

Foraging for Information in the Prefrontal Cortex

by

Geoffrey Keith Adams

Department of Neurobiology
Duke University

Date: _____

Approved:

Michael L. Platt, Supervisor

R. Alison Adcock, Chair

Jennifer Groh

Marc Sommer

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor of Philosophy
in the Department of Neurobiology in the Graduate School
of Duke University

2014

ABSTRACT

Foraging for Information in the Prefrontal Cortex

by

Geoffrey Keith Adams

Department of Neurobiology
Duke University

Date: _____

Approved:

Michael Platt, Supervisor

R. Alison Adcock, Chair

Jennifer Groh

Marc Sommer

An abstract of a dissertation submitted in partial
fulfillment of the requirements for the degree
of Doctor of Philosophy in the Department of
Neurobiology in the Graduate School of
Duke University

2014

Copyright by
Geoffrey Keith Adams
2014

Abstract

The ability to monitor, learn from, and respond to social information is essential for many highly social animals, including humans. Deficits to this capacity are associated with numerous psychopathologies, including autism spectrum disorders, social anxiety disorder, and schizophrenia. To understand the neural mechanisms supporting social information seeking behavior requires understanding this behavior in its natural context, and presenting animals with species-appropriate stimuli that will elicit the behavior in the laboratory. In this dissertation, I describe a novel behavioral paradigm I developed for investigating social information seeking behavior in rhesus macaques in a laboratory setting, with the use of naturalistic videos of freely-behaving conspecifics as stimuli. I recorded neural activity in the orbitofrontal and lateral prefrontal cortex of monkeys as they engaged in this task, and found evidence for a rich but sparse representation of natural behaviors in both areas, particularly in the orbitofrontal cortex. This sparse encoding of conspecifics' behaviors represents the raw material for social information foraging decisions.

Dedication

For Ronald Kermit Fahl, 1916 - 2011.

Contents

Abstract	iv
List of Tables	viii
List of Figures	ix
Acknowledgements	xi
1. Introduction	1
1.1 The utility of social information	4
1.1.1 Social information in nonhuman primates	5
1.2 Internal and external validity	5
1.2.1 The clinic and the community	6
1.2.2 The laboratory and the field.....	8
1.3 Video stimuli: bringing the field to the lab	10
1.4 The orbitofrontal and lateral prefrontal cortex	12
1.5 A social information foraging task.....	15
2. Neuroethology of decision-making.....	17
2.1 Introduction.....	17
2.2 Foraging decisions.....	20
2.3 Social information seeking	25
2.4 Discussion.....	32
3. Sparse coding of viewed natural behaviors in the prefrontal cortex.....	35
3.1 Materials & Methods.....	35

3.1.1 Behavioral paradigm.....	35
3.1.2 Video stimuli.....	37
3.1.3 Measurements of gaze during video presentation	45
3.1.4 Surgery	49
3.1.5 Electrophysiological recordings.....	50
3.1.6 Statistical models and regularization technique	51
3.1.7 Analysis of video selection decisions	55
3.1.8 Analysis of gaze during video presentation.....	56
3.1.9 Analysis of neuronal firing rates.....	58
3.2 Results	62
3.2.1 Monkeys' viewing preferences for naturalistic videos	62
3.2.2 Viewed behaviors influence monkeys' gaze behavior.....	64
3.2.3 Prefrontal neurons respond robustly to videos of conspecifics and sparsely represent viewed behaviors.....	72
4. Conclusion	84
4.1 Information foraging.....	84
4.2 Semantic representations of natural scenes	86
4.3 Representations of social information	88
4.4 Neuroeconomics and Neuroethology	90
References	95
Biography.....	106

List of Tables

Table 1: Ethogram used for coding videos.....	39
Table 2: Derived ethogram elements.....	43

List of Figures

Figure 1: Depending on the mechanisms present in the ancestral state, dissimilar behavioral problems can be solved by common mechanisms, and similar behavioral problems can be solved by disparate mechanisms	19
Figure 2: Rhesus macaques forage nearly optimally in a computerized patch-leaving task, and the rising value of leaving a patch is represented by single neurons in the macaque frontal cortex.	22
Figure 3: Social information is a valuable resource for macaque monkeys.....	27
Figure 4: The value of social information is signaled by neurons in the macaque visual orienting system.	31
Figure 5: A schematic diagram of a task allowing monkeys to view and make decisions about social information.	37
Figure 6: Schematic representation of ethogram values during video presentations from a selected section of one behavioral session.....	45
Figure 7: Illustration of coefficient selection under elastic net regularization.	54
Figure 8: Autocorrelation in neuronal firing rate during video presentation, LPFC example unit	61
Figure 9: Each subject exhibited clear preferences among viewing outcomes.....	62
Figure 10: Menu option utilities for each subject	63
Figure 11: Features of viewed video scenes, including conspecific behaviors, influenced monkeys' decisions to look at or away from the presented video at each frame.	65
Figure 12: Differences in the reliability of gaze across repeated viewings of the same video sequence associated with features of the video scene, including conspecific behaviors, were successfully captured by the gaze consistency metric and associated model.	69
Figure 13: Individual PFC units respond strongly to videos of conspecifics.	72
Figure 14: PFC population-averaged responses to task events.	73

Figure 15: PFC units' responses to task events were heterogeneous.....	74
Figure 16: PFC units' firing rates were modulated by viewed behaviors during video presentation.	75
Figure 17: The firing rate model reveals strong modulation of an LPFC neuron's firing rate in response to multiple elements of a video scene, including behaviors.	77
Figure 18: Sparse representation of ethogram entries and gaze behavior metrics in individual units in prefrontal cortex. Units are ordered within brain region by the mean of the absolute value of the log gain over all regressors.	79
Figure 19: The proportion of units identified as sensitive to each regressor.....	80
Figure 20: The proportion of regressors identified with nonzero coefficients for each unit.	81
Figure 21: Firing rates during video presentation of units in OFC tend to encode more regressors than those in LPFC.....	83

Acknowledgements

The work described in this dissertation would have been impossible without the contributions of the rich intellectual community I have been privileged to be a part of at Duke during my graduate studies. In particular, I would like to acknowledge the support of my advisor, Professor Michael Platt. The current and former members of his research group have also helped shape my scientific career, and I can only attempt to express my gratitude for their help and friendship: Karli Watson, John Pearson, Ben Hayden, Steve Chang, Lauren Brent, Jeff Klein, Arani Roy, Sarah Heilbronner, Amrita Nair, Becket Ebitz, Nick DeWind, David Barack, Jean-François Gariépy, Jessica Yorzinski, Amanda Utevsky, Caroline Drucker, Koji Toda, Seth Madlon-Kay, and Monica Carlson, who somehow keeps it all running. I would also like to thank the members of my committee, Professors Alison Adcock, Jennifer Groh, and Marc Sommer. This work also owes a tremendous debt to the outstanding veterinary and support staff in the Duke Division of Laboratory Animal Resources, whose contributions to the enterprise of biomedical research with expertise and compassionate care too often go unsung; in particular, I would like to thank Dr. Kyha Williams. And of course I owe so much to my family: my mother, Renée Adams, my father, Robert Adams, my brother, Robin Adams, and my wonderful partner, Jetta Peterkin, who devoted many hours to the tedium of counting monkeys. Finally, I thank Ernst, Oskar, Cajal, and Huxley, who have taught us all so much.

1. Introduction

All men by nature desire to know. An indication of this is the delight we take in our senses; for even apart from their usefulness they are loved for themselves; and above all others the sense of sight. For not only with a view to action, but even when we are not going to do anything, we prefer seeing (one might say) to everything else. The reason is that this, most of all the senses, makes us know and brings to light many differences between things.

Aristotle, *Metaphysics*, Book 1
Translated by W. D. Ross

Aristotle's belief that "all men by nature desire to know" is often quoted as a pithy description of human inquisitiveness, but for Aristotle it was part of a larger philosophical argument. If the desire for knowledge is an essential aspect of human nature, then it must serve some function (in Aristotle's terminology, a "final cause") toward human flourishing. That is, we should understand innate human curiosity at least in part through a teleological lens. Of course, Aristotle's phrasing and perspective also reveal his cultural assumptions and philosophical commitments to the primacy of the human male. But human curiosity knows no gender, and the motivation to seek information is found throughout the animal kingdom (Valone & Templeton 2002, Dall et al. 2005, Gottlieb 2012).

Since Darwin's theory of natural selection, modern biological theories of human and animal behavior also emphasize a teleological explanation of behavioral and cognitive traits, the concept of adaptive function. Indeed, the question of adaptive function from Tinbergen's "Four Questions" schema (Tinbergen 1963) is typically

identified with Aristotle's final cause (Hladký & Havlíček 2013). Though "hyperadaptationism" (Gould & Lewontin 1979) can lead to the adoption of implausible (and sometimes unfalsifiable) explanations for biological traits (Lloyd 2009, Sterelny 2007), the concept of adaptive function nevertheless frequently provides a useful guide in inquiring into the causes of animal behavior. The pitfalls of hyperadaptationism can be further avoided by pursuing along with adaptation explanations from alternative evolutionary mechanisms, including phylogenetic constraints, exaptation, developmental canalization, founder effects, and genetic drift. In neuroethology, these various evolutionary mechanisms can represent a rich source of hypotheses for the proximate mechanisms of behavior, and discovered proximate mechanisms can provide evidence in support of or against these various evolutionary explanations.

In this dissertation, I will undertake an investigation of the neural mechanisms involved in the exploration of the social environment, informed by a neuroethological perspective. While the acquisition and use of social information has clear adaptive significance (Dunbar 1998, Valone & Templeton 2002, Dall et al. 2005, Dubois et al. 2011, Archie et al. 2012), my approach will also be informed by an exaptive program (Andrews et al. 2002) proposing a mechanistic and possibly evolutionary link between foraging behavior and social exploration on the basis of two observations. First, there are formal similarities between the problems of foraging and information seeking (Stephens 1986, Pirolli 2007) suggesting that mechanisms which evolved to perform

computations for one of the two problems could have been subsequently exapted for the other. Second, the brain areas I will be investigating, the orbitofrontal cortex (OFC) and lateral prefrontal cortex (LPFC), have been proposed to play a role in foraging and decision-making for nutritive rewards (Carmichael & Price 1996, Ongür & Price 2000, Padoa-Schioppa & Assad 2006, Cai & Padoa-Schioppa 2014) as well as social valuation and decision-making (Sallet et al. 2011, Watson & Platt 2012, Noonan et al. 2014, Cai & Padoa-Schioppa 2014).

Chapter 2 was originally published as a review in *Current Opinion in Neurobiology* (2012), co-authored with Karli K. Watson, John Pearson, and Michael L. Platt. It further develops the neuroethological approach to the study of cognitive processes, with a particular emphasis on the relationship between foraging and social information seeking that motivates my original study.

Chapter 3 describes a novel experimental paradigm for investigating the neural mechanisms supporting social information foraging in rhesus macaques, and the behavioral and neurophysiological results of this study. In it, I will describe novel evidence for the sparse encoding of natural, biologically significant viewed behaviors in the prefrontal cortex of monkeys watching videos of conspecifics. This chapter represents the original research I performed during my doctorate studies.

Chapter 4 closes with a discussion of the implications of the results presented in Chapter 3. The role of prefrontal cortical areas in adaptive behavior and decision-

making is an exciting, rapidly-changing area of research, and the interpretation of my findings is both informed by and informative toward this body of literature.

1.1 The utility of social information

Humans actively seek out information about their social partners' emotions, beliefs, motivations, and relationships. This information is vital for producing appropriate behaviors during normal social interactions. Some sources of such information are immediately available during interaction with a social partner, such as the partner's facial expressions. Other sources of social information are available without direct interaction, through observing a social partner's behaviors and social interactions with others. In both cases, an individual must identify a potential source of social information, evaluate its likely significance, and experience motivation to track it. Abnormal social information seeking is a characteristic of a variety of neuropsychiatric disorders that affect social cognition, including autism spectrum disorders (ASD), schizophrenia, and social anxiety disorder, and the deficits to visual orienting behavior during viewing of naturalistic social scenes is predictive of the degree of overall social impairment (Klin et al. 2002). In children with ASD, joint attention training, which directly targets deficits in social information seeking, can improve overall social functioning (Whalen & Schreibman 2003). The fact that behavioral interventions during cognitive development which increase a patient's attention to social information can lead to overall improvement in social function demonstrates that abnormal social

information seeking behavior is not only an effect of disrupted social cognition, but can also contribute causally to disrupting social cognition by depriving an individual of social information that is vital for its normal development. Despite the importance of social information seeking to normal human behavior and the devastating consequences of impairments in this behavior, the manner in which social information is transformed into a motivational code to guide orienting remains largely unknown.

1.1.1 Social information in nonhuman primates

Like humans, monkeys can learn about their social partners not only through direct interaction, but also by observing them as a non-participating third party (Lorincz et al. 1999, Bovet & Washburn 2003). Furthermore, monkeys discriminate between sources of information on the basis of their likely utility (Shepherd et al. 2006). For example, wild baboons exhibit enhanced orienting responses toward sources of information that reflect a change in relationships in their troop relative to sources of information which are consistent with well-established relationships (Cheney et al. 1995). This and other studies provide good evidence that the social information seeking behavior of monkeys is fundamentally similar to that of humans.

1.2 *Internal and external validity*

Living systems are complex and contingent. Removed from their natural context, their properties frequently change in response to the new environment. It is this fact about biological systems that leads to a major issue in the methodology of the

biological sciences, the problem of balancing *internal* and *external validity*, which is also of great concern to clinicians, psychologists, and social scientists, for similar reasons. Internal validity refers to question of whether a study's design is such that conclusions about the system under study (e.g., the physiology of laboratory animals' brains) are sound, while external validity concerns the question of whether the system under study is representative of the population of interest (e.g., animals in the field) such that generalizing from the studied system is sound. These two goals are frequently seen as being in tension with each other, as for many questions, high internal validity is provided by techniques which are only available in a laboratory setting or entail manipulating the system under investigation in ways that may render it dissimilar to its natural counterparts, reducing external validity.

1.2.1 The clinic and the community

A particularly illuminating example of how investigators are confronted with considerations of internal and external validity may be found in the history of neuropsychological studies of human patients with lesions of the prefrontal cortex. In 1985, Eslinger and Damasio reported the case of a patient, EVR, who had suffered an orbitofrontal meningioma, leading to broad bilateral lesions of the prefrontal cortex. In their report, they wrote, "[t]he central finding was the dissociation between intact cognitive abilities measured by standardized tests, and the poor utilization of those abilities in the real environment." (Eslinger & Damasio 1985) Prior to his tumor, EVR

was described by friends and family as extremely responsible, with a successful marriage and career. Following the lesion, his behavior changed; he became unreliable, pursued risky business ventures which failed, lost his job, and eventually divorced his wife of 17 years; the full case report is a fascinating, if heartbreaking, story. However, despite these obvious changes in behavior which so impaired EVR's quality of life, the neuropsychological examination performed by Eslinger and Damasio revealed no obvious impairments. Indeed, EVR's intelligence was above average, and the clinicians appear to have been impressed with his conversational ability. This disparity between findings of mild or no impairments in clinical assessments versus significant impairments to "real life" behavior is relatively common in case reports from the time of patients with prefrontal lesions, particularly when the damage is restricted to orbital and medial areas (Bigler 1988). Indeed, although a number of clinical instruments to assess OMPFC function have been developed over the subsequent decades, accounting for the level of deficit in patients' behavior in the community outside of the clinic remains a challenge (Zald & Andreotti 2010). Perhaps the relatively simple environment of the clinic places weaker demands on the functions that these brain areas serve in natural behavior. In practice, this literature also illustrates how clinicians leverage both the internal validity offered by reproducible neuropsychological instruments in the controlled environment of the clinic and the external validity offered by patient histories

which include evidence from the more complex, ever-changing environment of the community.

1.2.2 The laboratory and the field

The use of animal models for understanding human physiology and disease also represents a compromise of external validity in favor of internal validity in order to apply techniques that would be impractical or unethical to apply to human subjects. For those aspects of human physiology which have been highly conserved over evolutionary time, the likely reduction of external validity imposed by studying the homologous physiological process in an animal model may be minimal. (If, that is, the trait has also been highly conserved in the model species since the common ancestor, a possibility which should not be overlooked.) However, if a researcher is interested in investigating aspects of human physiology which are highly specialized or unique, then the implications of employing a model species must be considered carefully. This is a problem which should be of particular interest to cognitive neuroscientists.

The rhesus macaque is a commonly studied model organism for cognitive and systems neuroscience. It represents an excellent choice for investigating the neural mechanisms underlying human cognitive processes. Macaques share with humans a similar neuroanatomy, relative phylogenetic proximity with a common ancestor approximately 30 Mya (Wilkinson et al. 2011), and a broad array of natural habitats that place demands on cognitive flexibility. Furthermore, a long history of comparative

neuroanatomical research (e.g., Ongür & Price 2000) serves to help clarify points of variation between the two species. These factors all contribute to investigators' ability to extrapolate from the results of studies in macaques to make inferences about neural mechanisms in humans.

However, the discrepancy seen in prefrontal lesion patients between clinical observations and "real life" behavioral deficits should serve as a cautionary note here as well. First, as a general principle, complex cognitive and behavioral phenomena may be difficult to characterize in the restrictive environment of the laboratory. This may be particularly true for social behaviors, which require a level of free interaction between individuals that is difficult to accommodate in the laboratory setting. Second, more specifically, the functions of the orbital and medial prefrontal cortex may be particularly difficult to characterize in the laboratory setting. If the functional deficits caused by damage to these areas in human patients is obscured in the relatively limited environment of the clinic, then the normal functions of these areas may also be difficult to elicit in animals in the limited environment of the laboratory. If this is indeed the case, then progress in understanding the function of prefrontal cortex in animal models may greatly benefit from considering in what ways the external validity of the experimental context can be improved.

If, then, the rhesus macaque is to be a valid model for the neurobiological study of complex cognitive phenomena, we can benefit from employing tasks and stimuli

which are inspired by the conditions that free-ranging monkeys encounter in their natural habitats. Such an approach has been successfully employed in capitalizing on monkeys' innate fear responses to snakes, a naturalistic stimulus (e.g., Izquierdo et al. 2005) and in simulated foraging tasks which are structured to resemble natural foraging problems (Hayden et al. 2011). To understand the neural mechanisms underlying social information seeking behavior, a similar approach is warranted.

1.3 Video stimuli: bringing the field to the lab

There is a successful history in the use of video stimuli in studies of animal behavior (D'Eath 1998). While there are a number of reasonable concerns about the external validity of video stimuli due to the fact that video recording and presentation devices are generally optimized for the human visual system, and non-human animals can differ significantly from humans in visual acuity, color perception, and temporal sensitivity, the success of video stimuli in driving animals to naturalistic responses in taxa as diverse as primates (Shepherd et al. 2010, Mosher et al. 2011), birds (Bird & Emery 2008, Rieucau & Giraldeau 2009), and arachnids (Tedore & Johnsen 2013) lends strong support to the notion that these stimuli are sufficiently naturalistic to have high external validity. In rhesus macaques, the close similarity of the visual system to that of humans also reduces concerns that video stimuli may fail to adequately simulate natural visual stimulation, although subtle differences in color perception may be important to consider (Higham et al. 2010). Furthermore, video stimuli in rhesus macaques

successfully elicit naturalistic visual orienting responses (Shepherd et al. 2010, Mosher et al. 2011, McFarland et al. 2013), suggesting that video stimuli presenting conspecifics engaged in species-typical behavior represents an excellent choice for studying social information seeking behavior.

In humans, video stimuli have also proved to be a powerful tool for the investigation of brain function (Hasson et al. 2008). Human subjects watching the same videos tend to have similar patterns of fMRI BOLD activity (Hasson et al. 2004), and by decoding BOLD activation patterns in visual cortex, the original video stimuli that elicited the neural responses can be recovered with some fidelity (Nishimoto et al. 2011). Furthermore, the semantic content of video stimuli appears to be represented broadly across the entire cortex (Huth et al. 2012). These findings endorse the use of video stimuli as a means of engaging the brain in a naturalistic fashion to explore numerous dimensions of neural function simultaneously.

In rhesus macaques, the use of videos of conspecifics as stimuli during electrophysiological recordings in the amygdala has found that individual cells in the amygdala respond to eye contact, that is, when the video-taped monkey looks directly at the camera, and the subject monkey looks at the eyes in the movie, simulating natural eye contact (Mosher et al. 2014). Importantly, these responses were much weaker when static images were presented. Presumably, the added naturalism of the videos, in which the observed monkeys exhibited biologically meaningful behaviors, encouraged the

more complete engagement of the subject animals' nervous systems in the task of processing the presented stimuli.

1.4 The orbitofrontal and lateral prefrontal cortex

The orbitofrontal cortex (OFC), covering the ventral surface of the prefrontal lobe, has characteristically been associated with motivation, valuation, and hedonic pleasure (Rolls 2000). However, its precise functions in producing behavior remain elusive.

In animal studies, one of the most characteristic behavioral deficits associated with lesions of the OFC is to performance in reversal learning tasks (Murray et al. 2007). For these tasks, animals are trained on a cue-outcome association, for example, that a tone precedes a food pellet. Both intact and OFC-lesioned animals can readily learn this initial association. Once the association has been learned, the associated outcome is changed so that it is no longer desirable, for example from a food pellet to a foot-shock. Intact animals rapidly alter their learned behavior, but animals with OFC lesions are slower to do so (Murray et al. 2007). A recent promising hypothesis for the role of OFC in learning suggests that the brain includes distinct learning systems for so-called model-free learning, in which learning occurs only by updating a value associated with an object or action, as well as model-based learning, in which animals learn the structure of relationships between cues and outcomes in the environment (Schoenbaum et al 2009, Walton et al. 2010). According to this model for OFC function, it serves as the model-

based critic, tracking the relative performance of the learned model's predictions against the observed outcomes. In the reversal learning paradigm, initial learning occurs with no prior model, and so only model-free learning is used. However, as the animal forms the cue-outcome association, intact animals also train their model-based learning system, permitting more rapid subsequent reversal. OFC-lesioned animals, lacking a properly-functioning model-based learning system, rely on the slower model-free system, thus perseverating on the initially-learned association.

Another approach to investigating the function of OFC has been to understand the information coding scheme of individual neurons. OFC units have been found to encode value (Tremblay & Schultz 1999) in a manner independent of the specific nature of the valuable stimulus (Padoa-Schioppa & Assad 2006), leading to the suggestion that OFC firing rates represent a "common currency" value code for decision-making (Padoa-Schioppa & Assad 2006, Landreth & Bickle 2008). According to this model, when animals have to make decisions between "goods," that is, objects in the world (or laboratory simulacra thereof) toward which they might direct consummatory actions, OFC allows the animal to make the decision by reducing the large space of features describing the good to a single abstract value, and the value of goods on the menu can be compared and the largest selected. Recent work has bolstered this model, and suggested that the process of converting the good-based decision signal into an action occurs in the lateral prefrontal cortex (LPFC) (Cai & Padoa-Schioppa 2014).

However, neural recordings in the OFC while monkeys made decisions about static social information found results not fully accounted for by this common-currency decision model (Watson & Platt 2012). In this study, Watson and Platt trained monkeys to make decisions between a single drop of juice versus a drop of juice paired with a social image. By varying the relative volumes of juice in each condition, an estimate of the common currency valuation of different categories of social images could be produced, in units of juice reward. While the study revealed representations of both juice and social image value, there was a surprising discrepancy between monkeys' decisions, which were driven primarily by the juice rewards, and the level of representation in OFC, with units encoding social value outnumber those encoding juice value by over a factor of two. From the perspective of revealed preference, monkey OFC appears to dramatically over-represent social information. However, from the perspective of ecological relevance, it is perhaps less surprising that social information, which can have tremendous utility over the life of a monkey, is well-represented in an area which may also play a role in learning a model of the environment (including the social environment), even if that information has little relevance to an immediate decision.

The model of Cai and Padoa-Schioppa (2014) that OFC represents the valuation of goods while LPFC either performs or represents the outcome of the good-to-action transformation should also be considered within the context of strong evidence for a key

role in social function for LPFC as well. Sallet and colleagues (Sallet et al. 2011) raised rhesus macaques in social groups of various sizes, and then performed volumetric analysis of structural MRIs of their brains, identifying a number of brain areas which were systematically associated with social group size and the individual's dominance rank. Intriguingly, among the identified areas was a region within the LPFC, on the ventral bank of the principal sulcus near its fundus (which they termed rostral prefrontal cortex, rPFC), which appears to correspond very closely to the region of LPFC investigated by Cai and Padoa-Schioppa. This region included both an area associated with social group size, closest to the sulcal fundus, and an adjacent, more lateral area associated with dominance rank. If this region of LPFC carries downstream value information from OFC in order to carry out a good-to-action transformation, then it seems plausible to conjecture that LPFC's increased volume with social group size is due to increased demands for converting rich representations of social information in OFC into adaptive information-seeking behaviors.

1.5 A social information foraging task

To better investigate the roles of OFC and LPFC in social information seeking behavior, I devised a task in which monkeys made viewing decisions about naturalistic videos of freely-behaving conspecifics, described in Chapter 3. To maximize the external validity offered by the use of such stimuli, the stimulus set was large and highly varied, such that monkeys saw the same sequence of video relatively infrequently,

typically at most once or twice per behavioral session. In order to identify reliable patterns in the data, I took as my fundamental unit of analysis during video viewing the nature of the behaviors on display, rather than the repeated video clip. That is, while a monkey saw an individual video sequence typically only once or twice in a session, there would be numerous examples of individual behaviors, such as allogrooming, on display, and by looking for regular responses to these viewed behaviors, patterns in the subject monkeys' behavioral and neurophysiological responses could be identified.

Decisions occurred on two scales. First, the "micro-decisions" during video presentation: whether to look at the screen, and where to direct gaze. Observations of monkeys' gaze behavior during video presentation provided a powerful window into understanding what features and behaviors drove information-seeking at this scale. Second, "macro-decisions" between video presentations: what kind of information would be most valuable to see? By allowing monkeys to make decisions between continuing to watch the video they had just seen versus switching to a new video, I expected to be able to make inferences about what types of social information were most valuable to the monkeys.

2. Neuroethology of decision-making

A neuroethological approach to decision-making considers the effect of evolutionary pressures on neural circuits mediating choice. In this view, decision systems are expected to enhance fitness with respect to the local environment, and particularly efficient solutions to specific problems should be conserved, expanded, and repurposed to solve other problems. Here, we discuss basic prerequisites for a variety of decision systems from this viewpoint. We focus on two of the best-studied and most widely represented decision problems. First, we examine patch leaving, a prototype of environmentally based switching between action patterns. Second, we consider social information seeking, a process resembling foraging with search costs. We argue that while the specific neural solutions to these problems sometimes differ across species, both the problems themselves and the algorithms instantiated by biological hardware are repeated widely throughout nature. The behavioral and mathematical study of ubiquitous decision processes like patch leaving and social information seeking thus provides a powerful new approach to uncovering the fundamental design structure of nervous systems.

2.1 Introduction

In the natural world, animals face many difficult decision-making problems, posed by the details of their habitat and social system. Such problems can include when, where, and for what to forage; with whom to mate and where to nest; whether to flee or

to ignore a potential predator; and with which conspecifics to form long-term associations. Evolutionary theory and behavioral ecology identify such decision-making problems, and the costs, benefits, and constraints associated with pursuing specific behavioral strategies for solving them.

Although the study of animal behavior has revealed a remarkable diversity of such problems and their solutions, mathematical analysis has demonstrated that dissimilar-seeming problems can be solved with similar strategies. For example, certain foraging problems can be generalized to a broad class of resource-seeking behaviors (Stephens & Krebs 1986). When the fitness impact of a decision is large, strong selective pressures should produce mathematically optimal behavior. Such a constraint means that the neurobiological mechanisms mediating decisions may be highly conserved or convergent across taxa for problems with similar mathematical formulations (Figure 1).

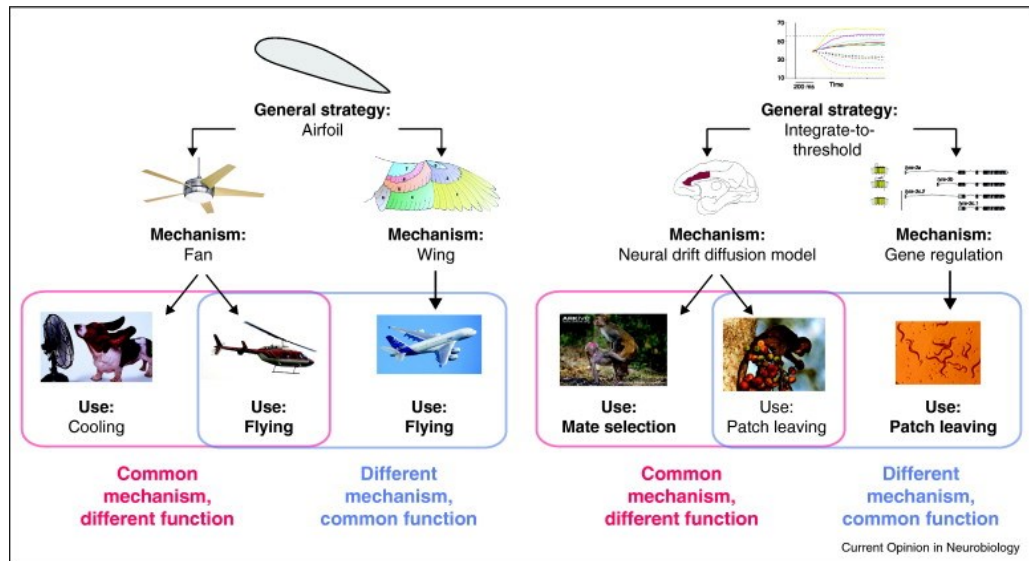


Figure 1: Depending on the mechanisms present in the ancestral state, dissimilar behavioral problems can be solved by common mechanisms, and similar behavioral problems can be solved by disparate mechanisms

(a) Analogy for the relationships between algorithm, mechanism, and function. A fundamental design, the airfoil, can be configured into two very different forms: the wing and the fan. These more elaborate configurations may be further modified for use in multiple, unrelated functions, as in the use of helicopter blades for flying or fan blades for cooling.

(b) We posit that the integrate-to-threshold algorithm, like the airfoil, is an evolutionary ‘building block’ that will be found frequently in various species and implementations. This algorithm can be implemented by several mechanisms, such as neural computational models as well as gene-regulatory mechanisms. As in (a), these specific mechanisms are integrated into the evolutionary toolkit, and made available for modification. Duplication and subsequent refinement of a particular mechanism can then adapt it to subserve a different behavior. For example, a neural mechanism that originally evolved to guide foraging behavior needs relatively small modifications before it can be implemented to enhance mate-seeking behavior.

To illustrate these ideas, we will discuss two types of decision problem in behavioral ecology and their neural implementations. First, we will examine foraging, a prototype of environmentally based switching between action patterns to optimize resource gathering. Second, we will examine social information seeking behaviors in

non-human primates. We conclude that this work portends a more general understanding of decision-making and, ultimately, motivates the unification of theoretical and experimental work in behavioral ecology and neuroscience.

2.2 Foraging decisions

A well-studied example of a natural foraging decision is the patch-leaving problem, mathematically analyzed by Charnov (Charnov 1976) and first tested in a series of experiments in birds (Krebs et al. 1974). This problem considers an animal foraging in an environment with food items distributed in sparse patches. As the animal forages in a patch, local resources are depleted, and the time required to find a new food item increases, reducing the rate of food intake. As a result, animals must balance the benefits of diminishing returns against the costs of searching for new patches. The optimal solution to this problem, the Marginal Value Theorem (MVT)—that foragers should abandon patches when the local rate of caloric return falls below the average for the environment as a whole—has been demonstrated to hold in a breathtakingly wide array of species, including worms, insects, fish, rodents, birds, nonhuman primates, and humans (Stephens & Krebs 1986, Stephens et al. 2007).

In a recent experiment investigating the neural basis of patch leaving decisions, Hayden, Pearson, and Platt (Hayden et al. 2011) designed a laboratory version of the problem in which monkeys chose between ‘stay’ and ‘leave’ options represented by visual stimuli on a computer monitor. The ‘stay’ option simulated foraging within a

gradually depleting patch by initially delivering a large juice reward, followed by a decreasing payoff on each sequential selection of 'stay.' The 'leave' option simulated travelling to a new, undepleted patch by incurring a 'travel time' before giving a juice reward and resetting the payoff associated with the 'stay' option. The travel time was explicitly cued to the monkeys by the height of a bar presented on the computer monitor, and was selected randomly for each new 'patch.' Monkeys readily optimized their patch residence times (Figure 2a) and their patch-leaving decisions also depended systematically on the 'travel time' to the next patch, a departure from the MVT consistent with observed behaviors in other species (Cuthill et al. 1990). Simultaneously, Hayden and colleagues recorded the firing patterns of single neurons in the dorsal anterior cingulate cortex (dACC), an area of the macaque and human brain linked to reward monitoring, error signaling, learning, and behavioral control. Neuronal firing rates during the peri-saccadic decision-making period rose with each successive decision to stay in a given patch, across multiple actions unfolding over tens of seconds (Figure 2b, c). An analysis of neuronal firing according to the monkeys' patch residence times demonstrated that these responses were more closely related to the decision to leave a patch than to the monkeys' patch residence times or reward rates (Hayden et al. 2011). This firing pattern revealed a strikingly simple implementation of a thresholded decision circuit. Using fMRI in humans, Kolling and colleagues (Kolling et al. 2012) investigated the function of ACC using a similar foraging task, and found that ACC

tracks the value of foraging (i.e. leaving the current patch) on a scale anchored to the average value of the environment. While dACC appears to encode a signal that is compared to a threshold to determine the patch leaving decision, the mechanism by which this threshold is set remains unknown, but according to the MVT should be flexibly sensitive to environmental conditions such as average patch richness.

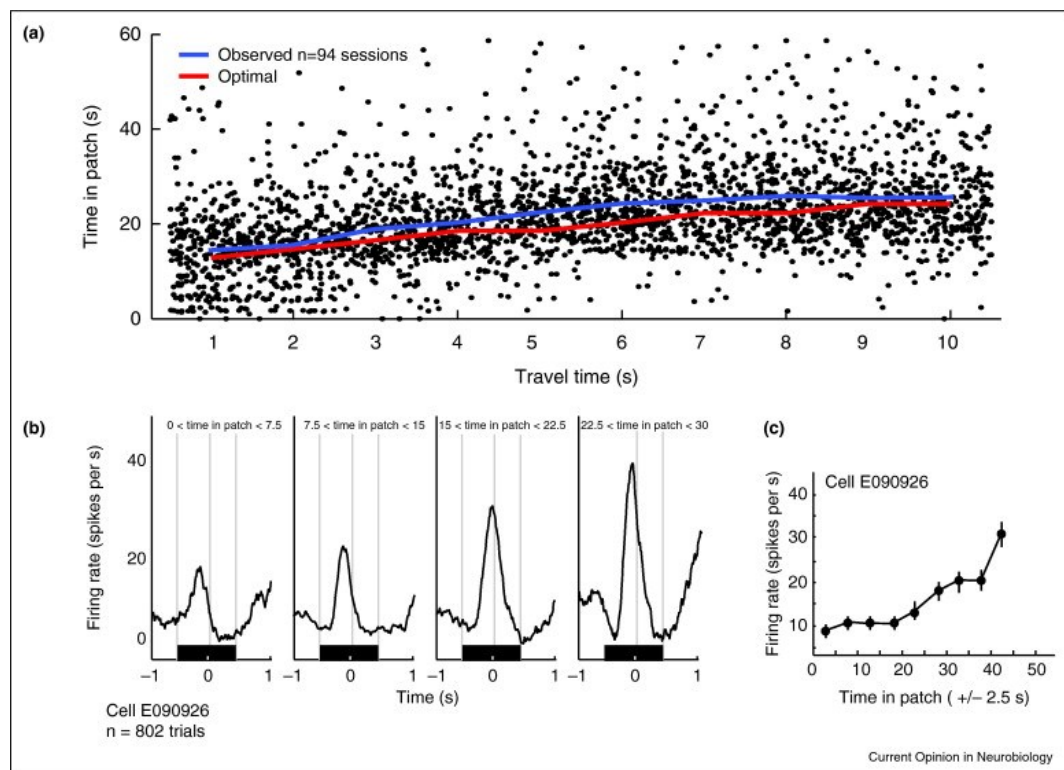


Figure 2: Rhesus macaques forage nearly optimally in a computerized patch-leaving task, and the rising value of leaving a patch is represented by single neurons in the macaque frontal cortex.

(a) Monkeys remain in the patch longer as travel time rises, as predicted by the marginal value theorem (MVT). Each dot indicates a single patch-leaving decision ($n = 2834$ patch-leaving events). The time at which the monkey chose to leave the patch (y axis) was defined relative to the beginning of foraging in that patch. Travel time was kept constant in a patch (x axis). Data from two monkeys are shown. Behavior (average is traced by the blue line) closely followed the rate-maximizing leaving time (red line), albeit delayed by 0–2 s. (b) Peri-stimulus time histograms (PSTHs) for an

example cell in anterior cingulate cortex (ACC). Neurons responded phasically around the time of decision-making saccades and then fell to a baseline level between trials. Time zero indicates end of saccade, indicating choice. Black rectangle indicates the average duration of the trial. The firing rate during the peri-saccadic decision-making period rose with each successive decision to stay in a given patch, across multiple actions unfolding over tens of seconds. Each panel indicates responses selected from one range of patch residence times. (c) Average responses of example neuron occurring in a series of 1-s peri-saccadic epochs. Firing rates increased as time in patch increased. Error bars represent s.e.m. Firing rates peaked with the decision to abandon a patch and move on to the next. Figure after Hayden et al. (2011), used with permission.

The rise-to-threshold process evident in the decision-related responses in ACC mirrors similar rise-to-threshold processes evident in the activity of neurons in parietal and prefrontal association cortex in monkeys rendering perceptual judgments (Gold & Shadlen 2007). While such perceptual judgments entail the integration of sensory evidence from a noisy environment over time until a threshold is reached, similar integration is not immediately apparent in ACC firing during this patch foraging task. Nevertheless, both algorithms entail threshold comparison, and variability in neural firing determines the monkeys' decisions. Further, such rise-to-threshold processes have been theorized to serve as implementations of the sequential probability ratio test (SPRT)—the optimal solution to this type of binary decision problem (Wald & Wolfowitz 1948, Ratcliff 1978, and Ratcliff & McKoon 2008). The observation that a similar process appears to govern both perceptual decisions and patch leaving decisions endorses the idea that the brain uses a small suite of similar, repeated mechanisms to solve diverse problems in multiple domains (Hayden et al. 2011).

Recent evidence also indicates that genes coding for neuromodulatory chemicals may mediate the impact of local environmental conditions on patch leaving decisions, perhaps by controlling the threshold for patch leaving. Bendesky and colleagues (Bendesky et al. 2011) compared patch-leaving decisions by a nematode (*C. elegans*) strain from Hawaii with one developed in the laboratory. Hawaiian worms abandoned algae lawns at ten times the rate that the laboratory strains did. The authors found that differences in patch-leaving threshold among distinct strains resulted from polymorphisms in promoter regions of the tyramine receptor gene (*tyra 3b*), which controls the expression of a G-protein coupled receptor analogue of vertebrate catecholamine receptors. Thus worms' patch leaving thresholds are regulated genetically, as opposed to monkeys' use of flexible thresholding, though both obey the MVT. In other words, different proximal mechanisms are capable of instantiating the same algorithm to solve a computationally equivalent biological problem.

Intriguingly, the invertebrate catecholamines tyramine and octopamine, which bind to the *tyra3* receptor, are closely related structurally to the vertebrate neuromodulator norepinephrine, which has been hypothesized to regulate exploration/exploitation tradeoffs in primates and humans (Aston-Jones & Cohen 2005, Cohen et al. 2007). It is tempting to speculate that individual differences in exploration behavior, an innate tendency to abandon the current behavioral strategy for another

potentially more profitable one, might also be mediated in humans through genetic influences on catecholaminergic neuromodulatory systems (Frank et al. 2009).

In a remarkable illustration of the idea that similar decision-making problems may be solved by widely disparate mechanisms following similar strategies, the nest site selection behavior of honeybees has been found to be governed by much the same integrate-to-threshold process as the foraging behavior of monkeys and worms (Seeley et al. 2012). In this case, it is the individual animal that is the basic unit of information processing rather than the single neuron. Individual bees dance to accumulate evidence in favor of a nest site, while other bees that suffered predator attacks at that site ‘head-butt’ the dancers, a stop signal that inhibits dancing. When the evidence in favor of a nest site reaches a threshold, the entire swarm mobilizes to the selected site. Notably, this research was motivated by predictions derived from findings in macaques, demonstrating the predictive potential of the comparative neuroethological approach (Marshall et al. 2009). Furthermore, the same dancing and head-butting mechanism is also used in the selection of foraging sites, lending support to the idea that a mechanism evolved to solve one decision-making problem may also be applied to solving a different but similar one.

2.3 Social information seeking

Social information seeking, like foraging, is a process of accessing a variably distributed resource, under the constraint of costs associated with the search behavior.

Information about social partners can be valuable in a wide variety of contexts, such as mate choice, foraging, and nest site choice, and the strategies used to solve one class of problem, such as foraging, may also be applied to another class of problems, such as social decision-making.

Thus the primacy of information in guiding decisions is abundantly evident in the social behavior of primates. Primates have frontally oriented, mobile eyes with a central fovea composed of a high density of cone photoreceptors, and thus are not capable of sampling all regions of the visual field simultaneously. This set of adaptations has led to the evolution of mechanisms that orient the visual system to objects with high information value via overt and covert attention (Moore et al. 2003). In the context of foraging, information is given value as a consequence of the value of the nutriment it may yield, as when an animal evaluates a tree for the presence of ripe fruit. But animals living in complex and dynamic societies can use the same attentional strategies to gather information about others (Klein et al. 2008), including rank (Bovet & Washburn 2003), identity (Parr et al. 2000), group membership (Mahajan et al. 2011), direction of gaze (Lorincz et al. 1999, Deaner & Platt 2003, and Ferrari et al. 2000), and emotional state (Sackett 1966).

While behavioral and physiological evidence demonstrates the value of information to primates in decision-making, information seeking also has associated costs. By contrast with foraging, for which locomotion imposes high energetic costs, the

metabolic costs of information seeking are quite low. However, information seeking can impose time costs, requiring animals to forego activities such as sleeping, drinking, or eating, which demand postures or behavioral states incompatible with attentive orienting (Figure 3a). Other costs are social: in the case of rhesus macaques, visual fixation on the face of another individual invites aggression (van Hooff 1967). Furthermore, inappropriate information seeking (e.g. directing attention to low-value information) can result in missed opportunities to gather more useful information elsewhere. These tradeoffs are part of the decision-making problem that social animals must solve in allocating their time and attention to social information seeking behaviors.

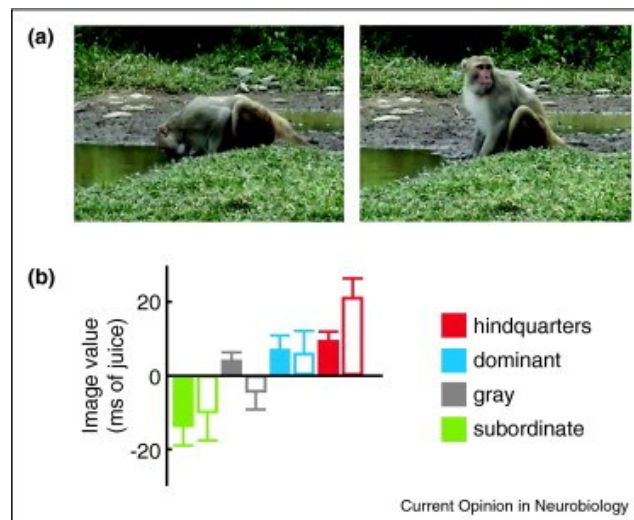


Figure 3: Social information is a valuable resource for macaque monkeys.

(a) A rhesus macaque on Cayo Santiago assumes a vulnerable posture (left) to drink from a puddle, but periodically interrupts this posture in order to visually scan the surrounding region for potential threats (right). There are no predators on the island, but aggressive social interactions are commonplace. **(b)** Values determined for different image classes for two male monkey subjects (open and closed bars), in ms of fluid delivery time. Positive deflections indicate the subject was willing to forgo fluid to view that image class. Negative deflections indicate the subject required fluid

overpayment to choose that image class. Hindquarters refer to the perineal sexual signals of familiar females. Dominant and subordinate refer to the faces of familiar dominant and subordinate males. Gray refers to a plain gray square matched for size and luminance to the other image classes. Behavioral data depicted here correspond to neural data depicted in Figure 4, below. Photographs by K.K. Watson.

Thus, as in foraging, adaptive choices during information seeking depend on the assessment and comparison of associated costs and benefits. Deaner, Khera and Platt (2005) demonstrated that thirsty male rhesus macaques will forego a small amount of juice in order to acquire specific types of social information, such as reproductive signals (i.e. female perinea) or the faces of dominant males, but will not do so for other types of social information, such as the faces of subordinates (Figure 3b). By contrast, monkeys look longest at reproductive signals but quickly avert gaze from both dominant and subordinate faces. These two observations invite the hypothesis that although both sexual signals and status-related signals contain high information value, there is a high cost associated with an extended period of direct eye contact. In support of this interpretation, a genetic polymorphism in the serotonin system associated with heightened anxiety elicits reduced attention to the faces of other monkeys, greater pupil dilation (a somatic index of elevated autonomic arousal) in response to faces of dominant males, and reduced reward value for viewing the faces of dominant males (Watson et al. 2009) in rhesus macaques.

Multiple cortical and subcortical brain areas are specialized for processing social information. Regions in the temporal and occipital lobes are specialized for processing

social information in humans and macaques, suggesting adaptation for the rapid assessment of visual social information. Humans and macaques both possess multiple brain regions, identified by functional imaging and confirmed by recordings from single neurons, exquisitely selective for faces (Kanwisher et al. 1997, Tsao et al. 2003). Neurons in other temporal lobe areas encode head direction (Perrett et al. 1985), face identity (Perrett et al. 1984), or biological motion (Oram & Perrett 1994). Cortical thickness in both temporal and prefrontal areas increases with increasing social group size in macaques (Sallet et al. 2011).

Subcortical structures play a particularly critical role in encoding affective properties of social stimuli. For example, the amygdala is known to encode stimulus salience, and damage to this region disrupts face processing and produces abnormal social approach behavior (reviewed in Adolphs 2010). In addition, fMRI studies reveal that many parts of the canonical subcortical reward circuit, most notably the nucleus accumbens (NAcc), are activated by visual social rewards (Smith et al. 2010), social approval (Rademacher et al. 2010), aversion to social inequity (Tricomi et al. 2010), and charitable donations (Carter et al. 2009). The special role of NAcc in social interaction is endorsed by animal research showing, for example, that D2 receptor manipulation in NAcc alters partner preference in mating voles (Gingrich et al. 2000), and that μ -opioid receptor manipulation in NAcc alters social play behavior in rats (Trezza et al. 2011).

These populations of neurons are interconnected and feed forward from the temporal into the frontal lobe, where various features of external stimuli, including their motivational value, are integrated (Ku et al. 2011). The architecture of this network suggests that decision-making mechanisms may have privileged access to social information in human and nonhuman primates alike. We conjecture that a common suite of algorithms, and perhaps mechanisms, subserves both 'social' and 'non-social' decision processes, with the former differentiated from the latter by the large number of specialized circuits used for detecting and processing information related to conspecifics.

Moreover, studies of neural activity in the primate brain support the notion that information itself has value for making decisions. Dopamine neurons, which respond to unpredicted primary reinforcers (such as nutritive rewards) and cues predicting reinforcers, also encode monkeys' preferences for advance information about the content of impending choices (Bromberg-Martin & Hikosaka 2009). Moreover, firing of neurons in lateral intraparietal area (LIP), a region of visuomotor cortex thought to encode a salience map of the visual world (Goldberg et al. 2006, Bisley & Goldberg 2010), varies with the value of social images displayed in the neurons' receptive fields. Neurons in LIP not only encode the value of the juice reward monkeys will gain for orienting to a particular location (Platt & Glimcher 1999), but also the intrinsic biological value of the social information they receive for orienting to the same location (Klein et al. 2008).

When monkeys are given the opportunity to either maximize a juice reward or display a social image, LIP neurons respond to both the amount of juice received and to the value of the social reward (Figure 4). Importantly, social and gustatory value are encoded independently along the same axis, suggesting that LIP plays a role in assigning value to a particular location in space, irrespective of the nature of the resource being sought. Thus, neurons in parietal cortex—a final common pathway for information seeking behavior—automatically integrate the intrinsic biological value of social agents into salience maps guiding orienting.

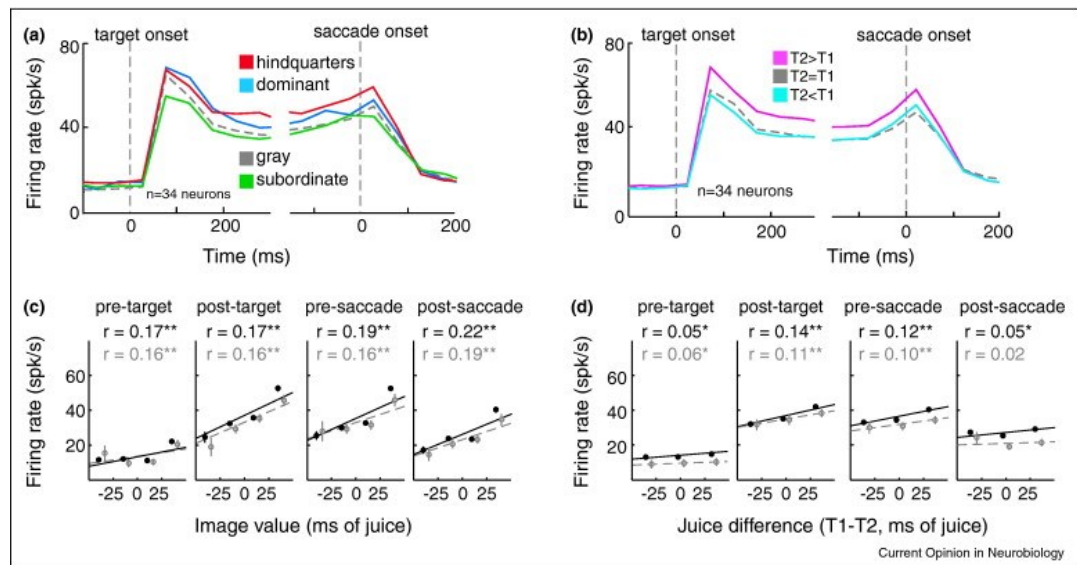


Figure 4: The value of social information is signaled by neurons in the macaque visual orienting system.

(a) Average firing rate for 34 LIP neurons plotted against time for all trials in which the subject chose to view the image (T2) in the neuron's response field, separated by image class. (b) Average firing rate of the population for all trials in which the subject chose to view the image (T2) in the neuron's response field, separated by fluid value relative to the non-chosen target (T1). (c) Firing rates during the 200 ms after target onset, plotted as a function of image value (left) and as a function of difference in fluid payoff between T2 and T1 (right). Regressions were performed on all data in

which the subject chose to view the image. The data in (c) were binned for display, but all regressions were performed on raw data. ****p < 0.001. Figure after Klein et al. (2008), used with permission.**

2.4 Discussion

The study of ecological decision-making suggests that natural selection has favored a set of simple, repeated design patterns: basic circuit elements capable of being implemented by many biological configurations. Instead of forming a single unified system for decisions, these local circuits might be capable of functioning independently for specialized subclasses of action planning and selection, as well as being recruited across regions for more complicated behaviors. Just as Gigerenzer has suggested that human decision processes draw on a well-stocked 'adaptive toolbox' filled with inexpensive, approximate heuristics, so evolution appears to favor repeating algorithms, despite disparate implementations (Gigerenzer & Selten 2001).

In this view, the study of a decision-making problem like patch leaving is valuable not only because it is ubiquitous, but also because the algorithm used to solve it—comparison of local returns to a fixed threshold—represents one of the simplest forms of a single-input/single-output control system, in which the controller implements a binary threshold operation (Brogan 1985). Any neural system capable of implementing such a circuit is likewise applicable to an extremely wide class of problems, one that extends far beyond foraging. In fact, a more general version of such a system, the Kalman filter, is known to be an optimal solution to the problem of predicting returns

under fairly general assumptions, and thus for fine-tuning behavior in response to changing environmental conditions (Brogan 1985, Anderson & Moore 1979). That such a system is linear, that it requires only a simple architecture, and that it is robust against noise all lead us to expect that it will not only appear in a wide diversity of species, but that it may be repeated and repurposed within a single brain to solve seemingly unrelated problems. In fact, such observations motivate a neural engineering viewpoint in which the unique classes of problems faced by an organism become primary, followed by the algorithms used to solve them, and only lastly their specific neural implementations (Marr 1982).

This does not mean, however, that comparative biology or neurophysiology become irrelevant for understanding decision-making. On the contrary, this evolutionary viewpoint suggests that algorithms implemented in simpler nervous systems are more generally applicable than we might have thought. Just as in vision or olfaction, the insights gained from studying flies or worms suggest possibilities at the algorithmic level in birds or primates. And though the details may differ as organisms become capable of more generalized and flexible behaviors, the same simple biological components, coupled like circuit elements, may likewise be expected to give rise to startlingly sophisticated generalizations (Brogan 1985).

Many models of decision-making, particularly those derived from economics, describe the decision process as a linear sequence: estimating the abstract utility of

several possible outcomes or behavioral plans, directly comparing these utilities, and finally selecting the goal or behavior associated with the highest utility (Glimcher 2004, Padoa-Schioppa 2007, Lee 2006, Glimcher et al. 2005, Sugrue et al. 2005). In some cases, such models leave the exact nature of the abstract utility undefined. Here, however, we have presented a 'bottom-up' as opposed to 'top-down' perspective, in which simple, reusable decision rules are implemented repeatedly to solve various decision-making problems across a diverse range of taxa. Owing to the rules' algorithmic similarity, their output could mistakenly be attributed to a single decision-making system. We argue that these design patterns, implemented by diverse suites of neural hardware, should nonetheless prove ubiquitous on evolutionary grounds, and that their simplicity and robustness should favor them both for convergent evolution and conservation within taxa. Such ideas represent a new opportunity for both systems theory and comparative biology, since the view of decision systems as evolving primarily to solve ecological problems demands renewed interest in both engineering disciplines and animal behavior. Indeed, the search for reusable design patterns in neural systems may provide a unifying framework for biological decision making in much the same way it has for vision and motor control.

3. Sparse coding of viewed natural behaviors in the prefrontal cortex

3.1 *Materials & Methods*

3.1.1 Behavioral paradigm

Three adult male rhesus macaques participated in this study. During an experimental session, subjects sat in a primate chair (Crist Instruments, Hagerstown, MD) in a darkened room with a computer monitor. Stimuli were displayed on the monitor under the control of custom scripts programmed in Matlab (The Mathworks, Natick, MA) with the Psychophysics Toolbox (Brainard 1997). Monkeys' eye movements were optically tracked at 1000 Hz via real-time pupil detection using an Eyelink 1000 system (SR Research, Mississauga, ON, Canada), and subjects were trained to interact with the task control system by making eye movements. The optical gaze tracker was calibrated prior to each behavioral run with a custom nine-point calibration routine. Juice rewards were delivered to a tube placed at the monkeys' lips from a reservoir by opening a solenoid valve, or by a peristaltic pump.

Monkeys initiated each trial by fixating a black circle positioned centrally on the computer monitor against a gray background for 500 ms. Upon successful fixation, the fixation target was removed, and either two (90% of trials) or one (10% of trials) eccentric colored choice targets were presented. Monkeys selected one of the targets by making a saccade, and the color of the target selected indicated a type of video presentation to follow. For each trial, the targets presented (the "menu") were randomly

selected from a larger set of possibilities. The complete set of options were targets yielding the following video presentation outcomes: Blank (indicating the presentation of a blank screen for 5 s), Switch (indicating the presentation of a 5 s video clip randomly chosen from the whole set), Continue (indicating the resumption of playback of the previous trial's video clip for an additional 5 s), and Repeat (indicating the repetition of the previous trial's 5 s video clip). In some sessions, Repeat was not included in the set of options. Upon registration of the monkey's choice, there was a brief (randomized between 100 and 600 ms) delay, followed by video playback for 5 s as determined by the selected target. Following video playback, there was a brief (randomized between 50 and 500 ms) delay, and then monkeys were rewarded with a drop of juice for successfully completing the trial. Juice rewards were fixed, not depending on monkeys' viewing decisions or gaze behavior. A visual schematic of the trial structure is illustrated in Figure 5.

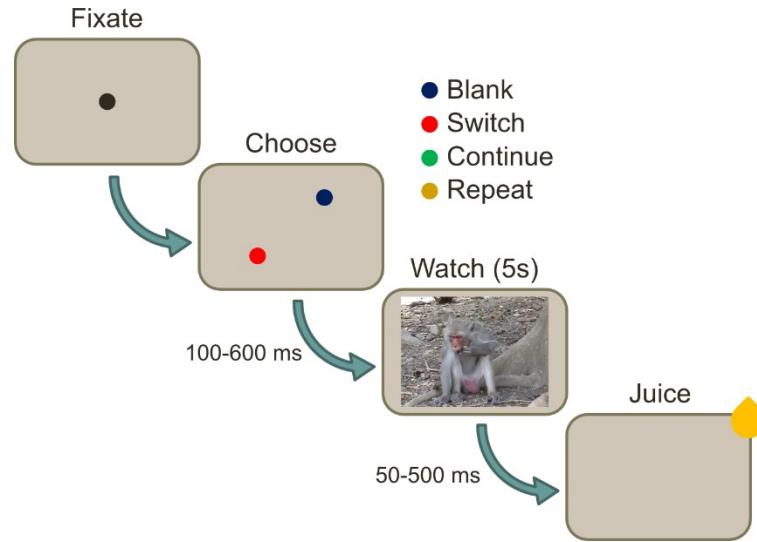


Figure 5: A schematic diagram of a task allowing monkeys to view and make decisions about social information.

3.1.2 Video stimuli

A stimulus set consisting of 4.87 hours of video in 429 files was recorded in August 2009 at the Caribbean Primate Research Center’s Cayo Santiago facility, a field station located on a small island populated by approximately one thousand rhesus macaques. Human researchers and technicians visit the island daily, and the monkeys are accustomed to the presence of human observers. During filming, the intended goal was to obtain a set of videos representing a wide range of natural behaviors. Videos included no cinematographic edits (e.g., cuts), but due to the highly dynamic and unpredictable environment of the island did include changes in zoom level and camera movements, in order to both capture behaviors of interest and to allow the videographer to safely move in response to potentially aggressive monkeys.

For compatibility with the experimental control software, all video files were converted to Apple Quicktime format at 610 width by 458 height resolution and 30 frames per second. Because the sound track's quality was highly variable, and frequently included vocalizations from animals out of frame which were unrelated to the visible events, audio was removed from all videos.

To assess subject animals' behavioral and neurophysiological responses to the conspecific behaviors displayed in the video stimuli, we coded all videos using an ethogram of rhesus macaque behavior. Our ethogram was developed based on a previously published ethogram of rhesus macaque behavior (Partan 2002), but tailored for the constraints of the videos and the demands of the study. Videos were coded using the software program Tinbergen, an open-source program written in Python by Geoffrey K. Adams to assist with the behavioral coding of videos of animals, available online at the URL <http://github.com/biogeol/tinbergen>. Tinbergen provides a simple video playback interface alongside a behavioral coding interface to allow a researcher to mark the starts, stops, and changes of state of behaviors seen in the videos in a project-specific ethogram. These events are stored in a simple plain-text file format which can be read and used by custom analysis software.

In addition to behaviors, additional descriptors of the video scenes were coded. Although the videos included no cinematographic edits (e.g., "cuts"), camera movements and changes in zoom level were included, and these events were coded

along with the behavioral events. As well as permitting the effect of exogenous changes in visual orientation to be investigated, camera movements are associated with large-scale optic flow. The number of monkeys present in the scene was also coded. Because the face and ano-genital regions are areas of particular biological significance, the view levels of face and hindquarters were also included in the coding scheme. The complete coding scheme, including behaviors and these view-related events (which in total we will refer to as the “ethogram” for simplicity despite its inclusion of non-behavioral observations) used for coding videos is shown in Table 1. Some ethogram entries were coded categorically, always in one of several mutually exclusive “levels,” while the remainder were coded in a binary fashion, always either occurring or not occurring.

Table 1: Ethogram used for coding videos

<i>Event name</i>	<i>Levels</i>	<i>Description</i>
Camera movement	holding	The camera is not moving much, holding a mostly stable image
	zooming in	The camera is holding a mostly stable image, but is zooming in
	zooming out	The camera is holding a mostly stable image, but is zooming out
	tracking	The camera is tracking the movements of a monkey, holding it largely stable against the moving background
	panning	The camera is moving without tracking anything, and the scene is still viewable
	slewing	The camera is moving wildly and the scene is difficult to view
Count	0	The number of monkeys visible in the scene, coded in approximately logarithmically-spaced levels
	1	
	2	
	3-5	

	6-10 >10	
Human visible		Whether human researchers are visible in the scene
Face visibility	none	No faces are visible with a view angle (the angle between the camera and the head axis) of less than ca. 100 deg.
	small	At least one face is visible with a view angle of less than ca. 100 deg., but none subtend a linear extent of greater than ca. 10% of the width of the video
	away	At least one face is visible with a view angle of between ca. 45 and 100 deg., and subtends a linear extent of greater than ca. 10% of the width of the video
	direct eye contact	As "away," but the view angle is less than ca. 45 deg. As "direct," but the monkey also appears to be looking directly at the camera
Ano-genital area visibility	none	The ano-genital area (AGA) is defined as the glabrous, red skin around the genitals, anus, and upper inner thigh (the "sex skin"), the scrotum on males, and the ischial callosities. The AGA of no monkeys is visible.
	small	At least one monkey's AGA is visible, but none are prominent on the screen.
	unsexed	At least one monkey's AGA is prominently visible, but the observer is unable to sex the animal
	male	At least one male monkey's AGA is prominently visible
	female	At least one female monkey's AGA is prominently visible
	both	At least one male and one female monkeys' AGA are visible
Drink		Any visible monkey is drinking water
Forage	no	No monkeys are engaged in foraging behavior.
	few	One or two monkeys are in a foraging state, characterized by various specific actions such as manipulating or ingesting food.
	many	Three or more monkeys are engaged in foraging behaviors. Due to the rapid increase in complexity of coding specific actions with this number of foraging monkeys, specific actions are only coded during the "few" condition.

Search		Searching for a food item by manipulating foliage, ground substrate, or a pile of provisioned food items
Grasp food		Reaching for a food item; begins at initiation of the reach movement and ends once the food item is in hand
Hold food		A food item is visibly held in hand or foot
Hold food in mouth		A food item is visibly held in the mouth; carrying food in cheek pouches is not included
Manipulate food		Active manipulate of a food item other than grasping or ingesting, e.g. "washing"
Ingest food		A food item is brought to the mouth. Begins at initiation of arm movement and ends when the food item is consumed or the hand leaves the mouth area. Includes movements of food item towards the mouth that do not actually result in consumption, e.g., "sniffing"
Ingest grooming manipuland		A hand is brought from a grooming target (self or partner) to the mouth, as if the monkey is consuming a parasite plucked from the skin
Chew		Rhythmic jaw movements; an entire bout of chewing is coded rather than individual bites
Retrieve from pouch		A food item is brought out from the cheek pouches
Heave		A full body movement involving tightening of the abdomen and straightening of the esophagus, as though the monkey is preparing to vomit or in the act of vomiting
Scratch		Rhythmic, vigorous movement of the hand or foot against the monkey's own body
Autogroom		Self-directed grooming behavior; involves more finely controlled hand and finger movements than scratching
Allogroom		Other-directed grooming behavior; allogroom and autogroom involve similar motor actions, differing only in the target
Solicit allogroom		A monkey approaches another monkey and sits or lies down, presenting itself for grooming
Aggression level	none	No aggression is present in the scene
	unidirectional	A single individual is displaying aggressive behaviors

	bidirectional	Two individuals are displaying aggressive behaviors toward each other
	joint	Two or more individuals are jointly displaying aggressive behaviors, the target(s) of which may or may not be visible
	intercoalition	Two "coalitions" are displaying aggressive behaviors toward each other
Strike		A monkey makes brief, aggressive physical contact with another, excluding shoving infants by adult females
Grapple		Two monkeys engage in prolonged aggressive physical contact
Lunge		A monkey makes a short, aggressive movement toward another monkey
Withdraw		A monkey backs away from another monkey while making aggressive or submissive displays
Charge		A monkey makes a rapid, prolonged aggressive movement toward another monkey
Flee		A monkey moves away from another monkey at high speed
Chase		One monkey charges another monkey as it flees
Threaten		A monkey performs a threat display, characterized by a round open mouth, prolonged staring, head bobbing, piloerection, and erect posture
Mounted threaten		One monkey mounts another while both threaten a third monkey
Submit		A monkey performs a submissive display, characterized by bearing teeth, squeaking, and withdrawn posture
Displace		One monkey walks directly toward or near another, which moves away
Lean away		A monkey posturally shifts away from an approaching conspecific without fully withdrawing or displacing
Avoid		A monkey pauses or alters course during movement to maintain greater distance to another monkey
Branch display		A monkey vigorously shakes a branch or branch-like object (e.g. a metal pole)

For use with quantitative analyses, the coding ethogram was then slightly altered to represent categorical entries as a set of related binary entries. Additionally, some behaviors in the original ethogram occurred insufficiently frequently to permit analysis with neurophysiological data (see below), but belonged to a natural semantic class of behavior, and were thus pooled for this purpose. Table 2 shows these derivative ethogram elements.

Table 2: Derived ethogram elements

<i>Event name</i>	<i>Definition</i>
Camera zooming in	"Camera movement" is "zooming in"
Camera zooming out	"Camera movement" is "zooming out"
Camera tracking	"Camera movement" is "tracking"
Camera panning	"Camera movement" is "panning"
Camera slewing	"Camera movement" is "slewing"
Count >= 1	"Count" is "1," "2," "3-5," "6-10," or ">10"
Count >= 2	"Count" is "2," "3-5," "6-10," or ">10"
Count >= 3	"Count" is "3-5," "6-10," or ">10"
Count >= 6	"Count" is "6-10" or ">10"
Count >= 11	"Count" is ">10"
Visible face	"Face visibility" is "small," "away," "direct," or "eye contact"
Prominent face	"Face visibility" is "away," "direct," or "eye contact"
Direct face	"Face visibility" is "direct" or "eye contact"
Eye contact	"Face visibility" is "eye contact"
Visible genitals	"AGA visibility" is "small," "unsexed," "male," "female," or "both"
Prominent genitals	"AGA visibility" is "unsexed," "male," "female," or "both"
Male genitals	"AGA visibility" is "male" or "both"
Female genitals	"AGA visibility" is "female" or "both"
Foraging	"Forage" is "few" or "many"
Group foraging	"Forage" is "many"
Any aggression	"Aggression level" is "unidirectional," "bidirectional," "joint," or "intercoalition"

Mutual aggression	"Aggression level" is "bidirectional," "joint," or "intercoalition"
Joint aggression	"Aggression level" is "joint" or "intercoalition"
Intercoalition aggression	"Aggression level" is "intercoalition"
Attack	Any of "Strike," "Grapple," "Charge," "Lunge," or "Chase"
Threaten	Any of "Threaten," "Mounted threaten," or "Branch display"
Submit	Any of "Withdraw," "Flee," "Chase," "Submit," "Displace," "Lean away," or "Avoid"

When used in the quantitative analyses below, ethogram entries were coded as "1" when present and "0" when absent in the video scene. A schematic of these values during a selected section of an example session is shown in Figure 6.

temporally downsampled to 30 Hz via boxcar averaging, minimizing the phase delay between the video frame drawing times and the centers of the boxcar windows, and spatially rescaled to a normalized coordinate system in which the vertical extent of the video is 1 and the origin is in the center of the video. We will refer to the reported position of gaze in this transformed coordinate system as the *gaze focus*, and note that each video frame in a presentation had exactly one associated gaze focus value, unless the optical tracker failed to detect the pupil due to the monkey's eye being closed. We further defined a criterion for gaze focus being *in frame*: it must be within a rectangle defined by the video frame with a buffer on each side of 0.05 normalized units (i.e., 5% of the height of the video). Numerically, the vertical coordinate must be in the interval $[-0.550, 0.550]$ and the horizontal coordinate in the interval $[-0.716, 0.716]$ for gaze focus to be in frame.

Videos of conspecifics frequently elicit characteristic gaze orientation patterns in rhesus macaques (Shepherd et al. 2010, Mosher et al. 2011, McFarland et al. 2013). Qualitatively, this result was borne out in our task as well. To quantify the extent to which monkeys repeated their gaze patterns across multiple viewings of the same video sequence, we compared all observations of gaze focus for individual frames. We introduce a novel metric summarizing the relationship between these observations, which we term *gaze consistency*. The metric is mathematically analogous to the potential energy of a collection of electrically-charged particles, with each “particle” representing

the focus of a monkey's gaze in a particular video frame from a single viewing. That is, for a particular video frame, we collect all observations of gaze focus which were in frame, each a vector denoted by \mathbf{x} with a subscript. The "energy" associated with a pair of gaze observations \mathbf{x}_i and \mathbf{x}_j is then given by

$$v_{ij} = \frac{d_0}{\max(\|\mathbf{x}_i - \mathbf{x}_j\|, d_0)}$$

where d_0 represents the smallest distance between two points considered to represent a meaningfully different focus of gaze (here, d_0 was 1% of the height of the video, 4.58 video pixels, or approximately 0.23 degrees of visual arc). The gaze consistency for a particular video frame is then the mean "energy" between all pairs of gaze points,

$$g = \frac{2}{n(n-1)} \sum_i \sum_{j < i} v_{ij}$$

Gaze consistency is a metric which has a theoretical maximum value of 1 if all observations of gaze were to fall within a disc of diameter d_0 , and decreases toward zero with greater disparity in the focus of gaze upon repeated viewings of the same video frame. Importantly, gaze consistency is only moderately sensitive to the presence of multiple clusters of gaze focus. A single tight cluster of gaze focus yields a greater consistency value than two such clusters, but two clusters yield a greater gaze consistency than a diffuse "cloud" of gaze focus, a feature which is not true for, e.g., the covariance of gaze focus. Thus gaze consistency values tend to track, in an ordinal fashion, the qualitative sense of "clustered-ness" in the gaze focus observations.

In addition to gaze consistency, we also introduce the related metric *gaze typicality*, which indexes how similar a particular observation of in-frame gaze focus is to other observations for the same video frame. This is defined as the average “energy” of all pairs involving that observation:

$$\tau_i = \frac{1}{n-1} \sum_{j \neq i} v_{ij}$$

Like gaze consistency, gaze typicality theoretically ranges between zero (highly atypical gaze) and one (highly typical gaze). The gaze typicality for a particular observation of gaze focus is high when there are many similar such observations for that frame, and low when it is distant from other observations of gaze focus.

Conceptually, gaze consistency is a property associated with a video frame, and periods of high gaze consistency indicate video sequences which present features or events which are potent attractors of gaze, irrespective of the particular nature of those features or events. By contrast, gaze typicality is a property associated with an individual gaze trace, and indicates the extent to which the animal’s gaze during a particular video presentation was attracted to the regions which, on average, tend to attract gaze. Gaze consistency and typicality are useful metrics for the interpretation of gaze behavior in a large, diverse stimulus set, offering the advantages of being location-independent (i.e., it does not matter where in the visual field a feature may be) and simple to calculate. However, it is important to interpret these values carefully, as they are highly sensitive to the relative size of features in the visual field. In particular, a

scene containing a feature which is a potent attractor of gaze but spans a large region of the visual field will tend to yield lower gaze consistency than one which contains a similar feature which spans a smaller region of the visual field.

Gaze consistency and gaze typicality were calculated only for frames with at least five in-frame observations of gaze focus, and gaze typicality was additionally only calculated for in-frame gaze focus. These criteria yielded gaze consistency values for 501,377 video frames (out of 527,480 total frames) and 11,520,724 gaze typicality observations (out of 12,619,650 frame presentations). In addition to these metrics, we also examined whether each gaze focus was in frame (*gaze onscreen*), and the overall proportion of in-frame gaze focus for all viewings of each frame (*proportion gaze onscreen*).

3.1.4 Surgery

All procedures were approved by the Duke University Institutional Animal Care and Use Committee and were designed and conducted in compliance with the Public Health Service's Guide for the Care and Use of Laboratory Animals. A small prosthesis for head restraint was implanted in each monkey using standard sterile surgical techniques. Following a six week recovery period, the monkeys were trained to perform eye movement tasks for fluid rewards. A second surgery was then performed to place a Cilux plastic recording chamber (Crist Instruments, Hagerstown, MD) above the prefrontal cortex. Monkeys received analgesics and antibiotics during recovery from

both surgeries. After implantation, the recording chamber was routinely treated with antiseptic flushes, including before and after each recording, and between recordings was kept sealed with sterile caps.

3.1.5 Electrophysiological recordings

We recorded action potentials from 63 single units from the lateral prefrontal cortex (LPFC) in the principal sulcus (monkey C: 31, monkey E: 32) and 56 single units from the orbitofrontal cortex (OFC) (monkey C: 29, monkey E: 27). Recordings were performed using either single-wire tungsten microelectrodes (Fred Haer Co., Bowdoin, ME) or an 8-channel multicontact probe (U-Probe, Plexon Inc., Dallas, TX) configured with two tetrode clusters. During recordings, the probes were guided to the intended recording site using a hydraulic microdrive (David Kopf Instruments, Tujunga, CA).

Target recording sites were determined through the conjunction of several techniques. All recordings were performed using a standard recording grid (Crist Instruments, Hagerstown, MD). Structural MRIs were obtained for monkeys E and C, and the distances between the dorsal surface of the brain below each grid hole and the intended recording sites were measured. Observations of changes in audible broadband power and multi-unit activity electrophysiological landmarks during the positioning of the electrode prior to recording (as determined by sending the amplified electrophysiological signal to a speaker) were used to further confirm the recording site. Finally, additional anatomical information was obtained via ultrasound imaging

through the recording chamber (Glimcher et al. 2001) using a hand-held digital ultrasound device (SonoSite 180, FUJIFILM SonoSite, Inc.).

3.1.6 Statistical models and regularization technique

Statistical analyses were performed using standard packages and custom software in Matlab (The Mathworks, Natick, MA), R (R Core Team 2014), and MySQL.

In investigating the effects of viewed behaviors on subject monkeys' behavioral and physiological responses, an unusual challenge is posed by the complexity of the viewed naturalistic scenes. Each item in the ethogram represents a potential regressor which could be included in a statistical model. Crucially, this set of regressors is highly correlated, and including all these regressors in a standard statistical model (e.g., a generalized linear model [GLM]) substantially increases the likelihood of over fitting, unstable fits due to multicollinearity, and reducing power due to over-partitioning the variance (a version of the so-called "large p, small n" problem). To address these issues, we modeled behavioral and physiological responses with GLMs with so-called "elastic net" regularization (Zou & Hastie 2005), a technique which has been successfully applied in several areas of biostatistics and which has empirically been demonstrated to perform well with large numbers of correlated regressors (Hastie et al. 2009). A full discussion of elastic net regularization is beyond the current scope, but briefly, elastic net regularization introduces a penalty term, λ , acting on both the sum of the absolute values of the regression coefficients (the L1-norm) and the sum of the squares of the

regression coefficients (the L2-norm). These two penalties are weighted relative to one another by another parameter, α . Due to the L1 penalty, as λ increases, coefficients for regressors which are only weakly associated with the model's response variable tend to be forced to exactly zero, effectively excluding "non-significant" regressors from the model completely. Due to the L2 penalty, coefficients for regressors which are highly correlated with one another tend to be included or excluded from the model as a group.

A key challenge with the elastic net model centers on finding an appropriate choice of the parameters α and λ , since these control the relative weight given to the model error and model complexity terms. The strategy most commonly used in such situations is cross-validation (Hastie et al. 2009). The model is trained using multiple parameter values on a subset of the original data, and its predictive accuracy for each set of parameters is then calculated based on the data not used for training (hold-out or test data). The values of the parameters selected are those which produce best performance on the held-out data. In k -fold cross-validation, this process of training and testing is repeated k times, each time with a different subset of data withheld. The test accuracy from these k steps is then averaged to give a more accurate estimate of model error and thus better parameter selection.

The performance of the elastic net model at various values of λ can be compared by calculating the deviance, a goodness of fit metric related to the log likelihood of the data, given the model and parameters. The distribution of the deviance function is

estimated by k -fold cross validation: the observations are partitioned into k test sets, and for each test set the model is first fit on the complement training set and then deviance is computed on the test set, yielding k estimates of the deviance function. Conventionally, the largest value of λ which produces a deviance within one standard error of the minimum deviance is selected (λ_{1SE}) (Friedman et al. 2010). Intuitively, this may be thought of as selecting the most conservative (i.e., most highly-constrained) model which is reasonable given the (more permissive) best-performing model. To illustrate the selection procedure, Figure 7 presents the relationship between λ and the coefficient fits and cross-validated deviance for the neuronal firing rate model fit result shown in Figure 17.




Figure 7: Illustration of coefficient selection under elastic net regularization. (a) As the elastic net constraint parameter λ increases, regressor coefficients are constrained towards zero, with weaker predictors of the response variable being forced to exactly zero at weaker constraint levels, and strong predictors being forced to zero only at strong constraint levels. Each trace represents the value of one coefficient at various constraint levels. (b) Model performance varies with λ . Model performance at various constraint levels is assessed by cross-validation, with lower deviance values indicating superior performance. Error bars represent standard error across the cross-validation test sets. For this illustrative model, a broad range of weak constraint levels performed similarly well, and the strongest of these (λ_{1SE}) was selected.

Because elastic net regression is a relatively new technique, hypothesis testing under it remains an active research topic (c.f. Lockhart et al. 2014). However, in cases where such tests are not available, we can use permutation tests to generate

distributions from the null hypothesis for comparison with observed data. Such tests entail comparing model performance on the true data set to model performance on artificial data sets formed by randomizing the relationship between the regressors and the response variable. These tests do not necessarily provide p-values, but do provide rigorous and meaningful control over Type I error rates.

We applied elastic net regularization to several models which included ethogram variables as regressors, using the R package “glmnet” (Friedman et al. 2010). In all cases we identified λ_{1SE} via 20-fold cross validation and report the coefficients associated with this constraint level.

3.1.7 Analysis of video selection decisions

We observed 86,893 viewing decisions (monkey C: 42,229, monkey E: 29,951, monkey H: 14,713), 74,874 of which occurred following the presentation of a movie rather than a blank screen on the previous trial (monkey C: 34,296, monkey E: 27,109, monkey H: 13,469). To determine the extent to which the offered menu and behaviors seen in the previous trial’s movie influenced decision-making, we constructed a model as follows. First, the menu options were assigned an arbitrary ranking: Blank=0, Repeat=1, Continue=2, Switch=3. This selection of rank ordering will have no influence on the inferences of the model; it simply defines a convention for coding its inputs and outputs. For each (non-forced) decision, the response variable was assigned 1 for selection of the higher-ranked option, and 0 for the lower-ranked option. Next, for each

of Repeat, Continue, and Switch, a regressor was assigned 1 if it was the higher-ranked option of the menu, -1 if it was the lower-ranked option of the menu, and 0 if it was not offered. (For example, if on a particular trial the menu was Repeat and Switch, then the regressors for Repeat, Continue, and Switch would be -1, 0, and 1, respectively.) Blank was treated as a baseline comparison condition and thus not assigned a regressor. Then, ethogram variables were assigned 1 if the entry occurred at all during the previous trial's video presentation, and 0 otherwise. Finally, we identified a model according to which each viewing decision d was a Bernoulli trial with parameter p_d , with p_d depending on the menu regressors and ethogram variables via the logistic function through the following equation (**Model 1**):

$$\ln\left(\frac{p_d}{1-p_d}\right) = \sum_i v_i m_{d,i} + \sum_i \sum_k \beta_{i,k} m_{d,i} x_{d,k}$$

where i indexes the three menu regressors, k indexes the ethogram variables, $m_{d,i}$ is the menu regressor for option i , $x_{d,k}$ is the ethogram variable for entry k , v_i is a coefficient representing the overall utility of option i , and $\beta_{i,k}$ is a coefficient representing the change in utility of option i when ethogram item k was previously viewed. We then fit the coefficients v_i and $\beta_{i,k}$ to the data using the previously described elastic net technique, with $\beta_{i,k}$ subjected to the elastic net penalty while v_i were not.

3.1.8 Analysis of gaze during video presentation

To assess subject monkeys' overall engagement with the presented videos as assessed by oculomotor orienting toward the stimuli, we pooled gaze data from 84,131

five-second video presentations (monkey C: 39,945; monkey E: 26,644; monkey H: 17,542), and for each frame in each presentation (521,059 frames total) identified whether gaze focus was within frame, “gaze onscreen” as defined above. We constructed a model treating each observation of gaze onscreen as a Bernoulli trial with parameter p_f , with p_f depending on the state of the video’s ethogram for frame f via the logistic function, (**Model 2**)

$$\ln\left(\frac{p_f}{1-p_f}\right) = \sum_k \beta_k x_{f,k}$$

where $x_{f,k}$ represents the value (zero or one) of the k ’th ethogram element for video frame f , and β_k is an associated coefficient. This model treats the monkey’s probability of looking at the presented video as a function of the sum of independent effects associated with each of the various ethogram events that may be occurring at any time. We then fit the coefficients β to the data using the previously described elastic net technique.

To assess how viewed behaviors might influence monkeys’ gaze behavior within the video frame, we also computed gaze consistency for this set of observations. To improve interpretability, we standardized gaze consistency by subtracting its mean and dividing by its standard deviation over the full set of observations. We then constructed a simple model treating each such observation as being drawn from a normal distribution with mean μ_f , with μ_f depending on the state of the video’s ethogram for frame f (**Model 3**):

$$\mu_f = \sum_k \beta_k x_{f,k}$$

(Note that β_k here is a distinct symbol from its appearance in the previous equation.)

The coefficients β were then fit to the data using the previously described elastic net technique.

3.1.9 Analysis of neuronal firing rates

To explore the relationship between task events and prefrontal neuronal firing rates, we constructed peri-stimulus time histograms with $33 \frac{1}{3}$ ms bin widths (i.e., 30 bins per second, chosen to match the video frame rate).

To investigate the relationship between viewed behaviors and neuronal firing rates during video presentation, we pursued a model wherein the number of spikes produced by a neuron followed a Poisson distribution with parameter $\phi_{n,t}$, where n indexed distinct video presentations and t the time within presentation. This rate parameter depended on an intrinsic baseline rate ϕ_0 modulated by a set of ethogram variables $x_{k,n,t-\tau}$ where we introduced the time delay τ to account for the latency for visual information to arrive in the prefrontal cortex. That is, spikes at time t occurred in response to visual stimuli at time $t - \tau$.

However, we also observed that two additional factors influenced neuronal firing: first, many units showed characteristic firing rate profiles over time during video presentation (Figure 13, Figure 14); and second, many units exhibited slow “drift” in their baseline firing rates which occurred on the timescale of minutes. The characteristic

within-presentation timecourse we labeled π_t , which was calculated using peri-stimulus time histograms over all video presentations. The drift factor we labeled δ_n , and calculated it according to the following algorithm. 1: For each trial n , we calculated the number of spikes observed in the 500 ms prior to the start of the trial, multiplied by two to give the “baseline spike rate.” 2: Smooth these values with a 21-trial moving average. 3: To compensate for very low firing rates in some units, add a “floor effect” of 0.0476 Hz, equivalent to observing one spike in 42 500-ms windows and representing a reasonable upper bound on the unit’s true baseline firing rate if no spikes were observed. 4: Normalize these values by dividing by the mean over all trials, and assign this as δ_n .

Finally, we constructed a model for the firing rate of prefrontal neurons during video presentation relating these various terms (**Model 4**):

$$\ln(\phi_{n,t}) = \ln(\phi_0 \pi_t \delta_n) + \sum_k \beta_k x_{k,n,t-\tau}$$

where β_k are coefficients representing the effect of each of the ethogram variables on the neuron’s firing rate. Note that by exponentiating, this equation can be rewritten as:

$$\phi_{n,t} = \phi_0 \pi_t \delta_n \prod_k \gamma_k^{x_{k,n,t-\tau}}$$

where $\gamma_k = \exp(\beta_k)$, which emphasizes that this model treats each item in the ethogram as having a gain-like effect on the firing rate of the cell.

Because some behaviors in the ethogram occurred relatively infrequently, it was possible for our observations of neural firing to include only a small number of presentations involving these behaviors. This was of concern not only because of the small sample size, but also because with only a small number of exemplars of a particular behavior, it is impossible to determine whether changes in neural firing rate were due to the behavior itself or an unrelated event that happened to co-occur with the behavior in one or two of the presentations. Therefore, for each item in the ethogram, we required that it be present in at least 5 distinct presentations for its regressor to be included in the model.

We were also interested in how PFC neurons' firing rates related to monkeys' gaze behavior. As estimators of intrinsic elements of the video scene that influenced monkeys' gaze, we included gaze consistency and the proportion of gaze-out-of-frame observations as regressors along with the ethogram variables. As estimators of per-view decision processes, we included gaze typicality and gaze-out-of-frame as regressors as well.

A final problem which remained was the principled selection of the "information latency" parameter τ and the window duration w used to count spikes. To select τ , we observed that those units in our population which exhibited changes in firing in response to video onset did so with a latency of approximately 100 ms (Figure 14), a value which is in accord with previous studies of these areas (Thorpe et al. 1983), and so

selected $\tau = 100$ ms. The selection of w offered greater challenges. Because both neuronal firing rates and ethogram regressors are temporally autocorrelated, selection of an analysis window shorter than the scale of neural firing rate autocorrelation would bias our model toward false positives, due to the fact that adjacent bins would be highly correlated in the data while treated as independent by the model. Conversely, selection of too large an analysis window would risk failing to adequately represent some of the briefer behaviors in the ethogram. Therefore we explored neuronal firing rate autocorrelations during video presentation in our data set (Figure 8).




Figure 8: Autocorrelation in neuronal firing rate during video presentation, LPFC example unit

On the basis of this exploratory analysis, we selected a 200 ms window as sufficiently conservative. Thus, each presentation was divided into twenty-five 200 ms observations of neuronal firing rate, each of which could be considered roughly independent for the purposes of statistical analysis. Because the observation window was larger than a single $33 \frac{1}{3}$ ms video frame, ethogram values were averaged over the

window to compute the regressors. We then fit the coefficients β to the data using the previously described elastic net technique.

3.2 Results

3.2.1 Monkeys' viewing preferences for naturalistic videos

All three monkeys displayed marked preferences among the four menu options (Figure 9). Notably, for all three monkeys, the least preferred option was Blank, and the most preferred was Switch, suggesting a general preference for video stimuli over blank screens, and for unpredictable video stimuli over stimuli predictable from the previous trial's presentation.

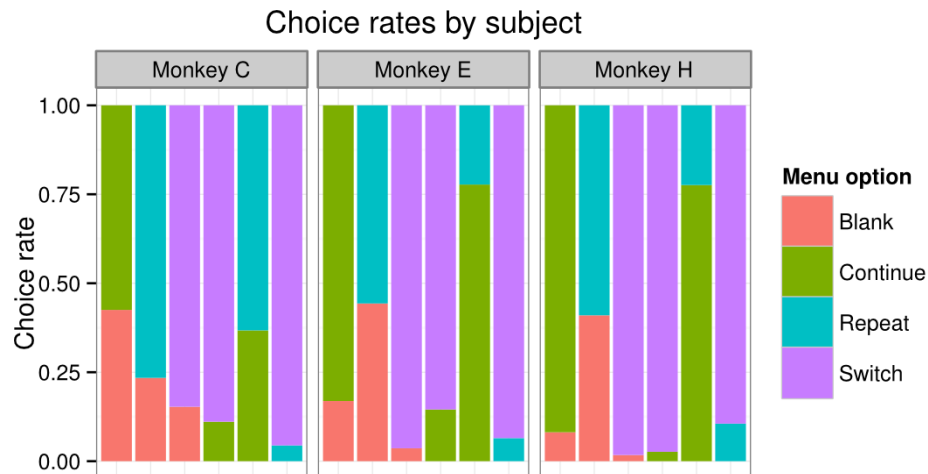


Figure 9: Each subject exhibited clear preferences among viewing outcomes. Each vertical pair of bars represents one “menu” of paired options presented to the subject. The relative height of each bar indicates the proportion of decisions for the associated option in that menu. (Monkey C: 42,229 decisions, Monkey E: 29,951 decisions, Monkey H: 14,713 decisions)

Contrary to our original expectations, our model of presentation decisions, Model 1, failed to discover any effect of viewed behaviors on subsequent decision-

making, yielding no non-zero coefficients for the option-ethogram interaction terms for any of the subject monkeys. However, the model did allow us to estimate utility coefficients for the menu options, v_i (Figure 10).

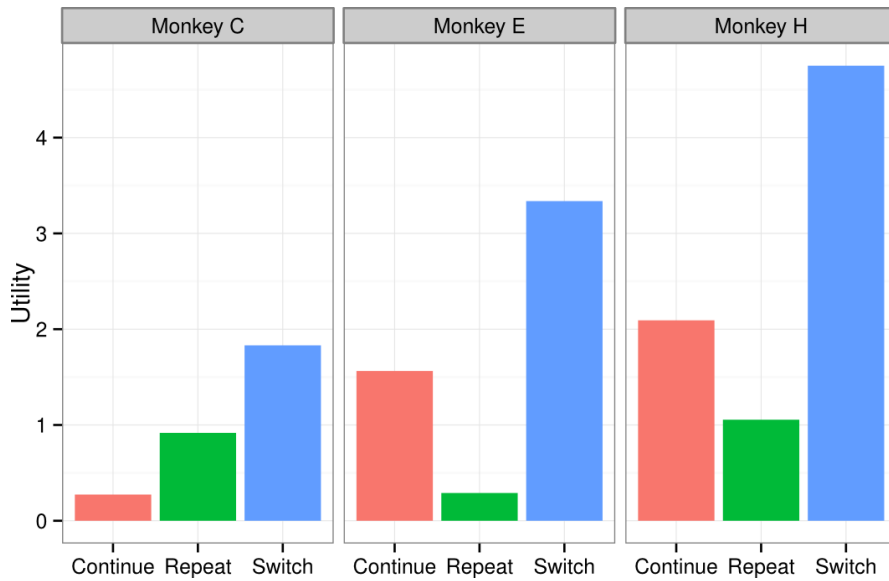


Figure 10: Menu option utilities for each subject

From the original construction of the model, the probability of a monkey selecting one option from a pair is given by the logistic function of the utility difference between the two options. Intuitively, the greater the utility difference between two options, the more likely the monkey is to select the larger-utility option. Had any effects of viewed behaviors been found, these would have been interpretable as additional (positive or negative) utility imparted to a viewing outcome by the previously-viewed behavior. The absence of such utility suggests that either monkeys' decision-making was driven solely by presentation outcome type rather than video subject matter, or the ethogram failed to capture those aspects of the viewed behaviors which were

biologically meaningful to subject monkeys. However, the relationship between subject monkeys' gaze behavior and the viewed behaviors suggests that the ethogram did in fact capture elements of the video presentations which were meaningful to subject monkeys.

3.2.2 Viewed behaviors influence monkeys' gaze behavior

The fit for the model of gaze onscreen, Model 2, against the data revealed that most of the ethogram entries had some influence on the likelihood that gaze focus was in frame (Figure 11).




Figure 11: Features of viewed video scenes, including conspecific behaviors, influenced monkeys' decisions to look at or away from the presented video at each frame.

The result of this model may be interpreted analogously to that of the decision-making model. At any time, the subject monkey is deciding whether to direct its gaze toward or away from the video frame, and that decision a function of the utility of the video at each time. Overall, the presented videos carry a high utility, as evidenced by the subject monkeys' overall decision to look in frame 91.7% of the time (out of 12,619,650 frame presentations) and indicated by the large intercept term fit by the model, but viewed features and behaviors add to or subtract from that utility.

Of particular note, camera movements in general seemed to enhance the visual interest of the video, with the exception of the wild "slewing" camera movements during which individual features in the video were difficult or impossible to distinguish. "Tracking" camera movements, in which the camera followed a moving animal to stabilize it in the frame, were particularly efficacious at attracting gaze to the scene. Monkeys' probability of looking at the video also increased with increasing numbers of conspecifics visible, with a particularly large increase associated with the second visible monkey ("Count ≥ 2 "). By contrast, and perhaps somewhat counter-intuitively, the prominent visibility of a face ("Face prominent") decreased the probability of looking at the video. However, this is consistent with previous findings that the utility of social information may be decoupled from looking-time in rhesus macaques for face stimuli (Deaner et al. 2005), perhaps because direct gaze is a threatening behavior in this species (van Hooff 1967).

In addition to these purely visual features, several behaviors significantly influenced the decision to look at the video. Aggressive behavior (“Any aggression”) increased the probability of looking at the video, with prolonged aggressive physical contact (“Grapple”) being particularly effective at attracting gaze to the scene. Submissive displays (“Submit”) were also effective attractors of gaze to the overall scene. Among the foraging behaviors, “Group foraging,” “Hold food in mouth,” and “Retrieve from cheek pouch” were particularly effective at attracting gaze to the video, the first of these likely reflecting the high information density when numerous conspecifics are behaving in the scene, and the latter two possibly indicating that close proximity of a food item to a conspecific’s mouth is highly salient for rhesus macaques. Somewhat surprisingly, “Allogroom” was associated with a reduction in the probability of looking at the video, despite its well-established importance for establishing and maintaining social relationships in macaques (Dunbar 2010). We speculate that because allogrooming is typically a lengthy behavior during which novel information for an observer is unlikely after the initial observation, the marginal utility of continued observation is low, decreasing the utility of continuous attention to the behavior for an external observer.

Beyond the decision to look in or out of frame, we were also interested in characterizing monkeys’ gaze behavior within the presented video frames, to better understand how viewed behaviors influence visual exploration. Theoretically, we might

have modeled this as a multinomial decision process analogous to the binomial decision to look within or away from the video screen, wherein the visual scene would be partitioned into regions identified by the features and behaviors occurring within them, each imparting a utility toward the subject's decision to direct its gaze to that region. However, hand-coding all such regions in over 500,000 video frames in order to build the model was impractical. Nevertheless, our ethogram gave us the time, if not the location in space, of our features of interest. Therefore we introduced the gaze consistency metric to attempt to index the reliability in the gaze behavior without explicit reference to spatial location within the video frame. Our inference was that the same process which would reliably drive gaze toward specific features in the scene would tend to reliably influence gaze consistency. Fitting our model of gaze consistency to the data revealed that it was modulated by a number of features and behaviors (Figure 12), demonstrating that despite its limitations it was successful in indexing some aspects of the subjects' reliability in gaze behavior with respect to the ethogram.




Figure 12: Differences in the reliability of gaze across repeated viewings of the same video sequence associated with features of the video scene, including conspecific behaviors, were successfully captured by the gaze consistency metric and associated model.

However, the nature of the gaze consistency metric means that the results of this model fit must be interpreted with some care. For example, observing the original videos with pooled observations of gaze focus makes it apparent that monkeys tend to consistently foveate male genitals when they are presented, yet the model reports a low gaze consistency associated with periods when male genitals are visible on screen. (Although note that because “Male genitals” is a nested regressor within “Prominent genitals” and “Visible genitals,” the total gaze consistency prediction for periods when male genitals are visible on screen is therefore the sum of these three values. The comparatively large negative value for “Male genitals” indicates a low gaze consistency compared to what would be expected for prominently visible genitals overall.)

However, this result is sensible when the nature of the feature is considered. Because of the scrotum, the adult male ano-genital area (AGA) is comparatively larger and visible at a wider range of viewing angles than that of females or juveniles (juveniles comprising the majority of the “unsexed” prominently visible AGAs in the video database). This means that the male AGA tends to subtend a larger area for visual exploration than that of females or unsexed individuals even at the same relative distance, camera zoom level, and posture. Thus, multiple gaze foci within the male AGA are more likely to be more widely separated compared to multiple gaze foci within a non-male AGA, hence the negative effect of the factor “Male genitals” on gaze

consistency. The result of the model for this feature reveals a limitation of gaze consistency as a metric.

Despite this important caveat, the model fit to the gaze consistency data reveals a number of intriguing patterns. As with gaze-in-frame probability, gaze consistency decreased with the presence of a prominently visible face in the scene. However, unlike gaze-in-frame probability, which showed little change as the view of a face moved from prominent, to direct, to making eye contact with the camera, the trend for gaze consistency reversed, rising with these increasingly “intense” views of a face. This is consistent with previous findings that monkeys’ scan patterns of faces tend to be highly stereotyped. Also notable is the comparatively large increase in gaze consistency associated with “Submit,” which parallels the increase in gaze onscreen probability, suggesting that submissive displays were of great visual interest to subject monkeys. By comparison, the increase in gaze consistency was much more modest for “Threaten,” and no effect was found for this behavior on gaze-in-frame probability. We speculate that this difference may exist because while both of these behaviors are highly informative social signals, subject monkeys may have been more prone to quickly avert their gaze from threat displays, as under natural conditions prolonged eye contact with a threatening conspecific could provoke an attack.

3.2.3 Prefrontal neurons respond robustly to videos of conspecifics and sparsely represent viewed behaviors

To explore the relationship between task events and prefrontal neuronal firing rates, we constructed peri-event time histograms for all main task events (fixation onset, choice target onset, choice registration, start of video playback, end of video playback, and juice delivery). A striking result was that many units in both OFC and LPFC displayed strong responses to video playback (Figure 13).

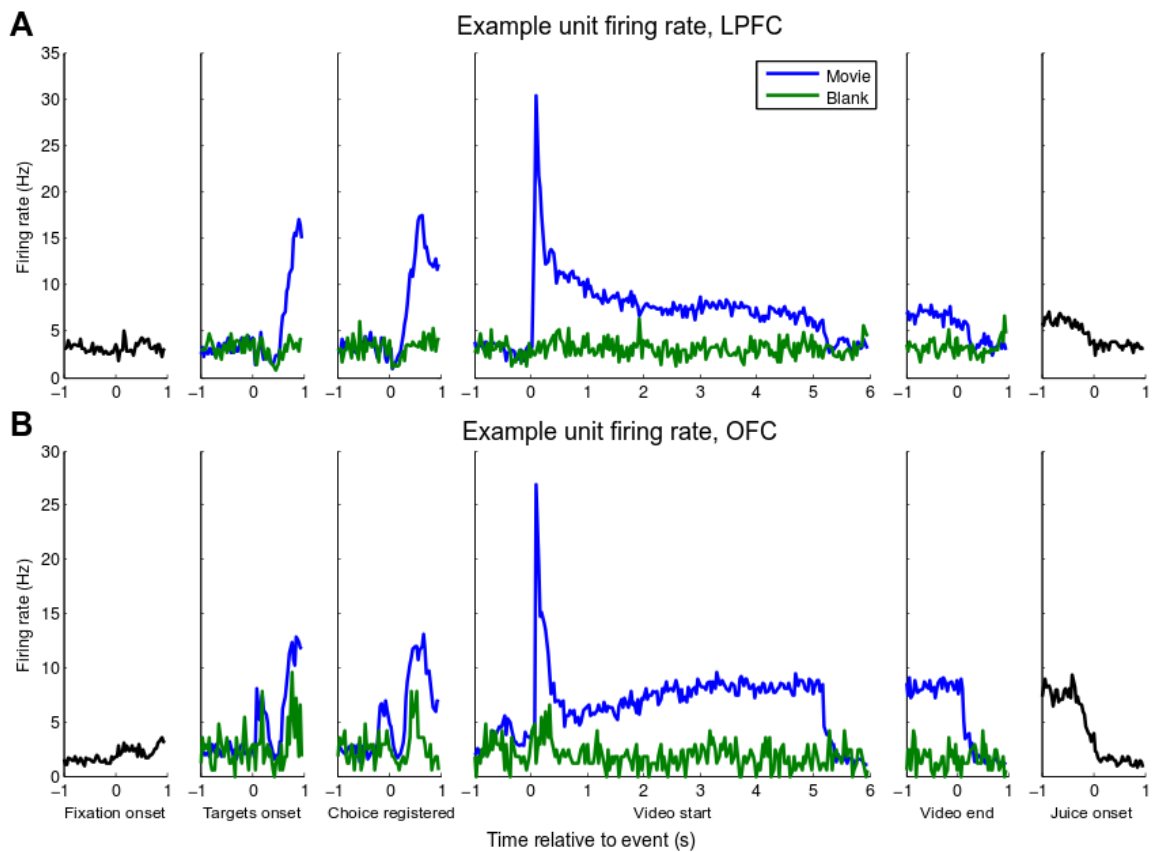


Figure 13: Individual PFC units respond strongly to videos of conspecifics. Each panel displays the mean firing rate of a unit in an epoch triggered by one of the major task events. Epochs are partially overlapping. (a) Example LPFC unit. (b) Example OFC unit.

To assess how the population of units in each area as a whole responded to each of the main events in the task, we also examined the mean normalized firing rate across the population of units within LPFC and OFC (Figure 14).

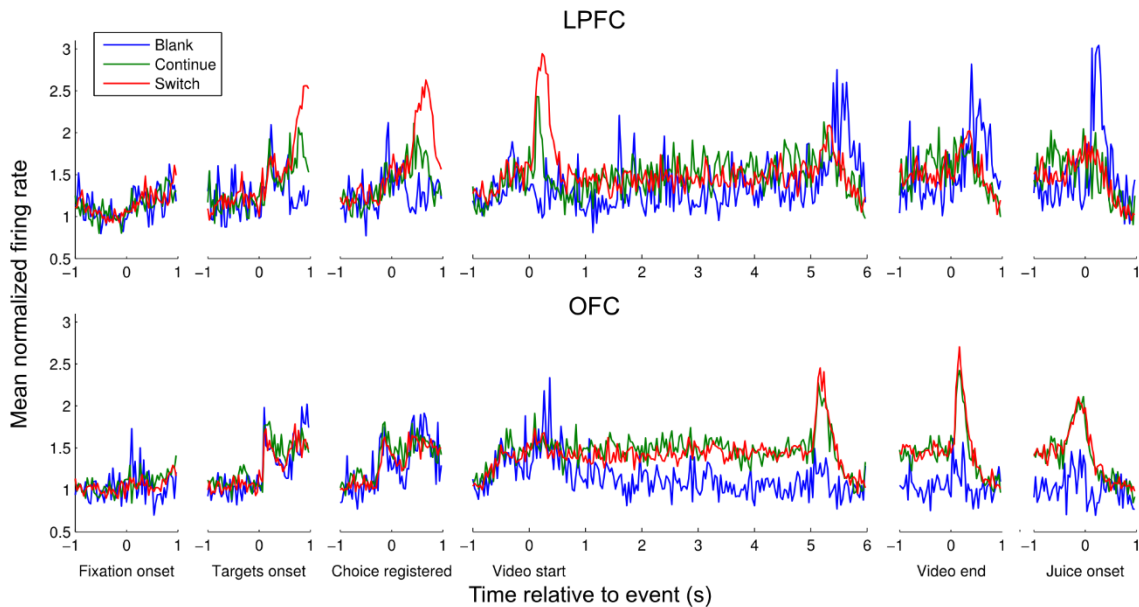


Figure 14: PFC population-averaged responses to task events.

For each unit, firing rates were normalized to the 500 ms prior to fixation onset. Each trace represents the mean normalized firing rate across the population for those trials in which the monkey selected “Blank” (blue), “Continue” (green), or “Switch” (red).

Because the option “Repeat” was included only for a subset of units, it is omitted from the data shown. (a) LPFC units. (b) OFC units.

The mean population firing rates reveal a number of interesting patterns. Both OFC and LPFC responded to the presentation of choice targets, but did not appear to discriminate among the selected options in the mean firing rate response. LPFC, as a whole, responded vigorously and phasically to the onset of video presentations, but not to the blank screen condition (which involved no change in the visual stimulus). OFC, by contrast, appeared to show a greater overall increase in activity during video

presentation, and a marked phasic response to the end of the video (at which time the screen returned to a neutral gray color). Intriguingly, LPFC exhibited a phasic response to the delivery of juice following a blank presentation, but not a video presentation.

However, the mean population response failed to capture the considerable heterogeneity across the population in the presence, magnitude, and sign of modulation in response to each of the task events, which was better seen by investigating the mean responses of each unit individually (Figure 15). Within both LPFC and OFC, many units deviated from the pattern shown by the population averages, suggesting that there was considerably more information contained within the activity of these ensembles as a whole than is reflected in the population mean.

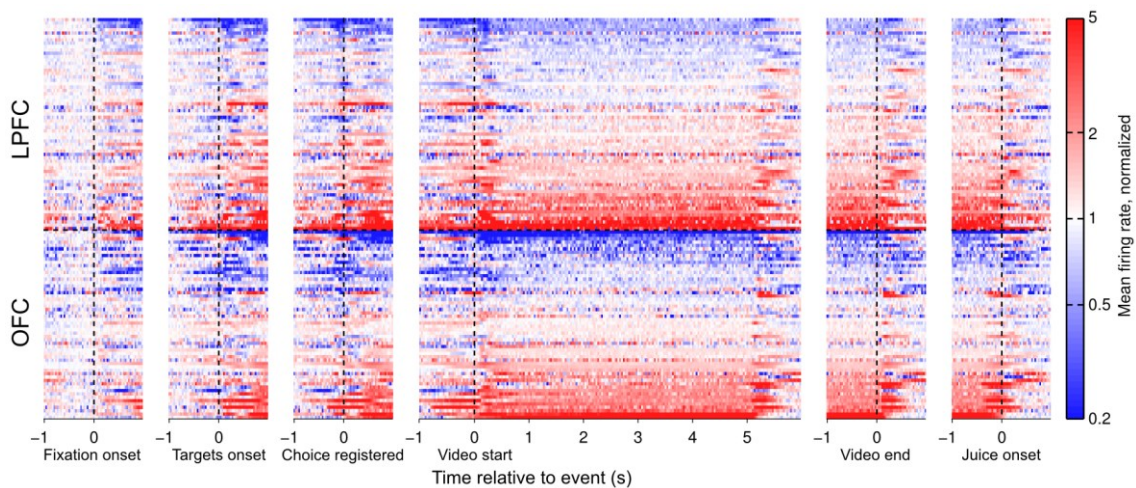


Figure 15: PFC units' responses to task events were heterogeneous. Each horizontal row represents one individual unit's peri-event activity, and the color of each cell indicates the normalized firing rate during a time bin. Units are ordered within each region by mean normalized firing rate during the video presentation epoch.

We were further interested in identifying whether behaviors viewed during video presentation had any impact on the firing rates of PFC neurons. We began with a simple approach of comparing the overall firing rate of PFC neurons during video presentation to the firing rate only when particular behaviors were occurring (Figure 16).

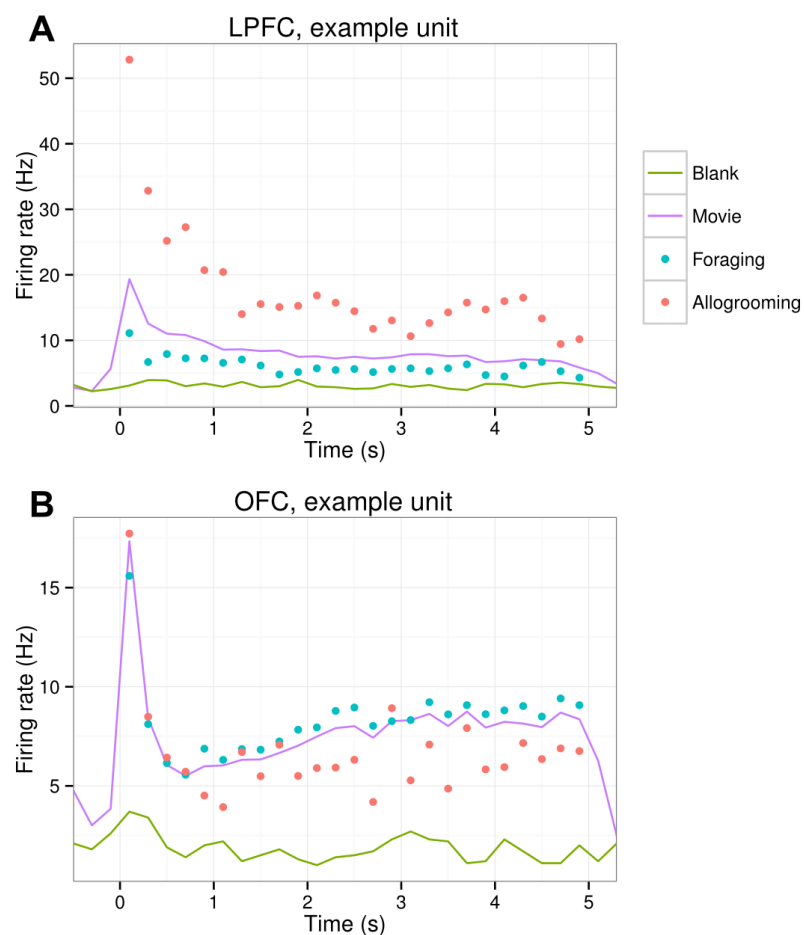


Figure 16: PFC units' firing rates were modulated by viewed behaviors during video presentation.

Connected lines represent PSTHs during presentation of a movie (purple) or blank screen (green), with 200 ms bins. Unconnected points represent mean firing rates for each bin across only those presentations in which foraging (cyan) or allogrooming

(red) was occurring. Because behavioral events did not necessarily span the entirety of a 5 second video presentation, each point may represent bins from a different subset of the total presentations. (a) LPFC example. This unit exhibited markedly greater firing rates when allogrooming was occurring, and lower firing rates when foraging was occurring, compared to the overall mean firing rates at each time during movie presentation. (b) OFC example. This unit exhibited lower firing rates when allogrooming was occurring, but did not appear to be notably modulated by foraging.

This approach yielded promising results, but has a number of limitations. First, it is incapable of addressing the problem of correlations between ethogram entries. Certain behaviors are more likely to co-occur than others, and because of the videotaped monkeys' movements through their environment as they behaved, camera movements were highly correlated with some behaviors. To avoid false inferences due to these correlations, we needed a model that would appropriately account for them. Furthermore, the simple PSTH-based method was successful for the two behaviors shown (allogrooming and foraging) in part because these are behaviors which tend to occur on relatively long timescales, meaning each time bin had a relatively large number of observations. For behaviors which typically occur more briefly (e.g., "Ingest food"), distributing their effects across the full 5 second presentation period made the signal difficult to extract from the noise even for a relatively large number of observations. Similarly, relatively small effects of viewed behaviors on firing rate are difficult to capture when the effect is distributed across 25 samples. These concerns led us to develop our firing rate model and fit it against the data. Figure 17 illustrates the result of the firing rate model for the unit shown in Figure 16a.




Figure 17: The firing rate model reveals strong modulation of an LPFC neuron's firing rate in response to multiple elements of a video scene, including behaviors.

The horizontal axis is logarithmically spaced.

The model fit confirmed the observation that this unit's firing rate was positively modulated by allogrooming and negatively modulated by foraging, and further identified numerous other factors that influenced firing rate, including the presence of monkeys in the scene ("Count ≥ 1 "), camera movements ("Camera zooming out," "Camera tracking," "Camera panning," and "Camera slewing"), other foraging-related behaviors (e.g. "Drink," "Ingest food," and "Chew"), and agonistic behaviors ("Joint aggression," "Submit"). Firing rate was also negatively modulated when the subject monkey was looking away from the screen ("Gaze offscreen"), which is consistent with the overall lower firing rate associated with less visual stimulation in the blank screen condition for this unit. However, this unit was also strongly negatively modulated by the parameter "Proportion gaze offscreen", suggesting that some features of the presented videos which made monkeys less likely to choose to direct their gaze at the screen (features which were not otherwise well described by the ethogram) also produced an overall lower firing rate for this cell.

As with the PSTHs, the model result revealed considerable heterogeneity in the population of PFC units (Figure 18).




Figure 18: Sparse representation of ethogram entries and gaze behavior metrics in individual units in prefrontal cortex. Units are ordered within brain region by the mean of the absolute value of the log gain over all regressors.

To estimate the false discovery rate associated with this result, we performed a permutation test, according to which for each unit, the relationship between the spike count in each bin and the ethogram regressors was randomly permuted, and the model then run on the resulting data. In terms of the model, this entailed substituting for the term $x_{k,n,t-\tau}$ some alternative $x_{k,n',t'-\tau}$ with n' and t' chosen by random permutation. This permutation test identified no non-zero regressors for any unit, implying a false discovery rate of exactly zero. While this cannot rule out that some regressors were

falsely identified as modulating firing rate due to their correlation with other regressors which were the true cause of the effect, the result of this permutation test speaks to the robustness of the elastic net regression technique.

To better understand the representation of ethogram and gaze variables in the population of PFC units, we counted the number of units exhibiting non-zero coefficients for each regressor (Figure 19) as well as the number of regressors with non-zero coefficients for each unit (Figure 20).




Figure 19: The proportion of units identified as sensitive to each regressor.

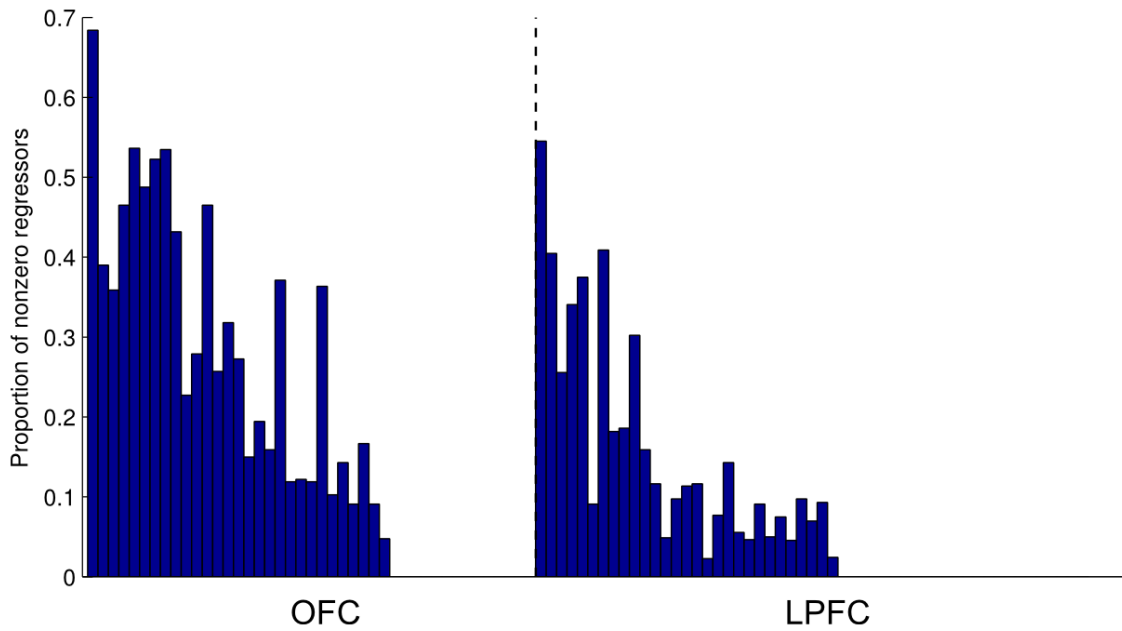


Figure 20: The proportion of regressors identified with nonzero coefficients for each unit.

The ordering of units (represented along the horizontal axis) within labeled area matches that of Figure 18.

These results reveal a few notable trends. Overall, the gaze-related regressors appear to be comparatively well represented in both OFC and LPFC. The overall proportion of units sensitive to each viewed behavior was fairly uniform across the set of behaviors that we investigated, with the exception of “Grasp food,” “Manipulate food,” “Ingest food,” and “Ingest grooming manipuland,” each of which appears to be less well-represented than most of the other behaviors. There are a few possible explanations for this. One possibility is raised by noting that these are all comparatively specific behaviors, and nested within “Foraging,” and their lesser degree of representation in the population may reflect the relative specificity of their definitions.

Another common feature among these behaviors is that they all tend to be relatively brief when they occur. With this in mind, the relatively low representation of these behaviors in the model results could reflect limitations in the assumptions that went into formulating the model. In particular, the model assumes that the units' firing rates will respond to each behavior with a boxcar-like profile, lasting exactly as long as the behavior is viewed, with a 100 ms lag. If the true temporal response profile differs from this assumption, the model is likely to under-detect true responses, and this effect would be more pronounced for shorter-duration behaviors. However, the proportions of responses to "Attack" and "Submit," behaviors which also tended to be relatively brief, were more in line with the rest of the ethogram.

OFC units in general appear to be more likely to be sensitive to ethogram regressors than units from the LPFC. The cumulative distribution functions for units having a certain proportion of nonzero regressors confirmed this impression (Figure 21).




Figure 21: Firing rates during video presentation of units in OFC tend to encode more regressors than those in LPFC.

Empirical cumulative distribution functions for the proportion of units encoding no more than a certain fraction of nonzero model (as in Figure 20) are shown for OFC and LPFC. The curve for OFC is shifted rightward relative to that for LPFC, indicating that at each level r , a greater proportion of OFC units were found to have at least r nonzero coefficients in the model.

4. Conclusion

4.1 Information foraging

Our study presented monkeys with an information foraging task, with a simulated “environment” consisting of naturalistic, information-rich videos of a diverse set of conspecific behaviors. During the task, monkeys made short-scale orienting decisions, “micro-decisions” about where to orient their gaze during video presentation, and longer-scale presentation decisions, “macro-decisions” about what kinds of videos to view. Monkeys’ presentation decisions revealed a strong preference for viewing videos of conspecifics over viewing a blank screen, and a strong preference for less-predictable videos over more-predictable videos. Contrary to our expectations, the nature of the viewed videos did not impact monkeys’ presentation decisions. There are several possible explanations that may account for this. First, monkeys simply may have had no preferences among the various behavioral categories in the ethogram. The striking effect of viewed behavioral categories on monkeys’ orienting decisions would seem to belie that explanation, but it is possible that these are fundamentally different decision processes, with distinct, unrelated preference functions. A related possibility is that monkeys did in fact make presentation decisions based on the viewed information, but our ethogram failed to capture the relevant dimensions of the stimuli. Again, this is at odds with the ethogram’s success in explaining at least some of the variability in orienting behavior and neural activity. However, monkeys’ very strong preference for

the Switch option over Continue or Repeat suggests that the lack of preferences for specific behaviors may have been a consequence of the nature of the “foraging environment” that we presented them with. Although our relatively large video database meant monkeys saw the same video sequence within the same behavioral session relatively infrequently, all of the subject monkeys performed this task regularly over many months, and consequently were eventually exposed to the entire video database numerous times. If monkeys were able to quickly recognize a given video sequence, and generally preferred the relative novelty (or at least unpredictability) offered by Switch, there may have been little utility in continuing or repeating a given video which the monkey already remembered well. Furthermore, even if monkeys had weak preferences for certain kinds of information, it is possible that the cognitive load of forming a decision on the basis of the viewed behavior incurred a cost greater than the value offered by the information, and so monkeys’ optimal strategy was to rely on simpler decision heuristics. These possibilities suggest a number of alterations to this information foraging paradigm. One option would be to employ a larger stimulus set to reduce monkeys’ familiarity with it, but given the challenges of coding such a large video database with such a complex ethogram, this may not be practical. Another possibility would be to adjust the menu options to reflect finer decision gradations. For example, offering “Similar” and “Different” options, which lead to presentations chosen

based on their ethogram's correlation with that of the previous presentation, would help to balance for the overall level of unpredictability offered by "Switch."

4.2 Semantic representations of natural scenes

The finding that the firing rates of single neurons in prefrontal cortex exhibit sparse representations of naturalistic, ethologically meaningful behavioral categories viewed in videos of conspecifics complements and extends upon previous findings from BOLD fMRI studies in humans. Huth and colleagues (2012) performed whole brain functional scans of human subjects while they viewed a varied set of video stimuli, and correlated BOLD responses to a large set of category labels describing the video scenes in each 1-second interval using a regularized GLM technique similar to the one described here. Perhaps unsurprisingly, their descriptive model performed best in areas typically associated with visual perception, occipital and temporal cortex (particularly the superior temporal sulcus and inferotemporal cortex), as well as medial parietal cortex. Intriguingly, they also identified the lateral prefrontal cortex as an area with BOLD responses relatively well-predicted by their model. However, the robustness of regularized GLM also allowed them to identify category labels which predicted BOLD responses throughout cortex, even in areas where their overall model explained less variability in the BOLD signal during movie viewing. They found sparse representation of these categories across cortex, with BOLD responses in spatially localized regions of the cortex exhibiting sensitivity to distinct subsets of their semantic category labels.

Our study adopted a similar approach in investigating the relationship between our ethogram categories and neural activity, but with some important methodological differences. Behaviorally, the human subjects in the study by Huth et al. maintained fixation during video presentation, and thus their neurophysiological responses may represent more “pure” perceptual representations, while our monkey subjects were free to orient their gaze in response to the stimulus, potentially engaging premotor and strategic information seeking decision-making systems as well. While Huth and colleagues applied an expansive set of 1,705 category labels to their video database, we elected for a more targeted set of 62 distinct ethogram labels selected for their known ethological importance to our study species. Huth and colleagues applied so-called “ridge regression” as their regularization technique, while we used the closely related but slightly more conservative elastic net technique, which more strongly penalizes nonzero terms in the model. The conservatism of our method likely accounts for our model’s high performance in excluding false positives under permutation testing, but may bias us toward underreporting the number of ethogram entries which had a nonzero influence on the firing rate of our cells. However, the most significant difference between the two studies is in the nature of the physiological response measured. The exact relationship between fMRI BOLD signals and the firing rate responses of individual cells or populations of cells remains poorly understood, but the best evidence suggests that there is not a simple correlation (Logothetis 2008, Goense &

Logothetis 2008, Maier et al. 2008), and thus our results represent an opportunity to compare these two signals under similar conditions and analysis regimes. First, we can note that our finding of sparse representations of viewed behavioral categories in the firing rates of neurons within OFC and LPFC mirrors the finding of sparse representations of viewed categories in BOLD response across cortex, but at a finer spatial and temporal scale. Second, we can consider the implications of the degree of information representation we observed in the neuronal firing rates of these areas with respect to the fMRI BOLD results.

4.3 Representations of social information

Our finding of a sparse representation of viewed behavioral categories in OFC and LPFC is highly suggestive of computational processing of this information in these areas. Sparse encoding appears to be a general computational principle in sensory systems (Waydo et al. 2006, Quiroga et al. 2008, Carlson et al. 2011, Palm 2013), and in this light our finding of sparse encoding of highly processed sensory information including viewed categories of behaviors is entirely consistent with work in other brain areas using more tightly controlled social stimuli (e.g., Perrett et al. 1984, Perrett et al. 1985, Friewald et al. 2009). Visual information arrives in OFC primarily via two pathways: one, a dense, specific projection from temporal cortical areas TE and TEO (the apex of the so-called “ventral visual stream” for extracting object and identity information) to ventrolateral prefrontal area 12l (Carmichael & Price 1995b), a

cytoarchitecturally-defined subregion which is tightly linked to both the orbital and ventrolateral prefrontal networks (Kadharbatcha et al. 2014); the other, a more diffuse, broadly-projecting connection from the basolateral amygdaloid nucleus (Carmichael & Price 1995a). Recent work suggests that even the amygdala sparsely extracts highly specific visual information of high behavioral importance such as conspecific eye contact (Mosher et al. 2014). Thus, our findings of sparse representations of natural behavioral categories in OFC and LPFC may reflect in part the nature of the visual inputs driving these areas.

An observation in our results that appears at first to be at odds with the neuroimaging result of Huth and colleagues (2012) is the finding that while BOLD signals in LPFC are more strongly associated with semantic categories than in OFC, we found neural firing in OFC to be more strongly associated with ethogram categories than in LPFC. Numerous factors may contribute to this difference, including species, stimulus, coding scheme, and task differences, and we cannot directly assess why this difference was observed. However, consideration of the relationship between these two regions and the nature of the signals being measured may indirectly resolve the conflict. First, Cai and Padoa-Schioppa (2014) proposed a feed-forward model for decision-making according to which good-based decisions are made in OFC, and the output of this is transformed into an action decision in or before LPFC (discussed in further detail below). Second, fMRI BOLD signals often reflect markers of synaptic activity more

strongly than spiking activity and thus in some cases reflect inputs to rather than spiking activity within a brain area (Logothetis 2008, Ekstrom 2010). It is possible, therefore, that the greater modulation of spiking activity in OFC with viewed categories may be reflected in a greater modulation of BOLD-associated physiological responses in LPFC, receiving these spiking inputs.

4.4 Neuroeconomics and Neuroethology

Historically, the sciences of economics and ecology have mutually benefitted from exchanging ideas (Røpke 2004), while often approaching similar questions with differing but complementary perspectives. Both disciplines have parallel behavioral (behavioral economics and behavioral ecology or ethology) and mechanistic (neuroeconomics and neuroethology) subdisciplines which can also benefit from a mutual exchange of ideas. Neuroeconomics is a highly successful approach to investigating brain function with respect to decision-making processes by providing a normative theory of behavior in well-defined contexts when rational actor theory applies, or when deviations from rationality are well-understood (Glimcher 2004). Neuroethology is a similarly successful approach to investigating brain function by providing a normative theory of natural behavior rooted in the notion of evolutionary adaptation (Ingle & Crews 1985, Carew 2000). Our study benefits from the productive dialogue between these two perspectives in understanding the roles of OFC and LPFC in adaptive behavior.

A proposal in the neuroeconomic perspective on the mechanisms of decision-making is the existence of “good-based” value representations, that is, that there exist in the brain populations of neurons which can represent the abstract value of any good (a stimulus or object in the world which may be targeted for consummatory behavior) in a common, unified fashion, and that decision-making occurs within this “goods space,” with systems for action selection being engaged only once a good has been selected (Padoa-Schioppa 2011). Such a mechanism offers numerous advantages for complex decision-making, including simplicity and scalability. Padoa-Schioppa and Assad (2006) found evidence that single neurons in OFC do in fact encode common currency value representations when monkeys make decisions between stimuli predicting juice rewards of different flavors and volumes. However, such a good-first decision process is not the only possible mechanism for making decisions between goods in the environment. From a neuroethological perspective, circumstances in which it is possible to make a decision purely between goods (which we might identify with the ecological concept of resources) may be relatively rare in the natural environment, and so both natural selection and phylogenetic conservatism may favor alternative decision processes in which both good values and action costs dynamically and interactively contribute to flexible decision-making (Cisek 2012). Nevertheless, Cai and Padoa-Schioppa (2014), by presenting monkeys with a decision-making task in which goods and actions were spatially and temporally decoupled, found that firing rates of OFC neurons signaled the

identity of a selected good, while those of LPFC tracked the subsequent action (a saccadic eye movement to a decision target) selected to consume the good, providing strong evidence for a goods-based decision process in OFC. However, it should also be considered that the task design of Cai and Padoa-Schioppa strongly favored the formation of an early good-based choice, as a delay between menu presentation and the appearance of the decision targets incurred a working memory requirement, and the consumption action costs were perfectly balanced in every decision, making good value the only relevant decision variable. In other contexts, good value and action costs must be tracked more dynamically, and a pure goods-based decision may not be possible (Cisek 2012).

Although our task was not designed to directly test the notion of goods-based decision making in OFC or LPFC, our results suggest an important point to consider when discussing the notion of value in decision-making tasks. In an economic framework, value is typically identified with revealed preference, that is, it may be measured by the subjects' decisions for or against a good (Samuelson 1948). From an ethological perspective, resources may have adaptive value even if an individual animal fails to display overt revealed preferences for them, either because of a maladaptive decision-making process, or because a particular context discourages the expression of a preference.

Our video stimuli presented monkeys with social information of high probable adaptive value, had our subjects been in a natural environmental context. However, in using presentation decisions to estimate value, we failed to find any evidence for revealed preferences among the various categories of behavior, possibly for contextual reasons as discussed above. Nevertheless, we found considerable evidence that this valuable (from an ethological perspective) information was represented in both OFC and LPFC. The representation of social information in OFC even in the absence of its strong influence of presentation decisions is consistent with previous findings that over twice as many units in OFC represent categories of social information compared to juice rewards even when decision-making is dominated by the volume of juice (Watson & Platt 2012). We propose that this apparent conflict can be resolved by considering that outside the laboratory context, resources of high ethological value, including social information, are much more likely to routinely influence decision-making.

Under this view, the model that OFC represents the value of “goods” for decision-making while LPFC represents the outcome of a good-to-action transformation (Cai & Padoa-Schioppa 2014) is consistent with our results, treating the various viewed behavioral categories as a set of potential features for decision-making analogous to “goods space.” In our case, although orienting decisions were sensitive to viewed behaviors, presentation decisions were not, and this reduction in the relevance of social information for producing actions would account for the reduced representation of

ethogram categories in LPFC. If this explanation is correct, we predict that an alternative information foraging task that elicits stronger behavioral responses to the viewed behaviors, either by altering the nature of the decision or the nature of the information environment, will lead to enhanced representation of behavioral categories in LPFC. In this view, the finding that cortical volume in LPFC increases in response to increased group size and dominance rank in rhesus macaques (Sallet et al. 2011) may be understood as a response to the demands of transforming the representation of social information in OFC and likely other brain areas into adaptive information foraging actions.

References

- Adolphs, R. (2010). What does the amygdala contribute to social cognition? *Annals of the New York Academy of Sciences*, 1191, 42-61.
- Anderson, B. D., & Moore, J. B. (1979). *Optimal Filtering*. Englewood Cliffs, NJ: Prentice-Hall.
- Andrews, P. W., Gangestad, S. W., & Matthews, D. (2002). Adaptationism - how to carry out an exaptationist program. *Behavioral and Brain Sciences*, 25(4), 489-553.
- Archie, E. A., Altmann, J., & Alberts, S. C. (2012). Social status predicts wound healing in wild baboons. *Proceedings of the National Academy of Sciences*, 109(23).
- Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annual Review of Neuroscience*, 28, 403-450.
- Bendesky, A., Tsunozaki, M., Rockman, M. V., Kruglyak, L., & Bargmann, C. I. (2011). Catecholamine receptor polymorphisms affect decision-making in *C. elegans*. *Nature*, 472, 313-318.
- Bigler, E. D. (1988). Frontal lobe damage and neuropsychological assessment. *Archives of Clinical Neuropsychology*, 3, 279-297.
- Bird, C., & Emery, N. (2008). Using video playback to investigate the social preferences of rooks, *Corvus frugilegus*. *Animal Behaviour*, 76(3), 679-687.
- Bisley, J. W., & Goldberg, M. E. (2010). Attention, intention, and priority in the parietal lobe. *Annual Review of Neuroscience*, 33, 1-21.
- Bovet, D., & Washburn, D. A. (2003). Rhesus macaques (*Macaca mulatta*) categorize unknown conspecifics according to their dominance relations. *Journal of Comparative Psychology*, 117, 400-405.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433-6.
- Brogan, W. L. (1985). *Modern Control Theory*. Prentice-Hall.
- Bromberg-Martin, E. S., & Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron*, 63, 119-126.

- Cai, X., & Padoa-Schioppa, C. (2014). Contributions of orbitofrontal and lateral prefrontal cortices to economic choice and the good-to-action transformation. *Neuron*, 81(5), 1140-1151.
- Carew, T. J. (2000). *Behavioral Neurobiology*. Sinauer Associates, Inc.
- Carlson, E. T., Rasquinha, R. J., Zhang, K., & Connor, C. E. (2011). A sparse object coding scheme in area V4. *Current Biology*, 21(4), 288-293.
- Carmichael, S. T., & Price, J. L. (1995). Sensory and premotor connections of the orbital and medial prefrontal cortex of macaque monkeys. *The Journal of comparative neurology*, 363(4), 642-664.
- Carmichael, S. T., & Price, J. L. (1996). Connectional networks within the orbital and medial prefrontal cortex of macaque monkeys. *Journal of Comparative Neurology*, 371(2), 179-207.
- Carter, R. M., Macinnes, J. J., Huettel, S. A., & Adcock, R. A. (2009). Activation in the VTA and nucleus accumbens increases in anticipation of both gains and losses. *Frontiers in Behavioral Neuroscience*, 3, 21.
- Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, 9(2), 129-136.
- Cheney, D. L., Seyfarth, R. M., & Silk, J. B. (1995). The responses of female baboons (*Papio cynocephalus ursinus*) to anomalous social interactions: evidence for causal reasoning? *Journal of Comparative Psychology*, 109(2), 134-140.
- Cisek, P. (2012). Making decisions through a distributed consensus. *Current Opinion in Neurobiology*, 22(6), 927-936.
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 362, 933-942.
- Cuthill, I. C., Kacelnik, A., Krebs, J. R., Haccou, P., & Iwasa, Y. (1990). Starlings exploiting patches: the effect of recent experience on foraging decisions. *Animal Behaviour*, 40, 625-640.

- D'Eath, R. B. (1998). Can video images imitate real stimuli in animal behaviour experiments? *Biological Reviews of the Cambridge Philosophical Society*, 73(3), 267-292.
- Dall, S. R., Giraldeau, L.-A., Olsson, O., McNamara, J. M., & Stephens, D. W. (2005). Information and its use by animals in evolutionary ecology. *Trends in Ecology & Evolution*, 20(4), 187-93.
- Deaner, R. O., & Platt, M. L. (2003). Reflexive social attention in monkeys and humans. *Current Biology*, 13, 1609-1613.
- Deaner, R. O., Khera, A. V., & Platt, M. L. (2005). Monkeys pay per view: adaptive valuation of social images by rhesus macaques. *Current Biology*, 15(6), 543-8.
- Dubois, F., Drullion, D., & Witte, K. (2011). Social information use may lead to maladaptive decisions: a game theoretic model. *Behavioral Ecology*(November).
- Dunbar, R. I. (1998). The social brain hypothesis. *Evolutionary Anthropology*, 6(5), 178-190.
- Dunbar, R. I. (2010). The social role of touch in humans and primates: behavioural function and neurobiological mechanisms. *Neuroscience and Biobehavioral Reviews*, 34(2), 260-8.
- Ekstrom, A. (2010). How and when the fMRI BOLD signal relates to underlying neural activity: The danger in dissociation. *Brain Research Reviews*, 62(2), 233-244.
- Eslinger, P. J., & Damasio, A. R. (1985). Severe disturbance of higher cognition after bilateral frontal lobe ablation: patient EVR. *Neurology*, 35(12), 1731-41.
- Ferrari, P. F., Kohler, E., Fogassi, L., & Gallese, V. (2000). The ability to follow eye gaze and its emergence during development in macaque monkeys. *PNAS*, 97, 13997-14002.
- Frank, M. J., Doll, B. B., Oas-Terpstra, J., & Moreno, F. (2009). Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nature Neuroscience*, 12, 1062-1068.
- Freiwald, W. A., Tsao, D. Y., & Livingstone, M. S. (2009). A face feature space in the macaque temporal lobe. *Nature neuroscience*, 12(9), 1187-96.
- Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software*, 33(1), 1-22.

- Gigerenzer, G., & Selten, R. (2001). *Bounded Rationality: The Adaptive Toolbox*. MIT Press.
- Gingrich, B., Liu, Y., Cascio, C., Wang, Z., & Insel, T. R. (2000). Dopamine D2 receptors in the nucleus accumbens are important for social attachment in female prairie voles (*Microtus ochrogaster*). *Behavioral Neuroscience*, *114*, 173-183.
- Glimcher, P. W. (2004). *Decisions, Uncertainty, and the Brain: The Science of Neuroeconomics*. MIT Press.
- Glimcher, P. W., Ciaramitaro, V. M., Platt, M. L., Bayer, H. M., Brown, M. a., & Handel, A. (2001). Application of neurosonography to experimental physiology. *Journal of Neuroscience Methods*, *108*, 131-144.
- Glimcher, P. W., Dorris, M. C., & Bayer, H. M. (2005). Physiological utility theory and the neuroeconomics of choice. *Games and Economic Behavior*, *52*, 213-256.
- Goense, J. B., & Logothetis, N. K. (2008). Neurophysiology of the BOLD fMRI Signal in Awake Monkeys. *Current Biology*, *18*, 631-640.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535-74.
- Goldberg, M. E., Bisley, J. W., Powell, K. D., & Gottlieb, J. (2006). Saccades, salience and attention: the role of the lateral intraparietal area in visual behavior.
- Gottlieb, J. (2012). Attention, learning, and the value of information. *Neuron*, *76*(2), 281-95.
- Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme. *Proceedings of the Royal Society of London. Series B*, *205*, 581-598.
- Hasson, U., Landesman, O., Knappmeyer, B., Vallines, I., Rubin, N., & Heeger, D. J. (2008). Neurocinematics: The Neuroscience of Film. *Projections*, *2*(1), 1-26.
- Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., & Malach, R. (2004). Intersubject synchronization of cortical activity during natural vision. *Science*, *303*, 1634-40.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York, NY: Springer.
- Hayden, B. Y., Pearson, J. M., & Platt, M. L. (2011). Neuronal basis of sequential foraging decisions in a patchy environment. *Nature Neuroscience*, *14*(7), 933-9.

- Higham, J. P., Brent, L. J., Dubuc, C., Accamando, A. K., Engelhardt, A., Gerald, M. S., . . . Stevens, M. (2010). Color signal information content and the eye of the beholder: a case study in the rhesus macaque. *Behavioral Ecology*, 21(4), 739-746.
- Hladký, V., & Havlíček, J. (2013). Was Tinbergen an Aristotelian? Comparison of Tinbergen's Four Whys and Aristotle's Four Causes. *Human Ethology Bulletin*, 4(4), 3-11.
- Huth, A. G., Nishimoto, S., Vu, A. T., & Gallant, J. L. (2012). A continuous semantic space describes the representation of thousands of object and action categories across the human brain. *Neuron*, 76(6), 1210-24.
- Ingle, D., & Crews, D. (1985). Vertebrate neuroethology: definitions and paradigms. *Annual review of neuroscience*, 8, 457-94.
- Izquierdo, A., Suda, R. K., & Murray, E. a. (2005). Comparison of the effects of bilateral orbital prefrontal cortex lesions and amygdala lesions on emotional responses in rhesus monkeys. *The Journal of Neuroscience*, 25(37), 8534-42.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The Fusiform Face Area: A module in human extrastriate cortex specialized for the perception of faces. *Journal of Neuroscience*, 17, 4302-4311.
- Klein, J. T., Deaner, R. O., & Platt, M. L. (2008). Neural correlates of social target value in macaque parietal cortex. *Current Biology*, 18, 419-424.
- Klin, A., Jones, W., Schultz, R., Volkmar, F., & Cohen, D. (2002). Visual fixation patterns during viewing of naturalistic social situations as predictors of social competence in individuals with autism. *Archives of General Psychiatry*, 59(9), 809-16.
- Kolling, N., Behrens, T. E., Mars, R. B., & Rushworth, M. F. (2012). Neural Mechanisms of Foraging. *Science*, 336, 95-98.
- Krebs, J. R., Ryan, J. C., & Charnov, E. L. (1974). Hunting by expectation or optimal foraging? A study of patch use by chickadees. *Animal Behaviour*, 22, 953-64.
- Ku, S. P., Tolia, A. S., Logothetis, N. K., & Goense, J. (2011). fMRI of the face-processing network in the ventral temporal lobe of awake and anesthetized macaques. *Neuron*, 70, 352-362.
- Landreth, A., & Bickle, J. (2008). Neuroeconomics, Neurophysiology and the Common Currency Hypothesis. *Economics and Philosophy*, 24(03), 419.

- Lee, D. (2006). Neural basis of quasi-rational decision making. *Current Opinion in Neurobiology*, 16, 191-198.
- Lloyd, E. A. (2009). *The Case of the Female Orgasm: Bias in the Science of Evolution*. Harvard University Press.
- Lockhart, R., Taylor, J., Tibshirani, R. J., & Tibshirani, R. (2014). A significance test for the lasso. *Annals of Statistics*, 42(2), 413-468.
- Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, 453(June), 869-878.
- Lorincz, E. N., Baker, C. I., & Perrett, D. I. (1999). Visual cues for attention following in rhesus monkeys. *Cahiers De Psychologie Cognitive*, 18, 973-1003.
- Mahajan, N., Martinez, M. A., Gutierrez, N. L., Diesendruck, G., Banaji, M. R., & Santos, L. R. (2011). The evolution of intergroup bias: perceptions and attitudes in rhesus macaques. *Journal of Personality and Social Psychology*, 100, 387-405.
- Maier, A., Wilke, M., Aura, C., Zhu, C., Ye, F. Q., & Leopold, D. a. (2008). Divergence of fMRI and neural signals in V1 during perceptual suppression in the awake monkey. *Nature neuroscience*, 11(10), 1193-1200.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (W. H. Free ed.).
- Marshall, J. A., Bogacz, R., Dornhaus, A., Planqué, R., Kovacs, T., & Franks, N. R. (2009). On optimal decision-making in brains and social insect colonies. *Journal of the Royal Society, Interface*, 6(40), 1065-74.
- McFarland, R., Roebuck, H., Yan, Y., Majolo, B., Li, W., & Guo, K. (2013). Social interactions through the eyes of macaques and humans. *PLoS ONE*, 8(2), 56437.
- Moore, T., Armstrong, K. M., & Fallah, M. (2003). Visuomotor origins of covert spatial attention. *Neuron*, 40, 671-683.
- Mosher, C. P., Zimmerman, P. E., & Gothard, K. M. (2011). Videos of conspecifics elicit interactive looking patterns and facial expressions in monkeys. *Behavioral Neuroscience*, 125(4), 639-52.

- Mosher, C. P., Zimmerman, P. E., & Gothard, K. M. (2014). Neurons in the monkey amygdala detect eye contact during naturalistic social interactions. *Current Biology*, *24*, 2459-2464.
- Murray, E. A., O'Doherty, J. P., & Schoenbaum, G. (2007). What we know and do not know about the functions of the orbitofrontal cortex after 20 years of cross-species studies. *Journal of Neuroscience*, *27*(31), 8166-9.
- Nishimoto, S., Vu, A. T., Naselaris, T., Benjamini, Y., Yu, B., & Gallant, J. L. (2011). Reconstructing visual experiences from brain activity evoked by natural movies. *Current Biology*, *21*, 1641-1646.
- Ongür, D., & Price, J. L. (2000). The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cerebral Cortex*, *10*(3), 206-19.
- Oram, M. W., & Perrett, D. I. (1994). Responses of anterior superior temporal polysensory (STPa) neurons to "biological motion" stimuli. *Journal of Cognitive Neuroscience*, *6*, 99-116.
- Padoa-Schioppa, C. (2007). Orbitofrontal cortex and the computation of economic value. *Annals of the New York Academy of Sciences*, *1121*, pp. 232-253.
- Padoa-Schioppa, C. (2011). Neurobiology of economic choice: a good-based model. *Annual review of neuroscience*, *34*, 333-359.
- Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature*, *441*(7090), 223-6.
- Palm, G. (2013). Neural associative memories and sparse coding. *Neural Networks*, *37*, 165-171.
- Parr, L. A., Winslow, J. T., Hopkins, W. D., & de Waal, F. B. (2000). Recognizing facial cues: individual discrimination by chimpanzees (*Pan troglodytes*) and rhesus monkeys (*Macaca mulatta*). *Journal of Comparative Psychology*, *114*, 47-60.
- Partan, S. R. (2002). Single and multichannel signal composition: facial expressions and vocalizations of rhesus macaques (*Macaca mulatta*). *Behaviour*, *139*(8), 993-1027.
- Perrett, D. I., Smith, P. A., Potter, D. D., Mistlin, A. J., Head, A. S., Milner, A. D., & Jeeves, M. A. (1984). Neurons responsive to faces in the temporal cortex: studies

- of functional organization, sensitivity to identity and relation to perception. *Human Neurobiology*, 3, 197-208.
- Perrett, D. I., Smith, P. A., 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. Series B*, 223, 293-317.
- Pirolli, P. L. (2007). *Information Foraging Theory: Adaptive Interaction with Information*. Oxford University Press.
- Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, 400, 233-238.
- Quiroga, R. Q., Kreiman, G., Koch, C., & Fried, I. (2008). Sparse but not 'Grandmother-cell' coding in the medial temporal lobe. *Trends in Cognitive Sciences*, 12(February), 87-91.
- R Core Team. (2014). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Rademacher, L., Krach, S., Kohls, G., Irmak, A., Gründer, G., & Spreckelmeyer, K. N. (2010). Dissociation of neural networks for anticipation and consumption of monetary and social rewards. *NeuroImage*, 49, 3276-3285.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59-108.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20, 873-922.
- Rieucau, G., & Giraldeau, L.-A. (2009). Video playback and social foraging: simulated companions produce the group size effect in nutmeg mannikins. *Animal Behaviour*, 78(4), 961-966.
- Røpke, I. (2005). Trends in the development of ecological economics from the late 1980s to the early 2000s. *Ecological Economics*, 55, 262-290.
- Rolls, E. T. (2000). The orbitofrontal cortex and reward. *Cerebral Cortex*, 10(3), 284-94.
- Sackett, G. P. (1966). Monkeys reared in isolation with pictures as visual input: evidence for an innate releasing mechanism. *Science*, 154, 1468-1473.

- Saleem, K. S., Miller, B., & Price, J. L. (2014). Subdivisions and connectional networks of the lateral prefrontal cortex in the macaque monkey. *The Journal of Comparative Neurology*, 522, 1641-90.
- Sallet, J., Mars, R. B., Noonan, M. P., Andersson, J. L., O'Reilly, J. X., Jbabdi, S., . . . Rushworth, M. F. (2011). Social network size affects neural circuits in macaques. *Science*, 334, 697-700.
- Samuelson, P. (1948). Consumption theory in terms of revealed preference. *Economica*, 15(60), 243-253.
- Schoenbaum, G., Roesch, M. R., Stalnaker, T. a., & Takahashi, Y. K. (2009). A new perspective on the role of the orbitofrontal cortex in adaptive behaviour. *Nature Reviews Neuroscience*, 10(12), 885-92.
- Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., & Marshall, J. A. (2012). Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064), 108-11.
- Shepherd, S. V., Deaner, R. O., & Platt, M. L. (2006). Social status gates social attention in monkeys. *Current Biology*, 16(4), 119.
- Shepherd, S. V., Steckenfinger, S. A., Hasson, U., & Ghazanfar, A. A. (2010). Human-monkey gaze correlations reveal convergent and divergent patterns of movie viewing. *Current Biology*, 20(7), 649-56.
- Smith, D. V., Hayden, B. Y., Truong, T.-K., Song, A. W., Platt, M. L., & Huettel, S. A. (2010). Distinct value signals in anterior and posterior ventromedial prefrontal cortex. *Journal of Neuroscience*, 30, 2490-2495.
- Stephens, D. W., & Krebs, J. R. (1986). *Foraging Theory* (Vol. 121). Princeton, NJ: Princeton University Press.
- Stephens, D. W., Brown, J. S., & Ydenberg, R. C. (2007). *Foraging: Behavior and Ecology*. University of Chicago Press.
- Sterelny, K. (2007). *Dawkins Vs. Gould: Survival of the Fittest*. Icon Books.
- Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2005). Choosing the greater of two goods: neural currencies for valuation and decision making. *Nature Reviews Neuroscience*, 6, 363-375.

- Tedore, C., & Johnsen, S. (2013). Pheromones exert top-down effects on visual recognition in the jumping spider *Lyssomanes viridis*. *Journal of Experimental Biology*, 216, 1744-1756.
- Thorpe, S. J., Rolls, E. T., & Maddison, S. (1983). The orbitofrontal cortex: neuronal activity in the behaving monkey. *Experimental Brain Research*, 49, 93-115.
- Tinbergen, N. (1963). On aims and methods of ethology. *Zeitschrift für Tierpsychologie*, 20, 410-433.
- Tremblay, L., & Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. *Nature*, 398(6729), 704-8.
- Trezza, V., Damsteegt, R., Achterberg, E. J., & Vanderschuren, L. J. (2011). Nucleus accumbens mu-opioid receptors mediate social reward. *Journal of Neuroscience*, 31, 6362-6370.
- Tricomi, E., Rangel, A., Camerer, C. F., & O'Doherty, J. P. (2010). Neural evidence for inequality-averse social preferences. *Nature*, 463, 1089-1091.
- 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), 989-95.
- Valone, T. J., & Templeton, J. J. (2002). Public information for the assessment of quality: a widespread social phenomenon. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 357(1427), 1549-57.
- Van Hooff, J. A. (1967). The facial displays of the catarrhine monkeys and apes.
- Wald, A., & Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, 19, 326-339.
- Walton, M. E., Behrens, T. E., Buckley, M. J., Rudebeck, P. H., & Rushworth, M. F. (2010). Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron*, 65(6), 927-39.
- Watson, K. K., & Platt, M. L. (2012). Social signals in primate orbitofrontal cortex. *Current Biology*, 1-6.
- Watson, K. K., Ghodasra, J. H., & Platt, M. L. (2009). Serotonin transporter genotype modulates social reward and punishment in rhesus macaques. *PLoS ONE*, 4(1), 4156.

- Waydo, S., Kraskov, A., Quian Quiroga, R., Fried, I., & Koch, C. (2006). Sparse representation in the human medial temporal lobe. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 26(40), 10232-10234.
- Whalen, C., & Schreibman, L. (2003). Joint attention training for children with autism using behavior modification procedures. *Journal of Child Psychology and Psychiatry*, 44(3), 456-68.
- Wilkinson, R. D., Steiper, M. E., Soligo, C., Martin, R. D., Yang, Z., & Tavaré, S. (2011). Dating primate divergences through an integrated analysis of palaeontological and molecular data. *Systematic Biology*, 60(1), 16-31.
- Zald, D. H., & Andreotti, C. (2010). Neuropsychological assessment of the orbital and ventromedial prefrontal cortex. *Neuropsychologia*, 48(12), 3377-3391.
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society*, 67(2), 301-320.

Biography

Geoffrey Adams was born in Alexandria, Virginia, in 1984. He studied Biological Sciences and Physics at Virginia Tech, receiving Bachelor of Science degrees in both fields in 2007. His undergraduate thesis, *Mechanical and osmotic water potentials of the eggs of the veiled chameleon, Chameleo calyptratus*, advised by Professor Robin Andrews, was awarded the Outstanding Undergraduate Research Award by the Virginia Tech Department of Biological Sciences. Following his undergraduate studies, he spent one year in a postbaccalaureate training program under the supervision of Dr. David Leopold, and subsequently matriculated into the Neurobiology graduate program at Duke University. His studies at Duke were financially supported by a James B. Duke Fellowship and a research grant to Professor Platt from the National Institute of Mental Health. He has authored several scientific publications (see below), presented his work at the Society for Neuroscience Annual Meeting and the Computational and Systems Neuroscience conference (COSYNE), and continues his scientific career as a postdoctoral fellow at Emory University and the Yerkes Primate Research Center, under the joint supervision of Professors Larry Young and Robert C. Liu.

Journal articles:

- DeWind, N. K., Adams, G. K., Platt, M. L., & Brannon, E. M. (under review) Modeling the approximate number system to quantify the contribution of visual stimulus features. *Cognition*
- Chang, S. W. C., Brent, L. J. N., Adams, G. K., Klein, J. T., Pearson, J. M., Watson, K. K., & Platt, M. L. (2013) Neuroethology of primate social behavior. *PNAS*
- Adams, G. K., Watson, K. K., Pearson, J. M., & Platt, M. L. (2012) Neuroethology of decision-making. *Current Opinion in Neurobiology*

- Platt, M. L., & Adams, G. K. (2012) Response to “Reading too much into baboon skills?” *Science*
- Platt, M. L., & Adams, G. K. (2012) Monkey see, monkey read. *Science*
- Maier, A., Adams, G. K., Aura, C., & Leopold, D. A. (2010) Distinct superficial and deep laminar domains of activity in the visual cortex during rest and stimulation. *Frontiers in Systems Neuroscience*
- Adams, G. K., Andrews, R. M., & Noble, L. M. (2010) Eggs under pressure: Components of water potential of chameleon eggs during incubation. *Physiological and Biochemical Zoology*

Book chapter:

- Adams, G. K., Watson, K. K., Pearson, J. M., & Platt, M. L. (2012) Neuroethology of decision-making. *Evolution and the Mechanisms of Decision-making*