the patterns in the STE to further reduce the dimensionality of the stimulus. The right panel of figure 3.4 illustrates the most informative direction of the stimulus cropped to encompass the receptive field located by iSTAC using the mask estimator. The reduction of dimensionality due to cropping did change the output of iSTAC, because the most informative direction is now different from the one illustrated by the middle panel. The full analysis (see Figure 3.5) suggests that this is because the masked estimator underestimated the importance of the first component in Figure 3.5, and picked the second component as the most informative dimension.

3.3.1 Recovering receptive field location

The white noise stimulus used in my experiments covered a large part of the visual field (50° to 60°) to accommodate the large number of receptive fields associated with the simultaneously recorded neurons. Applying analytical methods based on second-order statistics (such as STC and iSTAC) to

this data requires the accurate estimation of the covariance matrix with tens of thousands of dimensions, which is typically untenable given the amount of data available. Furthermore, even in noise-free conditions, the numerical calculation of eigenvectors of very large matrices can also be inaccurate.

The left panel of Figure 3.4 illustrates the problem: if data analysis is performed directly on the patterns that were displayed to the animal, it is difficult to extract meaningful information about the receptive field (either with STC or with iSTAC) because the estimators, with thousands of parameters, are contaminated by noise. In this particular case, the most informative direction estimated by iSTAC (giving a relatively large KL divergence of 0.0174 bit) does not have any clear structure.

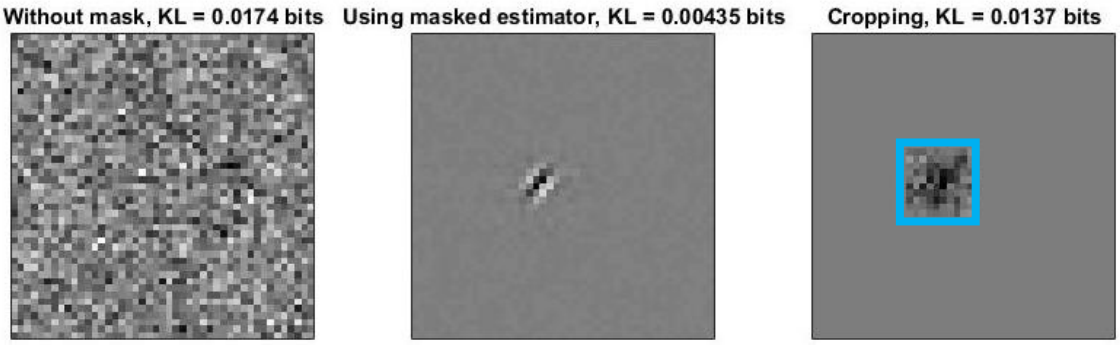

Figure 3.4. Locating receptive fields. The most informative vector dimension produced from iSTAC analysis on an exemplar DM neuron remapped to stimulus space (square elements with a width of 1° in the visual field). iSTAC analysis without using a masked estimator or cropping (left) returns a dimension in which the KL divergence between the spike triggered ensemble and the set of all stimuli is 0.0174 bits. However, this vector provides no information about the location of the neuron's receptive field. Using the masked estimator a dimension (centre) is found that has a lower KL divergence (0.0043 bits), but is sparser such that the most important dimensions are contained within a small area of the visual field. This vector was modelled as a 2D Gabor filter to estimate the location of the neuron's receptive field. The most informative dimension found within a cropped area of the stimulus space (right) has a KL divergence of 0.0137 bits.

If the location and the size of the receptive field is known, it is possible to crop the stimulus into a small spatial region that covers just the receptive field, and therefore reducing the dimensionality to a manageable degree. The stimulus included in analysis was cropped to a 10x10 stimulus element region (here 10°x10°), using the STA as an initial estimate of the location of the receptive field. However, this strategy proved to be impractical on a large scale, because the STA is not always a reliable estimator of the geometry of the receptive field, and in many cases it is close to zero.

To achieve this, I relied on prior knowledge that the receptive field is small (relative to the size of the monitor) and localised in space. I applied a "mask" to the estimated covariance matrix such that entries are weighted using a function that assigns smaller weights to pairs of pixels that are farther apart on the monitor according to a Gaussian distribution (see 3.2.1). This technique is equivalent to cropping the stimulus to all possible regions of a predefined size, because the covariance of all pixel pairs exceeding ten degrees on the monitor is effectively set to zero.

The method is related to a family of statistical techniques called the masked sample estimator, which was introduced in recent years for the estimation of very large sparse covariance matrices, (Cai et al., 2016; R. Y. Chen et al., 2012). The masked estimator increases the efficiency of estimating large covariance matrices by incorporating the prior knowledge that most of the entries are close to zero. Deviations from this prior knowledge are therefore mostly due to noise and are suppressed.

The middle panel of Figure 3.4 illustrates the most informative direction returned by iSTAC using the masked estimator, on an exemplar DM neuron. Although the KL divergence is significantly smaller

(0.00435 bits instead of 0.0174 bits), a clear receptive field was found, even though the covariance matrix had the same dimensionality as in the left panel.

This shows that by using a masked estimator, iSTAC can be used to locate the receptive fields of DM neurons where otherwise the structure in the covariance of the STE would be hidden by noise. This is most important for locating the receptive fields of neurons without a strong monotonic filter such as many of the ones found in DM. For these cells, the STA cannot be used to find the receptive field location and the use of the masked estimator becomes critical for further analysis.

3.3.2 Multiple linear-nonlinear components of DM feature selectivity

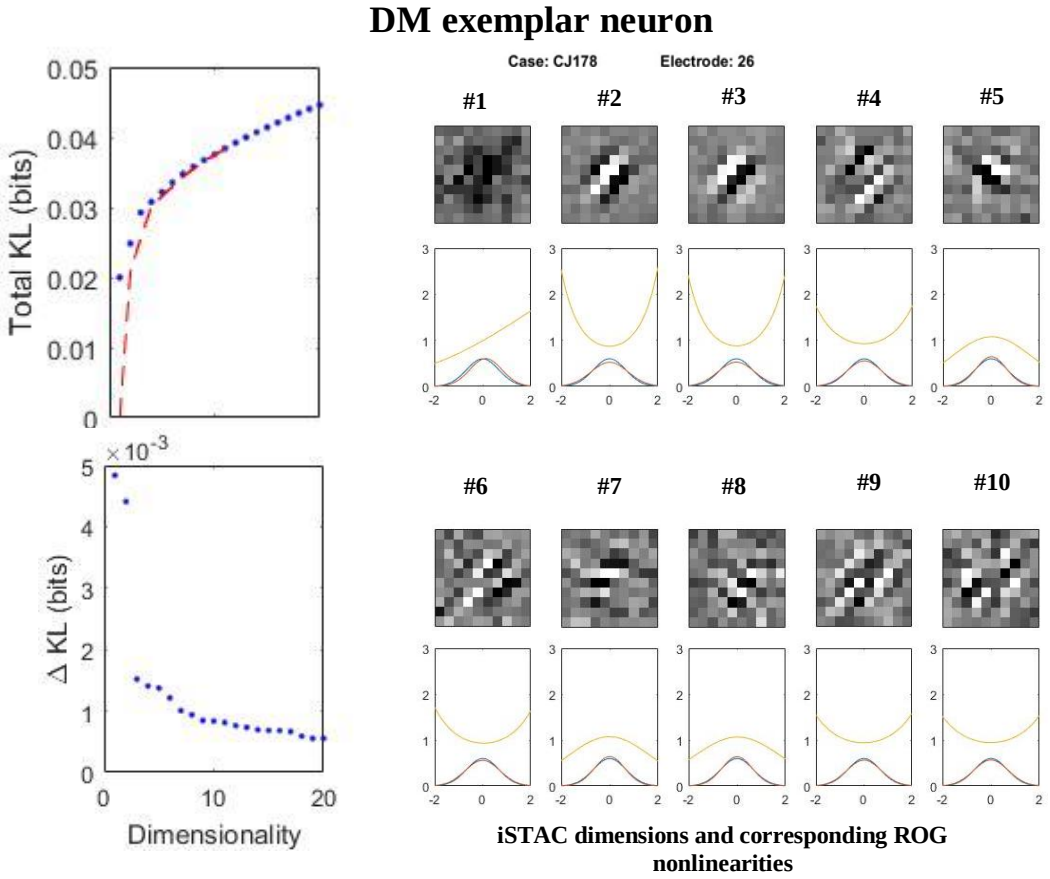

Figure 3.5. Right: The ten most informative vector dimensions produced from iSTAC analysis on an exemplar DM neuron remapped to stimulus space (square elements with a width of 1° in the visual field) and corresponding 1D non-linearity (underneath) for each vector as estimated by the Ratio of Gaussian's model (green curves) from the ratio of Gaussian's with mean and covariance of the spike triggered ensemble (red) and the set of all stimuli (blue). The model suggests that the first vector drives the neuron in a linear fashion, whereas the 2nd, 3rd,4th,6th, 9th and 10th are excitatory and that the 5th, 7th and 8th are suppressive. Top left: the total KL between a set of maximally informative orthonormal basis vectors and the set of raw stimuli increases as a function of dimensionality, bootstrapping confidence interval (very small) in red. The gradient of the total KL divergence increase quickly decays (bottom left) as there are few informative directions.

An example of the results produced by the modified iSTAC method is shown in Figure 3.5. These results are for the same exemplar DM neuron as in Figure 3.4. After finding the first most informative dimension, the modified iSTAC algorithm searches for a second vector, orthogonal to the first, such that the subspace defined by the two basis vectors maximises the KL divergence

between the STE and raw stimuli. It's important to note that the vectors are not found independently, instead additional orthonormal vectors are iteratively found such that together they form a basis for the "maximally informative" subspace. In the upper left panel of Figure 3.5 the total KL divergence (in blue) between the STE and Raw stimuli increases monotonically with each increase in dimensionality.

The results of the bootstrapping significance test are shown in red in the upper-left panel of Figure 3.5. The confidence interval was estimated by setting the mean and covariance in the previously found dimensions to be the same, and finding a window for the change in total KL for the next dimension due to randomly sampling noise, uninformed by spike times. This was accomplished by randomly shifting the spike times to remove the relationship between spike times and the stimulus. If the true value lies outside of a 95% confidence interval generated by noise, it is taken to be significant. For more details on the bootstrapping significance testing, see methods in (3.2.8). For this exemplar neuron, the bootstrapping significance testing found 11 dimensions to be significant.

Because the dimensions are found in order of most to least informative the addition of each subsequent dimension results in a smaller increase to the total KL. In other words, the gradient of the total KL as a function on dimensionality monotonically decreases. As such, it can be useful to consider the relative contribution of each dimension (lower left panel). The first two dimensions are very informative compared to the rest; the KL contribution of the third vector is less than half that of the second.

The vectors found by iSTAC analysis that form the maximally informative subspace can be remapped to stimulus space. In Figure 3.5, the panel labelled '#1' illustrates the first most maximally

informative iSTAC vector (same as in Figure 3.4, right). As discussed in the introduction, these vectors can be thought of as linear filters in an LNP model neuron. In addition to estimates of the linear filters, iSTAC analysis can be used to estimate the static non-linear functions associated with these filters, by taking the ratio of the Gaussian distributions estimated for the STE and the raw stimulus, restricted to each dimension. The process is illustrated by the blue and the red curves in Figure 3.5, which plot the estimated distributions of the raw stimulus and the STE. The ratio of the two Gaussian, plotted in yellow, is the estimated static-nonlinearity.

Three types of static nonlinearity were found: The static non-linearity for the most informative dimension is monotonic and close to linear, similar to the non-linear function of an idealised V1 simple cell. This suggests the first filter drives the neuron in a linear fashion. The second and third filters are associated with very similar excitatory nonlinearities. They increase with larger projection and they are symmetrical around 0 ("U"-shape). This implies that for this model an image similar to this filter, or its inverse in contrast, would excite behaviour in the modelled neuron. Both the nonlinearities and the spatial structure in pixel space of these pair of filters closely resemble a quadrature pair from Adelson and Bergen's energy model (Adelson and Bergen, 1985). This is similar to an idealised V1 complex cell with a spatially invariant orientation tuned response. The third type of nonlinearity is suppressive (dimension #5, #7, #8), which decreases with larger projected values (inverted "U"-shape). According to the shapes of the nonlinear functions, each dimension can be classified as monotonic, excitatory, and suppressive.

3.3.3 Multiple linear-nonlinear components of V1 feature selectivity

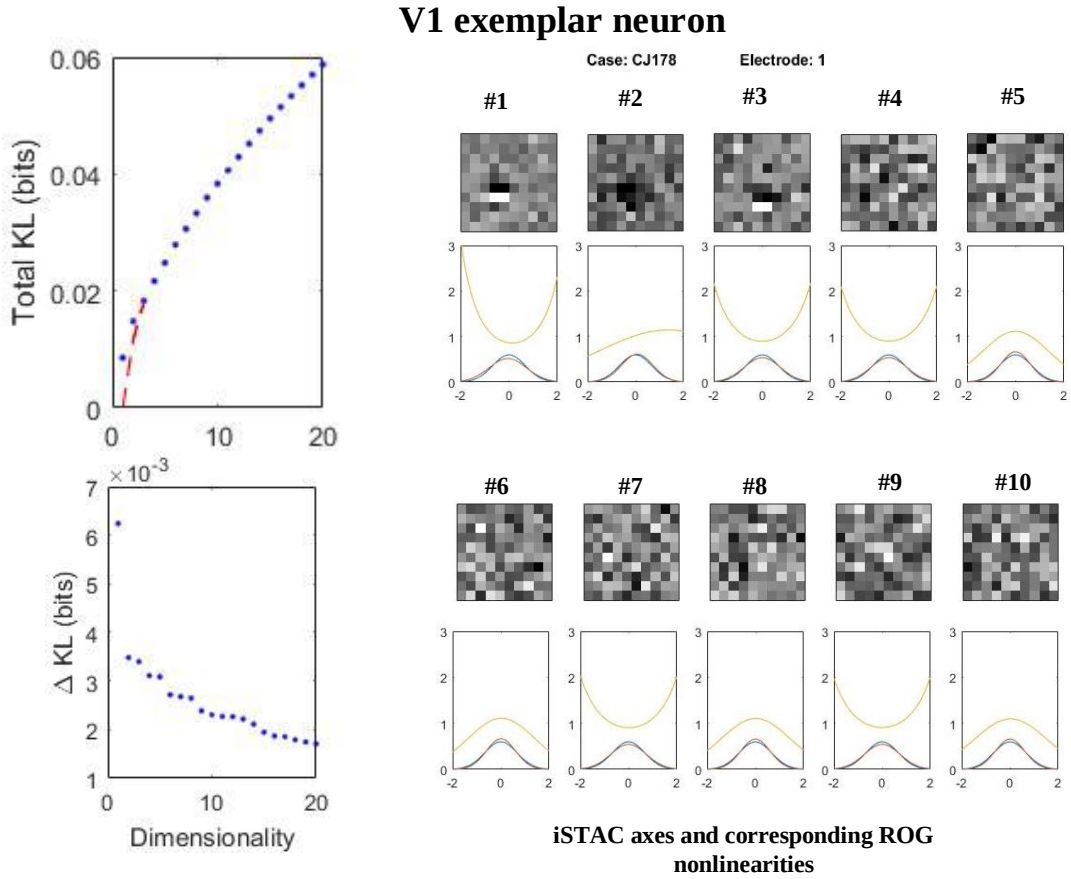

Figure 3.6. Same as Figure 3.5 for an exemplar V1 neuron where square elements had width of 0.4° in the visual field.

Figure 3.6 illustrates the results of the same analysis for an exemplar neuron in V1. The relationship between the first three iSTAC vectors for this V1 neuron are very similar to that of the first three iSTAC dimensions of the DM neuron in Figure 3.5. In both cases, the first three filters include an approximate quadrature pair as well as a filter that drives the response (approximately) linearly. In general for the first 10 vectors for both the exemplar V1 and DM neurons, the non-linear functions are also very similar, mostly either excitatory or suppressive and close to symmetrical around zero.

For this V1 neuron, the bootstrapping significance testing yielded only two significant iSTAC dimensions. This fits with the expectation that the response of V1 neurons should be driven by simpler visual features than DM neurons. However, despite not being statistically significant, the third iSTAC vector of the exemplar V1 neuron appears to be approximately a quadrature pair with the first vector.

3.3.4 ISTAC analysis over DM surface

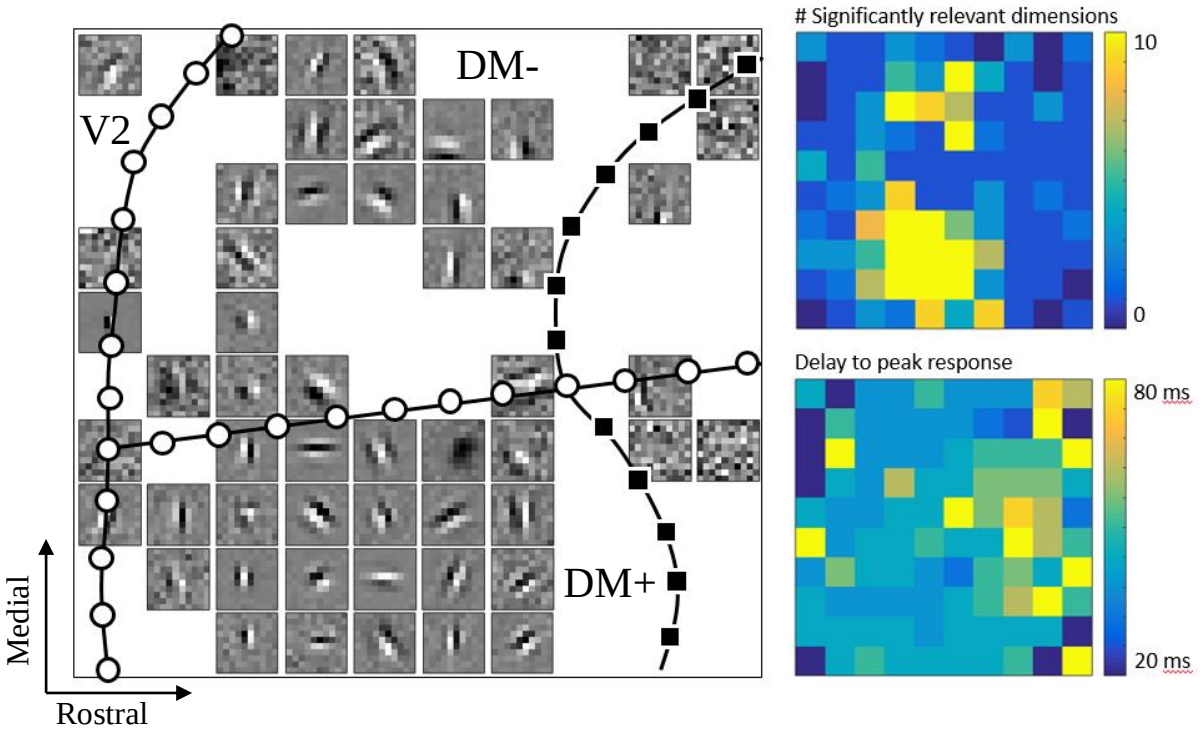

Figure 3.7. iSTAC analysis of neurons sampled with a 'Utah'-array reveals similar filters over upper and lower field DM. The number of significantly relevant dimensions (upper right) and optimal time-point (lower right)

3.3.5 ISTAC filters in V1 and DM

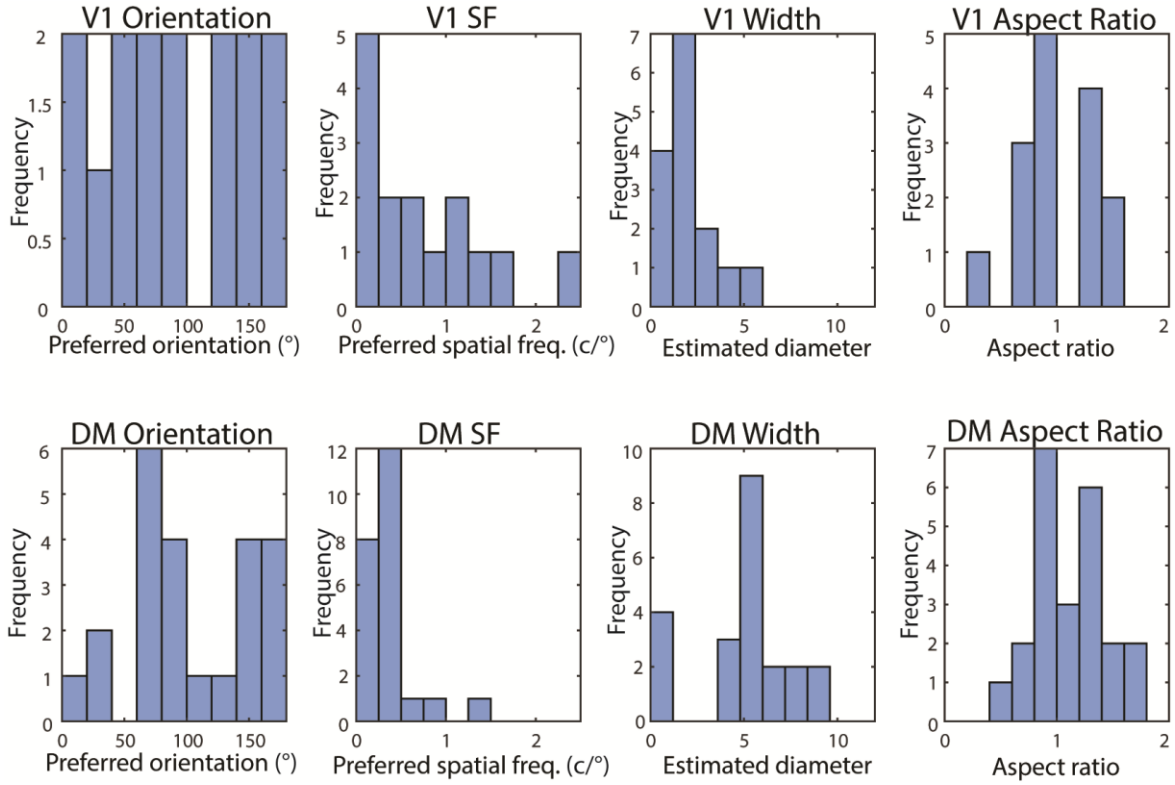

Figure 3.8. A comparison of filters for electrodes in V1 and DM. These are some of the parameters obtained from fitting a 2D Gabor filter to the first significant excitatory iSTAC vector where one exists. For this comparison, the data was collated from the three penetrations (each) in V1 (15 fitted filters) and DM (23 fitted filters) using the "Neuronexus" linear probe array. Distribution of parameters for orientation (left column), spatial frequency (middle-left column), diameter (middle-right column) and aspect ratio (right column) for V1 (top row) and DM (lower row) where these vectors were well fit by the Gabor model (sum of square errors>0.8). Diameter was taken as 4 times the sigma, 95%, of the Gaussian envelope of the Gabor filter.

Figure 3.8 shows comparisons between the distributions of some key parameters obtained from 2d Gabor models fitted to the first significant nonlinear excitatory iSTAC vectors on electrodes in V1 and DM. While many of these vectors are qualitatively similar in that they resemble and can be fit by models of 2d Gabor filters, the vectors from each area have different statistics reflecting the differences in their tuning. The orientation tuning results however, need to be understood in the context of how they were obtained. The linear probes were implanted as close to perpendicular to the surface of cortex as possible, however due to the curvature of cortex it is not possible to easily

sample from a single orientation column without many penetrations. In this case, we did not make a sustained effort to find orientation columns with the linear probes, firstly because this was not a primary aim of this experiment, but also because to avoid doing undue damage to the cortex with repeated insertions of the linear array. As such, the finding a roughly even distribution of preferred orientations in V1 and an uneven distribution is not unexpected but probably not representative of the overall distribution in the area.

Comparing the distributions of the other parameters of the Gabor models fitted to iSTAC vectors in V1 and DM, the main difference between the two areas is the scale of the filters. The spatial frequencies found for filters in V1 are on average much higher than those in DM (V1 mean: 0.7409 cycles/° with a standard deviation of 0.6796 cycles/°, DM mean: 0.3570 cycles/° with a standard deviation of 0.2712 cycles/°). The width of the filters along their primary axes are also much smaller in V1 than in DM (V1 mean: 2.08° with a standard deviation of 1.19°, DM mean: 5.31° with a standard deviation of 2.82°) and their aspect ratios are similar (V1 mean: 1.02 with a standard deviation 0.36, DM mean: 1.11 with a standard deviation of 0.32). This is consistant with the idea that a large part of what drives responses in DM is similar to what drives response of neurons in V1, but over a smaller spatial scale.

A potential future direction for extending this white noise analysis could be to investigate whether the receptive fields of DM cells can be modelled as the integration of many smaller V1 receptive fields. This would allow us to look at the potential transformations that DM neurons could be performing on the information from V1. The results presented here show models of receptive fields from both areas from the same type of stimuli under the same conditions which makes it perfect for comparisons between areas and creating models that encapsulate the interactions between them. This could potentially be achieved by simply extending the LNP model to include multiple layers, or

by using non-negative matrix factorization (Liu et al., 2017) to segment the DM receptive fields into subunits which can be compared to the V1 receptive fields modelled here.

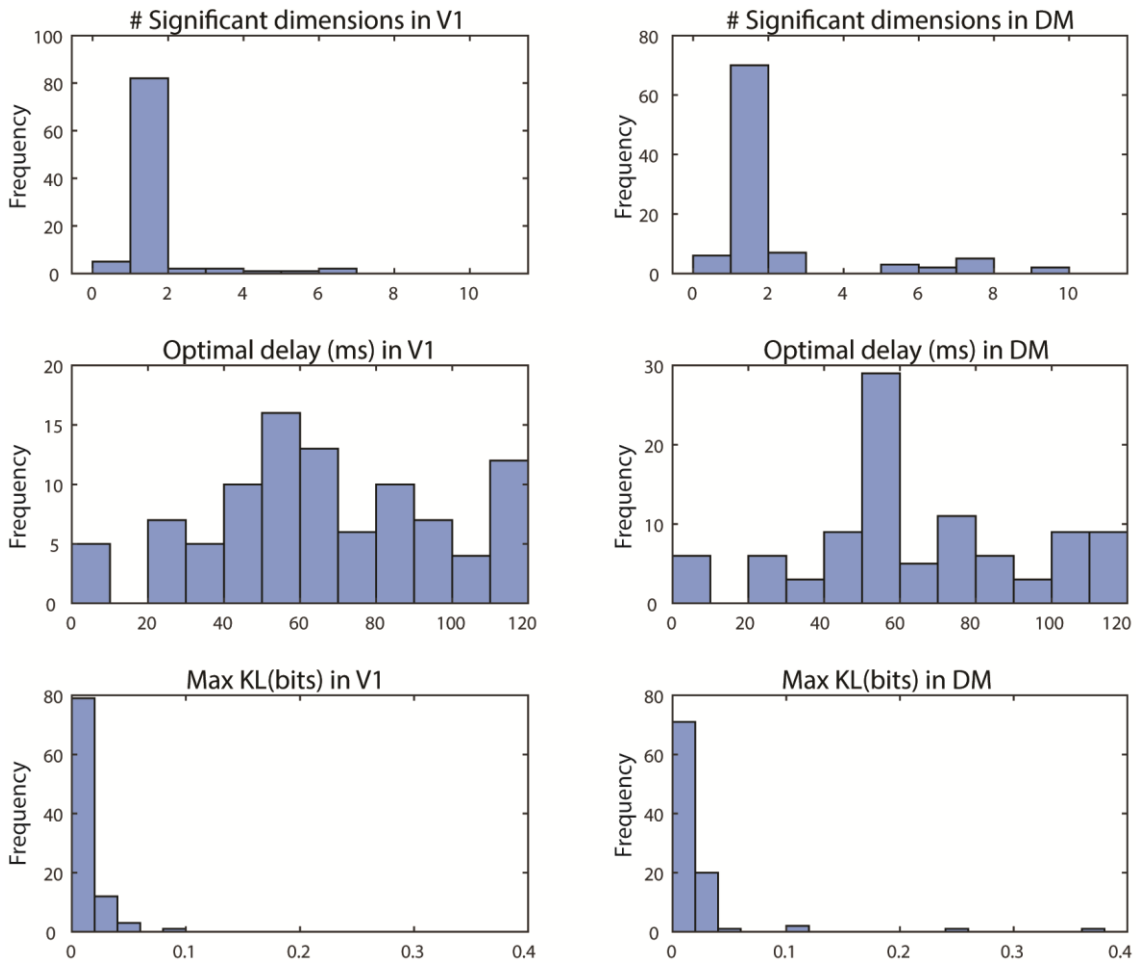

Figure 3.9. Neurons in DM have more complex responses than those in V1. Statistics of iSTAC vectors between DM and V1 for analyses on linear probes. Distribution of the number of significant filters (top row), as well as optimal delays (middle row) and KL_1 (bottom row) for the first most informative iSTAC vectors for V1 (left) and DM (right) for the same recording sessions as the data in Figure 3.8. Optimal time-point τ was set to the value that generated the maximal KL divergence for the first vector (KL_1).

Our analysis methods were successful across the set of data collected in DM, with many channels showing clear receptive field structures. For these neurons, the first four filters often include a

"blob" and pair of filters that resemble 2D Gabor filters approximately 90° out of phase. This is consistent with the filters found for the exemplar DM neuron in Figure 3.5. The KL divergence can be used as a measure of the discriminative power for each dimension. Although the dimensions with clear structures tend to have larger KL values, it should be noted that this is not always the case as some channels showed prominent spatial structure despite having KL values smaller than the neighbouring channels. Furthermore, although the channels that did not produce clear receptive field tend to have smaller KL divergence, there wasn't a threshold that separates those with spatial structure and those that do not. This is why the bootstrapping procedure is necessary for determining if a dimension is "just noise".

The distributions of τ and KL_1 across all electrodes for both areas are illustrated in Figure 3.9. Note that the plotted values did not exclude channels that did not produce clear RF structures. The bootstrap procedure will be able to establish the number of significant directions (0 for "dead channels"), the distributions of tau, and the distribution of the KL divergence of the subspace.

3.4 Discussion

The methods here show the power of reverse correlation white noise analysis. The masked covariance technique developed for locating receptive fields allowed for robust characterisation of the responses to local patterns of stimuli. To my knowledge this is the first time iSTAC analysis has been applied outside of V1, and the first time receptive fields in DM have been modelled in terms of their spatial structure. Reverse correlation analysis in DM revealed receptive field filters that suggest that neurons in DM encode information about long contours. Paired filters, similar to those of complex cells, were commonly found for DM cells as well as additional excitatory and suppressive filters. These results suggest that DM encodes information beyond just contours but further work is needed to characterise this fully.

3.4.1 Directions of invariance analysis

We were able to recover qualitative, but not quantitative results for the directions of invariance analysis. This method has promise for interpreting the visual stimuli that excite a neuron. The filter banks generated by iSTAC analysis describe the responses of neurons in terms of independent vectors. However, it's difficult to develop a hypothesis of what visual stimuli are actually relevant the cell's activity for these models alone. For this purpose, directions of invariance analysis could be a useful extension of the white noise analysis used here.

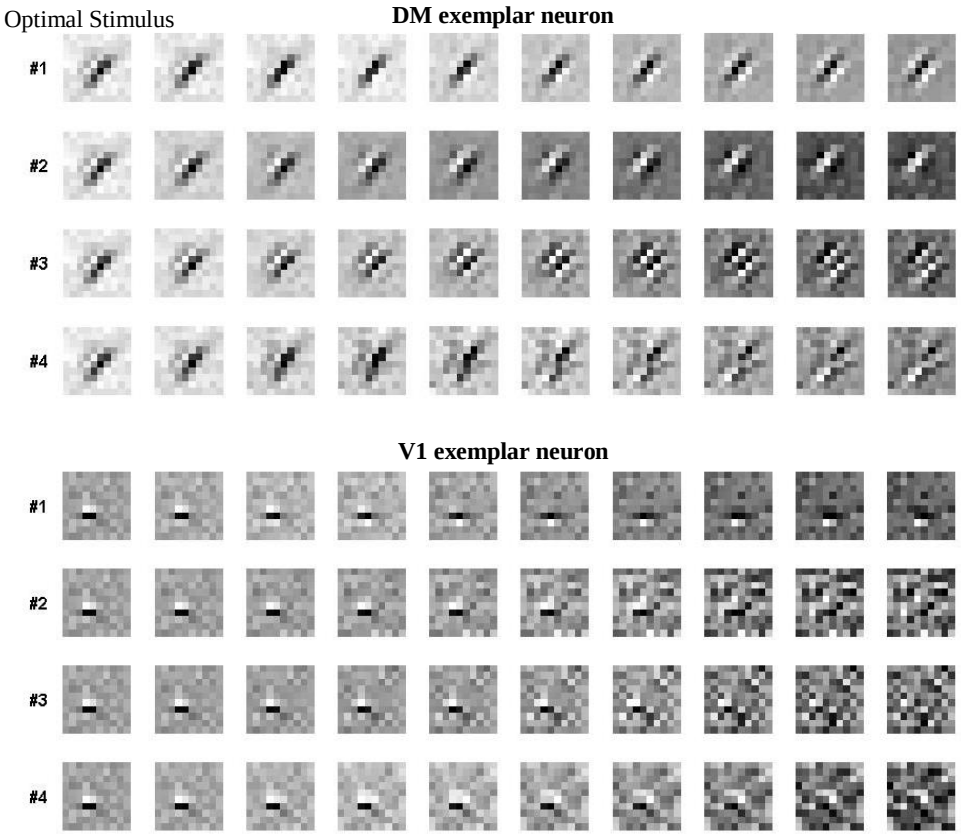

Figure 3.10. Transformations along the first four directions (rows) of invariance for the exemplar DM neuron from Figure 3.5 (top) and the V1 neuron in Figure 3.6 (below). In each row the optimal stimuli for the neuron is shown in the left-most column. Transformations are mapped across rows at 10° intervals along a geodesic of a surface where the energy of the stimulus is constant (Berkes & Wiskott, 2006).

Figure 3.10 details the results of direction of invariance analysis for the same DM (top) and V1 (below) exemplar neurons from previous figures. The first step in invariance analysis is to find the optimal stimuli for the neuron. Then, directions surrounding the optimal stimulus corresponding to transformations of the optimal stimulus, which change the responses of the neuron as little as possible, were estimated using the algorithm of Berkes and Wiskott (2007). Although the algorithm of Berkes & Wiskott (2007) was derived for model neurons assuming the quadratic form the same method can be applied to the ratio of Gaussian model produced by iSTAC without modification (see figure 3.3).

In Figure 3.10 the left-most image in each row is the optimal stimulus for the neuron. Each row (#1 to #4) gives 10 samples in a different direction of invariance along which the stimulus is gradually transformed away from the optimal image. The transformations are ordered in terms of how much of the optimal response can be preserved, with row #1 showing samples of the transformations to which the neuron is most invariant.

For the DM neuron each row appears to be sampling a different invariance. Row #1 and 2 are examples of phase invariance. If the alternating dark and light regions in the pattern are considered to be a sinusoidal function, then along row #1 the phase of this sinusoid will shift over 180° within the ten samples. Rows #3 and #4 are more difficult to interpret. Over the ten samples in row #3 the spatial frequency appears to increase. The transformation across row #4 suggests an invariance to orientation of some kind.

For the V1 neuron the transformation across the first row indicates a spatial phase invariance similar to the first two rows of the DM neuron. The remaining directions of invariance for the V1 neuron are probably not meaningful as the pattern within the centre of the receptive field stays the same or very close to the same as the optimal stimulus. This again suggests, that although there are many similarities between the receptive field properties of V1 and DM cells, DM cells are performing some more complicated computations.

The method of the direction of invariance analysis presented here is a starting point for what could be a useful tool for understanding what drives the responses of neurons. This is because the resulting set can give a more clear impression of what the neuron than the set of filters returned by iSTAC analysis. The set of excitatory stimuli are spanned by the set of linear filters, but what

combinations of filters mean in terms of pixels on a screen, or indeed objects in visual space can be difficult to deduce. The direction of invariance analysis however gives results in terms of transformations of images on screens which could be much easier to intuit.

However, much more work is needed to classify and describe the transformations beyond the first few. Here, we have restricted ourselves to qualitative descriptions of a few examples, but these transformations include complicated geometric changes including bands that split, change in orientation and in spatial frequency. In DM, the first couple of directions of invariance tend to simply show the phase invariance that we expect given the strength of the iSTAC filters resembling quadrature pairs. The directions of invariance beyond this are potentially much more interesting, if DM is really sensitive to these complex geometric transformations. The next step of this analysis is simply to see if the same transformations are seen throughout DM and what they might mean for the role DM is playing in processing visual information.

3.4.2 Extra-classical receptive fields

The white-noise analysis methods described above can be used to capture the properties of the receptive fields, but fail to estimate the influence of stimuli that fall outside the classical field. Although neurons don't respond directly to stimuli outside their receptive fields, for some neurons stimuli in the area adjacent to the RF (non-classical receptive field) can have a modulating effect on their response to stimuli within their receptive fields (H. E. Jones, Grieve, Wang, & Sillito, 2001; Knierim & van Essen, 1992; Series, Lorenceau, & Fregnac, 2003).

The size of the non-classical receptive field as well as the nature of the modulating effect can be investigated using drifting gratings in circular and annular apertures (Cavanaugh, Bair, & Movshon, 2002).

In V1, the modulation due to stimuli extending outside the classical RF is typically suppressive (Hubel & Wiesel, 1968) and orientation tuned (Bair, Cavanaugh, & Movshon, 2003). The modulating influence of the nCRF can also be feature selective. In order to extend white noise analysis to investigate feature selectivity in the nCRF, new methodology beyond what is presented in this report is required. This is because white noise analysis of classical receptive fields already requires recording for a long time. Searching for subtler statistical influences would likely take prohibitively long time to search the space of visual stimuli with random sampling.

We have collected pilot data in DM, using drifting gratings with varying diameters, and annulus gratings with varying inner diameter (Cavanaugh et al., 2002) to study the properties of non-classical receptive fields in DM. Preliminary analysis suggest that the integration fields of DM neurons are much larger than the size of the receptive fields revealed by iSTAC. Also, non-classical receptive fields are primarily excitatory in DM. This appears to be the most salient difference between the properties of V1 and DM neurons, but more work is needed to quantify this effect.

3.4.3 Function of area DM

A key motivation behind studying the brain in animal models is that the cortical area may be homologous to an area in the human brain. DM in marmosets has been proposed as a homolog to human visual area 6 (V6) (Rosa & Tweedale, 2001) although this is the subject of some debate (Lyon & Connolly, 2012). fMRI studies of human V6 show a strong response to moving stimuli, specifically

optic flow fields: coherent motion patterns that simulate the motion signals created in the visual field when the viewer moves in space (Pitzalis et al., 2010) . However, whether these properties are also present in marmoset DM remains to be determined. The results presented here support the conclusions of Lui et al. (2006), namely, DM is likely to play a role in the integration of contours and in the detection of continuity over extended regions of the visual field.

4 Multi-scale *in vivo* calcium imaging of the marmoset primary visual cortex

4.1 Introduction

While genetically encoded calcium imaging technology has been well established in rodents, its use in non-human primate models has only become practical in recent years (Heider, Nathanson, Isacoff, Callaway, & Siegel, 2010; M. Li, Liu, Jiang, Lee, & Tang, 2017; Sadakane et al., 2015). This technique offers several advantages over traditional *in vivo* recording techniques such as electrophysiology and intrinsic imaging. Electrophysiology directly measures the changes in electrical potential of or around a neuron, capturing the timing and dynamics of action potentials, but the number of neurons that can be studied at a time is limited. In addition, the spatial resolution is limited to the density of the electrodes that can be manufactured, with separations between channels being typically in the range of hundreds of μm. Optical imaging of intrinsic signals is used to observe changes in the activity of cortical regions over a large spatial scale. However, because the signal is driven by blood flow, the signals have poor spatial resolution (again, in the range of hundreds of μm) and temporal resolution (in the range of seconds). Moreover, the measurements obtained are only indirectly related to the underlying electrophysiological activity. Calcium imaging is an alternative method for monitoring the activities of neurons from relatively large sections of the cortex. Since the signal is driven by changes in intracellular calcium concentrations, which are related to the generation of action potentials, the recordings better reflect the signals recorded with electrophysiology.

Here, calcium signals were captured with two imaging modalities: one-photon fluorescent microscopy for wide-field imaging, and two-photon microscopy, which can penetrate several hundred microns deep into cortical tissue to resolve the activities of individual neurons (spatial resolution ~ 2μm). Together, these complementary techniques can reveal details of the functional

architecture of marmoset V1. Here, I validate the calcium imaging methodology for both single and two-photon recording using V1 orientation topology as a test bed.

4.1.1 Calcium imaging

Neurons at their resting membrane potential actively constrain their cytoplasmic Ca2+ concentrations at a level many thousands of times lower than that usually found in the extracellular space. Action potentials are accompanied by a large influx of calcium ions via voltage-gated channels due to the steep concentration gradient. Fluorescent calcium indicators allow for imaging of intracellular calcium transients related to spiking activity. By working backwards from the dynamics of the fluorescent signal, these indicators can be used to infer the functional behaviour of the neuron. Genetically Encoded Calcium Indicators (GECIs) are a set of indicators that can be non-invasively delivered to cells by either transgenesis or via viral vectors. GECIs have some advantages over regular calcium indicators in that they can be used for imaging multiple times over much longer periods of time provided that the expression level is maintained. The GCaMP series of GECIs were developed by combining a calcium-binding domain (calmodulin) fused to a fluorescent protein and M13 peptide fragment from myosin light chain kinase (Nakai, Ohkura, & Imoto, 2001). The fluorescent section of GCaMP is based on Green Fluorescent Protein (Shimomura, Johnson, & Saiga, 1962), which was originally derived from a jellyfish (*Aequorea victoria*).

4.1.2 Virus vectors

Transgenic lines expressing GECIs are well established for rodent models (Chen et al., 2013). But for non-human primates such as the marmoset monkey, transgenic lines are only recently being developed (Park et al., 2016). In lieu of a transgenic line, recombinant adeno-associated virus (AVV) vectors can be used to infect non-dividing neural cells and induce long-term stable expression of GECIs. We used a set of vectors recently developed in the Yamamori lab (Sadakane et al., 2015; see

also Ch 2. *Methods*), a group which has been very successful in adapting the technology to marmosets.

4.1.3 Two-photon microscopy

Two-photon imaging is a state of the art microscopy technique that is similar to confocal laser scanning microscopy (CLSM) in optical design, but with key differences that greatly improve its usefulness as a tool for neuroscience. Both methods use lasers to scan an area at a tightly constrained depth below its objective lens to collect images of very thin slices of tissue. CLSM uses a high-frequency laser for fluorescence excitation, but in order to deeply penetrate tissue a lower frequency laser can be used by applying the principle of two-photon absorption. Typically, low frequency (longer wavelength) light cannot be used as it carries less energy per photon than is needed to elicit a fluorescent response. However, at very high light intensities the usually negligible probability of a fluorescent molecule absorbing multiple photons simultaneously increases. When this occurs the energy from two or more photons are combined to excite the molecule to a higher energy state, from which it will release a photon of a higher frequency than the photons it absorbed. Two-photon microscopy exploits this phenomenon with the use of very high local instantaneous intensity, provided by tight optical focusing as well as the temporal concentration of a femtosecond pulsed laser (Denk, Strickler, & Webb, 1990).

There are two primary benefits of the technique for applications in neuroscience. Firstly, the lower frequency laser is far less damaging to the tissue it passes through. Only a small volume of tissue at the focal point is exposed to the high intensity light required to excite fluorescence, and only for very short durations. The use of the pulsed laser also works to prevent the photobleaching effect, in which fluorescent molecules are excited beyond the point of returning to their resting state. Furthermore, since the low frequency light is less likely to interact with the tissue outside of the

focal point, scattering effects and background noise are both reduced. Importantly for neuroscience applications, this means that it can be used for imaging several hundred microns deep into cortical tissue *in vivo* (Helmchen & Denk, 2005). This combined with its sensitivity and very high spatial resolution makes it a revolutionary new tool for neuroscience.

4.1.4 Wide-field imaging (one-photon microscopy)

Traditional fluorescent microscopy is sometimes referred to as one-photon microscopy, as the indicators are excited by single high-energy photons. This imaging modality cannot penetrate deep into the tissue due to scattering (unless a confocal microscope is used), and it cannot resolve the signals from individual cells (also due to scattering). For GECI calcium imaging, one-photon microscopy is typically used to gauge the expression level of the virus.

Wide field imaging has been recently gaining popularity as a tool in neuroscience. In rodents it has been used to simultaneously image from large regions of cortical surface (Homma et al., 2009; Mohammed et al., 2016). However, the use of wide-field calcium imaging in marmoset as a tool for neuroscience is still in its infancy. One of the goals of this project is to examine if this could be exploited as a way to perform multi-scale calcium imaging in combination with two-photon imaging in the visual cortex. The first goal of this project is to validate the calcium imaging methodology for both single and two-photon recording using V1 orientation topology as a test bed.

4.2 Methods

4.2.1 Data collection

The results presented in chapter 4 are from a series of experiments conducted on a single marmoset under both awake and anaesthetised experiments in 2017. We also have additional data for a further three animals, all from anaesthetised experiments, results of which are presented in the next chapter.

4.2.2 Image segmentation

For images acquired with two-photon microscopy, a matrix factorization algorithm (Pnevmatikakis et al., 2016) was used to isolate single neurons. This method automatically segments overlapping cell bodies (i.e., demixing) using the spatiotemporal correlations of the signals (Maruyama et al., 2014). After neurons were identified by the algorithm, signals from all pixels inside the identified neurons were averaged for the subsequent analyses.

4.2.3 Inferring neuronal activities from fluorescent traces

Blind deconvolution is a commonly used method for inferring neuronal activities from calcium-imaging data, which converts the fluorescent signal into a moment-by-moment estimate of the firing rate of the neuron by the application of an estimated deconvolution filter. The method is "blind" in the sense that it treats the fluorescent signal as a time series, ignoring the design of the experiment. Here, we used an alternative analysis technique proposed by Ganmor et al. (2016), which provides maximum likelihood estimates for each condition of the experiment, assuming that the firing rate of the neuron is constant for each condition. It can be understood as a deconvolution method that is aware of the design of the experiment. Since the method only needs to estimate a small number of parameters, it is more robust than the blind deconvolution method. It works by incorporating the intermediate steps between the firing rate in response to a stimulus condition and the resulting

measured fluorescence into a single loss function that can be collapsed across all repeats and all time points. The model assumptions are shown in Figure 4.1. The method can be understood as directly relating the fluorescence to the stimulus without introducing bias and error by calculating each intermediate stage.

4.2.4 Model implementation

The model was first tested on simulated data, an example of which can be seen in Figure 4.1, as well as combined electrophysiology and calcium imaging data from (Akerboom et al., 2012; T. W. Chen et al., 2013) to confirm that the method works as expected. An additional assumption was added to the model shown in Figure 4.1 to account for a delay between stimulus presentation and firing rate in order to relate the estimated rate to the stimulus on screen a short time before. The implementation used here fits a firing rate sequence for each stimulus type across all repeats. This allows for the firing rate to change over time during a stimulus presentation. In the results the mean rate for each stimulus type was taken, the changing histogram of rate over time was only used to estimate and set reasonable delay.

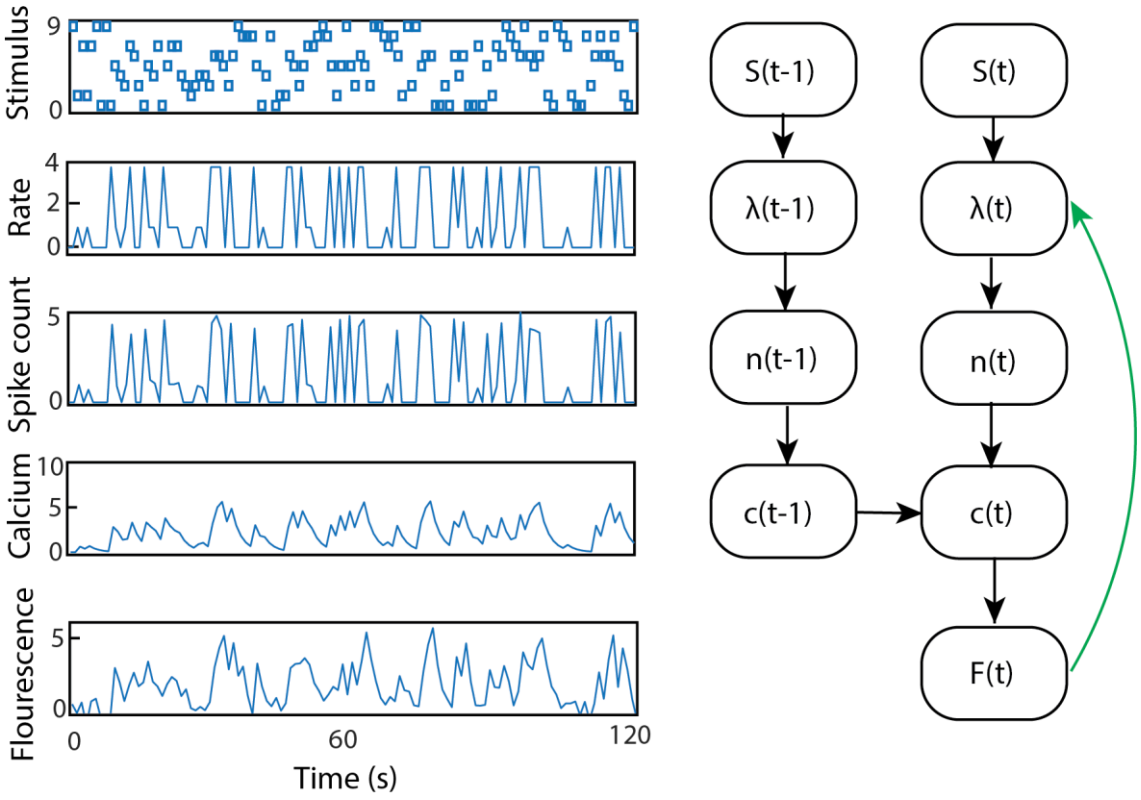

Figure 4.1. Direct Estimation. Graphical representation of model assumptions and dependencies described in Ganmor et al., (2016). Simulated fluorescence data used to test the direct estimation method was generated by setting a firing rate for each stimulus setting. Spike times were generated through a random Poisson process and convolved with an exponent to generate a calcium trace c(t) which gives the fluorescence F(t) after scaling and adding measurement noise. This was used to test if the rate for each stimulus condition λ(t) could be recovered directly by only observing the fluorescence (green arrow) rather than stepping back through each step of the model. The intermediate dependencies are instantaneous and simultaneous in time, except for the calcium levels which depend on values at the previous time step.

4.3 Results

In order to validate the combined one- and two-photon imaging protocol, we measured fluorescence levels from genetically encoded calcium indicators in V1 while presenting drifting gratings. Using both protocols, orientation selective responses were found. If these methods work as expected, the organisation of orientation tuning over cortex revealed by these methods should be consistent with what has been reported in V1 before. Additionally, agreement between the methods would confirm that the diffuse signal measured using wide field imaging is representative of the tuning of individual cells in the same region and can be used to estimate the tuning of cortical columns.

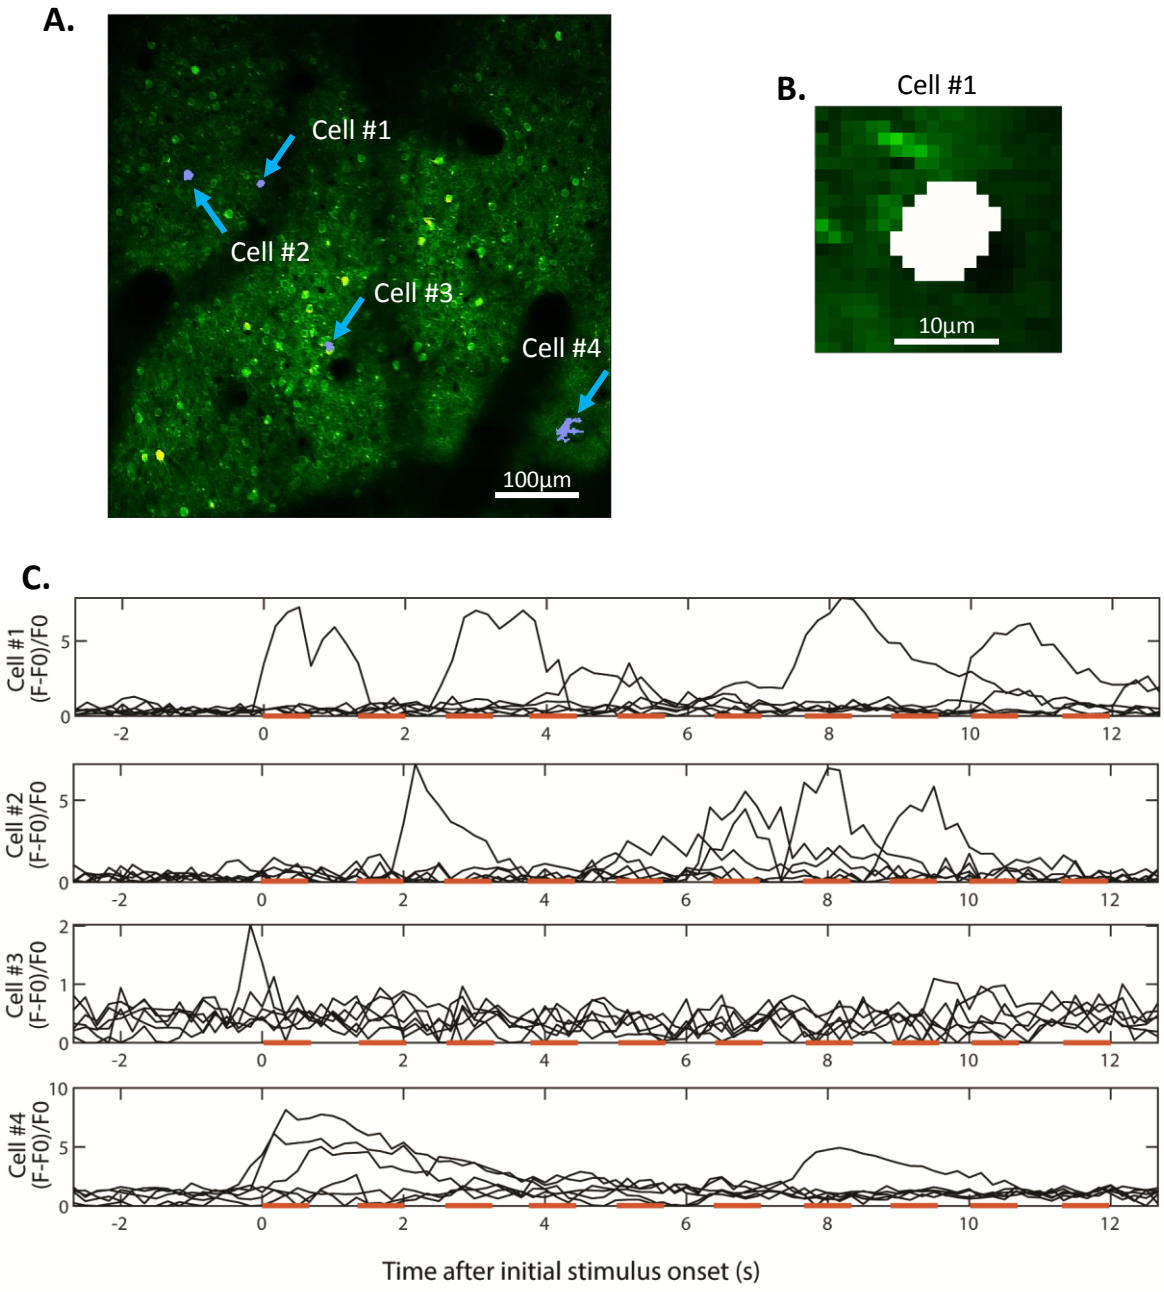

Figure 4.2. Single cell traces from two photon imaging data. A: Time averaged image of two photon imaging site. Location of four example cells in blue. B: A closer view of the isolation of cell #1 as determined by the matrix factorisation algorithm. 52 pixels were assigned to this cell. C: example traces for the cells labelled in A. Since many trials have little or no activity, the traces for the first 10 stimuli for 6 trials are shown on each axis.

4.3.1 Single cells imaged with two photon microscopy were stimulus selective

In the top left panel of Figure 4.3 is an example image created by averaging 2950 motion corrected frames of two photon imaging. In the 600x600µm region there are hundreds of cells lit up with fluorescence and dark areas where blood vessels occlude the image. The patterns of blood vessels are present in both the one-photon and two-photon images and may be used to align the two precisely. The matrix decomposition algorithm initially labelled 1312 cells, 785 of which were active over the recording time and the rest were rejected from further analysis. Four example cells are labelled in light blue with light blue arrows.

Cell #1 occupies 64.85µm^2 in the plane of the reconstructed image, with an approximate diameter of 10.95µm. Figure 4.2c shows the (F-F0)/F0 trace for the four labelled cells during 6 trials of recordings of 12 seconds with 10 stimuli onset times marked by red line segments on the horizontal axis. Changes in fluorescence extend over multiple trials due to the slow dynamics of CGaMP6s. This makes it difficult to judge whether the (F-F0)/F0 signal reflects stimulus-evoked or spontaneous firing. However, as can be seen in Figure 4.3a, averaging across all trials of best orientation (in blue) and worst orientation (in red) shows a clear preference (individual trials in dotted lines, means in bold). By applying Ganmor et al.'s (2016) method for directly estimating a single underlying rate for each stimulus type we can estimate a rate for each direction (Figure 4.3e).

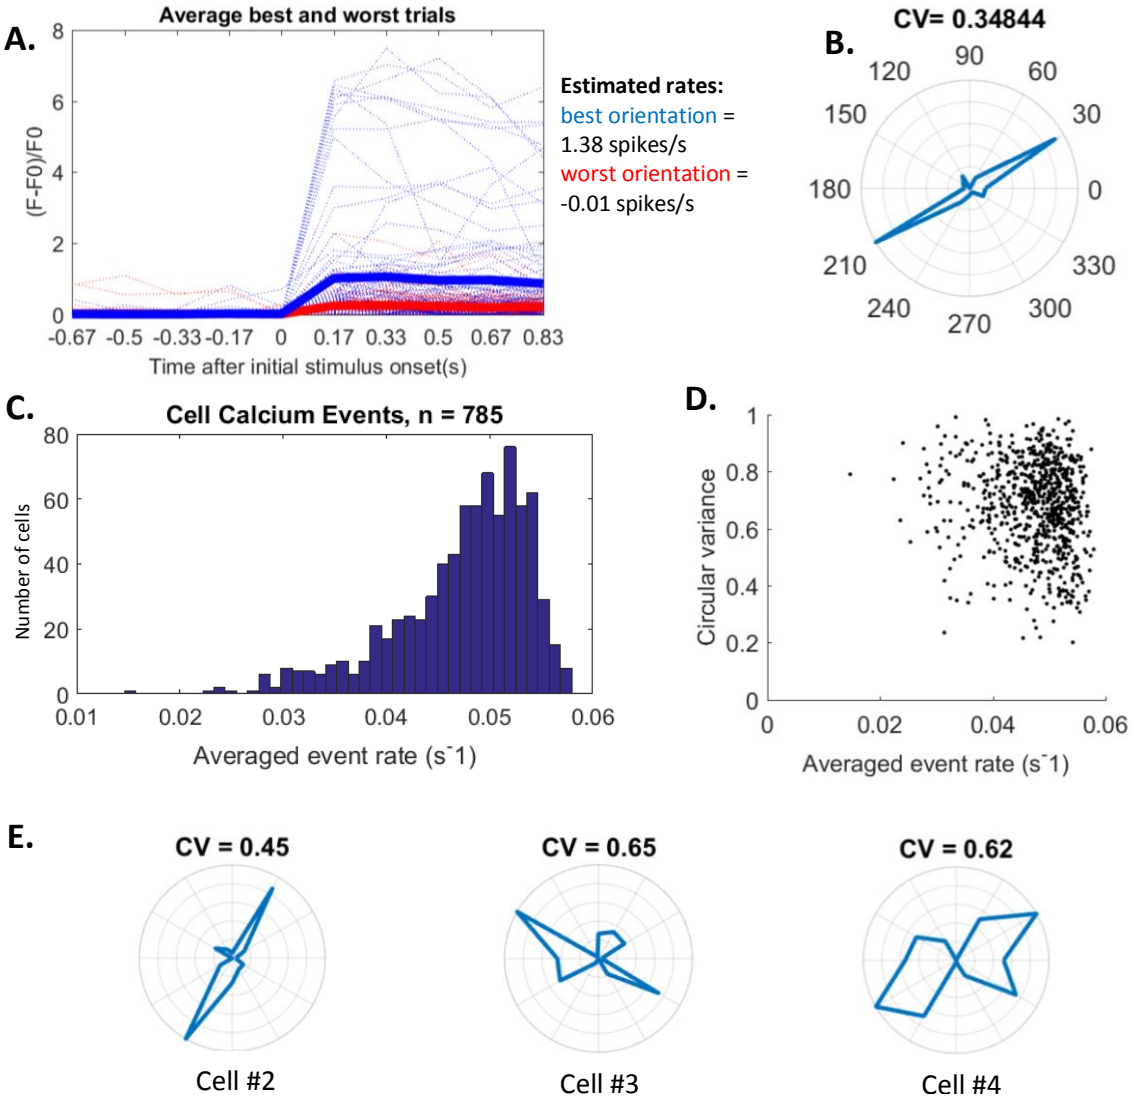

Figure 4.3. Single cell tuning from two photon imaging data. D: individual trials (dotted lines) and averages (in bold) for the best (blue) and worst (red) stimulus conditions. E: Polar plot of direction tuning of the same cell. F and G: population statistics for Calcium events and its relationship with circular variance. H: Direction tuning for the four cells labelled in figure 4.2a.

Large calcium events, as determined from thresholding the rising edge of the fluorescence trace, are very rare with the median event rate for all cells of 0.05 events per second. The small number of calcium events could be a sign of weak or suppressed neural activity in response to the visual stimuli. However strength of tuning (indicated by 1-circular variance) and event rate were not very correlated, suggesting that the low event rate was not a limit on tuning estimation. This suggests

that although the fluorescence due to single spikes get drowned out by larger slower signals, the direct estimation method can recover them through inference.

Recordings were made for 785 tuned neurons in the 600x600μm region. In order to demonstrate that imaging the same region of cortex with single photon imaging gives similar results, these cells were matched to the nearest pixel in the wide field imaging recordings. In comparison to single cell imaging, wide field imaging has an advantage as it can reveal patterns of tuning over a wider field of view.

4.3.2 Single-photon signals are stimulus selective

An example of the single photon imaging data recorded for a single pixel is shown in Figure 4.4. In panel A the raw traces for 4 repeats of the best condition (blue) and worst condition (red) are shown for the first two seconds after the onset of stimulus presentation. The signal from each trial is normalized against the average intensity of the fluorescent signal of 10 frames (0.42s) at the beginning of the trial, F0, by first subtracting then dividing. While the signals from individual trials are too noisy to see a difference in response when averaged across trials (Figure 4.4c) it becomes clear that the change in fluorescence at this pixel is stimulus selective. Figure 4.4d shows the values of the averaged responses for the time window from 0.7 to 1.12 seconds after stimulus onset and the mean across trials (large circles). This time window was chosen as the ten frames (0.43s) around the peak response from the mean of responses across all pixels to all trials.

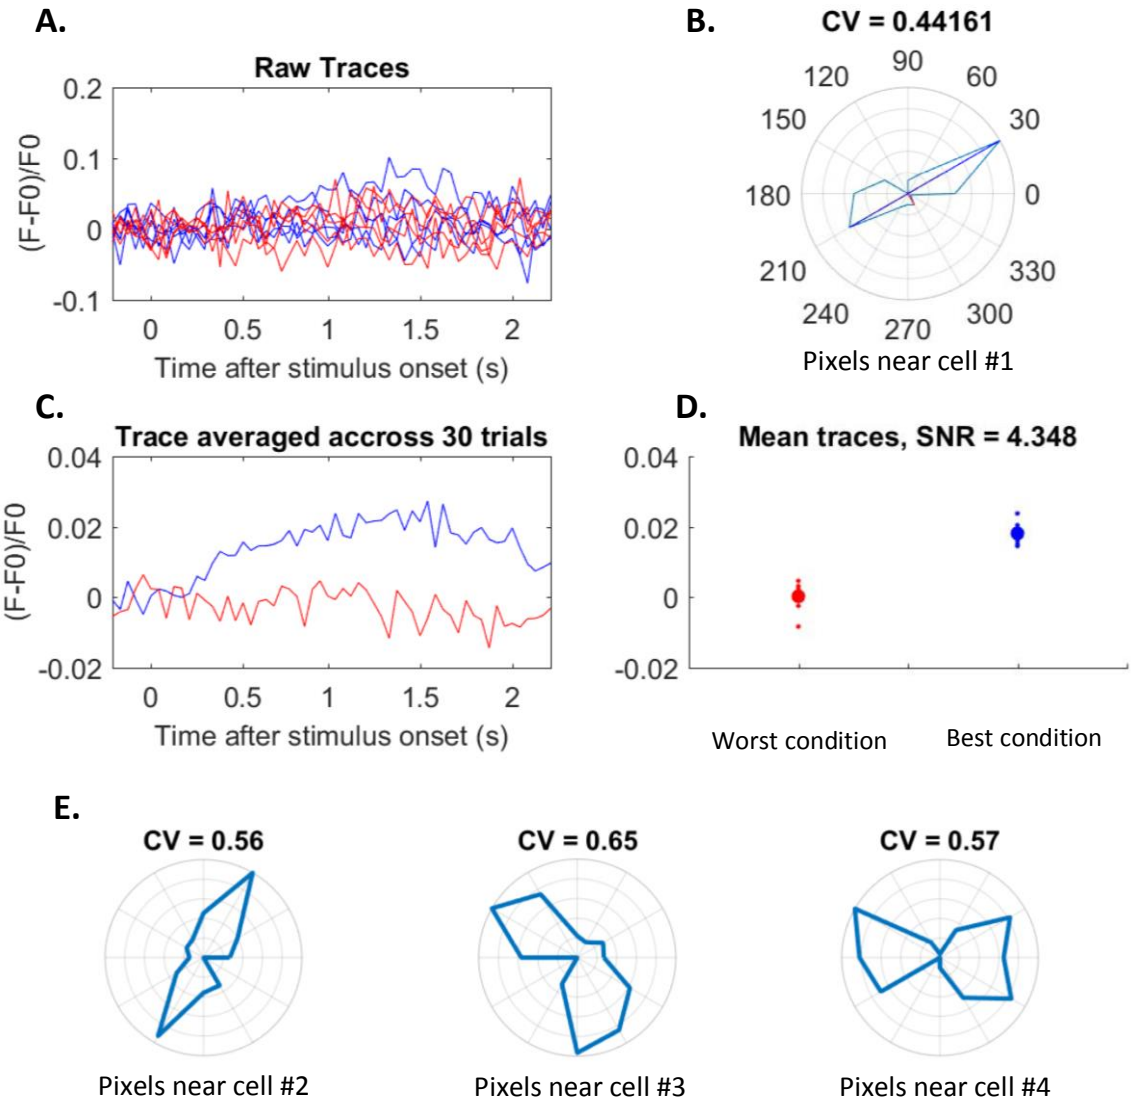

Figure 4.4. Orientation tuning of one-photon signals. Stimulus selectivity of an example pixel matched to the location of cell #1 in figure 4.3a, with spatial averaging over a 3x3 square centred on the pixel. A, the raw traces for 4 repeats of the best condition (blue) and worst condition (red) are shown for the first two seconds after the onset of stimulus presentation. B, polar plot of the tuning curve for this example pixel. C. Mean across all 30 trials. D. the values in C within the time window used for tuning estimation. E: Four example tuning curves generated from single photon imaging recordings displayed on a polar plot where the relative response to each direction is represented by the distance of the trace from the centre at that angle. The single photon trace in each row is matched as close as possible to the location of the example cells in figure 4.2.

We found the mean value for the normalized fluorescence (F-F0)/F0 for each pixel under every trial condition. We then compared the differences in response to orientation by collapsing across all other grating parameters. For instance, to generate the polar plot of direction tuning shown in B, the responses from trials with gratings drifting in that direction were averaged irrespective of the spatial frequency of the grating. This was necessary due to the large amount of noise in the signal. The tuning curves for four example pixels are shown in panel E of Figure 4.4. These pixels were chosen to match the location of the example cells in Figure 4.2 as closely as possible to allow for comparison.

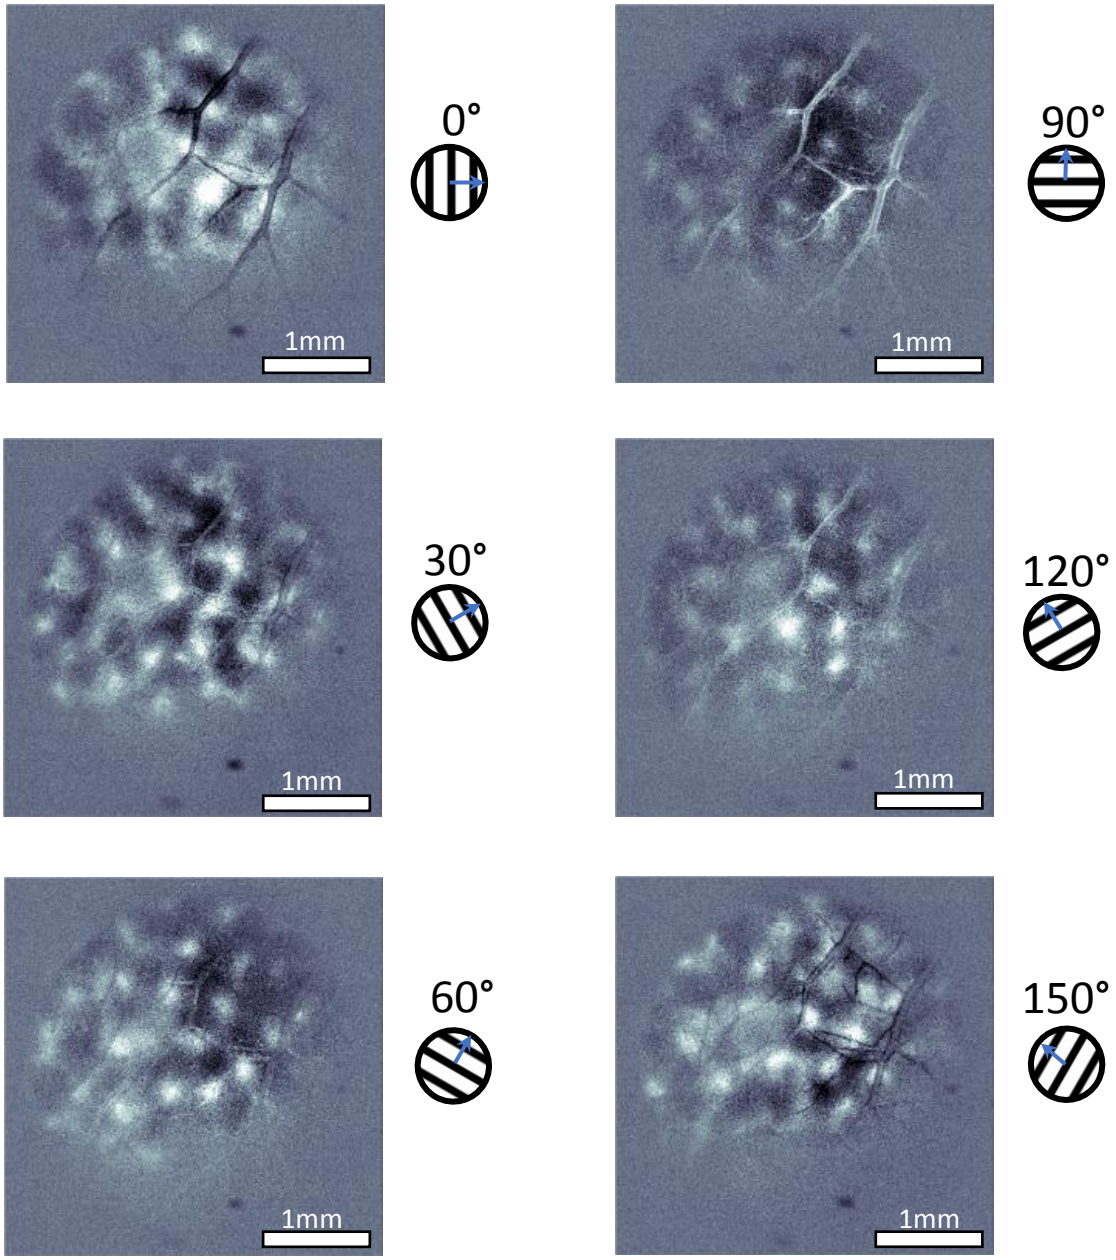

Figure 4.5. Averaged one-photon images reveal orientation column. Raw fluorescent trace for each orientation were averaged over 60 trials. Selectivity to each orientation over cortex is highly organized. Light patches in these images show regions of increased fluorescence due to activity in response to the drifting grating stimuli.

4.3.3 Single photon orientation maps

Wide scale single photon images were recorded while presenting randomized drifting gratings of six spatial frequencies and 12 directions. Fluorescence levels were averaged from five recordings for a total of 60 trials for each orientation by collapsing pairs of directions separated by 180° into six orientations. The images reveal distinct regions of cortex that respond preferentially to a single orientation, orientation columns (Blasdel, 1992; Hubel & Wiesel, 1968).

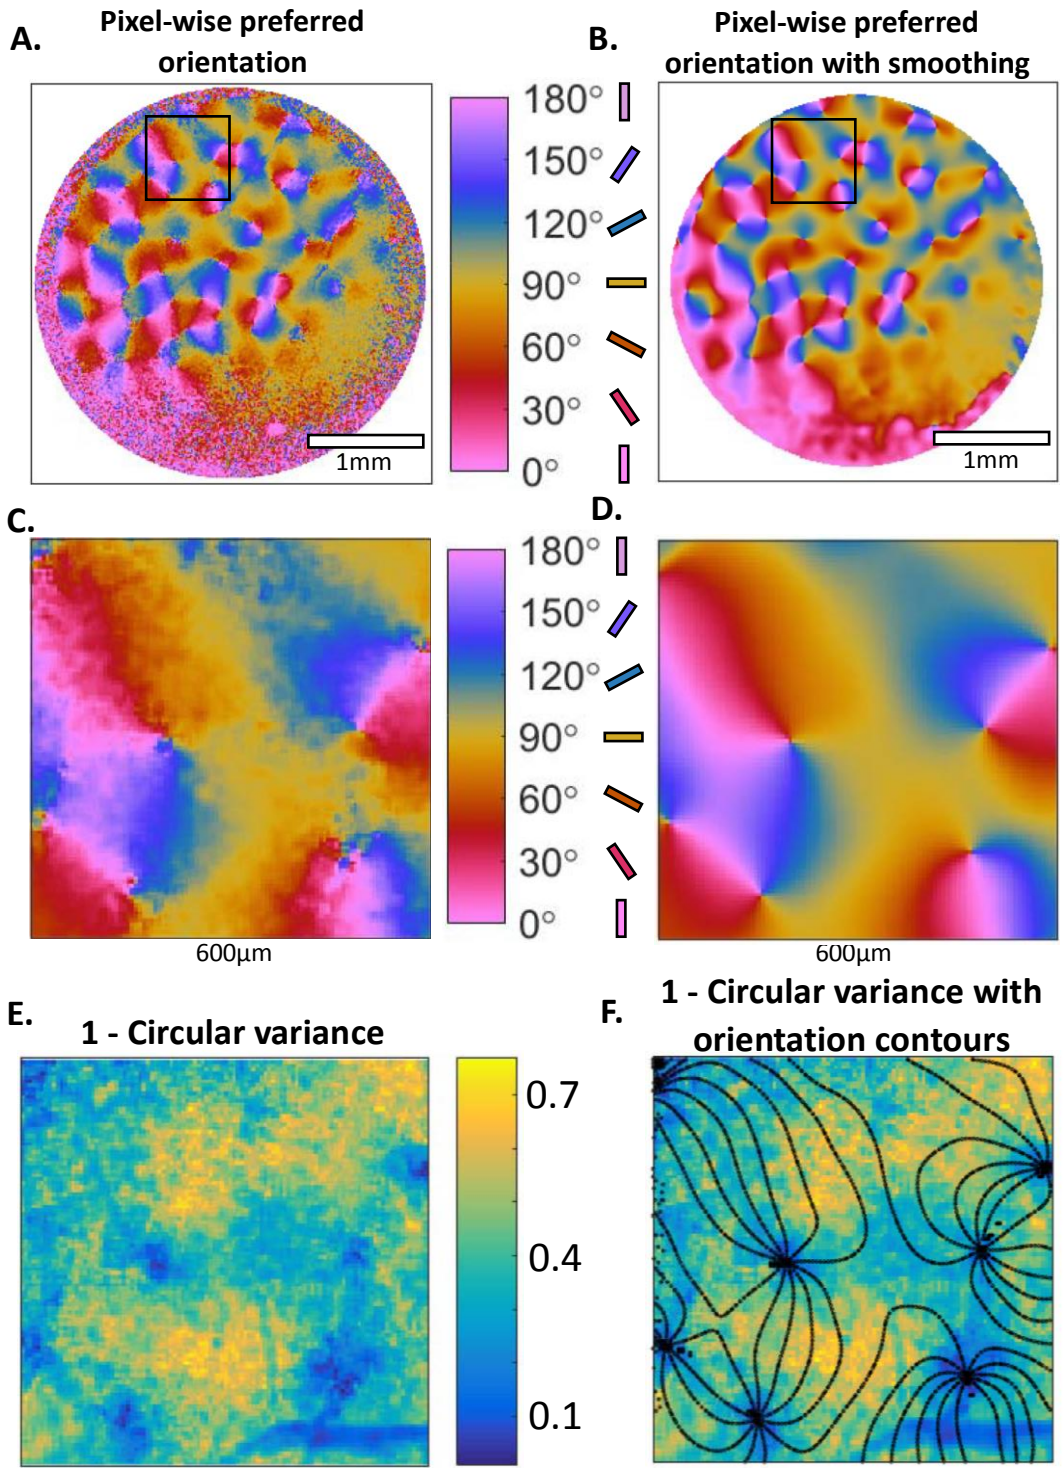

Figure 4.6. Pixel-wise preferred orientation. A: Unsmoothed pixel-wise preferred orientation calculated using one-photon data from the entire imaging window. C: Same for the 600x600µm region imaged in the two photon experiments. B&D: Same as A&C, but after smoothing with a Gaussian filter (5 pixel sigma). E: Pixel-wise values for 1- circular variance calculated using one-photon data over 600x600µm region imaged in the two photon experiments. F: By superimposing contours of equal orientation it appears that the local minima are aligned with the pinwheel centers.

The preferred orientation for each pixel was calculated using the responses to each orientation as in Figure 4.5. Given that single pixel tuning observed in the one-photon map reveals the orientation columns we combined the measurements for different orientations to a construct a map of preferred orientation over cortex. In order to find the best orientation a vector sum of vectors with magnitudes equal to the responses to each of the six orientations evenly distributed around 2π (Mazurek, Kager, & Van Hooser, 2014) such that the responses to orthogonal orientations cancel. This gives us the maps shown in Figure 4.6, where preferred orientation changes gradually over the surface of the cortex.

When orientation preferences for every pixel are combined, they form a map with a smooth, organised topology with distinct regions of cortex responding preferentially to a single orientation (orientation columns) and areas of rapid change in preference across cortex (pinwheels). These structures are well-known hallmarks of area V1 (Hubel & Wiesel, 1974) and their presence in these maps supports the idea that the epifluorescence signal could provide useful information about the underlying function of the cortex. The map was also smoothed with a Gaussian filter (5 pixel sigma) to more clearly identify these features (Figure 4.6b). A zoomed in view of the 600x600μm region imaged in the two photon experiments is shown in the panels in the middle row of Figure 4.6. In this small area of cortex there appears to be five pinwheels. This is the region of interest we selected for two-photon imaging.

We used 1 – circular variance of the same set of vectors (Mazurek et al., 2014) as a metric for the strength of tuning for each pixel. In areas where tuning was not expected, such as blood vessels and outside the imaging window, this metric gives values close to zero (blue in Figure 4.6e). The map of strength of tuning also reveals small regions of cortex with very low tuning strength. Superimposing

contours of equal orientation over the map (Figure 4.6f) shows that these low regions correspond to areas where the contours converge; i.e. the pinwheel centres. This is somewhat expected since the signals close to the pinwheel centre would contain the scattered fluorescence signal from many neurons around the pinwheel with different preferred orientations. Whether the strength of tuning of individual cells decreases as a function of proximity to a pinwheel centre is not something that this method has the resolution to answer.

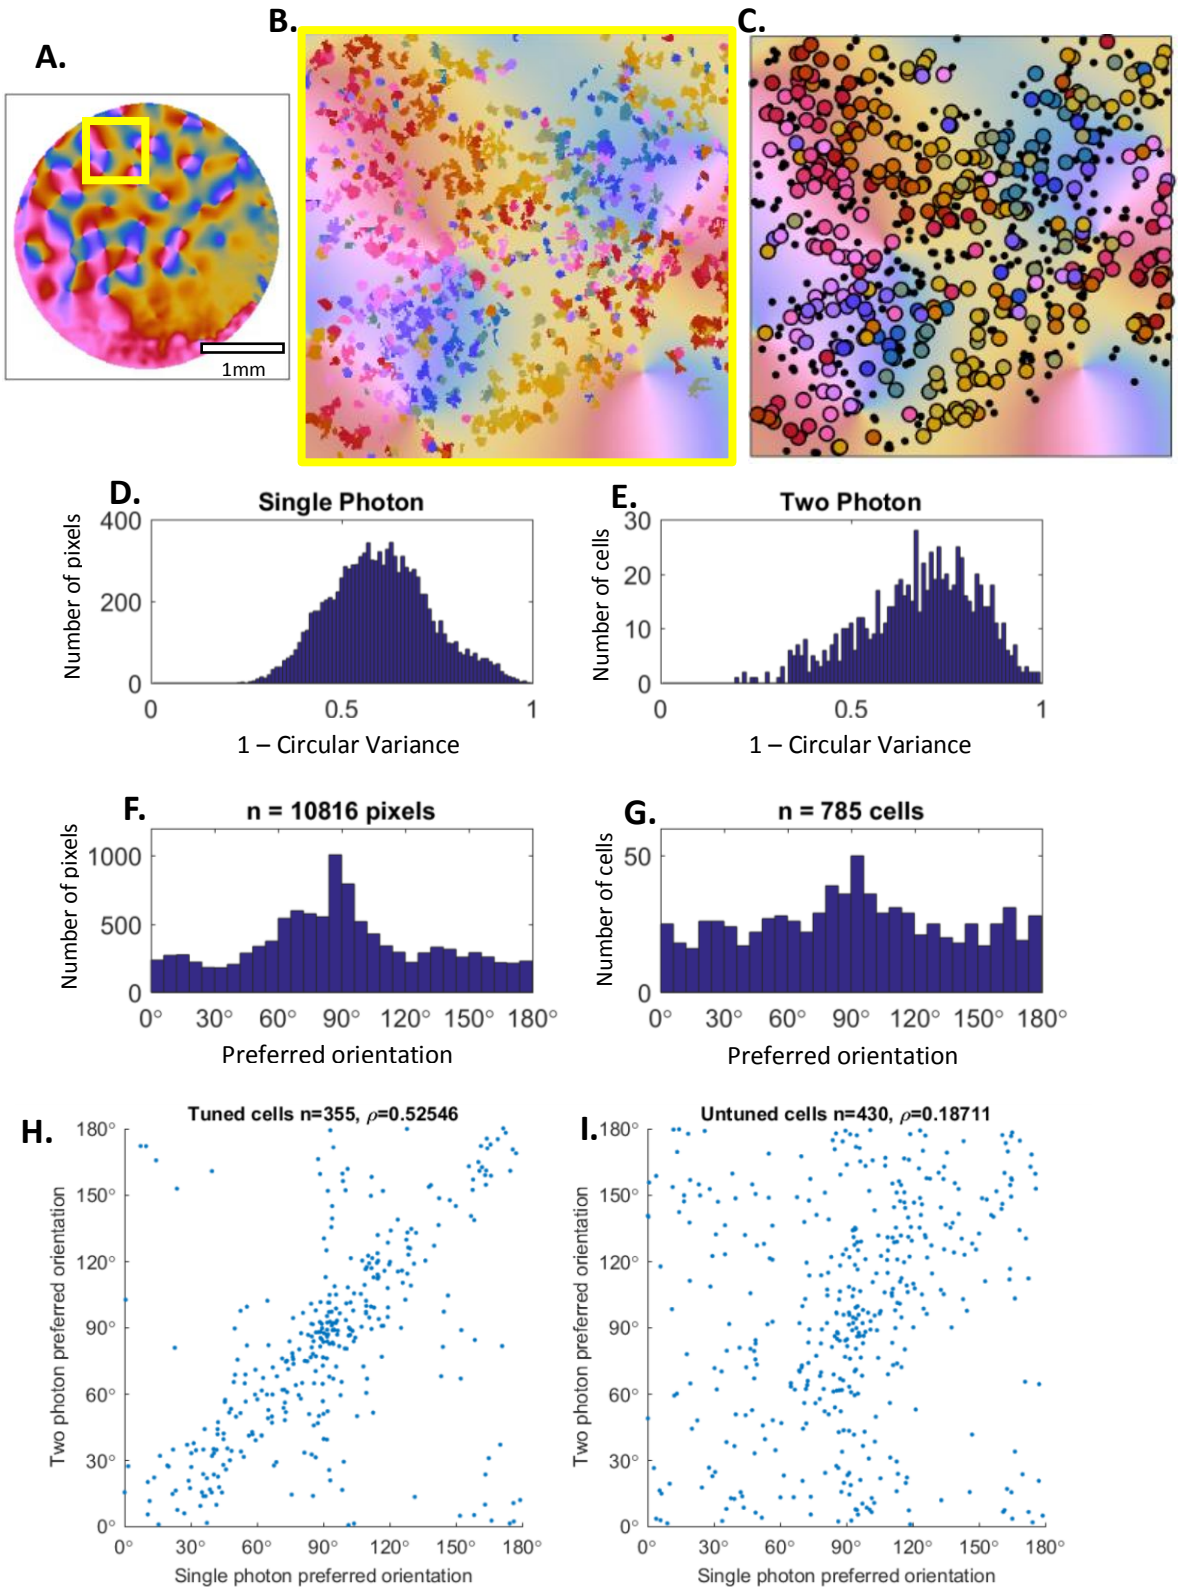

Figure 4.7. Alignment between one and two photon imaging. B: Tuning of individual cells versus the tuning from one photon imaging. C: 355 of 785 cells were "tuned" (1 - circular variance > 0.95), coloured circles, against the tuning from one photon imaging. D and E: The median circular variance was significantly different between tuning curves found from single and two photon imaging (Wilcoxon rank sum test p=

3.19e-53). H: The preferred orientation of tuned cells had a correlation coefficient of 0.53 with the corresponding preferred orientations from one photon data, whereas the untuned cells, I, had a coefficient of 0.19.

4.3.4 Alignment between one and two photon imaging results

The distributions of the tuning strength (as measured by circular variance) between the values calculated from the one photon and two photon imaging within the region are very different (Figure 4.7d&e), however this is not unexpected. The fluorescent signal for the one photon imaging is scattered as it passes through the cortex to a much higher degree than for the two photon. Nonetheless, the distributions of preferred orientations are very similar within the region and this is a good indication that the two imaging methods are both picking up the same underlying signal. Directly comparing the orientation tuning of (two-photon) cells with the closest (one-photon) pixel to their centre of mass shows good agreement for tuned cells (1 - circular variance > 0.95) with a correlation coefficient of 0.53.

4.4 Discussion

These results serve to validate calcium imaging as a tool for probing stimulus dependent activity in the visual cortex at multiple scales. Because the topology of orientation tuning in V1 as well as the behaviour of individual neurons in V1 are well described in the literature, V1 is the perfect testbed for validating cutting-edge techniques. The results here show that these methods produce maps for tuning which are consistent with both the understanding of V1 and also consistent across scales. This suggests that the methods are working as expected.

We found that the results for single cell two photon imaging were comparable to single electrode studies. One key difference was that the estimated spike rate was much lower than what is often reported for neurons in V1 (Yu et al., 2010). The discrepancy could represent a problem with the implementation of the direct estimation method used here. Alternatively, the difference between the rates between methods could be due to differences in the sampling between the two methods. In single electrode electrophysiology, searching for active neurons to record from may bias researchers to sample from neurons with higher average firing rates. Another explanation for this could be that we recorded from the supra-granular layers of V1, whereas electrophysiology experiments typically record from all layers, including layers 4 and 5, where firing rates are higher (Gur, Kagan, & Snodderly, 2005; Ringach, Shapley, & Hawken, 2002).

The wide field imaging results are comparable to results from intrinsic optical imaging. However, calcium imaging has some additional benefits over optical imaging methods. The signal recorded with optical imaging is typically very weak and in marmoset there has been difficulty in recovering maps for spatial frequency, colour and ocular dominance using optical imaging techniques. The wide field calcium imaging method used here may allow these maps to be robustly imaged in marmoset.

A major additional benefit of the calcium imaging methods described here is being able to easily compare between scales within the same region of cortex. Without this setup comparing between scales would usually require switching between imaging and electrophysiology during the experiment. Here, in contrast, the setup for the animal is the same between imaging methods. Finally, an important advantage of using two photon single cell imaging over electrophysiology is that the method is far less invasive and allows for repeated imaging over weeks and months without damaging the cortex. Although chronically implanted electrode arrays can also record over long time periods, once implanted they cannot be moved to sample different populations and signal tends to deteriorate over time (Szarowski et al., 2003).

4.4.1 Single cell imaging

One of the most important benefits of the techniques described here is that it can be used to describe the behaviour of single cells in the larger context of the functional architecture they are embedded in. The combination of wide field imaging and single-cell recordings allows for the organisation of cortex to be examined in unprecedented detail. Here, we simply verified that our methods give consistent results across scale, but in the future this technique could be used to probe questions about how the organisation of cortex is reflected in the tuning of individual cells. For example, are individual cells in colour sensitive regions all sensitive to changes in colour? Do cells near pinwheel centres have disturbed orientation tuning or discontinuous receptive fields? Of particular interest is understanding the relationship between orientation tuning and the progression of on and off receptive field subunits around pinwheel centres.

In the results presented here, direct estimation was used to estimate a firing rate for each stimulus type from the fluorescence trace and a series of assumptions about how the rate is reflected in the calcium concentration. The estimation of an underlying rate per stimulus can be extended to estimation of receptive fields via a 'rate map', simultaneously fitting the contributions towards firing rate for each part of the visual stimulus. Preliminary results for an example neuron are shown in figure 4.8. Sparse noise stimulus, many frames with one light and one dark square randomly placed on a grey background, was used to find the location of receptive fields. For this stimulus, deconvolution was more effective at recovering a structured receptive field than direct estimation of a rate map. This could be due to the large number of parameters in the rate map (441, 21x21, squares) over which the maximum likelihood function had to descend. In the future, parameterisation of the rate map, for example assuming a Gaussian function (5 parameters) or a two-dimensional Gabor function (9 parameters) in the model fitting could potentially dramatically improve performance of the direct estimation method.

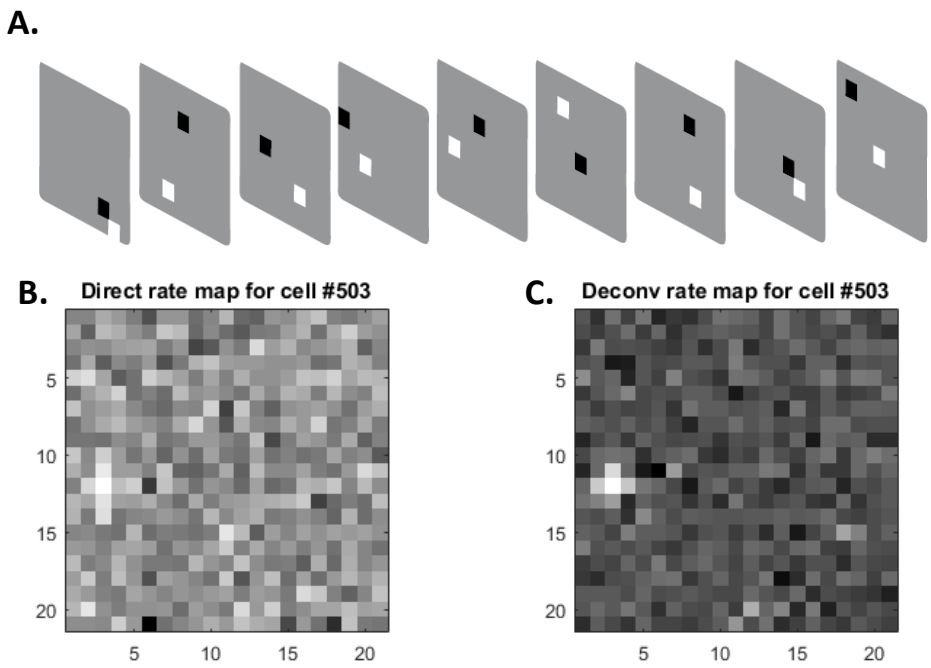

Figure 4.8. Receptive field mapping for an example cell. A. Stimuli were a pair of black and white squares at a random position in a 0.25 degree square grid (top). So far we have had limited success in recovering the receptive fields of the recorded neurons. Using deconvolution of spike timing followed by Spike-triggered averaging we can make estimates of a linear spatial filter underlying the responses of only a limited number of neurons (C). Surprisingly the direct estimation of a rate map (B) has failed to outperform the deconvolution based method in our analysis so far.

4.4.2 Future directions

Two-photon calcium imaging as a tool in primate neuroscience still has a lot of catching up to do relative to its use in rodents. In the future, it could be used to image active compartments such as dendrites. There has been some early success in the Yamamori lab in attempts to do this type of imaging in marmosets and this would be a natural next step after imaging from cell bodies of active neurons in V1 (Sadakane et al., 2015). One potential use would be to look at the information carried by long range feedback connections to interneurons in V1 and compare the tuning properties to the columns they connect to.

Another natural extension of the methods is to develop cell-type specific promotor to limit the vector expression to subpopulations in V1. There has been some work done towards this in rodents, making use of transgenic lines (Madisen et al., 2015), but so far the genetic tools in marmosets are still in development. A first step for this could be to label the interneurons and compare their tuning to the excitatory neurons that surround them as well as the maps revealed by wide field imaging. There is also the possibility of controlling the expression of vectors with labelling by tracer injections in other areas to probe how function relates to connectivity (Wachowiak & Cohen, 2001).

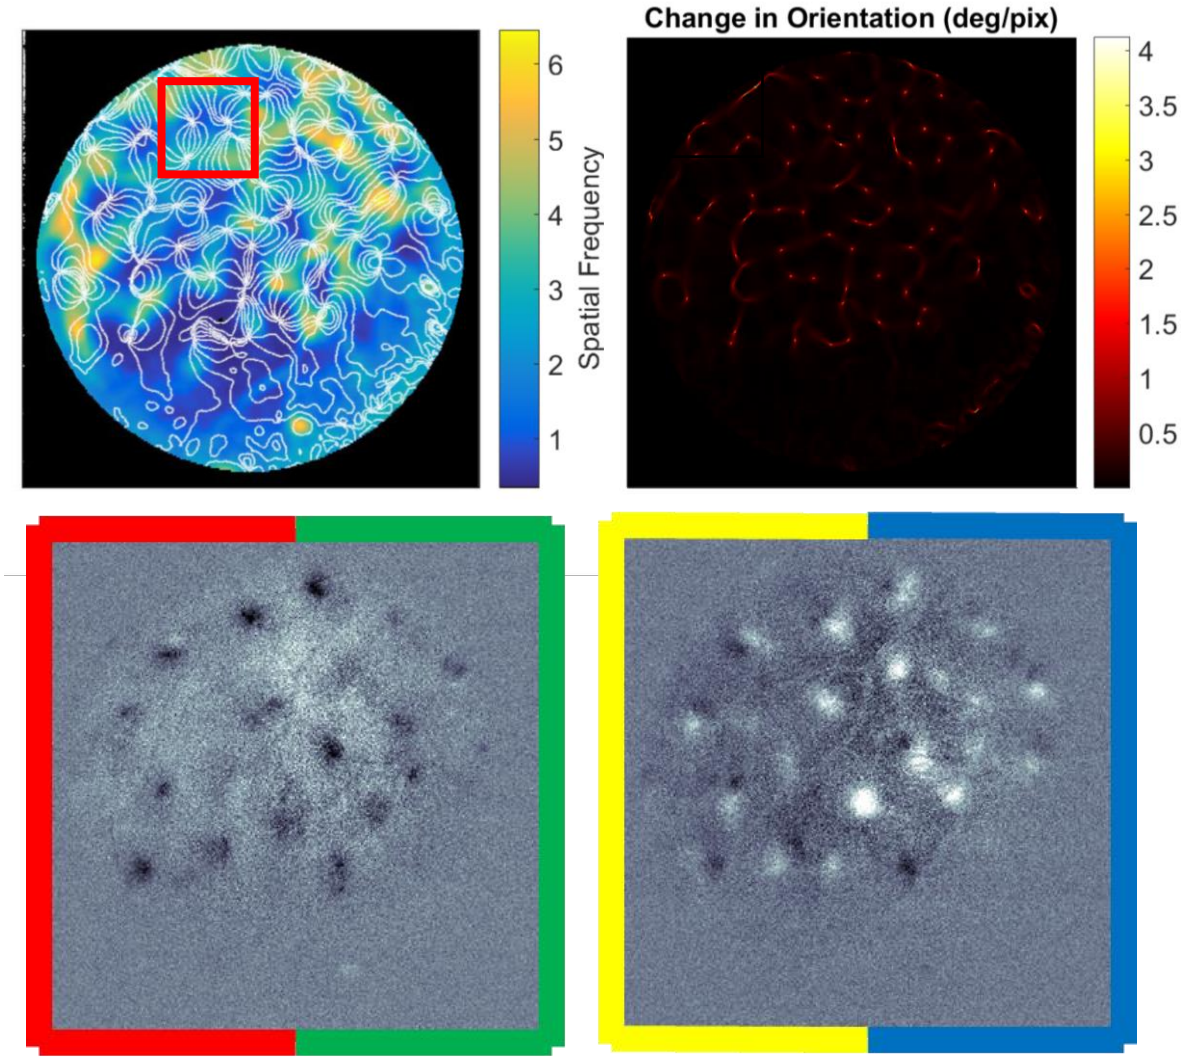

Figure 4.9. Additional maps. Top - left: Spatial frequency tuning map (cycles per degree) with overlaid orientation contours in white. Top-right: change in orientation per pixel. Below: Colour Opponency. The organization of sensitivity to a full screen coloured stimulus flickering between red and green (lower left), yellow and blue (lower right). Stimuli were luminance matched to avoid provoking a luminance related response. Colour responding regions are described more thoroughly in chapter 5.

This chapter has focused on the topology of orientation tuning in V1, however the data collected during these experiments also included many other stimulus parameters. Preferred spatial frequency also changed smoothly across the cortex, forming high and low spatial frequency selective domains. These can be seen in the upper left panel of Figure 4.9, where the pixel-wise preferred spatial frequency is shown with orientation contours superimposed. The prevailing hypothesis is that

the relationship between the topography of the orientation and spatial frequency selectivity maps is consistent with an architecture that is optimised for uniform coverage (Nauhaus, Nielsen, Disney, & Callaway, 2012). This suggests that the map gradients should tend towards orthogonality, with peaks and troughs of the preferred spatial frequency map and pinwheel centres of the preferred orientation maps being offset.

In the next chapter we continue to analyse the data across multiple animals and extend the methods to other features that V1 may have topological organisation of such as spatial frequency, colour, retinotopy and ocular dominance.

5 Functional maps in the marmoset primary visual cortex

5.1 Introduction

Neurons in primary visual cortex (V1) are selective for multiple visual features, such as orientation, spatial frequency, and colour (Hubel & Wiesel, 1968; Landisman & Ts'o, 2002b; Silverman, Grosof, De Valois, & Elfar, 1989). Tuning to each of these features varies smoothly across the cortical surface in what is known as a topographic map (Swindale, 1996). This chapter describes how multiple functional maps that encode diverse visual properties are configured to occupy the same space in primary visual cortex. Defining these relationships is important because it will give us insight into how feature selectivity is organised in cortex.

Spatial receptive fields in V1 are the smallest in the visual cortex, and in turn V1 has the smallest cortical area dedicated to the same point in visual space. The area of cortical surface dedicated to part of the visual field is termed the "cortical point image" (Capuano & McIlwain, 1981; Chaplin et al., 2013; Van Essen et al., 1984). The size of the cortical point image limits the numbers of feature combinations that be encoded for any part of the visual space. If one cortical 'minicolumn' (Buxhoeveden & Casanova, 2002; Mountcastle, 1997) is needed to represent each parameter combination, the number of minicolumns per cortical point image would put the upper limit of features that can be represented in V1 at around 8 to 10 (Swindale, 2000). Moving up the visual hierarchy the cortical point image increases in size (Gattass et al., 2005), the maximum number of features that can be represented increases and the complexity of receptive fields can expand into an increasing number of abstract dimensions. However, all this information must first pass through V1 and filter through the restricted feature space as dictated by the V1's physiology.

Given the competing constraints on the topographical organisation of V1 that stem from balancing a high-acuity representation with a range of fundamental visual features, we expect V1 to have highly optimised topographical maps with structured relationships between maps of tuning to different features. Elastic net models can distribute feature selectivity across cortical tissue in order to balance the smoothness, and coverage of the features (Swindale, 2004). However, the solutions an elastic net algorithm generates reveal strict geometrical archetypes that do not resemble the organisation of V1 (Keil & Wolf, 2011). This, and the fact that the number of features does not corrupt the structure up to a maximum number of feature representations defined by the structure of cortex, suggests that the organisation of V1 is suboptimal when the only constraints are coverage and smoothness (Swindale, 2004). Therefore, additional constraints outside of these parameters must be at play to drive the organisation of V1.

One possibility is that cortex is not a blank anatomical slate, as the elastic net model presumes. While organisation may develop towards an optimal solution in a way that resembles the elastic net algorithm, these maps need to first be 'seeded' genetically, with one or more maps constraining the others that develop around them. This chapter is focused on comparing overlapping maps in the same area of cortex. Although there have been many reports of multiple overlapping functional maps in the primate V1, it has been difficult to characterise these maps in the same animal at high spatial resolutions. Specifically, we can draw on the reported elements of V1 organisation to predict the structure of maps of previously studied features.

Notably, the cytochrome oxidase "blobs" are functionally distinct structures within V1, named for their cross-sectional appearance in cytochrome oxidase-stained sections, which could provide a framework for development of other maps (Tootell, Silverman, Hamilton, Switkes, & De Valois, 1988). Blobs are distinct from the interblob cortical tissue in anatomical terms, due to specialised

thalamic input (Casagrande et al., 2007). This localised specialisation may "seed" the development of V1 topography. We expect blobs to interact with the functional topography of other parameters because these regions may have predetermined tuning to other parameters bundled early in the colour processing pathway. For instance, blobs are typically associated with monocular and low spatial frequency tuned cells. (Hubener, Shoham, Grinvald, & Bonhoeffer, 1997), creating fixed points in the maps for spatial frequency tuning and ocular dominance. We expect the topography of spatial frequency tuning should be closely linked to the maps of blobs/interblobs as found through histological staining for cytochrome oxidase.

Additionally, we can make predictions about what to expect for the relationships between maps for a wide range of properties. Based on previous findings, we expect that orientation and spatial frequency will be orthogonally organised (Nauhaus et al., 2012). The eye/polarity grid model (Kremkow & Alonso, 2018; Kremkow, Jin, Wang, & Alonso, 2016) suggests that a key relationship underlying the organisation of V1 is an orthogonal relationship between the maps for ocular dominance and contrast polarity preference in a grid along the progression of increasing cortical magnification. Orientation is also orthogonally organised to contrast polarity in this model, due to a complex relationship between the progression of On and Off receptive field subfields.

Another model for V1 topology is the moiré interference pattern model (Paik & Ringach, 2011), also referred to as the 'statistical connectivity' model (Ringach, 2007). This model seeks to explain the tuning parameters of cells in V1 as a natural consequence of projecting the mosaic of retinal ganglion cells to V1 (Ringach, 2004; Soodak, 1987). The statistical connectivity model gives specific predictions about the organisation of maps. For example, iso-orientation domains should lie on a hexagonal lattice on the cortical surface. Additionally, it gives predictions about the relationships between multiple parameters, for example it predicts that pinwheels centres of the orientation

maps should align with spatial frequency tuning maps. Ringach (2007) goes on to suggest that instead of the cytochrome oxidase blobs "seeding" the topography, this process could happen in reverse. In his model, the patterns of connections can create regions of high metabolic demand with many highly-active, broadly tuned neurons. This high demand could drive these regions of V1 to increase the expression of cytochrome oxidase, creating blobs. This is supported by the fact that blobs have been reported in non-human primates that lack colour vision (*aotus*: Xu et al., 2004; *galago*: Condo & Casagrande, 1990).

While both the moiré interference pattern model and the 'statistical connectivity' model are based on experimental data, they were developed using cat V1 as the primary model and only later extended to primate models. In this chapter we will seek to compare model predictions with the results from our wide imaging from marmoset V1. Colour sensitive regions in marmoset V1 are predicted by anatomical studies (Casagrande, 1994; Livingstone & Hubel, 1984a). Intrinsic imaging work have recovered colour domains aligned to cytochrome oxidase blobs that are especially sensitive to red/green changes (Valverde Salzmann et al., 2012). However, so far the evidence for colour domains from single cell electrophysiology has been inconclusive (P. R. Martin, 2004). We expect with our more powerful imaging methods to be able to recover the properties of the colour sensitivity domains in marmoset V1 with more detail and precision.

This chapter describes how multiple functional maps that encode diverse visual properties are configured to occupy the same space in V1. Understanding maps of feature selectivity is important as these maps could be a general property of how information representation throughout cortex, beyond V1 and the visual system. A major benefit of using the wide field imaging techniques used here is its capacity to be used to investigate the relationships between the maps in the same animal

in the same tissue over a large field of view. The results presented here are probably the most comprehensive description of the organisation of multiplexed maps in marmoset V1 ever reported

5.2 Methods

To examine topography, we performed calcium imaging at both columnar resolution with wide-field one-photon imaging, and at single-neuron resolution with two-photon imaging. One month before the experiment a viral vector carrying the GCaMP6s calcium indicator was injected into the primary visual cortex (V1) of adult marmoset monkeys. Under a sufentanil anesthetised preparation, a suite of controlled visual stimuli was presented to the contralateral eye in order to provoke neural activity in V1. By correlating the evoked calcium signals to parameters of the stimuli we were able to quantify selectivity for the stimulus orientation, direction, spatial frequency and colour. All experiments were done in accordance to guidelines set out by the Institutional Animal Research Committee at RIKEN.

Preparation	Stimulus	Animal 1	Animal 2	Animal 3
Anaesthetised	Drifting Gratings	X	X	X
	Colour	X	X	X
	RFs	X	X	X
	Cone-isolating		X	X
	Ocular Dominance	X	X	X

Table1. Summary of the experiments that make up the results. .

5.2.1 Image processing

For single photon data registration was done in imageJ (Schindelin et al., 2012; Schneider et al., 2012) using a custom script to run Turboreg (Thevenaz et al., 1998) first over averaged sets of between 20-50 frames against a reference image and then within each subset to the mean. This provides a steady image even when there are large slow movements and also very small quick movements which are the main problems we faced due to animal moving in awake experiments. The same registration was done in both awake and anaesthetised experiments.

Each pixel of the single photon image frames were 'zero-mean' normalised in order to correct for differing levels of illumination, vector expression and cortex transparency. This procedure is similar to dividing each pixel by its time average or other representative frame, which are standard practices in optical imaging (Bonhoeffer & Grinvald, 1996). However, it differs slightly in that the mean is first subtracted (to normalise the absolute value) then divided (to normalise the variation).

5.2.2 Determining region of interest for analysis

A valid responding region was determined by taking a signal-to-noise (SNR) ratio over on the results for our most effective stimulus: orientated gratings. This was done by dividing the difference between the highest and lowest responses by their sum. This value was thresholded at 0.1 to create a mask to filter out blood vessels, areas outside of the window, areas lacking in illumination or virus expression and areas that for any reason do not respond to oriented gratings.

Over the responding areas, per pixel preference maps were created for orientation by calculating circular variance using the resultant vector method. This gives a resultant vector with a direction

that indicates the orientation tuning preference, and a magnitude that indicates the strength of the tuning. Spatial frequency tuning was estimated by calculating the centre of mass of the area under the curve and then refined and parameterised by fitting a 1D Gaussian function to the tuning curve.

5.2.3 Retinotopy

Retinotopy was estimated by fitting a 2D Gaussian function to the response profile created from the orientated bar stimulus. Reponses to the 7 orthogonal positions for each orientation were overlayed and averaged such that the intersections of highly exciting bars give a map of the receptive locations. This is a sparse white noise method that makes use of V1's robust response to oriented bar stimuli. In comparison, our initial attempts to obtain the retinotopy using a pair of black and white squares failed with the same number of trials. An estimate of the location of the centre of field of view, the foveal representation, was made from the progression of receptive field centres across the surface. Progressing over the cortex towards the foveal representation the distance between receptive field centres gets smaller and smaller as the receptive fields converge, this distance was matched to the retinotopic map of V1 described by Fritsches & Rosa (1996) and point that the receptive fields converged towards (in screen coorinates) was taken as an estimate of the centre of the visual field. This estimate was made to convert the receptive field centres from screen coordinates to degrees of visual field away from the fovea.

To quantify how strongly the signals were driven by changes in luminance, I compared the response to black and white bars with a ratio measure similar to an SNR, where the difference between responses is divided by their sum. This gives a simple parameter for how much of the variation in the

signal can be attributed to the change in the visual stimulus. This ratio was also used for comparing responses to coloured stimuli.

5.2.4 Colour

For coloured stimuli, we used the screen's native red, blue and green output channels as well as yellow (equal amounts red-green) and black and white. We adjusted the brightness of the full-screen colour stimuli so they were equal under all conditions according to a hand held photometer. All male and some female marmosets have dichromatic colour vision (Jacobs, Neitz, Deegan, & Neitz, 1996; Travis, Bowmaker, & Mollon, 1988), with a single short wave cone pigment expressed across the species. However, colour vision in marmosets is highly variable especially in females due to tri-allelic genetic expression of an additional single (dichromatic) or a pair (trichromatic) of cones with three possible long/medium wave pigments due to blood chimera (Kawamura, Hirai, Takenaka, Radlwimmer, & Yokoyama, 2001). To account for this, the RGB values of each stimulus condition were converted into cone contrast according to the cone types available to that animal (as determined by genotype) following the method in Valverde Salzmann et al. (2012). As indicated in Table 1, we also collected data for cone isolating stimuli which were designed to modulate the screen colour around a white point such that one or a pair of cone-types would be activated.

5.2.5 Map construction and comparison

After calculating the tuning of each pixel, maps of preferred orientation, spatial frequency, brightness and colour sensitivity were created by smoothing the pixelwise tuning with a Gaussian filter with a standard deviation of 5 pixels (28.85μm in our field of view). To compare between maps, the instantaneous gradients were determined for each pixel. The angles between the gradient

vectors of the different maps can be thought of as the angle of intersection between their contour maps and can reveal relationships between them.

5.2.6 Pinwheel centre location

As in chapter 4, for maps of preferred orientation drifting gratings of six spatial frequencies, between and 0.1 and 5 cycles/°, and 12 directions. Pinwheel centre locations are visible by eye, but they were quantitatively determined by finding singularities in the maps of preferred orientation where the preferred orientation changes rapidly over a small region of cortex. To do this, the peaks of the instantaneous gradient map over a threshold were found. Then for each peak, the 2D centre of mass of the clustered pixels above the threshold was taken as the pinwheel centre.

5.3 Results

To study the precise spatial relationships among functional maps, we conducted *in vivo* calcium imaging of neuronal activities using both single and two-photon imaging in the same tissue in V1. We expected that the relationships between topographic maps of different parameters should reflect that these representations are organised in a systematic way to optimise representations of all combinations of parameters. Animal 1 is used for illustration throughout the chapter. See Appendix for maps in Animal 2 and Animal 3.

5.3.1 Retinotopy

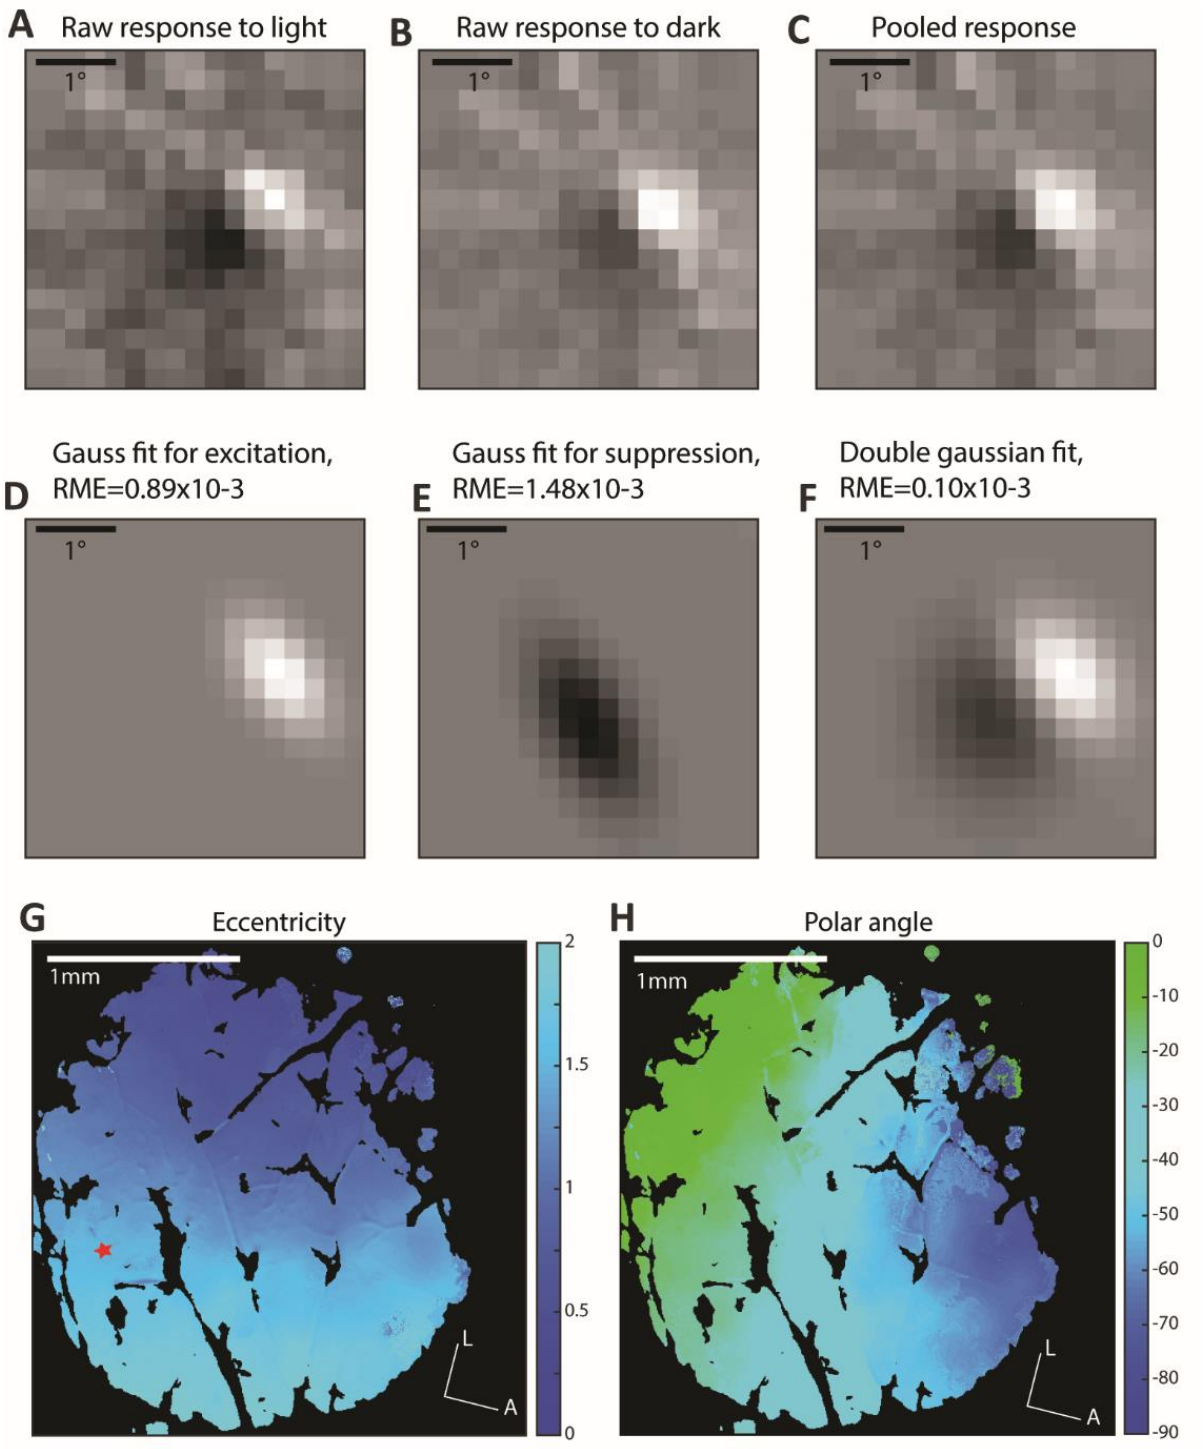

Figure 5.1. Mapping V1 retinotopy. For each single pixel recorded using one photon image we fitted a 2D Gaussians (lower row) to the smoothed raw RF (C) pooled from responses to light (A) and dark (B) bar stimuli. Retinotopic map for eccentricity (G), distance from the fovea, generated from the centre of the Gaussians fitted to the excitatory response line up with expected values on the dorsal surface of V1. Retinotopic map for polar angle (H), the angle subtended by the centre of the receptive field and the fovea where the horizontal meridian is 0° and lower vertical meridian is -90°.Red star in G indicates the location of the example pixel used for A-F.

To measure the receptive fields, we averaged the fluorescence responses for each flashed black and white bar across all repeats. Based on these responses, for each pixel in the imaging window, we obtained a spatial receptive field for both light and dark stimuli. We observed similar spatial structure of excitation and suppression for both contrast polarities. We pooled the responses across light and dark stimuli then fitted a 2d Gaussian to the excitatory response, and took the centre of the fit to be the centre of that pixel's receptive field in screen coordinates. We converted to eccentricity and polar angle using the location of the fovea (see Methods) to generate the maps in figure 5.1g and h.

Over the window we observe a lower field representation close to the centre of the visual field, with a smooth gradient over the imaged region. Its most central area (lowest eccentricity) is represented in the lateral side (top) of the figure 5.1g. In this view the V1-V2 border is just anterior (right) of the imaged region at which the receptive field centres approach the vertical meridian (-90 polar angle, dark blue in figure 5.1h). This is what we expect based on known marmoset V1 retinotopy on the dorsal surface (Fritsches & Rosa, 1996). Based on the retinotopy we have collected data that are relevant for foveal V1, but this may be less generalisable to the peripheral representation in V1.

5.3.2 Orientation columns have hexagonal symmetry

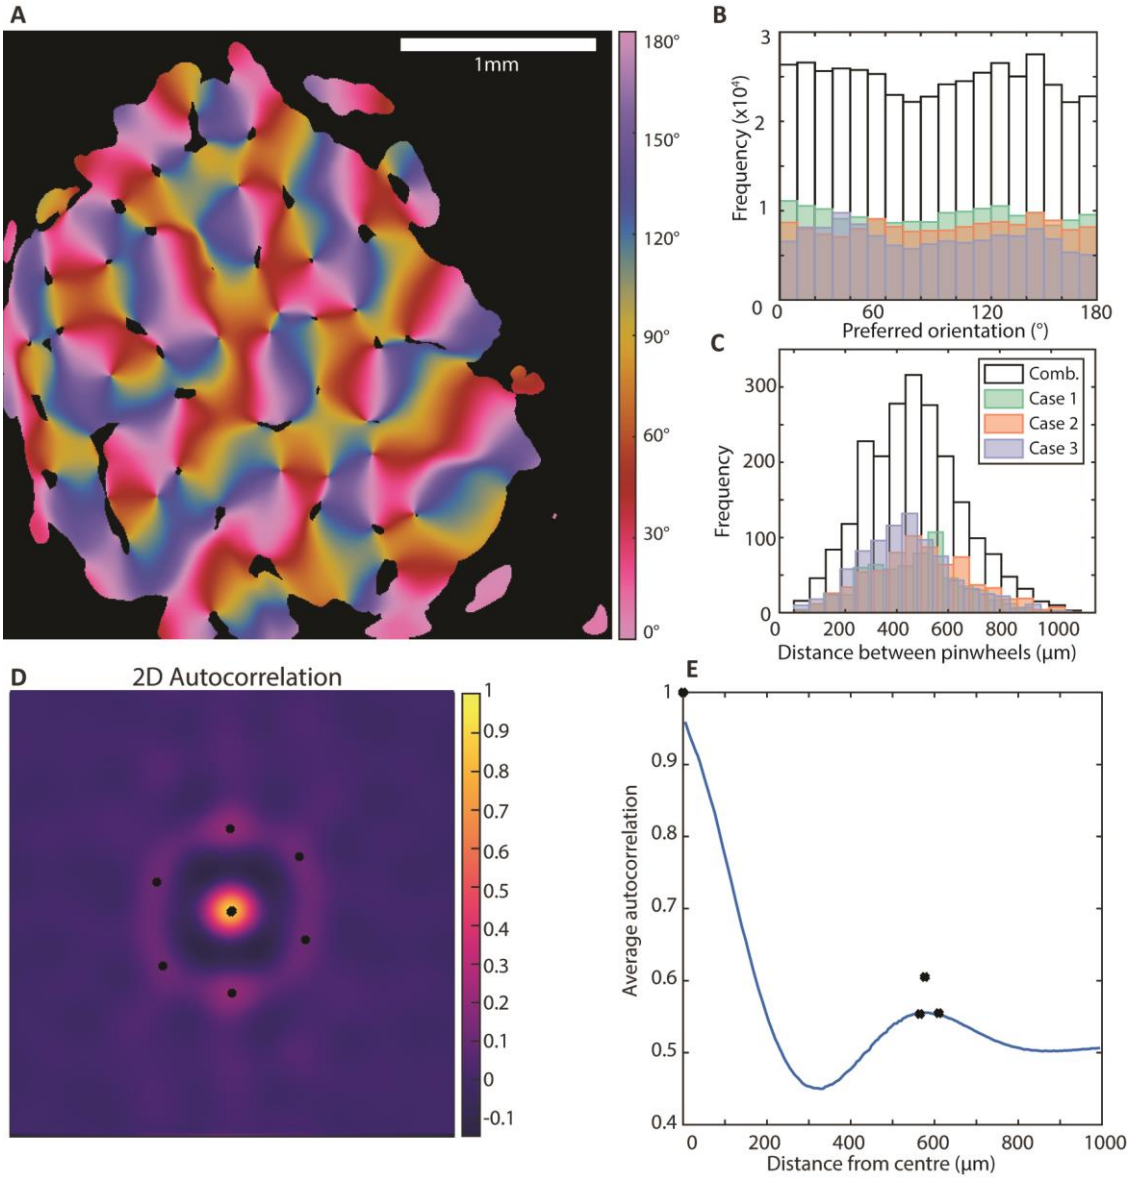

Figure 5.2. (A) Preferred orientation map. Preferred orientation is colour coded between 0 and 180 degrees. (B) Histogram of the orientation representation covered by these pixels (for all three animals). (C) A histogram of the distance between pairs of adjacent pinwheel centres (for all three animals). (D) Two dimensional autocorrelation of orientation tuning. Peaks in the autocorrelation show that the geometry of the orientation tuning is repetitive in a hexagonal lattice. (E) Collapsing over angles reveals a secondary peak at 570µm

We measured preferred orientation across the entire imaging window for each animal. In all cases, we observed the archetypical pinwheels and iso-orientation domains that we expected given the

well-established columnar organisation of V1 orientation tuning. We also found that geometry of the autocorrelation structure was generally hexagonal. This was done following a similar method to that of McLoughlin and Schiessl for intrinsic imaging (McLoughlin & Schiessl, 2006). The 2d autocorrelation was calculated for each direction condition of the stimuli and then combining with the strongest peak outside the centre at the top. This allows us to see the geometry of the repeating structure of the orientation columns (Figure 5.2d). Outside of the central peak, which simply shows that the orientation response is highly correlated to its neighbours, there is a ring of anti-correlation then six peaks in a ring where an orientation column of the same tuning is most likely to be. Collapsing over angles to plot correlations over radial distance there is a secondary peak at 570µm, with the three pairs of peaks at distances of 563, 571 and 607µm.

Interestingly, we do not see the overrepresentation of tuning to the horizontal and vertical orientations that has been widely reported previously in macaque (Obermayer & Blasdel, 1993). However, this lack of anisotropy is consistent with the literature of marmoset V1 (McLoughlin & Schiessl, 2006). The average distance between the pinwheel centres in this region was 466µm with a standard deviation of 190µm. This distance is close to, but lower than what has been reported previously for marmoset V1 (575µm, (McLoughlin & Schiessl, 2006)). Some of this difference may be due to the eccentricity of the region we imaged. The dorsal surface of marmoset V1 encodes visual information near to the centre of the visual field (Solomon & Rosa, 2014). The smaller orientation domains could be a consequence of tighter retinotopic constraints on a foveal representation. These differences could be partially accounted for by individual variability, as there is also evidence that the exact dimensions of columns can be variable within a species (Horton & Hocking, 1996).

5.3.3 Maps of preferred spatial frequency and preferred orientation are orthogonally organised

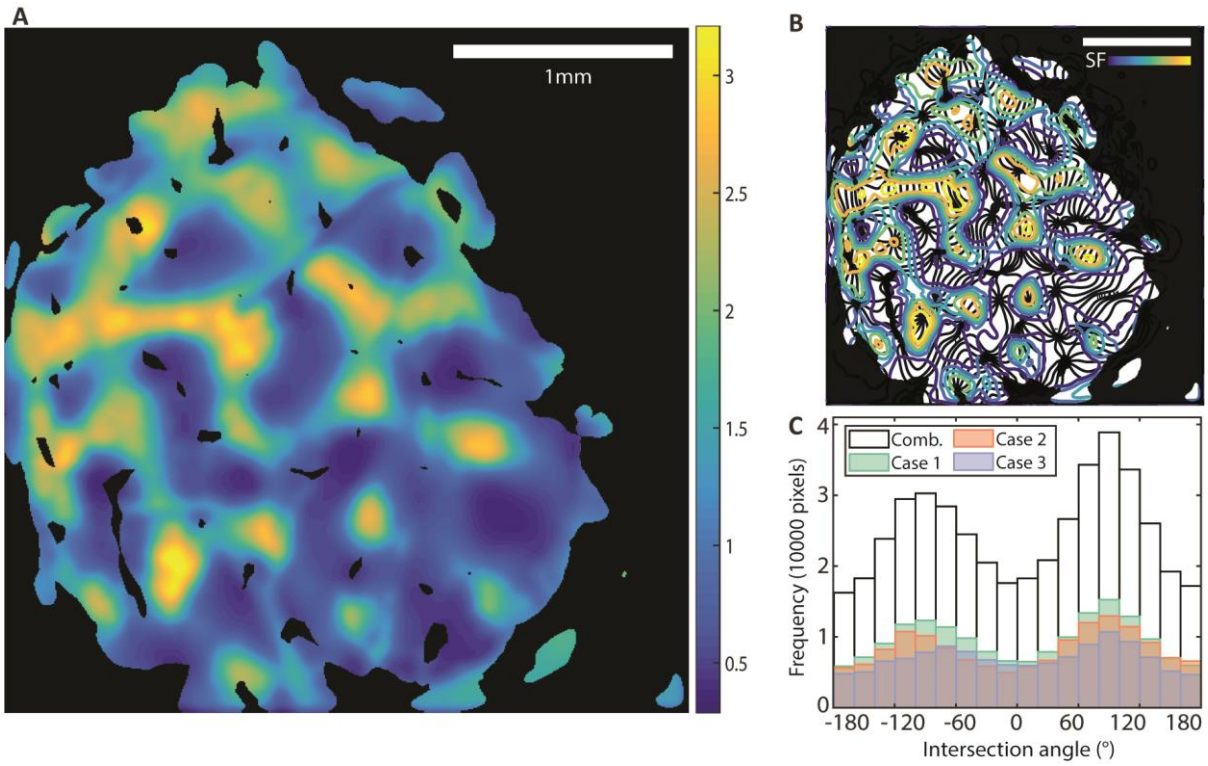

Figure 5.3. Maps of preferred spatial frequency and preferred orientation are orthogonally organized. (A) map of preferred spatial frequency with weakly responding pixels masked out. The resulting map creates surface with peaks and troughs rather than the bands previously reported in primates (Nauhaus et al., 2016). (B) Contour map with the same data as 2a overlapping contours from the preferred orientation map (1a). (C) Histogram of 'intersection angle' (for all three animals). The histogram of the angles between the gradients of the two maps reveals a tendency towards 'orthogonality'; the slopes of the surfaces tend to be perpendicular (-90 ° or 90 °) to each other.

The spatial frequency tuning curves extracted from the mean responses to each set of drifting gratings were generally weaker and noisier than the curves generated for orientation. In order to create a map for preferred spatial frequency from this data the preferred spatial frequency at each pixel was estimated from the tuning curve and smoothed across pixels using a two dimensional Gaussian filter. In previous work using intrinsic imaging, maps for spatial frequency are rarely reported due to the comparatively weaker signal to noise ratio (Roe et al., 2005). Due to these

technical limitations on previous imaging work, our map (Figure 5.3) may be one of the largest continuous examples for preferred spatial frequency in marmoset V1.

Smoothing may obscure some of the local rapid changes over the surface and blunt the peaks and troughs of the map, but it gives an unbiased estimate of the direction of the gradient vector anywhere in the imaged cortical region. The map shows clusters of high and low spatial frequency tuning as well as examples of bands of similarly tuned regions as reported in the macaque (Nauhaus et al., 2016). There is a tendency for the most foveal representation (top left in Figure 5.2a) to be tuned to higher spatial frequencies (Movshon, Thompson, & Tolhurst, 1978b; Yu et al., 2010).

From the contour map the centres of the orientation pinwheels, where the black contours converge, do not seem to have an obvious relationship with the peaks and troughs of the preferred SF map, in yellow and blue contours. However, by taking and comparing the angles between the instantaneous gradients of each map at every pixel we can determine that there is a tendency for these maps to progress in orthogonal directions. This can be seen in the histogram in Figure 5.2c with peaks at intersection angles of positive or negative 90°. An orthogonal relationship is an expected result for these maps, however the location of pinwheel centres in the spatial frequency tuning map is not clear from the contour maps. This will be addressed in more depth in subsequent sections.

All combinations of orientations were used to generate the map of preferred spatial frequency and vice versa in order to use the full dataset for each map. This may limit the interpretation of the relationship that we can gain from this data because tuning to orientation and spatial frequency are not completely independent. For example, cells that are tuned to low spatial frequencies generally have a broader orientation tuning for orientation. However the structure of the maps is unlikely to

be affected much by such interdependencies as these maps do not depend on changes in bandwidth of the tuning curves as the preferred value remains the same.

5.3.4 ON and OFF regions tile within low-sf domains

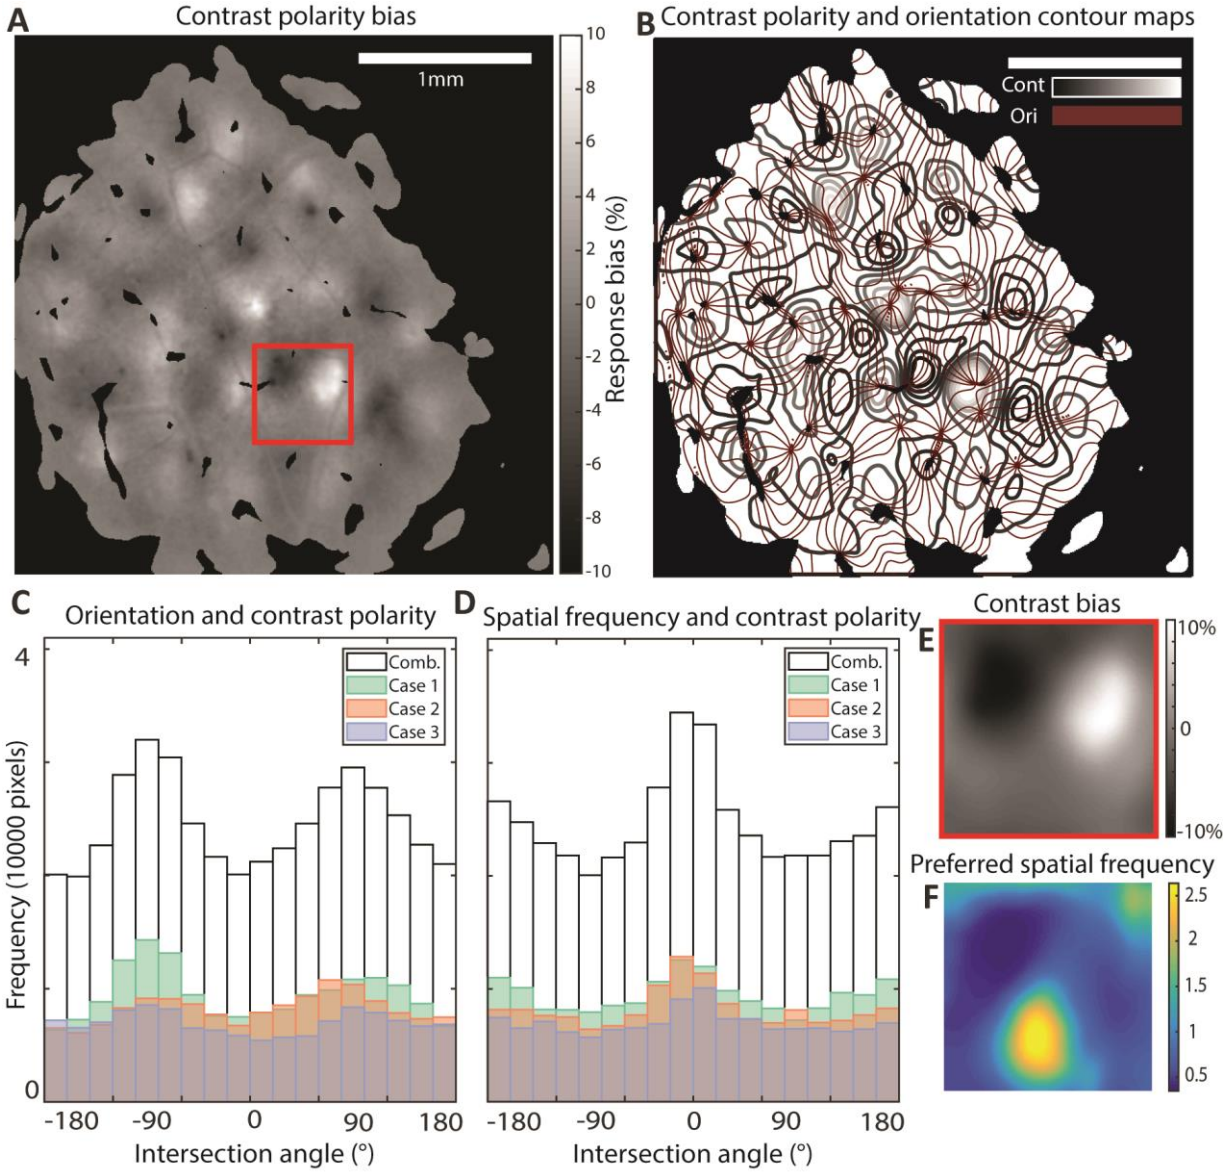

Figure 5.4. On and OFF regions tile within low-sf domains. (A) Map of preferred contrast polarity. (B) Contour for luminance (grey-scale, thicker) and orientation (red, thin lines) maps. Note that the pinwheels are offset to the extremes of the luminance map. The histogram of angles of intersection between gradients (for all three animals) reveals that the luminance map has an orthogonal relationship with orientation (C) and a parallel relationship with SF (D). An example region (red square in A) shows how on/off regions (E) tile within low SF (blue in F) selective regions.

Clusters of regions with a bias to a contrast polarity, either black or white, were found over the imaging region. Unlike maps for orientation and spatial frequency, the map of bias to contrast polarity (Figure 5.4a) does not show each pixel's preferred value after fitting a tuning curve. Instead, this map shows the pixel-wise average difference, as a percentage of maximum, between the fluorescence recorded during the presentation of white and black bars used for retinotopic mapping. Contrast selective regions in V1 have been reported in the ferret (Smith, Whitney, & Fitzpatrick, 2015), as well as cat (Wang et al., 2015), and some degree macaque (Kremkow et al., 2016). However, to our knowledge, this is the first time such regions have been reported in the marmoset.

Figure 5.4b shows the relationship between the organisation of orientation and contrast polarity bias contours. From simple observation it appears there is a slight tendency for orientation pinwheels to be near the regions with a bias to a contrast polarity, however it is difficult to tell from this alone what, if any, relationship these maps have to each other. However, we found that the angles of intersection between gradients for the map of contrast polarity and orientation tended towards orthogonality whereas for spatial frequency they tended to be parallel. This is in contrast with Smith et al. (Smith et al., 2015) who found orientation and bias to contrast polarity were not orthogonally organised in ferret V1, but instead coverage was maintained by different periodicities between the two maps.

We observed that the ON/OFF sensitive regions overlapped with regions of low spatial frequency tuning. This makes some intuitive sense in terms of receptive fields: responses that are dominated by one contrast polarity are equivalent to responses to a spatial wavelength larger than the width of the receptive field.

5.3.5 Maps of ocular bias and spatial frequency tuning progress in parallel

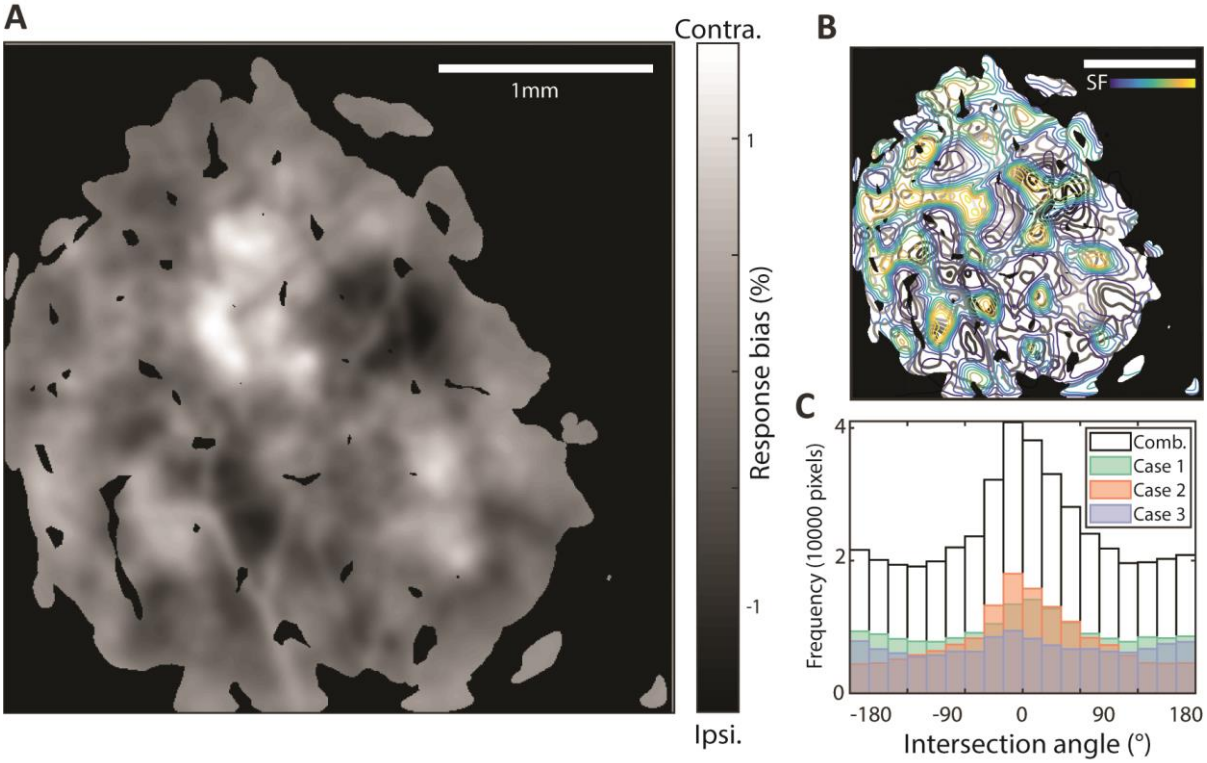

Figure 5.5. Map of ocular bias. (B) Contours for ocular bias (grey-scale, thicker) and spatial frequency (blue-yellow, thin lines) maps. Note peaks in the spatial frequency tuning map avoid the peaks and troughs of the ocular bias map. The intersection angles (for all three animals) between these two maps had a tendency to be parallel with a large peak at 0°

Organisation of tuning for ocular dominance into stripes has been well reported in cat, macaque and other primates including humans. However, in marmosets as well as squirrel monkeys, there is some debate about whether ocular dominance are present in adult V1 (Adams & Horton, 2003; Roe et al., 2005). Even when they are present in marmosets they are a lot less defined than in macaques (Chappert-Piquemal, Fonta, Malecaze, & Imbert, 2001; Sengpiel, Troilo, Kind, Graham, & Blakemore, 1996). We measured responses to drifting square wave gratings while one eye was covered. We alternated between displaying stimuli to the ipsilateral and contralateral eyes for each trial of 48 randomised repeats of the stimulus. The averaged difference in responses to matched stimuli for

each eye gives the map of eye preference in figure 5.5. Bands of cortex that respond preferentially to the contralateral (white) and ipsilateral (black) eye can be seen across the imaging window.

Our map show regions of biases in the strength of responses to the contra and ipsilateral eyes which we treat here as approximations of the ocular dominance columns, but in reality these biases are in the order of 1% in the difference in response. Patterns of weak ocular bias has been reported for marmoset previously (Roe et al., 2005), and it has been suggested that these regions could be the remains of ocular dominance columns left over from development but ablated in the adult cortex. For the purposes of comparison, we treat these biases here as approximations of the ocular dominance columns, but the generalisability of these results is not clear at point. Additionally the use of square gratings was chosen to evoke as much activity as possible in V1, but could be additional source of uncertainty in our ocular bias maps. Since this stimulus has a lot of orientation information, even though stimuli are matched between the eyes for comparison, small discrepancies in the orientation of the stimulus between the two eyes could be created by torsional eye movements. This may obscure the relationships between orientation maps and the maps for ocular bias, or may even be responsible for part of the differential response we see.

Overlaying the contours of the map to spatial frequency tuning. We can observe tuning to both high and low spatial frequency in bands of ocular dominance for both ipsi- and contra-lateral eyes. The moiré interference pattern model (Paik & Ringach, 2011) predicts that the highest spatial frequency tuning should align to ocular balanced regions. This has been reported in the macaque (Nauhaus et al., 2016). We found that although the gradients of the two maps were organised in parallel as reported in Nauhaus 2016, the relationship was the inverse of what the reported. We find that the most monocular biased regions corresponded to regions of higher SF (figure 5.6d) tuning compared to the unbiased regions.

5.3.6 Relationships between maps

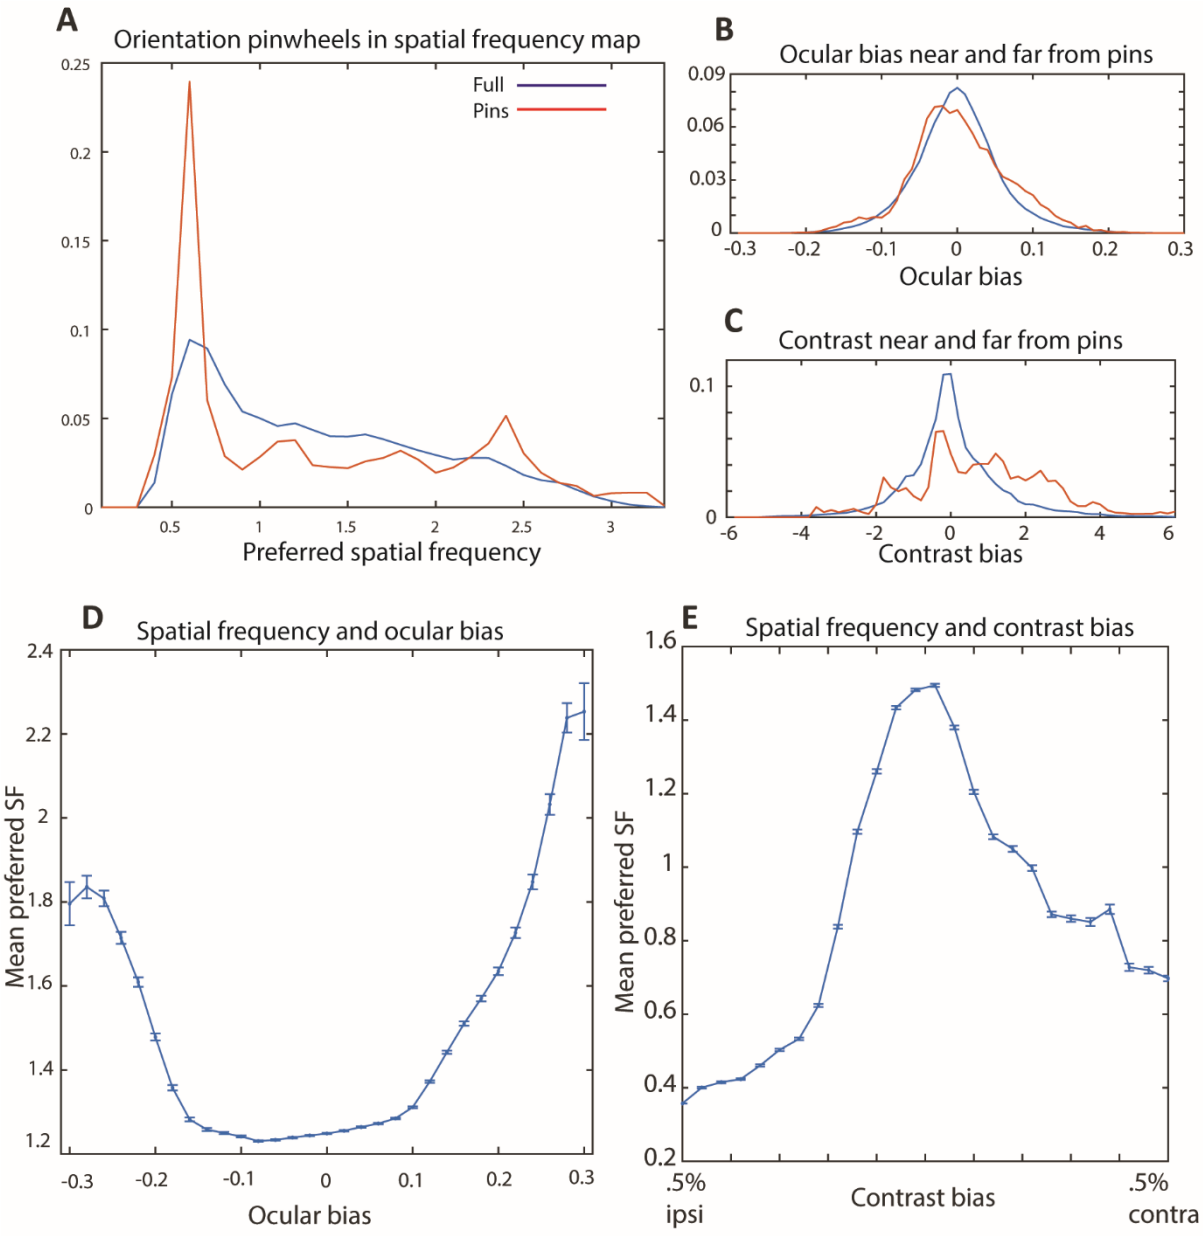

Figure 5.6. Relationships between maps for a single animal (Animal 1). Distributions, normalised to sum to one, of tuning for preferred spatial frequency (A) for pixels within 50µm of a pinwheel centre (red) against the distribution for pixels outside this distance reveal peaks in the distribution of tuning near pinwheels at both low (0.7) and high (2.4) ends of the spectrum. Differences in distributions of ocular bias (B) and contrast (C) tuning near and far from pinwheel are not obvious. Mean preferred spatial frequency for pixels binned according to ocular bias (D) and contrast bias (E). Pixels with a monocular bias to either eye tended to prefer higher spatial frequencies than regions with no ocular preference.

The distribution of angles of intersection between the gradients of two maps is a very useful tool for describing how the maps are organised relative to each other. However, this description does not include local properties such as alignment of key points of interest and correlated relationships between parameters. For this we looked directly at the tuning properties of regions of cortex immediately surrounding (within 50μm) of a pinwheel centre and compared them to the tuning properties of cortex outside these regions. The locations of pinwheel centres were estimated to the nearest pixel by finding the centres of peaks in the double derivative of the orientation tuning map.

In figure 5.6 histograms for the spatial frequency at pinwheel centres has peaks at 0.6 cyc/deg and 2.4 cyc/deg when compared to the distribution of tuning over the rest of the window. These values are at the extremes of the measured spatial tuning preferences over the window. This discounts our naive observations of the contour maps in figure 5.2 that the pinwheels are not aligned to the peaks and troughs of the spatial frequency map. With all pinwheels combined the relationship becomes clear. Additionally, this does fit in with what we expect based on moiré interference model (Ringach, 2007) which predicts that pinwheel centres would be located at the extremes of the SF tuning map.

We observe that the distribution of bias to contrast polarity at pinwheel centres is similar to the distribution of bias to contrast polarity for the rest of the window. This suggests that although orientation tuning and contrast polarity are organised orthogonally, the locations of the pinwheel centres are not aligned to regions of dominated by the responses to one of the contrast polarities. This is in contrast to the predictions of the eye/polarity grid idea put forward by Kremkow et al (Kremkow & Alonso, 2018; Kremkow et al., 2016). The model suggests that the smooth progression of receptive field centres is maintained by subregions of one contrast polarity (OFF in Figure 4 of

Kremkow et al. 2016) while the receptive field subregions for the other polarity rapidly move relative to the cortical surface to create the orientation pinwheel. They suggest that at a pinwheel centre the receptive field subregion for one contrast polarity is fixed to a point in visual space and dominates the response profile. Unfortunately, using these methods it is difficult to draw strong conclusions for the comparisons in the distribution of bias to contrast polarity. It may be that there is a relationship that is not apparent from simply comparing the distributions.

5.3.7 Colour sensitivity in low spatial frequency selective regions

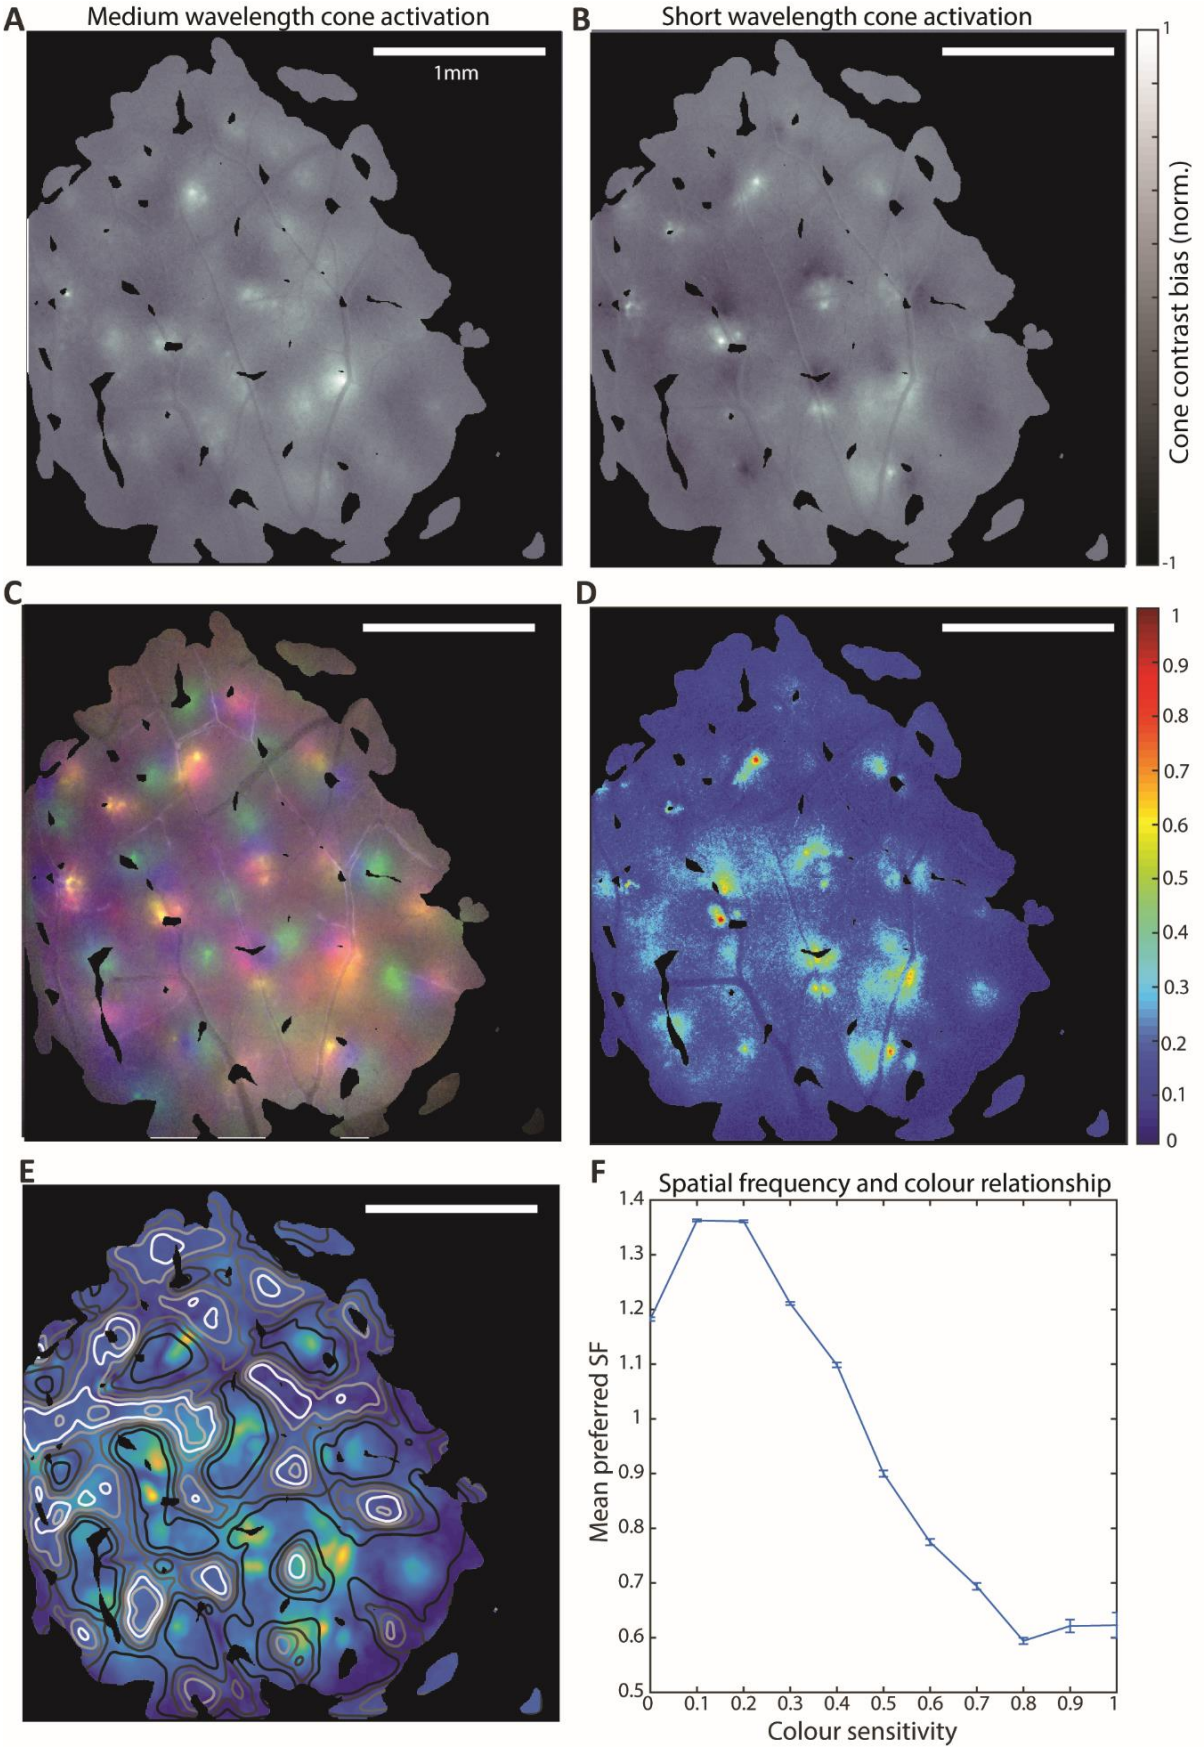

Figure 5.7. Color sensitive regions overlap with low spatial frequency selective regions. Sensitivity to contrast changes along the medium (A), and short wavelength (B) axes. 3C. Colour coded map to loosely simulate a map of colour preference. D) Normalised signal range driven by colour modulating stimuli. E) Same as D but with contours of the preferred spatial frequency map show the colocalisation of the low (darker lines) spatial frequency sensitive regions with the regions sensitive to colour changes.

Responses to isoluminant full screen colour stimuli reveal clusters of responsive regions of cortex for both types of cones in the retina of this animal (animal 1). For this animal, best estimates of cone-contrast was calculated for each of the colour stimuli used given the spectral properties of the monitor and the genotype of the animal. For the following animals cone-isolating stimuli were generated prior to the experiment and were used to create similar maps without the need for estimation of cone-contrast for coloured stimuli which would jointly activate multiple cone-types.

Combining responses from the full set of colour stimuli, we created a map of percentage signal range driven by colour. The analysis found the greatest activation for each pixel in response to changes between opponent colours as well as every individual colour tested alone against a luminance matched grey. In order to create a simple map to comparatively show the regions of cortex that are sensitive to colour, independent of the preferred colour, each pixel is normalised against the highest responding pixel in the imaging window. The regions of colour responding cortex tended to be aligned to the cortical areas that were tuned to low spatial frequency. Across animals this relationship is consistent (Figure 5.7f): the areas that responded most to the full screen colour contrast stimuli tended to be also sensitive to low spatial frequency. This relationship is expected from the anatomical structure of blobs, which are associated with the colour afferents from the retina as well as lower spatial frequency tuning.

Unfortunately our colour stimulus is a confounding factor for comparing with spatial frequency tuning as a full screen stimuli is necessarily a low spatial frequency stimulus that would (if not perfectly isoluminant) be expected to drive low sf responding regions. This makes it difficult to be certain whether the response we see is actually a response to colour or a response to luminance residuals by achromatic low spatial frequency preferring units. The ambiguity associated with using the full screen stimulus can be teased apart by using coloured gratings in the future using this type of stimuli, but we are not ready to present this data at this time.

This map is not the same as maps for other stimuli such as orientation and spatial frequency as the relative amplitude of the response, rather than the preference is shown across the imaging window. In the future, this analysis could be extended to include a map of estimates of preferred colour and its relationships to other maps. While this could be approximated with the current coloured stimuli, these estimate would probably benefit from a few more samples across the relevant colour space.

5.3.8 V1 function and histology

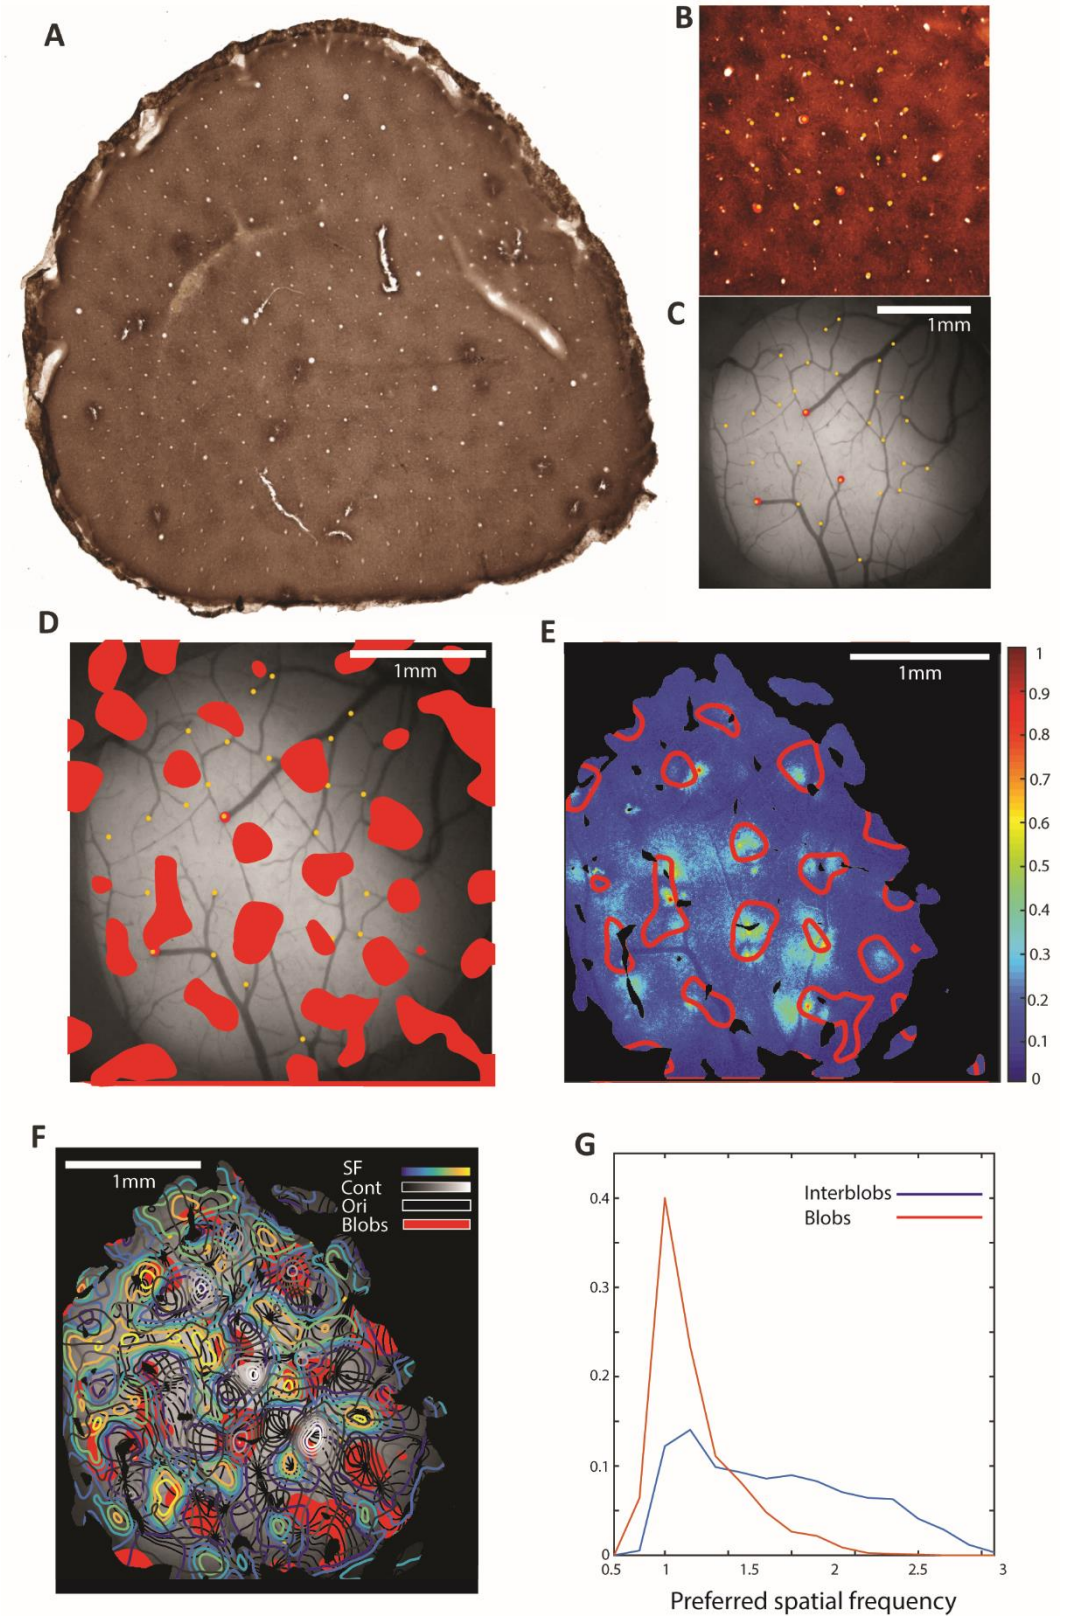

Figure 5.8. V1 function and blobs. A) Cytochrome oxidase staining shows blobs in tangential slices of V1. B) Registration was performed using the tracks left through depth by large vertical blood vessels, Yellow circles in B, C and D. From this, the location of the blobs in the imaging window (red in D) was estimated. E is the same as 7D; Normalised signal range driven by colour modulating stimuli, but with the outline

of the blobs overlayed in red. F) Same as D, but with contour maps for preferred spatial frequency (blue-yellow) and orientation (thin black) as well as contrast bias (thick grey-scale). G. The distribution of preferred spatial frequencies at blobs (red) overrepresented the lower end of the range compared to the interblob regions (blue) for animal one only.

Histological staining for Cytochrome oxidase (CO) revealed the stereotypical blob structure of marmoset V1 (figure 5.8a). After co-registration of the histological sections to our imaging window we observed that the regions tuned to high spatial frequencies avoided alignment with the CO blobs (figure 5.8g). This is consistent with previous reports (Tootell, Silverman, Hamilton, Switkes, et al., 1988) that the cells in blobs receive inputs from lower spatial frequency tuned thalamic afferents. Additionally we found that the colour sensitive regions, which are also associated with lower spatial frequency regions (5.3.7), were colocalised with the CO blobs. This is what we expected from previous reports that colour sensitive regions are aligned with CO blobs in marmoset as well as the greater literature on blobs in primate V1 which links them to colour processing (Landisman & Ts'o, 2002a; Valverde Salzmann et al., 2012).

5.4 Discussion

The results presented here are likely the most complete description of primate V1 topography at this scale. This is especially true for the maps for preferred spatial frequency, contrast polarity and ocular dominance which have been difficult recover from optical imaging methods. This is also the first time all of these parameters could be collected for the same part of cortex for the same animals for comparison. One qualifier is that all our recordings were made very close to the border between V1 and V2 (slightly outside the window, right side in figure 5.1g&h). The topology of this section of V1 is subject to additional restrictions in that it needs to maintain smoothness with representation across the border into V2 which is organised very differently to V1 (Rosa et al., 1997).

Nonetheless, we have shown that orientation, spatial frequency, colour and ocular bias and contrast bias are topographically organised in marmoset primary visual cortex, and that these maps are formed with systematic relationships to one another (see summary in Figure 5.9). While some of the maps were organised orthogonally to each other, many were not.

5.4.1 Intersection angles of maps

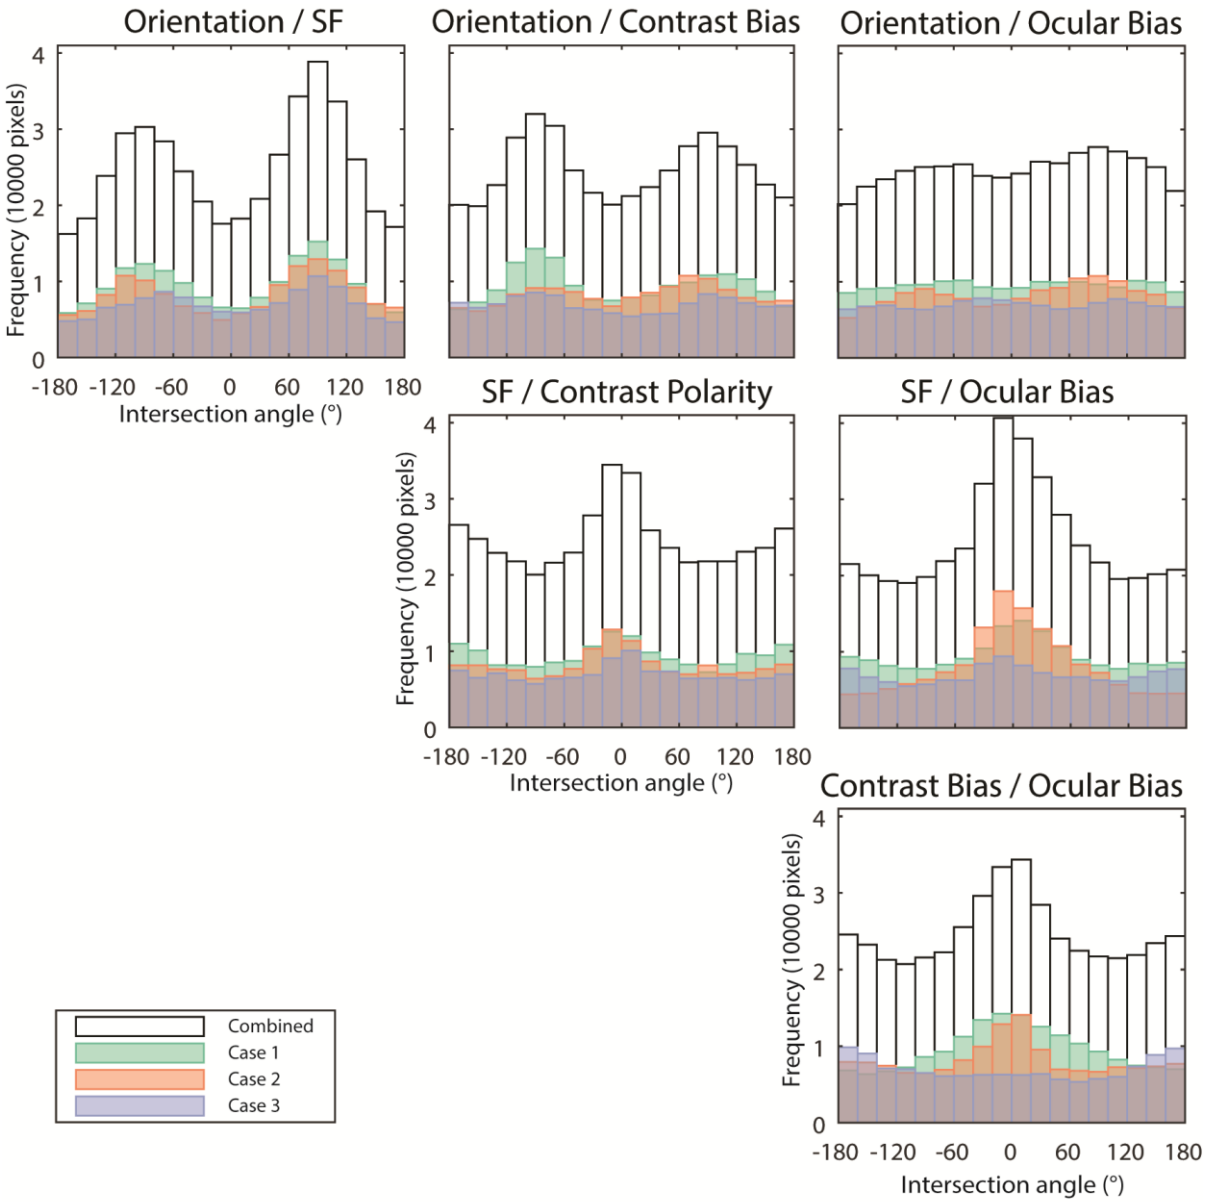

Figure 5.9. Co-organisation of V1 topographic maps. Combining histograms across three animals we can observe the relationships between each of maps that we recovered. Orientation was orthogonally organised relative to spatial frequency and contrast bias.

This study of intersection angle has revealed some unexpected relationships. Kremkow et al. (2016) reported observing orthogonal progressions of ocular dominance and contrast polarity bias in cat and to a limited extent macaque, but here we see the exact opposite; a parallel organisation. In fact,

we observe orthogonality only when looking at the relationships involving orientation maps, but not for any other pairs of parameters. It seems as though that the forces driving organisation of marmoset V1 are not using orthogonality as a way to achieve maximal coverage of these parameters. It should be noted that in one case (animal 3; blue in Figure 5.9 lower right panel), we did not see this parallel organisation which suggests that there is some variation in the organisation of these maps across individuals.

However, this does not mean that V1 is not optimising for coverage. Smith et al. (2015) found that polarity and orientation maps in ferret V1 show high coverage despite a lack of orthogonality. They attributed this high coverage to the difference in the sizes of hypercolumns for the two maps. Additionally, high coverage could also be achieved without orthogonality if hypercolumns are spatially isotropic (Swindale, 1991). For a full description of how well topographic maps cover parameters additional metrics for domain spacing and shape would need to be calculated for each of these maps as well as a metric for the coverage between each pair of maps.

5.4.2 Colour blobs in marmoset V1

We identified cytochrome oxidase blobs in histological slices and registered them to the imaged region of cortex to compare to our functional imaging. We observed that the blobs aligned with both the colour sensitive regions of the cortex and the low spatial frequency regions. Embedded colour blobs restrict the organisation of spatial frequency tuning. In marmoset the timeline of blobs formation is unclear, but in macaques, blobs form before birth and even if the animal has no eyes (Dehay, Giroud, Berland, Killackey, & Kennedy, 1996; Purves & LaMantia, 1993). Although the organisation of V1 will tend towards optimal organisation the blobs create fixed points of preferred

spatial frequency which the maps have to learn around. This could explain some of the sub-optimality of V1 topography compared to an elastic net algorithm without such constraints (Keil & Wolf, 2011).

For the purpose of comparing maps, we combined the results from animals with dichromatic and trichromatic colour vision. This could cause problems if the different genotypes have different connectivity of colour-processing centres. However, there is good evidence that the anatomy of early visual pathways are is similar for dichromats and trichromats (Chan, Martin, Clunas, & Grunert, 2001; Ghosh, Goodchild, Sefton, & Martin, 1996; Goodchild & Martin, 1998; Solomon, 2002; Wilder, Grunert, Lee, & Martin, 1996). For our analysis, we ignored the differences between the cone expression of the animals and cone contrast for all available cones was combined to give a simple single parameter for colour sensitivity. This remains as a potential area for further research.

5.4.3 Regions biased to contrast polarity and colour occupy low SF regions

In figure 5.4 we show that bias in response to contrast polarity can be seen in regions of cortex that tended to be tuned to low spatial frequencies. Additionally, we found that the colour sensitive regions also tended to be located in low spatial frequency selective areas (figure 5.7). However, the precise relationship between the locations of peak sensitive to contrast polarity and regions driven by cone contrast was beyond the scope of this experiment. We could speculate, however, that V1 would likely include regions for representations of as many combinations of contrast polarity and colour as possible in order to effectively code luminance and colour information. This could suggest another orthogonal mapping might exist between these two maps.

5.4.4 Moiré interference pattern model

The moiré interference pattern model is a model that relates the organisation of retinal ganglion cells to the topology of orientation tuning in V1 (Paik & Ringach, 2011). One of the key features of this model is the hexagonal geometry of the two dimensional autocorrelation of the preferred orientation map. This has been reported before (McLoughlin & Schiessl, 2006) in marmosets. In the cases where we see weaker hexagonal structure, the imaging window was over the dorsal surface of V1 very close to the border with V2. In this area the organisation of orientation tuning may become disturbed to accommodate a smooth transition to longer bands of orientation columns in V2 (Rosa et al., 1997). This could represent a larger limitation of our results given that the imaged regions were in all cases from the foveal representation near the border of V2.

In the future, it could be probably useful to align the autocorrelations of the orientation maps across cases in a more meaningful way. In Figure 5.2d, autocorrelations were aligned across cases by simply matching the position of the global maximum to the y axis based on the same technique used for optical imaging. Aligning the autocorrelation to the V1/V2 border could be a better solution for this problem, if the border could be accurately determined from the histology. In some preliminary testing of this, one of the cases had a slightly elongated autocorrelation along the axis in the direction of the V1/V2 border, but not the others. This could suggest that the imaging window for this animal might have been implanted closer to the rostral border of V1 than the others, close enough to experience some effects of the border on its orientation map. Investigating how the map for preferred orientation is organised across the areal border could be an interesting area for future study.

In addition to the hexagonal structure of the orientation maps, the moiré inference model makes a few key predictions about the location of pinwheel centres relative to the peaks and troughs of the preferred spatial maps. Our histograms of the distribution of preferred spatial frequency at and around pinwheels compared to the pixels further away reveals a tendency for pinwheel centres to sit at the peaks and troughs of the spatial frequency map. This finding is most apparent in a histogram (i.e. Figure 5.6a) rather than in the contour maps for the single imaging sites.

Additionally an updated version of this model seek to include ocular dominance in the model and predicts relationships between ocular dominance, spatial frequency and orientation (Song, Jang, Kim, & Paik, 2018). It predicts that orientation pinwheel centres will be located in the centre of ocular dominance columns. This has been reported previously in cats and some primates (Obermayer & Blasdel, 1993). As discussed above, the map for ocular bias presented here may not generalise to the ocular dominance columns in other animals due to the lack of well-defined ocular dominance columns in marmoset. This is important because the lack of ocular dominance columns across the surface of V1 would affect the organisation between these maps.

In the case of ocular dominance columns as in the macaque, if the orientation pinwheel centres were systematically located at the centres of monocular regions we would not expect to see a strong tendency toward orthogonality between these maps. In this organisation, orientation domains would be as likely to extend from the pinwheels in monocular regions along the ocular dominance column as they would to cross to the next column. Our results (figure 5.9) show no clear relationship between ocular bias and orientation, which is consistent with the organisation in the macaque. Additionally, the moiré inference model predicts for the high spatial frequency-tuning at regions with bias to one eye or the other that was reported by Nauhaus et al. (2016) in the macaque and is consistent with our results (Figure 5.6d).

5.4.5 Eye and polarity grid model

The eye and polarity grid model is based on the receptive field substructure revealed using linear probes inserted tangentially through layer 4 of V1 (Kremkow et al, 2016, 2018). They found that subregions for one polarity anchored retinotopy at a pinwheel and the subregions for the other polarity rapidly changed. However, we were not able to see systematic differences in the receptive fields we recovered for black and white bars at this scale. This could due the limitations of spatial resolution for our one photon signal. In order to fully understand the retinotopy of both the On and Off subfields we would need to deconstruct our RFs into On and Off and compare especially around pinwheel centers. This would be best left to high resolution single cell recordings with two-photon imaging rather than the single photon methods used here.

However, some predictions of the model can be tested with single photon imaging. For example, a key feature of this model is that the map of eye dominance and contrast polarity should be orthogonally organised. We found that this was not the case in our results (lower-right panel of figure 5.9). Instead, we found that these maps were organised in parallel for our imaged regions. Another prediction is that the centres of pinwheels should be located at the regions of that are biased towards one or other contrast polarity (Alonso, 2016). Unfortunately, our results for this were unclear (figure 5.6c), however we found that overall the organisation of the orientation tuning and contrast maps were orthogonal (figure 5.9) which supports the possibility of co alignment between pinwheel centres and regions biased to contrast polarity.

5.4.6 Speculative model for V1 topology

The speculative model for V1 topology is that the maps are seeded by initially with patterns from retinal ganglion cells as in the moiré inference model and are then optimised during development to maximally represent all combinations of parameters. However, for parameters other than orientation we do not observe orthogonal relationships, suggesting that another strategy is being used to optimise coverage. Additional constraints on V1 organisation could be created by blobs, which dedicate regions in V1 for the processing of colour information and broadly-tuned, low-spatial frequency vision (Casagrande, 1994; Livingstone & Hubel, 1984a). During development, the map of preferred spatial frequency may need to organise itself around these fixed points. In turn, the rest of the maps also need to warp to accommodate the constrained tuning at the blobs.

Alternatively, the cytochrome oxidase blobs could form later as a consequence of the development of the orientation tuning maps as in Ringach's (2007) 'statistical connectivity' model. In this interpretation, the orientation tuning map is the foundational map that constrains the development of all other maps. In either case, if the colour blobs are devoted to the processing of low spatial frequency and colour (Hubener et al., 1997), interblobs may be dedicated for parameters that rely on high spatial frequency such as disparity. Interestingly, we do not see evidence of eye/polarity orthogonality and while we see some evidence for a complex relationship between orientation tuning and bias to contrast polarity a full mapping of On and Off receptive field subregions is needed to fully understand this relationship.

6 Discussion

There were two main aspects to my PhD candidature, united by the theme of deciphering the functional organization of the visual cortex in the marmoset monkey.

Reverse correlation analysis in DM revealed receptive field filters that suggest that neurons in DM encode information about long contours. Paired filters, similar to those of complex cells, were commonly found for DM cells as well as additional excitatory and suppressive filters. These combined with the preliminary directions of invariance results suggest that DM encodes information beyond just contours but further work is needed to characterise this fully.

Calcium imaging in V1 revealed overlapping maps of selectivity to multiple parameters with organisations that are sometimes orthogonal but often not. The imaging results were consistent across scale and imaging methodology, and consistent with the literature on orientation tuning. The wide-field imaging methods were then used to more thoroughly describe the organisation of multiplexed maps in marmoset V1 than previously reported.

In both cases, the achievement of scientific objectives required a significant component of method development, which ultimately allowed me to interpret data of a new nature.

6.1 Determining and interpreting feature selectivity

The advantage of using white noise analysis is that it samples from the set of visual stimuli in a random and unbiased fashion. However, this randomness is also its most salient disadvantage. The

vast majority of the stimulus frames, and therefore recording time, are not relevant to the behaviour of a neuron. Also, in the space of all possible patterns sampled the space of relevant stimuli could be so small that relevant patterns would occur infrequently or not at all and any resulting excitation or suppression would be obscured by noise. For the white noise to be effective it needs to sample the stimulus space that is most relevant to the neurons behaviour, however there are many patterns that are simply too unlikely to occur by the random sampling to have a significant effect on the spike rate of the neuron.

This disadvantage could be mitigated by actively modelling the receptive field of a neuron in real time and updating the stimulus to search actively search stimulus space instead of randomly sampling (Okazawa, Tajima, & Komatsu, 2015; Ponce et al., 2019). This technique, adaptive stimulus generation, involves optimising the stimulus for probing the behaviour of a single neuron (Benda, Gollisch, Machens, & Herz, 2007; DiMattina & Zhang, 2013). However, the technique of actively searching for relevant stimulus space would have to be applied to cells one at a time. White noise stimulation, as we used it, can be run for many neurons simultaneously. Using an adaptive stimulus, the recording time needed to characterise each cell would be much less than that for white noise as the stimulus space would be restricted, but in the trade-off between displaying more stimuli to all cells in parallel and less stimuli serially to individual cells, it is unclear which would be more efficient.

Another challenge of white noise stimulation is in characterising surround effects such as surround excitation and suppression that have modulating effects on the response of a neuron to the stimuli in its receptive field. It is difficult to see how this effect could be captured by the white noise stimuli: the neuron needs to be driven by stimuli inside the receptive field before the modulation can be observed. One strategy to characterise the non-classical receptive field structure of a neuron using white noise analysis could be to fit a model of the classical receptive field, and then compare that

model's expected firing rate to the actual rate for each frame of the white noise stimulus. The reverse correlation procedure could then be repeated on the residuals. However, as both the expected effect and the space of relevant stimuli are very small, the number of stimulus frames would be prohibitively large for this approach to be feasible experimentally.

A better strategy may be to use what we have learned from the white noise analysis to control the response to the classical receptive field and then probe the effects of stimuli in the surround. One way of doing this is to use drifting gratings in circular and annular apertures to determine the size of the non-classical receptive field as well as the nature of the modulating effect (Cavanaugh et al., 2002). In V1, the modulation due to stimuli extending outside the classical receptive field is typically suppressive (Hubel & Wiesel, 1968) and orientation tuned (Bair et al., 2003). It was thought to be driven by lateral connections (Gilbert & Wiesel, 1983), but it has more recently become apparent that both feedforward and feedback connections contribute substantially (Angelucci et al., 2017; Nurminen, Merlin, Bijanzadeh, Federer, & Angelucci, 2018). However, the modulating influence of the non-classical receptive field is also feature selective (Walker, Ohzawa, & Freeman, 2002; Webb, Dhruv, Solomon, Tailby, & Lennie, 2005) and spatially complex (Hallum & Movshon, 2014; H. Tanaka & Ohzawa, 2009). For these reasons, characterising the non-classical receptive fields to the same degree as the classical receptive field may prove difficult. We know most about the circuitry of surround responses in primary visual cortex, but there is evidence of facilitatory surrounds in many cells in DM (Lui et al., 2006) that could be investigated in more detail in future studies.

The reverse correlation analysis used here is restricted to describing the relevant stimuli in terms of patterns of luminance. This disregards sets of potentially relevant features that are not described well by these patterns. For example, selectivity to isoluminant patterns of contrast modulation has been reported for cells in V2 and is thought to play a role in the encoding of textures (El-Shamayleh

& Movshon, 2011; Jeremy Freeman, Corey M. Ziemba, David J. Heeger, Eero P. Simoncelli, & J. Anthony Movshon, 2013). Second order stimuli such as these might be effective in driving DM neurons. We did not see any evidence of selectivity for second-order features in our results but our techniques may not have been sufficiently sensitive; DM receives input from V2 as well as V1, which could include texture information. This could be useful if DM neurons were sensitive to orientated contours in second order patterns as well as in luminance, creating a response that's invariant to the order of the stimulus (G. Li et al., 2014).

The white noise patterns used in this experiment are an artificial stimulus, and are not representative of the visual information that an animal would be exposed to in nature. The statistics of natural images are very different to white noise. Notably, they are sparse (Graham & Field, 2007), dominated by orientated contours and a wide range of spatial frequencies aligned across multiple phases (Hansen & Hess, 2007) with a 1/f amplitude falloff (Simoncelli & Olshausen, 2001; Tolhurst, Tadmor, & Chao, 1992). By using a set of naturalistic images instead of white noise to fit models of neural responses, one samples from a smaller, more relevant space of visual stimuli (David, Vinje, & Gallant, 2004; Jon Touryan, Felsen, & Dan, 2005). This means less time wasted on irrelevant stimuli and better sampling in the relevant stimulus space. However, the stimulus set is necessarily biased, so alternative model fitting procedures must be used (Prenger, Wu, David, & Gallant, 2004).

6.2 Neural network models of selectivity

Neural network models have some advantages over the models used here. For example, layered models mimic some of the hierarchical processing seen in the cortex. Convolutional neural networks especially have nodes that resemble linear non-linear model neurons as they apply a filter, by

performing a convolution operation, to either the input space or the output of the layer beneath them and put the output through a non-linear function, usually a rectified linear function. Such networks have been used to determine feature selectivity by finding a stimulus that the neuron responds strongly to (Benda et al., 2007; DiMattina & Zhang, 2013). One of the problems with this is that the inner workings of neural networks are opaque. In order to recover these representations another operation needs to be performed. Zeiler and Fergus (2014) introduced a deconvolutional neural network that runs the network backwards, top down to generate the images in input space that are most likely to activate a node in the network. This allows for visualisation of the input parameter driving the node.

For cells in V4, Bashivan, Kar, and DiCarlo (2019), used artificial neural network to synthesize images that acted as 'super-stimuli' for recorded neurons in V4. The network used imported weights from a network based on AlexNet that performs object recognition (Krizhevsky et al., 2012). A layer of this network was then trained to model the responses of the V4 neurons to a large set of natural images. This is like creating a weighted sum of filters (Rust et al., 2005), except that the filters are features learned by an artificial network. Next, synthetic stimuli were created to maximise the response of the model neurons. When displayed to recorded V4 neurons, these 'super-stimuli' drove the neurons to respond far beyond the firing rates observed in response to naturalistic stimuli. However, these stimuli are difficult to interpret, bear little obvious relationship to each other, and do not comprise any readily describable category (Batista & Kording, 2019).

It should not be surprising that the optimal stimuli are combinations of features that do not occur in real images, because for a cell to have a useful representation of multiple features it should not approach its maximal firing rate for any one feature. If it did, the response would saturate and the contribution of other features would be underrepresented in the firing rate. Indeed, it's possible this

happened in the case for the 'super-stimuli' as although the neurons were driven well beyond the normal rates of firing, the example cell didn't reach the levels predicted by the model. Because the model was generated from the neuron responding in its operating region, well below saturation, and it uses a rectified linear nonlinearity, which never saturates, it likely overestimates the maximum firing rate. In a natural system with biological constraints on firing rates, representation of information is unlikely to be as specific as what the 'super-stimuli' imply because it would be impossible for the system to create representations for each and every stimulus likely to be encountered in one individual's life. Thus the optimal strategy from the point of view of natural selection is likely to be to rely on an adaptive system that results in well above background activity in many neurons, and to rely on population coding.

An interesting part of the modelling in Bashivan et al.'s work is that they were able to perform a gradient descent of the models response, iteratively changing the input image to maximise the response of the model cell. This method could potentially be used for a directions of invariance analysis in a similar way: starting at the maximal 'super-stimulus' find the direction of lowest gradient in stimulus space and find what transformations you of the image the model cell is invariant to. This may help pull apart the different dimensions of feature space are combined and therefore obscured in the 'super-stimuli.' The development of new tools to combine the precision of actively searching stimulus space with the power of artificial neural networks models is a promising new direction that could lead to a better understanding of feature selectivity in mid-tier areas in the future.

Introducing convolutional layers to neural networks helped drive the recent renaissance of deep neural networks (LeCun & Bengio, 1998), which was realised when advancements in graphic card hardware made it very efficient to perform convolutional operations (Cox & Dean, 2014). Most

modern neural networks now employ multiple convolutional layers. For a convolution, a node will sample from either the input space or a set of nodes that are adjacent to each other in the lower nodes. This means that even in these artificial neural networks how nodes are arranged in their layers matters to weighted computations done in the next layer. The arrangement of nodes in a layer, a map of nodes, in terms of their learned function or receptive field properties has been missing from description of layered neural networks models. In the future, models that include maps may see benefits from imitating biology in how information is organised.

Parallels between the hierarchies in visual cortex and in artificial neural networks has helped us learn more about the processing behind feature selectivity (Cox & Dean, 2014; Kriegeskorte, 2015). For example, deep networks with convolutional layers often learn kernels that resemble two-dimensional Gabor functions similar to what is seen in V1 (Güçlü & van Gerven, 2015). This is because, just like neurons in the cortex, nodes in a neural network form optimised representations as they learn weights that most efficiently encode incoming information. This is important in a biological setting as most of the incoming information is irrelevant or redundant and needs to be compressed by the visual system. The most striking example of image compression happens in the retina, where signals from 125 million photoreceptors are compressed to around 1 million retinal ganglion cells (Hubel, 1995). The problem faced by the early visual system is how to take data from a very high dimensional input space and project it onto a more efficient lower dimensional representation space that facilitates adaptive perceptual behaviour. We may be able to learn more about what information is compressible by studying artificial networks trained to perform different visual tasks.

6.3 Extending our Ca imaging results to single cells

The information represented in a single cell's firing rate is naturally ambiguous because the firing can be driven by multiple features. For example, for a simple cell with a Gabor-like receptive field the same non-maximal response could be caused by a stimulus's orientation, spatial frequency or position not quite matching the receptive field. In order to unambiguously represent the image on the retina in terms of features a network of cells is needed (Rust & Stocker, 2010; Zemel, Dayan, & Pouget, 1998). Each cell acts as part of this network, columns learning their representation through competitive plasticity with the cortex around it (Hensch & Stryker, 2004). How well the responses of single cells match the larger maps they are embedded in an important question for understanding how the cortex manages the trade-offs in representing multiple parameters, or whether these problems are solved locally within one area at all. One of the major benefits of the calcium imaging setup used here is that comparisons can be made between recordings made at very different spatial scales. Of particular interest is how single cell responses recorded using two-photon imaging compare to the maps created with wide field single photon imaging.

6.3.1 On and Off

In our results, we observed the phenomenon of regions selective to a contrast polarity. This has been reported previously for cats (Wang et al., 2015), ferrets (Smith et al., 2015) and there is even some evidence in these regions in macaques (Supplementary figures in Kremkow et al. 2016). The widely accepted explanation of this is that they are caused by clustering of Off and On centred thalamic afferents the arrangement of which contributes to the coarse progression of orientation tuning over the cortical surface (Miller, 1994). The importance of the arrangements of off and on afferents is evident in how the orientation tuning of columns can be robustly predicted by the thalamic inputs (Kremkow & Alonso, 2018). How important the arrangement of clusters of ON/OFF

afferents in V1 are to the orientation tuning map is not yet known although there has been some work in this area.

One key area of interest is in how the receptive fields of individual cells change approaching a pinwheel centre. In cats, electrophysiological recordings taken from a linear probe inserted tangentially through V1 revealed that OFF receptive fields moved very little in the visual field as electrodes sampled cortex over pinwheels, anchoring the retinotopy. ON receptive fields, however, moved rapidly (Kremkow et al., 2016). It would be interesting to see if this finding holds true in marmosets as well as cats and whether there are also pinwheels that are anchored by ON receptive fields in V1. The data shown for single cell two-photon imaging around pinwheel centre's here shows that the orientation tuning of individual cells aligns with the results from wide-field single photon imaging. Determining the receptive fields for dark and light stimuli of cells around the pinwheel centre and observing how they progress over cortex is a direction for future analysis.

6.3.2 Colour

One direction to further our current results is to look for single cells that are tuned for colour. The single photon wide field imaging results for colour show alignment between colour sensitive regions of cortex and CO blobs. However, these results are at odds with electrophysiological recordings in marmoset which suggest that cells tuned for colour (Martin, 2004) can be found dispersed across V1 although they do resemble results from electrophysiology in the macaque (Landisman & Ts'o, 2002a). To reconcile this, the locations of colour tuned cells should be mapped and compared to the larger wide-field maps as well as the histological patterns of cytochrome oxidase. High resolution two-photon imaging is the perfect tool for this, and some data has already been collected to this end.

Additionally, the results here only consider regions of colour sensitivity of any type and no attempt has been made to quantify the tuning for the combinations of cone activations. There is some evidence for hue tuning and maps in V1 that have been reported previously in macaque (Xiao, Casti, Xiao, & Kaplan, 2007). This has yet to be attempted in marmoset, however building on the wide field colour sensitivity maps reported here to a more nuanced hue map is a natural next step. Marmoset colour vision is particularly unusual because of the different genetic expression of cone opsins among individuals. How differently colour is represented for dichromatic and trichromatic marmosets could be an interesting area for inquiry especially if the differences in colour maps correspond to changes in maps for other parameters.

Ocular dominance columns in marmoset V1 have been the subject of some controversy as methods that reveal obvious columns in cat (Shatz & Stryker, 1978), macaque (Obermayer & Blasdel, 1993) and even another new world monkey (Kaskan, Lu, Dillenburger, Roe, & Kaas, 2007) have not been clear in marmoset. It seems as though although there is evidence that ocular dominance columns are present early in development (Spatz, 1989), there is a high degree of variability in adults (Roe et al., 2005). This lack of clarity is seen in our results, as although regions with ocular preference can be observed, they are not as clear as in other species.

6.3.3 Multiplexed maps in marmoset V1

Some of the relationships between the organisations of tuning maps in the marmoset are consistent with findings from other animals. For instance, the two dimensional autocorrelation of the orientation maps reveals hexagonal patterns that have been reported previously in cats, ferrets and

macaques (Jang & Paik, 2017; Nauhaus & Nielsen, 2014; Nauhaus et al., 2016; Paik & Ringach, 2011). This is based on the hypothesis that orientation tuning maps in V1 are created by a moiré interference pattern between the patterns of on and off centred afferents in the retinal ganglion mosaic (Paik & Ringach, 2011). The orthogonal relationship between the representations of orientation and spatial frequency are also consistent with what has been reported in primates (Nauhaus et al., 2016).

However, some of the relationships between tuning maps in V1 reported here are quite different from those of other animals. For instance we found that many of the relationships between maps were not orthogonal. One interesting finding was the parallel organisation of the maps for spatial frequency and contrast polarity. Parallel maps are an interesting finding that could be created by tiling of regions inside larger regions with concentric parallel contours. Our results suggest this might be happening for ON and OFF regions inside regions with a preference to low spatial frequency, but more work is needed to say if this is true.

How can so much variation across animals be possible if the features being represented are the same and we assume many of the underlying mechanisms are preserved? Some differences may be a consequence of the deemphasised or degenerate ocular dominance map in marmosets. This may alter constraints on the spatial frequency map, for instance, its periodicity may be free to change to cover contrast if it's not locked to match ocular dominance columns. If the relationships between maps are very different for each animal model it suggests that there may be very little generalisability to the results found here for marmosets and their implications for the organisation of feature representation in other animals and humans. However, although a tendency towards orthogonal representations is thought to maximise the trade-off between smoothness and coverage

(Swindale et al., 2000) as it allows coverage of combinations of a pair of parameters smoothly, there are other ways maps can be organised to achieve this.

One way for two maps to have large coverage without orthogonal representations is for them to have different periodicities in their repetition across the two dimensional surface. In ferret V1, maps for orientation and contrast polarity are organised in this way (Smith et al., 2015). This is different to what was found here for those parameters in marmoset, but this way of achieving coverage could account for the relationships between some of the maps reported here, many of which are not orthogonal. To extend the analysis presented here, coverage should be measured for more combinations of parameters. The periodicities of multiple maps should also be measured and compared.

6.4 Maps throughout visual system

Early areas encode fundamental features of the visual scene such as orientation, spatial frequency and contrast. These parameters are easily split into semi-independent dimensions that can be mapped systematically. Do systematic maps of more complex visual features make sense? Higher tier visual areas are selective to a much more complicated set of features, should we expect higher level features to be organised in the same way? It's reasonable to expect some kind of maps throughout the visual cortex seeing as maps are a result of columnar organisation and wiring principles applicable everywhere in cortex. Evidence of this can be seen in the retinotopic organisation of the extrastriate visual cortex. However, mixed selectivity is the rule near the top of the cortical hierarchy (Rigotti et al., 2013), and may be important for maximising the encoding power of neuronal populations (Fusi, Miller, & Rigotti, 2016). It is possible that extrastriate visual and multi-sensory areas may fall somewhere in between.

6.4.1 Map formation

Initial formation of retinotopic maps in extrastriate cortex are created through molecular signals that guide the projection of afferents connections (Espinosa & Stryker, 2012; Sperry, 1963). Retinal waves, spontaneous activity in retinal ganglion cells, are critical for the retinotopic organisation of the LGN and visual areas, even before eye opening. Retinal waves are also suggested to be the driving force behind the organisation of orientation in V1, which has been shown to form around the time of eye opening without relying on visual experience (cats: Crair et al., 1998; mouse: Hagihara et al., 2015, ferret: Tiriac et al., 2018).

The formation of maps in V1 has been modelled using algorithms that learn optimal representations. Optimisation of the elastic net algorithm can give similar structures to the maps we reported here (Goodhill, 2007). However, continuing to optimise the elastic net model predicts that in the orientation tuning maps will eventually converge to an optimal geometric solutions rather than maps similar to those in physiology (Keil & Wolf, 2011). This suggests that relative to an optimal model, V1 is suboptimal, possibly due to anatomical constraints. The segregation of V1 into blob and interblob compartments may be responsible for some of this suboptimality. For instance, it may constrain the map for spatial frequency tuning by biasing the blob regions to low spatial frequency, and the rest of the map and other related maps may need to accommodate.

On the other hand, blobs could be the result of optimising the initial maps. Cytochrome oxidase staining reveals areas of high metabolic activity. It could be that any region of cortex with broad spatial frequency selectivity could be driven to higher average levels of activity due to the natural scenes being dominated by low spatial frequencies (Ringach, 2007) , however blobs form without visual experience, even if the eyes are removed before birth (Dehay et al., 1996). At this point, the

timing of the co-organisation of colour sensitive afferents and low spatial frequency into blobs is unclear. The complex relationships between maps that may constrain each other are not accounted for in simple algorithmic models to maximise coverage. In the future, results of multiple interacting maps may inform new developmental models that can be used to answer these questions. Part of the reason for this complexity could be the functional compartmentalisation of cortex.

6.4.2 Functional compartmentalisation in cortex

There are consequences to the compartmentalisation of representations, if parameters are independently represented in separate regions of cortex their combinations cannot be systematically mapped in the same surface as in the models for multiplexed maps. The blobs and interblob regions in V1 are sensitive to different types of information and are likely playing distinct roles in terms of information processing. This is further evident in the connections between the blobs and interblobs in V1 to pale and thin stripes in V2 respectively (Livingstone & Hubel, 1987; Sincich & Horton, 2005b; Sincich, Jocson, & Horton, 2007). The representations are no longer multiplexed, independent representations for each parameter. I.e.: colour, texture, motion in visual field encoded in separate streams in separate regions of cortex.

Similar to the compartmentalisation of V1 into blob and interblobs regions, V2 is organised into stripes which are thought to be responsible for the processing of different types of visual information (Federer et al., 2009; McLoughlin & Schiessl, 2006; Roe & Ts'o, 1997). This has ramifications for maps of tuning, the most clear being for retinotopy, which has to reconcile the need to repeat the representations of the visual field for each stripe with the need to smoothly and continuously cover the entire visual field. In order to accommodate this the retinotopic map is disturbed, with receptive centres progressing and then 'jumping back' on the borders between stripes (Roe & Ts'o, 1995; Rosa et al., 1997; Shipp & Zeki, 1989) or forming large regions where there

is little change in the retinotopic gradient, alternating with regions of fast change (Rosa, Sousa, & Gattass, 1988). In other visual areas, further compartmentalisation within areas is less clear. There is some evidence that parallel processing streams such as the magnocellular and parvocellular pathways are still partially separate at a columnar level in extrastriate cortex (Tootell & Nasr, 2017).

Although the retinotopy in V2 is not smooth like in V1, local continuity is still maintained as adjacent receptive fields are always overlapping. This is possible because the receptive fields in V2 are larger than those in V1 (Roe & Ts'o, 1995). The size of anatomical compartments and receptive field size are closely linked. In V2, cortical magnification is such that one repeat of stripes has one representation of the visual field and in V1 blobs are approximately the size of a cortical point image (Rosa et al., 1988; Sincich & Horton, 2005a). In V2, larger receptive fields allows receptive fields to progress further without creating gaps in the representation.

6.4.3 Magnification and complexity through the visual system

The idea that larger receptive fields can allow for less strictly smooth maps is backed up by the finding that continuous retinotopy can learn a twisted retinotopy. DM's strange retinotopy can be interpreted as a natural consequence of competing rules enforcing continuous receptive fields within the area vs continuous receptive over the border between areas (Yu et al., 2019). This is possible because despite the rapid change in retinotopy between lower and upper representations of the visual field the retinotopy (again) is not discontinuous. The receptive fields in DM and much larger than those in V1 and V2 which allows for progressions over the visual field between adjacent neurons that are much larger without gaps in the retinotopy. This is just one way that larger receptive fields ease constraints on maps.

As information gets passed up the visual hierarchy features become more abstract and the representation becomes more invariant to the exact retinal image. This coincides with increases in receptive field size, point image size (Gattass et al., 2005), spread of intrinsic connections (Amir, Harel, & Malach, 1993; Fujita & Fujita, 1996) and pyramidal cell dendritic arbor (Elston & Rosa, 1998a, 1998b; Elston, Rosa, & Calford, 1996) moving up through the hierarchy. When each part of the visual field is covered by more cortical tissue, more combinations of features can be encoded for each point in the visual space (Swindale, 2000). Unlike V1 and early vision, which likely performs dimensionality reduction (Olshausen & Field, 1997; Pehlevan, Hu, & Chklovskii, 2015), the increase in cortical point image allows for dimensionality expansion. Dimensionality reduction is necessitated by the fact that V1 has the smallest number of features to work with in the visual cortex. Beyond this, the maximum number of features that can be simultaneously represented by the same set of neurons increases and the complexity of receptive fields can expand into an increasing number of abstract dimensions. Beyond DM it could be expected that feature topography is less strict. Once the feature set of DM is fully characterised, determining the lawfulness of maps of this selectivity could be useful for understanding how maps and point image are related throughout the cortical hierarchy.

6.5 The dorsomedial area

So far the characterisation of receptive fields in DM is limited compared to what has been done for cells in V1. I, along with others (Lui et al., 2006), have compiled evidence that DM creates a more complex representation than V1. DM neuron's sensitivity for additional dimensions is evident in the larger number of significant linear filters when compared to V1 and is consistent with preliminary results that they have more directions of invariance. This supports the idea that complex representations are built up iteratively throughout the visual system hierarchy.

Our results suggest that neurons in DM are sensitive to extended contours with invariance to the spatial position. The information about contours in visual scenes could be useful in many computations. For example, identifying landmarks and horizons or the outlines of objects that extend over multiple lengths of V1 or V2 receptive field diameters. The receptive field structure that we see in DM can give us some clues about what functional role DM is playing in the visual system. However, the context of the area's place in the visual system hierarchy should be taken into consideration.

The results fits in well with the dorsal stream hypothesis that DM could be playing a role in encoding the locations of objects. Long contours could define the outlines of objects and anchor them to a location in visual field. Additional information about the objects internal fill in terms of colour and texture could be processed in the ventral stream. However, the receptive fields are spatially invariant, not sensitive to the exact location or "where" the contours are in visual space. This information is available in the population responses of V1 since neurons there have the finest spatial resolution, already. The resolution of the information from V1 doesn't need refinement, instead representations are formed from the generalisation of many specific receptive fields.

It's not difficult to imagine how information about contours could also be useful for identifying objects, a task ascribed to the ventral, "what", stream in the two stream hypothesis. For example the extended contours could be used to generated selectivity to larger objects in higher areas the same way contours encoded in V1 are thought to be combined to form sensitivity to objects. At this point it's not possible to say what the output of DM could be used for as beyond contours we don't yet have a full understanding of what processes DM is doing.

Part of the difficulty in studying mid-tier areas in primate is that there is significant controversy around where the borders for areas lie and what should be considered an area. Depending on the methods for parcellation, the area here called DM can be considered parts of different areas (Lyon & Connolly, 2012) or a single area with atypical retinotopy (Rosa, Angelucci, Jeffs, & Pettigrew, 2013). DM's atypical retinotopy can be accounted for by learning rules and the full retinotopic map while maintaining continuous representation with the rostral border of V2. However, just having a plausible retinotopic map is not sufficient evidence for this being a meaningful area. Our results suggest that the behaviour is maintained in the upper and lower field representations in DM. DM was first designated an area due to its distinctive anatomical markers and connections. Our results for function match the anatomical markers of DM identified by its dense myelination. The fact that the function and anatomical designation match and that the retinotopic weirdness can be explained by learning methods suggest that it is indeed one functional area.

Our results suggest that we should continue to parcellate the brain based on the ground truth of anatomy. Although retinotopic rules for parcellation such as the field sign (M. I. Sereno et al., 1994) have been used to for areal designations in mouse visual cortex (Garrett, Nauhaus, Marshel, & Callaway, 2014), primates and human fMRI (M. I. Sereno et al., 1995; Tootell et al., 1997) visual cortex experiments, the existence of twisted maps suggests it probably shouldn't be relied upon.

6.6 Conclusions

Feature selectivity and its organisation into maps is a general property of the visual cortex. But processing information in this way is likely to be a common strategy throughout the entire cortex, much of which relies on similar circuits and architecture. Understanding how feature selectivity is created and organised could lead to a better understanding of processing in the brain, better models and better artificial neural networks. So far maps, and lateral connections, are missing from layered neural networks and computational models. As shown here, a combination of complementary approaches can yield insight into how individual neurons function, and how that function relates to the structure of the network in which each neuron is embedded.

8 Bibliography

Adams, D. L., & Horton, J. C. (2003). Capricious expression of cortical columns in the primate brain. *Nat. Neurosci., 6*(2), 113-114. doi:10.1038/nn1004

Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *J Opt Soc Am A, 2*(2), 284-299.

Akerboom, J., Chen, T.-W., Wardill, T. J., Tian, L., Marvin, J. S., Mutlu, S., . . . Looger, L. L. (2012). Optimization of a GCaMP calcium indicator for neural activity imaging. *J. Neurosci., 32*(40), 13819-13840. doi:10.1523/JNEUROSCI.2601-12.2012

Albright, T. D., & Desimone, R. (1987). Local precision of visuotopic organization in the middle temporal area (MT) of the macaque. *Exp. Brain Res., 65*(3), 582-592. doi:10.1007/bf00235981

Allman, J. M., & Kaas, J. H. (1975). The dorsomedial cortical visual area: a third tier area in the occipital lobe of the owl monkey (Aotus trivirgatus). *Brain Res, 100*(3), 473-487.

Alonso, J. M. (2016). The Geometry of Visual Cortical Maps. *Neuron, 91*(4), 716-718. doi:10.1016/j.neuron.2016.08.001

Amir, Y., Harel, M., & Malach, R. (1993). Cortical hierarchy reflected in the organization of intrinsic connections in macaque monkey visual cortex. *J. Comp. Neurol., 334*(1), 19-46. doi:10.1002/cne.903340103

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

Angelucci, A., Roe, A. W., & Sereno, M. I. (2015). Controversial issues in visual cortex mapping: Extrastriate cortex between areas V2 and MT in human and nonhuman primates. *Vis. Neurosci., 32*, E025. doi:10.1017/S0952523815000292

Angelucci, A., & Rosa, M. G. P. (2015). Resolving the organization of the third tier visual cortex in primates: a hypothesis-based approach. *Vis. Neurosci., 32*, E010. doi:10.1017/S0952523815000073

Atkinson, J. (1992). Early visual development: differential functioning of parvocellular and magnocellular pathways. *Eye (Lond), 6 (Pt 2)*, 129-135. doi:10.1038/eye.1992.28

Attneave, F. (1954). Some informational aspects of visual perception. *Psychol Rev, 61*(3), 183-193.

Bair, W., Cavanaugh, J. R., & Movshon, J. A. (2003). Time course and time-distance relationships for surround suppression in macaque V1 neurons. *J Neurosci, 23*(20), 7690-7701.

Baker, J. F., Petersen, S. E., Newsome, W. T., & Allman, J. M. (1981). Visual response properties of neurons in four extrastriate visual areas of the owl monkey (Aotus trivirgatus): a quantitative comparison of medial, dorsomedial, dorsolateral, and middle temporal areas. *J Neurophysiol, 45*(3), 397-416.

Bashivan, P., Kar, K., & DiCarlo, J. J. (2019). Neural population control via deep image synthesis. *Science, 364*(6439). doi:10.1126/science.aav9436

Batista, A. P., & Kording, K. P. (2019). A Deep Dive to Illuminate V4 Neurons. *Trends Neurosci, 42*(9), 563-564. doi:10.1016/j.tins.2019.07.001

Benda, J., Gollisch, T., Machens, C. K., & Herz, A. V. (2007). From response to stimulus: adaptive sampling in sensory physiology. *Curr. Opin. Neurobiol., 17*(4), 430-436. doi:10.1016/j.conb.2007.07.009

Berkes, P., & Wiskott, L. (2006). On the analysis and interpretation of inhomogeneous quadratic forms as receptive fields. *Neural Comput, 18*(8), 1868-1895. doi:10.1162/neco.2006.18.8.1868

Berkes, P., & Wiskott, L. (2007). Analysis and interpretation of quadratic models of receptive fields. *Nat Protoc, 2*(2), 400-407. doi:10.1038/nprot.2007.27

Blasdel, G. G. (1992). Orientation selectivity, preference, and continuity in monkey striate cortex. *J Neurosci, 12*(8), 3139-3161.

Blasdel, G. G., & Salama, G. (1986). Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature, 321*(6070), 579-585. doi:10.1038/321579a0

Bonhoeffer, T., & Grinvald, A. (Eds.). (1996). *Optical imaging based on intrinsic signals: the methodology.* San Diego, CA: Academic.

Bourne, J. A., & Rosa, M. G. (2003). Preparation for the in vivo recording of neuronal responses in the visual cortex of anaesthetised marmosets (Callithrix jacchus). *Brain Res Brain Res Protoc, 11*(3), 168-177.

Brenner, N., Bialek, W., & de Ruyter van Steveninck, R. (2000). Adaptive rescaling maximizes information transmission. *Neuron, 26*(3), 695-702.

Bruce, C., Desimone, R., & Gross, C. G. (1981). Visual properties of neurons in a polysensory area in superior temporal sulcus of the macaque. *J Neurophysiol, 46*(2), 369-384.

Buxhoeveden, D. P., & Casanova, M. F. (2002). The minicolumn hypothesis in neuroscience. *Brain, 125*(Pt 5), 935-951. doi:10.1093/brain/awf110

Buzas, P., Szmajda, B. A., Hashemi-Nezhad, M., Dreher, B., & Martin, P. R. (2008). Color signals in the primary visual cortex of marmosets. *Journal of Vision, 8*(10), 7-7. doi:10.1167/8.10.7

Cadieu, C., Kouh, M., Pasupathy, A., Connor, C. E., Riesenhuber, M., & Poggio, T. (2007). A model of V4 shape selectivity and invariance. *J Neurophysiol, 98*(3), 1733-1750. doi:10.1152/jn.01265.2006

Cai, T. T., Ren, Z., & Zhao, H. H. (2016). Estimating structured high-dimensional covariance and precision matrices: Optimal rates and adaptive estimation. *Electron. J. Statist., 10*(1), 1-59.

Callaway, E. M. (2005). Structure and function of parallel pathways in the primate early visual system. *J. Physiol., 566*(Pt 1), 13-19. doi:10.1113/jphysiol.2005.088047

Cang, J., Rentería, R. C., Kaneko, M., Liu, X., Copenhagen, D. R., & Stryker, M. P. (2005). Development of precise maps in visual cortex requires patterned spontaneous activity in the retina. *Neuron, 48*(5), 797-809. doi:10.1016/j.neuron.2005.09.015

Capuano, U., & McIlwain, J. T. (1981). Reciprocity of receptive field images and point images in the superior colliculus of the cat. *J Comp Neurol, 196*(1), 13-23. doi:10.1002/cne.901960103

Carrasco, M., Ling, S., Gobel, J., Fuller, S., & Read, S. (2004). Attention alters appearance in early vision: Contrast sensitivity, spatial resolution, and color saturation. *J. Vis., 4*(8), 67-67. doi:10.1167/4.8.67

Casagrande, V. A. (1994). A third parallel visual pathway to primate area V1. *Trends Neurosci, 17*(7), 305-310.

Casagrande, V. A., Yazar, F., Jones, K. D., & Ding, Y. (2007). The morphology of the koniocellular axon pathway in the macaque monkey. *Cereb. Cortex, 17*(10), 2334-2345. doi:10.1093/cercor/bhl142

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

Chan, T. L., Martin, P. R., Clunas, N., & Grunert, U. (2001). Bipolar cell diversity in the primate retina: morphologic and immunocytochemical analysis of a new world monkey, the marmoset Callithrix jacchus. *J Comp Neurol, 437*(2), 219-239.

Chance, F. S., Nelson, S. B., & Abbott, L. F. (1999). Complex cells as cortically amplified simple cells. *Nat. Neurosci., 2*(3), 277-282. doi:10.1038/6381

Chaplin, T. A., Yu, H. H., & Rosa, M. G. (2013). Representation of the visual field in the primary visual area of the marmoset monkey: magnification factors, point-image size, and proportionality to retinal ganglion cell density. *J Comp Neurol, 521*(5), 1001-1019. doi:10.1002/cne.23215

Chappert-Piquemal, C., Fonta, C., Malecaze, F., & Imbert, M. (2001). Ocular dominance columns in the adult New World Monkey Callithrix jacchus. *Vis. Neurosci., 18*(3), 407-412.

Chen, G., Lu, H. D., & Roe, A. W. (2008). A map for horizontal disparity in monkey V2. *Neuron, 58*(3), 442-450. doi:10.1016/j.neuron.2008.02.032

Chen, R. Y., Gittens, A., & Tropp, J. A. (2012). The Masked Sample Covariance Estimator: An Analysis via Matrix Concentration Inequalities. *Information and Inference, 1*, 2-20.

Chen, T. W., Wardill, T. J., Sun, Y., Pulver, S. R., Renninger, S. L., Baohan, A., . . . Kim, D. S. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature, 499*(7458), 295-300. doi:10.1038/nature12354

Chichilnisky, E. J. (2001). A simple white noise analysis of neuronal light responses. *Network, 12*(2), 199-213.

Condo, G. J., & Casagrande, V. A. (1990). Organization of cytochrome oxidase staining in the visual cortex of nocturnal primates (Galago crassicaudatus and Galago senegalensis): I. Adult patterns. *J Comp Neurol*, 293(4), 632-645. doi:10.1002/cne.902930408

Conway, B. R., Moeller, S., & Tsao, D. Y. (2007). Specialized color modules in macaque extrastriate cortex. *Neuron, 56*(3), 560-573. doi:10.1016/j.neuron.2007.10.008

Cox, D. D., & Dean, T. (2014). Neural networks and neuroscience-inspired computer vision. *Curr. Biol., 24*(18), R921-R929. doi:10.1016/j.cub.2014.08.026

Crair, M. C., Gillespie, D. C., & Stryker, M. P. (1998). The role of visual experience in the development of columns in cat visual cortex. *Science, 279*(5350), 566-570. doi:10.1126/science.279.5350.566

Daugman, J. G. (1985). Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *J Opt Soc Am A, 2*(7), 1160-1169.

David, S. V., Vinje, W. E., & Gallant, J. L. (2004). Natural stimulus statistics alter the receptive field structure of v1 neurons. *J. Neurosci., 24*(31), 6991-7006. doi:10.1523/JNEUROSCI.1422-04.2004

De Boer, R., & Kuyper, P. (1968). Triggered correlation. *IEEE Trans Biomed Eng, 15*(3), 169-179.

De Ruyter van Steveninck, R., & Bialek, W. (1988). Real-Time Performance of a Movement-Sensitive Neuron in the Blowfly Visual System: Coding and Information Transfer in Short Spike Sequences. *Proc R Soc Lond B, 234*, 379-414.

De Valois, R. L., Albrecht, D. G., & Thorell, L. G. (1978). *Cortical Cells: Bar and Edge Detectors, or Spatial Frequency Filters?* Paper presented at the Frontiers in Visual Science.

De Valois, R. L., Albrecht, D. G., & Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Res., 22*(5), 545-559. doi:10.1016/0042-6989(82)90113-4

Dean, A. F., & Tolhurst, D. J. (1983). On the distinctness of simple and complex cells in the visual cortex of the cat. *J Physiol, 344*, 305-325. doi:10.1113/jphysiol.1983.sp014941

DeAngelis, G. C., Ohzawa, I., & Freeman, R. D. (1993). Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. II. Linearity of temporal and spatial summation. *J Neurophysiol, 69*(4), 1118-1135. doi:10.1152/jn.1993.69.4.1118

Dehay, C., Giroud, P., Berland, M., Killackey, H. P., & Kennedy, H. (1996). Phenotypic characterisation of respecified visual cortex subsequent to prenatal enucleation in the monkey: development of acetylcholinesterase and cytochrome oxidase patterns. *J. Comp. Neurol., 376*(3), 386-402. doi:10.1002/(SICI)1096-9861(19961216)376:3<386::AID-CNE3>3.0.CO;2-Z

Denk, W., Strickler, J., & Webb, W. (1990). Two-photon laser scanning fluorescence microscopy. *Science, 248*(4951), 73-76. doi:10.1126/science.2321027

Desimone, R., Albright, T. D., Gross, C. G., & Bruce, C. (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. *J Neurosci, 4*(8), 2051-2062.

Desimone, R., & Schein, S. J. (1987). Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. *J Neurophysiol, 57*(3), 835-868. doi:10.1152/jn.1987.57.3.835

DiMattina, C., & Zhang, K. (2013). Adaptive stimulus optimization for sensory systems neuroscience. *Front. Neural Circuits, 7*, 101. doi:10.3389/fncir.2013.00101

Ding, Y., & Casagrande, V. A. (1997). The distribution and morphology of LGN K pathway axons within the layers and CO blobs of owl monkey V1. *Vis. Neurosci., 14*(4), 691-704.

Dumoulin, S. O., Harvey, B. M., Fracasso, A., Zuiderbaan, W., Luijten, P. R., Wandell, B. A., & Petridou, N. (2017). In vivo evidence of functional and anatomical stripe-based subdivisions in human V2 and V3. *Sci. Rep., 7*(1), 733. doi:10.1038/s41598-017-00634-6

Durbin, R., & Mitchison, G. (1990). A dimension reduction framework for understanding cortical maps. *Nature, 343*(6259), 644-647. doi:10.1038/343644a0

El-Shamayleh, Y., & Movshon, J. A. (2011). Neuronal responses to texture-defined form in macaque visual area V2. *J. Neurosci., 31*(23), 8543-8555. doi:10.1523/JNEUROSCI.5974-10.2011

Elston, G. N., & Rosa, M. G. (1998a). Complex dendritic fields of pyramidal cells in the frontal eye field of the macaque monkey: comparison with parietal areas 7a and LIP. *Neuroreport, 9*(1), 127-131. doi:10.1097/00001756-199801050-00025

Elston, G. N., & Rosa, M. G. (1998b). Morphological variation of layer III pyramidal neurones in the occipitotemporal pathway of the macaque monkey visual cortex. *Cereb Cortex, 8*(3), 278-294. doi:10.1093/cercor/8.3.278

Elston, G. N., & Rosa, M. G. (2000). Pyramidal cells, patches, and cortical columns: a comparative study of infragranular neurons in TEO, TE, and the superior temporal polysensory area of the macaque monkey. *J. Neurosci., 20*(24), RC117.

Elston, G. N., Rosa, M. G., & Calford, M. B. (1996). Comparison of dendritic fields of layer III pyramidal neurons in striate and extrastriate visual areas of the marmoset: a Lucifer yellow intracellular injection. *Cereb Cortex, 6*(6), 807-813. doi:10.1093/cercor/6.6.807

Espinosa, J. S., & Stryker, M. P. (2012). Development and plasticity of the primary visual cortex. *Neuron, 75*(2), 230-249. doi:10.1016/j.neuron.2012.06.009

Federer, F., Ichida, J. M., Jeffs, J., Schiessl, I., McLoughlin, N., & Angelucci, A. (2009). Four projection streams from primate V1 to the cytochrome oxidase stripes of V2. *J. Neurosci., 29*(49), 15455-15471. doi:10.1523/JNEUROSCI.1648-09.2009

Federer, F., Williams, D., Ichida, J. M., Merlin, S., & Angelucci, A. (2013). Two projection streams from macaque V1 to the pale cytochrome oxidase stripes of V2. *J. Neurosci., 33*(28), 11530-11539. doi:10.1523/JNEUROSCI.5053-12.2013

Felleman, D. J., Lim, H., Xiao, Y., Wang, Y., Eriksson, A., & Parajuli, A. (2015). The Representation of Orientation in Macaque V2: Four Stripes Not Three. *Cereb. Cortex, 25*(9), 2354-2369. doi:10.1093/cercor/bhu033

Felleman, D. J., & Van Essen, D. C. (1987). Receptive field properties of neurons in area V3 of macaque monkey extrastriate cortex. *J. Neurophysiol., 57*(4), 889-920. doi:10.1152/jn.1987.57.4.889

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex, 1*(1), 1-47.

Firth, S. I., Wang, C.-T., & Feller, M. B. (2005). Retinal waves: mechanisms and function in visual system development. *Cell Calcium, 37*(5), 425-432. doi:10.1016/j.ceca.2005.01.010

Freeman, J., Ziemba, C. M., Heeger, D. J., Simoncelli, E. P., & Movshon, J. A. (2013). A functional and perceptual signature of the second visual area in primates. *Nat. Neurosci., 16*(7), 974-981. doi:10.1038/nn.3402

Freeman, J., Ziemba, C. M., Heeger, D. J., Simoncelli, E. P., & Movshon, J. A. (2013). A functional and perceptual signature of the second visual area in primates. *Nat Neurosci, 16*(7), 974-981. doi:10.1038/nn.3402

Friedman, R. M., Chen, L. M., & Roe, A. W. (2004). Modality maps within primate somatosensory cortex. *Proc. Natl. Acad. Sci. U. S. A., 101*(34), 12724-12729. doi:10.1073/pnas.0404884101

Frisén, J., Yates, P. A., McLaughlin, T., Friedman, G. C., O'Leary, D. D., & Barbacid, M. (1998). Ephrin-A5 (AL-1/RAGS) is essential for proper retinal axon guidance and topographic mapping in the mammalian visual system. *Neuron, 20*(2), 235-243. doi:10.1016/s0896-6273(00)80452-3

Fritsches, K. A., & Rosa, M. G. (1996). Visuotopic organisation of striate cortex in the marmoset monkey (Callithrix jacchus). *J Comp Neurol, 372*(2), 264-282. doi:10.1002/(SICI)1096-9861(19960819)372:2<264::AID-CNE8>3.0.CO;2-1

Fujita, I., & Fujita, T. (1996). Intrinsic Connections in the macaque inferior temporal cortex. *J. Comp. Neurol., 368*(4), 467-486. doi:10.1002/(SICI)1096-9861(19960513)368:4<467::AID-CNE1>3.0.CO;2-2

Fusi, S., Miller, E. K., & Rigotti, M. (2016). Why neurons mix: high dimensionality for higher cognition. *Curr. Opin. Neurobiol., 37*, 66-74. doi:10.1016/j.conb.2016.01.010

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

Ganmor, E., Krumin, M., Rossi, L. F., Carandini, M., & Simoncelli, E. P. (2016). Direct Estimation of Firing Rates from Calcium Imaging Data. *arXiv:1601.00364 [q-bio.NC]*.

Garg, A. K., Li, P., Rashid, M. S., & Callaway, E. M. (2019). Color and orientation are jointly coded and spatially organized in primate primary visual cortex. *Science, 364*(6447), 1275-1279. doi:10.1126/science.aaw5868

Garrett, M. E., Nauhaus, I., Marshel, J. H., & Callaway, E. M. (2014). Topography and areal organization of mouse visual cortex. *J. Neurosci., 34*(37), 12587-12600. doi:10.1523/JNEUROSCI.1124-14.2014

Gattass, R., Gross, C. G., & Sandell, J. H. (1981). Visual topography of V2 in the macaque. *J Comp Neurol, 201*(4), 519-539. doi:10.1002/cne.902010405

Gattass, R., Nascimento-Silva, S., Soares, J. G., Lima, B., Jansen, A. K., Diogo, A. C., . . . Fiorani, M. (2005). Cortical visual areas in monkeys: location, topography, connections, columns, plasticity and cortical dynamics. *Philos Trans R Soc Lond B Biol Sci, 360*(1456), 709-731. doi:10.1098/rstb.2005.1629

Gegenfurtner, K. R., Kiper, D. C., & Levitt, J. B. (1997). Functional properties of neurons in macaque area V3. *J. Neurophysiol., 77*(4), 1906-1923. doi:10.1152/jn.1997.77.4.1906

Ghodrati, M., Morris, A. P., & Price, N. S. (2015). The (un)suitability of modern liquid crystal displays (LCDs) for vision research. *Front Psychol*, 6, 303. doi:10.3389/fpsyg.2015.00303

Ghosh, K. K., Goodchild, A. K., Sefton, A. E., & Martin, P. R. (1996). Morphology of retinal ganglion cells in a new world monkey, the marmoset Callithrix jacchus. *J Comp Neurol, 366*(1), 76-92. doi:10.1002/(SICI)1096-9861(19960226)366:1<76::AID-CNE6>3.0.CO;2-H

Gilbert, C. D., & Wiesel, T. N. (1983). Clustered intrinsic connections in cat visual cortex. *J. Neurosci., 3*(5), 1116-1133. doi:10.1523/JNEUROSCI.03-05-01116.1983

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends Neurosci, 15*(1), 20-25.

Goodchild, A. K., & Martin, P. R. (1998). The distribution of calcium-binding proteins in the lateral geniculate nucleus and visual cortex of a New World monkey, the marmoset, Callithrix jacchus. *Vis Neurosci, 15*(4), 625-642.

Goodhill, G. J. (2007). Contributions of theoretical modeling to the understanding of neural map development. *Neuron, 56*(2), 301-311. doi:10.1016/j.neuron.2007.09.027

Goodhill, G. J., Bates, K. R., & Montague, P. R. (1997). Influences on the global structure of cortical maps. *Proc Biol Sci, 264*(1382), 649-655. doi:10.1098/rspb.1997.0092

Graham, D. J., & Field, D. J. (2007). Statistical regularities of art images and natural scenes: spectra, sparseness and nonlinearities. *Spat. Vis., 21*(1-2), 149-164. doi:10.1163/156856807782753877

Grinvald, A., Lieke, E., Frostig, R. D., Gilbert, C. D., & Wiesel, T. N. (1986). Functional architecture of cortex revealed by optical imaging of intrinsic signals. *Nature, 324*(6095), 361-364. doi:10.1038/324361a0

Güçlü, U., & van Gerven, M. A. J. (2015). Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream. *J. Neurosci., 35*(27), 10005-10014. doi:10.1523/JNEUROSCI.5023-14.2015

Gur, M., Kagan, I., & Snodderly, D. M. (2005). Orientation and direction selectivity of neurons in V1 of alert monkeys: functional relationships and laminar distributions. *Cereb Cortex, 15*(8), 1207-1221. doi:10.1093/cercor/bhi003

Hadjidimitrakis, K., Bakola, S., Chaplin, T. A., Yu, H.-H., Alanazi, O., Chan, J. M., . . . Rosa, M. G. P. (2019). Topographic Organization of the 'Third-Tier' Dorsomedial Visual Cortex in the Macaque. *J. Neurosci., 39*(27), 5311-5325. doi:10.1523/JNEUROSCI.0085-19.2019

Hagihara, K. M., Murakami, T., Yoshida, T., Tagawa, Y., & Ohki, K. (2015). Neuronal activity is not required for the initial formation and maturation of visual selectivity. *Nat. Neurosci., 18*(12), 1780-1788. doi:10.1038/nn.4155

Hallum, L. E., & Movshon, J. A. (2014). Surround suppression supports second-order feature encoding by macaque V1 and V2 neurons. *Vision Res., 104*, 24-35. doi:10.1016/j.visres.2014.10.004

Hanazawa, A., & Komatsu, H. (2001). Influence of the direction of elemental luminance gradients on the responses of V4 cells to textured surfaces. *J Neurosci, 21*(12), 4490-4497.

Hansen, B. C., & Hess, R. F. (2007). Structural sparseness and spatial phase alignment in natural scenes. *J. Opt. Soc. Am. A Opt. Image Sci. Vis., 24*(7), 1873-1885. doi:10.1364/josaa.24.001873

Heeger, D. J. (1992). Half-squaring in responses of cat striate cells. *Vis Neurosci, 9*(5), 427-443.

Hegde, J., & Van Essen, D. C. (2000). Selectivity for complex shapes in primate visual area V2. *J Neurosci, 20*(5), RC61.

Heider, B., Nathanson, J. L., Isacoff, E. Y., Callaway, E. M., & Siegel, R. M. (2010). Two-photon imaging of calcium in virally transfected striate cortical neurons of behaving monkey. *PLoS One, 5*(11), e13829. doi:10.1371/journal.pone.0013829

Helmchen, F., & Denk, W. (2005). Deep tissue two-photon microscopy. *Nature Methods, 2*(12), 932-940. doi:10.1038/nmeth818

Hendry, S. H., & Yoshioka, T. (1994). A neurochemically distinct third channel in the macaque dorsal lateral geniculate nucleus. *Science, 264*(5158), 575-577. doi:10.1126/science.8160015

Hensch, T. K., Fagiolini, M., Mataga, N., Stryker, M. P., Baekkeskov, S., & Kash, S. F. (1998). Local GABA circuit control of experience-dependent plasticity in developing visual cortex. *Science, 282*(5393), 1504-1508. doi:10.1126/science.282.5393.1504

Hensch, T. K., & Stryker, M. P. (2004). Columnar architecture sculpted by GABA circuits in developing cat visual cortex. *Science, 303*(5664), 1678-1681. doi:10.1126/science.1091031

Hinkle, D. A., & Connor, C. E. (2001). Disparity tuning in macaque area V4. *Neuroreport, 12*(2), 365-369. doi:10.1097/00001756-200102120-00036

Hinkle, D. A., & Connor, C. E. (2002). Three-dimensional orientation tuning in macaque area V4. *Nat Neurosci, 5*(7), 665-670. doi:10.1038/nn875

Homma, R., Baker, B. J., Jin, L., Garaschuk, O., Konnerth, A., Cohen, L. B., & Zecevic, D. (2009). Wide-field and two-photon imaging of brain activity with voltage- and calcium-sensitive dyes. *Philos. Trans. R. Soc. Lond. B Biol. Sci., 364*(1529), 2453-2467. doi:10.1098/rstb.2009.0084

Hubel, D. H. (1995). Eye, brain, and vision. *Scientific American library series, No. 22., 242*.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *J Physiol, 195*(1), 215-243.

Hubel, D. H., & Wiesel, T. N. (1974). Sequence regularity and geometry of orientation columns in the monkey striate cortex. *The Journal of Comparative Neurology, 158*(3), 267-293. doi:10.1002/cne.901580304

Hubel, D. H., & Wiesel, T. N. (1977). Ferrier lecture. Functional architecture of macaque monkey visual cortex. *Proc R Soc Lond B Biol Sci, 198*(1130), 1-59. doi:10.1098/rspb.1977.0085

Hubener, M., Shoham, D., Grinvald, A., & Bonhoeffer, T. (1997). Spatial relationships among three columnar systems in cat area 17. *J Neurosci, 17*(23), 9270-9284.

Jacobs, G. H., Neitz, M., Deegan, J. F., & Neitz, J. (1996). Trichromatic colour vision in New World monkeys. *Nature, 382*(6587), 156-158. doi:10.1038/382156a0

Jang, J., & Paik, S.-B. (2017). Interlayer Repulsion of Retinal Ganglion Cell Mosaics Regulates Spatial Organization of Functional Maps in the Visual Cortex. *J. Neurosci., 37*(50), 12141-12152. doi:10.1523/JNEUROSCI.1873-17.2017

Jeffs, J., Federer, F., & Angelucci, A. (2015). Corticocortical connection patterns reveal two distinct visual cortical areas bordering dorsal V2 in marmoset monkey. *Vis Neurosci, 32*, E012. doi:10.1017/S0952523815000097

Jones, H. E., Grieve, K. L., Wang, W., & Sillito, A. M. (2001). Surround suppression in primate V1. *J Neurophysiol, 86*(4), 2011-2028.

Jones, J. P., & Palmer, L. A. (1987). The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *J Neurophysiol, 58*(6), 1187-1211.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J Neurosci, 17*(11), 4302-4311.

Kaskan, P. M., Dillenburger, B. C., Lu, H. D., Roe, A. W., & Kaas, J. H. (2010). Orientation and Direction-of-Motion Response in the Middle Temporal Visual Area (MT) of New World Owl Monkeys as Revealed by Intrinsic-Signal Optical Imaging. *Front. Neuroanat., 4*, 23. doi:10.3389/fnana.2010.00023

Kaskan, P. M., Lu, H. D., Dillenburger, B. C., Roe, A. W., & Kaas, J. H. (2007). Intrinsic-signal optical imaging reveals cryptic ocular dominance columns in primary visual cortex of New World owl monkeys. *Front. Neurosci., 1*(1), 67-75. doi:10.3389/neuro.01.1.1.005.2007

Kawamura, S., Hirai, M., Takenaka, O., Radlwimmer, F. B., & Yokoyama, S. (2001). Genomic and spectral analyses of long to middle wavelength-sensitive visual pigments of common marmoset (Callithrix jacchus). *Gene, 269*(1-2), 45-51. doi:10.1016/s0378-1119(01)00454-1

Keil, W., & Wolf, F. (2011). Coverage, continuity, and visual cortical architecture. *Neural Syst Circuits, 1*, 17. doi:10.1186/2042-1001-1-17

Knierim, J. J., & van Essen, D. C. (1992). Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J Neurophysiol, 67*(4), 961-980.

Kobatake, E., & Tanaka, K. (1994). Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex. *J Neurophysiol, 71*(3), 856-867. doi:10.1152/jn.1994.71.3.856

Kotake, Y., Morimoto, H., Okazaki, Y., Fujita, I., & Tamura, H. (2009). Organization of color-selective neurons in macaque visual area V4. *J Neurophysiol, 102*(1), 15-27. doi:10.1152/jn.90624.2008

Kremkow, J., & Alonso, J. M. (2018). Thalamocortical Circuits and Functional Architecture. *Annu Rev Vis Sci, 4*, 263-285. doi:10.1146/annurev-vision-091517-034122

Kremkow, J., Jin, J., Wang, Y., & Alonso, J. M. (2016). Principles underlying sensory map topography in primary visual cortex. *Nature, 533*(7601), 52-57. doi:10.1038/nature17936

Kriegeskorte, N. (2015). Deep Neural Networks: A New Framework for Modeling Biological Vision and Brain Information Processing. *Annu Rev Vis Sci, 1*, 417-446. doi:10.1146/annurev-vision-082114-035447

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems 25* (pp. 1097-1105): Curran Associates, Inc.

Krubitzer, L. A., & Kaas, J. H. (1993). The dorsomedial visual area of owl monkeys: connections, myeloarchitecture, and homologies in other primates. *J Comp Neurol, 334*(4), 497-528. doi:10.1002/cne.903340402

Landisman, C. E., & Ts'o, D. Y. (2002a). Color processing in macaque striate cortex: electrophysiological properties. *J. Neurophysiol., 87*(6), 3138-3151. doi:10.1152/jn.00957.1999

Landisman, C. E., & Ts'o, D. Y. (2002b). Color processing in macaque striate cortex: relationships to ocular dominance, cytochrome oxidase, and orientation. *J. Neurophysiol., 87*(6), 3126-3137. doi:10.1152/jn.2002.87.6.3126

Laskowski, M. B., & Sanes, J. R. (1987). Topographic mapping of motor pools onto skeletal muscles. *J. Neurosci., 7*(1), 252-260.
Le, Q. V., Ngiam, J., Chen, Z., Chia, D., Koh, P. K., & Ng, A. Y. (2010). Tiled convolutional neural networks. *Advances in Neural Information Processing Systems (NIPS)*(23), 1279-1287.
LeCun, Y., & Bengio, Y. (1998). The Handbook of Brain Theory and Neural Networks. In M. A. Arbib (Ed.), (pp. 255-258). Cambridge, MA, USA: MIT Press.
Li, G., Yao, Z., Wang, Z., Yuan, N., Talebi, V., Tan, J., . . . Baker, C. L., Jr. (2014). Form-cue invariant second-order neuronal responses to contrast modulation in primate area V2. *J. Neurosci., 34*(36), 12081-12092. doi:10.1523/JNEUROSCI.0211-14.2014
Li, M., Liu, F., Jiang, H., Lee, T. S., & Tang, S. (2017). Long-Term Two-Photon Imaging in Awake Macaque Monkey. *Neuron, 93*(5), 1049-1057 e1043. doi:10.1016/j.neuron.2017.01.027
Li, P., Zhu, S., Chen, M., Han, C., Xu, H., Hu, J., . . . Lu, H. D. (2013). A motion direction preference map in monkey V4. *Neuron, 78*(2), 376-388. doi:10.1016/j.neuron.2013.02.024
Li, X., Zhu, Q., Janssens, T., Arsenault, J. T., & Vanduffel, W. (2019). In Vivo Identification of Thick, Thin, and Pale Stripes of Macaque Area V2 Using Submillimeter Resolution (f)MRI at 3 T. *Cereb. Cortex, 29*(2), 544-560. doi:10.1093/cercor/bhx337
Livingstone, M. S., & Hubel, D. H. (1983). Specificity of cortico-cortical connections in monkey visual system. *Nature, 304*(5926), 531-534. doi:10.1038/304531a0
Livingstone, M. S., & Hubel, D. H. (1984a). Anatomy and physiology of a color system in the primate visual cortex. *J Neurosci, 4*(1), 309-356.
Livingstone, M. S., & Hubel, D. H. (1984b). Specificity of intrinsic connections in primate primary visual cortex. *J. Neurosci., 4*(11), 2830-2835. doi:10.1523/JNEUROSCI.04-11-02830.1984
Livingstone, M. S., & Hubel, D. H. (1987). Connections between layer 4B of area 17 and the thick cytochrome oxidase stripes of area 18 in the squirrel monkey. *J Neurosci, 7*(11), 3371-3377.
Logothetis, N. K., & Sheinberg, D. L. (1996). Visual object recognition. *Annu Rev Neurosci, 19*, 577-621. doi:10.1146/annurev.ne.19.030196.003045
Lu, H. D., & Roe, A. W. (2008). Functional organization of color domains in V1 and V2 of macaque monkey revealed by optical imaging. *Cereb Cortex, 18*(3), 516-533. doi:10.1093/cercor/bhm081
Liu, J. K., Schreyer, H. M., Onken, A., Rozenblit, F., Khani, M. H., Krishnamoorthy, V., . . . Gollisch, T. (2017). Inference of neuronal functional circuitry with spike-triggered non-negative matrix factorization. *Nat Commun,* 8(1), 149. doi:10.1038/s41467-017-00156-9
Lui, L. L., Bourne, J. A., & Rosa, M. G. (2006). Functional response properties of neurons in the dorsomedial visual area of New World monkeys (Callithrix jacchus). *Cereb Cortex, 16*(2), 162-177. doi:10.1093/cercor/bhi094
Lyon, D. C., & Connolly, J. D. (2012). The case for primate V3. *Proc Biol Sci, 279*(1729), 625-633. doi:10.1098/rspb.2011.2048
Madisen, L., Garner, A. R., Shimaoka, D., Chuong, A. S., Klapoetke, N. C., Li, L., . . . Zeng, H. (2015). Transgenic mice for intersectional targeting of neural sensors and effectors with high specificity and performance. *Neuron, 85*(5), 942-958. doi:10.1016/j.neuron.2015.02.022
Martin, K. A. (1988). The Wellcome Prize lecture. From single cells to simple circuits in the cerebral cortex. *Q J Exp Physiol, 73*(5), 637-702.
Martin, P. R. (2004). Colour through the thalamus. *Clin Exp Optom, 87*(4-5), 249-257.
Maruyama, R., Maeda, K., Moroda, H., Kato, I., Inoue, M., Miyakawa, H., & Aonishi, T. (2014). Detecting cells using non-negative matrix factorization on calcium imaging data. *Neural Networks, 55*, 11-19. doi:10.1016/j.neunet.2014.03.007
Maunsell, J. H., & van Essen, D. C. (1983). The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *J. Neurosci., 3*(12), 2563-2586.
Mazurek, M., Kager, M., & Van Hooser, S. D. (2014). Robust quantification of orientation selectivity and direction selectivity. *Frontiers in Neural Circuits, 8*. doi:10.3389/fncir.2014.00092

McLaughlin, T., & O'Leary, D. D. M. (2005). Molecular gradients and development of retinotopic maps. *Annu. Rev. Neurosci., 28*, 327-355. doi:10.1146/annurev.neuro.28.061604.135714
McLoughlin, N., & Schiessl, I. (2006). Orientation selectivity in the common marmoset (Callithrix jacchus): the periodicity of orientation columns in V1 and V2. *NeuroImage, 31*(1), 76-85. doi:10.1016/j.neuroimage.2005.12.054
Mechler, F., & Ringach, D. L. (2002). On the classification of simple and complex cells. *Vision Res., 42*(8), 1017-1033. doi:10.1016/s0042-6989(02)00025-1
Miller, K. D. (1994). A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between ON- and OFF-center inputs. *J. Neurosci., 14*(1), 409-441.
Moeller, S., Crapse, T., Chang, L., & Tsao, D. Y. (2017). The effect of face patch microstimulation on perception of faces and objects. *Nat. Neurosci., 20*(5), 743-752. doi:10.1038/nn.4527
Mohammed, A. I., Gritton, H. J., Tseng, H.-A., Bucklin, M. E., Yao, Z., & Han, X. (2016). An integrative approach for analyzing hundreds of neurons in task performing mice using wide-field calcium imaging. *Sci. Rep., 6*, 20986. doi:10.1038/srep20986
Mountcastle, V. B. (1997). The columnar organization of the neocortex. *Brain, 120 (Pt 4)*, 701-722. doi:10.1093/brain/120.4.701
Movshon, J. A., Thompson, I. D., & Tolhurst, D. J. (1978a). Receptive field organization of complex cells in the cat's striate cortex. *J Physiol, 283*, 79-99.
Movshon, J. A., Thompson, I. D., & Tolhurst, D. J. (1978b). Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat's visual cortex. *J Physiol, 283*, 101-120. doi:10.1113/jphysiol.1978.sp012490
Nakai, J., Ohkura, M., & Imoto, K. (2001). A high signal-to-noise Ca2+ probe composed of a single green fluorescent protein. *Nature Biotechnology, 19*(2), 137-141. doi:10.1038/84397
Nakamoto, M., Cheng, H. J., Friedman, G. C., McLaughlin, T., Hansen, M. J., Yoon, C. H., . . . Flanagan, J. G. (1996). Topographically specific effects of ELF-1 on retinal axon guidance in vitro and retinal axon mapping in vivo. *Cell, 86*(5), 755-766. doi:10.1016/s0092-8674(00)80150-6
Nasr, S., Polimeni, J. R., & Tootell, R. B. H. (2016). Interdigitated Color- and Disparity-Selective Columns within Human Visual Cortical Areas V2 and V3. *J. Neurosci., 36*(6), 1841-1857. doi:10.1523/JNEUROSCI.3518-15.2016
Nassi, J. J., & Callaway, E. M. (2009). Parallel processing strategies of the primate visual system. *Nat. Rev. Neurosci., 10*(5), 360-372. doi:10.1038/nrn2619
Nauhaus, I., & Nielsen, K. J. (2014). Building maps from maps in primary visual cortex. *Curr. Opin. Neurobiol., 24*(1), 1-6. doi:10.1016/j.conb.2013.08.007
Nauhaus, I., Nielsen, K. J., & Callaway, E. M. (2016). Efficient Receptive Field Tiling in Primate V1. *Neuron, 91*(4), 893-904. doi:10.1016/j.neuron.2016.07.015
Nauhaus, I., Nielsen, K. J., Disney, A. A., & Callaway, E. M. (2012). Orthogonal micro-organization of orientation and spatial frequency in primate primary visual cortex. *Nature Neuroscience, 15*(12), 1683-1690. doi:10.1038/nn.3255
Newsome, W. T., & Paré, E. B. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J. Neurosci., 8*(6), 2201-2211.
Nurminen, L., Merlin, S., Bijanzadeh, M., Federer, F., & Angelucci, A. (2018). Top-down feedback controls spatial summation and response amplitude in primate visual cortex. *Nat. Commun., 9*(1), 2281. doi:10.1038/s41467-018-04500-5
Obermayer, K., & Blasdel, G. G. (1993). Geometry of orientation and ocular dominance columns in monkey striate cortex. *J Neurosci, 13*(10), 4114-4129.
Okazawa, G., Tajima, S., & Komatsu, H. (2015). Image statistics underlying natural texture selectivity of neurons in macaque V4. *Proc. Natl. Acad. Sci. U. S. A., 112*(4), E351-360. doi:10.1073/pnas.1415146112
Olshausen, B. A., & Field, D. J. (1997). Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vision Res., 37*(23), 3311-3325. doi:10.1016/s0042-6989(97)00169-7

Paik, S. B., & Ringach, D. L. (2011). Retinal origin of orientation maps in visual cortex. *Nat Neurosci, 14*(7), 919-925. doi:10.1038/nn.2824
Paninski, L. (2004). Maximum likelihood estimation of cascade point-process neural encoding models. *Network, 15*(4), 243-262.
Park, I., & Pillow, J. W. (2011). Bayesian spike-triggered covariance. *Advances in Neural Information Processing Systems (NIPS), 24*(eds. Shawe-Taylor, J.; Zemel, R.; Bartlett, P.; Pereira, F. & Weinberger, K.,), 1692-1700.
Park, J. E., Zhang, X. F., Choi, S.-H., Okahara, J., Sasaki, E., & Silva, A. C. (2016). Generation of transgenic marmosets expressing genetically encoded calcium indicators. *Scientific Reports, 6*(1), 34931. doi:10.1038/srep34931
Pasupathy, A., & Connor, C. E. (1999). Responses to contour features in macaque area V4. *J. Neurophysiol., 82*(5), 2490-2502. doi:10.1152/jn.1999.82.5.2490
Pehlevan, C., Hu, T., & Chklovskii, D. B. (2015). A Hebbian/Anti-Hebbian Neural Network for Linear Subspace Learning: A Derivation from Multidimensional Scaling of Streaming Data. *arXiv [q-bio.NC]*.
Pillow, J. W., & Simoncelli, E. P. (2006). Dimensionality reduction in neural models: an information-theoretic generalization of spike-triggered average and covariance analysis. *J Vis, 6*(4), 414-428. doi:10.1167/6.4.9
Pitzalis, S., Sereno, M. I., Committeri, G., Fattori, P., Galati, G., Patria, F., & Galletti, C. (2010). Human v6: the medial motion area. *Cereb Cortex, 20*(2), 411-424. doi:10.1093/cercor/bhp112
Pnevmatikakis, Eftychios A., Soudry, D., Gao, Y., Machado, T. A., Merel, J., Pfau, D., . . . Paninski, L. (2016). Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. *Neuron, 89*(2), 285-299. doi:10.1016/j.neuron.2015.11.037
Ponce, C. R., Xiao, W., Schade, P. F., Hartmann, T. S., Kreiman, G., & Livingstone, M. S. (2019). Evolving Images for Visual Neurons Using a Deep Generative Network Reveals Coding Principles and Neuronal Preferences. *Cell, 177*(4), 999-1009.e1010. doi:10.1016/j.cell.2019.04.005
Prenger, R., Wu, M. C. K., David, S. V., & Gallant, J. L. (2004). Nonlinear V1 responses to natural scenes revealed by neural network analysis. *Neural Netw., 17*(5-6), 663-679. doi:10.1016/j.neunet.2004.03.008
Priebe, N. J., Mechler, F., Carandini, M., & Ferster, D. (2004). The contribution of spike threshold to the dichotomy of cortical simple and complex cells. *Nat. Neurosci., 7*(10), 1113-1122. doi:10.1038/nn1310
Purves, D., & LaMantia, A. (1993). Development of blobs in the visual cortex of macaques. *J. Comp. Neurol., 334*(2), 169-175. doi:10.1002/cne.903340202
Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci., 2*(1), 79-87. doi:10.1038/4580
Reale, R. A., & Imig, T. J. (1980). Tonotopic organization in auditory cortex of the cat. *J. Comp. Neurol., 192*(2), 265-291. doi:10.1002/cne.901920207
Reid, R. C., & Alonso, J. M. (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature, 378*(6554), 281-284. doi:10.1038/378281a0
Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nat Neurosci, 2*(11), 1019-1025. doi:10.1038/14819
Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature, 497*(7451), 585-590. doi:10.1038/nature12160
Ringach, D. L. (2004). Haphazard wiring of simple receptive fields and orientation columns in visual cortex. *J Neurophysiol, 92*(1), 468-476. doi:10.1152/jn.01202.2003
Ringach, D. L. (2007). On the origin of the functional architecture of the cortex. *PLoS One, 2*(2), e251. doi:10.1371/journal.pone.0000251

Ringach, D. L., Shapley, R. M., & Hawken, M. J. (2002). Orientation selectivity in macaque V1: diversity and laminar dependence. *J. Neurosci., 22*(13), 5639-5651. doi:20026567

Roe, A. W., Fritsches, K., & Pettigrew, J. D. (2005). Optical imaging of functional organization of V1 and V2 in marmoset visual cortex. *The Anatomical Record Part A: Discoveries in Molecular, Cellular, and Evolutionary Biology, 287A*(2), 1213-1225. doi:10.1002/ar.a.20248

Roe, A. W., & Ts'o, D. Y. (1995). Visual topography in primate V2: multiple representation across functional stripes. *J. Neurosci., 15*(5 Pt 2), 3689-3715. doi:10.1523/JNEUROSCI.15-05-03689.1995

Roe, A. W., & Ts'o, D. Y. (1997). The Functional Architecture of Area V2 in the Macaque Monkey. In K. S. Rockland, J. H. Kaas, & A. Peters (Eds.), *Extrastriate Cortex in Primates* (pp. 295-333). Boston, MA: Springer US.

Rosa, M. G. (2002). Visual maps in the adult primate cerebral cortex: some implications for brain development and evolution. *Braz. J. Med. Biol. Res., 35*(12), 1485-1498.

Rosa, M. G., Angelucci, A., Jeffs, J., & Pettigrew, J. D. (2013). The case for a dorsomedial area in the primate 'third-tier' visual cortex. *Proc. Biol. Sci., 280*(1750), 20121372; discussion 20121994. doi:10.1098/rspb.2012.1372

Rosa, M. G., Fritsches, K. A., & Elston, G. N. (1997). The second visual area in the marmoset monkey: visuotopic organisation, magnification factors, architectonical boundaries, and modularity. *J Comp Neurol, 387*(4), 547-567.

Rosa, M. G., Palmer, S. M., Gamberini, M., Burman, K. J., Yu, H. H., Reser, D. H., . . . Galletti, C. (2009). Connections of the dorsomedial visual area: pathways for early integration of dorsal and ventral streams in extrastriate cortex. *J Neurosci, 29*(14), 4548-4563. doi:10.1523/JNEUROSCI.0529-09.2009

Rosa, M. G., & Schmid, L. M. (1995). Visual areas in the dorsal and medial extrastriate cortices of the marmoset. *J. Comp. Neurol., 359*(2), 272-299. doi:10.1002/cne.903590207

Rosa, M. G., Sousa, A. P., & Gattass, R. (1988). Representation of the visual field in the second visual area in the Cebus monkey. *J. Comp. Neurol., 275*(3), 326-345. doi:10.1002/cne.902750303

Rosa, M. G., & Tweedale, R. (2001). The dorsomedial visual areas in New World and Old World monkeys: homology and function. *Eur J Neurosci, 13*(3), 421-427.

Rust, N. C., & Dicarlo, J. J. (2010). Selectivity and tolerance ("invariance") both increase as visual information propagates from cortical area V4 to IT. *J Neurosci, 30*(39), 12978-12995. doi:10.1523/JNEUROSCI.0179-10.2010

Rust, N. C., Schwartz, O., Movshon, J. A., & Simoncelli, E. P. (2005). Spatiotemporal elements of macaque v1 receptive fields. *Neuron, 46*(6), 945-956. doi:10.1016/j.neuron.2005.05.021

Rust, N. C., & Stocker, A. A. (2010). Ambiguity and invariance: two fundamental challenges for visual processing. *Curr. Opin. Neurobiol., 20*(3), 382-388. doi:10.1016/j.conb.2010.04.013

Sadakane, O., Masamizu, Y., Watakabe, A., Terada, S., Ohtsuka, M., Takaji, M., . . . Yamamori, T. (2015). Long-Term Two-Photon Calcium Imaging of Neuronal Populations with Subcellular Resolution in Adult Non-human Primates. *Cell Reports, 13*(9), 1989-1999. doi:10.1016/j.celrep.2015.10.050

Schiessl, I., & McLoughlin, N. (2003). Optical imaging of the retinotopic organization of V1 in the common marmoset. *NeuroImage, 20*(3), 1857-1864. doi:10.1016/j.neuroimage.2003.07.023

Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., . . . Cardona, A. (2012). Fiji: an open-source platform for biological-image analysis. *Nat Methods, 9*(7), 676-682. doi:10.1038/nmeth.2019

Schneider, C. A., Rasband, W. S., & Eliceiri, K. W. (2012). NIH Image to ImageJ: 25 years of image analysis. *Nat Methods, 9*(7), 671-675.

Schwartz, O., Chichilnisky, E. J., & Simoncelli, E. P. (2001). *Characterizing Neural Gain Control using Spike-triggered Covariance.* Paper presented at the Advances in Neural Information Processing Systems 14 (NIPS 2001).

Schwartz, O., Pillow, J. W., Rust, N. C., & Simoncelli, E. P. (2006). Spike-triggered neural characterization. *J Vis, 6*(4), 484-507. doi:10.1167/6.4.13

Sengpiel, F., Troilo, D., Kind, P. C., Graham, B., & Blakemore, C. (1996). Functional architecture of area 17 in normal and monocularly deprived marmosets (Callithrix jacchus). *Vis. Neurosci., 13*(1), 145-160. doi:10.1017/s0952523800007197

Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., . . . Tootell, R. B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science, 268*(5212), 889-893.

Sereno, M. I., McDonald, C. T., & Allman, J. M. (1994). Analysis of retinotopic maps in extrastriate cortex. *Cereb. Cortex, 4*(6), 601-620. doi:10.1093/cercor/4.6.601

Sereno, M. I., McDonald, C. T., & Allman, J. M. (2015). Retinotopic organization of extrastriate cortex in the owl monkey—dorsal and lateral areas. *Vis. Neurosci., 32*. doi:10.1017/S0952523815000206

Sergent, J., Ohta, S., & MacDonald, B. (1992). Functional neuroanatomy of face and object processing. A positron emission tomography study. *Brain, 115 Pt 1*, 15-36. doi:10.1093/brain/115.1.15

Series, P., Lorenceau, J., & Fregnac, Y. (2003). The "silent" surround of V1 receptive fields: theory and experiments. *J Physiol Paris, 97*(4-6), 453-474. doi:10.1016/j.jphysparis.2004.01.023

Serre, T., Wolf, L., Bileschi, S., Riesenhuber, M., & Poggio, T. (2007). Robust object recognition with cortex-like mechanisms. *IEEE Trans Pattern Anal Mach Intell, 29*(3), 411-426. doi:10.1109/TPAMI.2007.56

Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal, 27*, 379–423 & 623–656.

Sharpee, T., Rust, N. C., & Bialek, W. (2004). Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput, 16*(2), 223-250. doi:10.1162/089976604322742010

Shatz, C. J., & Stryker, M. P. (1978). Ocular dominance in layer IV of the cat's visual cortex and the effects of monocular deprivation. *J. Physiol., 281*, 267-283. doi:10.1113/jphysiol.1978.sp012421

Shimomura, O., Johnson, F. H., & Saiga, Y. (1962). Extraction, Purification and Properties of Aequorin, a Bioluminescent Protein from the Luminous Hydromedusan,Aequorea. *Journal of Cellular and Comparative Physiology, 59*(3), 223-239. doi:10.1002/jcp.1030590302

Shipp, S., & Zeki, S. (1985). Segregation of pathways leading from area V2 to areas V4 and V5 of macaque monkey visual cortex. *Nature, 315*(6017), 322-325. doi:10.1038/315322a0

Shipp, S., & Zeki, S. (1989). The Organization of Connections between Areas V5 and V1 in Macaque Monkey Visual Cortex. *Eur. J. Neurosci., 1*(4), 309-332. doi:10.1111/j.1460-9568.1989.tb00798.x

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., . . . Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature, 529*(7587), 484-489. doi:10.1038/nature16961

Silverman, M. S. & Tootell, R. B. (1987). Modified technique for cytochrome oxidase histochemistry: increased staining intensity and compatibility with 2-deoxyglucose autoradiography. J *Neurosci Methods* 19(1): 1-10.

Silverman, M. S., Grosof, D. H., De Valois, R. L., & Elfar, S. D. (1989). Spatial-frequency organization in primate striate cortex. *Proc. Natl. Acad. Sci. U. S. A., 86*(2), 711-715. doi:10.1073/pnas.86.2.711

Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annu. Rev. Neurosci., 24*, 1193-1216. doi:10.1146/annurev.neuro.24.1.1193

Simoncelli, E. P., Pillow, J., Paninski, L., & Schwartz, O. (2004). Characterization of neural responses with stochastic stimuli. *The Cognitive Neurosciences, III, pages 327-- 338. MIT Press, Cambridge, MA*.

Sincich, L. C., Adams, D. L., & Horton, J. C. (2003). Complete flatmounting of the macaque cerebral cortex. *Vis. Neurosci., 20*(6), 663-686.
Sincich, L. C., & Horton, J. C. (2002). Divided by cytochrome oxidase: a map of the projections from V1 to V2 in macaques. *Science, 295*(5560), 1734-1737. doi:10.1126/science.1067902
Sincich, L. C., & Horton, J. C. (2005a). The circuitry of V1 and V2: integration of color, form, and motion. *Annu. Rev. Neurosci., 28*, 303-326. doi:10.1146/annurev.neuro.28.061604.135731
Sincich, L. C., & Horton, J. C. (2005b). Input to V2 thin stripes arises from V1 cytochrome oxidase patches. *J. Neurosci., 25*(44), 10087-10093. doi:10.1523/JNEUROSCI.3313-05.2005
Sincich, L. C., Jocson, C. M., & Horton, J. C. (2007). Neurons in V1 patch columns project to V2 thin stripes. *Cereb. Cortex, 17*(4), 935-941. doi:10.1093/cercor/bhl004
Smith, G. B., Whitney, D. E., & Fitzpatrick, D. (2015). Modular Representation of Luminance Polarity in the Superficial Layers of Primary Visual Cortex. *Neuron, 88*(4), 805-818. doi:10.1016/j.neuron.2015.10.019
Solomon, S. G. (2002). Striate cortex in dichromatic and trichromatic marmosets: neurochemical compartmentalization and geniculate input. *J Comp Neurol, 450*(4), 366-381. doi:10.1002/cne.10327
Solomon, S. G., & Rosa, M. G. (2014). A simpler primate brain: the visual system of the marmoset monkey. *Front Neural Circuits, 8*, 96. doi:10.3389/fncir.2014.00096
Song, M., Jang, J., Kim, G., & Paik, S.-B. (2018). Universality of the developmental origins of diverse functional maps in the visual cortex. *bioRxiv*.
Soodak, R. E. (1987). The retinal ganglion cell mosaic defines orientation columns in striate cortex. *Proc Natl Acad Sci U S A, 84*(11), 3936-3940. doi:10.1073/pnas.84.11.3936
Spatz, W. B. (1989). Loss of ocular dominance columns with maturity in the monkey, Callithrix jacchus. *Brain Res., 488*(1-2), 376-380. doi:10.1016/0006-8993(89)90734-8
Sperry, R. W. (1963). Chemoaffinity in the orderly growth of nerve fiber patterns and connections. *Proc. Natl. Acad. Sci. U. S. A., 50*, 703-710. doi:10.1073/pnas.50.4.703
Swindale, N. V. (1991). Coverage and the design of striate cortex. *Biol Cybern, 65*(6), 415-424.
Swindale, N. V. (1996). The development of topography in the visual cortex: a review of models. *Network: Computation in Neural Systems, 7*(2), 161-247. doi:10.1088/0954-898X_7_2_002
Swindale, N. V. (2000). How many maps are there in visual cortex? *Cereb Cortex, 10*(7), 633-643. doi:10.1093/cercor/10.7.633
Swindale, N. V. (Producer). (2004). Coverage, Polymaps and the Visual Cortex. *Mathematical Neuroscience Workshop, March 15, 2004 - March 19, 2004*. [[Online Lecture]]
Swindale, N. V. (2004). How different feature spaces may be represented in cortical maps. *Network, 15*(4), 217-242.
Swindale, N. V., Shoham, D., Grinvald, A., Bonhoeffer, T., & Hubener, M. (2000). Visual cortex maps are optimized for uniform coverage. *Nat Neurosci, 3*(8), 822-826. doi:10.1038/77731
Szarowski, D. H., Andersen, M. D., Retterer, S., Spence, A. J., Isaacson, M., Craighead, H. G., . . . Shain, W. (2003). Brain responses to micro-machined silicon devices. *Brain Res., 983*(1-2), 23-35. doi:10.1016/s0006-8993(03)03023-3
Tanaka, H., & Ohzawa, I. (2009). Surround suppression of V1 neurons mediates orientation-based representation of high-order visual features. *J. Neurophysiol., 101*(3), 1444-1462. doi:10.1152/jn.90749.2008
Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annu Rev Neurosci, 19*, 109-139. doi:10.1146/annurev.ne.19.030196.000545
Thevenaz, P., Ruttimann, U. E., & Unser, M. (1998). A pyramid approach to subpixel registration based on intensity. *IEEE Transactions on Image Processing, 7*(1), 27-41. doi:10.1109/83.650848
Thomas, J. A., & Cover, T. M. (1991). *Elements of information theory*. New York: Wiley.
Tiriac, A., & Feller, M. B. (2019). Embryonic neural activity wires the brain. *Science, 364*(6444), 933-934. doi:10.1126/science.aax8048

Tiriac, A., Smith, B. E., & Feller, M. B. (2018). Light Prior to Eye Opening Promotes Retinal Waves and Eye-Specific Segregation. *Neuron, 100*(5), 1059-1065.e1054. doi:10.1016/j.neuron.2018.10.011

Tolhurst, D. J., Tadmor, Y., & Chao, T. (1992). Amplitude spectra of natural images. *Ophthalmic Physiol. Opt., 12*(2), 229-232.

Tootell, R. B., & Hamilton, S. L. (1989). Functional anatomy of the second visual area (V2) in the macaque. *J. Neurosci., 9*(8), 2620-2644.

Tootell, R. B., Mendola, J. D., Hadjikhani, N. K., Ledden, P. J., Liu, A. K., Reppas, J. B., . . . Dale, A. M. (1997). Functional analysis of V3A and related areas in human visual cortex. *J. Neurosci., 17*(18), 7060-7078. doi:10.1523/JNEUROSCI.17-18-07060.1997

Tootell, R. B., & Nasr, S. (2017). Columnar Segregation of Magnocellular and Parvocellular Streams in Human Extrastriate Cortex. *J. Neurosci., 37*(33), 8014-8032. doi:10.1523/JNEUROSCI.0690-17.2017

Tootell, R. B., Silverman, M. S., Hamilton, S. L., De Valois, R. L., & Switkes, E. (1988). Functional anatomy of macaque striate cortex. III. Color. *J Neurosci, 8*(5), 1569-1593.

Tootell, R. B., Silverman, M. S., Hamilton, S. L., Switkes, E., & De Valois, R. L. (1988). Functional anatomy of macaque striate cortex. V. Spatial frequency. *J Neurosci, 8*(5), 1610-1624.

Tootell, R. B., Switkes, E., Silverman, M. S., & Hamilton, S. L. (1988). Functional anatomy of macaque striate cortex. II. Retinotopic organization. *J. Neurosci., 8*(5), 1531-1568. doi:10.1523/JNEUROSCI.08-05-01531.1988

Touryan, J., Felsen, G., & Dan, Y. (2005). Spatial structure of complex cell receptive fields measured with natural images. *Neuron, 45*(5), 781-791. doi:10.1016/j.neuron.2005.01.029

Touryan, J., Lau, B., & Dan, Y. (2002). Isolation of relevant visual features from random stimuli for cortical complex cells. *J Neurosci, 22*(24), 10811-10818.

Travis, D. S., Bowmaker, J. K., & Mollon, J. D. (1988). Polymorphism of visual pigments in a callitrichid monkey. *Vision Res, 28*(4), 481-490.

Ungerleider, L. G., & Haxby, J. V. (1994). 'What' and 'where' in the human brain. *Curr Opin Neurobiol, 4*(2), 157-165.

Valverde Salzmann, M. F., Bartels, A., Logothetis, N. K., & Schuz, A. (2012). Color Blobs in Cortical Areas V1 and V2 of the New World Monkey Callithrix jacchus, Revealed by Non-Differential Optical Imaging. *Journal of Neuroscience, 32*(23), 7881-7894. doi:10.1523/JNEUROSCI.4832-11.2012

Van Essen, D. C., Newsome, W. T., & Maunsell, J. H. (1984). The visual field representation in striate cortex of the macaque monkey: asymmetries, anisotropies, and individual variability. *Vision Res, 24*(5), 429-448.

Wachowiak, M., & Cohen, L. B. (2001). Representation of odorants by receptor neuron input to the mouse olfactory bulb. *Neuron, 32*(4), 723-735. doi:10.1016/s0896-6273(01)00506-2

Walker, G. A., Ohzawa, I., & Freeman, R. D. (2002). Disinhibition outside receptive fields in the visual cortex. *J. Neurosci., 22*(13), 5659-5668. doi:20026492

Wang, Y., Jin, J., Kremkow, J., Lashgari, R., Komban, S. J., & Alonso, J. M. (2015). Columnar organization of spatial phase in visual cortex. *Nat Neurosci, 18*(1), 97-103. doi:10.1038/nn.3878

Watanabe, M., Tanaka, H., Uka, T., & Fujita, I. (2002). Disparity-selective neurons in area V4 of macaque monkeys. *J Neurophysiol, 87*(4), 1960-1973. doi:10.1152/jn.00780.2000

Webb, B. S., Dhruv, N. T., Solomon, S. G., Tailby, C., & Lennie, P. (2005). Early and late mechanisms of surround suppression in striate cortex of macaque. *J. Neurosci., 25*(50), 11666-11675. doi:10.1523/JNEUROSCI.3414-05.2005

Wilder, H. D., Grunert, U., Lee, B. B., & Martin, P. R. (1996). Topography of ganglion cells and photoreceptors in the retina of a New World monkey: the marmoset Callithrix jacchus. *Vis Neurosci, 13*(2), 335-352.

Wong-Riley, M. (1979). Changes in the visual system of monocularly sutured or enucleated cats demonstrable with cytochrome oxidase histochemistry. *Brain Res, 171*(1), 11-28. doi:10.1016/0006-8993(79)90728-5

Xiao, Y., Casti, A., Xiao, J., & Kaplan, E. (2007). Hue maps in primate striate cortex. *NeuroImage, 35*(2), 771-786. doi:10.1016/j.neuroimage.2006.11.059

Xiao, Y., & Felleman, D. J. (2004). Projections from primary visual cortex to cytochrome oxidase thin stripes and interstripes of macaque visual area 2. *Proc. Natl. Acad. Sci. U. S. A., 101*(18), 7147-7151. doi:10.1073/pnas.0402052101

Xiao, Y., Wang, Y., & Felleman, D. J. (2003). A spatially organized representation of colour in macaque cortical area V2. *Nature, 421*(6922), 535-539. doi:10.1038/nature01372

Xiao, Y., Zych, A., & Felleman, D. J. (1999). Segregation and convergence of functionally defined V2 thin stripe and interstripe compartment projections to area V4 of macaques. *Cereb Cortex, 9*(8), 792-804. doi:10.1093/cercor/9.8.792

Xu, X., Bosking, W., Sary, G., Stefansic, J., Shima, D., & Casagrande, V. (2004). Functional organization of visual cortex in the owl monkey. *J Neurosci*, 24(28), 6237-6247. doi:10.1523/JNEUROSCI.1144-04.2004

Yacoub, E., Harel, N., & Ugurbil, K. (2008). High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences, 105*(30), 10607-10612. doi:10.1073/pnas.0804110105

Yarch, J., Federer, F., & Angelucci, A. (2017). Local Circuits of V1 Layer 4B Neurons Projecting to V2 Thick Stripes Define Distinct Cell Classes and Avoid Cytochrome Oxidase Blobs. *J. Neurosci., 37*(2), 422-436. doi:10.1523/JNEUROSCI.2848-16.2016

Yu, H. H., Rowley, D. P., Zavitz, E., Price, N. S. C., & Rosa, M. G. (2019). A “twisted” visual field map in the primate cortex predicted by topographic continuity. *bioRxiv*. doi:10.1101/682187

Yu, H. H., Verma, R., Yang, Y., Tibballs, H. A., Lui, L. L., Reser, D. H., & Rosa, M. G. (2010). Spatial and temporal frequency tuning in striate cortex: functional uniformity and specializations related to receptive field eccentricity. *Eur J Neurosci, 31*(6), 1043-1062. doi:10.1111/j.1460-9568.2010.07118.x

Zavitz, E., Yu, H. H., Rosa, M. G. P., & Price, N. S. C. (2019). Correlated Variability in the Neurons With the Strongest Tuning Improves Direction Coding. *Cereb Cortex, 29*(2), 615-626. doi:10.1093/cercor/bhx344

Zavitz, E., Yu, H. H., Rowe, E. G., Rosa, M. G., & Price, N. S. (2016). Rapid Adaptation Induces Persistent Biases in Population Codes for Visual Motion. *J Neurosci*, 36(16), 4579-4590. doi:10.1523/JNEUROSCI.4563-15.2016

Zeiler, M. D., & Fergus, R. (2014). *Visualizing and Understanding Convolutional Networks*. Paper presented at the ECCV, Zurich, Switzerland.

Zeki, S. (1980). The representation of colours in the cerebral cortex. *Nature, 284*(5755), 412-418. doi:10.1038/284412a0

Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural Comput., 10*(2), 403-430.

Zhou, Z. J., & Zhao, D. (2000). Coordinated transitions in neurotransmitter systems for the initiation and propagation of spontaneous retinal waves. *J. Neurosci., 20*(17), 6570-6577.

Zhu, Q., & Vanduffel, W. (2018). A new view of dorsal visual cortex in macaques revealed by submillimeter fMRI. *bioRxiv*.

Zhu, S., Allitt, B., Samuel, A., Lui, L., Rosa, M. G. P., & Rajan, R. (2019). Sensitivity to Vocalization Pitch in the Caudal Auditory Cortex of the Marmoset: Comparison of Core and Belt Areas. *Front. Syst. Neurosci., 13*, 5. doi:10.3389/fnsys.2019.00005

9 Appendix: Supplementary mapping results

9.1 Animal 2 results

9.1.1 Orientation tuning

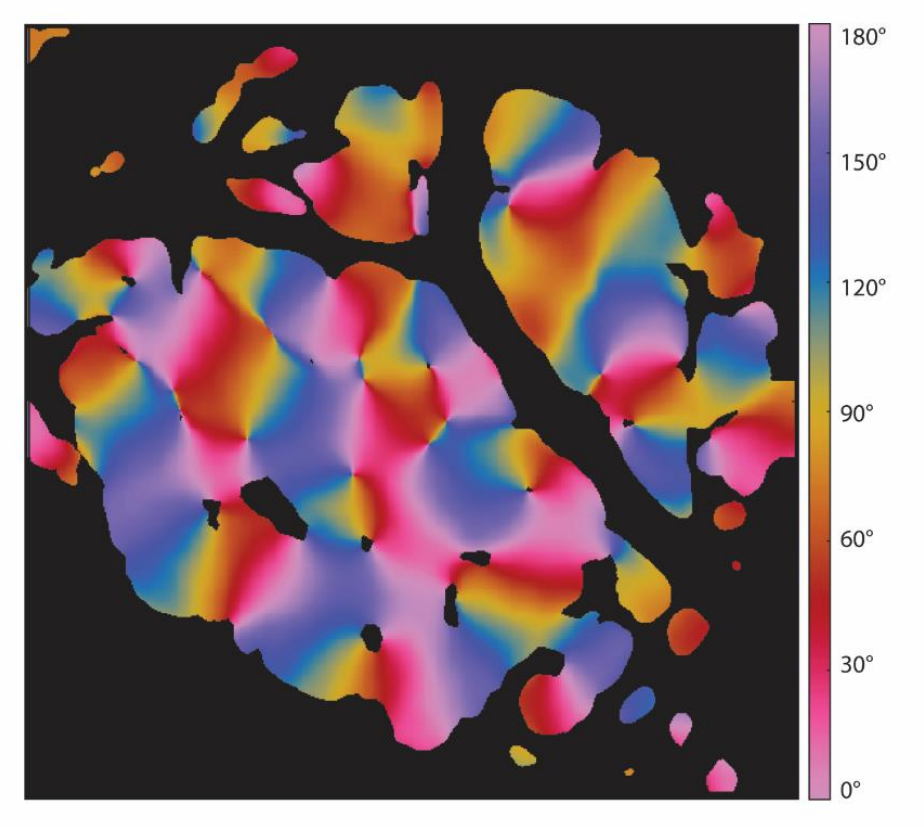

Figure 9.1. Preferred orientation map for animal 2. The preferred orientation was calculated for each pixel in the single photon image and smoothed using a 2d Gaussian filter with a 5 pixel standard deviation to give the map above. Pixels with a signal-to-noise ratio of less than 25%, or outside the imaging window, were blacked out. These regions correspond to large blood vessels that obscure the signal as well as areas of cortex with low fluorescence signal due to weak expression of the virus vector. Preferred orientation is colour coded using a repeating perceptually uniform colour scale (Kovesi, 2015) since orientation is periodic over 180 degrees.

9.1.2 Spatial frequency tuning

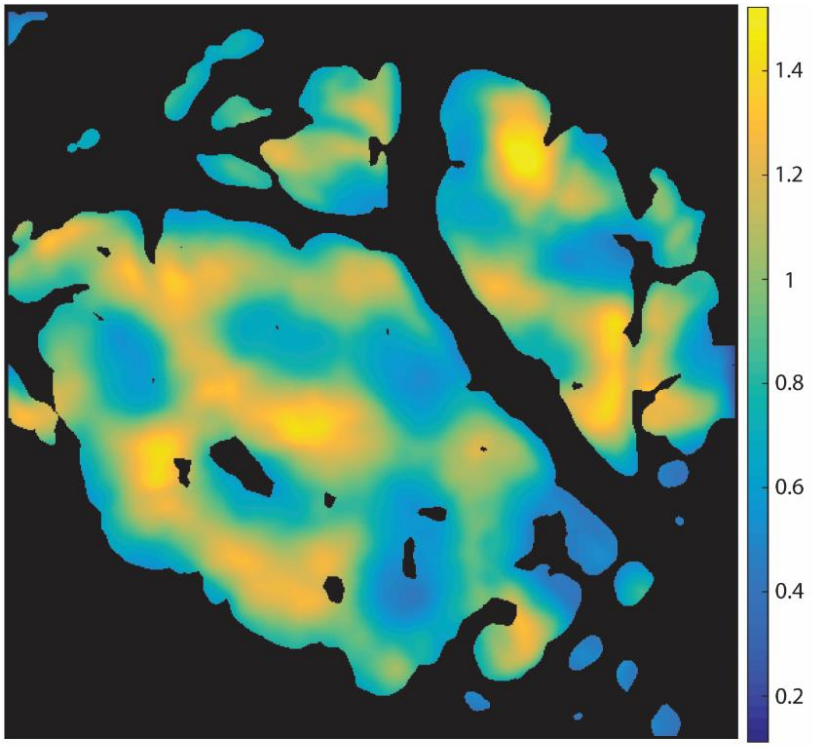

Figure 9.2. Preferred spatial frequency map for animal 2. Calculated for each pixel in the single photon image and smoothed using a 2d Gaussian filter with a 5 pixel standard deviation to give the map above. The same mask as was used for the orientation tuning map is used to black out weakly responding pixels. The resulting map creates surface with peaks and troughs rather than the bands previously reported in primates.

9.1.3 ON and OFF regions

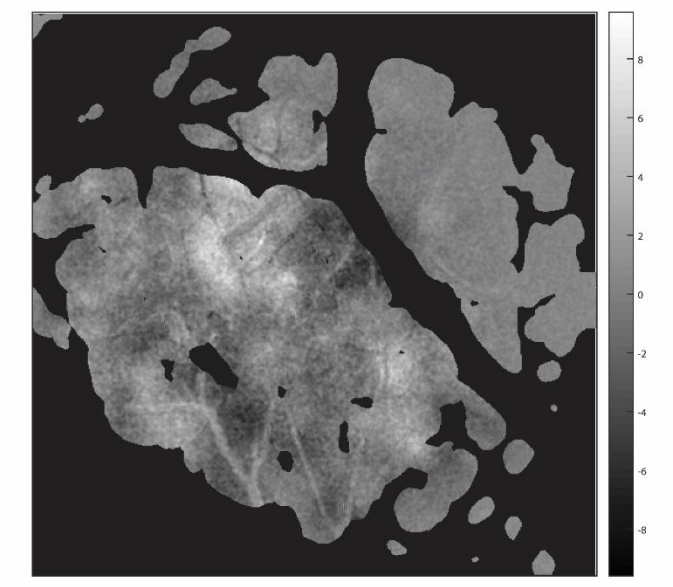

Figure 9.3. Bias to contrast polarity (1% of maximal response) for animal 2. Calculated for each pixel in the single photon image and smoothed using a 2d Gaussian filter with a 5 pixel standard deviation to give the map above. The same mask as was used for the orientation tuning map is used to black out weakly responding pixels.

9.1.4 Ocular bias

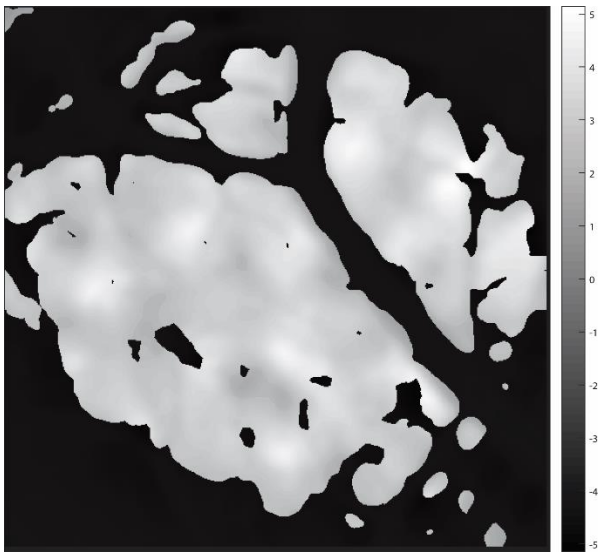

Figure 9.4. Ocular bias map (as 1% of maximal response) for animal 2. Calculated for each pixel in the single photon image and smoothed using a 2d Gaussian filter with a 5 pixel standard deviation to give the map above. The same mask as was used for the orientation tuning map is used to black out weakly responding pixels.

9.1.5 Colour sensitivity

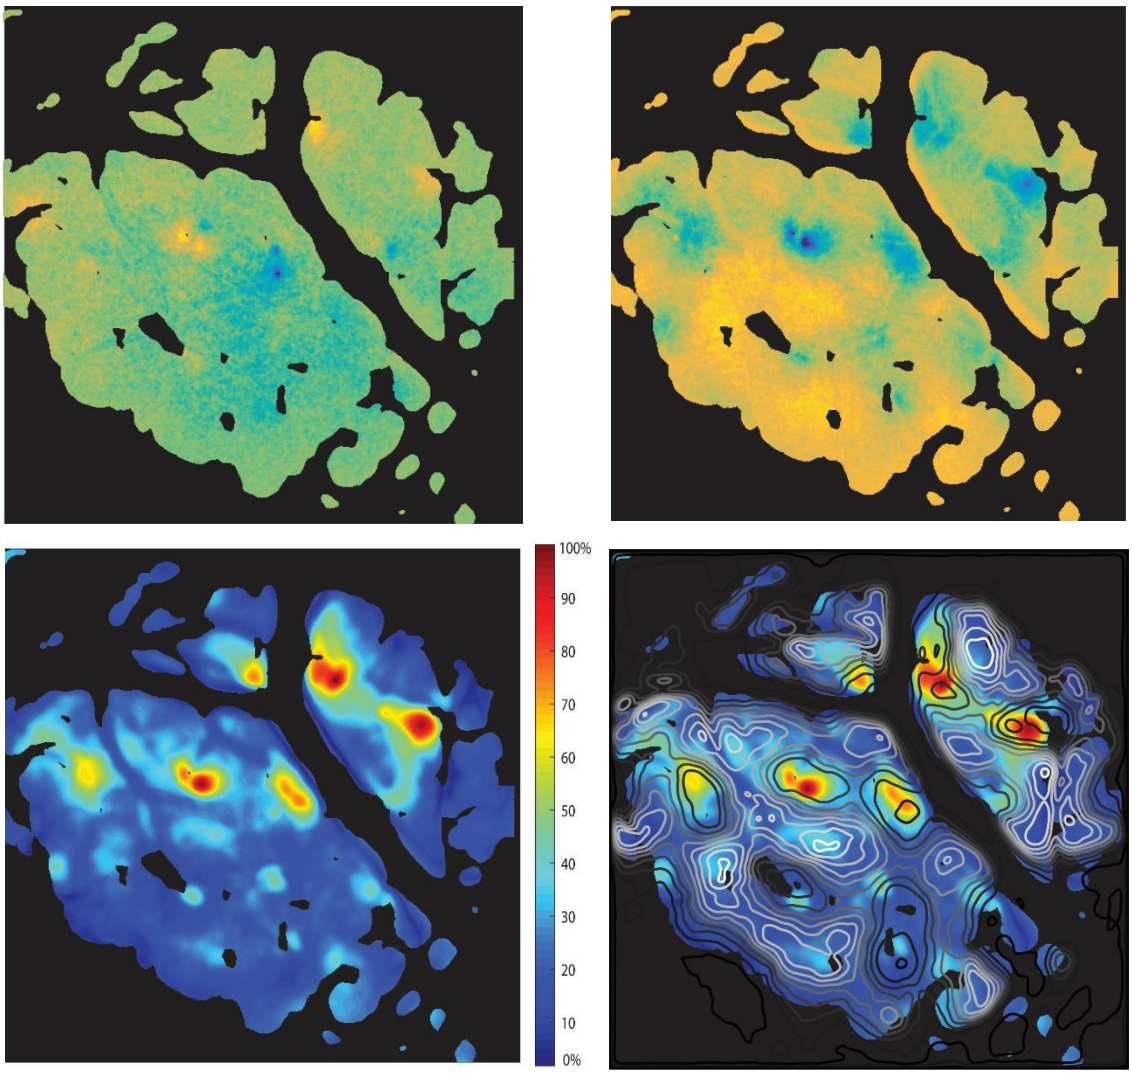

Figure 9.5. Cone contrast response maps for animal 2. Top-left: Sensitivity to S-cone modulating stimuli. For each pixel the sensitivity to S-cone isolating stimuli was calculated. Top-right: Sensitivity to L-cone modulating stimuli. Bottom-left. Signal range driven by cone modulating stimuli. This indicates the strength of colour tuning per pixel calculated by response to the best stimuli minus the response to the worst (or most suppressive) stimuli. This reveals a similar map of activity to the map of preferred spatial frequency suggesting that the population of cells that are sensitive to low spatial frequency gratings include or are colocalised with the population that receive inputs from the S-cones. Bottom-right: Contours of the preferred spatial frequency map show the colocalisation of the low (darker lines) spatial frequency sensitive regions.

9.2 Animal 3 results

9.2.1 Orientation tuning

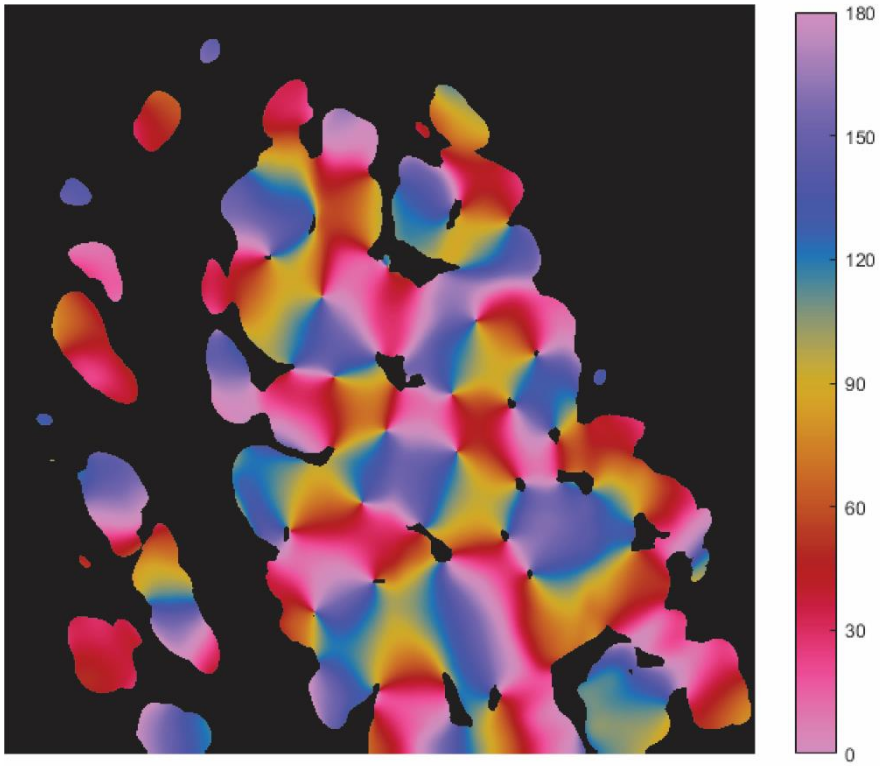

Figure 9.6. Preferred orientation map for animal 3. The preferred orientation was calculated for each pixel in the single photon image and smoothed using a 2d Gaussian filter with a 5 pixel standard deviation to give the map above. Pixels with a signal-to-noise ratio of less than 25%, or outside the imaging window, were blacked out. These regions correspond to large blood vessels that obscure the signal as well as areas of cortex with low fluorescence signal due to weak expression of the virus vector. Preferred orientation is colour coded using a repeating perceptually uniform colour scale (Kovesi, 2015) since orientation is periodic over 180 degrees.

9.2.2 Spatial frequency tuning

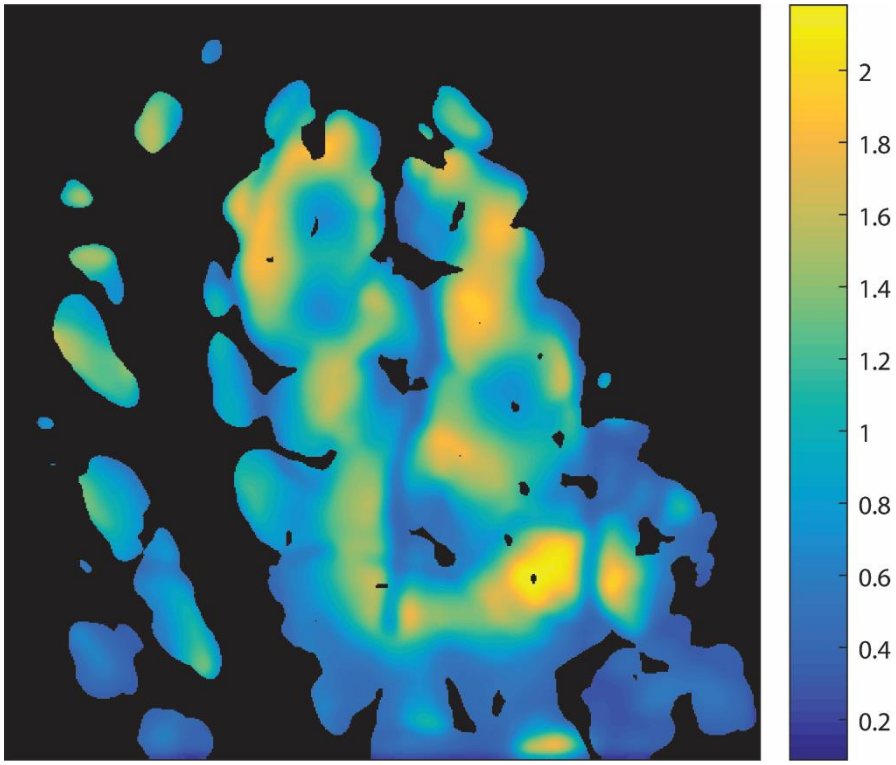

Figure 9.7. Preferred spatial frequency map (in cycle/degree) for animal 3. Calculated for each pixel in the single photon image and smoothed using a 2d Gaussian filter with a 5 pixel standard deviation to give the map above. The same mask as was used for the orientation tuning map is used to black out weakly responding pixels. The resulting map creates surface with peaks and troughs rather than the bands previously reported in primates.

9.2.3 ON and OFF regions

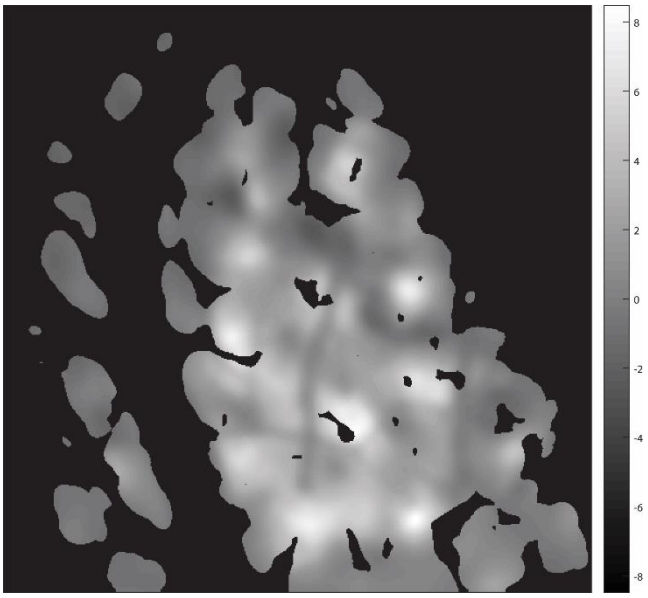

Figure 9.8. Bias to contrast polarity (1% of maximal response) for animal 3. Calculated for each pixel in the single photon image and smoothed using a 2d Gaussian filter with a 5 pixel standard deviation to give the map above. The same mask as was used for the orientation tuning map is used to black out weakly responding pixels.

9.2.4 Ocular bias

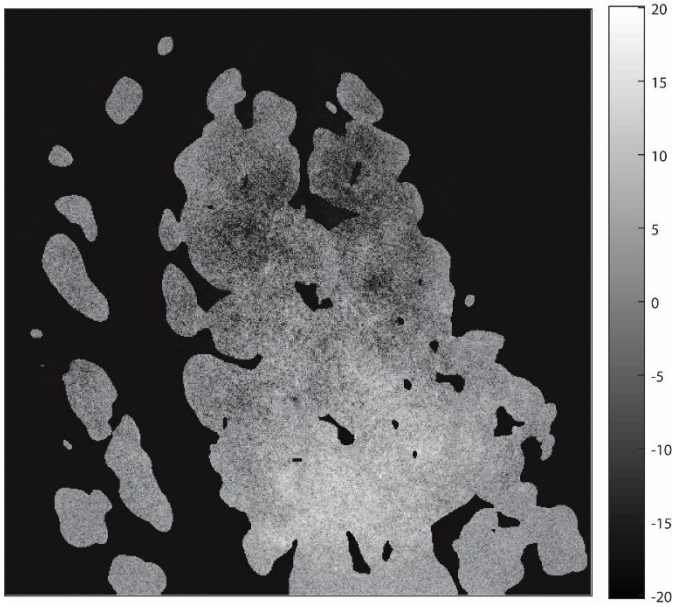

Figure 9.9. Ocular bias map (as 1% of maximal response) for animal 3. Calculated for each pixel in the single photon image and smoothed using a 2d Gaussian filter with a 5 pixel standard deviation to give the map above. The same mask as was used for the orientation tuning map is used to black out weakly responding pixels.

9.2.5 Colour sensitivity

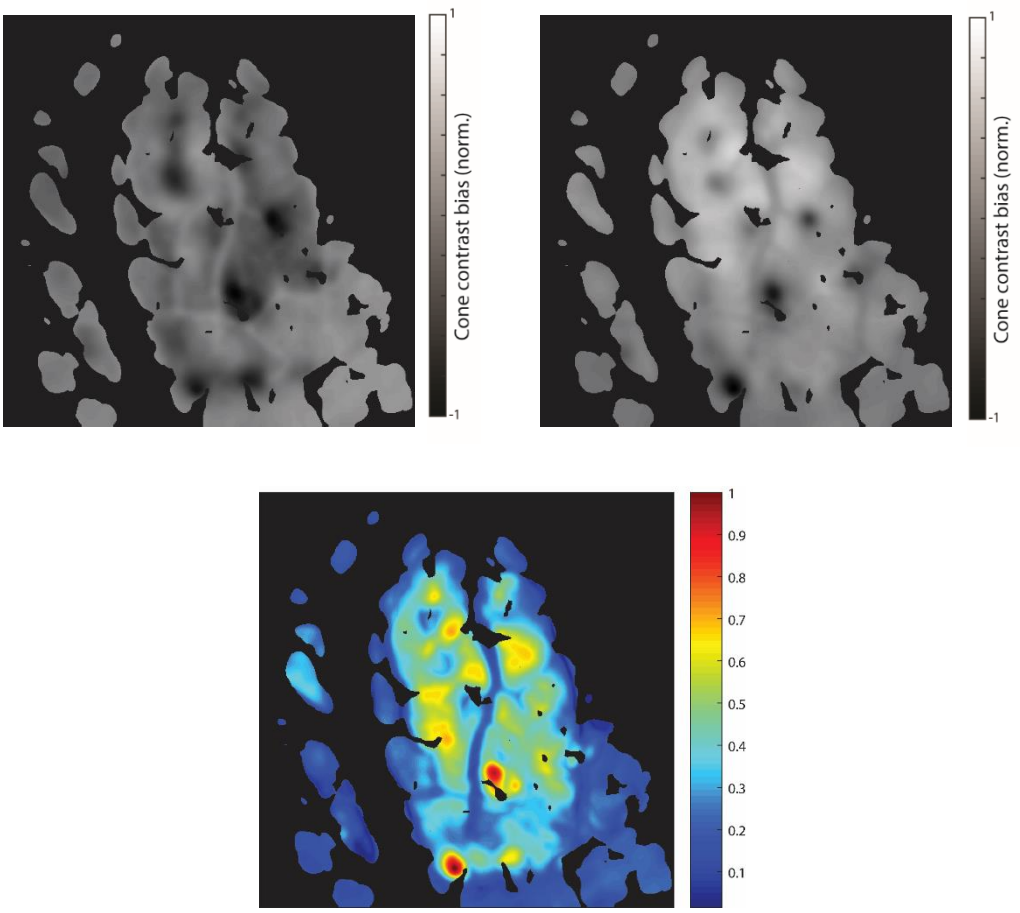

Figure 9.10. Cone contrast response maps for animal 3. Top-left: Sensitivity to S-cone modulating stimuli. For each pixel the sensitivity to S-cone isolating stimuli was calculated. Top-right: Sensitivity to L-cone modulating stimuli. Bottom. Signal range driven by cone modulating stimuli. This indicates the strength of colour tuning per pixel calculated by response to the best stimuli minus the response to the worst (or most suppressive) stimuli.