Title
Ventral Corticostriatal Mechanisms of Response Suppression: An Incentive Cue Responding Task

Author
Ghazizadeh, Ali

Publication Date
2011

Peer reviewed|Thesis/dissertation

Ventral Corticostriatal Mechanisms of Response Suppression: An Incentive Cue Responding Task

by

Ali Ghazizadeh Ehsaei

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Bioengineering

in the

GRADUATE DIVISION

To mom and dad, my sisters and my wife

Acknowledgements

I am dedicating this dissertation to my family. To my mother for instilling in me an insatiable desire for learning and her unwavering love, to my dad for his kind heart and unconditional support, to my sisters who have always been present in my happiest of their childhood memories and last but not least to my wife, Samin, who is and will remain the best thing that happened to me during my doctoral studies. I'd also like to extend my heartfelt thanks to my parents-in-law for their warmth and understanding.

For the work presented in this thesis, I am indebted first and foremost to my mentor, Professor Howard Fields. I'd like to thank him for investing in my scientific career, for teaching me how to think about and interpret scientific data and for shaping my approach to science. I also owe a great deal of my skills in electrophysiology and my knowledge of the literature to my friend and colleague Dr. Ferederic Ambroggi. His supervision and his day-to-day help with my projects were essential for the completion of this thesis. In addition, I'd like to thank Professor Saleem Nicola. I find myself heavily influenced by his ideas and behavioral models. His comments on my work during our meetings or occasional email exchanges have always been instrumental and inspiring.

I want to thank the members of my thesis committee, Profs. Loren Frank, Greg Hjemlstad, Patricia Janak and Jonathan Wallis for their help and useful comments on this manuscript. I also want to thank my academic advisor, Professor Christoph Schreiner, for his support and advice during my PhD studies in the Joint Bioengi-

neering Program in UCSF and UC Berkeley.

I have been lucky to enjoy the friendship of some special souls during my stay in the bay area. I extend my heartfelt thanks to Arash Ali Amini for being an invaluable friend. I am also grateful to him for helping me texify this manuscript and to proofread the deconvolution model. To him, I have to add Amin Aminzadedeh and Omid Etesami with whom I share memorable times. I cherish the intellectual conversations and the happy times we had together in Berkeley.

Finally, I hope that the result of my work benefits the society at large whose financial support has made it possible for a group of curious scientists to pose daring questions to nature.

Abstract

Ventral Corticostriatal Mechanisms of Response Suppression:

An Incentive Cue Responding Task

by

Ali Ghazizadeh Ehsaei

Adaptation to novel environments involves two classes of motivated behaviors; promoting actions that lead to positive outcomes and suppressing actions that do not lead to positive outcomes. For animals like mammals, it is likely that in any given situation many more behaviors are inhibited than are promoted. Importantly, diminished inhibitory control is a feature of many pathological conditions and maladaptive behaviors, including impulsivity.

Aspects of motivated behaviors, particularly those involving suppression of responses, in this thesis, are characterized in rodents using the DS task. In this task, presentation of a discriminative stimulus (DS) signals availability of sucrose reward upon a lever press. Importantly, presses outside the DS or during the neutral stimulus (NS) is not rewarded and is shown to be suppressed during training. Given the hypothesized role of Nucleus Accumbens (NAc) in guiding motivated behaviors, I first tested its influence on this cued appetitive task. Pharmacological inactivations of NAc differentiated a sub-region specific contribution to the DS task: NAc core promoted responding to rewarded cues while NAc shell suppressed unrewarded

responding, including NS responding. Surprisingly subsequent electrophysiological recordings in NAc, only showed moderate quantitative differences between the two NAc subregions. These results, therefore, indicate that the dissociable roles of NAc subregions are attributable at least partly to their distinct projection targets.

A growing body of evidence implicates the ventromedial prefrontal cortex (vmPFC) in suppression of unnecessary actions. Since vmPFC projects heavily to NAc shell, I hypothesized that this corticostriatal projection mediates the learned inhibitory control in the DS task. Simultaneous vmPFC inactivations and NAc shell recording confirmed this hypothesis. vmPFC inactivation led to a response disinhibition with temporal patterns matching that of NAc shell inactivations. Importantly electrophysiological results indicate that vmPFC mediates its effect by controlling a dual neural mechanism in NAc shell: suppression of NS excitations which can drive behavioral responding to NS and increased firing of gating neurons that are believed to inhibit spontaneous unreinforced responding.

Together the novel results presented, help reveal the ventral corticostriatal mechanism of action suppression with important ramifications for the study of decision making and its aberrations such as impulsivity or addiction.

Finally, two novel methodologies for construction and interpretation of peri-event time histograms (PETHs) are developed and used for analysis of scientific results discussed. For the first method, it is argued that optimal binning of PETHs, is necessary for accurate firing rate estimation and response detection. The search for the optimal binsize, is addressed by adapting the Akaike Information Criterion (AIC)

to binsize selection. The second method puts forth a novel iterative deconvolution technique that can parse out the true event related firings which are often distorted by conventional PETHs due to temporal proximity of experimental events. Both methods, thus, address longstanding issues and are instrumental in interpretation of neuronal firing.

Contents

List of Tables

List of Figures

Overview

This thesis is organized into three parts: part one presents an introduction to decision making and impulsivity which is followed by the main scientific results and includes the first three chapters. Part two, the next two chapters, presents the computational methodology that is developed and used on the data presented in this thesis to improve the analysis and interpretation of peri-event time histograms. The last part is a summary and discussion of the scientific findings presented in this thesis.

In chapter 1, first I review some of the psychological concepts related to decision making by discussing three of the elements that are important for impulsivity and goal directed behaviors namely: drives, attention and actions. Since impulsivity is known to be a multifaceted phenomenon (Evenden 1999b) various forms of impulsivity and various animal models that are developed to test each aspect are also discussed and differentiated. The behavioral review concludes by introducing the rodent behavioral task used in this thesis to study impulsivity and by describing the behavioral measures that are used to address various aspects of impulsive behavior. Chapter 1 ends by reviewing the anatomy of the nucleus accumbens (NAc) and its frontostriatal projections and what is known about the role of this circuitry in action selection and impulsivity.

Chapter 2 reports the results of my experiments to establish the causal role of the NAc in incentive cue responding and to dissociate its sub-region specific contributions

to impulsivity using pharmacological inactivations. This is accompanied by single unit recordings in the NAc showing similarities and differences in the encoding of various components of goal directed behaviors in NAc sub-regions core and shell.

Chapter 3 presents the result of my investigation into the role of ventromedial prefrontal projections to the NAc in suppression of impulsivity. Ventromedial prefrontal cortex (vmPFC) projects to NAc shell and is itself implicated in response suppression. Simultaneous inactivation of vmPFC and recording in NAc shell is used to study how the prefrontal influence in shell neural activity may mediate inhibition of impulsive actions. The changes in the neural encoding of various task elements in the NAc shell following inactivations of vmPFC is reported and shown to be consistent with the role of vmPFC in suppression of unreinforced impulsive actions.

I have developed a few mathematical tools that are important for the neural analysis presented in this dissertation and have a true potential to be useful for in-vivo electrophysiology. The next two chapters describe two of the novel methods developed and constitute the innovative method part of the thesis.

Chapter 4 lays out the theoretical foundations for a deconvolution method that can separate event related neuronal firings (ERFs) from peri-event time histograms (PETHs). This method is used to tease apart the components of neural firing in response to temporally adjacent experimental events. Furthermore, application of this method on simulated and actual data is presented and the improvement over PETH is demonstrated.

Chapter 5 lays out the theoretical foundations for an optimal binning technique

that is used throughout this thesis when constructing PETHs. This method is also tested on simulated and actual data and its advantages over equal binning of the whole neuronal population are discussed.

Chapter 6 presents my research conclusions and incorporates them into our current knowledge about the possible mechanisms of action selection and impulsivity. This chapter also discusses new hypotheses generated by my results and the future directions for impulsivity research.

Part I

First Part

Chapter 1

Introduction

The ability to select appropriate actions based on the association of environmental cues, internal drives and previous action outcomes is key for maximizing survival. At each moment of life, various needs which evoke motivation toward specific goals compete to attain the control of the actions of the organism. The reason for competition is that more often than not the actions required to satisfy these needs can be mutually exclusive. For example, for a rat the decision between scratching an itch or consuming food represents a conflict between two drives one of which tries to minimize a discomfort while the other aims for gaining much needed energy and nutrients[1]. Each need therefore signifies a goal to be satisfied. However, prioritizing the goals is only the first step in determining the future course of behavior. Once a goal has been selected, the animal needs to look for and focus its attention on the stimuli that are relevant for that particular goal. In the case of our example if the animal decides that the itch is the more urgent problem, it needs to determine the location of the itch based on the somatosensory inputs while ignoring the irrelevant sensations like the smell of the food. Afterwards, based on the information provided

[1]Let us assume here for simplicity that scratching behavior and food seeking can only be done in serial and that the rat cannot do both at the same time. This behavioral analogy is used throughout this thesis only as a toy model of decision making conflicts and does not make any claims as to the actual mental processes in a rat's brain.

by the sensory system various actions that are available to achieve a particular goal need to be assessed. The action and its timing should minimize the energy and opportunity costs while at the same time maximizing the degree to which the behavioral goal is achieved. Therefore a potentially beneficial action that is expressed at the wrong time can be worse than not performing an action at all. If our example rat is hungry and knows that food is going to be available only for a short period of time, optimal strategy dictates delaying the scratch action until it has secured the food. It is not hard to generalize based on these examples that all behavior can be viewed as expression of a continuum of choices that are being made simultaneously on at least three levels made up of drives, sensory processing and action expression. It also signifies that aberrations at each of the three levels can reduce the efficacy of the organism to deal with its environment and to ensure its survival.

This thesis concerns itself with one such behavioral aberration known as impulsivity. Impulsivity is generally defined as a tendency to perform 'actions that appear poorly conceived, prematurely expressed, unduly risky or inappropriate to the situation and that often result in undesirable consequences' (Daruna and P.A. Daruna and P.A. p.23). In humans strong impulsive tendencies are reported in a variety of pathological conditions such as obsessive compulsive disorder (OCD), mania and attention deficit hyperactivity disorder (ADHD). Impulsivity is also believed to be a risk factor and the consequence of drug abuse and relapse (Allen et al. 1998; Belin et al. 2008; Dalley et al. 2007; Jentsch and Taylor 1999; Kreek et al. 2005). Given the enormous social and financial costs associated with the impulsivity-related behavioral

mal-functions each year, it is important to understand the neural mechanisms underlying impulsivity. Furthermore since impulsive behavior involves many of the same brain structures that are involved in everyday decision making processes, studying it can be viewed as an important step for understanding decision making processes in the brain and with extension all behavior in general.

Before starting to study impulsivity, it is useful to develop a framework for action selection in goal directed behaviors. Broadly speaking, with the exception of involuntary behaviors such as breathing or reflexes involving the autonomic nervous system, all behaviors include choices happening at, at least, three levels which involve drives, sensory processing and action expression. Before expressing a 'goal-directed behavior', one needs to first prioritize and choose among various goals based on competing drives and motivations. Then, the sensory information that is relevant to the chosen motivated goal needs to be attended to while irrelevant information needs to be filtered out. Finally, given the prior experience of the world and the current sensory data a course of action needs to be selected and computed to achieve the goal (**FIG 1.1**). To understand what selection level entails the influential concepts and theories at each level are discussed below.

1.1 Goal Directed Behaviors: Three Levels of Selection

1.1.1 Drives and Incentives:

Neal Miller a key figure in behavioral neuroscience in 1940s-80s defined drives as the intervening variables that stand between causes and effects in behavioral responses

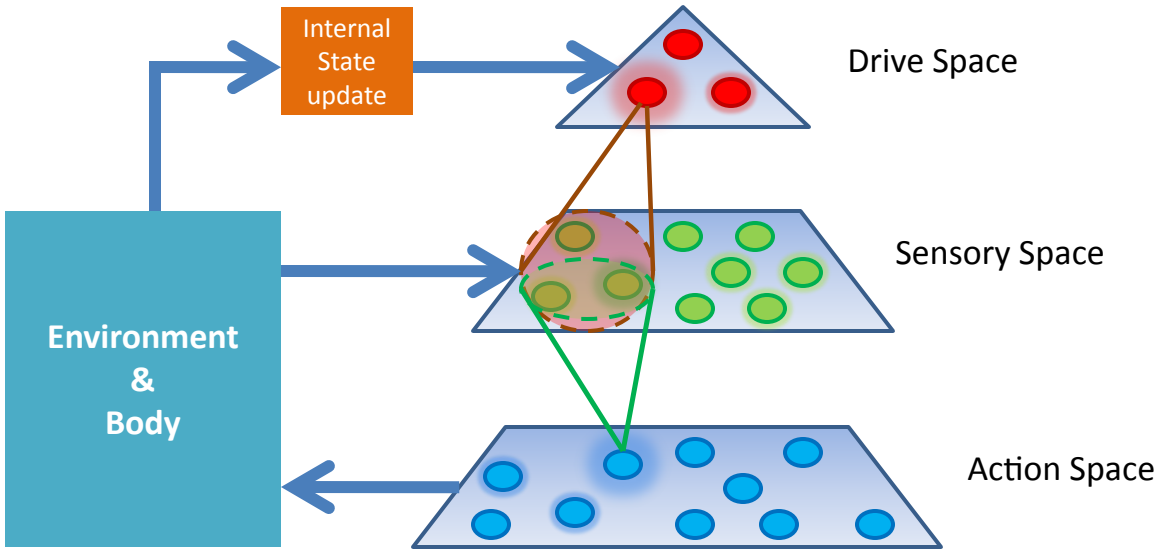

Figure 1.1: **A Diagram for the three levels of decision making and their interactions.** Drives, sensations and actions are conceptualized to exist in three different spaces (layers) that are all engaged in goal directed behaviors. Colored circles depict members of each space (e.g. hunger or reproductive drives in the drive space, smell of food or site of mate in sensory space ,etc). Internal state determines the relative strength of each drive. Each drive defines a region of focus in the sensory space. Sensory inputs arising from environments or the body mediate selection of appropriate behaviors in the action space. Actions eventually modify the current internal state to close the loop.

(Miller 1971) For instance, when deprived of water a rat has a drive to drink. Here the cause of drinking is water deprivation, the response is drinking and the intervening drive that is not observable directly is thirst. However drinking water can have other causes too such as eating salty food which is going raise plasma osmolarity. Osmodetectors in the hypothalamus sense this and cause the animal to increase drinking to maintain homeostasis. One of the properties of true drives according to Miller is that it can generate flexible behavior. That is a thirsty rat can be trained to press a lever or cross a barrier or even endure a costs like a shock if needed to access hypotonic liquid. Therefore thirst as a drive is a conceptual construct that integrates ideas about various causes that increase the need for water and links them to an array of flexi-

ble behaviors. Introduction of drive therefore simplifies thinking about cause/effect relationships in goal directed behaviors and provides a parsimonious explanation for actions.

Drives can be viewed as one aspect of another construct, motivation. This construct is central in defining 'goal-directed behaviors'. Truly goal directed behaviors are often thought to satisfy three criteria (Epstein 1976):(1) flexible goal directedness such as being able to learn new actions to achieve the goal or to overcome obstacles (2) goal expectations which manifest in behaviors such as incentive contrast seen in rats that run faster to the reward spout when the expected reward is greater (Crespi 1942) and (3) affect which can be thought as an emotional feedback after success or failure in achieving the goal. However an animal can have more than a single drive at any moment. In our example rat, the drive to respond to the itch versus the drive to eat might represent a conflict at the level of drives. If the actions needed to address the two drives are mutually exclusive, the animal has to choose which action to perform.

Currently there are two prominent theories regarding motivation: drive-reduction theories and incentive theories of motivation. Drive-reduction theories see goal directed actions as ultimately fulfilling a physiological need in the organism, resulting in a temporary reduction of the drives that promoted those actions in the first place. Drive theorists of motivation have developed concepts such as homeostasis and allostasis to explain the actions that are taken to rebalance the organism (Hull 1943; Mowrer 1980). Concepts such as the hydraulic view of drives (Lorenz and Leyhausen

1973) and opponent process theory (Koob 1996; Solomon and Corbit 1974) are example of drive related theories. More recent developments in motivation research have questioned the validity of simple drive and drive-reduction theories and have proposed instead that the incentive value of sensory pleasures and other conditioned reinforcers can independently promote the goal directed behaviors (Berridge 2001; Bindra 1978; Pfaffmann 1960). One of the favorite counter examples for drive reduction theories offered by incentive theory proponents is the case of a man with a gastrostomy fistula who insisted on chewing his food while having a meal even though he could not swallow it and had to spit it out (Wolf and Wolff 1943). Other experiments with rats and dogs have shown that when given the choice they still prefer to taste their food rather than just being fed intravenously or with an intragastric tube (Miller and Kessen 1952; Turner et al. 1975).

Simple physiological drives are not sufficient to explain the whole range of goal directed behaviors and the explanatory power of incentive/hedonic values of rewards are also needed to understand motivation in most behaviors, Some of the obvious examples of incentives that are not directly related to homeostatic regulation include exploration and novelty seeking, social interactions and in the case of our example rat, grooming. However, a pure version of incentive theories also fails to explain how certain sensory stimuli can become rewarding without impacting homeostatic drives. For example, the reinforcing effect of a sweet taste may be the evolutionary pairing of this taste with nutrients in fruit that satisfy the energy need of the organism. Therefore another possible explanation for the incentive value of sweetness that re-

mains faithful to drive theories can be that it is indirectly linked to behaviors which reduce a drive (Mowrer 1980). Furthermore behavioral effect such as 'alliesthesia' which point to modulation of incentive values by the physiological state such as in 'salt appetite' (Beauchamp et al. 1990; Cabanac 1979) as well as the anticipatory mechanisms that induce preemptive goal directed behaviors toward maintaining a 'homeostatic set-point' (Fitzsimons and Le Magnen 1969) seem to argue in favor of motivation constructs that incorporate physiological drive theories. Therefore motivation theories of future may in fact resemble a synthesis between both the incentive and drive reduction theories.

In any case it is clear that there should be a level of prioritization and selection of 'drives' or 'incentives'[2] before the animal can engage in a goal directed behavior and this represents in my mind a condition required for initiating a goal directed behavior. After choosing the 'drive' or 'incentive', the animal has to move on to planning stage of an appropriate behavior which involves sensorimotor processing.

1.1.2 Sensory Selection and Attention:

In psychology attention is often defined as selective concentration on some aspects of the environment while ignoring other aspects (Anderson 2010). Given the constant bombardment of sensory information a selective focus on the task relevant information is required for a successful goal directed behavior. Attention is known to enhance and accelerate the processing of relevant sensory information and the subsequent

[2]It can be argued that animals do not always choose 'incentives' but may come upon them inadvertently. However it is still plausible that the further processing of the encountered 'incentive' can be permitted or blocked therefore allowing for a degree of selection.

choice of appropriate actions. In fact, an influential theory of attention known as the early-selection model of attention posits that given that attentional resources are limited; not all sensory information can be processed equally even at the early stages of processing. Hence, selective focusing on the relevant information extends from the higher order cognitive processes down to primary sensory areas where task relevant sensory information is prioritized early in processing (Broadbent 1958). The early-selection theory of attention receives support from a large number of studies in the brain which show modulation of information in the primary sensory areas in a top-down modulation for various sensory modalities (Macaluso et al. 2000; Steinmetz et al. 2000; Woldorff et al. 1993). Even though the validity of a complete early filtering of some stimuli has been called into question based on other experimental evidence (Deutsch and Deutsch 1963), there is clearly enhanced processing of attended compared to non-attended sensory information as manifested by faster reaction times to attended stimuli (Posner 1980) and filtering of irrelevant stimuli in inattentional blindness (Rock et al. 1992).

For our example rat, when hungry, sensory selection means ignoring the somatosensory itch information and attending instead to say the auditory cues that can guide its behavior to food delivery. However, sensory selection does not only work between modalities, it can also result in increased focus for particular features of sensory information intra-modally. In fact, success in tasks such as Stroop test or Wisconsin card sorting requires that the participant selectively attend to a particular feature (e.g. meaning, color, number etc) of the visual stimuli. In an auditory ver-

sion of attention subjects can easily learn to 'shadow' the streaming sound to one ear while attending to and repeating the sounds played to the other ear (Broadbent 1958; Woldorff et al. 1993). This predicts that our hungry rat not only can filter out the itch sensation but should also provide preferential processing for the auditory stimuli that signal food availability over other neutral auditory cues.

It should be noted though that goal directed attentional selection for task or goal relevant sensory information can sometimes be interrupted by attention grabbing salient but task irrelevant external stimuli (Corbetta and Shulman 2002). This explains how goal directed behaviors could be aborted before completion. Our rat most probably will not finish its feeding behavior if a predator is detected nearby. This means that changes in the environment can cause the animal to override the chosen incentive or drive, switch its goals and update the rules for sensorimotor relevance. Both the endogenous and exogenous attention mechanisms that are involved in sensory selection therefore are important for promoting actions toward a goal as well as switching between various goals. Disruption of these attentional mechanisms can have adverse consequences for flexible goal directed behaviors.

1.1.3 Action Evaluation and Selection:

Once a goal is selected and its relevant sensory space is identified, various actions that can satisfy the goal need to be evaluated and the one that best achieves the goal with minimal cost needs to be selected. The value of each action in satisfying a goal can be learned through experience. Theories of reinforcement learning such as temporal difference and actor-critic model are widely used to explain both the

mechanism of action evaluation and selection in goal directed behaviors (Peters and Schaal 2008; Schultz et al. 1997). In the actor-critic model, the critic evaluates and updates the value of each action at a given state based on prior reward history and the actor then selects and expresses the best action based on the action values provided by the critic. These models and others in the category of reinforcement learning choose an action policy that maximizes the expected value of future rewards (Sutton and Barto 1998). Furthermore researchers in the field of neuroeconomics have introduced economical concepts such as maximizing expected utility of selected actions and have found correlates of subjective desirability of the outcomes in the brain (Breiter et al. 2001).

Another aspect of action selection relates to the timing and sequencing of action expression. To satisfy a behavioral goal it is not only enough to know what actions should be performed but also when they should be expressed in regards to the acquired behavioral and environmental contingencies. Current models of action selection also provide a way to include those temporal and behavioral contingencies by integrating time and other relevant statistics of environmental and behavioral events into an augmented time dependent state space at which each action can be evaluated.

It is worth pointing out that finding the optimal solution to the action selection problem can often require an exhaustive search in a high dimensional action space which itself can be costly and/or slow (thus violating the optimality) and thus lower the net value of the chosen action in real life situations. Therefore given the fact that most organisms including humans lack the strict criteria for goal maximization

and often have to rely on limited working memory and uncertain situations, they aim instead to find a solution to the action selection problem which is adequate but not necessarily always optimal. The idea of 'satisficing' instead of 'maximizing' put forward by the Herbert Simon, the renowned sociologist, in 1956 attempts to capture this adequacy criteria by proposing action selections that are just good enough (Simon 1956). On the other hand it is argued that if the computational and executional costs of all actions are taken into account then satisficing represents merely a new optimization problem with a modified cost/benefit function making this a more stylistic debate than a substantive one (Byron 2004). Even if representing a difference in taste, the idea of satisficing and the fact that humans as well as other animals do not always make choices based purely on optimizing the external utility of the selected action (Herrnstein 1961; Kahneman and Tversky 1979) should be kept in mind when studying and modeling the action selection behavior and can often have important consequences for interpretation of experimental results.

1.2 Goal Directed Behaviors: Revisited

The view presented in the previous section on goal directed behaviors as consisting of at least three elements namely drives, attention and action selection at the conceptual level favors a modular perspective on the behavior and the neural mechanisms that produce it. A fair point to ask here is 'does it even make sense to compartmentalize various components of goal directed behaviors or talk about selection mechanisms at each stage?' When we get thirsty and want to reach for a glass of water we do not usually feel that we are going through sequential stages of choice in the sensory

data and in the details of movement and grasp of our hand but rather this feels more like a continuous flow of sensations and actions. In fact it can be argued that the nervous system as a prime example of a complex systems shows a good deal of self organization phenomena and many of the observed behaviors can be interpreted as emergent properties of a system described by dynamic flows rather than by a classical mechanical interpretation which favors discrete modules with dissociable roles. This view is gaining ground especially in sensorimotor neuroscience as an alternative to the more classical hierarchical views (Kelso 1995). It can be shown that seemingly hierarchical behaviors can be generated from distributed nonhierarchical structures with attractors that determine the flow and sequence of behaviors (Botvinick 2007).

In our case a clear example of a nonhierarchical system that can produce goal directed behaviors can be found in swarm intelligence. Swarm intelligence describes the collective behavior of decentralized and self organized systems and has abilities that emerge only through the interaction of many simple interacting units (Kennedy et al. 2001). The defense behavior of a school of sardines that are attacked by a shark clearly demonstrate at the macro level an active avoidance, a well known example of a goal directed behavior. Yet this macro level conceptual construct is an emergent property of thousands of individual sardines following simple local rules of movement. The choice of fragmenting a larger school into two to evade the incoming shark and its subsequent rejoining is not made in any centralized decision center or through a deliberation of selection mechanisms but rather follows rules of self organization. The closely related concept of synergetics attempts to explain such self organizing

behaviors by using a few quantities known as order parameters. Order parameters are quantities that allow for a low dimensional description of the dynamical behavior of a complex high dimensional system (Haken 1983).

This new view on the organization of actions provides an exciting alternative to the current thinking on action selection and goal directed behaviors and at least can serve to caution us about the more mechanical and component based interpretations of brain and behavior. However it looks like as far as the brain is concerned there are strong arguments to be made about the modularity and functional specialization as well. Models of swarm intelligence usually assume a symmetric role for the individual agents of the system in generating the overall behavior. However, we know that various structures in the brain are organized rather hierarchically with regard to the nervous system's motor output and their role and impact on the final common pathway is not by any means equal[3]. Furthermore as the size of the brain increases the cost of rewiring and connection of neurons favors clusters of highly interconnected regions that are connected with each other through longer range projection in a small world scheme (Ringo 1991). This clustered structure favors modular organization of functions. Furthermore it can be argued that evolution may favor functional specialization of various clusters as each component can be optimized independent of the others, thus reducing the hurdles of interdependence (Wagner and Altenberg 1996). Finally it can be argued that while we can express behavior that can be explained by self-organization and neuronal swarm intelligence, we find ourselves equally capable

[3]Although it is still plausible that within a brain structure the activity of neurons can operate in a self organizing manner.

of behaving according to the more classical models of action selection. For instance our subjective interpretation of our decision making process in a game like chess conforms better to the modularized concepts of action selection while behaviors which require skillful sensorimotor coordination such as balancing ourselves on a bicycle maybe explained better by self-organization phenomena.

More research is needed to describe the range of behaviors that follow each model and to determine the mechanisms responsible for their generation. For this thesis due to novelty of concepts like complexity and emergent behavior and the proved usefulness of the modularized view of goal directed behaviors the interpretation of results will be made with emphasis mainly on the three previously discussed aspects of the goal directed behaviors. However the possibility of alternative interpretations based on self organization and emergence will also be recognized when appropriate.

1.3 Impulsivity:

Impulsivity is known to encompass a wide range of behaviors that have to do with reduced ability to inhibit proximal inappropriate or suboptimal responding. Impulsivity as a personality trait has been studied extensively in human populations and various personality and temperament questionnaires have been developed to identify the impulsive tendencies while dissociating it from other psychopathologies such as anxiety or depression (Cloninger 1987; Eysenck and Eysenck 1977; Patton et al. 1995). In addition to these tests that are largely based on introspection and self report there is a voluminous literature on impulsivity that has employed behavioral approaches in

humans and in animals.

Despite extensive studies there is still little agreement on how to categorize and group the various dimensions of impulsivity. There is not a unified terminology for aspects that seem similar even when referring to the same behavioral paradigm. One popular classification divides impulsivity into impulsive action (e.g. inability to inhibit premature responding) versus impulsive choice (e.g. preference of small immediate over larger but delayed reward). Similar impulsivity dimensions are referred to by others as motor and cognitive impulsivity. Cognitive impulsivity is also referred to as a failure to delay gratification and is proposed to represent a reduction in self control (Logue 1988). Still another form of cognitive impulsivity known as reflection impulsivity was proposed to describe the inability to collect and evaluate all necessary information before reaching decisions (Kagan 1966). In an important series of reviews on impulsivity and its varieties, John Evenden proposed a three component model for impulsive behavior: preparation to respond, response execution and assessment of outcome (Evenden 1999b; Evenden 1999a). Here the impulsivities related to the first component (i.e. preparation to respond) refer to responses that are made before all the relevant information is obtained (similar to reflection impulsivity). The second component, impulsive execution, was defined as a failure to waiting for one's turn and interrupting others. The third component, preference for impulsive outcome refers to intolerance for delayed rewards. Other groups have proposed similar three component models having to do with (1) input: accumulation of sensory information (2)processing: evaluation and selection of actions and (3) output: execution and

planning of actions (Chambers and Potenza 2003). There are yet other aspects to behavioral inhibition that are addressed in the literature which covers a wider range of executive, motivational and attentional inhibition processes. Some of those processes can be further classified based on whether they are mediated centrally (higher order cognitive processes) or more peripherally (sensorimotor processes) or whether they are effortful suppressions versus more automatic suppressive processes (Nigg 2000). Examples of effortful suppressions include resistance to distracter or proactive interference by simultaneously presented but irrelevant stimuli (Friedman and Miyake 2004). Examples of more or less automatic suppression can be seen, for instance, in an attentional phenomenon known as the inhibition of return where it is shown that the visual stimuli attended within 500-3000ms ago are suppressed (Posner et al. 1985).

Due to the richness of the impulsivity construct and the differences in how it is defined by different investigators, the most successful behavioral paradigms used to study this phenomenon are those that focus on a few aspects of it. Below I will discuss some of the popular tasks used for the study of impulsivity in rats that are most relevant to the kinds of impulsivity that my thesis research was designed to investigate.

1.3.1 Impulsive Choice:

This aspect of impulsivity is often assessed by using a *delay discounting task*. In this task two possible actions are offered to the animal, one of which results in an immediate reward while the other action leads to a larger delayed reward. Importantly,

this task is designed in such a way that the number of choices allowed in a session is the same regardless of action chosen, ruling out the possibility of earning a greater total reward by choosing multiple small rewards in a given session. The studies using this task both in primates and rodents agree that the subjective value of the reward follows a hyperbolic reduction as a function of delay (Ainslie 1975). However the steepness of the discount curve is highly variable in different individuals with high steepness corresponding to impulsive choice (tendency to choose small immediate rewards) (Bezzina et al. 2007; Cardinal et al. 2001). Similar tasks with manipulation of reward probabilities are also used where impulsive choice is defined as tendency to choose a less probable but greater magnitude reinforcement compared to a certain but smaller reward which can result in a lower total reward obtained (Cardinal 2006).

1.3.2 Response Inhibition:

Impulsivity in this category is recognized to contain at least two forms: action cancellation and action restraint (Schachar et al. 2007). Action cancellation involves successful inhibition of an action after it has been initiated while action restraint refers to delayed initiation of a pre-potent response until appropriate conditions for expression is reached.

Impulsivity in action cancellation is often measured using a *stop signal reaction time task* (SSRT). In this task the subject is required on the majority of trials to make short latency responses to presentation of a go cue. In some trials a stop cue is presented after the go cue and the subject has to stop the execution of the response after it has started. As the time interval between the go cue and stop cue increases

it becomes more difficult to abort the response. A theoretical race model is proposed and used for calculating the unobserved stop signal reaction time (Logan and Cowan 1984). This model assumes that there are two independent processes involved in this task namely the go and the stop. Once the go process is launched the stop process can only overtake it if it is itself launched within a short delay. Long stop signal reaction times are considered as an indication of compromised inhibitory control.

Response restraint on the other hand seem to be doubly dissociable from response cancellation (Robinson et al. 2009) and can be measured by tasks other than SSRT. One of the popular tasks in this category is the *go/no-go task*. In this task the subjects are supposed to respond to go cues and withhold responding to no-go cues to get rewarded but unlike SSRT go and no-go cues are presented on separate trials (although the percentage of no-go trials can be much smaller than the go cue, similar to SSRT). The main difference between these two tasks is that in the go/no-go one has to cancel a planned action (restraint)[4] while in SSRT an action that is already initiated is canceled. Another common task used to assess restraint is *differential reinforcement for low rates of responding* (DRL)(Dalley et al. 2008). In this procedure, rats are required to space their responding on a lever by a specific interval of time, pressing the lever within this interval will reset the interval and is taken as a failure of action restraint. A criticism of the DRL paradigm in measuring impulsivity has been that it cannot distinguish well between changes in the general levels of activity and specific premature responding (Evenden 1998). A related procedure, which is proposed to be

[4]The fact that action execution may start before completion of cue processing that reveals its nature (go/no-go) may weaken the distinction between SSRT and Go/No-Go tasks.

less sensitive to general changes of responding, is fixed consecutive number or paced fixed consecutive number schedule (FCN). In this task, rat has to do a predetermined number of consecutive lever presses on one lever (FCN lever) before pressing a second lever (reinforced lever) to obtain food. Premature abortion of FCN lever presses before moving to the reinforced lever is not rewarded and is considered to result from impulsivity. A second rationale for introducing this task is that it allows for studying another facet of impulsivity known as the lack of perseveration, which often presents itself as the inability to remain focused on the current task (Evenden 1998). However, another interpretation may be that premature abortions of FCN presses reflect intolerance to delayed reward. Finally, the *5 choice serial reaction time task* (5CSRT) has been used extensively by Robbins and colleagues to evaluate impulsivity (Belin et al. 2008; Robinson et al. 2009). This task, was originally developed as the rodent equivalent of continuous performance test to study sustained and divided attention and it has parameters that are suited for studying action restraint. In this task the rat has to attend to five holes in the wall waiting for a brief light cue to enter into one of the holes. A successful nose poke in the illuminated hole with a short latency (<5 sec) results in the delivery of a food pellet that can be retrieved from in a food receptacle on the opposite wall. Premature responses to any of the holes during the inter-trial interval (ITI) wait period (which usually lasts about 5 sec) are considered a failure in response restraint.

Another closely related concept to the response disinhibition aspect of impulsivity is the issue of compulsive over responding or perseverative responding. For instance,

in the 5CSRT task, perseverative responding to the recently rewarded nose poke hole occurs and is hypothesized to reflect a failure to disengage from a response driven by a predicted reward. Tasks such as *set shifting* or *reversal learning* have been used extensively to study the response suppression that would prevent perseveration and similar responses and the brain regions involved (Floresco et al. 2006; Ragozzino et al. 2002). In these tasks, a rat is trained to learn a particular set of rules or contingencies for obtaining a reward. Upon successful learning, the contingencies will be changed. The animal has to adapt to the new set of rules. In this context, errors that indicate animal's failure to adapt to the new contingencies are considered compulsive or perseverative. It should be noted, however, that perseverative behavior when framed as a problem of disengagement contrasts with the distractibility facet of impulsivity discussed previously as the inability to remain focused on a task. In fact compulsivity and distractibility can be viewed as two extremes in the engage/disengage framework with the latter implying inappropriate and frequent disengagements from a task while the former indicates remaining engaged in a task despite a reduction in its utility.

1.3.3 Impulsivity: Drives, Attention and Action Selection

Different aspects of impulsivity can be assigned to the three components of goal directed behaviors representing a dysfunction of selection mechanism at each level. For instance aspects of impulsivity such as intolerance to delayed gratification or compulsive behavior can be attributed to incentive sensitization which was originally proposed to explain drug addiction (Robinson and Berridge 1993). Incentive sensitization proposes that a particular drive can hijack the 'wanting' system such that any

cue or action that might satisfy that drive becomes hyper-salient. Incentive sensitization therefore can be conceptualized as an aberration in motivational weight such that one incentive becomes dominant in almost all situations. On the other hand the problem of distractibility, lack of preservation aspect of impulsivity as well as reduced resistance to other types of interference can be viewed as a dysfunction of sensory processing and attention. Here the selection and maintenance of relevant information in the face of interfering bottom-up exogenous influences is compromised such that a transient attentional shift due to distracters leads to a shift of the intentional network or a change of goals. Equivalently based on the context, various forms of impulsivity such as impulsive choice, premature responding and perseverative behaviors can result from problems and inefficiencies in the action selection component and the ability to engage/disengage in various action sequences. Therefore, correct classification of observed impulsive behaviors with respect to drives, attention and action selection requires carefully designed experiments to control for confounding interpretations and might even require studying the behavior across various tasks. On the other hand, the difficulty in disambiguating the confounding interpretations in impulsive behavior maybe taken as a sign that such behaviors are the emergent property of a fluid non-modularized nervous system that is intrinsically resistant to any modular compartmentalization at the level of either concept or brain structure.

1.4 DS task and The Study of Impulsivity:

In this thesis, we focus on concomitant promotion of a specific goal directed behavior and the inhibition of task irrelevant actions. This dual function is required to

improve efficiency on the task; i.e. optimizing reward receipt while minimizing the energy expended. The task used in this thesis to address this issue is known as the 'DS task' (**FIG 1.2**). Except for minor details and modifications between groups, the general description of the task is as following: The task is run in behavioral chambers that contain a food receptacle with two levers on either side, all installed on one of the chamber walls. The levers are assigned randomly as the active and inactive lever for each rat. A speaker installed close to the chamber roof is used to play two distinct auditory cues (a siren and an intermittent tone). The two tones are randomly assigned to each rat as the discriminative stimulus (DS) or the neutral stimulus (NS) both of which are played on a variable interval schedule with an average of 30s. NS is played for 10s and DS is played for up to 10s unless terminated by an active lever press. A sucrose pump containing a 10% sucrose solution is connected to the receptacle and is used to pump 60μl of liquid sucrose into the receptacle as reward. Rats, food restricted to $\sim$90% of their normal weight, are trained to press the active lever only during the DS in order to obtain sucrose reward. Presses on the active lever terminate the DS and result in reward delivery in the food receptacle. Pressing the active outside the DS or the inactive lever does not results in reward delivery but the animals are not punished by a time-out either (Nicola et al. 2004a).

The training of rats in the DS task for this thesis was typically done in the following order: first the rats were habituated to the DS chamber through a magazine training session. In this session reward was delivered after a fixed interval schedule upon the rat's entry to the food receptacle, no levers or cues were presented at this stage. The

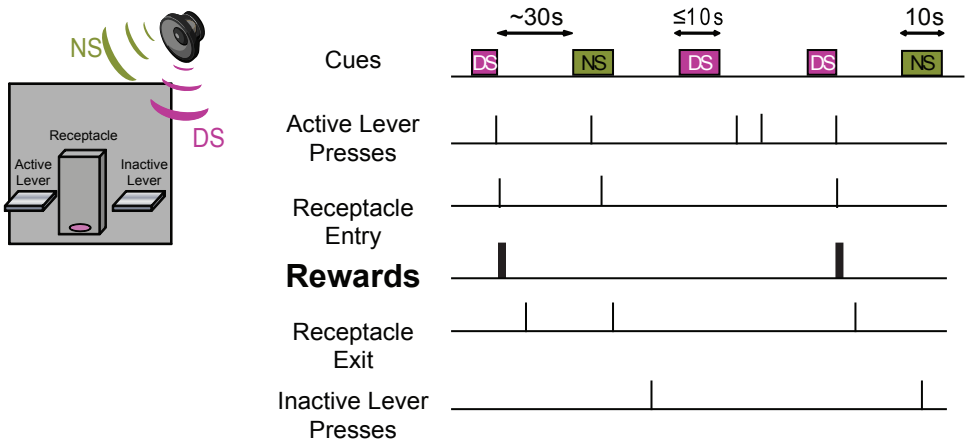

Figure 1.2: **Design of the DS task and relevant behavioral events.** Two cue tones (up to 10 s for the rewarded discriminative stimulus [DS]; 10 s for the unrewarded stimulus [NS]) were randomly presented on a variable-interval schedule with an average interval of 30 s. An active lever-press was required during DS presentation to terminate the DS and cause the delivery of a 10% sucrose reward into an adjacent receptacle. Inactive lever presses or active lever presses outside the DS was never rewarded. The timings of all the events shown were recorded for behavioral analysis.

session continued for one hour or until 99 rewards were obtained. From the second session levers were presented together with the DS. Only presses on the active lever when the DS was played resulted in the sucrose reward delivery. Reward delivery terminated the DS and initiated an inter trial interval (ITI) which was gradually increased over the sessions to 30s. Initially after the ITI expired the DS was played until terminated by an active lever press. This was changed for later groups of rats such that the DS was terminated after a set interval if no active lever press occurred. For these groups the DS duration shrank over sessions until reaching of 10s or less. The reason for this change was the observation that the DS termination after a set interval facilitated the association of DS with reward availability. Similar to the DS task, inactive lever presses did not have any programmed consequences. In addition, the NS was not presented during this stage of training. After the rats reached criterion for responding to the DS (latencies less than 5s and >80% DS followed by a response

if DS had a set interval), they were switched to the full DS task with both DS and NS. Neither the NS nor responses to the NS had any programmed consequences. Behavioral criteria in the DS task for the rats to be included in our studies was >80% DS ratio and <20% NS ratio (DS ratio and NS ratio are the percentage of DSs and NSs that are followed by a an active lever press respectively). This was usually accompanied by DS response latencies of less than 3s.

The use of the DS task to study impulsivity is novel. This task was originally developed to have experimental control over the rat's appetitive responses using stimuli that signaled reward availability[5]. Nevertheless the task schedule has a structure that makes it useful for studying various aspects of impulsivity. Learning the DS task involves two components. The first component, which is usually acquired faster, requires rats to learn that presses on the active lever can result in reward delivery (action-outcome). The second component requires the rats to learn that this action outcome contingency is only valid during the DS presentation. This learning, which usually develops later, involves suppression of lever pressing outside DS as well as its promotion during the DS. In fact, the total number of lever presses as well as the NS ratio is initially high for the rats just after they are switched to the DS task. These measures decrease over multiple sessions without affecting the high DS ratios, reflecting the development of inhibitory control over task irrelevant behavior. Below we discuss how various task events in the DS task can be used to address the different aspects of this inhibitory control.

[5]The design of the task (Cue-Action-Outcome) also helps with the interpretation of electrophysiological recordings to cues compared to a design using CSs. Since the response to the CS can be contaminated by approach and consumption.

1.4.1 General Response Inhibition:

The DS task contains ITIs which average around 60s between presentation of two consecutive DSs (DS and NS are presented with equal probability with variable intervals averaging 30s). Active lever presses that happen during this 60s period when the DS is not presented therefore reflect a failure in response control similar to the DRL schedule with two notable differences. First in the DRL task early responding resets the time-out and therefore that task imposes a greater cost for high levels of responding than the DS task. On the other hand one could argue that it should be easier to induce response disinhibition in the DS task compared to DRL exactly because there is a low cost in ITI responding in the DS task. The second difference is that unlike DRL the end of ITI is explicitly cued in the DS task which should minimize the influence of other confounding effects such as time perception in response initiation. Both of these differences therefore can be advantageous when studying changes in the general response levels in the DS task. The rate of lever pressing during the ITI can reflect the strength of general response inhibition in the task as well as strength of motivation. Food restriction and weight monitoring can be used to minimize the effect of motivational changes across sessions. Finally, increased inactive lever presses can be an indication of failed general response inhibition although significant increases can reflect a problem in action selection as well.

1.4.2 Specific Response Inhibition: Response to Neutral Stimuli

Distractibility is an important aspect of impulsivity. Presentation of a neutral auditory cue (NS) in the DS task provides an opportunity to measure the strength

of inhibitory control to similar but irrelevant sensory stimuli beyond the general response suppression during the ITI. Dysfunctions in the ability to inhibit actions in response to information that is irrelevant to the chosen goal is an important facet of impulsivity and is an example of interference control (Nigg 2000). However increased NS responding can have other equally valid interpretations. It could represent a problem in processing and discrimination of the two cues. If the rat's ability to discriminate between DS and NS is compromised or significantly slowed, an increase in NS responding should be observed. Increased NS responding could also result from a general reduction of inhibitory control on responding. This can be true even if the increase in NS responding is greater than the general increase in lever pressing during the ITI since it is possible that the mapping between motivation to respond and the actual responding follows nonlinear (say a sigmoidal) psychometric mapping. That means that if we assume the initial drive to respond is higher during the NS than during the ITI, it follows that a uniform reduction of drive inhibition would increase NS responding to a disproportionately higher degree than ITI responding. Examining each of these three possibilities would require modifications of the DS task as currently designed.

1.4.3 Specific Response Inhibition: Post Reward Perseveration

Perseverative or compulsive over responding is also an important concept related to impulsivity. Here in the DS task we observed that certain pharmacological manipulations can increase perseverative active lever presses or receptacle checkings following reward consumption. These phenomena, which can be reframed as a disengagement

problem from the goal directed epoch, may represent an enhancement of the motivation for a reward upon its receipt and might have similar mechanisms as the priming effects of reward on response initiation.

1.4.4 Specific Response Inhibition: Impulsive Entry

Another recent observation in the DS task on impulse control is that during the DS presentation rats sometimes enter the receptacle *before* pressing the active lever. In fact analysis of a subset of behavioral data showed that these premature receptacle entries occur in about 30% of rewarded trials. This type of failure of inhibitory control[6] may represent a form of motor impulsivity and taps into mechanisms usually assessed by tasks such as FCN where the rat has to complete a sequence before being rewarded. It is also related to the concept of impulsive choice as the rat has to choose the correct *sequence* of rewarded actions after DS presentation.

Finally sequences of available actions in the task during the cues or ITIs can be used to help with the interpretation of behavioral disinhibitions. For example increases in the active lever presses that are followed by receptacle entry might have a mechanism distinct from that resulting in increases in isolated lever presses.

1.5 Fronto-Striatal Basis of Impulsivity:

Impulsivity is hypothesized to depend upon a number of different brain structures. One of the most relevant circuits in this regard includes the nucleus accumbens in the

[6]It should be noted that many forms of discussed response disinhibition in this thesis may not necessarily reflect a failure in response suppression but an increase in the rate of exploration of alternative action routes which may lead to reward.

ventral striatum and its cortical inputs in the medial prefrontal cortex (Basar et al. 2010; Dalley et al. 2008). Here I will review some of the anatomical and functional data as they relate to the study of impulsivity.

1.5.1 Nucleus Accumbens: Anatomy and Histochemistry

Nucleus accumbens (NAc) is located in the ventral striatum which also includes the olfactory tubercle and is often used synonymously with it. On the rostral and lateral side it borders the external capsule and on the medial side it is bounded by the septum. Caudally it merges with the bed nucleus of the stria terminalis and ventrally with the olfactory tubercle (Zahm and Brog 1992). Even though the dorsal border of NAc and the rest of striatum cannot always be easily defined, a combination of histochemical criteria can be used for satisfactory delineation (Groenewegen et al. 1999). Being a part of ventral striatum, NAc is the input nucleus of the ventral stiatopallidal complex and is comprised largely of medium spiny GABAergic projection neurons (MSNs). In the rat striatum, the proportion of MSNs is estimated to be around 90-97% (Tepper and Bolam 2004). The rest of neurons (about 3-10%) belong to various groups of interneurons namely: aspiny cholinergic neurons, parvalbumin positive aspiny GABAergic neurons and the neurons co-expressing nitric oxide, somatostatin and neuropeptide Y which are also probably GABAergic (Tepper and Bolam 2004). There is some evidence linking these two GABAergic interneuron classes to the two physiological classes of fast spiking and low threshold spiking interneurons respectively (Kawaguchi et al. 1995). The cholinergic interneurons are also believed to correlate with the in-vivo class of tonically active neurons or TANs (Cragg 2006).

The NAc itself is further subdivided into at least two subregions core and shell. Core and shell can be distinguished on the basis of histochemical and connectivity differences. The core is distinguished histochemically from the surrounding shell (Jongenrelo et al. 1994; Zahm 1999). Most notably staining for calretinin and acetylcholinesterase are more strongly expressed in the shell than core while calbindin densely stains core and dorsomedial striatum but not shell (Zahm 1999; Zahm and Brog 1992). In some respects the internal structure of the core resembles the adjacent dorsomedial striatum. The morphology of its MSNs in terms of size and spine density is similar to the rest of the striatum. MSNs in shell on the other hand are smaller in size and less differentiated in terms of dendritic arbor (Meredith et al. 1992). Also similar to the rest of striatum, core projections target other basal ganglia nuclei. On the other hand, shell is unique among the striatal regions in having direct projections to structures outside the basal ganglia such as the hypothalamus ((Zahm and Brog 1992) but also see recent evidence showing weaker projections arising from core as well (Tripathi et al. 2010)).

The general architecture of dorsal basal ganglia involving direct and indirect pathways and their associated dopamine receptors and peptides seems to hold in the NAc with some exceptions. The output of core direct pathway in the rats reaches the entopenduncular nucleus (EP) and the dorsomedial segment of substantia nigra reticulata (SNr) (Deniau et al. 1994). The indirect pathway in core has reciprocal projections to dorsolateral section of ventral pallidum (VPdl) which is the ventral extension of globus pallidus (Heimer et al. 1991) VPdl projects to the subthalamic

nucleus (STN) and to GABAergic neurons in the dorsomedial SNr which is the same region targeted by the core direct pathway. Dorsomedial SNr and EP constitute the ventral striatal output nuclei and project to ventromedial and dorsomedial thalamic nuclei. These thalamic nuclei are in reciprocal connection with medial prefrontal and insular cortex which themselves project back to the medial and lateral core and lateral shell (Voorn et al. 2004) thus completing an open loop in the corticostriatothalamic pathway that involves NAc core (Alexander et al. 1986).

The shell direct pathway projects to EP (Heimer et al. 1991) but not to SNr. Instead, shell targets non-dopaminergic neurons (possibly including GABA neurons) in the ventral tegmental area (VTA) (Hjelmstad et al 2011) which is medial to SNr. The indirect pathway in shell consists of projections to the medial ventral pallidum (VPm) (Groenewegen et al. 1993). However this part of VP appears to not project to subthalamic nucleus or SNr. It instead has direct projections to mediodorsal nucleus in thalamus which raises the possibility that VP may have a combined functionality as both internal and external segments of globus pallidus (Nicola 2007; Zhou et al. 2003). Interestingly the thalamic projections of VPm seem to project to the cortical areas that send their axons to the core thereby suggesting a possible feed-forward mechanism from NAc shell to core (Zahm 2000). Another similar feed-forward mechanism has been proposed to run through the mesencephalic dopamine related regions where projections from shell to medial VTA affect the core projecting neurons in an ascending spiral (Haber et al. 2000). Such feed-forward mechanisms suggest a possible hierarchy in the organization and control of goal directed behaviors with

information flowing from the more ventral limbic regions to the dorsomedial association regions and finally to the dorsolateral sensorimotor regions in the basal ganglia (Although again see Tripathi et al. 2010 for evidence that core neurons collateralize more frequently in shell than vise versa). Finally VPm also has projections to lateral hypothalamus and pedunculopontine nucleus (PPtg) both of which also receive direct projections from shell (Groenewegen et al. 1993).

There is also a body of evidence suggesting that, similar to the dorsal striatum, D1 expressing neurons in the NAc co-express substance P and dynorphin and D2 neurons co-express enkephalin (Lu et al. 1998; Zhou et al. 2003). Similar to the dorsal striatum, about 6% of core neurons co-express both D1 and D2 receptors. However, this percentage is significantly higher in shell ($\sim$17%) (Bertran-Gonzalez et al. 2008). While in the dorsal striatum D1 and D2 expressing neurons constitute the direct and indirect pathways respectively; in the NAc the D1 and D2 expressing MSN populations projecting to the VP are about the same size. This deviation from the direct-indirect architecture is greater in shell than core (Lu et al. 1998;Zhou et al. 2003). However, due to the projections of VP to thalamus it is possible that the D1 dominant projections to the pallidum terminate on thalamic projection neurons which would be consistent with their having a role in the direct pathway. Furthermore, consistent with the direct-indirect architecture, the D2 dominant neurons project selectively to VP and projections to SNr/EP and VTA regions only arise from the D1 dominant population (Lu et al. 1998; Zhou et al. 2003).

Inputs to the NAc arise from a variety of cortical and subcortical structures.

Among the cortical inputs to the NAc in rats are the glutamatergic medial prefrontal cortex projections comprising the prelimbic (PL) and Infralimbic (IL) areas. The more ventral regions in mPFC including the IL have stronger projections to the ventromedial accumbens or shell while the more dorsomedial prefrontal areas project more strongly to the more dorsolateral regions in the core (Uylings et al. 2003; Voorn et al. 2004). The dorsal and ventral insular cortex, thought by some as the equivalent of orbitofrontal cortex in rats, also have projections that are more prominent in core and shell respectively. Among the subcortical projections to NAc are the ones from the hippocampal complex. These afferents contain ventral subiculum and ventral CA1 projections to shell, and medial and lateral entorhinal cortex projections to shell and core respectively (Swanson and Cowan 1977; Totterdell and Meredith 1997). Another notable limbic area projecting to NAc is the basolateral amygdala with the more caudal regions targeting the shell and more rostral regions targeting the core (Wright et al. 1996). Dopaminergic projections to NAc shell arise from medial VTA and to the core from the lateral VTA and medial SNc (Ikemoto 2007). Medial shell is also unique in that it is targeted by noradrenergic input possibly from caudal brain stem; this projection is mostly absent from the rest of the striatum (Berridge et al. 1997). Thalamic projections to the NAc arise mainly from intralaminar and midline divisions of thalamus with intralaminar mostly projecting to the core and the midline in particular the paraventricular nucleus targeting the shell (Berendse and Groenewegen 1990; Groenewegen et al. 1999).

1.5.2 Nucleus Accumbens and Impulsivity:

Nucleus accumbens has been the focus of many studies concerned with goal directed behaviors and in particular impulsivity related phenomena (Basar et al. 2010). NAc influences impulsivity measures across a large range of behaviors implicating it in the selection and expression of the three components of goal directed behaviors discussed before. Here I will review the role of NAc in control of impulsive behaviors and outline our predictions about how this control can be manifested in the DS task.

Response inhibition:

As previously discussed, response inhibition can be further divided into response cancelation, which is normally assayed by experimental designs such as SSRT, or response restraint, which is monitored through premature responses in tasks such as 5CSRT and DRL or error no-go trials in go/no-go tasks. Interestingly, even though lesions of medial dorsal striatum seem to impair inhibitory function in SSRT, lesions of NAc core do not impair performance of this task (Eagle and Robbins 2003a;Eagle and Robbins 2003b). On the other hand, there are many reports of disruption of response restraint following lesions or inactivation the whole NAc or its subregions core and shell. Lesions of the core increase premature responding in the 5CSRT task (Christakou et al. 2004) and increase mean lever press responses per reward in DRL tasks (Pothuizen et al. 2005; Reading and Dunnett 1995). Shell lesions on the other hand don't change premature responding or response rates in either 5CSRT or DRL tasks (Christakou et al. 2004; Pothuizen et al. 2005). Shell inactivations using $GABA_A$ receptor agonists increase Pavlovian conditioned approach (Blaiss and Janak 2009) as

well as lever presses following a food paired cue in reinstatement (Floresco et al. 2008). Both of these studies and others using either lesions or inactivations of the core have found a decrease in incentive cue or second order conditioning responses (Blaiss and Janak 2009; Floresco et al. 2006; Parkinson et al. 2000). In the DS task inactivations of the whole NAc by glutamate receptor antagonists increases active lever presses outside the cue and inactive lever presses (Yun et al. 2004). Even though the DS response ratio seems not to be affected by NAc inactivation, the response latency increases significantly suggesting that there might be two processes that are disturbed by the inactivation: a process that inhibits responding outside the DS and a process that promotes fast responding to incentive cues. Given the contrasting role of core and shell in incentive cue responding suggested by some of the previously mentioned studies, one of the aims of this thesis was to test the hypothesis that core is more involved in the promotion of rapid responses to reward predictive cues while shell primarily has a role in response inhibition (see **Ch. 2**).

Inactivation and cell body lesions of NAc also increases the rate of perseverative responding in some tasks such as 5CSRT (Christakou et al. 2004) or after reversal and extinction of matching to position (Reading and Dunnett 1991) but not always (Floresco et al. 2006). In the DS task, perseverative responses are observed after reward receipt but the role of NAc in this behavior has not been characterized. This behavior was analyzed in the context of NAc core and shell inactivations and will also be reported in **Ch. 2**.

It is worth noting that a possible confound in interpreting response disinhibition

after NAc inactivations or lesions is the spontaneous increase in locomotion that we and others have observed following NAc inactivation (Burns et al. 1994; Eagle and Robbins 2003b; Reading and Dunnett 1995; Yun et al. 2004). In fact it is well known that NAc has an inhibitory effect on locomotion (Pennartz et al. 1994) and a recent study in our lab showed that electrical stimulation of the NAc elicited consistent pauses in animal's ongoing licking and locomotion (Krause et al. 2010). The observed hyperactivity and increase in locomotion following NAc inactivation suggests that elevation of nonspecific responding in the DS task for example, might not arise solely from an increased rate of engagement in sucrose seeking behavior but might be due to an increased rate of switching between the various behaviors that are expressed by the animal in the chamber including the elements of the DS task. Such an interpretation can implicate NAc in maintaining the focus on the chosen goal and suppression of competing drives and will be discussed further in future chapters.

Attentional impulsivity:

Nucleus accumbens has been implicated in resistance to interference by either earlier associations or irrelevant stimuli. Specifically NAc shell is involved in latent inhibition (LI) where repeated non-reinforced presentations of a stimulus impairs an animal's ability to associate that stimulus with an outcome. Inactivation of shell but not core reduces LI in subsequent training (Jongen-Relo et al. 2002; Pothuizen et al. 2005). Also inactivation of shell during learning a response rule facilitates learning of a new rule while inactivation of core after the rule switch reduced its learning (Floresco et al. 2006). Inactivations of both core and shell increases perseverative responding

during extinction of a matching to position task (Reading and Dunnett 1991). In the same task NAc lesions impair the working memory for the matching to positions at long delays as well. In the DS task responding to the unrewarded stimulus, which is possibly under LI increases following inactivation of both core and shell. In the next chapter using selective inactivation of shell and core I will demonstrate that this effect is more prominent following shell inactivation. Also, as discussed earlier, I will attempt to separate the specific increase of responding to the NS from the general disinhibition that arises from inactivation.

Impulsive choice:

Nucleus accumbens is also involved in the selection of actions in goal directed behavior. In choices that involve temporal discounting of rewards, lesions of core in rats increases their preference for immediate small rewards (Bezzina et al. 2007; Cardinal et al. 2001; Pothuizen et al. 2005) and this effect seems to be independent of shell (Pothuizen et al. 2005). Although a recent study showed no effect of core inactivation on the choice of holding a nose poke longer for larger rewards suggesting that at least some forms of impulsive choice might not be dependent on NAc core (Gill et al. 2010). It has also been shown using a quantitative analysis of choices that the deficit following core inactivation is not due to a deficit in perceiving the value of the larger reward but of steepening of the delay discounting hyperbolic function (Bezzina et al. 2007). Interestingly another study has shown that individual rats that tend to engage in impulsive intertemporal choices might overlap with the ones that emit high premature responses during the ITI but not with the ones that are impaired

in SSRT response cancelation (Robinson et al. 2009). This is consistent with the fact that NAc is implicated in impulsive choice and premature response but not in SSRT as discussed previously.

Another aspect of impulsive choice involves effort instead of time. In fact, lesions of NAc core impair effort-based choice and biases the actions toward less effortful options (Hauber and Sommer 2009). This effect has been proposed to be mediated through the dopamine signaling in the NAc as dopamine depletion renders rodents unwilling to engage in effortful choices (Salamone et al. 2007)

The role of core and shell in the impulsivity manifest in the DS task as reward receptacle entry before lever press during DS was not addressed previously and will be assessed in this thesis. However, it should be noted that for the DS task both of the interpretations regarding the temporal versus effort-based choice are equally likely and a clear distinction between the two cannot be made as lever press before entry requires both effort and can be perceived to delay reward collection.

1.5.3 Medial Prefrontal Cortex:

Even though somewhat controversial it is generally accepted that rats do have a prefrontal cortex (PFC) with regions that parallel that of primates (Uylings et al. 2003). Following the Rose and Woolsey definition, PFC is the cortex reciprocally connected with mediodorsal thalamus (Rose and Woolsey 1948). In rats PFC is furthered divided into the medial, lateral and the ventral subdivisions all of which are important in executive functioning and in goal directed behaviors (Cardinal et al. 2002; Dalley et al. 2004; Miller 2000).

In particular, medial prefrontal cortex (mPFC) which has strong projections to the NAc is implicated in impulsivity and its related pathologies in rats and primates (Dalley et al. 2004; Kolb 1984; Winstanley et al. 2006). Medial PFC dorsally includes precentral and anterior cingulate (ACC) cortices and ventrally is divided into prelimbic (PL), infralimbic (IL) and medial orbital (MO) cortices. The dorsal part of mPFC which includes the ACC is involved in inhibition of impulsive choice in tasks requiring differential effort. On the other hand, the impulsivity in intertemporal choice is suppressed by orbitofrontal cortex in ventral PFC regions (Rudebeck et al. 2006). In the context of response inhibition lesions of medial PFC and especially its ventral division including IL increases premature and perseverative errors in the 5CSRT task (Chudasama et al. 2004; Muir et al. 1996) and this effect is mediated through its connections with the NAc (Christakou et al. 2004).

In the DS task inactivations of mPFC also result in behavioral disinhibition in responding to the NS and spontaneous lever pressing with the effect size being larger in the more ventral IL area compared to dorsal PFC areas (Ishikawa et al. 2008a). The temporal specificity of this increased responding in the DS task is not known and will be studied in more detail in later chapters of this thesis. Finally simultaneous inactivations of ventromedial PFC and single unit recordings in the NAc will be used to further investigate the role of ventral corticostriatal projection in controlling various aspects of impulsive behavior (see **Ch. 3**.

Chapter 2

Roles of nucleus accumbens core and shell in incentive cue responding and behavioral inhibition

2.1 Abstract

The nucleus accumbens (NAc) is involved in many reward-related behaviors. The NAc has two major components, the core and the shell. These two areas have different inputs and outputs suggesting that they contribute differentially to goal-directed behaviors. Using a discriminative stimulus (DS) task in rats and inactivating the NAc by blocking excitatory inputs with glutamate antagonists, we dissociated core and shell contributions to task performance. NAc core but not shell inactivation decreased responding to a reward predictive cue. In contrast, inactivation of either subregion induced a general behavioral disinhibition. This reveals that the NAc actively suppresses actions inappropriate to the DS task. Importantly, selective inactivation of the shell but not core significantly increased responding to the non-rewarded cue. To determine if the different contributions of the NAc core and shell depend on the information encoded in their constituent neurons, we performed electrophysiological recording in rats performing the DS task. Although there was no firing pattern unique to either core or shell, the reward-predictive cue elicited more frequent and

larger magnitude responses in the NAc core than in the shell. Conversely, more NAc shell neurons selectively responded to the non-rewarded stimulus. These quantitative differences might account for the different behavioral patterns that require either core or shell, and neurons with similar firing patterns could have different effects on behavior in this task.

2.2 Introduction

Nucleus accumbens (NAc) neurons contribute to behaviors elicited by reward predictive sensory cues (Cardinal et al. 2002; Humphries and Prescott 2010; Nicola 2007). Lesions or pharmacological manipulation of the NAc impair the ability of conditioned stimuli (CS) to promote operant responding for drugs of abuse or natural rewards (Cardinal et al. 2002;DiCiano et al. 2008; Floresco et al. 2008). However, reversible inactivation of the NAc failed to reduce cue responding in a discriminative stimulus (DS) task and instead increased unreinforced lever-pressing (Yun et al. 2004). In contrast to general inactivation of the NAc, blockade of dopamine receptors in the core strongly reduces incentive-cue responding (Ambroggi et al. 2008), whereas dopamine receptor blockade in the shell has minimal effects on similar behaviors (Nicola 2010). This indicates that NAc core and shell have distinct roles in task performance. Indeed, lesions of the core but not shell reduce performance in different incentive-cue responding tasks (Chaudhri et al. 2010; DiCiano et al. 2001; DiCiano and Everitt 2001; DiCiano et al. 2008; Floresco et al. 2008; Ito et al. 2004). We therefore investigated the differential contribution of NAc core and shell to behaviors in the DS task.

A subpopulation of NAc neurons responds to reward predictive cues. The magnitude of these responses correlates with both the predictive value of the cue and the action elicited by the cue (Ito and Doya 2009; Nicola et al. 2004a; Roesch et al. 2009; Taha et al. 2007) . Furthermore, inactivating structures that provide afferent input to the NAc, such as the ventral tegmental area (VTA), the basolateral amygdala (BLA) and the dorsal prefrontal cortex (PFC) disrupts both responses in NAc neurons and behaviors elicited by reward predictive cues (Ambroggi et al. 2008; Ishikawa et al. 2008a; b; Yun et al. 2004).These observations strongly support the hypothesis that that incentive-cue-elicited firing responses in NAc neurons promote reward-seeking behavior. However, it leaves open the question of which neuronal firing patterns normally inhibit reward-seeking actions during intervals when reward is not available.

The differential contribution of the NAc core and the shell to reward-seeking behavior raises the possibility that their constituent neurons encode different information. Consistent with this idea, NAc core and shell receive different afferent projections (Groenewegen et al. 1999). On the other hand, some brain regions project to both the NAc core and shell (Fields et al. 2007; Humphries and Prescott 2010; Voorn et al. 2004). Consequently, NAc core and shell neurons could encode similar information but have different effects on behavior depending upon the downstream regions they innervate. Indeed, NAc core and shell projection to downstream structures can be somewhat distinct and non-overlapping (Groenewegen et al. 1999; Humphries and Prescott 2010; Nicola 2007; Tripathi et al. 2010; Voorn et al. 2004). The difference of outputs raises the possibility that similar information encoded in the NAc core and

the shell could play markedly different roles in behavior. To address this issue, we recorded the activity of NAc core and shell neurons during the DS task.

2.3 Results

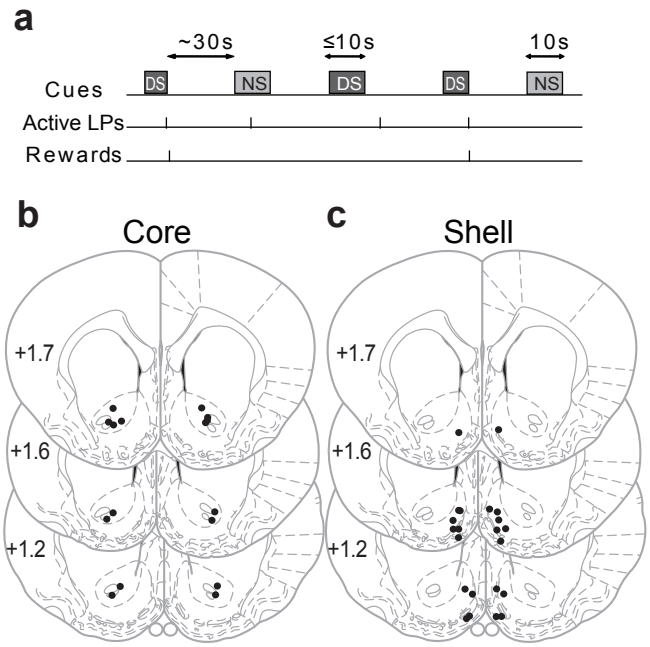

Figure 2.1: **Schematic of the DS task and histology.** (**a**) Two cue tones (up to 10 s for the rewarded discriminative stimulus [DS]; 10 s for the unrewarded stimulus [NS]) were randomly presented on a variable-interval schedule with an average interval of 30 s. A lever-press was required during DS presentation to terminate the DS and cause the delivery of a 10% sucrose reward into an adjacent receptacle. Histological reconstruction of injectors placements in (**b**) the core and (**c**) the shell.

All rats in this study were trained on the DS task (**FIG 2.1**). Two auditory cues, the rewarded discriminative stimulus (DS) and the non-rewarded stimulus (NS), were randomly presented on average every 30 s. During a DS presentation, rats pressed the active lever which turned off the auditory cue and delivered the reward. They then quickly entered the reward receptacle to retrieve the reward. After consumption of the sucrose reward rats would exit the reward receptacle to await the next cue

presentation. During the inter-trial interval and following the NS, rats also intermittently pressed the active and inactive levers or entered the reward receptacle, despite the fact that those actions were never rewarded.

2.3.1 Pharmacological inactivation of the NAc core and shell.

Two groups of animals with guide cannulae in either the NAc core or shell received micro-injections of CSF and a mixture of glutamate antagonists CNQX/AP5 (1 and 2 ug respectively in 0.5 μl) on separate days (**FIG 2.1b,c**).

Incentive cue responding: The DS response ratio (defined as the proportion of DSs the animals responded to), the number of rewarded active lever-presses and the latency to lever-press after DS presentation are related to the strength of the incentive-cue (the DS) to promote behavior to obtain the reward. All measures were significantly affected by the injections (**FIG 2.2a-c**, 2-way ANOVAs, Region $F(1,17)>6.7$, $P<0.02$, Injection $F(1,17)>39$, $P<0.001$, Region x Injection $F(1,17)>22$, $P<0.001$ for all variables). Inactivation of the NAc core but not the shell significantly reduced the DS ratio (**FIG 2.2a**, post-hoc, $P<0.001$) and increased the latency to lever-press after DS (**FIG 2.2b**, post-hoc, $P<0.001$).

Disinhibition of active lever-presses: The above results demonstrate that the NAc core but not shell is required for incentive-cue responding. In contrast, both NAc core and shell inactivation produced significant behavioral disinhibition. The NS ratio was increased following inactivation but did not depend on the region (**FIG 2.2a**, 2-way ANOVA Region $F(1,17)=2.2$, $P=0.15$, Injection $F(1,17)=7.4$, $P=0.01$, Region

x Injection F(1,17)=0.7, P=0.4). Similarly, the number of unrewarded (active) lever-presses (those occurring outside the DS presentations) was also increased after inactivation independent of the region (**FIG 2.2c**, 2-way ANOVA Region F(1,17)=1.7, P=0.2, Injection F(1,17)=15.2, P=0.002, Region x Injection F(1,17)=2.6, p=0.1).

The fact that unrewarded active lever-presses were increased after NAc inactivation raises the question of the temporal specificity of the increase during the NS. Do rats press the active lever in response to the NS or is the NS ratio increase due primarily to a non-specific general increase in lever-pressing? To address this question, we analyzed the frequency of lever-pressing during different time windows: during the inter-trial interval (spontaneous), during the DS, during the NS, and in the 10 sec after the rats consumed the reward and exited the receptacle (post-reward) (**FIG 2.3 a**). The injections affected the frequency of active lever-presses differently (**FIG 2.3 b**, 3-way ANOVA, Region F(1,68)=12.8, P=0.001, Injection F(1,68)=3.7, P=0.06, Window F(3,68)=211, P<0.001, Region x Injection x Window F(3,68)=7.2, P<0.001). NAc core inactivation reduced lever-pressing during the DS but did not significantly increase lever-pressing in the spontaneous or NS windows. Compared to CSF injection, inactivation of the core increased lever-pressing immediately after collection of the reward (post-reward window) to a greater degree than in the spontaneous window. NAc shell inactivation had greater disinhibitory effects: active lever-pressing was increased during the spontaneous, NS and post-reward windows. Furthermore, the increase in active lever-pressing in these two latter windows was stronger than that during the spontaneous window. This shows that the increase in NS ratio caused

by shell inactivation is the combination of a general increase in lever-pressing and the removal of a more specific NS-driven inhibition of responding.

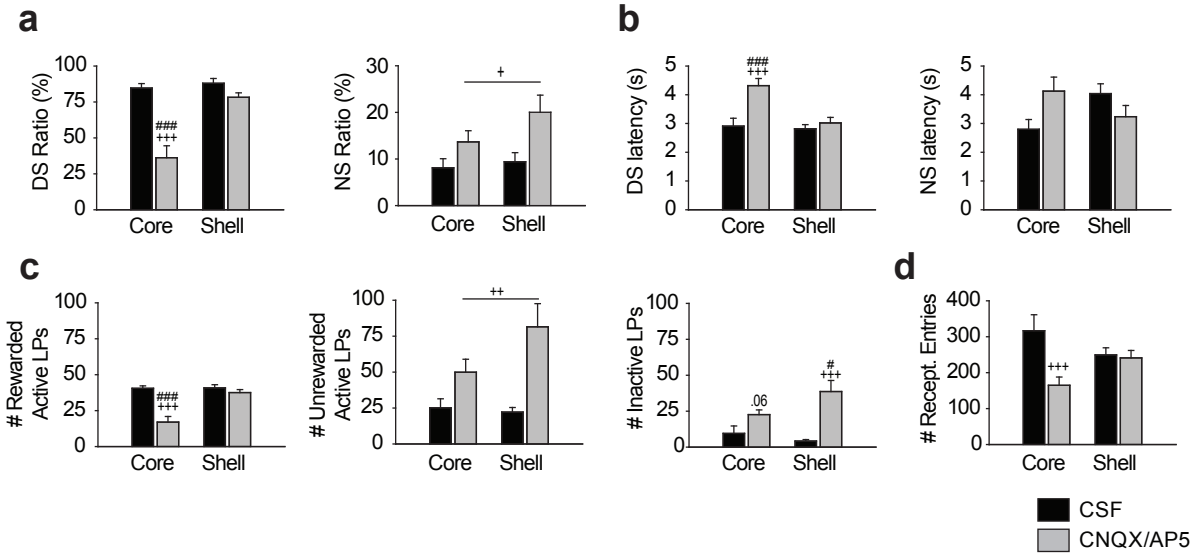

Figure 2.2: **Effect of pharmacological inactivation (CNQX/AP5, 1 and 2 μg/0.5 μl respectively) of the NAc core and shell on the DS task.** (**a**) Average DS (left panel) and NS (right panel) response ratios (proportion of DSs or NSs during which the animal pressed the active lever). (**b**) Average DS (left panel) and NS (right panel) response latencies (time from cue onset to the lever-press). (**c**) Average number of rewarded (i.e. occurring during DS presentations, left panel) and unrewarded (outside DS presentations, middle panel) lever-presses on the active and total number of lever-presses on the inactive lever (right panel). (**d**) Average number of reward receptacle entries. + p<0.05, ++ p<0.01, +++ p<0.001 compared to CSF injections and # p<0.05, ### p<0.001 compared to shell injections.

Because the disinhibitory effect of shell inactivation consisted, in part, of a specific increase in responding during the NS coupled with no effect on DS responding, one possibility is that shell inactivation reduces animals' ability to discriminate between the two cues. However, animals clearly continued to discriminate DS from NS after shell inactivation, as they responded to 80% of DSs and only 20% of NSs (**FIG 2.2a**). More importantly, shell inactivation increases responding in the absence of cues (**FIG 2.2c & 2.3b**). These results are more consistent with a contribution if the NAc shell to response inhibition than to cue discrimination.

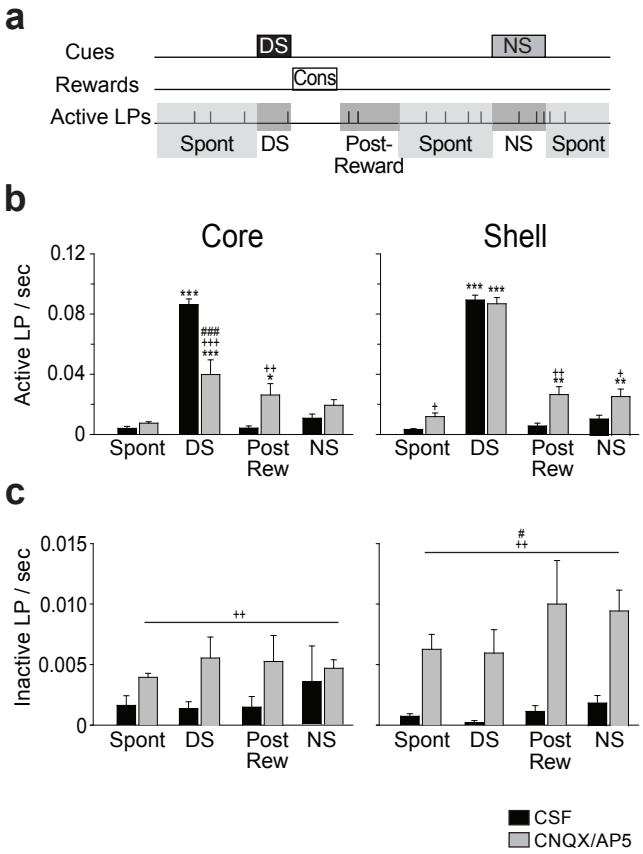

Figure 2.3: **Time window analysis of lever-pressing after pharmacological inactivation of the NAc core and shell** (**a**) Schematic illustrating the time windows in which the frequency of lever-pressing was analyzed. DS and NS windows correspond to the entire duration of the cues. The post-reward window corresponds to the 10 s following the consumption of the reward. The spontaneous window (Spont) corresponds to the remaining time. Frequency of (**b**) active and (**c**) inactive lever-pressing in the different time windows after CSF or CNQX/AP5 injections in the core (left panel) or the shell (right panel). * $p<0.05$, ** $p<0.01$, ***, $p<0.001$ compared to the spontaneous window; + $p<0.05$, ++ $p<0.01$, +++ $p<0.001$ compared to CSF injections and # $p<0.05$, ## $p<0.01$, ### $p<0.001$ compared to shell injections.

Disinhibition of inactive lever-presses: Inactivation of the NAc also increased the number of presses on the inactive lever (**FIG 2.2c**, 2-way ANOVA, Region $F(1,17)=0.8$, $P=0.4$, Injection $F(1,17)=25$, $P<0.001$, Region x Injection $F(1,17)=5.12$, $P=0.03$). NAc shell inactivation significantly increased inactive lever-presses (post-hoc $P<0.001$) whereas NAc core inactivation produced a non-significant trend toward an increase (Post-hoc, $P=0.06$).

When analyzed in different time windows, the number of presses on the inactive lever after NAc inactivation showed an injection effect and a Region x Injection effect (**FIG 2.3c**, 3-way ANOVA, Region $F(1,68)=1.5$, $P=0.2$, Injection $F(1,68)=35$, $P<0.001$, Window $F(3,68)=1.1$, $P=0.4$, Region x Injection x Window $F(3,68)=0.4$, $P=0.7$) indicating that the increase in inactive lever-pressing is stronger for shell than core inactivation (post-hoc on Region x Injection interaction effect, $P=0.01$). However, in contrast to active lever-presses, we found no time window effect. This indicates that the disinhibition of inactive lever-presses is independent of cues or reward; i.e. that there is no temporally specific relation of NAc mediated inhibition of inactive lever-pressing to any specific component of the task.

Reward receptacle entries: The total number of receptacle entries was affected by inactivations (**FIG 2.2d**, 2-way ANOVA, Region $F(1,17)<0.1$, $P=0.9$, Injection $F(1,17)=13$, $P=0.002$, Region x Injection $F(1,17)=10.5$, $p<0.005$) with a reduction after NAc core (post-hoc, $P<0.001$) but not shell inactivation (post-hoc, $P>0.05$). Because the number of receptacle entries is more than five times larger than the number of rewards earned (compare **FIG 2.2c&d**), the prominent reduction of entries after NAc core inactivation cannot be due solely to reduced DS responding, but also reflects lower motivation to check the reward receptacle in the absence of cues.

2.3.2 Electrophysiological activity of NAc core and shell neurons.

We recorded 983 neurons located in the NAc of 12 rats (Table 2.1) over 166 sessions during the DS task. After exclusion of the neurons with a low baseline firing

rate (<0.03 Hz), 456 and 484 neurons located respectively in the NAc core and shell remained and were selected for further analysis (**FIG 2.4a**). A small proportion of the neurons (n=44, dark gray dots in **FIG 2.4a**) were located in an area where the NAc shell and the ventral pallidum were indistinguishable both anatomically and electrophysiologically. Excluding neurons located in this area had no consequence of the effects reported below (data not shown). Therefore, since many of these neurons were probably shell neurons, we performed all our analyses on the entire shell data set.

Behaviorally, rats responded more and faster to the DS than the NS (**FIG 2.4b**). The basal firing rate in the NAc core and shell was similar (3.56 +/- 0.24 and 3.65+/- 0.34 Hz respectively, t(938)=0.21, P=0.83). As previously shown in the DS task, different but overlapping subpopulations of NAc neurons respond to cues, operant actions, reward consumption and exit of the reward receptacle (Nicola et al. 2004a). Here, we examine the distributions of each of these firing responses across NAc core and shell.

Cue-related responses in NAc neurons: We found cue-related responses in both the core and the shell. However, DS excitations were more frequent in the NAc core than in the shell (permutation test, P<0.007, **FIG 2.5 a**). This difference was not seen for NS excitations (permutation test, P>0.3). Neurons with cue-evoked inhibitions were equally common in the NAc core and shell (permutation test, P>0.26). For both DS and NS neuronal responses, onset latencies depended on the type of response (excitation/inhibition) and the structure (**FIG 2.5 b**, 2-way ANOVA, for DS struc-

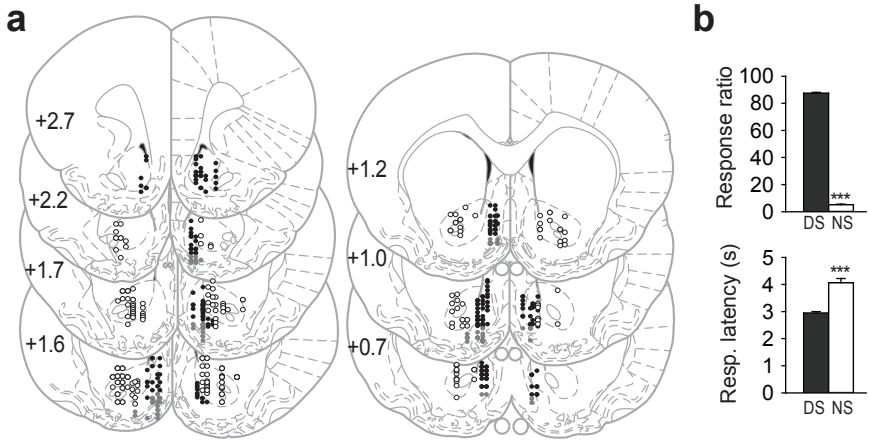

Figure 2.4: **Behavioral performance of the rats during electrophysiological recordings and histology** (**a**) histological reconstruction of electrode placements in the NAc core (open circles) and shell (black and gray dots). Gray dots represent the electrode sites close to the ventral pallidum that were included in the shell data. (**b**) Response ratio (top) and behavioral latency to respond (bottom) during recording sessions. *** P<<0.001 cue effect.

ture F(1,286)=0.35, P=0.55, Response Type F(1,286)=21.7, P<0.001, Region x Response Type F(1,286)=4.65, P=0.03; for NS, Region F(1,56)=4.82, P=0.03, Response Type F(1,56)=16.56, P=0.002, Region x Response Type F(1,56)=5.28, P=0.03). The onset latencies for both DS and NS excitations did not differ between core and shell (post-hoc P>0.27). Cue-evoked inhibitions however, appeared at longer latencies in the shell than in the core (for DS and NS, post-hocs P<0.05). Analysis of the duration of DS responses revealed a main effect of response type (F(1, 286)=14.5, P<0.005 and a Region x Response-type interaction (F(1,286)=9.2, P<0.005). Notably, DS inhibitions in the NAc core were significantly longer than all the other response types (post-hocs P<0.005). Indeed, some NAc core neurons were inhibited for over 10 sec (**FIG 2.5c**) while in the shell, the duration of inhibitions was similar to excitations (post-hoc P=0.58). Importantly, rats pressed the lever an average of 2.9+/-0.05 s af-

ter the DS was presented (**FIG 2.4a**). This implies that many inhibitions in the core encompassed the lever-press and reward consumption. This effect contrasts with the durations of other DS response types that usually terminated before or around the time of lever-press. We then sorted the neurons according to whether they responded to one of the cues or both (**FIG 2.6a,b**). Interestingly, despite similar proportions of NS excited neurons, the overlap between DS and NS excitations was larger in the NAc core than in the shell (permutation test, $P<0.05$).

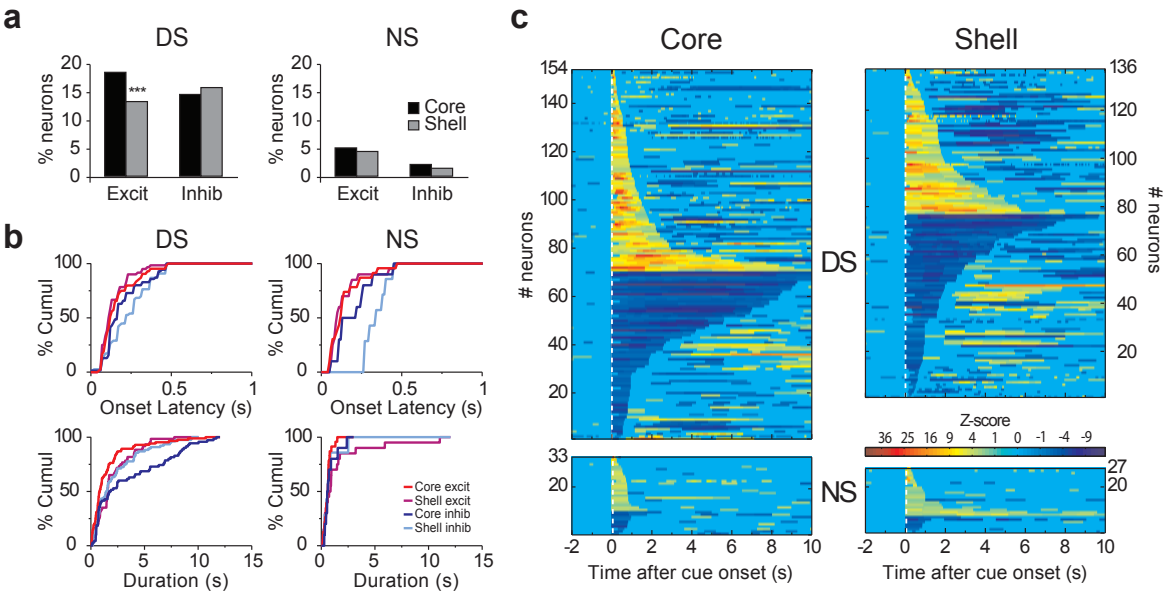

Figure 2.5: **Neuronal responses to the DS and NS in the NAc** (**a**) Percentage of excitation and inhibition in response to the DS (left) and NS (right). (**b**) Cumulative percentage of onset response latency (top) and response duration (bottom) for DS (left) and NS (right). (**c**) Normalized (as Z-score) and color-coded PETH showing all DS- (top) and NS-responsive (bottom) neurons. Each line represents the PETH of a single neuron. Cue-excited neurons are displayed at the top and cue-inhibited neurons at the bottom. Within each category, neurons are sorted by response durations (shorter responses at the top for excitations and at the bottom for inhibitions). Deconvolution of PETHs was not used for the analyses in this figure.

The fact that some neuronal responses span multiple task events is a potential confound for analyzing event-related neural responses. This is particularly true for responses to cues followed at short latency by actions. For example, most DSs were

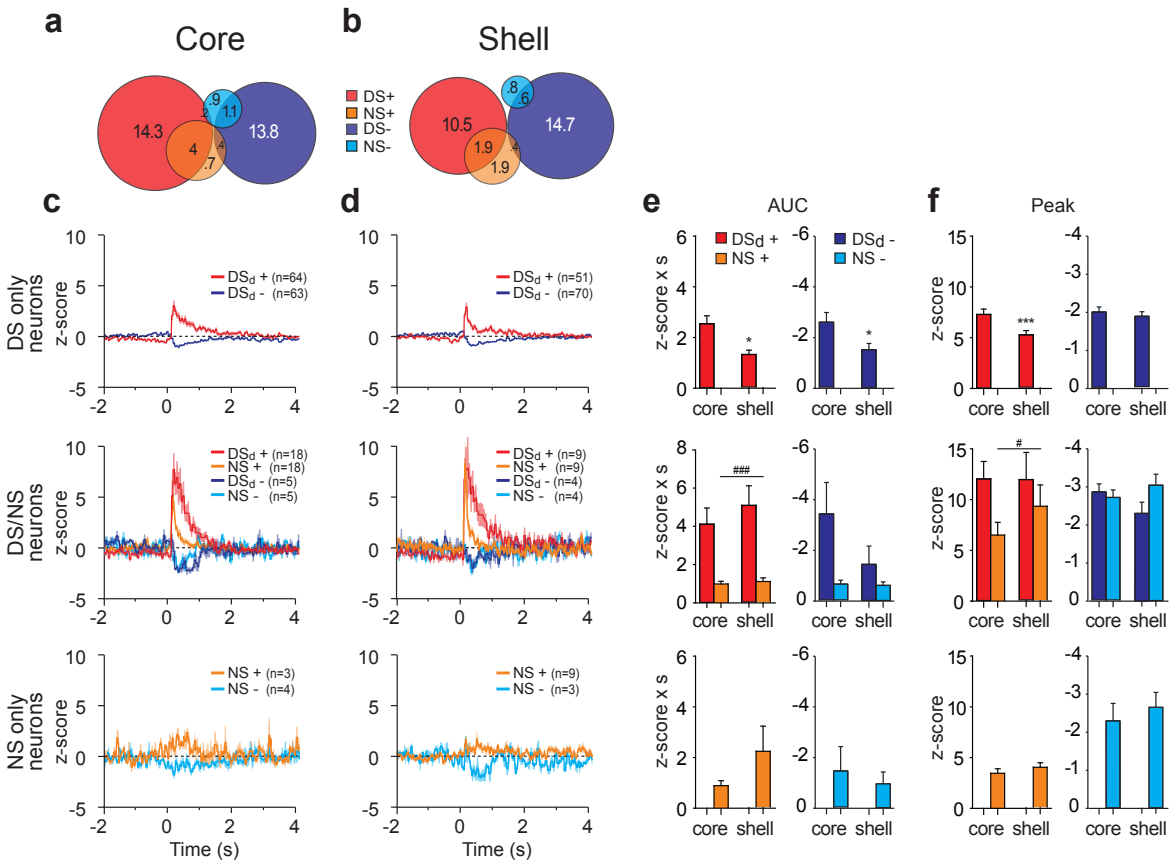

Figure 2.6: **Overlaps between DS and NS responses.** Venn diagram showing the proportion of DS-excited (red), DS-inhibited (dark blue), NS-excited (orange) and NS-inhibited (light blue) responses for (**a**) the core and (**b**) the shell. The overlaps represent the neurons having two types of responses. Average responses to DS and/or NS in (**c**) the core and (**d**) the shell. Top: neurons responding only to DS. Middle: neurons responding to both DS and NS with the same response direction (excitation or inhibition). Bottom: neurons responding only to NS. (**e**) AUC and (**f**) peak of DS and/or NS responses. *P<0.05 structure effect, ### P<0.001 cue effect.

rapidly followed by a lever-press while NSs typically were not. A different firing pattern following either DS or NS could therefore be primarily related to lever-press. We recently developed an analysis method for parsing the firing related to multiple sequential events. The procedure uses the natural variability in time between events to separate the firing that could be related to either or both events and thereby isolate the firing relative to a single event (Ghazizadeh et al. 2010 or see **Ch.4**). We applied this method to 3 events: DS, lever-press and reward receptacle exit. Because reward

receptacle entry occurred at relatively constant latencies after lever-presses (Nicola et al. 2004b; Nicola 2010), we did not include reward receptacle entry in our analysis. The lever-press related responses studied below could therefore be also related to receptacle entry.

Analysis of the neuronal population responsive only to the DS (**FIG 2.6c-f**, top row) revealed that excitations and inhibitions were larger in the NAc core than in the shell (**FIG 2.6e-f**, top row, excitations $t(114)<-2.0$, $P<0.05$ for both AUC and peak; inhibition $t(132)=-2.4$, $P<0.05$ for AUC and $t(132)=-0.7$, $P=0.47$ for peak). For the population of neurons excited by both cues, deconvolved DS (DSd) responses were considerably larger than NS responses in both the core and the shell (**FIG 2.6c-e**, middle row, 2-way ANOVA, cue $F(1,25)>4.2$, $P<0.05$, Region $F(1,25)>0.52$, $P>0.47$, cue x Region interaction $F(1,25)<0.51$, $P>0.48$ for both AUC and peak). We found no significant effect for inhibitions (2-way ANOVA, $F(1,6)<4$, $P>0.09$). The subpopulation of neurons responsive only to NS had small responses similar in core and shell (**FIG 2.6c-f**, bottom row, excitations $t(10)<-0.64$, $P>0.47$ for both AUC and peak; inhibition $t(6)<0.61$, $P>0.56$ for both AUC and peak) with less defined peaks. Excitation peaks appeared at considerably longer latencies than the other cue-related responses (excitations: 550 and 460 ms for NS responses of NAc core and shell respectively compared to 120 ms for all other cue responses in DS-only and DS/NS responsive neurons, $P>0.05$).

Overall, these results show that the DS is more robustly represented in NAc core than shell neurons. Indeed, the proportion of DS excitations and the magnitude of

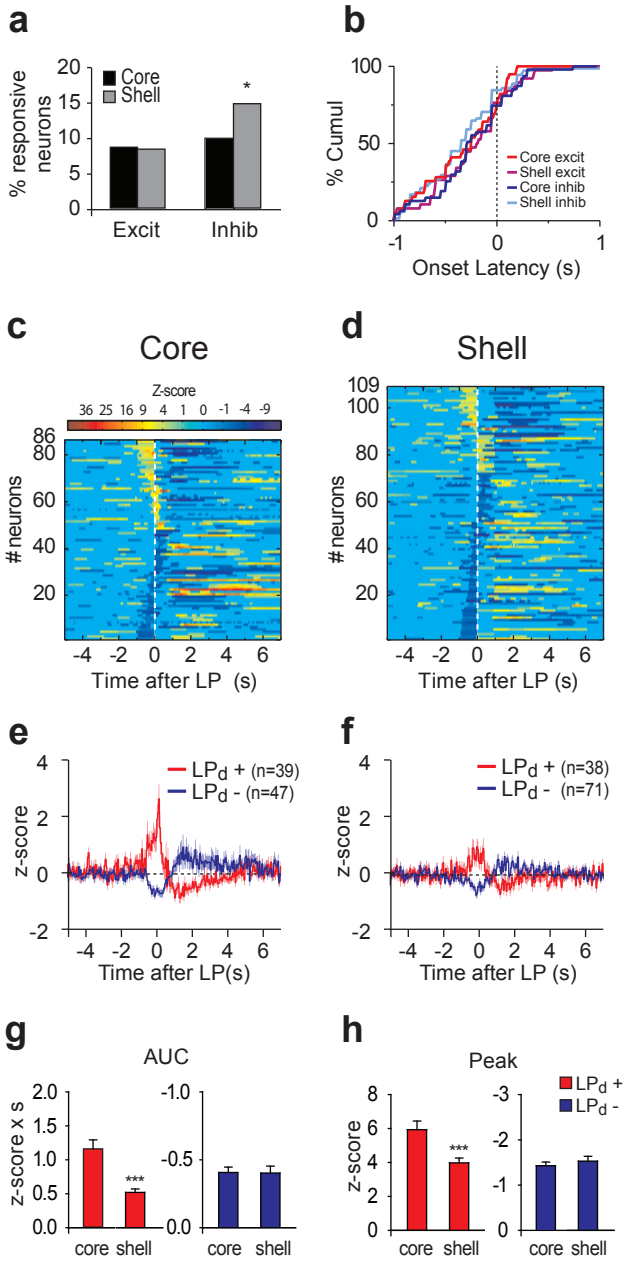

Figure 2.7: **Neuronal responses to rewarded lever-presses** (**a**) Percentage of excitation and inhibition in response to the lever-press in the core and the shell. (**b**) Cumulative percentage of onset latency of lever-press-related responses. Normalized and color coded PETHs aligned to the lever-press for (**c**) core and (**d**) shell. Each line represents the PETH of a single neuron. LP-excited neurons are displayed at the top and LP-inhibited neurons at the bottom. Within each category, neurons are sorted by response durations (shorter responses at the top for excitations and at the bottom for inhibitions). Average lever-press responses for (**e**) core and (**f**) shell neurons. (**g**) Peak and (**h**) AUC of lever-press responses. ***P<0.01 compared to core.

both DS-excitations and inhibitions are larger in the NAc core for neurons responding only to the DS. Notably, this is not the case for neurons co-excited by DS and NS, but this represents a considerably smaller population than those responsive to DS only. Although the proportion of NS responding neurons is similar in core and shell, more shell neurons respond selectively to the NS.

NAc neuronal correlates of rewarded lever-press responses: We then analyzed phasic responses occurring around rewarded active lever-presses (**FIG 2.7**, see methods). Neurons inhibited during the lever-press were more frequently found in the NAc shell than the core (**FIG 2.7a**, permutation test, P=0.027). Onset latencies of excitations and inhibitions were similar across structures (**FIG 2.7b**, 2-way ANOVA Region $F(1,188)<0.1$, P=0.8, Response type $F(1,188)=0.1$, P=0.7, Region x Response Type $F(1,188)<0.1$, P=0.8). Excitations were stronger in the NAc core than in the shell (**Figure 2.7c-h**, AUC, $t(74)=4.47$, $P<0.001$ and peak $t(74)=3.31$, $P<0.005$) but inhibitions were similar (AUC, $t(114)=0.69$, P=0.49 and peak $t(114)=0.94$, P=0.94). In conclusion, there is a reciprocal pattern of dominant neuronal representation of rewarded lever-presses in the NAc core and shell: stronger magnitude of excitations in the core but more frequent inhibitions in the shell.

NAc responses to reward consumption. Previous studies show that many NAc neurons respond during reward consumption (Krause et al. 2010; Nicola et al. 2004b; Roitman et al. 2005; Taha and Fields 2005) although whether their firing changes encode reward valence or promote responding for reward has been debated. Both

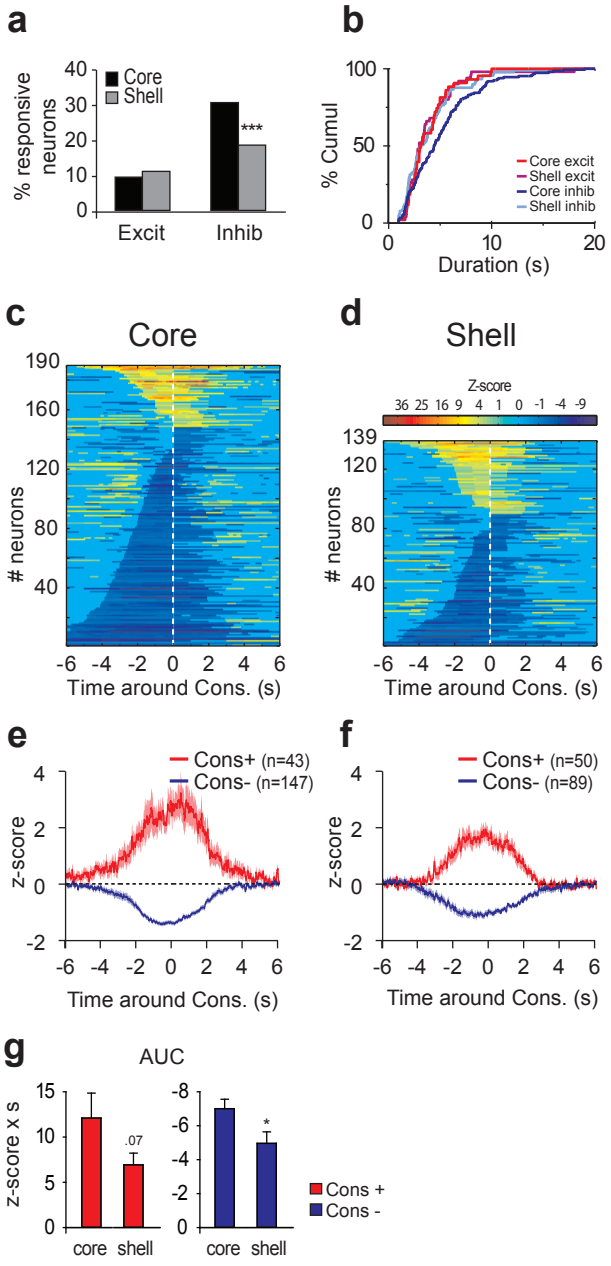

Figure 2.8: **Neuronal responses to reward consumption** (**a**) Percentage of excitation and inhibition in response to the reward consumption in the core and the shell. (**b**) Cumulative percentage of response duration of consumption-related responses. Normalized and color coded PETHs aligned to the middle of the reward consumption for (**c**) core and (**d**) shell. Consumption-excited neurons are displayed at the top and consumption-inhibited neurons at the bottom. Within each category, neurons are sorted by response durations (shorter responses at the top for excitations and at the bottom for inhibitions). Average response for consumption-responsive neurons in (**e**) the core and (**f**) shell. *P<0.05 compared to core.

NAc core and shell neurons were more often inhibited than excited during reward consumption (permutation test, P<0.01 for both core and shell, **FIG 2.8a**). Inhibitions in the NAc core during reward consumption were far more frequent than in the shell (permutation test, P<0.001, **FIG 2.8a,c,d**). Although inhibitions in the core did not last significantly longer than in the shell (**FIG 2.8b-d**, 2-way ANOVA Region F(1,325)=2.2, P=0.14, Response Type F(1,325)=3.2, P=0.07, Region x Response Type F(1,325)=2.4, P=0.12), they had a larger AUC (t(234)=2.3, P=0.02, **FIG 2.8e-g**). These strong inhibitions in the core are reminiscent of the long lasting inhibitions found after DS presentation (**FIG 2.5b,c and FIG 2.6**). Excitations in the NAc core displayed a trend toward a significantly greater firing rate increase relative to the shell (t(91)=1.8, P=0.07).

NAc responses to reward receptacle exit. Consistent with previous reports (Nicola et al. 2004b), NAc neurons respond to the exit of the reward receptacle. NAc core neurons were more often excited by receptacle exit (permutation test, P<0.05, **FIG 2.9a,c,d**). Response onset latency depended on the structure and the response direction (**FIG 2.9b-d**, 2-way ANOVA Region F(1,194)=2.9, P=0.09, Response Type F(1,194)=13.8, P<0.001, Region x Response Type F(1,194)=8, P=0.005). Inhibition onset in the NAc core occurred earlier than onset of all other response types (**FIG 2.9b**, post-hoc P<0.05). We found no difference between NAc core and shell neurons in term of response magnitude (AUC excitation, t(102)=0.58, P=0.56; peak excitation, t(102)=1.29 P=0.2; AUC inhibition, t(92)=-0.9, P=0.36; peak inhibition, t(92)=1.23, P=0.22).

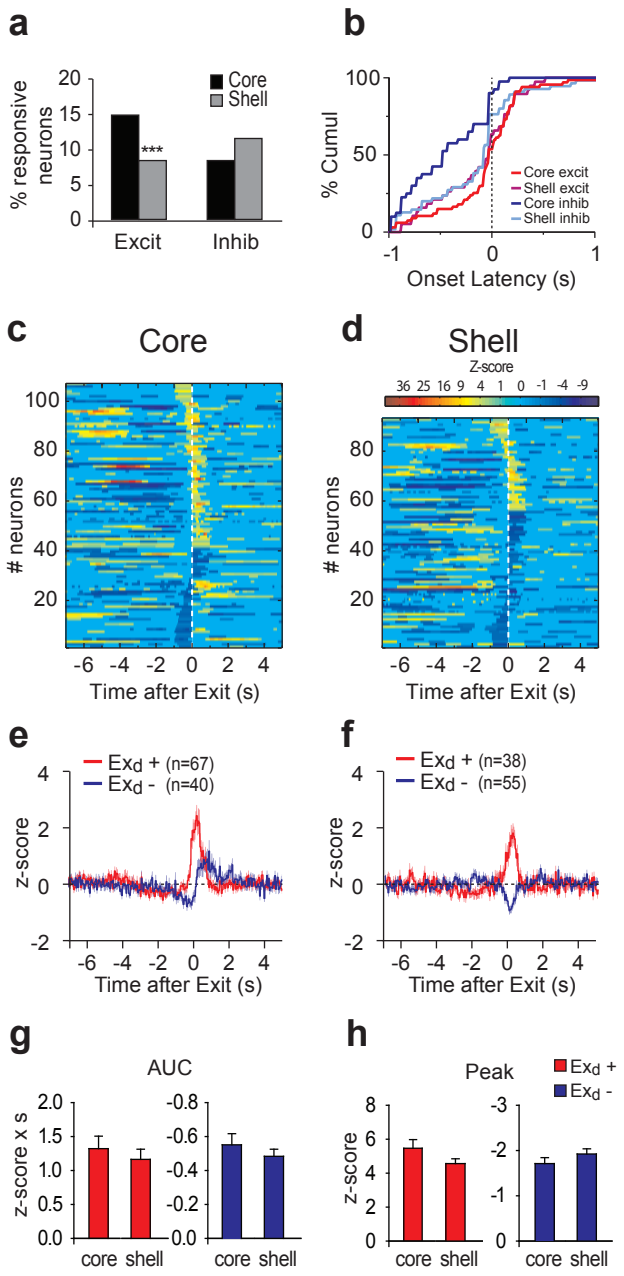

Figure 2.9: **Neuronal responses to reward receptacle exit** (**a**) Percentage of excitation and inhibition in response to the reward receptacle exit in the core and the shell. (**b**) Cumulative percentage of onset latency of reward receptacle exit-related responses. Normalized and color coded PETHs aligned to the middle of the reward receptacle exit for (**c**) core and (**d**) shell. Reward receptacle exit-excited neurons are displayed at the top and reward receptacle exit–inhibited neurons at the bottom. Within each category, neurons are sorted by response durations (shorter responses at the top for excitations and at the bottom for inhibitions). Average receptacle-exit responses for (**e**) core and (**f**) shell neurons. (**g**) Peak and (**h**) AUC of receptacle-exit responses. ***P<0.01 compared to core.

2.4 Discussion

Here we dissociate the effect of NAc core and shell inactivation in a discriminative stimulus task. Blockade of glutamate receptors in the core but not the shell impairs responding to a reward predictive cue. On the other hand, shell but not core inactivation produces a temporally specific increase in responding to the unrewarded cue. Although we did not find neuronal firing patterns specific to core or shell, we did find stronger neuronal representation of the DS in the core, and more selective representation of the NS in the shell. This is consistent with the hypothesis that DS responses in core promote behavioral responding to reward-predictive cues, whereas NS responses in the shell suppress responding to less relevant cues.

Role of the NAc core and shell in incentive-cue responding. Inactivation or lesion of the NAc core impaired performance in tasks where discrete sensory cues controlled operant behaviors (DiCiano et al. 2008; Floresco et al. 2008; Fuchs et al. 2004; Hall et al. 2001; Ito et al. 2004). However, inactivation of the entire NAc failed to decrease responding in the DS task and instead increased nonspecific responding (Yun et al. 2004). The present study resolves this issue by showing that selectively inactivating the core reduces responding to the DS. Shell inactivation, on the other hand, strongly increased unrewarded responding. The lack of reduction of DS-cued responding by general NAc inactivation (Yun et al. 2004) might therefore be due to opposing inhibitory and disinhibitory effects of core and shell inactivation, respectively, on cue responding.

The finding that core inactivation markedly reduces responding to the DS suggests that neuronal activity occurring between the DS and the lever-presses promotes responding. In line with this, NAc neuronal responses to the DS were greater when the rat made an operant response to the cue than if it did not (Nicola et al. 2004a). Overall, we found that both DS and lever-press neuronal responses were larger in the core than in the shell. This provides a potential explanation for the differential involvement of the core and shell in behavioral responding to the DS. On the other hand, many neurons in the shell did respond to the DS and lever-press with response characteristics similar to those of core neurons. Despite this fact, inactivating the shell had no effect on DS responding. Therefore, the differential contribution of core and shell to DS responding is likely to be due to differences in the projection targets of neurons with similar firing patterns. For instance the NAc core projects to the lateral ventral pallidum, the subthalamic nucleus and the medial substantia nigra pars reticulata, while the shell projects more frequently to the medial ventral pallidum, the VTA and the hypothalamus (Groenewegen et al. 1999; Heimer et al. 1991; Tripathi et al. 2010).

The similarity of DS representation in core and shell raises the question of the origin of the afferent input that drives the DS in these subregions. We previously showed that the VTA is necessary for DS-evoked responses in NAc neurons (Yun et al. 2004). Furthermore, cue-evoked dopamine release is similar in the core and shell (Wanat et al. 2010), consistent with the idea that dopamine promotes DS-evoked firing of both core and shell neurons. The BLA is also a critical source of input for

DS firing. BLA inactivation blocks DS-evoked firing in the core (Ambroggi et al. 2008; Jones et al. 2010). Although lesions of BLA reportedly failed to block shell DS responses (Jones et al. 2010), the rostral BLA, known to project preferentially to the core (Shinonaga et al. 1994) was targeted by Jones et al. (2010) . The caudal BLA, which projects to the shell, could therefore contribute to cue-evoked responses in the shell. On the other hand, it is also possible that distinct inputs contribute differentially to DS-driven responses in the core and shell. Indeed, the dorsal PFC, which projects preferentially to the core, promotes both behavioral and neuronal responses to incentive-cues (Ishikawa et al. 2008b). Whether shell responses are specific to other PFC inputs requires further investigation.

Role of the NAc core and shell in suppressing task-irrelevant behaviors. We found that NAc inactivation induced two different forms of behavioral disinhibition. One, produced by both core and shell inactivation, is a general increase in responding on both levers. Previous studies found that such effects were specific to shell inactivation (Floresco et al. 2008; Peters et al. 2008). The larger volume injected in our study and possible diffusion from one site to the other could explain this discrepancy. However, we also found a shell-specific form of behavioral inhibition: a temporally specific increase in active lever-pressing during the NS.

Our neuronal recordings cannot reveal potential physiological correlates of the non-specific increase in lever-pressing caused by NAc inactivation because it is impossible to determine when successful inhibition of a behavior occurred. However, the fact that more neurons are inhibited by lever-presses in the shell compared to the core

is consistent with fact that behavioral disinhibition is stronger after shell inactivation. This suggests that this pattern of activity allows for the occurrence of rewarded lever-pressing. On the other hand, the increased responding during the NS window following shell inactivation indicates that NAc neurons can generate temporally specific response suppression during the NS. In fact, we found a greater population of neurons excited by the NS (but not by the DS) in the shell than in the core. These NS-specific neuronal responses in the shell may inhibit operant responding promoted by the NS.

The population of NS responsive neurons in shell was relatively small, as was the increase in NS-induced lever-pressing following shell inactivation. What could explain this relatively weak effect? Rats initially lever-pressed to the NS when it was first introduced in training (see **FIG 3.1**) and had to learn to actively suppress such responding. Across several months of training, the degree of suppression of prepotent responding to the NS probably decreased over time. If the shell participates in suppressing responding, therefore, its involvement might be stronger at the beginning of extinction, when behavioral suppression is strongest. This hypothesis is supported by data obtained by Ghitza and collaborators (Ghitza et al. 2003): NAc neurons were recorded on the first day of extinction of a similar discriminative stimulus task where the cue was presented but the reward (cocaine) was not. They report that neurons in the shell responded more strongly to the cue than those in the core. We suggest that this was due to the fact that a population of shell neurons was recruited by extinction and that these neurons suppressed responding to a cue that was no longer associated

with the reward. Interestingly, the ventral PFC, which preferentially targets the shell (Voorn et al. 2004), has been implicated in extinction (Peters et al. 2008) and its inactivation induces behavioral disinhibition in the DS task (Ishikawa et al. 2008a). Therefore, activity related to action suppression in shell neurons may depend on ventral PFC inputs. Further experiments are needed to test this hypothesis.

Role of excitations versus inhibitions in promoting and stopping behavior

A critical unanswered question is whether excitation or inhibition of NAc neurons promotes or suppresses behavior. Because initial studies in primates did not report inhibitions in the striatum (Bowman et al. 1996; Cromwell and Schultz 2003; Hassani et al. 2001), most of the theories of the function of the ventral and dorsal striatum are related to excitations (Hikosaka 2007; Humphries and Prescott 2010; Nicola 2007). However, studies in rodents have shown that NAc neuronal inhibitions are prevalent in several appetitive behaviors (Janak et al. 2004; Krause et al. 2010; Taha and Fields 2006). The presence of long-lasting inhibition, such as those reported here in the core, have been proposed to act as a gate for extended goal-directed behaviors. This hypothesis states that a subpopulation of NAc neurons tonically suppresses behavior and that their inhibition is necessary for a behavior to occur. The fact that these inhibitions are more widespread and longer lasting in the core might appear to contradict to the stronger behavioral suppression exerted by the shell in this and other paradigms (Kelley and Swanson 1997; Pecina and Berridge 2000; Reynolds and Berridge 2002; Stratford and Kelley 1997). However, electrical stimulation of both NAc core and shell produces powerful inhibition of sucrose licking (Krause et al. 2010).

Furthermore, inhibitions (or excitations) in the core and shell could have opposing effects on behavior. Selective manipulation of excited and inhibited neurons will be required to determine their specific behavioral roles.

2.4.1 Conclusion

The NAc core and shell make distinct contributions to the DS task. Our recordings revealed that no event-related firing was specific to the core or the shell, suggesting that functionally similar populations of neurons are distributed across both regions. The fact that these sub-territories project to somewhat non-overlapping brain regions might therefore be more relevant to understand their opposing effects on behavior.

2.5 Materials and Methods

2.5.1 Animals

The subjects were male Long–Evans rats (Harlan Sprague Dawley, Indianapolis, IN) weighing ∼350 g on arrival and individually housed on a 12 hr light/dark cycle. Experiments were conducted during the dark phase. After receipt, rats were allowed at least 1 week of *ad libitum* food and water, followed by 1 week of restricted food before training. Throughout all experiments, food restriction was adjusted daily at the end of experimental manipulations to maintain the rats at ∼90% of their initial body weight. Animal handling and experiments conformed to National Institutes of Health and Ernest Gallo Clinic and Research Center animal care and use policies.

2.5.2 DS task

The behavioral pharmacology study was conducted in standard operant chambers (23.5 x 30.5 cm) and the electrophysiological study was conducted in larger cages (40.6 x 40.6 cm). All cages contained two retractable levers located on one wall of the chamber (one on each side of a reward receptacle), two house lights, a white noise speaker and a tone speaker (Med Associates, St. Albans, VT). Liquid sucrose reward was delivered into a well in the reward receptacle by a syringe pump. The DS task structure was similar to our previous experiments (Nicola et al. 2004a; Yun et al. 2004; Yun et al. 2004). Rats were run daily on the DS task for one or two hours (for behavioral and electrophysiological studies, respectively). Two tones, the discriminative stimulus (DS) and the non-rewarded stimulus (NS), were presented on a variable interval schedule with an average interval of 30 sec. Pressing one of the two levers (designated the active lever, randomly chosen across rats) during DS presentation resulted in the delivery of 50 μl of 10% sucrose into the reward receptacle and termination of the DS tone. Each DS lasted for up to 10 sec, and each NS lasted 10 sec. Responding on either lever during the NS or in the absence of the DS was never rewarded. The cues were either an intermittent 4 kHz tone (40 ms on and 50 ms off), or a siren tone (ramped from 4 to 8 kHz with a 400 ms period). Tones were randomly assigned to be the DS or the NS across rats.

Over the course of training, animals' responding to the DS increased while NS responding and uncued lever-pressing decreased. Surgeries were performed when the rats reached criterion performance of >80% DS response ratio and <20% NS response

ratio (defined as the proportion of all DSs or NSs in the session to which the animal responded).

2.5.3 Surgeries

For the behavioral study, rats were bilaterally implanted with microinjection guide cannulae (27 gauge, Plastics One, Roanoke, VA, USA) in the NAc Shell (AP +1.6, ML ±0.8, DV -6.0 mm relative to bregma) and NAc core (AP +1.6, ML ±2 and DV -6.0 mm).

For electrophysiological recordings, 2 arrays of 8 electrodes (NB Labs,Denison,TX, 50 μm stainless steel wires arranged in 2 rows of 4) were attached to a microdrive device that allowed the entire arrays to be lowered by 80 μm increments. Target coordinates of the medioposterior electrode of each array were as follows: core, AP +1.2, ML±2.0 and DV -6 to -8.5 mm and shell, AP +1.2, ML ±0.8, DV -6 to -8.5 mm).

Animals were anesthetized with isoflurane (5%) and placed in a stereotaxic apparatus. Anesthesia was maintained with isoflurane (0.5-2.0%) during surgery. Microdrives were secured to the skull with bone screws and dental acrylic, and wire obturators were inserted into the guide cannulae; the ends of the obturators were flush with the ends of the guide cannulae. Rats were given at least 7 days of recovery before being retrained on the DS task and habituated to the handling procedures.

2.5.4 Microinjections

To block excitatory transmission in the NAc, we used a cocktail of the α-amino-3-hydroxy-5-methylisoxazole-4-proprionic acid (AMPA) and N-methyl-d-aspartate (NMDA) antagonists 6-cyano-7-nitroquinoxaline-2,3-dione (CNQX) and 2-amino-5-phosphonopentanoic acid (AP5) which inhibit NAc medium spiny neuron activity (Horne et al., 1990; Pennartz & Kitai, 1991; Pennartz et al., 1991; Nicola et al., 1996; Hoffman & Lupica, 2001) because these neurons require glutamatergic input to fire (Wilson, 1998; Nicola et al., 2000). CNQX (4mg/ml) and AP5 (8mg/ml) were dissolved separately in 1M NaOH. After adjusting the pH to 7.0 with 1M HCl, the solutions were combined.

To inject animals, the obturators were removed and 30 gauge injector cannulae were inserted into the guides. Injectors extended 1.5 mm below the tip of the cannula. A volume of 0.5 μl of CSF or the cocktail CNQX/AP5 (1 and 2 μg respectively) was injected over 2 min. After a 1 min post-injection period, the injectors were gently removed, the obturators were replaced, the animal was immediately placed into the behavioral chamber, and the behavioral session began. CSF or the cocktail CNQX/AP5 were injected on different days (with at least one intervening day with no injection), in random order, in the same animals.

2.5.5 Electrophysiology

Electrophysiological recording was conducted as described previously (Ambroggi et al. 2008; Nicola et al. 2004a). Animals were connected to the recording apparatus

(Plexon, Inc.), which consisted of a headstage with operational amplifiers, cable and a commutator to allow the animal free movement within the chamber. Rats were run for 2 hr daily sessions of the DS task. The microdrive carrying the electrode arrays was lowered by 80 or 160 μm at the end of each session to get a new set of neurons every day.

Spike sorting. Isolation of individual units was performed off-line with Offline Sorter (Plexon, Inc., Dallas, TX) using principal component analysis. Only units with well defined waveforms were included in this study. Interspike interval distribution, crosscorrelograms and autocorrelograms were used to insure that single units were isolated. Additionally, only units in which waveform characteristics were constant over the entire recording session were included in this study.

Determination of the optimal bin size. The size of the bins greatly influences the information contained in peri-event time histograms (PETH). We have developed a principled way of finding optimal binsizes for each PETH using the Akaike Information Criterion (AIC) (**Ch.5**). Because the optimal bin size is usually rather large for most neurons ($\sim$0.1-1sec), and because the AIC commonly shows a fast reduction over small bin sizes followed by slow changes around the optimal bin size, we also made PETHs using the smallest possible bin size that showed less than a 10% change from the optimal AIC value. This bin size is referred to as the deflection point bin size and was on average 60 msec across our neuronal population.

Response detection. PETHs constructed around the behavioral events, with the optimal bin size, were used to detect excitations and inhibitions and the time at which they occurred. The 10 sec period before cue onset was used as a baseline period. Excitation and inhibition to each event was determined by the presence of at least one bin above (for excitations) or below (for inhibitions) the 99% confidence interval of the baseline during the analysis window for each event. Each unit was assigned, based on its firing pattern, to at least one of the subsets of neurons exhibiting the response types. Onset detection was performed using a telescopic approach. We selected the first significant bin using the optimal size and searched, within this time window for the first bin using the deflection point bin size that was beyond 99% confidence interval of the baseline made with this new bin size. If found, we then repeated the search within this deflection point bin, for the first 10ms bin that was beyond 99% confidence of the baseline made with 10ms bin size. With this method, the onset is defined by the highest valid resolution for each PETH. Response duration was computed by finding the first bin after the detected onset that fell into the 99% confidence interval of the baseline based on the optimal bin size only. The choice of different strategies used for onset and duration detection was based on our observation that the telescopic approach and the optimal bin size tend to be more accurate for onset and duration estimation respectively. For some neurons the response types, onsets and durations were also visually inspected afterwards to ensure accuracy.

Deconvolution. We have shown previously that PETHs made for temporally close events do not allow one to accurately account for the neuronal responses to each in-

dividual event. We have developed and validated an iterative deconvolution methodology (see **Ch.4**) that can parse out the neuronal responses of each individual event from the PETHs (Ghazizadeh et al. 2010). During DS trials, four events usually happen within a short period of time, namely: DS, lever-press, receptacle entry and receptacle exit. PETHs made for each of the events can potentially be distorted due the presence of the other events that are correlated with it in time. This can specifically pose a problem for comparing the DS to the NS firing which is not normally followed by other events. For each neuron, we deconvolved single event responses by using the maximum number of iterations that had a cross validation error lower than the PETHs (average iteration of 4.95 for the entire data set) (Ghazizadeh et al. 2010). Therefore instead of the PETHs, we used, when appropriate, deconvolved responses to report the average activity. The area under the curve (AUC, integrated from response onset to offset) and the peak were obtained from the deconvolved responses and used for statistical comparisons.

Data transformation and plotting. Color-coded maps were constructed with optimal bins and average PETHs across neurons were constructed with the deflection point bins. Prior to averaging, the firing rate of each neuron during each bin was transformed to a Z-score: $(F_i - F_{mean}) / F_{sd}$ where F_i is the firing rate of the i^{th} bin of the PETH, and F_{mean} and F_{sd} are, respectively, the mean and the SD of the firing rate during the 10 sec preceding cue onset.

2.5.6 Histology

Animals were deeply anesthetized with pentobarbital and perfused intracardially with saline and 4% formalin (plus 3% ferrocyanide for rats with electrode arrays). Brains were removed, sectioned (40μm), and stained for Nissl substance in order to locate injection or recording sites (labeled by passing a DC current through each electrode before perfusion).

2.5.7 Statistical analysis

T-tests or ANOVAs followed by Newman-Keuls post-hoc tests were used when appropriate. A permutation test was used to compare proportions. All results were considered significant at $P<0.05$.

2.5.8 Table

Rat		**Target**		**Core**		**Shell**
1		core and shell		57		39
2		core and shell		24		34
3		core and shell		144		34
4		core and shell		58		10
5		core and shell		14		15
6		core and shell		28		63
7		core		73		0
8		core		8		0
9		core		72		0
10		shell		0		198
11		shell		0		59
12		shell		0		53
total:				**478**		**505**
total>0.03Hz:				**456**		**484**

Table 2.1: **Distribution of neurons recorded per rat in the NAc core and shell**

Chapter 3

Ventromedial Prefrontal Cortex Suppresses Unreinforced Actions and Their Representations in the Nucleus Accumbens Shell: Corticostriatal Implications for Impulsivity

3.1 Abstract

Suppression of unnecessary or premature actions is typically essential for optimal performance of goal directed behaviors. Furthermore, diminished inhibitory control over unnecessary and/or competing actions is a feature of many pathological conditions and maladaptive behaviors, including impulsivity. The ventromedial prefrontal cortex (vmPFC) is involved in suppression of actions with extinguished outcome contingencies during both aversive and appetitive conditioning. Here our pharmacological inactivation data indicates that vmPFC activity is required for inhibition of unnecessary actions in a learned appetitive task. Furthermore simultaneous electrophysiological recordings in NAc shell, a major vmPFC projection target indicates that vmPFC exerts this behavioral inhibition by suppressing neurons that preferentially encode responses to unrewarded actions. We also present evidence for differential modulation by vmPFC of presumed action promoting and action suppressing shell

neurons. These results reveal a robust mechanism for suppression of unreinforced actions in the context of a learned operant task and identify vmPFC and NAc shell as critical elements of the circuit relevant for this suppression.

3.2 Introduction

In animals with flexible behavioral repertoires optimal action selection is critical for inclusive fitness. In adaptation to novel environments two major classes of behavioral change are required; increasing the probability of actions that promote a positive outcome and reducing the probability of actions that do not promote a positive outcome. For animals like mammals with relatively complex repertoires of possible actions, in any given situation it is likely that as they learn how to perform an operant task many more behaviors are inhibited than are promoted. How are unnecessary actions inhibited when animals perform a learned appetitive task?

There is a growing body of evidence that specifically implicates the prefrontal cortex in suppression of unnecessary actions. Over the course of mammalian evolution, the PFC has undergone a marked increase in size and functionality in species with increasingly complex behavioral repertoires. Dysfunctions or damage to the PFC is associated with loss of inhibitory control (Bechara and Van der Linden 2005; Dalley et al. 2004; Miller 2000). Previous studies in rodent found that the ventromedial part of the PFC (vmPFC, mostly the infralimbic area) is involved in suppression of behaviors such as conditioned fear or cocaine self-administration following extinction (Milad and Quirk 2002; Peters et al. 2008). Based on these observations, we hy-

pothesized that a major function of the PFC is to promote behavioral optimization by limiting the execution of unnecessary actions in tasks where reward availability is restricted to discrete well-signaled or predictable times. To test this possibility, we first studied the development of behavioral suppression during training as rats' performance improved in a learned appetitive behavior. We then inactivated the vmPFC after the rats had learned to suppress unnecessary actions.

How the vmPFC ultimately affects the motor system is unclear. The vmPFC is one of the major inputs to the NAc shell (NAcS) and the circuitry that includes these regions has been implicated in cocaine-seeking reinstatement (Peters et al. 2008). Consistent with the idea that vmPFC inhibits unnecessary or unreinforced actions through connections to the NAcS, we recently showed that electrical stimulation of the NAcS interrupts licking behavior at short latency (Krause et al. 2010) and that its inactivation induces behavioral disinhibition ((Ambroggi et al 2011) see also (Basar et al. 2010)). Therefore combining NAcS electrophysiology with inactivation of the vmPFC, we investigated the contribution of vmPFC inputs in shaping the firing of NAcS neurons that could mediate learned inhibitory control over unreinforced behavior.

3.3 Results

All rats in this study were trained in the DS task. **Figure 3.1a** shows a diagram of the chamber wall with levers and receptacle and the temporal sequence of events in the DS task. Briefly, two auditory cues, the rewarded discriminative stimulus (DS)

and the non-rewarded neutral stimulus (NS), were randomly presented on average every 30s. During the DS presentation, rats pressed the active lever which turned off the auditory cue and delivered sucrose into the reward receptacle.

3.3.1 DS Training

Learning of DS response promotion and unrewarded response suppression: Optimal performance in this task requires the acquisition of two conceptually separate components. The first component represents the instrumental contingency between the active lever press and reward delivery while the second restricts the validity of this instrumental contingency to the time that DS is presented. Therefore, the first component entails promotion of appetitive responding on the active lever while the second component leads to suppression of identical but unreinforced actions such as lever pressing or receptacle entries outside the DS. We analyzed several components of rat behavior throughout the course of training for the DS task (consisting of two phases, see Methods).

During the early phase of training in the cued FR1 task, the rates of both active lever presses and reward receptacle entries increased significantly as rats learned the action outcome contingency between active lever presses and reward (**FIG 3.1b** first column) (1-Way ANOVA Session $F(6,44)=3.35$ $P=0.008$ for lever press, Session $F(6,44)=3.03$ $P=0.014$ for entry). After learning this rats were switched to the full DS task and the overall rates of active lever pressing and receptacle entry decreased (**FIG 3.1b** second column) (ANOVA Session $F(6,47)=7.82$ $P<0.001$ for lever press, Session $F(6,47)=3.32$ $P<0.01$ for entry). This change occurred despite the fact that

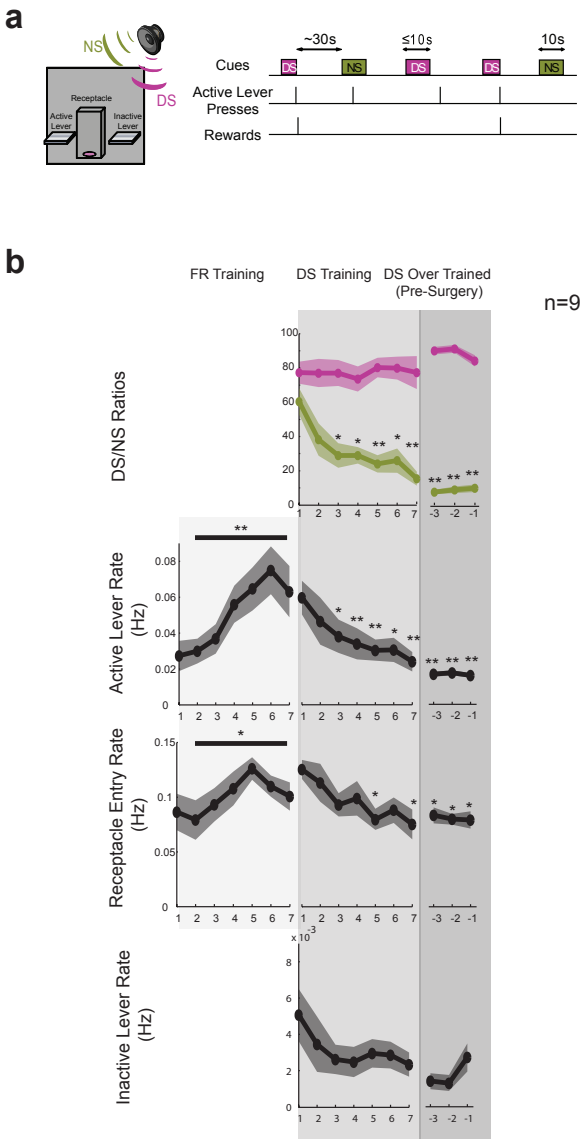

Figure 3.1: **Schematic of the DS task and performance during training** (**a**) Two cue tones (up to 10 s for the rewarded discriminative stimulus [DS]; 10 s for the unrewarded stimulus [NS]) were randomly presented on a variable-interval schedule with an average interval of 30 s. A lever press was required during DS presentation to terminate the DS and cause the delivery of a 10% sucrose reward into an adjacent receptacle. (**b**) The rate of active lever presses and receptacle entry significantly increases during the first 7 days of cued FR1 training (left column). Rats learn to suppress responding to NS and unnecessary lever pressing and receptacle entry is decreased significantly without affect the DS responding (middle column). Stabilization of performance with high DS and low NS ratios along with lower rate of lever presses and receptacle entry (right column).

there was no change in the rate of rewards earned as evident by the persistence of high DS ratios. Thus the reduction in the active lever presses was mainly due to reduced

responding outside the DS. Analysis of changes in the DS and NS ratios during the second phase showed similar trends of reduction for NS ratio from a high of about 60% (not different from DS ratio, post hoc $P>0.05$) to less than 20% over 7 sessions (2-Way ANOVA Cue $F(1,102)=318$ $P<0.001$, Session $F(6,102)=4.2$ $P<0.001$, Cue x Session $F(6,102)=4.8$ $P<0.001$). The same trend for reduction of inactive lever presses was observed although this effect did not reach significance (ANOVA $F(6,47)=1.3$ $P=0.27$).

All animals received DS training until reaching criterion (>80% DS ratio and <20% NS ratio) for at least 3 consecutive sessions. Prior to surgery the NS ratio as well as rates of lever pressing and receptacle entry remained significantly smaller than the first session of DS training (**FIG 3.1b** third column). These behavioral measures remained stable during this pre surgery stage indicating that the development of response suppression is completed.

Together the data shown in **Figure 3.1b** clearly shows the two components of learning: a response promoting component that develops initially and a response suppressing component that restricts active lever presses to times of reward availability and inhibits other task-irrelevant actions.

3.3.2 vmPFC Inactivation: Behavioral Results

If vmPFC mediates response suppression then its inactivation should result in disinhibition of the same responses that rats had learned to suppress during the second phase of training.

General disinhibition of unrewarded actions following vmPFC inactivation: Consistent with our previous findings (Ishikawa et al. 2008a) inactivation of vmPFC disinhibited NS responding (**FIG 3.2b** left column). On the other hand, the DS ratio was not affected. The DS ratio remained significantly higher than the NS ratio and the latency of response to the DS remained lower than that of the NS indicating that the animals still discriminated between DS and NS. (2-Way Ratio ANOVA Cue $F(1,114)=465$ $P<0.001$, Injection $F(1,114)=1.17$ $P=0.28$, Cue x Injection $F(1,114)=16$ $P<0.001$; 2-Way Latency ANOVA Cue $F(1,111)=4.37$ $P=0.04$, Injection $F(1,111)=0.03$ $P=0.85$, Cue x Injection $F(1,111)=0.25$ $P=0.6$).

Inactivation of vmPFC also resulted in significant increases in the overall rate of active lever presses, receptacle entries and inactive lever presses throughout the session (**FIG 3.2c** left three panels). Interestingly, similar to the NS ratio these three disinhibited actions were among those suppressed during DS training (see **FIG 3.1b** second column). This disinhibition of responding also led to a significant reduction of efficiency as shown by the increase in the number of active lever presses per reward (**FIG 3.2c** last panel on right).

Temporal specificity of behavioral disinhibition following vmPFC inactivation: Both the general disinhibition of active lever presses and the increase in NS ratio following vmPFC inactivation raises the question about whether the disinhibitory effect has temporal specificity. Is the disinhibition of lever pressing uniform over time or is it dependent on temporally specific task events such as presentation of cues or recent reward experience? A general increase in lever pressing could account for an apparent increase in NS responding, however, such an increase could be temporally correlated with NS presentation. We divided the DS task into 4 epochs to quantify the temporal specificity of active lever press changes (**FIG 3.2d**). These epochs include the DS and NS periods, the Post_reward period which includes all the active lever presses that happen in the 10 sec following rewarded receptacle exit and Spont period which includes spontaneous active lever presses happening outside the other 3 epochs. There was a strong modulation of active lever presses over epochs and injections in both CSF and B/M condition (2-Way CSF ANOVA Injection $F(1,182)=8$ $P<0.01$, Epoch $F(3,182)=1780$ $P<0.001$, Injection x Epoch $F(3,182)=10$ $P<0.001$, 2-Way B/M ANOVA Injection $F(1,217)=30$ $P<0.001$, Epoch $F(3,217)=398$ $P<0.001$, Injection x Epoch $F(3,217)=25$ $P<0.001$). Presentation of the DS induced a rate of active lever pressing higher than any other epoch in both CSF and B/M conditions (post hoc $p<0.001$) (**FIG 3.2e**). However, after inactivation of vmPFC there were significant temporally specific increases in the lever pressing during NS and Post_reward epochs (post hoc $p<0.001$ compared to pre-injection period) (**FIG 3.2e** right column). Inactivation of vmPFC also resulted in a trend toward increase in active lever presses during the Spont epoch. Furthermore the increase in active

lever pressing during the NS and Postreward epochs were significantly higher than the Spont epoch (post hoc $p<0.001$). This supports the hypothesis that the observed increased NS ratio after vmPFC inactivation is not due to a general increase in lever pressing across all epochs. It should be pointed out that in both injection conditions the active lever press rate during the DS showed a moderate yet significant reduction (post hoc $P<0.001$ CSF and $P<0.01$ B/M).

Analysis of receptacle entries in the same time windows revealed overall similar effects (**FIG 3.3a**) suggesting that the vmPFC controls both active lever press and receptacle entry components of of reward-seeking behavior. Interestingly, vmPFC inactivation also increased the rate of premature receptacle entries during the DS that happen before active lever presses indicating that these two components of reward approach can be emitted independently and that some aspects of impulsivity such as inappropriate choice or sequencing of actions might also be manifested during the DS epoch (**FIG 3.3b**).

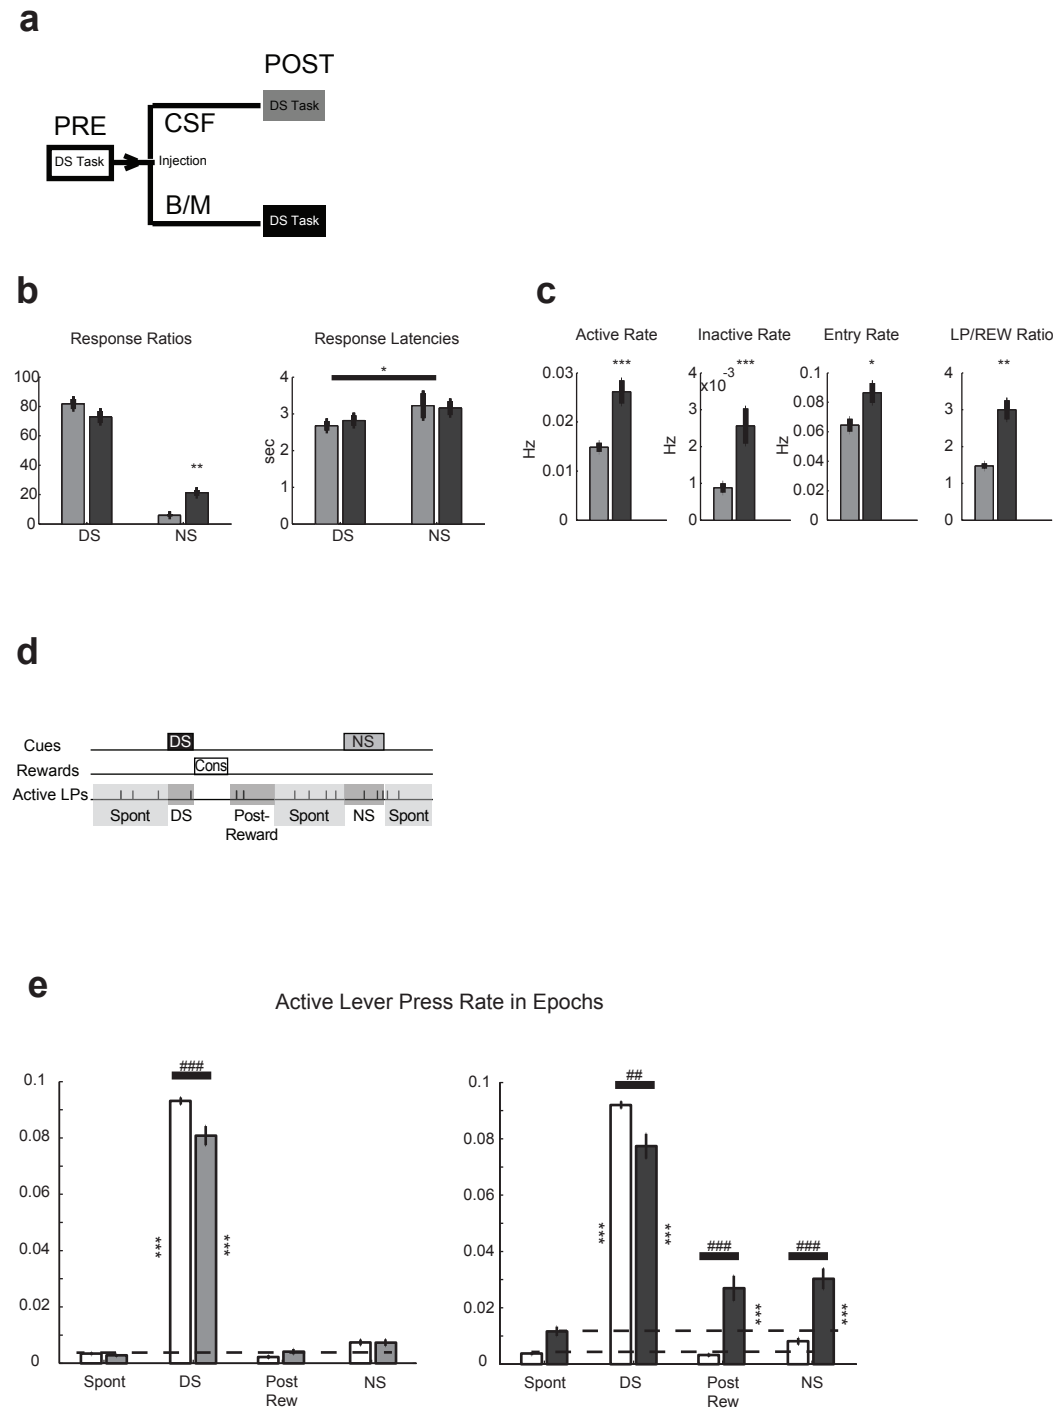

Figure 3.2: **Injection paradigm and behavioral results** (**a**) rats were run in the DS task for 1-2 hours before cortical microinjection (PRE) then were taken out randomly and received injection of aCSF or B/M in the vmPFC and immediately put back for another 1-2 hours following injection (POST) (**b**) Rates of responding to DS and NS and latency of responses in the post aCSF (gray) or B/M (black) injections. (**c**) Rates of active and inactive lever presses, receptacle entry and active lever presses per reward respectively from left to right for aCSF and B/M post-injection. (d) Depiction of epochs used for behavioral microanalysis. (**e**) Rates of active lever presses during each epoch for pre- and post-injection periods for both aCSF (left) and B/M (right) conditions (Total sessions: n=27 for aCSF and n=32 for B/M).

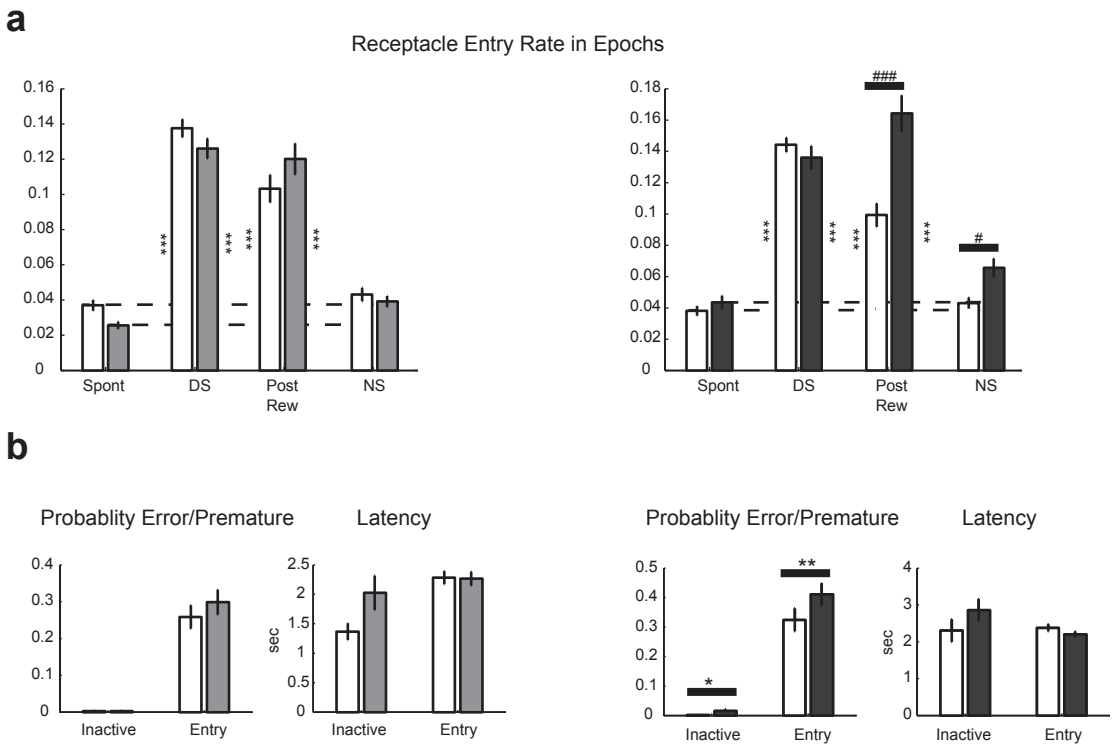

Figure 3.3: **Behavioral disinhibition as measured by receptacle entry** (**a**) Rates of receptacle entries during each epoch for pre- and post-injection periods for both CSF (left) and B/M (right) conditions. (Total sessions: n=27 for CSF and n=32 for B/M) (**b**) Probability and latency of inactive lever press or premature entry before active lever press during DS presentation for CSF (left two columns) and for B/M (right two columns) conditions.

Gating model or impulse model? We earlier posited that the optimal responding in the DS task requires both promoting relevant and inhibiting irrelevant actions. An important question is how these two processes interact throughout the task epochs to promote accurate and efficient acquisition of the desired outcome. There are at least two distinct models that could generate the learned pattern of optimal responding we observed in the DS task (**FIG 3.4a**). The first model here called the "gating model" assumes a constant positive motivational process that increases the probability of responding throughout the task. The actual temporal pattern of responding posited in this model is produced by timing of periods when the suppression strength is reduced, similar to opening a gate that allows responding. The second model we call the "impulse model" proposes that motivated actions are produced when the process that promotes responding rises to exceed the constant background levels of suppression. Both models assume that responding is proportional to the magnitude of the difference between opposing processes and only happens when the difference is positive.

If we assume that the inactivation of vmPFC effectively eliminates the suppressive process the temporal rate of responding in this condition should specifically reveal the shape of the promotive process. Even though we cannot yet directly measure the dynamics of the action suppressive process its shape can be inferred by subtracting the rate of responding in the B/M condition which is affected solely by the promotive process from the CSF condition which is affected by both promoting and suppressing processes (**FIG 3.4b**). The temporal dynamics of promotion and suppression suggests

a combination of gating and impulse model in motivated responding in the DS task. This analysis showed that suppression was negligible during the DS but increases during Post_reward and NS. Therefore these results suggested that the reduction of suppressive process following vmPFC inactivation should not affect the neural representation of events during the DS epoch in NAcS but should have considerable consequences during other epochs. We tested this hypothesis using our simultaneously recorded neural data.

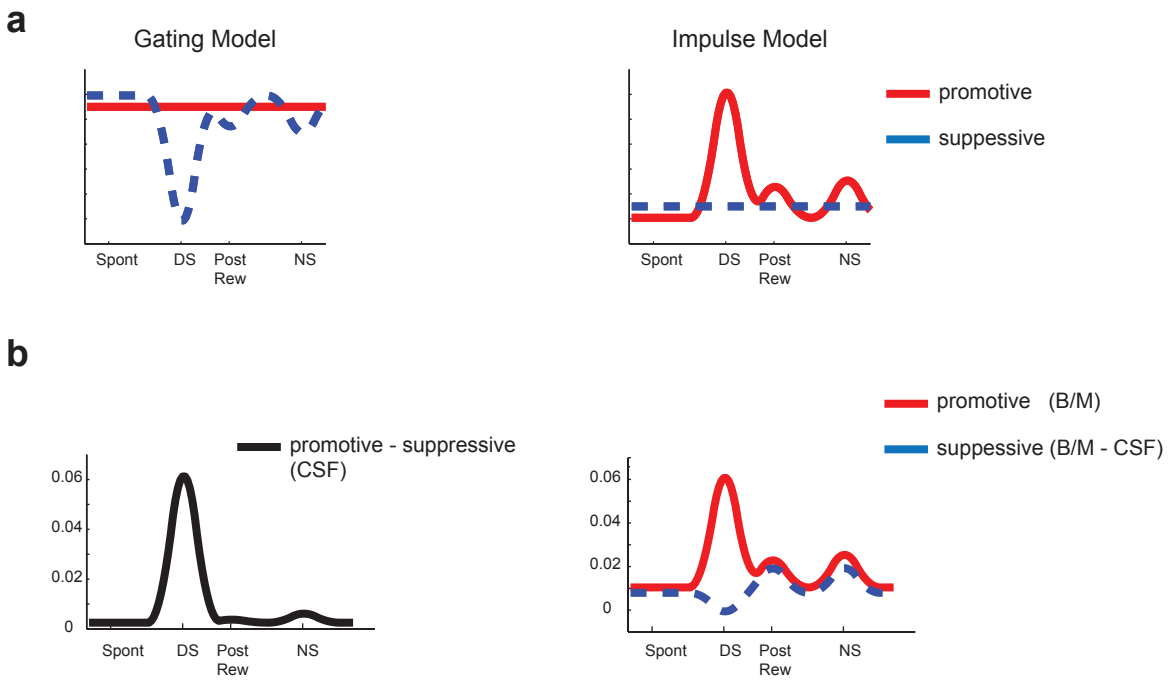

Figure 3.4: **Models of motivated responding in the DS task** (**a**) schematics of gating model (left) and impulse model (right) of interactions between the promoting and suppressive process. (**b**) Interpolated and smoothed rate for active lever pressing during various task epochs in the CSF condition (left) and B/M condition (right red trace). Response rate during the B/M is assumed to be solely due to promoting component. Suppressing component can be inferred by comparing the CSF and B/M conditions (right blue dashed line).

3.3.3 vmPFC Inactivation: electrophysiological recordings in NAc shell

We recorded a total of 178 neurons in the B/M condition and 128 neurons in the CSF conditions during the DS task in 32 and 27 sessions, respectively. After exclusion of neurons with very low baseline firing rate or the neurons that were not present in both pre- and post-injection epochs, 131 and 98 neurons, respectively, in B/M and CSF condition remained in our sample for further analysis. Here we examine the responses of these neurons in each condition before and after injection.

No net change in the representation of the reward related events: On the majority of trials after the DS was presented the rats pressed the active lever and entered the receptacle to consume the sucrose reward. Immediately following reward consumption rats exited the receptacle. Inactivation of vmPFC did not affect this aspect of the DS task. At the neuronal level, DS correlated activity in the NAcS was also unaffected by vmPFC inactivation. Such inactivation did not produce significant changes in either the percentage of DS responsive neurons or the magnitude of their response to the DS, rewarded lever press, receptacle entry, consumption or receptacle exit (**FIG 3.5**).

This result is interesting in the light of our behavioral results suggesting that the contribution of vmPFC to NAcS neuronal firing in the DS epoch would be negligible. However, despite the lack of change in the phasic firing that is temporally correlated with rewarding events, we did observe a significant reduction of the baseline activity in neurons showing phasic inhibitions to rewarded actions (post hoc rewarded lever press

P=0.02, rewarded entry P=0.008, consumption P=0.02, rewarded exit P=0.003). This means that the firing of these neurons is normally increased by vmPFC inputs during periods outside the DS epoch.

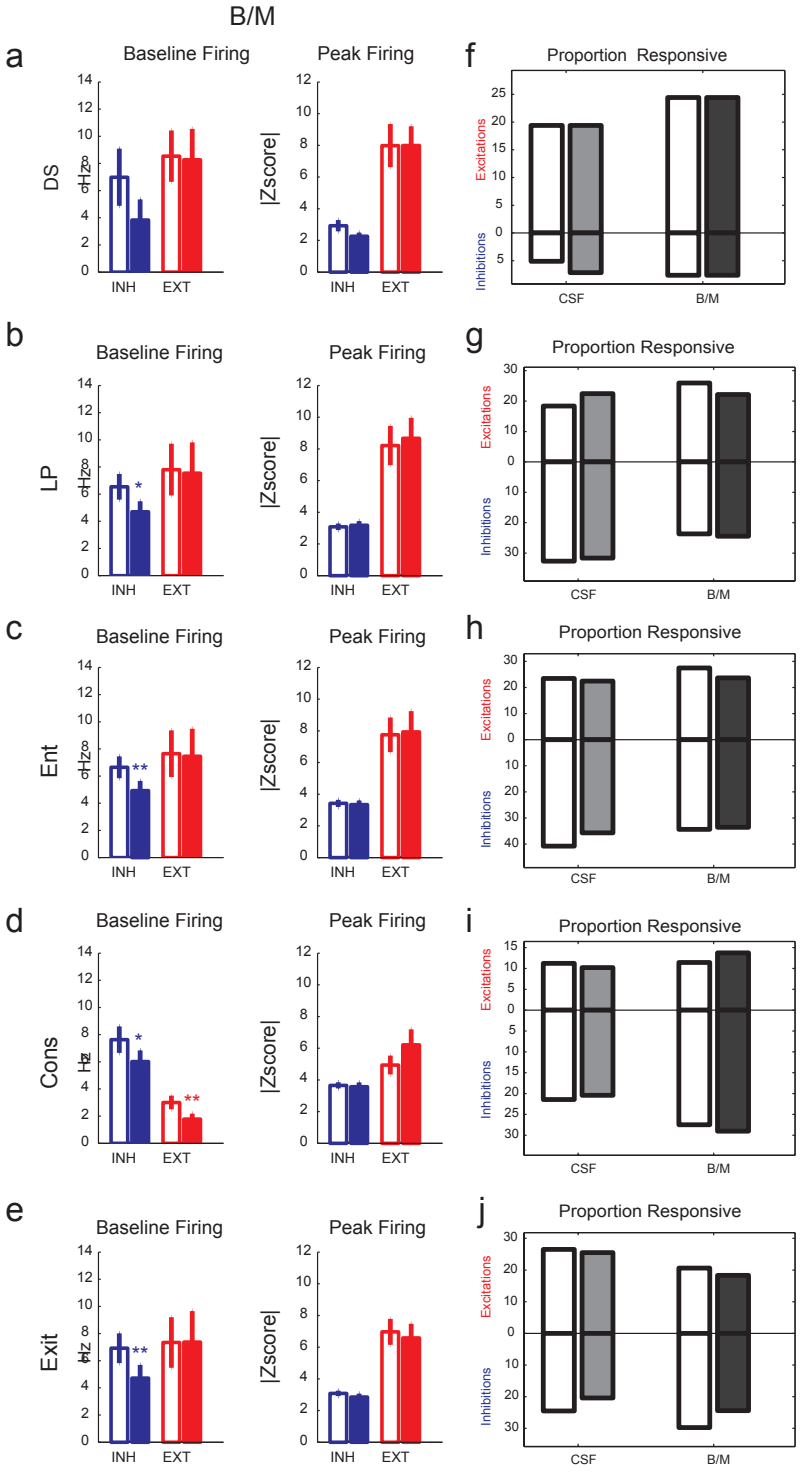

Figure 3.5: **Encoding of sequence of actions in the rewarded trials following DS presentation before and after vmPFC inactivation** (**a-e**) Changes in the baseline (first column) and the peak firing Z-score (second column) of the neurons responsive to rewarded DS, rewarded LP, rewarded receptacle entry, reward consumption and rewarded exit in either pre- or post-injection period in the B/M condition. Blue is inhibition and red is excitation. (Total neurons in first and second column: n=37,14 for DS+,-; n=35,40 for LP+,-; n=40,56 for Entry+,-; n=27,44 for Cons+,-; n=32,47 for Exit+,-) (**f-j**) Percentage of neurons responsive to rewarded DS, rewarded LP, rewarded receptacle entry, reward consumption and rewarded exit in the pre- and post-injection periods in CSF and B/M conditions.

vmPFC inactivation enhances the neural representation of unreinforced events: Unlike the neural activity correlated with the DS, we observed significant changes in encoding of unrewarded actions outside the DS. There was a significant increase in Z-scores of the peaks of both excitations and inhibitions correlated with the unrewarded lever press following injection of B/M in vmPFC (**FIG 3.6a**) (post hoc $P<0.05$ and $P<0.001$ for excitation and inhibition respectively). In the neurons that encoded unrewarded receptacle entry a significant reduction in the baseline firing for inhibited neurons following vmPFC inactivation was observed (**FIG 3.6b**) (post hoc $P<0.001$). Similar trends toward reduction of basal firing was also present for unrewarded lever press inhibited neurons (post hoc $P=0.052$). There was also a large increase in the percentage of unrewarded lever press inhibited neurons (**FIG 3.6c**) (+238%, permutation test $P<0.001$, compared to the CSF condition). vmPFC inactivation produced no significant change in percentage of neurons responsive to unrewarded entry in B/M compared to CSF condition (permutation test $P>0.05$) (**FIG 3.6d**).

Importantly, the increase in lever-pressing during the NS following vmPFC inactivation (**FIG 3.2b,e**) was associated with a significant increase in NS-evoked firing in a subset of NAcS neurons. With vmPFC intact, the proportion of NS responsive neurons was low. Inactivating the vmPFc inactivation recruits a new population of NS excited neurons. Overall, the percentage of NS-excited neurons increased by 58% after inactivation (permutation test $p=0.019$, **FIG 3.7b**). Furthermore the peak firing Z-score of NS excitations showed a significant increase after vmPFC inactivation

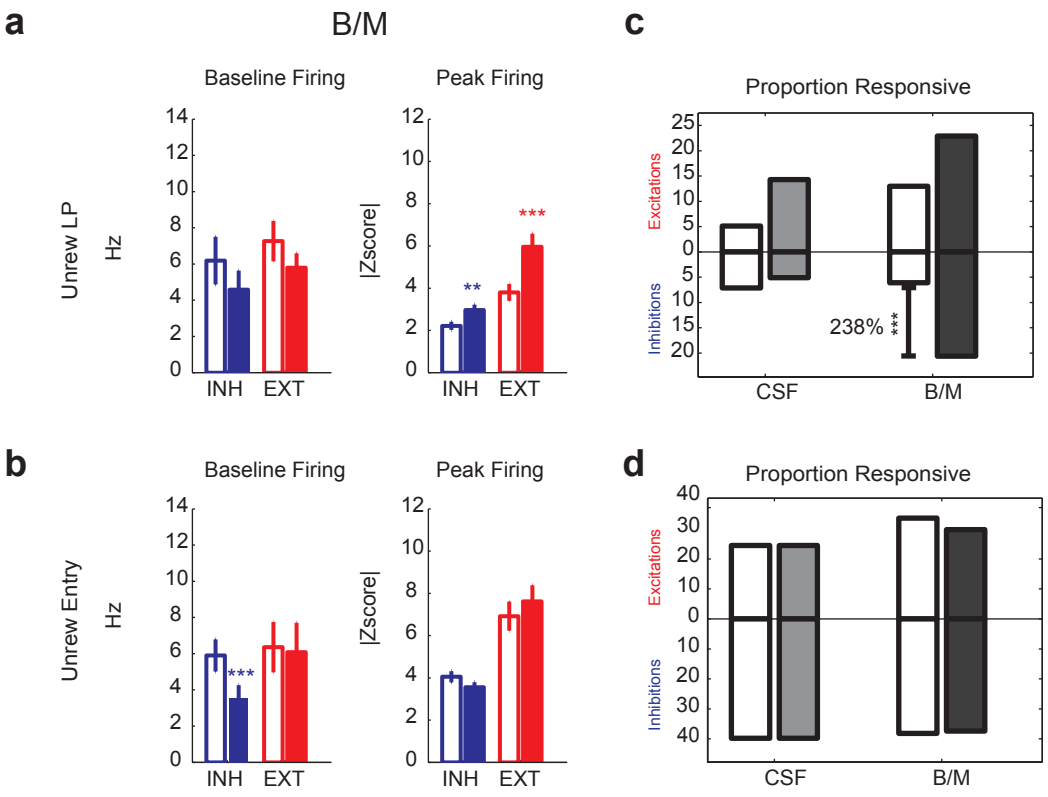

Figure 3.6: **Encoding of unreinforced actions in NAcS neurons before and after vmPFC inactivation** (**a-b**) Changes in the baseline (first column) and the peak firing Z-score (second column) of the neurons responsive to unrewarded lever press (top row) or unrewarded receptacle entry (bottom row) in either pre- or post-injection period in the B/M condition. Blue is inhibition and red is excitation. (Total neurons: n=32,29 for Unrew LP+,-; n=45,54 for Unrew Entry+,-) (**c-d**) Percentage of neurons responsive to unrewarded lever press (top row) or unrewarded receptacle entry (bottom row) in the pre- and post-injection periods in CSF and B/M conditions.

(post hoc p=0.025 **FIG 3.7a**). **Figure 3.7c** shows an example neuron with emergent phasic short latency NS excitation following vmPFC inactivation. Interestingly, all the neurons with emergent NS excitations were also excited by the DS suggesting that their activation promotes reward predictive cue responding. However, despite the clear enhancement of NS-excitations by vmPFC inactivation in neurons showing either emergent (n=9) or enhanced (n=3) NS excitations, the average DS response for this same population of neurons was not significantly changed by vmPFC inactivation, suggesting that vmPFC exerts an inhibitory effect on NS excitations that

is specific to both the sensory qualities and timing of the unrewarded cue neurons (**FIG 3.7d**). NS-responsive neurons in our task were distributed throughout the medial shell (**FIG 3.8**) but there was no relationship between the magnitude of the firing changes in the B/M condition and the rostrocaudal or dorsomedial axes or their interactions (R^2=0.2, P=0.26).

Since the rate of lever-pressing increased after vmPFC inactivation, the emergence of NS excitations could be coincidental and not strictly correlated in time with the occurrence of the NS. However, some observations ruled out this possibility. First, NS behavioral responses appeared at short latency to cue onset (<0.5 sec compared to >2sec average lever press latency to NS). Second, NS emergent neural responses remained even if we only included trials when no lever pressing occurred during the NS (**FIG 3.7c** overlaid red line,**FIG 3.8**). Furthermore, in some sessions, we performed unilateral vmPFC inactivations to investigate the laterality of NS emergence. Overall 17 and 22 neurons were recorded ipsi- and contra-lateral to the injection site. Consistent with bilateral vmPFC projection to the NAcS, emergent NS excitations were found both ipsi and contra-lateral to the side of inactivation (**FIG 3.7e**).The increase in the total number of NS excitations following unilateral inactivations was significant compared to the CSF condition after combining the results of both ipsi and contralateral conditions (the percentage of emergent NS excitations were similar in ipsi- and contra-lateral sites; 2 out of 17 for ipsi and 3 out of 22 for contra, permutation test P=0.73) (**FIG 3.10b** top row). Again similar to the bilateral inactivation there was no significant change in the percentage of DS encoding neurons following

unilateral inactivations (**FIG 3.10b** bottom row). Interestingly, unilateral vmPFC inactivations did not induce significant behavioral disinhibition (n=7 sessions, **FIG 3.10a**). This further strengthens our previous argument, that the changes in the NS firing are not secondary to increased lever pressing.

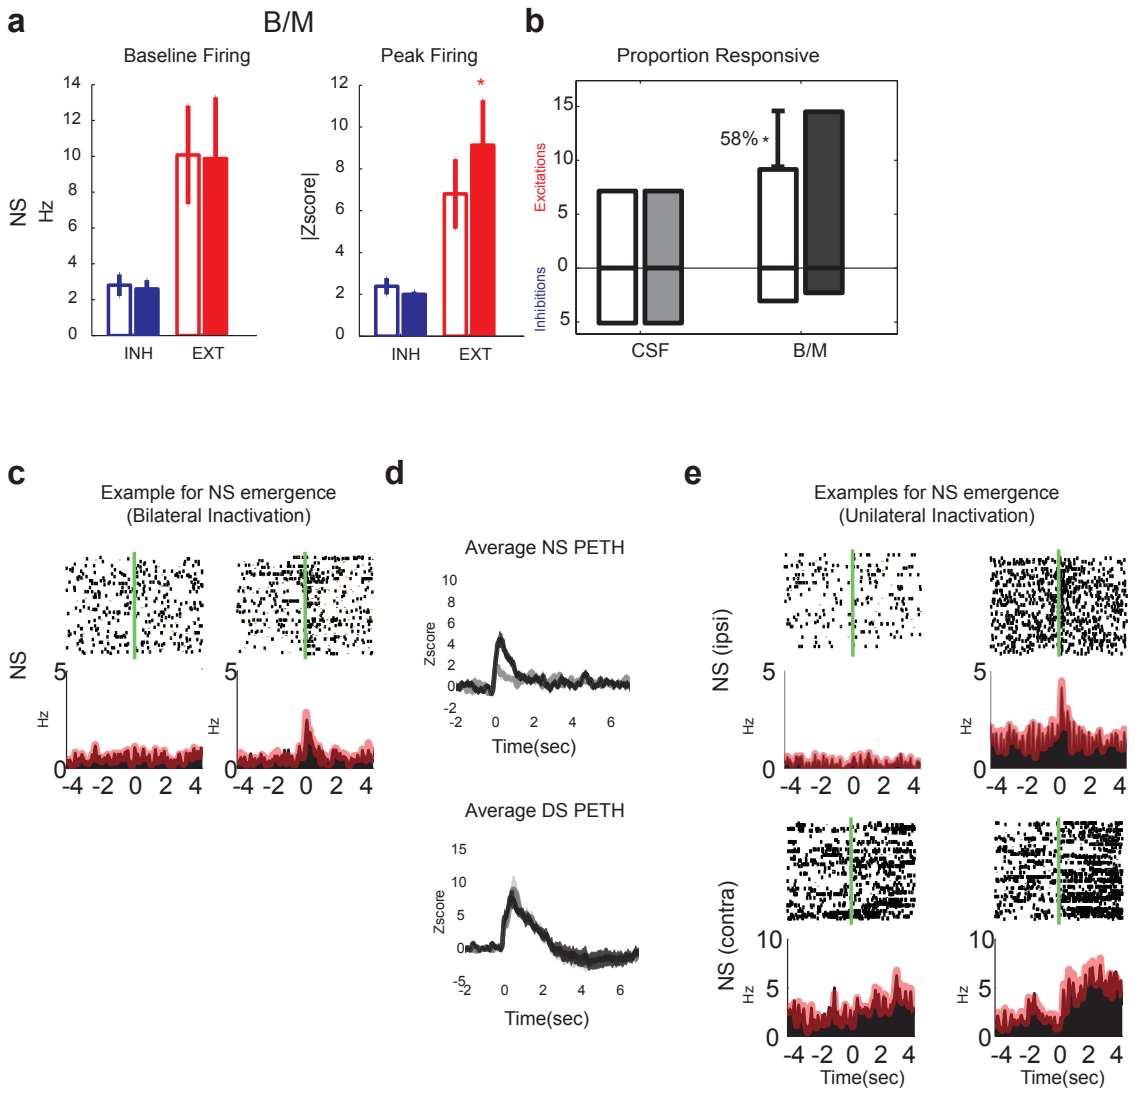

Figure 3.7: **Encoding of neutral stimuli in NAcS before and after vmPFC inactivation** (**a**) Changes in the baseline (first column)and the peak firing Z-score (second column) of the neurons responsive to NS (top row) or DS (bottom row) in either pre- or post-injection period in the B/M condition. (Total neurons: n=21,7 for NS+,-;) (**b**) Percentage of neurons responsive to NS (**c**) Example neuron showing emergent NS excitation following bilateral vmPFC inactivation. Overlaid red line shows the PETH made with NSs not followed with a behavioral response (top row) (**d**) Average Z-score of NS PETHs for neurons showing emergence or significant increase in excitations to NS before (gray) and after (black) vmPFC inactivation (top row) and average Z-score of DS PETHs for the same neurons as above (bottom row). (**e**) Example neurons showing emergent NS excitation following (top row) ipsi- or (bottom row) contra-lateral vmPFC inactivations. Overlaid red line shows the PETH with NSs not followed with a behavioral response.

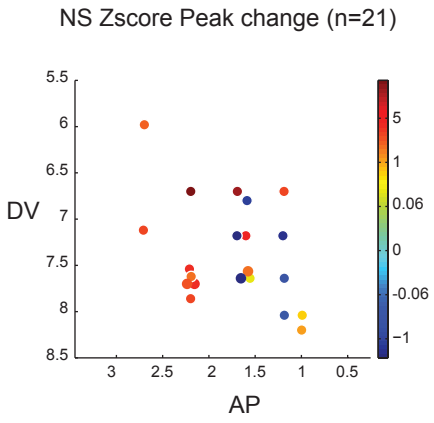

Figure 3.8: **location of the NS responsive neurons.** Pre- or post-injection for B/M condition color-coded for changes in the peak Z-score firing to NS following vmPFC inactivation (red increases, blue decreases).

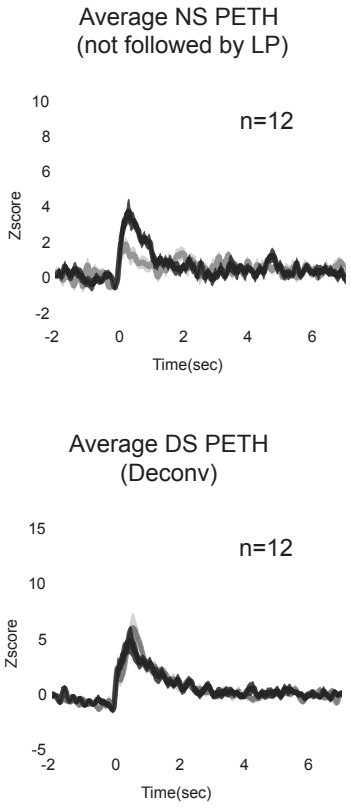

Figure 3.9: **Increased NS excitation is not secondary to increased lever pressing.** Average Z-score of NS PETHs for neurons showing emergence or significant increase in excitations to NS before (gray) and after (black) vmPFC inactivation using only NSs not followed by behavioral responding (top) and average Z-score of deconvolved DS PETHs for the same neurons (bottom).

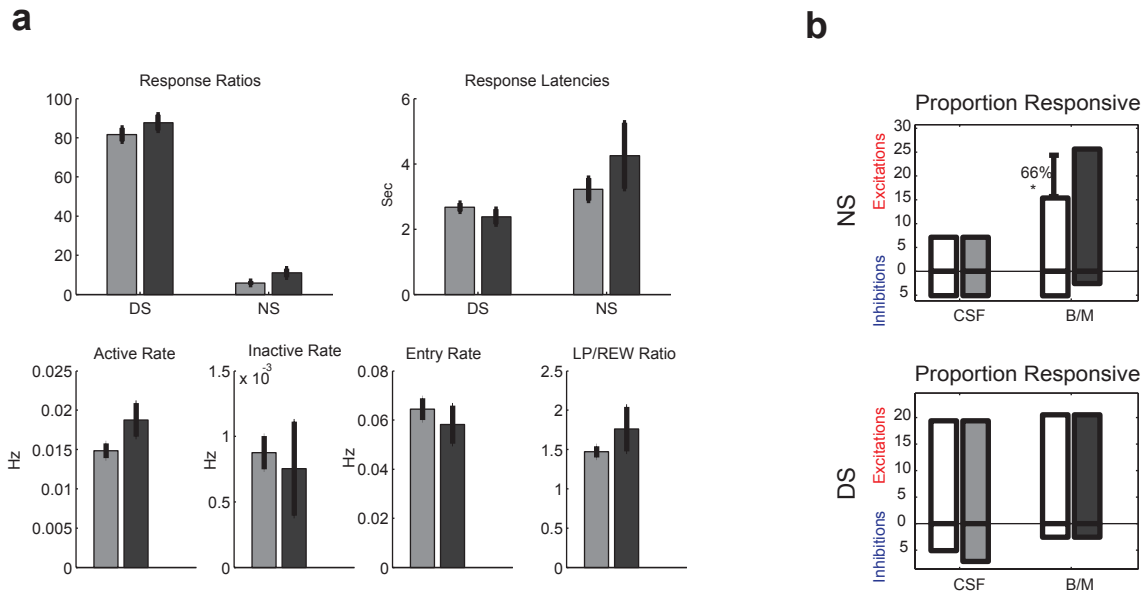

Figure 3.10: **Comparison of unilateral B/M and bilateral CSF injection in vmPFC**(**a**) Rates of responding to DS and NS and latency of responses in the post CSF (gray) or unilateral B/M (black) injections (top row) and rates of active and inactive lever presses, receptacle entry and active lever presses per reward respectively from left to right for CSF and unilateral B/M post-injection (Total sessions: n=27 for CSF and n=7 for unilateral B/M) (**b**) Changes in the percentage of neurons responsive to NS (top row) or DS (bottom row) in the pre- and post-injection periods in CSF and unilateral B/M conditions.

Differential modulation of action encoding neurons: The reduction of the basal firing rate of action inhibited neurons and the emergence of unrewarded lever press inhibited neurons is consistent with a reduction of excitatory input to the NAcS following vmPFC inactivation. However, the increased magnitude of NS excitations indicates that some NAcS neurons are normally inhibited by vmPFC during the NS and that there is differential modulation of action excited neurons and action inhibited neurons by vmPFC. Consistent with this hypothesis, the neurons whose basal firing increased following vmPFC inactivation were, on average, excited prior to unrewarded receptacle entries while those with decreased basal firing were inhibited (**3.11** top row). Also neurons with significant increases in baseline firing on average showed excitation to NS while neurons with significantly reduced baseline were not NS responsive. This was consistent with inhibition by vmPFC of NS excited neurons that presumably drive NS responding (**3.11** bottom row).

One simple hypothesis is that the firing of action excited neurons promotes responding while the firing of action inhibited neurons suppresses responding. If this hypothesis is true then the magnitude of the difference between the firing of action excited minus action inhibited neurons should also be similar to the changes in the response rate during the DS task epoch and following vmPFC inactivation. The average response of action excited and action inhibited neurons in the both CSF and B/M conditions are shown (**FIG 3.12a** top row left and right respectively). In fact plotting the firing difference between neurons that were action excited ($\sim$30%) and the neurons that were action inhibited ($\sim$50%) over the 4 behaviorally relevant epochs

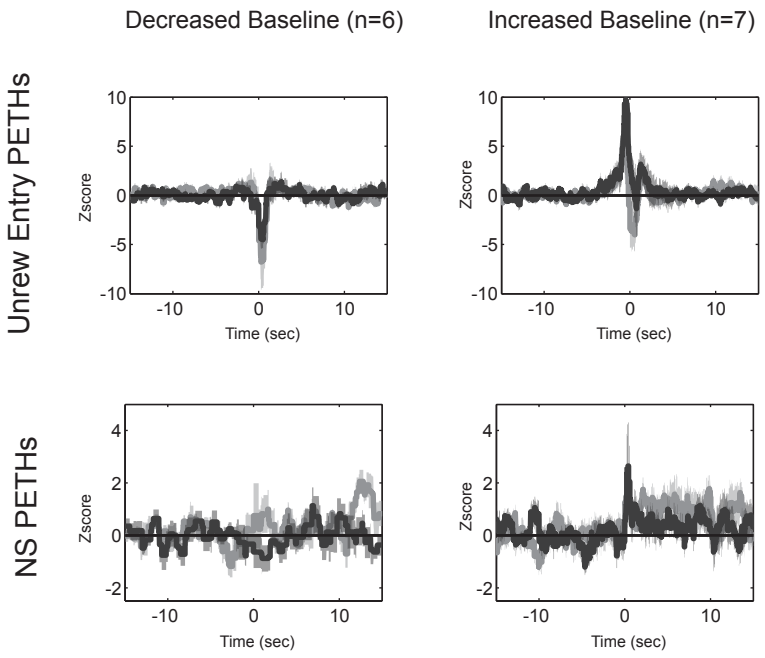

Figure 3.11: **Neurons with significant increase and decrease in baseline are on average excited and inhibited to unrewarded events** Responses of neurons with significant decreases (left column) and increases (right column) in their baseline following vmPFC inactivation to unrewarded receptacle entry (top row) and NS (bottom row). PETHs during pre- (gray) and post-injection (black) are overlaid.

revealed a pattern similar to the behavioral lever pressing data in **Figure 3.2e** in both injection conditions (**FIG 3.12b**). In the CSF condition this difference was significantly higher than zero only during the presentation of DS in both pre- and post-injection periods (post hoc $P<0.001$). This was consistent with high rate of lever pressing during the DS. The same pattern was observed in the pre-injection period in B/M condition. However following the vmPFC inactivation there was a significant increase in the magnitude of difference between firing of action excited and action inhibited neurons similar to the behavioral increase in active lever pressing (2-Way CSF ANOVA Injection $F(1,462)=1.08$ $P=0.29$, Epoch $F(3,462)=18$ $P<0.001$, Injection x Epoch $F(3,462)=0.13$ $P=0.9$, 2-Way B/M ANOVA Injection $F(1,749)=21$ $P<0.001$, Epoch $F(3,749)=7.3$ $P<0.001$, Injection x Epoch $F(3,749)=0.2$ $P=0.9$)).

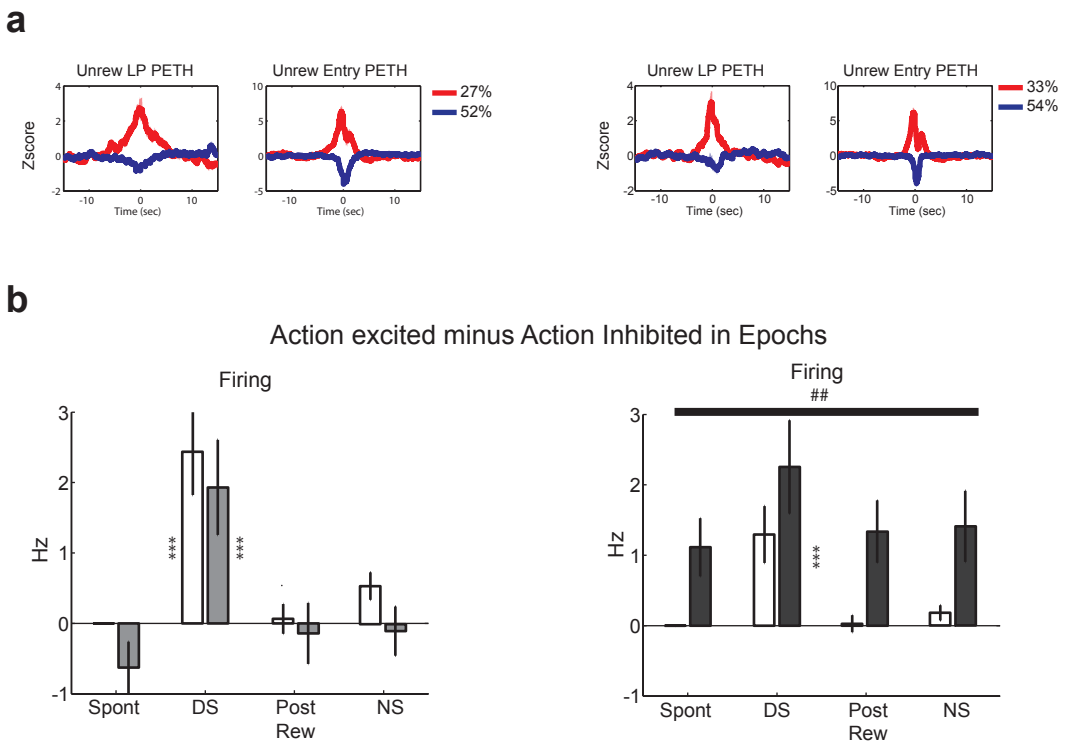

Figure 3.12: **Firing difference between action excited and action inhibited neurons during task epochs** (**a**) Average PETH Z-scores for neurons that respond to either unrewarded lever press or unrewarded entry by excitation (action excited: red) or inhibition (action inhibited: blue) in CSF or B/M conditions (left and right respectively) (**b**) modulation of the firing difference between action excited and action inhibited neurons across trial epochs before and after injection of CSF or B/M (left and right respectively).

3.4 Discussion

Here we demonstrate that successful performance in a learned appetitive task requires at least two distinct processes mediated by the nucleus accumbens: a process that promotes and invigorates performance of actions leading to reward (promotive process) and a complementary process that inhibits responding that is irrelevant to reward acquisition and consumption (suppressive process). In the DS task used in this study the promotive process is learned early during training while the suppression of active lever presses outside the DS presentation emerges later (**FIG 3.1b**). We previously demonstrated that the promotive process in the DS task requires projections to the NAc core from the ventral tegmental area (VTA), basolateral amygdala (BLA) (Ambroggi et al. 2008; Yun et al. 2004) and dorsomedial prefrontal cortex including the prelimbic area (PL) (Ishikawa et al. 2008b) .

In contrast, inactivation of the vmPFC (largely the infralimbic area (IL) disinhibits task irrelevant actions indicating that it normally inhibits NS responding in the DS task (Ishikawa et al. 2008a). Consistent with this finding, vmPFC also inhibits premature and perseverative errors in a 5-choice serial reaction time task (Chudasama et al. 2004; Muir et al. 1996). Furthermore, in the fear conditioning paradigm phasic activation of vmPFC neurons including IL inhibits freezing to a tone previously paired with foot-shock (Milad and Quirk 2002). The current study extends understanding of the mechanisms by which vmPFC can inhibit actions that rats have learned are irrelevant and potentially in conflict with to the goals of the task in which they are engaged (**FIG 3.1b**).

Analysis of active lever pressing following vmPFC inactivation reveals disinhibition of responding in all epochs of the task except the DS period. However the increased rate of active lever pressing is not uniform over time. For instance we find a greater increase in responding during NS compared to the Spont epoch. This result suggests that the promotive process is stronger during NS than Spont periods. In fact, our analysis of the temporal dynamics of promoting and suppressing processes during the DS task reveals little evidence for a suppressive process during the DS while an active suppressive process powerfully opposes the promotive process for the active lever press during unrewarded epochs. The overall pattern of responding indicates the existence of promoting tendencies along with temporally specific reduction of suppression (gating) during the reward consummatory period (**FIG 3.4**).

Our previous studies indicate that a subpopulation of NAc neurons promotes responding ('cue incentive neurons' see (Ambroggi et al. 2008; Yun et al. 2004)) while another, larger subset (gating neurons) show pauses in their firing beginning just prior to reward approach and consumption (Taha and Fields 2006) in learned appetitive tasks. We recently reported that both types of neurons are found in NAc core and shell but that while shell inactivation produces robust disinhibitory phenomena; unlike core inactivation, it has no effect on DS responding (Ambroggi et al 2011). In the current study we chose to focus on the suppressive process and therefore studied the influence of vmPFC on NAcS neuronal activity. In fact, anatomical and functional data show that striatal projections arising from the IL mainly target the medial NAcS and that these projections are required for suppression of unrewarded responding (Christakou

et al. 2004; Peters et al. 2008; Voorn et al. 2004). The current work extends those studies by showing how different subsets of NAcS neurons contribute differentially to inhibition of task irrelevant actions.

Given the role of shell in suppression of unrewarded actions and the glutamatergic projections from the IL we hypothesized that inactivation of IL would reduce the firing of a subpopulation of shell neurons whose firing is inversely correlated with actions. Consistent with this proposal there is a significant reduction in the basal firing of neurons that are inhibited with unrewarded receptacle entry and a trend toward reduction for unrewarded lever press inhibited neurons (**FIG 3.6a,b**). Furthermore, following vmPFC inactivation there is a significant increase in the number of NAcS neurons showing inhibition when the rat makes unrewarded lever presses (**FIG 3.6c**). The strength of phasic inhibitory responses correlated with unrewarded lever press as measured by peak Z-score firing is also significantly increased in the post- versus pre-injection period in the B/M condition (**FIG 3.6a**). This is consistent with our recent result that implicates firing pauses in NAc neurons in gating (or permitting) appetitive actions (Krause et al. 2010; Taha and Fields 2006).

We also discovered a significant increase in the average peak of unrewarded lever press excitations following inactivations (**FIG 3.6d**). This argues that inactivation of vmPFC concomitantly increases the phasic firing of unrewarded lever press excitation and decreases the firing of unrewarded lever press inhibitions. Consistent with this differential modulation of these two subgroups of NAcS neurons by vmPFC an analysis of changes in the basal firing rates compared to the CSF condition shows

that neurons that significantly increase their basal firing after vmPFC inactivation are on average excited prior to unrewarded receptacle entry while the population with significant reduction of baseline on average are inhibited around receptacle entry (**FIG 3.11**). Further analysis confirms that vmPFC inactivation increases the difference between the firing of neurons showing action excitations and those showing action inhibitions over 4 distinct behaviorally relevant epochs and that this firing difference resembles the increased occurrence of active lever presses (compare **FIGs 3.12b&3.2e**).

Another aspect of the observed behavioral disinhibition following vmPFC inactivation has to do with the increase in responding to the NS above that of uncued behavioral responding. Importantly, the increase in behavioral responding to the NS after vmPFC inactivation is accompanied by an increase in the magnitude of NS excitations in the NAcS through a combination of emergent and increased NS excitations. The short latency of these NS neuronal excitations as well as the fact that the increase is present even in the trials when no NS elicited lever pressing occurred argues against the emergent firing being secondary to the increase in NS elicited active lever presses. Rather, this is consistent with the hypothesis that these NS excitations promote behavioral responding to the NS (**FIG 3.7 & 3.9**). The observed emergence of NS excitations following both ipsilateral and contralateral inactivations (**FIG 3.7e & 3.10b**) is consistent with anatomical data showing a bilateral projection from vmPFC to NAc (Sesack et al. 1989).

Although there are other possible pathways by which vmPFC can affect the firing

of NAcS neurons, there are significant direct projections from vmPFC to NAcS. Since these direct projections are glutamatergic, the emergence of NS and unrewarded lever press excitations after inactivation implies the existence of at least one layer of intervening GABAergic neurons that receive the direct excitatory vmPFC input and inhibit the emergent responses under normal conditions. Within the striatum this could be achieved through fast spiking interneurons by fast feedforward inhibitions (Mallet et al. 2005) or by lateral inhibitions from the collaterals of other medium spiny projection neurons (MSNs) (Taverna et al. 2004). For instance, what we observed as a significant reduction of baseline firing in many action inhibited neurons ($\sim$50%) might be responsible for the emergence of excitatory responses to unrewarded cue and actions. Future studies will be required to address this issue.

To explain how vmPFC control of NAcS neurons contributes to action promotion and inhibition we propose a circuit model illustrated in **Figure 3.13**. In this model the action inhibited neurons receive direct excitatory input from vmPFC and act to suppress responding. We further propose that these action inhibited neurons inhibit action excited neurons via local axon collaterals. Based on this model and the behavioral data, this lateral inhibition is reduced normally only during the DS presentation when the vmPFC derived inhibition decreases (**FIG 3.4b**). Removal of this vmPFC derived inhibition on the action excited neurons allows excitatory inputs from other structures such as the VTA and BLA to promote responding during the DS. This model also proposes that inactivation of vmPFC results in emergent excitations to events (e.g. the NS) that were previously suppressed through firing of action inhibited

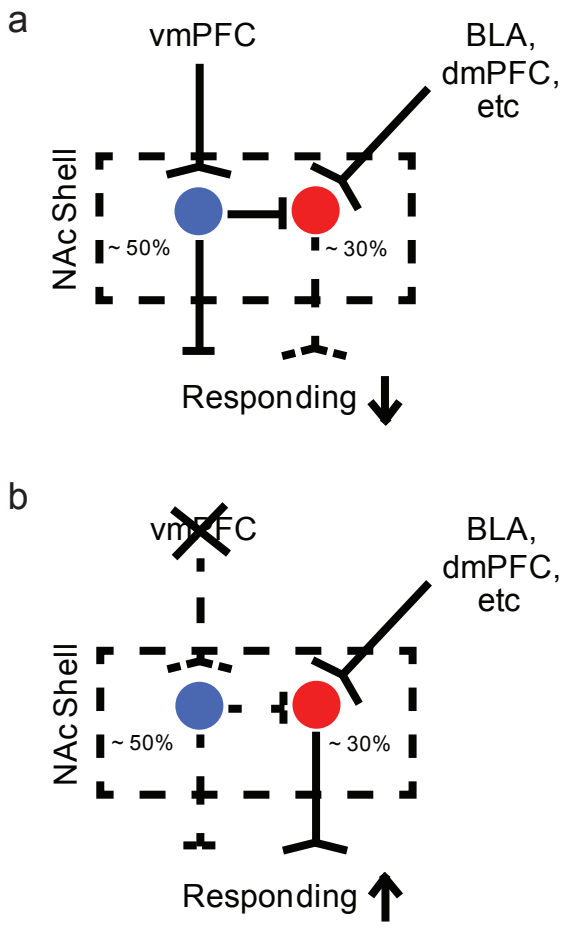

Figure 3.13: **Proposed circuit model for shell neurons that suppress or promote responding** (**a**) vmPFC input excites action inhibited neurons (blue) to suppress responding. This input also indirectly inhibits the firing of action excited neurons (red) to unreinforced cues or actions (**b**) following vmPFC inactivation firing of action inhibited neurons is reduced. Action excited neurons can now be driven by other inputs to promote responding.

neurons.

3.4.1 Conclusion

Here we reveal a promoting and a suppressive process that develop early and late in training respectively. We describe the temporal dynamics of each process during the learned appetitive task and show the suppressive process to be dependent on the intact function of vmPFC. We also demonstrate that vmPFC inactivations change

the balance of neurons that may promote and suppress actions in the NAcS which is a major vmPFC projection target and is itself implicated in response suppression. Specifically vmPFC inactivations increase representation of unrewarded events but not rewarded events. Furthermore consistent with an overall role of NAcS in action suppression we observe a reduction of basal firing in the action inhibited but not action excited neurons. We propose a circuit model in which action inhibited neurons receive direct excitatory input from vmPFC and inhibit responding directly by projecting to downstream areas as well as indirectly by inhibiting action excited projection neurons.

3.5 Materials and Methods

3.5.1 Animals

The subjects were male Long–Evans rats (Harlan Sprague Dawley, Indianapolis, IN) weighing ∼350 g on arrival and individually housed on a 12 hr light/dark cycle. Experiments were conducted during the dark phase. After receipt, rats were allowed at least 1 week of ad libitum food and water, followed by 1 week of restricted food before training. Throughout all experiments, food restriction was adjusted daily at the end of experimental manipulations to maintain the rats at ∼90% of their initial body weight. Animal handling and experiments conformed to National Institutes of Health and Ernest Gallo Clinic and Research Center animal care and use policies.

3.5.2 Operant Chambers

The behavioral pharmacology studies and the electrophysiological studies were conducted in standard operant chambers (23.5 x 30.5 cm). All cages contained two

retractable levers and one reward receptacle located on one wall of the chamber (one on each side of a reward receptacle) (**FIG 3.1a** left). There were also two house lights, a white noise speaker and a tone speaker installed in the chambers (Med Associates,St. Albans, VT). 50μl of 10% of liquid sucrose delivered into a well in the reward receptacle by a syringe pump was counted as a reward and an infrared beam installed across the reward receptacle detected receptacle entry and exits.

3.5.3 Magazine and Cued FR1 Training

The first training session was designed to habituate the rats to the experimental chambers and to introduce them to the liquid sucrose reward delivered into the reward receptacle. After the session started, an entry into the receptacle resulted in delivery of sucrose reward into the receptacle. This would be followed by a 10 sec time out during which additional receptacle entries had no programmed consequences. The magazine session continued until each rat earned 100 rewards which on average took 40 $\pm$ 4 (mean + SD) minutes.

The rats were then switched to a cued FR1 task on the same day or the next. At the beginning of a session one of the two levers was inserted into the chamber and an auditory cue was presented. The randomly inserted lever was designated as the active lever and remained as such during all other phases of training for each rat. The inactive lever was not presented in this stage. The auditory cue was randomly chosen for each rat to be either an intermittent 4 kHz tone (40 ms on and 50 ms off), or a siren tone (ramped from 4 to 8 kHz with a 400 ms period). A press on the lever during the cue resulted in cue termination and delivery of reward into the

receptacle. An ITI would follow the reward delivery during which active lever presses had no programmed consequences. The ITI duration were gradually increased over 7 sessions from 10 seconds to 30 seconds. After ITI expiration the auditory cue would be played signaling reward availability upon active presses. For 4 of the 9 rats there was no maximal cue presentation length while for 5 rats cue duration was gradually shortened across 7 sessions from 60 seconds to 10 seconds. The animals were trained in the cued FR1 task every day or every other day for 2-4 hours.

3.5.4 DS task training

The DS task structure was similar to our previous experiments (Nicola et al. 2004a; Yun et al. 2004). After the rats finished training in the cued FR1 task they were switched to the DS task. Here both the active and inactive levers were inserted into the chamber. Also both the siren and the intermittent tones were presented on a variable interval schedule with an average interval of 30 sec. The tone that was initially used for each rat in the cued FR1 task continued to signal reward availability upon active lever press and is referred to as the discriminative stimulus (DS). The novel auditory cue is referred to as the neutral stimulus (NS). Pressing the active lever during the DS presentation resulted in the reward delivery and termination of the DS tone similar to cued FR1 task. Each DS lasted for up to 10 sec if no active lever press occurred, and each NS lasted 10 sec regardless of the rat actions. Responding on either lever during the NS or in the absence of the DS was never rewarded in this study.

The animals were trained in the DS task every day for about 2 hours. Surgeries

were performed when the rats reached criterion performance of >80% DS response ratio and <20% NS response ratio (defined as the proportion of all DSs or NSs in the session to which the animal responded).

3.5.5 Surgeries

Bilateral guide cannulae (27 gauge, Plastics One, Roanoke, VA, USA) aimed at the ventromedial PFC were implanted. A 14 degree rostral angle was used for the cannulae to allow enough space for the electrode array and the microdrive. The target cannulae coordinates were AP +4.2, ML ±0.8 and DV -2.8 mm. The injectors were cut to extend 2.5 mm beyond the cannulae therefore the target injection coordinates were AP +2.9, ML ±0.8 and DV -5.15 mm in the vmPFC.

For electrophysiological recordings, both bilateral (4 rats) and unilateral (5 rats) electrode arrays were used. In 2 animals 32 electrodes (16 in each side) were used while all other animals had 16 electrodes, The electrodes (NB Labs, Denison, TX, 50 μm stainless steel wires in 5 animals or Innovative Neurophysiology, Durham, NC, 50 μm tungsten wires in 4 animals) were attached to a homemade microdrive that allowed the entire array to be lowered by 80 μm increments. Target coordinates of the medioposterior electrode of each array for NAcS during surgery was AP +1.2, ML ±0.8, DV -6.5 or -7 mm.

Animals were anesthetized with isoflurane (5%) and placed in a stereotaxic apparatus. Anesthesia was maintained with isoflurane (0.5-2.0%) during surgery. Microdrives were secured to the skull with bone screws and dental acrylic, and wire obturators were inserted into the guide cannulae; the ends of the obturators were

flush with the ends of the guide cannulae. Rats were given at least 7 days of recovery before being retrained on the DS task and habituated to the handling procedures.

3.5.6 Microinjections

To inactivate ventromedial PFC a cocktail of Bacolfen and Muscimol (B/M, $GABA_B$ and $GABA_A$ receptor agonists respectively) was used. To inject animals, the obturators were removed and 30 gauge injector cannulae were inserted into the guides (for unilateral inactivations only one side was used for injection). Injectors extended 2.5 mm below the tip of the cannulae 0.5 μl of artificial CSF (hereafter CSF) or B/M (12.5 ng) was injected over 2 minutes. After a 1 min post-injection period, the injectors were gently removed, the obturators were replaced back.

3.5.7 DS task: vmPFC inactivation and NAc shell recording

After the rats regained the performance criteria (>80% DS response ratio and <20% NS response ratio) following surgery they were introduced to the final phase for data collection. In this phase the rats were run in the DS task for 1-2 hours. They were then removed from the chamber and microinjected in the vmPFC with CSF or B/M according to a randomized schedule for each animal. There was at least a day interval between two consecutive injections. Following the injection the animals were immediately put back in the behavioral chamber and were run in the DS task for 1-2 more hours (**FIG 3.14a**). The behavioral data and electrophysiological recordings were obtained before and after injection during the task performance. Each rat received multiple injections of B/M and CSF between which the electrodes were

lowered to record a new set of units. Therefore all behavioral data presented are from injections that both within and between animals (for histological locations of the injections sites see **FIG 3.14a**). All the bilateral injection data reported here are from 27 CSF and 32 B/M sessions. There were 3 sessions with CSF injections and 8 sessions with B/M injections which are excluded from our analysis due to absence or low rate of behavioral responding following injections (DS<30%).

Sessions with unilateral injections of B/M were also done in a select few sessions to investigate the laterality of the effects and to help establish the direction of causality between neuronal changes and changes in the behavior. It should be mentioned that almost all of the unilateral injections were done following sessions were NS emergence was observed after bilateral inactivations and before driving down the electrodes. This strategy was taken to increase chances of recording neurons with emergent NS responses. A total of 4 sessions with ipsilateral injections and 3 sessions with contralateral injections are reported here.

3.5.8 Electrophysiology

Electrophysiological recording was conducted as described previously (Ambroggi et al. 2008; Nicola et al. 2004a). Animals were connected to the recording apparatus (Plexon, Inc.), which consisted of a head-stage with operational amplifiers, cable and a commutator to allow free movement within the chamber. The microdrive carrying the electrode arrays was lowered by 160 μm or more at the end of each session to get a new set of neurons every day. Histological analysis was done to verify the location of vmPFC cannulae and electrodes in the NAcS. We excluded neurons with a low

baseline firing rate (<0.03Hz). 131 neurons out of 178 recorded in B/M condition and 98 out of 128 neurons in the CSF condition met these criteria (**FIG 3.14b)**.

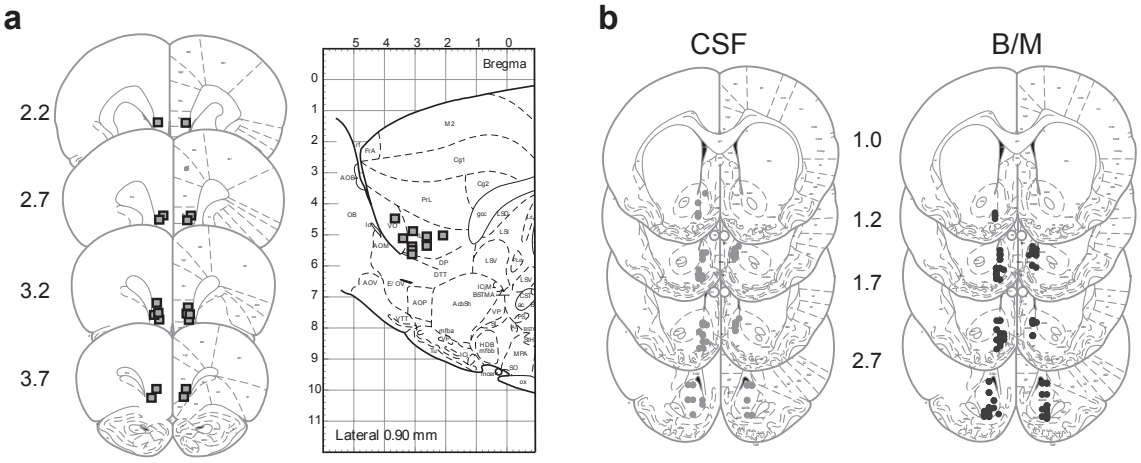

Figure 3.14: **Coronal sections showing the histological sites of injection and recording** (**a**) vmPFC injections and (**b**) NAcS recordings in CSF (left column) and B/M (right column) conditions.

Spike sorting: Isolation of individual units was performed off-line with Offline Sorter (Plexon, Inc., Dallas, StateTX) using principal component analysis. Only units with well defined waveforms were included in this study. Interspike interval distribution, cross-correlograms and autocorrelograms were used to insure that single units were isolated. Additionally, only units in which waveform characteristics were constant over the entire recording session were included in this study.

Determination of the optimal bin size. The size of the bins greatly influences the information contained in peri-event time histograms (PETH). I have developed a principled way finding optimal binsizes for each PETH using Akaike Information Critera (AIC) (**Ch.5**). Because the optimal bin size is usually rather large for most neurons (∼0.1-1sec), and because the AIC commonly shows a fast reduction over

small bin sizes followed by slow changes around the optimal bin size, we also made PETHs using the smallest possible bin size that showed less than a 10% change from the optimal AIC value. This bin size is referred to as the deflection point bin size and was on average 60 msec across our neuronal population.

Response detection: PETHs constructed around the behavioral events, with the optimal bin size, were used to detect excitations and inhibitions. The 10 sec period before cue onset was used as the baseline period. Excitation and inhibition to each event was determined by the presence of at least one bin above (for excitations) or below (for inhibitions) the 99% confidence interval of the baseline during the analysis window for each event. Onset detection was performed using a telescopic approach. We selected the first significant bin using the optimal size and searched, within this time window for the first bin using the deflection point bin size that was beyond 99% confidence interval of the baseline made with the deflection point bin size. If found, we then repeated the search within this deflection point bin, for the first 10ms bin that was beyond 99% confidence of the baseline made with the 10ms bin size. With this method, a high resolution for onset is achieved for each PETH. Response duration was computed by finding the first bin after the detected onset that fell into the 99% confidence interval of the baseline based on the optimal bin size only. For some neurons the response types, onsets and durations were also visually inspected afterwards to ensure accuracy. For a neuron to be included as responsive to an event it should have had an onset latency of no later than 0.5s to that event. With the exception of consumption PETHs there was no restriction on the duration of responses

and therefore no distinction is made between phasic versus sustained responses to events. Consumption PETHs were required to have be at least 1s in duration. This is based on our observation that consumption related responses are often sustained from receptacle entry to receptacle exit (often longer than 3s).

Data transformation and plotting: For average PETHs across neurons, the firing rate of each neuron during each bin was transformed to a Z-score: (F_i – F_{mean}) / F_{sd} where F_i is the firing rate of the i^{th} bin of the PETH, and F_{mean} and F_{sd} are, respectively, the mean and the SD of the firing rate during the 10 sec preceding cue onset.

3.5.9 Histology

Animals were deeply anesthetized with pentobarbital and perfused intracardially with saline and 4% formalin (plus 3% ferrocyanide for rats with stainless steel electrode arrays). Brains were removed, sectioned (50 μm), and stained with Nissl substance in order to locate injection or recording sites (labeled by passing a DC current through each electrode before perfusion).

3.5.10 Statistical analysis

T-tests or ANOVAs followed by post-hoc tests corrected for multiple comparisons by Bonferroni or Bonferroni-Holms methods were used when appropriate. Bootstraps with permutation were used to compare proportions. All results were considered significant at $P<0.05$.

Part II

Second Part

Chapter 4

Isolating Event Related Neuronal Responses by Deconvolution

4.1 Abstract:

Although many studies in neuroscience are based on comparing neuronal responses to single, isolated sensory or motor events, multiple events frequently occur in close temporal proximity in freely moving animals. This often obscures the precise temporal correlation between each event and the relevant brain activity. By simulating neuronal responses in multi-event tasks, we show that peri-event time histograms (PETHs) greatly distort the underlying true responses. We propose a multi-event deconvolution method that can separate the contribution of each event to the overall neuronal activity. The improvements over PETH in analyzing real data are demonstrated using simulated data and a sample electrophysiological recording obtained from rats in a task involving responses to a reward predictive cue.

4.2 Introduction:

An important goal in neuroscience is to understand how brain activity generates sensation and behavior. A major step toward achieving this goal is to determine how changes in neuronal activity patterns correlate with transient stimuli or motor

actions. In many situations, several events occur closely in time, making it difficult to determine to which event an observed change in neuronal activity is correlated. In some situations, this problem can be addressed by experimentally separating the events in time or by using experimental conditions where only one event is present. However, in many other situations, the behavioral response to one event as well as its related activity is contingent on other events and therefore cannot be separated experimentally. (Hollerman et al. 1998; Lau and Glimcher 2008; Lauwereyns et al. 2002; Nicola et al. 2004b; Samejima et al. 2005). For instance, to study how the brain translates sensory information into immediate actions one has to determine (in each relevant brain region) whether the neuronal activity between the stimulus and the action codes for the sensory stimulus or the action that follows it. It is also possible that in some regions the observed activity specifically encodes some unique function of the relation of the two events.

The ambiguities that arise from response overlap are known in the field of electrophysiology and some methods have been put forward to address it (Jin et al. 2009; Lau and Glimcher 2008; Nicola et al. 2004a; Taha et al. 2007). Most of these methods assume that if the average neural activity is more correlated to one event than another, the response to that event should be of shorter duration and larger magnitude compared to the response to the other event. Some have extended this concept by comparing the rise time of the responses to each event to determine the most correlated event (Taha et al. 2007). Another group has used a shuffling of event timings between trials for each event to separate significant changes in the firing pat-

terns that indicate response correlation with one event compared to others (Jin et al. 2009). Though intuitively accurate in many cases these methods can only report which of the events correlates better with the neural response. They cannot rule out the dependency of the response to the event with the lower correlation nor can they determine the precise relation of each event to the neuronal response. Here we develop a method to dissociate these activities and test it on extracellular recording of single neurons obtained in behaving rats. Importantly, this method could in theory apply to any type of time varying brain activity such as voltammetry, event-related potentials or functional-MRI.

4.3 Results

4.3.1 Deconvolution Intuition:

In electrophysiology, the conventional method to estimate a neuron's response to a stimulus is to shift and time-lock its spike train to that stimulus and then make a peri-event time histogram (PETH) by averaging the time-locked activity over many trials. An example PETH for an actual neuron is shown in **Figure 4.1**. The firing of this neuron is time-locked to both a sensory cue and the motor action that follows it (**FIG 4.1a**). The PETHs in **Figure 4.1b** show a broad excitation after the cue and a sharp peak of activity around the motor action. However the firing rasters show that the activity of this neuron is more tightly locked to the motor action compared to the cue, but because the motor actions happen in temporally close yet variable intervals after the cue a broad excitation appears shortly in the cue PETH following cue onset. Therefore based on the PETH results, one could hypothesize that this

neuron is excited to the cue presentation as well as the motor action which can lead to an error in response classification. Our proposed deconvolution method can be used to address this issue by teasing apart the motor activity from the cue activity as the following. If the neuron is only responsive to the action and not the cue, one should be able to account for the cue PETH by using the action PETH in each trial shifted to the time of the action in that trial and then adding the shifted versions of the action PETH over all the trials as shown in **Figure 4.1c**. A simple overlay of the predicted cue PETH constructed by this technique on the raw cue PETH shows a good degree of overlap between the two responses (**FIG 4.1d** left); indicating that the broad excitation appearing on the cue PETH is likely the result of the neuron's response to the action shifted over many trials. Subtracting the predicted cue PETH from the raw cue PETH then successfully eliminates the influence of motor action on cue related firing (**FIG 4.1e** left, baseline is added back for comparison).

On the other hand, since most of the neuron's response is due to the action and not the cue, the same procedure performed on the raw motor PETH (subtracting the predicted PETH made by shifting the cue PETH in each trial from the motor PETH) does not account for the activity of the neuron around the motor action (**FIG 4.1c-e** right, **FIG 4.1e** baseline is added back for comparison**).** Therefore this neuron is more accurately categorized as action-encoding but not cue-encoding after we account for the response to the motor action in the raw cue PETH.

In general if we define an event-related firing (ERF) to be the true firing of a neuron to one event independent of other event timings, then it follows that PETHs

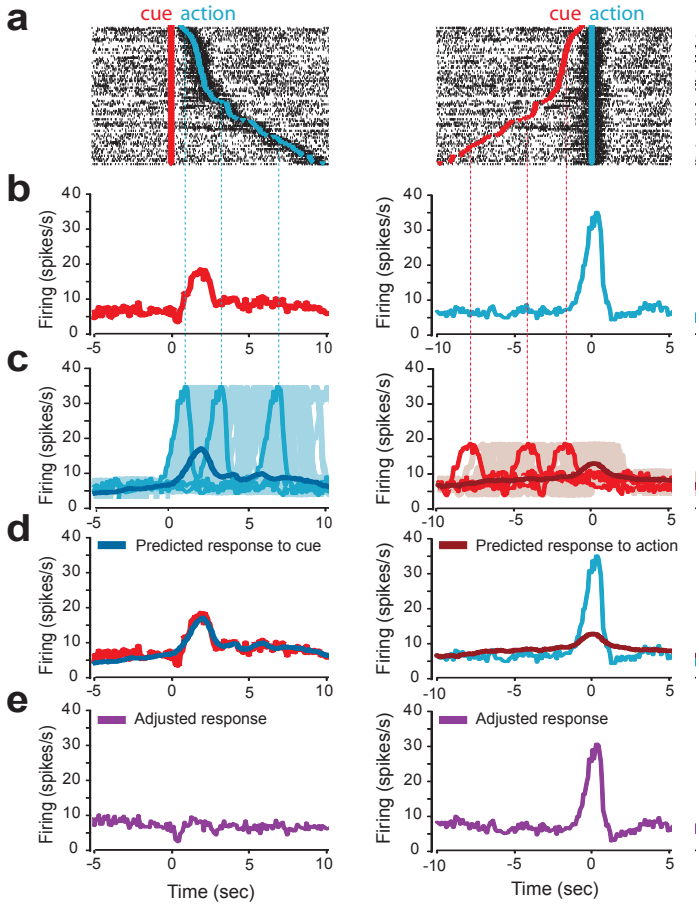

Figure 4.1: **Deconvolution intuition using activity of an example neuron time-locked to a cue and the action it elicits.** (**a**) Firing rasters time-locked to the cue (left) or to the action (right). (**b**) Raw PETHs made for each event showing the increase in firing after both the cue (left) and the action (right).(**c**) Construction of the predicted PETH for each event based on the firing estimate of the other event. Raw PETHs are shifted to the time of occurrence for each event in a trial and a mean predicted PETH is made over all trials. (**d**) The predicted and the raw PETH for each event are overlaid for comparison. (**e**) The subtraction of the predicted PETH from the raw PETH for each event results in an estimate for the actual event-related firing or the ERF after firing overlaps are accounted for. This neuron is correctly categorized as responsive to the motor action but not the cue based on the adjusted responses.

themselves are a superimposition of the ERFs that are time shifted and mixed based on the relative timings of the events. In a linear and time independent regime, PETHs represent a convolution of the ERFs with the event timings. To recover the ERFs, one needs to reverse the convolution or in other words deconvolve the PETHs. This can be done either by trying to solve the inverse problem of convolution in one step

(complete deconvolution) or iteratively where the current ERF estimates in each step are used together with the event timings to improve the ERF estimates in the next step (partial deconvolution). The general formulation for the deconvolution is laid out in the Methods section where the procedure presented above in **Figure 4.1** is shown to constitute only the first iteration in using raw PETHs and timing of all events repeatedly for partial deconvolution of the ERFs. This iterative approach is carried out in the time domain and will be referred to as deconvolution in time or temporal deconvolution (see Methods EQ (4.6)). The complete deconvolution which aims to solve the deconvolution in one step through a direct matrix inversion is carried out in the frequency domain and will be referred to as deconvolution in frequency or frequency deconvolution (see Methods EQ (4.8)). It is important to note that due to the correspondence of time and frequency domain the choice of performing partial deconvolution in time and complete deconvolution in frequency is rather arbitrary and therefore the true comparison in the following sections is between partial and complete deconvolutions not between time and frequency domains per se. The following sections will examine the deconvolution results first for simulated data and then for actual single unit recordings.

4.3.2 Simulation Results:

In this section we first show how response overlaps can distort ERFs obtained for an event by simulating two hypothetical experiments. One of the examples involves multiple events happening closely in time and another involves high frequency repetitions of a single event. We show that in both cases deconvolution in time or in

frequency can recover the ERFs successfully.

Multi-event Deconvolution:

For the first experiment, we assume to have four events A, B, C and D. For simplicity we assume that each event happens only once during each trial and always in the order ABCD with a total of 40 trials. The inter-event intervals between events in a given trial are drawn randomly from Poisson distributions. We assume to have a neuron with ERFs that show excitation to event A, excitation then inhibition to event B, inhibition to event C and inhibition then excitation to event D (**FIG 4.2c**, dotted lines). The firing of this neuron in two different trials in the presence of all 4 events is shown in **Figure 4.2a**. The PETHs are obtained by convolving the ERFs with the event timings (EQ (4.1) with Gaussian noise, sd=1). **Figure 4.2c** shows the event-locked PETHs overlaid on their corresponding ERFs. This figure clearly demonstrates that the raw PETHs built by simply averaging the firing over trials can deviate considerably from the true ERFs and as such provide poor estimates of the true ERFs.

It is important to note that in this example and in other multi-event experiments where the number of each event per trial does not change between trials, the DC contents of the ERFs are not deconvolvable. This can be verified in the current example by looking at the determinant of the convolution matrix and noticing that it is zero (and therefore not invertible) at or near DC frequencies (**FIG 4.2b**). This is because the DC contents of the ERFs are not affected by the timing jitters that are used to deconvolve the higher frequency contents and as such different trials here

do not provide additional information on the value of the ERFs DC content. This leaves us with an ill posed problem of one equation for the total DC content and four unknowns corresponding to each ERF. Nevertheless after correcting for DC shifts in the results the application of deconvolution method successfully recovers the ERFs with high fidelity in time or in frequency as shown in **Figure 4.2d** and in **Figure 4.2e** respectively.

The ERF estimation error can be measured by the average mean squared error (hereafter MSE) for all 4 events integrated over time and as a function of deconvolution iterations (**FIG 4.2f**). The ERF estimation error shows more than 70% improvement just after the first deconvolution iteration in time compared to the raw PETHs (MSE=0.7 versus 2.6). The ERF estimate continues to converge to the true ERFs and the error goes down by another order of magnitude after 1000 iterations comparable to average PETH firing noise (MSE=0.03 versus 0.025 for noise variance in 40 trials). Our ability to fit the raw firing PETHs in the presence of all 4 events based on the deconvolved ERFs also greatly improves over deconvolution iterations with errors as low as 10^{-7}after 1000 deconvolution iterations (**FIG 4.2g**). A final deconvolved signal (equivalent to infinite iterations in time domain) can be obtained from frequency deconvolution with errors that are not distinguishable from the average firing noise (MSE=0.027 versus 0.025 for noise variance in 40 trials) (**FIG 4.2f**).

Fast Occurring Single Event Deconvolution:

For the second experiment, let us assume an event A happening 10 times in each trial with randomly spaced inter-event intervals from a Poisson distribution. We

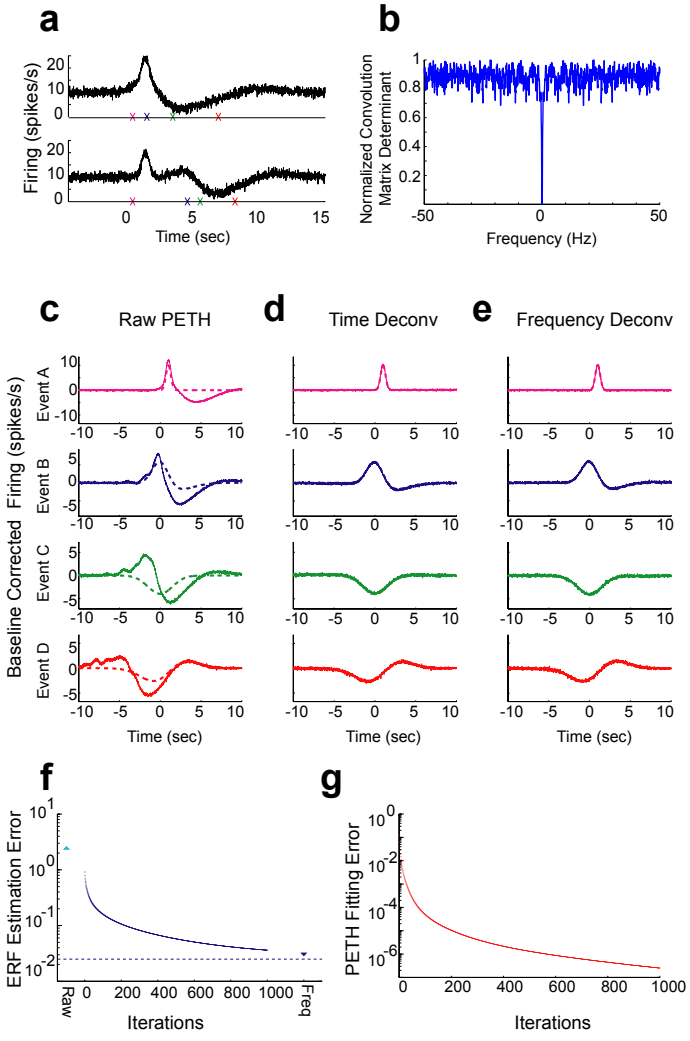

Figure 4.2: **Deconvolution of event-related firings in a simulated four event task.**(**a**) Firing rate of a hypothetical neuron in two randomly chosen trials with different inter-event intervals. Each event in a trial is shown with a cross with four different colors (pink, blue, green and red indicating four different events A, B, C and D respectively in the task). The following distributions were used to draw the event timings for each trial:
$t_A = 0 \quad t_B = 1 - \log(1-\tau) \quad t_C = t_B + 0.5 - \log(1-\tau) \quad t_D = t_C + 2 - 1.5 \times \log(1-\tau) \quad \tau \sim U(0,1)$
(**b**) Determinant of the convolution matrix normalized by its maximum value over a range of frequencies from 0 to 50 Hz. (**c**) ERFs for each event (dashed lines) and their corresponding raw PETHs (solid line, derived from convolving the respective ERFs) after averaging 40 trials. (**d, e**) The deconvolved ERFs after 1000 iterations in the time domain (D, solid line) and the frequency domain (E, solid line) overlaid on the actual ERFs. Deconvolution in both time and frequency shows considerable improvement over the raw PETH estimates and is successful in recovering the true ERFs. (**f**) Log of the mean squared error between the estimated ERFs and the true ERFs before deconvolution (cyan triangle), over iterations in time domain and after frequency deconvolution (blue triangle) corresponding to infinite iterations in time. Raw PETH noise variance is indicated by the dashed line. (**g**) Log of the mean squared fitting error between the raw PETH and the PETH constructed from the temporally deconvolved ERFs as a function of deconvolution iterations.

assume to have a neuron with an ERF that shows excitation then inhibition to this event and inter-event intervals that are shorter than the ERF duration. The firing of this neuron in two different trials in the presence of 10 event repetitions is shown in **FIG 4.3a**. The raw PETH is again obtained by convolving the ERF with the event timings (EQ (4.1) with Gaussian noise, sd=1). The close temporal spacing between events in this case results in a broad distortion of ERF estimate based on the raw PETH. **Figure 4.3c** demonstrates a gross overestimation of both amplitude and duration of the ERF. Applying the deconvolution method in time or in frequency successfully recovers the actual neuronal response to a single occurrence of event A (**FIG 4.3d** and **FIG 4.3e**). Although small jitters in the event timings tend to reduce the determinant of the convolution matrix in low frequencies, in this case the convolution matrix is invertible over all frequencies making the deconvolution in the frequency domain feasible for all frequencies (**FIG 4.3b**).

ERF estimation error measured by MSE again shows considerable improvement close to an order of magnitude only after the first deconvolution iteration compared to the raw PETH (MSE=1.5 versus 9.8). This error continues to go down after some 500 iterations to levels undistinguishable from average firing noise. The MSE resulting from the frequency domain deconvolution (equivalent to infinite iterations in time domain) is also comparable to the average PETH firing noise (MSE=0.0022 for both time and frequency deconvolution versus 0.0025 for noise variance in 40 trials containing 10 event repetitions each) (see **FIG 4.3f).** Also the raw PETH fitting accuracy using the deconvolved ERF improves rapidly over deconvolution iterations

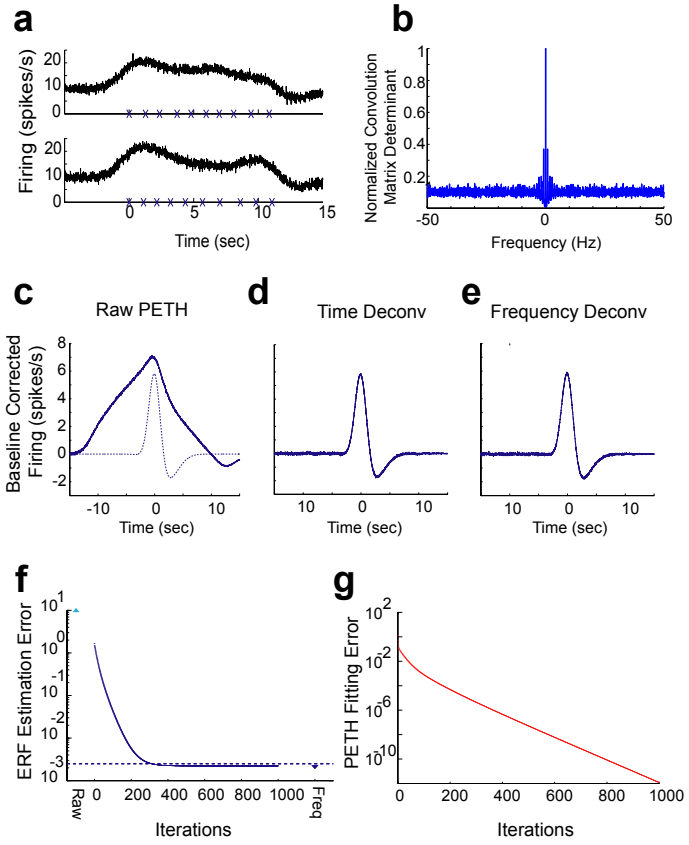

Figure 4.3: **Deconvolution of ERFs in a simulated repetitive event task.**(**a**) Firing rate of a hypothetical neuron in two randomly chosen trials with different inter-event intervals. Timing of each event occurrence in a trial is shown with a blue cross. The following distribution was used to draw the timings of the 10 repetitions of event A in each trial:
$t_A^1 = 0 \quad t_A^i = t_A^{i-1} + 1 - \frac{\log(1-\tau)}{4}$ for $1 < i \leqslant 10 \quad \tau \sim U(0,1)$ where t_A^1 indicates time of the first occurrence of event A and so on.
(**b**) Determinant of the convolution matrix normalized by its maximum value over a range of frequencies from 0 to 50 Hz. (**c**) ERF (dashed line) and the corresponding raw PETH (solid line, derived from convolving the ERF) after averaging 40 trials. (**d, e**) The deconvolved ERFs after 1000 iterations in the time domain (D, solid line) and the frequency domain (E, solid line) overlaid on the actual ERF. Deconvolution in both time and frequency shows considerable improvement over the raw PETH estimate and is successful in recovering the true ERF. (**f**) Log of the mean squared error between the estimated ERF and the true ERF before deconvolution (cyan triangle), over iterations in time domain and after frequency deconvolution (blue triangle) corresponding to infinite iterations in time. PETH noise variance is indicated by the dashed line. (**g**) Log of the mean squared fitting error between the raw PETH and the PETH constructed from the temporally deconvolved ERF as a function of deconvolution iterations.

where after 1000 iterations the MSE for predicting PETHs goes down to a negligible 10^{-12} (**FIG 4.3g**).

4.3.3 Cross-validation of Deconvolution:

The simulation results in the previous section demonstrate the feasibility of complete deconvolution of the ERFs when the data is generated based on the deconvolution model assumptions, with the ERF estimation error bounded only by a linear combination of the raw PETH noise variances (see Methods EQ (4.7) and EQ (4.12)).

In dealing with the actual firing data the assumptions of the model can be violated to varying degrees. For instance, the ERFs might not be stationary over time or might show dependence on the inter-event intervals (a.k.a. ERF interactions). Furthermore in dealing with the actual data we might be limited by the number of available trials to average out the trial noise which would raise the lower bound of deconvolution error. These factors can adversely affect our ability to fully recover the true ERFs. If the ERFs were known one could simply measure the MSE of the deconvolved ERFs in predicting the true ERFs and obtain a quantitative measure of how the assumption violations or trial noise are affecting the deconvolution performance over the iterations. Such a quantitative performance measure is in fact shown for the simulation results in **Figures 4.2f** and **4.3f** where the ERF estimation error was calculated and plotted over deconvolution iterations. In dealing with real data, obviously the true ERFs are not known and therefore another approach is needed to monitor the model performance over the deconvolution iterations.

One way of characterizing the effectiveness of the deconvolution for real firing data is by cross-validation of deconvolved ERFs. Cross-validation is a simple but powerful statistical technique that can be used to choose between models (Stone 1974) which

in our case would be between performing or not performing the deconvolution and also between different levels of deconvolution for a given neuron. Cross-validation can be done by randomly dividing the trials into two groups, one used for finding the ERFs using a certain deconvolution iteration and the other one for validating the estimated ERFs to see if they can convolve back to predict the PETHs made from the test trials (see Methods EQ (4.13)).

We first tested the sensitivity of the deconvolution to noise by decreasing the signal to noise ratio in the 4 event deconvolution experiment shown in **Figure 4.2**. The signals to noise values tested were 22dB, 12 dB, 6 dB and 4 dB. Increasing the noise level resulted in increasing levels of cross-validation error as expected. Although in all cases some levels of partial deconvolution outperformed the raw PETH in cross-validation, the deconvolution iteration resulting in the lowest cross-validation error decreased in conjunction with the signal-to-noise ratio. (see **FIG 4.4a** left). Since in the case of simulation we a priori know the true ERFs, MSE for ERF prediction error at each noise level can be calculated as well. One can see that the trend of changes for the ERF prediction error is qualitatively the same as that of cross-validation (corrcoef=0.85, $P<0.001$) though the iteration with minimum error for each noise level is not necessarily the same for the two error measures (**FIG 4.4a** right). This confirms that in the absence of the ERF prediction error, cross-validation error can provide a comparable criterion for finding the appropriate level of deconvolution in a given neuron (for a derivation of the relationship between ERF and cross-validation MSEs look at EQ (4.14) and EQ (4.15) in the Methods).

It is important to note that the changes in the cross-validation and ERF MSE reflect changes in both the estimates bias and variance. In cases of low signal two noise ratio the amount of noise that is added after a few iterations integrated over time more than cancels out the improvement in the ERF estimate bias which results in the rising of the cross-validation MSE after only a few trials as can be seen in **Figure 4.4a** both right and left column (also see Methods EQ (4.7) for added noise terms in each trial).

We also tested the sensitivity of the deconvolution to increasing levels of ERF interactions in one of the many possible ways that such interactions may occur. In this case we assumed that the amplitude of the ERFs would decrease and their duration would increase with the interval between the events. The performance of the deconvolution model was compared when there were no interactions versus when there was interactive modulation of the first ERF, the first and second ERFs, up to the point that all four ERFs were modulated by the spacing between their events and the next event except for the last event's ERF that was modulated by the interval between itself and the previous event. Again Increasing levels of ERF interactions resulted in increased cross-validation error as expected. And although in all cases deconvolution outperformed the raw PETH in cross-validation, the deconvolution iteration resulting in the lowest cross-validation error decreased by increasing ERF interactions (**FIG 4.4b** left). Here too the ERF prediction error showed a similar trend of changes as a function of ERF interaction and deconvolution level to the cross-validation (corrcoef=0.6, $P<0.001$) but the iteration with minimum ERF prediction error was not

necessarily the same as that of the cross-validation error (**FIG 4.4b** right).

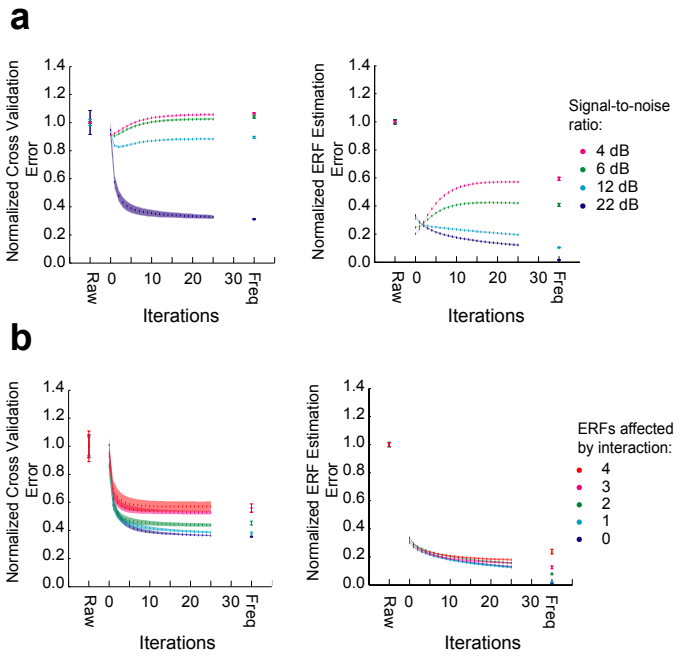

Figure 4.4: **Changes in cross-validation error depend on the level of noise and the degree of ERF interactions.** The trend of changes in the cross-validation error matches that of the ERF prediction error. Cross-validation is done by constructing a PETH prediction based on the deconvolved ERFs made from 90% of trials (36 training trials) to be compared with the raw PETH in the remaining 10% of trials (4 validation trials). (**a**) Mean squared error for cross-validation between the predicted PETHs and the validation PETHs normalized by the raw PETH cross-validation error (left) and mean squared error between the predicted ERFs and the true ERFs (right) normalized by the raw PETH ERF estimation error over a range of signal–to-noise ratios (4dB, 6dB, 12dB, 22dB) (**b**) Mean squared error for cross-validation between the predicted PETHs and the validation PETHs normalized by the raw PETH cross-validation error (left) and mean squared error between the predicted ERFs and the average of true ERFs (right) normalized by the raw PETH ERF estimation error over increasing levels of ERF interactions (interaction for 0, 1, 2, 3 and 4 ERFs).

It is also worth pointing out that in both cases of low signal-to-noise ratio or high degree of ERF interactions, the estimate provided by frequency deconvolution (equivalent to complete deconvolution) can be inferior to a partial deconvolution (obtained after some deconvolution iterations). A fact that is demonstrated both by the ERF estimation error and the cross-validation error for varying levels of noise or ERF interactions. This can be partially due to the singularity of the convolution matrix

(**FIG 4.3b**) in this experiment for the lowest frequencies as well as the discussed effect of noise in low signal to noise ratio signals in raising the estimate errors

In using the cross-validation as a tool for selecting the appropriate deconvolution iteration we define the "optimal iteration" to be the iteration resulting in the lowest significant cross-validation error compared to the raw PETH error and the "maximal iteration" to be the highest iteration having a cross-validation error significantly lower than that of the raw PETH (paired t-test $P<0.01$ for significance). Optimal and maximal iterations therefore define the most and least conservative criteria respectively for the degree of deconvolution that can be safely applied to the data.

Finally it is worth mentioning that we could have chosen an alternative way of characterizing the effectiveness of the deconvolution by dividing the data randomly into two equal parts and using the complete deconvolution from one half to estimate the ERF MSE directly based on the plug in principle. This direct method which normally yields similar results to cross-validation (see Methods EQ (4.16)) requires double the amount of deconvolution compared to cross-validation and therefore is more computationally intensive. Furthermore in many cases complete deconvolution of ERFs might not be available or might show large fitting errors (see Experimental Results **FIG 4.6b** and **FIG 4.7a**). Therefore we have used cross-validation as the faster and more robust approach in this paper thought when applicable all of our results with regard to cross-validation should be replicable using this direct method as well.

4.3.4 Experimental Results:

In this section we demonstrate the use of the deconvolution method to parse out actual neuronal responses and will use the cross-validation technique introduced in the previous section to calibrate the number of deconvolution iterations applied on each neuron given its firing pattern and the degree of noise.

We used a sub-set of equal to 106 neurons that were recorded in the nucleus accumbens (NAc) in a discriminative stimulus task (Ambroggi et al. 2008). Briefly, two auditory cues, discriminative stimuli (DS) and neutral stimuli (NS) were randomly presented (with durations of up to 10 s for the DS and 10 s for the NS) on a variable interval schedule with an average interval of 30 s. A lever press was required during DS presentation to cause the delivery of a reward into an adjacent reward receptacle and responding during NS was never rewarded. For simplicity, only the DS is displayed in **Figure 4.5** with intervals between the two DS's shown on average to be 60s.

Firing of neurons on trials when the animal successfully responded to the cue can be modulated by four events which happen shortly one after the other. These events include presentation of the DS, lever press (LP), receptacle entry (Ent) and receptacle exit (Ext). Many of the neurons recorded in this task show changes that are time-locked to more than one event. In fact complex responses to multiple events are quite common in the rat ventral striatum (Humphries and Prescott 2010; Nicola et al. 2004a; Nicola 2007; Schultz et al. 1992). **Figure 4.6a** shows the firing raster of a putative medium spiny neuron recorded in rat NAc during the rewarded tri-

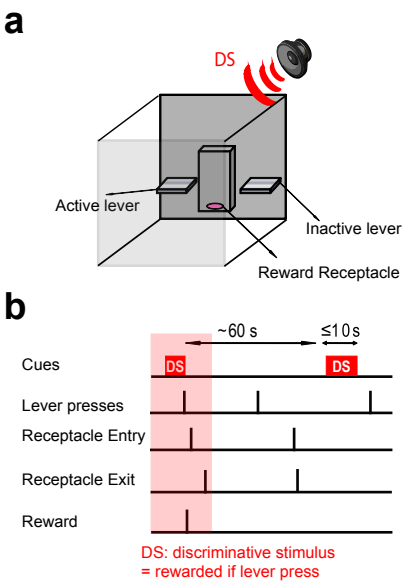

Figure 4.5: **Experimental procedure** (**a**) Rats were placed in 40 x 40cm chambers containing two levers (active and inactive) on each side of a reward receptacle. A speaker played the auditory cues above the chamber. (**b**) Behavioral paradigm showing hypothetical sequence of events. Only active lever presses during the cue presentation resulted in reward delivery in the receptacle after which the rat would normally enter the receptacle to consume the reward and exit the receptacle after the consumption. The event sequence consisting the cue, active lever press, receptacle entry and exit in the presence of reward is counted as one rewarded trial (shaded area).

als. This neuron is excited after the DS, shows inhibition during the consummatory period between receptacle entry and exit and then has a brief excitation following receptacle exit. Although the PETH fitting error decreases rapidly over deconvolution iterations, cross-validation for this neuron reaches a minimum after about 3 deconvolution iterations (optimal deconvolution) but then increases back with increasing degree of deconvolution with the maximal deconvolution happening at iteration 22 (**FIG 4.6b** and **FIG 4.6c**). The result of partial deconvolution with optimal and maximal iterations is shown together with the raw PETHs in **Figure 4.6a.** One can see that the excitation to DS based on the deconvolved ERF has a shorter duration compared to what was previously estimated by the raw PETH. This can be verified by looking at the firing raster time locked to the DS where part of the excitation can

be seen to be due to lever press preparation. Deconvolution also removes the broad consumption-related inhibition from the DS PETH as well as the lever press and cue related excitations from the Ent and Ext PETHs. The presence of the excitation following the receptacle exit in the Ext ERF shows that in fact this activity is correlated with exiting the receptacle and not a contamination by other events.

The fact that the partial deconvolution outperforms the raw PETHs in cross-validation shows that the neuron's activity obeys at least partially the deconvolution model (see Methods EQ (4.1)). But possibly due to interactions, nonlinearities and moderate signal-to-noise ratio (∼10dB, **FIG 4.4a**), complete deconvolution of responses is not feasible. Another fact that further complicates the complete deconvolution of this neuron's firing is the fact that the determinant of the convolution matrix made from event timings shows a rather wide notch for low frequencies which will make complete deconvolution at those frequencies erroneous. This results in high PETH fitting errors for frequency deconvolution in **Figure 4.6b**. (see **FIG 4.6d** for the convolution determinant)

Among the 106 neurons recorded, 85 were responsive to at least one event in rewarded trials. As shown in **Figure 4.7a** the average PETH fitting error for responsive neurons decreases rapidly over deconvolution iterations while on average the frequency deconvolution shows large errors in deconvolving the ERFs. Once more the poor performance of the frequency deconvolution can partially be explained by average determinant of the convolution matrix which shows a wide notch over low frequencies (see **FIG 4.7c** for the average convolution determinant). As discussed

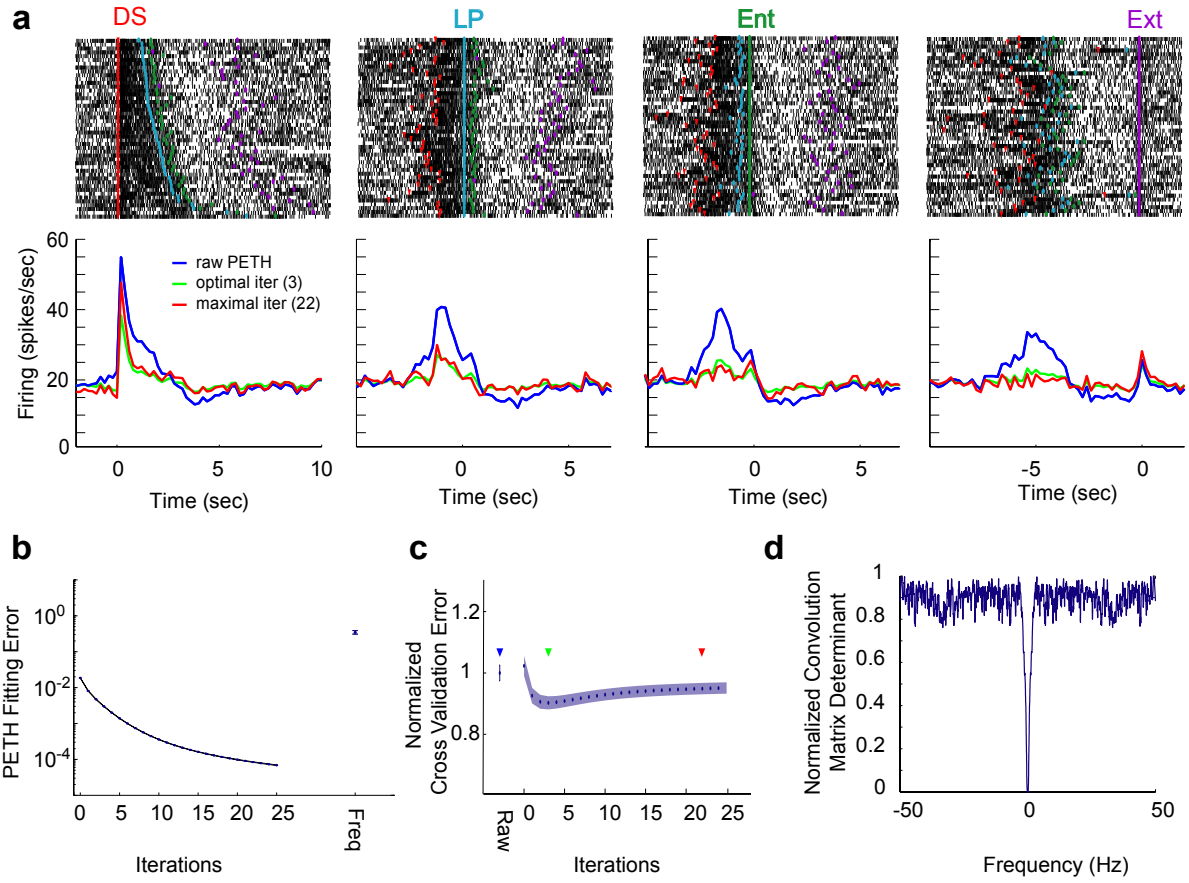

Figure 4.6: **Optimal and maximal deconvolution in an example neuron responding to all the task events.**(**a**) The firing rasters from 46 trials (top) and the corresponding raw PETHs (bottom, blue lines) time-locked to DS, LP, Ent and Ext. The deconvolved ERFs after the optimal (green line) and maximal (red line) iterations are displayed. A bin size of 200 ms was used for all the histograms. (**b**) Log of the mean squared fitting error gradually decreases as a function of deconvolution iterations. (**c**) Normalized cross-validation mean squared error reduces after partial deconvolution with three iterations having the lowest error (optimal) and the twenty second iteration being the highest iteration with significantly lower error than raw PETH cross-validation error (maximal). (**d**) Determinant of the convolution matrix normalized by its maximum value over a range of frequencies from 0 to 50 Hz.

previously this is a direct consequence of multi-event experimental paradigms with unchanging number of occurrences for each event per trial in which case a partial temporal deconvolution usually outperforms frequency deconvolution.

The optimal deconvolution for each of the responsive neurons can be determined by finding the minimum of the cross-validation error term. The population average

of the cross-validation error for the responsive neurons indicates that on average partially deconvolved ERFs outperform raw PETHs in predicting the validation trial firing patterns with the average cross-validation error being significantly lower than the average raw PETH errors for up to 6 iterations (multiple comparison $P<0.01$ Bonferroni) and a maximum improvement of about 8% in the average cross-validation MSE (**FIG 4.7b**).

For 25% of the neurons, the minimum of the cross-validation error happens for one or more iterations. For the other neurons, this minimum happens at the zero level deconvolution. The deconvolved ERFs at zero iteration are simply the raw PETHs from the training trials normalized by a factor which in this case is the number of event types (see Methods EQ (4.5)). The fact that in all the responsive neurons partial deconvolution outperforms raw PETHs in predicting the validation trial firings means that the raw PETHs time-locked to each event are influenced by the timing of the surrounding events. The optimal deconvolution found for each neuron is correlated with the signal-to-noise ratio (corrcoef=0.71 $p<0.01$) with higher degrees of deconvolution corresponding to higher signal–to-noise ratios (**FIG 4.7d**). The maximal deconvolution iteration for each neuron shows a wider distribution over iterations with an average of about 6 iterations (mode equal to 4 iterations). The maximal iteration also shows dependence on the level of signal to noise with higher iterations corresponding to higher levels of signal to noise ratio. (corrcoef=0.40 $p<0.01$) (**FIG 4.7e**).

The neuronal responses shown in **Figure 4.8a** demonstrate that in fact in com-

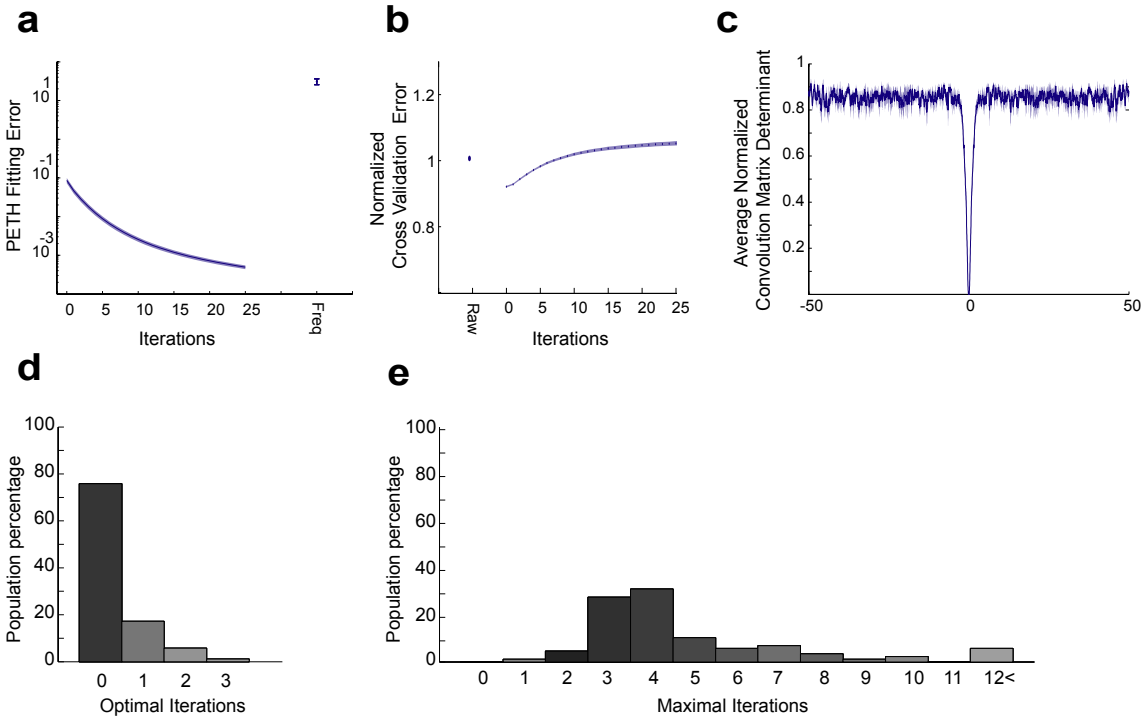

Figure 4.7: **Deconvolution error and distribution of optimal and maximal iterations for all responsive neurons.**(**a**) Log of the average (for all neurons) mean squared fitting error as a function of deconvolution iterations. Complete deconvolution in frequency deconvolution does not provide a good estimate on average. (**b**) Average normalized cross-validation mean squared error as a function of deconvolution iterations. (**c**) Average determinant of the convolution matrix normalized by its maximum value for all the recorded sessions over a range of frequencies from 0 to 50 Hz. Shaded area shows the standard error of the mean. (**d, e**) Distribution of the optimal (**d**) and maximal (**e**) deconvolution iterations among all the responsive neurons. Color coding of the bars correspond to the average signal to noise ratio for all neurons in each optimal deconvolution (Black: low signal to noise, White: high signal to noise) The signal to noise is calculated here and for previous simulations as:$10\log(\frac{var(PETH)}{var(baseline\ PETH)})$
instead of the more usual formulation: $10\log(\frac{var(PETH)-var(baseline\ PETH)}{var(baseline\ PETH)})$
since the argument of log in this formulation can become negative and the signal to noise undefined when dealing with the actual data.

parison with the raw PETHs the degree of contamination by non time locked events is reduced for all of the deconvolved ERFs after partial deconvolution (results of maximal iteration shown). For instance the duration of excitations following the DS is reduced significantly based on the deconvolved DS ERF compared to the raw DS PETH (average 1.51 sec vs. 2.18 sec, paired t-test $P<0.05$), pointing to the fact that the post-DS excitations in these cells were probably contaminated by excitations pre-

ceding the lever presses (**FIG 4.8b**). Interestingly, some responses were not affected by deconvolution. Indeed, receptacle exit ERFs maintained post-exit excitations after deconvolution ($P>0.05$ paired t-test for average response from -1 to 1 sec around receptacle exit) while the cue and lever press contaminations on the exit ERFs are attenuated after deconvolution ($P<0.01$ paired t-test for average response from -7 to -5 sec around receptacle exit corresponding to average time of cue presentation and lever press respectively). Also inhibitions during the consumption period are shown to be mainly correlated with the receptacle entry and not with the receptacle exit since the sustained inhibition prior to the exit is removed after deconvolution from Ext ERFs. On the other hand Ent ERFs still show inhibition during the consumption even after the deconvolution (**FIG 4.8b**). However it is important to be cautious when interpreting the appearance or disappearance of broad sustained changes compared to the faster phasic changes in the ERFs following the deconvolution due to the problematic nature of low frequency deconvolution in this experiment (**FIG 4.7c**).

Besides changes in response dynamics there is a considerable change in the percentage of responsive neurons to each event after deconvolution. The total number of neurons responsive to each event was reduced by 36% on average (averaged over all 4 events) (see **FIG 4.8a** right versus left columns). This reduction happens while the total number of neurons that are responsive to at least one event stays roughly the same before and after deconvolution (85 neurons before versus 79 neurons after). This indicates that the deconvolution method is effective in reducing the neurons that are falsely categorized as responsive to multiple events (**FIG 4.8c**).

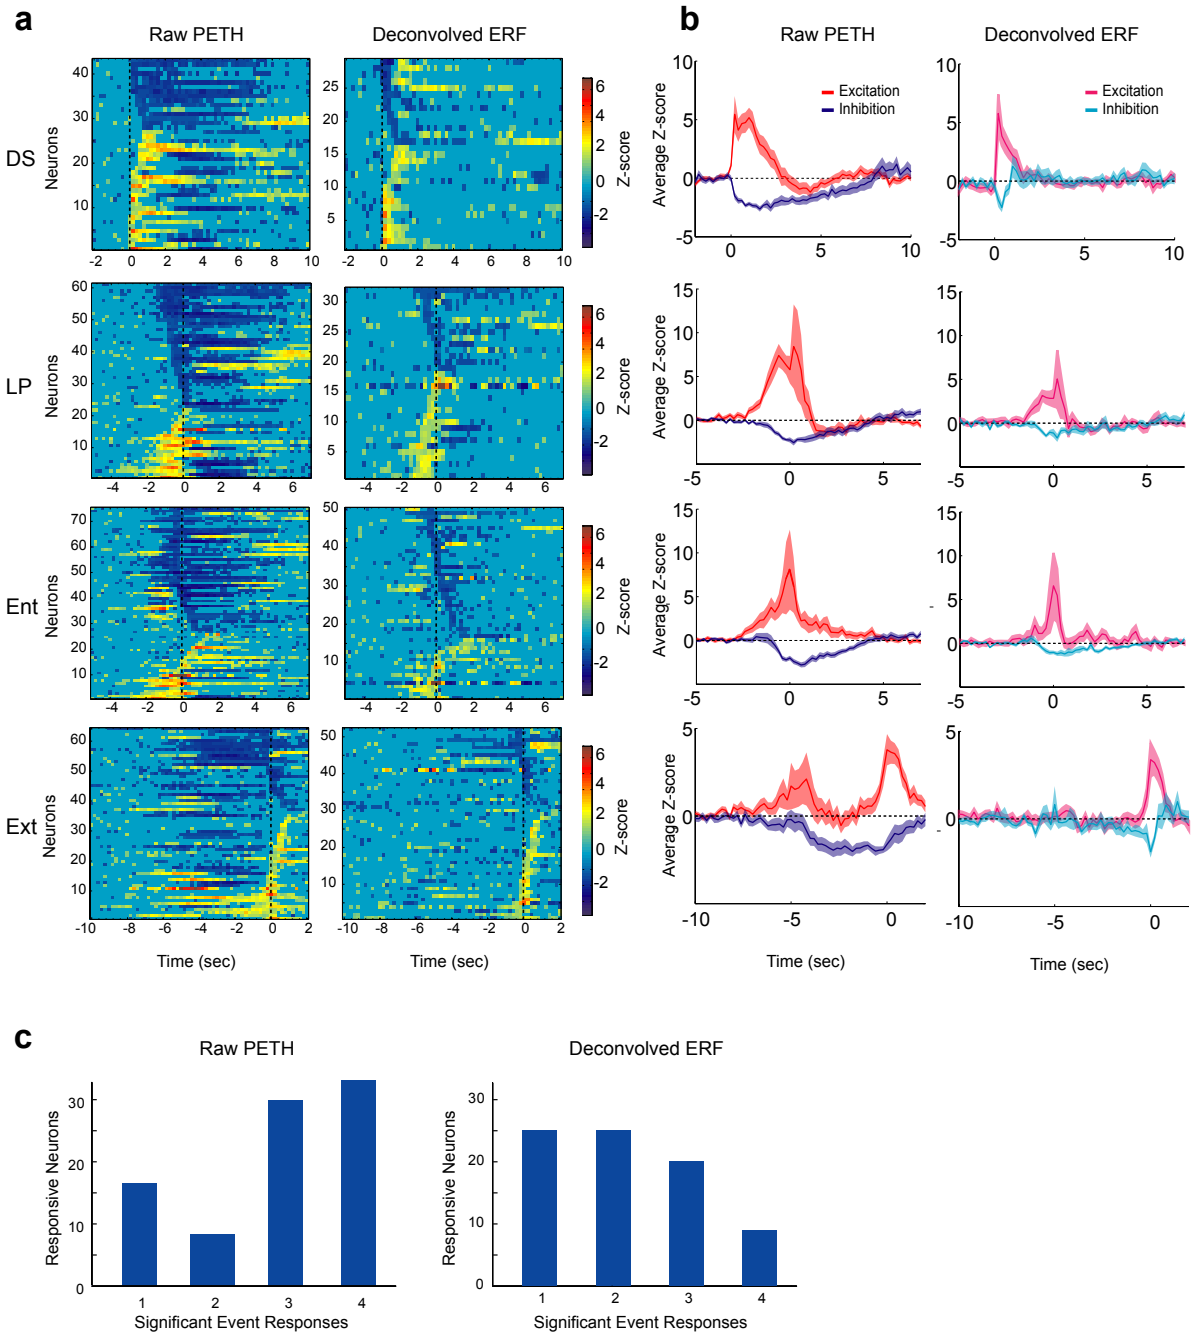

Figure 4.8: **Comparison of population profile of responses to each event based on raw PETHs and deconvolved ERFs.** (**a**) Color coded raw PETH Z-scores (left column) and deconvolved ERF Z-scores (right column) for all the responsive neurons to each event. The four rows from top to bottom correspond to DS, LP, Ent and Ext responsive neurons. Each line in the color coded plots corresponds to the firing of one neuron in Z-score. Only Z-scores outside the 90% confidence interval of baseline (pre-cue firing) are color coded. Values within the confidence interval are all colored at zero.(**b**) Average raw PETH Z-scores versus deconvolved ERF Z-scores. Red and blue correspond to excitatory and inhibitory responses based on raw PETHs respectively. Magenta and cyan correspond to excitatory and inhibitory responses based on deconvolved ERFs respectively.(**c**) Number of neurons being categorized as responding to one, two, three or four events before deconvolution based on raw PETHs (left) and after partial maximal deconvolution based on deconvolved ERFs (right). The comparison shows a reduction in the number of neurons categorized as responding to multiple events.

4.4 Discussion:

The ambiguity that typically occurs in the analysis of neuronal coding of temporally close events can be viewed as a problem with the PETHs being the convolved readout of several underlying ERFs that are time-shifted with respect to each other over trials. We first demonstrated that the raw PETHs in some cases provide a poor estimate of the underlying ERFs due to their overlap. We then presented a deconvolution formulation that is suited to the specific problem of finding ERFs in a multi-event task based on raw PETHs. We derived a pair of theoretically equivalent formulations for deconvolution in the time and in the frequency domains and showed through simulation that the formulation proposed can successfully parse out individual ERFs in multi-event and in fast occurring single event tasks. Since convolution in time is equivalent to multiplication in frequency (Oppenheim et al. 1997), deconvolution in the frequency domain theoretically simplifies the inversion of the convolution matrix and so is used for complete deconvolution. On the other hand, the temporal approach involves a step-by-step deconvolution to approximate this inverse. Although the two methods are theoretically equivalent, the iterative approach in the temporal deconvolution is more robust when applied to actual data. This is due to the fact that the deconvolution model entails validity of certain assumptions such as linearity and stationarity of responses and non-singularity of the convolution matrix which are not necessarily true when dealing with actual event timings and neuronal responses. In dealing with neuronal recordings, the temporal approach has the advantage of permitting a partial deconvolution given the level of the noise and

the degree to which the data deviates from model assumptions. In this sense, the iterative temporal approach can be more flexible than the complete deconvolution approach in the frequency domain.

Noise can also play a limiting role on the degree of deconvolution that can be successfully applied to the actual PETHs (**FIG 4.4**&**4.7**). We therefore recommend using a low pass filter on the data by applying Gaussian smoothing or Wiener filters (Papoulis 1977) on the data or have bin sizes that are large enough to limit the noise power before applying the deconvolution. Noise levels can also be lowered proportional to the square root of the number of trials. As shown in the result section, neurons with higher signal–to-noise ratios tend to have higher optimal and maximal deconvolution iterations.

It is also important to note that the experimental design and the amount of jitter in the event timings can influence the effectiveness of the deconvolution procedure, especially in deconvolving the broad tonic changes in the ERFs (which correspond to low frequency content of the signals). In fact there are two channels of information that are used for deconvolution, one is the changes in the number of occurrences of each event per trial and the other is the changes in the inter-event intervals between different events over trials. Of these two factors the former is more important when deconvolving the tonic changes as they are less affected by the event timing jitters due to their slow temporal dynamics. For example in the discussed experimental paradigm where each event occurred only once per trial the deconvolution of the lowest frequencies was not well defined. In such cases caution is recommended when drawing

conclusions about changes in tonic components of the ERFs. The best case scenario for successful deconvolution ERFs therefore includes variable number of occurrences of events per trial with extensive event timing jitter between trials.

Cross-validation can be used as an effective criterion to choose between using raw PETHs without deconvolution or a partial deconvolution to obtain the best possible estimates for the ERFs. Depending on the error tolerance a range of partial deconvolutions from optimal to maximal can be chosen for each neuron. The result of the partial deconvolution should provide ERF estimates that are closer to the true ERFs. This in turn can inform the response detection algorithms to provide a better estimate of the responsive neurons to each event.

Deconvolution on real data obtained from extracellular recordings provides the proof-of-concept that this method can modify the conclusions obtained from such data. First it allows better categorization of the response types. Indeed, as shown in **Figure 4.1** and **4.6**, based on raw PETHs, neurons can be erroneously categorized as responsive to a given event if they respond strongly to another close event. The fact that cross-validation finds the iteration zero on average to fit the validation data better than the raw PETH (**FIG 4.7b**) supports the idea that many of the raw PETHs are dependent on the responses to the surrounding events. Categorizing the neurons based on their deconvolved ERFs addresses this issue. As expected using the deconvolved ERFs instead of the raw PETHs reduces the number of events the neurons are responsive to (**FIG 4.8c**). Furthermore, the deconvolution allows a better estimation of the response dynamics, both in duration and amplitude. Notably,

if a neuron is responsive to two temporally close events, the responses add up. For example we find that the duration of the neuronal responses to reward predictive cues (DS) is reduced after deconvolution. This results from correcting for the preparatory responses to the lever press following the cue onset. Therefore deconvolution provides a better estimate of duration in the presence of temporally proximal events.

The ambiguities that arise from response overlap are not only relevant in electrophysiology but in neuroscience as a whole. In fact there are studies in functional MRI and EEG literature that have identified this issue and provided methods for deconvolution of responses (Bohrquez and Ozdamar 2006; Dale 1999; Glover 1999; Hinrichs et al. 2000) However all of these methods only allow for a complete deconvolution (like the one presented here in frequency domain) and as such are prone to large errors in the presence of interactions and convolution matrix singularities as discussed and demonstrated in this paper. Furthermore some of these studies are only appropriate for a single event deconvolution and are not general enough to allow for multiple event types (Bohrquez and Ozdamar 2006; Glover 1999). To our knowledge the method presented here is the most general formulation of deconvolution for trial based experimental paradigms. It is also the first to allow for a partial deconvolution with a criterion based on cross-validation to address the success of varying degrees of deconvolution in dealing with the real data.

Finally our formulation of deconvolution assumes the timings of all the events in a task to be available. It may be the case in some experiments that the existence and/or the timings of some events are not known or recorded. These unknown event

occurrences especially when correlated in time with the recorded events can introduce significant errors in the ERF estimations. In these cases other techniques such as blind deconvolution or source separation (Comon 1994; Haykin 1994) can be used. These techniques not only estimate the original signals (in our case the ERFs) but also try to approximate the mixing or convolution matrix. Methods such as independent component analysis (ICA) that assume statistical independence between the original signals are successfully applied in many source separation problems. The number of unknown events that needs to be estimated in this way might in turn be decided by a model selection technique such as cross-validation. Such an extensive model which includes known events and optimizes for unknown events and their timings however would be heavily parameterized and computationally intensive. The usefulness of such models in deconvolving the ERFs remains to be tested in the future.

In conclusion, we have described and tested a novel technique that allows separation of the components of the neuronal activity that are associated with temporally close but distinct behavioral events. We have demonstrated the value of the deconvolution method using simulated data showing that it can significantly improve estimates of the event related neuronal responses compared to the raw PETHs. We then applied our method to single unit recordings of nucleus accumbens neurons obtained in behaving rats showing that the deconvolution of the neuronal responses disambiguates the classification of responsive neurons to each event and adjusts each event-related response by reducing the response component that is due to the overlap of temporally adjacent events. We believe that this novel analytical tool can greatly

improve the interpretability of data and as such is essential for understanding the temporal relationship of neuronal activity to behavioral events.

4.5 Methods:

4.5.1 Deconvolution Model:

The deconvolution model assumes that the total firing rate of a neuron in each trial is equal to the linear sum of the contributions of each ERF delayed by the event latencies in that trial. It further assumes that the shapes of the ERFs are independent of inter-event intervals and of time. This means that properties of ERFs such as duration or amplitude do not change by changing the spacing between events and are also stationary over trials. Given these assumptions, the firing of a neuron in trial m with a total of K event types and T time samples can be written as:

$$f_m(t) = \sum_{k=1}^{K} \sum_{\tau=1}^{T} z_{mk}(\tau - t) h_k(\tau) + \eta_m(t) \tag{4.1}$$

Here $f_m(t)$ represents the firing of the neuron in trial m as a function of time and $h_k(t)$ is the neuron's ERF for event k. The term $z_{mk}(t)$ is set to one when event k occurs and is zero otherwise. Furthermore $\eta_m(t)$ represents the noise in trial m and can model both the inherent randomness of neuronal firing as well as the contributions of all other potential events in the task that are not explicitly included in the model.

If the firing rates and noise terms are collected over time and over all the **M** trials in column vectors $F = [f_1(1) \dots f_1(T) \dots f_M(1) \dots f_M(T)]^T$ and $N = [\eta_1(1) \dots \eta_1(T) \dots \eta_M(1) \dots \eta_M(T)]^T$ and the ERFs to be estimated are also collected into another

column vector $H = [h_1(1) \dots h_1(T) \ \dots h_K(1) \dots h_K(T)]^T$ and correspondingly the event timings z_{mk} are turned into a $\mathbf{MT} \times \mathbf{KT}$ matrix referred to as Z, the convolution equation can be written as linear regression:

$$F = ZH + N \tag{4.2}$$

EQ (4.2) shows the full model for the convolution of ERFs over all the trials for a given neuron. The goal of deconvolution is to find an unbiased estimator $\hat{H}$ to approximate the true ERFs in H given that the firing of a neuron in all trials (a.k.a F) as well as the timing of all events (a.k.a Z) in each trial are known. Raw PETHs aligned to each event can be derived from EQ (4.2) by multiplying both sides of the equation with Z^T and normalizing (dividing) by the number of trials.

$$\bar{H} = \frac{1}{M} Z^T F = \frac{1}{M} V H + \frac{1}{M} Z^T N \tag{4.3}$$

where $\bar{H} = [\bar{h}_1(1)...\bar{h}_1(T)....\bar{h}_k(1)...\bar{h}_k(T)]$ is a $\mathbf{KT} \times \mathbf{1}$ column vector containing the raw PETHs for all events. V called hereafter the convolution matrix is defined as and is a $\mathbf{KT} \times \mathbf{KT}$ square matrix that is mapping the K ERFs to the corresponding K PETHs.

EQ (4.3) shows that PETHs are in general themselves only a convoluted version of true ERFs. The further the convolution matrix V is from identity matrix the more distorted $\bar{h}$ will be from h . This in turn will depend on the distribution of event timings in each trial. EQ (4.3) not only confirms our main argument that PETHs can

be very different from the actual neuronal ERFs but also provides a way to formulate this difference simply as a multiplication between a matrix containing all the event timings and the ERFs. To find the best estimate for the ERFs, one needs to undo all the time shifts that are applied to H by the convolution matrix $\frac{1}{M}V$. This can be done either iteratively (see EQ (4.6)) or in a single pass using the least squares technique (see EQ (4.11)). For computational reasons the least squares technique is more efficiently implemented in the frequency domain while for the iterative approach we find it more intuitive to stay in the time domain. However the iterative approach and the least square approach are theoretically applicable in either domain.

4.5.2 Iterative Deconvolution in Time:

To undo the effect of convolution matrix V on the ERFs, an inverse matrix V^{-1} needs to found. It can be shown that for an invertible and properly scaled matrix B, $B^(-1)$ can be calculated using the Neumann series (Stewart 1998):

$$B^{-1} = \sum_{i=0}^{\infty} (I - B)^i \tag{4.4}$$

One can verify that the matrix V which is equal to $Z^T Z$ is a symmetric positive semi-definite matrix and when V^{-1} exists, it can almost always be properly scaled for the Neumann series to converge. To properly scale the convolution matrix $\frac{1}{M}V$ in EQ (4.3) one can divide it by any real positive number $S > \frac{\lambda_{\max}}{2}$ where $\lambda_{\max}$ is the largest eigenvalue of $\frac{1}{M}V$.

The normalized PETH formed can be represented as:

$$\bar{H}_{/S} = \frac{1}{MS} VH + \frac{1}{MS} Z^T N \tag{4.5}$$

Here we use $\bar{H}_{/S}(t)$ to represent the normalized vector of raw PETHs by S . The inverse of convolution matrix $\frac{1}{MS}V(t)$ can then be found using EQ (4.4). The result can be convolved with both sides of EQ (4.5), after which we have:

$$\begin{aligned}
\hat{H}^{(\infty)} &= \sum_{i=0}^{\infty} (I - \frac{1}{MS}V)^i \bar{H}_{/S} = \bar{H}_{/S} + (I - \frac{1}{MS}V)\bar{H}_{/S} + (I - \frac{1}{MS}V)^2 \bar{H}_{/S} + \dots \\
\hat{H}^{(0)} &= \bar{H}_{/S} \\
\hat{H}^{(1)} &= \bar{H}_{/S} + (I - \frac{1}{MS}V)\bar{H}_{/S} \\
\hat{H}^{(2)} &= \bar{H}_{/S} + (I - \frac{1}{MS}V)\bar{H}_{/S} + (I - \frac{1}{MS}V)^2 \bar{H}_{/S} \\
&\vdots
\end{aligned} \tag{4.6}$$

$$\begin{aligned}
E^{(n)} &= \frac{1}{MS} \sum_{i=0}^{n} (I - \frac{1}{MS})^i Z^T N - \frac{1}{MS} \sum_{i=n+1}^{\infty} (I - \frac{1}{MS})^i VH \\
E^{(\infty)} &= \hat{H}^{(\infty)} - H = V^{-1} Z^T N
\end{aligned} \tag{4.7}$$

EQ (4.6) shows that can be constructed by adding the terms of the Neumann series in each iteration with superscripts representing the iteration number and the zero iteration being the normalized PETHs themselves. EQ (4.7) shows an estimate of the final deconvolution error in predicting the true ERFs and its dependence on the trial noise terms.

4.5.3 Complete Deconvolution in Frequency:

Given the fact that convolution in time domain is equivalent to multiplication in frequency domain (Oppenheim et al. 1997) we can rewrite EQ (4.1) this time in the frequency domain, we have:

$$f_m(\omega) = \sum_{k=1}^{K} z_{mk}(\omega)h_k(\omega) + \eta_m(\omega) \tag{4.8}$$

One of the benefits of representing the deconvolution problem in the frequency domain is that as can be seen in EQ (4.8) the value of $f_m(\omega)$ at $\omega = \omega_0$ depends on the values of $z_{mk}(\omega), h_k(\omega)$and $\eta_m(\omega)$ only at $\omega = \omega_0$ and as such is decoupled from all other frequencies and can be studied independent from them. Therefore if the firing rates and noise terms over all the M trials are collected for each frequency in column vectors $F(\omega) = [f_1(\omega)....f_M(\omega)]^T$ and $N(\omega) = [\eta_1(\omega)....\eta_M(\omega)]^T$ and the ERFs to be estimated are also collected into another column vector for each frequency $H(\omega) = [h_1(\omega)....h_K(\omega)]^T$ and correspondingly the event timings z_{mk} are turned into a $\mathbf{M} \times \mathbf{K}$ matrix referred to as $Z(\omega)$, the convolution equation can again be written as linear regression this time in the frequency domain:

$$F(\omega) = Z(\omega)H(\omega) + N(\omega) \tag{4.9}$$

Similar to the time domain the raw PETHs can be obtained by multiplying both sides of EQ (4.9) by $Z^T(-\omega)$:

$$\bar{H}(\omega) = \frac{1}{M}V(\omega)H(\omega) + \frac{1}{M}Z^T(-\omega)N(\omega) \tag{4.10}$$

Where $V(\omega)$ is a $\mathbf{K} \times \mathbf{K}$ the convolution matrix and is equal to $Z(-\omega)^T Z(\omega)$. To deconvolve we need to find the inverse of $\frac{1}{M}V(\omega)$ at each frequency. Multiplying both sides of EQ (4.10) with this inverse results in the deconvolved ERF vector in the frequency domain:

$$\hat{H}(\omega) = M \times V^{-1}(\omega)\bar{H}_{/k}(\omega) \tag{4.11}$$

Assuming white noise, the power density of the error term will be:

$$S_{ee}(\omega) = \sigma^2 V^{-1}(\omega) \text{ where } \sigma^2 \text{ is the noise variance} \tag{4.12}$$

Equations EQ (4.6) and EQ (4.11) should yield equivalent results for deconvolving the ERFs. In theory the frequency domain approach is more efficient for complete deconvolution due to decoupling the convolution effect on each frequency. However in practice due to singularities of the convolution matrix at some frequencies or violations of the model assumptions the partial deconvolution in time domain might provide more accurate results.

The deconvolution requires the convolution matrix to be invertible whether in time or in frequency domain. This means that the data should have enough number of trials with some minimum amount of jitter in the event timings so as to provide

the deconvolution model with enough power to parse out ERFs over all frequencies. This is especially true for the lower frequencies where larger jitters are needed to reveal their amplitude changes over time. In fact in the case of single occurrence of each event in a given trial it can be shown that the DC mapping of the ERFs to the PETHs is not invertible since is a KxK matrix with all its elements set to one. In such cases where deconvolution is not well defined for low frequencies, the deconvolved ERFs may show spurious upward or downward DC shifts that can be corrected back to zero. Having a variable repetition of at least some events during each trial can resolve the ambiguity of the ERFs DC content since the changing number of events in each trial can provide another dimension of information beside the temporal jitters for the deconvolution of the ERFs.

4.5.4 Cross-Validation:

In cross-validation we divide the trials into two groups, one which usually includes the majority of trials is used for training and one with the remaining trials for validation. We have:

$$\bar{H}_t = V_t H + Z_t^T N_t \quad \hat{H}_t^{(n)} = \frac{1}{MS}\sum_{i=0}^{n}(I - \frac{1}{MS}V_t)^i \bar{H}_t$$

$$\bar{H}_v = V_v H + Z_v^T N_v$$

where subscript t represents the training trials and subscript v represents the validation trials. The cross-validation MSE for iteration n is defined as:

$$MSE_{CV} = E((\bar{H}_v - V_v\hat{H}_t^{(n)})^T(\bar{H}_v - V_v\hat{H}_t^{(n)})) \tag{4.13}$$

We can show that this error is correlated with the ERF estimation error and therefore can be used in its place when the true ERFs are not known.

We have:

$$\begin{aligned} MSE_{cv} &= E((\bar{H}_v - V_v\hat{H}_t^{(n)})^T(\bar{H}_v - V_v\hat{H}_t^{(n)})) \\ &= tr(E((V_vH + Z_v^TN_v - V_v\hat{H}_t^{(n)})(V_vH + Z_v^TN_v - V_v\hat{H}_t^{(n)})^T)) \\ &= tr(V_v^2E((H - \hat{H}_t^{(n)})(H - \hat{H}_t^{(n)})^T)) + tr(V_v\sigma^2) \end{aligned} \tag{4.14}$$

$$MSE_{ERF} = E((H - \hat{H}_t^{(n)})^T(H - \hat{H}_t^{(n)})) = tr(E((H - \hat{H}_t^{(n)})(H - \hat{H}_t^{(n)})^T)) \tag{4.15}$$

As can be seen in EQ (4.14) and EQ (4.15) the cross-validation error is similar to ERF error except for an added noise term and scaling by V_v. It can be shown that if ERF estimation error goes down to zero then the cross-validation error goes down to its lower bound set by noise.

We can also divide the data randomly into two equal halves and use the ERFs estimated by one part and the plug in principle to try to directly estimate the ERF estimation error as the following:

$$\bar{H}_1 = V_1 H + Z_1^T N_1 \quad \hat{H}_1^{(n)} = \frac{1}{MS}\sum_{i=0}^{n}(I - \frac{1}{MS}V_1)^i \bar{H}_1$$

$$\bar{H}_2 = V_2 H + Z_2^T N_2 \quad \hat{H}_2^{(\infty)} = V_2^{-1}\bar{H}_2$$

where subscripts 1 and 2 represent the first half and second half of the trials. The MSE for this direct method is defined as:

$$\begin{aligned} MSE_{direct} &= E((\hat{H}_2^{(\infty)} - \hat{H}_1^{(n)})^T(\hat{H}_2^{(\infty)} - \hat{H}_1^{(n)})) \\ &= tr(E((H - \hat{H}_1^{(n)})(\hat{H} - \hat{H}_1^{(n)})^T)) + tr(V_2^{-1}\sigma^2) \end{aligned} \tag{4.16}$$

This shows that the MSE of the direct method is comparable to both the ERF and cross-validation MSEs and in fact where the results of the complete deconvolution is well behaved can be preferred over the cross-validation as being more similar to ERF MSE shown in EQ (4.15).

4.5.5 Behavior and electrophysiology:

Rats were trained with reward predictive auditory cues (up to 10s) randomly presented on a variable interval schedule with an average interval of 60s. A cartoon of the experimental apparatus is shown in **FIG 4.5.a** (For a detailed description of the experimental paradigm see (Ambroggi et al. 2008; Nicola et al. 2004a). Active lever presses only during presentation of the cue resulted in the termination of the cue and delivery of a 10% sucrose reward into an adjacent reward receptacle. The timing of cue presentations, lever-presses and the entries into and exits out of the reward

receptacle were all simultaneously recorded together with the firing of ventral striatal neurons using chronically implanted micro-array electrodes. A sub-set of randomly chosen neurons (n=106) from a previous study (Ambroggi et al. 2008) is used here.

4.5.6 Simulation and data analysis:

All simulations and data analyses were done in MATLAB 7.9.0. The algorithms used for simulation and analysis are available upon request.

Chapter 5

Optimal Binning of Peri-Event Time Histograms Using Akaike Information Criterion

5.1 Abstract

Peri-event time histograms (PETH) are widely used to study correlations between experimental events and neuronal firings. The quality of information that is communicated via PETHs critically depends on the choice of binsize used. We show that the choice of the optimal binsize for a PETH depends on factors such as the trial numbers and the temporal dynamics of the firing rate. These factors warn against use of a one-size-fit-all binsize when analyzing PETHs from various neurons coming from sessions with different trial numbers. Here we propose a binsize selection mechanism by adapting Akaike information Criterion (AIC) to this problem. Simulations show that optimal binsizes estimated by AIC closely match the true optimal binsizes. In real data optimal binning is shown to improve detection of responses and their dynamics. Furthermore it is shown that different response types (excitations versus inhibitions) may require different binsizes depending on the response strength. Together our analysis strongly supports the use of optimized binning of PETHs and proposes an AIC based method as a computationally efficient approach for this optimization process.

5.2 Introduction

The rate coding theory of the nervous system proposes that neurons encode information via changes in their average firing rate. To estimate a neuron's firing rate, its spike trains are usually aligned to a given event and the number of spikes that occur in short time bins are counted. These count histograms are then normalized by the number of trials and by the bin durations resulting in what is known as a peri-event time histogram or PETH. One of the parameters that can affect the accuracy of a PETH in estimating the firing rate is the chosen binsize. However, due to the lack of a straight forward and principled method, assigning the appropriate binsize is often left to the experimenter's subjective assessment. Furthermore, in spite of the fact that noise levels and firing dynamics in different neurons may warrant using variable binsizes for each PETH, it is a common practice to use the same binsize for all the neurons recorded in a task. Importantly, a one-size-fits-all procedure ignores the inherent differences in the temporal resolution of PETHs, which can negatively affect the conclusions drawn from the population analysis of the neuronal activity by increasing response detection error and errors in estimating the response dynamics such as response onset and duration.

Here, we propose that the search for the optimal binsize can be viewed as a special case of a model selection problem where the number of bins (and therefore the binsize) chosen to estimate the firing represent the number of model parameters. As for all model selection problems having a large number of parameters, or in our case small binsizes, results in over-fitting the data while lack of enough parameters

(or large binsizes) fails to capture the variance of the data. Here we develop and propose a method based on Akaike information criterion (AIC) to determine the optimal binsize for each PETH. Our simulation results indicate that in most cases this method provides an accurate estimate of the optimal binsize which minimizes the mean squared error (MSE) between the firing rate estimate and the true firing rate.

5.3 Results

5.3.1 Optimal Binsize Motivation:

It is known that changing the bin duration can affect the quality of the firing rate estimation when making PETHs. The trade off here is between two desirable qualities, one is high temporal precision which favors smaller binsizes and the other one is low noise levels which favor large binsizes. To demonstrate this trade off let us consider two rather extreme examples. For the first example we assume to have a neuron with a constant firing rate during a trial with spikes counts that follow a Poisson distribution. Although in this case the firing rate is constant, the timing of spikes in each trial will be variable as shown in **FIG 5.1a**. Therefore one expects that using larger binsizes can help eliminate the local variation in the inter spike intervals while small binsizes overfit the observed spike data resulting in noisy estimates of the firing (**FIG 5.1b**). The MSE between the PETH estimates using different binsizes and the true firing rate, shows that in this case the estimation error decreases by increasing the binsizes (**FIG 5.1c**). Therefore in this example the largest binsize possible is the optimal binsize to be used. However this conclusion does not hold for our second example neuron. Here we have a neuron that changes its firing rate

rapidly during a trial as shown in **FIG 5.1d**. Unlike the previous example the choice of excessively large binsizes for this neuron results in loss of temporal resolution in estimating the firing rate for this neuron. In this case one might be tempted to use the smallest possible binsize. However choosing too small of a binsize results in noisy estimates by over fitting the local jitters in firing (**FIG 5.1e**). This is in fact what we observe after plotting the MSE between the estimate and the true firing rate. As **FIG 5.1f** shows in this case the optimal binsize takes an intermediate value in the range of binsizes from the smallest to the largest.

These results clearly demonstrate the importance of choosing the right binsize for accurate estimation of firing dynamics for neurons. It also shows that the optimal binsize can be different for different neurons and therefore questions the validity of populational analysis of data with fixed binsizes without regards to the needs of individual PETHs. In the next section we investigate the role of parameters such as the number of trials and the temporal dynamics of firing on the optimal binsize using simulated data. We also show in dealing with real data when the true firing rate is not known the optimal binsize can be approximated by minimizing Akaike Information Criterion (AIC) as a cost function that behaves similarly to the MSE (see Methods).

5.3.2 Simulation Results:

Given certain regularity conditions, any signal can be approximated as an integral of sinusoids over a range of frequencies (i.e. frequency decomposition). Therefore studying the optimal binsize for sinusoidal firing patterns as basis functions should provide insights that are generalizable to other arbitrary firing patterns.

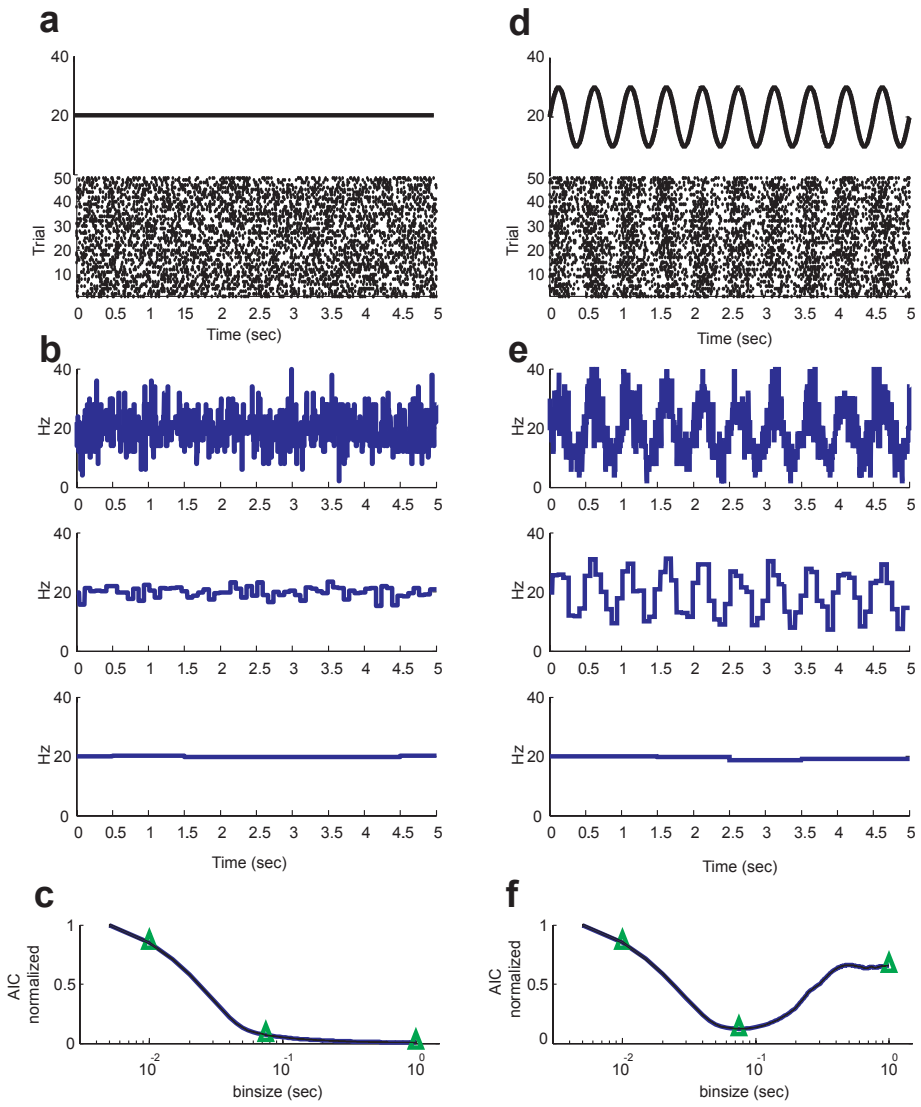

Figure 5.1: **Dependence of firing rate estimate on the binsize choice**(**a**) Firing rate and simulated spike train raster with constant Poisson rate of 20Hz. (**b**) Firing rate estimate using 10 ms binsize (1st row) 75 ms (2nd row) and 1 sec (3rd row). (**c**) Mean sum of squared error for the constant Poisson firing monotonically decreases by increasing the binsize. Green triangles mark the 10 ms, 75 ms and 1 sec binsizes used in the estimates shown in part A. (**d**) Firing rate and simulated spike raster with a sinusoidal oscillating firing rate with an average of 20Hz and oscillation frequency of 2Hz. (**e**) Firing rate estimate using 10 ms binsize (1st row) 75 ms (2nd row) and 1 sec (3rd row). (**f**) MSE for the oscillating Poisson firing reaches a minimum at 75 ms. Green triangles mark the 10 ms, 75 ms and 1 sec binsizes used in the estimates shown in part **c**.

Let us start with the hypothetical neuron shown in **FIG 5.2a.** As shown in the figure this neuron has an oscillatory firing with an average of 20 spikes/sec and a sinusoidal modulation of ± 10 spikes/sec at 1Hz. Let us further assume that we have

recorded the activity of this neuron in 50 trials of 5 second durations each. As can be seen over trials the neurons activity follows the general rise and fall of the underlying sinusoidal rate while the details of spike timings vary from trial to trial as is typically seen with real neuronal activity.

To estimate the underlying firing rate we can construct firing rate histograms using varying binsizes and calculate their MSE in estimating the true sinusoidal firing rate. As expected too large or too small of a binsize fail to capture the true fluctuations of the firing rate over time as shown by the changes of MSE over binsize values in **FIG 5.2c**. In this case the minimum of the MSE happens at 115ms. If we now use the spike timings shown in **FIG 5.2a** and calculate and plot the AIC cost (see **EQ** (5.3)) over the same binsize range on one can see that in fact the changes of AIC and MSE as function of binsize are quite similar (**FIG 5.2c**). The optimal binsize estimate based on AIC also closely matches the estimate base on MSE (AIC 110ms vs MSE 115ms). The significance of this result lies in he fact that it shows the possibility of accurately predicting the effect of binsize change on firing rate estimation error without using or knowing the true firing rate. The PETH made with the optimal binsize is shown in **FIG 5.2b** as the best approximation of the true sinusoidal firing.

In general the value of the optimal binsize is a function of number of trials available. One can expect that increasing the number of trials should decrease the optimal binsize as more trials improve the signal to noise ratio by cancelling the noise over trials thus allowing smaller binsizes to be used. For the example neuron shown in **FIG 5.2** simulations show that this in fact is the case. The optimal binsize which

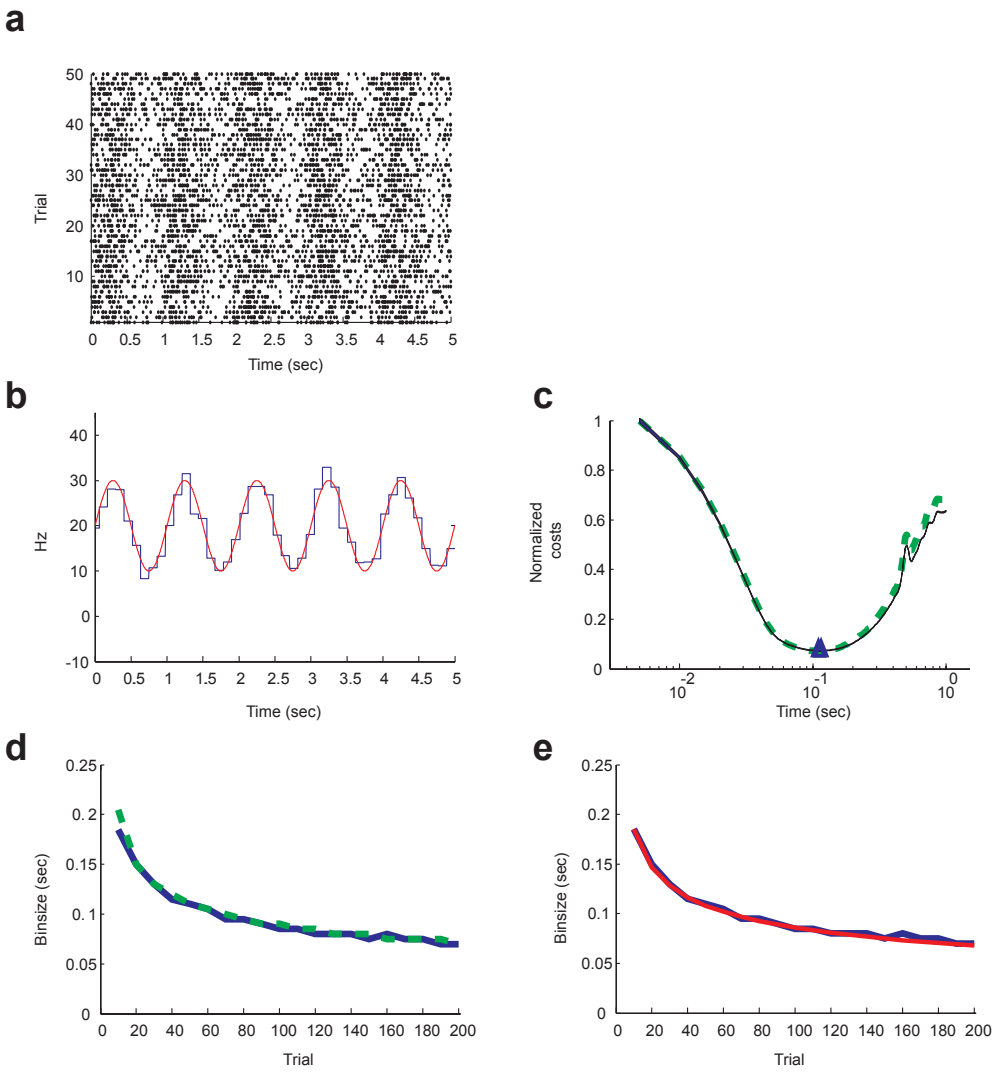

Figure 5.2: **An example of finding optimal binsize using AIC measure and its dependence on the number of trials** (**a**) Simulated spike train raster with a sinusoidal oscillating firing rate with an average of 20Hz and oscillation frequency of 1Hz. (**b**) Firing rate estimate using the optimal binsize (blue) and the true sinusoidal firing rate (red) overlaid. (**c**) MSE (solid blue) and AIC costs (dashed green) as a function of the binsizes overlaid. AIC changes over the binsizes closely match that of MSE. (**d**) Decrease of the optimal binsize as estimated by MSE (solid blue) and by AIC (dashed green) by increasing the number of trials from 10 to 200. Changes of the optimal binsize as predicted by MSE and AIC closely match each other. (**e**) The decrease of the optimal binsize by increasing the number of trials as predicted theoretically (solid red) matches the estimates made by MSE (solid blue) over a range of trials from 10 to 200.

minimizes the MSE decreased from 180ms with only 10 trial to lower than 70ms with 200 trials (**FIG 5.2d** solid blue). Once more we can see that over this moderate range of trial numbers the optimal binsize based on the AIC closely follows the trend

of changes in the optimal binsize based on MSE (**FIG 5.2e** dashed green).

In fact the dependence of optimal binsize of data histograms on the sample size has been studied theoretically in the field of nonparametric density estimations as far back as 1979 (Izenman 1991; Scott 1979). The result of those studies which can be generalized to firing rate histograms mostly suggest that the optimal binsize decreases proportional to $\sim N^{-\frac{1}{3}}$ with N being the number of trials (or samples in their case). This is slower than the rate of the reduction of noise power which changes as $\sim N^{-\frac{1}{2}}$ but in fact matches the reduction of optimal binsize seen in **Figure 5.2** for the simulation data (**FIG 5.2d,e**). Therefore the cube root law can provide a convenient tool to extrapolate the number of trials that is needed in a given neuron to reach a certain binsize precision.

Besides the dependence on the trial number, the optimal binsize should also show dependence on the rate of firing rate changes during a trial. Again the intuition here is that the faster the firing rate changes over time the smaller the optimal binsize should be in order to capture the temporal changes. **Figure 5.3** shows the dependency of the optimal binsize on the frequency of the sinusoid used in our simulations. The optimal binsize based on MSE and AIC (solid and dashed lines, respectively) show an inverse relation to the firing rate frequency. Furthermore over the range of frequencies used AIC based optimal binsize is in good agreement with MSE based optimal binsize (**FIG 5.3a**).

The dependence of optimal binsize to frequency of firing rate changes can also be

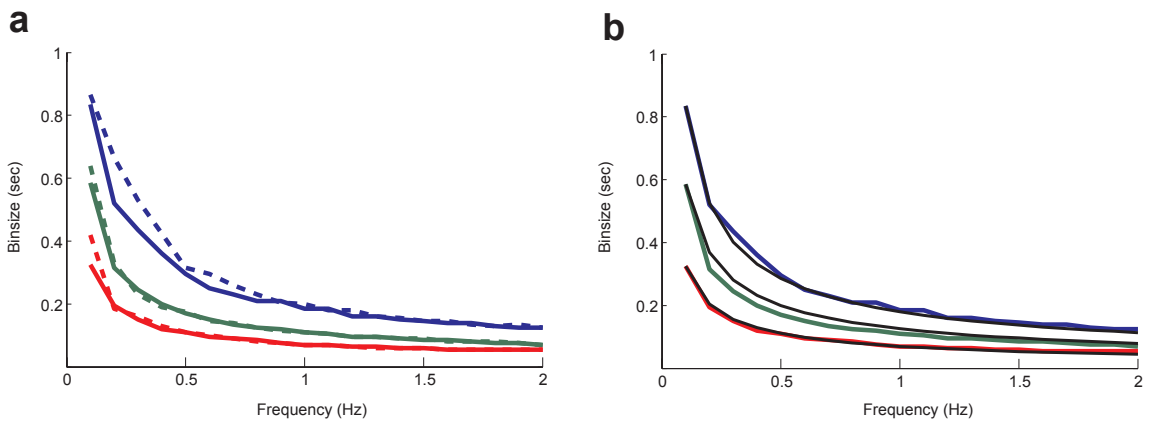

Figure 5.3: **Codependence of the optimal binsize on the firing rate bandwidth and the trial number**(**a**) The optimal binsize decreases by increasing the frequency of the firing rate changes (0.1 to 2 Hz) and also by increasing the trial numbers (10, 50 and 200 shown by blue, green and red respectively). Solid lines are based on the MSE and the dashed lines are based on the AIC. (**b**) The decrease of the optimal binsize by increase of either the trial numbers or the firing rate frequency as predicted theoretically (black) matches the estimates made by MSE over a range of bandwidths and trial numbers.

derived theoretically and can be shown to be proportional to $\sim \omega_0^{-\frac{2}{3}}$ where ω_0 represents the frequency of the sinusoid (see Methods **EQ** (5.5)). This can in fact be verified in **Figure 5.3b** where the power law is shown to match the changes of MSE over the range of simulated frequencies.

5.3.3 Experimental Results:

In this section we first demonstrate the effect of varying the binsizes in detection of responses to a reward predictive cue. We used a sub-set of 106 neurons recorded in the nucleus accumbens (NAc) in a discriminative stimulus task. Briefly, two auditory cues, discriminative stimuli (DS) and neutral stimuli (NS) were randomly presented (with durations of up to 10 s for the DS and 10 s for the NS) on a variable interval schedule with an average interval of 30 s. A lever press was required during DS presentation to cause the delivery of a reward into an adjacent reward receptacle and

responding during NS was never rewarded. For simplicity, only the DS is displayed in **Figure 5.4** with intervals between the two DS's shown on average to be 60s. (For a detailed description (Nicola et al. 2004a))

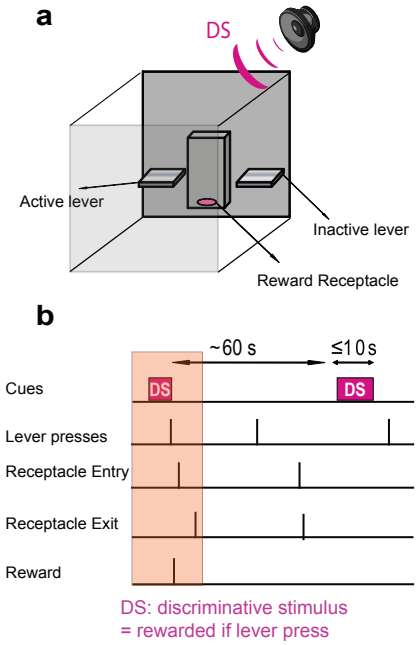

Figure 5.4: **Experimental procedure**(**a**) Rats were placed in 40 x 40cm chambers containing two levers (active and inactive) on each side of a reward receptacle. A speaker played the auditory cues above the chamber. (**b**) Behavioral paradigm showing hypothetical sequence of events. Only active lever presses during the cue presentation resulted in reward delivery in the receptacle after which the rat would normally enter the receptacle to consume the reward and exit the receptacle after the consumption. The event sequence consisting the cue, active lever press, receptacle entry and exit in the presence of reward is counted as one rewarded trial (shaded area).

Responses of neuron to the DS can be obtained by making PETHs that are time locked to the cue presentation. Significant changes in the firing rate of the neuron were detected whenever the firing rate of the neuron fell outside the 95% percentile of the baseline (pre-cue) firing rate. Thus excitations are defined as significant increases and inhibitions are defined as significant decreases in the firing rate. **Figure 5.5a** shows the spike rasters of 5 example neurons that show changes in their spiking pattern following the DS presentation. It can be seen that using small binsize of only

10ms for the neurons results in noisy estimates of their firing rate (**FIG 5.5b**). The noisiness of the estimate seem to especially affect the detection of inhibitory responses that follow the cue (neurons #1,2,3 and 5). This is because the depression in the firing rate fails to reach the threshold for significance due to the noisiness of low baseline firing (floor effect). In some of these cases the threshold for inhibition is at zero which makes it impossible for detect any significant inhibitions (neurons #2,3 and 5). The noisiness of the estimate not only reduces the detection of inhibitions, thus increasing the false negative error rate, but can also increase detection of spurious activity and thus introduce false positive errors. For example in **FIG 5.5b** neurons 1,2,3 and 4 all show seemingly short excitatory or inhibitory activity in the baseline or long after the cue onset. These are most likely spurious based on a cross examination of the firing raster and by using less noisy PETHs made by larger binsizes. On the other hand using the optimal binsize for each PETH reveals the inhibitory responses to the cue. It also eliminates many of the falsely detected activities thus reducing the false positive rate (**FIG 5.5d**). Looking at the AIC curve over a range of binsizes from 10ms up to 1 sec for each neuron shows that estimation error to decrease rapidly by for binsizes $<\sim$60ms. This is followed by a deflection point (DP) in the AIC curve after which the changes of the AIC curve over the binsizes become small (see **FIG 5.5e**). The optimal binsize which corresponds to the AIC minimum in these cases happens for values larger than the DP binsize.

Although using the optimal binsize results in the best estimate of the firing rate as shown in **Figures 5.5d** and **5.5e**, due to rather large sizes (>500ms for neurons

3,4 and 5) it results in insufficient temporal resolution when a more precise detection of changes in the firing is required. Since the AIC gain of using optimal binsize over the DP is rather small (<10% increase in AIC cost) in cases where temporal precision is important like detecting the response onset one might choose to use the DP instead of the optimal binsize. However it should be noted that using DP can yield suboptimal results in response detection. For example while it outperforms the minimum binsize in revealing the inhibitions for neurons 1,2 and 3, in the more noisy cases for neurons 4 and 5 it fails to detect the excitation and inhibition to the cue (**FIG 5.5c**). Cross examination of rasters in **Figure 5.5a** show the presence of a weak excitation in neuron 4 and an a weak inhibition in neuron 5 following the cue onset which is correctly detected if optimal binsizes are used instead (**FIG 5.5d**)

AIC curves shown in **FIG 5.5e** also demonstrate that that the optimal binsize is different for each of the neurons shown. Neurons that have a more robust and less noisy response have a smaller optimal binsizes (neurons 1 and 2 with 155ms and 265ms optimal bins respectively) while for neurons with more noisy firings larger optimal binsizes are required to reveal the significant fluctuations in their rate (neurons 3, 4 and 5 with 685ms, 835ms and 940ms optimal bins respectively). On the other hand the DP binsize which turns out to be roughly the same for all neurons in the task ($\sim$60 ms) marks a lower bound on the level of temporal resolution that is possible for this population of neurons before the AIC cost gets prohibitively large.

We know from theoretical and simulation results that the optimal binsize for each neuron depends on the number of trials. If neuronal response is stationary over trials

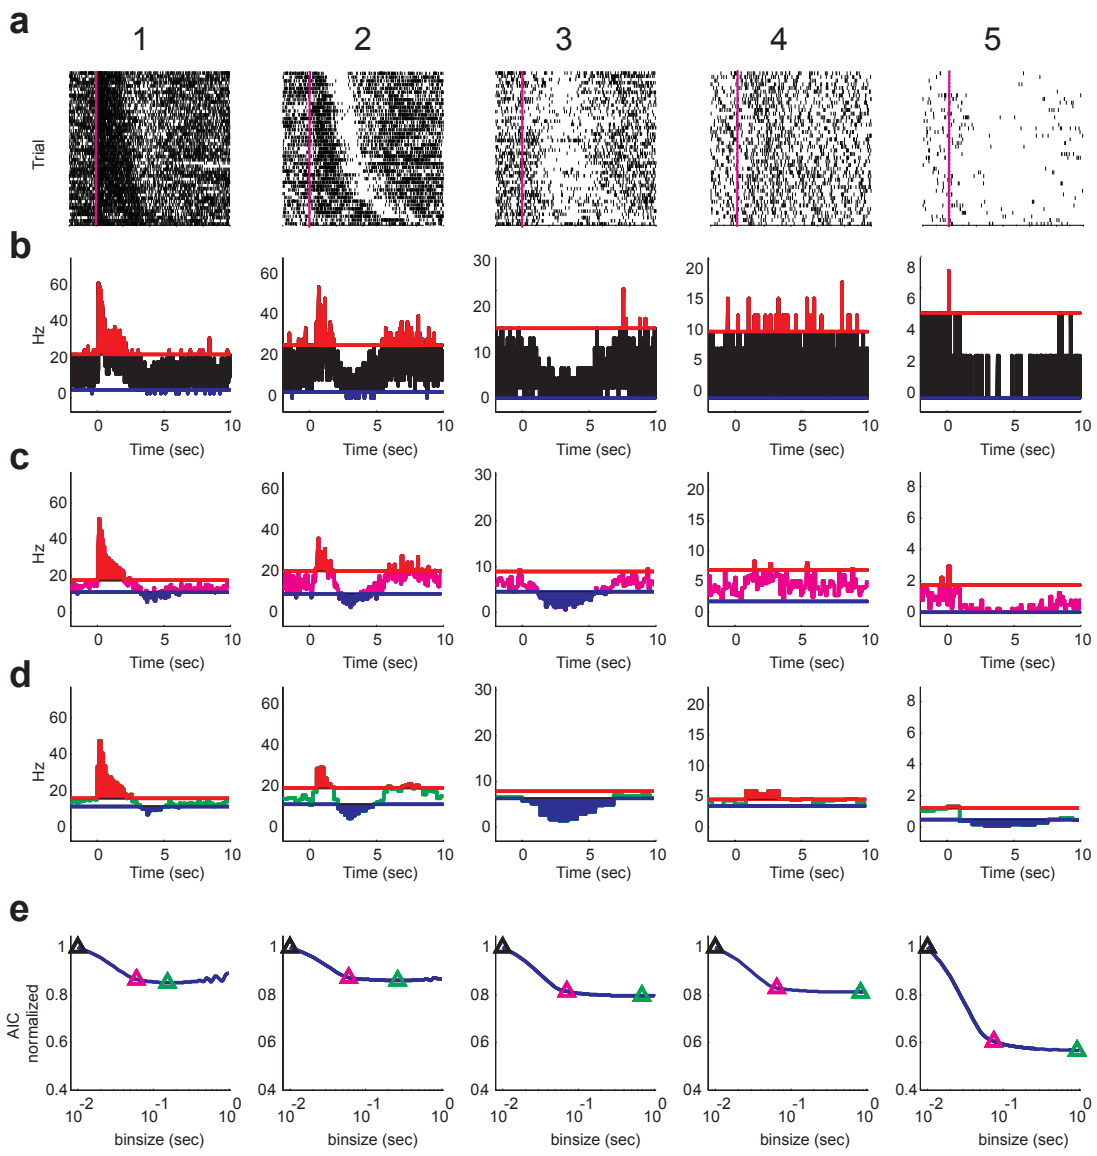

Figure 5.5: **Example PETHs using small, DP and optimal binsizes.**(**a**) Firing raster of five example neurons and their corresponding PETHs with (**b**) binsize of 10ms, with (**c**) binsize equal to the AIC deflection point and with (**d**) optimal binsize for each neurons. (**e**) The AIC cost function for each neuron over a range of binsizes from 10ms to 1 sec. The AIC at 10ms binsize, deflection point and the optimal binsize are marked by black, magenta and green triangles, respectively (columns 1-5 correspond to neurons 1-5).

(i.e. its response only depend on the time locked event not trial per se) the optimal binsize should be a decreasing function of the trial number following the previously discussed $\sim N^{-\frac{1}{3}}$ power law. However there are often other behavioral events that happen with variable timing in trials relative to the time locked event that can

cause significant non-stationarities. For instance DS PETHs like the ones shown in **Figure 5.5** are often modulated by other temporally close events such as lever press and entry into the receptacle which vary in their latency in each trial. Therefore one should predict that the reduction of optimal binsize with the real firing data should be much lower than the theoretical rate. One way to estimate the dependence of optimal binsize for the PETHs of a given neuron is by choosing various random subsets of the available trials and calculate the resulting optimal binsize for a given trial number. Performing this analysis on our experimental data shows that on average, the optimal binsize decreases by adding more trials to the experiment. However, as expected the rate of this decrease is slower than predicted for stationary rates (**FIG 5.6a**). The distribution of the decrease rates of optimal binsize over trial numbers for all the neurons also shows that while some neurons follow the theoretical value of $\sim N^{-0.33}$ most neurons show a slower decrease as a function of trial number while others show no change or even an slight increase of the optimal binsize as a function of trial number (**FIG 5.6b**). This results in average rate being $\sim N^{-0.1}$ for this population of neurons.

Another parameter that is known to affect the optimal binsize is the bandwidth of the firing rate. Firing rates with higher bandwidths that have larger and more frequent changes require smaller binsizes to capture their temporal variation compared to firing rates with slower changes. Although we do not have access to the true firing rate of our neurons we should be able to get a relatively accurate estimate of their bandwidth by using the firing rate estimates made with the optimal binsize. Plotting

the optimal binsize of each neuron as a function of its firing bandwidth in fact shows a relatively good agreement with the theoretical values as the optimal binsizes are shown to be a decreasing function of the bandwidth with a rate proportional to $\sim \omega_0^{-0.66}$ is close to theoretical rate of $\sim \omega_0^{-1}$ (**FIG 5.6c**).

The optimal binsize shows a bimodal distribution over in our neuronal population with two peaks one around 0.2 sec and another around 0.9 sec. Further separation of the binsize distribution between the neurons excited to the cue and inhibited to the cue shows that in fact the peak with the smaller binsize corresponds to excitatory responses while the second peak with the larger binsizes corresponds to the inhibitory responses (**FIG 5.6d**) The distribution of the DP binsize on the other hand shows a much narrower spread around 50ms for excitations and 70ms for inhibitions (**FIG 5.6e**). The average AIC measure for excitatory and inhibitory responses also demonstrate the fact that the optimal binsize of the average AIC for inhibitions are larger than for excitation with the difference diminishing for the DP binsize (**FIG 5.6f**).

As discussed previously for the examples in **Figure 5.5** the choice of binsize can greatly affect the detection of responses and their temporal dynamics including the magnitude, onset and offset of excitations and inhibitions. To further examine this effect we detected response to DS using both the optimal binsize and a fixed 60ms (DP-like) binsize. Since the value of DP binsize was very similar for all neurons and was shown to be around 60ms (**FIG 5.6e**) the choice of this binsize gives similar results to DP binning while having the added benefit of being the same for all neurons. While similar number of responsive neurons is detected using 60ms or optimal

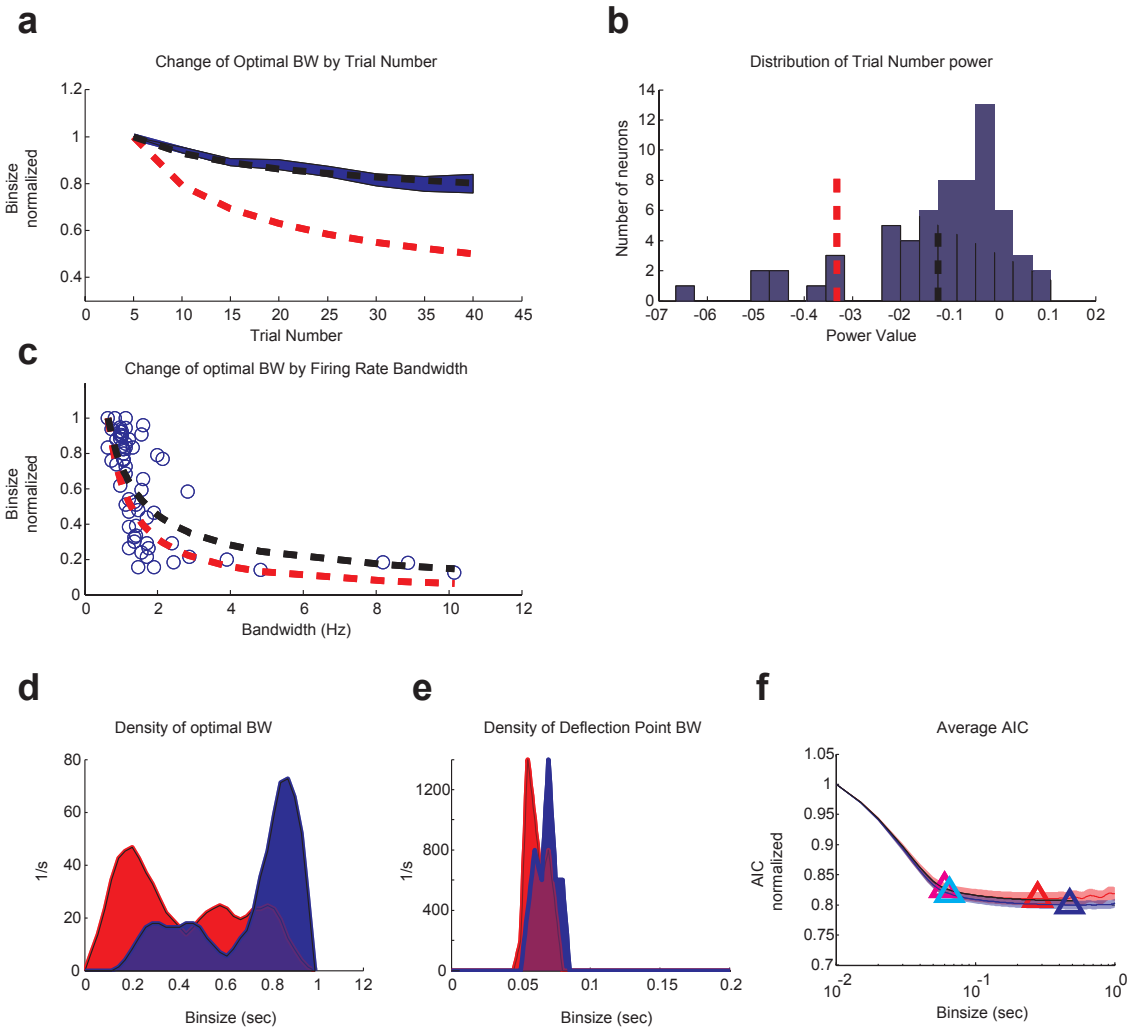

Figure 5.6: **Distribution of optimal binsize in the recorded neurons and its codependence on the trial number and firing rate bandwidth** (**a**) The optimal binsize decreases on average for the recorded neurons as the number of trials increases (solid blue: average optimal binsize, blue shade: standard error of the mean). The theoretical rate of decrease (dashed red) and the polynomial fit to the data (dashed black) showing a slower rate of decay for recorded data compared to what is predicted for an idealized stationary data. (**b**) The distribution of the rate of decay of the optimal binsize by the number of trials. Negative values indicate decreasing and positive values indicate increasing with larger trial numbers. Theoretical rate of decay (dashed red) and the average rate from the recorded data (dashed black). (**c**) The optimal binsize decreases on average for the recorded neurons as bandwidth of the firing rate increases. Each open circle corresponds to one neuron. The theoretical rate of decrease (dashed red) and the polynomial fit to the data (dashed black) showing a similar rate of decay for recorded data compared to what is predicted. (**d**) Density of the optimal binsizes among all cue responsive neurons overlaid for excitations (red) and inhibitions (blue).

binsize (49 and 54 for OP and fixed 60ms, respectively) the profiles of detected cue responses look different in each condition (**FIG 5.7a**). Using the optimal binsize

generally results in much higher Z-scores in the responsive neurons compared to the baseline mainly due to reduction of baseline noise. This is true for both excitatory and inhibitory responses. The average Z-score for each response type shows a marked increase by using the optimal binsize over 60ms binsize (**FIG 5.7b**). Inhibitory responses are specifically enhanced in their Z-score after using the optimal binsize resulting in a more symmetrical look in the amplitude of excitatory and inhibitory Z-scores using optimal binsizes for each neuron.

The distribution of response durations also shows that using the optimal binsizes results in detecting an array of durations from less than 1 sec to more than 8 sec while most of the response durations detected by the DP are detected to be less than 0.5 sec (**FIG 5.7c**). This is the result of the many discontinuities observed in the profile of responses using DP binsize for each neuron as shown in **FIG 5.7a**. On the other hand due to the relatively large values of the optimal binsizes the distribution of the response onsets using optimal binsizes lacks the resolution to show the natural variation seen when using smaller binsizes. Therefore using a small DP like binsize to make PETHs has a higher resolution for the onset detection and results in more accurate estimates of the response onset for this dataset (**FIG 5.7d**).

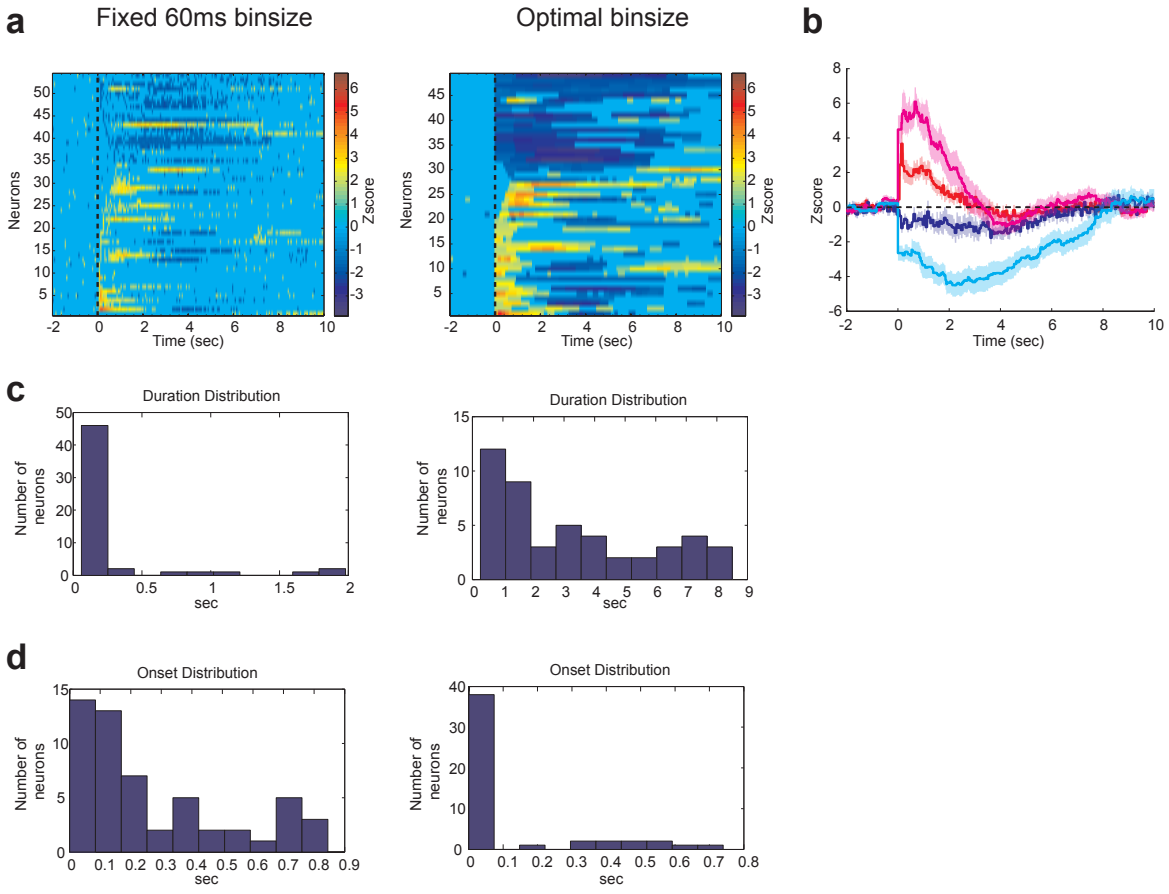

Figure 5.7: **Comparison of population profile of responses to the cue based on PETHs made with optimal versus fixed DP-like binsizes** (**a**) Color coded PETH Z-scores made with fixed 60ms (left column) and with optimal binsize (variable) (right column) for all the responsive neurons to the cue. Each line in the color coded plots corresponds to the firing of one neuron in Z-score. Only Z-scores outside the 90% confidence interval of baseline (pre-cue firing) are color coded. Values within the confidence interval are all colored at zero. (**b**) Average PETH Z-scores made with fixed 60ms binsize versus Z-scores made with optimal binsize. Red and blue correspond to excitatory and inhibitory responses made with DP binsize respectively. Magenta and cyan correspond to excitatory and inhibitory responses made with optimal binsize respectively. (**c**) Distribution of the cue response durations made with fixed 60ms (left column) and with optimal binsize (variable) (right column) for all the responsive neurons to the cue. (**d**) Distribution of the cue response onsets made with fixed 60ms (left column) and with optimal binsize (variable) (right column) for all the responsive neurons to the cue.

5.4 Discussion

One of the main goals of systems neuroscience is to relate the firing of neurons in the brain to the behavioral or environmental events that are generated by it or influence its rate. PETHs are widely used in the field for exactly this purpose and have contributed in a major way to our functional understanding of the nervous system. The usefulness of PETH depends on at least two qualities. First is how well the PETH manages to cancel the firing variability across time and trials to best portray the firing rate modulation that is related to the event of interest. Second is how well it can illustrate the temporal changes of this firing modulation with high resolution. Maximizing both of these qualities is often not possible since noise cancellation often requires use of larger binsizes for a PETH that in turn reduces the temporal resolution (**FIG 5.1**).

We have addressed this trade off between noise and temporal resolution by proposing a principled way for choosing the optimal binsize for each PETH by using Akaike information criteria or AIC. Briefly this method defines a cost with one term that favors smaller binsizes to minimize a MSE cost and another term that favors larger binsizes to minimize the total number of bins (considered as the number of model parameters) used for PETH construction. The simulation results show that PETHs made with the AIC optimal binsize often have the lowest MSE errors when compared to the true firing rates.

We investigated two of the parameters that can affect the optimal binsize for a PETH. The first parameter affecting the optimal binsize is the number of trials used

to make a PETH. Intuitively the higher the number of trials is the more confident one get about the quality of the firing rate estimation and lower binsizes can be used because a good portion of noise cancellation is done over trails. Theoretical results estimate the inverse relation between trial number and optimal binsize to follow a cube power law. Our simulation results is again in good agreement with this result (**FIG 5.3**). An important assumption for this relationship is that the true neuronal firing to the event of interest should not change from trial to trial (stationarity). However, the mere presence of other variably occurring events that can modulate the firing of the neuron during a trial nullifies this assumption when dealing with actual neural data. In line with the violation of stationarity we find that the actual decrease rate of the optimal binsize as a function of trial number is often much slower than cube power with an average power of about -0.1 (**FIG 5.6a,b**).The second parameter is how much and how fast the firing of a neuron shows changes to the event of interest. The faster these changes happen in the firing the smaller the binsize required to capture their temporal dynamics. Theoretical results show that the optimal binsize is almost inversely correlated with the frequency content of the firing rate modulation. We find good agreement with this theoretical result in both simulations (**FIG 5.2&5.3**) and in analysis of actual data (**FIG 5.6c**).

We also showed that binsize has a dramatic impact on the response detection. Many false positives or false negatives can result if binsizes that are too small or too large are used for a PETH. The optimal binsize for each PETH differs according to the number of trials and its firing dynamics. This difference in the optimal binsizes

means that using the same binsize for analysis of neural population can lead to misclassification or even misinterpretation of the data. In our population of neurons in the NAc which are considered to be mainly low firing medium spiny neurons this misinterpretation can particularly occur when comparing excitations and inhibitions. Due to the floor effect of low baseline firing and the slower dynamics of inhibitions it is often useful to make inhibitory PETHs with larger binsizes (>0.5). In fact using binsizes that are often sufficient for detection of excitations often result in missing weak inhibitions in the neuronal population (**FIG 5.5**). In fact analysis of optimal binsize across excitations and inhibitions clearly show this separation of optimal binsizes across the two groups with inhibitions having a mode around 0.8s while the mode for excitations happen much earlier at 0.2s. Interestingly use of optimal binsize for both excitations and inhibitions show that the average magnitude of inhibitions can be comparable to excitations in the NAc, a fact that is easily missed if small 60ms binsizes are used for all neurons (**FIG 5.7b**).

One draw back using optimal binsizes can be the loss of high resolution estimates of the response onset times since optimal binsize is often large (>0.2). A solution that we put forward is to use the binsize at the deflection point of AIC cost. This binsize (DP binsize ∼60ms) happens after the sharp drop of the AIC cost but is usually much smaller than the optimal binsize (**FIG 5.5e**). We demonstrated that the use of DP binsize can improve the onset estimation while having little effect on the percentage of detected neurons (**FIG 5.7d**). However, for the response duration optimal binsize still seem to provide more accurate results compared to short binsizes (**FIG 5.7c**).

In conclusion we demonstrated the importance of the binsize choice in accuracy of firing rate estimations using PETHs. We proposed that optimal binsize for a PETH can be estimated by using an AIC-based cost and we investigated the dependence of this optimal binsize on other parameters such as firing rate modulation and the trial number. Furthermore our results showed that for neurons with low basal firing rate like NAc medium spiny neurons small binsizes can miss many of the inhibitory responses to the events. Accordingly optimal binsizes for inhibitions turn out to be much larger than excitations. There we propose use of optimal binsize for PETH based response detection.

5.5 Methods

5.5.1 Model Selection Approach to Binning:

Let us assume that we have multiple trials of a neuron's firing with the same underlying rate $\lambda(t)$ in each trial. Our goal is to find $\hat{\lambda}(t)$ which minimizes the mean sum of squared error or MSE defined as following:

$$MSE = E(\frac{1}{T}\sum_{t=1}^{T}(\hat{\lambda}(t) - \lambda(t))^2) \tag{5.1}$$

Where E refers to expectation over all trials and T represents the trial duration in samples where time is considered to be a discrete variable which depends on the sampling rate. Minimizing MSE with respect to $\hat{\lambda}(t)$ should result in a minimum variance estimate of the true firing rate during a trial which is usually unbiased.

The difficulty of minimizing MSE arises from the fact that in dealing with actual

data $\lambda(t)$ in not known. Therefore one needs to find an approximate way to minimize the MSE.

Here our model space includes the PETHs that are used to estimate $\lambda(t)$. We have:

$$\hat{\lambda}(t,\Delta) = \sum_{i=1}^{P}(\frac{K(i\Delta) - K(i\Delta - \Delta)}{N\Delta}\Pi(t - i\Delta, \Delta)) \tag{5.2}$$

Where Δ refers to the binsize and $P \equiv \frac{T}{\Delta}$ is the total number of bins during a trial. We define $K(t)$ as the total number of spikes in all the Ntrials from the beginning of the trials up-to time t and finally $\Pi(t)$ is a rectangular function defined as following:

$$\Pi(t,\Delta) = \begin{cases} 1 & 0 < t < \Delta \\ 0 & elsewhere \end{cases}$$

Using the PETH formulation presented in EQ (5.2), we observe that the total number parameters that are used to construct the PETH is equal to the number of bins. Therefore the problem of choosing the best binsize which minimizes MSE can be viewed as a model selection problem with different number of parameters being the number of bins.

Among the various model selection methods in statistics, Akaike Information Criteria or AIC is one that is probably the most popular and at the same time the most effective for a vast range of model selection problems. Therefore we have adapted this method for choosing the optimal binsize. In deriving the AIC formulation for

our current problem we assumed that the estimation errors are close to normally distributed in which case we have:

$$AIC = T\ln(\frac{SSE}{T}) + 2P \text{ where } SSE = \sum_{t=1}^{T}(\hat{\lambda}(t,\Delta) - \hat{\lambda}(t,1))^2 \tag{5.3}$$

Here $\hat{\lambda}(t,1)$ represents the firing rate estimate with the smallest binsize equal to 1 sample unit. $\hat{\lambda}(t,1)$ is taken to be the model free raw data that is being estimated by $\hat{\lambda}(t,\Delta)$. SSE is minimized when Δ is goes to one in which case the number of parameters P is as large as the number of samples or T. However since AIC includes a term that penalizes large parameter numbers in most cases AIC is only minimized at intermediate binsizes that result in lower number of parameters compared to T. This is not accidental since very small binsizes can result in over fitting the data and are not necessarily optimal.

In the simulation results we will demonstrate that in fact the AIC formulation presented in EQ (5.3) behave similarly to MSE as a function of binsize over a wide range of experimental conditions including number of trials and firing rate changes.

5.5.2 Dependence of optimal bin to the frequency content:

This dependency is also theoretically worked out using a similar measure of optimality as MSE (Scott 1979). Adapted to our problem and assuming a continuous time firing rate the dependency of the optimal binsize to the true firing rate looks like the following:

$$\Delta^* \sim (\int_{-\infty}^{+\infty} (\frac{d\lambda(t)}{dt})^2 dt)^{\frac{1}{3}} \tag{5.4}$$

Using Parseval's theorem and Fourier transform properties we can rewrite EQ (5.4) in the frequency domain as following:

$$\Delta^* \sim (\int_{-\infty}^{+\infty} \omega^2 \lambda(\omega)^2 d\omega)^{\frac{1}{3}} \tag{5.5}$$

For the sinusoidal firings considered here EQ (5.5) can be solved explicitly and shows the optimal binsize to be proportional to $\sim \omega_0^{\frac{2}{3}}$ where ω_0 represent the frequency of the sinusoid. This can in fact be verified in **Figure 5.3b** where the power law matches the changes of MSE over the range of simulated frequencies. More generally for firing rates that have relatively constant power upto a cutoff frequency the optimal binsize will be simply proportional to the inverse of the cutoff frequency or alternatively $\sim \omega_c^{-1}$ with ω_c refereeing to the cutoff.

5.5.3 Other Methods:

In an elegant study by Shimazaki and Shinomoto, the authors mathematically derived a cost function based on the MSE for firing rates that follow Possion distribution (Shimazaki and Shinomoto 2007). This cost function adapted to our notation is:

$$C(\Delta) = \frac{2}{N\Delta} < \hat{\lambda}(t,\Delta) > - \frac{1}{T} \sum_{t=1}^{T} (\hat{\lambda}(t,\Delta) - < \hat{\lambda}(t,\Delta) >)^2$$

where $< \hat{\lambda}(t, \Delta) >= \frac{1}{T} \sum_{t=1}^{T} \hat{\lambda}(t, \Delta)$

However as they point out their method diverges for slow fluctuating rates if the number of trials is not large enough. This is in part due to the fact that their derivation is an approximate solution to the problem of minimizing the MSE where the expectation in EQ (5.1) is replaced by the trial average. Therefore at its current formulation Shimazaki-Shinomoto method is the same as a maximum likelihood estimate based on the observed data which should converge to the true MSE as long as the law of large numbers applies given a large number of trials. AIC on the other hand tries to minimize the MSE over all possible observations and as such provides a superior estimate compared to any maximum likelihood approach. As can be seen in **Figure 5.8** the Shimazaki-Shinomoto cost function tends to overestimate the optimal bin size for slow changing firing rates and low trial numbers. Increasing the number of trials brings this error down for a given frequency (**FIG 5.8a**). This overestimation of the binsize for slow changing firing rates (lower frequencies) translates into a larger MSE when using the Shimazaki-Shinomoto cost function compared to the AIC although both methods seem to perform equally well for fast changing firing rates (**FIG 5.8b,c**).

In another study authors have suggested to use an information based technique for optimal firing rate estimation (Paulin and Hoffman 2001). Although they have used their method to find the optimal level of smoothing, their method can readily be generalized to finding an optimal binsize. The measure of optimality used here is closely related to an information theoretic concept known as "cross-entropy" and

aims to minimize the disparity between the predicted probability of spiking and the observed probability of spiking in each moment of time. This is done by dividing the data into two halves one used to make predictions and the other for testing. Therefore this approach is equivalent to using a cross validation technique using an information based measure. In fact it has been shown that AIC, itself is equivalent to a cross validation method (Stone 1974) Although we can verify that their method with some modification should be equivalent to AIC ,however, given the relative complexity of the cross entropy compared to AIC and the computational extensiveness of cross validation we propose AIC as the simpler and more straightforward measure to be used for finding the optimal binsize.

5.5.4 Behavior and electrophysiology:

Rats were trained with auditory cues (up to 10s) randomly presented on a variable interval schedule with an average interval of 60s. A cartoon of the experimental apparatus is shown in **FIG 5.4a** (For a detailed description of the experimental paradigm see (Nicola et al. 2004a)). Active lever presses only during presentation of the cue resulted in the termination of the cue and delivery of a 10% sucrose reward into an adjacent reward receptacle. The timing of cue presentations, lever-presses and the entries into and exits out of the reward receptacle were all simultaneously recorded together with the firing of ventral striatal neurons using chronically implanted micro-array electrodes. A sub-set of randomly chosen neurons (n=106) from a previous study is used here. From these neurons only 64 that had basal firing rate of higher than 1Hz are used for analysis shown in **Figure 5.6a-c**.

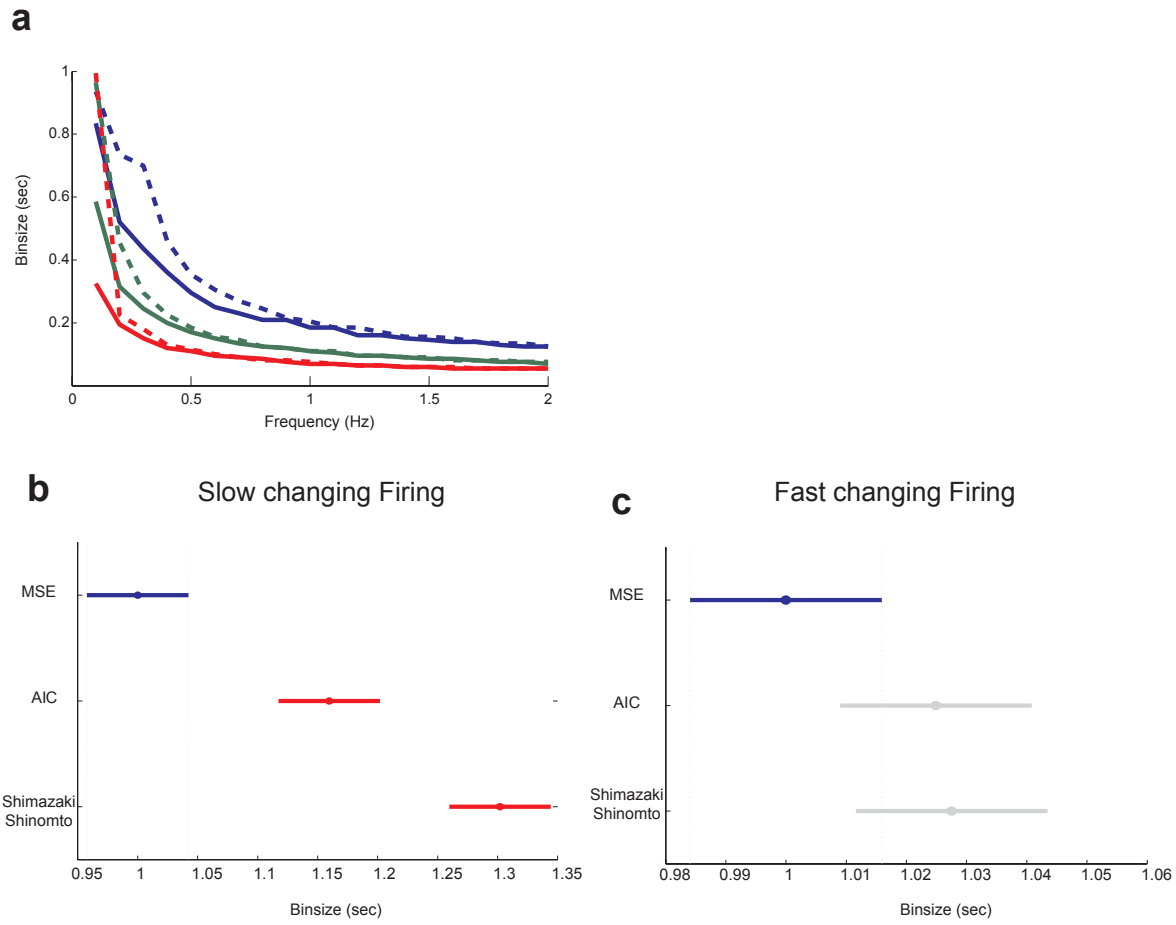

Figure 5.8: **Comparison of AIC with Shimazaki-Shinomoto method**(**a**) Comparison of the optimal binsizes as estimated by MSE (solid lines) and Shimazaki-Shinomoto method (dashed lines) over a range of firing rate frequencies and trial numbers. Shimazaki-Shinomoto method tends to overestimate the optimal binsize for low frequency (slowly changing) firing rates.Comparison of the average sum of squared error between the true firing rate and the estimated firing rate using optimal binsize estimated by minimizing MSE, AIC and Shimazaki-Shinomoto method for (**b**) slowly changing firing rates <0.5Hz or low bandwidths and (**c**) fast changing firing rates>0.5Hz or high bandwidths.

5.5.5 Simulation and data analysis:

All simulations and data analyses were done in MATLAB 7.9.0 by the first author. For all the firing rate simulations, the spike trains are generated using a Poisson point process with a temporally changing rate. The algorithms used for simulation and analysis are available upon request.

Part III

third Part

Chapter 6

Summary and Discussion

6.1 Background and Summary of Aims:

Success in goal directed behaviors often requires not only learning to express the reinforced actions but also to avoid engaging in irrelevant or harmful behaviors. One of the key results of this thesis is the demonstration that the action suppressive component in a learned appetitive task lags behind the action promotive component during training (**FIG 3.1b**). The animals used in my experiments, Long Evans rats, tend to initially express actions at high rates. Some of those actions are rewarded and some are never rewarded but nevertheless expressed. After many days of training animals learn to weed out the unnecessary actions and fine tune the rewarded actions to a minimal rate sufficient for most efficiently earning a maximum number of rewards. The presence of an active suppressive component even after many weeks of training is revealed following inactivation of NAc or vmPFC. Inactivating either region disinhibits some of the same actions that rats have learned to suppress during training (**Chs.2&3**). This is consistent with the idea that response extinction is a general feature of operant conditioning. However, my results indicate that the response extinction does not have to be learned only following extinction of action-outcome associations (see (Peters et al. 2008; Quirk et al. 2006)) but can happen

spontaneously when no explicit outcome extinction is present.

Previously in **Ch.1**, I proposed a conceptual framework which included three layers of selection that are engaged prior to and during goal directed behaviors namely: drive selection, selection and attention to the relevant sensory space and selection and expression of appropriate actions that satisfy the current drive. As I argued previously, the assumption of a selection mechanism necessitates the existence of a dual process, one that promotes selection and the other that suppresses the unselected but potential competition. Hence, the suppressive component in goal directed behaviors can be manifested at any or all the three levels. This may partially explain why it is hard to provide a simple model of impulsivity and why there are many impulsivity-related constructs and behavioral models within the field of experimental psychology and cognitive/systems neuroscience (Evenden 1999a; Nigg 2000).

The DS task used in this thesis provides a unique opportunity to probe into a variety of impulsivity-related measures specially the aspects that have to do with response inhibition and to some degree impulsive choice (e.g. premature entry) (**Ch. 1& FIGs 3.2&3.3**). The work in this thesis was launched after earlier work in our lab implicated NAc and vmPFC in inhibition of NS responding and unrewarded lever pressing (Ishikawa et al. 2008a; Yun et al. 2004). For my dissertation I focused on three specific aims: 1) what are the details of behavioral disinhibition that is observed in the DS task following NAc or vmPFC inactivations and how would they relate to the selection levels discussed previously? 2) How much of the similarities and differences in the behavioral role of NAc core and shell, revealed through aim 1, is

reflected in their neuronal representation of the DS task? 3) Given the anatomical and some functional disconnection evidence implicating vmPFC projections to NAc shell in action suppression, how do vmPFC projections influence the activity of neurons in the NAc shell for action suppression?

For the first aim, I developed an assay for analysis of behavioral microstructure in the DS task. This assay looks at the rate of responding (e.g. lever pressing, receptacle entry, etc) in four behaviorally relevant epochs during a session. These epochs include responding during the DS, responding during the NS, responding followed by the reward receipt (Post_reward) and spontaneous responding (Spont) which includes all responses outside the other three epochs. This behavioral assay was an outgrowth of the engage-disengage perspective (**FIG 6.1**) of the DS task that I developed when I was originally introduced to the DS task and was influenced by views put forth in a paper by Saleem Nicola (Nicola 2007). I used rats that had cannulae targeting NAc core or shell as well as vmPFC to inactivate each region during the DS task. I then analyzed the behavioral changes using both aggregate and microstructural measures to dissect the roles of core and shell as well as vmPFC in each element of the DS task.

For the second aim, I recorded neuronal activity from rats implanted with microarray electrodes in the NAc shell in the DS task and to further elucidate the role of NAc subregions in the DS task,. I then used the neural data recorded in this batch of rats and the data that was available from a previous experiment in the NAc core in the DS task (Ambroggi et al. 2008) to look for differences and similarities in the neural code in each sub-region that may explain their role in the DS task.

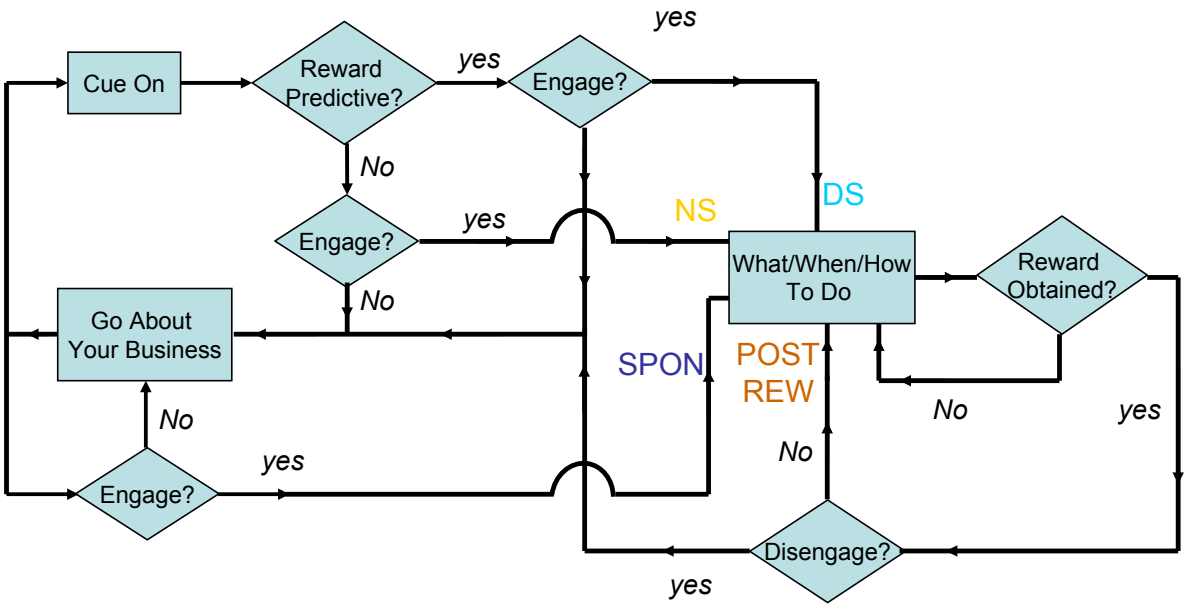

Figure 6.1: **Engage-disengage paradigm** This figure demonstrates a possible algorithm for decision making in the DS task. The proposed paradigm provides a rationale for microstructural analysis of responding in four epochs: DS, NS, Spont and Post_reward which is used in this thesis.

Finally for the third aim, rats were implanted with both cannulae in the vmPFC and electrode arrays in the NAc shell to allow for simultaneous inactivations of vmPFC inputs to the NAc shell which allowed concomitant analysis of behavioral and neuronal changes following inactivation. To my knowledge, this method provided for the first time the opportunity to probe the influence of vmPFC on the NAc neurons mediating suppression of actions in awake behaving animals.

It is also worth mentioning that I have developed a few novel methodologies for the analysis of neural data that mainly have to do with construction and interpretation of PETHs. Implementation and discussion for each method and its mathematical foundations can be found in **Chs.4&5** and are not further discussed here.

6.2 Aim 1: Pharmacology and Behavioral Analysis

Our data indicate that the inactivation of either NAc shell or vmPFC results in significant disinhibition as measured by aggregate behavioral measures such as NS ratio, rates of active and inactive lever pressing and receptacle entry. Also both inactivations result in specific increases in active lever presses during the NS and Post_reward epochs beyond the increase in the Spont epoch. There is a significant increase and a trend toward significant increase in the spontaneous lever pressing (Spont epoch) following NAc shell and vmPFC inactivations respectively. Compared to CSF no significant change in the DS ratio or DS response latency is observed after shell or vmPFC inactivations. This situation is in stark contrast with NAc core inactivation which results in a significant reduction of DS ratio (by about 50%) as well as a significant increase in the latency to respond to the DS. Mild disinhibitory effects are also observed after NAc core inactivation as measured by increased responding during Post_reward. But many of the disinhibitory effects observed after shell inactivation are shown to be smaller or absent with core inactivations (e.g. inactive lever rate increase is higher in shell than core and NS and Spont active lever rate increases are specific to shell) (**FIG 2.3&3.2**).

Our behavioral results suggest that NAc core is mainly involved in the promotive process of the DS task consistent with other studies that showed it to be invigorating incentive cue responding (Blaiss and Janak 2009; Floresco et al. 2008). Inactivations of excitatory inputs to this sub-region was shown in a previous study to also reduce DS ratio dramatically (Ishikawa et al. 2008b). On the other hand NAc shell and the

prefrontal region projecting to it namely, vmPFC are shown here to powerfully inhibit various unreinforced behaviors (e.g. unrewarded active and inactive lever presses) and to exert this inhibition over at least three of the four epochs studied (i.e. SPON, NS, Post_reward). In particular vmPFC is known to inhibit actions after their outcome associations are extinguished over training in both fear conditioning and appetitive conditioning paradigms (Peters et al. 2008; Peters et al. 2008; Quirk et al. 2006) but see (Bossert et al. 2011). One of the novel results in this thesis is that it expands the tentative role of vmPFC and NAc shell to suppression of all unnecessary actions whether or not they have gone through explicit outcome extinction training. (In the DS task none of the suppressed actions such as NS responding or inactive lever pressing was ever explicitly paired with reward)

I have previously discussed (**Ch.1**) how behavioral disinhibition in each behavioral epoch might be interpreted as and find parallels in other behavioral paradigms that are designed to probe various aspects of impulsivity. My results indicate that inactivation of vmPFC or NAc shell increases impulsive responding in at least four distinct categories: 1) The inactivations reduce the ability of the animals to exercise action restraint during the ITI as measured by increased lever pressing during the Spont epoch. I argued before that this aspect is similar to the response disinhibition seen in tasks such as DRL (**Ch.1**). 2) The inactivations reduce the ability of the animals to suppress responding to the interfering stimuli (i.e. the NS). This ability was discussed under the general category of attentional impulsivity and resistance to distracters in **Ch.1** and has potential links to the phenomena of latent inhibition

3) The inactivations increase the rate of perseverative responding following reward receipt with potential links to perseveration and compulsivity 4) Finally the inactivation of vmPFC is shown to increase premature entry during the DS. This effect may address a similar phenomenon measured by tasks such as FCN or impulsive choice and may reveal inefficiencies in accurate sequencing of actions.

Together the task impairments that are seen following vmPFC or NAc shell or core inactivations in the DS task implicate this corticostriatal circuitry at all the three selection levels involved in goal directed behaviors: Increased spontaneous lever pressing and post reward perseverative responding can be interpreted as problems in drive selection or switching in and out of an appetitive behavior. Reduction of responding to DS (core) and increased responding to NS (shell, vmPFC) can be related to compromised attentional and sensory selection and increased premature entry or inactive lever presses can reflect suboptimal action selection. These interpretations are consistent with role of mPFC and NAc in many other tasks that indicate multiple roles for this circuitry in goal directed behaviors such as drug self-administration (drive or incentive sensitization), latent inhibition (sensory suppression) and set-shifting (action selection).

6.3 Aim 2: NAc Core and Shell Neuronal Responses

Surprisingly, a comprehensive analysis of neural firing in core versus shell shows that virtually all the response types that are present in one structure are also present in the other despite the differences in the behavioral role of each region in the DS task.

However there are several notable quantitative differences in the neural representation of task events in each region. Consistent with the role of NAc core in invigorating the DS response, a significantly larger percentage of neurons are found to be DS excited in the core compared to shell (18% vs. 12%). Furthermore the Z-score peak and area under the curve (AUC) of DS excitations in neurons that selectively responded to DS but not NS are larger in core compared to shell. Shell, on the other hand, shows a significantly higher proportion of active lever press inhibited neurons. If one assumes that these inhibitions gate the expression of active lever press similar to other gating responses reported in the NAc (Taha and Fields 2006), then this result is consistent with the stronger disinhibitory effect that is observed following NAc shell inactivation compared to NAc core. On the other hand, lever press excitations in core are shown to have a higher Z-score peak and AUC compared to shell consistent with their presumed role in the promotive component of goal directed behaviors.

Overall, the presence of many event-related firings in both core and shell suggests that their different behavioral effects in the DS task may be instead mediated through their different projection targets. I previously (**Ch.1**) reviewed the differences in the efferent projections of NAc core and shell. Most notably there is a mediolateral divide between core and shell projection targets in the VP and midbrain dopamine areas (i.e. shell projecting to more medial and core to more lateral subregions). Also shell is believed to be unique among striatal regions in having direct projections to structures outside basal ganglia such as lateral hypothalamus (Zahm and Brog 1992 but see Tripathi et al. 2010). Furthermore there is a feedforward connection from NAc

shell to core both through thalamocortical loops and mesencephalic ascending spirals (Haber et al. 2000; Zahm 2000) which argues for a possible hierarchy in behavioral control. Further experiments looking at the basal ganglia nuclei downstream to core and shell are needed to test whether these differences in the projection targets can better account for the dissociable behavioral role of NAc core and shell.

6.4 Aim 3: NAc Shell Firing Modulation Following vmPFC Inactivation

Simultaneous inactivation of vmPFC and recordings in NAc shell reveal significant changes in the tonic and phasic responses of neurons that encode unrewarded actions in the NAc shell consistent with their hypothesized role in promotion or suppression of responding along with behavioral disinhibition of unrewarded actions. There is a significant increase in the Z-score peak of both unrewarded lever press excitations and inhibitions. As I argued previously in comparing the neural firing in core and shell if we assume the excitations to be promoting and inhibitions to be gating actions, this result points to one possible neural mechanism that can mediate the increased rate of lever pressing following vmPFC inactivation. Furthermore, I find a significant increase in the Z-score peak of all NS excitations that include emergent NS excitations following vmPFC inactivation. Again, if the role of cue excitations is taken to be promotion of responding like the argument for DS excitations (Ishikawa et al. 2008b; Yun et al. 2004), the emergence of NS excitations in the shell following vmPFC inactivation can point to a possible explanation for the elevated NS ratio beyond spontaneous increases in responding. On the other hand, on average none of reward

related neural firings show a significant change following vmPFC inactivation. This is important as it matches the behavioral results that show disinhibition only in unreinforced actions but no significant change in responding to DS following vmPFC inactivation.

Moreover, the data show both significant increases and decreases in baseline (tonic firing) in the B/M condition compared to the CSF. Interestingly neurons with increased baseline tend to be on average excited to the receptacle entry and NS while neurons with decreased baseline tend to be on average inhibited to the receptacle entry and not responsive to NS. In fact many action inhibited neurons show significant (or a trend toward significance) decreases in their baseline. If these action inhibited neurons which represent a large proportion of NAcS neurons, are in fact causally suppressing unrewarded actions (i.e. gating) the reduction in their baseline is consistent with a higher rate of unrewarded action expression. In fact analysis on the modulation of magnitude of firing difference between action excited and action inhibited neurons show that not only can this difference recreate the pattern of lever pressing in the CSF condition during the behavioral epochs discussed before, but it can reflect to some degree the increase in the rate of lever presses after vmPFC inactivations in those same epochs.

Finally, given the fact that the projections from vmPFC to NAc shell are supposed to be excitatory, the emergence of excitations requires an intervening GABAergic influence. Currently some of the likely explanations are thought be feedforward inhibitions by fast spiking interneurons in the NAc, lateral inhibitions of MSN or feedback

or feedforward projections from other structure such as VP or BLA. Investigating the mechanism of these emergent excitations awaits future experiments.

6.5 Future Directions:

Given the results presented in this thesis new questions can be raised about the neural circuitry in the NAc and mPFC that mediates behavior in goal directed behaviors such as in DS task. Some of these questions are discussed below:

1) What is the neural representation of DS task in the vmPFC? Also how does the encoding of the DS task evolve in the vmPFC and NAc circuitry?

To further explore the corticostriatal circuitry involved in suppression of unnecessary actions in the DS task, it is essential to have data on how this task is represented in the vmPFC. There is already evidence of emergent cue excitations in the IL area following extinction in the fear conditioning (Milad and Quirk 2002). It would be tempting to speculate that the NS representation in the IL might also emerge following the suppression seen during training. Representation of other event related firings in the vmPFC and their influence on the NAc activity in the DS task are not known and beg to be studied in the future.

Moreover since our behavioral training data argue for development of both promotive and suppressive processes during learning, it would be interesting to have neuronal recordings in the vmPFC and NAc shell and core during learning of the DS task to study the time course of the evolution of event related responses and to compare it with behavioral improvements observed. Electrophysiological parallels to

behavior during learning can be conducive to finding neural correlates of promotive and suppressive components in goal directed behaviors.

2) What are the causal roles of various event related responses in NAc core and shell?

I find robust encoding of DS in both NAc core and shell regardless of their differential role in DS responding. Specifically, the pharmacological results show that NAc core has a key role in promotion of DS responding while NAc shell does not seem to affect DS responding. It was previously shown in our lab that both phasic DS excitations and inhibitions decrease dramatically in NAc core following VTA or dmPFC inactivations (Ishikawa et al. 2008b; Yun et al. 2004). This reduction of DS coding is accompanied by a significant reduction of DS responding. These results, combined with core inactivation results, argue strongly for a causal role for DS responsive neurons in NAc core in behavioral DS responding. But we currently do not have a clear understanding of the role of DS responsive neurons in shell. For example, it is not known what happens to DS responsive neurons in NAc shell after the inactivation of dmPFC or VTA when DS responding is reduced. There are data suggesting that at least some cue excited neurons in NAc might have a role in response suppression instead of response promotion (Ambroggi et al, unpublished results in a go/no-go task). An interesting possibility that is consistent with shell inactivation data is that some cue excitations in this region might be similar to no-go neurons. This possibility is strengthened if experiments that manipulate the rates of responding (e.g. inactivations, devaluations or contingency degradations) show that cue responses in core

and shell are modulated in the opposite direction after devaluation or contingency degradation (e.g. core DS excitations decrease while shell DS excitations increase following devaluation). However, the emergence of NS excitations in shell after vmPFC inactivation that accompanied behavioral disinhibition suggests that at least some cue excitations in shell are involved in response promotion as well. Furthermore neural recordings as well as pharmacological inactivation in areas downstream from NAc core and shell can be performed in the DS task to test the idea that the differences in the projection targets of core and shell may explain differential or similar functions for cue responses in each sub-region.

Similar questions regarding the causal roles of the neurons that encode lever presses, receptacle entry, exit and consumption are still not fully answered. My vmPFC inactivation results argue for instance that unrewarded lever press related excitations and inhibitions might have opposite roles in the DS task (promote and suppress responding respectively). However until there is a way to selectively manipulate the firing of these neurons in the NAc the question of causality is not fully resolved. Addressing the causality issue is in fact one of the main motivating factors for the next future direction.

3) What are the histochemical cell type characteristics and anatomical connections of various event related neurons in the NAc?

Almost all the events or states in the DS task are encoded by a population of neurons in the NAc. One of the key questions at this stage is to be able to identify whether neurons belonging to various response types also segregate in terms of their

anatomical afferent and efferent connections or their histochemical and receptor expressions. This is an extremely crucial piece of data if one wants to understand the neural circuitry that mediates goal directed behaviors in the NAc. Recent developments in the optogenetic field along with the promise of transgenic rats with cell type specific promoters have created high hopes for resolving some of the longstanding gaps in our understanding about the circuitry underlying these goal directed responses in the NAc. Future experiments are needed to systematically identify the cell type and connection characteristics of various event related populations that are involved in response promotion and suppression. Specifically in relation to vmPFC-shell circuitry and its role in response suppression, further details on the cell type characteristics and anatomical connections of the emergent responses (e.g NS excitations, lever press inhibitions etc) await future experimentations. The discoveries made in this area can also be crucial for exploring the question of causality by allowing selective manipulation of cells that share similar firing patterns in the behavior if they turn out to have similar immunohistochemistry or anatomical connections.

6.6 Conclusion:

In conclusion I have established the role of vmPFC and NAc in particular its shell sub-region in suppression of unreinforced actions. My results implicate this corticostriatal circuitry in multiple aspects of impulsivity. vmPFC is proposed to suppress unrewarded behavioral responding by differentially modulating neurons that promote and suppress responding in NAc shell through inhibition and excitation of each population respectively. My results also implicate NAc core but not shell in

the promotion of rewarded actions. The presence of similar firing patterns in the neuronal in NAc core and shell in the DS task suggests that the differences in the neuronal representation alone may not fully explain the dissociable role of these two subregions in response promotion and suppression and differences in the functional roles of downstream projection targets for each sub-region should be investigated in the future.

Finally I have developed novel methodologies for construction and interpretation of PETHs that are widely used for analysis of neural activity. Notably a deconvolution method is presented that efficiently eliminates PETH contaminations by temporally close events with an iterative approach that is sensitive to signal to noise ratios in each PETH. Furthermore an optimal binning technique is presented which takes into account the number of trials as well as the firing dynamics of each neuron to choose the best binsize for a given PETH. I have demonstrated that these methods significantly improve the analysis and interpretation of data. Both methods are used for the analysis of the results presented in this thesis.

Bibliography

Ainslie, G. (1975). Specious reward - behavioral theory of impulsiveness and impulse control. *Psychological Bulletin 82*(4), 463–496.

Alexander, G. E., M. R. Delong, and P. L. Strick (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience 9*, 357–381.

Allen, T., F. Moeller, H. Rhoades, and D. Cherek (1998). Impulsivity and history of drug dependence. *Drug Alcohol Depend 50*(2), 137–45.

Ambroggi, F., A. Ishikawa, H. Fields, and S. Nicola (2008). Basolateral amygdala neurons facilitate reward-seeking behavior by exciting nucleus accumbens neurons. *Neuron 59*(4), 648–61.

Anderson, J. R. (2010). *Cognitive psychology and its implications* (7th ed.). New York: Worth Publishers.

Basar, K., T. Sesia, H. Groenewegen, H. Steinbusch, V. Visser-Vandewalle, and Y. Temel (2010). Nucleus accumbens and impulsivity. *Prog Neurobiol*.

Beauchamp, G. K., M. Bertino, D. Burke, and K. Engelman (1990). Experimental sodium depletion and salt taste in normal human volunteers. *American Journal of Clinical Nutrition 51*(5), 881–889.

Bechara, A. and M. Van der Linden (2005). Decision-making and impulse control after frontal lobe injuries. *Current Opinion in Neurology 18*(6), 734–739.

Belin, D., A. Mar, J. Dalley, T. Robbins, and B. Everitt (2008). High impulsivity predicts the switch to compulsive cocaine-taking. *Science 320*(5881), 1352–5.

Berendse, H. W. and H. J. Groenewegen (1990). Organization of the thalamostriatal projections in the rat, with special emphasis on the ventral striatum. *Journal of Comparative Neurology 299*(2), 187–228.

Berridge, C. W., T. L. Stratford, S. L. Foote, and A. E. Kelley (1997). Distribution of dopamine beta-hydroxylase-like immunoreactive fibers within the shell subregion of the nucleus accumbens. *Synapse 27*(3), 230–241.

Berridge, K. C. (2001). Reward learning: Reinforcement, incentives, and expectations. *Psychology of Learning and Motivation: Advances in Research and Theory, Vol 40 40*, 223–278.

Bertran-Gonzalez, J., C. Bosch, M. Maroteaux, M. Matamales, D. Herve, E. Valjent, and J. A. Girault (2008). Opposing patterns of signaling activation in dopamine d-1 and d-2 receptor-expressing striatal neurons in response to cocaine and haloperidol. *Journal of Neuroscience 28*(22), 5671–5685.

Bezzina, G., T. H. C. Cheung, K. Asgari, C. L. Hampson, S. Body, C. M. Bradshaw, E. Szabadi, J. F. W. Deakin, and I. M. Anderson (2007). Effects of quinolinic acid-induced lesions of the nucleus accumbens core on inter-temporal choice: a quantitative analysis. *Psychopharmacology 195*, 71–84.

Bindra, D. (1978). How adaptive-behavior is produced - perceptual-motivational alternative to response-reinforcement. *Behavioral and Brain Sciences 1*(1), 41–52.

Blaiss, C. A. and P. H. Janak (2009). The nucleus accumbens core and shell are critical for the expression, but not the consolidation, of pavlovian conditioned approach. *Behavioural Brain Research 200*(1), 22–32.

Bohrquez, J. and O. Ozdamar (2006). Signal to noise ratio analysis of maximum length sequence deconvolution of overlapping evoked potentials. *J Acoust Soc Am 119*(5 Pt 1), 2881–8.

Bossert, J. M., A. L. Stern, F. R. Theberge, C. Cifani, E. Koya, B. T. Hope, and Y. Shaham (2011). Ventral medial prefrontal cortex neuronal ensembles mediate context-induced relapse to heroin. *Nat Neurosci 14*(4), 420–2.

Botvinick, M. M. (2007). Multilevel structure in behaviour and in the brain: a model of fuster's hierarchy. *Philosophical Transactions of the Royal Society B-Biological Sciences 362*(1485), 1615–1626.

Bowman, E. M., T. G. Aigner, and B. J. Richmond (1996). Neural signals in the monkey ventral striatum related to motivation for juice and cocaine rewards. *Journal of Neurophysiology 75*(3), 1061–1073.

Breiter, H. C., I. Aharon, D. Kahneman, A. Dale, and P. Shizgal (2001). Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron 30*(2), 619–639.

Broadbent, D. (1958). *Perception and communication. New York: Pergamon Press.* New York: Pergamon Press.

Burns, L. H., B. J. Everitt, A. E. Kelley, and T. W. Robbins (1994). Glutamate-dopamine interactions in the ventral striatum - role in locomotor-activity and responding with conditioned reinforcement. *Psychopharmacology 115*(4), 516–528.

Byron, M. (2004). *Satisficing and maximizing : moral theorists on practical reason.* Cambridge ; New York: Cambridge University Press.

Cabanac, M. (1979). Sensory pleasure. *Quarterly Review of Biology 54*(1), 1–29.

Cardinal, R. N. (2006). Neural systems implicated in delayed and probabilistic reinforcement. *Neural Networks 19*(8), 1277–1301.

Cardinal, R. N., J. A. Parkinson, J. Hall, and B. J. Everitt (2002). Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neuroscience and Biobehavioral Reviews 26*(3), 321–352.

Cardinal, R. N., D. R. Pennicott, C. L. Sugathapala, T. W. Robbins, and B. J. Everitt (2001). Impulsive choice induced in rats by lesions of the nucleus accumbens core. *Science 292*(5526), 2499–2501.

Chambers, R. A. and M. N. Potenza (2003). Neurodevelopment, impulsivity, and adolescent gambling. *Journal of Gambling Studies 19*(1), 53–84.

Chaudhri, N., L. L. Sahuque, W. W. Schairer, and P. H. Janak (2010). Separable roles of the nucleus accumbens core and shell in context- and cue-induced alcohol-seeking. *Neuropsychopharmacology 35*(3), 783–791.

Christakou, A., T. W. Robbins, and B. J. Everitt (2004). Prefrontal cortical-ventral striatal interactions involved in affective modulation of attentional performance: Implications for corticostriatal circuit function. *Journal of Neuroscience 24*(4), 773–780.

Chudasama, Y., F. Passetti, S. E. V. Rhodes, D. Lopian, A. Desai, and T. W. Robbins (2004). Dissociable aspects of performance on the 5-choice serial reaction time task following lesions of the dorsal anterior cingulate, infralimbic and orbitofrontal cortex in the rat: differential effects on selectivity, impulsivity and compulsivity (vol 146, pg 105, 2003). *Behavioural Brain Research 152*(2), 453–453.

Cloninger, C. R. (1987). A systematic method for clinical description and classification of personality variants - a proposal. *Archives of General Psychiatry 44*(6), 573–588.

Comon, P. (1994). Independent component analysis, a new concept? *Signal Processing 36*(3), 287–314.

Corbetta, M. and G. L. Shulman (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience 3*(3), 201–215.

Cragg, S. J. (2006). Meaningful silences: how dopamine listens to the ach pause. *Trends in Neurosciences 29*(3), 125–131.

Crespi, L. P. (1942). Quantitative variation of incentive and performance in the white rat. *American Journal of Psychology 55*, 467–517.

Cromwell, H. C. and W. Schultz (2003). Effects of expectations for different reward magnitudes on neuronal activity in primate striatum. *Journal of Neurophysiology 89*(5), 2823–2838.

Dale, A. (1999). Optimal experimental design for event-related fmri. *Hum Brain Mapp 8*(2-3), 109–14.

Dalley, J., T. Fryer, L. Brichard, E. Robinson, D. Theobald, K. Lne, Y. Pea, E. Murphy, Y. Shah, K. Probst, I. Abakumova, F. Aigbirhio, H. Richards, Y. Hong, J. Baron, B. Everitt, and T. Robbins (2007). Nucleus accumbens d2/3 receptors predict trait impulsivity and cocaine reinforcement. *Science 315*(5816), 1267–70.

Dalley, J., A. Mar, D. Economidou, and T. Robbins (2008). Neurobehavioral mechanisms of impulsivity: fronto-striatal systems and functional neurochemistry. *Pharmacol Biochem Behav 90*(2), 250–60.

Dalley, J. W., R. N. Cardinal, and T. W. Robbins (2004). Prefrontal executive and cognitive functions in rodents: neural and neurochemical substrates. *Neuroscience and Biobehavioral Reviews 28*(7), 771–784.

Daruna, J. and B. P.A. *A neurodevelopmental view of impulsivity. The impulsive client: Theory, research, and treatment.* Washington, DC, US: American Psychological Association.

Deniau, J. M., A. Menetrey, and A. M. Thierry (1994). Indirect nucleus-accumbens input to the prefrontal cortex via the substantia-nigra pars reticulata - a combined anatomical and electrophysiological study in the rat. *Neuroscience 61*(3), 533–545.

Deutsch, J. A. and D. Deutsch (1963). Attention - some theoretical considerations. *Psychological Review 70*(1), 80–90.

DiCiano, P., R. N. Cardinal, R. A. Cowell, S. J. Little, and B. J. Everitt (2001). Differential involvement of nmda, ampa/kainate, and dopamine receptors in the nucleus accumbens core in the acquisition and performance of pavlovian approach behavior. *Journal of Neuroscience 21*(23), 9471–9477. Times Cited: 118.

DiCiano, P. and B. J. Everitt (2001). Dissociable effects of antagonism of nmda and ampa/ka receptors in the nucleus accumbens core and shell on cocaine-seeking behavior. *Neuropsychopharmacology 25*(3), 341–360. Times Cited: 137.

DiCiano, P., T. W. Robbins, and B. J. Everitt (2008). Differential effects of nucleus accumbens core, shell, or dorsal striatal inactivations on the persistence, reacquisition, or reinstatement of responding for a drug-paired conditioned reinforcer. *Neuropsychopharmacology 33*(6), 1413–1425. Times Cited: 24.

Eagle, D. M. and T. W. Robbins (2003a). Inhibitory control in rats performing a stop-signal reaction-time task: Effects of lesions of the medial striatum and d-amphetamine. *Behavioral Neuroscience 117*(6), 1302–1317.

Eagle, D. M. and T. W. Robbins (2003b). Lesions of the medial prefrontal cortex or nucleus accumbens core do not impair inhibitory control in rats performing a stop-signal reaction time task. *Behavioural Brain Research 146*(1-2), 131–144.

Epstein, A. N. (1976). Physiology of thirst. *Canadian Journal of Physiology and Pharmacology 54*(5), 639–649.

Evenden, J. (1999a). Impulsivity: a discussion of clinical and experimental findings. *J Psychopharmacol 13*(2), 180–92.

Evenden, J. (1999b). Varieties of impulsivity. *Psychopharmacology (Berl) 146*(4), 348–61.

Evenden, J. L. (1998). The pharmacology of impulsive behaviour in rats ii: The effects of amphetamine, haloperidol, imipramine, chlordiazepoxide and other drugs on fixed consecutive number schedules (fcn 8 and fcn 32). *Psychopharmacology 138*(3-4), 283–294.

Eysenck, S. B. G. and H. J. Eysenck (1977). Place of impulsiveness in a dimensional system of personality description. *British Journal of Social and Clinical Psychology 16*(FEB), 57–68.

Fields, H., G. Hjelmstad, E. Margolis, and S. Nicola (2007). Ventral tegmental area neurons in learned appetitive behavior and positive reinforcement. *Annu Rev Neurosci 30*, 289–316.

Fitzsimons, T. J. and J. Le Magnen (1969). Eating as a regulatory control of drinking in the rat protein carbohydrate diet. *Journal of Comparative and Physiological Psychology 67*(3), 273–283.

Floresco, S. B., S. Ghods-Sharifi, C. Vexelman, and O. Magyar (2006). Dissociable roles for the nucleus accumbens core and shell in regulating set shifting. *Journal of Neuroscience 26*(9), 2449–2457.

Floresco, S. B., R. J. McLaughlin, and D. M. Haluk (2008). Opposing roles for the nucleus accumbens core and shell in cue-induced reinstatement of food-seeking behavior. *Neuroscience 154*(3), 877–884.

Friedman, N. P. and A. Miyake (2004). The relations among inhibition and interference control functions: A latent-variable analysis. *Journal of Experimental Psychology-General 133*(1), 101–135.

Fuchs, R. A., K. A. Evans, M. C. Parker, and R. E. See (2004). Differential involvement of the core and shell subregions of the nucleus accumbens in conditioned cue-induced reinstatement of cocaine seeking in rats. *Psychopharmacology 176*(3-4), 459–465.

Ghazizadeh, A., H. Fields, and F. Ambroggi (2010). Isolating event-related neuronal responses by deconvolution. *J Neurophysiol 104*(3), 1790–802.

Ghitza, U. E., A. T. Fabbricatore, V. Prokopenko, A. P. Pawlak, and M. O. West (2003). Persistent cue-evoked activity of accumbens neurons after prolonged abstinence from self-administered cocaine. *Journal of Neuroscience 23*(19), 7239–7245.

Gill, T. M., P. J. Castaneda, and P. H. Janak (2010). Dissociable roles of the medial prefrontal cortex and nucleus accumbens core in goal-directed actions for differential reward magnitude. *Cerebral Cortex 20*(12), 2884–2899.

Glover, G. (1999). Deconvolution of impulse response in event-related bold fmri. *Neuroimage 9*(4), 416–29.

Groenewegen, H. J., H. W. Berendse, and S. N. Haber (1993). Organization of the output of the ventral striatopallidal system in the rat - ventral pallidal efferents. *Neuroscience 57*(1), 113–142.

Groenewegen, H. J., C. I. Wright, A. V. J. Beijer, and P. Voorn (1999). Convergence and segregation of ventral striatal inputs and outputs. *Advancing from the Ventral Striatum to the Extended Amygdala 877*, 49–63.

Haber, S. N., J. L. Fudge, and N. R. McFarland (2000). Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *Journal of Neuroscience 20*(6), 2369–2382.

Haken, H. (1983). *Synergetics : an introduction : nonequilibrium phase transitions and self-organization in physics, chemistry, and biology* (3rd rev. and enl. ed.). Springer series in synergetics. Berlin ; New York: Springer.

Hall, J., J. A. Parkinson, T. M. Connor, A. Dickinson, and B. J. Everitt (2001). Involvement of the central nucleus of the amygdala and nucleus accumbens core in mediating pavlovian influences on instrumental behaviour. *European Journal of Neuroscience 13*(10), 1984–1992.

Hassani, O. K., H. C. Cromwell, and W. Schultz (2001). Influence of expectation of different rewards on behavior-related neuronal activity in the striatum. *Journal of Neurophysiology 85*(6), 2477–2489.

Hauber, W. and S. Sommer (2009). Prefrontostriatal circuitry regulates effort-related decision making. *Cerebral Cortex 19*(10), 2240–2247.

Haykin, S. S. (1994). *Blind deconvolution.* Prentice Hall information and system sciences series. Englewood Cliffs, N.J.: PTR Prentice Hall.

Heimer, L., D. Zahm, L. Churchill, P. Kalivas, and C. Wohltmann (1991). Specificity in the projection patterns of accumbal core and shell in the rat. *Neuroscience 41*(1), 89–125.

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior 4*(3), 267–272.

Hikosaka, O. (2007). Gabaergic output of the basal ganglia. *160*, 209–26.

Hinrichs, H., M. Scholz, C. Tempelmann, M. Woldorff, A. Dale, and H. Heinze (2000). Deconvolution of event-related fmri responses in fast-rate experimental designs: tracking amplitude variations. *J Cogn Neurosci 12 Suppl 2*, 76–89.

Hollerman, J., L. Tremblay, and W. Schultz (1998). Influence of reward expectation on behavior-related neuronal activity in primate striatum. *J Neurophysiol 80*(2), 947–63.

Hull, C. L. (1943). *Principles of behavior, an introduction to behavior theory.* The Century psychology series. New York,: D. Appleton-Century Company.

Humphries, M. and T. Prescott (2010). The ventral basal ganglia, a selection mechanism at the crossroads of space, strategy, and reward. *Prog Neurobiol 90*(4), 385–417.

Ikemoto, S. (2007). Dopamine reward circuitry: Two projection systems from the ventral midbrain to the nucleus accumbens-olfactory tubercle complex. *Brain Research Reviews 56*(1), 27–78.

Ishikawa, A., F. Ambroggi, S. M. Nicola, and H. L. Fields (2008a). Contributions of the amygdala and medial prefrontal cortex to incentive cue responding. *Neuroscience 155*(3), 573–584.

Ishikawa, A., F. Ambroggi, S. M. Nicola, and H. L. Fields (2008b). Dorsomedial prefrontal cortex contribution to behavioral and nucleus accumbens neuronal responses to incentive cues. *Journal of Neuroscience 28*(19), 5088–5098.

Ito, M. and K. Doya (2009). Validation of decision-making models and analysis of decision variables in the rat basal ganglia. *J Neurosci 29*(31), 9861–74.

Ito, R., T. W. Robbins, and B. J. Everitt (2004). Differential control over cocaine-seeking behavior by nucleus accumbens core and shell. *Nature Neuroscience 7*(4), 389–397.

Izenman, A. J. (1991). Recent developments in nonparametric density-estimation. *Journal of the American Statistical Association 86*(413), 205–224.

Janak, P. H., M. T. Chen, and T. Caulder (2004). Dynamics of neural coding in the accumbens during extinction and reinstatement of rewarded behavior. *Behavioural Brain Research 154*(1), 125–135.

Jentsch, J. and J. Taylor (1999). Impulsivity resulting from frontostriatal dysfunction in drug abuse: implications for the control of behavior by reward-related stimuli. *Psychopharmacology (Berl) 146*(4), 373–90.

Jin, D., N. Fujii, and A. Graybiel (2009). Neural representation of time in cortico-basal ganglia circuits. *Proc Natl Acad Sci U S A 106*(45), 19156–61.

Jones, J. L., J. J. Day, R. A. Wheeler, and R. M. Carelli (2010). The basolateral amygdala differentially regulates conditioned neural responses within the nucleus accumbens core and shell. *Neuroscience 169*(3), 1186–1198.

Jongen-Relo, A. L., S. Kaufmann, and J. Feldon (2002). A differential involvement of the shell and core subterritories of the nucleus accumbens of rats in attentional processes. *Neuroscience 111*(1), 95–109.

Jongenrelo, A. L., P. Voorn, and H. J. Groenewegen (1994). Immunohistochemical characterization of the shell and core territories of the nucleus-accumbens in the rat. *European Journal of Neuroscience 6*(8), 1255–1264.

Kagan, J. (1966). Reflection-impulsivity - generality and dynamics of conceptual tempo. *Journal of Abnormal Psychology 71*(1), 17–&.

Kahneman, D. and A. Tversky (1979). Prospect theory - analysis of decision under risk. *Econometrica 47*(2), 263–291.

Kawaguchi, Y., C. J. Wilson, S. J. Augood, and P. C. Emson (1995). Striatal interneurons - chemical, physiological and morphological characterization. *Trends in Neurosciences 18*(12), 527–535.

Kelley, A. E. and C. J. Swanson (1997). Feeding induced by blockade of ampa and kainate receptors within the ventral striatum: a microinfusion mapping study. *Behavioural Brain Research 89*(1-2), 107–113.

Kelso, J. A. S. (1995). *Dynamic patterns : the self-organization of brain and behavior*. Cambridge, Mass.: MIT Press.

Kennedy, J., R. C. Eberhart, and Y. Shi (2001). *Swarm intelligence*. The Morgan Kaufmann series in evolutionary computation. San Francisco: Morgan Kaufmann Publishers.

Kolb, B. (1984). Functions of the frontal-cortex of the rat - a comparative review. *Brain Research Reviews 8*(1), 65–98.

Koob, G. F. (1996). Drug addiction: The yin and yang of hedonic homeostasis. *Neuron 16*(5), 893–896.

Krause, M., P. W. German, S. A. Taha, and H. L. Fields (2010). A pause in nucleus accumbens neuron firing is required to initiate and maintain feeding. *Journal of Neuroscience 30*(13), 4746–4756.

Kreek, M., D. Nielsen, E. Butelman, and K. LaForge (2005). Genetic influences on impulsivity, risk taking, stress responsivity and vulnerability to drug abuse and addiction. *Nat Neurosci 8*(11), 1450–7.

Lau, B. and P. Glimcher (2008). Value representations in the primate striatum during matching behavior. *Neuron 58*(3), 451–63.

Lauwereyns, J., K. Watanabe, B. Coe, and O. Hikosaka (2002). A neural correlate of response bias in monkey caudate nucleus. *Nature 418*(6896), 413–7.

Logan, G. D. and W. B. Cowan (1984). On the ability to inhibit thought and action - a theory of an act of control. *Psychological Review 91*(3), 295–327.

Logue, A. W. (1988). Research on self-control - an integrating framework. *Behavioral and Brain Sciences 11*(4), 665–678.

Lorenz, K. and P. Leyhausen (1973). *Motivation of human and animal behavior; an ethological view*. Behavioral science series. New York,: Van Nostrand Reinhold Co.

Lu, X. Y., M. B. Ghasemzadeh, and P. W. Kalivas (1998). Expression of d-1 receptor, d-2 receptor, substance p and enkephalin messenger rnas in the neurons projecting from the nucleus accumbens. *Neuroscience 82*(3), 767–780. Times Cited: 110.

Macaluso, E., C. D. Frith, and J. Driver (2000). Modulation of human visual cortex by crossmodal spatial attention. *Science 289*(5482), 1206–1208.

Mallet, N., C. Le Moine, S. Charpier, and F. Gonon (2005). Feedforward inhibition of projection neurons by fast-spiking gaba interneurons in the rat striatum in vivo. *Journal of Neuroscience 25*(15), 3857–3869.

Meredith, G. E., R. Agolia, M. P. M. Arts, H. J. Groenewegen, and D. S. Zahm (1992). Morphological differences between projection neurons of the core and shell in the nucleus-accumbens of the rat. *Neuroscience 50*(1), 149–162.

Milad, M. R. and G. J. Quirk (2002). Stimulation of infralimbic cortex simulates memory for extinction of conditioned fear. *Society for Neuroscience Abstract Viewer and Itinerary Planner 2002*, Abstract No. 83.12.

Miller, E. K. (2000). The prefrontal cortex and cognitive control. *Nature Reviews Neuroscience 1*(1), 59–65.

Miller, N. (1971). selected papers.

Miller, N. and M. Kessen (1952). Reward effects of food via stomach fistula compared with those of food via mouth. *J Comp Physiol Psychol 45*(6), 555–64.

Mowrer, O. H. (1980). Citation classic - learning-theory and behavior. *Current Contents/Social & Behavioral Sciences* (5), 16–16.

Muir, J. L., B. J. Everitt, and T. W. Robbins (1996). The cerebral cortex of the rat and visual attentional function: Dissociable effects of mediofrontal, cingulate, anterior dorsolateral, and parietal cortex lesions on a five-choice serial reaction time task. *Cerebral Cortex 6*(3), 470–481.

Nicola, S., I. Yun, K. Wakabayashi, and H. Fields (2004a). Cue-evoked firing of nucleus accumbens neurons encodes motivational significance during a discriminative stimulus task. *J Neurophysiol 91*(4), 1840–65.

Nicola, S., I. Yun, K. Wakabayashi, and H. Fields (2004b). Firing of nucleus accumbens neurons during the consummatory phase of a discriminative stimulus task depends on previous reward predictive cues. *J Neurophysiol 91*(4), 1866–82.

Nicola, S. M. (2007). The nucleus accumbens as part of a basal ganglia action selection circuit. *Psychopharmacology 191*(3), 521–550.

Nicola, S. M. (2010). The flexible approach hypothesis: Unification of effort and cue-responding hypotheses for the role of nucleus accumbens dopamine in the activation of reward-seeking behavior. *J. Neurosci. 30*(49), 16585–16600.

Nigg, J. T. (2000). On inhibition/disinhibition in developmental psychopathology: Views from cognitive and personality psychology and a working inhibition taxonomy. *Psychological Bulletin 126*(2), 220–246.

Oppenheim, A. V., A. S. Willsky, and S. H. Nawab (1997). *Signals & systems* (2nd ed.). Prentice-Hall signal processing series. Upper Saddle River, N.J.: Prentice Hall.

Papoulis, A. (1977). *Signal analysis*. New York: McGraw-Hill.

Parkinson, J. A., P. J. Willoughby, T. W. Robbins, and B. J. Everitt (2000). Disconnection of the anterior cingulate cortex and nucleus accumbens core impairs pavlovian approach behavior: Further evidence for limbic cortical-ventral striatopallidal systems. *Behavioral Neuroscience 114*(1), 42–63.

Patton, J. H., M. S. Stanford, and E. S. Barratt (1995). Factor structure of the barratt impulsiveness scale. *Journal of Clinical Psychology 51*(6), 768–774.

Paulin, M. G. and L. F. Hoffman (2001). Optimal firing rate estimation. *Neural Networks 14*(6-7), 877–881.

Pecina, S. and K. C. Berridge (2000). Opioid site in nucleus accumbens shell mediates eating and hedonic 'liking' for food: map based on microinjection fos plumes. *Brain Research 863*(1-2), 71–86.

Pennartz, C. M. A., F. H. L. Dasilva, and H. J. Groenewegen (1994). The nucleus-accumbens as a complex of functionally distinct neuronal ensembles - an integration of behavioral, electrophysiological and anatomical data. *Progress in Neurobiology 42*(6), 719–&.

Peters, J., R. T. LaLumiere, and P. W. Kalivas (2008). Infralimbic prefrontal cortex is responsible for inhibiting cocaine seeking in extinguished rats. *Journal of Neuroscience 28*(23), 6046–6053.

Peters, J. and S. Schaal (2008). Natural actor-critic. *Neurocomputing 71*(7-9), 1180–1190.

Peters, J., J. Vallone, K. Laurendi, and P. W. Kalivas (2008). Opposing roles for the ventral prefrontal cortex and the basolateral amygdala on the spontaneous recovery of cocaine-seeking in rats. *Psychopharmacology 197*(2), 319–326.

Pfaffmann, C. (1960). The pleasures of sensation. *Psychological Review 67*(4), 253–268.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology 32*(FEB), 3–25.

Posner, M. I., R. D. Rafal, L. S. Choate, and J. Vaughan (1985). Inhibition of return - neural basis and function. *Cognitive Neuropsychology 2*(3), 211–228.

Pothuizen, H. H. J., A. L. Jongen-Relo, J. Feldon, and B. K. Yee (2005). Double dissociation of the effects of selective nucleus accumbens core and shell lesions on impulsive-choice behaviour and salience learning in rats. *European Journal of Neuroscience 22*(10), 2605–2616.

Quirk, G. J., R. Garcia, and F. Gonzalez-Lima (2006). Prefrontal mechanisms in extinction of conditioned fear. *Biological Psychiatry 60*(4), 337–343.

Ragozzino, M. E., K. E. Ragozzino, S. J. Y. Mizumori, and R. P. Kesner (2002). Role of the dorsomedial striatum in behavioral flexibility for response and visual cue discrimination learning. *Behavioral Neuroscience 116*(1), 105–115.

Reading, P. J. and S. B. Dunnett (1991). Response disinhibition on a delayed matching to position task induced by amphetamine, nicotine and age. *Psychopharmacology 104*(1), 137–139.

Reading, P. J. and S. B. Dunnett (1995). Embryonic striatal grafts reverse the disinhibitory effects of ibotenic acid lesions of the ventral striatum. *Experimental Brain Research 105*(1), 76–86.

Reynolds, S. M. and K. C. Berridge (2002). Positive and negative motivation in nucleus accumbens shell: Bivalent rostrocaudal gradients for gaba-elicited eating, taste "liking"/"disliking" reactions, place preference/avoidance, and fear. *Journal of Neuroscience 22*(16), 7308–7320.

Ringo, J. L. (1991). Neuronal interconnection as a function of brain size. *Brain Behavior and Evolution 38*(1), 1–6.

Robinson, E. S. J., D. M. Eagle, D. Econornidou, D. E. H. Theobald, A. C. Mar, E. R. Murphy, T. W. Robbins, and J. W. Dalley (2009). Behavioural characterisation of high impulsivity on the 5-choice serial reaction time task: Specific deficits in 'waiting' versus 'stopping'. *Behavioural Brain Research 196*(2), 310–316.

Robinson, T. E. and K. C. Berridge (1993). The neural basis of drug craving - an incentive-sensitization theory of addiction. *Brain Research Reviews 18*(3), 247–291.

Rock, I., C. Linnett, P. Grant, and A. Mack (1992). Perception without attention: results of a new method. *Cogn Psychol 24*(4), 502–34.

Roesch, M., T. Singh, P. Brown, S. Mullins, and G. Schoenbaum (2009). Ventral striatal neurons encode the value of the chosen action in rats deciding between differently delayed or sized rewards. *J Neurosci 29*(42), 13365–76.

Roitman, M. F., R. A. Wheeler, and R. M. Carelli (2005). Nucleus accumbens neurons are innately tuned for rewarding and aversive taste stimuli, encode their predictors, and are linked to motor output. *Neuron 45*(4), 587–597.

Rose, J. E. and C. N. Woolsey (1948). Structure and relations of limbic cortex and anterior thalamic nuclei in rabbit and cat. *Journal of Comparative Neurology 89*(3), 279–347.

Rudebeck, P. H., M. E. Walton, A. N. Smyth, D. M. Bannerman, and M. F. S. Rushworth (2006). Separate neural pathways process different decision costs. *Nature Neuroscience 9*, 1161–1168.

Salamone, J. D., M. Correa, A. Farrar, and S. M. Mingote (2007). Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. *Psychopharmacology 191*(3), 461–482.

Samejima, K., Y. Ueda, K. Doya, and M. Kimura (2005). Representation of action-specific reward values in the striatum. *Science 310*(5752), 1337–40.

Schachar, R., G. D. Logan, P. Robaey, S. Chen, A. Ickowicz, and C. Barr (2007). Restraint and cancellation: Multiple inhibition deficits in attention deficit hyperactivity disorder. *Journal of Abnormal Child Psychology 35*(2), 229–238.

Schultz, W., P. Apicella, E. Scarnati, and T. Ljungberg (1992). Neuronal activity in monkey ventral striatum related to the expectation of reward. *J Neurosci 12*(12), 4595–610.

Schultz, W., P. Dayan, and P. R. Montague (1997). A neural substrate of prediction and reward. *Science 275*(5306), 1593–1599.

Scott, D. W. (1979). Optimal and data-based histograms. *Biometrika 66*(3), 605–610.

Sesack, S. R., A. Y. Deutch, R. H. Roth, and B. S. Bunney (1989). Topographical organization of the efferent projections of the medial prefrontal cortex in the rat - an anterograde tract-tracing study with phaseolus-vulgaris leucoagglutinin. *Journal of Comparative Neurology 290*(2), 213–242.

Shimazaki, H. and S. Shinomoto (2007). A method for selecting the bin size of a time histogram. *Neural Computation 19*(6), 1503–1527.

Shinonaga, Y., M. Takada, and N. Mizuno (1994). Topographic organization of collateral projections from the basolateral amygdaloid nucleus to both the prefrontal cortex and nucleus-accumbens in the rat. *Neuroscience 58*(2), 389–397.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review 63*(2), 129–138.

Solomon, R. L. and J. D. Corbit (1974). Opponent-process theory of motivation .1. temporal dynamics of affect. *Psychological Review 81*(2), 119–145.

Steinmetz, P. N., A. Roy, P. J. Fitzgerald, S. S. Hsiao, K. O. Johnson, and E. Niebur (2000). Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature 404*(6774), 187–190.

Stewart, G. W. (1998). *Matrix algorithms*. Philadelphia: Society for Industrial and Applied Mathematics.

Stone, M. (1974). Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society. Series B (Methodological) 36*(2), 111–147.

Stratford, T. R. and A. E. Kelley (1997). Gaba in the nucleus accumbens shell participates in the central regulation of feeding behavior. *Journal of Neuroscience 17*(11), 4434–4440.

Sutton, R. S. and A. G. Barto (1998). *Reinforcement learning : an introduction*. Adaptive computation and machine learning. Cambridge, Mass.: MIT Press.

Swanson, L. W. and W. M. Cowan (1977). Autoradiographic study of organization of efferent connections of hippocampal formation in rat. *Journal of Comparative Neurology 172*(1), 49–84.

Taha, S. and H. Fields (2005). Encoding of palatability and appetitive behaviors by distinct neuronal populations in the nucleus accumbens. *J Neurosci 25*(5), 1193–202.

Taha, S., S. Nicola, and H. Fields (2007). Cue-evoked encoding of movement planning and execution in the rat nucleus accumbens. *J Physiol 584*(Pt 3), 801–18.

Taha, S. A. and H. L. Fields (2006). Inhibitions of nucleus accumbens neurons encode a gating signal for reward-directed behavior. *Journal of Neuroscience 26*(1), 217–222.

Taverna, S., Y. C. van Dongen, H. J. Groenewegen, and C. M. A. Pennartz (2004). Direct physiological evidence for synaptic connectivity between medium-sized spiny neurons in rat nucleus accumbens in situ. *Journal of Neurophysiology 91*(3), 1111–1121.

Tepper, J. M. and J. P. Bolam (2004). Functional diversity and specificity of neostriatal interneurons. *Current Opinion in Neurobiology 14*(6), 685–692.

Totterdell, S. and G. E. Meredith (1997). Topographical organization of projections from the entorhinal cortex to the striatum of the rat. *Neuroscience 78*(3), 715–729.

Tripathi, A., L. Prensa, C. Cebrin, and E. Mengual (2010). Axonal branching patterns of nucleus accumbens neurons in the rat. *J Comp Neurol 518*(22), 4649–73.

Turner, L. H., R. L. Solomon, E. Stellar, and S. N. Wampler (1975). Humoral-factors controlling food-intake in dogs. *Acta Neurobiologiae Experimentalis 35*(5-6), 491–498.

Uylings, H. B. M., H. J. Groenewegen, and B. Kolb (2003). Do rats have a prefrontal cortex? *Behavioural Brain Research 146*(1-2), 3–17.

Voorn, P., L. Vanderschuren, H. J. Groenewegen, T. W. Robbins, and C. M. A. Pennartz (2004). Putting a spin on the dorsal-ventral divide of the striatum. *Trends in Neurosciences 27*(8), 468–474.

Wagner, G. P. and L. Altenberg (1996). Perspective: Complex adaptations and the evolution of evolvability. *Evolution 50*(3), 967–976.

Wanat, M. J., C. M. Kuhnen, and P. E. M. Phillips (2010). Delays conferred by escalating costs modulate dopamine release to rewards but not their predictors. *Journal of Neuroscience 30*(36), 12020–12027.

Winstanley, C. A., D. M. Eagle, and T. W. Robbins (2006). Behavioral models of impulsivity in relation to adhd: Translation between clinical and preclinical studies. *Clinical Psychology Review 26*(4), 379–395.

Woldorff, M. G., C. C. Gallen, S. A. Hampson, S. A. Hillyard, C. Pantev, D. Sobel, and F. E. Bloom (1993). Modulation of early sensory processing in human auditory-cortex during auditory selective attention. *Proceedings of the National Academy of Sciences of the United States of America 90*(18), 8722–8726.

Wolf, S. and H. G. Wolff (1943). Human gastric function: An experimental study of a man and his stomach. *Human gastric function: An experimental study of a man and his stomach*, 195p.

Wright, C. I., A. V. J. Beijer, and H. J. Groenewegen (1996). Basal amygdaloid complex afferents to the rat nucleus accumbens are compartmentally organized. *Journal of Neuroscience 16*(5), 1877–1893.

Yun, I., K. Wakabayashi, H. Fields, and S. Nicola (2004). The ventral tegmental area is required for the behavioral and nucleus accumbens neuronal firing responses to incentive cues. *J Neurosci 24*(12), 2923–33.

Yun, I. A., S. M. Nicola, and H. L. Fields (2004). Contrasting effects of dopamine and glutamate receptor antagonist injection in the nucleus accumbens suggest a neural mechanism underlying cue-evoked goal-directed behavior. *European Journal of Neuroscience 20*(1), 249–263.

Zahm, D. S. (1999). Functional-anatomical implications of the nucleus accumbens core and shell subterritories. *Advancing from the Ventral Striatum to the Extended Amygdala 877*, 113–128.

Zahm, D. S. (2000). An integrative neuroanatomical perspective on some subcortical substrates of adaptive responding with emphasis on the nucleus accumbens. *Neuroscience and Biobehavioral Reviews 24*(1), 85–105.

Zahm, D. S. and J. S. Brog (1992). On the significance of subterritories in the accumbens part of the rat ventral striatum. *Neuroscience 50*(4), 751–767.

Zhou, L., T. Furuta, and T. Kaneko (2003). Chemical organization of projection neurons in the rat accumbens nucleus and olfactory tubercle. *Neuroscience 120*(3), 783–798.

Publishing Agreement

It is the policy of the University to encourage the distribution of all theses, dissertations, and manuscripts. Copies of all UCSF theses, dissertations, and manuscripts will be routed to the library via the Graduate Division. The library will make all theses, dissertations, and manuscripts accessible to the public and will preserve these to the best of their abilities, in perpetuity.

Please sign the following statement:

I hereby grant permission to the Graduate Division of the University of California, San Francisco to release copies of my thesis, dissertation, or manuscript to the Campus Library to provide access and preservation, in whole or in part, in perpetuity.

Author Signature

Apr-25-2011

Date