CODING OF SEQUENTIAL BEHAVIORS BY ANTERIOR CINGULATE CORTEX ENSEMBLES

by

Liya Ma

B.Sc., Carleton University, 2007

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES

(Neuroscience)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

June 2014

Abstract

The anterior cingulate cortex (ACC) has been implicated in a myriad of different functions. Converging evidence suggests that the ACC continuously monitors and evaluates actions and their consequences. Such functions are essential in representing action sequences which are the building blocks of all complex behaviors. This dissertation seeks to delineate *how* ACC neuronal ensembles represent different types of information with special emphasis on action sequences.

Chapter 2 shows that the ACC ensembles represents different action sequences via unique activity patterns that change if the order of the actions are altered or if the locations of the actions is changed. Interestingly such shifts are achieved when overall levels of activity remain fixed.

Chapter 3 reveals a very different arrangement in which progression through a sequence of actions towards a goal is associated with a change in the overall level of neural activity without a significant change in the patterns of activity. Specifically, ACC ensembles display a smooth progressive change in overall activity over three lever press actions that culminate in a reward. In contrast, the dorsal striatal (DS) ensembles recorded simultaneously from the same animals display fluctuations in activity level that are tightly linked to each action. Together these two chapters show that the ACC may use two different firing rate-related codes to convey categorical versus continuous forms of information.

Chapter 4 provides a further examination of the mechanisms which allows the ACC ensembles to encode multiple types of categorical information. While the DS neurons encode both the sequence and the location of the levers in a somewhat synchronized fashion, ACC

neurons encoded both of these types of information but kept them functionally segregated. As a result, even though ACC single neurons were no better than the DS in sequence decoding, sequence decoding by ACC ensembles was far superior to DS ensembles.

The last chapter attempts to produce a unified theory of ACC function based on its coding properties. I will argue that the ACC monitors many aspects of experience while evaluating the current state with reference to a goal. Its multiple coding schemes efficiently serve both monitoring and evaluating functions.

Preface

The studies reported in this dissertation were designed together by my supervisor Dr. Jeremy Seamans and me, with important contributions from my colleague Dr. James Hyman and my co-supervisor Dr. Anthony Phillips. I conducted all the experiments and performed all data analysis under their close guidance.

A version of Chapter 3 has been published in: Ma L, Hyman JM, Phillips AG, Seamans JK (2014) Tracking Progress Toward a Goal in Corticostriatal Ensembles. J Neurosci 34:2244-2253. I designed the study together with Dr. Jeremy Seamans and conducted all the experiments. I performed the data analysis and created the figures under the close guidance of Drs. Jeremy Seamans and James Hyman. I wrote the manuscript together with Drs. Jeremy Seamans, James Hyman and Anthony Phillips.

A version of Chapter 4 has been accepted and will be published on June 29th, 2014: Ma L, Hyman JM, Lindsay AJ, Phillips AG, Seamans JK (2014) Differences in the Emergent Coding Properties of Cortical and Striatal Ensembles. Nat Neurosci, DOI: 10.1038/nn.3753. I designed the study together with Dr. Jeremy Seamans and conducted all the experiments. I performed the data analysis and created the figures under the close guidance of Drs. Jeremy Seamans and James Hyman, and with the help of Adrian Lindsay. I wrote the manuscript together with Drs. Jeremy Seamans, James Hyman and Anthony Phillips.

The experiments using animals were carried out in accordance with the Canadian Council of Animal Care, and with the approval of the Animal Care Committee at the University of British Columbia. The animal use protocol number was A10-0055.

Table of contents

List of figures

List of abbreviations

ACC	Anterior cingulate cortex
ANOVA	Analysis of variance
AP	Anterior-posterior
AUC	Area under curve
BOLD	Blood-oxygen-level dependent
dACC	Dorsal anterior cingulate cortex
D_{Euc}	Euclidean distance
dlPFC	Dorsolateral prefrontal cortex
DLS	Dorsolateral striatum
D_{Mah}	Mahalanobis distance
df	Degree of freedom
dmPFC	Dorsomedial prefrontal cortex
DMS	Dorsomedial striatum
DS	Dorsal striatum
EEG	Electroencephalography
FI	Fixed interval schedule
fMRI	Functional magnetic resonance imaging
FR	Fixed ratio schedule
GLM	Generalized linear model
iFR	Instantaeous firing rate
IL	Infralimbic cortex
ITC	Inferotemporal cortex
LP	Lever-press
MDA	Mahalanobis discriminant analysis
MDS	Multi-dimensional scaling
ML	Medial-lateral
MLR	Multi-linear regression
mPFC	Medial prefrontal cortex
MSUA	Multiple single-unit activity
NP	Nose-poke
OCD	Obsessive-compulsive disorder
PCA	Principal component analysis
PET	Positron emission tomography
PFC	Prefrontal cortex
PL	Prelimbic cortex
pre-SMA	Pre-supplementary motor area
RI	Random interval schedule
ROC	Receiver operating characteristic
SMA	Supplementary motor area
V1	Primary visual cortex
vmPFC	Ventromedial prefrontal cortex
VTA	Ventrotegmental area
Wheel-turn	WT

Acknowledgements

First I would like to acknowledge those who were involved in reviewing this thesis, including my committee members Drs. Stan Floresco and Yu-Tian Wang. I am deeply grateful for the tremendous amount of mentoring, support and help from my primary supervisor, Dr. Jeremy Seamans. I would also like to thank my co-supervisor, Dr. Anthony Phillips, for his invaluable mentoring and support in all aspects, and all the inspirations he provided me. I could not have hoped for supervisors that were more helpful, patient or tolerant than Jeremy and Tony, and could not have made even the smallest step towards this thesis without them. Their teaching has opened me up to a much greater world of intellectual possibilities than I could ever have imagined.

I would also like to give my enduring gratitude to Dr. James Hyman, who provided me with extensive and in-depth training on a great variety of techniques and skills, as well as numerous insightful comments and wise suggestions to my project. Without his tireless help and immeasurable patience, this project would not have been possible to complete.

I would like to thank Dr. Natalia Gorelova, from whom I not only acquired guidance and skills, but also reassurance and confidence to continue with my work. I was extremely lucky to receive her mentoring and support. I would also like to thank Drs. Carine Dias and Kelly Butts for their friendship and many thoughtful questions and comments on my project.

I would also like to acknowledge all the help I received from my fellow students, Barak Caracheo and Adrian Lindsay. I thank them for alleviating my struggle with the project and for making life in the lab so enjoyable.

I would like to give my deepest gratitude to my parents, for making me who I am and for encouraging me to travel the path of my choice. I thank my father, Qiang Ma, who has always expected and trusted me to fulfill new achievements in life, and my mother, Yanyun Li, who has always held unreserved faith and confidence in me.

Lastly, I would like to thank my best friend, soul mate and husband, Nan Zhou, for believing in my abilities and strength, and for accepting my many faults. Had he not been providing me with continuous encouragement and support, there would not have been a thesis.

Dedication

This dissertation is dedicated to my parents, Ms. Yanyun Li and Mr. Qiang Ma.

Chapter 1: General introduction

The mammalian brain relies on a vast network of interconnected neurons to obtain key information from the environment. Given the richness of information in the environment, differentiating the critical information (or 'signals') from the irrelevant (or 'noise') can be a daunting task. Moreover, given the constantly changing world and the multitude of biological needs to be satisfied, the 'signals' have to be frequently updated to guide the selection and alteration of the optimal behavioral strategy. Lastly, multiple sensory systems are likely to produce multiple 'signals' that need to be integrated, held in parallel and/or suppressed, depending on the current goal. Despite of the glorious feat they achieve collectively, the activity and functional role of each neuron is constrained by its synaptic input and output, its intrinsic properties (e.g. its morphology and receptor distributions) and its chemical environment.

As the evolutionarily newest structure, the prefrontal cortex (PFC) receives converging inputs from both the external world and the internal states of the organism, and sends widespread outputs to cortical and subcortical regions. What's more, the drastic increase in the volume of the PFC during evolutionary history has long been associated with the emergence of highly complex cognitive functions such as behavioral flexibility, abstract reasoning and creativity (Clarke et al., 2004, Greene et al., 2004, Aziz-Zadeh et al., 2009, Takeuchi et al., 2012). Neurons in the PFC, therefore, can be expected to have the most diverse functionality, as well as the widest spectrum of behavioral targets. Such complexity has severely limited our understanding of the functional properties of the PFC. This dissertation explores the properties of population coding in the rat medial PFC which is homologous to the anterior cingulate

cortex (ACC) of humans and other primates (Uylings and van Eden, 1990, Groenewegen and Uylings, 2000).

This introduction opens with an overview of the anatomy of the ACC across species, followed by a general discussion of its functions in the primates (including humans) as well as in the rodents. Then by focusing on one aspect of behavior, namely sequential behaviors, the role of the ACC is reviewed in greater details. The involvement of the ACC in sequence tasks is then compared with the role of the dorsal striatum (DS), which is a major output of the dmPFC as well as a key player in sequential behaviors. Following the section on functionality, a general discussion of the basic properties of information processing in the mammalian brain ensues, with an emphasis on the ACC. Finally, major mathematical and statistical methods used in the project are briefly reviewed as will their pros and cons. The chapter concludes with an outline of the overall objectives of this thesis.

1.1 Functional anatomy of regions of interest

1.1.1 The rat dmPFC as a model of primate ACC

As a popular species for studying aspects of cognitive neuroscience, the rat has a medial prefrontal cortex (mPFC), which encompasses the rostral anterior cingulate cortex (area 24), the prelimbic cortex (area 32) and infralimbic cortex (area 25), and together are homologous to the anterior cingulate cortex in primates (Uylings and van Eden, 1990, Groenewegen and Uylings, 2000). Additionally, although rats and mice do not have a dorsolateral prefrontal cortex (dlPFC), their mPFC may participate in a number of cognitive functions commonly subserved by the dlPFC in primates (Uylings et al., 2003, Seamans et al., 2008).

In both primates and rodents, area 24 and 32 are situated to receive a convergence of sensory signals from the external world as well as information about the organism's internal states, via reciprocal connections with the cortical areas involved in the visual, auditory, somatosensory, gustatory, nociceptive and visceral senses, as well as the motor and other prefrontal regions (Baleydier and Mauguiere, 1980, Van Eden et al., 1992, Van Eden and Buijs, 2000, Hoover and Vertes, 2007). Area 24 and dorsal area 32 are connected with the caudate nucleus (dorsal striatum in rodents), the midline thalamic nuclei, hippocampus, claustrum, basolateral amygdala, brain stem nuclei (e.g. the periaqueductal gray), and the spinal cord (Baleydier and Mauguiere, 1980, Sesack et al., 1989, Gabbott et al., 2005, Hoover and Vertes, 2007). Additionally, these areas are reciprocally connected with neuromodulatory regions such as the ventral tegmental area (VTA), dorsal raphe nuclei and locus coeruleus (Ongur et al., 2003, Hoover and Vertes, 2007).

In the rodent the ventral prelimbic and infralimbic cortex (IL, area 25) are anatomically quite distinct from the dorsal mPFC (Heidbreder and Groenewegen, 2003). As the focus of this thesis, the term dorsomedial prefrontal cortex (dmPFC) is used to refer to the anterior cingulate cortex (Cg1, or area 24b) and the dorsal part of the prelimbic cortex (PL, or area 32) in rat (Paxinos and Watson, 2005), which are homologous to areas 24b and dorsal area 32 in primates (Vogt and Paxinos, 2014). However, it should be noted that all subdivisions of the mPFC are inter-connected mostly via excitatory synapses both ipsi- and contra-laterally (Heidbreder and Groenewegen, 2003, Hoover and Vertes, 2007), which is in line with coordinated signal processing rather than a sharp division between parallel processes. The current project explores the functions of the rodent dmPFC, but the findings are discussed in relation to those on the ACC in humans and non-human primates.

1.1.2 ACC functions in human and non-human primate: Functional networks

The list of tasks that activates the ACC (areas 24/32/25) according to the blood-oxygen-level dependent (BOLD) response using functional magnetic resonance imaging (fMRI) is extensive. The ACC is a key hub in at least five heterogeneous functional networks as assessed with fMRI (Cole et al., 2009, Hutchison et al., 2012):

First, area 25 is part of the 'limbic' network (Yeo et al., 2011). This limbic network is related to visceromotor functions, because its constituent regions send projections to all hypothalamic areas that control feeding behavior (Ongur et al., 1998) as well as brainstem structures such as the periaqueductal gray, which regulate visceral and skeletomotor responses to stress or threatening stimuli (Price, 1999). Via its projections to the hypothalamus, area 25 may use the sensory (including viscerosensory) inputs received by the adjacent orbital area to regulate visceromotor output (Ongur and Price, 2000). Area 25 is also part of the ventromedial PFC (vmPFC), the damage of which results in impaired visceral responses to affective stimuli as well as the evaluation of anticipated outcome (Bechara et al., 1997). Such observations have led some to proposed that the visceral output from area 25 is involved in the evaluation of the expected outcome of a given decision (Nauta, 1971, Damasio, 1996).

Second, the pre- and subgenual area 32 and posterior cingulate cortex are part of the 'default' mode network (Raichle et al., 2001, Greicius et al., 2003). The default mode network shows significant decreases in BOLD activity during various cognitive and goal-directed tasks (Raichle et al., 2001, Greicius et al., 2003). The default network may be involved in internal mental functions such as thinking about the past and the future (Raichle et al., 2001, Buckner et al., 2008).

Third, the ‘executive’ or ‘frontoparietal’ control network involves the dorsal ACC (dACC, including supragenual area 24 and 32), as well as the lateral PFC and inferior parietal lobule (Dosenbach et al., 2007, Vincent et al., 2008). In particular, the dACC and the anterior insula/frontal operculum set apart from the other structures in this control network, as they are the only areas showing a generalized role in task control and are activated during the initiation and maintenance of responses, as well as during the feedback and adjustment phases (Dosenbach et al., 2007). Thus distinct from the ‘task-negative’ or ‘default’ network involving the pre- and subgenual area 32 and posterior cingulate cortex, the dACC appears to turn on during cognitive tasks (Dosenbach et al., 2007, Buckner et al., 2008).

Fourth, the ACC is known to be part of the the ‘ventral attention’ network, especially the part of dACC that lies immediately posterior to the sub-area that belongs to the frontoparietal control network (see the third point above; Yeo et al., 2011). The ‘ventral attention’ network is involved in re-attending to salient and unexpected events (Fox et al., 2006). Consistent with this view, human neuropsychological studies have long shown that the ACC plays an important role in attentional processes (Mesulam, 1981, Petersen and Posner, 2012). It is also noteworthy that an additional ‘dorsal attention’ network has been identified, distinct from the ACC involving instead the intraparietal sulcus and frontal eye field (Corbetta and Shulman, 2002, Fox et al., 2006). This network is implicated in sustained attention as required in working memory (Corbetta and Shulman, 2002). However, both the dorsal and the ventral attention systems are part of the task-positive network (Fox et al., 2005a).

Taken together, it is evident that the ACC is involved in networks that are functionally highly diverse. Furthermore, different subregions within the ACC respectively belong to the task-negative (i.e. the default network) and task-positive networks (i.e. the control and

attentional networks). Moreover, it is worth noting that activation of the default network does not necessitate the inhibition of the ventral attention network (Corbetta et al., 2008). In other words, albeit functionally distinct, subregions of the ACC do not necessarily oppose each other. But can these diverse functions be integrated into an unifying theory of ACC?

Both the control network and the ventral attention network appear to be activated by high attentional or cognitive demands, in contrast to the default network which becomes deactivated as demand increases (Paus et al., 1998, Gusnard and Raichle, 2001, Raichle et al., 2001). This 'demand-based recruitment' theory has led some to suggested that both the task-control and attentional functions of the ACC are part of overall 'cognitive control', defined as the ability of the cognitive system to adaptively reconfigure to meet the demand of different tasks (Botvinick et al., 2001). This flexible configuration may be achieved in several functions including response biasing and online maintenance of contextual information (Botvinick et al., 2001). Recently Carter and Krug (2012) refined the purpose of the system from meeting demands in discrete tasks to 'dynamically changing goals and environmental demands'.

But what specifically does the ACC contribute to 'cognitive control'? Based on single-unit recording data in non-human primate, it has been argued that the dACC computes the cost-benefit ratio, or net utility of available options in support of decision making (Walton et al., 2007a, Walton et al., 2007b, Croxson et al., 2009, Talmi et al., 2009, Hillman and Bilkey, 2010, 2012, Hosokawa et al., 2013). This is because ACC neurons have been found to encode a combination of a certain action and outcome (Matsumoto et al., 2003, Rushworth et al., 2004, Luk and Wallis, 2009, 2013), a combination of the amount of effort required and the outcome received (Kennerley et al., 2009, Hillman and Bilkey, 2010, 2012, Hosokawa et al., 2013), the delay until the outcome and the outcome itself (Hosokawa et al., 2013) or the

quantity and probability of both anticipated and experienced outcomes (Hayden et al., 2009, Kennerley et al., 2009). In this way, a critical function of the ACC may be as a cost-benefit 'outcome calculator'.

This focus on the role of the ACC on multi-modal outcome processing is consistent with the well-established role of this region in error or conflict detection. These theories are based largely on the results of EEG and neuro-imaging techniques in humans. Specifically, it has been suggested that the ACC detects errors (Niki and Watanabe, 1979, Holroyd et al., 2002, Swick and Turken, 2002, Ito et al., 2003, Holroyd et al., 2004, Yeung et al., 2004) or conflicts in response streams (Pardo et al., 1990, Botvinick et al., 1999, Carter et al., 1999, Botvinick et al., 2001, Gehring and Fencsik, 2001, van Veen et al., 2009, Yeung and Nieuwenhuis, 2009). Notably, the dACC is thought to contribute to conflict detection in a proactive manner: first, it exhibits increased activation during task epochs that are *expected to* contain conflict based on the subjects past history (Carter et al., 1999, MacDonald et al., 2000); second, although on error trials the dACC responded only when error occurred, on correct trials it became engaged *prior to* the motor response due to timely detection of conflict (van Veen and Carter, 2002, 2006). In rats, PL lesions impair response conflict resolution (de Wit et al., 2006) and disrupt context-guided disambiguation of conflicting information (Haddon and Killcross, 2006). Together these findings suggest that the monitoring of error and/or conflict by the dACC may occur both before and after the motor response.

Despite their differences, the error/conflict detector theory and the cost-benefit 'outcome calculator' theories point towards the more general notion that the dACC specifically prepares for and monitors actions and compares the outcomes of actions with predictions. Shenhav et al. (2013) have summarized this broad spectrum of functions as the 'expected value of control',

in order to describe the prediction of potential outcomes followed by the calculation of the net benefit of allocating control to a given task. In other words, the ACC evaluates the motivational and cognitive need for allocating control, as well as the potential consequence of following each option, and computes their difference. This theory implies that the ACC is not necessarily the structure exerting direct control over specific actions, but instead is focused on determining whether the current situation is important and difficult enough to require further input as well as and the degree of control required. The idea that the ACC is not required to exert 'cognitive control' is underscored by the fact that relatively large dACC lesions cause little impairment to the patients' attention and executive control (Cohen et al., 1999b, Swick and Turken, 2002, Stemmer et al., 2004, Critchley, 2005, Fellows and Farah, 2005). Typically patients receiving anterior cingulotomies for the treatment of OCD or chronic pain act and think normally, other than having fewer obsessions or less pain (Brotis et al., 2009). While the ACC may indeed exert an influence on cognitive control, it would appear to play a evaluative role in decision making.

The ACC may bias cognitive evaluations based on autonomic, motivational or emotional cues (Paus 2001). Accordingly dACC activity was positively correlated with sympathetic activation in tasks involving cognitive efforts (Critchley et al., 2000, Williams et al., 2000, Patterson et al., 2002, Critchley et al., 2003, Critchley, 2005) whereas subgenual ACC activity correlates with parasympathetic activation (Nagai et al., 2004). Critchley (2005) theorized that while the default network (including the subgenual ACC) sets the baseline 'feeling state', the dACC represents the intensity of personal salience, which may be emotional, homeostatic, cognitive or social in nature. These activities of the ACC in turn coordinate autonomic responses with the behavioral responses and their expected consequences (Critchley, 2005).

This theory is supported by a recent study examining the effect of mild electrical stimulation of the ACC in awake patients. When the dACC was stimulated, the patients reported both autonomic activitation as well as the feeling that efforts were required to survive a challenge to which they had to face (Parvizi et al., 2013). This finding may be considered as an extention to the role of the dACC in effort-based decision-making that not only evaluates the effort required, but also makes the evaluation available to consciousness, and perhaps even prepares and/or invigorates the organism via sympathetic activation. Interestingly, in the study by Parvizi et al. (2013), the perceived difficulty evoked by ACC stimulation was a 'feeling' that did not come with sensory information from the external world; instead, the context of the feeling was only described metaphorically by the patients. Thus, the interconnected subregions of the ACC (and other regions in their respective functional networks) may keep track of a baseline 'state of being'. When a significant environmental change occurs, it may then evoke a somatic or autonomic response which changes the 'state of being', in anticipation of the outcome of a behavioral response (Critchley, 2005).

According to this view subregions of ACC may produce a *continuous*, subjective and evaluative experience of life which is available to consciousness. Although most studies emphasize the discrete states of being 'on-task' or 'off-task', 'baseline' or 'activation', given the interwoven and dynamically changing goals in life (Carter and Krug 2012), the ACC likely functions with a higher degree of continuity.With these complex functions in mind, I wish to argue that rather than utilizing the discrete metabolic changes of large swaths of the ACC provided by fMRI, the field would be better served by focusing on the highly continuous nature of ensemble recordings in animal models as an ideal way of elucidating the coding properties of the ACC.

1.1.3 dmPFC functions in rodents

As noted above, the dmPFC in rodents includes the anterior cingulate cortex (ACC, area 24) and the prelimbic cortex (PL, area 32), which I shall discuss separately.

Monitoring actions and outcomes

In rats and mice, lesions to the ACC have a greater impact on cognitive functions than that in humans. In rats, ACC lesion impair response flexibility (Seamans et al., 1995, Aultman and Moghaddam, 2001), as well as difficult forms of stimulus-reward learning and autoshaping (Bussey et al., 1997a, Bussey et al., 1997b), disrupt temporal sequencing of response in both spatial and non-spatial delayed alternation (Delatour and Gisquet-Verrier, 2001), impair the retrieval of spatial sequence but not isolated locations (Kesner and Holbrook, 1987, Chiba et al., 1997) and cause fragmentation of operant action sequences (Ostlund et al., 2009). From these studies it seems that although ACC lesions does not impair the animal's ability to form simple associations between two behavioral elements such as a stimulus, response or reward, they have a profound impact when several of these elements need to be flexibly combined and organized coherently. Thus, the ACC may monitor actions and outcomes through time and may exert a critical influence on the manner in which these temporally discontinuous events are organized into coherent action sequences or plans.

Remote memory

The recent decade has also seen a number of studies suggesting that the rodent ACC plays an important role in a completely different and essentially unrelated process, namely, remote memory. No such function has been widely ascribed to the human ACC however. Remote memory is operationally defined for rats as memories formed over the previous 30 days. ACC

lesions affect these remote memories but not memories formed within the past few days. The ACC has been implicated in the consolidation of remote memory (Wiltgen et al., 2004, Restivo et al., 2009, Vetere et al., 2011) and storage/retrieval of remote memories (Takehara et al., 2003, Frankland et al., 2004, Frankland and Bontempi, 2005, Wheeler et al., 2013). Yet these data may be explained by other theories of ACC function described above. Given that most studies consistent with a role of ACC in remote memory processing involved the learning of aversive associations, these finding may be explained equally well in terms of the specific role for the ACC in aversive long-term memory. Alternatively, because mPFC lesions impair the guidance of action by the expected value of outcome, rather than the association between action and outcome per se (Corbit and Balleine, 2003, Tran-Tu-Yen et al., 2009), the same mechanism may explain the animals' lack of freezing or eye-blinking in the context or stimulus paired with aversive stimuli. In a study involving associations between a specific location and reward, Maviel et al. (2004) reported that ACC inactivation can impair the retrieval of remote spatial memory. Moreover, as a memory trace is typically weakened by time, the retrieval of a 'remote' memory is an effortful process, and this may in turn be related to the engagement of the ACC in cognitive effort (Endepols et al., 2010, Shenhav et al., 2013). Thus it is unclear whether the ACC is activated specifically by the retrieval of remote memory, or by the effort engendered by this process. Future studies are needed to disambiguate these two potential roles of the rodent ACC.

Acquisition of operant actions

While the mPFC may not store remote memories, it may nevertheless impact the way that remote memories are formed. For instance, PL lesions impede acquisition of operant actions (Balleine and Dickinson, 1998, Corbit and Balleine, 2003, Killcross and Coutureau, 2003) and

even when the PL-lesioned animals learn to acquire an operant action, their responses were supported by a stimulus-response association rather than an action-outcome association (Corbit and Balleine, 2003, Killcross and Coutureau, 2003). Interestingly, although PL lesions affected the guidance of action by outcome expectancy, they did not affect the storage or the retrieval of the association between the outcome, its paired stimulus, and the action (Corbit and Balleine, 2003, Ostlund and Balleine, 2005). A presentation of either the outcome or the stimulus was sufficient for the animals to choose the correct action; this is consistent with the hypothesis that the PL plays a role in forming the expectation of the outcome to guide action selection rather than in maintaining the memory of the action-outcome association (Corbit and Balleine, 2003). Thus, these data are also consistent with a major role of the dmPFC in monitoring actions and outcomes.

Rather than using lesions to gain insights into dmPFC function, numerous studies have investigated how dmPFC neurons respond during various behaviors. Neurons in the dmPFC in general are responsive to a broad range of behavioral events and cognitive processes. Some processes involve the experience 'here and now', such as representing the sensory contexts (Hyman et al., 2012), movement trajectories (Euston and McNaughton, 2006), the immediate availability of reward (Caracheo et al., 2013) and its consumption (Horst and Laubach, 2013). While other neural responses appear to tie the current state to future and past outcomes, such as neuronal responses to operant actions that predict only certain outcomes (Hyman et al., 2013), neuronal responses predicting previous or upcoming goals (Baeg et al., 2003), neuronal responses associated with unique body postures that can be used as mnemonic cues for forthcoming actions (Cowen and McNaughton, 2007) and neuron responses associated with post-error processing (Narayanan and Laubach, 2008). Therefore, similar to the main

conclusions of numerous studies of ACC function in humans and non-human primates, these data are consistent with the idea that the dmPFC of the rat monitors and evaluates actions and outcomes continuously through time.

1.1.4 Role of the ACC in sequential behavior

Assuming that the dmPFC monitors actions and outcomes continuously through time, it makes sense to study a sequence of actions that evolve through time and culminate in a reward. The ability to organize isolated actions into sequences is the foundation to all complex behaviors and while some behavioral sequences may be instinctive (such as the grooming and nesting behavior), others are acquired through learning. The flexible selection and execution of a sequence from a learned repertoire of behaviors requires the representation of all steps within the sequence, the consecutive activation of each step and the inhibition of all other steps until the right moment. Additionally, completion of each step of the sequence brings the organism closer to the final goal yet may not be rewarded right away, and may be followed by a shift in the immediate action or target.

A region such as the primate ACC or the rodent dmPFC, which monitors the ongoing actions and expected consequence with reference to the goal, and which is involved in both attention and executive control, would appear to be an ideal structure to be engaged by sequence tasks. Indeed, there is already a literature regarding the involvement of the ACC and other frontal regions in sequence tasks across species and action sequencing has also proven to be challenging for frontal patients (Penfield and Evans, 1935, Kolb and Milner, 1981, Milner, 1982). In humans, both the dlPFC and the ACC are preferentially activated by the acquisition and/or execution of motor sequences (Grafton et al., 1995). The absence of such activation may also contribute to the motor-sequencing deficit in Parkinsonian patients (Rowe et al.,

2002). Critically, dlPFC is involved in planning and preparation of action sequences that specify both movements and timing (Sakai et al., 2002), and is also activated when a response sequence is held in working memory (Pochon et al., 2001). The mid-dlPFC also displays sustained activity throughout the execution of new as well as over-learned (habitual) action-sequences, presumably reflecting active monitoring and maintenance processes, which is distinctive from other frontal regions (Sakai et al., 2002). Likewise in non-human primates, the local field potentials of the dlPFC and the dorsal striatum show task phase-dependent coordination even when performing over-learned sequences (Fujii and Graybiel, 2005). Lesion to the dlPFC also result in impaired self-order sequencing performance and increased perseveration over already-taken actions (Collins et al., 1998). Collectively, these data highlight the importance of the human ad non-human primate PFC in both the acquisition of novel sequences and well in the performance of over-learned sequences.

While there is little debate about the role of various frontal regions in the acquisition of novel sequences, the role of the lateral prefrontal cortices in the execution of extensively trained sequences is more equivocal. On one hand, positron-emission tomography (PET) and fMRI studies suggested a disengagement of the PFC and a transfer of sequence information to other cortical areas (Jenkins et al., 1994, Jueptner et al., 1997, Shadmehr and Holcomb, 1997, Petersen et al., 1998) but see (Sakai et al., 2002). On the other hand, electrophysiological studies routinely found sequence-related activities in the PFC of well-trained animals (Averbeck et al., 2002, 2003a, Averbeck et al., 2003b, Fujii and Graybiel, 2005, Mushiake et al., 2006). Together these findings indicate that even though the total amount of prefrontal activity during action-sequence performance may decrease over training, the representation of sequence may continue to evolve in the PFC.

Frontal neurons may represent various aspects of action sequences. In single unit studies, cells responsive to the serial order of the actions within a sequence were found in the primate ACC (Procyk et al., 2000, Shidara and Richmond, 2002), lateral PFC (Barone and Joseph, 1989, Ninokura et al., 2004, Ryou and Wilson, 2004, Mushiake et al., 2006, Averbeck and Lee, 2007, Berdyyeva and Olson, 2010), and supplementary motor area (SMA) and pre-SMA (Clower and Alexander, 1998, Berdyyeva and Olson, 2010). The prevalence of such responses was only moderately different across regions (Berdyyeva and Olson, 2010). Additionally, neurons in the primate dorsolateral or dorsomedial frontal cortex fire prior to the performance of specific sequences of actions (Nakamura et al., 1998, Procyk et al., 2000, Ninokura et al., 2004, Mushiake et al., 2006, Averbeck and Lee, 2007, Shima et al., 2007). These properties have led some to suggest that there may be dedicated frontal neurons whose function is to represent a specific behavioral sequence (Ninokura et al., 2004, Averbeck and Lee, 2007, Shima et al., 2007).

In a unique study, monkeys were required to complete a series of 4 trials, each of which contained identical events including the onset of fixation cue, wait cue, bar-release cue and the subject's release of the bar, in order to obtain reward (Shidara and Richmond 2002). A group of dorsal ACC units were found to respond to one or more of these events in every trial, and their amplitude of the response increased incrementally until either the last or the second last trial in the sequence (Shidara and Richmond 2002). This signal may encode reward-expectancy during action sequences (Shidara and Richmond 2002), or alternatively, may be used as a means to represent progress towards a goal in the ACC. ACC neurons also display some other unique responses during action sequencing tasks. Specifically instead of targeting the actions only, ACC cells also responded to the states of peripheral stimuli which provided

context for the actions (Barone and Joseph, 1989), to the 'subgoals' of each action (Mushiake et al., 2006), and to impending error responses (Averbeck and Lee, 2007). Collectively these data suggest that ACC neurons encode actions in light of their cognitive and emotional context, effort and outcome with reference to the goal. These characteristics make the ACC an ideal structure for supporting the selection and performance of action-sequences, where representations of individual actions need to be made and updated with reference to their evolving goal and sequential context. Additionally, actions in a sequence are often associated with distinct locations, whereby carrying spatial information in addition to their sequential identity. Therefore, it seems that sequence tasks provides a suitable platform to study the coding properties of the ACC with regard to different types of information, and how its coding schemes meet the need of monitoring and supporting complex behaviors.

1.1.5 Role of DS in sequential behaviors

Of course the ACC does not generate action by itself. The DS is a downstream structure that is also intensively studied for its role in action-sequencing tasks. I will briefly review basic characteristics of the DS and its role in action sequencing and attempt to contrast this with what is known about how the ACC encodes sequence-related information.

As the input structure of the basal ganglia, the striatum receives converging projections from the neocortex, and sends feedback via the pallidum and substantia nigra pars reticulata, and then the thalamus to the neocortex (Groenewegen et al., 1997). In both rodents and primates, different cortical regions and their corresponding striatal targets form part of different "loops": the vmPFC-ventral striatal "limbic" loop is implicated in motivational and reward processing, whereas the other corticostriatal loops involving the dmPFC and sensorimotor cortex and the dorsal striatum are suggested to substantiate cognitive and

sensorimotor processes (Alexander and Crutcher, 1990, Uylings et al., 2003, Pennartz et al., 2009).

The DS is not only pivotal in motor functions, but is also important in higher cognition (Pennartz et al., 2009). In line with its role in motor and procedural learning, and its pathology in Parkinson's disease (Antonini et al., 2001), Huntington's disease (Joel, 2001) and in Tourette's syndrome (Marsh et al., 2004), the DS has been most extensively studied with respect to its role in bottom-up regulation of movements (Mink, 1996, Balleine and O'Doherty, 2010). Recent years have seen numerous studies revealing an involvement of the DS in the executive functions such as working memory and reward prediction (Chang et al., 2002, Lauwereyns et al., 2002, Tanaka et al., 2006), most likely through its feedback projections to the prefrontal cortex and the orbitofrontal cortex via the mediodorsal nucleus of the thalamus (Alexander and Crutcher, 1990, Balleine and O'Doherty, 2010). The DS is believed to contribute to both the cognitive and the motor aspects of cortical processing in a goal-directed task, via 2 parallel pathways: 1. the feed-forward and feed-back dopaminergic and GABAergic projections spiraling between the VTA, nucleus accumbens, substantia nigra pars compacta, dorsomedial striatum (DMS) and dorsolateral striatum (DLS) may provide the prefrontal cortex with motivational information (Alexander and Crutcher, 1990, Balleine and O'Doherty, 2010); 2. The input from the dorsolateral striatum via the substantia nigra pars reticulata and the cerebellum converging in the ventral anterior and ventrolateral nuclei of thalamus may provide the primary and association motor cortices with motor information (Alexander and Crutcher, 1990, McFarland and Haber, 2000). Although the PL projects mainly to the dorsomedial striatum (DMS), the anterior cingulate cortex projects to the central

striatum (Sesack et al., 1989, Mailly et al., 2013), an area that borders upon both the DMS and the DLS, which will be investigated in this thesis.

Recently the DLS has been shown to be necessary for the acquisition of a sequence but not non-sequential actions (Yin, 2010). The DS also supports the initiation and accuracy of learned operant action sequences, independent of the course of improvement in the motor skill per se (Bailey and Mair, 2006). One distinction between an isolated action and one in a sequence, is that at each step in a given sequence, the choices must be evaluated and the correct action chosen. Thus a critical aspect of action-sequence acquisition or performance is the retention of representations of both actions and sequences and their activations at the correct time. Ideally, the representations of an action and the sequence in which it is currently embedded should become active in an integrated fashion. Although DS neurons encode actions strongly when they occur in a sequence but not in isolation (Aldridge and Berridge, 1998, Schmitzer-Torbert and Redish, 2004), it is unclear whether they also reflect a specific sequence-action combination, i.e. encoding actions only when they occur in sequence A but not sequence B. Interestingly, such an integrated fashion of encoding action within its context along with its cost and consequence is exactly what characterized the primate ACC or rodent dmPFC (see sections above). Therefore the circuitry involving both the ACC and the DS appears to be ideal for supporting the selection and execution of actions in sequences. Simultaneous recordings from the dmPFC and the DS during a sequence task may provide insight to their different roles in sequential behaviors, supported by different coding properties.

1.2 Information coding by neuronal ensembles

While single unit studies have provided valuable information about the neural basis of sequence coding in corticostriatal circuits as discussed above, a limitation of many past

studies is that their focus on one neuron at a time. The fact that no neuron exists in isolation, and instead is under constant influence from other neural elements, strongly impacts the information it processes. Therefore, it is of paramount importance not only to understand how dmPFC and DS neurons encode sequence information in isolation but how they do so as a collective. While there have been few attempts to characterize the ensemble coding of sequence representations in corticostrital circuits, ensembles have been studied mainly in a variety of other contexts. This section is an overview of the theoretical and methodological challenges in the study of information coding by neuronal ensembles.

1.2.1 Properties of neural codes

In the mammalian brain, it is highly unlikely that the output of a group of neurons sub-serving an identical function is projected collectively onto a single downstream cell to control behavior. Rather, each presynaptic neuron innervates more than one cell, and each postsynaptic neuron receives inputs from heterogeneous sources. Different representations are produced by combining the activity coming from different groups of neurons. Having multimodal neurons that can combine in many different ways provides a high degree of flexibility. Yet with heterogeneity also comes noise (which may potentially be signal in a different context). To overcome noise and ensure signal transmission, a critical number of responsive neurons are required at each step of processing. Indeed, the majority of single-unit recording studies make an attempt at estimating the proportion of neurons involved in the specific cognitive functions of interest. Two interesting questions then arise: 1. how exclusive is the actions of such a groups of cells related to the function(s) under study? 2. Does each individual neuron contain the entirety of the information, or only some small part of it? In other words, the first question asks whether different cognitive processes recruit overlapping

groups of cells (or ‘cell assemblies’), and the second asks how distributed the information is within the cell group.

Although often confused, these two aspects of coding can be quite independent of each other: a cell can hold multiple pieces of information by varying its firing rate, or respond only to a single target. This property may be called the ‘dedicated’ neuron hypothesis. Simultaneously, it may encode the entirety of the situation, or only some small part of it. This may be called the ‘completeness’ of the representation. Assuming that the brain never relies on a single neuron to encode a given target, ‘completeness’ is equivalent to ‘redundancy’ from the perspective of the target being encoded. For either ‘dedication’ or ‘completeness’, the potential measurements of all neurons in the brain constitute a continuum, and the theoretical concept of a ‘grandmother cell’ is located at the origin of the coordinates, where the extreme of the two continua meet (green circle, **Fig. 1.1**). The observation that has come the closest to the ‘grandmother cell’ is the hippocampal ‘concept’ cells, which responds to two or three closely related familiar movie characters and captures several modes of the same concept—be it a side view of the character, or their name either read or heard (Quiroga et al., 2005). In terms of dedication (red horizontal axis, **Fig. 1.1**), on one extreme, there is again the ‘grandmother cell’ which responds to a grandmother only and no one else (Gross, 2002, Bowers, 2009). On the other end are the theoretical omnipotent cells responding to every behavioral target encountered. In between the extremes, one may find neurons with varying degree of functional selectivity. In terms of the completeness axis (blue vertical axis, **Fig. 1.1**), on one end there is again the ‘grandmother cell’ that responds to every defining feature of grandmother: her appearance from all angles, her name, her voice, etc. (Barlow 1994, Gross 2002). On the other end of the spectrum, there are neurons that respond to each piece of the

representation and the complete 'picture' (or target) can only be decoded from an ensemble. Between the extremes, one can find neurons with varying degrees of overlap in their representations.

Cells in the primary visual cortex have receptive fields (RFs) that overlap somewhat (Hubel and Wiesel, 1962) but are limited by local inhibition (Angelucci et al., 2002) within 'columns' (Horton and Adams, 2005). Thus the redundancy across groups or modules is relatively low ('V1 cells' with solid black ellipse, **Fig. 1.2**). They are also finely tuned to specific orientation and speed of movement within a small portion of the visual field (Hubel and Wiesel, 1962), thus their level of dedication is high. Of course if the reference target is not isolated features but a series of large and complex stimuli each covering the receptive field of a given V1 neuron, then the cell is expected to respond to all of them (Quian Quiroga and Kreiman, 2010)—hence the low level of dedication ('V1 cells*' with dashed black ellipse, **Fig. 1.2**). Along the ventral stream of visual processing, the information encoded by single neurons becomes more and more complete. A good example is the face processing patches in the primate temporal lobe: neurons in the middle lateral and middle fundus patches respond to specific views of a face; those in the anterior lateral patch are tuned towards a pair of symmetrical angles of the face, whereas those in the anterior medial patch responded reliably to a face regardless of the view (Freiwald and Tsao, 2010).

Given the functional anatomy of the rodent dmPFC, it is not surprising that cells in this region have been found to respond to a wide variety of cognitive targets (Jung et al., 1998, Lapish et al., 2008, Durstewitz et al., 2010, Hyman et al., 2012, Hyman et al., 2013, Rigotti et al., 2013). As many have pointed out, single units within the dmPFC may be expected to encode a wide variety of tasks (Jung et al., 1998, Seamans et al., 2008). Given such versatility

of dmPFC function, neurons in this area cannot be easily characterized by complete sets of 'tuning profiles' that resemble cells in the sensory systems. Nevertheless, the responsiveness of some dmPFC neurons towards discrete events such as an operant response or reward-delivery, is frequently modulated by a number of major changes in the task, including sensory context (Hyman et al., 2012, Chapter 2), response rules (Durstewitz et al., 2010, Cowen et al., 2012), available cost-benefit options (Hillman and Bilkey, 2010) or different sequencing of the actions (Chapters 2-4). In primates, a given prefrontal neuron may be engaged in abstract categorization along one or multiple dimensions (Cromer et al., 2010). Thus while some PFC neurons may maintain their functional correlates despite major changes in the task rule or context, others 'remap' significantly. When considered together, PFC cells have limited dedication to any given target, but are likely to have some degree of overlap in their behavioral correlates ('PFC cells' with black ellipse). This is consistent with their role in information integration (Walton et al., 2007a, Cowen et al., 2012).

An additional variable in coding schemes is how 'wide-spread' the information is in a given brain region. When target information only exists within a few cells, the brain region is said to employ a 'sparse distributed code' or simply 'sparse code' (Bowers, 2009, Wohrer et al., 2013). When target information penetrates a large population of cells, it is said that the brain region encodes this information with a 'dense/fully distributed code', or simply 'distributed code' (Wohrer et al., 2013, Bowers, 2009). Again, this 'sparse-distributed' distinction is more of a continuum than a dichotomy. In this section I have so far avoided the term 'distributed code' but instead used 'redundancy' and 'spread' (for want of a better word) of information, because by saying a code is distributed among many cells, i.e. widely spread,

researchers do not always specify how much overlap or redundancy there is among the cells (detailed above). I will outline 3 combinations of different levels of redundancy and spread.

1. When information is wide-spread and redundancy is low, one may expect information encoding/decoding to improve greatly with the addition of each single unit, at least when the ensemble is small. Consider the process of constructing a jigsaw puzzle—which has no information overlap among its pieces—without knowing the complete picture. By adding one new piece, we learn more about the picture than just the features in that piece; instead at some point the whole pattern will occur to us before the puzzle is complete. Of course, when the picture is nearly complete, the addition of each puzzle piece has less benefit than at the beginning of the process. In the same way, the decoding performance-ensemble size relationship may be likely to resemble a power function with exponent <1.
2. When information is wide-spread with high redundancy, it is complete with a few neurons, beyond which point adding more cells is unlikely to boost information representation dramatically. Thus a logarithmic function which plateaus faster than does a power curve is more likely to illustrate this scenario. In this case, the advantage of having more target-related neurons lies more in a boosting of the signal-to-noise ratio.
3. When information is localized to a few neurons, the redundancy is also low. This is the case of a sparse code. The relationship between information decoding and ensemble size is likely to be positive initially as more cells within the small group are added, but quickly turn negative as non-target-related cells join and add noise to the system. In

other words, one may expect the success rate in encoding/decoding to be like a second-order polynomial function of ensemble size, i.e. the trajectory of a projectile.

Different benefits and costs are associated with sparse, dedicated coding versus distributed redundant coding. If a group of neurons are dedicated to the target and encode the complete information, one only needs to sample one of them to decode the information perfectly. Indeed this is the rationale for most single-unit studies where the response of a few single neurons is taken as evidence for the functional property of that brain region. If information is distributed across an ensemble with little redundancy (i.e. a jigsaw puzzle model), then it would be necessary to record from all the constituting units to complete the decoding. If the information is distributed redundantly among many neurons, whether it can be decoded from a single unit depends on the unit's functional fidelity, or 'dedication'. A trade-off exists between dedication and coding capacity: dedicated coding scheme provides a reliable 'labeled line' for information transmission, but limits the coding capacity of the region. For a region involved in behavioral spontaneity, flexibility and adaptability such as the prefrontal cortex, distributed codes with moderate redundancy are clearly more favorable than a dedicated, 'hard-wired' coding scheme.

Additional costs are associated with either sparse, dedicated or distributed, versatile coding schemes. The cost for a sparse code focuses on ensuring the influence of the few single units on the downstream neurons and effectors, whereas the cost for a distributed code centralizes on the downstream convergence of the key ensemble's influence. Either way, signal transmission and decoding involves a battle of signal against noise. Noise may come from failed (or omitted) responses, from synaptic inputs unrelated to the target, or from the

neurons' own responses towards alternative behavioral correlates when they are not dedicated to a single target.

Using terms from signal detection theory, sparse coding scheme could minimize the false alarm rate, but may fail to limit the number of misses given the unreliable nature of individual neurons (Sakai, 2001). Thus a critical number of single unit participants are required in this type of code. On the other hand, having redundantly distributed and non-dedicated coding scheme maximizes hit rate because of the large ensemble size, but the false alarm rate of the single units is also expected to be high, because more than one event can trigger a response. This problem can be solved by expanding an ensemble's size. In comparison to the sparse coding scheme, this type of distributed code is metabolically more expensive and less stable in the long term, but is also more flexible and versatile.

For an example of the different utilities of sparse, dedicated coding and distributed, versatile coding, consider the encoding of abstract concepts and its individual exemplars or constituent elements. It was demonstrated that categories are encoded by many different brain regions for different purposes, utilizing different coding schemes (Seger and Miller, 2010). Neurons in the inferotemporal cortex (ITC) do not generalize their responses to all members in a category (Freedman et al., 2003, Jiang et al., 2007); instead their role lies more in the analysis of features contributing to categorization (Seger and Miller, 2010). Neurons in the PFC, on the other hand, showed different responses to categories—both perceptual and abstract and arbitrary—that generalized to all exemplars in the category (Freedman et al., 2001, Wallis et al., 2001, Freedman et al., 2002, 2003, Wallis and Miller, 2003, Muhammad et al., 2006), and such responses developed faster than categorical response in the ITC (Freedman et al., 2003, Meyers et al., 2008). The flexibility and adaptability of the distributed

PFC neural code to arbitrary coexists with its lack of fidelity to any individual stimulus, which is a function better served by the sparser code in the ITC.

1.2.2 Forms of neural codes

Fundamentally there are two broad types of coding mechanisms in the brain: 1. a rate code, whereby neurons encode information by changing their firing frequency; and 2. a temporal code, whereby neurons firing in synchrony with each other resulting in a sharpened response (Ainsworth et al., 2012). The rate code has two flavors: information may be encoded in the presence/lack of deviation from baseline activity in a given cell, or it may be encoded and discriminated from other information via a difference in the intensity of the cell's response. The former is also referred to as a 'place code' (e.g. as in the cochlea) or a 'labelled-line' code (e.g. the gustatory system), where the identity of the active cell is all that critical. It is likely that most neurons in the brain can employ at least one type of rate code. A type of temporal code is often observed during the 'phase-locking' phenomenon, where single neurons tend to fire at specific phases—e.g. the peak rather than the trough—of local oscillations in a certain frequency band. Information may be encoded as neurons switch from out-of-phase to in-phase with the local oscillation, and such switch may or may not coincide with a change in their firing rates. Hence these two means of information encoding are independent of each other, even though they may be employed simultaneously. The following discussion will focus on the rate code.

Importantly, here the 'rate code' is defined in a broad sense, which extends beyond a change in the mean firing rate in a subgroup of statistically selected neurons. Firstly, given a constant mean firing rate, the temporal structure of a neuron's spike train may vary in a functionally relevant fashion. Characteristics such as 'burstiness' differ between brain regions

(Maimon and Assad, 2009, Shinomoto et al., 2009) and neuron types, and often change with the demand of information encoding, and the timing of such change may be critical (Compte et al., 2003, Ponce-Alvarez et al., 2010, Hamaguchi et al., 2011). Variations in the structure of the spike times gives rise to a 'temporal pattern code', which has been observed in many brain regions and across multiple species (Theunissen and Miller, 1995, Martinez-Conde et al., 2002, Krahe and Gabbiani, 2004, Alitto et al., 2005, Oswald et al., 2007, Eyherabide et al., 2009). Furthermore, when multiple cells or cell assemblies fire in a recurring sequence, they generate a 'spatiotemporal pattern code', where the timing of patterns indicates when stimulus features occur. Subcortical regions, especially those involved in movement selection and coordination, often employ such spatiotemporal pattern code, as in the case of the 'central-pattern generators' in the brain stem (Grillner et al., 2005b).

Secondly, the assumption that changes in the mean firing rate is the only possible manifestation of a 'rate code' often leads to over-simplified and inaccurate classification of neurons. As pointed out by Wohrer et al. (2012), the responses towards a stimulus in a neuronal population often fall on a continuum with a single peak, and that the difference between the significant and non-significant responders is not itself significant (Nieuwenhuis et al 2011). Highly distributed population code has been researched extensively in the motor cortex (Georgopoulos et al., 1982, Georgopoulos et al., 1986). Additionally, it is common for neurons in the association and prefrontal cortices to have multiple behavioral targets which may or may not interact with each other (Hung et al., 2005, Kennerley et al., 2006, Meyers et al., 2008, Cowen et al., 2012), thus classifying them based on their mean firing rates around a single target fails to capture their complete functional profiles. Across a large ensemble of such multi-tasking neurons, however, each behavioral target is likely associated with what

may be termed the 'firing-rate pattern code' (or FR-pattern code). It should be clarified that this FR pattern code is different from a classical 'labeled-line code', where subsets of neurons were defined in binary terms to each given behavioral target: responders vs. non-responders. An FR pattern includes both cells that altered their activities significantly from its baseline or mean firing rate, as well as those that didn't.

This FR pattern code has unique features, such as the fact that it does not assume any neuron to be task-irrelevant even though it may not respond significantly, nor does it assume that neurons have quasi-permanent behavioral correlates. Take the encoding of sensory context for an example: although the ACC does not contain counterparts of the hippocampal place cells which reliably tracks location, cells in this region of the frontal cortex provide a much more distinct representation for each unique sensory context on the ensemble level than did the hippocampus (Hyman et al. 2012). What's more, when animals make several alternating visits to contexts A and B, even though there are only ~11% ACC neurons that consistently preferred one context over another, context decoding continues to improve until half of all neurons are added to the ensemble (Hyman et al. 2012). Therefore, despite few neurons dedicated to a given behavioral correlate, strong and reliable encoding of sensory contexts can still *emerge* from the FR patterns of ACC ensembles, which included neurons with non-significant or unreliable responsiveness to context. Such properties may be well-suited to a functionally versatile region like the ACC for the following reasons:

1. A postsynaptic neuron's receptive field is the weighted temporal and spatial summation of the response profiles of all its presynaptic neurons, including both the excitatory and inhibitory portions of their receptive fields. This subtle distinction becomes particularly important when the presynaptic neurons have complex

responding profiles: the encoding of the spatial world by the medial entorhinal grid cells and hippocampal place cells are excellent examples (Derdikman and Moser, 2010, Schmidt-Hieber and Hausser, 2013). Given the roles of the PFC in new learning and cognitive flexibility, it is expected that the PFC neurons not only have complex responding profiles (Rigotti et al., 2013), but the 'on' and 'off' portions of their receptive fields are often abstract and situation-dependent (Durstewitz et al., 2010, Lapish et al., 2008). The temporal distribution of sequence information among ACC neurons is analyzed in detail in Chapter 4.

2. Unique and critical characteristics of the ACC neurons include the fact that they receive multi-modal inputs, and appear to be active even without a behavioral task (Raichle et al., 2001, Margulies et al., 2007, Hutchison et al., 2010) and therefore respond by 'tuning up/down' rather than 'turning on/off'. It is therefore hard to judge whether a statistically insignificant change from their background levels of activity reflect a summation of signals that cancelled out each other or merely noise. Rather than setting an arbitrary cut-off line, a more rigorous and informative method is to consider all the neurons in at least some analyses and check whether a prominent change is visible or not.
3. Although a mean firing-rate code is effective in overcoming uncorrelated noise that is averaged out across neurons, the problem of correlated noise is more serious. Because the noise is correlated across cells, i.e. cells fire above or below their respective mean FR together in a given trial, averaging the cells response cannot cancel out the noise. In an ensemble consisting of 2 cells, if they have positive noise correlation, and if they both respond to an event with an increase in FR (i.e. having a positive signal

correlation), then the boundary between signal and noise would be blurred. Noise correlations can have an effect on ensemble decoding which increases with ensemble size (Averbeck and Lee, 2006, Averbeck et al., 2006), although larger ensembles may also contain more information embedded in the firing rate variance (Shamir and Sompolinsky, 2004). Therefore, given the large number of neurons involved in complex cognitive tasks, along with their functional heterogeneity, correlated signal and noise could potentially be an issue, in that information may either be enhanced or compromised by these forms of neuron-neuron interaction.

4. In addition to noise correlation, higher-order properties of neuronal ensembles such as the co-variance and cross-correlation among neurons may also contain information (Kara et al., 2005, Baeg et al., 2007). Again, such properties may not matter much in the case of a highly sparse rate code, e.g. 'concept neurons' in the hippocampus, but are important in the case of a highly distributed code involving versatile neurons such as those in the ACC.

In short, for a region such as the dmPFC, it is important to look beyond the mean firing rates to consider higher-order statistics such as the covariance and the irregularity of the mean and the variance (Vogels et al., 1989, Britten et al., 1992, Shamir and Sompolinsky, 2004). On the seemingly simple basis of a change in the firing rates of single neurons, the 'rate code' may manifest as different subtypes of coding schemes, whereby information may be encoded and decoded from the mean rate, the signal correlation, the noise correlation (Scaglione et al., 2011), and the relationship between mean and variance (Compte et al., 2003).

1.3 Mathematical and statistical techniques

1.3.1 Multiple single unit activity (MSUA) space and Mahalanobis distance (D_{Mah})

A simple way of measuring unit's response during a behavioral epoch would be to compare its mean firing rate before versus during the epoch. If there is a significant difference by a one-sample t-test, one may conclude that the behavioral epoch has triggered a response in the neuron. This method, however, fails to consider the details in trial-to-trial variation in the single-unit response.

Furthermore, when multiple neurons are recorded simultaneously, their synchronicity and functional connectivity may change with behavioral epochs, the pattern of which can provide insights into how the ensemble (or even the brain region) encodes information essential to the task. If the firing rates across multiple single units were simply averaged, then such information would be lost. For an over-simplified instance, if a third of the neurons become activated by a stimulus and increased their firing rates, another third become suppressed, while the remainder stayed at their previous level of firing, the ensemble would not show any response to the stimulus. Dividing the ensemble up into 3 groups of neurons and then taking average helps characterize the effect of this stimulus better, but if another stimulus causes a similar 3-way divided responding pattern in the ensemble, with each third now including different cells, then the grouping will have to be altered. In other words, if one chooses to use 3 dimensions—cells that fired more, those that fired less and those that stayed the same—along which to describe the effect of one stimulus, another 3 dimensions will have to be used for a different stimulus. It is always possible to report the percentage of responding cells and the average responsiveness, but then the identities of the cells are lost and potentially

interesting phenomena—such as multi-tasking and frequent change in functional "receptive field", which are often observed in prefrontal neurons—go unnoticed.

Therefore it seemed only prudent to allow each neuron to have its own firing rate axis, by constructing an N-dimensional (N=ensemble size) multiple single-unit activity (MSUA) space. In the MSUA space, the firing rates of each neuron are represented along a unique axis or dimension, and there are as many dimensions as there are neurons. Therefore, in each time bin a population activity vector is obtained from N firing rate numbers contributed by the N neurons. This population vector can then be represented as a single dot in the N-dimensional space. In this high-dimensional analysis, three aspects are considered: 1. the firing rates from all N neuron in each time-bin during that epoch, determining the location of a single dot in the N-dimensional space; 2. the variance in each neuron's activities from time-bin to time-bin, determining how spread apart the cluster is along that neuron's axis; 3. the covariance among the neurons. When the firing rates of a pair of neurons' have high covariance relative to their respective variance, the cluster becomes elongated along a line in the plane defined by these two neurons' axes—that is, the two cells have correlated activities, or encodes information in a redundant fashion.

How to quantify the spread of the dots in the MSUA space: Mahalanobis distance (D_{Mah}) is a multivariate metric for the separation between a vector and a group of vectors or between two groups of vectors. It is more appropriate than the Euclidean distance (D_{Euc}) when the groups of population vectors have differential variances and correlations between the vectors (that is, the dimensions or axes) (Krzanowski, 2000). In other words, D_{Euc} requires normal data, whereas D_{Mah} only requires the data to be multinormal. In geometrical terms, to use D_{Euc} the clusters have to be spherically symmetrical, whereas the D_{Mah} takes into account the

difference in density gradients (i.e. variance) along different dimensions in the data, as long as approximate normality can be assumed within each dimension (**Fig. 1.3**). For the separation between two groups of multinormal population vectors, D_{Mah} is calculated as the square-root of the difference between the two population vectors divided by their covariance matrix:

$$d(\vec{x}, \vec{y}) = \sqrt{(\vec{x} - \vec{y})^T S^{-1} (\vec{x} - \vec{y})} \quad eq.1$$

Two types of Mahalanobis distances may be calculated for different purposes: the distance between 2 clusters, and the distance from a single dot to a cluster. The cluster-cluster distance is computed for example when one is interested in how distinct the ensemble activity patterns are during 2 different behavioral epochs. In this case, $\vec{x} - \vec{y}$ is the distance between the 2 centres of masses and $S^{-1}$is the inverse matrix of the pooled covariance matrix of the 2 clusters. The dot-cluster distance is computed to verify the assignment of the dot—representing ensemble activities during a single time-bin—to the correct cluster, in which case $\vec{x} - \vec{y}$ is the distance between the dot and the cluster's centre of mass, and $S^{-1}$is the inverse matrix of cluster's covariance matrix. If a dot is assigned correctly, then its distance should be smaller to the assigned cluster than the other (or any other) cluster in question.

Figure 1.2 demonstrates a scenario where D_{Mah} is suitable: because the black dot lies beyond the boundary of the blue cluster and within that of the red cluster, it belongs to the red cluster. However, it also has longer Euclidean distance to the centre of mass of the red cluster than to that of the blue cluster, or $D_{Euc\text{-}Blue} < D_{Euc\text{-}Red}$. Therefore, using the D_{Euc}, one would have reached an incorrect conclusion. Because both clusters (especially the red) have different density gradients along different dimensions, e.g. x and y, the assumption for D_{Euc} has been violated. When using D_{Mah}, the D_{Euc} from the black dot to the centre of the red cluster is

divided by a greater covariance than that of the black dot to the centre of the blue cluster, giving: $D_{Mah\text{-}Blue} > D_{Mah\text{-}Red}$, supporting the correct conclusion that the black dot indeed belongs to the red cluster.

The conclusion of 2 "distinct" clusters, however, requires more than calculating the distance *per se*—some control distance has to be computed using bootstrapping or random permutation in order to test the statistical significance of the actual distance, or the "distinctiveness" or "separateness" of the clusters. In this thesis, this was done by randomly selecting half of all dots from each cluster to construct "pseudo-clusters" and calculating the distance between them. This process was repeated 100 times to produce a distribution of control distances. The actual distance has to be outside this distribution to support the distinctiveness of the 2 clusters. The bootstrapping method was also used to control for the different sizes of ensembles, such that each distance calculated was always between clusters in spaces with Nmin being the minimal ensemble size across all sessions considered. This enables valid grouping of distances from different sessions for statistical comparison.

1.3.2 Principal component analysis (PCA)

PCA is a technique that finds and describes the "trend lines" along which the data contain the largest variance, thereby describing the key characteristics of the data. Given its power in dimension-reduction with minimal loss of information, PCA is commonly used in data (e.g. image) compression and pattern recognition/detection. In neuroscience, it is frequently used in the analysis of DNA microarray data, of BOLD signals, and has become increasingly popular for the analysis of large neuronal ensembles.

PCA finds mutually independent (or 'orthogonal') patterns, i.e. eigenvectors or 'principal components', in the data matrix by decomposing the covariance matrix. It should be pointed out that PCA used in this project (Chapter 3) will *not* be a decomposition of the neuron-by-neuron covariance matrix but rather on the time-bin by time-bin covariance matrix. In other words, this type of 'temporal PCA' characterizes whether different time bins contain similar or opposite firing patterns across the entire ensemble. Once decomposed, the principal components (PCs) show the dominant pattern in the ensemble throughout all time bins as well as the percentage of total variance accounted for by each PC. One additional advantage of doing PCA in this fashion is the 'loadings' obtained for each neuron on each PC give an estimate of how strongly each neuron displays that temporal pattern. Indeed, neurons with high loadings on a given PC display a firing pattern that strongly resembles the PC itself. It is worth mentioning that loadings can be positive and negative, and the contribution of each neuron to the dominance of the PC depends on the absolute value of the loading. This is because, for a given PC, the sum of all neurons' loadings squared equals the eigenvalue of that PC. And the percentage of variance accounted for by a given PC is simply its eigenvalue divided by the sum of all eigenvalues.

PCA performed in this fashion is useful in two ways: it finds the dominant pattern of activity in the ensemble, and decomposes each neuron's activity into different behaviourally relevant 'layers'. These results would not be very fruitful if the dominant pattern turned out to be strictly defined by isolated behavioral elements, such as a lever press or a reward, in which case one can define categories of neurons a priori and use a number of methods (e.g. ANOVA and linear regression) to group them. It becomes interesting, however, when the top PCs

display a pattern systematically weaving together multiple behaviors, thereby revealing the presence of abstract cognitive processes (Chapter 3).

PCA also has its drawbacks. For one thing, it is hard to compare results across separate PCAs as they don't produce the same PCs. For instance while it would be useful to examine how the dominance of a given activity pattern varies from structure to structure, or group to group this is not possible. This issue may be partially solved by projecting a second set of data onto the existing PCs and find their loadings and variance accounted for by these PCs, which can then be compared to the results of the original PCA as I did in chapter 3.

A related challenge is that a single behavior or cognitive process appears in multiple PCs. Because PCA separates independent sources of variance, this may happen if a single cognitive process is served by a highly heterogeneous group of neurons. The versatility of mPFC neurons may further complicate this issue, since it is not uncommon for a subgroup of the cells serving process A to be also involved in process B, in which case these processes may co-occur in a single PC, and each process may become 'fragmented' in more than one PC. Therefore, in order to characterize the functional categories of neurons, I chose to use techniques other than PCA (such as linear regression).

1.4 Overview and objectives

In a nutshell, sequence tasks engage several aspects of dmPFC functions, including the monitoring of actions and their consequences with reference to the goal, along with the sequential and sensory contexts. Given that the dmPFC contains neurons responsive to a variety of behavioral correlates, and is most likely to have a distributed coding scheme, it follows that an integrated 'sequence' signal can emerge at the ensemble level. If this is the

case, then the dmPFC would be likely to play an essential role in a task requiring flexible switching among sequences, which reflects many real life situations. Additionally, given the role of the ACC in error and conflict detection, the dmPFC may be expected to generate a unique response to errors when they arise.

So far most studies of information processing in the primate ACC or rodent dmPFC using sequence tasks have focused on the behavioral correlates of single units. On one hand, as mentioned above, the function of the ACC appears to be task-general, with its activities encompassing several epochs in a given task (Dosenbach et al., 2007). On the other hand, ACC single units are often specifically responsive to one or more individual behavioral epochs (Barone and Joseph, 1989, Procyk et al., 2000, Shidara and Richmond, 2002). Therefore, an abstract signal such as 'sequence', which is defined as the collective of all its actions and their sequential organization, is likely to be the emergent property of an ensemble of neurons responsive to individual epochs. Of course, it is also possible that a learned sequence can be sparsely encoded by a small and dedicated group of neurons, although such a coding scheme is inconsistent with the versatility and flexibility displayed by the dmPFC neurons (see section 1.2). This thesis will examine the coding scheme of the signals of both concrete behaviors such as presses on a certain physical lever as well as more abstract concept such as sequence.

As discussed briefly above, the DS is another critical region involved in sequence acquisition and performance. It terms of its anatomy, the DS is a very different structure from the ACC and therefore provides a good contrast with regards to coding properties. Given the differences in anatomy, major cell types and the cytoarchitecture of the two regions, I hypothesize important differences in the way the same information is encoded in the two

regions simultaneously. Even if single neurons from both regions respond to similar aspects of the tasks, I also expect that their ensemble level representations will be quite different.

This thesis examines the functional properties of dmPFC ensembles in the specific setting of multi-action sequence tasks, which require organizing multiple operant actions into one of many possible orders to obtain a final reward. The specific objectives are as follows:

1. To determine how dmPFC ensembles encode instrumental actions, given the sensory context, the sequence and the stage of training in which they occur. As demonstrated by Hyman et al. (2012), the dmPFC generates highly distinct representation for different sensory contexts. Given the flexible nature of dmPFC representations, I hypothesize that the same actions may be represented differently in two different sequences or sensory environments. Also given that the dmPFC may play some role in remote memory or habit learning, the same sequence may be represented differently throughout training and with over-training.
2. To determine how sequential actions are represented in relation to their outcome(s). As reviewed above, the dmPFC is involved not only in sequence tasks but also in evaluating outcomes and in detecting errors. Given this, the second series of experiments examine how representations of sequential actions change depending on the actions position in the sequence relative to the outcome and whether or not the action was correct or an error. I hypothesize that given the extremely robust responses of ACC neurons to errors, the same action will be encoded much more strongly if it is performed in the incorrect sequence than if it is performed correctly. On the other hand, given the strong representation of outcomes, I also predict that actions closer to the goal should be associated with greater firing. Given the lack of

the same type of error and outcome processing in the DS, these types of differential modulation should be absent.

3. To compare and contrast the coding schemes in the most general sense in the dmPFC versus the DS. This will again involve comparing responses of ACC and DS neurons on sequencing tasks, and the goal here is to attempt to extract some general rules about how the two regions represent information. Given the fact that the dmPFC is a vastly interconnected excitatory network while the DS is also not so recurrently connected and is an inhibitory network, I predict the coding scheme should be distributed in the dmPFC and sparse in the DS.

Experimental data will be presented and discussed in the following three chapters, each focusing on testing a hypothesis listed above. Each chapter will include a short introduction providing rationale for the experimental design, as well as a discussion section over the findings specifically from the experiment. The last chapter will be a general discussion, in an attempt to place the findings into the literature and speculate on the insights they provide to the functions and coding properties of the dmPFC.

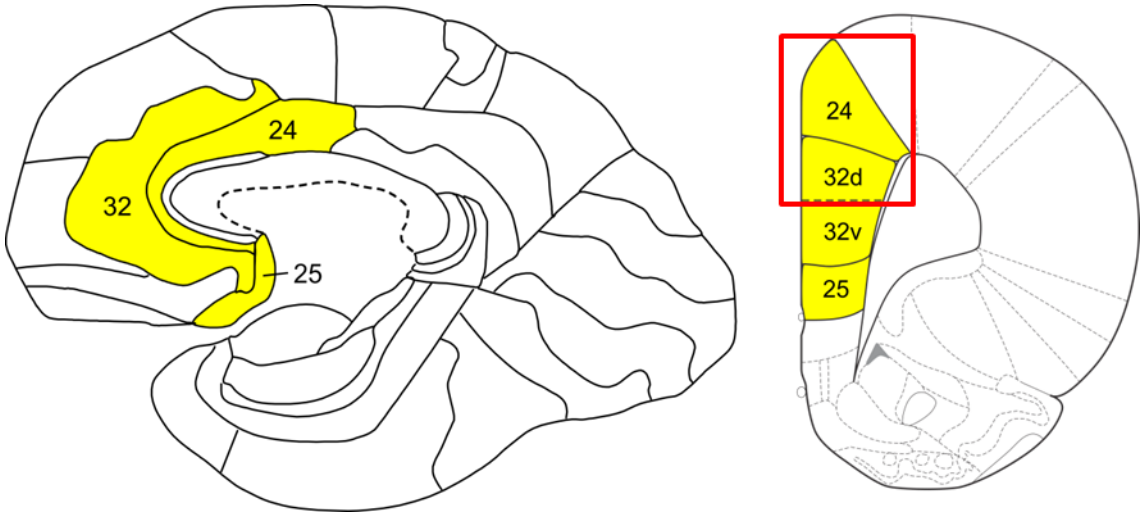

Figure 1.1 Comparative anatomy of the anterior cingulate cortex

In the human brain (left, mid-sagittal view) includes areas 24, 32 and 25 (yellow), and the same applies to the rat brain (right, coronal view at 3mm anterior to Bregma). In this thesis, the dmPFC is includes both area 24 and area 32d (dorsal prelimbic cortex), as indicated by the red box. Recordings were conducted within area 24 for Chapters 3 and 4, but also extended into the dorsal edge of area 32 in Chapter 2.

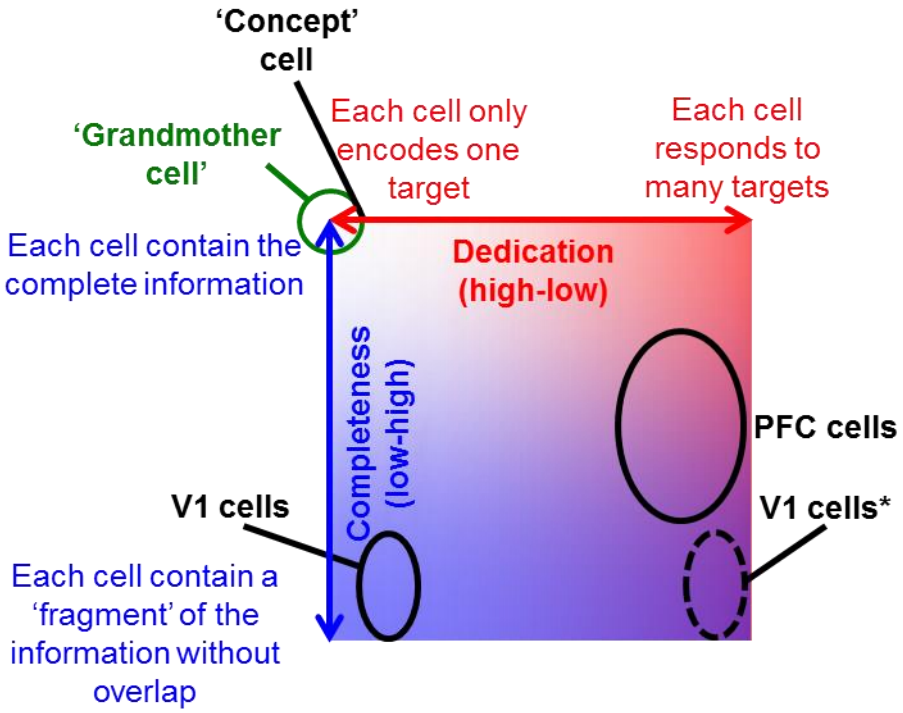

Figure 1.2 Two independent aspects of coding properties

Dedication (blue vertical axis) and completeness (red horizontal axis) are two mutually independent coding properties. Dedication characterizes the variety of information the neurons may encode. Completeness gauges whether a neuron carries the entirety of the signal, or a larger or small proportion of it. Assuming a given target is always encoded by more than one neuron, greater completeness of information within each participating neuron equals greater redundancy in the encoding of the target. In this theoretical framework, a hypothetical 'grandmother cell' would occupy the origin, where the highest level of both completeness and dedication are achieved (green circle), and the 'concept cell' found in the hippocampus comes close to it. At the far end of the dedication axis there are the theoretical omnipotent cells responsive to virtually everything. Prefrontal cells come close to this scenario (solid black circle to the right). In terms of the completeness of information encoding, however, PFC cells are unlikely to go as high as the concept cells, because they were found to encode abstract rules based on one or more characteristics shared across many stimuli or actions, rather than all aspects of one specific stimulus or action. Also, the PFC is known for its distributed coding scheme whereby decoding performance improves quickly as ensemble size increases, suggesting that information encoded in different neurons may complement each other to produce a more *complete* picture. V1 neurons have even lower completeness in information encoding. When a simple visual feature towards which a V1 neuron is tuned is considered the 'target', then V1 neurons have high level of dedication (solid black circle on the left). But because numerous complex and large visual stimuli can elicit response from a given V1 neuron, the neuron is considered to have low level of dedication when such stimuli are defined as the targets of interest.

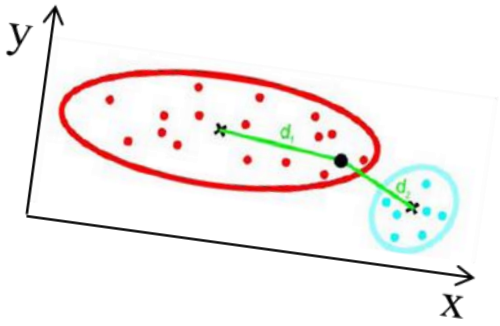

Figure 1.3 The use of Mahalanobis distance
The Mahalanobis distance (D_{Mah}) is suitable when the clusters have different density gradients along different dimensions. Because the black dot lies beyond the boundary of the blue cluster and within that of the red cluster, it belongs to the red cluster. However, it also has shorter Euclidean distance to the centre of mass of the red cluster than to that of the blue cluster, or $D_{Euc\text{-}Blue} < D_{Euc\text{-}Red}$. Therefore, using the D_{Euc}, one would have reached an incorrect conclusion. Because both clusters (especially the red) have different density gradients along different dimensions, e.g. x and y, the assumption for D_{Euc} has been violated. When using D_{Mah}, the D_{Euc} from the black dot to the centre of the red cluster is divided by a greater covariance than that of the black dot to the centre of the blue cluster, giving: $D_{Mah\text{-}Blue} > D_{Mah\text{-}Red}$, supporting the correct conclusion that the black dot indeed belongs to the red cluster.

Chapter 2: Flexible encoding of actions in different contexts and sequences by dmPFC ensembles

2.1 Introduction

Flexible behavior is most commonly achieved by combining a limited repertoire of basic actions in different ways depending upon the current context. While the same action can take on a different meaning when performed in different contexts or sequences, the fact we do not need to relearn each action in every new circumstance points as well to an implicit consistency in our action representations. Thus the key to efficient behavioral adaptation lies in achieving a balance between the ability to update the higher-order meaning of actions while maintaining some consistency in the way they are represented at a more basic level.

There is an extensive literature highlighting the contribution of the frontal cortex to flexible action representations .Within the primate or rodent frontal cortex, individual neurons modulate their firing in response to novel action sequences (Procyk et al., 2000, Mulder et al., 2003, Averbeck and Lee, 2007), to changes in task rules (Mansouri et al., 2009, Rich and Shapiro, 2009, Durstewitz et al., 2010, Vallentin et al., 2012) or changes in task context (Bossert et al., 2011, Hyman et al., 2012). Although such changes are apparent at the single neuron level, a series of rodent studies have investigated in greater detail how populations of frontal neurons respond during flexible modes of behavior. Kargo et al. (2007) showed that while many medial frontal neurons exhibit differential firing to different action sequences, the firing rate of the overall population remained relatively constant (Kargo et al., 2007). Likewise, during a switch from an egocentric to an allocentric memory strategy on a plus maze, the overall population firing pattern was found to differ yet the overall population firing

rate remained constant (Rich and Shapiro, 2009). A similar result was obtained using an operant set-shifting task as the overall ensemble activity state pattern switched abruptly as rats changed strategies even though the overall firing rate of the population remained about the same (Durstewitz et al., 2010). This was also true when the animal was physically moved from one context to another (Hyman et al., 2012). It is presently unclear how or why dramatic pattern reconfigurations occur without significant changes in overall population activity. It is also not clear whether the representations of multiple distinct actions shift independently or coherently with a change in sequence or context.

To gain insight into these questions, the animals were trained to perform three distinct operant actions, namely, nose-poke (NP), lever-press (LP) and wheel-turn (WT), in the correct sequence to obtain reward. Neuronal activities were recorded as the animals performed the task at different stages of training, in different contexts, and in different sequences. As expected based on past studies, unique ensemble activity states consistently emerged in association with the performance of each of the three unique actions. Furthermore, a change in sequence or context strongly shifted overall ensembles representations. Importantly, the amount of shift in the ensemble representations was the same for the three actions, such that the relative differentiation of action representations was preserved. This occurred by virtue of remarkably well balanced increases/decreases in the firing rates of individual neurons such that any increase in the firing of one neuron was mirrored by an equivalent decrease in another neuron or neurons. As a result the overall population rate remained constant in spite of the dynamic morphing of the patterns.

2.2 Methods and material

2.2.1 Apparatus

The Box (25”×18”) was made of transparent Plexiglas, lined with white wallpaper from the exterior. A corrugated plastic main panel was installed with a nose-poke port, a lever, and a wheel installed, from right to left, approximately 2” apart (**Fig. 2.1B**), with a 3W cue light located above each manipulanda. An area of 25”×13” was left for the freely moving rat. On the opposing-side wall, a food cup was located at the centre, with each delivery of reward accompanied by a tone.

The Maze consisted of 4 platforms (coded 1-4) connected by 4 two-foot long passages into a diamond shape (**Fig. 2.1A**), with platforms 1 and 3 at the sharp tips of the diamond, and platforms 2 and 4 connected by an additional shortcut. The individual platforms differed in size, shape, odor, floor texture and wall patterns. Each of Platform 1 to 3 contained a unique manipulanda: a nose-poke port in the right platform; a lever in the middle platform; and a response wheel in the left platform. Above each of these manipulandae was a 3W signal light. Beside each manipulanda, a food-cup could be inserted for food-pellet delivery, which was always accompanied by a 0.5s pure tone at 1.5 KHz. All cue lights, tone-generator, manipulandae and pellet dispensers were operated by a MedPC IV system (Med Associates, Georgia, VT). Doors located at the start of each passage could be controlled from outside of the maze.

2.2.2 Behavioral tasks

Pre-training on the Maze: Training on the maze started with single-action instrumental conditioning. The animals were restricted to each of the NP (right), LP (middle)

and WT (left) platforms to receive daily training on each of the 3 instrumental actions. They progressed through FR1, fixed-interval 10s (FI10s) and random-interval 15s (RI15s) schedules with a performance criterion of 30pellets/20min.

Ten daily sessions of sequence shaping started the day after they attained criterion performance on all three individual operant actions. The animals were placed in the reward platform at the beginning of the session and the only open passage led to the rightmost NP platform, where the light above the NP port was illuminated. Once the animals reached this platform, two doors blocked both exits and the light stayed on until a NP response was emitted. At that point a door opened to allow access to the LP platform 2 and the light above the lever was illuminated. The animals could only leave the LP platform after they performed a LP. At that point the light above the lever was extinguished and the door was opened and the light above the WT on platform 3 was illuminated. Once the animal reached the WT platform and turned on the wheel in a full circle, the light above the WT was extinguished and the last door opened to allow entrance to the reward platform where 4 food pellets were delivered accompanied by a 0.5-s pure tone at 1.5 KHz. After 4 seconds, another door opened so they could move to the NP platform and start the next trial. These sessions lasted for no less than 45min and ended either when the animal stopped responding for at least 3min or 60min had elapsed.

Maze sequence task: On the following day after 10 daily shaping sessions, all doors were removed. Rats were still required to perform the 3 actions in the aforementioned order, but could commit errors such as running in the wrong direction whereby respond out of sequence. Performance was evaluated by the number of out-of-sequence errors committed per trial. Repeated responses on the manipulandae immediately after (within 4s of) the correct

response were not considered as “errors”. The session typically lasted for 60min, but extra time (up to 10min) may be given if the animal was in the middle of the sequence at 60min. The animals continued to receive this self-paced sequence training for a total of 23 days. Only the trials completely free of out-of-sequence errors were used in the analyses of neural data.

Box pre-training and sequence task: Similar to the pre-training on the maze, the animals progressed through FR1, fixed-interval 10s (FI10s) and random-interval 15s (RI15s) schedules, with only one manipulanda available in a 20-min session. Upon reaching the criterion of 30pellets/20min, the animals were trained on the self-paced 3-action sequence using the same procedures as in the maze sequence task (see above). Because no door or partition was present in the Box, the animals were guided only by the cue lights.

Reversed-sequences training and sequence-switch task in the Box: After training on the original sequence (sequence A), the animals were trained to perform the action-sequence in the reversed order (i.e., wheel-turn→lever-press→nose-poke→reward: sequence B). Rats were first guided through sequence B with doors, then without, until they reached the same level of efficiency as exhibited previously on sequence A. At that point the sequence-switch sessions commenced. For sequence-switch sessions they were required to complete at least 20 trials on sequence B within 20min at which point they were removed from the Box for a minute. They were then placed back in the Box on the reward platform and the light above the NP was illuminated instead of the light above the WT (as was the case for sequence B). This prompted them to perform sequence A. Four daily sequence-switch sessions were acquired from all but two animals. Again, only the trials completely free of out-of-sequence errors were used in the analyses of neural data.

Maze-Box context-switch task: After learning the first sequence in the Box task, but before learning the second sequence, 4 rats performed the Maze-Box context-switch task. The rats first performed the original sequence for at least 12 trials within 50min on the maze, followed by performing the same sequence in the Box for at least 12 trials within 15min. This task was administered for 4 consecutive days. Neural data were obtained on a total of 16 sessions.

2.2.3 Subjects

Eight experimentally naïve male Long-Evans rats (450-550g) were used in the Maze task of prolonged sequence training (n=8), 5 out of the 8 went on to receive training on the Box sequence task, and subsequently performed the sequence-switch task in the Box (n=4) and the Maze-Box context-switch task (n=4). The animals were housed in a facility with 12hr light-dark cycle, with all training and recording taking place during the light cycle. For the duration of the behavioral experiments, the rats were food-restricted to just below 90% of their free-feeding weights. Feeding took place in the home cage after their daily training/recording sessions, and water was available *ad libitum* in the cages at all times. All procedures were carried out in accordance with the Canadian Council of Animal Care and the Animal Care Committee at the University of British Columbia.

2.2.4 Surgery

Stereotaxic surgeries were performed on naive rats with sterilized-tip procedures. NSAIDs analgesic, antibiotic, and a local anesthetic, were given before incision. An elliptical-shaped craniotomy was made, centered at: AP: +3.2mm, ML: ±0.5mm. Once the dura mater was retracted, the bottoms of the two bundles of 8 30-gauge tubes, containing a total of 16 tetrodes, were placed bilaterally immediately beside the central sinus, touching the cortical

surface. Each of the bundle had a cylindrical shape with bottom radius ~0.4mm, and were angled by 3.5~5 degrees. The implants were fixed with bone screws and dental acrylic. All tetrodes were extended by ~0.7mm into the brain at the end of the surgery. After 10d of recovery, the tetrodes were advanced ventrally into the mPFC. Once all tetrodes were placed into the dmPFC according to lowering records and atlas coordinates, small adjustments were made with hyperdrives to maximize the number of neurons recorded.

2.2.5 Acquisition of electrophysiological data

For data acquisition, EIB-36TT with pre-amplifier (Neuralynx Inc., Bozeman, MT, USA), connected to the extracellular electrodes, were plugged into HS-36 headstages and tether cables (Neuralynx Inc., Bozeman, MT). Signals were converted by a Digital Lynx 64 channel system (Neuralynx Inc., Bozeman, MT) and sent to a PC workstation, where electrophysiological and behavioral data were read into Cheetah 5.0 software (Neuralynx Inc., Bozeman, MT). Files were then read into Offline Sorter (Plexon Inc., Dallas, TX) for spike sorting, based on visually dissociable clusters in 3D projections along multiple axes for each electrode of a tetrode (peak and valley amplitudes, peak-to-valley ratio, principal components and area). Sorting was confirmed by examining auto- and cross-correlations, and ANOVAs were conducted from the 2D and 3D projections. Spike timestamps were then read into Matlab (Mathworks Inc., Natick, MA) for all further analysis.

2.2.6 Histology

At the end of the studies, the animals were deeply anesthetized using urethane i.p. injection, and a 100μA current was passed through the electrodes for 30s. Animals were then perfused with a solution containing 250ml 10% buffered formalin, 10ml glacial acetic acid, and 10 g of potassium ferrocyanide. This solution causes a Prussian blue reaction, which

marks with blue the location of the iron particles deposited by passing current through the electrodes. The brains were then removed and stored in a 10% buffered formalin/20% sucrose solution for at least 1 week, before being sliced and mounted to determine precise electrode locations. Since multiple sessions were recorded from individual animals the precise recording locations could not be derived from electrode lesions, but all electrode tracks were inferred between the entrance point and the dyed spot. Representative recording sites are shown in **Fig. 2.1C**.

2.2.7 Data analyses

Instantaneous firing rate (iFR): A total of 60 large ensembles were collected from 4 rats that acquired all 3 sequences and successfully switched among them within a given session (mean ensemble size: n=39.4). To obtain an estimate of the neural firing rate for each isolated cell i as a function of time bin *t*, $r_i(t)$, for each spike train in each 200-ms bin, the instantaneous firing rates (iFRs) were calculated as the reciprocals of the inter-spike intervals, convolved with 20-ms Gaussian kernels and then averaged (Durstewitz et al., 2010, Hyman et al., 2012). Neurons firing less than 0.14 Hz were excluded from further analysis, because the sample of spikes was too small (250 or less) to be reliably representative of the cell's activity in relation to behavior. Each lever-press epoch included the 1-s period centered at the moment of lever-press, while the reward-approach period was the 1-s period immediately after the 3rd lever-press epoch, and the reward-consumption period was the 2-s period after the reward-approach period.

Multiple-single unit activity (MSUA) analysis: For population analysis, population vectors $\mathbf{r}(t) = [r_1(t) \ldots r_N(t)]$ were constructed, with N equal to the number of single units isolated from a given recording session. The term MSUA space refers to the N-dimensional

space spanned by all recorded units and populated by these vectors **r***(t)*. Each dot in the MSUA space represents the state of the entire recorded ensemble within one 200 ms bin. All points corresponding to different 200ms bins within the epochs of the same behavior are shown in the same color. All statistical analyses were performed in the full space of all recorded units. For the purpose of visualization, multi-dimensional scaling (MDS) was applied to reduce dimensionality.

To quantify the effects of action and sequence on network activity, the Mahalanobis distances (D_{Mah}) were computed between the sets of N-dimensional vectors associated with task epochs of interest. To control for differences in MSUA space dimensionality (i.e. ensemble size) in D_{Mah} comparisons, a normalization procedure was employed: N_{min} was the minimum number of units recorded in any of the data sets to be compared, and K_{min} was the minimum number of time bins. For data sets with N and K greater than N_{min} and K_{min}, N_{min} units and K_{min} data points were selected at random and D_{Mah} was computed. This procedure was repeated 100 times and the results averaged to make full use of all units and data points recorded. The resultant normalized D_{Mah} averages were used in various statistical analyses. In order to determine the significance level of a given D_{Mah} value, between-sequence separation was compared to within-sequence separation (**Fig. 2.4C**). To calculate average D_{Mah} within a sequence block, bootstrap surrogate blocks were created by randomly shuffling 1-s blocks of the iFR matrices. The distance between the 2 shuffled blocks therefore represents the separation between activities during random behavioral events. The process was repeated 100 times and the D_{Mah} values averaged.

Multiple linear regression (MLR): A model containing 3 mutually independent and uncorrelated factors were constructed to analyze individual neurons' involvement in various

aspects of the task. The factor S was designed to capture the effect of sequence in sequence-switch sessions, or context in context-switch sessions. The second and third factors, A1 and A2 were made to capture the 3 different operant actions:

$F(t) = b_0 + b_1S(t) + b_2A1(t) + b_3A2(t) + \varepsilon(t)$, *(eq. 1)*

For a behavioral epoch *t*, $F(t)$ represented the neuron's averaged normalized iFR within a 1-s window of that epoch. $S(t)$ differentiated the sequence/context identity of each epoch by assigning 1 or -1 to designate from which sequence/context block that epoch belong, i.e. when that operant action occurred. $A1(t)$ and $A2(t)$ differentiated between the actions regardless of sequence/context blocks. $A1(t)$ contrasts between nose-pokes and the other actions by assigning 2 to each nose-poke and -1 to the others, while $A2(t)$ assigns 0 to each nose-poke, 1 and -1 to lever-presses and wheel-turns, respectively. $\varepsilon(t)$ was the error term, and b_0, b_1, b_2 and b_3 were the regression coefficient. Significance on each factor was assessed with the Bonferroni corrected t-statistic ($\alpha=0.05/3=0.0167$). From each session, the beta values were obtained for each cell, for both the sequence factors and the lever factors for both brain regions to perform further statistical analysis.

Generalized Linear Model (GLM): Because the LR analysis assumes a normal distribution and the iFR distributions were not always normal, GLM was used as a check on the results obtained using the LR. The GLM used similar models as the LR but employed a binomial link function. In this case, iFRs were transformed into spike counts/200ms bin. Significance on each factor was again assessed with the bonferroni corrected t-statistic. The same neurons were consistently selected as being responsive to serial position or sequence by both the LR and GLM but the GLM tended to be more liberal and detected more neurons overall. Therefore when in conflict, results from the LR were used.

2.3 Results

2.3.1 Different actions in a sequence are represented distinctively

All experiments in this study involved the completion of action sequences to obtain a final reward. Within any given session, each successful trial consisted of a single lap around the maze or the box where the rats had to perform a nose poke (NP), lever press (LP) and wheel turn (WT) in the correct order to attain food reward. The animals were given ample opportunities (23 daily sessions) become skillful at the task. Before going into the specifics of each experiment, however, a phenomenon common to all tasks in the study is worth notice.

Fig. 2.2A, B, D, E, G and **H** are 3-dimensional visualizations of the multiple single unit activity (MSUA) spaces of entire ACC ensembles recorded from representative sessions. Each plot is limited to activities recorded during the performance of a single type of action sequence on a single apparatus. As illustrated in these plots, ensemble activity states during NPs (red), LPs (green) and WTs (blue) each formed a unique cluster in the multiple single unit activity (MSUA) space. Such between-action separation was not only visible, but also significantly greater than the separation within action, as quantified by the Mahalanobis distance (D_{Mah}) (in the case of **Fig. 2.2C**, Kruskal-Wallis test, $\chi^2(3,176) = 124.18$, $p = 9.7 \times 10^{-27}$; post hoc test: early training vs. control: p=0; late training vs. control: p=0; for statistics on the following experiments see sections below). Given that each dot represents the ensemble activity state during a given action in a single trial, as the animal completed one trial followed by another, the ensemble activities rotated through these three states in a cyclical manner. These plots demonstrate the degree of loyalty with which the ACC ensembles followed every step within each action sequence during every trial.

2.3.2 Separation among action representations was maintained after prolonged training

The first set of experiments analyzed whether the relationship between the representations of the three different actions changed as the animals performed these actions repeatedly over hundreds of trials in a fixed sequence. Within a session, each successful trial on this self-paced sequence task consisted of a single lap around the maze where the rats were required to perform a nose poke (NP) →lever press (LP) →wheel turn (WT) in a sequence to attain food reward (**Fig. 2.1A**).This stage of training lasted for 23 days, and the neural activities from the first 4 training days were compared to the last 4 training days (**Fig. 2.1C**). The behavioral performance during late training days showed a significant improvement from early training (paired-sample t-test, t_{19}=3.43, p=0.0028). **Fig. 2.1D** shows an example of typical recording sites. The white arrows in the dorsal edge of the prelimbic cortex marks the last recording sites in the prolonged study, hence it may be inferred that the vast majority of neurons were recorded in the ACC, and all neurons were recorded from the dmPFC.

Comparing **Fig. 2.2A** and **B**, there is no visible difference in the separation among the 3 action clusters between the early and the late stages of training (Kruskal-Wallis test and post hoc test: early vs. late training: p=1). Thus while one might naively hypothesize that actions should in some sense become 'linked' or otherwise associated if performed repeatedly in a sequence, this did not appear to be the case as the separation between the action representations was not impacted by training or experience. Hence the ACC ensembles appeared to keep reliable track of each step within the action sequence, not only from trial to trial, but also from day to day, at least up to 23 days of training.

2.3.3 Separation among action representations was maintained after context switch

After 23 days on the 3-action maze task, the same animals were retrained on the same type of NP→LP→WT sequence task but this time within a single operant box (**Fig. 2.1B**). Once the animals achieved criterion performance on this 3-action box task, 'context-switch' sessions commenced. Context-switch sessions required the rats to complete a block of trials on the 3-action maze task followed by a block of trials on the 3-action box task. During these context-switch sessions, the animals displayed similar levels of performance on both tasks (paired t-test, t_{13}=1.49, p=0.16). Although the two tasks were performed in very distinct contexts, the separation among the action representations was nevertheless similar (Kruskal-Wallis and post hoc test, maze vs. box task: p=0.22; **Fig. 2.2F**). Thus a change in context does not appear to affect the differentiation between the three different actions, given that the abstract task rule (i.e. the 3-action sequence) remained unchanged.

2.3.4 Separation among action representations was maintained after a switch in action order

Given that a switch in the environment where the sequence was conducted did not alter the distinctness of the action representations, it was worth investigating whether altering the relationship of the 3 actions within the sequence may lead to a different result. For these experiments a new set of animals was first trained on the maze to perform the original NP→LP→WT sequence and then to perform the sequence in the reverse: WT→LP→NP. On the 'sequence-switch' sessions, the animals first completed a block of trials on the more recently acquired reversed sequence and were then switched back to the original NP→LP→WT sequence. They were cued to the sequence switch while on the reward platform by the illumination of the light above the NP rather than the light above the WT.

Performance on each sequence was equal during sequence-switch sessions (paired t-test, t_{13}=-0.88, p=0.39). The D_{Mah} between the 3 actions was also similar for the 1st (**Fig. 2.2G**) and the 2nd sequence (**Fig. 2.2H**) (Kruskal-Wallis and post hoc test, 1st vs. 2nd sequence: p=1; **Fig. 2.2I**).

Collectively, these results show that on a task where a sequence of 3 different actions are performed in order to obtain reward, dmPFC ensembles generate equally distinct representations for each action regardless of the level of prior task experience, the contexts where the actions are performed or the ordering of the actions themselves.

2.3.5 Context-switch or sequence-switch resulted in an equivalent amount of shift in the representations of all actions

The analyses above focused exclusively on the action periods themselves (i.e. the 5 time bins surrounding each action). While the three actions always maintained a consistent relationship to each other what was not taken into account in these analyses was the issue of whether the entire block of 3 actions was changing. In other words, whether or not the higher-order representation of 'sequence' or 'context' produced a wholesale shift in the way any or all actions were represented. To address this question, the movement in the representation of the 3 actions were examined as a single sequence block in the MSUA space, as the animals experienced a change in context or sequence.

As shown in **Fig 2.3A**, a clear shift in the overall ensemble activity state occurred when an animal performed the same 3-action sequence in the two contexts (Fig 3A). Statistically, the D_{Mah} between the points associated with a single action were larger if this

action was performed in two contexts than if it were performed a single context (Wicoxon signed rank test: ranksum=3313, $p=5.4\times10^{-13}$, **Fig. 2.3B**).

The context-dependent shift in ensemble activity state was not related to a change in the overall level of ensemble activity (unpaired t-test: $t_{3562}=0.39$, p=0.69; **Fig. 2.3C**) or any change in the level of performance (paired t-test: $t_{13}=1.49$, p=0.16; **Fig. 2.3D**).

Perhaps even more striking, the amount of shift in ensemble activity was remarkably similar across the 3 actions (1-way ANOVA, $F_{2,45}=0.25$, p=0.78, **Fig. 2.3E**). Furthermore, the proportions of neurons that underwent a significant increase versus decrease in firing rate during the 3 actions were also equivalent (Pearson's Chi-squared test for goodness of fit: $\chi^2(2)=5.22$, p=0.073, **Fig. 2.3F**). This means that a change in context shifted produced a coordinated change in the entire block of action representations without altering the relationships between the actions themselves.

Similar to a switch in the task context, a reversal in the sequence of actions also resulted in a significant shift in the overall ensemble activity states encompassing the 3 actions (from black dots to gray dots, **Fig. 2.4A**; Wicoxon signed rank test: ranksum=2499, $p=1.7\times10^{-10}$, **Fig. 2.4B**). Once again this occurred even though the separation between the actions within each sequence (**Fig. 2.2G-I**) as well as the overall level of activity remained unperturbed (unpaired t-test: $t_{3610}=-0.89$, p=0.37; **Fig. 2.4C**), and the performance remained the same (paired t-test: $t_{13}=-0.89$, p=0.39; **Fig. 2.4D**). Also similar to the context-switch task, the distance by which the activity states shifted was equivalent across the 3 actions (Kruskal-Wallis test, $\chi^2(2,39)=1.95$, p=0.38, **Fig. 2.4E**) as was the proportions of neurons that underwent significant increase or decrease in firing rate (Pearson's Chi-squared test for

goodness of fit: $\chi^2(2)=2.60$, p=0.27, **Fig. 2.4F**). Therefore, a change in the context or sequence was encoded by way of an almost perfectly balanced reorganization of the action representations across the population.

2.3.6 Properties of the cells that produce the balanced shifts in action representations

The results presented above showed that sequence (or context) dependent shifts in ensemble activity state patterns are associated with an equal movement in the MSUA space for all 3 actions. This could arise if there were a dedicated group of 'sequence' (or 'context') selective neurons that changed their firing in different sequences (or contexts) plus a separate group of dedicated 'action' selective neurons that did not. An alternative possibility is that sequence (or context) signal permeated throughout all neurons such that a change in sequence (or context) collectively reorganized the entire ensemble. To gain insights into this issue a MLR model was employed to first categorize the neurons based on whether they responded only to sequences, only to actions or to both. A similar MLR analysis was performed on the context-switch sessions but in this case cells were categorized based on whether they responded only to contexts, only to actions or to both. Once classified, it was then investigated how the activity of each group was affected by a change in sequence (or context).

Fig. 2.5A shows the proportions of the three types of neurons in the sequence-switch sessions, while **Fig. 2.5B** and **C** gives the absolute beta values on each of the factors. **Fig. 2.5D, E** and **F** show the parallel results from the context-switch sessions. One surprising outcome from the MLR analysis was the large proportion of dmPFC neurons that multiplexed information about sequences and actions or contexts and actions. Also of note was the considerable strength of such neurons on the sequence or context factors which actually overshadowed that of the neurons attaining significance on the pure sequence (unpaired t-test:

t_{127}= -3.23, p=0.0016, **Fig. 2.5B**) or context factors themselves (Wilcoxon signed rank test, ranksum=6682, p=2.3×10^{-6}, **Fig. 2.5E**).

Regardless of which group was analyzed, a remarkably similar proportion of cells exhibited firing rate increases and decreases following a shift in sequence (**Fig. 2.6A**) or context (**Fig. 2.6E**). Even more striking was that the magnitude of the change in firing from one sequence or context to the next was also balanced (**Fig. 2.6B,E**). This was true even for neurons that exhibited the most stable action representations (i.e. 'action-only' neurons in **Fig. 2.6A-D**). However, it was the neurons that multiplexed sequences and actions or contexts and actions that exhibited the greatest change in firing upon a switch (**Fig. 2.6B,E**). **Fig. 2.6C** (left) illustrates this phenomenon in a pair of representative neurons: Neuron A had an increase in firing rate from the 1st to the 2nd sequence whereas neuron B showed an equivalent amount of decrease. Together the change was balanced thereby producing near-zero change in the overall activity of the 2-cell ensemble. In spite of this, their ensemble activity state has shifted from point X to Y in the 2-dimensional space (**Fig. 2.6C**, right). This distribution of balanced firing rate changes suggests that a switch in sequence or context truly reorganized the firing properties of large proportions of neurons and the shift in MSUA activity was not attributed to specialized subsets of neurons that only encoded 'sequence' or 'context'.

On the other hand, the reorganization was not apparent in all of the recorded neurons. For instance, even though there was a large proportion of neurons that failed to attain significance on any MLR factors (**Fig. 2.6A,E**) these neurons showed far less of a change in firing in either direction upon a sequence switch (Kruskal-Wallis test: $\chi^2(3,590)=279.73$, p=2.4×10^{-60}, post hoc test: interactive vs. other cells: p=0, sequence-only vs. other cells: p=0; **Fig. 2.6B**) or context switch (Kruskal-Wallis test: $\chi^2(3,598)=249.98$, p=6.6×10^{-54}, post hoc

Tukey's test: interactive vs. other cells: p=0, context-only vs. other cells: p=0; **Fig. 2.6F**). This observation suggests that while the switch may have not have impacted all neurons in the population, it most strongly impacted those that functionally participated in the task.

2.4 Discussion

The results of the present study showed that the differentiation of specific action representations in the dmPFC remains very stable across training as well as with changes in the context or the sequence in which the actions were performed. On the other hand, actions representations move collectively as unified blocks following a change in context or sequence. This arrangement may provide both the flexibility and familiarity required for adapting behavior in a fast yet efficient manner.

There is a long-standing literature showing that the cingulate cortex and the primate ACC in particular have an important role in action initiation. Microstimulation of the cingulate cortex in humans elicits various types of actions (Talairach et al., 1973) while classic neuroanatomical studies in primates revealed that the cingulate cortex has strong connections with cortical and subcortical motor regions (Muakkassa and Strick, 1979, Barbas and Pandya, 1987, Vogt and Pandya, 1987, Vogt et al., 1987) including the pontine nuclei (Vilensky and van Hoesen, 1981) as well as a topographical projection spinal cord itself (Hutchins et al., 1988). Moreover, cingulate cortex neurons, especially within the 'cingulate motor area' respond to either self-initiation or cue guided actions (Niki and Watanabe, 1979, Isomura and Takada, 2004) and the elevated firing in the ACC preceding movement can be as robust as the pre-movement activity found in M1 (Shima et al., 1991).

Nevertheless, the relation between the rodent dmPFC and action generation is not direct but is more consistent with a role in determining how actions are used or controlled. In

rodent dorsomedial frontal cortex, neurons often fire robustly preceding movements (Jung et al., 1998, Narayanan and Laubach, 2006, Cowen et al., 2012, Hyman et al., 2013) and transient inactivation of this region disrupted activity in motor cortex preceding movements (i.e. during the delay period of a delayed response task), but not firing directly time-locked to the responses themselves (Narayanan and Laubach, 2006). On tasks involving sequences of similar actions, neurons in the primate medial frontal cortex fire prior to different action sequences or fire in the sequence to which the actions will be performed (Hoshi et al., 2005). Since in these studies, neurons typically exhibited sequence specificity prior to the initiation of the action sequence itself it suggests that sequence coding and action coding can be independent at a neural level in the medial frontal cortex. Indeed, primate ACC neurons can encode the serial order of movements independent of the component movements, their kinematics or the arm used to perform them (Procyk et al., 2000, Hoshi et al., 2005). More abstract forms of information can also impact the way ACC neurons respond to actions. For instance, if the rules of a task change on a set-shifting task, there is a tightly correlated change in the way the same responses are represented by the anterior cingulate cortex of rats both at the level of single neurons and ensembles (Rich and Shapiro, 2009, Durstewitz et al., 2010). In addition Shidara & Richmond (2002) showed that ACC neurons responded differently to a given action depending on its serial position relative to reward. I obtained a similar result in rat dmPFC neurons as described in Chapter 3.

Collectively these data are consistent with the idea that the ACC represents actions and stimuli with reference to the task being performed (Hoshi et al., 2005). Accordingly it has been proposed that the ACC is an integral part of a network that formulates task-sets, i.e. the

dynamic configuration of perceptual, attentional, mnemonic, and motor processes necessary to accomplish a particular task (Weissman et al., 2005, Dosenbach et al., 2006, Sakai, 2008).

The present study presents a potential mechanism through which task- sets could impact action representations. A unique pattern of activity was associated with each action, consistent with many past studies in the rat. Both Kargo et al. (2007) and Rich and Shapiro (2009) found that mPFC can switch between different patterns yet the overall level of activity in the ensemble remains essentially the same. Here these findings were expanded to show that not only can ensemble patterns switch in the absence of a change in the overall levels of activity but that the balance in the changes in firing rates of the individual neurons was almost perfect. Furthermore, the magnitude of change in each class was balanced in terms of the relative size of the increases and decreases across the group (**Fig. 2.6B,D**). In other words, if a group of neurons was very sensitive to sequence or context changes, the neurons would exhibit large changes in firing yet the large increase in firing in half of the group was balanced by an equivalent large decrease in other half of the neurons. Likewise, this may occur across groups of neurons showing small changes in activity upon a change in context or sequence. It was the reconfiguration of the patterns rather than overall activity levels that transferred information about a change in context or sequence. Because the reconfiguration was so well balanced it ensured that the relationships between the three actions remained constant as did the overall levels of activity.

In a sequence task where rats made nose-poke responses at sequential locations, Euston and McNaughton (2006) compared identical segments of movement in different sequential contexts. They reported a significantly larger proportion of cells preferring the first to the second segment, which was accounted for by variation in the movement trajectories

rather than representing different sequential contexts (Euston and McNaughton 2006). In our study, when sequential context of the actions was changed, the movement trajectories were essentially reversed. However an equal proportion of cells were observed to favor either sequence (**Fig. 2.6A**), which could not have been the case if the reorganization in firing pattern was associated with movement trajectories alone. Instead, our results seem to suggest that both sequences, albeit involving different movement trajectories, were treated as equally important by the ACC ensembles, because they each was the dominant rule of the task at the time and thus had equivalent motivational valence. This discrepancy in finding may be due to a number of differences in the design of the two studies, such as reward delivery occurring after each movement segment (Euston and McNaughton 2006) versus only at the end of all actions (present study), and the use of a single (Euston and McNaughton 2006) versus multiple types (present study) of operant actions in the sequence.

Pattern reconfiguration with a net zero change in overall activity raises fundamental questions about what overall changes in gross activity actually mean. One factor that produces profound changes in overall activity is the presence or absence of reward (Pratt and Mizumori, 2001, Cowen et al., 2012, Caracheo et al., 2013). For instance, dmPFC neurons change their firing rates from baseline vastly more for actions that are closer in time to reward (Chapter 3) (Ma et al., 2014). While the dynamic change was still largely balanced, neurons showing an increase in activity outweighed those showing a decrease. Therefore, one possibility is that discrete or declarative information are encoded by through balanced pattern based reconfigurations while information that varies along a continuum such as the size of the reward, the delay before the reward, the effort required or the number of actions to be completed before reward delivery, is encoded in changes in overall levels of activity. From

this perspective it would seem that changes in the BOLD responses detected using fMRI, are not indicative of whether the ACC is “activated” or not, since it is “activated” in a more or less equivalent manner under a variety of conditions but rather whether the subject likes a certain situation or is anticipating a particular outcome. In this way, the ACC may use two different flavors of rate coding to convey two different types of information: different patterns of activity with fixed overall levels are used to represent declarative information while variance in the levels of activity associated with a given pattern might convey the motivational significance of the declarative information carried by these patterns. This concept will be elaborated upon in subsequent chapters.

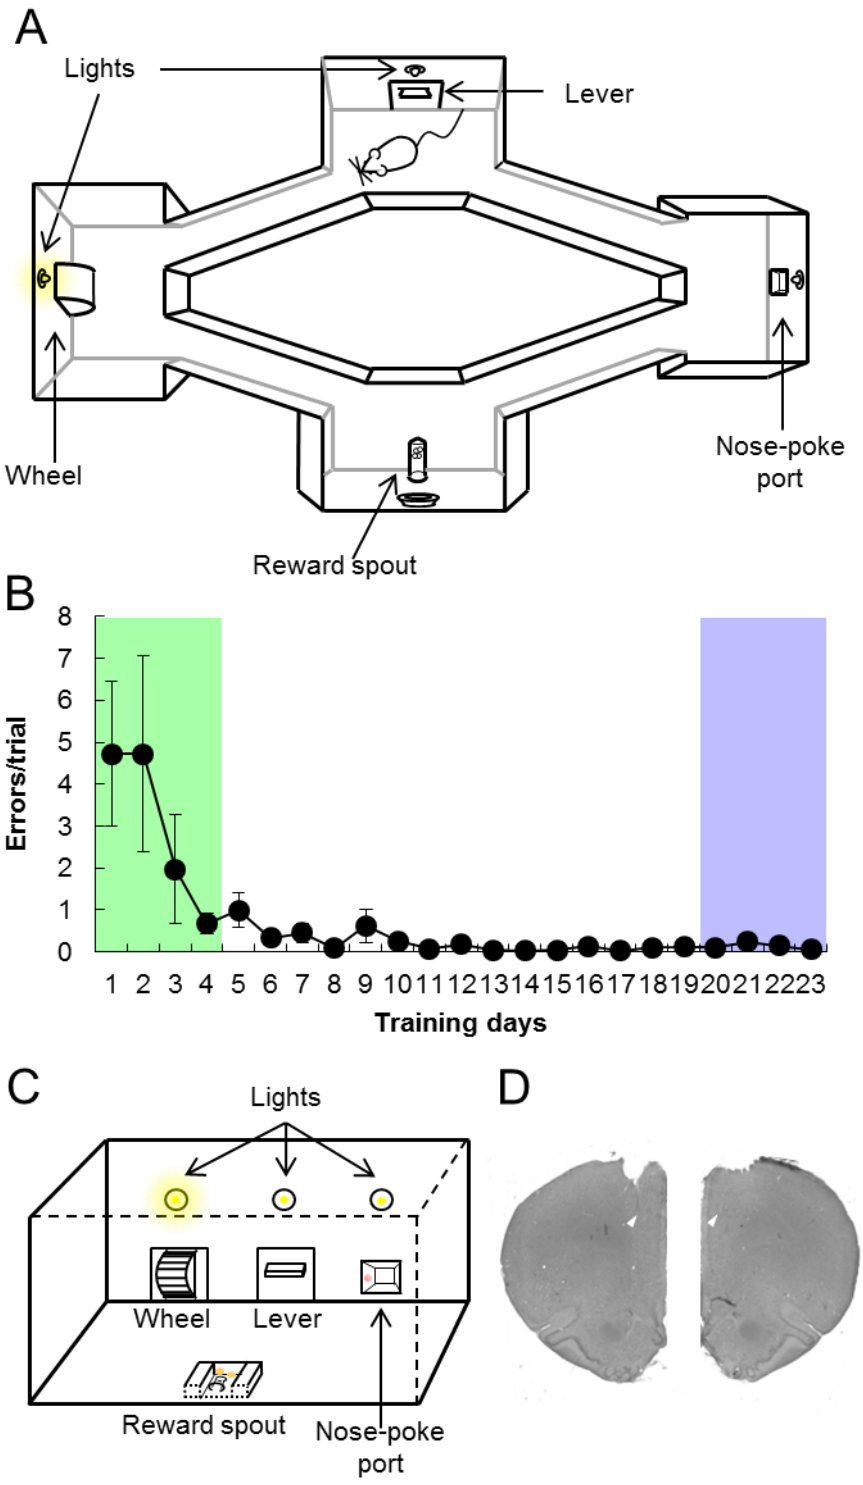
A
Lights
Lever
Wheel
Nose-poke port
Reward spout
B
Errors/trial
0 1 2 3 4 5 6 7 8
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
Training days
C
Lights
Wheel
Lever
Reward spout
Nose-poke port
D

Figure 2.1 Task apparatus and histology of recording sites

A) A schematic of the 4-Platform Maze task. The maze was a diamond shaped with 3 unique platforms each containing a unique manipulandae on 3 of the tips of the maze. The starting or reward platform was the fourth tip. The rat was cued as the correct direction to lap the maze using lights above each manipulandae. Upon the completion of each action, the light above the manipuldandae of the next platform in the sequence was illuminated. This light indicated that next manipulandae was now active. **B)** In the operant Box sequence task, the animals performed the same 3 actions in the same sequence(s) as in the maze. The difference was that now all responses occured in the same sensory context. In this case, all manipulandae were located on the same main panel, and the reward spout was located at the centre of the opposing wall. **C)** Representative electrode tracks (white arrows) in the medial frontal cortex through the ACC.

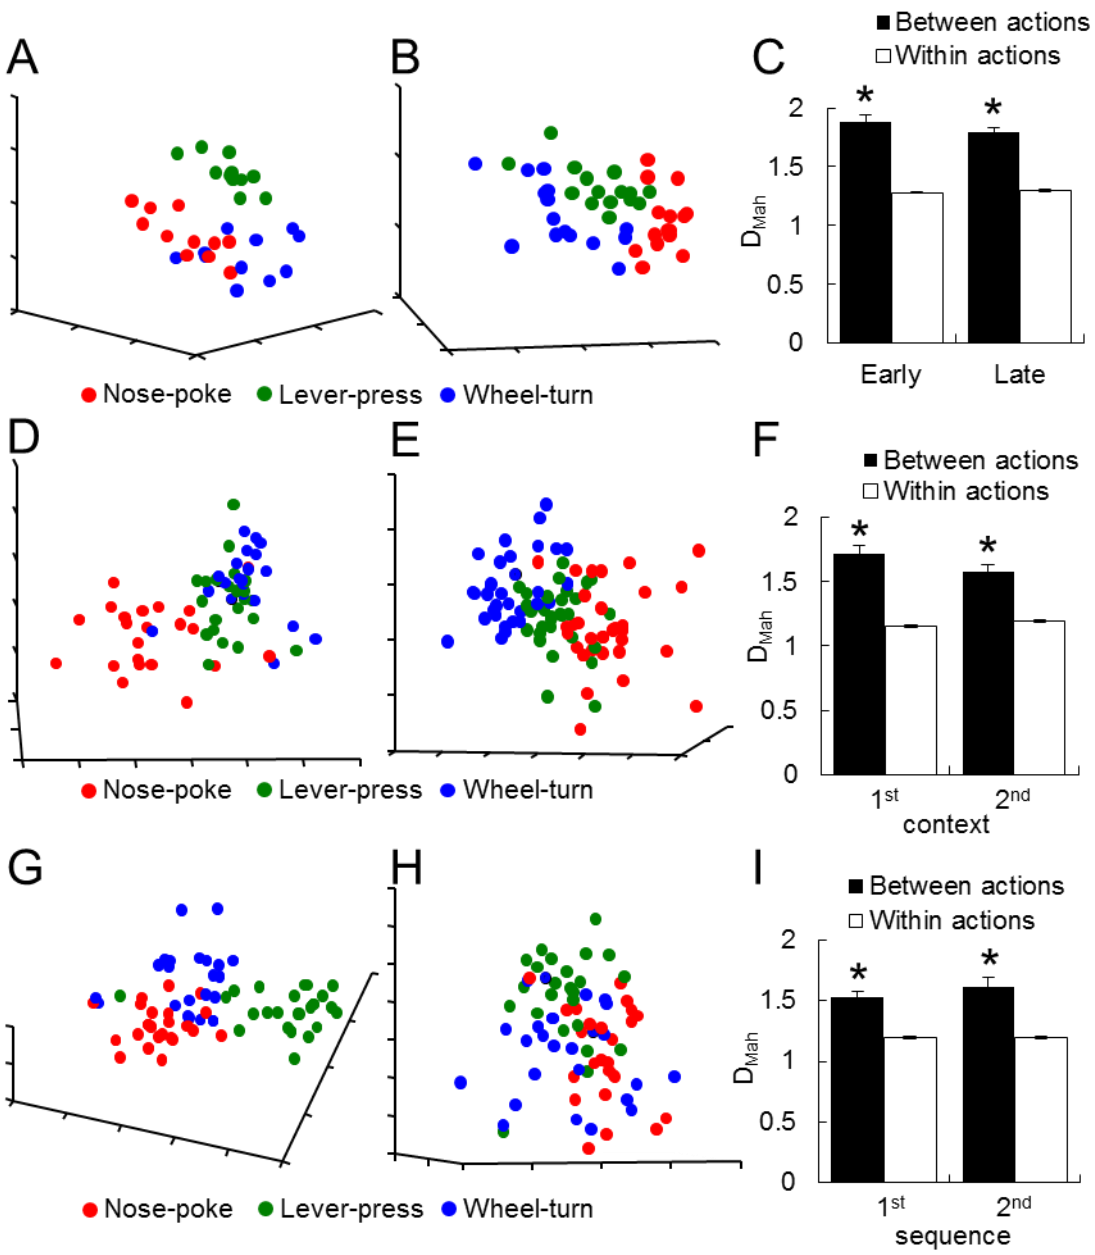
A
B
C
Between actions
Within actions
2
1.5
1
0.5
0
D_{Mah}
*
*
Early
Late
Nose-poke
Lever-press
Wheel-turn
D
E
F
Between actions
Within actions
2
1.5
1
0.5
0
D_{Mah}
*
*
1st
2nd
context
Nose-poke
Lever-press
Wheel-turn
G
H
I
Between actions
Within actions
2
1.5
1
0.5
0
D_{Mah}
*
*
1st
2nd
sequence
Nose-poke
Lever-press
Wheel-turn

Figure 2.2 Action representation in different training stages, contexts and sequences

Actions are represented as distinct ensemble activity state patterns in the mPFC, across different stages of training, different contexts and different sequences. **A)** An example MSUA space constructed from the iFRs of all 64 mPFC neurons recorded during a single session on day 3 of self-paced sequence training in the Maze task. The full space of all 64 dimensions—one for each neuron—was reduced to 3 dimensions using multidimensional scaling for the purpose of visualization. Each dot is a population vector representing the activities of the entire ensemble during the periods when the rat produced the correct responses. Dots are colored red if they were time bins associated with nose-pokes, green if associated with lever-presses and blue if associated with wheel-turns. **B)** An example MSUA space constructed from the iFRs of all 60 ACC neurons recorded on day 22 of self-paced sequence training of the same rat. The same behavioral epochs and color schemes were used as in **(A)**. **C)** From all early-training (D1-4) and late-training sessions (D20-D23) with a minimum of 8 error-free trials, the distances (D_{Mah} in MSUA space) among the activity states associated with the 3 operant actions were calculated in 19-dimensional space, and compared to control distances within action-states. The average separation among different action representations (black bars) was significantly larger than within-action control distances (white bars), both in the early and late training stages. **D)** An example of a reduced MSUA space constructed from the iFRs of 46 ACC neurons recorded during the first part of a single context-switch session when the rat performed the 3-action sequence in the 1st context (Maze). The same behavioral epochs and color schemes were used as in **(A)**. **E)** The same MSUA space as D) is shown (in a slightly different view), but with the action-periods from the second part of the same session, when the rat performed the 3-action sequence in the 2nd context (Box). **F)** Using data from all the context-switch sessions, the distances (D_{Mah} in MSUA space) among the activity states associated with the 3 operant actions were calculated and compared to control distances within action-states. The average separation among different actions (black bars) performed in either context was significantly larger than within-action control distances (white bars). The between-action separations did not differ between the 2 contexts. **G)** An example of a reduced MSUA space constructed from the iFRs of 58 ACC neurons recorded during the first part of a single sequence-switch session when the rat performed the WT→LP→NP sequence B in the Box. The same behavioral epochs and color schemes were used. **H)** The same MSUA space as **G)** is shown (in a slightly different view), but with the action-periods from the second part of the same session, when the performed the NP→LP→WT sequence A. **I)** The distances (D_{Mah} in MSUA space) among the activity states associated with the 3 operant actions were calculated and compared to control distances within action-states. The average separation among different actions (black bars) performed in either sequence was significantly larger than within-action control distances (white bars). The between-action separations did not differ between the 2 sequences. *$p<0.00001$

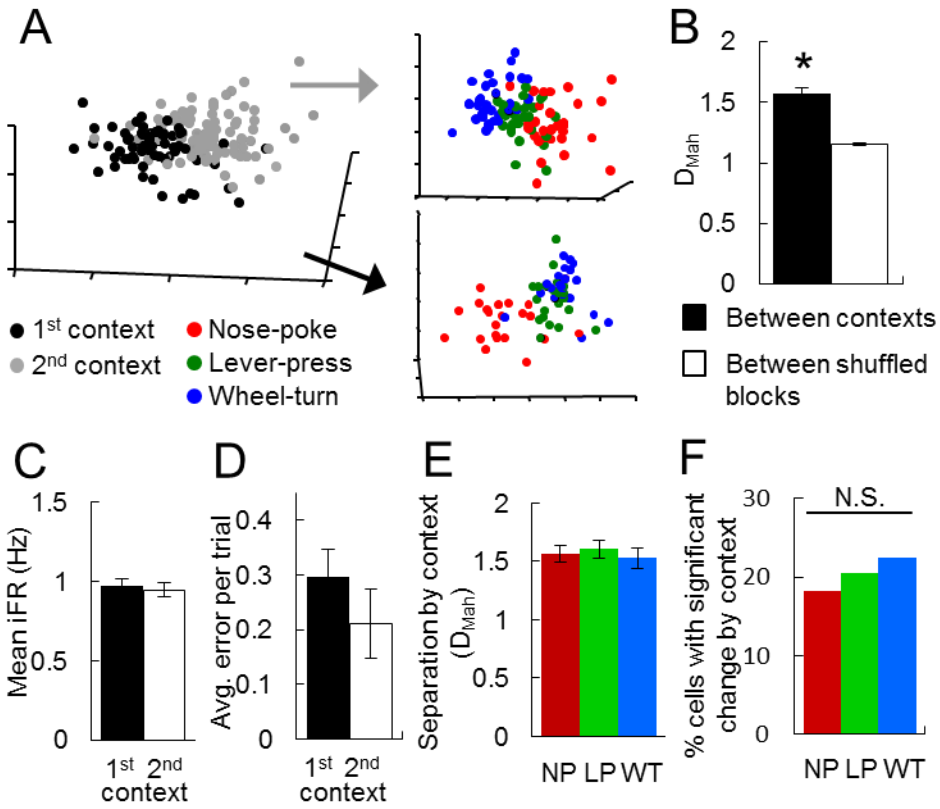
A
1st context
2nd context
Nose-poke
Lever-press
Wheel-turn
B
*
DMah
Between contexts
Between shuffled blocks
C
Mean iFR (Hz)
1st 2nd context
D
Avg. error per trial
1st 2nd context
E
Separation by context (DMah)
NP LP WT
F
% cells with significant change by context
N.S.
NP LP WT

Figure 2.3 Action representations shifted when context changed

The same actions were represented as distinct activity states by the mPFC ensembles in different contexts. **A)** The same MSUA space from **Fig. 2D** and **E** with the combined activity state correlated with the 3 actions in the 1st context (black dots, left panel), which was distinct from the combined state of the 3 actions in the 2nd context (gray dots). Lower right panel shows the action correlates in the 1st context, and upper right panel shows those in the 2nd context. The right panels are the same plots as **Fig. 2D** and **E**. **B)** For all context-switch sessions, the distances (D_{Mah} in MSUA space) between the activity states associated with the same operant actions in the 2 contexts, e.g. between nose-pokes in the 1st context and the nose-pokes in the 2nd context, were calculated in 19-dimensional space, and compared to control distances between shuffled blocks within context. The average separation between different contexts (black bars) was significantly larger than within-context control distances (white bars). **C)** The averaged level of activity across the ensembles remained unchanged from the 1st to the 2nd context within session. **D)** The level of performance remained equivalent whether the animal was in the 1st or the 2nd context within session. **E)** The distance between the activity state of NP in the 1st context and that of NP in the 2nd context (red) is similar to the distance between the activity states associated with LP (green) and the distance between the activity states associated with WT (blue). In other words, all 3 actions underwent the same amount of 'shift' in their activity states as the animal switched from the 1st to the 2nd context. **F)** The percentage of cells underwent significant change in firing rates during NP periods (red) from the 1st to the 2nd context was equivalent to the percentage of cells responded significantly differently to LP (green) and to those responded differently to WT (blue). That is, similar amounts of cells' activity were modulated by context-switch during the 3 actions. $*p<1\times10^{-12}$

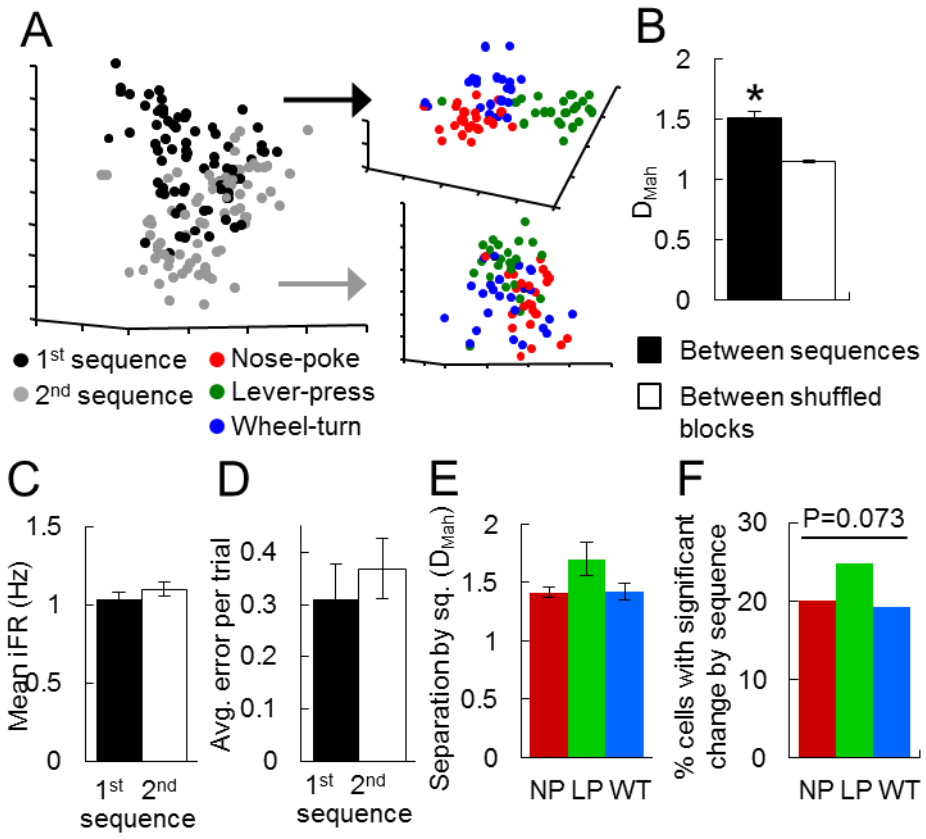

Figure 2.4 Action representations shifted when sequence switched

The same actions were represented as distinct activity states by the mPFC ensembles in different sequences. **A)** The same MSUA space from **Fig. 2.2G** and **H** with the combined activity state correlated with the 3 actions in the 1st sequence (black dots, left panel), which was distinct from the combined state of the 3 actions in the 2nd sequence (gray dots). Upper right panel shows the action correlates in the 1st sequence, and lower right panel shows those in the 2nd sequence. The right panels are the same plots as **Fig. 2.2G** and **H**. **B)** For all sequence-switch sessions, the distances (D_{Mah} in MSUA space) between the activity states associated with the same operant actions in the 2 sequences were calculated in 19-dimensional space, and compared to control distances between shuffled blocks within sequence. The average separation between different sequences (black bars) was significantly larger than within-sequence control distances (white bars). **C)** The averaged level of activity across the ensembles remained unchanged from the 1st to the 2nd sequence within session. **D)** The level of performance remained equivalent whether the animal was in the 1st or the 2nd sequence within session. **E)** All 3 actions underwent the same amount of 'shift' in their activity states as the animal switched from the 1st to the 2nd sequence. Color scheme is the same as **Fig. 2.3E**. **F)** Statistically equivalent percentages of cells' activity were modulated by sequence-switch during the 3 actions. Color scheme is the same as **Fig. 2.3F**. *$p<2\times10^{-10}$

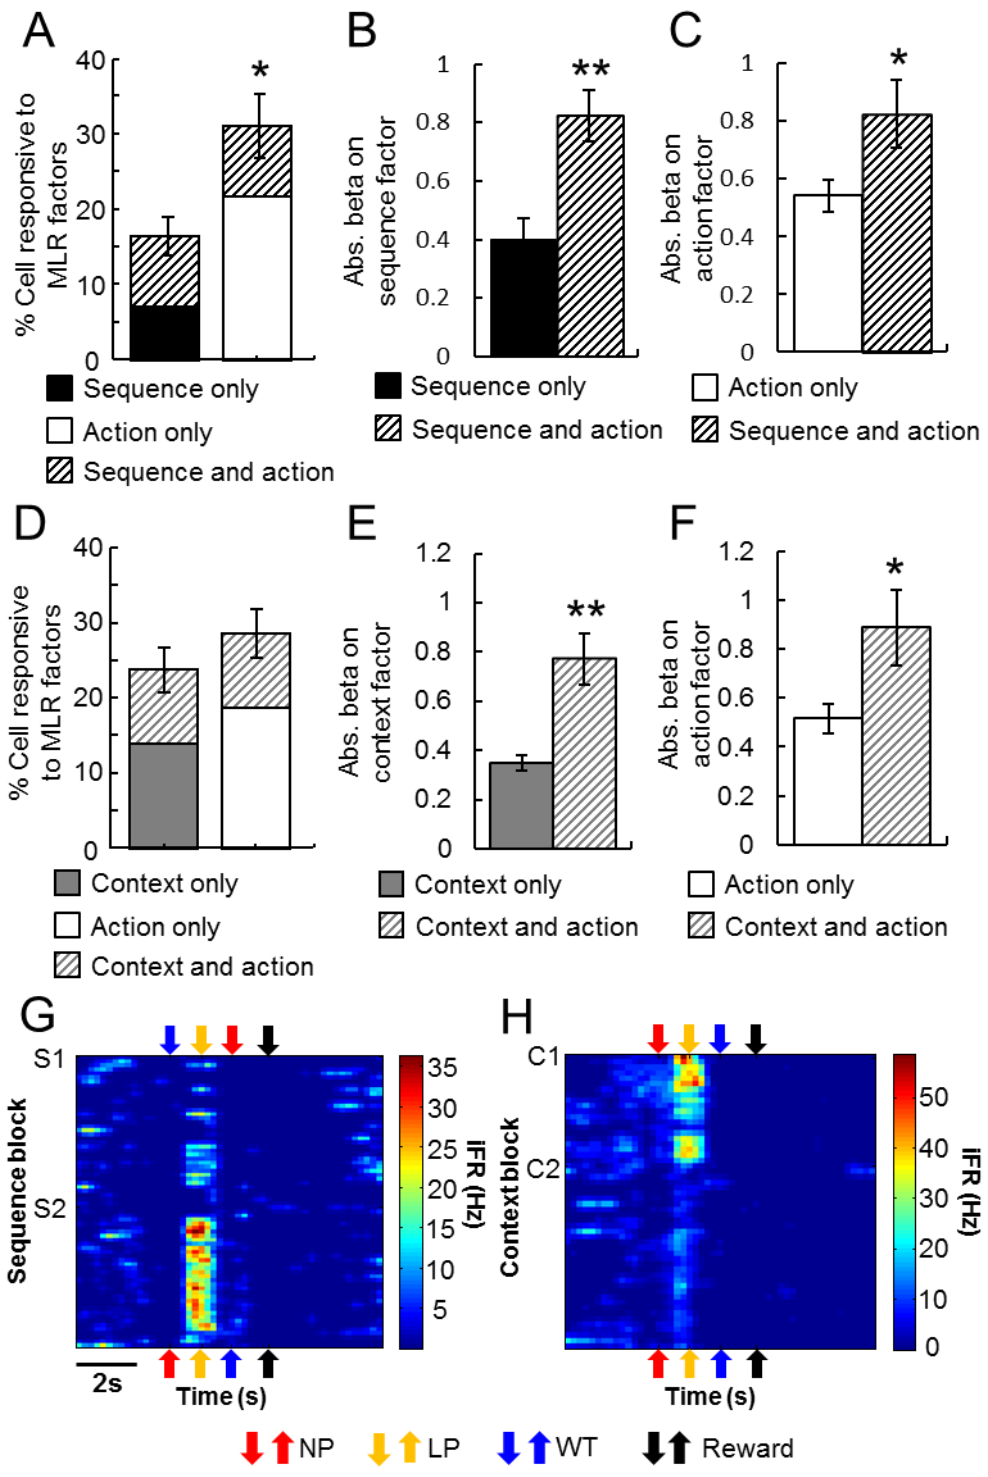
A
% Cell responsive to MLR factors
40
30
20
10
0
*
Sequence only
Action only
Sequence and action
B
Abs. beta on sequence factor
1
0.8
0.6
0.4
0.2
0
**
Sequence only
Sequence and action
C
Abs. beta on action factor
1
0.8
0.6
0.4
0.2
0
*
Action only
Sequence and action
D
% Cell responsive to MLR factors
40
30
20
10
0
Context only
Action only
Context and action
E
Abs. beta on context factor
1.2
1
0.8
0.6
0.4
0.2
0
**
Context only
Context and action
F
Abs. beta on action factor
1.2
1
0.8
0.6
0.4
0.2
0
*
Action only
Context and action
G
S1
S2
Sequence block
35
30
25
20
15
10
5
IFR (Hz)
2s
Time (s)
H
C1
C2
Context block
50
40
30
20
10
0
IFR (Hz)
Time (s)
NP
LP
WT
Reward

Figure 2.5 Proportion of cells responsive to sequence and context switch

The proportions of cells sensitive to different factors in the sequence-switch and context-switch tasks according to MLR, and their degrees of sensitivity to these factors. **A)** The proportions of cells responsive only to the sequence factor (black bar), only to the action factors (white bar), and to both types of factors (hatch black bar). **B)** Cells responsive to both types of factors, or 'interactive' cells (hatch black bar), were on average more sensitive to the sequence factor than the cells only responsive to sequence (black bar). **C)** Sequence-action interactive cells (hatch black bar) are also more sensitive to the action factors than cells responsive only to these factors (white bar). **D)** The proportions of cells responsive only to the context factor (gray bar), only to the action factors (white bar), and to both types of factors (hatch gray bar). **E)** and **F)** show that similarly in the context-switch task, cells responsive to both context and action factors (hatch gray bar) had higher sensitivity to the sequence factor than the sequence-only cells (gray bar **,E**), and higher sensitivity to the action factors than the action-only cells (white bar, **F**). **G)** is an example neuron responsive to both sequence and action factors. Specifically, it responded strongly and reliably to lever-presses only in the 2^{nd} (bottom yellow arrow) but not the 1^{st} sequence (top yellow arrow). **H)** shows an example neuron responsive to both context and action factors. Specifically, it responded strongly and reliably to lever-presses only in the 1^{st} (top yellow arrow) but not the 2^{nd} context (bottom yellow arrow). *$p<0.02$, **$p<0.002$

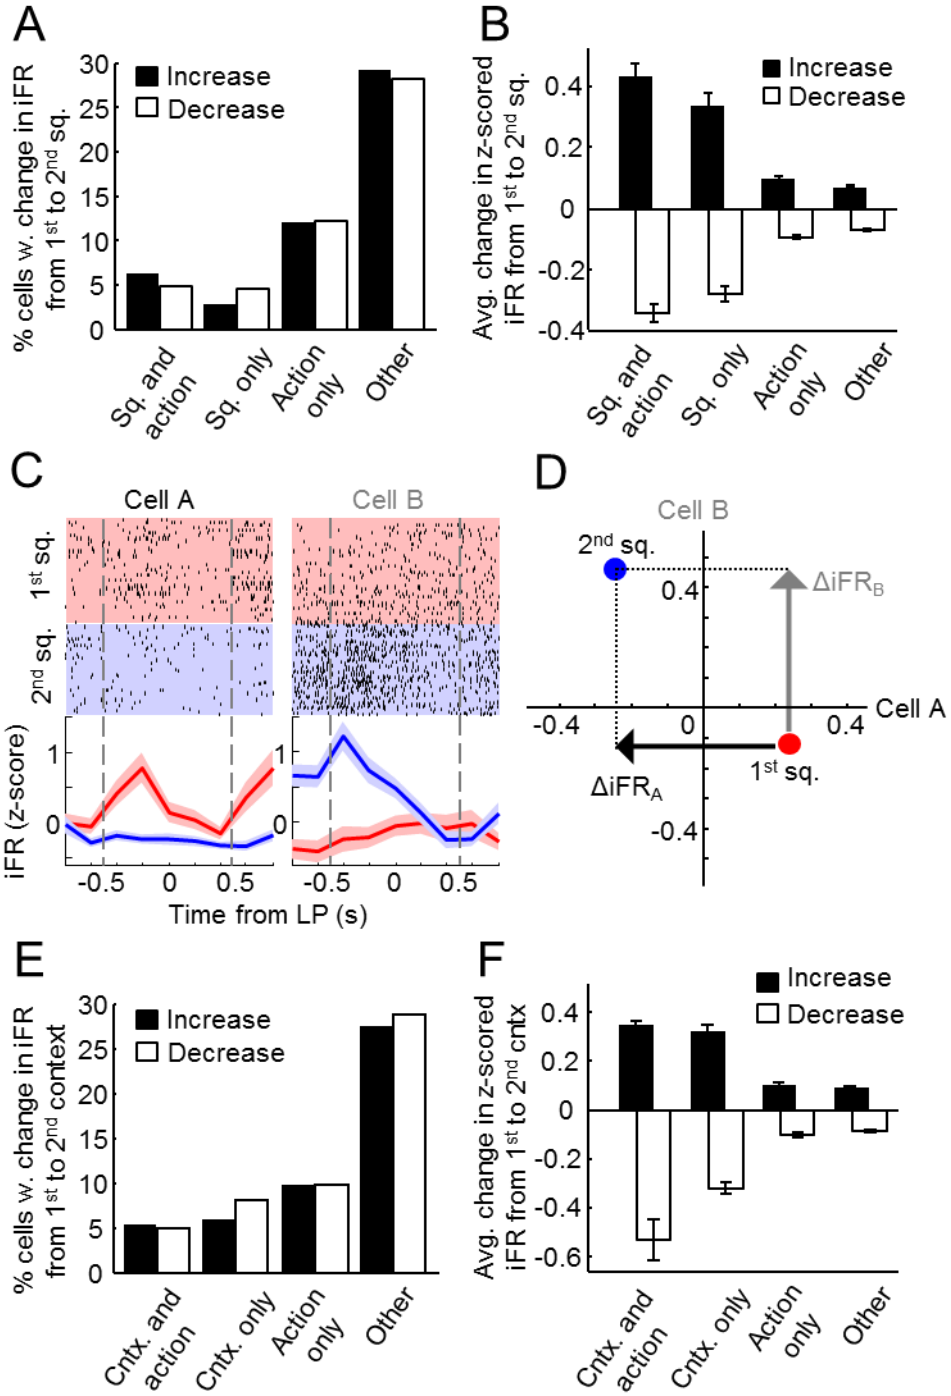
A
% cells w. change in iFR from 1st to 2nd sq.
Increase
Decrease
Sq. and action
Sq. only
Action only
Other
B
Avg. change in z-scored iFR from 1st to 2nd sq.
C
Cell A
Cell B
1st sq.
2nd sq.
iFR (z-score)
Time from LP (s)
D
Cell B
2nd sq.
ΔiFR_B
Cell A
1st sq.
ΔiFR_A
E
% cells w. change in iFR from 1st to 2nd context
Cntx. and action
Cntx. only
F
Avg. change in z-scored iFR from 1st to 2nd cntx

Figure 2.6 Balanced remapping of action-responsive neurons

When sequence or context change, each functional category of mPFC cells displayed remapping of action sensitivity that were well balanced both in number and in strength. **A)** When sequence was switched, approximately half of each type cells displayed increase in firing rates (black bars) while the other half had a decrease in activity (white bars). This was true to each of all cell groups including sequence-action interactive cells sequence-only cells, action-only cells and all the other cells. **B)** Within each cell type, the cells that increased their firing rate did so by a similar amount as those that decreased their firing rate. Together this observation across all cell groups amounted to the lack of change in overall level of activity shown in **Fig. 2.4C**, and indicated that the cells were perfectly balanced in the change in their level of activity as their sensitivity reorganized during sequence switch. Although the sequence-sensitive cells accounted for a relatively small (~17%) proportion of the entire population, they underwent a significantly greater change than the action-only and other cells during sequence-switch. **C)** The rasters and line plots of a pair of cells that balanced each other in the amount of activity reduction (Cell A, left panel) and enhancement (Cell B, right panel) from the 1st sequence (red) to the 2nd sequence (blue). Gray dashed lines mark the 1-s periods of interest in all analysis. **D)** A schematic illustrating the effective coding of sequence information in the mPFC ensemble with zero change in overall activity during sequence switch. In an ensemble that consists only of a pair of well-balanced cells, Cell A and B shown in **C)**, the activity state of the responses in the 1st sequence is defined by the activities of both cells during these periods (red dot). When sequence switched, Cell A's activity decreased by an amount (gray arrow) similar to the amount of increase in Cell B's firing rate (black arrow). Although total change in activity is close to zero as the 2 cells' changes could cancel out each other, the ensemble activity state has changed to represent the 2nd sequence (blue dot). **E)** When context was switched, approximately half of each type cells displayed increase in firing rates (black bars) while the other half had a decrease in activity (white bars). This was true to each of all cell groups including context-action interactive cells context-only cells, action-only cells and all the other cells. **F)** Within each cell type, the cells that increased their firing rate did so by a similar amount as those that decreased their firing rate. Together this observation across all cell groups amounted to the lack of change in overall level of activity shown in **Fig. 2.3C**, and indicated that the cells were perfectly balanced in the change in their level of activity as their sensitivity reorganized during sequence switch. Accounted for a relatively small (~24%) proportion of the entire population, the context-sensitive cells underwent a significantly greater change than the action-only and other cells during context-switch.

Chapter 3: Tracking progress in corticostriatal ensembles

3.1 Introduction

When performing a sequence of actions, one must continually track progress in order to know which actions are still required to attain the goal. As actions progress in the sequence there is a growing anticipation or expectancy about receiving the reward. This anticipation can be dissociated from the actions themselves and helps one stay on track in the face of distractions or errors. The present study investigated the neural correlates of sequence progression simultaneously in the anterior cingulate cortex (ACC) and dorsal striatum (DS).

Actions are encoded in numerous brain regions, however circuits involving the frontal cortex and striatum play a particularly important role in the flexible encoding of action sequences. In the DS neurons with strong action encoding tend to fire when the rat is performing multiple actions, but only when the actions occur within a specific sequence and not when they occur in isolation (Aldridge and Berridge, 1998, Schmitzer-Torbert and Redish, 2004). Most frontal regions contain neurons that selectively encode the serial position of each action in a sequence (Barone and Joseph, 1989, Clower and Alexander, 1998, Procyk et al., 2000, Ninokura et al., 2004, Ryou and Wilson, 2004, Mushiake et al., 2006, Averbeck and Lee, 2007, Berdyyeva and Olson, 2010) , with a particularly interesting form occurring in the ACC. Specifically, some ACC neurons fire progressively more following the completion of each action or subtask that brings the animal closer to a goal (Shidara and Richmond, 2002, Toda et al., 2012). This progressive increase in activity is referred to as a 'reward expectancy' signal.

A reward expectancy signal must be both highly dynamic and integrative since at any point in time the degree of expectancy depends on what actions have been performed as well as actions yet to be performed. A reward expectancy signal is also likely to be critical for proposed functions of the ACC centered on reward processing (Pratt and Mizumori, 2001, Matsumoto et al., 2003, Rushworth et al., 2011, Cowen et al., 2012) or assigning value to actions (Kennerley et al., 2009). In parallel, dopamine signals in the striatum may be related to sustained motivation towards the final goal (Howe et al., 2013). Reward expectancy representations in the ACC and possibly the DS could ultimately be used by the brain to track progress and could help keep the organism stay focused on task at hand in the face of set-backs or distractions.

To gain further insights into the neural basis of reward expectancy signals occurring during action sequencing, ensembles of neurons were recorded simultaneously from the ACC and the portion of the DS receiving afferents from the ACC (Sesack et al., 1989) while rats performed different sequences of actions that lead to a food reward. The task was constructed such that the first, second and third actions relative to a goal were performed on different physical levers in different action sequences. As a result it was possible to disambiguate signals related to the encoding of relative progress towards a goal versus signals related to pressing specific physical levers or moving to specific locations. Using this task, neural responses to unique serial positions were discovered in both the ACC and DS but their properties differed in important ways.

3.2 Material and methods

3.2.1 Subjects

Four experimentally naïve male Long-Evans rats (450-550g) were housed in a facility with 12hr light-dark cycle, with all training and recording taking place during the light cycle. For the duration of the behavioral experiments, the rats were food-restricted to just below 90% of their free-feeding weights. Feeding took place in the home cage after their daily training/recording sessions, and water was available *ad libitum* in the cages at all times. All procedures were carried out in accordance with the Canadian Council of Animal Care and the Animal Care Committee at the University of British Columbia.

3.2.2 Apparatus

Within a large opaque Plexiglas box (25”×18”), a main panel was installed with 3 levers designated Lever1 to 3 from right to left (**Fig. 3.1A**). On any given day of sequence training, a unique tactile object (Velcro, cardboard or soft foam) was stuck to the lever panel (but not on the lever itself) and the area on the floor immediately in front of the lever panel to symbolize the order in which the 3 levers should be pressed. An area of 25”×13” was left for the rat to move freely. On the opposing-side wall, a food cup was located at the centre, with each delivery of reward accompanied by a pure tone. Retractable levers and pellet dispenser were controlled and recorded with a PC via a Med Associate interface system (St.Albans, VT, USA).

3.2.3 Behavioral task

The naïve subjects were first trained on an FR1 schedule to press each of the 3 levers. A minimum of 60 presses within 0.5hrs, with no less than 15 presses on each lever was required

prior to the rat moving on to the next stage of training. After 3-5d of FR1 training, the rats learned three 3-lever sequences: Sequence A, B and C, in 3 consecutive stages of training. The order of lever presses in each sequence was given by tactile objects placed on the panel and floor in front of the levers. For a given animal, each object consistently designated a single serial position. The order of lever presses in Sequence A was Right Lever→Middle Lever→Left Lever (**Fig. 3.1B**). A lever retracted only when it was pressed in the correct order and remained extended in the event of an error. Therefore the level of performance on the right lever in Sequence A was necessarily 100%, whereas the chance level of performance on the middle lever was 50% (i.e. incorrect if pressed prior to the right lever) and that on the left lever was 33%. For training on each sequence, the percentage of correct responses on the 3rd item of the sequence had to reach 75% before moving on to the next stage of training. Sequence B consisted of 3 lever-presses in the order of Middle Lever→Left Lever→Right Lever, and Sequence C in the order of Left Lever→Right Lever→Middle Lever. At any one of the 3 stages of single-sequence training, if after 3 days of training, the animal still hadn't reached criterion and if day-to-day improvement stopped, a delay-punishment protocol was introduced to extinguish errors made on the 3rd lever of the given sequence. Specifically, if the 3rd lever was pressed before the 1st lever, all levers retracted and a 10-s time-out period ensued. This training continued until the animal reached criterion performance. When the criterion had reached criteria on all 3 sequences, the rat was surgically implanted and allowed 10d to recover. After recovery, 2-3 refresher sessions on each sequence were given before the first multi-sequence block test day.

On the multi-sequence block test days, the animals had to perform a block of at least 10 trials on each sequence at or above criterion, before switching to the next sequence in one of

three possible pseudorandom orders: Sequence 3→Sequence 2→Sequence 1, Sequence 2→Sequence 1→ Sequence 3, or Sequence 1→Sequence 3→Sequence 2. In-between sequence blocks, the animals were taken out of the box to allow for rearrangement of the tactile objects.

3.2.4 Surgery

Stereotaxic surgeries were performed with sterilized-tip procedures under anesthesia by isoflurane. NSAIDs analgesic, antibiotic, and a local anesthetic, were given before incision. One elliptical-shaped craniotomy was made centered at: AP: +3.2mm, ML: +1.0mm, and another craniotomy was made centered at AP: +1.2mm and ML: +3mm (Paxinos and Watson, 2005). Once the dura mater was retracted, the bottoms of the two bundles of 8 30-gauge tubes, containing a total of 16 tetrodes, were placed on the cortical surface. The bundles were of cylindrical shape with a bottom radius of ~0.4mm, and were angled medially by ~15 degrees. The implants were fixed with bone screws and dental acrylic. At the end of the surgery, tetrodes in the anterior bundle were extended by ~1.4mm into the brain to enter the anterior cingulate cortex (ACC), and tetrodes in the posterior bundle were extended by ~3mm to enter the dorsal striatum (DS) (Paxinos and Watson, 2005). Animals were given 10d to recover. Prior to each recording session, small adjustments were made with the hyperdrives to maximize the number of neurons recorded.

3.2.5 Acquisition of electrophysiological data

Data acquisition and offline spike sorting were conducted using the same methods, equipment and software as described in Chapter 2 and previously published in Hyman et al. (2012).

3.2.6 Histology

At the end of the studies, the animals were deeply anesthetized using urethane i.p. injection, and 100μA electrical current was passed through the electrodes for 30s. Animals were then perfused with a solution containing 250ml 10% buffered formalin, 10ml glacial acetic acid, and 10 g of potassium ferrocyanide. This solution causes a Prussian blue reaction, which marks with blue the location of the iron particles deposited by the electrode lesion. The brains were then removed and stored in a 10% buffered formalin/20% sucrose solution for at least 1 week, before being sliced and mounted to determine precise electrode locations. Since multiple sessions were recorded from individual animals the precise recording locations could not be derived from electrode lesions, but all electrode tracks were inferred between the entrance point and the dyed spots. **Fig. 3.1E** and **F** show representative recording sites for ACC and DS, respectively.

3.2.7 Data analyses

A total of 33 large ensembles (DS: N_{min}=19, ACC: N_{min}=21) were collected. Neurons firing less than 0.14 Hz were excluded from further analysis, because the sample of spikes was too small (250 or less) to be reliably representative of the cell's activity in relation to behavior. To obtain an estimate of the neural firing rate for each isolated cell i as a function of time bin *t*, $r_i(t)$, firing rates (FRs) were calculated for each spike train in each 50-ms bin (note FR was different from instantaneous firing rate (iFR) used in our past studies; Durstewitz et al 2010; Hyman et al 2012). In each trial of the task, each behavioral epoch—a lever-press or reward-approach period—includes a 1-s period centered at the timestamp of that behavior. Thus each epoch is 20 bins long, and 80 bins from each trial were analyzed, including 3 lever-presses

and a reward-approach period. Further analyses were performed using custom-made MatLab (Mathworks Inc., Natick, MA, USA) algorithms.

Principal Component Analysis (PCA): In the behavioral task, because each trial contained 4 epochs (3 actions plus reward-approach period), each of which lasted for 1s. Thus each trial contained 80 bins. Each cell's normalized FRs were averaged across all trials to produce a single FR vector containing 80 numbers. Thus in a matrix containing FRs from all cells recorded, there are 80 variables, for each of which there are 637 observations (i.e. total cell number) for the ACC and 351 for the DS. Principal component analysis (PCA) was performed on these 80 variables. The top PCs represent the most prevalent firing patterns among all cells. For each PC, each cell has a unique coefficient or loading, indicating the extent to which this PC represents this cell's firing pattern. Ranking cells based on their loadings resulted in 2 groups of cells: those with positive loadings and those with negative loadings. The averaged firing rates of these 2 groups analyzed with 2-way ANOVA using a 2 (# of groups) -by-3 (# of actions) design. *Post-hoc* Tukey's test was then used to test for differences among groups.

Correlation analysis: In order to determine whether the 'ramping' and 'action-linked' patterns exist within individual ACC and DS cells, two models were constructed to capture these features. The correlations between single-unit activities and the models were examined. The 'ramping' model was a 60-bin vector containing numbers from 0 to 1, stepped at 1/59. The 'action-linked' model was a 60-bin vector containing numbers from 0 to 1 stepped at 1/9, from the 1^{st} to the 10^{th} bin, from the 21^{st} to the 30^{th} bin, as well as from the 41^{st} and the 50^{th} bin (red line, **Fig. 3.4A**). The 10^{th}, 30^{th} and 50^{th} bin where when the 3 lever presses respectively occurred. Additionally, from the 11^{th} to the 20^{th} bin, from the 31^{st} to the 50^{th} bin, as well as from the 51^{st} and the 60^{th} bin the vector contained numbers from 1 to 0 stepped at -

1/9, thereby forming a saw-tooth shape (green line, **Fig. 3.4A**). Spearman's correlation and p-values were calculated between each cell's averaged FRs and these models ($\alpha<0.05$). Bonferroni correction was not performed, because the purpose of this analysis was not to test significance, but to categorize cells and examine their collective action.

Error analysis: There were two types of error trials observed in almost all sessions. In the first case, the animal started a trial by pressing the wrong lever for that sequence block, then went on to press the 3 levers in the correct order. In the second case, the animal responded correctly on the 1st lever for that sequence block, then pressed on the incorrect lever (i.e. pressed on the lever which should have been the 3rd item in the sequence rather than the 2nd), and then went on to complete the trial correctly. The activities of ramping model-correlated ACC cells and action-linked model-correlated DS cells during these 2 types of error trials were averaged and compared with those recorded during the error-free trials, one type at a time. The averaged firing rates of these 2 groups analyzed with 2-way ANOVA using a 2 (error trials vs. error-free trials) -by-3 (# of actions) design. *Post-hoc* Tukey's test was then used to test for difference among groups.

3.3 Results

Rats were trained in an experimental apparatus which contained a panel with three retractable levers located on one wall (**Fig. 3.1A**). Each lever was distinguished by cues (velcro, cardboard or foam) temporarily affixed to the area immediately surrounding the levers (but not the levers themselves). For any given sequence block, each cue indicated the serial position in which individual levers had to be pressed in order to obtain food reward. This sequence of cues was always the same for a given rat, but the cues were moved to different lever locations for each of the two or three sequence blocks (preliminary testing revealed that

the cues were neutral as neither the rats nor the recorded neurons systematically preferred one cue; **Fig. 3.1B**). On the multi-sequence block test days, behavioral performance within and across sequence blocks did not differ (within sequences: t_{1102}=0.89, p=0.38, **Fig. 3.1C**; across sequences: $F_{2,48}$=0.22, p=0.80, **Fig. 3.1D**).

The first serial position was investigated in isolation by combining all responses on the right lever from sequence block A, the middle lever from sequence block B and left lever from sequence block C and so on for second and third serial positions. At the end of testing, the animals were euthanized and recording sites in ACC (**Fig. 3.1E**) and DS (**Fig. 3.1F**) were located by histology.

3.3.1 Reward expectancy/proximity signals in ACC ensembles

Single neurons in the medial prefrontal cortex (mPFC), including the ACC, can exhibit virtually any type of response profile (Jung et al., 1998, Baeg et al., 2003, Hyman et al., 2005, Lapish et al., 2008, Hyman et al., 2010, Rigotti et al., 2013), therefore it was important to examine the overall signal emitted by large groups of neurons in order to reveal what specific aspects of our sequence task dominated this signal. Accordingly, Principal Component Analysis (PCA) was used to identify the main sources of firing rate variance across all 637 ACC neurons recorded in this study. When lever presses were organized according to serial position across the three sequences, the first PC, which accounted for 17.7% of total firing rate variance, exhibited a smooth progressive increase consistently across the three lever presses prior to declining steeply following the third lever response as the animal approached the goal port (**Fig. 3.2A**). The eigenvector values of PC1 remained positive throughout the three lever presses, indicating that each of the ACC neurons in this network displayed similar changes in firing rates (excitation or inhibition) for all three responses and

the magnitude of these changes increased with each successive lever press. The second PC (accounting for 9.46% of variance) increased slightly across the three lever presses, but exhibited an abrupt acceleration during goal approach at the point where PC1 began to decline (**Fig. 3.2B**). PC2 started trials with negative eigenvector values that suddenly shifted into positive values signifying that neurons in this network had opposing responses during the lever presses and goal approach periods (i.e. from inhibition to excitation or vice-versa).

Since neurons can load positively or negatively on a given PC, their loadings on PC1 were simply rank-ordered, and the group average of all neurons with positive loadings was plotted (left inset, **Fig. 3.2C**) and so was the group average of all neurons with negative loadings (right inset, **Fig. 3.2C**). Neurons with positive loadings increased their overall firing rate smoothly across the three serial positions whereas the firing rate of neurons with negative loadings decreased smoothly, as one would predict. The average aggregate firing rates of all neurons that loaded positively on PC1 exhibited significantly higher firing rates for the second serial position relative to the first (Tukey's *post-hoc* test: $p=4\times10^{-12}$) and significantly higher firing rates for the third serial position relative to the second (Tukey's *post-hoc* test: $p=7\times10^{-6}$, **Fig. 3.2D, left**). The neurons with negative loadings on PC1 exhibited a similar serial position linked change in response magnitude but in the opposite direction (Tukey's *post-hoc* test: 1^{st} action>2^{nd} action: $p=4\times10^{-5}$, 2^{nd} action marginally higher than 3^{rd} action: $p=0.07$, **Fig. 3.2D, right**). PCA therefore provided a means with which to group the firing activities of all 637 neurons based on the most dominant pattern of firing rate variance (PC1), but by so doing revealed a clear, robust and relatively uniform activity pattern during the performance of a sequence of actions that led to a reward. An example of a neuron with strong positive loadings on both PC1 and PC2 is shown in **Fig. 3.3A**. This neuron exhibited both the smooth increase

in firing consistent with PC1 and also a further increase that emerged abruptly during reward approach consistent with PC2. To demonstrate the prevalence of these activity patterns, two more examples of cells with strong positive loadings on PC1 are shown in **Fig. 3.3B, C. Fig. 3.3D, E** and **F** show examples of neurons loaded negatively on PC1. The strong correlation between the responses of these neurons with the patterns identified by both of the top PCs illustrates the ability of PCA to detect neurons that contribute to multiple distinct network signals during different cognitive/behavioral epochs.

3.3.2 The reward expectancy/proximity signal is produced by a smooth firing rate progression in ACC neurons

The smooth progression in firing rate revealed by PCA and shown in the overall average firing rates could be the product of two distinct patterns of spiking activity across the network. It is possible that individual neurons' firing rates ramp smoothly over the course of the three serial positions, or it is also possible that this smooth progression was produced by the combined activity of neurons that fired to each action stochastically, but asynchronously. In order to determine whether one or both response types were present in the ACC, correlations between single-unit activities and models designed to capture these features (**Fig. 3.4A**) were calculated. A large group of neurons correlated positively the 'ramp' factor (n=107, **Fig. 3.4B**) while a smaller group exhibited significant positive correlation with the 'action-linked' factor (n=37, **Fig. 3.4C**). Although these 37 neurons fired in conjunction with each action, on average their firing rate did not progress with each subsequent action. However, 11/637 neurons had firing activities which were correlated with both factors. These neurons exhibited an action linked responses as well as a progressive increase for later actions (**Fig. 3.4D**). Finally, neurons without any correlation between average firing activity and

either factor are shown in **Fig. 3.4E** (n=376). These neurons exhibited relatively flat firing rates throughout the trial. Given that this analysis categorized all neurons exhaustively, one can conclude that the ramping activity, captured by PC1 and PC2, was produced mainly by a network of neurons that themselves exhibited a ramping pattern.

So far neurons were categorized based on their responses over an entire trial, but it is also informative to illustrate how neurons responded specifically during each lever press period. **Fig. 3.5A** shows, for a single example session, the average change in firing rates (relative to their individual session-wide means) during a 1s period surrounding presses on the same physical lever, grouped according to whether the lever was pressed as the 1^{st}, 2^{nd} or 3^{rd} action in a sequence. The patterns of activity in the three cases were quite similar because the action and the physical lever was the same. Yet it is noteworthy that the pattern became more differentiated for each subsequent press because neurons that fired above their average rate in response to presses on a given lever fired more for later serial positions whereas neurons that fired below their average rate in response to presses on that lever fired progressively less for later serial positions. This illustrates that progression towards a goal tended to enhance the distinctness of action-linked activity state patterns.

3.3.3 The reward expectation signal ramps only in association with correct choices

If the firing rate progression associated with movement towards a goal served as a type of reward expectancy (Shidara and Richmond, 2002) or reward proximity signal, it was unclear whether it should scale equally for all actions or only for correct actions that actually brought the animal closer to the goal. In the present study, well-trained rats would sometimes press the incorrect lever as their first or second choice of a trial (there could be no errors on the 3^{rd} action in a trial because there was only one remaining lever as the levers retracted after

each correct press). The effects of these two types of errors on the firing rates of the 107 ramping neurons shown in **Fig. 3.4B** were analyzed in detail below.

When an error was committed on the first lever press, the overall average firing rates of the neurons during the next correct lever press was the same size as during the first correct lever press on error-free trials (2-way ANOVA main effect of action: $F_{2,114}$=108.3, p=5×10^{-27}, but not trial type: $F_{1,114}$=2.94, p=0.08; Tukey's *post-hoc* test: 1st correct actions in error trials vs. those in error-free trials: p=0.95) yet was significantly different from the second correct lever press on error-free trials (Tukey's *post-hoc* test: p=0.0002, **Fig. 3.6A, B**). Likewise, for trials where an error occurred on the second lever press, the firing rate on the next correct lever press was the same size as it was on the 2nd lever press for error-free trials (2-way ANOVA main effect of action: $F_{2,114}$=134.2, p=1×10^{-30}, but not trial type: $F_{1,114}$=1.13, p=0.28; Tukey's *post-hoc* test: p=0.87) and significantly different from what it was during the 3rd lever press on error free trials (Tukey's *post-hoc* test: p=5.5×10^{-11}) (**Fig. 3.6C, D**). To illustrate this visually the average response of the neurons on error-free trials were overlaid with their average responses on error trials, assuming an error had not been committed. Specifically, the average response on error-free trials was shifted by one lever press and then overlaid (red line, **Fig. 3.6A**) with trials where the first lever press was incorrect (orange line, **Fig. 3.6A**) or was cut after the first lever press and then shifted by one lever press and overlaid (red lines, **Fig. 3.6C**) with trials where the second action was incorrect (purple line, **Fig. 3.6C**). The firing rates were not altered or scaled in this plotting scheme. From these figures, it is clear that only correct choices initiated the firing rate progression (**Fig. 3.6A**) and once initiated it appeared to progress only in association with correct choices (**Fig. 3.6C**). Thus the

smooth ramping signal in the ACC appears to track actual progress towards a goal rather than simply counting actions.

Next, activities arising during different portions of the error trials were examined. Errors committed prior to the 1st correct action could involve presses on the lever that was supposed to be the 2nd or the 3rd correct action in the sequence block. Behavioral analysis revealed that the rats were more likely to press the lever associated with the 3rd correct action than the one associated with the 2nd correct action (Wilcoxon ranksum test: ranksum=257.5, p=0.00086). Given the abundance of trials where an initial error involved the lever that was supposed to be the 3rd correct action in the sequence block, it was explored whether ACC activity during these types of errors resembled activity during the 1st correct action or the 3rd correct action. It was found that activity during these errors were actually significantly different from both (1-way ANOVA, $F_{2,59}$=53.42, p=8.5×10^{-14}). Specifically, these errors were associated with higher levels of activity than the 1st correct response (Tukey's HSD test: error>1st correct, p=0.00012), but with lower activity than the 3rd correct response (Tukey HSD test: error <3rd correct: p=0.0016). This later effect is exactly consistent with what one would expect for a system encoding serial position. The former effect arose because the first correct response tended to be below baseline (Fig 4b) yet when an error was committed firing was essentially at baseline levels. As a result, the firing was higher when the first response was an error than when it was correct.

3.3.4 Firing rate progression in DS neurons is tightly tied to actions

The same techniques used to examine ACC ensembles were applied to DS ensembles recorded simultaneously from the same animals. PCA performed on all DS neurons revealed that the first PC (**Fig. 3.7A**) identified a group of DS cells whose activity varied closely in

association with the lever presses themselves. The second PC (**Fig. 3.7B**) showed variation around each lever press, but like PC2 in the ACC also exhibited a pronounced acceleration during reward approach. The firing rates of neurons that loaded positively versus negatively on PC1 were plotted separately and revealed that the change in firing rates for later serial positions occurred during the time of the actions themselves (**Fig. 3.7C**). Once again, all DS neurons were included in this figure but were simply split into the two subgroups based on their loadings on PC1 and shown in the insets. Positively loaded cells fired significantly higher during the 2^{nd} action than the 1^{st} action (Tukey's *post-hoc* test: p=0.007), but did not fire differently during the 3^{rd} action (Tukey's *post-hoc* test: 3^{rd} vs. 1^{st} action: p=0.93, 3^{rd} vs. 2^{nd} action: p=0.10; **Fig. 3.7D**, left). Negatively loaded cells responded similar to all 3 actions (**Fig. 3.7D**, right). **Fig. 3.7E, F** show examples of neurons loaded strongly on PC1 and PC2, respectively.

Although this analysis suggested that DS activity was tightly linked to the lever presses, it remains a possibility that neurons similar to those ACC neurons that ramped more smoothly were buried within the DS population. In an attempt to extract these neurons, the correlation between firing rates of all DS neurons and the 'ramp' and 'action-linked' factors were examined and shown in **Fig. 3.4A**. In this case, more neurons were significantly positively correlated with the 'action-linked' factor (n=32, **Fig. 3.8A**) than the 'ramp' factor (n=19, **Fig. 3.8B**). Even though these 19 DS neurons were significantly correlated with the 'ramp' factor, when plotted it was evident that the degree of ramp-like activity was very weak (**Fig. 3.8B**). To further test this possibility, the DS ensemble activities were projected onto PC1 derived from ACC ensembles (i.e. the principal component associated with smooth ramping activity, Fig. 2A), and tested the difference between the absolute values of the

loadings of DS neurons and those of the ACC neurons on this PC. Absolute loadings were used because they reveal the amount of variance accounted for by the ACC PC in both ACC and DS neurons. The resulting loadings were significantly larger in the ACC neurons than in the DS neurons (independent-sample t-test, t_{986}=2.40, p= 0.0167). In other words, the smooth ramping pattern is indeed much more dominant in ACC than in DS. Finally, the firing rates of the DS neurons that failed to attain significance on either factor were also largely flat throughout the trial (n=211, **Fig. 3.8C**). Based on these analyses, it would appear that firing rate progression in the DS relative to the goal was weak but very tightly linked to the actions themselves. This conclusion based on PCA and correlation analysis of the entire ensemble was consistent with the firing properties of single neurons (**Fig. 3.7E**).

3.3.5 Errors do not alter serial position signalling in the DS

Given the relatively weak firing rate progression and the tight association between DS neuron activity and the actions themselves, it was unclear whether these neurons would fire differently to errors or whether they would continue to faithfully track all actions indiscriminately. Since in the DS the bulk of the progressive change in firing for subsequent actions occurred during the actions, for this analysis the focus was on the responses of the 32 DS neurons shown in **Fig. 8A** that exhibited a significant positive correlation to the 'action-linked' factor.

When an error was committed as the first lever press, the FR on the next correct lever press was not different from the first correct lever press on error-free trials (2-way ANOVA, main effect of action: $F_{2,114}$=4.91, p=0.01, but not trial type: $F_{1,114}$=0.74, p=0.39; Tukey's *post-hoc* test: 1st correct actions in error trials vs. those in error-free trials: p=0.99) and was also not different from the second correct lever press on error-free trials (Tukey's *post-hoc* test:

p=0.63) (**Fig 3.9A, B**). When the error occurred on the second lever press, the FR for the next correct lever press was the same size as the 2nd lever press on error free trials and was also not different from what it was on the 3rd lever press on error free trials (2-way ANOVA, no main effect of action: $F_{2,114}$=2.57, p=0.08, nor of trial type: $F_{1,114}$=0.01, p=0.91) (**Fig 3.9C, D**). This lack of differentiation between correct and error trials is evident when the average firing rates of the neurons on correct and error trials were overlaid (**Fig. 3.9A, C**). There was little difference in the DS signal on correct versus error trials because the DS neurons fired so strongly in association with errors. DS activity arising during different portions of the error trials themselves was also explored in greater detail. On the trials that began with the rat pressing the lever that should have been the 3rd-in the sequence, firing rates during these initial errors were not different from firing during either the 1st correct response (1-way ANOVA, $F_{2,59}$=4.01, p=0.024, Tukey's HSD test: p=0.088) or from when the lever was pressed correctly as the 3rd response in the sequence (Tukey's HSD test: p=0.88). These results are different from what was observed in the ACC but are fully consistent with a generic action-related change in firing rate for DS neurons.

3.4 Discussion

Previous studies have shown robust representations of various aspects of sequential behaviors by single neurons in the frontal cortex (Barone and Joseph, 1989, Clower and Alexander, 1998, Nakamura et al., 1998, Procyk et al., 2000, Ninokura et al., 2004, Ryou and Wilson, 2004, Fujii and Graybiel, 2005, Mushiake et al., 2006, Averbeck and Lee, 2007, Shima et al., 2007, Berdyyeva and Olson, 2010). In the present study, when trials were organized according to serial position of the lever presses relative to reward, the dominant pattern of activity accounting for the largest portion of variance across all recorded ACC

neurons was a smooth 'ramp-like' change in firing rate as animals progressed through the three actions to the reward. Firing rate progression appeared to occur only in association with correct actions and not errors. The ability to represent one's progress in a sequence relative to a goal could be expected in a brain region such as the ACC that is involved in monitoring actions and outcomes (Alexander and Brown, 2011). In contrast to the ACC, the largest portion of variance across all DS neurons, recorded simultaneously from the same animals, was a change in firing rate linked to a lever press that had a slight tendency to become stronger for lever presses occurring in later serial positions. Unlike ACC neurons, DS neurons tended to fire the same way regardless of whether the actions were correct or not suggesting that this region mainly encodes the actions.

It is becoming increasingly evident that neurons in the frontal cortex have multiple, diverse and dynamic firing rate correlates (Jung et al., 1998, Rigotti et al., 2013). Neurons in the rat medial prefrontal cortex (mPFC; ACC and prelimbic regions) exhibit potent responses to both actions (Jung et al., 1998, Hyman et al., 2005, Lapish et al., 2008, Durstewitz et al., 2010) and rewards (Pratt and Mizumori, 2001, Hyman et al., 2011, Caracheo et al., 2013, Horst and Laubach, 2013). mPFC neurons fire during reward approach with the strength of the response being dependent on the reward magnitude (Pratt and Mizumori, 2001, Kargo et al., 2007). Furthermore many mPFC neurons fire differently during identical action sequences if the sequences resulted in the delivery of rewards of a different magnitude (Kargo et al., 2007). In addition, modulation of mPFC neurons that respond in anticipation of a forthcoming reward was also based on the route taken to acquire the reward, suggesting that mPFC neurons multiplex information about actions with expectations of future rewards and their magnitude (Cowen et al., 2012). In the primate ACC, a subset of neurons fire in response to the

completion of different actions or subtasks but fired more robustly for later actions that were closer to the delivery of reward (Shidara and Richmond, 2002, Toda et al., 2012). The present results resemble those of Shidara & Richmond (2002) with the exception that only 2 of 106 neurons in their study showed a smooth ramp-like increase in firing. Instead most neurons fired upon lever release with further increases in firing rate for those lever releases performed later in the multi-trial schedule. In addition to obvious differences in species, one reason for these different profiles may be attributed to our use of self-paced and continuous responding rather than a series of discrete subtasks. In our case, neuronal responses seemed to decay more slowly after each action, resulting in a smoother overall response profile.

The continuous response profile most commonly observed in ACC in the present study, was similar to climbing or ramping activity previously reported for subsets of dorsolateral PFC neurons during the delay period of a memory task prior to a response, as well as in primate and rat ACC on tasks requiring interval timing (Niki and Watanabe, 1979, Narayanan and Laubach, 2009). Ramping activity has been proposed to reflect the neural processes associated with withholding a temporally inappropriate response (Narayanan et al., 2006, Narayanan and Laubach, 2009, Hyman et al., 2013), with timing of an interval (Durstewitz, 2003) or with the timing of rewards (Horst and Laubach, 2013). Hence ramping activity may be a general mechanism related to expectation and in the present case specifically reward expectation. Importantly, it was observed that ramping was absent during errors, suggesting that the expectation signal is not simply an action or time accumulator, but rather may process abstract knowledge about which actions actually move the animal forward towards the reward as compared to those actions that are ineffective. Furthermore, the firing rate progression did not reset to baseline levels following an error but rather resumed at a level where it would

have been had the error not occurred (**Fig. 3.6**). Accordingly, this pattern of firing rate appeared to maintain a running tally of the animal's location relative to the reward, a property that remained intact even when temporarily disrupted by an error. Collectively, these considerations suggest that while progressive changes in the firing activity of ACC neurons during approach to a goal could reflect expectation of a forthcoming reward as previously proposed (Shidara and Richmond, 2002), this change may also reflect an internalized representation of progress.

Previously the ACC and prelimbic neurons have been reported to represent discrete types of information using a distributed pattern based coding scheme (Lapish et al., 2008, Durstewitz et al., 2010, Hyman et al., 2012, Hyman et al., 2013). Evidence was also found for this type of code in the present study as illustrated in **Fig. 3.5** where the lever press was associated with a pattern of activity as indicated by increases or decreases in firing relative to each neuron's overall grand mean. Presses on individual levers were associated with different activity patterns much in the same way that distinct activity patterns effectively differentiated each unique epoch in a radial arm maze working memory task (Lapish et al., 2008), individual rules on a set-shifting task (Durstewitz et al., 2010), different stimuli and actions in an operant delayed alternation task (Hyman et al., 2013) as well as distinct environmental contexts (Hyman et al., 2012). This type of pattern-based coding scheme involving large groups of general purpose neurons may subserve the capacity of the frontal cortex to efficiently parse any arbitrary task and hence may be what endows this region of the brain with considerable flexibility. In contrast to such a patterned-based code, the reward expectancy signal is quite different. As shown in **Fig. 3.5**, the firing-rate pattern is similar for all three types of lever presses because the animal performs the identical action on the exact same physical lever.

However, what does vary is the magnitude of the differentiation in firing rates, which is progressively increased when that lever is pressed as the first, second or third action in the sequence. As a result, the pattern associated with presses on a given physical lever is progressively amplified as the animal approached the reward.

Although the mechanisms responsible for this phenomenon are unknown, dopamine levels in the mPFC increase in a manner that is directly related to the approach to a reward (Ahn and Phillips, 2002, Phillips et al., 2004, Rossetti and Carboni, 2005). Furthermore, based on biophysical data obtained by patch-clamp recordings in vitro, dopamine (as well as serotonin) could produce exactly this type of pattern-based amplification (Di Pietro and Seamans, 2011). Therefore, as noted in previous chapters, discrete declarative forms of information may be represented in the frontal cortex using a patterned-based coding scheme while motivation variables, such as the proximity to reward may act via neuromodulators to amplify or modify the strength of these activity state patterns.

The reward expectation signal in the DS was quite different from what was observed in the ACC. Even on this self-paced task, there was no evidence for the same type of continuous smooth ramp-like activity observed in ACC ensembles recorded simultaneously from the same animals. Instead, DS activity was tightly locked to the lever presses (PC1) and/or the reward approach (PC2). PC1 in DS ensembles was characterized by strong variation during the lever presses periods which increased slightly as the animal progressed through the sequence. While several neurons were correlated with the ramping factor, not all exhibited the progressive increase in firing for subsequent lever presses and any progression was lost in the overall average. Recently Howe et al. (2013) showed that the extracellular dopamine level in the striatum ramps as rats run towards the goal in a T-maze. Yet in many cases, this effect did

not appear in the averaged firing rates of DS cells. The current study differs from that of Howe et al. (2013) in that it involved several operant actions and it remains to be seen whether DS dopamine would also ramp up smoothly in our task. Collectively, it seems that while DS neurons are sensitive to the sequence in which actions are performed, they do not seem to code progress within a sequence in the same dynamic and integrated manner as neurons in the ACC. Furthermore, DS neurons did not respond to correct lever presses as distinct from incorrect ones. When viewed together, the main signal present in ACC ensembles on this task is associated with tracking progression in a manner that is largely abstracted from the encoding of the actual lever presses, while the DS acts as the compliment, tracking lever presses in a more literal fashion. The integrated signal produced by these two interconnected regions would effectively keep the animal on track in its progression towards a goal, even in the face of errors, delays or distractions.

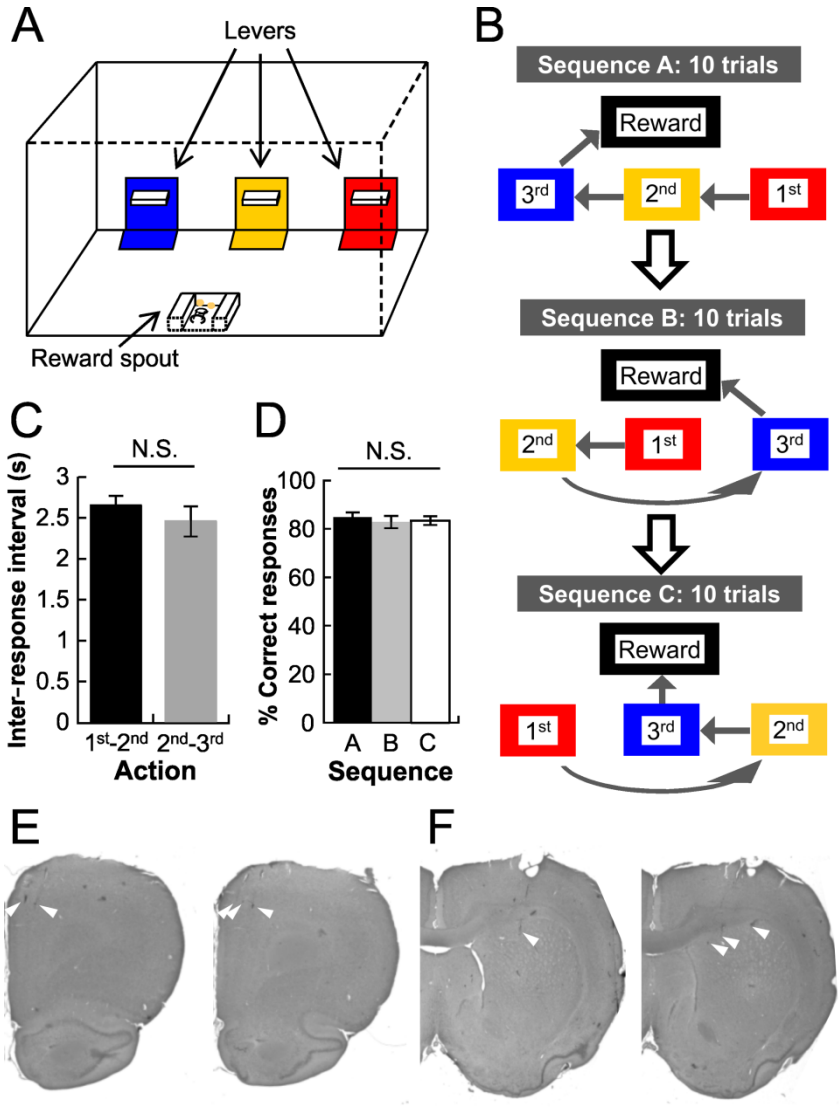
A
Levers
Reward spout
B
Sequence A: 10 trials
Reward
3rd
2nd
1st
Sequence B: 10 trials
Reward
2nd
1st
3rd
Sequence C: 10 trials
Reward
1st
3rd
2nd
C
N.S.
Inter-response interval (s)
3
2.5
2
1.5
1
0.5
0
1st-2nd
2nd-3rd
Action
D
N.S.
% Correct responses
100
80
60
40
20
0
A
B
C
Sequence
E
F

Figure 3.1 Task description and performance

A) The operant chamber contained 3 levers installed on the front panel and a food-cup on the opposing wall. A unique sensory cue was attached to the floor immediately in front of each lever as well as on the surrounding wall. The 1st lever (color-coded red in the schematic) to be pressed in a given sequence block could be surrounded by Velcro, the 2nd lever (yellow in schematic) by cardboard and the 3rd lever (blue in schematic) by soft foam. **B)** Example of a test day where the rat had to perform a minimum of 10 trials on each of the 3 sequence blocks, which were given in a pseudorandom order. Sequence block A required the rat to respond on the right lever, followed by the middle lever and then the left lever, before reward pellets were delivered to the food-cup on the opposite wall. The serial order of the 3 sensory cues remained constant for a given rat but they were moved to different levers for each of the sequence blocks. **C)** Latencies to travel between the 1st and the 2nd levers were equivalent to the latencies to travel between the 2nd and the 3rd levers across all sequence blocks. **D)** The animals' performance levels were equivalent across the 3 sequence blocks, as measured by the percentage of correct response on the 3rd lever of each sequence block. **E)** Histology showing representative electrode track endings (white arrows) in the ACC. **F)** Representative electrode track endings (white arrows) in the DS.

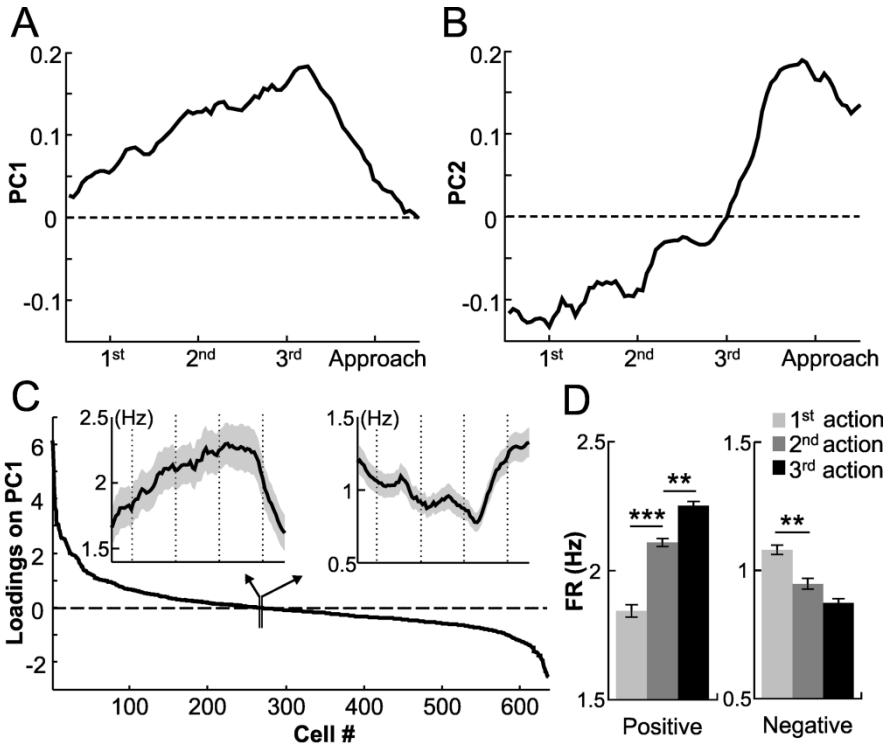

Figure 3.2 The main patterns of firing rate variance in ACC ensembles

PCA was performed on all neurons recorded across all animals and sessions. Trials were arranged such that the 1st, 2nd, 3rd lever presses and reward-approach period, each occupying a 1-s interval, were aligned across the three sequence blocks, regardless of which actual physical levers were pressed in each case. **A)** The PC1 eigenvector was characterized by a progressive increase during the lever presses followed by an abrupt decline during reward-approach. **B)** The PC2 eigenvector started negative and smoothly increased with an abrupt acceleration into positive values during reward-approach. **C)** Factor loadings for all 637 ACC neurons on PC1. *Left inset:* average firing rate of all neurons with positive factor loadings on PC1. *Right inset:* Average firing rate of all neurons with negative factor loadings on PC1. **D)** The average firing rates during the 1s lever press periods for all neurons shown in (C) with positive factor loadings on PC1 (left), or negative factor loadings on PC1 (right). The average firing of the neurons during lever presses in the first serial position are shown in light gray, the second serial position in dark gray and the third serial position in black. *** denotes significance at $p<1\times10^{-10}$, ** denotes significance at $p<0.0005$.

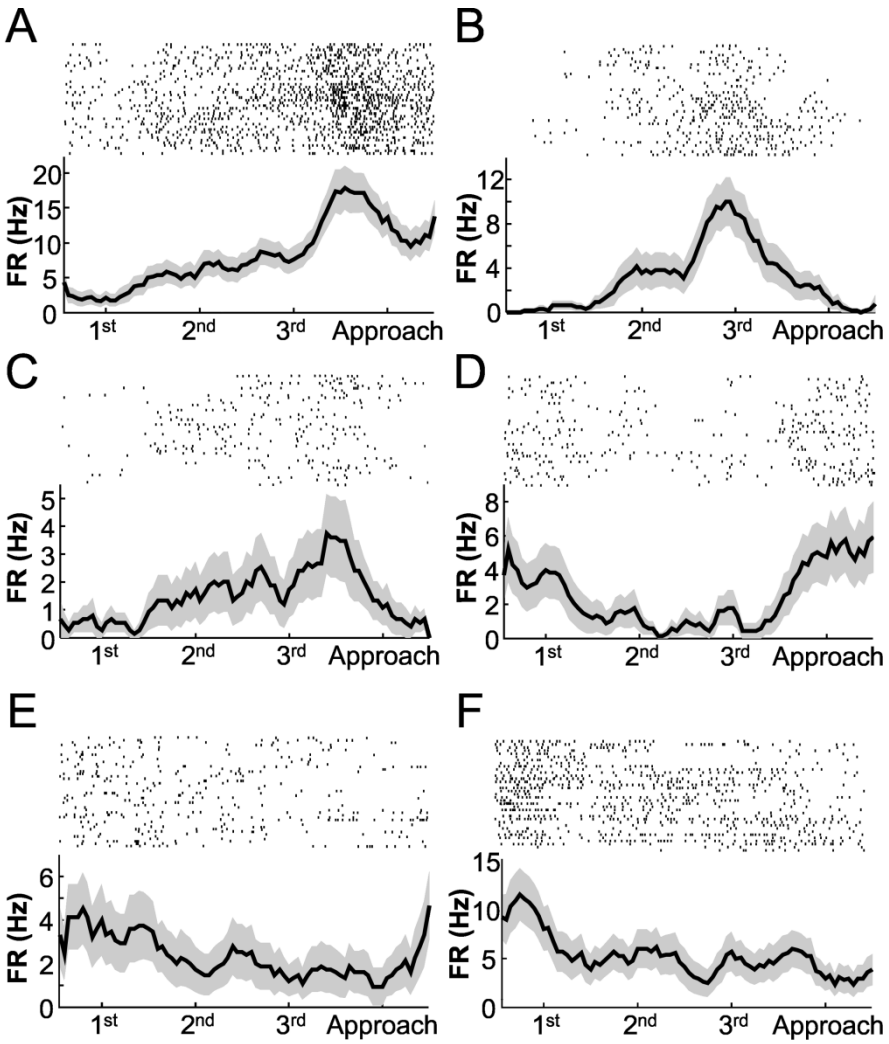

Figure 3

Figure 3.3 The progressive firing rate patterns exhibited by ACC neurons

A-C) Examples of single neurons strongly positively loaded on both PC1 and PC2 (A) or on PC1 only (B, C). Top: raster plot; bottom: average FR across all trials (mean ± s.e.m.). **D-F)** Examples of single neurons negatively loaded on PC1. Top: raster plot; bottom: all-trial average FR (mean ± s.e.m.).

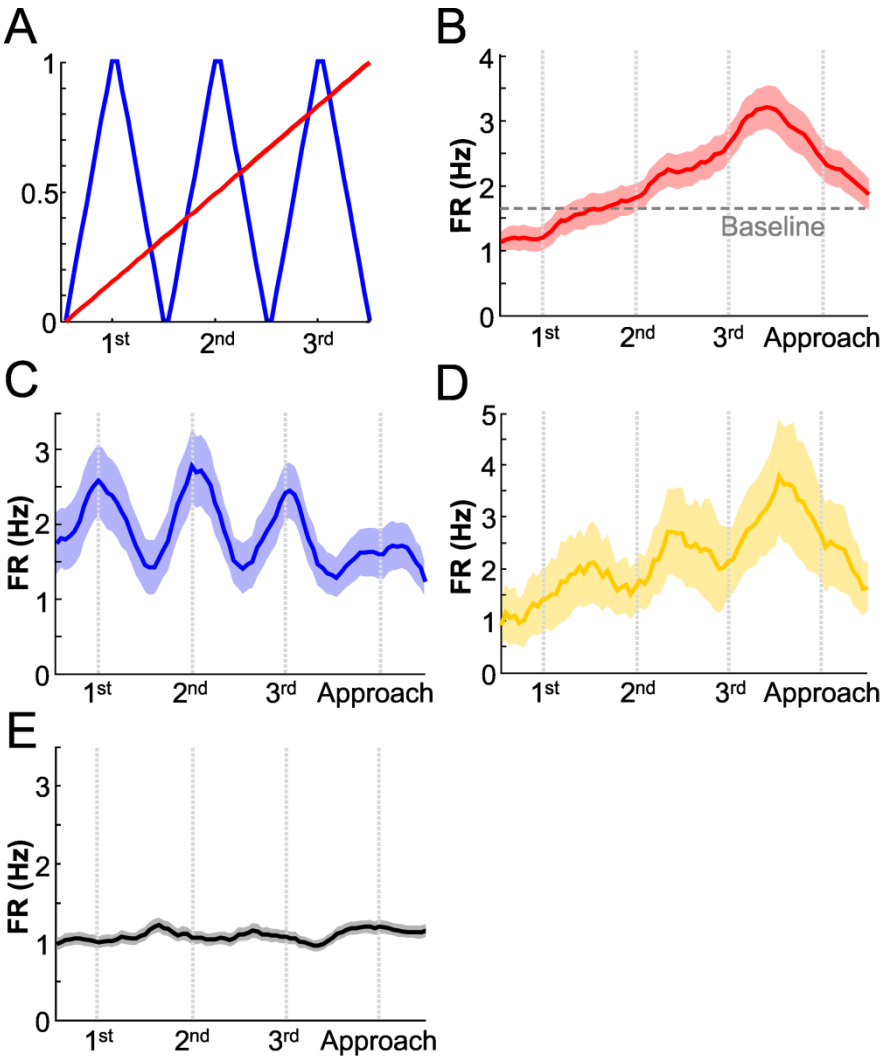

Figure 3.4 Possible single-unit firing patterns underlying the progressive firing-rate increase

Possible single-unit firing patterns that underlay the progressive firing-rate increase across the three actions leading to reward were tested. **A)** Two potential models of firing rate variance across a trial were a smooth progressive change (red) or a lever press-linked change (blue). **B)** Average firing rate for all 107 neurons showing a positive correlation with the ramping model (mean ± s.e.m.). The horizontal line represents the overall average baseline firing rate of the neurons. Note the firing during the first correct response was below baseline while subsequent responses rose above baseline. **C)** Average firing rate for all 37 neurons with a positive correlation with the action-linked model (mean ± s.e.m.). **D)** The average firing rate of the 11 neurons that attained significance on both factors. Note how these neurons exhibited transient action linked responses that increased for actions closer to the reward. **E)** The average firing rate for all neurons whose firing did not correlate with either model (n=376).

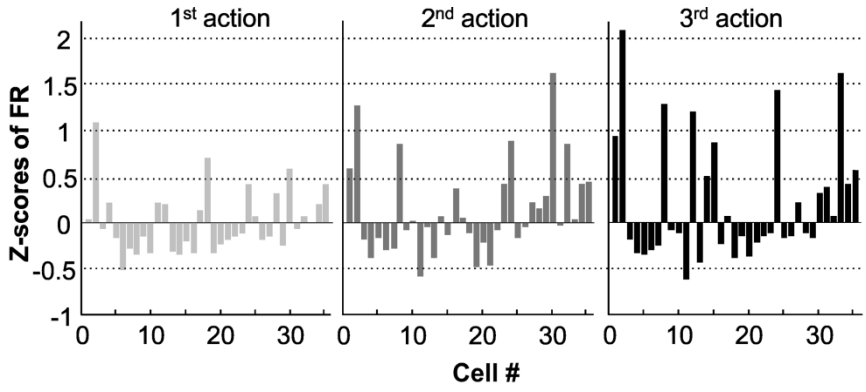

Figure 3.5 Prototypical activity state pattern in ACC ensembles

Prototypical activity state pattern observed in the ensemble in the 1s period surrounding presses on the right lever in a single session. Each bar represents the firing rate z-scores for a single neuron averaged across the 1s interval surrounding all presses on that lever within a sequence block. Average FR z-score values when the right lever was the 1st lever pressed in one sequence block (Left: light gray bars), the 2nd lever pressed in another sequence block (Middle: dark gray bars) or the 3rd lever in a third sequence block (Right: black bars). The divergence in z-scores was related to differences in the serial positions of the lever press relative to the reward since the physical lever was identical in each case.

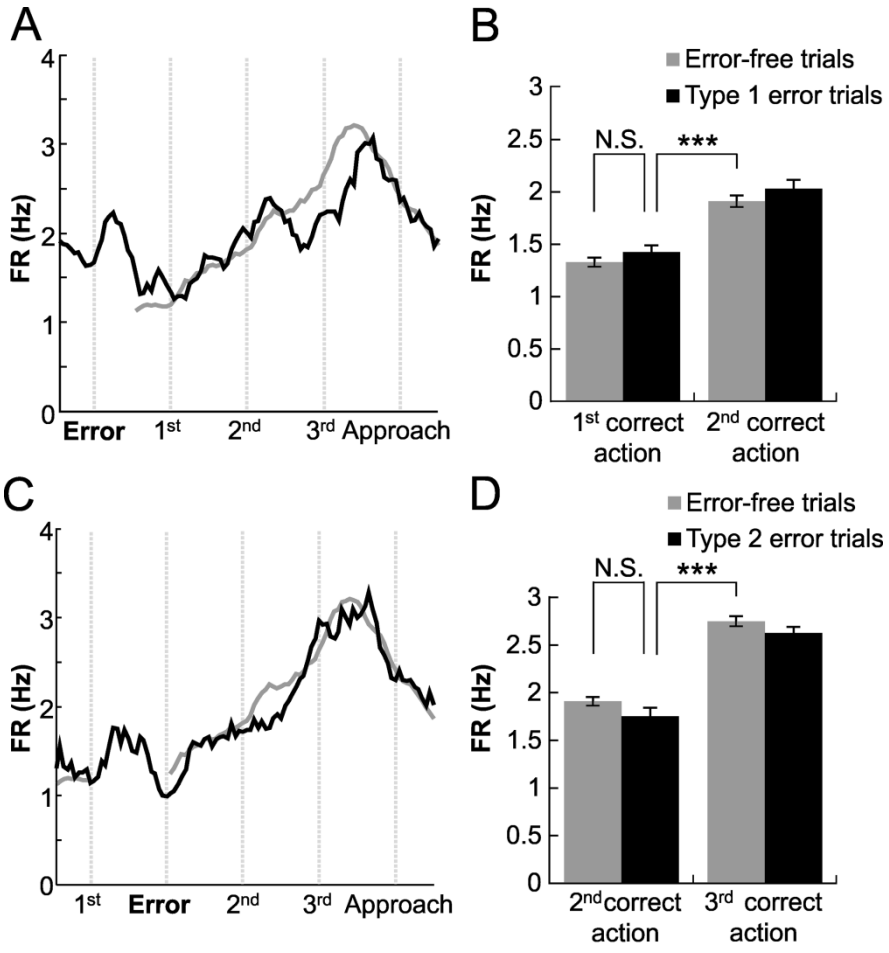
A
FR (Hz)
Error
1st
2nd
3rd Approach
B
Error-free trials
Type 1 error trials
N.S.
***
FR (Hz)
1st correct action
2nd correct action
C
FR (Hz)
1st
Error
2nd
3rd Approach
D
Error-free trials
Type 2 error trials
N.S.
***
FR (Hz)
2nd correct action
3rd correct action

Figure 3.6 Firing rate progression in the ACC during correct choices versus errors

A) Neurons' average firing rate on error-free trials shown in Fig. 4B was re-plotted in gray, with the average firing rate of the same neurons on trials starting with an error (Type 1 error trials, black line). Note that the firing rate of the neurons increased but then returned to baseline following the initial error. Once correct choices resumed, the firing rate progression was essentially identical to trials without errors. **B)** Statistical analysis of firing rate changes of these neurons on error-free trials (gray bars) versus Type 1 error trials (black bars). Left: no significant differences (n.s.) in average firing rates in the 1s period surrounding the first *correct* lever press, whether that lever press was preceded by an initial error (black bar) or not (gray bar). Right: while average firing rates for the first versus second *correct* lever presses were different, this was true regardless of whether the trial began with an error (black bars) or not (gray bars). **C)** As in (A), neurons' the average firing rate on error-free trials shown in **Fig. 3.4B** was re-plotted in gray, with the average firing rate of the same neurons on trials where the 1st correct choice was followed by an error are shown in gray (Type 2 error trials). Note that the firing rate of the neurons increased but then returned to baseline following the error. Once correct choices resumed, the firing rate progression recommenced at the same level as on trials without errors. **D)** Statistical analysis of firing rate changes of these neurons on error-free trials (gray bars) versus Type 2 error trials (black bars). Left: there were no significant differences (n.s.) in average firing rates during the first correct lever press, whether that lever press was followed by an error (black bar) or not (gray bar). Right: the average firing rates for the first versus second *correct* lever presses were different, and this was true regardless of whether or not an error occurred after the first correct choice (black bar) or not (gray bar). Since each subsequent lever press was associated with a significant change in average firing rate and this was true regardless of whether or not the trial contained an error, it suggests that the firing rate progression from one lever press to the next occurred only during correct choices. *** denotes significance at $p<1\times10^{-10}$, ** denotes significance at $p<0.0005$.

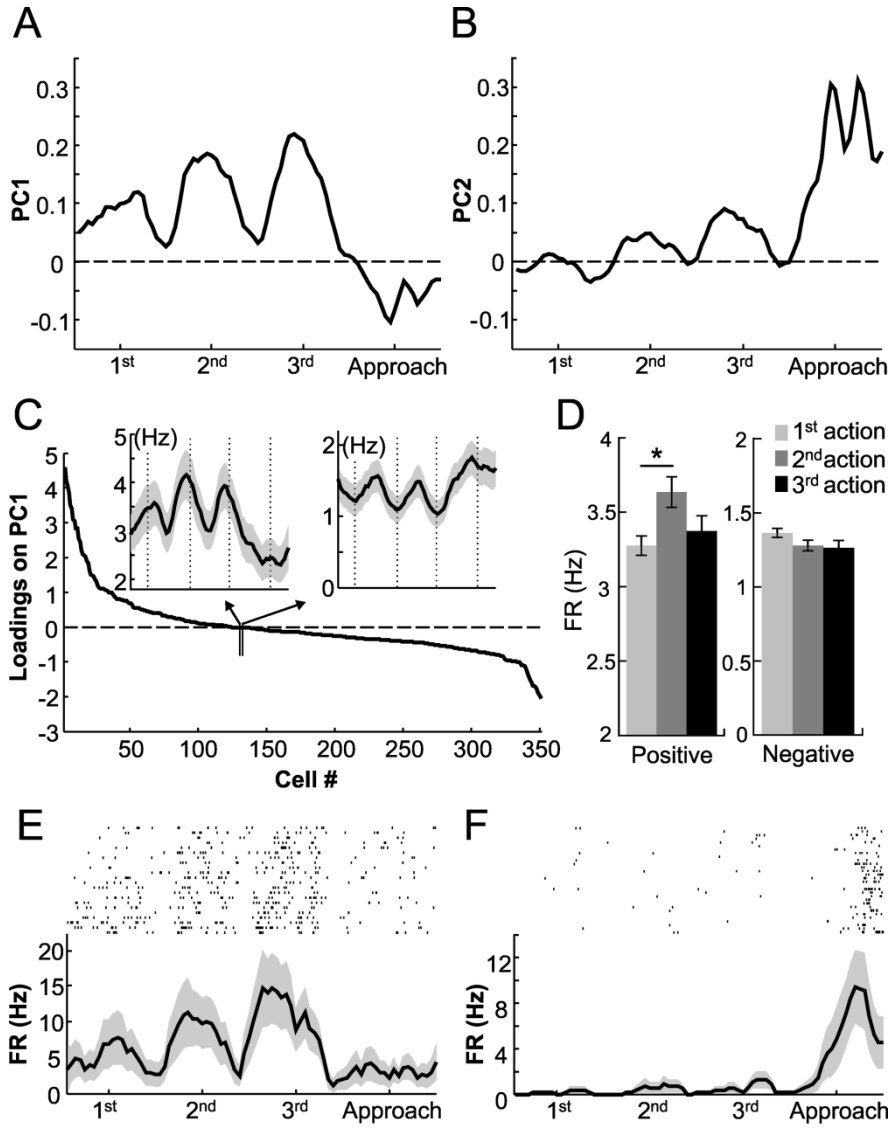
A
B
PC1
PC2
0.3
0.2
0.1
0
-0.1
1st
2nd
3rd
Approach
C
D
Loadings on PC1
Cell #
(Hz)
FR (Hz)
1st action
2nd action
3rd action
Positive
Negative
*
E
F
FR (Hz)
20
15
10
5
12
8
4

Figure 3.7 The main patterns of firing rate variance in DS ensembles as revealed by PCA

Plots were arranged as in Figure 2. **A)** The PC1 eigenvector was characterized by strong variations in association with each lever presses followed by an abrupt decline during the reward approach epoch. **B)** The PC2 eigenvector showed smaller variations in association with the lever presses followed by an abrupt acceleration in the reward approach epoch. **C)** Factor loadings for all DS neurons on PC1. *Left inset:* average firing rate of all neurons with positive factor loadings on PC1. These neurons exhibited peaks in firing tightly locked to lever presses. *Right inset:* average firing rate of all neurons with negative factor loadings on PC1. These neurons exhibited troughs in firing tightly locked to lever presses. **D)** Average firing rates during the 1s period surround each of the three lever presses and reward-approach for neurons that were loaded positively (left) or negatively (right) on PC1. Although there was a significant increase in DS firing for positively loaded cells from the first to the second action, this increase did not continue to the third action even though it was the action closest to the reward. No change in average firing was observed for the negatively loaded cells. **E)** Example of a single neuron that was strongly positively loaded on both PC1. Top: raster plot; bottom: average FR across all trials (mean ± s.e.m.). **F)** Example of a single neuron that was strongly loaded on PC2. Top: raster plot; bottom: all-trial average FR (mean ± s.e.m.). * denotes significance at $p<0.01$.

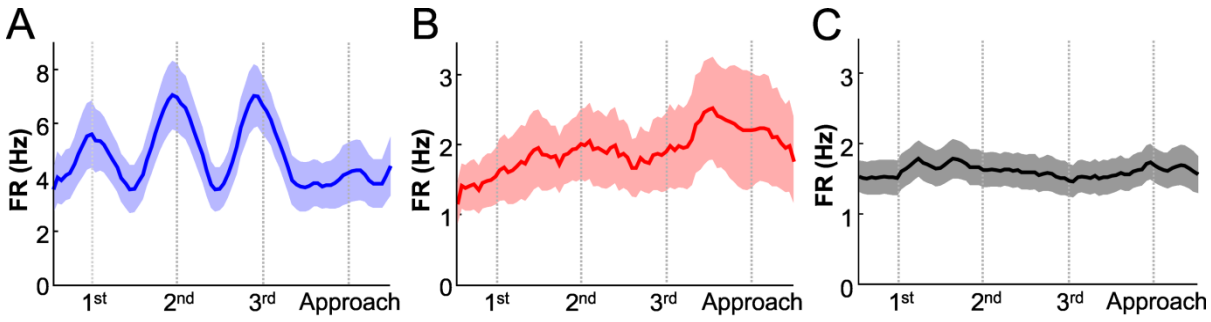

Figure 3.8 Possible single-unit firing patterns during the three actions leading to reward in the DS

All 351 DS neurons were tested for correlations with the model features shown in Figure 4A. **A)** The average firing rate for all DS neurons showing a positive correlation with the action-linked model (n=32). **B)** The average firing rate for all DS neurons with a positive correlation with the ramp model (n=19). **C)** The average firing rate for DS neurons whose firing did not correlate with either model (n=211).

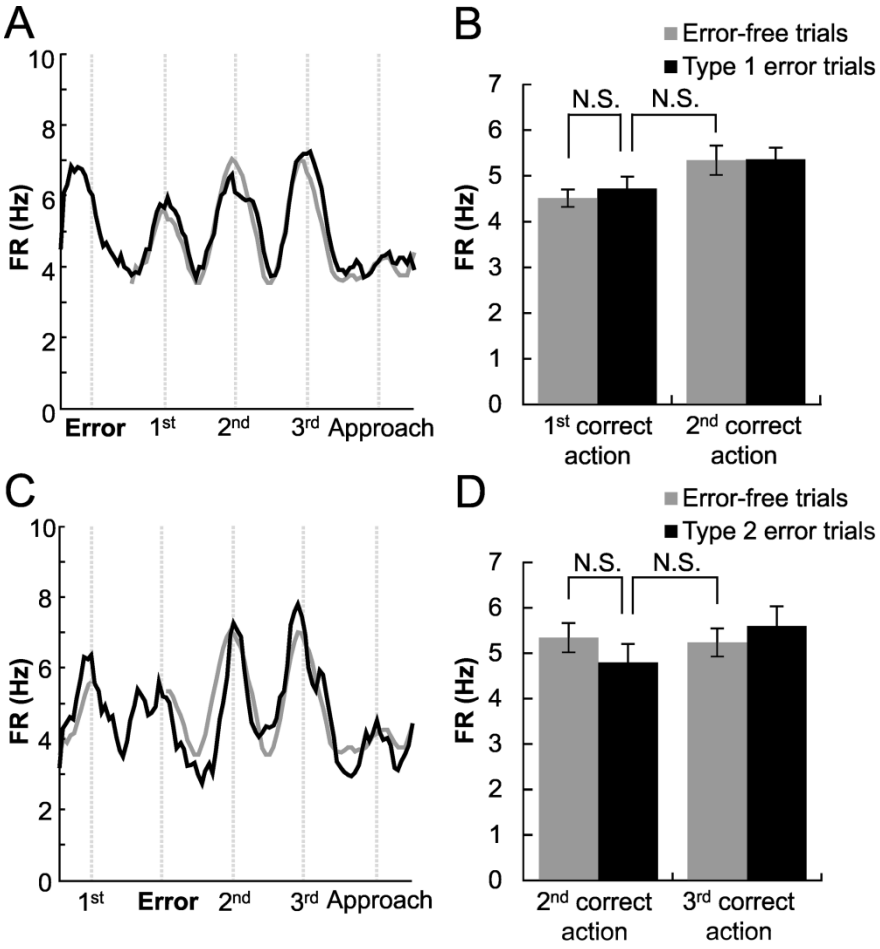
A
B
C
D
FR (Hz)
Error
1st
2nd
3rd Approach
Error-free trials
Type 1 error trials
Type 2 error trials
N.S.
1st correct action
2nd correct action
3rd correct action

Figure 3.9 Firing rate progression in the DS during correct choices versus errors

A) Neurons' average firing rate on error-free trials shown in Fig 4B was re-plotted in green, with the average firing rate of the same neurons on trials starting with an error in black (Type 1 error trials). **B)** Statistical analysis of firing rate changes of these neurons on all correct trials (gray bars) versus trials where the first response was an error (black bars). *Left:* there were no significant differences (n.s.) in average firing rates in the 1s period surrounding the first *correct* lever press depending on whether that lever press was preceded by an initial error (black bar) or not (gray bar). *Right:* importantly, and in contrast to the ACC, for DS neurons the first and second *correct* lever presses were not significantly different. This was true regardless of whether the trial began with an error (black bar) or not (gray bar). **C)** The average firing rate change of the same neurons examined in (A-B) but on trials where the second lever press in the sequence block was an error (Type 2 error trials, black line). **D)** Statistical analysis of firing rate changes of these neurons on all correct trials (gray bars) versus trials where the second response was an error (black bars). *Left:* there were no significant differences (n.s.) in average firing rates in the 1s period surrounding the first correct lever press depending on whether that lever press was followed by an error (black bar) or not (gray bar). *Right:* however, more critically, the average firing rates for the first versus second *correct* lever presses were not different, and this was true regardless of whether or not an error occurred after the first correct choice (black bar) or not (gray bar). Since there was no difference in firing for the first versus second correct choice or the second versus third correct choice, the DS could not be considered to track action progression towards the goal. Rather the DS tracked all actions equally regardless of serial position or whether they were correct or not.

Chapter 4: Differences in emergent coding properties in cortical and striatal ensembles

4.1 Introduction

A fundamental goal in neuroscience is to determine the specific function and information processing capabilities of unique brain regions. A common approach to this problem is to infer what type of information is being processed by a given region based on what its neurons respond to. While this approach has proven invaluable for our understanding of sensory areas, for regions like the frontal cortex, it can be problematic since frontal neurons are capable of responding to practically anything (Jung et al., 1998, Duncan, 2001). Furthermore, at least in the case of action representations, the single neuron correlates can be quite similar between the frontal cortex and its efferent target, the striatum. In this case, perhaps it would be informative to consider *how* neurons collectively respond rather than *what* they respond to.

How a neuron responds is determined by many factors including the differences in the biophysical properties of the individual neurons, the input it receives as well as the structure of the local microcircuit. In the neocortex, pyramidal neurons form massive, connected networks with interneurons exerting relatively weak inhibition (Fino and Yuste, 2011, Packer et al., 2012). This arrangement favors recurrent excitation and with it, persistent activity patterns through time. Overall, cortical ensemble patterns are usually well-balanced with equal numbers of neurons slightly increasing or decreasing their response to any given event (Rich and Shapiro, 2009, Durstewitz et al., 2010). By contrast, local regions of the striatum receive strong excitatory inputs from the cortex, which generate responses that are temporally

restricted by powerful and widespread inhibition produced by local interneurons(Parthasarathy and Graybiel, 1997, Koos and Tepper, 1999, Gage et al., 2010). Accordingly, when recorded from behaving animals, striatal neurons tend to be activated transiently and synchronously (Graybiel et al., 1994, Berke et al., 2004). This synchronized activity may be an asset in the context of movement generation and learning (Graybiel, 1995, Berke et al., 2004), but may also be a detriment if it concurrently amplifies variability or 'noise' in the neural responses over time(Averbeck and Lee, 2006). By understanding how these differences might influence the way in which neural ensembles within the frontal cortex versus striatum encode the same information, I hope to gain new insight into their unique functions of these two interrelated brain regions.

To address this issue, bundles of tetrodes were used to record simultaneously from neurons in the anterior cingulate cortex (ACC) and the central portion of the dorsal striatum (DS)(Sesack et al., 1989, Zheng and Wilson, 2002)while rats performed a sequential action task. Although there is already a rich literature on the responses of individual frontal cortex and striatal neurons on such tasks (Barone and Joseph, 1989, Aldridge and Berridge, 1998, Nakamura et al., 1998, Procyk et al., 2000, Shidara and Richmond, 2002, Fujii and Graybiel, 2003, Ninokura et al., 2004, Ryou and Wilson, 2004, Schmitzer-Torbert and Redish, 2004, Fujii and Graybiel, 2005, Lu and Ashe, 2005, Mushiake et al., 2006, Averbeck and Lee, 2007, Shima et al., 2007, Seo et al., 2012), our aim was to better understand how coherent ensemble representations emerged from the activity of single neurons. The task required rats to presses three levers in different temporal orders (termed 'sequence blocks' below) to receive reward. It was found that even though single neurons in the ACC and DS represented information about sequences and lever presses with similar overall accuracy, the unique way this

information was combined across the neurons through time in the ACC yielded far superior overall ensemble sequence decoding accuracy than I observed in the DS.

4.2 Methods and material

4.2.1 Subjects

Four experimentally naïve male Long-Evans rats (450-550g) were housed in a facility with 12hr light-dark cycle, with all training and recording taking place during the light cycle. For the duration of the behavioral experiments, the rats were food-restricted to just below 90% of their free-feeding weights. Feeding took place in the home cage after their daily training/recording sessions, and water was available *ad libitum* in the cages at all times. All procedures were carried out in accordance with the Canadian Council of Animal Care and the Animal Care Committee at the University of British Columbia.

4.2.2 Apparatus

Within a large Plexiglas box (25”×18”), a main panel was installed with 3 levers designated Lever1 to 3 from right to left (**Fig. 4.1A**). On any given day of sequence training, a unique tactile object (velcro, cardboard or soft foam) was stuck to the lever panel (but not on the lever itself) and the area on the floor immediately in front of the lever panel to symbolize the order in which the 3 levers should be pressed. An area of 25”×13” was left for the rat to move freely. On the opposing-side wall, a food cup was located at the centre, with each delivery of reward accompanied by a pure tone. Retractable levers and pellet dispenser were controlled and recorded with a PC via a Med Associate interface system (St.Albans, VT, USA).

4.2.3 Sequence task

The naïve subjects were first trained on an FR1 schedule to press each of the 3 levers. A minimum of 60 presses within 0.5hrs, with no less than 15 presses on each lever was required prior to the rat moving on to the next stage of training. After 3-5d of FR1 training, the rats were trained on sequence A: Right Lever→Middle Lever→Left Lever. Thereafter they were trained on sequence B (Middle Lever→Left Lever→Right Lever) and finally sequence C (Left Lever→Right Lever→Middle Lever). In each case food reward was given after the correct lever press in the sequence. A lever retracted only when it was pressed in the correct order, and remained extended in the event of an error. For training on all sequences, the percentage of correct response on the 3rd item of the sequence had to reach 75% before moving on to the next trained sequence. The order of lever presses in each sequence was given by tactile objects placed on the panel and the floor in front of the levers. For a given animal, each object consistently designated a single serial position. At any one of the 3 stages of single-sequence training, if after 3 days of training, the animal still hadn't reached criterion and if day-to-day improvement ceased, a delay-punishment protocol was introduced to extinguish errors made on the 3rd lever of the given sequence. Specifically, if the 3rd lever was pressed before the 1st lever, all levers retracted and a 10-s time-out period ensued. This training continued until the animal reached criterion performance. When the criterion had reached criteria on all 3 sequences, the rat was surgically implanted and allowed 10d to recover. After recovery, 2-3 refresher sessions on each sequence were given before the first multi-sequence test day.

On the multi-sequence test days, the animals had to perform a minimum of 10 trials on each sequence at or above criterion, and switch from sequence to sequence in one of three

possible orders pseudorandomly: Sequence C→Sequence B→Sequence A, Sequence B→Sequence A→ Sequence C, or Sequence A→Sequence C→Sequence B. In-between sequences, the animals were taken out of the box to allow for rearrangement of the tactile objects. Because the task was self-paced, and because it always took the animals more than 1s between lever presses, each lever-press epoch was defined with reference to the time stamp of the lever press. In all but the ROC analysis, six 200-ms time bins were used, including four bins prior to the LP bin, the LP bin itself, plus one bin after the LP. In the ROC analysis, one more bin was included after the 6-bin epoch to have 7 numbers for the purpose of having meaningful and reliable correlation coefficients. In addition, after the last lever-press epoch in each trial, a 1-s period was defined as the 'reward approach' period, during which the animals ran from the last lever towards the food cup, prior to consuming the reward. Neural activities during reward approach and consumption were shown in single neuron examples (**Fig.4.3D-F**) but were not used in any analyses.

4.2.4 Surgery

Stereotaxic surgeries were performed with sterilized-tip procedures under anesthesia by isoflurane. NSAIDs analgesic, antibiotic, and a local anesthetic, were given before incision. One elliptical-shaped craniotomy was made centered at: AP: +3.2mm, ML: +1.0mm, and another craniotomy was made centered at AP: +1.2mm and ML: +3mm. Once the dura mater was retracted, the bottoms of the two bundles of 8 30-gauge tubes, containing a total of 16 tetrodes, were placed on the cortical surface. The bundles were of cylindrical shape with a bottom radius of ~0.4mm, and were angled medially by ~15 degrees. The implants were fixed with bone screws and dental acrylic. At the end of the surgery, tetrodes in the anterior bundle were extended by ~1.4mm into the brain to enter the anterior cingulate cortex (ACC), and

tetrodes in the posterior bundle were extended by ~3mm to enter the dorsal striatum (DS). Animals were given 10d to recover. Prior to each recording session, small adjustments were made with the hyperdrives to maximize the number of neurons recorded.

4.2.5 Acquisition of electrophysiological data

Data acquisition and offline spike sorting were conducted using the same methods, equipment and software as described in Chapter 2 and previously published in Hyman et al. (2012).

4.2.6 Histology

At the end of the studies, the animals were deeply anesthetized using urethane i.p. injection, and a 100μA current was passed through the electrodes for 30s. Animals were then perfused with a solution containing 250ml 10% buffered formalin, 10ml glacial acetic acid, and 10 g of potassium ferrocyanide. This solution causes a Prussian blue reaction, which marks with blue the location of the iron particles deposited by passing current through the electrodes. The brains were then removed and stored in a 10% buffered formalin/20% sucrose solution for at least 1 week, before being sliced and mounted to determine precise electrode locations. Since multiple sessions were recorded from individual animals the precise recording locations could not be derived from electrode lesions, but all electrode tracks were inferred between the entrance point and the dyed spot. Representative recording sites are shown in **Fig. 4.8A,C**, and the ranges of recording are shown in **Fig. 4.8B,D**.

4.2.7 Data analysis

Instantaneous firing rate (iFR): A total of 33 large ensembles were collected from 4 rats that acquired all 3 sequences and successfully switched among them within a given session. To

obtain an estimate of the neural firing rate for each isolated cell *i* as a function of time bin *t*, $r_i(t)$, for each spike train in each 200-ms bin, the instantaneous firing rates (iFRs) were calculated as the reciprocals of the inter-spike intervals, convolved with 20-ms Gaussian kernels and then averaged (Durstewitz et al., 2010, Hyman et al., 2012). Neurons firing less than 0.14 Hz were excluded from further analysis, because the sample of spikes was too small (250 or less) to be reliably representative of the cell's activity in relation to behavior. Each lever-press epoch included the 1-s period centered at the moment of lever-press, while the reward-approach period was the 1-s period immediately after the 3rd lever-press epoch.

Multiple Linear regression (MLR): A model with 5 mutually independent and uncorrelated factors that characterized the animals' spatial location and movements were used to isolate the effect of behavioral variability on neural activities during the 'common-segment' pairs. A linear model was employed simply because it was impossible to specify a type of consistent non-linear relationship between the firing rate of the neurons and changes in the behavioral variables. Nevertheless linear regression was appropriate since for each neuron the residuals and the predicted values produced by the model were uncorrelated, indicating that the iFRs indeed had a linear relationship with the behavioral predictors. All behavioral analyses were performed using the common segment LPs. Common segments included instances where the same two levers were pressed in the same order, but as part of two different sequences. In 13 out of the 19 sessions, the rats performed 3 sequences, and all 3 possible pairs of common-segments were used. In 6 out of 19 sessions the rats performed 2 sequences, so only one pair of common-segments existed in each of these sessions, and was used in the analysis.

Behavioral variables were constructed based on video tracking data. Since the video rate was 1 frame/33.3ms, 6 frames were averaged to synchronize the video with the 200ms

iFR bins. The first two factors extracted from the video tracking data were the 'X' and 'Y' position of the animals in the chamber. Approach angle and velocity were calculated based on the animal's current position and its position 200ms prior. The third factor 'A' was the approach angle, or the angle between this vector and the lever. The fourth factor 'V' was the bin-by-bin velocity of the animal during lever approach and was simply the length of this vector at each time step (converted into cm/s). The fifth factor 'T' was the time (in seconds) since the previous lever press:

$$F(t) = b_0 + b_1X(t) + b_2Y(t) + b_3A(t) + b_4V(t) + b_5T(t) + \varepsilon(t), \text{ (eq. 1)}$$

For a given time bin *t*, *F(t)* represented the neuron's normalized iFR within that bin. For each neuron, the percentage of variance accounted for by the model (the R^2 statistic) was obtained and entered into an overall distribution (**Fig. 4.2B**). The residual matrix resulting from the regression in each session was used to perform a series of control analyses, including D_{Mah}, D_{Mah}-based leave-one-out error, and ensemble variance and covariance (see sections below). By comparing the results from these residual-based analyses with those from the full iFR matrices, the effect of behavioral variability on sequence decoding was elucidated.

Selectivity index: To examine whether individual units were responsive to sequence, the selectivity for each unit *i* with respect to each pair of common segments—both of which were associated with the same lever but belonged to two different sequence blocks—was obtained by grouping the firing rates into two classes: the iFRs of a given neuron during one common segment were assigned to Class A, whereas the iFRs of the same neuron during the other common segment in the pair were assigned to Class B. The index was then computed as:

$$d_i' = \frac{\left|\left\langle \{r_i(t) \mid t \in A\} \right\rangle - \left\langle \{r_i(t) \mid t \in B\} \right\rangle\right|}{\sqrt{\sigma^2_{i,t\in A} + \sigma^2_{i,t\in B}}}$$ where $\langle \cdot \rangle$ denotes the *mean*.

Receiver operating characteristic (ROC) curve and statistics: The ROC method was used to test classification performance of single neurons and ensembles. ROC analysis has an advantage in situations where the distributions are unknown because it assesses performance over a range of threshold values rather than being forced to evaluate differences at a single threshold level relative to some theoretical distribution. In a simple example, if a detector's sensitivity level is set to 0.5, signals (or observations) stronger than this level are reported as positives (or '1') and below as negatives (or '0'). Out of all positives, the proportion of incidents when the target was truly present would be the 'hit rate' and the remaining proportion would be 'false-alarm rate'. Each threshold level yields a single dot in the ROC curve. In the present study, an ROC analysis was run on correlation scores as described in the text, for the classification of both sequences and lever locations. Because correlation coefficients are on a continuous scale, a large number of thresholds were used to produce the detailed ROC curves. The area-under-curve (AUC) from each individual curve was used for statistical analysis. Because a 45-degree straight-line with an AUC=0.5 indicates a lack of signal-noise differentiation, the AUCs can be tested for significance by a simple one-sample t-test against a normal distribution with mean = 0.5 and unknown variance. In the case where AUCs from ensembles and from single-units were compared, an independent-sample t-test was used. In all of the cases involving multiple groups, a 2-way ANOVA was also used with 2 control groups—random normal distributions with the same sample size and the same variance and mean = 0.5. The relative numbers of neurons in each region were not different from those obtained using the Bonferroni-corrected one-sample t-test.

In order to analyze the effect of ensemble size on the between-sequence separation of activity states, ensemble AUCs were calculated between pairs of common-segments (i.e. lever responses that differed only in their sequence identity) based on randomly selected ensembles with varying sizes (n=4, 7, 10, 13, 16 and 19). For each ensemble size, 100 random draws of n neurons were performed for each session and the resulting AUC averaged, for a complete representation of the whole data set. For both the ACC and the DS, the average AUCs at each ensemble size were plotted (**Fig. 4.4D**) and a power function was fitted for each region. The power function was selected as it accounted for more than 99% of the variance.

Multiple-single unit activity (MSUA) analysis: For population analysis, population vectors $\mathbf{r}(t) = [r_1(t) \ldots r_N(t)]$ were constructed, with N equal to the number of single units isolated from a given recording session. The term MSUA space refers to the N-dimensional space spanned by all recorded units and populated by these vectors $\mathbf{r}(t)$. Each dot in the MSUA space represents the state of the entire recorded ensemble within one 200 ms bin. All points corresponding to different 200ms bins within the epochs of the same behavior are shown in the same color. All statistical analyses were performed in the full space of all recorded units. For the purpose of visualization, multi-dimensional scaling (MDS) was applied to reduce dimensionality.

To quantify the effects of sequence and lever location on network activity, the Mahalanobis distances (D_{Mah}) were computed between the sets of N-dimensional vectors associated with task epochs of interest. To control for differences in MSUA space dimensionality (i.e. ensemble size) in D_{Mah} comparisons, a normalization procedure was employed: N_{min} was the minimum number of units recorded in any of the data sets to be compared, and K_{min} was the minimum number of time bins. For data sets with N and K greater

than N_{min} and K_{min}, N_{min} units and K_{min} data points were selected at random and D_{Mah} was computed. This procedure was repeated 1000 times and the results averaged to make full use of all units and data points recorded. The resultant normalized D_{Mah} averages were used in various statistical analyses. In order to determine the significance level of a given D_{Mah} value, between-sequence separation was compared to within-sequence separation (**Fig. 4.5C**)**,** and between-lever separation was compared to within-lever separation. To calculate average D_{Mah} within a sequence block, bootstrap surrogate blocks were created by randomly shuffling 1-s blocks of the iFR matrices. The distance between the 2 shuffled blocks therefore represents the separation between activities during random behavioral events. The process was repeated 100 times and the D_{Mah} values averaged. The residual matrices were also used from the linear regression (see above) to calculate the D_{Mah}, and compared the results to the D_{Mah} calculated from the full iFR matrix in order to reveal the influence of the behavioral variables on sequence decoding (**Fig. 4. 6C, E**).

The calculation of D_{Mah} incorporates 3 aspects of ensemble activity: the difference in mean firing rates (i.e., Euclidean distance or D_{Euc}), the variance in each neuron's activity and covariance between any 2 cells in the ensemble. The focus was on the bin-by-bin variance and covariance, calculated among neurons after averaging across all trials for each time bin within the lever-press epochs. In order to better understand the difference between ACC and DS in sequence encoding on the ensemble level observed in D_{Mah}, D_{Euc}, variance and covariance was also analyzed separately. From each session, 19 cells were randomly drawn and the total covariance were calculated and summed, and the process was repeated for 100 times before the results were averaged for each session. Thus the ensemble bin-to-bin variance was the average of summed variance in a typical 19-cell (i.e. Nmin) ensemble (**Fig. 4.6A**), and the

ensemble covariance shown in **Fig. 4.6B** was the summed absolute value of covariance between each cell pair in an ensemble. The absolute value of the covariance were used as the focus was on its magnitude rather than its direction. The residual matrices from the linear regression (see above) were also used to calculate the ensemble variance and covariance in order to examine the effect of behavioral variation (**Fig. 4.6A** and **B**).

Leave-one-out Mahalanobis discriminant analysis (MDA): In the analysis above, D_{Mah} was calculated between clusters. The leave-one-out prediction uses another variant of D_{Mah}: dot-to-cluster distance. In each distance calculation, a 'dot' is the ensemble iFR vector of a single time-bin recorded during a common-segment action. When all the dots for each sequence block are plotted together, there are two clusters in the MSUA space. If the D_{Mah} from the dot to its home sequence-cluster is shorter than that to the alternative sequence-cluster, then a correct classification is counted. The final performance is shown in percentage of correct classifications out of all time-bins tested (**Fig. 4.5D**). In order to have a control for the classification performance, bootstrap surrogate blocks were created by randomly shuffling 1-s blocks of the iFR matrices. The distance between the 2 shuffled blocks therefore represents the separation between activities during random behavioral events. The process was repeated 100 times and the leave-one-out errors averaged. Two-way ANOVA was then used to test the performance of ensembles from both regions (**Fig. 4.5D**). Additionally, the residual matrices from the linear regression (see above) were used to calculate the leave-one-out errors and their results were compared the results to those calculated from the full iFR matrix (**Fig. 4.6D** and **F**).

F-statistic for sequence discrimination and for lever discrimination: The F-statistic was calculated for 2 different purposes: sequence classification and lever discrimination. To

characterize the cells' temporal profile of sequence discrimination (**Fig. 4.7A**), for each neuron, the F-statistic was calculated between the iFRs during each pair of common segments (from 2 different sequence-blocks) in each of the 6 time bins, covering the interval from 900ms prior to a given action to 300ms after the action. In other words, the F-statistic for sequence was the between-sequence variance divided by the within-sequence variance in a given time bin. Subsequently, the frequency distribution of the time bin in which a cell achieved its maximum F-statistic was plotted for each region (**Fig. 4.7A**).

To characterize the level of lever discrimination throughout the lever-press epoch, F-statistics were also calculated among each cell's activities associated with responses on the 3 levers within each sequence block. The F-statistic for action was the between-lever variance divided by the within-lever variance in a given time bin, in a given sequence. The frequency distribution of the time bin in which a cell achieved its maximum F-statistic is plotted in **Fig. 4.7B**.

4.3 Results

4.3.1 Behavior

The experimental apparatus is depicted schematically in **Fig. 4.1A** and the task in **Fig. 4.1B**. Different sequences of actions (referred as 'sequence blocks') consisted of presses on the same physical levers in different temporal orders. Each lever was distinguished by specific cues temporarily affixed to the area immediately surrounding the levers. The sequence of cues was always the same for a given rat, but the cues were moved to different lever locations for each of the two or three sequence blocks. This task design permitted an examination of the manner in which the two regions encoded objective information about discrete lever press

actions versus more abstract information about 'sequences' of actions. The rats were required to perform three different sequences of operant responses in a single session, and did so with a high degree of accuracy (% correct response: mean ± SD: 89.2 ± 5.4%). Overall behavioral performance across different sequences in all sessions did not differ (1-way ANOVA, $F_{2,48}$=0.22, p=0.80; **Fig. 4.1C**). Across all trials in all sequence blocks, the latencies from the first to the second lever, and those from the second to the third lever were also equivalent (unpaired t-test, t_{1102}=0.89, p=0.38).

The neural analysis described below will compare neural activity associated with presses on the same physical lever in different sequence blocks. Even though the focus will be on presses on the identical physical lever, the lever may be approached at a different body angle, along a different trajectory or at a different velocity in each sequence block. This could be problematic because differences in movements or movement paths leading up to an action can affect ACC activity(Euston and McNaughton, 2006, Cowen and McNaughton, 2007). In an attempt to minimize the behavioral variability across sequence blocks, I focused on 'common segment' elements. 'Common segments' consist of presses on the same lever when approached from the same preceding lever in two sequence blocks. As shown in **Fig. 4.2A**, the average X,Y trajectories, the angle of approach to the lever and the approach velocity were similar for common segment lever presses performed in different sequence blocks. None of the behavioral measurements differed significantly across sequence blocks (repeated-measures ANOVA, no effect of sequence for x-position: $F_{1,78}$=1.24, p=0.27; for y-position: $F_{1,78}$=3.37, p=0.07; for approach angle: $F_{1,78}$=0.0063, p=0.94; for approach velocity: $F_{1,78}$=0.055, p=0.82; no interaction between time and sequence: for x-position: $F_{5,390}$=0.061, p=1; for y-position:

$F_{5,390}$=0.49, p=0.79; for approach angle: $F_{5,390}$=1.91, p=0.093; for approach velocity: $F_{5,390}$=1.18, p=0.32).

While similar, they were not completely overlapping and in order to determine whether this remaining behavioral variability may have impacted the neuronal responses, the behavioral variables shown in **Fig. 4.2A** were used as factors in a multi-linear regression model performed individually on each neuron (see Methods). The analysis revealed that the total amount of variability in firing rate that could be accounted for by the 5 factors was on average 3.04% for ACC neurons and 3.39% for DS neurons. In fact, in only 3.99% of the ACC neurons and 5.64% of DS neurons did the behavioral variables collectively account for more than 10% of firing rate variance during the common segment lever press periods (**Fig. 4.2B**). Among these neurons, the percentages of variance accounted for by the model did not differ (unpaired t-test, t_{39}= –1.21, p=0.23). While this percentage would likely be larger if all task periods had been considered, at least for the common segment lever press periods, any impact of the differences in behavior across sequence blocks was relatively small for most individual neurons.

4.3.2 Single neuron correlations of sequence differentiation

The selectivity of single ACC and DS neurons for sequences was assessed using signal detection approach by calculating for each neuron a selectivity index (see Methods). The SI for 'sequence' was calculated by comparing the firing rates in the 6 bins surrounding common segment lever presses performed in different sequence blocks. Even though the effects of behavioral variability on instantaneous firing rates (iFRs) were small for the common segment periods (**Fig. 4.2B**), I nevertheless tested whether they could impact SI-based sequence discrimination. I re-calculated the SI values using the residual firing rate matrix generated

from the multi-linear regression analysis performed above since in the residual matrix, the impact of these variables had theoretically been regressed off.

Across all neurons recorded (N_{ACC}=637, N_{DS}=351), 24.0% of individual ACC neurons and 24.0% of individual DS neurons were selective for sequence (i.e. had absolute SI values >=0.5), a difference which was not statistically significant (2-way ANOVA, no effect of region: $F_{1,60}$=0.11, p=0.74, unequal N HSD test: ACC vs. DS: p=1; **Fig. 4.3A**). When the calculation of SIs was repeated using the residual matrices, 24.1% of ACC neurons and 25.2% of DS neurons were found to be sequence selective. The differences in the number of sequence-selective neurons detected using the full versus residual matrices were not significant for the ACC (2-way ANOVA, no effect of regression: $F_{1,60}$=0.17, p=0.69, unequal N HSD test: ACC original vs. residuals: p=1), or the DS (unequal N HSD test: p=0.96). In order to test the strength of the sequence signals, I compared the SIs of all the neurons possessing SI values >=0.5. The average SIs of these putatively sequence-selective neurons did not differ between the two regions (ACC=0.79; DS=0.73; 2-way ANOVA, $F_{1,916}$=1.34, p=0.25, unequal N HSD test: ACC vs. DS: p=0.34; **Fig. 4.3B**). Furthermore, the mean SI values of these neurons were similar if they were calculated using the full or residual matrices (unequal N HSD test, original vs. residuals for ACC: p=1; for DS: p=0.24; **Fig. 4.3B**). These results implied that the ACC and DS had remarkably similar numbers of neurons that were equally selective for sequence differentiation. The extraordinary similarity in the distribution of SI sequence values for the two regions is evident from **Fig. 4.3C**. Examples of two ACC neurons and a DS neuron that exhibited significant sequence selectivity are shown in **Fig. 4.3D**, **E** and **F** respectively.

4.3.3 The representation of sequence information by ACC and DS neurons versus ensembles

The next question was how single neurons compared to ensembles in terms of sequence differentiation and whether this relationship varied between the two regions. To compare the sequence decoding properties of single neurons versus ensembles on an equal footing, a modified Receiver Operator Characteristic (ROC) analysis approach was employed. The ROC analysis has an advantage in situations where the distributions are unknown because it assesses performance over a range of threshold values rather than being forced to evaluate differences at a single threshold level relative to some theoretical distribution. The signal detection characteristics of individual ACC and DS neurons were investigated by creating for each neuron a sequence- specific template for the 1.4s period leading up to and including a common segment lever press in half of the trials in one sequence block. As the data were binned at 200ms, the template was a vector of 7 instantaneous firing rate (iFRs) values for each single neuron. The template was then moved through both sequence blocks and bin by bin correlations were calculated. This process was repeated using templates created from different groups of trials and the results averaged. A 'true positive' or 'hit' occurred when a correlation score larger than a threshold value was found in the lever press interval of the remaining half of trials of the sequence block from which the original template was constructed. A 'false alarm' occurred when a correlation score larger than the threshold value was found between the template and the 'common segment' lever presses performed in the alternate sequence block. If hits and false alarms occur at the same rates at all thresholds, the ROC curve would be a straight-line with a 45° slope, with an area-under-the-curve (AUC) of 0.5.

In order to compare ensembles to single neurons using this approach, a similar analysis was performed, except in this case rather than the vectors being 7 iFR values, they were N x 7 iFR values, with N being the number of recorded neurons per session. Because the templates were much larger than those of any individual neuron, one would expect a low probability of recurrence of a similar ensemble activity state pattern, thereby reducing false alarm rates. Concurrently, the larger templates would be equally disadvantageous as it is equally unlikely that a match between the template and activity patterns during the remaining half of the trials in the same sequence block (i.e. hits) would ever occur. Because the ROC analysis weighs hits versus false alarms in this way, it is an effective means to directly compare signal detection of single neurons versus ensembles. It is important to emphasize that simply having a larger template confers no advantage on its own. As proof of this, the performance of the ensembles following random time block shuffles is given by the thick dashed lines in each panel (**Fig. 4.4A, B**) which are essentially at 45°.

For sequence decoding, as shown in **Fig. 4.4A**, the hit:false alarm ratios for ACC ensembles (solid red curve) were higher than 96.77% of individual neurons (pink curves) during the 'common segment' periods of the two sequence blocks. In contrast, only 9.12% of individual DS neurons (light blue curves, **Fig. 4.4B**) were superior to the DS ensembles for the same periods (solid blue curve, **Fig. 4.4B**). Overall, single DS neurons were superior to DS ensembles significantly more often than was the case for single ACC neuron compared to ACC ensembles (Pearson's chi-square: χ^2 =12.1, df=1, p=0.0005). It is clear from **Fig. 4.4A, B** that the reason for this difference is not that DS neurons were superior to ACC neurons on a single neuron basis, in fact they were not (independent-sample t-test: t_{941}=1.59, p=0.11; **Fig. 4.4C**), instead, when DS neurons were combined into ensembles, their performance was

significantly worse than when ACC neurons were combined into ensembles (independent-sample t-test: t_{31}=4.44, p=1.1×10^{-4}).

To examine the relative superiority of ACC ensembles for sequence decoding in greater detail, neurons were selected randomly from the population of 637 ACC and 351 DS neurons to create ensembles of different sizes and then the ROC analysis was performed on each ensemble. Every data point shown in **Fig. 4.4D** is the average of 100 draws. This analysis reveals that the ACC ensembles achieved superior signal detection across all ensemble sizes. Additionally, the steeper ACC curve, in comparison to the DS curve, indicated that signal detection improved progressively as more ACC neurons were added. Based on the functions fit to the data, it may be extrapolated that on average a randomly drawn ensemble of 112-neurons would be required for the DS to achieve the same level of sequence-signal detection as an ensemble of 19 neurons drawn randomly from the ACC. This is quite striking when one considers that individual neurons in the two areas performed equally on a per neuron basis (**Fig. 4.3A-C**). Paradoxically, the DS is at a disadvantage when its neurons are combined into ensembles for sequence decoding.

4.3.4 Unique activity state patterns represent sequence information

In order to understand how ACC ensembles achieved higher ROC performance, it was necessary to probe more deeply into the nature of the ensemble codes themselves. **Fig. 4.5A** and **B** show 3-dimensional representations of the multiple single unit activity (MSUA) spaces from representative ACC and DS ensembles respectively. As illustrated in these plots, ensemble activity states during lever presses in one sequence block are tightly clustered but shift to another region of the MSUA space when the same lever is pressed as part of a

different sequence block. A shift in the MSUA space means that the ensembles entered a distinct activity state pattern for each sequence block.

In order to quantify the differences in firing patterns associated with different sequence blocks, the Mahalanobis distance (D_{Mah}) (Krzanowski, 2000, Durstewitz et al., 2010) was calculated between population vectors in the MSUA spaces. The D_{Mah} between the activity states associated with common-segment lever presses in different sequence blocks was significantly larger than the D_{Mah} between shuffled control blocks (Kruskal-Wallis test: $F_{3,140}$=96.34, p=9.5×10^{-20}), for both the ACC (Tukey's test, p=0) and the DS (Tukey's test, p=1.0×10^{-6}; **Fig. 4.5C**). Once again the ACC was superior, as the D_{Mah} between the neural patterns associated with presses on the same lever occurring in different sequence blocks was significantly larger for ACC ensembles than for DS ensembles (Tukey's test, p=0.023).

In order to show that these activity state patterns were actually important for sequence decoding, a leave-one-out variant of the D_{Mah} analysis was applied to the common-segment lever press periods. Using this form of Mahalanobis discriminant analysis (MDA), the correct sequence-block could be accurately predicted in 66.6% of cases for ACC ensembles and in 61.9% of cases for DS ensembles both of which were significantly better than when the procedure was repeated using shuffled sequence-block assignments (2-way ANOVA, main effect of sequence $F_{1,140}$=204.89, p=1.8×10^{-29}, main effect of region $F_{1,140}$=20.53, p=1.3×10^{-5}; **Fig. 4.5D**). More importantly, ACC ensembles classified each specific sequence with higher accuracy than did DS ensembles (Tukey's test, p=0.00022).

4.3.5 Differences between the regions in ensemble variance/covariance

The sequence classification and decoding measures used above weigh differences between the patterns (i.e. the 'signal') relative to the variance and covariance in the patterns across time. Therefore, ACC ensembles could be superior to DS ensembles in terms of the strength of their 'signals' or because they exhibited less variance/covariance through time. In order to disambiguate these two possibilities, each was examined separately.

In order to assess potential difference in 'signal' strength the Euclidean distance (D_{Euc}) was calculated rather than the D_{Mah} between points in the MSUA spaces associated with each sequence block. D_{Euc} and D_{Mah} both measure the distance between the sets of points in the MSUA space but differ in that D_{Euc} is a simple measure of the geometric distances between the centers of clusters, whereas D_{Mah} weighs these distances relative to the individual variances and pooled covariance of the two sets of points. Therefore, D_{Euc} represents a pure measure of the ensemble signal independent of variance/covariance. Unlike D_{Mah}, D_{Euc} between common-segments in different sequence blocks did not differ between the ACC and the DS ensembles (independent-sample t-test: $t_{70} = -1.09$, p=0.28) suggesting that the difference between the ACC and DS was not due to differences in the sequence 'signal'. This was not unexpected given the similarities in sequence differentiation by single ACC and DS neurons described above.

Therefore, the differences in decoding accuracy must be related to the differences in the variance/covariance between the two areas. Accordingly, the variance across bins within a trial was indeed higher in the DS than the ACC (post hoc Tukey's test, p=0.00015; **Fig. 4.6A**, dark bars) as was the absolute covariance (post hoc Tukey's test, p=0.00017; **Fig. 4.6B**, dark bars).

In addition, it was found that a significant portion of the variance and covariance was related to the coordinated firing in response to behavioral events in the DS because variance (repeated measures ANOVA, main effect of regression: $F_{1,70}$=18.3, p=5.8×10^{-5}, Tukey's test: p=0.00017; **Fig. 4.6A**, right bars) and covariance (repeated measures ANOVA, effect of regression: $F_{1,70}$=19.7, p=3.4×10^{-5}; Tukey's test: p=0.00017; **Fig. 4.6B**, right bars) were both lower in the residual matrices than in the full iFR matrices. In contrast, this was not true for the ACC, as the variance (Tukey's test: p=0.88) and covariance (Tukey's test: p=0.77) were both similar in the full and residual matrices from the behaviorally derived multi-linear regression discussed previously (**Fig. 4.6A,B**, left bars). Thus the responses of DS neurons to the behavioral variables contained more coordinated variation throughout a given action than did ACC neurons.

The higher degree of covariance in the responses of DS neurons to the behavioral variables may have contributed to the poorer performance of the DS on the ensemble measures of sequence decoding. To test this possibility the sequence discrimination analysis above was repeated using the residual matrices. This indeed improved the performance of DS ensembles as it led to a significant increase in the D_{Mah} between common segment lever presses (repeated-measures ANOVA and post hoc Tukey's test: residuals>full matrix: p=0.00039; **Fig. 4.6C**) and improved sequence decoding accuracy using MDA (repeated-measures ANOVA and post hoc Tukey's test: residuals>full matrix: p=0.015; **Fig. 4.6D)**. In contrast, repeating the same analysis using the residual ACC matrices had no effect on the ability of ACC ensembles to separate common segments lever presses in different sequence blocks (repeated-measures ANOVA and post hoc Tukey's test: residuals=full matrix: p=0.82; **Fig. 4.6E**) nor did it have any effect on MDA based decoding accuracy (repeated-measures ANOVA and

post hoc Tukey's test: residuals=full matrix: p=0.99; **Fig. 4.6F**). Therefore, the firing rate variance of DS neurons associated with the measures of behavioral variability shown in **Fig. 4.2** contributed to its inferior sequence decoding at the ensemble level. This makes intuitive sense given that the levers being compared in the two sequence blocks were identical and therefore the stronger the neurons encoded this commonality, the less likely they would be able to differentiate between the two sequence blocks.

4.3.6 The superiority of single DS neurons but not ensembles in separating lever locations

The previous analysis suggested that relative to the ACC, DS ensembles more strongly tracked the commonalities associated with pressing the same physical levers in different sequence blocks. In order to test whether the spatial lever decoding ability, the same type of ROC-based analysis described above was again employed. In this case, a template was created by randomly selecting half of all presses on a given physical lever. A 'hit' occurred when a correlation score larger than a threshold value was found in the remaining half responses on the same lever, while a 'false alarm' occurred when a correlation score larger than the threshold value was found between the template and the presses performed on a different lever. This procedure was performed both on single units as well as ensembles as outlined above for the case of sequence decoding.

In support of the prediction, the decoding of spatial lever identity by DS neurons was indeed superior to that of ACC neurons on a single-neuron basis (2-way ANOVA, main effect of region: $F_{1,2853}$=29.78, p=5.3×10^{-8}). Yet surprisingly, in spite of this, ACC ensembles nevertheless matched DS ensembles in lever-decoding ROC performance (2-way ANOVA, no effect of region: $F_{1,93}$=0.70, p=0.40). Likewise the D_{Mah} between population vectors associated

with responses on the 3 different levers were similar between regions (2-way ANOVA, no effect of region: $F_{1,194}$ =1.08, p=0.30; both regions discriminated among the 3 levers significantly better than shuffled controls, 2-way ANOVA, $F_{1,194}$ =497.3, p=2.0×10^{-55}). Similar results were also obtained using MDA (2-way ANOVA: $F_{1,194}$=770.4, p=1.8×10^{-69}; post hoc Tukey's test: DS=ACC, p=0.63). Therefore, even though DS neurons were indeed better than ACC neurons in terms of their ability to differentiate unique lever press actions, this advantage did not translate to the ensemble level. It would appear that on this task, the signals carried by individual neurons tended to be less synergistic when combined in the DS relative to the ACC.

4.3.7 Differences in the timing of sequence and lever press signals in the ACC and DS

What is it about the DS that resulted in this type of counteractive ensemble effect? One factor may relate to the timing of different signals. To explore this issue, a rolling F-statistic was used to find the proportions of neurons that exhibited their maximal firing differentiation of sequence versus lever identity within each trial. The proportions of ACC neurons differentiating sequences (Pearson's Chi-square: χ^2=9.28, df=5, p=0.098; **Fig. 4.7A**, left bars) or lever presses (Pearson's Chi-square: χ^2=4.79, df=5, p= 0.44; **Fig. 4.7B**, left bars) was equal for all time bins within a trial. Furthermore, there was no relationship between the bins in which the maximal differentiation among levers and the bins where the maximal discrimination between sequences occurred within a trial (Spearman's rho: r=0.024, p=0.36). In other words, the sequence and lever press signals evolved independently in the ACC within each trial.

In marked contrast, the proportions of DS neurons maximally differentiating between sequences (Pearson's Chi-square: χ^2=21.62, df=5, p=0.00062; **Fig. 4.7A**, right bars) or

between levers (Pearson's Chi-square: χ^2=17.98, df=5, p=0.0030; **Fig. 4.7B**, right bars) were highly non-uniform across the bins within each trial. More importantly, there was a small but significant correlation between the bins where the greatest proportions of DS neurons maximally differentiated sequences and lever press actions (Spearman's rho: *r*=0.097, p=0.013). Specifically, the maximal point of differentiation for both sequences and lever press actions occurred concurrently at ~600ms prior to the actions in the DS (**Fig. 4.7A, B**). This was also close to the point where generalized activity was also highest (**Fig. 4.7C**). This effect can be seen in the example DS neuron responses shown in **Fig. 4.7D, E**.

Collectively, these results can be summarized as follows. In order to make a correct choice on this task a rat must maintain a representation of the overall sequence block as well as the actual physical lever to be pressed. DS neurons collectively exhibited large firing rate fluctuations in response to the specific lever press actions and this coherent fluctuation appeared at the same point within each trial where the ensembles maximally differentiated sequence blocks. Because the two streams of information were at odds during common segment lever presses of this task, the decoding accuracy for each stream suffered. In contrast, ACC ensembles excelled not because they had more neurons selectively coding the correct action or sequence block, but because the two streams of information remained independent across neurons through time.

4.4 Discussion

Previous studies have observed robust responses of single neurons in the frontal cortex and striatum during actions (Barone and Joseph, 1989, Clower and Alexander, 1998, Nakamura et al., 1998, Procyk et al., 2000, Ninokura et al., 2004, Ryou and Wilson, 2004, Fujii and Graybiel, 2005, Mushiake et al., 2006, Averbeck and Lee, 2007, Shima et al., 2007,

Berdyyeva and Olson, 2010). The present study focused on how information about sequential actions carried by single neurons was combined into ensemble codes and whether this differed in the ACC versus the DS. The results showed that the key factor that separated the two regions was the manner in which information was combined across the population through time rather than specific responses of single neurons.

At the single neuron level, ACC and DS neurons consistently performed similarly across various measures including single neuron sequence selectivity and single neuron ROC performance. However, DS neurons performed more poorly on almost all ensemble-based measures of sequence differentiation, including the separation of the sequence activity state patterns based on D_{Mah}, MDA and ROC performance. The one exception was the separation of the activity state clusters using D_{Euc}. This was the only measure of pure sequence differentiation (i.e. 'signal') that was completely independent of variance/covariance across time. DS neurons tended to exhibit larger variance and covariance and respond more as a collective. The point where the largest proportion of DS neurons tended to maximally differentiate sequence identity was also the point where most neurons differentiated lever identity and this coincided with the point of maximal overall activity. Since in our task, the sequences were different yet the actual physical levers pressed were identical across the two sequence blocks, if sequence and lever signals emerged simultaneously but transiently across the ensemble, the decoding of either would suffer. By contrast in ACC ensembles, information about the physical lever was functionally independent from the information about the overall sequence (**Fig. 4.6E, F**) and as a result, sequence and lever differentiation in the ACC remained high across all 6 bins rather than being concurrently maximal in just one or two. It is also important to emphasize that this uniformity through time was not a result of persistent

firing in individual ACC neurons that has been well characterized in the frontal cortex. Rather it was because the neurons had mutually independent time courses and were thereby able to tile all seven time bins as a population. The asynchrony in the ACC population allowed multiple conflicting sources of information to co-exist, while the synchronous nature of DS activity pitted different representations against each other within a small time window. Of course these differences between the two areas were not absolute, but were nevertheless large enough to significantly impact the accuracy of the representations that emerged at the ensemble level.

While generally disadvantageous in terms of sequence and lever decoding on this particular task, the inherent synchrony in the DS is vitally important for functions mediated by the DS such as movement generation and learning (Graybiel, 1995, Grillner et al., 2005a). The differences in persistence and synchrony between the ACC and DS are therefore functionally important and likely reflected the general properties of the two regions. The larger synchrony of DS over ACC neurons (Berke et al., 2004) may be related to the local circuitry within the striatum itself that is able to transform tonic excitatory cortical drive into alternating and synchronized activity patterns (Carrillo-Reid et al., 2008). In addition, because a small pool of interneurons are able to exert powerful control over many medium spiny neurons (MSNs), it is possible that changes in a single interneuron can effectively turn off a large group of MSNs (Koos and Tepper, 1999). A second important difference I observe between the ACC and DS is the long-tailed firing rate distributions exhibited by populations of DS neurons (**Fig. 4.9**, blue curve). Long-tailed distributions are consistent with firing at low rates interspersed with brief periods of very high activity. This activity profile is likely a consequence of the biophysical properties of MSNs that require large synchronous inputs from the cortex in order

to fire since they possess hyperpolarized resting potentials and strong K+ currents that suppress all but the strongest synchronous inputs (Calabresi et al., 1987, Wilson and Kawaguchi, 1996). The firing characteristics I observed in the DS population are therefore consistent with the known physiology and anatomy of the region.

The increased variance and covariance across DS neurons may help tune DS neurons to promote a single coherent but intermixed signal at a very specific point in time prior to an action. In contrast the ACC may persistently maintain a conceptual representation of the general plan or action strategy that can be independent from the actual actions involved. Therefore both regions are endowed with features that work together to allow the animal to consider the entirety of ongoing experience and yet respond in a decisive manner.

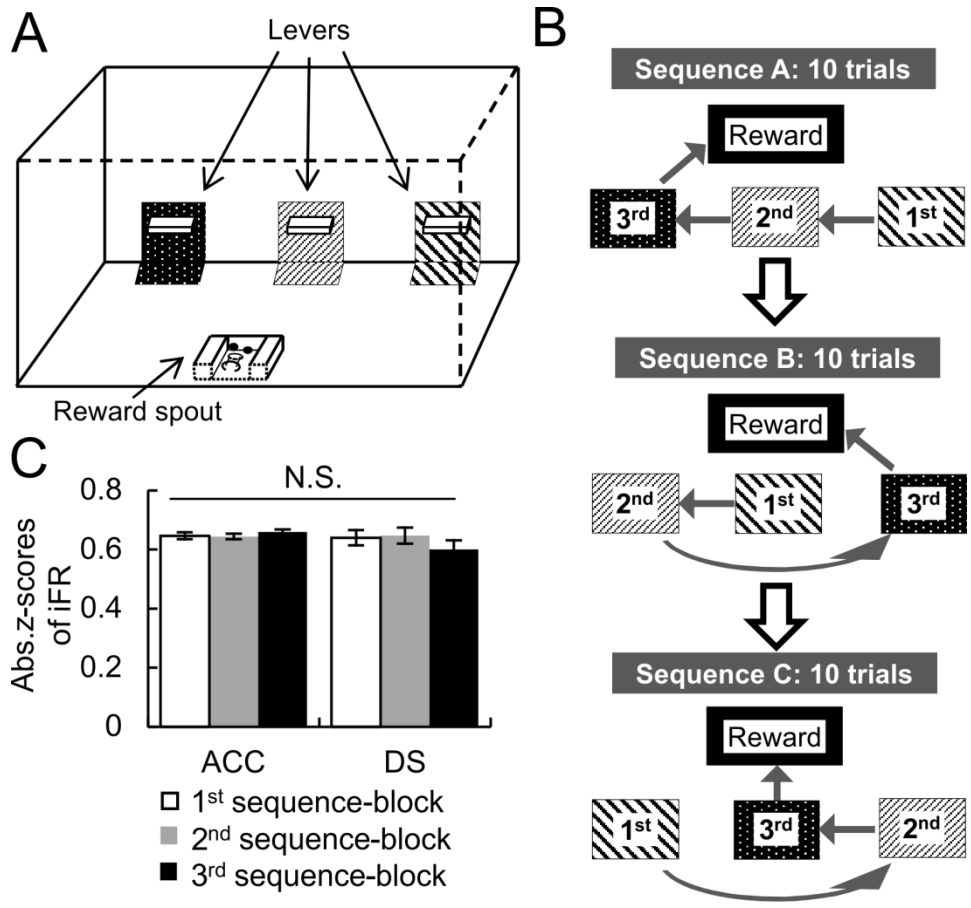

Figure 4.1 Task description and performance

A) The operant chamber contained 3 levers installed on the front panel and a food-cup on the opposing wall. A unique sensory cue was attached to the floor immediately in front of each lever as well as on the surrounding wall. The 1st lever (diagonally-striped in schematic) to be pressed in a given sequence block could have Velcro attached, the 2nd lever (hatched in schematic) cardboard attached and the 3rd lever (dotted in schematic), soft foam attached. **B)** Example of a typical test day where the rat had to perform a minimum of 10 trials on each of the 3 sequence blocks, which were given in a pseudorandom order. On this session, sequence block A required the rat to respond on the right lever, followed by the middle lever and then the left lever, before reward pellets were delivered to the food-cup on the opposite wall. The serial order of the 3 sensory cues (velcro, cardboard, foam) remained constant for a given rat but were moved to different levers for each of the sequence blocks.

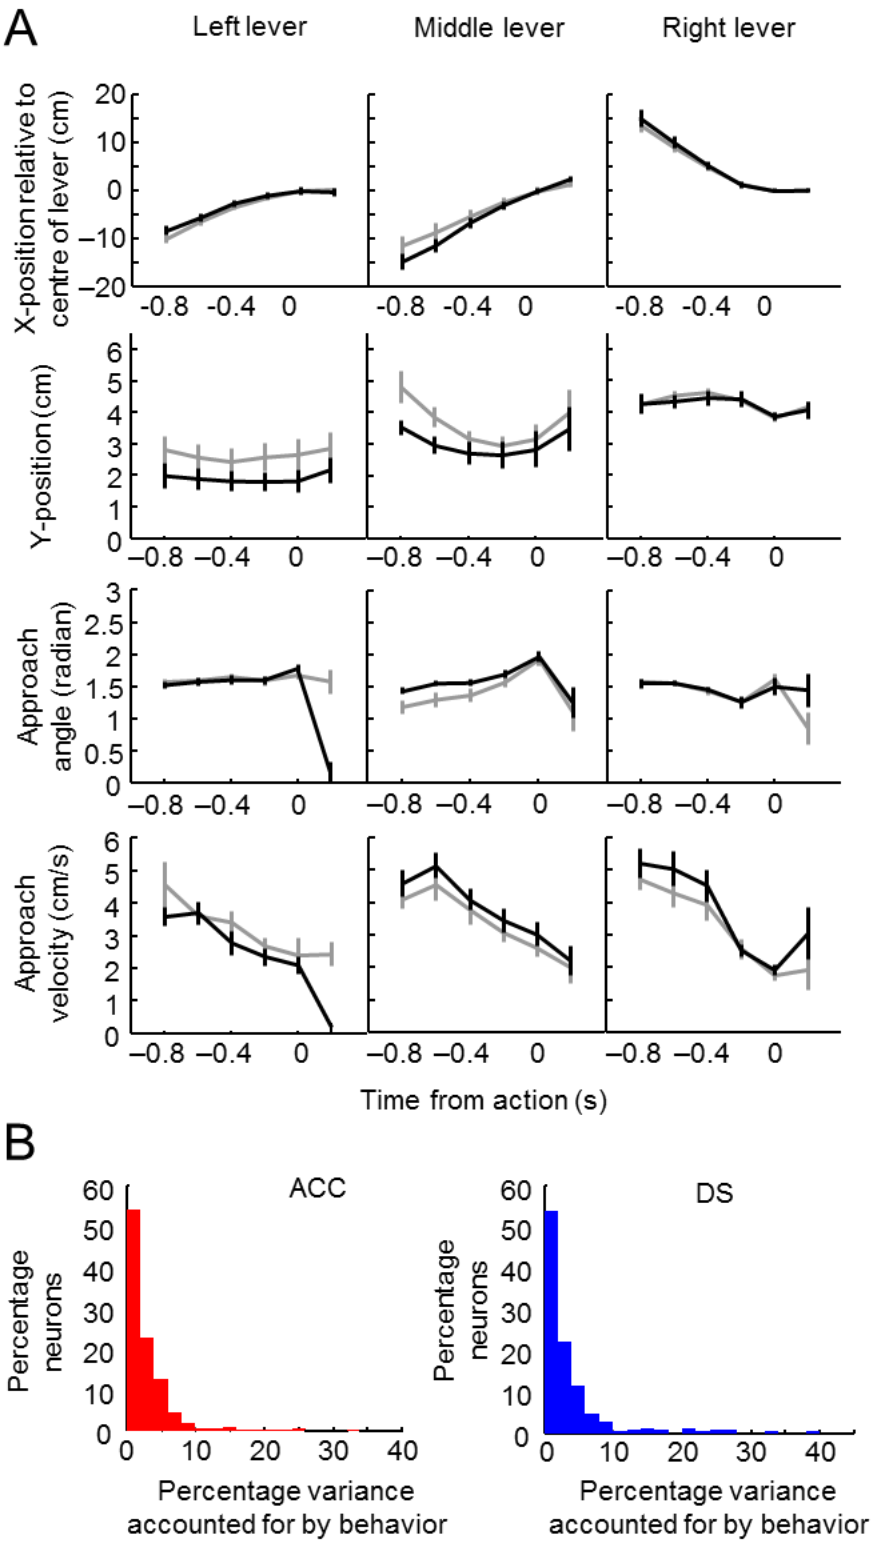
A
Left lever
Middle lever
Right lever
X-position relative to centre of lever (cm)
Y-position (cm)
Approach angle (radian)
Approach velocity (cm/s)
Time from action (s)
B
ACC
DS
Percentage neurons
Percentage variance accounted for by behavior

Figure 4.2 Behavior variables did not vary systematically across sequence blocks

A) Values of the behavioral variables for matched pairs of common-segments lever presses. Common-segments involved presses on the same two levers occurring as part of two different sequence blocks (black vs. gray). Each column represents a different physical lever while the 4 rows display the animal's x-coordinate, y-coordinate, lever approach angle and lever approach velocity in the 6 bins surrounding each lever press (the lever press occurred at time 0 in each panel). The two lines in each panel correspond to the value of each of these variables in one of the two sequence blocks. During these intervals, the animals' location and movement were highly similar across sequence blocks but were quite different across the three physical levers. **B)** Distributions of the percentage of firing rate variance accounted for when the behavioral variables shown in (A) were used as factors in a multiple linear regression analysis, performed on all ACC neurons (left panel) or all DS neurons (right panel). The residual matrices represent the iFR values that remain after the impact of these behavioral variables have been removed and will be used in the analyses shown in subsequent figures.

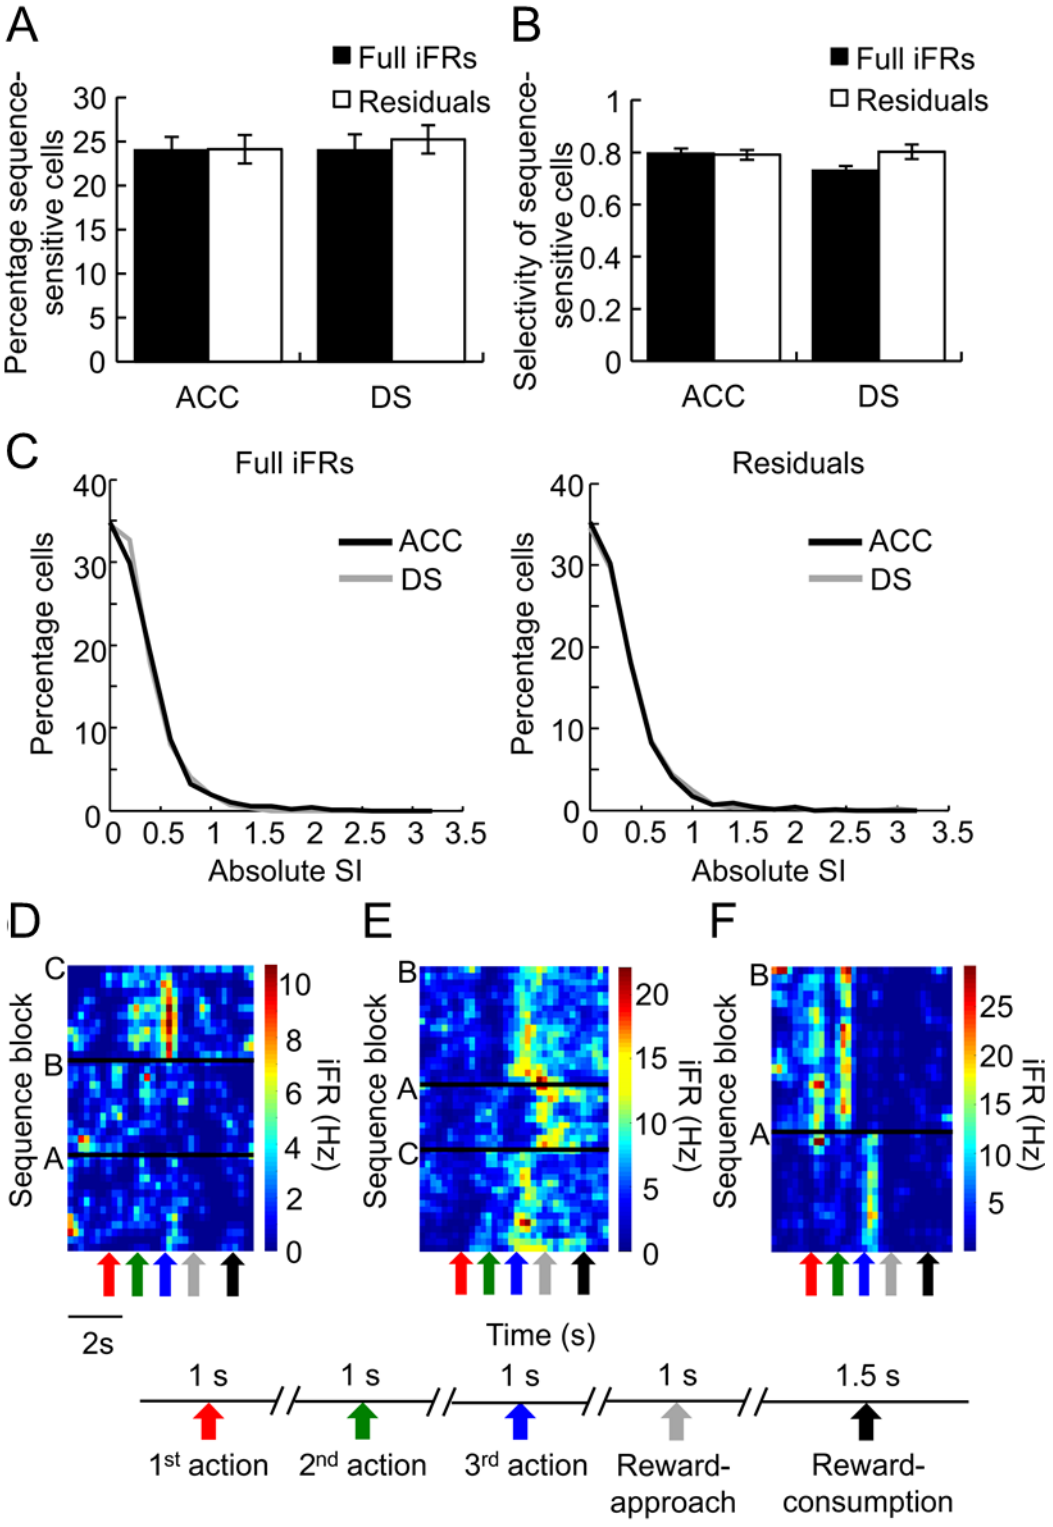
A
Full iFRs
Residuals
Percentage sequence-sensitive cells
30
25
20
15
10
5
0
ACC
DS
B
Full iFRs
Residuals
Selectivity of sequence-sensitive cells
1
0.8
0.6
0.4
0.2
0
ACC
DS
C
Full iFRs
Residuals
ACC
DS
Percentage cells
40
30
20
10
0
0 0.5 1 1.5 2 2.5 3 3.5
Absolute SI
D
E
F
Sequence block
C
B
A
iFR (Hz)
10
8
6
4
2
0
20
15
10
5
0
25
20
15
10
5
2s
Time (s)
1 s
1 s
1 s
1 s
1.5 s
1st action
2nd action
3rd action
Reward-approach
Reward-consumption

Figure 4.3 ACC and DS neurons respond equivalently to different sequences as well as different physical levers

A) The ACC and DS had a similar proportion of neurons showing significant responsiveness for sequences calculated using the original iFR matrices (left, ACC: black bar, DS: white bar) or the residual iFR matrices (right, ACC: black bar, DS: white bar). Error bars indicate s.e.m. **B)** Sequence-sensitive neurons from the ACC (black) and the DS (white) had similar absolute selectivity indices (SI), calculated either using the full iFR matrices (left, ACC: black bar, DS: white bar) or the residual iFR matrices (right, ACC: black bar, DS: white bar). Error bars indicate s.e.m. **C)** The distributions of absolute SIs based on either the full iFR matrices (left) or the residual iFR matrices (right) across the neuronal populations in the ACC (black) and DS (gray) overlap. **D)** Heat relief iFR plot of an example neuron from the ACC that responded to only one type of lever press in one sequence block. **E)** An example ACC neuron that responded to the 3rd lever press (and reward) in a sequence-dependent manner. **F)** Heat relief iFR plot of an example neuron from the DS that responded to only one type of lever press in one sequence block. Note that presses on the middle lever in one sequence block (as the 1st action) consistently elicited responses from the neuron, whereas presses on the same lever in the other sequence block (as the 2nd action) did not.

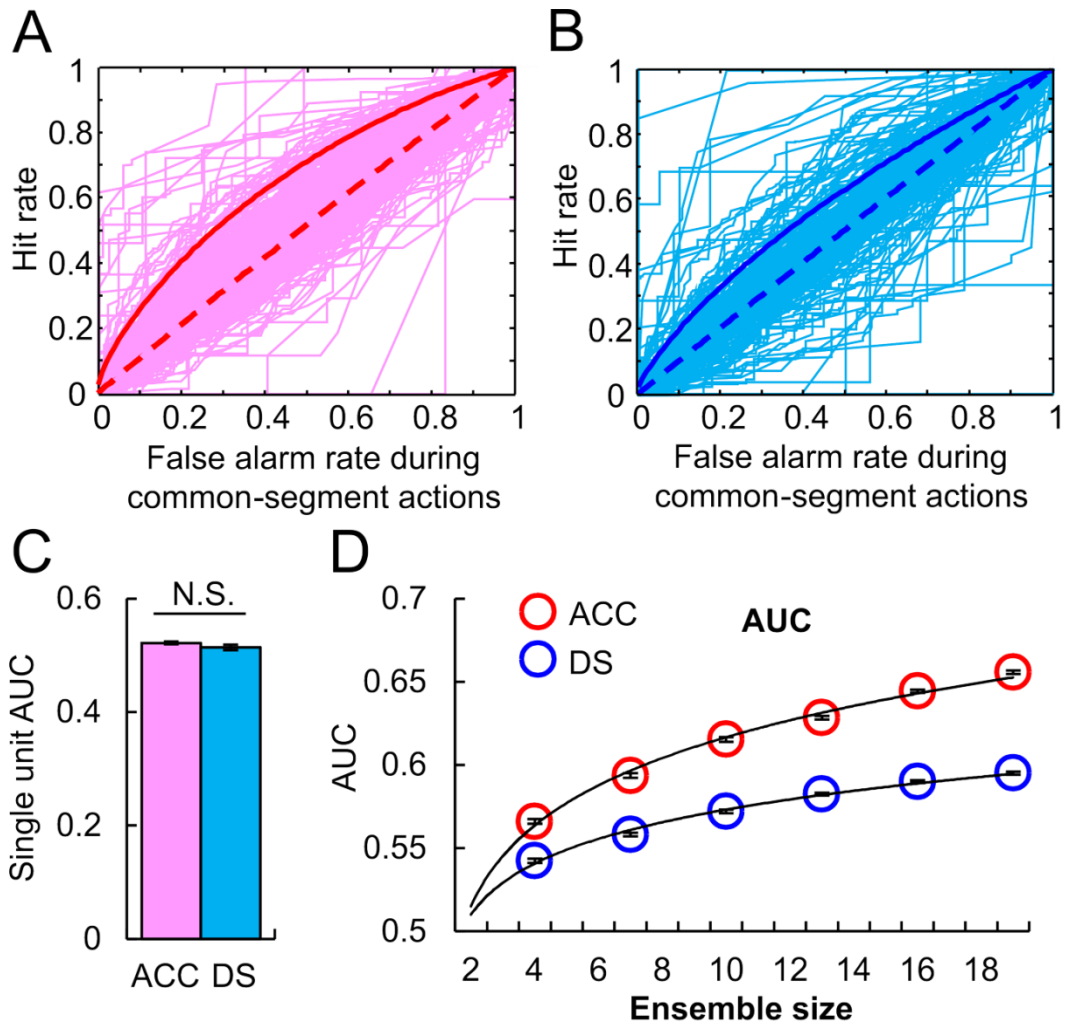
A
B
Hit rate
False alarm rate during
common-segment actions
C
D
N.S.
Single unit AUC
ACC DS
AUC
ACC
DS
Ensemble size

Figure 4.4 Comparison of the signal detection characteristics of single neurons versus ensembles

A) ACC single unit ROCs (pink curves) were constructed from hit and false-alarm rates during the 'common-segment' periods for each neuron. Hits and false-alarms were based on running correlations (1-bin/step) between a 7-bin iFR vector template derived from half of the trials in one sequence block and 7-bin iFR vectors recorded during the remaining half trials in the same sequence-block ('hits') or with 7-bin iFR vectors recorded during the alternate sequence-block ('false alarms'). The ensemble ROC (solid red line) detected the correct sequence better than the majority of the single units or the time bin shuffled ensemble controls (dashed red line). **B)** DS single unit ROCs (light blue curves) were constructed from hit and false-alarm rates during the common-segments using the same methods as in (A). The ensemble ROC (solid dark blue line) detected the correct sequence better than its shuffled control (dashed dark blue line) but was eclipsed by many more single units than in the ACC. **C)** ACC (pink) and DS (blue) single-units classified sequence equally well, as measured by the area-under-curve (AUC) of their respective ROCs. **D)** Ensemble decoding performance improved with increasing ensemble size in both the ACC and DS, but the rate was much higher in the ACC. Ensembles of different sizes were randomly drawn from ACC and DS neuronal populations, and the areas-under-curve (AUCs) were calculated from the sequence signal-detection ROCs. This process was repeated for 100 times at each ensemble size for each region, and the mean and s.e.m. are plotted (ACC: red circles, DS: blue circles). The best fitting trend lines were power functions (ACC: top, DS: bottom), which explained more than 99% of the variance.

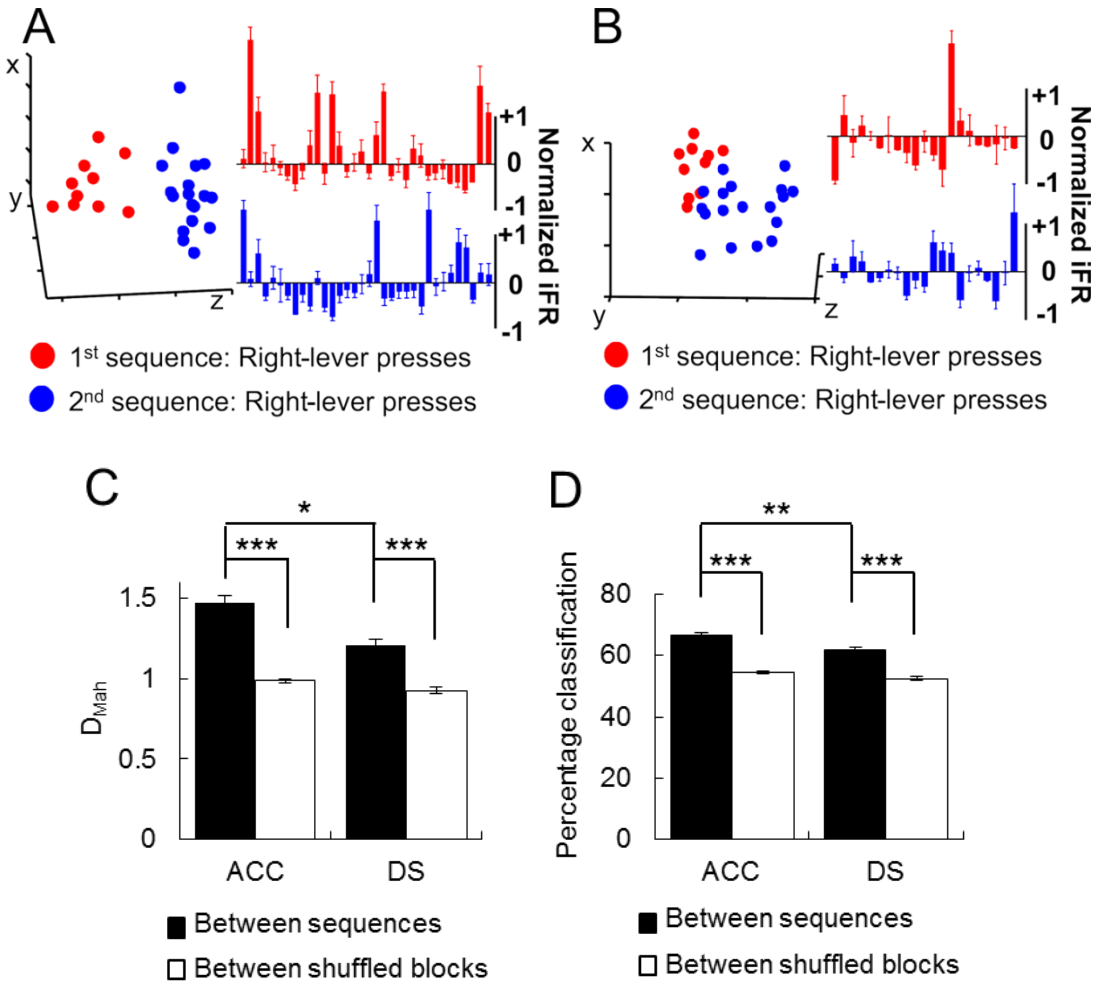
A
x
y
z
+1
0
-1
+1
0
-1
Normalized iFR
1st sequence: Right-lever presses
2nd sequence: Right-lever presses
B
x
y
z
+1
0
-1
+1
0
-1
Normalized iFR
1st sequence: Right-lever presses
2nd sequence: Right-lever presses
C
*
***
***
1.5
1
0.5
0
DMah
ACC
DS
Between sequences
Between shuffled blocks
D
**
***
***
80
60
40
20
0
Percentage classification
ACC
DS
Between sequences
Between shuffled blocks

Figure 4.5 Sequence information is represented as differences in ensemble activity state patterns in both the ACC and DS

A) Ensemble variance calculated based on activities during the time bins within each pair of common-segment lever presses was higher in the DS than in the ACC. Ensemble variance calculated on the residual iFR matrices (right, white) was smaller than the variance calculated from full iFR matrices (right, black) in the DS, but the two were equivalent in the ACC (left bars). **B)** Ensemble covariance was also greater in the DS (right) than in the ACC (left). Ensemble covariance calculated on the residual iFR matrices (white) was smaller than the covariance calculated from full iFR matrices (black) in the DS, but the two were equivalent in the ACC. **C)** In ACC ensembles, between-sequence D_{Mah} calculated using the residuals from the initial behavioral regression was equivalent to that calculated using the full iFRs, and both were greater than shuffled controls. **D)** In ACC ensembles, between-sequence leave-one-out MDA using the residual iFR matrices (right) was equivalent to that calculated using the full iFR matrices (left). **E)** In contrast to the ACC, in DS ensembles, D_{Mah} calculated using the residual iFR matrices from the initial behavioral regression (right) was greater than that calculated using the full iFR matrices (left), although both the full iFRs and the residuals still contained robust sequence information. **F)** In DS ensembles, sequence classification using the residual iFR matrices (right) was more accurate than when calculated using the full iFR matrices (left). Error bars indicate s.e.m. *$p<0.05$, ** $p<0.001$, *** $p<0.00001$

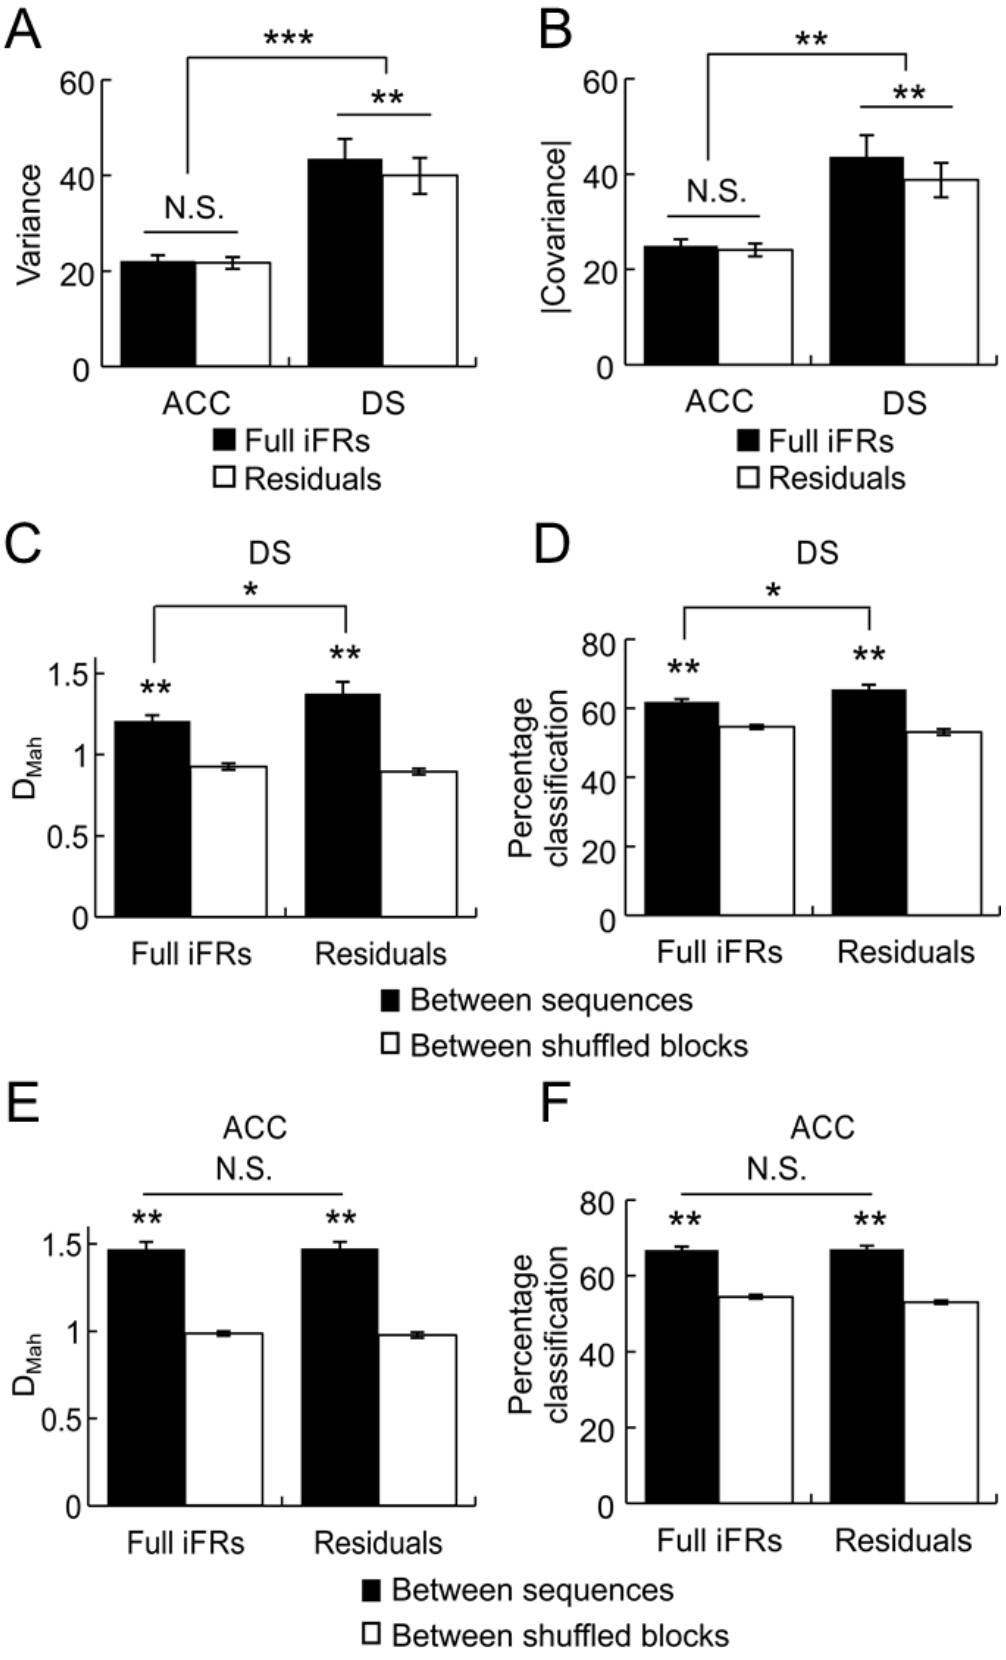
A
***
**
N.S.
Variance
60
40
20
0
ACC
DS
Full iFRs
Residuals
B
**
**
N.S.
|Covariance|
60
40
20
0
ACC
DS
Full iFRs
Residuals
C
DS
*
**
**
D_{Mah}
1.5
1
0.5
0
Full iFRs
Residuals
D
DS
*
**
**
Percentage classification
80
60
40
20
0
Full iFRs
Residuals
Between sequences
Between shuffled blocks
E
ACC
N.S.
**
**
D_{Mah}
1.5
1
0.5
0
Full iFRs
Residuals
F
ACC
N.S.
**
**
Percentage classification
80
60
40
20
0
Full iFRs
Residuals
Between sequences
Between shuffled blocks

Figure 4.6 Comparison of the timing of maximal sequence and lever differentiation in ACC and DS ensembles

A) Ensemble variance calculated based on activities during the time bins within each pair of common-segment lever presses was higher in the DS than in the ACC. In the DS, ensemble variance calculated on the residual iFR matrices (white) was smaller than the variance calculated from full iFR matrices (black). In the ACC, ensemble variance was the same for full iFR or residual iFR matrices. **B)** Ensemble covariance was greater in the DS (right) than in the ACC (left). In the DS, ensemble covariance calculated on the residual iFR matrices (white) was smaller than the covariance calculated from full iFR matrices (black). In the ACC, ensemble covariance was the same for full iFR (left) or residual iFR matrices (right). **C)** In the ACC ensembles, between-sequence D_{Mah} calculated using the residuals from the initial behavioral regression was equivalent to that calculated using the full iFRs. Both the full iFRs and the residuals contained robust sequence information. **D)** In the ACC ensembles, between-sequence leave-one-out MDA performed using the residual iFR matrices (right) was equivalent to that calculated using the full iFR matrices (left). **E)** In contrast to the ACC, in the DS ensembles, D_{Mah} calculated using the residual iFR matrices from the initial behavioral regression (right) was greater than that calculated using the full iFR matrices (left), although both the full iFRs and the residuals contained robust sequence information. **F)** In the DS ensembles, sequence classification using the residual iFR matrices (right) was more accurate than that when calculated using the full iFR matrices (left). *$p<0.05$, ** $p<0.001$, *** $p<0.00001$

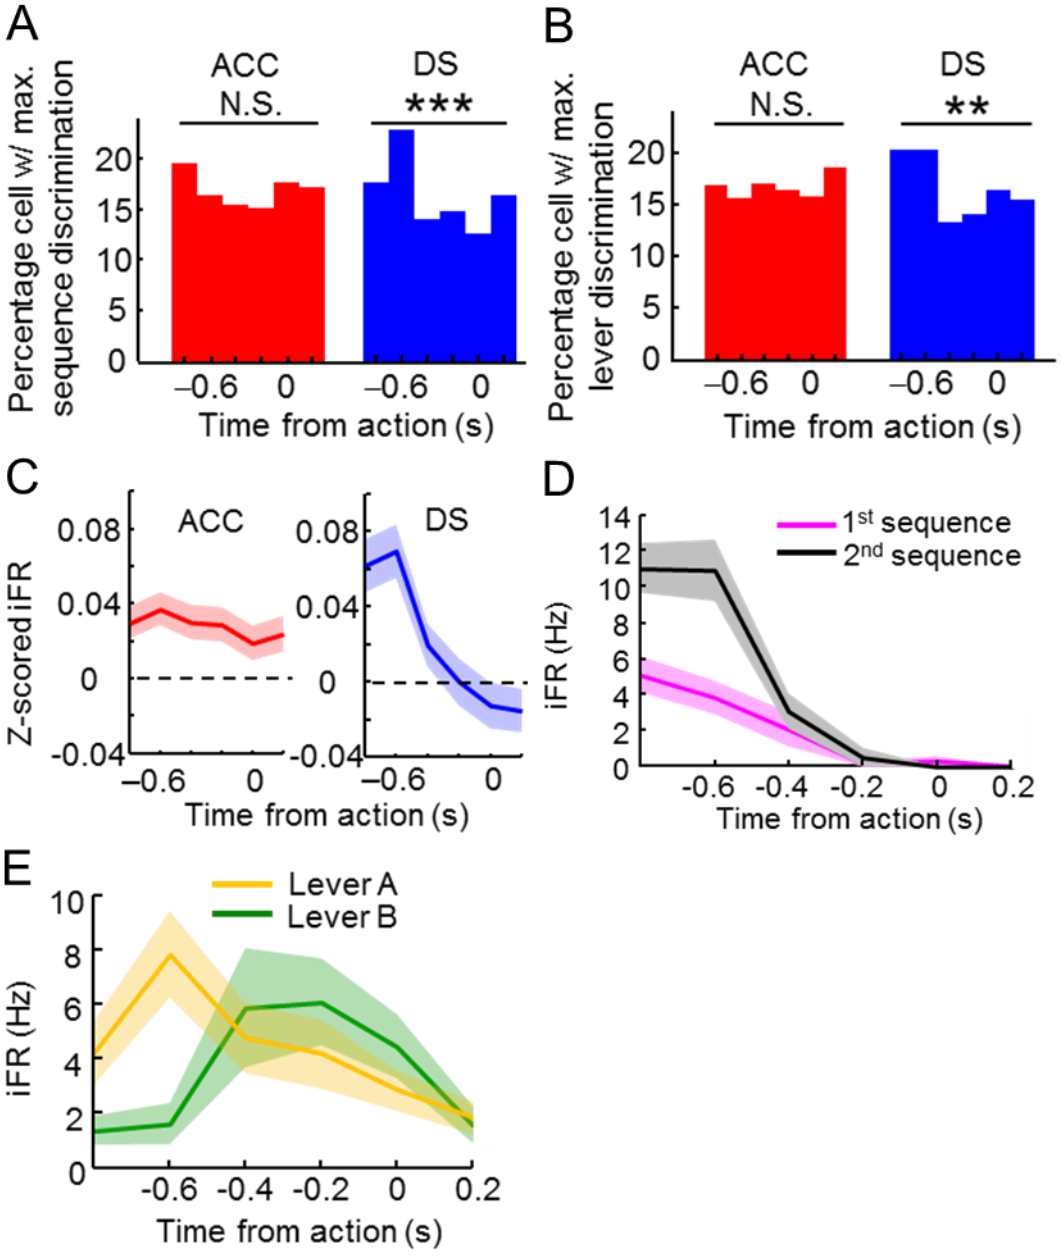
A
ACC
N.S.
DS
***
Percentage cell w/ max. sequence discrimination
20
15
10
5
0
–0.6
0
–0.6
0
Time from action (s)
B
ACC
N.S.
DS
**
Percentage cell w/ max. lever discrimination
20
15
10
5
0
–0.6
0
–0.6
0
Time from action (s)
C
ACC
DS
Z-scored iFR
0.08
0.04
0
-0.04
–0.6
0
–0.6
0
Time from action (s)
D
1st sequence
2nd sequence
iFR (Hz)
14
12
10
8
6
4
2
0
-0.6
-0.4
-0.2
0
0.2
Time from action (s)
E
Lever A
Lever B
iFR (Hz)
10
8
6
4
2
0
-0.6
-0.4
-0.2
0
0.2
Time from action (s)

Figure 4.7 Information is encoded differently through time in the DS but not the ACC

A) A rolling F-statistic was used to determine where neurons maximally separated the two sequences based on firing rate differences in each time bin that constituted the common segment lever press period across sequence blocks. The strength of sequence discrimination fluctuated across bins in the DS (blue bars), where the largest proportion of neurons exhibited maximal sequence specific firing differentiation around 0.6s prior to lever press while the smallest proportion exhibited sequence specific firing differentiation at around the actual lever press response itself. Such was not the case in the ACC (red bars), where an equal percentage of neurons exhibited maximal sequence specific firing differentiation in each of the 6 bins. **B)** A rolling F-statistic was used to determine where neurons maximally differentiated the levers within each sequence block, based on firing rate differences in each of the 6 time bins in each lever-press epoch. The strength of lever discrimination fluctuated across bins in the DS (blue bars), where the largest proportion of neurons exhibited maximal lever differentiation 0.6s prior to lever press. Such was not the case in the ACC (red bars), where an equal percentage of neurons exhibited maximal sequence specific firing differentiation in each of the 6 bins. **C)** Average level of activity across time bins leading up to all leverpresses in all sequence blocks and sessions, in all ACC neurons (left) and the DS neurons (right) respectively. **D)** An example of a DS neuron displaying maximal differentiation between *sequences* at approximately 0.6s prior to presses on the same lever. **E)** An example of a DS neuron reaching maximal differentiation between *levers* at approximately 0.6s prior to lever presses within the same sequence block. *$p<0.01$, **$p<0.001$, *** $p<0.0001$.

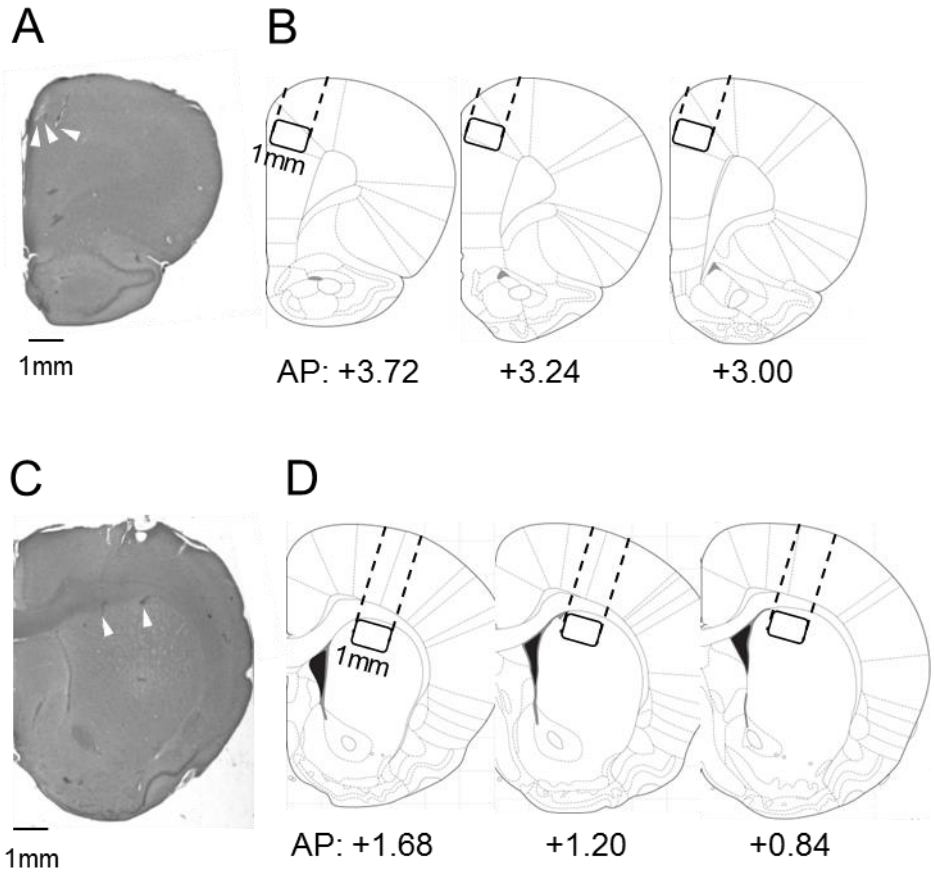

Figure 4.8 Histology showing recording sites

A) Histology showing representative electrode track endings (white arrows) in the ACC. **B)** Schematics showing the location and range of the recording sites in the ACC. Given that not all electrodes left a track, the recording range was inferred based on the visible tracks and the size of the electrode arrays. **C)** Representative electrode track endings (white arrows) in the DS. **D)** Schematics showing the location and range of the recording sites in the DS.

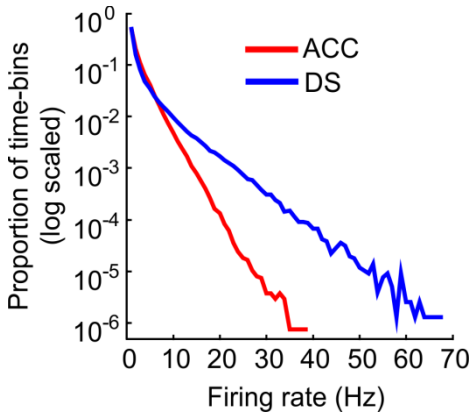

Figure 4.9 Firing properties in the ACC and DS

The iFRs from all bins and neurons were combined for each region, and their probability density functions were plotted on a logarithmic scale. Although both ACC and DS neurons usually have low firing rates, DS neurons exhibited brief periods of high activity more often than did ACC neurons, giving rise to the long-tailed firing rate distribution (DS: blue, ACC: red).

Chapter 5: General discussion: what does the dmPFC encode and how is this achieved?

The function of the primate ACC or rodent dmPFC appears to be heavily intertwined with the functions of other brain regions, and permeates both internally-oriented and externally-directed cognitive domains (Raichle et al., 2001, Fox et al., 2005b). Accordingly, discerning the function of the ACC has been a great challenge. How can one use a single theory to synthesize the many proposed functions of the ACC that include autonomic regulation (Critchley, 2005, Parvizi et al., 2013), nociception and affect (Vogt, 2005, Parvizi et al., 2013), executive functions (Carter et al., 1998, Gehring and Fencsik, 2001, Gasquoine, 2013), attention (Passingham, 1996, Petersen and Posner, 2012) and motor functions (Picard and Strick, 1996, Dum and Strick, 2002)?

Accumulating evidence suggests that the ACC may support several cognitive processes common to many tasks, including their retention in the absence of a defined task (Raichle et al., 2001, Fox et al., 2005b, Dosenbach et al., 2006, Dosenbach et al., 2007). The ACC may monitor and evaluate the progress towards various dynamic goals. Given the multiple coding schemes that a brain region such as the ACC may utilize, keeping track of actions and contexts are not necessarily at conflict with evaluating action and outcome, or responding to conflict or errors.

Based on the findings described in Chapters 2-4, two different coding schemes appear to be employed by the dmPFC: one based on the pattern of firing across the ensemble, and another based on the overall level of activity averaged across the ensemble. While the pattern-based code is suitable for representing discrete events or concepts, the averaged level of

activity is suitable for encoding psychological constructs that can be evaluated along a continuum, such as reward expectation or progress towards a goal. The sequence tasks used in this project provided the unique opportunity to demonstrate the utilization of the two coding schemes simultaneously. But even in the absence of well-defined task events, it is likely that the ACC monitors ongoing experience using these two coding schemes. One reason for making this statement is that the dmPFC is always active and is part of the rodent resting-state functional network recorded in both awake and anesthetized animals (Liang et al., 2011, 2012, Keilholz et al., 2013, Liang et al., 2013).

Self-paced free exploration provides a naturalistic setting in which to study the organization of cognitive space in the dmPFC without imposing a task structure. Intriguingly, the dmPFC displays constant spontaneous activities which encode meaningful information during free exploration or rest in novel or familiar environment (Hyman et al., 2012). Such information may include sensory context and the time spent in the context (Hyman et al., 2012). The time information is manifested as a context-specific and continuously evolving firing-rate pattern in dmPFC ensembles, even when the animals was absent from the environment and when they subsequently revisited the place (Hyman et al., 2012). Thus even though abrupt change in ensemble firing pattern indeed occurred with a change in context (although still without change in overall level of activity), this change is overlaid on the background of constantly evolving encoding of experience (Hyman et al., 2012).

Yet if the ACC is constantly encoding experience, how does this occur given that each neuron has a limited range of firing rates? The constant change in ensemble firing pattern observed in Hyman et al (2012) appears to be due to relatively small but *asynchronous* changes in activity among all neurons. These small changes result in a moment by moment

'refreshing' of on-going representation drawing out of a potentially infinite pool of patterns. This enormous number of patterns can be achieved because of the lack of strong excitatory or inhibitory influence that would synchronize large groups of neurons (Fino and Yuste, 2011, Packer and Yuste, 2011). Physiologically, for the purpose of tracking on-going experience, such a property is much more cost-efficient than altering the activity of large groups of neurons synchronously, or assigning 'responsibility' to separate groups of neurons, which would require active inhibition on all the groups that were not currently 'on duty'.

Therefore, even in the absence of a structured task, the ACC continually refreshes its activity patterns, and this could provide a means to track the entirety of an ever-evolving experience. This constant tracking is useful in the real world situations which are seldom divided into 'on-task' and 'off-task' periods.

5.1 The dmPFC tracks discrete events in structured tasks

The ACC's timely engagement during events of effort, cost, error or conflict indicates that it monitors all task events essential to the achievement of the goal (Carter et al., 1998, Botvinick et al., 1999, Sohn et al., 2007, Emeric et al., 2008, Ichikawa et al., 2011, Apps et al., 2012). Although sustained activation spanning all task epochs is rarely seen in individual dmPFC neurons, by working together, these neurons can 'relay' information between themselves and provide a continuous representation of the task. In this way, representing task events on the ensemble level does not require overall change in activity as the animal transitions from one sequence or context to another (**Fig. 2.3A-C, Fig. 4.1C, Fig. 4.4**). As shown in Chapter 2 (**Fig. 2.6B,F**), this pattern shift without net change in activity levels ('Total ΔFR=0') was found in every functional group of cells examined. Our data indicate that the following events can be encoded with patterned based representations involving no-net

changes in overall activity levels: individually-rewarded operant actions and their respective contexts (**Fig. 5.1**), operant actions linked into a sequence (**Fig. 2.2, 2.3A** and **2.4A**), different overall sequences (**Fig. 2.4A**), sequences associated with different reward amplitudes (Kargo et al., 2007), choices, outcomes and different phases of a spatial delayed win-shift task (Lapish et al., 2008), prior and future outcomes (Baeg et al., 2003), abstract task rules (Rich and Shapiro, 2009, Durstewitz et al., 2010), stimuli and actions in delayed alternation task (Hyman et al., 2013) and delay activity in a reaction time task (Narayanan and Laubach, 2009).

From these studies it is clear that different activity patterns at the ensemble level are able to encode a vast array of different types of events. This is not to be confused with a 'labeled-line' code, where one and only one group of cells is dedicated to a specific type of signal. The pattern-based code in the ACC is distinct from a labeled-line code in at least 3 important ways:

1) The dmPFC neurons are known for their highly mixed selectivity, i.e. dedicated 'sequence A' or 'nose-poke' neurons are rare (Jung et al., 1998, Kennerley and Wallis, 2009, Hayden and Platt, 2010, Cowen et al., 2012, Rigotti et al., 2013). As shown in the previous chapters, whenever a dmPFC neuron responded to an action, its responsiveness was often strongly modulated by the sensory context (**Fig.2.5E**), the overall sequence to which the action belonged (**Fig. 2.5A, Fig. 4.2D,E**) and in the serial position of that action or how close the action was to the goal (**Fig. 3.5**). Furthermore, the responsive of these 'interactive' cells to more than one factor can have high sensitivity to both the preferred action and the preferred overall sequence/context that they overshadowed the 'pure' action cells (**Fig. 2.5C, F**), 'pure' sequence cells (**Fig. 2.5B**) or 'pure' context cells (**Fig. 2.5E**). This kind of information

multiplexing causes individual mPFC neurons to be highly unreliable as isolated action encoders, as they cannot respond to a given action in every sequence or every context. However, collectively they never failed to generate unique action-related ensemble activity states depending on the task context. These observations support the idea that mPFC neurons are capable of binding together related information, such as an action and its task context or its outcome.

2) The majority of dmPFC neurons 'remap' their behavioral sensitivity when the task changes, e.g. a cell sensitive to lever-press is unlikely to remain so in a different task (Jung et al., 1998). Even from the beginning to the end of a period of stable performance without change in strategy, sequence or path of movement, approximately 12.6% of dmPFC neurons significantly altered their firing rate (Rich and Shapiro, 2009). Additionally, this remapping was well balanced, such that the number of neurons losing the correlate during a sequence or context switch was fully compensated by the number of neurons gaining the correlate (**Fig. 2.6**). As a result, on the ensemble level, the relationship among the representations of task components (e.g. different actions) was maintained (**Fig. 2.3F, Fig. 2.4F**), even though they were supported by a different group of neurons after the switch (**Fig. 2.6**). Thus it seems reasonable to suggest that the ACC may have a 'reserve pool' of neurons, which can be flexibly recruited to one of a great variety of behavioral targets (Bernacchia et al., 2011). Such quick shift in firing pattern need not necessarily engage long-term synaptic plasticity, instead it can be achieved via the activation of dopamine D1 receptors, which promote sustained firing while dampen transient depolarization, whereby may stabilize selected behavioral representations (Williams and Goldman-

Rakic, 1995, Durstewitz et al., 2000, Seamans et al., 2001). These findings are consistent with the fact that in any arbitrary task, researchers never fail to find neurons responsive to most of the events, including those which are novel to the subjects (Duncan, 2001). In other words, the dmPFC neurons are 'multi-purpose' computational units, as they can process information from any task that the organism may engage in.

3) Even cells not sensitive to sequence according to multiple linear regression (MLR) during the actions still held detectable sequence information in Chapter 4. As Wohrer et al. (2013) pointed out, the difference in the cells with significant and non-significant behavioral sensitivity is usually not significant itself. Hence neurons that failed to reach the cut-off line may still encode a weak sequence signal. When the responsiveness distribution of cells is bimodal, dividing cells into responsive and non-responsive groups is accurate. However more than often this distribution of responsiveness is unimodal or a simple descending function as is often the case in the dmPFC, the use of arbitrary cut-offs may be practical but also less accurate. This difficulty in defining functional groups reflects an important coding property of the dmPFC cells: information is widely distributed and to a variety of degrees across the population, thus dividing it into responsive and non-responsive groups results in an oversimplified coding scheme.

The firing pattern code in the dmPFC is highly distributed, which gives rise to cognitive and behavioral flexibility, without compromising the stability of each behavioral strategy. Distributed coding has been reported in association sensory cortex (Hung et al., 2005), motor cortex (Georgopoulos et al., 1982, Georgopoulos et al., 1986), as well as in the

frontal cortices (Abe and Lee, 2011, Hyman et al., 2012). In different areas they serve a variety of cognitive functions, namely, perceiving and identifying object, guiding movements, and monitoring and evaluating choices and outcomes. In any case, ensembles employing distributed patterned-based coding contain neurons that can be 'tuned' to any one of a broad range of targets and share a significant portion of these targets across neurons (i.e., information redundancy). Such properties not only allow for resilience against neuronal loss, but also engender an exponentially expanded repertoire of firing patterns to be used for recognizable objects, alternative movements and possible associations that can be made between stimuli, actions and outcomes. Clearly, distributed codes are highly versatile and can be used for both concrete information such as visual forms (Hung et al., 2005) and orientations of movements (Georgopoulos et al., 1986), as well as for abstract concepts such as rules (Rich and Shapiro, 2009, Durstewitz et al., 2010) or the expected outcome of a choice (Sul et al., 2010). Last but not least, distributed pattern-based codes also inherit flexibility from their constituent neurons, which may remap greatly as the task changes (see point 2 above). Such a property (as it occurs in the dmPFC) further expands the cognitive capacity afforded by a distributed coding scheme.

Flexible switching between behaviors or behavioral strategies also requires reliable representations of all the options. Specifically, the organism needs to know which strategy it is abandoning and which it is adopting, and not to confuse the two. Thanks to information redundancy, the distributed coding scheme in the dmPFC does not depend on the response reliability or consistency of any individual neuron. More importantly, asynchrony within the ensemble also contributes to reliable coding, if information permeates widely in the ensemble. As demonstrated in Chapter 4, both dmPFC and DS single neurons held the same amount of

sequence information with similar level of reliability (**Fig. 4.3A,B,C**). Because the DS ensembles had stronger synchrony, their signal tended to weaken or the noise tended to strengthen in the same time bin(s) within trial (blue bars or line, **Fig. 4.7A-C**), they did not reach as high a level of sequence decoding as the dmPFC ensembles (**Fig. 4.4D, Fig. 4.5C-D**). On the other hand, in the dmPFC ensembles, neurons encoded sequences independently through time, whereby maintained sequence information as a population (red bars or line, **Fig. 4.7A-C**). This difference in coding property is not only reflected in sequence encoding, but was also observed in the representation of spatial lever information associated with different actions. Although on the single unit level the DS encoded lever location more strongly than the dmPFC, this difference did not translate into any advantage in the DS on the ensemble level (Chapter 4). Similar to the case of sequence encoding, stronger synchrony and the fact that more than one signal (i.e. sequence and lever location) occurred simultaneously in the DS undermined its ability to encode either at the ensemble level (**Fig. 4.7A-C**). Overall, such differences in coding properties are consistent with the cognitive functions of the two regions. While independence among the sources of information (i.e. single neurons) in the dmPFC is conducive to both diversity and stability of representations, synchronization among DS neurons ensures the output of only the chosen response (Berke et al., 2004, Gage et al., 2010). In short, the expanded coding capacity and robustness against single-neuron failure are some of the critical advantages of distributed codes over sparse codes.

In summary, a firing-rate pattern-based coding scheme involving large groups of multi-purpose neurons may substantiate the capacity of the dmPFC to efficiently parse any arbitrary task, endowing it with considerable flexibility without compromising reliability. Additionally, the proposed continuity in the monitoring and valuation of task events by the

ACC is not in conflict with the event-related 'sharp peaks' in the mean firing rate of selected cells (Quilodran et al., 2008, Kennerley et al., 2009): apart from keeping a running tally of the discrete events in the on-going experience, the ACC may simultaneously use overall activity level to encode information of a different nature, as detailed in section 5.2 below.

Lastly, although not the focus of this thesis, the activity patterns of the DS provide an important insight to the complexity of neural coding which may be conveyed back to the dmPFC and other regions: DS ensembles may alternate between a relatively sparse or more distributed coding schemes by altering its degree of synchrony as well as temporal firing pattern (Carrillo-Reid et al., 2008). This may explain the limited advantage of combining DS neurons into ensembles in sequence decoding (**Fig. 4.3D**), in comparison to the dmPFC. This duality in DS ensemble coding may afford the DS its roles in both decision-making and movement generation (Graybiel, 1995, Pennartz et al., 2009). Therefore, it is not only possible, but indeed highly likely for a given brain region to encode information using different coding schemes to subserve different functions. Neuromodulators such as dopamine and acetylcholine seem to play a pivotal role in such transitions (Durstewitz and Seamans, 2008, Carrillo-Reid et al., 2009, Carrillo-Reid et al., 2011).

5.2 dmPFC continuously monitors progress towards the goal

5.2.1 dmPFC monitors progress rather than actions per se

The sections above dealt with the unique properties of the pattern-based code used by ACC ensembles. I will now deal with the second type of coding scheme detailed in this thesis, the one that relies on variance in the overall level of activity across the ensembles.

The dmPFC not only monitors actions and their consequences, but also evaluates these events with reference to the current goal of the organism. However these two types of information differ in nature: discrete task events are qualitatively distinct from each other, whereas the progress towards the goal must be encoded quantitatively. It follows that the reorganization in firing patterns cannot be the only coding scheme employed by the dmPFC ensembles. When motivational valence changes as the animals make progress towards the goal, dmPFC ensembles are capable of encoding such quantitative change in expected timing/sequential order of reward obtainment along a continuum of averaged level of overall activities (**Fig. 3.2A,C** and **Fig. 3.4B**), which is consistent with previous findings (Shidara and Richmond, 2002). This finding complements previous discoveries that ACC neurons showed anticipatory responses related to reward magnitude (Pratt and Mizumori, 2001, Kargo et al., 2007), and the involvement of ACC activity in both delay and effort during decision making (Hosokawa et al., 2013).

The mechanisms responsible for such continuous change connecting actions to the outcome is likely driven by dopaminergic input from the ventral tegmental area (VTA) (Ahn and Phillips, 2002, Phillips et al., 2004, Rossetti and Carboni, 2005). This process may be initiated by an input of reward-expectancy information from the mPFC to the VTA (Jo et al., 2013), whereby commanding neuro-modulatory influence back to the mPFC providing motivational tone. Additionally, the error analyses in Chapter 3 ruled out the possibility that the ramping in firing rates were driven either by a temporal accumulator model (Niki and Watanabe, 1979, Narayanan and Laubach, 2009), or by an 'action-counting' effect. Rather the progression was halted momentarily by an erroneous action but then continued after the disruption (**Fig. 3.6**). While this progressive changes in ACC ensemble activity could reflect

heightened expectation of a forthcoming reward or an increase in the value of later actions (Shidara and Richmond, 2002, Quilodran et al., 2008, Kennerley et al., 2009, Hayden and Platt, 2010, Sul et al., 2010, Forster and Brown, 2011), it may also reflect an internalized representation of progress. Comparing to previous study by Shidara and Richmond (2002), the 'progress' signal shown in Chapter 3 was smoother and more continuous. This may reflect the self-paced nature of this task design, although future studies are needed to determine whether self-pacing indeed leads to 'smoothness' in the evolution of single cell or ensemble firing rates.

A comparison between **Fig. 3.3** and **Fig. 3.7E,F** reveals a potential difference between the task-related cells in dmPFC and DS: dmPFC cells appeared to respond in a much more continuous and smoother fashion than the DS cells. This directly led to the phenomenon in **Fig. 3.2C** and **Fig. 3.7C**, where an averaged curve of the dmPFC population was smooth whereas that of the DS population was locked to actions. Thus while DS neurons are sensitive to the sequence in which actions are performed, they do not seem to code progress within a sequence in the same dynamic and integrated manner as dmPFC neurons. Furthermore, unlike the dmPFC, DS neurons did not respond to correct lever presses as distinct from incorrect ones. When viewed together, the main signal present in dmPFC ensembles is associated with tracking progression in a manner that is largely abstracted from the encoding of the actual lever presses, while the DS acts as the compliment, tracking lever presses in a more literal fashion. The integrated signal produced by these two interconnected regions would effectively keep the animal on track in its progression towards a goal, even in the face of errors, delays or distractions.

In summary, the dmPFC tracks the cumulative effect of task events, which may allow it to evaluate actions with reference to the current goal. This conclusion is consistent with the literature on ACC's role in evaluating the current choices based on both the predicted outcome (Schall et al., 2002, Kennerley et al., 2009, Luk and Wallis, 2009, Hayden and Platt, 2010, Cowen et al., 2012, Luk and Wallis, 2013) and accumulated experience from the recent past (Baeg et al., 2007, Behrens et al., 2007, Quilodran et al., 2008, Kennerley and Wallis, 2009, Sharot et al., 2010, Sul et al., 2010). Even though published studies mostly focused on the 'sharp' cellular responses with well-defined epochs such as 'instruction', 'choice' and 'outcome', it is more parsimonious to hypothesize that online information is held and processed throughout the task. This is supported by findings in Chapter 3 as well as the continuous spontaneous activities observed in the ACC across species (Kennerley et al., 2009, Cowen et al., 2012).

5.2.2 Activity-level code is independent of firing pattern code but the two coexist

The conclusion that the overall activity-level code is independent of the firing pattern code requires the satisfaction of two premises: 1. Information encoding by activity level can occur without change in ensemble firing pattern. 2. Information encoding by firing pattern can occur without change in the overall level of activity in the ensemble. In support of the first premise, information encoding by overall activity level resulted in little reorganization in the pattern of excitation and inhibition across the ensemble and specific to the action (**Fig. 3.5**). Instead, this pattern remained stable for that particular action, while its *amplitude* was modulated by the serial position, or how close the action was with reference to the goal (**Fig. 3.5**). In support of the second premise, as described in section 5.1, a shift in firing pattern due to sequence or context switch (**Fig. 2.3A-B, Fig. 4.4**) was not accompanied by any change in

the overall level of activity (**Fig. 2.3C , Fig. 4.1C**). Firing pattern change in the absence of change in overall activity has also been observed to encode other types of information, including shifts between place vs. egocentric strategy of T-maze navigation (Rich and Shapiro, 2009) and the choice of different paths associated with high vs. low reward (Kargo et al., 2007). Thus these two types of coding schemes indeed seem to be independent of each other, but exist simultaneously.

Therefore, while the dmPFC keeps track of discrete events (e.g. action, sequence, context) using a patterned-based coding scheme, motivation variables such as the proximity to reward may amplify or modify the strength of these activity state patterns, presumably via the influence of neuromodulators. Albeit heavily innervated by dopaminergic neurons, which may release dopamine in the DS following a ramping pattern in anticipation to reward (Howe et al., 2013), this modulation was not found in the activities of DS ensembles themselves (**Fig. 3.7A,C** and **Fig. 3.8B**), hence was quite unique to the dmPFC.

Together the changes in firing pattern as well as in overall firing rate demonstrate the ACC's enormous capacity for ongoing representation of experience either in structured tasks or during free exploration. Such dynamic representation may be conducive to a number of computations, including but not limited to evaluating an action with its predicted outcome given the task context, calculating each set of cost-benefit trade-off, detecting an error, and anticipating an error-prone phase in the task. When shifts in firing pattern are used to encode new information, the organism's experience can be continuously tracked without consuming extra energy on driving the population excitatory/inhibitory (E/I) ratio away from the baseline. However a change in the E/I ratio may become necessary in the case of an error (Niki and Watanabe, 1979, Pardo et al., 1990, Bernstein et al., 1995, Holroyd and Coles, 2002, Ito et al.,

2003), a conflict (Botvinick et al., 2001, Haddon and Killcross, 2006, van Veen et al., 2009) or a re-direction of attention (Mesulam, 1981, Petersen and Posner, 2012), giving rise to a metabolic signal detectable by imaging techniques.

5.3 Towards a unifying theory of dmPFC function

Researchers have postulated several integrative theories to synthesize ACC functions: 1. The ACC 'assesses the motivational content of internal and external stimuli and regulates context-dependent behaviors' (Devinsky et al., 1995); 2. The ACC is necessary for 'cognitive control' (Shenhav et al., 2013) or monitors behavior and detects when such control or intervention is necessary (Carter et al., 1999, Carter et al., 2000, Cohen et al., 2000, Gehring and Knight, 2000, Mansouri et al., 2009); 3. The ACC is an evaluator of effort/cost and the associated expected outcome for directing behavioral choice (Walton et al., 2003) or for sustaining the already-made choices (Cowen et al., 2012, Parvizi et al., 2013); 4. The ACC keeps track of a baseline state and detects environmental change by producing a bodily response which brings the 'feeling' away from the baseline, and helps elicit a behavioral response (Critchley, 2003, Critchley, 2005). All theories have received support, but there is no reason that they should remain separate.

On the basis of all evidence discussed so far, ACC ensembles appear to have a richness of potential activity patterns which may enable the region to keep an uninterrupted and evaluative representation of the organism's experience with reference to one or more dynamic goals. Coloring this running account of the current experience is the motivational tone from neuromodulator(s). By applying this 'color' according to immediate feedback and the anticipated goal, the ACC can differentiate between the representation of a correct,

effective response and an error. In short, the ACC appears to be a structure that monitors the experience while actively evaluating the current state with reference to the goal state.

In addition to monitor experience, there is also the need to flexibly shift focus or even to alternate between streams of thoughts as required to attain a goal. This process requires a reconstructing and/or reorganization of the current events in light of the new task rule, which was observed in the dmPFC as changes in ensemble activity patterns, even though the actions or paths of movement per se remained the same (Durstewitz et al., 2010, Kargo et al., 2007). Specifically, the dmPFC ensembles reconstructed the representation of the same action via a shift in ensemble activity pattern when the light cues no longer indicated the correct lever, but spatial location of the lever now did (Durstewitz et al., 2010). This applies to the real world, where animals often use different cues to navigate in different weather or day/night conditions. Alternatively, as the value of the outcome of a given set of behaviors or movements increase or decrease, the ensemble activity pattern shifted accordingly to reflect this change (Kargo et al., 2007). Again this process is highly adaptive as for an animal a water source may be drained or replenished within days, and for a driver the traffic on a highway may alternate between heavy and light several times a day.

Furthermore, not only must the behavioral strategy adapt to the changing environment, but the goal itself, derived from the need of the individual, also changes frequently for animals in their natural habitat and for people in the real world. In order to survive, an organism must simultaneously satisfy several fundamental needs, including but not limited to nutrients, warmth, safety and social interaction. Thus while one may focus on one goal at a given time, the relative urgency among all potential goals must be constantly monitored. Although rarely does a study involve the choice among goals of different nature (e.g. food vs. water), given

that the ACC keeps track of stimuli and events in the environment with motivational salience attached, it may be speculated that it also monitors the ranking of the relative urgency among one's many needs.

Of course, the ACC accomplishes these functions in concert with many other brain regions. For instance, human dlPFC and orbitofrontal cortex (OFC) were both involved in effort-based but not delay-based decision making, whereas ACC neurons responded to both (Hosokawa et al. 2013). Additionally, while the dlPFC may integrate response with expected outcome, the OFC integrated multiple variables contributing to the value of the outcome (Wallis, 2007), and updated values based on current motivational state (Rudebeck et al., 2013) and recent reward history (Wallis and Kennerley, 2011). Meanwhile the ACC was suggested to monitor and integrate information regarding motivationally salient events such as physical and psychological pain, anxiety, as well as all the variables influencing the value of the outcome. In fact, some have suggested that the ACC may facilitate the dedication of additional resource from other areas such as the dlPFC (Gehring and Knight, 2000, Carter and van Veen, 2007).

Not only does the ACC work in concert with other brain regions but subregions within the ACC may serve different functions (Vogt et al., 1992). Recent studies on functional connectivity identified different sub-regions in the ACC that respectively belongs to task-positive and task-negative networks (Dosenbach et al., 2007, Vincent et al., 2008). Therefore, it seems that although the ACC supports both off-task, self-referential and introspective cognition (Raichle et al., 2001) and on-task, goal-directed effortful behavior (Dosenbach et al., 2006), different sub-regions are in fact in charge of them respectively (Fox et al., 2005b, Golland et al., 2007). These findings may explain why surgical ablation of different

subregions within the ACC was effective for different psychiatric disorders, for instance, the effective target area is more caudal for treating chronic pain than for treating obsessive-compulsive disorder (Brotis et al., 2009). Additionally, because other areas within the prefrontal cortex, such as the dlPFC and OFC, can also provide cognitive resources to achieve behavioral flexibility and resolve conflicts, removing part of the ACC did not result in any apparent loss of function. However, the small impairment in focused and selective attention, and the increased reaction time in cognitively demanding tasks (Cohen et al., 1999a, Cohen et al., 1999b) in cingulotomy patients may reflect deficits in the active and continuous monitoring function, and the timely dedication of extra cognitive resources usually mediated by the ACC.

5.4 The role of the ACC in psychiatric conditions

Abnormal activity within the ACC has been implicated in a number of psychiatric disorders, including unipolar depression (Pizzagalli et al., 2001, Hamilton et al., 2011, Nejad et al., 2013, Yoshimura et al., 2013), OCD (Anticevic et al., 2013, Cheng et al., 2013, Kaufmann et al., 2013, de Wit et al., 2014), schizophrenia (Carter et al., 2001, Kerns et al., 2005, White et al., 2011, Ikuta et al., 2012, Kraguljac et al., 2012), borderline personality disorder (Minzenberg et al., 2007, Silbersweig et al., 2007, Mak and Lam, 2013, Koenigsberg et al., 2014), autism spectrum disorders (Thakkar et al., 2008, Tesink et al., 2009, Assaf et al., 2010, Eilam-Stock et al., 2014), among others. In some cases, structural changes in the ACC were also observed in patients with schizophrenia (Velakoulis et al., 2002, Marquardt et al., 2005, Byun et al., 2012, Shepherd et al., 2012) and major depression (Sacher et al., 2012, Grieve et al., 2013). The pathological processes within the ACC are likely to be different depending on the disorder, although there appears to be common themes in the functional

impairment across these conditions. One of such themes is the impaired cognitive and behavioral flexibility (Gu et al., 2008, Floresco et al., 2009, Geurts et al., 2009, Lederbogen et al., 2011, Remijnse et al., 2013), and another is the impaired coping strategy, or cognitive resource allocation when faced with stress, anxiety and psychological pain (Matheson and Anisman, 2003, Southwick et al., 2005, Jones and Fernyhough, 2007).

Another intriguing aspect of schizophrenia is the impairment in an interconnected set of cognitive functions including self-referential processing (Subramaniam et al., 2012), source/agency attribution (Werner et al., 2014), and in Sense of Agency (SoA) (Bulot et al., 2007, Jeannerod, 2009, Schimansky et al., 2010, Maeda et al., 2012, Maeda et al., 2013, Moore et al., 2013). Given that the ACC has a pivotal role in monitoring responses and their outcome as well as self-related cognition (Qin et al., 2010, Qin and Northoff, 2011, Qin et al., 2012), it is likely that the ACC is an important element of the circuitry supporting SoA (Raos et al., 2007). Additionally, abnormal activities in the ACC both at rest (Greicius, 2008, Whitfield-Gabrieli and Ford, 2012) and during cognitive tasks (Tu et al., 2010, White et al., 2010, Hasenkamp et al., 2011) are observed in schizophrenia patients. Thus ACC dysfunction may contribute to SoA symptoms in schizophrenia. Support for this comes from a recent clinical study in which schizophrenia patients were trained for 80hrs on a self-referential processing (Subramaniam et al., 2012). The task involved determining whether a word was generated by oneself or by the experimenter and this engaged the ACC in healthy participants but not in patients with schizophrenia. Importantly however, following training the on-task activation in the ACC of patients resembled that seen in healthy controls and this activation was correlated with their long-term improvement (Subramaniam et al., 2012). Therefore, further consideration of the role of ACC in self-related source attribution, and specifically

SoA may shed light on the pathological cognitive processes underlying psychosis and other forms of mental illness leading eventually to the development of more effective therapies.

5.5 Conclusion

The tools and techniques used in this thesis provide a new perspective on how ACC neurons encode information and how the information carried by these neurons generates highly flexible ensemble level representations. More importantly, by studying ACC ensemble activities one can begin to move beyond theories of the ACC based on a limited set of tasks or processes and in turn arrive at a clearer picture of ACC function in both normal and disease states.

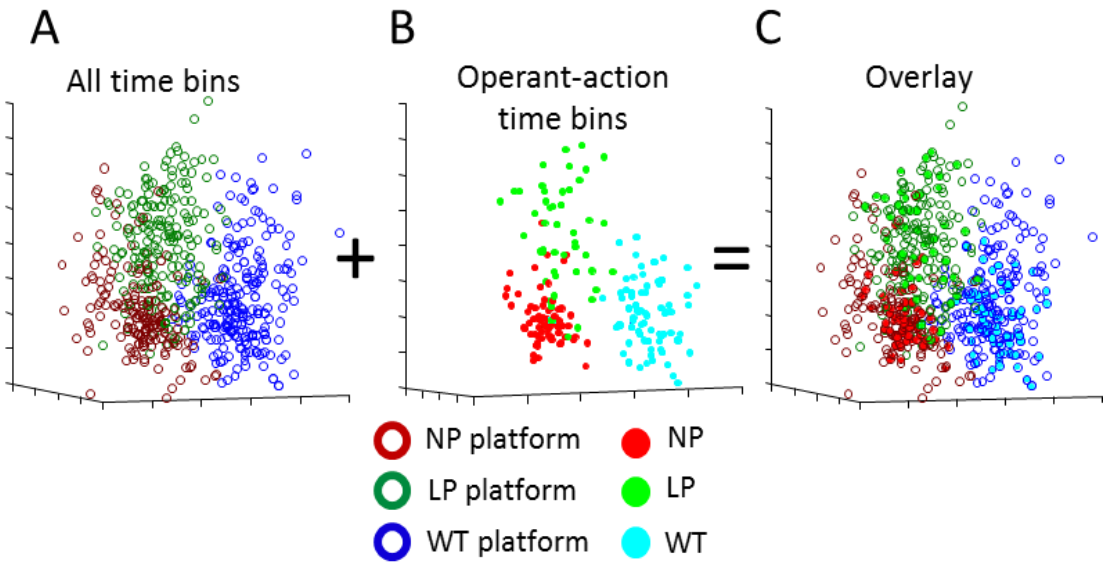

Figure 5.1 Ensemble activity states of independently-rewarded operant actions, each in a unique context

A) In 3 different contexts on the 4-platform Maze, animals spent 3 consecutive task periods, respectively. A sample of the ensemble activity was taken every 5s, plus the activities during the action periods were plotted in the MSUA space (1-s time bins from nose-poke task period—dark red circles, from lever-press task period—dark green circles, from wheel-turn task period—dark blue circles). Activities during the 3 task-periods formed 3 distinct states in the MSUA space. **B)** Activities during the 3 different operant actions also formed separate states (nose-pokes—red dots, lever-presses—green dots, wheel-turns—blue dots). **C)** An overlay of A and B shows that the action-related activity states appeared to locate within the "domains" demarcated by their respective context-states.

References

Abe H, Lee D (2011) Distributed coding of actual and hypothetical outcomes in the orbital and dorsolateral prefrontal cortex. Neuron 70:731-741.

Ahn S, Phillips AG (2002) Modulation by central and basolateral amygdalar nuclei of dopaminergic correlates of feeding to satiety in the rat nucleus accumbens and medial prefrontal cortex. J Neurosci 22:10958-10965.

Ainsworth M, Lee S, Cunningham MO, Traub RD, Kopell NJ, Whittington MA (2012) Rates and rhythms: a synergistic view of frequency and temporal coding in neuronal networks. Neuron 75:572-583.

Aldridge JW, Berridge KC (1998) Coding of serial order by neostriatal neurons: a "natural action" approach to movement sequence. J Neurosci 18:2777-2787.

Alexander GE, Crutcher MD (1990) Functional architecture of basal ganglia circuits: neural substrates of parallel processing. Trends Neurosci 13:266-271.

Alexander WH, Brown JW (2011) Medial prefrontal cortex as an action-outcome predictor. Nat Neurosci 14:1338-1344.

Alitto HJ, Weyand TG, Usrey WM (2005) Distinct properties of stimulus-evoked bursts in the lateral geniculate nucleus. J Neurosci 25:514-523.

Angelucci A, Levitt JB, Lund JS (2002) Anatomical origins of the classical receptive field and modulatory surround field of single neurons in macaque visual cortical area V1. Prog Brain Res 136:373-388.

Anticevic A, Hu S, Zhang S, Savic A, Billingslea E, Wasylink S, Repovs G, Cole MW, Bednarski S, Krystal JH, Bloch MH, Li CS, Pittenger C (2013) Global Resting-State Functional Magnetic Resonance Imaging Analysis Identifies Frontal Cortex, Striatal, and Cerebellar Dysconnectivity in Obsessive-Compulsive Disorder. Biol Psychiatry.

Antonini A, Moresco RM, Gobbo C, De Notaris R, Panzacchi A, Barone P, Calzetti S, Negrotti A, Pezzoli G, Fazio F (2001) The status of dopamine nerve terminals in Parkinson's disease and essential tremor: a PET study with the tracer [11-C]FE-CIT. Neurol Sci 22:47-48.

Apps MA, Balsters JH, Ramnani N (2012) The anterior cingulate cortex: monitoring the outcomes of others' decisions. Soc Neurosci 7:424-435.

Assaf M, Jagannathan K, Calhoun VD, Miller L, Stevens MC, Sahl R, O'Boyle JG, Schultz RT, Pearlson GD (2010) Abnormal functional connectivity of default mode sub-networks in autism spectrum disorder patients. Neuroimage 53:247-256.

Aultman JM, Moghaddam B (2001) Distinct contributions of glutamate and dopamine receptors to temporal aspects of rodent working memory using a clinically relevant task. Psychopharmacology (Berl) 153:353-364.

Averbeck BB, Chafee MV, Crowe DA, Georgopoulos AP (2002) Parallel processing of serial movements in prefrontal cortex. Proc Natl Acad Sci U S A 99:13172-13177.

Averbeck BB, Chafee MV, Crowe DA, Georgopoulos AP (2003a) Neural activity in prefrontal cortex during copying geometrical shapes. I. Single cells encode shape, sequence, and metric parameters. Exp Brain Res 150:127-141.

Averbeck BB, Crowe DA, Chafee MV, Georgopoulos AP (2003b) Neural activity in prefrontal cortex during copying geometrical shapes. II. Decoding shape segments from neural ensembles. Exp Brain Res 150:142-153.

Averbeck BB, Lee D (2006) Effects of noise correlations on information encoding and decoding. J Neurophysiol 95:3633-3644.
Averbeck BB, Lee D (2007) Prefrontal neural correlates of memory for sequences. J Neurosci 27:2204-2211.
Averbeck BB, Sohn JW, Lee D (2006) Activity in prefrontal cortex during dynamic selection of action sequences. Nat Neurosci 9:276-282.
Aziz-Zadeh L, Kaplan JT, Iacoboni M (2009) "Aha!": The neural correlates of verbal insight solutions. Hum Brain Mapp 30:908-916.
Baeg EH, Kim YB, Huh K, Mook-Jung I, Kim HT, Jung MW (2003) Dynamics of population code for working memory in the prefrontal cortex. Neuron 40:177-188.
Baeg EH, Kim YB, Kim J, Ghim JW, Kim JJ, Jung MW (2007) Learning-induced enduring changes in functional connectivity among prefrontal cortical neurons. J Neurosci 27:909-918.
Bailey KR, Mair RG (2006) The role of striatum in initiation and execution of learned action sequences in rats. J Neurosci 26:1016-1025.
Baleydier C, Mauguiere F (1980) The duality of the cingulate gyrus in monkey. Neuroanatomical study and functional hypothesis. Brain 103:525-554.
Balleine BW, Dickinson A (1998) Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. Neuropharmacology 37:407-419.
Balleine BW, O'Doherty JP (2010) Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. Neuropsychopharmacology 35:48-69.
Barbas H, Pandya DN (1987) Architecture and frontal cortical connections of the premotor cortex (area 6) in the rhesus monkey. J Comp Neurol 256:211-228.
Barone P, Joseph JP (1989) Prefrontal cortex and spatial sequencing in macaque monkey. Exp Brain Res 78:447-464.
Bechara A, Damasio H, Tranel D, Damasio AR (1997) Deciding advantageously before knowing the advantageous strategy. Science 275:1293-1295.
Behrens TE, Woolrich MW, Walton ME, Rushworth MF (2007) Learning the value of information in an uncertain world. Nat Neurosci 10:1214-1221.
Berdyyeva TK, Olson CR (2010) Rank signals in four areas of macaque frontal cortex during selection of actions and objects in serial order. J Neurophysiol 104:141-159.
Berke JD, Okatan M, Skurski J, Eichenbaum HB (2004) Oscillatory entrainment of striatal neurons in freely moving rats. Neuron 43:883-896.
Bernacchia A, Seo H, Lee D, Wang XJ (2011) A reservoir of time constants for memory traces in cortical neurons. Nat Neurosci 14:366-372.
Bernstein PS, Scheffers MK, Coles MG (1995) "Where did I go wrong?" A psychophysiological analysis of error detection. J Exp Psychol Hum Percept Perform 21:1312-1322.
Bossert JM, Stern AL, Theberge FR, Cifani C, Koya E, Hope BT, Shaham Y (2011) Ventral medial prefrontal cortex neuronal ensembles mediate context-induced relapse to heroin. Nat Neurosci 14:420-422.
Botvinick M, Nystrom LE, Fissell K, Carter CS, Cohen JD (1999) Conflict monitoring versus selection-for-action in anterior cingulate cortex. Nature 402:179-181.
Botvinick MM, Braver TS, Barch DM, Carter CS, Cohen JD (2001) Conflict monitoring and cognitive control. Psychol Rev 108:624-652.

Bowers JS (2009) On the biological plausibility of grandmother cells: implications for neural network theories in psychology and neuroscience. Psychol Rev 116:220-251.
Britten KH, Newsome WT, Saunders RC (1992) Effects of inferotemporal cortex lesions on form-from-motion discrimination in monkeys. Exp Brain Res 88:292-302.
Brotis AG, Kapsalaki EZ, Paterakis K, Smith JR, Fountas KN (2009) Historic evolution of open cingulectomy and stereotactic cingulotomy in the management of medically intractable psychiatric disorders, pain and drug addiction. Stereotact Funct Neurosurg 87:271-291.
Buckner RL, Andrews-Hanna JR, Schacter DL (2008) The brain's default network: anatomy, function, and relevance to disease. Ann N Y Acad Sci 1124:1-38.
Bulot V, Thomas P, Delevoye-Turrell Y (2007) A pre-reflective indicator of an impaired sense of agency in patients with Schizophrenia. Exp Brain Res 183:115-126.
Bussey TJ, Everitt BJ, Robbins TW (1997a) Dissociable effects of cingulate and medial frontal cortex lesions on stimulus-reward learning using a novel Pavlovian autoshaping procedure for the rat: implications for the neurobiology of emotion. Behav Neurosci 111:908-919.
Bussey TJ, Muir JL, Everitt BJ, Robbins TW (1997b) Triple dissociation of anterior cingulate, posterior cingulate, and medial frontal cortices on visual discrimination tasks using a touchscreen testing procedure for the rat. Behav Neurosci 111:920-936.
Byun MS, Kim JS, Jung WH, Jang JH, Choi JS, Kim SN, Choi CH, Chung CK, An SK, Kwon JS (2012) Regional cortical thinning in subjects with high genetic loading for schizophrenia. Schizophr Res 141:197-203.
Calabresi P, Misgeld U, Dodt HU (1987) Intrinsic membrane properties of neostriatal neurons can account for their low level of spontaneous activity. Neuroscience 20:293-303.
Caracheo BF, Emberly E, Hadizadeh S, Hyman JM, Seamans JK (2013) Abrupt changes in the patterns and complexity of anterior cingulate cortex activity when food is introduced into an environment. Front Neurosci 7:74.
Carrillo-Reid L, Hernandez-Lopez S, Tapia D, Galarraga E, Bargas J (2011) Dopaminergic modulation of the striatal microcircuit: receptor-specific configuration of cell assemblies. J Neurosci 31:14972-14983.
Carrillo-Reid L, Tecuapetla F, Tapia D, Hernandez-Cruz A, Galarraga E, Drucker-Colin R, Bargas J (2008) Encoding network states by striatal cell assemblies. J Neurophysiol 99:1435-1450.
Carrillo-Reid L, Tecuapetla F, Vautrelle N, Hernandez A, Vergara R, Galarraga E, Bargas J (2009) Muscarinic enhancement of persistent sodium current synchronizes striatal medium spiny neurons. J Neurophysiol 102:682-690.
Carter CS, Botvinick MM, Cohen JD (1999) The contribution of the anterior cingulate cortex to executive processes in cognition. Rev Neurosci 10:49-57.
Carter CS, Braver TS, Barch DM, Botvinick MM, Noll D, Cohen JD (1998) Anterior cingulate cortex, error detection, and the online monitoring of performance. Science 280:747-749.
Carter CS, Macdonald AM, Botvinick M, Ross LL, Stenger VA, Noll D, Cohen JD (2000) Parsing executive processes: strategic vs. evaluative functions of the anterior cingulate cortex. Proc Natl Acad Sci U S A 97:1944-1948.

Carter CS, MacDonald AW, 3rd, Ross LL, Stenger VA (2001) Anterior cingulate cortex activity and impaired self-monitoring of performance in patients with schizophrenia: an event-related fMRI study. Am J Psychiatry 158:1423-1428.
Carter CS, van Veen V (2007) Anterior cingulate cortex and conflict detection: an update of theory and data. Cogn Affect Behav Neurosci 7:367-379.
Chang JY, Chen L, Luo F, Shi LH, Woodward DJ (2002) Neuronal responses in the frontal cortico-basal ganglia system during delayed matching-to-sample task: ensemble recording in freely moving rats. Exp Brain Res 142:67-80.
Cheng Y, Xu J, Nie B, Luo C, Yang T, Li H, Lu J, Xu L, Shan B, Xu X (2013) Abnormal resting-state activities and functional connectivities of the anterior and the posterior cortexes in medication-naive patients with obsessive-compulsive disorder. PLoS One 8:e67478.
Chiba AA, Kesner RP, Gibson CJ (1997) Memory for temporal order of new and familiar spatial location sequences: role of the medial prefrontal cortex. Learn Mem 4:311-317.
Clarke HF, Dalley JW, Crofts HS, Robbins TW, Roberts AC (2004) Cognitive inflexibility after prefrontal serotonin depletion. Science 304:878-880.
Clower WT, Alexander GE (1998) Movement sequence-related activity reflecting numerical order of components in supplementary and presupplementary motor areas. J Neurophysiol 80:1562-1566.
Cohen JD, Botvinick M, Carter CS (2000) Anterior cingulate and prefrontal cortex: who's in control? Nat Neurosci 3:421-423.
Cohen RA, Kaplan RF, Moser DJ, Jenkins MA, Wilkinson H (1999a) Impairments of attention after cingulotomy. Neurology 53:819-824.
Cohen RA, Kaplan RF, Zuffante P, Moser DJ, Jenkins MA, Salloway S, Wilkinson H (1999b) Alteration of intention and self-initiated action associated with bilateral anterior cingulotomy. J Neuropsychiatry Clin Neurosci 11:444-453.
Cole MW, Yeung N, Freiwald WA, Botvinick M (2009) Cingulate cortex: diverging data from humans and monkeys. Trends Neurosci 32:566-574.
Collins P, Roberts AC, Dias R, Everitt BJ, Robbins TW (1998) Perseveration and strategy in a novel spatial self-ordered sequencing task for nonhuman primates: effects of excitotoxic lesions and dopamine depletions of the prefrontal cortex. J Cogn Neurosci 10:332-354.
Compte A, Constantinidis C, Tegner J, Raghavachari S, Chafee MV, Goldman-Rakic PS, Wang XJ (2003) Temporally irregular mnemonic persistent activity in prefrontal neurons of monkeys during a delayed response task. J Neurophysiol 90:3441-3454.
Corbetta M, Patel G, Shulman GL (2008) The reorienting system of the human brain: from environment to theory of mind. Neuron 58:306-324.
Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. Nat Rev Neurosci 3:201-215.
Corbit LH, Balleine BW (2003) The role of prelimbic cortex in instrumental conditioning. Behav Brain Res 146:145-157.
Cowen SL, Davis GA, Nitz DA (2012) Anterior cingulate neurons in the rat map anticipated effort and reward to their associated action sequences. J Neurophysiol 107:2393-2407.
Cowen SL, McNaughton BL (2007) Selective delay activity in the medial prefrontal cortex of the rat: contribution of sensorimotor information and contingency. J Neurophysiol 98:303-316.

Critchley H (2003) Emotion and its disorders. Br Med Bull 65:35-47.
Critchley HD (2005) Neural mechanisms of autonomic, affective, and cognitive integration. J Comp Neurol 493:154-166.
Critchley HD, Elliott R, Mathias CJ, Dolan RJ (2000) Neural activity relating to generation and representation of galvanic skin conductance responses: a functional magnetic resonance imaging study. J Neurosci 20:3033-3040.
Critchley HD, Mathias CJ, Josephs O, O'Doherty J, Zanini S, Dewar BK, Cipolotti L, Shallice T, Dolan RJ (2003) Human cingulate cortex and autonomic control: converging neuroimaging and clinical evidence. Brain 126:2139-2152.
Cromer JA, Roy JE, Miller EK (2010) Representation of multiple, independent categories in the primate prefrontal cortex. Neuron 66:796-807.
Croxson PL, Walton ME, O'Reilly JX, Behrens TE, Rushworth MF (2009) Effort-based cost-benefit valuation and the human brain. J Neurosci 29:4531-4541.
Damasio AR (1996) The somatic marker hypothesis and the possible functions of the prefrontal cortex. Philos Trans R Soc Lond B Biol Sci 351:1413-1420.
de Wit S, Kosaki Y, Balleine BW, Dickinson A (2006) Dorsomedial prefrontal cortex resolves response conflict in rats. J Neurosci 26:5224-5229.
de Wit SJ, Alonso P, Schweren L, Mataix-Cols D, Lochner C, Menchon JM, Stein DJ, Fouche JP, Soriano-Mas C, Sato JR, Hoexter MQ, Denys D, Nakamae T, Nishida S, Kwon JS, Jang JH, Busatto GF, Cardoner N, Cath DC, Fukui K, Jung WH, Kim SN, Miguel EC, Narumoto J, Phillips ML, Pujol J, Remijnse PL, Sakai Y, Shin NY, Yamada K, Veltman DJ, van den Heuvel OA (2014) Multicenter voxel-based morphometry mega-analysis of structural brain scans in obsessive-compulsive disorder. Am J Psychiatry 171:340-349.
Delatour B, Gisquet-Verrier P (2001) Involvement of the dorsal anterior cingulate cortex in temporal behavioral sequencing: subregional analysis of the medial prefrontal cortex in rat. Behav Brain Res 126:105-114.
Derdikman D, Moser EI (2010) A manifold of spatial maps in the brain. Trends Cogn Sci 14:561-569.
Devinsky O, Morrell MJ, Vogt BA (1995) Contributions of anterior cingulate cortex to behaviour. Brain 118 (Pt 1):279-306.
Di Pietro NC, Seamans JK (2011) Dopamine and serotonin interactively modulate prefrontal cortex neurons in vitro. Biol Psychiatry 69:1204-1211.
Dosenbach NU, Fair DA, Miezin FM, Cohen AL, Wenger KK, Dosenbach RA, Fox MD, Snyder AZ, Vincent JL, Raichle ME, Schlaggar BL, Petersen SE (2007) Distinct brain networks for adaptive and stable task control in humans. Proc Natl Acad Sci U S A 104:11073-11078.
Dosenbach NU, Visscher KM, Palmer ED, Miezin FM, Wenger KK, Kang HC, Burgund ED, Grimes AL, Schlaggar BL, Petersen SE (2006) A core system for the implementation of task sets. Neuron 50:799-812.
Dum RP, Strick PL (2002) Motor areas in the frontal lobe of the primate. Physiol Behav 77:677-682.
Duncan J (2001) An adaptive coding model of neural function in prefrontal cortex. Nat Rev Neurosci 2:820-829.
Durstewitz D (2003) Self-organizing neural integrator predicts interval times through climbing activity. J Neurosci 23:5342-5353.

Durstewitz D, Seamans JK (2008) The dual-state theory of prefrontal cortex dopamine function with relevance to catechol-o-methyltransferase genotypes and schizophrenia. Biol Psychiatry 64:739-749.
Durstewitz D, Seamans JK, Sejnowski TJ (2000) Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex. J Neurophysiol 83:1733-1750.
Durstewitz D, Vittoz NM, Floresco SB, Seamans JK (2010) Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. Neuron 66:438-448.
Eilam-Stock T, Xu P, Cao M, Gu X, Van Dam NT, Anagnostou E, Kolevzon A, Soorya L, Park Y, Siller M, He Y, Hof PR, Fan J (2014) Abnormal autonomic and associated brain activities during rest in autism spectrum disorder. Brain 137:153-171.
Emeric EE, Brown JW, Leslie M, Pouget P, Stuphorn V, Schall JD (2008) Performance monitoring local field potentials in the medial frontal cortex of primates: anterior cingulate cortex. J Neurophysiol 99:759-772.
Endepols H, Sommer S, Backes H, Wiedermann D, Graf R, Hauber W (2010) Effort-based decision making in the rat: an [18F]fluorodeoxyglucose micro positron emission tomography study. J Neurosci 30:9708-9714.
Euston DR, McNaughton BL (2006) Apparent encoding of sequential context in rat medial prefrontal cortex is accounted for by behavioral variability. J Neurosci 26:13143-13155.
Eyherabide HG, Rokem A, Herz AV, Samengo I (2009) Bursts generate a non-reducible spike-pattern code. Front Neurosci 3:8-14.
Fellows LK, Farah MJ (2005) Is anterior cingulate cortex necessary for cognitive control? Brain 128:788-796.
Fino E, Yuste R (2011) Dense inhibitory connectivity in neocortex. Neuron 69:1188-1203.
Floresco SB, Zhang Y, Enomoto T (2009) Neural circuits subserving behavioral flexibility and their relevance to schizophrenia. Behav Brain Res 204:396-409.
Forster SE, Brown JW (2011) Medial prefrontal cortex predicts and evaluates the timing of action outcomes. Neuroimage 55:253-265.
Fox AS, Oakes TR, Shelton SE, Converse AK, Davidson RJ, Kalin NH (2005a) Calling for help is independently modulated by brain systems underlying goal-directed behavior and threat perception. Proc Natl Acad Sci U S A 102:4176-4179.
Fox MD, Corbetta M, Snyder AZ, Vincent JL, Raichle ME (2006) Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. Proc Natl Acad Sci U S A 103:10046-10051.
Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, Raichle ME (2005b) The human brain is intrinsically organized into dynamic, anticorrelated functional networks. Proc Natl Acad Sci U S A 102:9673-9678.
Frankland PW, Bontempi B (2005) The organization of recent and remote memories. Nat Rev Neurosci 6:119-130.
Frankland PW, Bontempi B, Talton LE, Kaczmarek L, Silva AJ (2004) The involvement of the anterior cingulate cortex in remote contextual fear memory. Science 304:881-883.
Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2001) Categorical representation of visual stimuli in the primate prefrontal cortex. Science 291:312-316.
Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2002) Visual categorization and the primate prefrontal cortex: neurophysiology and behavior. J Neurophysiol 88:929-941.

Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2003) A comparison of primate prefrontal and inferior temporal cortices during visual categorization. J Neurosci 23:5235-5246.
Freiwald WA, Tsao DY (2010) Functional compartmentalization and viewpoint generalization within the macaque face-processing system. Science 330:845-851.
Fujii N, Graybiel AM (2003) Representation of action sequence boundaries by macaque prefrontal cortical neurons. Science 301:1246-1249.
Fujii N, Graybiel AM (2005) Time-varying covariance of neural activities recorded in striatum and frontal cortex as monkeys perform sequential-saccade tasks. Proc Natl Acad Sci U S A 102:9032-9037.
Gabbott PL, Warner TA, Jays PR, Salway P, Busby SJ (2005) Prefrontal cortex in the rat: projections to subcortical autonomic, motor, and limbic centers. J Comp Neurol 492:145-177.
Gage GJ, Stoetzner CR, Wiltschko AB, Berke JD (2010) Selective activation of striatal fast-spiking interneurons during choice execution. Neuron 67:466-479.
Gasquoine PG (2013) Localization of function in anterior cingulate cortex: from psychosurgery to functional neuroimaging. Neurosci Biobehav Rev 37:340-348.
Gehring WJ, Fencsik DE (2001) Functions of the medial frontal cortex in the processing of conflict and errors. J Neurosci 21:9430-9437.
Gehring WJ, Knight RT (2000) Prefrontal-cingulate interactions in action monitoring. Nat Neurosci 3:516-520.
Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J Neurosci 2:1527-1537.
Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. Science 233:1416-1419.
Geurts HM, Corbett B, Solomon M (2009) The paradox of cognitive flexibility in autism. Trends Cogn Sci 13:74-82.
Golland Y, Bentin S, Gelbard H, Benjamini Y, Heller R, Nir Y, Hasson U, Malach R (2007) Extrinsic and intrinsic systems in the posterior cortex of the human brain revealed during natural sensory stimulation. Cereb Cortex 17:766-777.
Grafton ST, Hazeltine E, Ivry R (1995) Functional mapping of sequence learning in normal humans. J Cogn Neurosci 7:497-510.
Graybiel AM (1995) Building action repertoires: memory and learning functions of the basal ganglia. Curr Opin Neurobiol 5:733-741.
Graybiel AM, Aosaki T, Flaherty AW, Kimura M (1994) The basal ganglia and adaptive motor control. Science 265:1826-1831.
Greene JD, Nystrom LE, Engell AD, Darley JM, Cohen JD (2004) The neural bases of cognitive conflict and control in moral judgment. Neuron 44:389-400.
Greicius M (2008) Resting-state functional connectivity in neuropsychiatric disorders. Curr Opin Neurol 21:424-430.
Greicius MD, Krasnow B, Reiss AL, Menon V (2003) Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. Proc Natl Acad Sci U S A 100:253-258.
Grieve SM, Korgaonkar MS, Koslow SH, Gordon E, Williams LM (2013) Widespread reductions in gray matter volume in depression. Neuroimage Clin 3:332-339.

Grillner S, Hellgren J, Menard A, Saitoh K, Wikstrom MA (2005a) Mechanisms for selection of basic motor programs--roles for the striatum and pallidum. Trends Neurosci 28:364-370.
Grillner S, Markram H, De Schutter E, Silberberg G, LeBeau FE (2005b) Microcircuits in action--from CPGs to neocortex. Trends Neurosci 28:525-533.
Groenewegen HJ, Uylings HB (2000) The prefrontal cortex and the integration of sensory, limbic and autonomic information. Prog Brain Res 126:3-28.
Groenewegen HJ, Wright CI, Uylings HB (1997) The anatomical relationships of the prefrontal cortex with limbic structures and the basal ganglia. J Psychopharmacol 11:99-106.
Gross CG (2002) Genealogy of the "grandmother cell". Neuroscientist 8:512-518.
Gu BM, Park JY, Kang DH, Lee SJ, Yoo SY, Jo HJ, Choi CH, Lee JM, Kwon JS (2008) Neural correlates of cognitive inflexibility during task-switching in obsessive-compulsive disorder. Brain 131:155-164.
Gusnard DA, Raichle ME (2001) Searching for a baseline: functional imaging and the resting human brain. Nat Rev Neurosci 2:685-694.
Haddon JE, Killcross S (2006) Prefrontal cortex lesions disrupt the contextual control of response conflict. J Neurosci 26:2933-2940.
Hamaguchi K, Riehle A, Brunel N (2011) Estimating network parameters from combined dynamics of firing rate and irregularity of single neurons. J Neurophysiol 105:487-500.
Hamilton JP, Chen G, Thomason ME, Schwartz ME, Gotlib IH (2011) Investigating neural primacy in Major Depressive Disorder: multivariate Granger causality analysis of resting-state fMRI time-series data. Mol Psychiatry 16:763-772.
Hasenkamp W, James GA, Boshoven W, Duncan E (2011) Altered engagement of attention and default networks during target detection in schizophrenia. Schizophr Res 125:169-173.
Hayden BY, Pearson JM, Platt ML (2009) Fictive reward signals in the anterior cingulate cortex. Science 324:948-950.
Hayden BY, Platt ML (2010) Neurons in anterior cingulate cortex multiplex information about reward and action. J Neurosci 30:3339-3346.
Heidbreder CA, Groenewegen HJ (2003) The medial prefrontal cortex in the rat: evidence for a dorso-ventral distinction based upon functional and anatomical characteristics. Neurosci Biobehav Rev 27:555-579.
Hillman KL, Bilkey DK (2010) Neurons in the rat anterior cingulate cortex dynamically encode cost-benefit in a spatial decision-making task. J Neurosci 30:7705-7713.
Hillman KL, Bilkey DK (2012) Neural encoding of competitive effort in the anterior cingulate cortex. Nat Neurosci 15:1290-1297.
Holroyd CB, Coles MG (2002) The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. Psychol Rev 109:679-709.
Holroyd CB, Coles MG, Nieuwenhuis S (2002) Medial prefrontal cortex and error potentials. Science 296:1610-1611 author reply 1610-1611.
Holroyd CB, Nieuwenhuis S, Yeung N, Nystrom L, Mars RB, Coles MG, Cohen JD (2004) Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. Nat Neurosci 7:497-498.
Hoover WB, Vertes RP (2007) Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. Brain Struct Funct 212:149-179.

Horst NK, Laubach M (2013) Reward-related activity in the medial prefrontal cortex is driven by consumption. Front Neurosci 7:56.
Horton JC, Adams DL (2005) The cortical column: a structure without a function. Philos Trans R Soc Lond B Biol Sci 360:837-862.
Hoshi E, Sawamura H, Tanji J (2005) Neurons in the rostral cingulate motor area monitor multiple phases of visuomotor behavior with modest parametric selectivity. J Neurophysiol 94:640-656.
Hosokawa T, Kennerley SW, Sloan J, Wallis JD (2013) Single-neuron mechanisms underlying cost-benefit analysis in frontal cortex. J Neurosci 33:17385-17397.
Howe MW, Tierney PL, Sandberg SG, Phillips PE, Graybiel AM (2013) Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. Nature 500:575-579.
Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J Physiol 160:106-154.
Hung CP, Kreiman G, Poggio T, DiCarlo JJ (2005) Fast readout of object identity from macaque inferior temporal cortex. Science 310:863-866.
Hutchins KD, Martino AM, Strick PL (1988) Corticospinal projections from the medial wall of the hemisphere. Exp Brain Res 71:667-672.
Hutchison RM, Mirsattari SM, Jones CK, Gati JS, Leung LS (2010) Functional networks in the anesthetized rat brain revealed by independent component analysis of resting-state FMRI. J Neurophysiol 103:3398-3406.
Hutchison RM, Womelsdorf T, Gati JS, Leung LS, Menon RS, Everling S (2012) Resting-state connectivity identifies distinct functional networks in macaque cingulate cortex. Cereb Cortex 22:1294-1308.
Hyman JM, Hasselmo ME, Seamans JK (2011) What is the Functional Relevance of Prefrontal Cortex Entrainment to Hippocampal Theta Rhythms? Front Neurosci 5:24.
Hyman JM, Ma L, Balaguer-Ballester E, Durstewitz D, Seamans JK (2012) Contextual encoding by ensembles of medial prefrontal cortex neurons. Proc Natl Acad Sci U S A 109:5086-5091.
Hyman JM, Whitman J, Emberly E, Woodward TS, Seamans JK (2013) Action and outcome activity state patterns in the anterior cingulate cortex. Cereb Cortex 23:1257-1268.
Hyman JM, Zilli EA, Paley AM, Hasselmo ME (2005) Medial prefrontal cortex cells show dynamic modulation with the hippocampal theta rhythm dependent on behavior. Hippocampus 15:739-749.
Hyman JM, Zilli EA, Paley AM, Hasselmo ME (2010) Working Memory Performance Correlates with Prefrontal-Hippocampal Theta Interactions but not with Prefrontal Neuron Firing Rates. Front Integr Neurosci 4:2.
Ichikawa N, Siegle GJ, Jones NP, Kamishima K, Thompson WK, Gross JJ, Ohira H (2011) Feeling bad about screwing up: emotion regulation and action monitoring in the anterior cingulate cortex. Cogn Affect Behav Neurosci 11:354-371.
Ikuta T, Szeszko PR, Gruner P, DeRosse P, Gallego J, Malhotra AK (2012) Abnormal anterior cingulate cortex activity predicts functional disability in schizophrenia. Schizophr Res 137:267-268.
Isomura Y, Takada M (2004) Neural mechanisms of versatile functions in primate anterior cingulate cortex. Rev Neurosci 15:279-291.

Ito S, Stuphorn V, Brown JW, Schall JD (2003) Performance monitoring by the anterior cingulate cortex during saccade countermanding. Science 302:120-122.
Jeannerod M (2009) The sense of agency and its disturbances in schizophrenia: a reappraisal. Exp Brain Res 192:527-532.
Jenkins IH, Brooks DJ, Nixon PD, Frackowiak RS, Passingham RE (1994) Motor sequence learning: a study with positron emission tomography. J Neurosci 14:3775-3790.
Jiang X, Bradley E, Rini RA, Zeffiro T, Vanmeter J, Riesenhuber M (2007) Categorization training results in shape- and category-selective human neural plasticity. Neuron 53:891-903.
Jo YS, Lee J, Mizumori SJ (2013) Effects of prefrontal cortical inactivation on neural activity in the ventral tegmental area. J Neurosci 33:8159-8171.
Joel D (2001) Open interconnected model of basal ganglia-thalamocortical circuitry and its relevance to the clinical syndrome of Huntington's disease. Mov Disord 16:407-423.
Jones SR, Fernyhough C (2007) A new look at the neural diathesis--stress model of schizophrenia: the primacy of social-evaluative and uncontrollable situations. Schizophr Bull 33:1171-1177.
Jueptner M, Stephan KM, Frith CD, Brooks DJ, Frackowiak RS, Passingham RE (1997) Anatomy of motor learning. I. Frontal cortex and attention to action. J Neurophysiol 77:1313-1324.
Jung MW, Qin Y, McNaughton BL, Barnes CA (1998) Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. Cereb Cortex 8:437-450.
Kara IO, Gonlusen G, Sahin B, Ergin M, Erdogan S (2005) A general aspect on soft-tissue sarcoma and c-kit expression in primitive neuroectodermal tumor and Ewing's sarcoma. Is there any role in disease process? Saudi Med J 26:1190-1196.
Kargo WJ, Szatmary B, Nitz DA (2007) Adaptation of prefrontal cortical firing patterns and their fidelity to changes in action-reward contingencies. J Neurosci 27:3548-3559.
Kaufmann C, Beucke JC, Preusse F, Endrass T, Schlagenhauf F, Heinz A, Juckel G, Kathmann N (2013) Medial prefrontal brain activation to anticipated reward and loss in obsessive-compulsive disorder. Neuroimage Clin 2:212-220.
Keilholz SD, Magnuson ME, Pan WJ, Willis M, Thompson GJ (2013) Dynamic properties of functional connectivity in the rodent. Brain Connect 3:31-40.
Kennerley SW, Dahmubed AF, Lara AH, Wallis JD (2009) Neurons in the frontal lobe encode the value of multiple decision variables. J Cogn Neurosci 21:1162-1178.
Kennerley SW, Wallis JD (2009) Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. Eur J Neurosci 29:2061-2073.
Kennerley SW, Walton ME, Behrens TE, Buckley MJ, Rushworth MF (2006) Optimal decision making and the anterior cingulate cortex. Nat Neurosci 9:940-947.
Kerns JG, Cohen JD, MacDonald AW, 3rd, Johnson MK, Stenger VA, Aizenstein H, Carter CS (2005) Decreased conflict- and error-related activity in the anterior cingulate cortex in subjects with schizophrenia. Am J Psychiatry 162:1833-1839.
Kesner RP, Holbrook T (1987) Dissociation of item and order spatial memory in rats following medial prefrontal cortex lesions. Neuropsychologia 25:653-664.
Killcross S, Coutureau E (2003) Coordination of actions and habits in the medial prefrontal cortex of rats. Cereb Cortex 13:400-408.

Koenigsberg HW, Denny BT, Fan J, Liu X, Guerreri S, Mayson SJ, Rimsky L, New AS, Goodman M, Siever LJ (2014) The neural correlates of anomalous habituation to negative emotional pictures in borderline and avoidant personality disorder patients. Am J Psychiatry 171:82-90.

Kolb B, Milner B (1981) Observations on spontaneous facial expression after focal cerebral excisions and after intracarotid injection of sodium amytal. Neuropsychologia 19:505-514.

Koos T, Tepper JM (1999) Inhibitory control of neostriatal projection neurons by GABAergic interneurons. Nat Neurosci 2:467-472.

Kraguljac NV, Reid M, White D, Jones R, den Hollander J, Lowman D, Lahti AC (2012) Neurometabolites in schizophrenia and bipolar disorder - a systematic review and meta-analysis. Psychiatry Res 203:111-125.

Krahe R, Gabbiani F (2004) Burst firing in sensory systems. Nat Rev Neurosci 5:13-23.

Krzanowski WJ (2000) Principles of multivariate analysis : a user's perspective. Oxford Oxfordshire ; New York: Oxford University Press.

Lapish CC, Durstewitz D, Chandler LJ, Seamans JK (2008) Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex. Proc Natl Acad Sci U S A 105:11963-11968.

Lauwereyns J, Watanabe K, Coe B, Hikosaka O (2002) A neural correlate of response bias in monkey caudate nucleus. Nature 418:413-417.

Lederbogen F, Kirsch P, Haddad L, Streit F, Tost H, Schuch P, Wust S, Pruessner JC, Rietschel M, Deuschle M, Meyer-Lindenberg A (2011) City living and urban upbringing affect neural social stress processing in humans. Nature 474:498-501.

Liang Z, King J, Zhang N (2011) Uncovering intrinsic connectional architecture of functional networks in awake rat brain. J Neurosci 31:3776-3783.

Liang Z, King J, Zhang N (2012) Anticorrelated resting-state functional connectivity in awake rat brain. Neuroimage 59:1190-1199.

Liang Z, Li T, King J, Zhang N (2013) Mapping thalamocortical networks in rat brain using resting-state functional connectivity. Neuroimage 83:237-244.

Lu X, Ashe J (2005) Anticipatory activity in primary motor cortex codes memorized movement sequences. Neuron 45:967-973.

Luk CH, Wallis JD (2009) Dynamic encoding of responses and outcomes by neurons in medial prefrontal cortex. J Neurosci 29:7526-7539.

Luk CH, Wallis JD (2013) Choice coding in frontal cortex during stimulus-guided or action-guided decision-making. J Neurosci 33:1864-1871.

Ma L, Hyman JM, Phillips AG, Seamans JK (2014) Tracking progress toward a goal in corticostriatal ensembles. J Neurosci 34:2244-2253.

MacDonald AW, 3rd, Cohen JD, Stenger VA, Carter CS (2000) Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. Science 288:1835-1838.

Maeda T, Kato M, Muramatsu T, Iwashita S, Mimura M, Kashima H (2012) Aberrant sense of agency in patients with schizophrenia: forward and backward over-attribution of temporal causality during intentional action. Psychiatry Res 198:1-6.

Maeda T, Takahata K, Muramatsu T, Okimura T, Koreki A, Iwashita S, Mimura M, Kato M (2013) Reduced sense of agency in chronic schizophrenia with predominant negative symptoms. Psychiatry Res 209:386-392.

Mailly P, Aliane V, Groenewegen HJ, Haber SN, Deniau JM (2013) The rat prefrontostriatal system analyzed in 3D: evidence for multiple interacting functional units. J Neurosci 33:5718-5727.
Maimon G, Assad JA (2009) Beyond Poisson: increased spike-time regularity across primate parietal cortex. Neuron 62:426-440.
Mak AD, Lam LC (2013) Neurocognitive profiles of people with borderline personality disorder. Curr Opin Psychiatry 26:90-96.
Mansouri FA, Tanaka K, Buckley MJ (2009) Conflict-induced behavioural adjustment: a clue to the executive functions of the prefrontal cortex. Nat Rev Neurosci 10:141-152.
Margulies DS, Kelly AM, Uddin LQ, Biswal BB, Castellanos FX, Milham MP (2007) Mapping the functional connectivity of anterior cingulate cortex. Neuroimage 37:579-588.
Marquardt RK, Levitt JG, Blanton RE, Caplan R, Asarnow R, Siddarth P, Fadale D, McCracken JT, Toga AW (2005) Abnormal development of the anterior cingulate in childhood-onset schizophrenia: a preliminary quantitative MRI study. Psychiatry Res 138:221-233.
Marsh R, Alexander GM, Packard MG, Zhu H, Wingard JC, Quackenbush G, Peterson BS (2004) Habit learning in Tourette syndrome: a translational neuroscience approach to a developmental psychopathology. Arch Gen Psychiatry 61:1259-1268.
Martinez-Conde S, Macknik SL, Hubel DH (2002) The function of bursts of spikes during visual fixation in the awake primate lateral geniculate nucleus and primary visual cortex. Proc Natl Acad Sci U S A 99:13920-13925.
Matheson K, Anisman H (2003) Systems of coping associated with dysphoria, anxiety and depressive illness: a multivariate profile perspective. Stress 6:223-234.
Matsumoto K, Suzuki W, Tanaka K (2003) Neuronal correlates of goal-based motor selection in the prefrontal cortex. Science 301:229-232.
Maviel T, Durkin TP, Menzaghi F, Bontempi B (2004) Sites of neocortical reorganization critical for remote spatial memory. Science 305:96-99.
McFarland NR, Haber SN (2000) Convergent inputs from thalamic motor nuclei and frontal cortical areas to the dorsal striatum in the primate. J Neurosci 20:3798-3813.
Mesulam MM (1981) A cortical network for directed attention and unilateral neglect. Ann Neurol 10:309-325.
Meyers EM, Freedman DJ, Kreiman G, Miller EK, Poggio T (2008) Dynamic population coding of category information in inferior temporal and prefrontal cortex. J Neurophysiol 100:1407-1419.
Milner B (1982) Some cognitive effects of frontal-lobe lesions in man. Philos Trans R Soc Lond B Biol Sci 298:211-226.
Mink JW (1996) The basal ganglia: focused selection and inhibition of competing motor programs. Prog Neurobiol 50:381-425.
Minzenberg MJ, Fan J, New AS, Tang CY, Siever LJ (2007) Fronto-limbic dysfunction in response to facial emotion in borderline personality disorder: an event-related fMRI study. Psychiatry Res 155:231-243.
Moore JW, Cambridge VC, Morgan H, Giorlando F, Adapa R, Fletcher PC (2013) Time, action and psychosis: using subjective time to investigate the effects of ketamine on sense of agency. Neuropsychologia 51:377-384.

Muakkassa KF, Strick PL (1979) Frontal lobe inputs to primate motor cortex: evidence for four somatotopically organized 'premotor' areas. Brain Res 177:176-182.
Muhammad R, Wallis JD, Miller EK (2006) A comparison of abstract rules in the prefrontal cortex, premotor cortex, inferior temporal cortex, and striatum. J Cogn Neurosci 18:974-989.
Mulder AB, Nordquist RE, Orgut O, Pennartz CM (2003) Learning-related changes in response patterns of prefrontal neurons during instrumental conditioning. Behav Brain Res 146:77-88.
Mushiake H, Saito N, Sakamoto K, Itoyama Y, Tanji J (2006) Activity in the lateral prefrontal cortex reflects multiple steps of future events in action plans. Neuron 50:631-641.
Nagai Y, Critchley HD, Featherstone E, Trimble MR, Dolan RJ (2004) Activity in ventromedial prefrontal cortex covaries with sympathetic skin conductance level: a physiological account of a "default mode" of brain function. Neuroimage 22:243-251.
Nakamura K, Sakai K, Hikosaka O (1998) Neuronal activity in medial frontal cortex during learning of sequential procedures. J Neurophysiol 80:2671-2687.
Narayanan NS, Horst NK, Laubach M (2006) Reversible inactivations of rat medial prefrontal cortex impair the ability to wait for a stimulus. Neuroscience 139:865-876.
Narayanan NS, Laubach M (2006) Top-down control of motor cortex ensembles by dorsomedial prefrontal cortex. Neuron 52:921-931.
Narayanan NS, Laubach M (2008) Neuronal correlates of post-error slowing in the rat dorsomedial prefrontal cortex. J Neurophysiol 100:520-525.
Narayanan NS, Laubach M (2009) Delay activity in rodent frontal cortex during a simple reaction time task. J Neurophysiol 101:2859-2871.
Nauta WJ (1971) The problem of the frontal lobe: a reinterpretation. J Psychiatr Res 8:167-187.
Nejad AB, Fossati P, Lemogne C (2013) Self-Referential Processing, Rumination, and Cortical Midline Structures in Major Depression. Front Hum Neurosci 7:666.
Niki H, Watanabe M (1979) Prefrontal and cingulate unit activity during timing behavior in the monkey. Brain Res 171:213-224.
Ninokura Y, Mushiake H, Tanji J (2004) Integration of temporal order and object information in the monkey lateral prefrontal cortex. J Neurophysiol 91:555-560.
Ongur D, An X, Price JL (1998) Prefrontal cortical projections to the hypothalamus in macaque monkeys. J Comp Neurol 401:480-505.
Ongur D, Ferry AT, Price JL (2003) Architectonic subdivision of the human orbital and medial prefrontal cortex. J Comp Neurol 460:425-449.
Ongur D, Price JL (2000) The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. Cereb Cortex 10:206-219.
Ostlund SB, Balleine BW (2005) Lesions of medial prefrontal cortex disrupt the acquisition but not the expression of goal-directed learning. J Neurosci 25:7763-7770.
Ostlund SB, Winterbauer NE, Balleine BW (2009) Evidence of action sequence chunking in goal-directed instrumental conditioning and its dependence on the dorsomedial prefrontal cortex. J Neurosci 29:8280-8287.
Oswald AM, Doiron B, Maler L (2007) Interval coding. I. Burst interspike intervals as indicators of stimulus intensity. J Neurophysiol 97:2731-2743.
Packer AM, Peterka DS, Hirtz JJ, Prakash R, Deisseroth K, Yuste R (2012) Two-photon optogenetics of dendritic spines and neural circuits. Nat Methods 9:1202-1205.

Packer AM, Yuste R (2011) Dense, unspecific connectivity of neocortical parvalbumin-positive interneurons: a canonical microcircuit for inhibition? J Neurosci 31:13260-13271.
Pardo JV, Pardo PJ, Janer KW, Raichle ME (1990) The anterior cingulate cortex mediates processing selection in the Stroop attentional conflict paradigm. Proc Natl Acad Sci U S A 87:256-259.
Parthasarathy HB, Graybiel AM (1997) Cortically driven immediate-early gene expression reflects modular influence of sensorimotor cortex on identified striatal neurons in the squirrel monkey. J Neurosci 17:2477-2491.
Parvizi J, Rangarajan V, Shirer WR, Desai N, Greicius MD (2013) The will to persevere induced by electrical stimulation of the human cingulate gyrus. Neuron 80:1359-1367.
Passingham RE (1996) Attention to action. Philos Trans R Soc Lond B Biol Sci 351:1473-1479.
Patterson JC, 2nd, Ungerleider LG, Bandettini PA (2002) Task-independent functional brain activity correlation with skin conductance changes: an fMRI study. Neuroimage 17:1797-1806.
Paus T, Koski L, Caramanos Z, Westbury C (1998) Regional differences in the effects of task difficulty and motor output on blood flow response in the human anterior cingulate cortex: a review of 107 PET activation studies. Neuroreport 9:R37-47.
Paxinos G, Watson C (2005) The rat brain in stereotaxic coordinates. Amsterdam ; Boston: Elsevier Academic Press.
Penfield W, Evans J (1935) The frontal lobe in man: A clinical study of maximum removals. Brain 58:115-133.
Pennartz CM, Berke JD, Graybiel AM, Ito R, Lansink CS, van der Meer M, Redish AD, Smith KS, Voorn P (2009) Corticostriatal Interactions during Learning, Memory Processing, and Decision Making. J Neurosci 29:12831-12838.
Petersen SE, Posner MI (2012) The attention system of the human brain: 20 years after. Annu Rev Neurosci 35:73-89.
Petersen SE, van Mier H, Fiez JA, Raichle ME (1998) The effects of practice on the functional anatomy of task performance. Proc Natl Acad Sci U S A 95:853-860.
Phillips AG, Ahn S, Floresco SB (2004) Magnitude of dopamine release in medial prefrontal cortex predicts accuracy of memory on a delayed response task. J Neurosci 24:547-553.
Picard N, Strick PL (1996) Motor areas of the medial wall: a review of their location and functional activation. Cereb Cortex 6:342-353.
Pizzagalli D, Pascual-Marqui RD, Nitschke JB, Oakes TR, Larson CL, Abercrombie HC, Schaefer SM, Koger JV, Benca RM, Davidson RJ (2001) Anterior cingulate activity as a predictor of degree of treatment response in major depression: evidence from brain electrical tomography analysis. Am J Psychiatry 158:405-415.
Pochon JB, Levy R, Poline JB, Crozier S, Lehericy S, Pillon B, Deweer B, Le Bihan D, Dubois B (2001) The role of dorsolateral prefrontal cortex in the preparation of forthcoming actions: an fMRI study. Cereb Cortex 11:260-266.
Ponce-Alvarez A, Kilavik BE, Riehle A (2010) Comparison of local measures of spike time irregularity and relating variability to firing rate in motor cortical neurons. J Comput Neurosci 29:351-365.

Pratt WE, Mizumori SJ (2001) Neurons in rat medial prefrontal cortex show anticipatory rate changes to predictable differential rewards in a spatial memory task. Behav Brain Res 123:165-183.
Price JL (1999) Prefrontal cortical networks related to visceral function and mood. Ann N Y Acad Sci 877:383-396.
Procyk E, Tanaka YL, Joseph JP (2000) Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. Nat Neurosci 3:502-508.
Qin P, Di H, Liu Y, Yu S, Gong Q, Duncan N, Weng X, Laureys S, Northoff G (2010) Anterior cingulate activity and the self in disorders of consciousness. Hum Brain Mapp 31:1993-2002.
Qin P, Liu Y, Shi J, Wang Y, Duncan N, Gong Q, Weng X, Northoff G (2012) Dissociation between anterior and posterior cortical regions during self-specificity and familiarity: a combined fMRI-meta-analytic study. Hum Brain Mapp 33:154-164.
Qin P, Northoff G (2011) How is our self related to midline regions and the default-mode network? Neuroimage 57:1221-1233.
Quian Quiroga R, Kreiman G (2010) Measuring sparseness in the brain: comment on Bowers (2009). Psychol Rev 117:291-297.
Quilodran R, Rothe M, Procyk E (2008) Behavioral shifts and action valuation in the anterior cingulate cortex. Neuron 57:314-325.
Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. Nature 435:1102-1107.
Raichle ME, MacLeod AM, Snyder AZ, Powers WJ, Gusnard DA, Shulman GL (2001) A default mode of brain function. Proc Natl Acad Sci U S A 98:676-682.
Raos V, Evangeliou MN, Savaki HE (2007) Mental simulation of action in the service of action perception. J Neurosci 27:12675-12683.
Remijnse PL, van den Heuvel OA, Nielen MM, Vriend C, Hendriks GJ, Hoogendijk WJ, Uylings HB, Veltman DJ (2013) Cognitive inflexibility in obsessive-compulsive disorder and major depression is associated with distinct neural correlates. PLoS One 8:e59600.
Restivo L, Vetere G, Bontempi B, Ammassari-Teule M (2009) The formation of recent and remote memory is associated with time-dependent formation of dendritic spines in the hippocampus and anterior cingulate cortex. J Neurosci 29:8206-8214.
Rich EL, Shapiro M (2009) Rat prefrontal cortical neurons selectively code strategy switches. J Neurosci 29:7208-7219.
Rigotti M, Barak O, Warden MR, Wang XJ, Daw ND, Miller EK, Fusi S (2013) The importance of mixed selectivity in complex cognitive tasks. Nature 497:585-590.
Rossetti ZL, Carboni S (2005) Noradrenaline and dopamine elevations in the rat prefrontal cortex in spatial working memory. J Neurosci 25:2322-2329.
Rowe J, Stephan KE, Friston K, Frackowiak R, Lees A, Passingham R (2002) Attention to action in Parkinson's disease: impaired effective connectivity among frontal cortical regions. Brain 125:276-289.
Rudebeck PH, Saunders RC, Prescott AT, Chau LS, Murray EA (2013) Prefrontal mechanisms of behavioral flexibility, emotion regulation and value updating. Nat Neurosci 16:1140-1145.
Rushworth MF, Noonan MP, Boorman ED, Walton ME, Behrens TE (2011) Frontal cortex and reward-guided learning and decision-making. Neuron 70:1054-1069.

Rushworth MF, Walton ME, Kennerley SW, Bannerman DM (2004) Action sets and decisions in the medial frontal cortex. Trends Cogn Sci 8:410-417.
Ryou JW, Wilson FA (2004) Making your next move: dorsolateral prefrontal cortex and planning a sequence of actions in freely moving monkeys. Cogn Affect Behav Neurosci 4:430-443.
Sacher J, Neumann J, Funfstuck T, Soliman A, Villringer A, Schroeter ML (2012) Mapping the depressed brain: a meta-analysis of structural and functional alterations in major depressive disorder. J Affect Disord 140:142-148.
Sakai K (2008) Task set and prefrontal cortex. Annu Rev Neurosci 31:219-245.
Sakai K, Ramnani N, Passingham RE (2002) Learning of sequences of finger movements and timing: frontal lobe and action-oriented representation. J Neurophysiol 88:2035-2046.
Sakai Y (2001) Neuronal integration mechanisms have little effect on spike auto-correlations of cortical neurons. Neural Netw 14:1145-1152.
Scaglione A, Moxon KA, Aguilar J, Foffani G (2011) Trial-to-trial variability in the responses of neurons carries information about stimulus location in the rat whisker thalamus. Proc Natl Acad Sci U S A 108:14956-14961.
Schall JD, Stuphorn V, Brown JW (2002) Monitoring and control of action by the frontal lobes. Neuron 36:309-322.
Schimansky J, David N, Rossler W, Haker H (2010) Sense of agency and mentalizing: dissociation of subdomains of social cognition in patients with schizophrenia. Psychiatry Res 178:39-45.
Schmidt-Hieber C, Hausser M (2013) Cellular mechanisms of spatial navigation in the medial entorhinal cortex. Nat Neurosci 16:325-331.
Schmitzer-Torbert N, Redish AD (2004) Neuronal activity in the rodent dorsal striatum in sequential navigation: separation of spatial and reward responses on the multiple T task. J Neurophysiol 91:2259-2272.
Seamans JK, Durstewitz D, Christie BR, Stevens CF, Sejnowski TJ (2001) Dopamine D1/D5 receptor modulation of excitatory synaptic inputs to layer V prefrontal cortex neurons. Proc Natl Acad Sci U S A 98:301-306.
Seamans JK, Floresco SB, Phillips AG (1995) Functional differences between the prelimbic and anterior cingulate regions of the rat prefrontal cortex. Behav Neurosci 109:1063-1073.
Seamans JK, Lapish CC, Durstewitz D (2008) Comparing the prefrontal cortex of rats and primates: insights from electrophysiology. Neurotox Res 14:249-262.
Seger CA, Miller EK (2010) Category learning in the brain. Annu Rev Neurosci 33:203-219.
Seo M, Lee E, Averbeck BB (2012) Action selection and action value in frontal-striatal circuits. Neuron 74:947-960.
Sesack SR, Deutch AY, Roth RH, Bunney BS (1989) Topographical organization of the efferent projections of the medial prefrontal cortex in the rat: an anterograde tract-tracing study with Phaseolus vulgaris leucoagglutinin. J Comp Neurol 290:213-242.
Shadmehr R, Holcomb HH (1997) Neural correlates of motor memory consolidation. Science 277:821-825.
Shamir M, Sompolinsky H (2004) Nonlinear population codes. Neural Comput 16:1105-1136.
Sharot T, Shiner T, Dolan RJ (2010) Experience and choice shape expected aversive outcomes. J Neurosci 30:9209-9215.

Shenhav A, Botvinick MM, Cohen JD (2013) The expected value of control: an integrative theory of anterior cingulate cortex function. Neuron 79:217-240.
Shepherd AM, Laurens KR, Matheson SL, Carr VJ, Green MJ (2012) Systematic meta-review and quality assessment of the structural brain alterations in schizophrenia. Neurosci Biobehav Rev 36:1342-1356.
Shidara M, Richmond BJ (2002) Anterior cingulate: single neuronal signals related to degree of reward expectancy. Science 296:1709-1711.
Shima K, Aya K, Mushiake H, Inase M, Aizawa H, Tanji J (1991) Two movement-related foci in the primate cingulate cortex observed in signal-triggered and self-paced forelimb movements. J Neurophysiol 65:188-202.
Shima K, Isoda M, Mushiake H, Tanji J (2007) Categorization of behavioural sequences in the prefrontal cortex. Nature 445:315-318.
Shinomoto S, Kim H, Shimokawa T, Matsuno N, Funahashi S, Shima K, Fujita I, Tamura H, Doi T, Kawano K, Inaba N, Fukushima K, Kurkin S, Kurata K, Taira M, Tsutsui K, Komatsu H, Ogawa T, Koida K, Tanji J, Toyama K (2009) Relating neuronal firing patterns to functional differentiation of cerebral cortex. PLoS Comput Biol 5:e1000433.
Silbersweig D, Clarkin JF, Goldstein M, Kernberg OF, Tuescher O, Levy KN, Brendel G, Pan H, Beutel M, Pavony MT, Epstein J, Lenzenweger MF, Thomas KM, Posner MI, Stern E (2007) Failure of frontolimbic inhibitory function in the context of negative emotion in borderline personality disorder. Am J Psychiatry 164:1832-1841.
Sohn MH, Albert MV, Jung K, Carter CS, Anderson JR (2007) Anticipation of conflict monitoring in the anterior cingulate cortex and the prefrontal cortex. Proc Natl Acad Sci U S A 104:10330-10334.
Southwick SM, Vythilingam M, Charney DS (2005) The psychobiology of depression and resilience to stress: implications for prevention and treatment. Annu Rev Clin Psychol 1:255-291.
Stemmer B, Segalowitz SJ, Witzke W, Schonle PW (2004) Error detection in patients with lesions to the medial prefrontal cortex: an ERP study. Neuropsychologia 42:118-130.
Subramaniam K, Luks TL, Fisher M, Simpson GV, Nagarajan S, Vinogradov S (2012) Computerized cognitive training restores neural activity within the reality monitoring network in schizophrenia. Neuron 73:842-853.
Sul JH, Kim H, Huh N, Lee D, Jung MW (2010) Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. Neuron 66:449-460.
Swick D, Turken AU (2002) Dissociation between conflict detection and error monitoring in the human anterior cingulate cortex. Proc Natl Acad Sci U S A 99:16354-16359.
Takehara K, Kawahara S, Kirino Y (2003) Time-dependent reorganization of the brain components underlying memory retention in trace eyeblink conditioning. J Neurosci 23:9897-9905.
Takeuchi H, Taki Y, Hashizume H, Sassa Y, Nagase T, Nouchi R, Kawashima R (2012) The association between resting functional connectivity and creativity. Cereb Cortex 22:2921-2929.
Talairach J, Bancaud J, Geier S, Bordas-Ferrer M, Bonis A, Szikla G, Rusu M (1973) The cingulate gyrus and human behaviour. Electroencephalogr Clin Neurophysiol 34:45-52.
Talmi D, Dayan P, Kiebel SJ, Frith CD, Dolan RJ (2009) How humans integrate the prospects of pain and reward during choice. J Neurosci 29:14617-14626.

Tanaka SC, Samejima K, Okada G, Ueda K, Okamoto Y, Yamawaki S, Doya K (2006) Brain mechanism of reward prediction under predictable and unpredictable environmental dynamics. Neural Netw 19:1233-1241.
Tesink CM, Buitelaar JK, Petersson KM, van der Gaag RJ, Kan CC, Tendolkar I, Hagoort P (2009) Neural correlates of pragmatic language comprehension in autism spectrum disorders. Brain 132:1941-1952.
Thakkar KN, Polli FE, Joseph RM, Tuch DS, Hadjikhani N, Barton JJ, Manoach DS (2008) Response monitoring, repetitive behaviour and anterior cingulate abnormalities in autism spectrum disorders (ASD). Brain 131:2464-2478.
Theunissen F, Miller JP (1995) Temporal encoding in nervous systems: a rigorous definition. J Comput Neurosci 2:149-162.
Toda K, Sugase-Miyamoto Y, Mizuhiki T, Inaba K, Richmond BJ, Shidara M (2012) Differential encoding of factors influencing predicted reward value in monkey rostral anterior cingulate cortex. PLoS One 7:e30190.
Tran-Tu-Yen DA, Marchand AR, Pape JR, Di Scala G, Coutureau E (2009) Transient role of the rat prelimbic cortex in goal-directed behaviour. Eur J Neurosci 30:464-471.
Tu P, Buckner RL, Zollei L, Dyckman KA, Goff DC, Manoach DS (2010) Reduced functional connectivity in a right-hemisphere network for volitional ocular motor control in schizophrenia. Brain 133:625-637.
Uylings HB, Groenewegen HJ, Kolb B (2003) Do rats have a prefrontal cortex? Behav Brain Res 146:3-17.
Uylings HB, van Eden CG (1990) Qualitative and quantitative comparison of the prefrontal cortex in rat and in primates, including humans. Prog Brain Res 85:31-62.
Vallentin D, Bongard S, Nieder A (2012) Numerical rule coding in the prefrontal, premotor, and posterior parietal cortices of macaques. J Neurosci 32:6621-6630.
Van Eden CG, Buijs RM (2000) Functional neuroanatomy of the prefrontal cortex: autonomic interactions. Prog Brain Res 126:49-62.
Van Eden CG, Lamme VA, Uylings HB (1992) Heterotopic Cortical Afferents to the Medial Prefrontal Cortex in the Rat. A Combined Retrograde and Anterograde Tracer Study. Eur J Neurosci 4:77-97.
van Veen V, Carter CS (2002) The anterior cingulate as a conflict monitor: fMRI and ERP studies. Physiol Behav 77:477-482.
van Veen V, Carter CS (2006) Error detection, correction, and prevention in the brain: a brief review of data and theories. Clin EEG Neurosci 37:330-335.
van Veen V, Krug MK, Schooler JW, Carter CS (2009) Neural activity predicts attitude change in cognitive dissonance. Nat Neurosci 12:1469-1474.
Velakoulis D, Wood SJ, Smith DJ, Soulsby B, Brewer W, Leeton L, Desmond P, Suckling J, Bullmore ET, McGuire PK, Pantelis C (2002) Increased duration of illness is associated with reduced volume in right medial temporal/anterior cingulate grey matter in patients with chronic schizophrenia. Schizophr Res 57:43-49.
Vetere G, Restivo L, Cole CJ, Ross PJ, Ammassari-Teule M, Josselyn SA, Frankland PW (2011) Spine growth in the anterior cingulate cortex is necessary for the consolidation of contextual fear memory. Proc Natl Acad Sci U S A 108:8456-8460.
Vilensky JA, van Hoesen GW (1981) Corticopontine projections from the cingulate cortex in the rhesus monkey. Brain Res 205:391-395.

Vincent JL, Kahn I, Snyder AZ, Raichle ME, Buckner RL (2008) Evidence for a frontoparietal control system revealed by intrinsic functional connectivity. J Neurophysiol 100:3328-3342.
Vogels R, Spileers W, Orban GA (1989) The response variability of striate cortical neurons in the behaving monkey. Exp Brain Res 77:432-436.
Vogt BA (2005) Pain and emotion interactions in subregions of the cingulate gyrus. Nat Rev Neurosci 6:533-544.
Vogt BA, Finch DM, Olson CR (1992) Functional heterogeneity in cingulate cortex: the anterior executive and posterior evaluative regions. Cereb Cortex 2:435-443.
Vogt BA, Pandya DN (1987) Cingulate cortex of the rhesus monkey: II. Cortical afferents. J Comp Neurol 262:271-289.
Vogt BA, Pandya DN, Rosene DL (1987) Cingulate cortex of the rhesus monkey: I. Cytoarchitecture and thalamic afferents. J Comp Neurol 262:256-270.
Vogt BA, Paxinos G (2014) Cytoarchitecture of mouse and rat cingulate cortex with human homologies. Brain Struct Funct 219:185-192.
Wallis JD (2007) Orbitofrontal cortex and its contribution to decision-making. Annu Rev Neurosci 30:31-56.
Wallis JD, Anderson KC, Miller EK (2001) Single neurons in prefrontal cortex encode abstract rules. Nature 411:953-956.
Wallis JD, Kennerley SW (2011) Contrasting reward signals in the orbitofrontal cortex and anterior cingulate cortex. Ann N Y Acad Sci 1239:33-42.
Wallis JD, Miller EK (2003) From rule to response: neuronal processes in the premotor and prefrontal cortex. J Neurophysiol 90:1790-1806.
Walton ME, Bannerman DM, Alterescu K, Rushworth MF (2003) Functional specialization within medial frontal cortex of the anterior cingulate for evaluating effort-related decisions. J Neurosci 23:6475-6479.
Walton ME, Croxson PL, Behrens TE, Kennerley SW, Rushworth MF (2007a) Adaptive decision making and value in the anterior cingulate cortex. Neuroimage 36 Suppl 2:T142-154.
Walton ME, Rudebeck PH, Bannerman DM, Rushworth MF (2007b) Calculating the cost of acting in frontal cortex. Ann N Y Acad Sci 1104:340-356.
Weissman DH, Gopalakrishnan A, Hazlett CJ, Woldorff MG (2005) Dorsal anterior cingulate cortex resolves conflict from distracting stimuli by boosting attention toward relevant events. Cereb Cortex 15:229-237.
Werner JD, Trapp K, Wustenberg T, Voss M (2014) Self-attribution bias during continuous action-effect monitoring in patients with schizophrenia. Schizophr Res 152:33-40.
Wheeler AL, Teixeira CM, Wang AH, Xiong X, Kovacevic N, Lerch JP, McIntosh AR, Parkinson J, Frankland PW (2013) Identification of a functional connectome for long-term fear memory in mice. PLoS Comput Biol 9:e1002853.
White T, Hongwanishkul D, Schmidt M (2011) Increased anterior cingulate and temporal lobe activity during visuospatial working memory in children and adolescents with schizophrenia. Schizophr Res 125:118-128.
White TP, Joseph V, Francis ST, Liddle PF (2010) Aberrant salience network (bilateral insula and anterior cingulate cortex) connectivity during information processing in schizophrenia. Schizophr Res 123:105-115.

Whitfield-Gabrieli S, Ford JM (2012) Default mode network activity and connectivity in psychopathology. Annu Rev Clin Psychol 8:49-76.
Williams GV, Goldman-Rakic PS (1995) Modulation of memory fields by dopamine D1 receptors in prefrontal cortex. Nature 376:572-575.
Williams LM, Brammer MJ, Skerrett D, Lagopolous J, Rennie C, Kozek K, Olivieri G, Peduto T, Gordon E (2000) The neural correlates of orienting: an integration of fMRI and skin conductance orienting. Neuroreport 11:3011-3015.
Wilson CJ, Kawaguchi Y (1996) The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons. J Neurosci 16:2397-2410.
Wiltgen BJ, Brown RA, Talton LE, Silva AJ (2004) New circuits for old memories: the role of the neocortex in consolidation. Neuron 44:101-108.
Wohrer A, Humphries MD, Machens CK (2013) Population-wide distributions of neural activity during perceptual decision-making. Prog Neurobiol 103:156-193.
Yeo BT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zollei L, Polimeni JR, Fischl B, Liu H, Buckner RL (2011) The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J Neurophysiol 106:1125-1165.
Yeung N, Botvinick MM, Cohen JD (2004) The neural basis of error detection: conflict monitoring and the error-related negativity. Psychol Rev 111:931-959.
Yeung N, Nieuwenhuis S (2009) Dissociating response conflict and error likelihood in anterior cingulate cortex. J Neurosci 29:14506-14510.
Yin HH (2010) The sensorimotor striatum is necessary for serial order learning. J Neurosci 30:14719-14723.
Yoshimura S, Okamoto Y, Onoda K, Matsunaga M, Okada G, Kunisato Y, Yoshino A, Ueda K, Suzuki SI, Yamawaki S (2013) Cognitive behavioral therapy for depression changes medial prefrontal and ventral anterior cingulate cortex activity associated with self-referential processing. Soc Cogn Affect Neurosci.
Zheng T, Wilson CJ (2002) Corticostriatal combinatorics: the implications of corticostriatal axonal arborizations. J Neurophysiol 87:1007-1017.