

ABSTRACT

Title of Document: LONGITUDINAL SINGLE-UNIT
RECORDING IN THE MACAQUE FACE
PATCH SYSTEM: IDENTITY AND
PLASTICITY IN THE ANTERIOR FUNDUS
FACE PATCH

Adam P. Jones, Doctor of Philosophy, 2015

Directed By: Dr. David A. Leopold, National Institute of
Mental Health

Face perception, a fundamental component of primate social behavior, is supported by a network of specialized visual regions within the ventral visual stream of humans and macaques. Discrete regions, or “patches” within this network respond preferentially to face images over non-face object images, with the majority of visually responsive neurons within these regions responding selectively to faces. In recent years, the functional specialization of neurons within particular fMRI-defined face patches has been studied intensively. In this series of studies, we have investigated one such patch (AF) located in the anterior fundus of the superior temporal sulcus using a new method longitudinal of electrophysiological single unit recording, where individual neurons can be isolated and monitored for several months using chronically implanted electrodes. We made use of this unique opportunity to study the responses of face-selective neurons in two different ways. Our first approach was the longitudinal observation of neurons over time scales relevant to behavioral learning. By observing single-units longitudinally, over several weeks of natural visual experience, as well as throughout the course of an intensive face-learning paradigm, we asked whether individual neurons would change their response selectivity. In the case of the learning paradigm, we asked whether there might be

changes in face-selective neural responses that occur while the animal gradually acquired greater sensitivity for the distinctive features of novel face stimuli. For this, we taught the animal to recognize individual monkeys and humans using morphed faces of diminished identity levels. We found AF neurons were remarkably stationary in their responses over time scales of weeks and months, and even for periods as long as one year. Even during periods of intensive training using the stimuli for which the neurons appeared specialized, AF neurons demonstrate very limited change in their response patterns and selectivity. In our second approach- taking advantage of the fact that the neural response selectivity is unchanging across sessions- we accumulated neural responses to a large number of visual stimuli presented over multiple sessions, which allowed for both a broader sampling of the visual stimulus space and an enhanced sensitivity to the preferences of low-firing rate neurons. In this approach, we focused on visual responses to faces, asking how a conjunction of identity, species, viewing angle, and spatial scale shaped single-unit responses in face patch AF. Neurons in AF showed selective responses to each of these manipulations to different degrees. Within the domain of faces, most striking determinants of neural firing were face scale and the systematic, norm-based tuning for face identity. Many neurons were also strongly selective for macaque faces over human faces. Other factors, such as 3-dimensional face rotation, were less important. In addition to insight into the functional compartmentalization of the face-processing network, these results from face patch AF also provide a unique view into the relationship between high-level neural selectivity in the brain and the inherent trade-off between stability and plasticity that accompanies learning-induced changes in behavior.

LONGITUDINAL SINGLE-UNIT RECORDING IN
THE MACAQUE FACE PATCH SYSTEM:
IDENTITY AND PLASTICITY IN THE ANTERIOR FUNDUS FACE PATCH

By

Adam P. Jones

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2015

Advisory Committee:

Professor David A. Leopold (NIMH/NIH) - Research Adviser
Professor Daniel A. Butts (Biology) - Academic Adviser, Chair
Professor Jonathan Z. Simon (Biology)
Professor Bruno B. Averbeck (NIMH/NIH) - Special Member
Professor Donald J. Bolger (Human Development) - Dean's Rep.

© Copyright by
Adam P. Jones
2015

Preface

All electrophysiological and behavioral experiments were performed by Adam P. Jones. Chronic electrode implantation surgeries were performed by David B. T. McMahon, Brian E. Russ, and David A. Leopold.

One publication resulting from this work has been published to date:

David B. T. McMahon, Adam P. Jones, Igor V. Bondar, and David A. Leopold

Face-selective neurons maintain consistent visual responses across months.

Proceedings of the National Academy of Sciences, 111(22), 8251-8256.

The brains of humans and other primates contain specialized regions dedicated to the perception of socially important objects such as faces. It is not known whether face representations remain stable over time, or alternatively whether they are subject to slow change in response to ongoing experience. By using implanted microwires to monitor the activity of single cells longitudinally across several weeks and up to one year, we demonstrate that face-selective neurons maintain the same distinctive selectivity patterns for as long as they were followed. The long-term consistency of these neurons might reflect a division of labor within the face-processing network for aspects of social perception that require stability as opposed to plasticity.

This work is described in Chapter 3.

A publication is currently being prepared from data presented in Chapter 4.

Adam P. Jones, David B. T. McMahon, Brian E. Russ, David A. Leopold

Neurons in macaque face patch AF exhibit norm-based tuning for face identity. *In preparation.*

Dedication

I dedicate my dissertation work to my family; my parents who have always supported me (through the good and the bad), and my sister who has always inspired me and kept me in touch with what's important in life.

I also want to give special thanks to my advisers; David, who taught me to expect more from myself, and Dan, who has always come through for me.

Finally, my lab-mates; I could not have ever asked for a better group to teach me how to be a good scientist, as well as a good person.

That said, dedications are silly. lol

Acknowledgements

The research presented in this volume would not have been possible without the contributions of the following people: David B.T. McMahon, Heba D. Elnaiem, Luba N. Vasilyeva, and Brian E. Russ.

Table of Contents

Preface.....	ii
Dedication.....	iii
Acknowledgements.....	iv
Table of Contents.....	v
List of Tables.....	vii
List of Figures.....	viii
Chapter 1: Introduction – Face processing in the macaque.....	1
I. Psychophysical observations.....	5
1. Configural processing.....	5
2. Special dimensions in “face space”.....	6
3. Advantages of norm-based coding of faces.....	11
II. Physiological observations.....	13
1. Neurophysiological properties of face-selective neurons.....	13
2. Face-selective patches in the macaque temporal lobe.....	14
3. Electrophysiological evidence for norm-based processing.....	16
4. IT plasticity.....	17
5. Experimental limitations to studying plasticity in IT.....	19
Chapter 2: Methods – Microwires for long-term recordings.....	21
I. Introduction.....	21
1. Tracking changes in response selectivity of individual cells over weeks and months.....	21
2. Using stable isolation to accumulate repeated responses over extended periods.....	22
II. Chronic recording with microelectrode bundles.....	23
1. Subjects.....	23
2. Apparatus.....	23
3. MRI.....	24
4. Surgical implantation.....	25
5. Assessing stability of single-units over consecutive recording sessions....	27
6. Analysis.....	28
III. Measuring responses.....	29
1. Stimuli.....	30
2. Recording Procedure.....	33
3. Stimulus presentation.....	33
4. Spike-sorting.....	34
Chapter 3: Stability of stimulus selectivity in AF neurons.....	35
I. Introduction.....	35
II. Results.....	39
1. Response significance and response latency.....	39
2. Waveform stability.....	40
3. Selectivity patterns over time.....	41
4. Response consistency assessments.....	43
5. Channel-by-channel isolation.....	44
III. Discussion.....	45

Chapter 4: Representations of faces in AF	48
I. Introduction	48
II. Results	48
1. Identity	49
2. Tuning for other properties of faces in AF	70
III. Discussion	76
1. Identity encoding	76
2. Tuning for natural transformations of the face and head	83
Chapter 5: Effects of learning on face-selective neurons	88
I. Introduction	88
I. Methods	91
1. Behavioral training.....	91
2. Analysis.....	93
II. Results	94
1. Behavioral results.....	95
2. Neurophysiological results.....	98
III. Discussion	104
Chapter 6: Conclusions	108
I. Longitudinal recordings.....	109
II. Norm-based encoding of face identity	110
1. Familiarity effects	110
2. Species-selectivity.....	111
III. Selectivity for size.....	112
IV. Learning	113
V. Future experiments.....	114
Appendices.....	115
Bibliography	118

List of Tables

2.1 – Summary of study design

5.1 – Population analyses results of face-training paradigm

List of Figures

Chapter 1

- 1.1 – Mars Face
- 1.2 – Models of “Face space”
- 1.3 – Face Trajectories
- 1.4 – Face Patch Schematic

Chapter 2

- 2.1 – Microdrive
- 2.2 – MRI
- 2.3 – Repetition Example
- 2.4 – Face Average
- 2.5 – Fiduciary Point Example

Chapter 3

- 3.1 – Young and Old Faces
- 3.2 – Latency
- 3.3 – Waveform Stability
- 3.4 – Visual Response Stability
- 3.5 – Longitudinal Visual Selectivity
- 3.6 – Stability Index Examples
- 3.7 – Population Stability Summary
- 3.8 – Channel Isolation

Chapter 4

- 4.1 – Face Identity
- 4.2 – Monkey Face Trajectory Response
- 4.3 – Monkey Face Identity Responses
- 4.4 – Monkey Face Dodecahedrons
- 4.5 – Species Effect
- 4.6 – Population Species Effect
- 4.7 – Population Identity Trajectories
- 4.8 – Adaptation Results
- 4.9 – Caricature Rasters
- 4.10 – Caricature Trajectory Averages
- 4.11 – Population Caricature Responses
- 4.12 – Population Tangential Responses
- 4.13 – Orientation Selectivity
- 4.14 – Population Scale Selectivity

Chapter 5

- 5.1 – Saccade Task
- 5.2 – Experimental Timeline
- 5.3 – Behavioral Performance
- 5.4 – Psychometric Curve

- 5.5 – Longitudinal Rasters
- 5.6 – Longitudinal Population Responses

Appendix

- 7.1 – Category Selectivity Stimuli
- 7.2 – Identity Selectivity Stimuli
- 7.3 – Training Stimuli
- 7.4 – Caricature Stimuli
- 7.5 – Orientation Selectivity Stimuli

Chapter 1: Introduction – Face processing in the macaque

Faces contain information regarding the identity, age, sex, and emotional status of an individual. Social primates such as humans have evolved to be experts in social communication through facial signaling. Perhaps not surprisingly then, the visual system of social primates seems to be adapted to process faces as “special” stimuli, to the point of occasionally perceiving faces when they aren't even present (e.g. the face on Mars, see Figure 1.1). It has been suggested that this predisposition toward seeing faces is a side effect of possessing a highly specialized neural system that is dedicated specifically for face processing.




Figure 1-1 Seeing faces that aren't. A 2km mesa in the Cydonia region of Mars with a striking resemblance to a humanoid face was imaged in 1976 by the Viking 1 Orbiter. This landmark has been widely regarded as an example of the psychological phenomenon of pareidolia, in which a stimulus triggers the perception of a familiar pattern where none actually exists. Common examples of which include faces, animals, and body parts.

Faces seem to be special stimuli in that they are processed in a fashion distinct from other visual objects. A primary piece of evidence for this is the face inversion effect, in which a face becomes especially difficult to recognize once turned upside down, relative to the same manipulation of objects (Yin, 1969). It's been suggested

that this ‘diagnostic’ feature of face recognition is a result of encoding faces as whole template rather than as a sum of independent features (i.e. “holistically”), a property not observed for non-face objects (Tanaka & Farah, 2003). In addition, individual face parts (eyes, mouth, etc.) are better discriminated when presented within the context of an intact face, than when presented alone (Tanaka & Farah, 1993).

However, the inversion effect has been shown to exist for non-face objects, as well, when subjects possess a high-level of training or familiarization with those objects (Diamond and Carey, 1986). For example, it has been shown that dog experts, with sufficient knowledge of a particular breed, exhibit similar holistic processing properties such as the inversion effect. Thus, one issue that has not been concretely resolved is whether these properties are a consequence of innate, specialized neural machinery or the extensive experience (overtraining) with faces that primates undergo during development.

From a behavioral perspective, the human brain does seem to be “pre-wired” to process faces and regard them as important stimuli. Humans are born with an innate tendency to observe faces -- newborns have been shown to track schematic faces longer than their scrambled counterparts (Johnson et al., 1991). Face recognition abilities continue to develop for many years beyond this, increasing dramatically between 6 and 16 years of age (Mondloch et al., 2003). Not surprisingly then, face-selective regions of cortex develop more slowly and over a longer period than object-selective regions. Correspondingly, improvements in face recognition memory also continue into adulthood (Golarai et al., 2007). This developmental trajectory allows the average adult to become familiar with thousands of faces

throughout their lifetime and to continue to identify them for years following relatively brief exposure (Bahrnick et al., 1975).

At the same time, facial expertise undergoes continuous modification throughout one's lifetime. The highest rate of learning takes place during early development, when there is a "perceptual narrowing" of face expertise for those types of faces most commonly seen – typically those from one's own race ("other-race effects"; Malpass & Kravitz, 1969) as well one's own species ("other-species" effects;. Keating & Keating, 1993). However, extensive exposure to different faces at a young age can shift expertise to other races and even other species (Sugita, 2008; Pascalis et al., 2005). In the adult brain, new faces are still readily learned and remembered, with the sensitivity for individual face recognition continuing to improve until the fourth decade of life for humans (Germine, 2011). Healthy individuals can retain the ability to learn and recognize complex visual objects, such as faces, for the rest of their life (Standing, 1973).

As children, as well as adults, we are regularly required to learn and remember the identity of faces. We have all had experiences in situations in which we can remember a face, but not the name, of an individual we've encountered in the past. This is presumably a reflection of our specialized ability, as a highly visual species, to encode and retrieve visual information. But how, and when, does this specialized ability become available to us? The process of developing our unique ability to process faces, along with our ability to learn new faces and other types of objects, presumably reflect neural plasticity in the brain. Where and how might such plasticity be expressed, and how does learning-induced plasticity affect the existing

long-term representations in the brain? Abundant research points to the inferotemporal cortex as a critical site for high-level visual processing in primates. Recent work in the macaque has emphasized a network of 4-6 face-selective regions, hereafter referred to as “face patches”.

This dissertation outlines a series of studies investigating the representation of faces and objects, as well as the plasticity of these representations, within a face-selective region in the macaque temporal cortex. These studies examine the responses of single neurons in face patch AF to variations in the semantic category of visual stimuli, including faces, and how manipulations in the orientation or identity of both faces and objects are represented by these so-called, “face neurons”. The remainder of Chapter One provides an overview of certain psychophysical and physiological aspects of face processing that will be relevant to the experiments described in Chapters 3 through 5. Chapter Two includes a detailed description of the methodological aspects of those experiments. The results in each chapter following that have been organized in the following format: Chapter 3 contains results concerning the stability of visual neurons in face-patch AF following repeated presentation of the same stimuli. Chapter 4 focuses on the unique response properties of AF neurons without regard for their stability over time. Chapter 5 describes the behavior of these neurons throughout the course of an intensive training paradigm, which was intended to modify behavioral expertise with select monkey faces. Finally, Chapter 6 primarily concerns the interpretation of the results presented in this volume, and makes recommendations regarding future experiments that are called for in order to continue this line of inquiry.

I. Psychophysical observations

A wealth of evidence indicates that humans and other primates are able to readily identify faces, and that this skill requires a level of perceptual processing that is fundamentally different from most types of visual recognition. What follows is a summary of some unique aspects of face perception as well as the description of a theoretical model, norm-based encoding, which attempts to account for them.

1. Configural processing

Compared to other forms of visual recognition, face recognition is unique in that it entails discrimination among stimuli that all possess the same basic features. All faces have a nose, two eyes, etc., laid out in the same basic (“first order”) arrangement, meaning that recognition of an individual using the appearance of the face must rely on higher order visual information. Further, such recognition must be accomplished despite changes to the retinal image that arise because of alternate viewing angles, lighting, and facial expressions. Thus to reliably distinguish among individuals, a face recognition system must be sensitive to distinctive aspects of individual faces and at the same time sufficiently flexible to overcome irrelevant image transformations that commonly take place. This distinctiveness can take the form of an extreme feature (e.g. very large ears), or a combination of individual features as well as their configuration. While the first-order organization of features is highly similar across all faces, their second-order configuration (e.g. the distance between the eyes relative to the size of the nose) and shape, vary markedly from individual to individual and are thus useful for identity recognition.

Considerable psychophysical work has shown that the cognitive representation of faces is holistic, encoding faces as whole template rather than as a sum of independent features (Farah et al., 1998). One experiment demonstrating the importance of holistic processing in face perception demonstrated that subjects are better at distinguishing two face parts in the context of a whole face than in isolation (Tanaka and Farah, 1993). This has been taken as evidence that faces are perceived as whole entities as this advantage seems to be small or absent when identifying non-face objects, for subjects without prolonged exposure to test stimuli (Tanaka and Gauthier, 1997).

However, it is possible to develop similar holistic processing after gaining perceptual expertise, even among stimuli that do not closely resemble faces. For example, Gauthier et al. (1999) provided extensive training to some of their subjects, which they termed "experts," in distinguishing among a family of "greebles". The greeble stimuli are a homogenous set of computer-generated non-face stimuli that are similar to faces along several dimensions. Following training, expert subjects were able to recognize individual greebles based on configural cues, similar in many ways to the perceptual expertise surrounding faces. These experiments raised several questions including where and how the brain changes when confronted with new visual problems, and how acquired perceptual expertise for non-face objects relates to the specialized processing observed with faces.

2. Special dimensions in “face space”

Since the 1970's, it has been well established that the distinctiveness of a face is an important factor in our ability to recognize it (e.g. Going & Read, 1974).

Numerous studies have shown that a distinctive or unusual face is generally more easily remembered than a typical face in a recognition task. This has been interpreted by some as evidence for the role of prototypes, or “norms” in the recognition of faces (Valentine & Bruce, 1986).

One useful heuristic that is often used to discuss formal models that might account for this is the multi-dimensional “face space”, defined by the dimensions that allow us to discriminate faces most easily. The center of this proposed space is assumed to represent the average value of the population on each dimension. Within this framework, the quantity and quality of dimensions are issues are left open to debate. However, it is assumed that a high number of dimensions are required to accurately define this space. It is also assumed that faces will form a normal distribution on each dimension and so extreme faces will be outliers. .

In a norm-centered model of face space, the identity of each face is represented as a vector within this space, describing the position of each face relative to the population average (see Figure 1.2). Many such trajectories intersect at the average face. When faces are transformed along each vector, this is represented by increases in the deviation of a face from the norm (“radial” trajectories; see Figure 1.3A). In contrast, in an exemplar-based model of face space, faces are represented by their similarity to exemplars of previously learned faces, rather than the average (Ross et al., 2013).

Perceptual studies have shown that certain face manipulations within this face space model are particularly influential with regard to the perceived face identity. Transformations corresponding to changes in the deviation of a face, relative to the

population average, are typically recognized as belonging the same individual, even for faces very near to the norm (Valentine, 1991). This effect is not observed when faces are transformed along vectors joining two different individuals (“tangential” trajectories), in which faces are perceived as changing dramatically from one identity to another.

Within the same norm-based framework, caricatures represent an extrapolation of the vector joining the average face to a particular individual. As predicted by this model, caricature faces are more easily recognized than their original counter-parts, presumably owing to the increased distinctiveness separating caricatures from the norm. Likewise, those lying on the same axis, but leading toward the opposite side of the mean (so-called ‘anti-faces’), are perceived as an entirely different individual (Leopold et al., 2001).

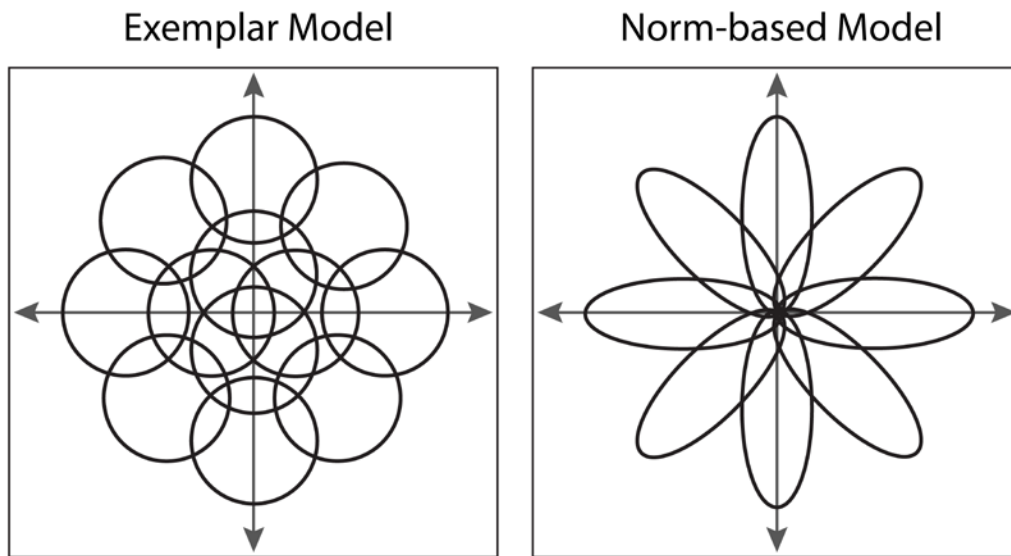


Figure 1-2 Prevailing models of multi-dimensional “face space”. Shown here are exemplar-based (left) and norm-based (right) models of face space. In each model, the perceptual identity of different faces are each represented by a circle. In the exemplar based model, all faces are represented in terms of the unique geometry of their individual facial features. Alternatively, in the norm-based model, all faces are represented as vectors, describing their position relative to the population average.

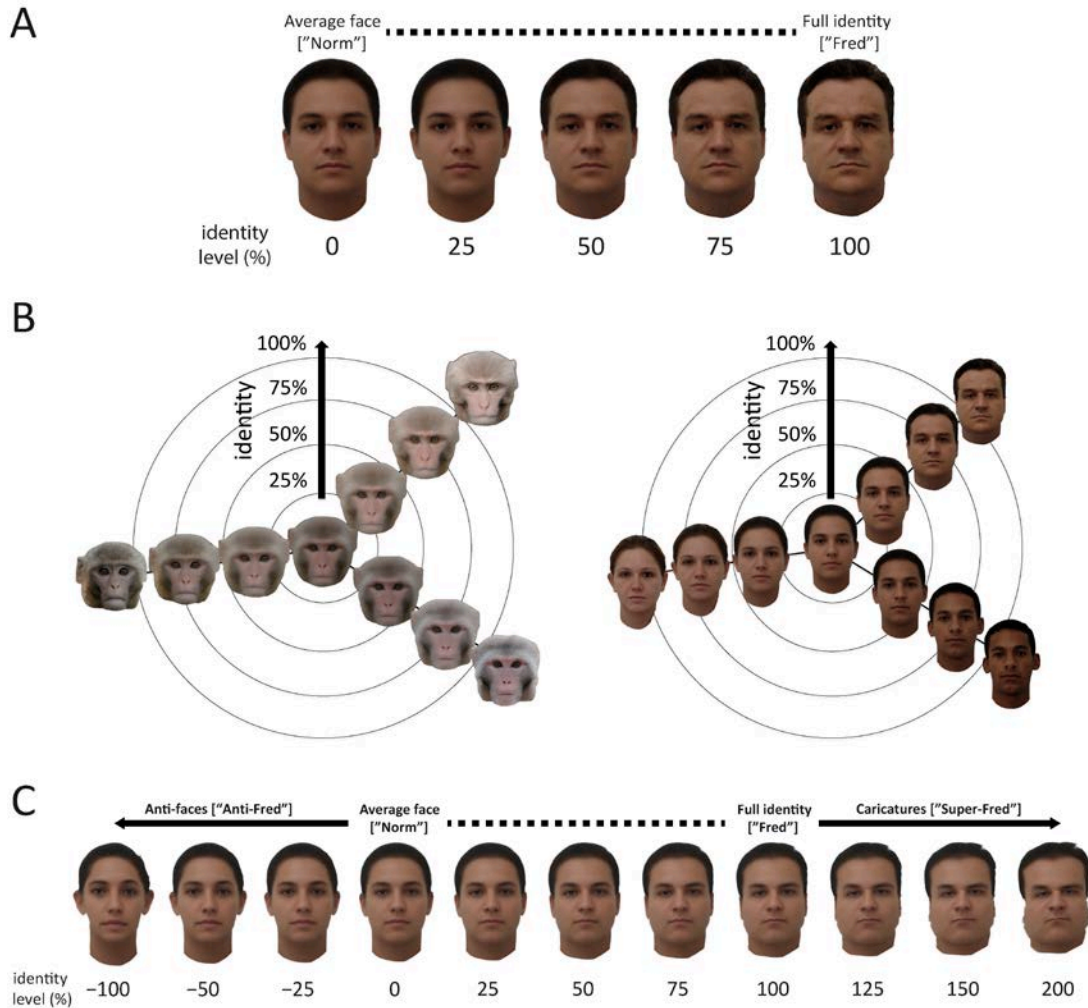


Figure 1-3 Norm-centered face space- (a) An example “identity trajectory”, in which a face is gradually “morphed” between a population super-imposed average (0%, left) and the original (100%, right). As the identity level of each face increases, so does its distinctiveness from the average. (b) A simplified (two-dimensional) face-space composed of three identity trajectories, in which all converge on the shared average of the population. (c) An expanded version of a single identity trajectory, in which faces are morphed between the “anti-face” and the caricature of the original, full identity face. Anti-faces (e.g. Anti-Fred) have been geometrically transformed to contain the opposite features of the full identity face (e.g. Fred), relative to the average. Similarly, caricature faces (e.g. Super-Fred) have been transformed to increase the features distinguishing Fred from the average, thereby exaggerating the distinctive features of the face.

This framework, which is supported by a wide range of additional psychophysical experiments discussed below, essentially suggests that individual recognition hinges on the internal representation of the average face of a given category, which acts as a prototype or “norm” against which the structure of incoming faces is compared. This notion stands somewhat in contrast to alternative models, wherein faces could be encoded as a single point in multi-dimensional space,

represented by tuning functions that do not depend upon a particular point of space as a reference. In this scenario, faces could be encoded jointly by neural populations exhibiting bell-shaped tuning curves whose peaks collectively span the entirety of face space (Tsao and Freiwald, 2006).

Using digital morphing, it is possible to simulate trajectories spanning certain sections of face space. Note that the morphing transformation can be extended beyond either end of this trajectory, exaggerating the features of the original individual to create a caricature, or creating a new identity face with nominally opposite features, sometimes called an “anti-face” (see Figure 1.3C).

Encoding of identity trajectories might be achieved by a form of opponent coding, analogous to that found in color vision. In color vision, opponency can be demonstrated by simple visual adaptation between complementary colors, for example red and green. For example, prolonged exposure to a red patch shifts the perception of a neutral (white) patch to appear green (see Webster and Mollon, 1991). Such “aftereffects” are a psychophysical means to inferring coding principles in the central nervous system.

Numerous facial characteristics have been shown to exhibit aftereffects in a similar manner to those of color vision, for example: sex, race, expression, normality, attractiveness, and eye gaze direction (Rhodes and Leopold, 2011; Webster et al., 2004). In recent years, aftereffects following adaptation to faces have been used to lend further support to norm-based models of face encoding (Webster and MacLin, 1999). For example, a study by Leopold et al., (2001) reported perceptual aftereffects that were larger for identities at the opposite end of a norm-centered face space than

for non-opposite, but equally dissimilar identities. This result has been used to suggest that opponent processes are employed to represent face/anti-face pairs, the trajectory between which constituting a single, identity-based dimension. Such effects are generally not observed when subjects are adapted to average faces, suggesting that it may occupy an especially neutral position on the axis of adaptation.

There may be a relationship between the effects of short-term perceptual adaptation and the long-term plasticity associated with the acquisition of facial expertise. But, to date there exists very little evidence linking these processes, which exist on very different time scales. In one study, Leopold et al. (2005) found that, similar to previously described visual aftereffects, the face identity aftereffect is modulated as a function of test/adaptation time over periods up to tens of seconds. These authors found the identity aftereffect grew logarithmically stronger as a function of adaptation time and exponentially weaker as a function of test duration. There is also evidence for longer lasting after effects (up to 24 hours) that have been taken to suggest a role for ongoing adaptation processes in the maintenance of these properties throughout our lifetimes (Carbon et al., 2007). While far from definitive, these experiments suggest a possible mechanism for shaping one's perceptual machinery with experience, in which the accumulation of adaptive mechanisms over time gradually shapes the neural representation of one's internal prototype or norm.

3. Advantages of norm-based coding of faces

A norm-based face processing system has the potential to offer several advantages with regard to representation of faces by mechanisms within the brain. Since faces all share the same general configuration, one efficient way to individuate

them would be to encode how each face deviates from a common configuration or norm. In this framework, faces could be represented in terms of their deviation from a mean, rather than their explicit structure, much of which is shared (and thus redundant) across the larger population of faces (Giese and Leopold 2005). In addition, by emphasizing the response to distinctive faces, responses to more frequently encountered faces would be reduced relative to extreme, and therefore less frequently encountered, stimuli.

A norm-based system for processing faces could also offer greater flexibility than feature-based alternatives. Studies of dissociable aftereffects imply that the face-perception system possesses multiple norms for different viewing angles (Jeffery et al., 2007). Such category-specific face norms may be useful in that they can be updated following experience with stimuli belonging to a particular category. A parallel representation system consisting of multiple, adaptive norms could enable the viewer to perform expert-level discrimination of a variety of face categories, regardless of the group to which they belong (Rhodes and Leopold 2011). It is also possible that such a scheme, in elucidating the expected image transformations for a given facial identity, can aid in the reading of subtle facial expression. It remains unclear, however, to what extent a norm should be considered an explicit, neural mechanism or a theoretical or psychological heuristic. If the norm is to be considered a neural entity, it should be possible to characterize it in terms of electrophysiological responses. It should also be possible to track its changes over time, such as during development or the acquisition of perceptual expertise.

II. Physiological observations

1. Neurophysiological properties of face-selective neurons

A wide range of results suggest face and object processing relies on specialized, high-level areas within the ventral visual pathway of the primate brain, and particularly on the inferior temporal cortex. This idea is contributed in part by the following findings: 1) lesions of the inferior temporal (IT) lobe in both humans (Meadows, 1974) and monkeys (Mishkin & Pribram, 1954) can result in an inability to recognize faces or objects, and 2) a subset of neurons in IT respond selectively to faces, and other biologically relevant stimuli (Gross et al., 1972). Such results led some to speculate that individual neurons might have dedicated roles in the processing of certain types of object stimuli, predicting, for example the existence of "gnostic units" (Konorski, 1967). Nonetheless, when Gross and colleagues' reported neurons with electrophysiological response properties fitting this concept (1969, 1972), the results were not initially well accepted. However, over the next years, myriad studies in the monkey confirmed and elaborated high-level neural specialization for objects, and particularly for biologically relevant features such as faces, bodies and hands. Moreover, neurons selective for faces were shown to possess a high degree of image invariance, responding with similar stimulus specificity across different image sizes, positions, and orientations (Perrett et al., 1982), leading to the somewhat hasty conclusion that selectivity for complex objects was simply an extension of the selective pooling mechanisms already proposed by Hubel and Wiesel, following their studies of primary and secondary regions of the occipital cortex (Hubel and Wiesel, 1962). Within the proposed hierarchy, output from

occipitotemporal cortex would lead to increased complexity of receptive-fields in IT (and later pre-frontal cortex) through feed-forward pooling of the output of afferent neurons, themselves possessing relatively simple receptive-fields, resulting in selectivity for complex visual objects (such as hands or bodies) in a position invariant manner. Ultimately, object recognition would be accomplished by transforming retinal input into increasingly abstract representations, which become more invariant to the physical properties of the stimulus (Farah, 2004). As support for this view, receptive-field size and response latency increase as one progresses along the ventral visual pathway towards IT (Thorpe & Fabre-Thorpe, 2001). Although anatomical and physiological evidence for such a simplified framework is lacking, this idea has dominated much of the thinking on the role of face-cells within the visual system for several decades (e.g. Gross, 1994).

More recently however, a richer picture of the ventral visual pathway has begun to emerge, implicating the same regions as part of a recurrent and highly interactive occipitotemporal network linking early visual areas and anterior IT (Kravitz et al., 2013). Within this network, multiple routes of visual information processing exist, connecting at least six subcortical and cortical regions from different parts of the brain, which themselves support many different aspects of visual object processing.

2. Face-selective patches in the macaque temporal lobe

One obvious advantage of single-unit recording, both past and present, is the extremely fast temporal resolution relative to alternative recording and imaging techniques. This increased resolution provides us with the opportunity to distinguish

responses to systematic manipulation of same-category stimuli. Imaging methods such as fMRI, on the other hand, allow mapping of much larger regions of the brain without the invasiveness of single-unit recording. This ability to observe widespread responses throughout the brain has made it possible to describe neural activity in humans spanning multiple, functionally-distinct regions of the brain. By combining the advantages of both these techniques, Tsao and colleagues discovered a series of face-selective regions within the ventral visual pathway of the macaque, which are often observed by contrasting selective responses between object categories such as faces and non-face objects (Tsao et al., 2003). Further, by targeting single-unit recordings within the patches they observed, they found that these regions (or “patches”- see Figure 1.4) were composed of dense clusters of predominantly face-selective neurons (Tsao et al., 2006).

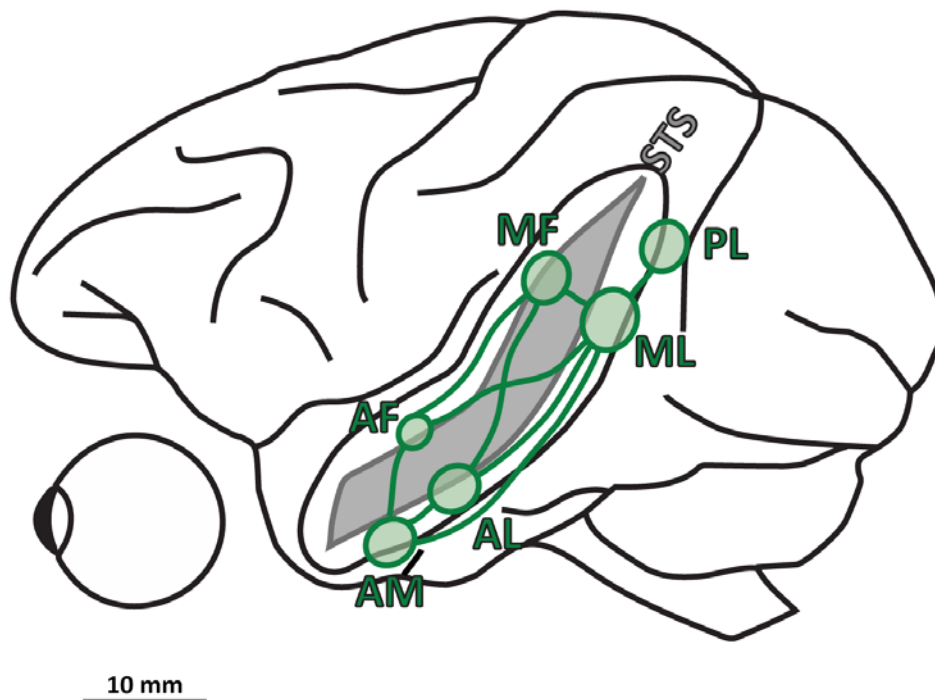


Figure 1-4 Schematic of the temporal lobe face patch network in macaque. A lateral view of the monkey eye and brain highlighting six face-selective regions (light green) and the connectivity between them (dark green). The superior temporal sulcus (STS) is partially spread, in order to highlight the position of face selective regions within it. Abbreviations: PL, posterior face patch; ML, middle face patch; MF, middle fundus face patch; AL, anterior face patch; AF, anterior fundus face patch; AM, ventral anterior face patch

3. Electrophysiological evidence for norm-based processing

Several recent studies have examined the degree to which multi-dimensional face space can be considered a useful framework for understanding how individual cells represent faces. As mentioned previously, the prediction of norm-based encoding would be that faces are represented in terms of their relationship with the average for a given population. Such a norm-based model would predict an uneven distribution of faces within a hypothetical face-space (a densely populated center consisting of more common faces). In this scenario one possible prediction might be that tuning for a given face feature would follow a linear, monotonic curve as each feature moves either towards or away from the mean. This is in perfect agreement with the findings of Leopold et al., (2006). They used a face-morphing technique to show that responses of some face cells increased linearly as faces are shifted away from the average of the entire stimulus set. This reinforces the concept of a norm-based face space established previously by studies of face identity aftereffects.

More recent recordings by Freiwald and colleagues (2009) used fMRI-targeted recording in one of six temporal lobe face patches that have been identified (face patch MF- posterior to those in Leopold et al.), and found neurons to be sensitive to the manipulation of particular face feature dimensions within a large cartoon face space. In this study cartoon faces were used in order to allow the independent manipulation of different structural parameters of the face such as the position of the eyes, nose, and mouth. They also found neurons that exhibited monotonic tuning functions, though their response minima or maxima less commonly occurred near the average face. Together, these findings make it clear that neurons in

the inferotemporal cortex systematically code individuating aspects of face identity, with more extreme features eliciting more extreme responses.

4. IT plasticity

The precise nature of physical changes in the brain following different types of visual learning is unclear. What changes occur when we learn a new face, or when we improve our capacity to discriminate individuals from new race following immersion into a different population? Given the accumulated evidence for the role of face-modules in the processing of face stimuli, and since we retain the ability to learn novel faces throughout the majority of our lifetimes, it seems reasonable to anticipate that visual learning related to faces would take place in face patches, with individual neurons exhibiting plasticity not only during development but also in adulthood.

While it has been argued that learning-induced plasticity could occur in a distributed fashion throughout the visual system (Sigman et al., 2005; Hochstein & Ahissar, 2002), the majority of studies investigating the development of expertise with visual objects have pointed to changes occurring in high-level visual regions (for review see: Gilbert et al., 2001). Some such studies have used arbitrary and parametric stimuli in order to isolate and manipulate particular feature dimensions (e.g. Sigala & Logothetis, 2002). However, few have looked at the learning with respect to realistic faces over long time scales, and as a result not much is known about learning-induced changes among face-selective cells.

How might face-specific learning be expressed at the single-neuron level? Theoretically, we can predict several types of changes that could occur following visual object learning (Op de Beeck & Baker, 2010a). For one, learning might

modulate responses the optimal stimulus (that which evokes the greatest response) of the neuron. For example, experience with certain stimuli might lead neurons that were naively selective for other stimuli to shift their optimal stimulus towards the trained stimulus. Alternatively, learning might affect the entire tuning function, modulating responses not only to an optimal stimulus but to a wider range of stimuli that bear some semblance to it. For example, training within a given stimulus dimension may lead to sharpening of tuning along that dimension, but a broadening of tuning for a different, less relevant dimension. Finally, learning might be expressed at the level of the entire population, slightly altering the responses of many neurons in a manner that does not correspond to either of the above alternatives.

Existing evidence from single-unit studies of IT plasticity has led to conflicting interpretations regarding the effects of learning on the functional properties in high-level areas of the ventral visual stream (e.g. Op de Beeck & Baker, 2010a, Gauthier et al., 2010, Op de Beeck & Baker, 2010b). To begin with, the relationship between changes to selectivity for experienced objects and their role in the training task is somewhat murky. Whereas, some studies have found a strong increase in selectivity for trained, but not untrained objects (Logothetis et al., 1995; Kobatake et al., 1998), others found selectivity increases to be specific only for relevant stimulus dimensions (Sigala and Logothetis, 2002). Yet other studies have provided evidence suggesting learning might also modify the optimal stimuli of affected neurons rather than the shape of the tuning function. For example, Baker et al., (2002) trained monkeys using stimuli which consisted of batons that were made up of a combination of parts, and found greater responses to the specific combination

of parts, than to the presence of the parts in isolation. Such nonlinear interaction between tuning for individual parts may reflect the generation of new dimensions of tuning.

5. Experimental limitations to studying plasticity in IT

Learning-induced changes in visual perception are likely to transpire over periods of days or weeks. Although “instantaneous” learning effects have been reported, whereby explicit training induces neurons to change their response within a single recording session (see Tovee, 1996), the majority of changes that affect perceptual systems are likely to take place over a much greater period of time. If the time course of learning is much longer than the maximal duration that neurons can be monitored, such learning effects cannot be studied directly as changes in the response patterns of individual neurons.

For that reason, studies examining the effects of learning on individual neurons in the temporal lobe have typically relied on comparisons between neural responses to learned objects and an unlearned baseline, recorded either previous to learning or in other naïve animals (Miyashita et al., 1988; Logothetis et al., 1995; Baker et al., 2002). Similarly, several forms of in-vivo neuronal changes that have been observed, have occurred over periods of days or weeks (Merzenich et al., 1984, Shatz & Stryker, 1978). Thus much of what we know about the plasticity of neural responses in the temporal lobe is inferred based on statistical comparison between responses to different groups of stimuli with distinct training histories. Understanding how changes in individual neurons are related to the learning of new

faces will require longitudinal studies that track single-unit responses for days, weeks, or perhaps longer.

In an effort to provide such a tool, our lab has developed a novel technique for longitudinal recording of single-units via a MRI-compatible microdrive (Elnaiem et al., 2013, Jones et al., 2013, McMahon et al., 2013, Leopold et al., 2013). In addition to increased recording duration, MR-localization enables precise electrode placement within functionally-defined brains regions (Ohayon et al., 2012). This emerging technology provides a new window through which to study the acquisition and refinement of neural selectivity associated with the development of visual expertise in the brain of nonhuman primates. Additional details regarding this technique are described in the following section.

Chapter 2: Methods – Microwires for long-term recordings

I. Introduction

The work in this volume relies heavily on the considerable effort expended by the members of my lab aimed at developing a microwire bundle array that permits the stable isolation, and hence longitudinal monitoring, of individual cells. In the present work, we exploit two unique features of this longitudinal recording approach, which are each described in the Introduction of this chapter.

1. Tracking changes in response selectivity of individual cells over weeks and months

Measuring responses to predetermined stimuli during brief encounters with individual neurons has taught us much about the behavior of neurons within a relatively short portion of their entire lifespan. While much progress towards understanding the brain has been made using this approach, it has provided little information as to the behavior neurons over periods exceeding several hours -- a time frame over which several experience-dependent behavioral changes are thought to take place. For example, studies have demonstrated that behavioral improvements in perceptual expertise occur are only observable after multiple days of training (Karni and Sagi, 1993; Schoups et al., 2001). Likewise, critical periods in visual system development have been shown to span several weeks of an animal's early life (Hubel and Wiesel, 1970). Changes in cortical organization following injury during this period also occur over a timespan of several weeks (Chang and Merzenich, 2003).

Due to limitations on recording duration, studies attempting to investigate such plasticity have typically employed one of two designs; sampling separate

neurons before and after training (Sakai & Miyashita, 1991) and comparing between responses to trained and untrained stimuli after training has taken place (Tanaka, 1996; Logothetis & Pauls, 1995). Current recording techniques are very limited in their ability to follow the activity of single neurons in order to observe changes firsthand, rather than by inference from sparsely sampled time points. By expanding the duration of electrophysiological recordings to span greater time scales, we will be able to observe changes that accompany behavioral training, firsthand. The first main advantage of this technique: the observation of responses over weeks and months.

2. Using stable isolation to accumulate repeated responses over extended periods

Extended observation of single units offers additional advantages as well. Theoretically, the way to completely describe the encoding of stimuli would be to present as close to every level of every dimension of stimulus space as possible. One attempt to overcome this has involved the use of a closed-circuit design, in which stimuli are generated based on responses observed during previous trials (Yamane et al., 2008). However, these approaches are still constrained in the number of stimuli they can present within a limited period of interaction with each individual neuron. So, by expanding the window across which we can observe neurons, we also increase the level of detail at which we can explore their behavior. Further, this enhanced resolution allows us to observe neurons that would otherwise be discarded due to lower overall firing rates. Therefore, the second main advantage afforded by this technique is the accumulation of repeated responses over several days, rather than several hours.

In order to overcome these limitations, we developed a chronically implanted electrode bundle that permits the observation of single neurons for periods exceeding several months (Elnaiem et al., 2013, Jones et al., 2013, McMahon et al., 2013, Leopold et al., 2013). A MRI-compatible microdrive permits fMRI-guided stereotaxic positioning of electrode relative to identified areas of selectivity. This emerging technology provides a new window through which to study the acquisition and refinement of neural selectivity associated with the acquisition of visual expertise in the brain of nonhuman primates.

II. Chronic recording with microelectrode bundles

In these experiments, we attempted to leverage the unique strengths afforded by this technique, which were described in the preceding section. What follows is a methodological summary of those steps leading up to those experiments, which are then described in detail in the section that follows.

1. Subjects

Three electrode bundles were chronically implanted in three rhesus macaques (2 females: *m1*, 5.6 kg and *m3*, 5 kg and 1 male: *m2*, 8.7 kg). The data presented in Chapter 3 was collected from neurons in M1 and M3, whereas the remaining Chapters contain data collected from M1 and M2.

2. Apparatus

A. Microwire electrode bundle

Electrodes bundles were designed based on previous reports (Porada et al., 2000; Bondar et al., 2009). Each bundle consisted of 64 platinum/iridium wires with polyimide insulation (12-um in diameter). This was encased in a polyimide-insulated

silica shaft, which was laser-cut to end 5 mm above the end of the microwire bundle. On the opposite end, those wires connected to four 16-pinn connectors (Omnetics, Minneapolis, MN). The mean impedance of the PtIr wires was 0.7 M Ω , measured before implantation. Each bundle was advanced through the brain within a 45-mm polyimide guide tube lined with PTFE (Micro-Lumen, Oldsmar, FL).

B. Chronic microdrive

The microdrive consisted of 3 horizontal Ultem disks, with three vertical G10 resin rods and 1 vertical PEEK screw drive held between the outermost plates (see Figure 2.1). By clamping the electrode bundle to the middle disk, the position of the electrode could be manipulated by via rotation of the screw drive. This permitted 7.5-mm of total travel distance, once implanted. This was fastened above a cylindrical recording chamber with a 0.89-mm hole, through which the electrode shaft was passed before filling with silicone sealant. Both chamber and microdrive were covered by a protective cap, which was removed only to when advancing the electrode bundle.

3. MRI

Due to the MRI-compatibility of the electrode bundle and microdrive, functional scans of blood oxygenation level-dependent (BOLD) signals were able to be collected on 4.7T vertical scanner (Bruker BioSpin, Rheinstetten, Germany). Anatomical scans were collected separately and co-registered to functional scans.

T1-weighted images were acquired before surgery using a 1.5T General Electric Signa scanner. Functional scans were then obtained pre-surgically on a vertical 4.7T-scanner (Bruker, Rheinstetten, Germany).




Figure 2-1 Chronic recording chamber permitted periodic advancement via rotation of a screw, but was otherwise covered by a protective cap. The microwire bundle was advanced into the brain through a transdural polyimide guide tube, which was embedded within a silicone well, thus eliminating the need for routine cleaning to prevent infection.

MRI scans were performed 1 week post-surgery to localize the electrode position and identify potential signs of infection in the brain (see Figure 2.2).

Microwire bundles were placed within fMRI-defined patches, mapped by contrasting responses to different visual stimulus categories. For example, contrasting faces with objects is a standard and robust method that permits the localization face patches.

4. Surgical implantation

Four micro-electrode bundles were implanted in three monkeys, 1 unilaterally in *m1* and *m3*, and 2 bilaterally in *m2*. T1-weighted images were acquired using a 1.5T General Electric Signa scanner. Functional scans were then obtained pre-surgically on a vertical 4.7T-scanner (Bruker, Billerica, MA) before being fitted with a MR-compatible head post. Areas of functional selectivity were identified and targeted during electrode implantation.




Figure 2-2 Functional MRI patterns, overlaying anatomical scans, revealing face-selective regions (warm colors) in MF (lateral view, left patch) and AF (lateral view, right patch; sagittal view) in the STS of monkey 1. Pre-surgical scanning permitted precise placement of the chronic recording bundle (green arrows) within face-selective regions.

Monkeys were sedated with ketamine, diazepam, and glycopyrrolate on the morning preceding surgery. Vital signs were monitored continuously while maintaining anesthesia under isoflurane. After drilling a craniotomy in the skull, the dura was pierced and the guide tube lowered into the cortex via stereotactic micromanipulator (Kopf Instruments, Tujunga, CA) to the depth determined appropriate following pre-surgical MR-imaging. The guide tube was then glued to the recording chamber, securing it in place. The exposed craniotomy was filled with bone wax, and the area surrounding the recording chamber was secured to the surrounding skull with acrylic cement and reinforced with ceramic bone screws. Each chamber was uniquely contoured to match the shape of the underlying skull, before being secured using dental acrylic.

The guide tube was cut approximately 2 mm below the top of the chamber, leaving room to be covered with silicone sealant. After attaching the electrode bundle

to the microdrive, it was then threaded through the guide tube using the micromanipulator. The microdrive was then secured to the chamber using nylon screws, before being used to advance the electrode 1 mm past the end of the guide tube. This left the electrode tips within the white matter of the STS. The 2mm well above the guide tube and electrode bundle was then filled with silicone (Kwik-Cast, World Precision Instruments, Sarasota, FL), which acted as a barrier to prevent infection in the brain. The electrode bundle wires extended from the silicone well before running through a hole in the chamber wall and terminating in a 4 rows of connector blocks secured to the acrylic base of implant.

Table 2-1 Summary of study design

	M1	M2	M3
Neurons tested for stable selectivity	59		85
Neurons tested with face identity stimuli	19	5	
Neurons tested with caricature stimuli	17	9	
Neurons tested with 3-d view stimuli	30		
Neurons tested with face-training stimuli	23		

5. Assessing stability of single-units over consecutive recording sessions

Single-units isolated on the same channel on consecutive days were considered the same neuron if the average spike waveform, patterns in spiking activity, and stimulus selectivity were maintained between days.

A. Spike waveforms

During spike-sorting, each time the voltage signal exceeds a certain threshold, based on the signal-to-noise ratio, 40 ms before and after each event is recorded, then averaged together to acquire an average waveform for the spikes recorded on each day. Information contained in this spike shape was quite unique and could be included in the features used to establish the stability of our recordings.

B. Stability index

If spike-shape and spiking statistics were sufficiently reliable across recording days, it was considered as likely being the same neuron held across sessions, based in part on previous work from the laboratory. Stably isolated neurons were then evaluated for the stability in their visual responses across sessions. For this, we used a stability index that compared the reliability of stimulus-selective visual responses across sessions to the same measure within session, according to the formula:

$$\textit{stability index} = (\mathbf{r}_{\text{across}} - \mathbf{r}_{\text{shuffle}}) / (\mathbf{r}_{\text{within}} - \mathbf{r}_{\text{shuffle}})$$

In this index, \mathbf{r} represents the correlation coefficient between vectors consisting of the responses to each stimuli within a given day. As such, $\mathbf{r}_{\text{across}}$ represents the correlation between response-vectors on consecutive days, $\mathbf{r}_{\text{within}}$ represents the correlation between even and odd trials on the same day, and $\mathbf{r}_{\text{shuffle}}$ represents the same correlation after shuffling randomly shuffling the stimulus identities.

6. Analysis

Data were analyzed using custom software written in Matlab (Mathworks Inc,

Natick, MA). Unique peri-stimulus time histograms were constructed for each face by aligning the spikes elicited by that face for each trial to the stimulus onset. Spike counts in a time interval from 100 to 400 milliseconds post-stimulus-onset were summed and used as basis for the generation of tuning curves and the analysis of population responses. The population analysis was restricted to neurons that responded with at least 5 spikes/s above baseline to at least one of the faces.

Due to the enhanced resolution afforded by accumulating many responses over consecutive recording sessions (see Figure 2.3), many neurons could be included in our analysis that would otherwise be discarded due to lower overall firing rates. This allowed us to examine the response properties of low-firing-rate neurons, which might otherwise have been overlooked.




Figure 2-3 Responses of a single neuron, accumulated over different periods in time. Each panel contains a spike-density function throughout a 300 msec presentation of a single stimulus (dark bar), shaded regions represent standard error. The response in each panel is derived from an increasing number of repetitions, from left to right.

III. Measuring responses

The following section describes the details regarding the measurement of responses to the stimuli presented in this study, including a description of the stimuli and how they were generated and the details of day-to-day procedures involved in recording of single-unit responses to those stimuli.

1. Stimuli

A. Face and object images

Human face images were selected from the FEI face database (Thomaz & Giraldi, 2010). Monkey faces and bodies were provided by Dr. Olga Dal Monte. All other stimulus categories were selected from the results of internet searches.

For the experiments described in this volume, four separate stimulus sets were used. The first stimulus set was intended to identify visually responsive neurons and assess their selectivity for semantic categories, including faces (see Figure 7.1). The second set was intended to assess the selectivity of visual neurons for faces of different identities, and so included 24 different faces (both human and macaque) as well as their morphed intermediaries (see Figure 7.2 and Figure 7.4). The third and final set included a large array of human faces (200 total), as well as a variety of 2-dimensional manipulations to the most preferred face from that set (size and orientation – see Figure 7.5 and Figure 7.6). Visual stimuli were presented on a black background. Face photographs were cropped to only include the face, head, and neck, and then adjusted to a common size (GNU Image Manipulation Program, www.gimp.org).

B. Face morphing

Manipulation of identities were performed using a picture morphing software program (Abrosoft Fantamorph), which allowed us to create morph sequences.

In this volume, “morph” refers to a sequence of intermediate images, which when put together with the original images, represent the transition from one image to the other. The specific changes that occur during the morphing process are defined by

points placed on the major, shared elements of both faces (e.g. eyes, nose, and mouth; See Figure 2.4).





Figure 2-4 Fiduciary point example. Intermediary images, or “morphs” were drawn from a linear continuum combining the geometric features of each original faces. A fixed number of fiduciary points (stars) were places on the shared landmarks of each original face, and a vector representing the transition between each pair (green dotted line) was established. The area surrounding each point was transitioned along the vector for each pair of points, resulting in “morphed” face blends of the originals (right).

C. Creating a “face-space” centered on the average face

Using either human or monkey faces described above, we created 3 (2 human, 1 macaque) multidimensional “face spaces”, within which each point represented one of 12 individual faces, with the average for each species in the center (See Figure 2.5).

We explored several types of trajectories within this proposed arrangement; we focused our sampling on the trajectory connecting each veridical (100% identity) face and the respective norm for that species. This trajectory will be referred to as “radial”. In addition, we presented faces from morph sequences between 2 veridical faces. This will be referred to as “tangential”.

Later in this study, caricature faces were generated in a similar fashion, but in this instance, the average texture of each species was applied to all faces within that



Original faces



Averages faces





Figure 2-5 Face averages. Initial face stimuli (left column) were selected semi-randomly from the FEI face database (humans) or provided by Dr. Olga Dal Monte (monkeys). Face averages for each species (right column) were created by superimposing the faces of each group, and then averaging the position and color/tone of a standardized number of fiduciary points on each face. This resulted in a single face that contained the average configuration and texture of all of the faces within a group.

group. In this way, we meant to control for responses to changes in the texture of the face, rather than face identity. With these stimuli, morph sequences were extrapolated both beyond the average face and the veridical face, thereby creating “anti-faces” and “caricatures” (See Figure 1.2C), respectively.

2. Recording Procedure

At the beginning of each recording session, the monkey’s head was restrained and a single 16 channel pre-amplifier (Tucker-Davis Technologies, Alachua, FL) was plugged into each row of the connector block. Recording took place within an RF-shielded room. Electrophysiological signals were digitized using a 128-channel RS4 Bioamp Processor. Signals were broadband filtered between .5 Hz and 5 kHz before spike-sorting. One channel of the microwire bundle was designated as ground, and signals were referenced relative to it.

3. Stimulus presentation

Stimuli were presented on a ViewSonic 18” CRT display placed 91 cm away from the monkey. Stimuli in this volume were presented within a 7 degree window, unless specified otherwise. A single trial refers to a 900 ms period in which stimuli are presented for 300 ms and both preceded and followed by 300 ms periods in which the screen is black. Four stimuli were presented in succession on each trial. A typical recording day included 25-30 trials per stimulus.

Stimulus presentation, monitoring of eye position, and triggering of reward delivery were controlled by custom QNX software from Dr. David Sheinberg. Stimuli were presented by a slave computer running Psychtoolbox (Mathworks, Natick, MA). Eye position was observed using infrared cameras and Eyelink software (SR

Research, Kanata, Canada) to compute horizontal and vertical coordinates of gaze. During “passive” presentation of stimuli, a series of 4 successive trials was initiated following fixation by the monkey, within a 3° window on and around a 0.3° white dot in the center of the visual field. Stimuli were then presented in behind and around the fixation spot. The monkey was rewarded with a drop of water following completion of four consecutive trials.

4. Spike-sorting

Individual spikes were isolated using custom Matlab (Mathworks, Natick, MA) scripts, including WaveClus (Quiroga et al., 2004). This stochastic algorithm implements an automatic threshold based on the standard deviation of the noise in order to detect spikes. Once detected, this algorithm uses the wavelet transform of those spikes to isolate their component features, thereby reducing the dimensionality of the data in a manner that out-performs principal components analysis. Clustering is accomplished using superparamagnetic clustering; a stochastic method for grouping spikes with similar shapes and assigning them to the same unit.

In the next section, we will review the first group of experimental results- specifically, those related to the stability of stimulus selectivity exhibited by AF neurons, when comparing responses over consecutive recording sessions.

Chapter 3: Stability of stimulus selectivity in AF neurons

I. Introduction

Human infants will automatically orient their visual attention towards conspecifics in their immediate environment. Throughout development, therefore, we are exposed to faces more frequently than other stimuli. To some extent, the primate brain could be described as “pre-programmed” to practice and acquire expertise in the processing of faces. Our ability to recognize the identity of those faces is thought to be extracted from “invariant” aspects of facial structure (e.g. distance between the eyes or aspect ratio of the face) that remain relatively stable over a wide variety of changes in face expressions and viewing conditions (Haxby et al., 2000). Since healthy adults are able to continue recognizing certain faces effectively for most of their lifetime, it remains possible that the same set of neurons that once responded to these invariant aspects of certain faces, such as your grandmother’s, as an infant, may still be responsible for encoding those aspects of her image today. A scenario such as this might favor extremely stable encoding of features by neurons encoding complex visual objects, such as faces, within the brain.

On the other hand, even so-called invariant aspects of the face tend to change dramatically throughout an individual’s lifetime (see Figure 3.1). Further, while our ability to learn new face identities is greatest during development, our expertise at interpreting subtle facial changes continues to improve throughout our lifetime. As such, we are able to continuously improve our ability to send and receive communication signals using the face, while also maintaining a robust representation of the identity of each face under the wide variety of viewing conditions in which we

typically encounter them. In order to support this, neurons encoding face identities may need to be highly plastic in order to continue successfully discriminating between members of an ever-expanding library of face expressions and identities. So, how then do face-selective neurons behave over time, in order to support these demands for both robust stability and adaptive plasticity?




Figure 3-1 Diamonds (but, not faces) are forever. An example of the dramatic and somewhat unpredictable changes that may occur in an individual face throughout the course of a human lifetime.

As discussed in Chapter 1, we know that face cells in the adult brain are organized into distinct functional regions or “patches”. The following chapter will discuss competing theories regarding the encoding of neurons in these patches, and how this might contribute to our specialized visual abilities. In this chapter, however, we will begin to explore how this sophisticated system is developed and maintained, and how the behavior of individual cells play a role in this development.

Our understanding of the role of how face- and object-selective neurons change over time remains extremely limited. Studies have reported face-selective cells in monkeys as young as six weeks old (Rodman et al., 1993), and that these neurons exhibit “adult-like” properties by 6 months of age. But, it remains unclear how the response properties of such neurons change over time and what factors are

involved in influencing their development. It has been suggested that visual experience plays a role, particularly that involving face stimuli (Logothetis & Scheinberg, 1996), but observations of these neurons remain limited to single, isolated recording sessions spanning several hours at a time.

Some studies have succeeded in exploring the anatomical connectivity of cortical visual neurons of the adult macaque in a longitudinal fashion, and found strong evidence for changes in the synaptic connectivity occurring over a period of several weeks. In this study, *in-vivo* 2-photon microscopy of genetically-labeled axons was used to monitor the branching pattern of axonal connections, in order to infer changes in the weight of connectivity between neighboring neurons. Results suggested that large-scale branching patterns were stable, but that synaptic boutons demonstrate constant transition from week to week (Stettler et al., 2006). Rather than representing the result of experience, this evidence has been taken to suggest that V1 neurons are involved in a process of continuous transition in their connections with their proximal neighbors.

Another study, recording from IT neurons, and revealed experience-dependent changes to the stimulus selectivity within just an hour of their initial exposure to the statistics of a particular stimulus set (Li & DiCarlo, 2008). In this study, neurons exhibited tolerance in their responses to changes in the position of a limited subset of “preferred” stimuli. Similar tolerance to stimulus position could be elicited by previously non-tolerant stimuli, following their temporal association with preferred stimuli. This could be interpreted as suggesting that the experience-dependent improvements in view-invariant identification of visual objects, as described at the

beginning of this chapter, may be a consequence of changes in the selectivity of object-selective neurons that occur following experience with a particular object. Perhaps most important, regarding the results in the section that follows, is that the relationship between experience and changes in the selectivity of IT neurons was significant within just 1 hour of exposure.

Similarly, another study exploring plasticity in IT face-cells, found changes in their responses to stimuli over several hours of exposure to a given stimulus set (Rolls et al., 1989). In this study, authors found that 6 out of 22 neurons exhibited response degradation as a function of familiarity (number of presentations), whereas this same relationship was much less prevalent when the monkey had been previously exposed to those particular stimuli. Finally, responses to familiar stimuli were altered following the addition of novel stimuli to the stimulus set, suggesting these representations undergo continual, immediate change as a consequence of the likelihood of observing certain stimulus properties.

Taken together, these studies might lead to the prediction that object-selective IT neurons should exhibit clear plasticity in their response properties within several hours of exposure to a particular stimulus set. In order to explore the stability of response properties of face-cells in AF, we observed populations of visual neurons throughout repeated presentation of both face and non-face stimuli over a period far greater than that of previous studies. By making use of the uniquely stable recording of single-units made possible by our recently developed chronic, longitudinal recording technique, we were able to monitor single-unit responses throughout

repeated presentations of both face and non-face object stimuli for an initial period of over one month in the life of 2 macaques.

In short, we found the response properties of single object-selective visual neurons to be extremely stable from day-to-day -- exhibiting no evidence for any systematic changes that might correspond to stimulus familiarity, even following more than 1,000 presentations of each stimuli.

II. Results

Given the substantial investment of time and energy they required for the longitudinal recordings presented in this volume, the overall number of neurons recorded is less than that of conventional methods. However, since each neuron was typically recorded from for several weeks or more, this is shortcoming is counter-balanced by an extremely novel and intricate data set, permitting us to pose new questions that were not possible to answer using conventional techniques.

1. Response significance and response latency

Neurons were defined as visually responsive to any one stimuli if the SDF of the response between 100 and 300 msec post-stimulus onset was significantly different than the response between 200 and 0 seconds pre-stimulus presentation (one-sampled t test). Measured in this way, 55 out of 66 neurons in Monkey 1 and 34 out of 35 neurons in Monkey 2 were defined as visually responsive.

The response latency to a particular stimulus was defined as the first sample after the z-score of the spike-density function exceeds 1.96, which corresponds to a pre-stimulus spike rate confidence interval of approximately 95%. Mean response latency averaged across all of 206 face-identity stimuli (including objects). The

distribution of response latencies across all neurons was similar in the two monkeys (see Figure 3.2; M1: 169.93 ± 6.71 msec; M2: 171.38 ± 10.96 msec) and did not exhibit a significant difference between animals (un-paired t-test, $p = 0.87$).

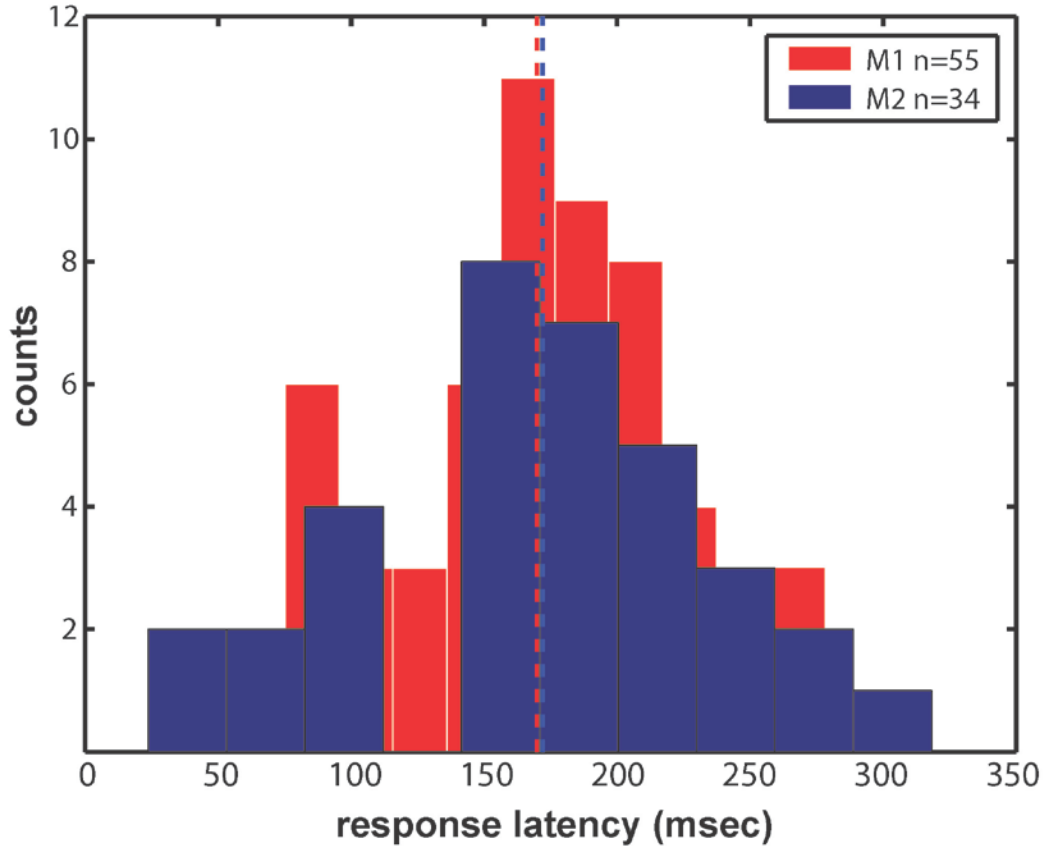


Figure 3-2 Average post-stimulus response latencies for neurons exhibiting significant responses in monkey1 and monkey2. There was no significant different in average response latencies for each monkey (one-sampled t test, $p = .87$). The mean latency for each group of neurons was 169.91 msec in monkey1 (red, dotted line) and 171.97 msec in monkey 2 (blue, dotted line).

2. Waveform stability

To begin characterizing the response stability of visual neurons in AF, the shape of the average waveform of each neuron was compared from day-to-day. As the examples in Figure 3.3 portray, AF neurons recorded on different channels often demonstrated remarkable stability in their wave-form shape between successive days. This stability, in spike amplitude and timing, indicated that the isolation of these neurons was not changing substantially. This suggests that whatever changes to the

synaptic inputs do occur, they do not result in observable changes in the spike-shape of these neurons over several weeks of daily recordings.

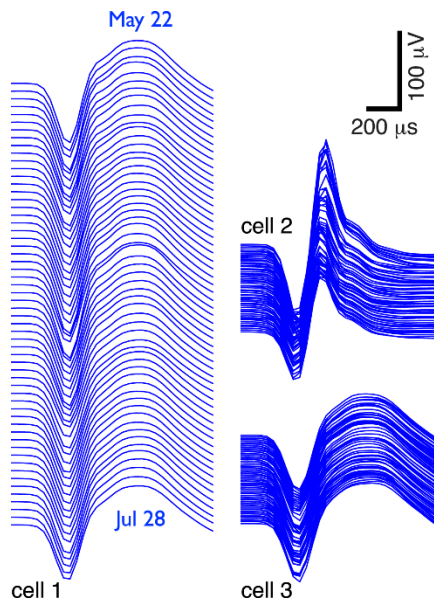


Figure 3-3 Average spike waveforms of three AF neurons throughout the course of 68 daily recording sessions. Each stack represents the responses of a single neuron over consecutive recording spanning over two months.

3. Selectivity patterns over time

In addition to exhibiting stable waveforms, this population also demonstrated surprising constancy in their selectivity for visual stimuli. Figure 3.4 portrays the response of a single AF neuron to both face and non-face stimuli, over consecutive daily recording sessions. As shown in this example, these neurons exhibited a variety of responses to both face and non-face stimuli. However, as was common across the population, the overall pattern of stimulus selectivity, and dynamics of responses to individual stimuli (both timing and amplitude of the response) was virtually unchanging throughout several weeks of daily observation.

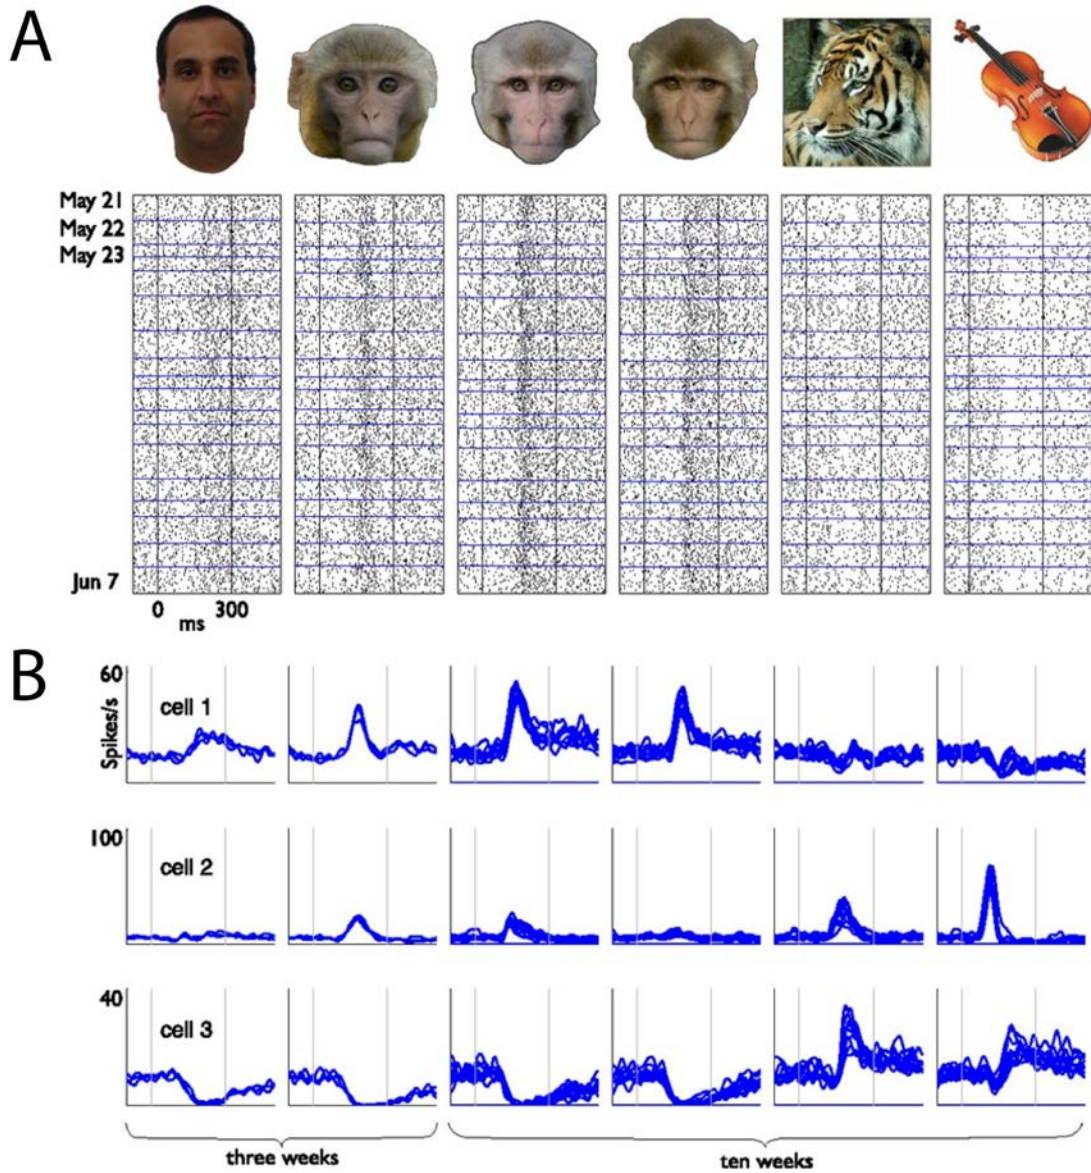


Figure 3-4 (a) Stimuli and corresponding responses from an example neuron. Each panel contains the aggregated responses over consecutive recording sessions (indicated by horizontal blue lines). (b) Average weekly responses of three example neurons. Rows: individual neurons, columns: stimuli corresponding with those in A. Responses are presented as spike-density histograms, which are superimposed across three weeks (left two columns) and 10 weeks (right four columns).

In order to explore the behavior of these neurons over a more substantial portion of the animal's lifespan, we also compared the selectivity of AF cells over a period exceeding one year. We found a similarly consistent pattern of diverse, but stable response properties for a variety of naturalistic images. As indicated by the example in Figure 3.5, AF neurons observed over this time, continued to exhibit

stable response properties, in both waveform shape (Figure 3.5A) and stimulus selectivity (Figure 3.5B).

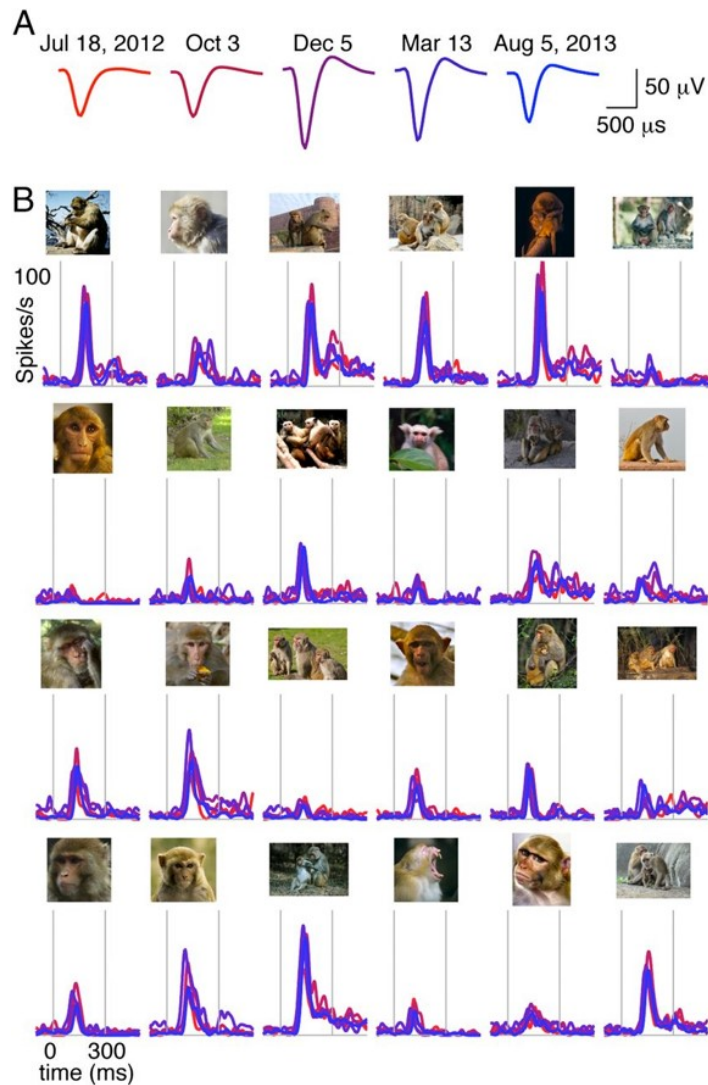


Figure 3-5 Longitudinal visual selectivity over one year of recording. (a) Average spike waveform from daily recordings in a single neuron spanning over one year (7/18/2012-8/18/2013). (b) Visual responses of the same neuron to 24 stimuli consisting of the faces and bodies of new- and old-world monkeys. Daily response histograms are overlaid and color-coded corresponding to the date of each waveform in A.

4. Response consistency assessments

To further explore the stability in the response properties in this population of neurons, the correlation of average responses within the same day to all stimuli presented on odd vs. even trials (r_{within}) was compared to the same correlations between consecutive days (r_{across}) and between the same values after shuffling the

stimulus identity (r_{shuffle}). One might predict, that should the variability from day-to-day exceed that of variability within recording sessions, (which are generally assumed to be stable) the correlation within the same day (r_{within}) should noticeably exceed the correlation across days (r_{across}). As demonstrated by the examples in Figure 3.6, selectivity correlations assessed across days was only slightly less correlated than for the same selectivity patterns measured across days. (mean, 0.64 and 0.59, respectively; see Figure 3.7), whereas both correlations were significantly greater than the stimulus-shuffled control (mean $r = 0.006$; $P \ll 0.01$, two-sampled t test; Fig. 3.7 inset).

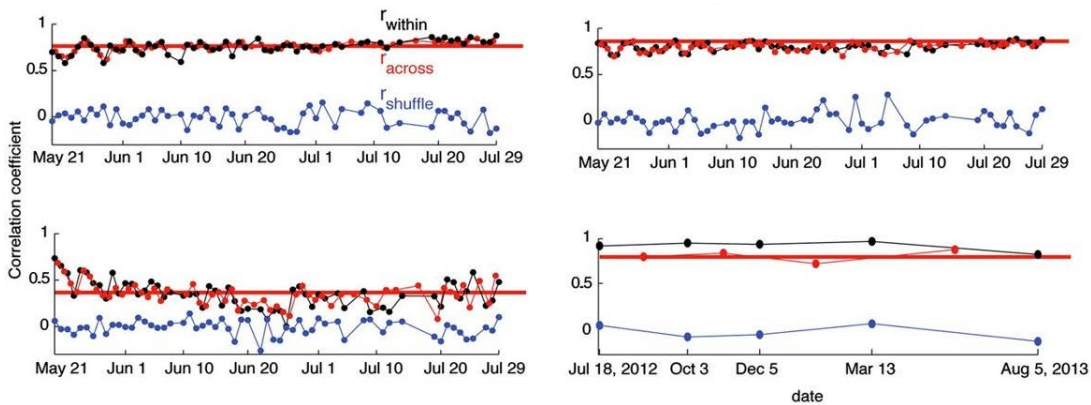


Figure 3-6 Correlation coefficients of split-halves analysis (odd vs. even trials) comparing visual responses recorded within a single day (black), between consecutive days (red), and after shuffling the stimulus identity (blue). Each panel corresponds to the responses of a single neuron. The thick horizontal red line in each indicates the r_{across} correlation between the first and last days of recording.

5. Channel-by-channel isolation

Due to the flexible, “paint-brush” design of the microwire array, it remained possible that the similar response properties of these neurons was a result of repeated recordings from proximal neurons on multiple channels. In order to investigate this possibility, we measured the within-cell consistency of responses by computing the correlation between responses to all stimuli on odd versus even trials (see Figure 3.8; median r , 0.94; interquartile range, 0.74–0.98; $n = 144$). We then repeated this

calculation for odd versus even trial of different neurons screened with the same stimuli (median $r = 0.10$; interquartile range, -0.09 – 0.37 ; $n = 1,561$ pairs). Although the correlation between different neurons was much less than that for the same neurons, the correlation was significantly greater than zero ($P \ll 0.01$, one-sampled t -test), affirming that neurons in face patch AF do exhibit a significant level of similarity in their selectivity for visual stimuli.

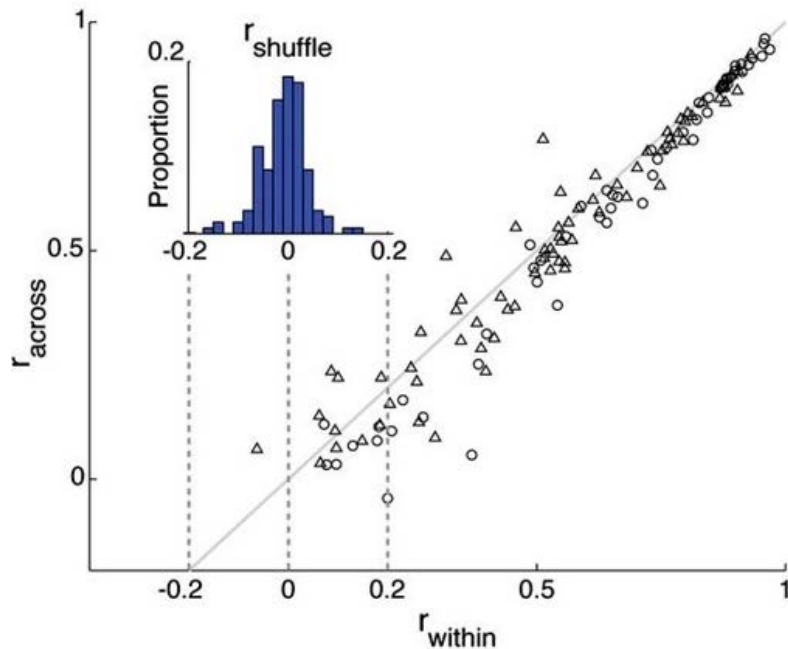


Figure 3-7 Population stability summary of 144 neurons (59 from m1, triangles; 85 from m3, circles) showing correlations computed between odd and even trials recorded on the same day (r_{within} , x axis) and on different days (r_{across} , y axis). (Inset) Distribution of $r_{shuffle}$ values for the same 144 neurons.

III. Discussion

Previous studies, which have observed selectivity patterns of IT neurons (TEa and TEav; Bondar et al., 2009) for over 2 weeks found similarly consistent responses in IT neurons. In the current study, we focused our recording within a more recently identified, fMRI-defined face patch -- AF. In addition to what this reveals regarding the stability of AF object-selective neurons throughout development, in particular, these results also confirm the possibility of observing single-neurons for periods

exceeding one year, using this technique. By isolating single-neurons in this region for a period as long as one year, these findings extend the timeframe of previous studies by an order of magnitude.

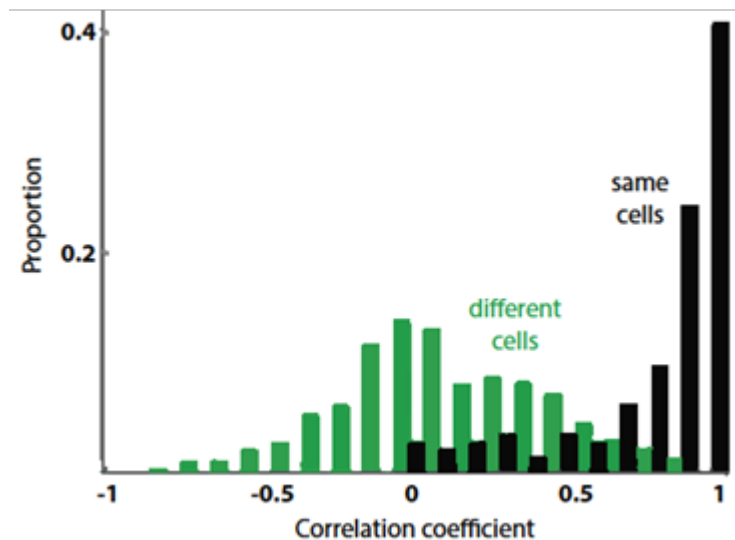


Figure 3-8 Population isolation summary comparing split-halves correlations for trials from the same neuron versus different neurons. Shown in black are coefficients for odd versus even trials from the same neuron (144 cells); and in green are odd versus even trials from all pair-wise combinations of different neurons for which responses to the same stimuli were available (1,561 pairs, eight stimulus sets).

These results challenge earlier findings, which have described changes in the response properties of IT neurons following repeated exposure with stimuli, presumably reflecting changes in familiarity (Li et al., 1993). The discrepancy between these previous findings and those described here may indicate a unique functional specialization within this particular region, relative to other components of the ventral visual pathway. Such specializations may be necessary to support consistent identification of visual objects between presentations, which are sometimes separated by intervals of months, years, or decades.

An analogous specialization in stability within distributed populations of hippocampal place cells have been revealed recently. In Mankin et al., (2012), authors

observed neuronal responses during repeated behaviors over intervals of from several hours to multiple days. It was found that hippocampal CA1 populations, in particular, differ in their responses to repeated experiences from CA3 neurons, which were highly reproducible over time. It could be possible that CA3 neurons are responsible for providing a stable memory code over time. Likewise, the findings presented in this sections leave open the possibility that AF neurons are similarly specialized to support consistent recognition of objects throughout an animal's lifetime.

Finally, these results may suggest that the stimulus-selective responses of AF neurons in the adult macaque are relatively fixed by this stage in development. Numerous examples of critical windows exist, which suggest that the brain is particularly adaptable to new sensory input during certain "critical periods" that occur throughout the course of development (e.g. Wiesel & Hubel, 1963). In some cases, if the appropriate stimulation is not received during this time, it can be impossible to develop certain functions later in life. It remains possible that the stability described in this chapter is a reflection of the maturational stage of the adult macaque. While popular opinion would suggest that face individuation abilities continue to develop very slowly, this is not well explored in the macaque. Perhaps similar investigations, at various stages of development, would improve our understanding of the role of age in the stability of these neurons.

Chapter 4: Representations of faces in AF

I. Introduction

This chapter will describe the results of an experimental investigation into the representation of faces, in particular, by single-neurons in the AF face patch. In this section, we utilize the unique opportunity afforded by longitudinal recording to accumulate a large number of responses to different stimuli over consecutive recording sessions. Two aspects of face representations seemed to exert particularly strong influences on the response properties of these neurons: face identity-level and size. In our exploration of the influence of face identity-level, we've extended the results of previous studies of seemingly norm-based encoding by face-selective neurons, by testing for the presence of these tuning properties over a much greater variety of intervals in time, and individual face identities. The latter half of this chapter will address the conspicuous sensitivity of face-selective neurons in AF to the size of a presented face, as compared to changes following other 3-dimensional manipulations.

II. Results

The results that follow are grouped into 2 main sections. First, those results concerning the influence of changes in face identity-level, and the dimensions of those changes the seemed to dictate responses of AF neurons most reliably. The second part of this results section concerns the influence of a variety of manipulations in the 3-dimensional orientation of certain face identities, which were identified in the first section as being “preferred” by the largest majority of neurons included in this study.

1. Identity

Background

In a previous study of face cells in anterior temporal cortex, Leopold and colleagues (2006) made use of a norm-based face caricaturization framework to explore the correspondence between perceptual models of face-processing and single-unit responses to face identity in this area. In this study, they found firing rate to rely on the magnitude of deviation from the average of all face stimuli. This suggests that these responses are sensitive to the similarity of a face to the rest of the population -- in other words, the distinctiveness of a face. This conceptualization fits well with evidence from psychophysical models of face identity perception, as well.

However, recent studies exploring the effects of short-term exposure on the development of the face-norm have found that the position of the face norm is modified by very recent exposure to particular faces, rather than over the long-term. Van Rensbergen and Op de Beeck (2014) found that while systematic manipulation of the average face within a set of 10 preceding trials resulted in shifting the norm toward the average, no influence was found for the order or recency among these trials; thus, no evidence that the last faces mattered more than the first. Results such as these have been taken to suggest that previously reported norm-based tuning in IT neurons may simply be a result of within-session adaptation to a stimulus set in which the majority of faces are similar to the average (Kahn & Aguirre, 2012). For this reason, understanding the influence of within-session adaptation on the identity-tuning of these neurons remains an important step in understanding the relationship

between perceptual models of face encoding and identity-selectivity in face-selective neurons.

In order to better clarify the role of norm-based encoding of faces in the AF face patch, we recorded responses to both human and monkey faces, morphed along radial and tangential axes of a norm-based face space (see Figure 1.3A & B). The sensitivity of our measurements of spiking responses was greatly enhanced by accumulating responses to 196,244 stimulus presentations from the same individual neurons over 76 days of recording. This dataset provided a unique opportunity to ask fundamentally new questions about stimulus responses. In this chapter, we use this method to investigate the contribution of within-session adaptation by only considering the first presentation of a stimulus for each session. In the next chapter, we use the longitudinal aspect of the recordings to examine changes in individual cells' response during periods of initial familiarization and intensive perceptual training.

Neural tuning for face identity

In order to acquire an initial estimate of the sensitivity of neurons within this particular face patch to changes in face identity, we began estimating the selectivity of AF face-cells for faces of different identities. Raster plots of the responses to 24 human and 12 monkey faces were compared. The responses to the apparently preferred face identities were often substantially greater than that for the least preferred identities (see Figure 4.1). A notable aspect of this data is that the facial features (e.g. mouth width or nose length and head shape or face size) of the preferred identities (green stars- Figure 4.1) was often very different. For example, by

considering the faces in the left-most column of Figure 4.1, we see very little in common between the faces of an adult male (top) and a developing young female (bottom). This result suggests that neurons in AF are highly sensitive to the identity of a given face, but that the selectivity for individual faces might be difficult to account for by a similar geometry of facial features alone.

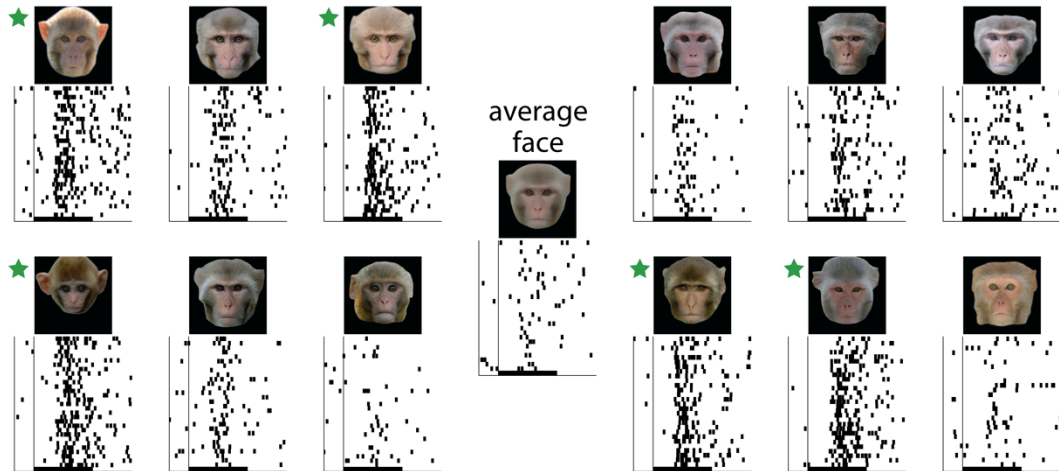


Figure 4-1 Responses of a single AF neuron to each of the 12 monkey face identities included in this study. Each panel corresponds to the raster in response to a 300-ms presentation of each face stimulus (black bar) beginning at time 0 (vertical gray line). The most strongly preferred faces are denoted by a green star.

Identity trajectory

As the next step toward estimating the most relevant dimensions of encoding of face-identity in AF neurons, we then explored how the identity-level, or distinctiveness, of a particular face identity would influence the responses of AF face-cells. As mentioned previously, recordings in other face patches (namely, AM – Leopold et al., 2006) have found distinctiveness to be an important determinant in the responses of face-selective neurons. In order to explore the extent to which this might be the case in AF, or whether this was a unique property of neurons in AM, we compared responses, in the form of spike-density-functions, across six levels of increasing distinctiveness between face and the average (i.e. an “identity trajectory”).

The example in Figure 4.2 shows a typical result, for two different faces which have been modulated at six identity levels from 0% (average face) through to 100% (full identity face). This neuron, like several that were observed, exhibited gradual increases in the neural response as a function of identity level, with the highest identities triggering the strongest responses and the average face (0% identity level, left-most column) triggering the lowest response. In this example, in particular, it is especially difficult to identify facial features shared between these two face identities (A and B), whereas they both evoke similar responses at the greatest distance from the mean.

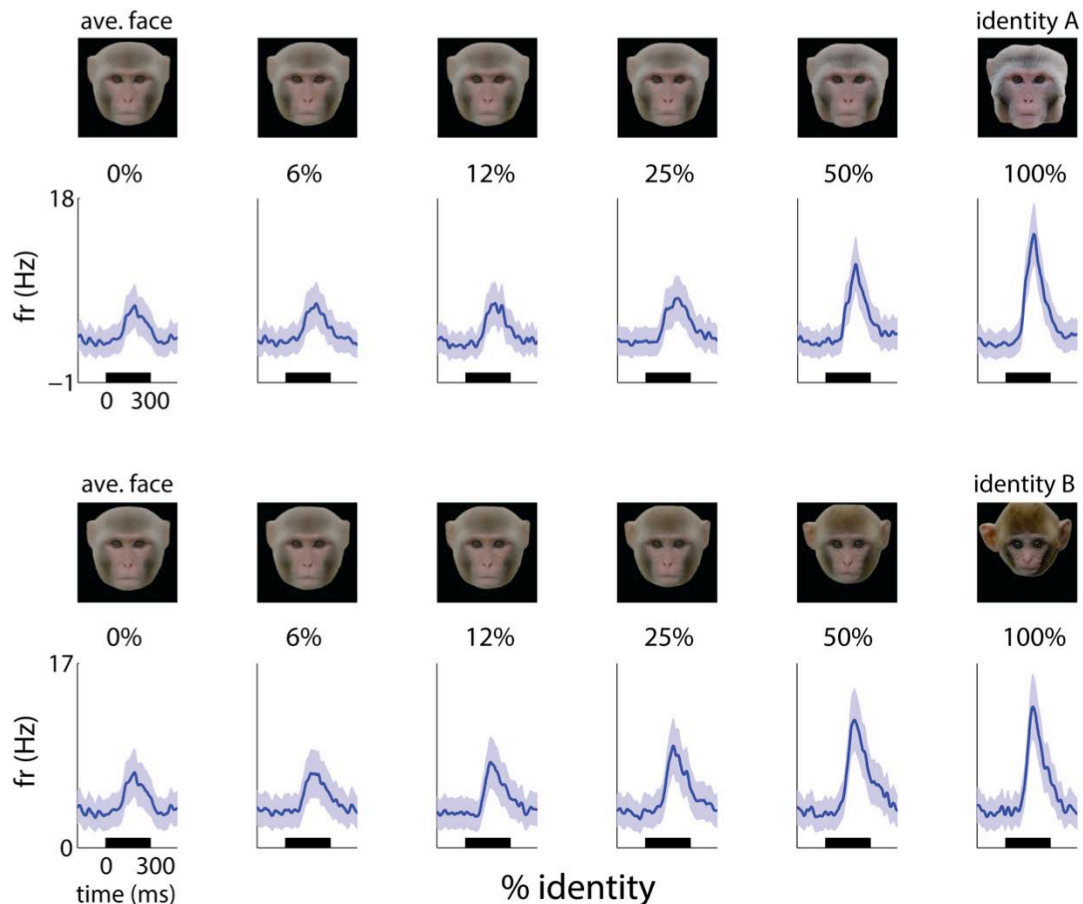


Figure 4-2 Each panel shows the responses of a single AF neuron to faces along the identity trajectory. Each panel contains the smoothed peristimulus time histogram in response to a 300-ms presentation (black bar) of the face stimulus shown above. Each row represents a radial axis in face space, from the average (0% - left column) through to the full identity face (100% - right column). Light blue: s.e.m.

This can be compared across multiple identity trajectories in another example neuron shown in Figure 4.3, wherein the responses to each stimuli along multiple different identity trajectories are represented by color-coded spheres. AF neurons exhibited variable response-selectivity for faces of different identities. However, responses were typically greater for full (100%) identity faces than for the average face (shown in the center), regardless of the identity they belonged to. This suggests that although AF face-neurons seem to prefer faces belonging to certain identities over others, overall neural responses seemed to be more consistently modulated by the identity level (i.e. distinctiveness) of a particular face than the unique geometry of its individual features.

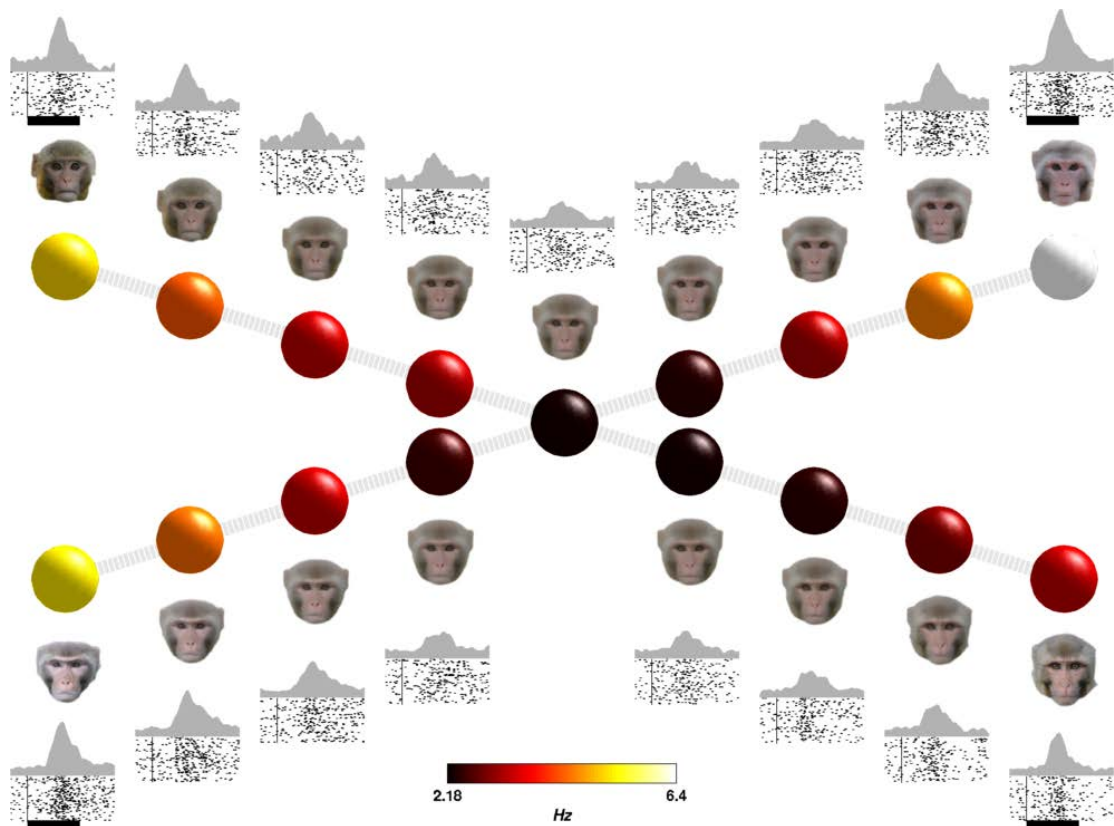


Figure 4-3 Rasters and peristimulus time histograms show responses to 300-ms presentation (dark bar) of four different monkey face trajectories. Each sphere corresponds to the response following the presentation (black bar) of the stimulus shown above or below it. The color of each sphere indicates the mean spike rate of the response within the peristimulus time histogram shown above or below it. Below each histogram is raster for the same stimulus, collected on the first recording day. Identity-level of each face increases from the average (center) on each radial axis.

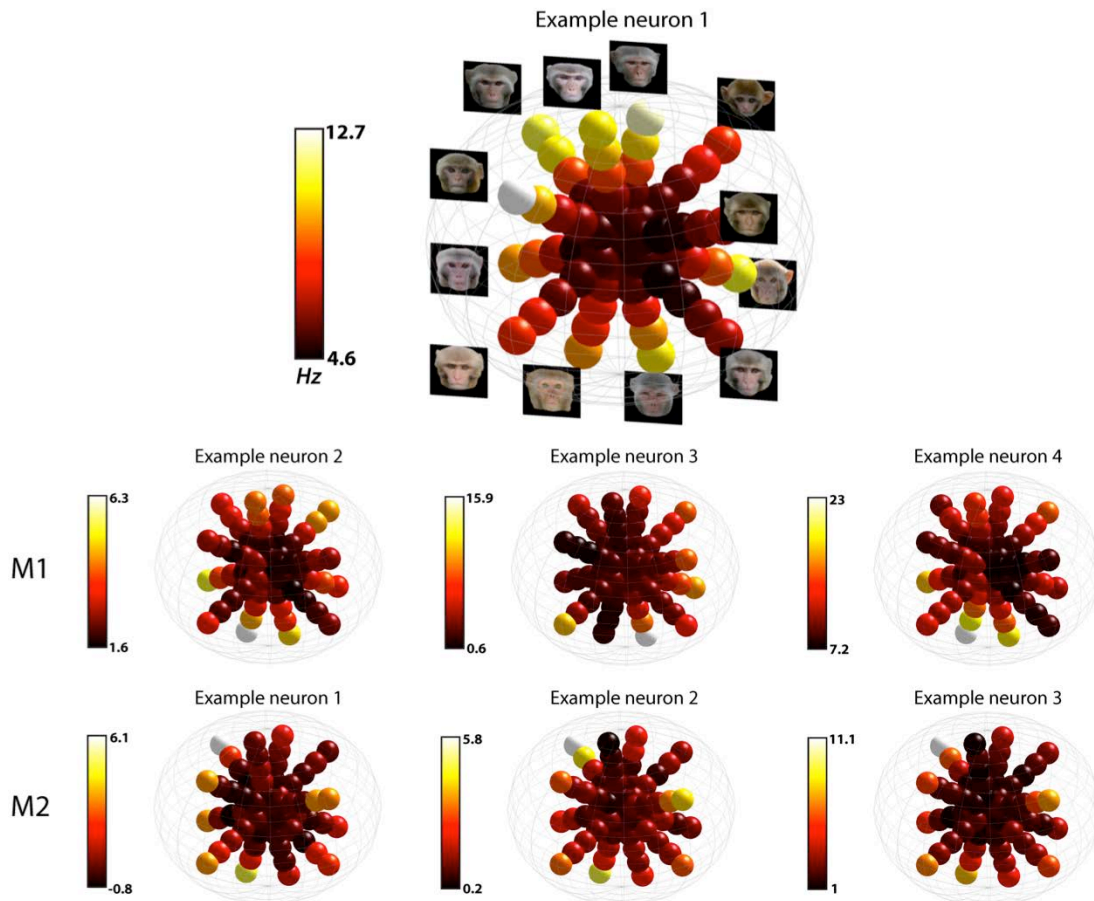


Figure 4-4 Response magnitudes for all monkey face identity trajectories. Each sphere of the dodecahedron corresponds to the response along the radial axis for each of the face identities shown in A (colored to indicate the mean spike rate of the response), representing a single point within the face space model. Each dodecahedron indicates the responses of a different example AF neuron.

In order to explore the correspondence between the responses of AF neurons to different face identities and identity-levels, and norm-based models of face perception, the responses of each cell to each identity trajectory were projected within the previously mentioned “face space”, in which the average face occupies the center. Figure 4.4 shows multiple examples of the selectivity of individual neurons to all 12 of the monkey face-identity trajectories we tested. As in the previous figure, responses are represented by color-coded spheres. However, in this case, they are arranged in a geometry similar to that of the norm-based face space, with each trajectory emanating from the centrally-located average face. We found that AF

neurons commonly exhibited “ramp-shaped” tuning between the average and the full identity face for a subset of the face identities presented. The monotonicity of this tuning varied in its slope for different face identities, as the number and identity of preferred faces varied between neurons. But, as mentioned above, there was typically no obvious resemblance in the superficial features of the faces that were most preferred by each neuron. Thus above all, the average face was the principal determinant the neurons’ tuning, typically eliciting a response that was at or near the minimum, with responses dropping off from the full identities.

Stronger responses for monkey vs. human faces

Given that we presented a large number of stimuli from both human and monkey face spaces, we next examined whether there were any species-specific differences in the responses of AF neurons. We found that morphed human faces showed tuning with the minimal responses at the center, but that the responses to the morphed monkey faces was usually much stronger. An example of this observation can be seen in the AF neuron in Figure 4.5. This species specific effect is described in greater detail at the population level in the next section.

Population results

The species-selectivity we observed seemed to be consistent across the population of AF face-cells. As indicated by the results of the population analysis shown in Figure 4.6, the average response of neurons in both animals (M1 and M2) to the full identity (100%) monkey faces were significantly greater than those for human faces ($P \ll 0.01$, two-sampled t test; see inset- Figure 4.6). In addition, when comparing the role of identity level in this relationship, the average of each identity

trajectory for macaque faces is larger than the equivalent average trajectory for human faces in both animals (see main plot- Figure 4.6).




Figure 4-5 Responses of an example AF neuron to human versus monkey faces, projected within norm-based face space. Each sphere of the dodecahedron corresponds to the response along the radial axis for either human (left) or monkey (right) face identities.

The population analysis described thus far did not rule out the possibility that the shape of this tuning was being driven largely by the responses to a small number of faces that were for some reason special. In order to explore this, we examined the consistency of this tuning for the individual face-identity trajectories included within this stimulus set across the population. Prior to averaging across cells, we normalized the responses of each neuron. If the neurons in this population were not consistently sensitive to the distance of each face from the mean, then the peak of each averaged trajectory should fall somewhere other than at the full identity level. If these neurons are consistently sensitive to the distance of each face, relative to the mean, then each averaged trajectory should possess exhibit gradual, monotonic increases in normalized firing rate when proceeding from the mean to the full identity face. As exhibited by Figure 4.7A, we consistently found selectivity in the form of ramp-shape tuning for identity level within each of the 12 monkey (see Figure 4.7A- top) and

human (see Figure 4.7A- bottom) face trajectories presented, after averaging across the population of both animals.

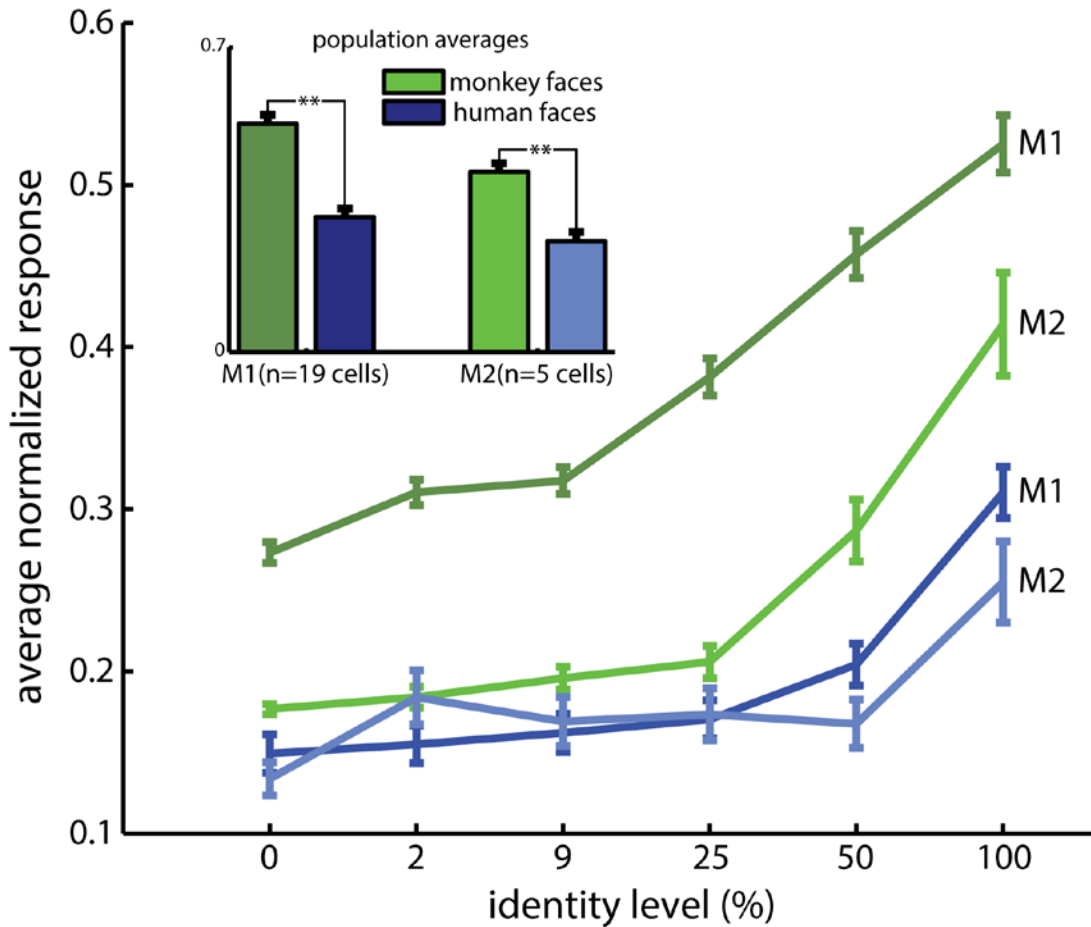
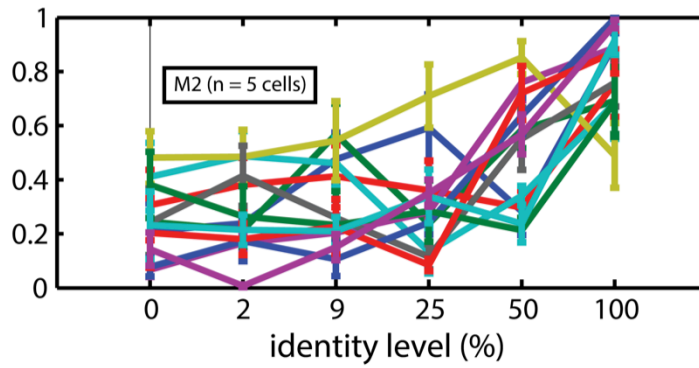



Figure 4-6 Population responses as a function of identity level for monkey versus human faces **Inset.** Population average responses for all neurons in each monkey (M1- darker, M2- lighter), for full identity (100%) monkey (green) vs. human (blue) faces. **Main plot.** Averaged identity trajectories for faces of each species as a function of identity level for the population of AF neurons in each monkey. Error bars: s.e.m.

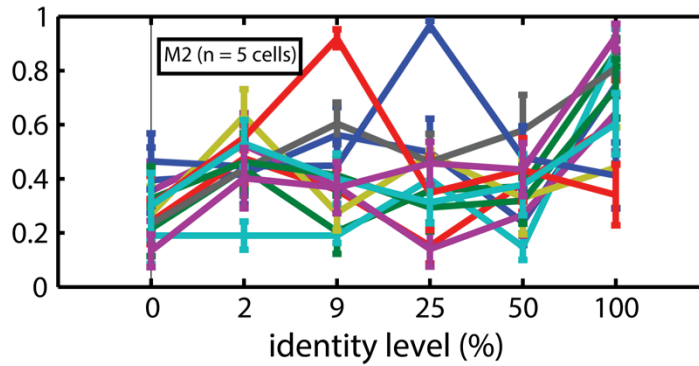

We next asked how many identity trajectories, on average, tuning functions in the same direction away from the average face. To do this, we computed the rank-ordered population average of each of the 12 monkey (see Figure 4.7B- top) and human (see Figure 4.7B- bottom) face trajectories presented in this study. In this analysis, the responses to each of the 12 individual faces were ranked based on the maximum (normalized) response for any face along the trajectory. For each response

A

population responses - monkey faces


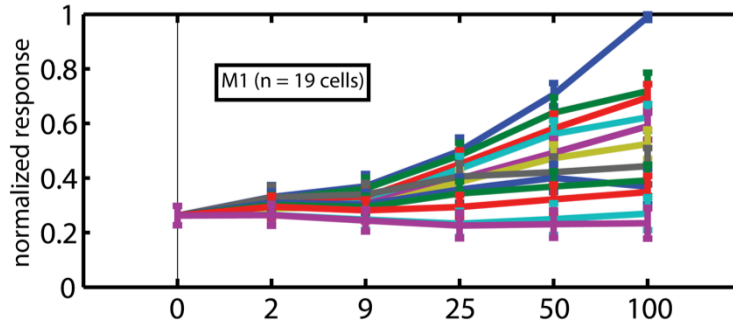


population responses - human faces



B

rank ordered responses - monkey faces



rank ordered responses - human faces

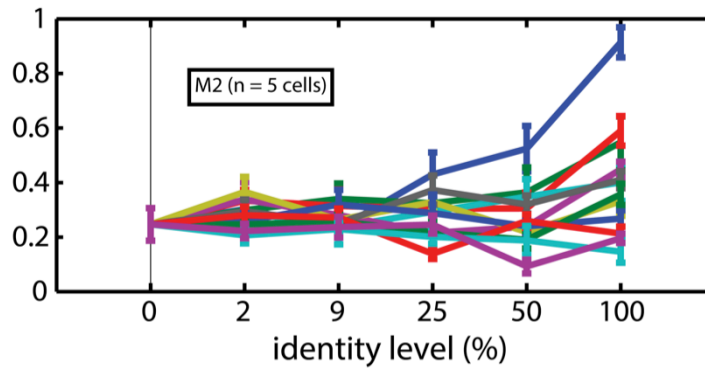



Figure 4-7 Population responses as a function of identity level- a. In this analysis, each identity trajectories for all 12 faces of each species (monkeys-above; humans-below) are averaged across all neurons in each monkey. Each panel corresponds to the population responses in one monkey. b. In this analysis, each identity trajectory is rank-ordered from maximum to minimum for each neuron, and then averaged separately for all 12 faces of each species, across the population for each monkeys (M1 and M2).

rank, an average tuning curve was then computed for each rank-position (1st, 2nd, 3rd, etc.) over all neurons as function of the identity level. This analysis reveals the extent to which monotonic tuning was shared, with the same valence, across identity trajectories. If, for example, AF neurons tended to respond in this manner to only 25% of faces, one should expect to see ramp-shaped tuning along the identity trajectory for only the 1st, 2nd, and 3rd ranked trajectories, but flat tuning for the remainder. If however, this tuning was consistent, to some degree, within each face in the sample, all ranked trajectories should show similarly ramp-shaped tuning. As exhibited by Figure 4.7B, with the exception of human faces in monkey one, the most preferred identity trajectory exhibits a much greater slope than the remaining identity trajectories, across the population, while the remaining identity trajectories exhibited less positive, and even slightly negative slopes away from the norm. This suggests that the majority of different monkey faces, and a smaller portion of preferred human faces, exhibit ramp-shaped tuning for increasing levels of distinctiveness in AF neurons.

Does norm-based tuning arise because of within-session adaptation?

One suggestion that has been made is that the identity-level encoding, which was described in the previous section, is not an accurate reflection of the neurons' identity tuning, but rather a result of within session adaptation (see Kahn & Aguirre, 2012). According to this view, the main observation of the lowest response for the average face, here and in the previous electrophysiology study (Leopold et al., 2006),

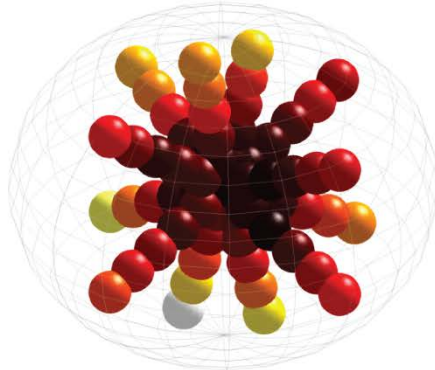
could be attributed to the repeated presentation of morphed faces of low identity. In other words, during each day of testing the norm-based tuning would have been *created* by the presentation of the stimuli used to evaluate the tuning. Now, armed with the unique capacity for recording the same neurons over many recording sessions, we were able to address this concern explicitly by focus our analysis on just the *first* response to a given stimuli on each recording day. With this approach, the within-session adaptation is minimized is thus unlikely to have had any effect by this point, and should reflect the “de-novo” selectivity of these neurons, before adaptation might have an effect.

Figure 4.8 compares the population average dodecahedron from M1, with the same depiction when we exclude responses outside of a particular time window. As shown within the middle panel, when we exclude all but the very first responses to each stimuli each day in our analysis of identity selectivity – we continue to observe very similar identity-level tuning, with marginally increased noise as a result of the reduced sample size (75 trials rather than the original 1,350).

Does norm-based tuning emerge gradually over sessions?


Given the longitudinal recording method, another question we could ask is whether is that norm-based selectivity develops slowly or is there from the beginning. Here we investigated this possibility by comparing the tuning of individual neurons during the first day of testing to that that emerged later. As shown by the bottom panel of Figure 4.8, tuning on the first day, for the very first trials that the monkey saw this stimulus set, closely resembled the tuning for the remainder of the stimulus

Population Average (n = 21 cells)
All trials, all days




.1 .7
Ave. normalized resp.

First trial, all days



.3 .7

All trials, first day



.2 .7

Figure 4-8 Time-limited dodecahedrons. Population average dodecahedron arranged as in Figs 4-4 and 4-5, and averaged across population of AF neurons in M1. Middle and bottom panel represent responses within limited time windows. Responses are limited to either the first trial for each stimuli, spanning all 75 recording days (middle) or the first recording day, including all presentations of each stimuli (bottom).

set. Thus there was no detectable phase of learning or familiarity in the expression of norm-based tuning in our experiments.

Responses to along extended identity trajectory

While these results demonstrate that the abstract representation of the average face, presumably represented internally following sufficient experience, seems to influence neural responses to faces in AF, it remains unclear how AF face-cells would respond to manipulations within additional regions of the proposed norm-based face space. The results in this section, therefore, focus on more extreme stimuli, along the same dimensions as the previous section, in which the distinctive features have been greatly exaggerated. These include face manipulations extending beyond both the original, full identity face (caricature faces) as well as manipulations beyond the norm (“anti-faces”).

To avoid potential confounds related to changes in the texture of morphed faces, the texture of the stimuli described in this section were equalized throughout the entire stimulus set. As a result, the faces included in this stimulus set (unlike those of the previous section) could be described as sharing the same “hair” and “skin”.

Previous psychophysical studies of facial recognition have found that the identity of faces possessing exaggerated distinctive features (i.e. caricatures) are more quickly and easily perceived than the original face. If AF neurons are participating in an analogous encoding process, as suggested by the results of the previous section, it might be expected that responses to caricatures would exceed those of their original counterparts as well. Further, AF face-neurons exhibiting ramp-shaped responses to

faces morphed between norm and full-identity faces, might also exhibit linear increases in response to extrapolation of the distinctive features of caricaturized faces.

Another finding from psychophysical studies of norm-based face processing is that faces extrapolated in the opposite direction of a particular identity trajectory, relative to the norm, (i.e. “anti-faces”) are perceived as “opposite” identities from that of the full-identity face (Rhodes and Jeffery, 2006). Previous studies have reported on a small number of instances in which face-selective neurons demonstrate similar, ramp-shaped tuning in their responses to anti-faces, which is horizontally-reflected at the average face (Leopold et al., 2006). Such “v-shaped” tuning, centered about the average, would correspond perfectly with norm-based perception of face identity. But this area of norm-centered “face space” has not previously been explored systematically.

As mentioned previously, recent recordings in the middle face patch (ML/MF), have found that the response of face-neurons exhibit linear, monotonic relationships with the geometry of particular face components. In other words, response variation was monotonic in that it followed linear transformations of particular face features in a ramp-shaped manner, without any sign of sensitivity to the feature’s distinctiveness. However, if neurons in AF are, in fact, sensitive to the distinctiveness of faces relative to the norm, we might anticipate that they will exhibit a reversal in their responses to faces within the anti-faces regime, resulting in “v-shaped” tuning, centered about the norm.

When we tested the responses of AF neurons to caricature faces, we found neurons exhibiting roughly linear increases for a subset of faces extrapolated beyond

the full identity face (see Figure 4.9, top and bottom rows). When testing faces manipulated into the opposite region of this theoretical face-space (i.e. “anti-faces”), we also observed similarly linear increases in response to a subset of faces extrapolated beyond the mean-face (see Figure 4.9, middle row).

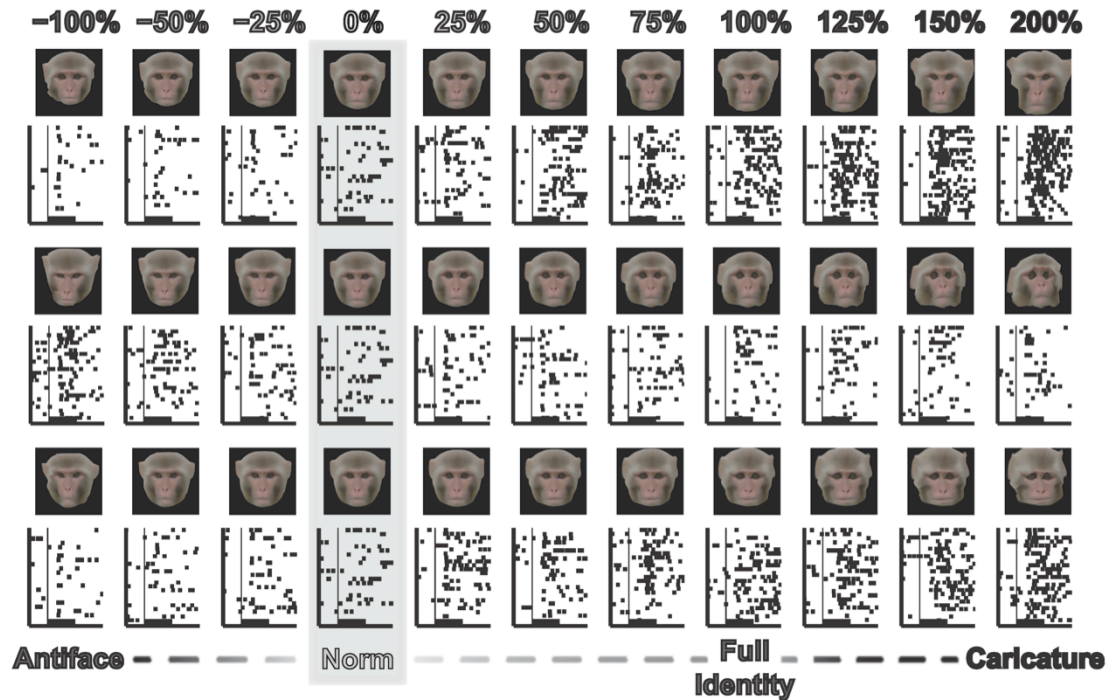


Figure 4-9 Responses of a single AF neuron to faces along the caricature/anti-face trajectory. Each panel corresponds to the rasters in response to a 300-ms presentation (black bar) of the stimuli shown above each. Rows represent distinct trajectories, each incorporating a different full-identity (100%) faces. Columns represent different identity levels within each trajectory (anti-faces-left-most column; norm face- 4th column; full-identity faces-8th column; caricatures-right-most column).

However, no clear evidence of “v-shaped” tuning (centered about the norm) was observed in the responses to changes along individual face trajectories (Figure 4.10- gray lines). Although, after averaging the responses of single neurons across either human (n = 12 faces) or monkey (n = 6 faces) faces (black lines), several examples of “v-shaped” tuning (reversing at or near to the mean) were observed. The majority of neurons that exhibited ramp-shaped tuning in their average responses between the average and all full-identity faces, also exhibited greater averaged responses for caricature faces, and reversals in averaged responses for anti-faces,

reflected by response increases in both directions away from the average face in the averaged response for all trajectories. Notice that certain cells share a reversed relationship with the norm, such that responses to the norm are highest, and decrease in either direction away from the norm (e.g. Neuron 3). Both conditions still imply a unique role for the average face, relative to those faces occupying the remainder of our theoretical face space.

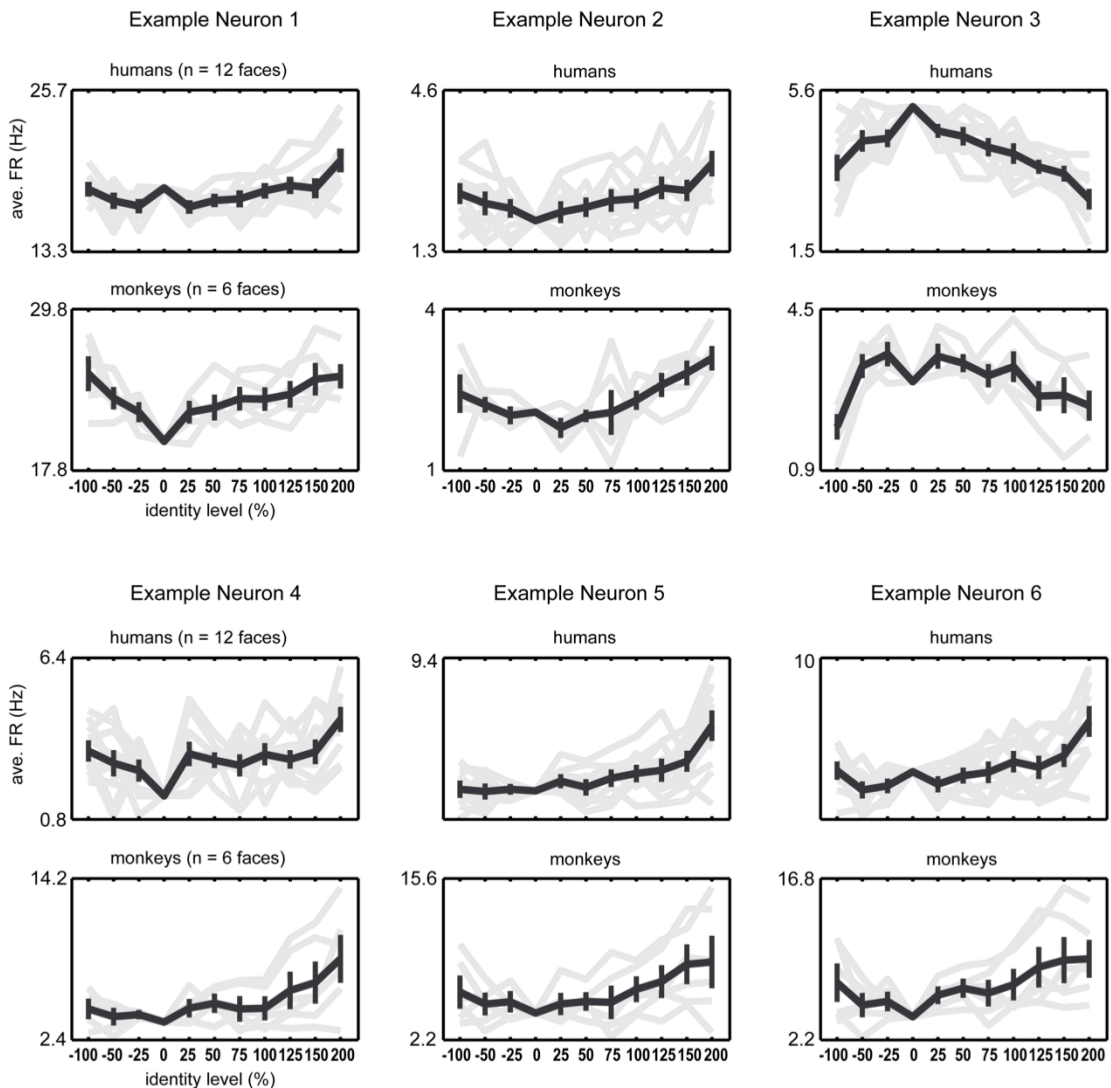


Figure 4-10 Caricature trajectory averages The responses of 6 example AF neurons (columns) to identity trajectories including caricatures and anti-faces, which were averaged across all face identities in each species (rows).

Since the precise nature of the mental representation of population norms are not yet clear, some assumptions were made regarding the resemblance of our face average of the true combined average of all experienced faces for each of these animals. While these results suggest that our estimation may not be far from correct, certain examples of caricature/anti-face trajectories exhibiting response minima at locations near to 0% identity levels (i.e. -25 and 25%), may be accounted for by discrepancies in this process of approximating the norm.

Population responses for extended identity trajectory

In order to explore how well the activity AF face-neurons can be accounted for by a norm-based theory of face-processing, responses to changes in the distinctiveness of all faces (both human and monkey), the tuning curve for each caricature/anti-face trajectory was averaged across the population of neurons in each animal. As in the analysis described in the previous sub-section, the tuning curves of neurons were first normalized to the highest responses for each trajectory and inverted in the cases showing highest responses for the average face, permitting the inclusion of neurons that exhibited ramp-shaped tuning either toward or away from the average face.

As can be observed in Figure 4.11, the trends observed were representative of the population of AF face-cells in both animals. When comparing the average of each neuron's average caricature/anti-face trajectories across all neurons in a single monkey, we see that the population average for each monkey seems to follow a highly similar trend of monotonicity in either direction moving away from the norm. The highest and lowest responses across the population were at 0% and 200% identity

levels, respectively, for both M1 and M2. Both animals demonstrated a clear monotonicity for face identity moving away from norm to the caricature, which was imperfectly reflected in the responses to the anti-face.

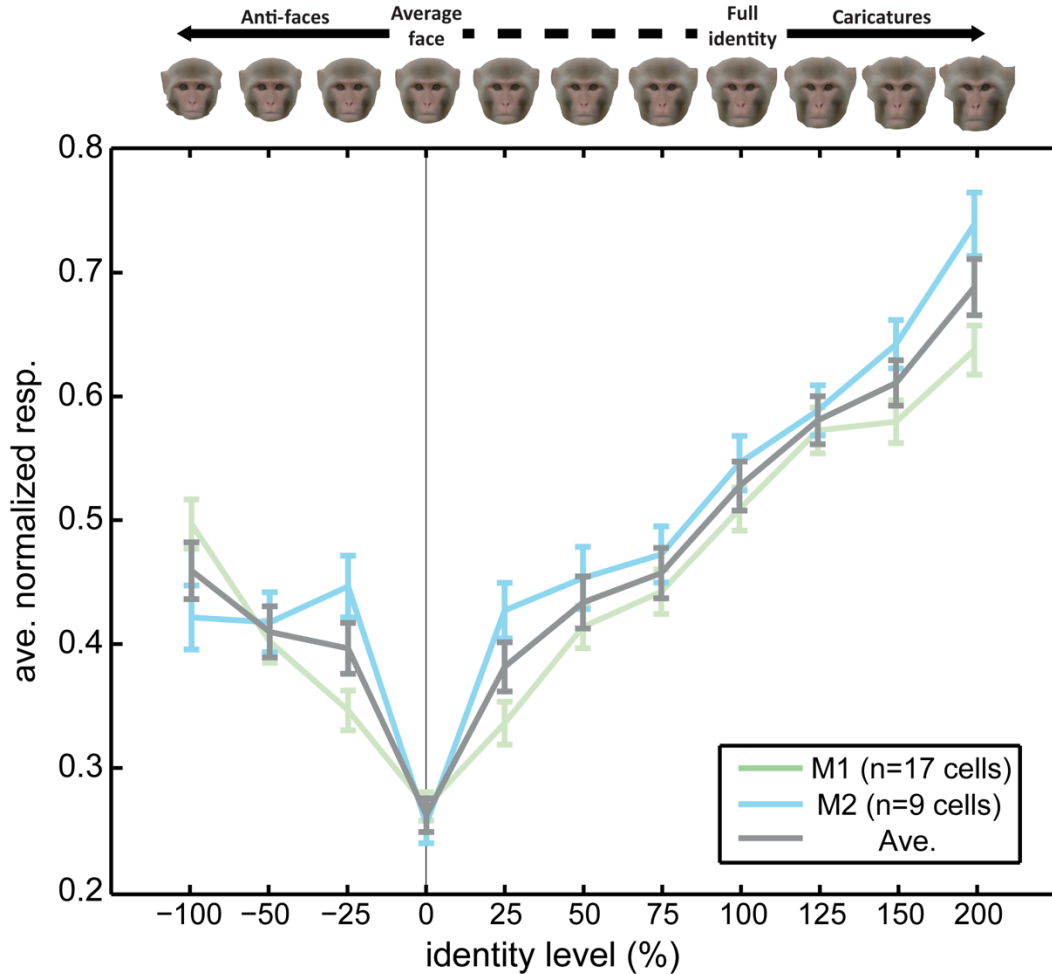


Figure 4-11 Population caricature responses. The responses of each neuron to identity trajectories including caricatures and anti-faces were averaged across all face identities (both species) and then averaged independently across the population of neurons in each monkey (M1-green, M2- blue).

The final component of a theoretical norm-based face space that was explored by in this study were trajectories linking different identities to each other (i.e. “tangential” trajectories). Comparing the responses to “radial” (i.e. between the norm and the full identity face) with responses to “tangential” (i.e. between two full identity faces, without respect to the mean face) allowed us to test an important prediction made by the norm-based hypothesis- that simultaneous manipulation of multiple face

features, which do not converge toward the mean, should influence the responses of presumed norm-based face-neurons in a random, non-systematic fashion. In order to explore this possibility, we presented faces along tangential directions in face space, corresponding to morphs between different facial identities, but without transitioning through the average face.

Responses to 12 (6 human and 6 monkey) individual tangential trajectories were normalized and averaged in the same manner as those of radial trajectories. If the changes occurring within each of these trajectories were influencing the responses of AF neurons in a purely random fashion, and responses to a sufficient number of trajectories were sampled, we should expect the average of the tuning curves for each of these trajectories to be flat, with a value of .5 (since responses were normalized between 0 and 1). As shown in Figure 4.12, no systematic trend was observed for faces morphed along these tangential trajectories and the average normalized response value for each step along this trajectory was consistently near to .5. This result could be very different if, for example, AF neurons were responding to changes occurring during the morphing of images (e.g. percentage of high vs. low spatial frequencies contained in each face). In this example, v-shaped tuning might be expected for the averaged responses along the tangential trajectory, since this should consistently reverse at 50% for all tangential morph trajectories between full-identity faces. The relatively flat tuning that was observed, however, suggests that the sensitivity of these neurons to changes within face morph sequences is not an artifact of the morphing procedure that was employed here.

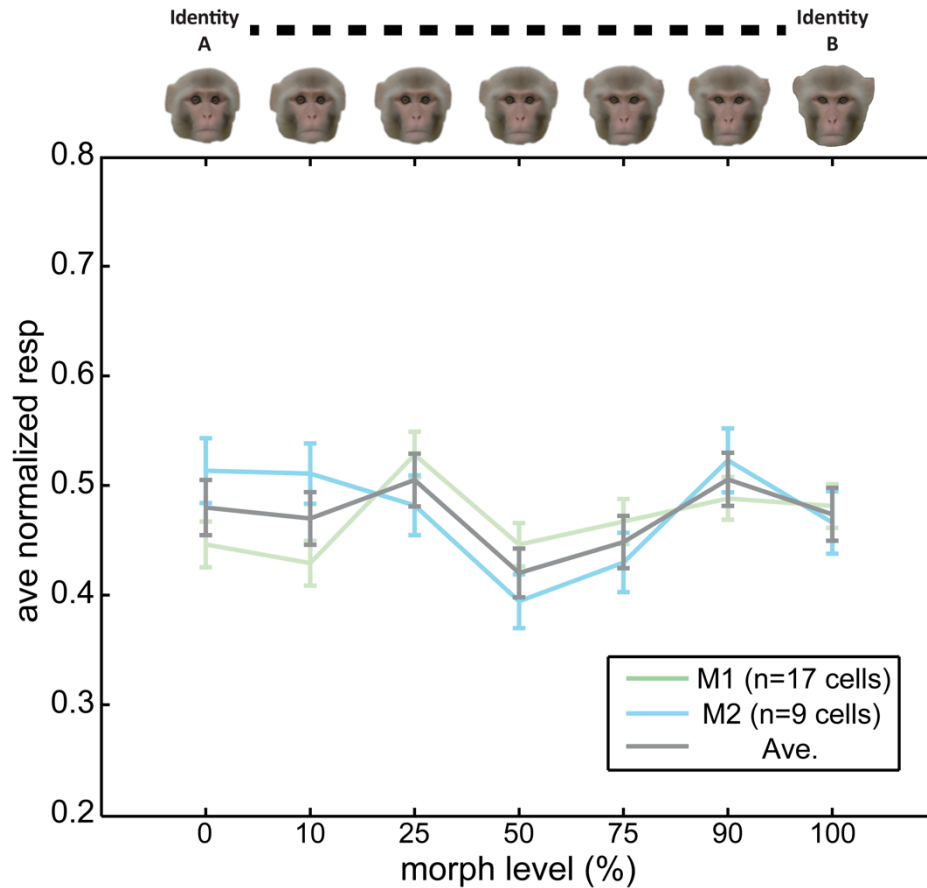


Figure 4-12 Population responses to tangential morphs The responses of each neuron to tangential morph trajectories were averaged across all face identities of both species, independently across the population of neurons in each monkey (M1- green, M2- blue).

2. Tuning for other properties of faces in AF

One of the major difficulties in describing the selectivity properties of face and object-selective visual neurons is the complexity of their tuning. The responses of face-cells are known to be sensitive to changes along multiple feature dimensions simultaneously (e.g. face versus object and static versus dynamic; Polosecki et al., 2013). In the previous section, we reviewed the selectivity of AF neurons for faces of a particular identity, and how this selectivity corresponded to the changes which might occur within a theoretical norm-based face space. In this section, we explore how the responses of the same AF face-cells changes following 3-dimensional

manipulations of a preferred face identity. In doing so, we take advantage of an additional advantage uniquely afforded by the longitudinal recording technique- the ability to observe the response properties of the same neuron using multiple stimulus sets. This makes it possible to investigate the selectivity of neurons for a particular parameter, and then make use of that information in the design of subsequently presented stimuli.

Background

Previous studies utilizing fMRI-localized single-unit recording have shown that face-cells in the posterior ML/MF patches exhibit sensitivity to the 3-D orientation of face stimuli (Freiwald et al., 2010). Whereas, responses in more anterior AL patches exhibit less selectivity to such changes. Ultimately, neurons in the anterior-most AM face patch exhibit robust responses to faces of particular identities, without regard for their 3-D orientation.

Based on these results, a conception of the face patch network has emerged in which strongly inter-connected patches (as evidenced by microstimulation studies- Moeller et al., 2008), possess increasingly view-invariant representations at successively higher levels in the face patch hierarchy spanning from extrastriate to PFC. Similar specialization has been found in fMRI studies of face patches in the human brain- Haxby and colleagues (2000) found that early face processing regions exhibited greater sensitivity to variations in face stimuli than later ones within the visual pathway of the human brain.

In this study, we sought to explore the sensitivity of neurons in face patch AF to face orientation, in order to contribute to an emerging picture of the functional specialization within the face patch network.

In order to do so, we first screened the responses of AF neurons in one animal to a collection of 250 human face images. The responses of an example neuron to this initial screening is shown in Figure 4.13A. Through a subjective assessment process, the identity evoking the greatest response among the population as a whole was selected (see red box, Figure 4.13A). This face was then used to create a much larger stimulus set that included multiple 3-dimensional manipulations of the preferred face identity (see Figure 4.13B). Three different manipulations to the 3-dimensional viewing angle of the preferred-face identity were presented: 16 levels of rotation in the picture plane (rotations about the clock-face) spanning all 360° of rotation, 11 levels of head rotation spanning 90° (head turned to the person's left) to -90° (head turned to the person's right), and at 4 sizes of retinal angle: 1.5° , 3° , 6° , and 12° . This dense-sampling of "orientation-space" was only possible due to the extended period of observation afforded by the longitudinal recording technique. In order to accomplish this, 176 images were presented at 5 different sizes 32 times each – requiring 28,160 trials to record over 12 consecutive days.

In order to visualize the responses of AF neurons to simultaneous manipulations along multiple dimensions in "orientation-space", we created heatmaps for each neuron, in which the x-axis is defined by changes in the picture-plane and the y-axis is defined by changes in the rotation of the head. In this projection, each pixel represents the response of each neuron to a combination of picture-plane rotation and

viewing angle values. An example from a representative AF neuron is shown in Figure 4.13C.

Picture-plane rotations

Although this dimension was sampled more broadly than was gaze angle rotation (360° versus 180°), these neurons still exhibited more examples of selectivity to changes in the picture-plane (x-axis- Figure 4.13C) than in the head rotation (y-axis- Figure 4.13C). Response areas exhibited a variety of shapes, most commonly columnar, centered about the upright face (0° picture-plane rotation). Responses then decreased substantially outside of $\pm 60^\circ$ rotation away from the upright face. This property was not observed when mapping responses to non-face objects within an analogous “orientation-space”. This observation would seem to suggest that these neurons (or their afferent partners), may be performing a form of template-matching during the encoding of face stimuli, such that their sensitivity is greatest to faces in an upright, “natural” position.

3-D Head rotations

Tuning of AF neurons was much less selective for changes in head rotation than for changes in picture-plane rotation (see Figure 4.13C). As can be seen in the example provided, select AF neurons responded robustly to drastic changes in the viewing angle of a face (left-facing to right-facing), despite those images undergoing substantial changes in the relative position of their individual facial features. This may also suggest that AF neurons respond robustly to changes in face stimuli, which are likely to occur during “natural” behavior.

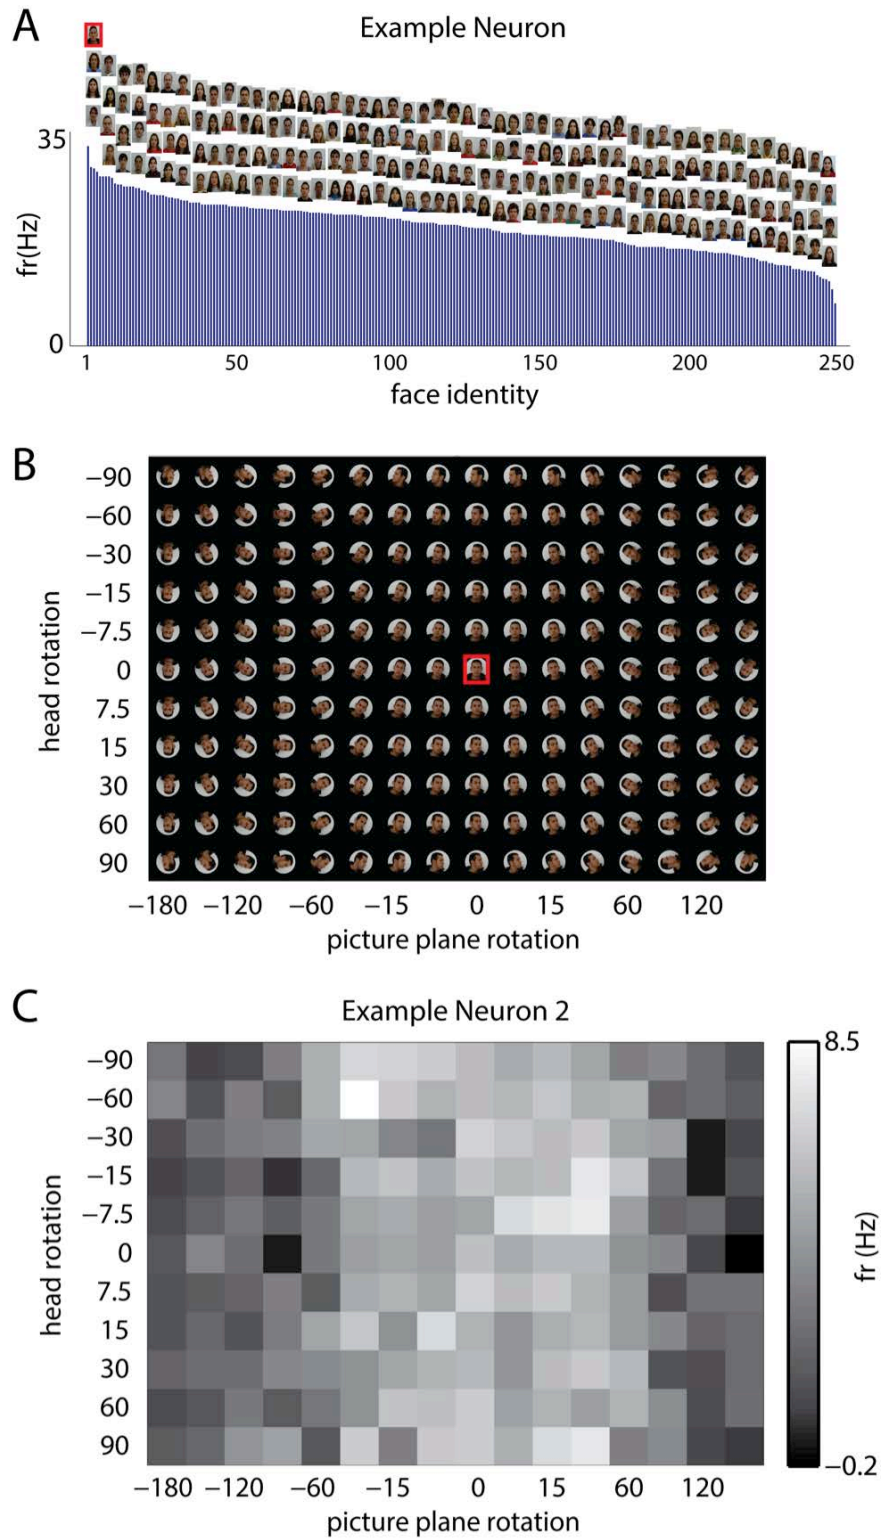


Figure 4-13 Responses of an example AF neuron to 3-dimensional manipulation of a preferred face identity. Responses to 250 different face identities were compared (A) in order to identify the face most “preferred” by this population (highlighted in red). This image was then transformed along multiple dimensions, including picture-plane rotation and head rotation (as shown in B), as well as image size. The response to these manipulations are shown in C, with each pixel representing the response of an example AF neuron to the corresponding stimuli in B.

Image size-

In order to look more closely at the role of image scale in moderating the responses of AF face-cells, faces were presented at 1.5°, 3°, 6°, and 12° visual angle. Response amplitudes were averaged across all stimuli and all neurons in this population (see Figure 4.14). While this population exhibited with relatively variable selectivity in their tuning to picture-plan or head rotation, neural responses in AF were far more reliably moderated by image size. Responses to images presented at 6 and 12 degrees were significantly larger than those for smaller image sizes.




Figure 4-14 Population responses as a function of image scale for the “preferred” face identity. The average responses for all neurons in M1 are compared. Responses to images of 6 and 12 degrees were significantly greater than those for smaller image sizes (*t*-test, $p < .05$).

The reliability of size as a moderator of responses contradicts previous descriptions of face-cells in which responses were invariant with regard to changes in

the size and/or position of a presented face (Ito et al., 1995). Given the behavioral relevance of face size in determining the proximity of face, and the significance role that face-size information could be expected have on social behavior, AF face neurons may play a specialized role in encoding face proximity information within this network.

III. Discussion

1. Identity encoding

A growing body of literature exists regarding the most accurate representation of face encoding in IT. Models describing how faces are represented in the brain loosely belong to one of two groups – exemplar models or axes models (Tsao & Freiwald, 2006). The predications associated with each model can be differentiated based on the most relevant dimensions that influence face discrimination within a theoretical, multidimensional “face space” encompassing the entire range of variation that exists within the enormous diversity of faces that we are likely to encounter.

An exemplar model, as the name implies, predicts that each face is represented in terms the unique combination of geometric features it possesses. The prediction of axes-based models, on the other hand, would be that all faces are represented in terms of their relationship with a prototype, representative of a given stimulus category. Both models predict an uneven distribution of faces within a hypothetical face-space (a densely populated center consisting of more common faces). But, in the norm-based scenario we could predict that tuning for a given face feature would follow a linear, monotonic curve as each feature moves either towards or away from the mean, in particular. This is in perfect agreement with the findings of

Leopold et al., (2006). They used a face-morphing technique to show that responses of some face cells increased linearly as faces are shifted away from the average of the entire stimulus set. This suggests that previous conceptions of a norm-based face space, established previously by studies of face perception and face identity aftereffects, may do a good job of accounting for the responses of individual face-selective neurons.

fMRI-localized electrophysiological recordings by Freiwald et al. (2009) in face patches posterior to those in Leopold et al., made use of cartoon faces in order to independently manipulate the structural parameters of the face such as the position of the eyes, nose, and mouth. The important difference between these studies being that, in the latter study, changes were made without regard for similarity with any categorical prototype or average. Despite this difference, Freiwald et al. also found neurons that exhibited monotonic tuning functions, similar to those reported in the first section of this chapter, though their response minima or maxima were not positioned near the average face.

In the studies reported here, we made use of the enhanced sensitivity afforded by chronic recording to test the extent to which norm-based encoding of faces is employed by AF face-cells. We found that a large portion of face-cells demonstrate a strong tendency towards tuning relative to a norm, frequently exhibiting more consistent sensitivity to the distinctiveness of a face, relative to the norm, than for changes between different face identities.

These findings demonstrate that the distinctiveness of a face relative to the mean, a property known to be an important dimension regarding the perceived

identity of a face, is also a major influence on the firing rate of face-selective neurons in the AF face patch.

This sensitivity to changes in the distinctiveness of a face relative to the norm is likely to represent the outcome of a comparative process between incoming faces and the population average. The specific nature of this average – for which categories it exists, how many there are, how they develop – remains unclear. The results described here suggest that it is of particular relevance for certain stages of the face-processing pathway.

Potential role of adaptation-

As mentioned previously, a point of contention regarding the “specialness” of the average face, and its influence on neural encoding, is that the tuning observed is simply the product of adaptation to the statistics of a prototype-based stimulus set. Numerous studies of adaptational after-effects have suggested that this form of selectivity is plastic, and seems to be influenced by the presentation of recent stimuli (Van Rensbergen & Op de Beeck, 2014; Leopold et al., 2005).

But, there is ample evidence in the psychophysical literature for contextual effects for a variety visual properties, including motion, brightness, orientation, as well as faces (Schwartz et al., 2007). However, at some level, we still assume that distinct cortical regions are accurately described as representing certain aspects of our visual environment. Using the unique advantages of the longitudinal recording technique, we have attempted to parse the influence of the temporal context of visual stimuli, from the underlying computation being performed within the AF face patch. By focusing our analysis on distinct windows within 75 days of nearly-consecutive

recordings from the same neurons, we have attempted to show that this particular computation (the distinctiveness of faces) seems to be present throughout the time period we observed. Most importantly, perhaps, was the observation of this selectivity on the first presentation of each stimuli. These results suggest that if these neurons are encoding face identity in a relativistic fashion, the prototype against which faces are compared is established either earlier in life (perhaps during a limited critical period in development).

Species sensitivity

Despite the relative abundance of similarities shared by macaques and humans in processing of conspecific faces, which might be taken to argue in favor of a common neural mechanism, it remains unclear whether a common face processing template exists for different primates. While previous studies have employed both human and monkey faces in experimental stimuli, very few reports have directly compared the differences in the responses elicited by each in individual face-selective neurons in the brain.

In this study, the responses of most cells were distinctly higher following the presentation of conspecific faces than for those of humans, while still maintaining an overall encoding scheme consistent with the norm-based tuning described in the previous section.

Although these results make sense intuitively, they are somewhat surprising in light of previous findings within the face patch system. Previous reports, using both single-unit recording (Freiwald et al., 2009) and functional MRI (Tsao et al., 2006), have made use of macaque, human, and even greatly simplified cartoon faces and

treated the responses to each of these categories as equivalent. In one study, which compared these responses directly, fMRI was used to closely compare the responses of face-selective regions in the brains of humans versus macaques, during the presentation of both human and macaque faces (Tsao et al., 2003). Authors found that while humans exhibited greater responses to human faces, the responses in macaques exhibited no preference for macaque over human faces.

This result may offer additional insight into physiological basis of the “other-species effect”, and more specifically, whether it is a result of the response properties of face-cells such as these. As were all comparisons of neural selectivity over time, the preference for macaque faces in the responses of AF face-cells was stationary throughout the entire time they were observed. This suggests that these properties are fixed for intervals of at least several months in the adult macaque. Whether these properties are innate or are the result of exposure and experience during a certain developmental window earlier in development remains uncertain.

Due to the increased likelihood of encountering them, an enhanced ability to discriminate subtle details in the faces of conspecifics is likely to afford social animals certain behavioral advantages. Increased selectivity for the faces of conspecifics, for example, may make it easier to discriminate faces of relatives from non-relatives. These findings leave open the possibility that these cells are involved in conveying social information extracted from the faces of conspecifics.

Realistically, it’s likely to be a combination of genetic-hardwiring to be sensitive to the faces of conspecifics at birth followed by learning to associate the faces of relatives with positive rewards, which allow these neurons to maintain

optimal responsiveness to the most ethologically-relevant stimuli (the faces of conspecifics).

Caricatures

Most of us are regularly exposed to caricatures in the media, and as a result, can attest to their remarkable ability to exaggerate the distinctive or unusual features of the face to sometimes unrealistic proportions, and yet still be recognizable as the individual the artist intended to portray. Likewise, although caricatures differ from the true arrangement of their original counterparts, studies have shown that caricatures are recognized more quickly and accurately than their true, veridical originals (Benson & Perrett, 1994). In this way, caricatures provide improved access to memories of faces, which has been suggested to support to models of face-processing based on norm-based coding (Rhodes & Leopold, 2011) -- although the caricatures differ from their originals, they do so along a dimension of “face-space” which is consistent, relative to the norm.

In the beginning of this chapter, we reported on the ‘ramp-shaped’ tuning for identity levels between the original and the average face that were observed in the neurons in this study. As the next component in this study, we created morphed-face stimuli, in which the distinctive features, relative to the norm, were exaggerated beyond the level of their original faces. In many instances, we found AF face-cells that increased continuously for caricature faces, as well- exhibiting a monotonic increase in firing rate for caricature faces projected beyond the original. In addition to the results of preceding sections, this result further suggests that the responses of AF

face-neurons are particularly sensitive to the perceived distinctiveness of an incoming face, relative to the norm.

Anti-faces-

Following the “face-space” trajectory in the opposite direction, beyond the average face, each face identity is paired with a corresponding anti-face identity. Psycho-physical studies of face-identity adaptation have shown that adaptation to anti-faces lowers the threshold for recognizing the original face (Rhodes & Jeffery, 2006), highlighting a particular importance for face manipulations relative to the norm, in particular. A similar effect for anti-face adaptation is not observed for the recognition of faces with different identities than the original, counterpart face. This result could be thought of as further support for a norm-based model of face-processing, since adaptation is specific to morph trajectories that passing through the average face.

To further explore this relationship, the faces in this study were also manipulated in the opposite direction of “face-space”, as well. In many instances, we found that AF face-cells exhibited ramp-shaped tuning, after averaging over multiple faces, for increasing distinctiveness of anti-face stimuli, relative to the average face. This tuning was often inverted in its slope, relative to the slope of tuning for the norm-to-caricature trajectory. Findings such as these are difficult to reconcile with an exemplar-based model of face processing, which would predict that in most instances, responses should exhibit a linear decrease as faces were morphed away from the original face. In other words, exemplar-based models would not have predicted how reliably tuning would be reflected at or around the average face.

The apparent influence of the norm in dictating the responses of these neurons leads to important questions regarding how the average face changes with learning and task requirements, which are addressed in the fifth chapter of this volume.

2. Tuning for natural transformations of the face and head

Studies using a closed-loop design generate novel stimuli, based on the online neural responses to previous stimuli (Hung et al., 2012). Using machine learning algorithms, descendant stimuli are probabilistically generated on each iteration. Stimuli generally evolve through multiple generations under the guidance of neural feedback, and are gradually transformed into the shape predicted to elicit the greatest response. This approach offers much denser sampling of the stimulus dimensions that most effectively dictate the responses of neural population, and thus increase the efficiency of time-constrained recordings of single neurons *in-vivo*. Such online approaches for neural characterization have just become feasible in the last few years, due to advances in parallel computing and machine learning theory. The unique advantages afforded by the longitudinal recording technique may be ideally suited for use in studies such as these. Given the greatly enhanced stability and duration of single-unit isolation offered by this technique, it may permit the exploration of stimulus selectivity in neurons at a new level of resolution.

We attempted to explore the possibility of using this technique in a similarly exploratory fashion by screening the selectivity of AF face-cells to a library of 250 human faces of different identities, before then generating an additional stimulus set in which we manipulated the 3-dimensional orientation of the most preferred face

identity. Using this approach, we asked how AF face-neurons respond to changes in the 3-dimensional viewing angle of a preferred face identity.

Following screening of the 250 face identities, AF face-cells typically responded at different intensities to a range of different identities, rather than strictly one identity. As was observed in previous sections, there was no apparent similarity in the faces that were preferred by any particular neuron. Given that these neurons seem to be carrying information involved in differentiating between face identities, and each neuron plays a role in encoding several different identities, these results support the notion that each identity is represented by a distributed ensemble of face cells in the brain.

Picture-plane rotations

Studies have shown that the face-inversion effect- the dramatic drop in our face perception proficiency when faces are inverted- was reversed in prosopagnosic (“face-blind”) patients, who could more accurately match inverted faces than upright ones (Farah et al., 1995). Normal subjects in this task are particularly good at matching upright faces and, conversely, perform worse with the inverted face stimuli. Clinical data, such as these, support the view that the face-processing system must rely on configurational processing at particular stages of the visual pathway, since they can be selectively impaired.

In order to determine whether AF participates in such configurational processing, we asked how changes in rotation (from upright to upside-down) changed the responses of face-cells to faces known to elicit a substantial response in the neurons we recorded. We found that neurons in this area varied in their tolerance to

picture-plane rotation, but usually decreased substantially in their response to faces outside a limited range of being rotated in either direction (+ or -) away from the upright (0° rotation) position. The majority of AF face-cells, therefore, seem to prefer faces in their upright, “natural” position. This sort of template-matching has been observed previously in IT (see Perrett 1985, but also see Bruce 1982), but these results suggest that AF face-cells, in particular, may participate in configurational processing of face stimuli.

Head rotations

Much more frequent during behavior, than the rotation of a face in the picture plane, are changes in the viewing angle of the head. Frequent changes in the gaze or in the relative position of two animals require that we maintain a robust representation of identity, despite substantial changes in the viewing angle of the face.

In this study, we asked how changes in viewing angle (turning of the head or “headshaking”) affected the responses of AF face-cells to faces known to elicit their response. We found that neurons in this area exhibited very little sensitivity in their responses to faces across the entire range of viewing angles presented, relative to their sensitivity to changes in the picture plane. This is despite substantial changes in the visual information present in the image of the face at each angle.

Such invariability of responses to changes in viewing angle may be highly adaptive as prolonged interaction commonly involves movements in viewing angle, but less so for changes in the picture plane of a face. If so, the general shape of the selectivity we have observed may be ideally suited to correctly extract the identity of

a face, throughout a wide range of changes in the view, which are likely to occur throughout the course of social interaction.

Recent studies have shown greater viewpoint-dependent selectivity in posterior regions of the temporal lobe, which are posited to define the local features of more invariant selectivity in anterior regions (Freiwald & Tsao, 2010). Viewpoint-dependent selectivity for local features in posterior face patches could provide the building blocks for the view-invariant identity tuning seen in anterior face patches. This may point to a division of labor within the face network between regions specialized for the perception of face identity.

Face size

Another manipulation that is likely to occur under ethologically-relevant conditions is variation in the size of the retinal image of a face, which reflect changes in its distance from the viewer. Such information would be highly relevant during behavior as proximity partially dictates the probability of many forms of social interaction.

In order to explore the sensitivity of AF face-cells to changes in the size of a preferred face identity, responses of AF face-cells were compared for faces presented at 1.5, 3, 6, and 12° retinal angle. Neural responses in AF exhibited a particularly reliable, positive relationship with the size of the presented face image.

Fitting with the results of previous studies (Ito et al., 1995), the overall shape of selectivity for other stimulus dimensions (head rotation and picture plane rotation) was similar across all face size levels. While the selectivity of these neurons for changes in face size was apparent, it did not seem to disrupt the simultaneous tuning

for additional viewing dimensions. This may allow AF neurons to simultaneously carry signals useful for correctly identifying the identity of a face and perceiving one's relative position among conspecifics.

Chapter 5: Effects of learning on face-selective neurons

1. Introduction

The “other-race effect”, in which subjects have more difficulty identifying faces of ethnicities other than their own, can be overcome following sufficient experience with faces of a foreign ethnicity (Pascalis et al., 2005). This is just one example of the malleability of face-perception as a consequence of learning. Such learning-induced changes in face perception occur not only during development, but also well into adulthood. Healthy adults, even in the late stages of their lifetime, continue to retain the ability to remember the faces they have encountered, and make use of this memory during subsequent facial recognition. The basis of these changes in recognition abilities must rely on changes in connectivity within the brain. The face-processing system must presumably alter its selectivity to previously unfamiliar faces in order to accommodate improvements in face discrimination. But, how such learning is related to changes in the face patch system of the macaque remains completely unknown.

In the third chapter of this volume, the stability of AF face patch neurons was explored over time frames that were previously impossible using standard electrophysiological techniques. It was shown the stimulus selectivity of neurons in the AF face patch showed little or no spontaneous changes over time. In Chapter 4, it was shown that neurons in the same area were conspicuously responding to faces based on their relationship with the average face. In this chapter, however, we will explore the stability of these same neurons, throughout the course of training in a learning

task specifically designed to engage the neurons within this system. The results of this section allow us to ask: if face processing is inherently plastic, and if the neurons in AF show norm-based tuning that is thought to reflect an animal's experience with faces, then might it be possible to change the nature of that tuning with experience?

Several examples of learning-induced changes in sensory neurons within the ventral visual pathway exist. As an example, Sakai and colleagues (1991) engaged monkeys in a paired-association learning task in which they were trained to associate pairs of arbitrary visual objects (Fourier descriptors) over roughly one month. During subsequent single-unit recording, they found a surprising proportion of neurons that responded selectively for visual object pairs, at a rate significantly more frequent than would be expected by chance. From this, authors concluded visual neurons in IT were participating in coding the relationship between associated objects. Another study found a greater proportion of neurons selective for arbitrary, non-face objects following behavioral training with those objects (Tanaka, 1996). In this study, monkeys were engaged in a delayed-match-to-sample task for a period of one year, before recording from single-units in IT. After comparing responses to trained v. untrained stimuli, authors found that about 25% of neurons responded maximally to stimuli included in the training – whereas, only 5% of IT neurons responded maximally in untrained monkeys. Another study, which recorded from IT in monkeys trained to identify a particular class of novel objects from two novel object classes, has shown that neurons in TE respond to a limited subset of views of trained objects, specifically -- as face-neurons do for views of the face (Logothetis & Pauls, 1995).

Based on our finding that a large portion of AF neurons seem to discriminate face identity based on its distinctiveness (Chapter 3), we engaged the monkey in a learning task that required the discrimination of faces identities. In the final phase of the task, the distinctiveness of each face was reduced following each successful trial, thereby increasing the difficulty as the monkey succeeded. By maintaining the difficulty of the task near the monkey's perceptual threshold, we hoped to increase the animal's engagement in the task. We looked for two forms of experience-dependent plasticity in neurons recorded before, during, and after behavioral training. In order to explore the first form, familiarization, we tracked changes in neural responsiveness over consecutive presentations of the same stimuli, in the absence of any demands imposed by training. In order to explore the second form- learning-induced plasticity- we tracked changes throughout face-discrimination training.

Here we show that, despite marked improvements in the monkey's perceptual capacity during this intensive training, neural responses to both the conditioned and unconditioned stimuli remained remarkably stable. As suggested by the results described in Chapter 4, these results also leave open the possibility that different regions within the face processing network exhibit distinct forms of plasticity from one another following learning. This might allow certain regions to be specialized for permanence, and others for plasticity. Such functional specialization may increase the capability of the entire system to meet the wide variety of demands imposed on it. Such examples exist in basal ganglia-thalamocortical circuit of songbirds, for instance, in which the entire pathway is responsible for song production, but plasticity following learning is limited to discrete components contained within that pathway

(Boettiger and Doupe, 2001). Likewise, it seems possible that while the entire face-processing system may be involved in processing faces, plasticity following learning related to faces may be limited to discrete components, which are yet to be identified.

I. Methods

The following section will review the details of the associational learning task that was performed.

1. Behavioral training

Behavioral training consisted of 4 different phases, each a variation of a conditional ocular-motor associational learning task and are described below. For all phases, each trial was initiated following the monkey's fixation on a small dot in the center of the screen, followed by the appearance of both the stimuli and saccade targets. Stimuli were presented in the middle of the screen and flanked by targets positioned above, below, and to either side of the stimuli.

Phase 1

In the first phase of expertise training, the monkey was presented with an object image and four surrounding circles that served as saccade "targets" (see Figure 5.1). If, following the presentation of a particular stimuli, the monkey made a saccade to the correct target, it was rewarded with water. After saccading in the correct direction on more than 80% of trials, additional stimuli were introduced and the conditioning process described above was repeated with a new stimulus. This was repeated with novel set of non-face images every ~3 days, for 20 training sessions, spanning 33 days total.

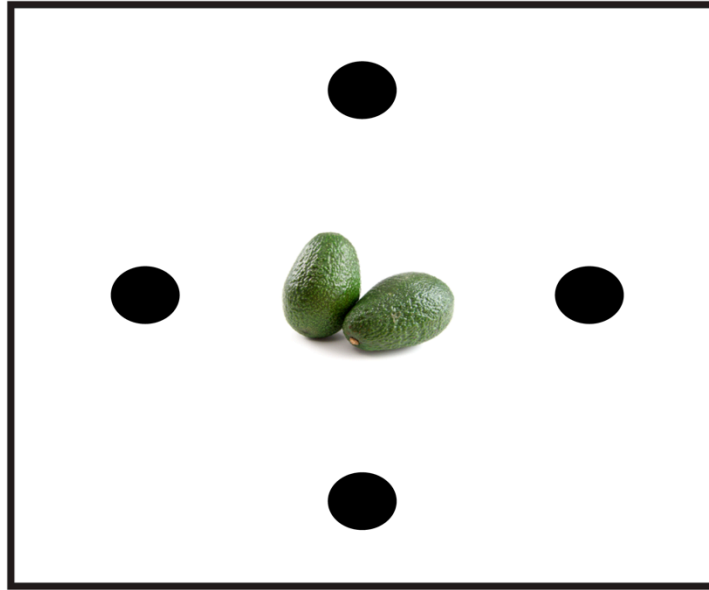


Figure 5-1 An example stimulus within the layout of the saccade-learning task. In this example, the image presented is that of a lovely pair of avocados. During each round of training, the monkey is challenged with learning and recalling the correct saccade direction associated with each image. Mmmm... avocados.

Phase 2

In the next phase (see Figure 5.2), no training took place and responses to norm-based face space stimuli (see Figure 7.2) were recorded for 18 consecutive days. This served as the familiarization phase, during which responses to repeated presentation of the same stimuli were observed. Because there was very little evidence of changes during this period, phase 2 responses also served as a baseline, against which changes following training could be compared.

Phase 3

In the third phase, we continued to record responses to the same face stimuli each day. But, recording on each day was then followed by conditioning training (as described in Phase 1) using a subset of faces contained within the (now familiar) face stimuli (see Figure 7.3). This took place over 27 recording sessions, spanning 35 days total.

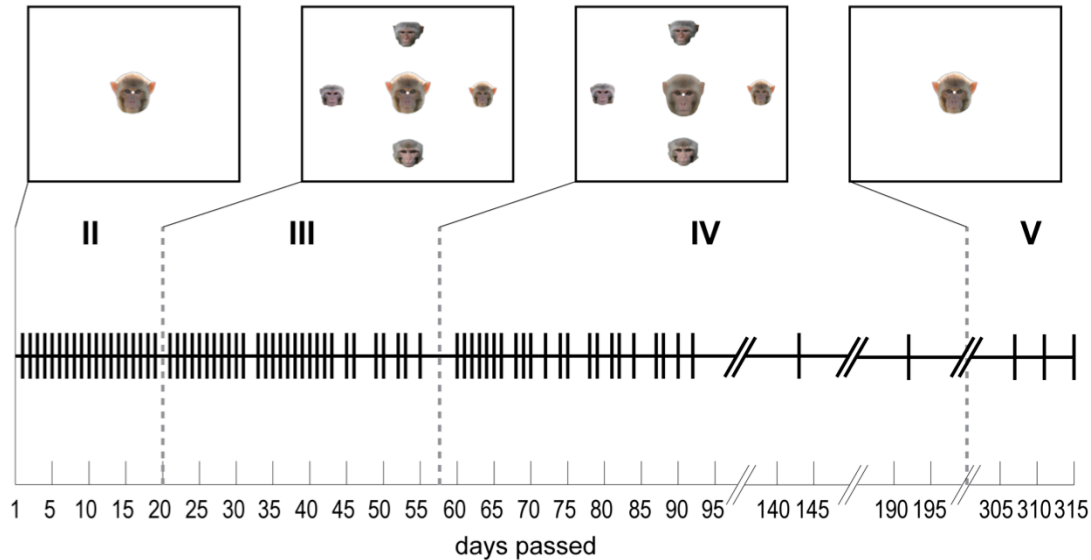


Figure 5-2 Experimental timeline of face-learning study. The transitions between phases of the study is demarcated by the dotted gray line. The layout of the saccade learning task for each phase is shown in the box above each phase. Each vertical black line represents a single day of recording (phases II and V) or training/recording (phases III and IV).

Phase 4

In the fourth phase of expertise training, each face stimulus was adaptively morphed closer to the average face, thus suppressing the distinctive features of the face (adaptive staircase rules: 2-down, 1-up rule; Levitt, 1971). In this way, the difficulty of the task was adaptively maintained as the monkey became increasingly skilled at the task, providing a measure of the monkey's perceptual threshold for identity. This took place over 22 recording sessions spanning 33 days total.

Phase 5

In three additional follow-up recording sessions, responses to norm-based face space stimuli (as in Phase 2) were recorded 11 months and 9 days after the study began to look for changes over this extended time scale.

2. Analysis

The following section will discuss the specific analyses that were performed in order to compare neural responses to conditioned versus un-conditioned stimuli.

These analyses were performed on the responses recorded on each day, before behavioral training. Consequently, the results in the next section will not include the responses recorded during behavioral training.

Latency- The response latency for a particular stimuli was defined as the first sample after the z-score of the spike-density function exceeds 1.96, which corresponds to a confidence interval of approximately 95%.

Linear de-classifications- The accuracy of a linear classifier in correctly predicting the stimulus presented, based on the population response, was measured on each day. The models were cross-validated using the leave-one-out approach. With this approach, the model is trained with $(N-1)/N$ responses recorded in a given day, and then validated with the response that was left out. This process was repeated, or “folded”, many times over until a reliable indicator of validity is achieved.

Noise correlations- Residual response vectors were created by subtracting the mean response to each stimuli from the trial-by-trial responses to those stimuli for each day. Noise correlation matrices measured the correlation between the residual response vectors for each neuron with that of every other neuron. The variance of the noise correlation matrices were measured on each recording day.

Signal correlation- Signal correlation matrices, on the other hand, were created by comparing the mean response vectors for each neuron with every other neuron, and measuring the mean of the signal correlation matrix measured on each recording day.

II. Results

The results that follow are grouped into two main sections. The first section includes the behavioral results of the conditioning paradigm described in the previous

section. The second part of this results section concerns the influence of the conditioning paradigm on the responses of AF face-neurons to the conditioned versus un-conditioned face stimuli, the analysis of which were outlined in the previous section.

1. Behavioral results

In order to elicit changes in their ability to distinguish the faces of particular individuals, we engaged the monkey in the face-training task for 51 days. It was expected that over this time, behavioral performance would gradually improve as the relationship between a given face identity and the associated saccade direction was repeatedly reinforced. In stage III (Fig. 5-3, middle), the task required the monkey to make a visual saccade to the correct target location on the screen. The face was presented in its original (full-identity) form, without any manipulation of its distinctiveness relative to the average. The number of possible target locations was 2 for the first 6 days (chance performance – shown in grey: 50%), and 4 for the remainder of stages III and IV (chance performance: 75%). While there were fluctuations in performance, particularly earlier in the process of learning the task, there is also a clear increase in the percentage correct during stage III.

In stage IV, the task was changed to modulate the identity level (distance from the norm in norm-based face space) based on the performance during each trial (2-down, 1-up rule). Essentially, the percentage correct was being controlled by manipulating the distinctiveness (i.e. morph level) relative to the average face, of the presented face. As the monkey's expertise continued to increase, distinctiveness (i.e.

morph level) continued to decrease slowly throughout stage IV following an initially rapid increase in performance.




Figure 5-3 Behavioral performance throughout phases II-IV of associational learning task. Phase II (left-most portion) served as the baseline phase, against which later responses could be compared, and so no training occurred. In phase III (center portion), the visual saccade task was begun with first 2, then 4 full-identity monkey faces. Chance performance level for this task is shown in grey, and the actual performance of the monkey on this task (% correct) is shown in blue. In phase IV, task difficulty was made dependent on performance by decreasing the identity level (red) of the presented face following each successful trial.

These results suggest that this task was successful in evoking consistent improvements in face-identification behavior over a period of 51 days. Such consistent improvements are likely to reflect a continuous improvement in expertise, with the specific distinctive features associated with the faces being conditioned in this task.

In order to more carefully describe the relationship between the level of distinctiveness of a presented face and the perception of face identity by the monkey, a mean psychometric response function for all trials in stage IV was derived (see Figure 5.4A) in addition to the mean response function for each day of training (see Figure 5.4B).

As a consequence of previous training, the monkey in this study developed an idiosyncratic preference for right-ward saccades, when performing this task. As such,

all right-ward saccades (both correct and incorrect), were removed before performing this analysis in order to correct for this bias. As expected for a 3-alternative choice task, the lower asymptote of the mean function is very near to the predicted chance level of 33.3% (33.89%). For the lowest identity levels of the presented face (0% identity), the monkey's performance remained very close to the level predicted by chance alone. This suggests that at this level, the correct face identity is almost never correctly identified and therefore the probability of correct responses is very near chance level. On the other hand, when the presented identity level exceeds 60%, there was an approximately 85% chance of responding with a saccade to the correct target location. The sensory threshold, defined as the point at which the function reaches the middle between chance (33.3%) and 100%, to be 14.3% for the monkey in this study, roughly similar to thresholds observed previously (17.3%; Leopold et al., 2006).

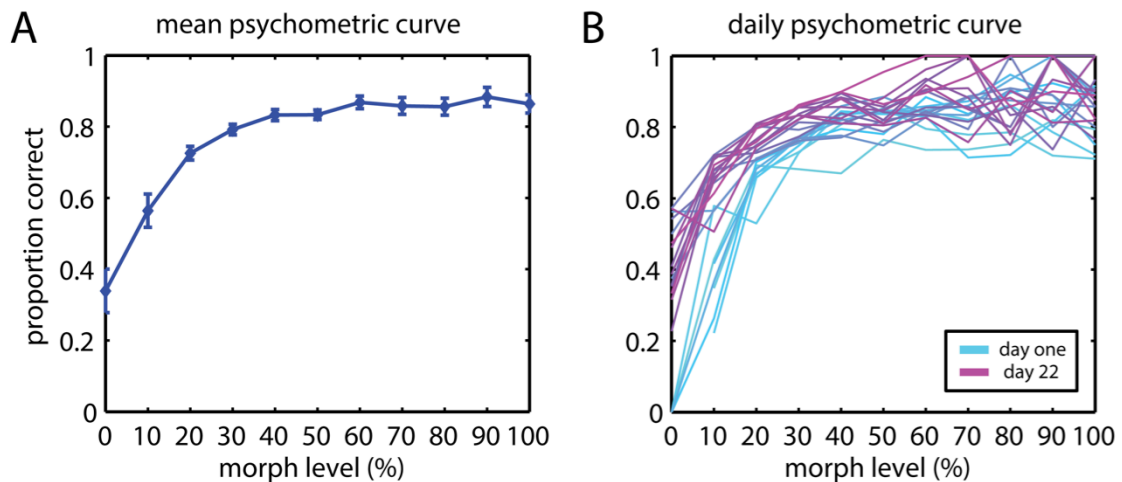


Figure 5-4 Psychometric curves showing relationship between the face stimulus and the behavioral outcome of the oculomotor learning task. (A) The relationship between identity level of the presented face and the proportion of trials in which the correct saccade target was chosen over all days. Error bars: s.e.m. (B) The same relationship, shown for all 22 days of training.

When comparing the psychometric curve for each day, the curve was fit to a psychometric function using a maximum-likelihood criterion. As shown in Figure 5.4B, the gradual improvement in this task is evident in the down-ward shift of the

daily psychometric curve over the duration of training. This further supports the conclusion that this task brought about gradual, steady improvement in this monkey's ability to discriminate the distinctive features of the conditioned face-stimuli within this paradigm.

2. Neurophysiological results

Single-unit-

It could be expected that changes in selectivity would be evident in the evoked responses of individual neurons to conditioned versus un-conditioned stimuli over time, as a result of either familiarization or training. In order to look for changes in the responses of AF face-neurons, which might be associated with either of these influences, the evoked firing rate and response latency for several stimulus categories (object stimuli, human faces, unconditioned or “control” monkey face stimuli, and conditioned or “experimental” monkey face stimuli) were assessed for each recording day. Initial comparison of the response to several groups of stimuli revealed no obvious, systematic change associated with the onset or transition between each phase of training paradigms (vertical, dotted lines- Figure 5.5A top). The background-subtracted responses of each neuron to the 4 groups of visual stimuli, were averaged for each recording day (see Figure 5.5A top). Aside from a slow, incremental decay in the response amplitude that occurred over the span of 315 days, no apparent changes in the relative responses to any of these groups occurred throughout this time. When we compare this pattern to the responses to individual stimuli (see Figure 5.5A bottom, left) no systematic changes in evoked firing rate are observed, and

appears to be similar for responses to the entire stimulus set (see Figure 5.5A bottom, right).

When comparing the slope of the evoked firing over time for each of the 26 stimuli presented, the slopes were typically negative (see Figure 5.5A bottom, right). To look more closely at the responses properties of AF face-cells to individual stimuli throughout the course of behavioral conditioning, the stimuli associated with the greatest changes in response rate over time (i.e. greatest slope for evoked response rate over time) was examined for each neuron (see Figure 5.5B). Although, a distinct drop off is apparent in the final days of experiment, these would seem to be well accounted for by the slow decay in response amplitude that was evident throughout the lifetime of the electrode array implantation. As such, changes in selectivity brought about by either familiarization or learning were not evident in the evoked responses to conditioned versus un-conditioned stimuli over time in individual neurons.

It might also be expected that either familiarization- or learning-induced changes in might be observed in the timing of responses to conditioned versus un-conditioned stimuli over time. In order to investigate this, a similar approach was then used to examine the latency of responses to visual stimuli.

Similarly, no obvious changes were evident in the latency of responses to either control (objects, human faces, or control monkey faces) or experimental (trained monkey faces) stimuli, throughout the course of the experiment (see Figure 5.5C top). When comparing the slope of the latency over time for each of the 26 stimuli presented, the slopes were typically distributed around zero (see Figure 5.5C

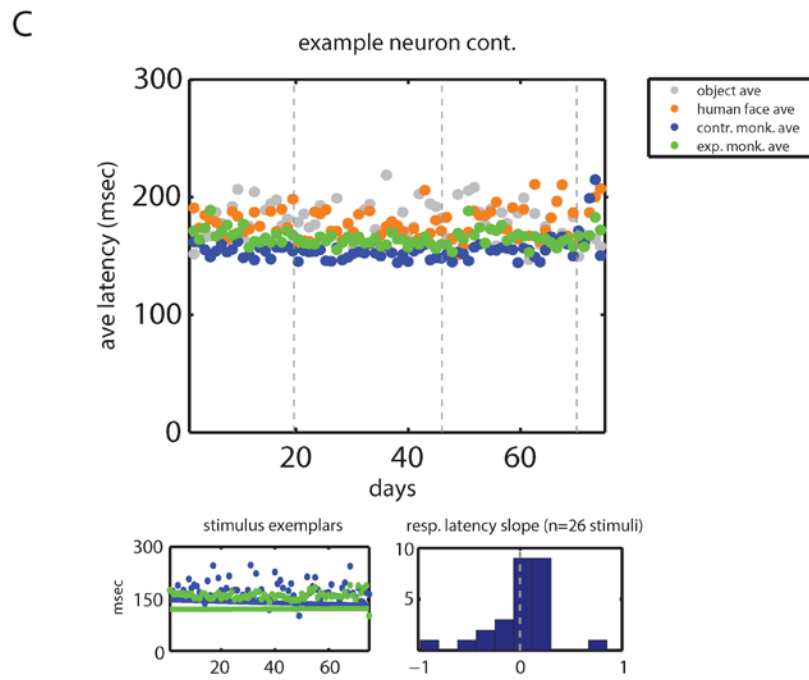
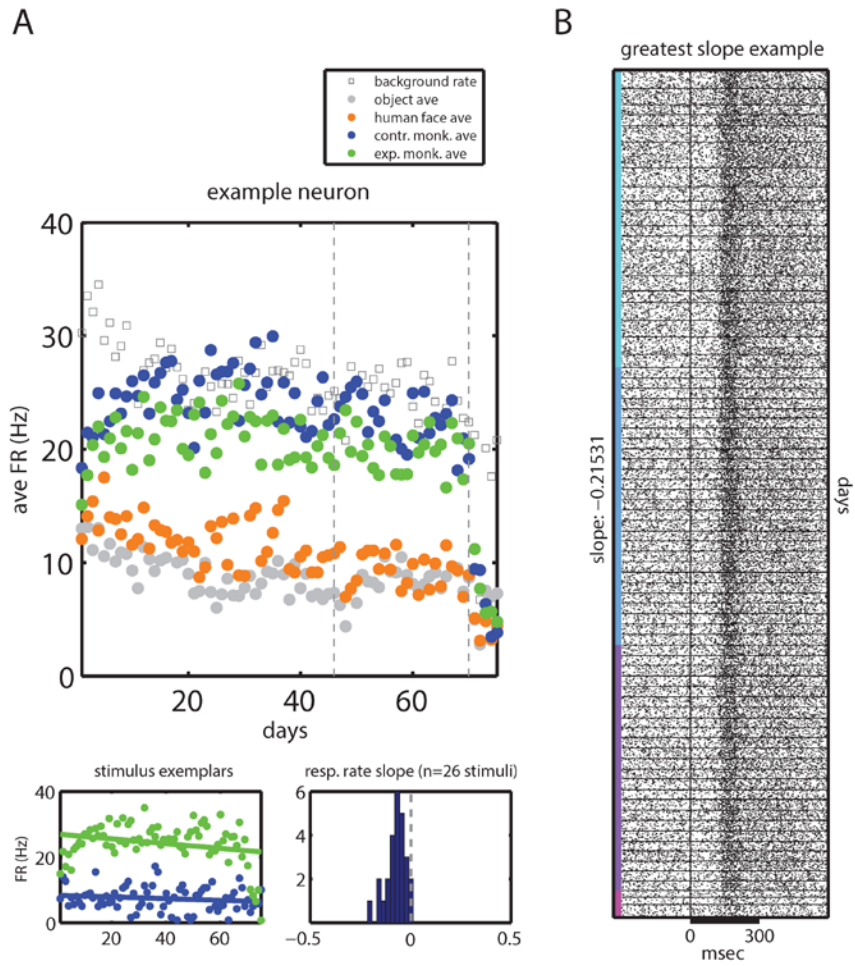


Figure 5-5 Got any change? Screening for plasticity throughout learning. (A) Top: response rates of an example AF neuron throughout the course of this study. Bottom, left: responses to individual stimuli (two monkey faces) plotted over the same time. Bottom, right: slope of responses to each stimulus as a function of the same time (26 total stimuli). (B) Response of responses to the stimuli that the greatest absolute slope value in A. Rasters show responses to 300-ms presentation of stimuli (black bar) on each day (thin black lines). The study phase of raster plot is indicated by the colored bar on the left of each panel. Moving from top-to-bottom: phases II-IV. (C) Top: response latencies of an example AF neuron throughout the course of this study. Bottom, left: responses latencies for individual stimuli (two monkey faces) plotted over the same time. Bottom, right: slope of responses latencies for each stimulus as a function of the same time (26 total stimuli).

bottom, right). Despite steady behavioral improvements over several weeks, it seems that the single-unit representations of stimuli, in terms of response amplitude and latency, remain stable throughout the course of the face training paradigm.

Population: It has been suggested that neuronal changes related to learning may take the form of distributed changes among a population of neurons, rather than being limited to a few isolated neurons. In order to explore this possibility, several methods of population analysis, which are discussed within the following section, were used to compare the response properties of the entire population of AF face-neurons on each recording day.

i. Linear declassifications – A linear classifier allows one to categorize the responses of a population of neurons, and attempt to predict the stimuli that evoked those responses, assuming that those responses are linearly separable. If changes were occurring at the population level response, as a result of learning, it might be expected that those changes would be evident in the performance of a linear classifier, in distinguishing responses to trained versus untrained stimuli. In order to look for changes in the separability of population responses to trained versus untrained stimuli, we tested the performance of a linear classifier on each recording day (see Figure 5.6C). Initial classification accuracy of a linear classifier, based on population responses (previous to training) was .65, and remained near this level throughout the remainder of training. There were no substantial changes in the classification of

trained versus untrained stimuli, relative to the particular stage of training. This suggests that, while a large portion- approximately 65%- of responses could be accurately predicted by a linear de-classifier, this factor does not seem to change noticeably as a consequence of learning.

ii. Noise correlations – Noise correlations reveal relationships in the fluctuations of responses across the population of neurons (Averbeck et al., 2006). If changes were occurring in the population-level response, as a result of learning, it might be expected that those changes could be reflected in the fluctuation of the population response. In order to probe for population-level changes in neuronal variability, noise correlations were derived from the responses to trained versus un-trained stimuli for each recording day (see Figure 5.6D). Pre-training correlation coefficients values were .05 (\pm .003), and remained near this level throughout the remainder of training (see Table 5.1). There were no substantial changes in the noise correlation of responses to trained versus untrained stimuli, relative to the onset or offset of a particular stage of training. This suggests that while this group of AF face-cells do exhibit a similar synchronization to that seen in other cortical regions, this factor also does not exhibit noticeable change as a consequence of learning.

iii. Signal correlations – Signal correlations reflect correlations in the average response of a population of neurons (Averbeck et al., 2006). In order to look for changes in the similarity of tuning curves that might be associated with learning, signal correlations were derived from the responses to trained versus un-trained stimuli for each recording day (see Figure 5.6E). It might be expected that changes in selectivity, that might occur as a result of learning, could be reflected as changes in

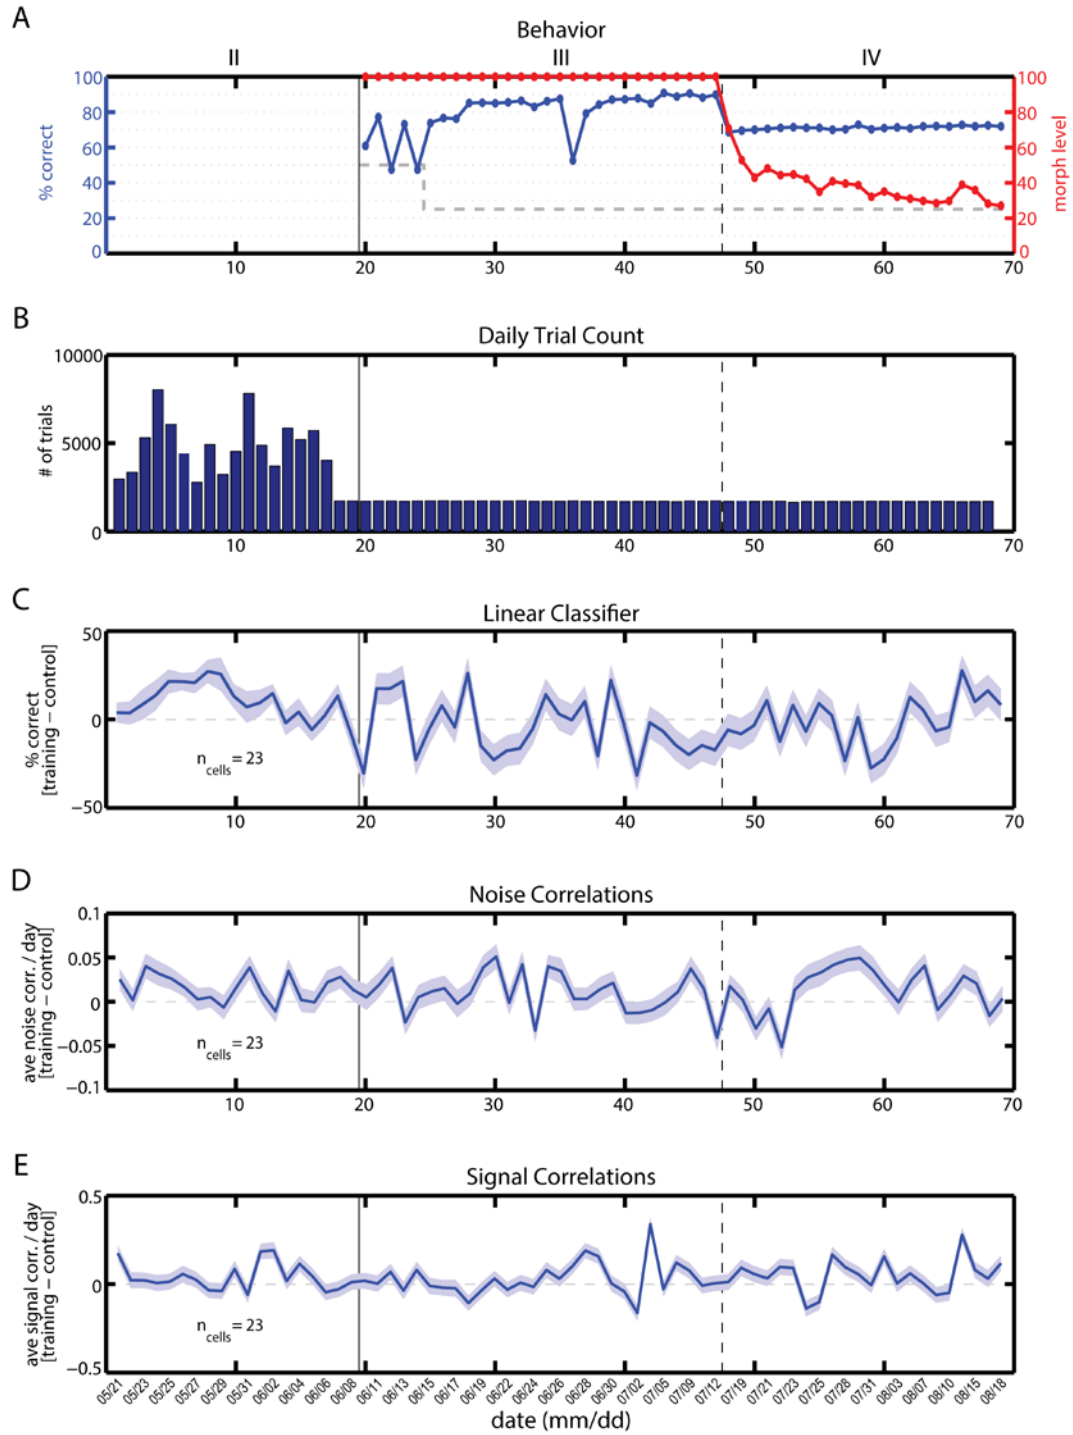


Figure 5-6 Longitudinal population responses (A) Behavioral performance throughout phases II-IV of associational learning task (as in Figure 5-3). (B) Number of daily trials is shown for each day of the same time. Trial count relied on subject performance during phase II, but remained fixed throughout phases III-IV. (C) Performance of linear declassifier on categorization of trained versus untrained stimuli based on population responses, as a function of recording day (shaded area represents the standard error of the difference between means [training vs. control]). (D) Noise correlations derived from population responses to trained versus un-trained stimuli for each recording day. (E) Signal correlations derived from population responses to trained versus un-trained stimuli for each recording day.

the similarity of tuning curves across a population. Pre-training correlation coefficients values were .05 (\pm .01), and remained near this level throughout the remainder of training (see Table 5.1). There were no substantial changes in the signal correlation of responses to trained versus untrained stimuli, relative to the onset or offset of a particular stage of training. This suggests that while this group of AF face-cells do exhibit a level similarity of tuning similar to that seen in other cortical regions, this factor also does not exhibit noticeable change as a consequence of learning.

Table 5-1 Population analyses results of face-training paradigm

		Phase II [mean (std. error)]	Phase III	Phase IV
Linear declassifier performance (% corr)	Training stimuli (T)	.68 (\pm .02)	.47 (\pm .01)	.55 (\pm .02)
	Control stimuli (C)	.59 (\pm .02)	.49 (\pm .01)	.56 (\pm .02)
Noise correlations (r)	T	.05 (\pm .003)	.05 (\pm .003)	.05 (\pm .004)
	C	.04 (\pm .002)	.04 (\pm .004)	.03 (\pm .004)
Signal correlations (r)	T	.05 (\pm .01)	.05 (\pm .01)	.06 (\pm .01)
	C	.02 (\pm .01)	.03 (\pm .009)	.02 (\pm .01)

III. Discussion

Healthy adults maintain the ability to learn and remember new visual objects, particularly faces, throughout most of their lives. Our ability to make new associations between certain outcomes and novel visual stimuli must rely on plasticity in some part of the brain. But, due to limitations in the timespan of current *in-vivo* recording technique, changes such as these have not been well explored in individual neurons. Previous studies attempting to investigate such plasticity have typically

employed one of two designs; sampling separate neurons before and after training (Sakai & Miyashita, 1991) and comparing between responses to trained and untrained stimuli after training has taken place (Tanaka, 1996; Logothetis & Pauls, 1995). By expanding the duration of electrophysiological recordings to span greater time scales, we were able to observe single neurons throughout the course of several months of behavioral training.

In the current study, we focused on the responses of AF face-neurons throughout the course of an intensive behavioral conditioning paradigm, intended to simulate the process of meeting a new person and developing a robust representation of their identity through continued exposure. Despite evidence of a consistent increase in the perceptual skill of this animal, observation for more than 300 days provided no indication of experience-dependent plasticity in the neurons that were observed here.

This outcome may have been influenced by a variety of factors. As suggested previously, it remains possible that learning induced plasticity of face representations occur in a different face patch than we explored. Perhaps the changes in the brain that underlie the observed changes in behavior weren't even attributable to the face patch system. Although a recent optogenetic study of face-neurons in nearby IT have shown that deactivating this population leads to detriments in face identification performance (Afraz et al., 2015), which suggests an important role for these cells in supporting face-processing behavior. However, until similar studies are performed in this particular region of the brain, we cannot know the role of that these neurons play in this process, for sure.

In addition, it may be that the changes associated with learning are limited to a very small subset of neurons, and our recordings were too sparsely-sampled to observe them. For example, if there is any reality to the notion of a “grandmother cell” (or a “Halle Berry neuron” - Quiroga et al., 2005), in which single neurons are dedicated to representing a particular person in a variety of different situations, we may simply be failing to target the right neurons sparsely representing the individual faces presented within this study.

Finally, it remains possible that our design was flawed, and that the behavior did not sufficiently represent the natural behavior of the animal when becoming acquainted with a new individual. Our finding, regarding larger responses for conspecific faces, may suggest that this area is highly sensitive to the behavioral-relevance of a face. Perhaps more realistic, or dynamic stimuli may have been optimal to observe more naturalistic learning in this population, such as telemetric recording in freely behaving monkeys.

This study established, for the first time, that the response patterns of AF neurons remain stable throughout the training that has been described. But, this stability would also seem to indicate a similar level of stability in the afferent visual areas, as well. Otherwise, we would expect the variability of afferent neurons to be reflected in the behavior of AF neurons. This final result suggests that if population-level changes are responsible for the improvements in behavior that were observed here, they must be occurring further along in the face-processing pathway such that those changes are not reflected in single neurons in the AF face patch. Possible candidates include efferent regions such as prefrontal cortex and medial-temporal lobe,

both of which have demonstrated plasticity on time scales appropriate for these behaviors (Squire et al., 2004).

Chapter 6: Conclusions

Accurately reading and remembering faces is a key component to many primate social behaviors. As such, networks of face-specialized clusters, or patches, have been found in the high-level visual systems of a growing number of primate species (humans- Kanwisher et al., 1997, macaques- Tsao et al., 2003, marmosets- Hung et al., 2015). These specialized networks are likely to be responsible for our unique use of faces as our primary communicative organ (Weiner & Grill-Spector, 2015; but see Leopold & Rhodes, 2010 for non-primate exceptions). Due to the complex, multi-dimensional selectivity of neurons in these areas, characterizing their precise selectivity has been a slow and difficult process. However, a picture of the functional organization of the components within this network is starting to emerge, which both conflicts and agrees with previous observations and assumptions. In order to include the properties of the AF face patch into this developing conception of the face processing network, we used a novel technique that also afforded the ability to track neurons over longer periods of time than previously possible. The benefits of such an investigation include: 1) observing the dynamics of stimulus selectivity over timeframes that have not yet been explored, and 2) comparing and contrasting the functional properties of neurons in this area, with those of regions that have already been explored. The remaining sections review and interpret the most substantial findings from the previous sections, and address how each finding compares with what might have been expected, based on the conclusions of previous studies.

1. Longitudinal recordings

It might be argued that while substantial investments have been made by neuroscientists in order to permit the recording of large numbers of neurons simultaneously, comparatively little has been invested in technology to compare their response properties over time. In Chapter 3, the feasibility of addressing new types of experimental questions with longitudinal recordings is established. It was shown that using this technique, it is possible to stably isolate single neurons for weeks and months at a time, thus permitting the exploration of questions otherwise beyond the reach of conventional single-unit recording methods. This represents an important advance as it opens a new avenue for the investigation of a wide-variety of dynamic neural properties over behaviorally-relevant timescales.

In the same chapter, we exploit this novel approach by examining the consistency of stimulus-selectivity exhibited by neurons in the anterior temporal lobe face patch AF. It was shown that these neurons exhibit a high degree of day-to-day consistency throughout the course of repeated presentations of stimuli, even over timescales of many months. The stable response patterns that were observed contrasts with the findings of acute-recording studies, which have found that the responses of temporal lobe neurons diminish as stimuli grow more familiar (Fahy et al., 1993; Freedman et al., 2006; Li et al., 1993; Woloszyn & Sheinberg 2012).

One possible explanation for the discrepancy in the results of each approach may be differences in the stability of recordings from individual neurons. It may be that the decrement in responses observed using acute recording is actually a result of unavoidable drifting of the electrode position, relative to the neuron under

examination. As mentioned previously, another explanation may be that this stability represents a specialized function of the particular face patch we happened to be recording from. Although clustering of face-selective neurons had been observed by previous studies of temporal cortex (Perrett et al. 1992), the organization of the face patch network was not revealed until 2003. As a result, previous studies generally targeted neurons located comparatively broadly throughout IT and STS. It remains possible that familiarization-effects, similar to those described by previous studies, may be present within other regions within the face patch network. Until similar longitudinal recordings are performed in additional areas, we cannot know for certain.

II. Norm-based encoding of face identity

1. Familiarity effects

Numerous psycho-physical studies have found that the perception of a face's identity relies on an internal, relativistic comparison between it and the prototypical average face (Leopold et al., 2001; Rhodes et al., 2005, but see Ross et al., 2014). One potential criticism that has been leveled against this conclusion is that these observations are simply the result of within-session adaptation of these neurons to the statistics of the recently-presented stimuli. Due to the limited recording time permitted by traditional recording techniques, it has not been possible to explore this possibility by observing responses from the first to last presentation of a given stimulus.

In Chapter 4, we constrained our analysis to focus on responses within limited temporal windows and showed that the responses of AF face-cells within each temporal window that was analyzed show a similar sensitivity in their responses to

distinctive faces – such that responses to faces were sensitive to a the relationship between each face and the average for that species- even when comparing only the responses to the first presentation of stimuli on each recording day. This suggests that this pattern of norm-based selectivity is not easily accounted for by the effects of either within- or between-session adaptation.

2. Species-selectivity

In addition, AF neurons were found to be more sensitive to changes in the distinctiveness of the faces of monkeys than for the faces of humans. This runs counter to the findings of previous studies, which have found similar responses to the faces of humans, monkeys, or even greatly simplified cartoon faces (Freiwald et al., 2009).

Such a species-specific effect might offer distinct advantages in terms of the efficiency and flexibility of encoding the faces of different species. Previous electrophysiological studies have shown prototype-based tuning in the preference of neurons for multiple object categories (Kayaert et al., 2005), suggesting that multiple population averages may be represented simultaneously. By relying on a species-specific average face, greater flexibility may be achieved by having the ability to shift the prototype for each species independently following experience with members of different species within one's environment (conspecifics, predators, prey, etc.), rather than relying on a single population average for all faces, in general.

While the precise nature of the norm – how many there and for which categories they exist – remains unclear, these results suggest that it's an important factor that shapes the responses of face-neurons in the anterior temporal lobe. By

linking perceptual and neurophysiological models, such as these, we are uncovering the cellular mechanisms behind the unique ability of the primate brain to efficiently extract information from faces.

III. Selectivity for size

Face-neurons in the temporal lobe have previously been shown to respond similarly across different image sizes, positions, and orientations (Perrett et al., 1982). Such “invariant” response properties fit with the notion that selectivity for complex objects is simply an extension of the hierarchy of visual selectivity that was already proposed by Hubel and Wiesel following their studies of primary and secondary regions of the occipital cortex (Hubel & Wiesel, 1962). Within such a hierarchy, increasing receptive-field complexity would be inherited by pooling the output of cells with relatively simple receptive-fields. Ultimately, complex object recognition would be accomplished by transforming retinal input into increasingly abstract representations, which become more invariant to the physical properties of the stimulus (Farah, 2004). This hierarchical view is longstanding, and supported by observations such as the gradually increasing receptive field size, onset latencies, and complexity of stimulus selectivity observed in visual neurons sampled along on the ventral visual stream (Kravitz et al., 2013).

In Chapter 4 it was shown that the neurons we observed, in addition to being sensitive to identity, were also highly sensitive to the sizes of faces and bodies. When considered together, the results of Chapter 4 suggest that AF face-neurons participate in the discrimination of not only face identity, but also in the size of the retinal image projected by that face. Since, under ecologically-valid conditions, variation in the size

of a face is mainly associated with the perception of distance rather than head size, our results demonstrate that AF neurons carry signals that may be useful for extracting spatial relations between an animal and their conspecifics nearby. In the context of natural viewing, therefore, this joint sensitivity to face identity and size may enable AF neurons to act as “face-distance” encoding cells. It may be reasonable to speculate that AF face-neurons may play a role in representing spatial aspects of biological vision -- a property more traditionally associated with the forms spatial visual processing present in the dorsal visual stream (Mishkin et al., 1983).

IV. Learning

Understanding how the “unique” properties of the face processing system develop remains a critical goal in efforts to provide useful interventions in the case of psychiatric disorders affecting social processing (e.g. autism and schizophrenia). Historically, substantial debate has existed regarding whether stereotyped face processing modules are innately wired to represent faces (see Kanwisher, 2000) or are adapted, through learning, to represent any set of over-trained stimuli, including but not limited to faces (see Gauthier et al., 2000). While the correct answer is likely to include some combination of both, the dynamics of the neural circuitry underlying this process have yet to be observed over the timescales relevant to such development.

In Chapter 5 it was shown that the stimulus selectivity of AF neurons was insensitive to behavioral training intended to mimic the development of such perceptual expertise. Despite steady improvement in discriminating the distinctiveness of conditioned faces, neurons in this region exhibited stable selectivity

for both conditioned and unconditioned stimuli over a period of nearly one year in the adult macaque. As suggested by the results of Chapter 3, these results may indicate a specialized feature of AF face-neurons. But, further investigation within additional face patch regions is necessary, before this can be known for sure.

The observation of learning-induced changes in individual face-cells still remains a critical milestone in our progress towards understanding the development of such highly-specialized processes in the primate brain. This knowledge will allow us to challenge current assumptions regarding how we, as primates, observe, identify, and eventually categorize objects such as faces, and how learning play a role in shaping our subsequent experiences and interactions with visual objects, such as faces, throughout our lifetime. More broadly, these results will offer insights into the core mechanisms by which we, as humans, look upon others and how measurable changes in the brain impact our perceptions over time.

V. Future experiments

Studies of the face patch system in macaques have provided clear evidence of functional specialization within the individual patches of the face-processing network. In order to explore the extent to which these results extend to other regions in the face patch system, similar longitudinal recording will need to be performed in other face patch regions using comparable methodology. By placing the results described in this volume in context with observations from other face-selective patches, it should be possible to identify whether the stability of AF neurons is “special”, or whether this a more general feature of visual neurons that has not yet been identified, due to limitations in more traditional neurophysiological recording techniques.

Appendices




Figure 7-1: Category Selectivity Stimuli




Figure 7-2: Identity Selectivity Stimuli




Figure 7-3: Training Stimuli




Figure 7-4: Caricature Stimuli



Figure 7-5: Category Selectivity Stimuli

Bibliography

- Afraz, A., Boyden, E. S., & DiCarlo, J. J. (2015). Optogenetic and pharmacological suppression of spatial clusters of face neurons reveal their causal role in face gender discrimination. *Proceedings of the National Academy of Sciences*, 112(21), 6730-6735.
- Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature reviews: neuroscience*, 7(5), 358-366.
- Bahrick, H. P., Bahrick, P. O., & Wittlinger, R. P. (1975). Fifty years of memory for names and faces: A cross-sectional approach. *Journal of experimental psychology: General*, 104(1), 54.
- Baker, C. I., Behrmann, M., & Olson, C. R. (2002). Impact of learning on representation of parts and wholes in monkey inferotemporal cortex. *Nature neuroscience*, 5(11), 1210-1216.
- Benson, P. J., & Perrett, D. I. (1994). Visual processing of facial distinctiveness. *Perception*, 23, 75-93.
- Bondar, I. V., Leopold, D. A., Richmond, B. J., Victor, J. D., & Logothetis, N. K. (2009). Long-term stability of visual pattern selective responses of monkey temporal lobe neurons. *PloS one*, 4(12), e8222.
- Boettiger, C. A., & Doupe, A. J. (2001). Developmentally restricted synaptic plasticity in a songbird nucleus required for song learning. *Neuron*, 31(5), 809-818.
- Bruce, C. (1982). Face recognition by monkeys: absence of an inversion effect. *Neuropsychologia*, 20(5), 515-521.
- Carbon, C. C., Strobach, T., Langton, S. R., Harsányi, G., Leder, H., & Kovács, G. (2007). Adaptation effects of highly familiar faces: Immediate and long lasting. *Memory & Cognition*, 35(8), 1966-1976.
- Chang, E. F., & Merzenich, M. M. (2003). Environmental noise retards auditory cortical development. *Science*, 300(5618), 498-502.
- de Beeck, H. P. O., & Baker, C. I. (2010a). The neural basis of visual object learning. *Trends in cognitive sciences*, 14(1), 22-30.
- de Beeck, H. P. O., & Baker, C. I. (2010b). Informativeness and learning: Response to Gauthier and colleagues. *Trends in cognitive sciences*, 14(6), 236.
- Diamond, R., & Carey, S. (1986). Why faces are and are not special: an effect of expertise. *Journal of experimental psychology*, 115(2), 107.

- Elnaiem, H. D., McMahon, D. B. T., Russ, B. E., & Leopold D. A. (2013). Longitudinal investigation of IT cortex: Responses to naturalistic movie stimuli. *SfN Abstracts*, 555.18/VV22
- Fahy, F. L., Riches, I. P., & Brown, M. W. (1993). Neuronal activity related to visual recognition memory: long-term memory and the encoding of recency and familiarity information in the primate anterior and medial inferior temporal and rhinal cortex. *Experimental brain research*, 96(3), 457-472.
- Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. N. (1998). What is "special" about face perception? *Psychological review*, 105(3), 482.
- Farah, M. J. (2004). *Visual agnosia*. MIT press.
- Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2006). Experience-dependent sharpening of visual shape selectivity in inferior temporal cortex. *Cerebral cortex*, 16(11), 1631-1644.
- Freiwald, W. A., Tsao, D. Y., & Livingstone, M. S. (2009). A face feature space in the macaque temporal lobe. *Nature neuroscience*, 12(9), 1187-1196.
- Freiwald, W. A., & Tsao, D. Y. (2010). Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science*, 330(6005), 845-851.
- Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform 'face area' increases with expertise in recognizing novel objects. *Nature neuroscience*, 2(6), 568-573.
- Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature neuroscience*, 3(2), 191-197.
- Gauthier, I., Wong, A. C. N., & Palmeri, T. J. (2010). Manipulating visual experience: comment on Op de Beeck and Baker. *Trends in cognitive sciences*, 14(6), 235.
- Germine, L. T., Duchaine, B., & Nakayama, K. (2011). Where cognitive development and aging meet: Face learning ability peaks after age 30. *Cognition*, 118(2), 201-210.
- Giese, M. A., & Leopold, D. A. (2005). Physiologically inspired neural model for the encoding of face spaces. *Neurocomputing*, 65, 93-101.
- Gilbert, C. D., Sigman, M., & Crist, R. E. (2001). The neural basis of perceptual learning. *Neuron*, 31(5), 681-697.

- Ghirlanda, S., Jansson, L., & Enquist, M. (2002). Chickens prefer beautiful humans. *Human nature*, 13(3), 383-389.
- Going, M., & Read, J. D. (1974). Effects of uniqueness, sex of subject, and sex of photograph on facial recognition. *Perceptual and motor skills*, 39(1), 109-110.
- Golarai, G., Ghahremani, D. G., Whitfield-Gabrieli, S., Reiss, A., Eberhardt, J. L., Gabrieli, J. D., & Grill-Spector, K. (2007). Differential development of high-level visual cortex correlates with category-specific recognition memory. *Nature neuroscience*, 10(4), 512-522.
- Gross, C. G., Bender, D. B., & Rocha-Miranda, C. E. (1969). Visual receptive fields of neurons in inferotemporal cortex of the monkey. *Science*, 166(910), 1303-1306.
- Gross, C. G., Rocha-Miranda, C. E., & Bender, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the Macaque. *Journal of neurophysiology*, 35(1), 96-111.
- Gross, C. G. (1994). How inferior temporal cortex became a visual area. *Cerebral cortex*, 4(5), 455-469.
- Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in cognitive sciences*, 4(6), 223-233.
- Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, 36(5), 791-804.
- Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of physiology*, 160(1), 106.
- Hubel, D. H., & Wiesel, T. N. (1970). The period of susceptibility to the physiological effects of unilateral eye closure in kittens. *Journal of physiology*, 206(2), 419-436.
- Hung, C. C., Carlson, E. T., & Connor, C. E. (2012). Medial axis shape coding in macaque inferotemporal cortex. *Neuron*, 74(6), 1099-1113.
- Hung, C. C., Yen, C. C., Ciuchta, J. L., Papoti, D., Bock, N. A., Leopold, D. A., & Silva, A. C. (2015). Functional mapping of face-selective regions in the extrastriate visual cortex of the marmoset. *Journal of neuroscience*, 35(3), 1160-1172.
- Ito, M., Tamura, H., Fujita, I., & Tanaka, K. (1995). Size and position invariance of neuronal responses in monkey inferotemporal cortex. *Journal of neurophysiology*, 73(1), 218-226.
- Jeffery, L., Rhodes, G., & Busey, T. (2007). Broadly tuned, view-specific coding of face shape: Opposing figural aftereffects can be induced in different views. *Vision research*, 47(24), 3070-3077.

- Johnson, M. H., Dziurawiec, S., Ellis, H., & Morton, J. (1991). Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition*, 40(1), 1-19.
- Johnson, S. (1991). Micturition in the canid family: the irresistible pull of the hydrant. *Physics quarterly*, 33, 203-220.
- Jones, A. P., Leopold D. A., & McMahon, D. B. T. (2013). Longitudinal investigation of IT cortex: Probing category selectivity with 10,000 stimuli. *SfN Abstracts*, 555.17/VV21.
- Kahn, D. A., & Aguirre, G. K. (2012). Confounding of norm-based and adaptation effects in brain responses. *Neuroimage*, 60(4), 2294-2299.
- Karni, A., & Sagi, D. (1993). The time course of learning a visual skill. *Nature*, 365(6443), 250-252.
- Kanwisher, N. (2000). Domain specificity in face perception. *Nature neuroscience*, 3, 759-763.
- Kayaert, G., Biederman, I., Op de Beeck, H. P., & Vogels, R. (2005). Tuning for shape dimensions in macaque inferior temporal cortex. *European journal of neuroscience*, 22(1), 212-224.
- Keating, C. F., & Keating, E. G. (1993). Monkeys and mug shots: cues used by rhesus monkeys (*Macaca mulatta*) to recognize a human face. *Journal of comparative psychology*, 107(2), 131.
- Kobatake, E., Wang, G., & Tanaka, K. (1998). Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. *Journal of neurophysiology*, 80(1), 324-330.
- Konorski, J. (1967). *Integrative activity of the brain: An interdisciplinary approach*. Chicago: University of Chicago.
- Kravitz, D. J., Saleem, K. S., Baker, C. I., Ungerleider, L. G., & Mishkin, M. (2013). The ventral visual pathway: an expanded neural framework for the processing of object quality. *Trends in cognitive sciences*, 17(1), 26-49.
- Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature neuroscience*, 4(1), 89-94.
- Leopold, D. A., Rhodes, G., Müller, K. M., & Jeffery, L. (2005). The dynamics of visual adaptation to faces. *Proceedings of the Royal Society of London B: Biological sciences*, 272(1566), 897-904.

- Leopold, D. A., Bondar, I. V., & Giese, M. A. (2006). Norm-based face encoding by single neurons in the monkey inferotemporal cortex. *Nature*, 442(7102), 572-575.
- Leopold, D. A., & Rhodes, G. (2010). A comparative view of face perception. *Journal of Comparative Psychology*, 124(3), 233.
- Leopold D. A., Jones, A. P., & McMahon, D. B. T. (2013). Longitudinal investigation of IT cortex: Delayed emergence of learning-induced plasticity. SfN abstracts, 555.15/VV19.
- Levitt, H. C. C. H. (1971). Transformed up-down methods in psychoacoustics. *Journal of the acoustical society of America*, 49(2B), 467-477.
- Li, L., Miller, E. K., & Desimone, R. (1993). The representation of stimulus familiarity in anterior inferior temporal cortex. *Journal of neurophysiology*, 69(6), 1918-1929.
- Li, N., & DiCarlo, J. J. (2008). Unsupervised natural experience rapidly alters invariant object representation in visual cortex. *Science*, 321(5895), 1502-1507.
- Logothetis, N. K., & Pauls, J. (1995). Psychophysical and physiological evidence for viewer-centered object representations in the primate. *Cerebral cortex*, 5(3), 270-288.
- Logothetis, N. K., Pauls, J., & Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. *Current biology*, 5(5), 552-563.
- Logothetis N. K., & Sheinberg D. L. (1996). Visual object recognition. *Annual review of neuroscience* 19:577–621.
- Malpass, R. S., & Kravitz, J. (1969). Recognition for faces of own and other race. *Journal of personality and social psychology*, 13(4), 330.
- Mankin, E. A., Sparks, F. T., Slayyeh, B., Sutherland, R. J., Leutgeb, S., & Leutgeb, J. K. (2012). Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences*, 109(47), 19462-19467.
- McMahon, D. B., Jones, A. P., Bondar, I. V., & Leopold, D. A. (2014a). Face-selective neurons maintain consistent visual responses across months. *Proceedings of the National Academy of Sciences*, 111(22), 8251-8256.
- McMahon, D. B., Bondar, I. V., Afuwape, O. A., Ide, D. C., & Leopold, D. A. (2014b). One month in the life of a neuron: longitudinal single-unit electrophysiology in the monkey visual system. *Journal of neurophysiology*, 112(7), 1748-1762.
- McMahon, D. B., Russ, B. E., Elnaiem, H. D., Kurnikova, A. I., & Leopold, D. A. (2015). Single-Unit Activity during Natural Vision: Diversity, Consistency, and Spatial

- Sensitivity among AF Face Patch Neurons. *Journal of neuroscience*, 35(14), 5537-5548.
- Meadows, J. C. (1974). The anatomical basis of prosopagnosia. *Journal of Neurology, Neurosurgery & Psychiatry*, 37(5), 489-501.
- Merchant, H., Naselaris, T., & Georgopoulos, A. P. (2008). Dynamic sculpting of directional tuning in the primate motor cortex during three-dimensional reaching. *Journal of neuroscience*, 28(37), 9164-9172.
- Merzenich, M. M., Nelson, R. J., Stryker, M. P., Cynader, M. S., Schoppmann, A., & Zook, J. M. (1984). Somatosensory cortical map changes following digit amputation in adult monkeys. *Journal of comparative neurology*, 224(4), 591-605.
- Mishkin, M., & Pribram, K. H. (1954). Visual discrimination performance following partial ablations of the temporal lobe: I. Ventral vs. lateral. *Journal of comparative and physiological psychology*, 47(1), 14.
- Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: two cortical pathways. *Trends in neurosciences*, 6, 414-417.
- Miyashita, Y. (1988). Neural correlate of visual associative long-term memory in the primate temporal cortex. *Nature*, 335, 817-820.
- Moeller, S., Freiwald, W. A., & Tsao, D. Y. (2008). Patches with links: a unified system for processing faces in the macaque temporal lobe. *Science*, 320(5881), 1355-1359.
- Mondloch, C. J., Geldart, S., Maurer, D., & Le Grand, R. (2003). Developmental changes in face processing skills. *Journal of experimental child psychology*, 86(1), 67-84.
- Ohayon, S., Freiwald, W. A., & Tsao, D. Y. (2012). What makes a cell face selective? The importance of contrast. *Neuron*, 74(3), 567-581.
- Pascalis, O., Scott, L. S., Kelly, D. J., Shannon, R. W., Nicholson, E., Coleman, M., and Nelson, C. A. (2005). Plasticity of face processing in infancy. *Proceedings of the National Academy of Sciences*, 102(14), 5297-5300.
- Perrett, D. I., Rolls, E. T., & Caan, W. (1982). Visual neurones responsive to faces in the monkey temporal cortex. *Experimental brain research*, 47(3), 329-342.
- Perrett, D. I., Smith, P. A. J., Potter, D. D., Mistlin, A. J., Head, A. S., Milner, A. D., & Jeeves, M. A. (1985). Visual cells in the temporal cortex sensitive to face view and gaze direction. *Proceedings of the Royal Society of London B: Biological sciences*, 223(1232), 293-317.

- Perrett, D. I., Hietanen, J. K., Oram, M. W., Benson, P. J., & Rolls, E. T. (1992). Organization and functions of cells responsive to faces in the temporal cortex [and discussion]. *Philosophical Transactions of the Royal Society B: Biological sciences*, 335(1273), 23-30.
- Polosecki, P., Moeller, S., Schweers, N., Romanski, L. M., Tsao, D. Y., & Freiwald, W. A. (2013). Faces in motion: selectivity of macaque and human face processing areas for dynamic stimuli. *Journal of neuroscience*, 33(29), 11768-11773.
- Porada, I., Bondar, I., Spatz, W. B., & Krüger, J. (2000). Rabbit and monkey visual cortex: more than a year of recording with up to 64 microelectrodes. *Journal of neuroscience methods*, 95(1), 13-28.
- Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., & Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045), 1102-1107.
- Rhodes, G., & Jeffery, L. (2006). Adaptive norm-based coding of facial identity. *Vision research*, 46(18), 2977-2987.
- Rhodes, G., & Leopold, D. A. (2011). Adaptive norm-based coding of face identity. *Oxford handbook of face perception*, 263-286.
- Rodman, H. R., Scalaidhe, S. P. O., & Gross, C. G. (1993). Response properties of neurons in temporal cortical visual areas of infant monkeys. *Journal of neurophysiology*, 70, 1115-1115.
- Rolls, E. T., Baylis, G. C., Hasselmo, M. E., & Nalwa, V. (1989). The effect of learning on the face selective responses of neurons in the cortex in the superior temporal sulcus of the monkey. *Experimental brain research*, 76(1), 153-164.
- Ross, D. A., Deroche, M., & Palmeri, T. J. (2014). Not just the norm: Exemplar-based models also predict face aftereffects. *Psychonomic bulletin & review*, 21(1), 47-70.
- Sakai, K., & Miyashita, Y. (1991). Neural organization for the long-term memory of paired associates. *Nature*, 354(6349), 152-155.
- Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, 412(6846), 549-553.
- Schwartz, O., Hsu, A., & Dayan, P. (2007). Space and time in visual context. *Nature reviews neuroscience*, 8(7), 522-535.
- Sigala, N., & Logothetis, N. K. (2002). Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature*, 415(6869), 318-320.

- Sigman, M., Pan, H., Yang, Y., Stern, E., Silbersweig, D., & Gilbert, C. D. (2005). Top-down reorganization of activity in the visual pathway after learning a shape identification task. *Neuron*, 46(5), 823-835.
- Shatz, C. J., & Stryker, M. P. (1978). Ocular dominance in layer IV of the cat's visual cortex and the effects of monocular deprivation. *Journal of physiology*, 281(1), 267-283.
- Squire, L. R., Stark, C. E., & Clark, R. E. (2004). The medial temporal lobe. *Annual review of neuroscience*, 27, 279-306.
- Standing, L. (1973). Learning 10,000 pictures. *Journal of experimental psychology*, 25(2), 207-222.
- Stettler, D. D., Yamahachi, H., Li, W., Denk, W., & Gilbert, C. D. (2006). Axons and synaptic boutons are highly dynamic in adult visual cortex. *Neuron*, 49(6), 877-887.
- Sugita, Y. (2008). Face perception in monkeys reared with no exposure to faces. *Proceedings of the National Academy of Sciences*, 105(1), 394-398.
- Quiroga, R., Nadasdy, Z., & Ben-Shaul, Y. (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural computation*, 16(8), 1661-1687.
- Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *The Quarterly journal of experimental psychology*, 46(2), 225-245.
- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual review of neuroscience* 19:109-39.
- Tanaka, J. W., & Gauthier, I. (1997). Expertise in object and face recognition. *Psychology of learning and motivation*, 36, 83-125.
- Tanaka, J. W., & Farah, M. J. (2003). The holistic representation of faces. *Perception of faces, objects, and scenes: Analytic and holistic processes*, 53-74.
- Thomaz, C. E., & Giraldi, G. A. (2010). A new ranking method for principal components analysis and its application to face image analysis. *Image and vision computing*, 28(6), 902-913.
- Thorpe, S. J., & Fabre-Thorpe, M. (2001). Seeking categories in the brain. *Science*, 260-262.
- Tovee, M. J., Rolls, E. T., & Ramachandran, V. S. (1996). Rapid visual learning in neurones of the primate temporal visual cortex. *Neuroreport*, 7(15-17), 2757-2760.
- Tsao, D. Y., Freiwald, W. A., Knutsen, T. A., Mandeville, J. B., & Tootell, R. B. (2003). Faces and objects in macaque cerebral cortex. *Nature neuroscience*, 6(9), 989-995.

- Tsao, D. Y., & Freiwald, W. A. (2006). What's so special about the average face? *Trends in cognitive sciences*, 10(9), 391-393.
- Tsao, D. Y., Freiwald, W. A., Tootell, R. B., & Livingstone, M. S. (2006). A cortical region consisting entirely of face-selective cells. *Science*, 311(5761), 670-674.
- Tsao, D. Y., Moeller, S., & Freiwald, W. A. (2008). Comparing face patch systems in macaques and humans. *Proceedings of the National Academy of Sciences*, 105(49), 19514-19519.
- Valentine, T., & Bruce, V. (1986). The effects of distinctiveness in recognising and classifying faces. *Perception*, 15(5), 525-535.
- Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *The quarterly journal of experimental psychology*, 43(2), 161-204.
- Van Rensbergen, B., & de Beeck, H. P. O. (2014). The role of temporal context in norm-based encoding of faces. *Psychonomic bulletin & review*, 21(1), 121-127.
- Webster, M. A., & Mollon, J. D. (1991). Changes in colour appearance following post-receptoral adaptation. *Nature*, 349(6306), 235-238.
- Webster, M. A., & Maclin, O. H. (1999). Figural aftereffects in the perception of faces. *Psychonomic bulletin & review*, 6(4), 647-653.
- Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, 428(6982), 557-561.
- Weiner, K. S., & Grill-Spector, K. (2015). The evolution of face processing networks. *Trends in cognitive sciences*, 19(5), 240-241.
- Wiesel, T. N., & Hubel, D. H. (1963). Effects of visual deprivation on morphology and physiology of cells in the cat's lateral geniculate body. *Journal of neurophysiology*, 26(978), 6.
- Woloszyn, L., & Sheinberg, D. L. (2012). Effects of long-term visual experience on responses of distinct classes of single units in inferior temporal cortex. *Neuron*, 74(1), 193-205.
- Yamane, Y., Carlson, E. T., Bowman, K. C., Wang, Z., & Connor, C. E. (2008). A neural code for three-dimensional object shape in macaque inferotemporal cortex. *Nature neuroscience*, 11(11), 1352-1360.

Yin, R. K. (1969). Looking at upside-down faces. *Journal of experimental psychology*, 81(1), 141.