Striatal dynamics represent subjective time

A psychophysical study of the neural representation of time by striatal populations

Thiago Santos Gouvêa

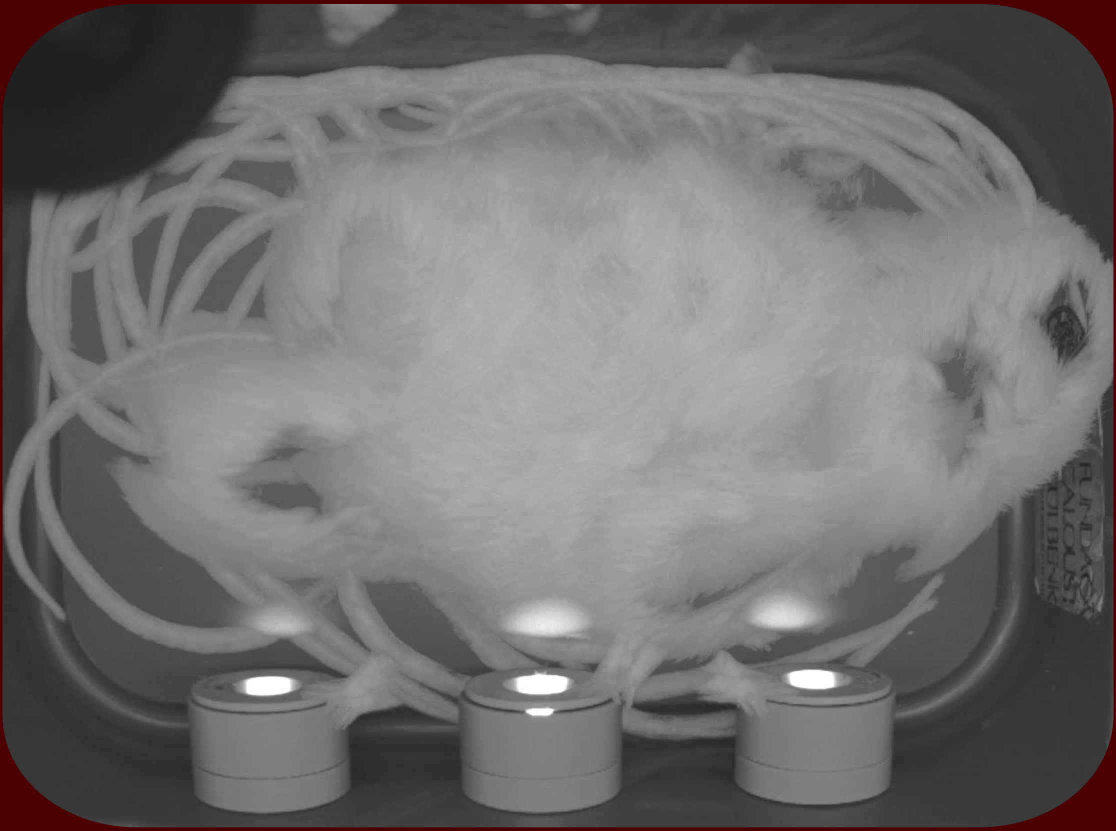

Dissertation presented to obtain the Ph.D degree in Biology/Neuroscience

Instituto de Tecnologia Química e Biológica António Xavier | Universidade Nova de Lisboa

Oeiras,
May, 2016

INSTITUTO
DE TECNOLOGIA
QUÍMICA E BIOLÓGICA
ANTÓNIO XAVIER /UNL

Knowledge Creation

Striatal dynamics represent subjective time

A psychophysical study of the neural representation of time by striatal populations

Thiago Santos Gouvêa

Dissertation presented to obtain the Ph.D degree in Biology/Neuroscience
Instituto de Tecnologia Química e Biológica António Xavier I Universidade Nova de Lisboa

Research work coordinated by:

Champalimaud Foundation

Oeiras, May, 2016

FCT Fundação para a Ciência e a Tecnologia
MINISTÉRIO DA EDUCAÇÃO E CIÊNCIA

INSTITUTO DE TECNOLOGIA QUÍMICA E BIOLÓGICA ANTÓNIO XAVIER /UNL
Knowledge Creation

Cover image Composition displaying, for each pixel, the brightest value over a set of snapshots taken of a rat performing the interval categorization task described in section 2.4.1 on page 35 of this monograph.

STRIATAL DYNAMICS REPRESENT SUBJECTIVE TIME

A PSYCHOPHYSICAL STUDY OF THE NEURAL REPRESENTATION OF TIME BY STRIATAL POPULATIONS

THIAGO SANTOS GOUVÊA

A DISSERTATION PRESENTED TO THE
INSTITUTO DE TECNOLOGIA QUÍMICA E BIOLÓGICA
OF THE UNIVERSIDADE NOVA DE LISBOA
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

JOSEPH J. PATON

2016

que sejas ainda mais vivo
no som do meu estribilho,
tempo,
tempo,
tempo,
tempo

Caetano Veloso

Acknowledgments

This monograph marks the culmination of an unforgettable period.

I am thankful to each of the members of the jury who, by assessing critically the outcome of my work, have honored me with their time and effort. I am thankful in particular to Dean Buonomano and Elliot Ludvig who offer me their fresh perspectives.

I am thankful to my thesis committee, Megan Carey and Christian Machens, who were very generous and helpful whenever besought.

I am thankful to Joe Paton, an outstandingly caring supervisor, and to the Paton Lab, the nest and the family where I seasoned.

I am also thankful to Zach Mainen and the Mainen Lab, for providing the food for thought an eager young PhD student required.

I am thankful to the INDP-2009 class, the partners with whom I took the first steps in this enterprise, and to the larger CNP community for turning a certain indie/synth-pop hit into such a deep mark on me.

I am also thankful to multicultural Lisbon and the Republic of Portugal for having accepted, incorporated, and sponsored me during these formative years.

While often seeking to understand my social environment at the level of collective structures, I ought to express my gratitude to some individuals of outstanding importance [in rough chronological order]: Dr. Dušica Radoš, whose generosity and timely friendship made for a smooth landing; Ali Özgür Argunşah, the Ottoman spirit lifter; Leila Shirai, whose spacious inner universe was a relief to the daily claustrophobia; Joe Paton, who from early on secured his place on this list too; Eric DeWitt, the most generous intellect I've ever encountered; Pietro Vertechi, whose sharp intellect is to be handled carefully; and Nadja Oellrich, who brought a new light to this last stretch, bore with me while I wrote this monograph and, most importantly, was there to say the phrase that made all the difference: “*When I felt like that, I did* not *go sleep.*”

Finally, I feel indebted with my family and friends from home in Brazil who payed the price of distance along with me. To them, in particular, I dedicate the joy of bringing this work to completion.

Título

Dinâmica estriatal representa o tempo subjetivo — Um estudo psicofísico da representação do tempo por populações neuronais no estriado

Resumo

O tempo é uma das dimensões fundamentais do ambiente. A abilidade de estimar a passagem do tempo é essencial tanto para a aprendizagem como para a performance de comportamento adaptativo em situações naturais. Apesar da importância desta função cognitiva, a forma como esta é implementada no cérebro ainda não é bem compreendida.

Estudos mostram que a dinâmica de populações neuronais em diversas áreas cerebrais contém informação sobre a passagem do tempo. Entretanto, não se sabe na maioria dos casos se estes padrões neurais dinâmicos são de fato utilizados na formação dos julgamentos temporais, ou se meramente covariam com a passagem do tempo.

O estriado é uma estrutura dos gânglios da base implicada em várias funções sensíveis à passagem do tempo tais como aprendizagem associativa, tomada de decisão, e formação de estimativas temporais. Com o objetivo de estudar a representação do tempo por populações neuronais do estriado, nós treinamos ratos em uma tarefa psicofísica de discriminação da duração de intervalos. Enquanto os animais desempenhavam a tarefa, nós monitoramos o comportamento com uma câmera de alta velocidade, assim como a atividade de neurônios individuais no estriado.

Os animais desenvolveram sequências comportamentais altamente consistentes durante o intervalo sendo julgado. A variabilidade dessas sequências se mostrou preditiva dos julgamentos comportamentais desde muito antes do momento do julgamento, dando suporte à ideia de que os animais podem estar utilizando padrões comportamentais aprendidos para estimar a duração de intervalos de tempo.

Nós também observamos que a dinâmica da atividade neural de populações estriatais contém informação sobre a passagem do tempo, e que a velocidade com a qual esses padrões neurais progrediram é refletida na expressão comportamental do julgamento acerca da duração do intervalo. Esse resultado não pôde ser explicado pelas sequências comportamentais apresentadas durante o intervalo, sugerindo que a dinâmica da população neuronal estriatal forma uma espécie de relógio interno que está na base da abilidade dos animais de estimar a passagem do tempo.

Por fim, nós desenvolvemos variações da tarefa que permitirão testar se e até que ponto a representação estriatal de julgamentos temporais é ação-específica.

Abstract

Time is a fundamental dimension of the environment. The ability to estimate the passage of time is essential for both learning and performance of adaptive behavior in natural situations. Yet, how this ability is implemented in the brain is poorly understood.

Prior studies have shown that population dynamics in a number of brain areas carry information about the passage of time. However, it is not known whether such time-varying activity patterns inform subjects' judgments of duration or merely covary with elapsing time.

The striatum is an input structure of the basal ganglia implicated in several time-dependent functions such as reinforcement learning, decision making, and interval timing. In order to study the representation of time by striatal neuronal populations, we trained rats in an duration categorization task while continuously monitoring their behavior with a high speed camera, as well as striatal spiking activity.

Animals developed highly reproducible behavioral sequences during the interval being timed. Those sequences were often predictive of perceptual report since early in the trial, providing support to the idea that animals may use learned behavioral patterns to estimate the duration of time intervals.

We also found that the dynamics of neural activity in populations of striatal neurons encode the passage of time, and that the speed with which time-encoding neural activity progressed reflected on the behavioral report of duration judgment. These results could not be explained by ongoing behavior, suggesting that striatal dynamics form an internal "neural population clock" that supports the fundamental ability of animals to judge the passage of time.

Lastly, we developed task variants that will allow testing whether and to what extent the striatal representation of temporal judgments is action-specific.

Author Contributions

Experiments contained in this monograph were designed by Thiago Gouvêa, Tiago Monteiro and Joseph Paton.

Data were acquired by Thiago Gouvêa, Tiago Monteiro, Joseph Paton, and Filipe Rodrigues.

Data were analyzed by Thiago Gouvêa, Tiago Monteiro, Asma Motiwala, Christian Machens, and Joseph Paton, with contributions of Sofia Soares, Serkan Sülün, João Frazão, and Gonçalo Lopes.

Chapters 2 and 3 largely recapitulate the articles reprinted in appendix A (Gouvêa, Monteiro, Soares, Atallah, & Paton, 2014; Gouvêa et al., 2015a, respectively). These articles were written by Thiago Gouvêa, Tiago Monteiro, Asma Motiwala, and Joseph Paton, with contributions of all authors.

Financial Support

The work contained in this monograph was funded by the Portuguese FCT - Fundação para a Ciência e a Tecnologia, Fundação Champalimaud, Fundação Bial, and Simons Foundation.

Contents

List of Tables

List of Figures

Acronyms

2AFC two alternatives forced-choice. 9, 10

AI Artificial Intelligence. 83–85
AIC Akaike information criterion. 32
ANN Artificial Neural Network. 84
AP anterior-posterior. xix, 57–59
auROC area under the ROC curve. 46

BeT Behavioral Theory of Timing. 15, 21
BIC Bayesian information criterion. 32, 33, 38, 63

CSC Complete Serial Compound. 90

DV dorsal-ventral. xix, 57–59

EAB Experimental Analysis of Behavior. 80–83, 89

FI Fixed Interval. 7, 8, 19

ISI inter-stimulus interval. 90
ITI inter-trial interval. 90

LDA linear discriminant analysis. 30, 37, 41, 48, 60, 61
LeT Learning to Time. 15, 16, 21

ML medial-lateral. 59

MTS Multiple Time Scale. 16

PCA principal component analysis. 38, 59

PFC prefrontal cortex. 17, 19

PI Peak-Interval. 8

PSE point of subjective equivalence. 9, 10

PSTH peri-stimulus time histogram. xviii, 45–47, 59

RNN Recurrent Neural Networks. 17–20

ROC receiver operating characteristic. 29, 37, 46, 60

SET Scalar Expectancy Theory. 11, 12

STM Spectral Timing Model. 16

TD Temporal-Difference. 90–92

Contents

Chapter 1

Introduction

1.1 Biological autonomous machines

Intelligent behavior, as displayed by a variety of biological species, emerges from the organization of biological matter. What is special about that form of organization? An answer to this question should describe how do cells, the fundamental units composing biological agents, compose devices able to convert different forms of energy emanating from the environment into usable signals, and how those signals combine to generate appropriate behavioral responses.

1.1.1 Sensors and effectors

A fair amount is known about how animal cells can function as sensors (e.g., photoreceptor cells in the retina, hair cells in the cochlea) or motor effectors (e.g., muscle fibers). Perhaps less satisfactory though still extensive is our understanding about how do physiological signals originated in the sensory periphery organize to give rise to representational systems in the brain. Knowledge currently available in any standard neuroscience textbook (e.g., Kandel, Schwartz, & Jessell, 2000) allows descriptions such as what follows:

Stunt I

> Consider, as an example case, a professional soccer player about to strike a volley kick. The match is being played under abundant light. As

the ball approaches, crossed from the corner arc, photons flying off a light source — either the sun or artificial floodlights — hit the ball and bounce back toward the player's eyes. Beams of photons, focused on his retina by his eyes' lenses, are absorbed by opsin proteins sitting on the membrane of photoreceptor cells that compose the light-sensitive epithelium in the back of his eyes. The absorption of photons by pigment molecules nested inside opsin proteins induces conformational changes that lead to the activation of enzymes that modify the cytoplasmic levels of second-messenger molecules, beginning a signaling cascade. Elements of the signaling cascade alter the conductance of the photoreceptor cells' membrane, generating electric signals that propagate to the axon terminals where they modulate neurotransmitter release at synaptic contacts with other cells composing a feedforward network that ultimately conveys the signal to the brain in a topologically organized manner (as if through *labeled lines*).

At this stage, electromagnetic energy from the environment has been detected by the biological machinery composing his eyes, and the signal has been conveyed to his central nervous system. On its path through different processing stages in the brain, the signal reaches the primary visual cortex in an organized manner — that is to say that different features of the visual scene induce predictable activity patterns in his cortex, thus constituting a neural *representation* of the visual world.

Image copyright holder unknown; fair use according to World Trade Organization (1995, Art. 13^{o}).

Figure 1.1. **A volley kick.** Bebeto strikes a volley kick to score the first goal of Brazil 2 × 0 Argentina during Copa America 1989 (Dieguez, n.d.).

The visual scene happens to be a soccer ball flying over the goal area where an opponent defender deflects it toward our striker. He fixates his gaze on the approaching ball while the goal shows in the periphery of his field of view — as reflected by the neural representation in his visual cortex. Electric signals propagate from there: information flows across the brain, ultimately influencing endogenous neural dynamics in some other brain system from where motor commands are issued. The signal now travels down the axons on his spinal chord via multiple parallel paths targeting different muscle groups. After reaching motor neurons in the ventral horn of the spinal chord, the signal leaves the central nervous system to reach the muscles where acetylcholine release modulates cytosolic levels of calcium inside muscle fibers thus changing the conformation of structural proteins, contracting and relaxing muscles and repositioning bones relative to each other. From the perspective of the camera, the player is adjusting his posture and unleashing the volley in the precise moment the ball comes to reach, his instep hitting the center of the ball in a firm and precise kick. His brain has successfully routed the signals captured by his retina to his musculoskeletal system so as to score that goal.

Stunt II

On another green patch not distant from there, a hungry bird sits on a tree branch. Landing on it was an effortlessly acrobatic move. He has been sitting there poking at the tree holes for many seconds already. It is probably the first time he comes to this particular tree, but he has generally had good luck with trees of such dense foliage in the past. There is a large number of dried fruits, which is not an encouraging sign in itself, but he still expects insects to be relatively abundant here. In fact, given how long he has been exploring this tree, he should expect to have caught some bugs by now. He trusts his senses: up to that point, no prey was there to be detected. At least a full minute has passed. Is it time to acknowledge a minor defeat, spread the wings, and change his bet? Or would that be an impatient and deleterious move?

The bird must make a binary decision faced with dangers at both flanks — i.e., it risks amassing insufficient energy to survive either by exploring under-supplied patches for too long, or by not staying long enough on reasonably supplied ones thus spending too much time flying between branches. However, in order to make this decision he can't rely solely on stimuli that are presently available to his senses. Instead, he needs to take into account estimates of variables such as the probability of capturing preys per unit time spent on a given tree.

The neural representations and computations supporting this temporal decision could not be found in the reference textbook.

1.1.2 Intelligent behavior is adaptive

Finely sensing the environment and performing dexterous motion acts are indeed remarkable achievements of some biological forms. These abilities, however, do not suffice to explain the appearance of intelligence conveyed by their behavior. Instead, it is the appropriateness of behavior to different situations that compels the observer to infer intelligence — had the striker invested his dexterity into scoring an own go, fans wouldn't deem him as much respect. In fact, intentionally emitting unfit behaviors is at the heart of certain types of humor, supposedly for violating expectations about intelligence. Similarly, the bird must deploy its flying apparatus at the appropriate occasions. Beyond producing actions that are biomechanically fit, agents need to be able to identify what behaviors are appropriate, and when to emit them.

How does an agent learn to emit behaviors in their appropriate contexts? They certainly must be able to discern among contexts — e.g., while the bird takes note of the density of the tree crown and the state of its fruits, the striker should attend to the color of the uniform of other players to distinguish team mates from opponents, and to how they move and where they stand in space to know which goal to attack and which one to defend. But figuring out what aspects of a context are relevant in guiding behavior is a difficult problem the animal has to solve. The space in which the solution is contained should be formed by several features of the environment that must include sensory stimuli, spatial locations, and time. Furthermore, after assessing its environment the agent will need to decide among the behaviors available in its repertoire. Ultimately, the decision should be based on estimates of the consequences of the different behaviors, conditional on context. In other words, agents should take note of the co-occurrences of contexts, behaviors, and desirable outcomes, and use that information to select the behaviors that, given the present circumstances, will yield the most desirable outcome.

1.1.3 Temporal structure holds valuable information

In order to decide what behaviors to emit on a given context, a subject ought to I) characterize the context, and II) estimate the consequences of the different behaviors in its repertoire. Time plays a double role in this process. On the one hand, it is one of the features defining context; on the other hand, it forms a dimension along which events organize in ways that allow the inference of causal links between behavioral responses and their consequences. This topic is expanded in chapter 5 (see section 5.3 on page 89).

1.2 Interval timing behavior

Behavior is sensitive to temporal regularities in the environment. Temporal features of the environment across different time scales are of behavioral importance. The relevance of timing can be seen, for example, on the dynamics of sensory events (e.g., speech recognition requires fine discrimination of subsecond fluctuations in sound), the dynamics of motor behavior (as in complex movements composed of a well timed sequence of simpler actions), and in the way behaviors are distributed over larger time scales (e.g., foraging behavior).

Behaviors relying on estimation of time in the range of seconds are collectively known as *interval timing*, and a number of behavioral paradigms have been employed to study it. While interval timing paradigms have been summarized elsewhere (Grondin, 2010; Merchant & de Lafuente, 2014), a non-exhaustive list containing some popular ones is presented next:

1.2.1 Interval production

Interval production tasks are those in which the behavioral correlate of the subjective estimate of interval duration consists of the timing of a self paced, observable response, relative to the occurrence of a reference event (e.g., the last reward). These tasks are referred to as *prospective timing* tasks.

Falling within this class of tasks, operant training under the Fixed Interval (FI) reinforcement schedule is among the most well established procedures for studying interval timing (e.g., Skinner, 1938; Ferster & Skinner, 1957; Cumming & Schoenfeld, 1958; Treisman, 1963; Mello, Soares, & Paton, 2015). The procedure consists of reinforcing a behavioral response (usually key-pecking for birds, or lever-pressing for rodents) contingent on the elapse of an interval of fixed duration relative to the last reinforcement event.

Under this contingency, animals show temporal discrimination — i.e., they tend to emit the behavioral response preferentially during a specific time window that ends with delivery of the reinforcer, even in the absence of any external contextual cues signaling the occasion. Specifically, animals tend to display a low rate of responding during the beginning of the FI, and change abruptly to a high response rate at around two-thirds of the interval — it is the so called *break-and-run* pattern (Cumming & Schoenfeld, 1958; Schneider, 1969). The breakpoint at which response rate increases is taken to indicate the animal's estimate of proximity to the end of the FI.

In a variant known as the Peak-Interval (PI) procedure, the reinforcer is omitted in a small fraction of trials, so that a second breakpoint occurs in which responding returns to a low rate (Meck & Buhushi, 2010).

These tasks are attractive due to the simplicity of the contingencies employed. However, animals performing them do not have a strong incentive to be precise — by responding at times distant from the reinforced occasion, subjects incur no cost other than that of executing the response itself, which is supposedly low compared to the cost of missing an available reward. As a consequence, the link between time estimation ability and observable behavior is weakened.

1.2.2 Interval reproduction

Interval reproduction tasks combine the prospective component of production tasks with a *retrospective* component: the timing of the behavioral response to be produced by the subject must match the duration of an interval presented by the experimenter (e.g., Treisman, 1963; Michon, 1967; Wing & Kristofferson, 1973; Pastor, Artieda, Jahanshahi, & Obeso, 1992; Jazayeri & Shadlen, 2010, 2015).

This type of task has the advantage over production tasks that intervals of different duration can be elicited flexibly across trials. Perhaps due to the cognitive flexibility demanded, these tasks most commonly employ primates as subjects. One downside of these tasks is the difficulty in teasing apart behavioral variability between the retrospective and prospective components — i.e., variability in estimating/remembering the reference interval, as opposed to variability in producing the desired interval.

1.2.3 Interval discrimination

Discrimination refers to the process through which a behavioral response is maintained by reinforcers contingent on properties of both the response itself *and* the context in which it occurs (Skinner, 1965; Catania, 1999). In other words, a subject is said to discriminate between contexts when the probability of emitting a given behavioral response is different for different contexts, as shaped by its context-specific reinforcement history. In this case, contextual stimuli are said to *set the occasion* for that behavioral response, and behavior is said to be under stimulus control (Skinner, 1965; Catania, 1999).

In interval discrimination tasks, the contextual trait setting the occasion for different behaviors is *time* — i.e., the duration of a stimulus interval that has been presented to the subject prior to the behavioral response. Usually, these tasks are set to reinforce two distinct behavioral responses differentially and non-probabilistically (e.g., Stubbs, 1968; Platt & Davis, 1983; Leon & Shadlen, 2003; Balcı et al., 2008; Gouvêa et al., 2014), and are therefore described as binary-choice, or two alternatives forced-choice (2AFC) tasks. That is the case for all tasks presented in this section.

Interval discrimination tasks are purely retrospective timing tasks in the sense that, unlike in prospective timing tasks, the timing of the emitted behavior is irrelevant — with the trivial exception that some deadline for responding is often imposed. Instead, a given behavioral response will or will not be reinforced contingent solely on the duration of the interval that preceded it.

Temporal bisection

In temporal bisection tasks (Stubbs, 1976; Church & Deluty, 1977; Platt & Davis, 1983; Maricq & Church, 1983), two distinct stimulus intervals — one short and one long — set the occasions in which either of the two behavioral responses is reinforced (i.e., choosing to emit one of the responses is reinforced following presentation of the short stimulus, while the alternative choice is reinforced following presentation of the long stimulus). Importantly, subjects are not allowed to respond before stimulus termination. The behavioral correlate of interval estimation is the choice probability conditional on stimulus, and generally subjects will choose the reinforced option in a nearly deterministic manner.

A distinctive trait of temporal bisection tasks is that, on a fraction of trials, stimuli of intermediate duration are presented although they do not alter reward probability (e.g., choice is never rewarded, or is rewarded randomly, upon presentation of intermediate stimuli). In these cases, it is often observed that the stimulus duration eliciting the two choices equiprobably (as calculated by interpolation) lies close to the geometric mean between the two extreme, reinforced stimuli (Stubbs, 1976; Church & Deluty, 1977; Platt & Davis, 1983). This specific stimulus duration is termed the point of subjective equivalence (PSE).

The unconstrained contingencies for intermediate stimuli do not incentivize subjects to discriminate time to the best of their ability, and therefore are not well suited for study of the precision of temporal perceptual abilities. Instead, the interest elicited by temporal bisection tasks comes from observing where do subjects place their PSE spontaneously, or how is that affected by neurobiological manipulations (e.g., Maricq & Church, 1983). A phenomenon of particular interest is the regularity of the behavior that emerges spontaneously under such unconstrained contingencies — the fact that the PSE is reproducibly well approximated by the geometric mean of the reinforced stimuli is often taken as evidence that subjective time evolves in a relative scale, in agreement with the scalar-timing hypothesis (Gibbon, 1977).

Switching task

The switching task is a simplified variant of the bisection task in which only extreme, reinforced stimuli are presented, and animals are allowed to respond during the stimulus period instead of being forced to wait for stimulus completion (Balcı et al., 2008). In this case, the behavioral correlate of interval estimation is the time at which the subject switches away from the short response given presentation of a long stimulus. One advantage of the switching task over the bisection and categorization tasks is that it allows a read out of the PSE as a continuous variable on single trials.

Interval categorization

Binary categorization tasks, often referred to as simply 2AFC tasks within the sensory neuroscience community, have been widely employed in the study of perceptual decision making (Parker & Newsome, 1998). In these tasks, stimuli vary along some unique dimension — e.g., vibration frequency of a somatosessnory stimulus (Mountcastle, Talbot, Sakata, & Hyvarinen, 1969; Romo, Brody, Hernández, & Lemus, 1999); relative concentration of a binary mixture (Uchida & Mainen, 2003); motion coherence in a random-dot kinetogram (Salzman, Britten, & Newsome, 1990; Britten, Shadlen, Newsome, & Movshon, 1992; Britten, Newsome, Shadlen, Celebrini, & Movshon, 1996) — and reinforcement of each of the two available responses is contingent on whether the stimulus lies above or below a reference value along the variable dimension. The reference value can be presented explicitly as a comparison stimulus (e.g., Creelman, 1962; Mountcastle et al., 1969; Romo et al., 1999), or as an implicit constant rule (e.g., Salzman et al., 1990; Uchida & Mainen, 2003; Znamenskiy & Zador, 2013)

Because in this paradigm every stimulus presentation sets the occasion for discrimination, animals are incentivized to utilize sensory information to the best of their ability, and failure in emitting the reinforced response is taken to indicate that the subject is operating in the vicinity of its perceptual limits.

Within the interval timing literature, interval categorization tasks have been employed with different model organisms such as pigeons (Stubbs, 1968), human- (Creelman, 1962), and non-human primates (Leon & Shadlen, 2003). In the current monograph, this paradigm is employed to study the neural representation of temporal judgments in the rodent striatum. A description of the task can be found in section 2.4.1 of this volume, as well as in Gouvêa et al. (2014). Additionally, two new variants of this task are introduced in chapter 4.

1.3 Mathematical models of timing behavior

The array of behavioral paradigms listed above generates a large amount of data on interval timing behavior. A number of mathematical models exist that attempt to quantitatively capture it, some of which are listed below:

1.3.1 Scalar Expectancy Theory

As noted by Gibbon (1977), a very conspicuous regularity found across a large number of studies is that the variability of time estimates, as measured by the standard deviation, scales linearly with the mean duration of the estimated interval across a large range of interval durations. In other words, the coefficient of variation (= standard deviation / mean) is constant across a range of durations. This phenomenon is known as the *scalar property* of interval timing, and constitutes the cornerstone of his Scalar Expectancy Theory (SET).

In that same work, Gibbon proposed that interval timing behavior comes about by a *comparison* mechanism based on the *ratio* between the inverse of the estimated time until reinforcement ("an instantaneous rate-of-reinforcement estimate", Gibbon, 1977, p. 281), and the baseline rate of reinforcement estimated over an entire session.

Of crucial importance in this proposition is the role played by the manner by which the time until reinforcement is estimated. While refraining from proposing any mechanistic or otherwise structural hypothesis about how time estimates are produced (SET is "a theory properly described as a discrimination theory of temporal control", Gibbon, 1977, p. 281), the author does put to test four different models that generate estimates x

under different reference intervals T — elicitation (x identically distributed irrespective of T), absolute timing (x distribution is centered on T and has a fixed shape), Poisson timing (x is Poisson distributed with mean and variance T), and scalar timing (x is normally distributed with mean T, and standard deviation γT where γ is the constant coefficient of variation). As can be noted already from their descriptions, only the latter generative model is consistent with the scalar property.

Once scalar timing has been established as the only reasonable form of generating estimates of time until reinforcement, the "expectancy ratio comparator" decision mechanism is then shown to fit well different behavioral paradigms. Behavior under temporal production contingencies are postulated to be generated by undergoing a break-and-run transition whenever the ratio between instantaneous and baseline reinforcement rates (a.k.a. *expectancies*) cross a threshold that is arbitrary but fixed. Interval categorization, on the other hand, is implemented by comparing the ratio between the expectancies estimated assuming the stimulus belongs to either of the two categories.

While scalar timing coupled with expectancy based decision making mechanims can describe a wide range of behavioral results, an attempt to link SET to an algorithmic implementation of interval timing behavior was only made later, when Gibbon proposed what he calls an information processing model of interval timing (first introduced in Gibbon & Church, 1984). This model is introduced below.

1.3.2 Pacemaker-Accumulator models

Pacemaker-accumulator models are perhaps the most widely used models to explain behavioral data on a number of paradigms in the interval timing field. Their origins, however, point to a different class of temporally structured biological phenomena.

As noted by Crozier and Stier (1925) for a number of rhythmic neuromuscular phenomena in arthropods, frequency increases monotonically with temperature — from the velocity of progression in ants to the frequency of chirping in crickets, abdominal respiratory movements in dragonfly larvae, and even flashing in fireflies. Taking the known relation between rate of chemical reactions and temperature to a surprising limit, the authors apply the Arrhenius equation[1] to fit the rhythmic neuromuscular arthropod behaviors.

They found a surprising level of similarity in the single free parameter of the model, the temperature characteristic μ, across behaviors and species, suggestive of a shared

---

[1] The Arrhenius equation describes the monotonic relation between temperature and the rate of a chemical reaction. It can be formulated as $K_2 = K_1 \exp\left[\mu\left(\frac{1}{T_1} - \frac{1}{T_2}\right)\right]$, where K_i, T_i are the rate and temperature of process i, and μ is the temperature characteristic, a free parameter.

catalyst (Rice, 1923 apud Crozier & Stier, 1925). On this regard, Crozier and Stier (1925, p. 429) put forth the following observation:

> The surprising degree of quantitative concordance evident in this list led to the assumption that not only do such rhythmic activities reflect the determination of the frequency of neuromuscular movements through the agency of chemical transformations, but that the *governing* chemical process must be of identical type in the several instances — at least in the sense that there is involved a common catalyst. The conclusion is the more permissible because the activities compared may be taken to involve "central nervous discharge" as controlling element, and are thus truly homologous.

The idea that the rate of discharge of processes taking place in the central nervous system controls the speed of temporally structured behavior extends from that day to the present. Whether temperature can be a significant determining factor in modulating interval timing, however, was yet to be tested. Hoagland took that step after the following reasoning (Hoagland, 1933, p. 268):

> If our judgments of time depend upon an underlying chemical "master reaction," in, let us say, cells of the brain, modification of the internal body temperature might be expected to alter judgments of time intervals in a way consistent with the Arrhenius equation and might yield a significant temperature characteristic. Such a finding might even imply a specifically catalyzed, irreversible chemical mechanism controlling the consciousness of duration.

Using as subjects his feverish, influenza-infected wife and a student who "agreed to submit himself to diathermy treatment"[2] (Hoagland, 1933, p. 270), Hoagland found that the relation between body temperature and behavior in an interval production task are well fit by the Arrhenius equation, albeit with higher cross-subject variability in temperature characteristic.

Taking Hoagland's suggestion that subjective time might arise from a rhythm generating "chemical clock" that is susceptible to global variables such as temperature, Treisman (1963) proposed an algorithmic model "devised in an attempt to explain and relate the psychophysical findings" (p. 18) of a series of interval timing experiments presented in the same article. In his own words (Treisman, 1963, p.19):

---

[2] "The subject, insulated from heat loss by wrapping, was exposed to a high frequency alternating current until his temperature (by mouth) registered 38.8"C" (Hoagland, 1933, p. 270)

> Hoagland (1933) compared his "chemical clock" to pacemaker neurons. These produce impulses at rates which may vary and transmit them along axons whose speeds of conduction are relatively fixed. Analogous assumptions are made here, but no hypotheses about the neural identity of the components of the model are intended

Treisman's model is composed of a pulse generator (pacemaker), a counter, a storage (memory) mechanism, and a comparator (decision mechanism). Additionally, an arousal center controls the pace of pulses. The pulse generator has a basic interpulse interval with mean t_0 and standard deviation σ_t. Pulses are emitted with a mean interpulse interval t' defined as

$$t' = t_0 q \tag{1.1}$$

, where q is the arousal level. During an interval to be estimated, the counter keeps track of the number of pulses emitted. The pulse count can be transferred to the storage, or it can be compared to a stored value to produce a decision (e.g., increase the rate of responding in an interval production task). Additionally, several noise terms are introduced at different stages of the process.

The impact of the different noise terms on variability in behavioral output is the focus of the article in which Gibbon and Church (1984) introduce their information processing model of interval timing. Briefly, the authors show that, in order for the output of the pacemaker-accumulator model to display the scalar property, variability needs to be postulated at the level of two variables that get multiplied — within trial-interpulse intervals and cross-trial pulse rate (t_0 and q in equation 1.1, respectively). In the absence of slow, cross-trial noise, variance increases with the square of the mean of the estimated interval, implying that the coefficient of variation is gets smaller for longer intervals — in other words, estimates of longer intervals should be relatively more precise than that of short ones, in conflict with empirical knowledge (Gibbon, 1977).

Pacemaker-accumulator models propose multiplicative variance as a way of producing the scalar property. However, as pointed out by Beck, Kanitscheider, and Pouget (2012), the multiplicative variance proposed by models such as the one at hand does not imply, as Gibbon and Church (1984) suggest, addition of noise at different processing stages. Rather, multiplicative variance in the pacemaker-accumulator model can be interpreted as correlated fluctuations on the components of the pulse generating process as imposed by some global variable. According to this interpretation, slow fluctuations in the global variables would account for the cross-trial variability in clock speed postulated by the models. In fact, such global variables that can influence the pace of the internal clock were proposed by both Hoagland and Treisman — respectively temperature and arousal, although to what extent the latter can be independently verified is arguable.

1.3.3 Behavioral-state based models

Under environments with strong temporal regularities, animal behavior is known to become temporally structured (e.g., Skinner, 1948; Hodos, Ross, & Brady, 1962; Anderson & Shettleworth, 1977; Haight & Killeen, 1991; Machado & Keen, 2003; Balcı et al., 2008; Ölveczky, 2011). Stemming from this observation, some authors have posited that animals might estimate the passage of time directly from transitioning through a series of behavioral states — the data presented in chapter 2 of this monograph speaks to this idea.

Behavioral Theory of Timing

Observing that under interval timing tasks animals tend to display reproducible sequences of interim behaviors, Killeen and Fetterman (1988) noted that these behaviors could themselves serve as discriminative stimuli. The authors then propose a simple model, known as Behavioral Theory of Timing (BeT), in which each interim behavior is associated with a different state. The transition between states is postulated to be triggered by the emission of a pulse from a Poisson pulse emitter. The rate of the pulse emitter is controlled by the rate of reinforcement, the fluctuation of which introduces the multiplicative variance that guarantees the scalar property — making BeT formally very similar to the pacemaker-accumulator models of Treisman (1963) and Gibbon and Church (1984).

Learning to Time

An alternative model based on the same idea that sequential interim behaviors can serve as discriminative stimuli was proposed by Machado (1997). This model, known as Learning to Time (LeT), is reminiscent of feed-forward neural networks. In fact, it is defined in terms of linear differential equations describing the change in activation level of a given state as a function of the activation levels of both the same and preceding states. In addition, an emphasis on learning is expressed by the presence of a variable that "[describes] how the behavioral states come to regulate the operant response" (Machado, 1997, p. 243).

The learned variable in LeT is, effectively, a linear readout vector, and the equations governing readout learning are reminiscent of methods applied to artificial neural networks — namely, minimization by gradient descent of a convex error function such as the sum of squared errors.

1.3.4 Interval timing as a combination of time-series

The idea of estimating the duration of an interval by combining (e.g., linearly summing) a number of different time-varying processes is shared between LeT and several other models, some of which preceded it. This idea is akin to the mathematical notion of *basis functions*: a new function (e.g., the degree of belief that a given time point matches true current time, as a function of a range of time points) can be expressed as a sum of pre-existing functions. Alternatively, the different time series can be interpreted as forming a feature space well suited to represent/approximate arbitrary temporal patterns.

These models differ on how the time-varying processes (or basis functions) are generated. In the Multiple Time Scale (MTS) model proposed by Staddon and Higa (1999) the basis functions are monotonically decreasing functions that, after being triggered by an input stimulus, decay at different rates. Similarly, the basis time-series in the Spectral Timing Model (STM) of Grossberg and Schmajuk (1989) are produced by providing a range of time scales, but their form and nature differ: they are not functions, but the time evolution of a set of dynamic variables governed by coupled differential equations that result in a set of progressively widening unimodal patterns. Time series of similar shape are proposed by Ludvig, Sutton, and Kehoe (2008), though in this work the time series are defined by functions describing the occupancy of linearly spaced temporal receptive fields by an exponentially decaying trace. In the model presented by Shankar and Howard (2012), the time series are generated by leaky integrators with different leak rates that effectively perform a Laplace transform of the stimulus history; importantly, the transformation can be (approximately) inverted by a linear read out unit.

1.3.5 Coincidence detection over oscillatory processes

The present class of models explores a different form of reading out time from the combination of a set of time-series: by detecting phase coincidences over a set of sinusoidal oscillations.

The brain is composed of neurons whose activity is known to be short-lived, i.e., following a disturbance, activity in single neurons reaches equilibrium following a few milliseconds. Any system that needs to generate reliable dynamics on a time scale much longer than that of its components, as is the case in the brain, is faced with a difficult problem. Miall (1989) noted that a solution could lie in the combinatorics of oscillatory processes: if a number of sine wave oscillators with slightly different periods are triggered in phase synchrony, the time it will take until the initial phases synchronize again is much longer than the period of the average oscillator[3]. Miall proposes that real neurons in the brain could implement a clock with time scales much longer than that of single neurons if neurons could function as oscillators. Bearing on the literature on the neurobiology of timing, Matell and Meck (2004) proposed that the oscillatory processes could be implemented by neurons in the cortex, while the oscillation detectors could lie in the striatum.

1.3.6 Time as intrinsic property of network dynamics

On a largely parallel historical thread, artificial neural networks have been explored as useful computational architectures loosely inspired in biological brains. For a long time the field was dominated by feed-forward architectures (Palnitkar & Cannady, 2004), perhaps due to the success of the *perceptron* (Rosenblatt, 1958) as both a cognitive theory of discrimination and a classification tool with applications in artificial intelligence. In these types of network, neurons are organized in layers, and connections are established exclusively unidirectionally between neurons in adjacent layers. This results in a process in which time is not a relevant dimension, as the computation of interest (e.g., the categorization of a static input pattern) is often performed on a single feed-forward sweep through the network.

An alternative architecture of interest is that of Recurrent Neural Networks (RNN). Unlike feedforward networks, in RNN any neuron can be connected to any other. As a result, activity reverberates giving rise to interesting dynamics that are in closer resemblance with neurobiological patterns. RNNs displaying attractor dynamics have been proposed to implement the *persistent activity* pattern thought to be used by neural circuits in the prefrontal cortex (PFC) to implement working memory (Wang, 2001). While useful to store the memory of a static input for an extended period of time after termination of stimulus presentation, persistent activity patterns lack the dynamics necessary to encode information about the passage of time.

Maass, Natschläger, and Markram (2002) proposed that a RNN could also be used to encode the memory of a stimulus in a dynamic way so that both stimulus identity and

---

[3] A related strategy has been suggested to be implemented by hippocampal grid cells in supporting the formation of a cognitive map (Tran & Fiete, 2015)

timing could be inferred. As in attractor networks, memory of stimuli persists in RNNs displaying dynamic trajectories for a period longer than the time constant governing the dynamics of individual neurons. In line with this observation, it has been noted that the dynamics of RNNs could also be used to solve the interval timing problem (Buonomano & Merzenich, 1995; Buonomano & Maass, 2009; Buonomano, 2014; Goel & Buonomano, 2014). On this regard, the activity of single neurons composing a RNN can be seen as separate time series that can be combined to represent/approximate arbitrary temporal patterns, as in the models introduced in section 1.3.4.

Recent work showing RNNs can be trained to support temporal warping (Goudar & Buonomano, 2015) suggests a phenomenon analogous to multiplicative variance, responsible for the scalar property in pacemaker-accumulator models, might emerge from RNN dynamics.

1.4 Encoding time with neurons: a brief review

A representation of time can be implemented by artificial neural networks such as RNNs. But how is the passage of time represented in the activity of biological brains? Insights into the neurobiology of interval timing is provided by a number of studies ranging from lesion and pharmacology to direct observation of neural activity under temporally regular contingencies. Potentially time-encoding neural activity has been observed across a number of brain areas — although evidence of their use in informing behavior is often missing. A number of these studies addressing the question of how does the animal nervous system keep track of time are presented next.

1.4.1 The basal ganglia

The dopaminergic system Patients of Parkinson Disease underestimate time intervals, as assessed on interval reproduction and verbal estimation (Pastor et al., 1992; Malapani et al., 1998) tasks, suggesting time as encoded by the nervous system runs slower under chronically low dopamine levels. Additionally, when trained to reproduce two different intervals in blocks of trials, Parkinson patients off medication display a "regression toward the mean" effect (i.e., they show overproduction of short-, and underproduction of long intervals, Malapani et al., 1998), consistent with a previously described impairment in interval discrimination (Artieda, Pastor, Lacruz, & Obeso, 1992). In line with these observations, the administration of dopaminergic agonists has been found to produce effects on interval timing behavior consistent with an increase on the speed of subjective time estimation (e.g., Maricq & Church, 1983; Balcı et al., 2008).

Selectively lesioning dopaminergic cells in ventral midbrain abolished completely the temporal structure in rate of responding in an interval production task (Meck, 2006). The effect could be reproduced by lesioning only the cell terminals in striatum, suggesting the effects of dopaminergic manipulations on interval timing are mediated by computations taking place in the striatum.

Striatum Striatal activation has been reported in humans performing interval timing task (Hinton & Meck, 2004; Harrington, Zimbelman, Hinton, & Rao, 2009), thought the content of striatal activity cannot be inferred from these studies.

Spiking activity recorded in the striatum of monkeys performing tasks containing a delay period of fixed duration revealed groups of neurons becoming active at different points in time realtive to delay onset, thus allowing experimenters to decode time from neural activity (Jin, Fujii, & Graybiel, 2009; Adler et al., 2012), In both these studies, however, activity was recorded during performance of tasks for which time estimation was irrelevant given the reinforcement contingencies. Therefore, whether the observed neural dynamics constitute a representation of time used to guide behavior or they merely covary with the passage of time cannot be assessed.

A comparable pattern of activity was found in the rodent striatum (Mello et al., 2015) during an interval production task. This study presents the advantage of having being recorded under contingencies for which time matters — animals were trained to lever press under FI reinforcement schedules, the duration of the interval varying over blocks of tens of trials. Interestingly, the authors found that the time window tiled by the firing of striatal neurons re-scaled to match the duration of the FI, and that, during block transitions, the number of trials needed for behavior to adjust was matched by the time course of neural adaptation, suggesting this neural code for time might be used to inform behavior.

1.4.2 Cortices

Prefrontal cortex During performance of working memory tasks, the PFC is thought to keep memories by means of its recurrent connectivity (Wang, 2001). Networks of similar architecture (i.e., RNNs) are capable of encoding the passage of time, thus suggesting PFC might be implicated in interval timing. In fact, a number of studies have provided corroborating evidence. Reversible inactivation of PFC caused a significant impairment in interval discrimination performance (Kim, Jung, Byun, Jo,

& Jung, 2009), and time encoding neural activity have been observed both in primates (Machens, Romo, & Brody, 2010) and rodents (Kim, Ghim, Lee, & Jung, 2013; Xu, Zhang, Dan, & Poo, 2014).

Parietal cortex The parietal cortex has been implicated in functions such as category representation (Freedman & Assad, 2006), working memory (Harvey, Coen, & Tank, 2012), and accumulation of sensory evidence (Shadlen & Newsome, 1996, 2001; Hanks et al., 2015).

In the working memory task, the parietal cortex has been shown to display sequential activity patterns that encode both stimulus identity and timing (Harvey et al., 2012), precisely as in the RNNs mentioned above. Furthermore, its involvement in interval timing has been probed in a number of different paradigms. On an interval discrimination task, parietal activity reflected the category the monkeys were most likely to chose at any moment in time — irrespective of whether the choice was correct or incorrect (Leon & Shadlen, 2003). This can be interpreted as meaning that parietal activity makes the same errors as the animal, suggesting it might be guiding choice. On an anticipation task (Janssen & Shadlen, 2005), parietal activity was shown to reflect readiness to respond, perhaps in agreement with the interval discrimination results. Lastly, during an interval reproduction task, parietal activity showed a ramp with constant ramp during presentation of the sample interval; it then returned to baseline, and ramped again, now at a rate inversely proportional to the duration of the sample interval, in such a way that activity would be always at a similar level at the time of the reproduction response (Jazayeri & Shadlen, 2015).

1.4.3 Hippocampus

The hippocampus is known for its importance in spatial navigation (O'Keefe & Nadel, 1978; Hafting, Fyhn, Molden, Moser, & Moser, 2005) and episodic memory (Scoville & Milner, 1957). More recently, it has been demonstrated that, during tasks involving a delay period of predictable duration, the hippocampus also exhibits sequential, potentially time encoding activity patterns (Pastalkova, Itskov, Amarasingham, & Buzsáki, 2008; MacDonald, Lepage, Eden, & Eichenbaum, 2011). Whether these activity patterns are used to guide interval timing behavior, however, is yet unclear.

1.5 Our addition

The literature reviewed above reveals a wealth of knowledge about regularities in interval timing behavior, as well as in neural patterns in different brain areas that lend themselves as candidate implementations of timing abilities.

By combining highly sensitive behavioral paradigms and neurobiological recording methods, the present work aims at constituting a relevant contribution toward a satisfactory mechanistic understanding of interval timing behavior.

In chapter 2 we introduce the behavioral paradigm employed — an interval discrimination task in line with the psychophysical tradition (Parker & Newsome, 1998). By monitoring behavior during task performance with high-speed video, we directly test — perhaps for the first time in a psychophysical task — the prediction made by BeT and LeT theories that subjective time estimation should correlate with the dynamics of behavioral sequences. The work presented in this chapter has been published (Gouvêa et al., 2014). A copy of the published article can be found as an appendix in section A.1 on page 95.

The basal ganglia has been implicated in interval timing by a number of clinical, pharmacological, lesion, and physiological studies. However, evidence directly linking striatal activity and subjective time estimation is missing. In chapter 3 we establish a link between striatal activity at the single cell and population levels, and perceptual decisions on a psychophysical task. Additionally, we employ video recordings to explore the relation between the striatal representation of time and behavioral sequences. The work presented in this chapter has been published (Gouvêa et al., 2015a), and a copy of the published article can be found as an appendix in section A.2 on page 107.

Lastly, in chapter 4 a new interval discrimination behavioral paradigm designed to probe the action-specificity of the striatal representation of temporal categories is presented. A general discussion is presented in chapter 5.

Chapter 2

Structured behavior predicts duration judgments

Under temporally structured reinforcement schedules, animals tend to develop temporally structured behavior, and interval timing has been suggested to be accomplished by learning sequences of behavioral states. If this is true, trial to trial fluctuations in behavioral sequences should be predictive of fluctuations in time estimation. We trained rodents in an duration categorization task while continuously monitoring their behavior with a high speed camera. Animals developed highly reproducible behavioral sequences during the interval being timed. Moreover, those sequences were often predictive of perceptual report since early in the trial, providing support to the idea that animals may use learned behavioral patterns to estimate the duration of time intervals.

2.1 Introduction

Animals live in naturally stochastic environments where apprehending environmental regularities is extremely important. In particular, being able to identify temporal regularities in the environment enables animals to predict future events such as the presence of mates, food or danger (Balsam & Gallistel, 2009), or to decide between alternative courses of action, e.g. deciding when to switch from exploiting a depleting food patch to exploring for new ones so as to optimize energy balance (Kacelnik & Brunner, 2002; Bateson, 2003). Behaviorally relevant temporal regularities in the environment are often on the scale of multiple seconds, therefore understanding how organisms handle time durations on this scale is extremely important to understand behavior itself.

Traditional sensory modalities such as vision, audition or tactile sensation are processed by known sensory organs and brain areas. The implementation of time perception within the nervous system, on the other hand, is still a matter of debate and ongoing research (viz. Wittmann, 2013).

Neurally inspired models for interval timing include those that involve coincidence detection among oscillations of varying frequencies (Miall, 1989; Matell & Meck, 2000, 2004), integration of the noisy firing of neural populations (Simen, Balcı, de Souza, Cohen, & Holmes, 2011) and variable firing dynamics within a population of neurons (Grossberg & Schmajuk, 1989; Buonomano & Merzenich, 1995; Meck, Penney, & Pouthas, 2008; Shinomoto et al., 2011) as encoding schemes for time related information. Additionally, several abstract models of how animals track the passage of time have been proposed, many of which fall in one of two categories: accumulator models tell time by counting pulses emitted by a pacemaker and comparing it to a remembered value (Treisman, 1963; Gibbon & Church, 1984), while state based models represent time as a trajectory progressing through a sequence of states (Killeen & Fetterman, 1988; Machado, 1997; Ludvig et al., 2008). A subset of sequential state timing models posit that states reflect behavior (Killeen & Fetterman, 1988; Machado, 1997), stemming from the widely replicated observation that structured behavioral chains emerge under temporally structured reinforcement contingencies (e.g. Skinner, 1948; Hodos et al., 1962; Anderson & Shettleworth, 1977; Haight & Killeen, 1991; Machado & Keen, 2003; Balcı et al., 2008; Ölveczky, 2011).

Suggesting that interval timing is driven by behavioral state transitions implies a clear prediction that, to our knowledge, has not yet been tested: variation of behavioral chains should correlate with variations in time estimation. In the current work we continuously monitored the behavior of rodents as they categorized interval durations as longer or shorter than a learned standard interval. Idiosyncratic behavioral sequences displayed during the interval being timed were highly reproducible across trials and sessions. Moreover, the small variation present in the behavioral trajectories was often predictive of temporal judgments from very early in the trial, sometimes in advance of trial onset. These results revealed a correlation between learned motor behavior and perceptual report of elapsed time.

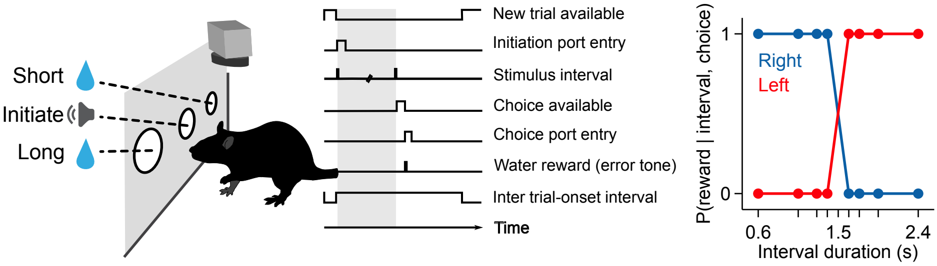

Figure 2.1. **Task schematic and reinforcement contingency.**

2.2 Results

2.2.1 Animals learned to categorize time intervals

We trained three rats to categorize time intervals as either long or short by making left/right choices (figure 2.1; see section 2.4.1 on page 35 for full description). In brief, animals initiated trials by inserting their snout in a centrally located initiation nose port, thus triggering the presentation of a stimulus. Stimuli consisted of two brief sound tones separated in time by an interval randomly selected from the set $I = \{0.6, 1.05, 1.26, 1.38, 1.62, 1.74, 1.95, 2.4\}$ seconds. Judgments about interval duration were reported at two laterally located nose ports: choosing the left side was reinforced with a drop of water after intervals longer than 1.5 seconds, and the right side otherwise. Incorrect choices were punished with an error tone and a time out. An-

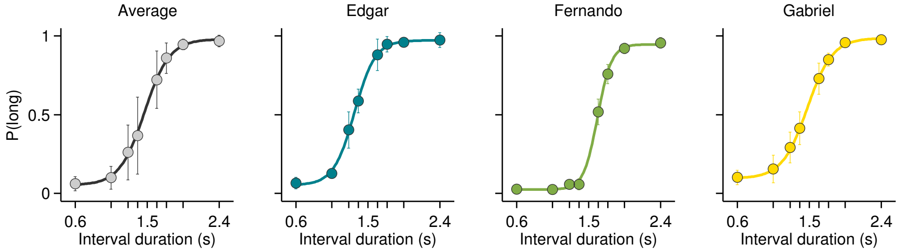

Figure 2.2. Psychometric functions (mean and standard deviation across sessions and logistic fit).

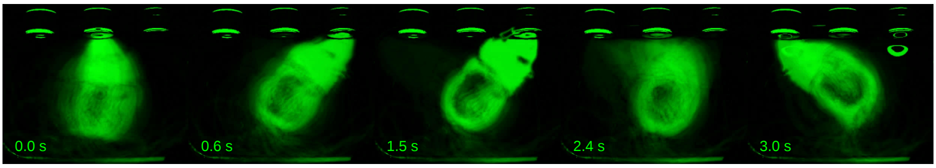

Figure 2.3. **Behavior displayed during stimulus interval is highly reproducible**. A series of video frames taken from a representative session of rat Fernando at specific time points within trials were averaged across all presentations of the longest interval. Times when frames were taken are indicated in seconds relative to interval onset. $n = 62$ trials.

imals were free to move during stimulus presentation, as long as they withheld choice until interval offset.

Each interval presentation followed by a choice constitutes one trial, and rats performed on average 453 trials per daily session (range = 346 to 558, standard deviation = 49.4). As revealed by their psychometric functions (figure 2.2), animals correctly categorized the easiest (i.e., shortest and longest) stimulus intervals at a rate of 95.8±0.03%, while performance declined as intervals came closer to the 1.5 seconds categorical boundary reaching 69.4±5.58% for the most difficult stimuli; over all eight stimuli, 84.7±3.57% of trials were categorized correctly (mean ± standard deviation).

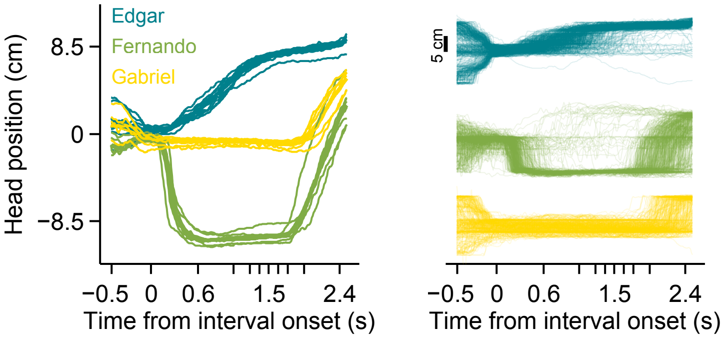

Figure 2.4. **Head trajectories are reproducible and idiosyncratic**. (*left*) Head trajectories around presentations of longest interval. Thin lines are single session means. Thick lines are means across session means. (*right*) Trajectories at all single trials in which the longest interval was presented.

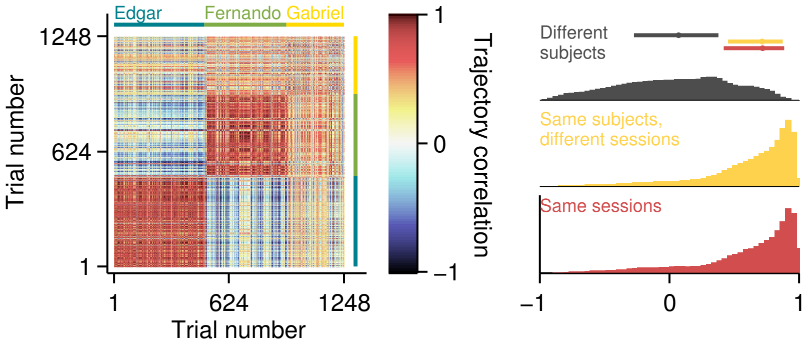

Figure 2.5. **Pairwise correlations between head trajectories**. (*left*) Matrix of pairwise correlations between trajectories shown in figure 2.4 on the previous page. Trials are ordered by subject (color bars framing top and right margins), then session, then trial number. (*right*) Normalized histograms of correlation coefficients between single trial trajectories from different subjects (dark gray), or from the same subject and in the same (yellow) or different (red) sessions. Dots and bars above histograms are medians and interquartile ranges. $n = \{491, 445, 312\}$ trials in $\{10, 8, 6\}$ sessions from the three rats.

2.2.2 Animals developed temporally structured behavior

Apart from the discrete, temporally sparse behavioral measurements obtained from the nose ports, we monitored behavior continuously with a high-speed camera (120 fps). The camera was located at the top of the behavioral box, thus only detecting motion occurring within planes parallel to the floor. Videos taken around presentation of stimulus intervals revealed highly consistent body motion patterns. To illustrate this, we overlaid video excerpts of a representative session of rat Fernando time-aligned at the onset of presentations of the longest interval. This interval duration was chosen for the analysis because it allows behavioral sequences to unfold for as long as possible before being disrupted by the interval offset tone, which acts as a go signal prompting the animal to move to a choice port. As the behavioral sequences are highly reproducible, the resulting averaged images are surprisingly sharp (figure 2.3 on the preceding page).

In order to quantify this effect and extend our analysis to a number of subjects and sessions, we used the aid of computer vision tracking algorithms to follow the position of each animal's head in time (see section 2.4.3 on page 36 for details). In agreement with the example video in figure 2.3 on page 26, head trajectories revealed body motion patterns that were very consistent across trials, as well as across sessions (figure 2.4 on page 26). Interestingly, each subject developed its own distinct trajectory.

To assess the within and between subject variability in head trajectory, we computed a correlation matrix comparing all pairwise combinations of trials wherein the longest interval duration was delivered (Figure 2.5, left). Correlations between trajectories produced by a given subject were highly and consistently positive, whether or not they occurred in the same session (Figure 2.5, right; distributions of correlation coefficients are indistinguishable, Kolmogorov-Smirnov test, $p = 0.84$). Trajectories from different subjects, on the other hand, were uncorrelated on average, and the distribution of coefficients differed significantly from both within subject distributions (K-S tests, $p = 9.7 \times 10^{-4}$ in both cases).

Given the observation that trajectories are idiosyncratic and consistent from session to session, we pooled data across sessions within subjects for the following analyses.

2.2.3 Ongoing behavior bears information about unfolding perceptual decisions

Next, we asked whether trial-to-trial variability in body trajectories carried information about the perceptual decisions being forged; were this the case, different categorizations of the same stimuli should be accompanied by distinct behavioral trajectories. For a representative session of rat Edgar, we selected all presentations of the stimulus for which choice variance (σ^2_{choice}) was highest ($I = 1.38$ seconds; $n = 56$ trials; $\sigma^2_{choice} = 13.3571$). The different color channels in the video were used to parse trials by choice: long choice

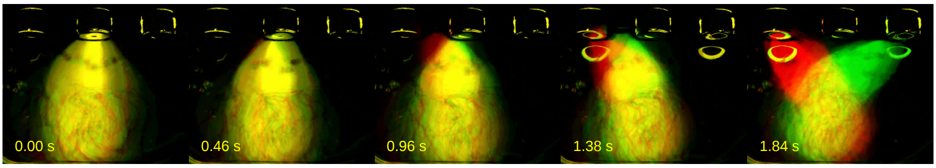

Figure 2.6. **Distinct behaviors accompany different categorizations of same stimulus.** A series of video frames were taken from a representative session of rat Edgar at specific time points during presentations of a near-boundary stimulus interval ($I = 1.38$ seconds). Short choice trials were put on the green, and long choice trials on the red channel. Times when frames were taken are indicated in seconds relative to interval onset.

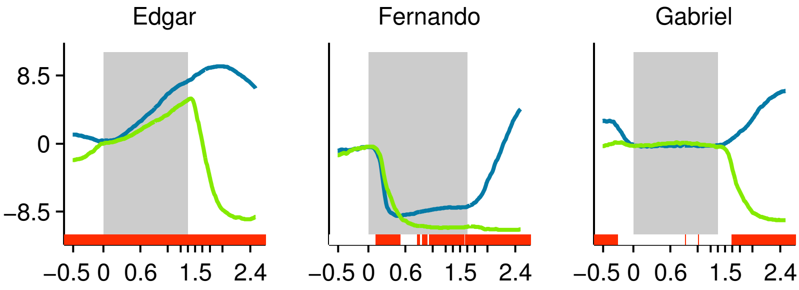

Figure 2.7. **Head trajectory is predictive of choice**. Average head trajectories leading to long (blue) and short (green) categorizations of same near boundary durations. Gray shaded area indicates stimulus interval period. For each subject, the stimulus of highest choice variance across sessions was selected. Red bars indicate moments when head position is significantly predictive of choice (95% bootstrap confidence intervals on area under the receiver operating characteristic (ROC) curve). $n = \{678, 475, 300\}$ trials.

trials were put in the red, and short choice trials in the green channel. The resulting video reveals a separation in body position since the first few hundreds of milliseconds (Figure 2.6 on the preceding page).

The differences in behavioral trajectories leading to different categorizations of the same stimuli imply that it should be possible to predict choice from behavioral trajectories. To quantify this effect, we employed a metric commonly used in sensory neuroscience known as *choice probability* (Britten et al., 1996; Nienborg, R. Cohen, & Cumming, 2012). Choice probability is defined as the degree to which fluctuations of a variable during repeated presentations of a stimulus are predictive of perceptual judgments about that stimulus. This metric is commonly applied to the firing of neurons in sensory brain areas in order to estimate their involvement in the formation of percepts (Parker & Newsome, 1998). We extend its use to assess whether body trajectories carried information about unfolding perceptual judgments of time intervals.

We started by calculating choice probability from head position at individual time steps within a period extending from 0.5 seconds before to 2.5 seconds after trial initiation (figure 2.7). Choice probability from head position was quantified as the area under a ROC curve, a metric that can be interpreted as the probability that a sample taken from one of two univariate distributions will be correctly assigned to its source

(Green, Swets, et al., 1966). This curve was calculated using head position distributions for data from short versus long choice trials (see section 2.4.4 on page 37 for details). In agreement with the example video, this analysis revealed that overt behavioral sequences often allow perceptual judgment to be predicted above chance. The profile of choice probability over time differed for each individual subject (figure 2.7 on the previous page). Edgar displayed a monotonically increasing profile that was significant from before stimulus onset and throughout stimulus presentation. Fernando displayed a more complex profile that was significant early in the stimulus, lost significance, and then regained small but generally significant separability from ≈ 0.7 seconds onward. Gabriel did not display overt head trajectories during the interval period, staying at the initiation port throughout presentation of the stimulus interval instead (see figure 2.4 on page 26). However, Gabriel's choice probabilities were significantly greater than chance prior to trial initiation. The absence of appreciable change in Gabriel's head position during the interval period made it impossible to extract any information from this variable. However, close inspection of individual videos suggested that this rat may have produced smaller scale movements around the initiation port in the axis normal to the image plane. We were not able to quantify such movements using the current setup, and thus likely underestimated the degree to which this animal's movement may have related to choice.

The analysis of head position presented in figure 2.7 on the previous page represents a behavioral analogue of instantaneous neuronal choice probability, calculated during the presentation of near boundary stimuli. The following analysis differs in that it assesses the predictive power of a 1-second long segment of the head trajectory with regard to choice across different difficulty levels.

Briefly, the trajectory described by head position on trial k during a time window extending from 0.5 seconds before trial initiation to 0.5 seconds after was represented as a vector $\mathbf{h}_k$. Since the shortest possible stimulus duration is 0.6 seconds, this period is identical across all trials in the sense that no information about stimulus duration has been presented. Intuitively, if behavioral trajectories are systematically related to perceptual decisions, head position vectors $\mathbf{h}$ observed on long and short choice trials should form separable clouds.

Choice probability was estimated from single trials with LDA. Next, we grouped trials based on choice probability in five bins containing equal number of trials, and plotted the mean trajectory and psychometric curve for each bin (Figure 2.8 on page 41; see section 2.4.4 on page 37 for details).

All psychometric curves asymptoted near 0 and 1, but biases were strikingly different and ordered. As noted, this analysis took as input behavior occurring before stimulus identity could possibly be known by the subject. Importantly, the ordered psychometric functions do not suggest that the animal has already made its decision prior to stimulus presentation. Were this the case, performance should be at chance level. Rather, the animal's head trajectory exerted a bias on choice, as the difference in psychometric functions was mainly captured by the bias parameter. The same pattern would be expected if we could bin trials with respect to internal decision variables such as clock speed or decision criterion.

2.2.4 Behavioral trajectory improves choice prediction beyond trial history

The average head trajectories during the period preceding trial initiation were strikingly related to choice probability (Figure 2.8 on page 41, top panels). This suggests that the predictive power of trajectories reflected events preceding the current trial, such as choices made and rewards received on recent trials (e.g., Sugrue, Corrado, & Newsome, 2004; Lau & Glimcher, 2005). To test whether behavioral trajectories significantly improved our ability to predict choice beyond the information provided by trial history, we fit four logistic regression models to the choice data that differed in the combination of predictor variables included. We allowed different combinations of subject, current trial stimulus, recent trial history of choices, stimuli and rewards, and head trajectory to be weighed in predicting choice (see section 2.4.5 on page 38 and Table 2.2 on page 39 for details).

Model #1 captured the effect of stimuli, expected to be strong if animals learned the duration discrimination rule inherent in the task. Dummy variables standing for individual subjects were used to capture cross subject differences in psychometric functions. As expected, model #1 predicted choice at a high success rate (Table 2.1 on page 33), and was strongly significant as compared to a constant model (log likelihood ratio test, $\chi^2 = 6.01 \cdot 10^3, df = 5, p$ within rounding error of 0).

Next we assessed the contribution of behavioral trajectories during the time window used to calculate choice probability (Figure 2.8). Model #2 maintained the predictors present in model #1, to which it added two variables describing behavioral trajectories (i.e., projections on the first and second principal components; see section 2.4.5 for details). This model showed a modest improvement in prediction success rate over model #1. Better predictions are expected from models with more free parameters. We therefore employed formal model comparison methods by calculating Bayesian information criterion (BIC) and Akaike information criterion (AIC). In summary, these methods impose a cost for adding free parameters. The improvement achieved by adding extra free parameters for trajectories outweighed the costs imposed by both methods (Table 2.1). In addition, a log likelihood ratio test indicated that the improvement of model #2 over model #1 is highly significant ($\chi^2 = 445, df = 5, p = 3.19 \cdot 10^{-94}$).

In order to assess the contribution of information regarding trial history, model #3 added to predictors in model #1 variables describing stimulus, difficulty and reward on the preceding trial. Similar results were obtained for model #3 as for model #2, both compared to model #1 (Table 2.1 on the following page; log likelihood ratio test, $\chi^2 = 135, df = 1, p = 1.59 \cdot 10^{-31}$). A direct comparison between BIC and AIC values of models 2 and 3 indicated that trajectories were in fact slightly more informative than trial history. Furthermore, neither of the simpler models were better than model #4, a full model incorporating variables relative to subject, stimulus interval, trajectories and trial history (Table 2.1 on the next page; log likelihood ratio tests relative to model #4; model #2: $\chi^2 = 51, df = 1, p = 5.52 \cdot 10^{-13}$; model #3: $\chi^2 = 360, df = 5, p = 4.93 \cdot 10^{-76}$);

2.3 Discussion

Under environments with strong temporal regularities, animal behavior is known to become temporally structured (e.g. Skinner, 1948; Hodos et al., 1962; Anderson & Shettleworth, 1977; Haight & Killeen, 1991; Machado & Keen, 2003; Balcı et al., 2008; Ölveczky, 2011). Stemming from this observation, trajectories through behavioral states have been proposed to implement interval timing (Killeen & Fetterman, 1988; Machado, 1997). A prediction implied by this rationale is that trial to trial variations in the flow of behavioral sequences should correlate with trial to trial variations in temporal estimation. Consistent with this prediction, we found that behavioral trajectories differed between cases where the same intervals were categorized differently, allowing choice to be significantly predicted from ongoing behavior.

While some sequential state timing models map states directly onto behaviors, many models posit a more abstract sequential state representation of time. These include neural network models that produce dynamics in response to input (Buonomano & Merzenich, 1995), successively broadening temporal basis functions (Grossberg & Schmajuk, 1989; Suri & Schultz, 1999; Ludvig et al., 2008), or models that time intervals through specific phase relationships amongst a diversity of oscillatory processes

Model #	Predictor groups	Prediction success	Deviance	BIC	BIC
1	Subject, interval	83.8%	7103	7158	7115
2	Subj, int, trajectory	84.9%	6658	6759	6680
3	Subj, int, trial history	84.6%	6968	7032	6982
4	Subj, int, traj, trial hist	85.0%	6607	6717	6631

Table 2.1. **Behavioral trajectories improve choice prediction beyond trial history**. Outcome of multivariate logistic regression models fit to predict choice from subject identity, stimulus, trial history, and behavioral trajectory.

(Miall, 1989; Matell & Meck, 2000, 2004). The data presented in this chapter provides evidence in favor of state transition timing models, but it does not speak to the underlying neural mechanisms, and thus is consistent with all of the above sequential state models for timing.

Furthermore, our findings do not rule out the existence of dedicated timing mechanisms such as the pacemaker-accumulator model posited by scalar expectancy theory (Gibbon & Church, 1984) In other words the quantitative relationship between continuous behavior and subsequent choice we observed does not prove that behavior directly drives the perceptual process of time estimation. Behavior may instead simply reflect a central timing process. Alternatively, the rodents in our study may have indeed used behavioral sequences to estimate time in our task, with such a strategy being a useful but non-unique solution to the problem of how to estimate duration. Hence, our results alone merely suggest embodied strategies as one of a number of possible solutions to timing. Lastly, the brain is presumably responsible for the production of stereotyped motor sequences, and so in a trivial sense, the brain must be part of the system that estimates duration, even if this computation is to some degree dependent on the inertia of ongoing behavior. Why might organisms include ongoing behavior in the process of timing? Limbs and muscles have mass and inertia, which may increase time constants present in movement to a degree that is harder to achieve within the nervous system alone (cf. Vogels, Rajan, & Abbott, 2005). Additional experiments that manipulate the environment dynamically would help reveal whether behavior is upstream of the perceptual process of time estimation.

A number of brain structures have been implicated in interval timing. These areas include parietal cortex (Leon & Shadlen, 2003; Janssen & Shadlen, 2005; Jazayeri & Shadlen, 2015), prefrontal cortex (Fuster, 2001; Kim et al., 2009, 2013; Xu et al., 2014), the basal ganglia (Maricq & Church, 1983; Matell & Meck, 2000, 2004; Meck et al., 2008; Fiorillo, Tobler, & Schultz, 2003; Jin et al., 2009; Adler et al., 2012; Mello et al., 2015; Gouvêa et al., 2015b), and the cerebellum (Ivry & Keele, 1989). Interestingly from the perspective of embodied solutions to timing, these areas represent some of the most highly integrative territories in the mammalian brain, processing information concerning a broad range of sensory modalities and effector systems.

As is often the case, the full nature of interval timing likely reflects a mixture of various mechanisms, embodied and non-embodied alike (viz. Wittmann, 2013). Animals and brains have evolved to be opportunistic, and are capable of employing a variety of strategies to solve cognitive tasks depending on the scenario they find themselves. The method of computing choice probabilities from behavior presented here represents a path forward for neuroscientists seeking to disambiguate embodied versus non-embodied components of cognitive acts such as perceptual decision making. In other decision-making contexts, it has been shown that information about an unfolding decision continually flows to the motor system, such that continuously observing behavioral output provides information about the dynamics of a decision process (Selen, Shadlen, & Wolpert, 2012). In this way, ongoing behavior can provide a readout of the current state of decisions, which can then be compared to neural signals that are thought to be involved. We suspect that, whether studies involve human, non-human primate, rodent, or other species as experimental subjects, some neural correlates of decision variables are explained by changing motor output or sensory input resulting from ongoing behavior. We also suspect that some neural correlates of decision variables precede the emergence of decision variables in ongoing behavior. We test some of these ideas in chapter 3 on page 43.

2.4 Materials & Methods

All experiments were approved by the Champalimaud Centre for the Unknown Bioethics Committee and the Portuguese Direção Geral Veterinária.

2.4.1 Duration categorization task

Three adult male wild type Long-Evans hooded rats were trained to categorize time intervals as either long or short by making left/right choices. Animals learned to trigger stimulus intervals by inserting their snout in the centrally located initiation nose port when it was illuminated. Triggering a stimulus would immediately turn off the initiation port light, and cause a pair of sound tones to be played separated in time by an interval randomly selected from the set $I = \{0.6, 1.05, 1.26, 1.38, 1.62, 1.74, 1.95, 2.4\}$ seconds. Intervals in the set are symmetric around the 1.5 seconds categorical boundary, and make up four difficulty levels in geometric progression. Animals were required to withhold poking at the choice ports until the interval offset tone was played, but were otherwise unrestrained during stimulus presentation. The duration of the presented interval governed reinforcement of nosepoking at the laterally located choice ports, and responses were interpreted as the animal's perceptual judgment regarding interval duration. For intervals longer than the categorical boundary, a water reward became available for delivery upon choice of the left nose port, or at the right nose port for intervals shorter than the boundary. Wrong choices were punished with a brief white noise sound and a time out. Nine seconds after the initiation of the previous trial, or twenty seconds following errors, the initiation port would become illuminated again indicating that another trial could be initiated. Twenty four sessions were selected based on categorization performance ($n = \{10, 8, 6\}$ sessions from rats Edgar, Fernando, and Gabriel, respectively).

On figure 2.2 the probability of a long choice as a function of stimulus duration (psychometric function) was fit with a 4 parameter-logistic regression of the form

$$f(x) = d + \frac{a - d}{1 + e^{-b(x-c)}} \tag{2.1}$$

where b controls the slope, c is the inflection point and a and d are the maximum and minimum values of the curve respectively.

2.4.2 Behavioral set up

Behavioral boxes consisted of a plastic bucket (IKEA, Alfragide, Portugal) containing one speaker (Cover Industrial Co., Guangdong, China) and three nose ports (Island Motion). Each nose port contained one infra-red beam/sensor pair for detecting nosepoking and one visible LED. The choice ports, contained in addition a water tube connected to a solenoid valve for reward delivery. Valves were calibrated to deliver $25\mu l$ of water per reward event.

Except for the video camera, all sensors and effectors in the behavioral box were read and controlled by an Arduino Mega 2560 microprocessor (additional information and free software available at `http://www.arduino.cc/`) via a custom circuit board. The microprocessor implemented the behavioral task, and, through a serial communication port, outputted data to a desktop computer running custom software based on Python's pySerial module (freely available at `http://pyserial.sourceforge.net/`).

2.4.3 Video acquisition and tracking

Video was acquired with a high speed camera (Flea3 FL3-U3-13S2C-CS, Point Grey Research Inc., Richmond, Canada) at 120 frames per second with a resolution of 1280 x 960 pixels in grayscale at 8 bits. Video acquisition and offline tracking were performed using the in-house developed software Bonsai (Lopes et al., 2015, freely available at `http://bitbucket.org/horizongir/bonsai`). To extract position of the head from the raw videos, images were background subtracted and thresholded so that the animal's body appeared as a distinct blob. For each frame, the blob's largest axis was found, and the spatial position of the axis tip corresponding to the animal's head was tracked in both x and y dimensions. All analyses were carried on position along the axis in which nose ports are lined up, while motion along the orthogonal axis was discarded.

2.4.4 Estimating choice probability from ongoing behavior

Momentary head position

Choice probability given head position at a particular time, $P(c \mid h_t)$, was calculated by applying an ROC analysis (Green et al., 1966) to assess the degree of overlap between the two known distributions $P(h_t \mid c)$ for $c \in \{long, short\}$. The area under the ROC curve was calculated and rectified about 0.5. This number can be interpreted as the probability that an ideal observer (i.e., one with full knowledge of the distributions) would correctly categorize a new sample. Effectively, it provides an instantaneous metric of the degree to which head position is informative about unfolding perceptual decisions. A 95% confidence interval around chance level was estimated by calculating choice probability over randomly shuffled data ($n = 1000$ iterations). The analysis was applied separately for each subject. Trials used were those in which the stimulus interval of highest choice variance was presented.

Behavioral trajectories

Choice probability was also estimated from behavioral trajectories. For this, we employed a LDA. Behavioral trajectory on trial k was defined as a vector $\mathbf{h}^k$ whose elements are instantaneous head positions recorded during a one second long time window centered on interval onset, as denoted by $\mathbf{h}^k = \left(h^k_{t=-0.5s}, \cdots, h^k_{t=0.5s}\right)$. This time window was chosen because during this period subjects had no information about the stimulus being presented, thus allowing us to combine trials of different stimulus types in the same analysis. Intuitively, one possible scenario in which trajectory early in the trial is predictive of choice is one in which vectors $\mathbf{h}^k$ from long and short choice trials form linearly separable clouds. That is the scenario we test here. We proceeded by finding the direction $\mathbf{w}$ that maximized variance across, while minimizing variance within, categories:

$$\begin{aligned}\mathbf{w} &= \operatorname{argmax}\frac{\mathbf{w}^\top S_B \mathbf{w}}{\mathbf{w}^\top S_W \mathbf{w}} \\ \mathbf{w} &= S_W^{-1}(\mu_1 - \mu_2) \qquad (2.2)\end{aligned}$$

where S_B is the between class covariance, S_W the within class covariance and μ_1 and μ_2 are the means of all points in class 1 and class 2 respectively. Trials were binned based on choice probability to generate the trajectories and psychometric functions presented in Figure 2.8 on page 41. Similar results were observed when trajectories were fit using flat weights (data not shown). The analysis was applied separately for each subject.

2.4.5 Generalized linear models

The four logistic regression models fit to choice data included different combinations of predictor variables referring to subject, stimulus, behavioral trajectory and trial history (Table 2.1 on page 33). Subjects were represented by categorical variables, while stimulus duration was represented as a continuous variable. Behavioral trajectories were considered within the same time window used to calculate choice probability (i.e. a one second time window centered on interval onset). Given the rate at which videos were acquired, trajectory excerpts of this length are variables of 120 dimensions. To avoid adding unnecessary free parameters to the logistic models, principal component analysis (PCA) was run across all trials. The first two principal components explained 89.8% of trajectory variance, and single trial projections onto these were fed to the logistic model. Trial history was represented by three continuous variables referring to trials prior to choice: stimulus duration, stimulus difficulty (defined as the unsigned distance from the 1.5 seconds categorical boundary), and reward (defined as 1 for reward after long choices, 0 for no reward, and -1 for reward after short choices). Initially, models were specified with all linear terms and no interactions (Table 2.2 on the next page, middle column). Models were then modified by a stepwise procedure that added interactions or removed terms so as to minimize BIC. Final model specifications are shown in the last column of Table 2.2. Models described following notation by (Wilkinson & Rogers, 1973).

Acknowledgments

The ideas presented in this chapter benefited significantly from discussions with the Champalimaud Neuroscience Program community. Experiments were planned by Joseph Paton and Thiago Gouvêa. Data were acquired by Tiago Monteiro and Thiago Gouvêa. Thiago Gouvêa drafted the manuscript on which this chapter is based (Gouvêa et al., 2014); Tiago Monteiro and Joseph Paton contributed to the final text. Sofia

Model #	Initial specification	Final specification
1	$S + I$	$S * I$
2	$S + I + T_{PC1} + T_{PC2}$	$S * I + T_{PC1} + S * T_{PC2} + I * T_{PC2}$
3	$S + I + R_{t-1} + d(I_{t-1}) + I_{t-1}$	$S * I + R_{t-1}$
4	$S + I + T_{PC1} + T_{PC2} + R_{t-1} + d(I_{t-1}) + I_{t-1}$	$S * I + T_{PC1} + S * T_{PC2} + I * T_{PC2} + R_{t-1}$

Table 2.2. **Specification of logistic regression models used to predict choice**. The four models were initially specified will all linear terms and no interactions. Models were then modified by a stepwise procedure that added interactions or removed terms so as to minimize BIC. S = subject; I = interval; T_{PC1} and T_{PC2} = trajectory of head position on a one second time window centered on interval onset, projected on the first and second principal components calculated across trials; R_{t-1} = reward on previous trial; $d(i_{t-1})$ = difficulty of the stimulus presented on previous trial, defined as the unsigned distance from the 1.5 seconds boundary.

Soares and Bassam Atallah contributed data on a similar experiment on a different animal species to the manuscript (Gouvêa et al., 2014).

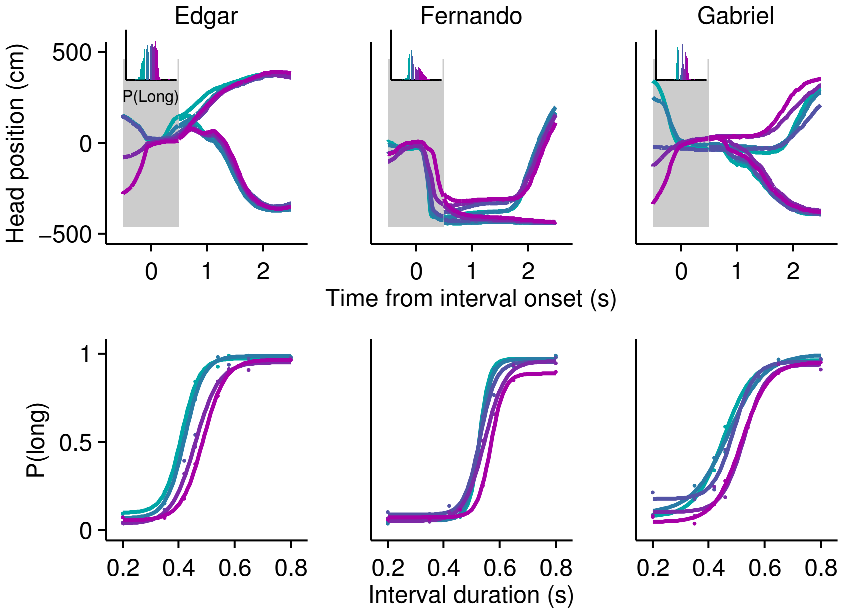

Figure 2.8. **Head trajectory reveals categorization bias.**. Single trial trajectories of head position in a 1-second time window centered on interval onset were used to predict choice. Choice probability, estimated with linear discriminant analysis (LDA), was used to bin trials. (A) Trajectories averaged within bins are shown for the time window used for prediction (shaded area). For the remaining time, mean trajectories were further split by choice. Color code indicates bins. Inset: histogram of choice probabilities estimated from trajectories. (B) Psychometric curves for trials comprised in each bin. $n = \{4542, 3680, 2641\}$ trials grouped in 5 bins.

Chapter 3

Striatal dynamics represent subjective time

Time is a fundamental dimension of the environment, yet how time is processed in the brain is poorly understood. Prior studies have shown that population dynamics in a number of brain areas encode information about the passage of time. However, it is not known whether such temporal representations inform subjects' judgments of duration or merely covary with elapsing time. The striatum is an input structure of the basal ganglia implicated in several time-dependent functions such as reinforcement learning, decision making, and interval timing. Here we show that the dynamics of neural activity in populations of striatal neurons predict judgments about the duration of time intervals. Simultaneously recorded ensembles of around 20 neurons allowed stimulus categorization at a performance level comparable to animal behavior. Furthermore, striatal neurons were necessary for duration judgments, as muscimol infusions produced a specific impairment in animals' duration sensitivity. Lastly, we show that the speed of time as encoded by striatal populations correlated with behavioral report, with encoded time progressing faster in long choice, and slower in short choice trials. These results suggest that striatal dynamics form an internal "neural population clock" that supports the fundamental ability of animals to judge the passage of time.

3.1 Introduction

Time, like space, is a fundamental dimension of the environment, yet how it is processed in the brain is poorly understood. A number of recent studies have identified dynamics that allow for robust representation of time by populations of neurons in multiple areas including the hippocampus (Pastalkova et al., 2008; MacDonald et al., 2011), prefrontal (Matell, Meck, & Nicolelis, 2003; Jin et al., 2009; Kim et al., 2013; Xu et al., 2014), parietal (Leon & Shadlen, 2003; Janssen & Shadlen, 2005; Jazayeri & Shadlen, 2015) and motor (Lebedev, O'Doherty, & Nicolelis, 2008; Knudsen, Powers, & Moxon, 2014) cortices, and the striatum (Matell et al., 2003; Jin et al., 2009; Adler et al., 2012; Mello et al., 2015). However, any dynamics that result in a continuously-evolving and non-repeating population state can be used to encode time (Buonomano, 2014), and it is not known whether such temporal representations would inform subjects' judgments of duration or merely covary with elapsing time. The striatum, a brain structure known to be involved in reinforcement learning and decision making (Samejima, Ueda, Doya, & Kimura, 2005; Lau & Glimcher, 2008; Lee, Tai, Zador, & Wilbrecht, 2015), has been implicated in interval timing by several lines of evidence (Malapani et al., 1998; Hinton & Meck, 2004; Meck, 2006; Harrington et al., 2009; Wencil, Coslett, Aguirre, & Chatterjee, 2010). However, whether dynamics in striatal activity can explain the perceptual performance of behaving subjects is unknown. To determine whether striatal ensembles drive subjects' judgments of duration, we manipulated and recorded from striatal neurons in rats performing a duration categorization psychophysical task.

3.2 Results

To measure the duration sensitivity of subjects' timing judgments, we trained rats on the previously described duration categorization task (section 2.4.1 on page 35, Gouvêa et al., 2014; figure 2.1 on page 25). Animals performed on average 380 trials per daily session (range = 255 to 559, standard deviation = 76.0). As was the case with the data set in chapter 2, animals made virtually no errors when categorizing the easiest (i.e., shortest and longest) intervals, but categorization performance declined as intervals approached the 1.5 seconds categorical boundary. As revealed by their psychometric functions (figure 3.1 on the following page), animals correctly categorized the easiest (i.e., shortest and longest) stimulus intervals at a rate of $94.9 \pm 3.7\%$, while performance declined as intervals came closer to the 1.5 seconds categorical boundary

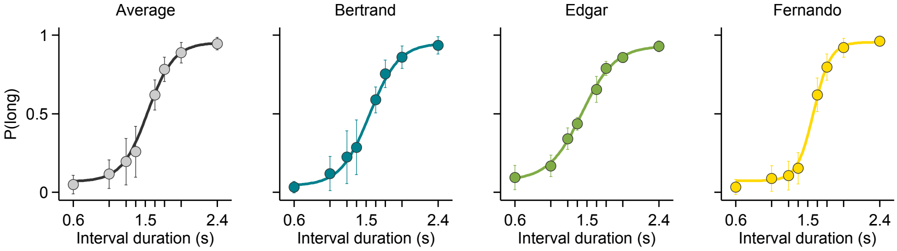

Figure 3.1. **Categorization performance during recording sessions.** Dots and whiskers are mean and standard deviation across sessions. Lines are psychometric functions fitted with equation 2.1 on page 36. Data averaged across sessions of all animals (left), or for each animal separately (animal indicated in panel title). n = (10, 9, 18) sessions from 3 animals.

reaching 68.1 ± 8.2% for the most difficult stimuli; over all eight stimuli, 82.8 ± 5.1% of trials were categorized correctly (mean ± standard deviation).

3.2.1 Striatal neurons show diverse temporal firing patterns

Several lines of evidence implicate the striatum in interval timing, but whether striatal neural activity can explain the perceptual performance of behaving subjects is unknown. We recorded action potentials from populations of single striatal neurons during task performance (see figure 3.16 and section 3.4.3 on page 59). We observed that striatal neurons displayed diverse firing patterns, with different units firing at different times within the interval period (Figure 3.2). Can such firing patterns support duration judgments? To determine whether and the degree to which individual neurons could

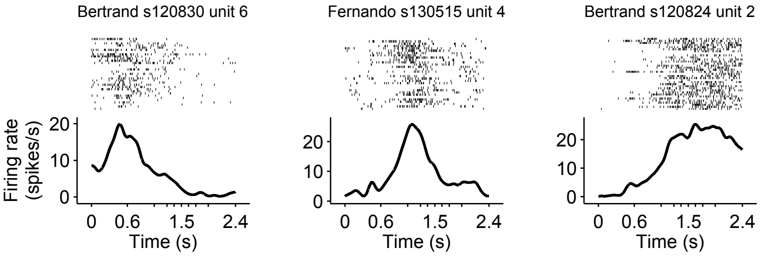

Figure 3.2. **Example neurons.** Raster plots and PSTH of activity from three example units during trials in which the longest stimulus interval (2.4 s) was presented.

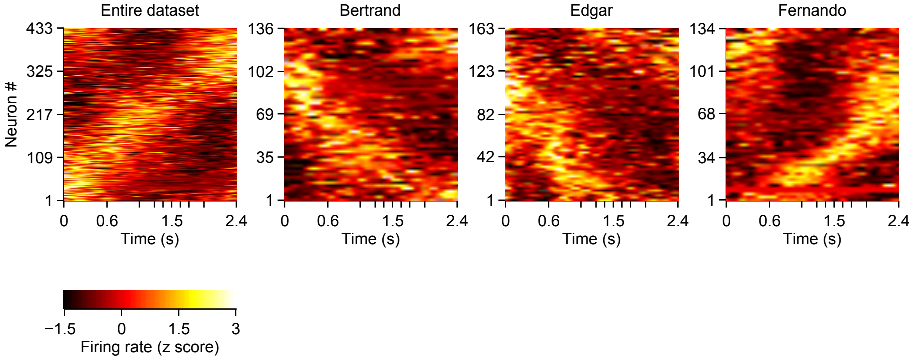

Figure 3.3. **PSTH of all neurons in the dataset.** Normalized PSTH of all neurons in the dataset for trials in which the longest stimulus interval (2.4 s) was presented. Time 0 = interval onset. n = 433 units.

contribute to duration judgments, for each trial, we counted spikes in the last $500ms$ of the interval period and compared spike count distributions of short vs long stimulus trials using a ROC analysis (see Methods). We found that the majority of neurons (57%) showed significant preferences (permutation test on area under the ROC curve (auROC), $p < 0.05$) for either short or long stimuli (figure 3.4; short-preferring: $n = 159$ in 433, 36.7%; long-preferring: $n = 87$ in 433, 20.1%). As expected, short-preferring neurons displayed higher firing on average prior to the 1.5 s category boundary, after which long-preferring neurons displayed higher firing (figure 3.5, left panel).

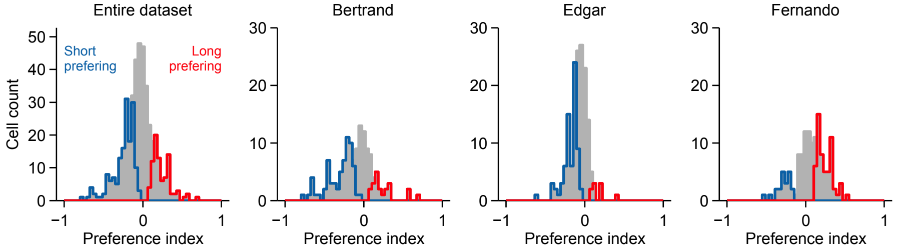

Figure 3.4. **Histogram of preference indices**. Blue and red outlines indicate subpopulations with significant short and long preferences, respectively.

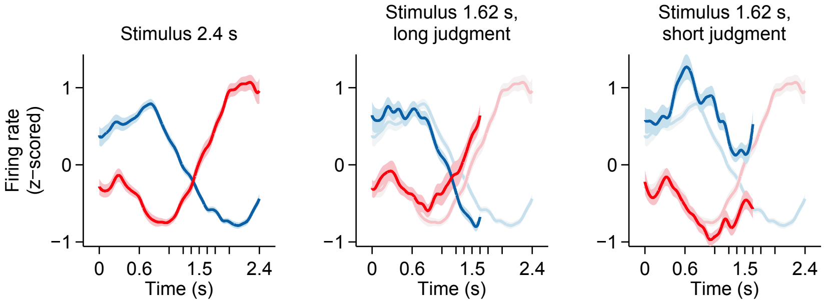

Figure 3.5. **Averaged PSTH of striatal subpopulations** Averaged, normalized PSTH of the two subpopulations outlined in figure 3.4. (*left*) Data from trials trials in which the longest stimulus interval (2.4 s) was presented. (*middle, right*) Data from trials trials in which a near boundary stimulus interval (1.62 s) was presented and the animal judged it as long (middle) or short (right). Curves from the longest stimulus are repeated as a watermark for comparison. In all panels, curves and shaded areas are mean±SEM.

3.2.2 Striatal subpopulations show distinct dynamics for different temporal judgments

The averaged activity patterns depicted in figure 3.5 resemble the likelihood of receiving reward on a moment-by-moment basis should the animal choose short or long (compare with reward contingency in figure 2.1 on page 25). Such signals, previously observed in the parietal cortex of monkeys performing a similar timing task (Leon & Shadlen, 2003) and in the striatum in a value based decision task (Lau & Glimcher, 2008), are potentially useful for guiding choice. However, were animals' judgments indeed guided by such signals, it should be possible to predict choices reported later in the trial using neural activity collected during interval stimuli. Indeed, in trials wherein a near boundary interval was judged as long, firing of the short- and long-preferring subpopulations evolved faster than on average (compare with curve for an easy stimulus, in watermark), so that the two curves crossed before the 1.5 s boundary (figure 3.5, middle panel). Conversely, in trials wherein the same interval was judged as short, the two curves evolved more slowly so that at the time of interval offset the short-preferring subpopulation was still firing at a higher level and a crossing point had not yet been reached (figure 3.5, right panel).

3.2.3 Perceptual sensitivity is predicted by neural separability

We observed in the striatum large proportions of short- and long-preferring neurons whose averaged dynamics predicted how near boundary stimuli would be categorized. This is suggestive that duration judgments are guided by the state of striatal populations. However, this analysis consisted of combining data across multiple trials and even across different sessions, benefiting from a privileged standpoint that is never accessible to the brain.

To test whether duration judgments could be supported by activity occurring in single trials, we analyzed the activity of small populations of simultaneously recorded striatal neurons. Our analysis consisted of measuring the separability of activity states at the offset of short versus long intervals, and comparing session to session fluctuations in this metric with fluctuations in behavioral performance. Briefly, for each trial in a session we characterized neural population state as a vector $\mathbf{r} = (r_1, r_2, ..., r_N)$, where r_n is the number of spikes fired by neuron $n \in N$ within the last $500ms$ of the interval period. Next, for each session we found the linear discriminant that best separated population state vectors according to whether they came from a long or a short interval trial (LDA, projection 1 in figure 3.6, left panel). A threshold placed along the linear discriminant was then used as a decision rule (black line in the left panel of figure 3.6) to generate a *neural duration judgment* for each trial. This procedure allowed us to obtain,

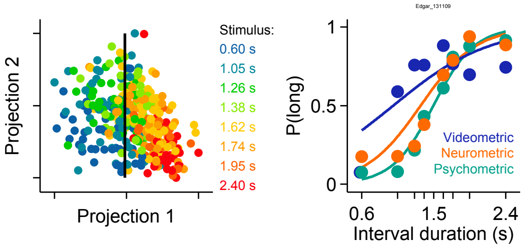

Figure 3.6. **Population state at interval offset allows stimulus categorization.** (*left*) Population state at interval offset during an example session. Each dot is a low dimensional representation of population state on a single trial at the last moment of the color-coded interval. Black line is the decision rule (see text and methods at subsection 3.4.7). (*right*) Probability that the given intervals would be classified as long by the animal (psychometric) or by a LDA applied to neural (neurometric) or video (videometric) data acquired just before interval offset.

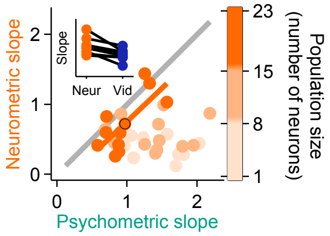

Figure 3.7. **Slopes of psychometric and neurometric curves correlate.** Slopes of psychometric and neurometric curves for all sessions. Color indicates tercile of population size in number of neurons. Highlighted data point corresponds to the session in figure 3.6. Inset: regression slope of neurometric and videometric curves for sessions in the upper tercile. See Figure 3.8 for psychometric-videometric comparison at different time points.

for each session, a quantitative description of how well simultaneously recorded neurons could categorize stimuli, i.e., a neurometric function comparable to the behavioral psychometric function (for an example session, see orange curve in figure 3.6, right panel). Consistent with duration information being encoded at the population level, we found that for sessions in which greater numbers of neurons were recorded simultaneously (i.e., upper tercile of sessions with regard to population size) psychometric and neurometric performances were similar and strongly correlated (slope $= 0.95$, $r^2 = 0.726$, $p = 0.0114$; figure 3.7). These results demonstrate that a read out of stimulus category from even modestly-sized ensembles of striatal neurons was in many cases sufficient to explain the pattern of duration judgments produced by behaving subjects.

3.2.4 Striatal neural separability is not explained by ongoing behavior

It has been previously reported that duration judgments could be predicted by animals' ongoing behavior during the interval period (Matthews & Lerer, 1987; Killeen & Fetterman, 1988; Machado, 1997; Machado & Keen, 2003; Gouvêa et al., 2014). In addition, it is well known that striatal neurons can fire around movements (Alexander & Crutcher, 1990; Jin & Costa, 2010). Could the categorization performance of striatal ensembles reflect activity related to movements the animal might be making during the task? To test to what degree ongoing behavior could explain the categorization performance of striatal neural activity, we applied an analogous classification analysis to video images taken of the animal just before interval offset (see subsection 3.4.8 on page 61). We found that our ability to categorize intervals using video frames was consistently poorer as compared to neural data collected at the analogous time periods during the task (blue curve in figure 3.6, inset in figure 3.7). In contrast, we were able to categorize stimuli as well as the animal using video frames taken at the point when animals expressed their choice at one of the reward ports (figure 3.8). Furthermore, movement related responses in the striatum are known to occur both pre- and post-movement onset, much later than in other motor areas such as pre-motor and motor cortex(Alexander & Crutcher, 1990). Thus, if purely movement-related activity were responsible for the categorization performance of striatal ensembles, we would expect ensemble performance to display a similar time course to that of video frames. Applying the same analyses at multiple points in time ranging from 500 ms preceding to 500 ms following stimulus offset revealed a strikingly different profile of categorization performance for video frames as compared to neural ensembles (figure 3.9). Specifically, the time course of duration categorization by neural ensembles was best correlated with the duration categorization by video frames when using spikes collected between 400 ms and 200 ms preceding a reference video frame. These indicate that the categorization performance of striatal neurons was not simply related to the immediate sensorimotor

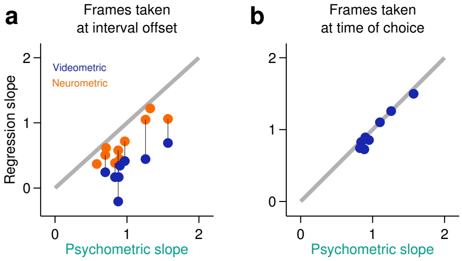

Figure 3.8. **Neural activity predicts stimulus better than ongoing behavior at stimulus offset.** (**a**) Neurometric (orange data points) or videometric (purple data points) logistic slope plotted against the psychometric slope for each session in the upper tercile with respect to simultaneously recorded population size. (**b**) Videometric slope plotted against the psychometric slope where the videometric curve was built using image frames taken at the time that animals expressed their choice.

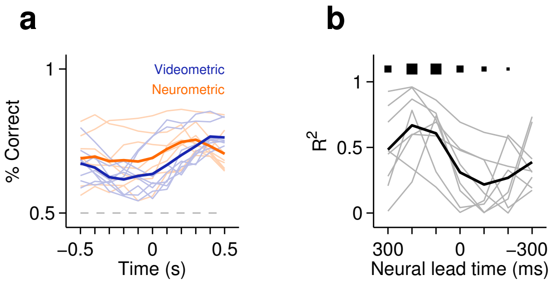

Figure 3.9. **Information about stimulus category is present in neural activity before it appears in ongoing behavior.** (**a**) Performance of an ideal observer analysis in predicting stimulus category, applied to neural (orange) and video (blue) data obtained at different times relative to interval offset. Thin lines corresponds to individual sessions. Thick lines are averages. (**b**) The orange and blue curves (shown in panel (a)) for corresponding sessions were regressed against each other at different time shifts. The regression R^2 values for each session are shown in thin grey lines. The average over all sessions is shown in black. Sizes of black squares indicate the number of sessions with significant positive correlations (largest squares at 200 and $100ms$ correspond to 5 sessions and smallest one at $-200ms$ to 1, out of a total of 8 sessions).

state of the animal, and instead likely reflects that striatal neurons encode an internal neural representation of the state of animals' categorical decisions.

3.2.5 Striatum encodes subjective time

We have shown thus far that categorical information about interval duration contained in the firing of striatal populations at the time of stimulus offset can explain the precision of animals' judgments about duration. However, in the task employed here, categorical judgments must be derived from a continuously evolving decision variable that represents how much time has elapsed since the onset of the stimulus. As can be seen in the diversity of firing patterns in figure 3.3, the state of population activity evolved continuously during interval stimuli (figure 3.12a), a feature not captured by binary classification. Might trial to trial variations in population state predict the apparent speed of animals' internal representation of elapsed time? To test this possibility, we performed two additional analyses.

First, we projected the state of simultaneously recorded neuronal populations at stimulus offset in individual trials onto the mean trajectory within each session, noted the fraction of the mean trajectory traversed for each trial, and pooled the data for each stimulus over all sessions within a given subject. Indeed, when population state at stimulus offset had advanced relatively more or less along the mean trajectory, animals were more likely to judge intervals as long or short respectively (figure 3.10,figure 3.11). This effect was observed most consistently for interval stimuli that were closer to the category boundary, and thus variations in projected population state led to horizontal shifts in the psychometric curves (see Methods). These data are consistent with striatal population state encoding a representation of elapsed time that is used by animals to determine categorical judgments. Indeed, such a pattern of population activity has been proposed as a suitable neural code for elapsing time (Karmarkar & Buonomano, 2007; ?, ?).

However, if such a representation encodes elapsed time, and not only subjects' judgments in this task, neural activity should continuously traverse a non-repeating trajectory in state space in a manner that predicts duration judgments during presentation of particular stimuli. Indeed, even in a low dimensional projection of population activity, we found that network state ran ahead or behind depending on whether the animal

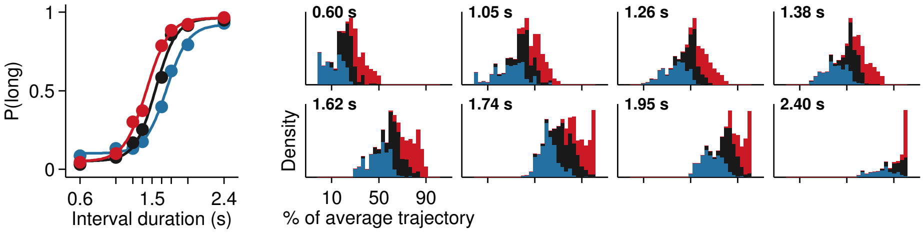

Figure 3.10. **Distance traveled in neural space along reference trajectory is reflected in decision bias** (*left*) Psychometric curves constructed from trials separated in three groups according to how far the population state at stimulus offset had advanced along the mean trajectory. Color indicates trials in the panels to the right. (*right*) Histograms of projections of population state at stimulus offset onto normalized mean trajectory. Colors indicate terciles calculated within single sessions (blue is 1st, black 2nd, and red 3rd tercile). Histogram bars of different colors are stacked. Each panel shows data for a different stimulus interval. Data from all rats. See figure 3.11 on the next page for single subject results.

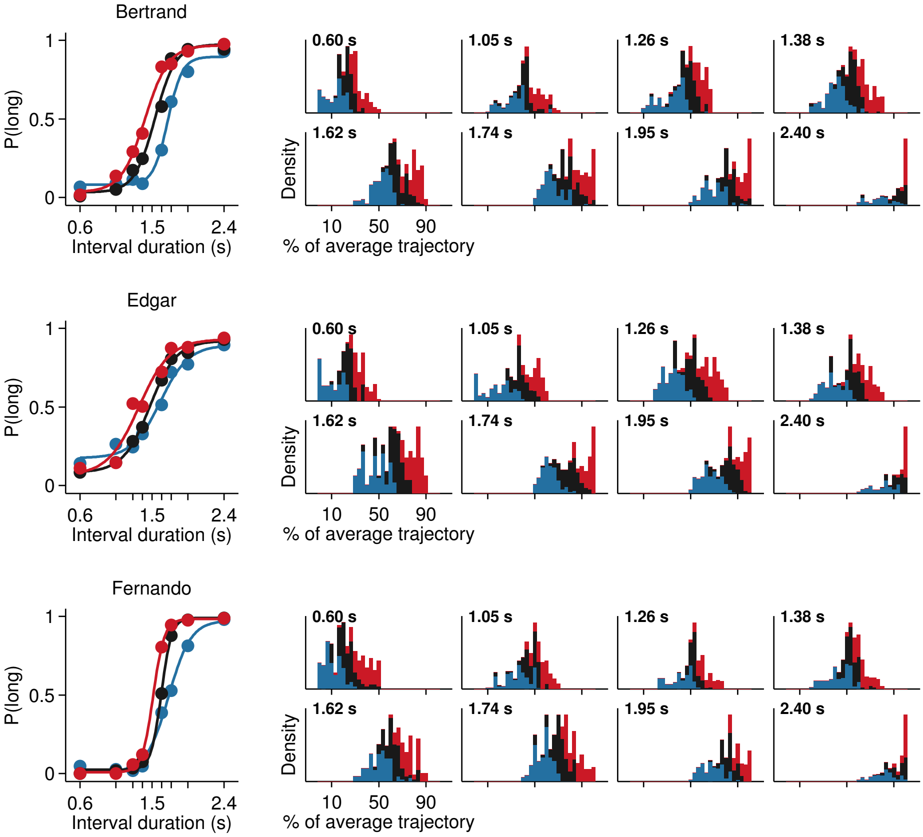

Figure 3.11. **Effect of distance traveled in neural space on decision bias is evident within single subjects** Same as in figure 3.10 for each of the three individuals.

judged a near boundary stimulus as long or short (figure 3.12b-c). The correspondence between population trajectory and duration judgments further suggests that striatal dynamics may form an internal representation of elapsed time that informed categorical decisions about duration. To directly test this hypothesis, we focused on stimuli near the category boundary. We decoded time from the population using a naive Bayes decoder and asked whether such a representation correlated with animals' judgments. We found that decoded estimates of time ran faster or slower when animals judged a given stimulus as long or short, respectively (figure 3.13, figure 3.14). This indicates that striatal activity contains information about elapsing time, the continuously varying decision variable necessary to inform judgments in the task. Furthermore, if this information were read out and used to guide judgments, those judgments would match those of the rats.

3.2.6 Intact striatum is necessary for task performance

If the striatal activity we describe above directly supported task performance, manipulating the striatum should modify duration judgments. To test whether this was the case, we bilaterally injected the GABAa receptor agonist muscimol. As a result, the duration sensitivity of animals' judgments dropped significantly as compared to interleaved saline control sessions (figure3.15; psychometric slope on saline sessions = [1.031.20] vs on muscimol sessions = [0.280.67]; 95% confidence intervals), yet animals otherwise performed normally. These results, by demonstrating that duration categorization in this task was dependent on a normally functioning striatum, suggest that the neural

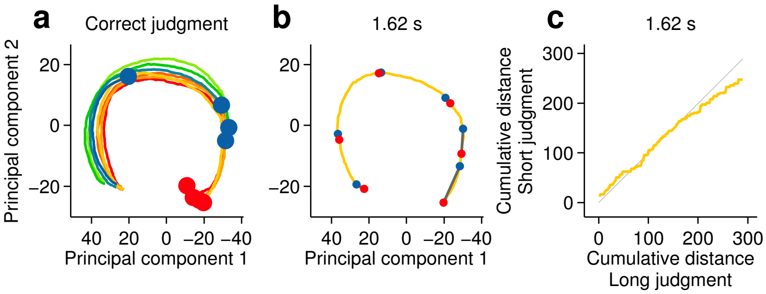

Figure 3.12. **Average speed of smoothly changing population state varies with perceptual report** (**a**) Low dimensional representation of population state during entire interval period of correct trials. Line colors indicate interval duration (warmer colors are longer intervals, as in figure 3.6). Dots are placed at the interval offset end, and their color indicates judgment (blue: short; red: long). (**b**) Yellow curve is same as in (a) (stimulus 1.62s). Red dots are six time points evenly spaced between interval onset and offset. Blue dots are projections of population state during short judgment trials. Grey lines link population states at equivalent time points. (**c**) Average cumulative distance traveled in full neural space along trajectory represented in (b) on long versus short judgment trials. $n = 433$ neurons from 3 rats.

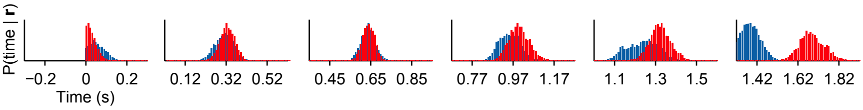

Figure 3.13. **Time decoded from striatal activity at different time points within stimulus period** Posterior probability of time given population state observed at six time points evenly spaced between interval onset and offset (as in 3.12b), averaged within trials of each judgment type (short judgment in blue, and long in red). $n = 433$ neurons from 3 rats.

signals we observed directly supported duration judgments. However, the possibility that muscimol infusions changed other functions important for task performance such as reward processing or memory for the mapping between time and choice can not be ruled out.

3.3 Discussion

Attempts to understand the neural mechanisms of time estimation have often focused on continuously evolving population dynamics as a general mechanism for time encoding across the brain (Buonomano & Mauk, 1994; Buonomano & Merzenich, 1995; Mauk & Buonomano, 2004; Karmarkar & Buonomano, 2007; MacDonald et al., 2011; Gershman,

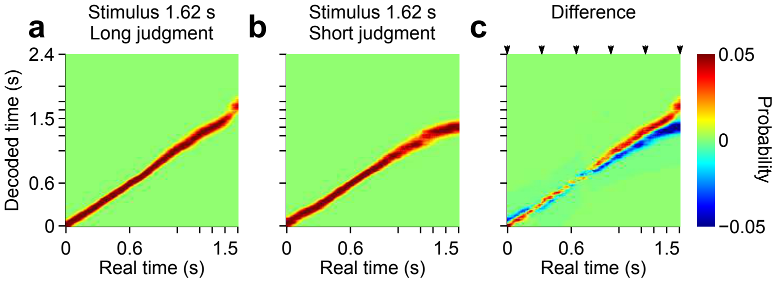

Figure 3.14. **Time decoded from striatal activity during entire stimulus period** (**a**,**b**) Same as in figure 3.13 for the entire interval period. (**c**) Difference between posteriors for long and short judgment trials. Arrowheads indicate same time points used in figures 3.12b and 3.13. n = 433 neurons from 3 rats.

Moustafa, & Ludvig, 2013; Goel & Buonomano, 2014). According to this view, time may be encoded by any reproducible pattern of activity across a population of neurons for as long as the pattern is continuously changing and non-repeating (Buonomano, 2014). However, no study to date has directly compared the speed of such "population clocks" with the duration judgments of the behaving subjects in which they are found. We show that as rats judged the duration of interval stimuli, striatal neurons displayed dynamics in firing rate that contained information about elapsed time. Furthermore, this information was sufficient to account for the animals' perceptual decisions, and was not accompanied by systematic differences in outwardly expressed behavior over time. Combined with the observation that striatal inactivation caused a specific decrement in timing performance, these data suggest that striatal dynamics form a central neural representation of time that guides animals' decisions about duration. Such a coding mechanism in the striatum is well situated to inform the appropriate selection of actions through downstream circuitry involving the globus pallidus, substantia nigra, and various extrinsic connections between the basal ganglia and brainstem, thalamic, and cortical motor areas (Steiner & Tseng, 2010). However, the coding properties tested here could be generally tested in other brain areas where timing signals have been identified such as the hippocampus(Pastalkova et al., 2008; MacDonald et al., 2011), medial prefrontal(Kim et al., 2013; Xu et al., 2014), parietal(Leon & Shadlen, 2003; Janssen & Shadlen, 2005; Jazayeri & Shadlen, 2015) and motor(Lebedev et al., 2008) cortices, among others. Such an approach promises to elucidate where and how time information is used by the brain to support the myriad time-dependent functions we and other organisms rely on for survival.

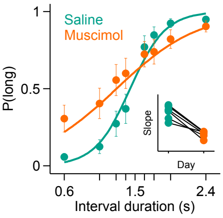

Figure 3.15. **Duration judgments are impaired following muscimol inactivation of the striatum** Averaged psychometric curves following bilateral muscimol or saline injections in dorsal striatum (mean ± standard deviation across session means, and logistic fit; n = 3 rats, 4 sessions each). Inset: slope of psychometric curves on consecutive saline and muscimol sessions.

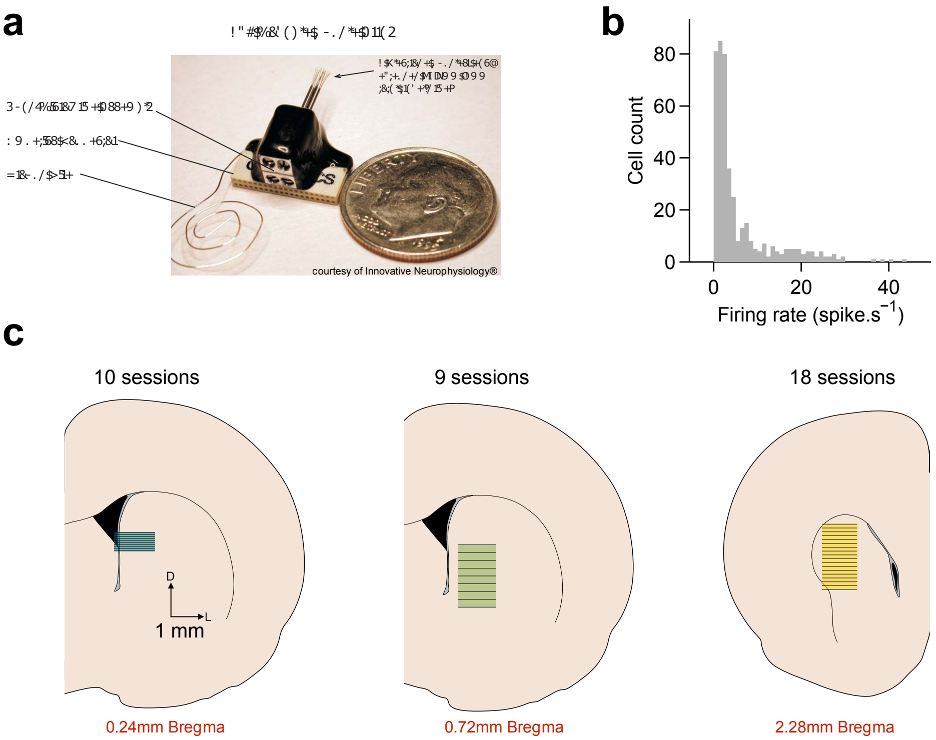

Figure 3.16. (**A**) Movable microwire bundle array (Innovative Neurophysiology) used for all neural recordings. (**B**) Histogram of firing rates for all selected cells (bin size 1 spike/s). (**C**) Schematic representation of the striatal recording sites. Coronal slices at intermediate AP positions are show for reference (left to right, rats Bertrand, Edgar and Fernando). Colored rectangles show the approximate DV position of the wire bundles across recording sessions and horizontal black lines represent session-by-session recording sites, for 10, 9 and 18 recording sessions, respectively.

3.4 Methods

3.4.1 Subjects

Five male Long-Evans hooded rats (*Rattus norvegicus*) between the ages of 6 and 24 months were used for this study. Three rats were used for neural recordings and two rats for pharmacological manipulations. All experiments were in accordance with the European Union Directive 86/609/EEC and approved by the Portuguese Veterinary General Board (Direcção-Geral de Veterinária, project approval 014303 - 0420/000/000/2011).

3.4.2 Behavior

Rats were trained to perform a two-alternative forced choice timing task as previously described (Gouvêa et al., 2014; also in section 2.4.1). Briefly, animals had to categorize time intervals as either long or short by making left/right choices. For each session the animals were placed in a custom made behavioral box containing 3 nose ports and a speaker. Each trial was self-initiated by entry into the central nose port and was followed by a pair of brief auditory tones (square pulses at 7,500 Hz, 150 ms) separated by an interval selected randomly with equal probability out of 8 possible durations (0.6, 1.05, 1.26, 1.38, 1.62, 1.74, 1.95 and 2.4 s). Judgments were reported at two laterally located nose ports. Left responses were reinforced with a drop of water (solenoid valves, Lee Company) after intervals longer than 1.5 seconds, and right responses otherwise. Incorrect responses were punished with a brief white noise burst (150 ms) and a time out. High speed video (120 fps) was collected from above during task performance. Psychometric functions were fitted using equation 2.1 on page 36.

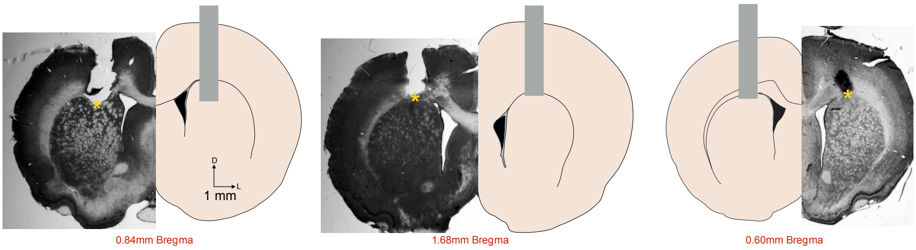

Figure 3.17. (**A**) Histology slices and schematic representation of the location of saline and muscimol injections. Coronal slices at intermediate AP positions are shown for reference at $+0.84mm$ (left, rat Albert), $+1.68mm$ (center, rat Yuri) and $+0.60mm$ (right, rat Zack) from Bregma. Vertical grey bars represent the location of the cannula placements. Yellow asterisks show the approximate DV position from where the injectors extended $1.5mm$ bellow the cannulae.

3.4.3 Electrophysiology

Rats were implanted with 32-channel tungsten microwire moveable array bundles (figure 3.16a, Innovative Neurophysiology) under isoflurane anaesthesia. All recordings targeted dorsal striatum with coordinates centred at $+0.2mm$ AP and $\pm3mm$ medial-lateral (ML) (rat Bertrand), and $+0.84mm$ AP and $\pm2.5mm$ mm ML (rats Edgar and Fernando), from Bregma. Rats were given a week of post-surgical recovery and array placements were confirmed with histology (figure 3.16c). Neural signals were recorded at $30kHz$ during behavior, amplified and band-pass filtered at 250-750 Hz (Cerebus - Blackrock Microsystems). Each independent bundle was moved $50 - 100\mu m$ after every recording session to ensure that independent neural populations were sampled across recording sessions. Waveforms corresponding to action potentials from single neurons were sorted offline using PCA (offline sorter, Plexon). All remaining analysis were run in Matlab (Mathworks). We selected all isolated units with a mean session firing rate $> 0.5Hz$ and from sessions with $> 70\%$ correct performance (averaged across all stimuli) and a minimum of 250 trials ($n = 433$ units). To build PSTH, spikes were counted in 2-ms bins and convolved with a gaussian kernel with 25-bin standard deviation. PSTH in figure 3.3 were ordered by angular position in the space formed by the first 2 principal components describing firing dynamics (i.e., dimensions are all time bins within interval period, samples are each neuron's mean PSTH). This method orders cells with respect to their dynamics while taking into consideration the full response profile over the relevant temporal window, and not just a single response feature such as peak response time (Geffen, Broome, Laurent, & Meister, 2009).

3.4.4 Pharmacology

We implanted 3-mm 20-gauge stainless steel guide cannulas (Belany) bilaterally into the striatum of 2 rats [$+0.84mm$ AP,$\pm2.5mm$ ML, from Bregma, and $-3mm$ DV (from cortex surface) under isoflurane anesthesia. After one week of post-surgical recovery and 4 days of training, rats were injected with either vehicle (saline, PBS 1x) or muscimol (GABA-A agonist, 20 mg/L, Sigma®) solutions in alternate days. Two 1-μL syringes (Hamilton), attached to an injection pump (Harvard Apparatus) through 20-gauge internal cannulas that extended 1.5 mm bellow the guide cannulas, injected $0.6\mu L$ of solution per site during 2.5 min. The internal cannulas were left in place for an additional 1.5 min and the rats were given a 45-min recovery period in their home-cage before starting the task. Cannula placements were confirmed by histology (Figure 3.17).

3.4.5 Preference index

We counted spikes during the last $500ms$ of the stimulus period, and built two separate spike count distributions for short and long judgment trials. Next, we used a ROC analysis to measure the separation between distributions (95% bootstrap confidence interval, 1000 iterations). We then linearly transformed the area under the ROC curve ($a \in [0, 1]$) into a preference index ($i = 2a - 1$; $i \in [-1, 1]$). We adopted the convention that neurons with positive preference indices fired preferentially for long stimuli (figure3.4).

3.4.6 Low dimensional representations of population state

We refer to the vector describing instantaneous firing rates (measured within an integration window) across a population of neurons as the population state. The population state vector is a high dimensional variable (i.e., it has as many dimensions as neurons). With the purpose of visualizing population state in 2d plots, we employed standard dimensionality reduction techniques. In Figure 3.6, left panel, we chose to represent in the abscissa a direction that emphasizes the separability between short and long stimulus trials (i.e. the direction that maximizes variance between groups while minimizing variance within groups; LDA; see below), and in the ordinate the axis of maximal variance that is also orthogonal to the abscissa (i.e., first principal component calculated in the null space of the linear discriminant). In Figure 3.12, population state is represented in the space formed by the first two principal components describing population state, calculated during presentation of the interval for which choice variance is maximal (i.e. dimensions are neurons, samples are averaged spike counts for the time bins within that interval).

3.4.7 Neurometric curves

For each trial in a session we characterized neural population state as a vector $\mathbf{r} = (r_1, r_2, ..., r_N)$, where r_n is the number of spikes fired by neuron $n \in N$ within the last $500ms$ of the interval period in that trial. Next, for all trials but one from each session (training set; leave-one-out cross-validation procedure), we found the linear discriminant that best separated population state vectors according to whether they came from long or short interval trials (LDA). The linear discriminant is given by

$$\mathbf{w} = \text{argmax}\frac{\mathbf{w}^\top S_B \mathbf{w}}{\mathbf{w}^\top S_W \mathbf{w}}$$
$$\mathbf{w} = S_W^{-1}(\mu_1 - \mu_2) \tag{3.1}$$

where $\mathbf{w}$ is the vector of coefficients for the linear discriminant, S_B is the between class covariance, S_W the within class covariance and μ_1 and μ_2 are the means of all points in class 1 and class 2 respectively. A threshold placed along the linear discriminant was then used as a decision rule applied to neural data from the remaining trial (test set). Figure 3.6, left panel, depicts population vectors from an example session (black line: decision rule). We iterated over this procedure until all trials had been tested, thus obtaining for each trial a ‘neural duration judgment’. In analogy with behavioral judgments, we used two parameter logistic fits (equivalent to equation 2.1 with $d = 0$ and $a = 1$) to obtain a quantitative description of the performance of simultaneously recorded neurons in categorizing stimuli — the neurometric function (figure 3.6, orange curve in the right panel).

3.4.8 Videometric curves

Full session videos (256x192 pixels resolution) were cut into 3-s long clips with Bonsai (Lopes et al., 2015). Individual frames from approximately 75 ms before interval offset were used for this analysis. This buffer was added to ensure that all frames used preceded stimulus offset. Images were first represented as vectors composed of individual pixel luminance values. Given that image sequences tend to lie on curved low dimensional manifolds in pixel space (Pless, 2003), any slight differences in behavioral state reflected in images collected at the offset of short and long interval categories are not necessarily expected to be linearly separable. Thus, we employed isomap (Tenenbaum, De Silva, & Langford, 2000), a non-linear dimensionality reduction method, to obtain an information rich yet low dimensional representation of animals’ ongoing behavior. This approach has the advantage over tracking methods that it does not make assumptions as to what part of the animals’ movements might provide information about stimulus category. The neighborhood size, used to compute the shortest paths between data points, was set to 25 frames to minimize, on average, the dimensionality at which the reconstruction error elbow occurred. In analogy with the neurometric curves, for each stimulus type, we then trained a linear discriminant (leave-one-out cross-validation procedure) to classify frames into those that were recorded during trials where a ‘short’ or ‘long’ stimulus interval was presented. The classification was performed in the reduced space determined by isomap. As a positive control for the method, we repeated the same analysis for frames captured at the moment animals expressed their judgment by inserting their snout at one of the two choice ports. Here, the neighborhood size was chosen to be the minimum for which all frames (from a single session) could be included in a single embedding. This analysis was done for all usable videos (8 out of 11) of sessions in the upper tercile with regard to population size.

3.4.9 Time course of classification performance from neural and video data

To compare how the decoding performance using neural and video data evolved over time, the classification analyses described in Neurometric curves and Videometric curves was performed every 100 ms within a one second window centered around stimulus offset. Video frames at the each time point and neural data in a 200 ms time bin terminating at each time point were used for the analysis. This generated neural and video classification curves that described the ability of simultaneously recorded neural ensembles and video frames to correctly classify interval stimuli as long or short (figure 3.9, left panel). To determine the relative timing of classification ability in neural ensembles and behavior, we regressed the neural classification curve against the video classification curve for shifts ranging from -300 ms to 300 ms in 100 ms steps (figure 3.9, left panel).

3.4.10 Psychometric curves split by population state at interval offset

We projected neural activity (composed of simultaneously recorded neurons) on individual trials in high dimensional neural space onto the mean trajectory of those neurons during the delay period. We normalized these projections by the length of the mean trajectory of that group of neurons for the longest interval. Pooling normalized projections over all sessions for each animal, we plotted, for each stimulus, distributions of normalized projections at interval offset. To test whether distance traversed along the mean trajectory is predictive of animals’ perceptual report, we separated the distribution of pooled projections for each stimulus into 3 bins. Psychometric curves were constructed using trials from each bin. To quantify the key differences between each of these psychometric curves, we performed model comparison using the 4 parameter logistic function in equation 2.1. For two of three animals (Bertrand and Edgar), the model that best accounted for the differences between the three curves (based on BIC scores) was one with only horizontal shifts between the curves. In the third animal (Fernando), the model that best fit the data was one in which the fit to the three curves differed in both horizontal shift and slope.

A trial's projection on the mean trajectory can be interpreted as a method for decoding time from neural state. Hence, trials that are outliers in the distribution of projections on the mean could potentially correspond to trials where the animal was disengaged. To remove such trials we defined a fraction (60%) of normalized trajectory around the mode of the distribution of pooled projections for each stimulus and excluded trials with projections outside this window.

3.4.11 Population decoder

We decoded elapsed time from striatal activity using a cross validated, flat prior naive Bayes decoder. First, spikes were counted in 500-ms wide, 10-ms apart overlapping time bins (time referring to the right edge of the bin). For each neuron, we captured cross-trial variability in spike counts at each time bin by building empirical distributions. We did it by computing, for each neuron and time bin, a weighted histogram of spike counts across all correct trials. We defined the weight applied to the spike count observed at a given trial as the choice variance associated with the stimulus presented in that trial. Specifically, weights were defined for each stimulus value s as the product of the probabilities of long and short judgments, i.e., $P(\text{long judgment} \mid \text{stimulus} = s) \cdot P(\text{short judgment} \mid \text{stimulus} = s)$. Histograms were then smoothed using local linear regression (lowess), and normalized to unit area. As a result, near boundary trials had a greater contribution to the final shape of the histograms. Iterating this procedure over all time bins within the interval period produced the conditional probability distribution $P(r_n \mid \text{time})$, a.k.a. the likelihood function. Whenever appropriate (i.e. when decoding from correct trials), leave-one-out cross validation was performed by recomputing the likelihood function with all correct trials but the one being decoded from. Populations of neurons were built by concatenating together trials of same stimulus and judgment type. For each of 100 such trials, posteriors were computed for each neuron with a flat prior, then multiplied across neurons and renormalized to unit area to generate the population posterior.

Acknowledgments

Joseph Paton and Thiago Gouvêa conceived and designed the experiments. Tiago Monteiro and Thiago Gouvêa acquired the data from electrophysiology sessions; Gustavo Mello helped with surgeries. Tiago Monteiro carried the pharmacology experiment. Serkan Sülün and Sofia Soares contributed for data pre-processing. Gonçalo Lopes and João Frazão provided technical support with video acquisition/processing. Joseph Paton, Asma Motiwala, Tiago Monteiro, and Thiago Gouvêa analyzed the data; in particular Asma Motiwala, under the supervision of Christian Machens, led the video analyzes, as well as those presented in figures 3.10 and 3.11. Joseph Paton drafted the manuscript on which this chapter is based (Gouvêa et al., 2015b, 2015a); Tiago Monteiro, Thiago Gouvêa, Asma Motiwala, Sofia Soares, and Christian Machens revised and/or contributed to the final text.

Chapter 4

Is the striatal representation of time action specific?

We have shown in chapter 3 that neural dynamics in the striatum contain a representation of temporal judgments. The task design employed, however, does not allow the disambiguation of two scenarios: I) that striatum represents the passage of time per se, or II) that striatum contains time-varying representations of specific actions. In the current chapter, task variants developed to overcome this limitation are presented, and predictions are made about how striatal activity should behave in the new tasks given scenarios I and II.

4.1 Introduction

We have shown in chapter 3 that neural dynamics in the striatum represents temporal judgments. However, in the interval categorization task employed, each of the two categorical judgments were mapped onto a single action (see task description in section 2.4.1 on page 35). As a consequence, the observations presented in chapter 3 accommodate two distinct scenarios that we would like to distinguish: that striatal activity represents the passage of time per se, or that it contains time-varying representations of specific actions. Disambiguating between these scenarios would require a timing task in which the one-to-one mapping between categorical judgment and action is broken. In the current chapter, two variants of the interval categorization task employed in the previous chapters are presented that break that one-to-one mapping.

4.2 A strategy for breaking the one-to-one mapping between categorical judgment and action

There are certainly multiple behavioral strategies that would enable probing the action-specificity of the striatal encoding of temporal judgments. Figure 4.1 represents the particular way in which this is achieved in all tasks presented in the current chapter: in half of the trials (subtask A), the actions (e.g., left- or rightward orienting movements toward choice ports) and reinforcement contingency (i.e., reward probability contingent on [*stimulus*, *action*] pair) were identical to those in the task presented in chapter 2. In the remaining trials reinforcement was contingent on the same pair of actions, but the mapping between reinforced action and interval category was reversed (subtask B in figure 4.1).

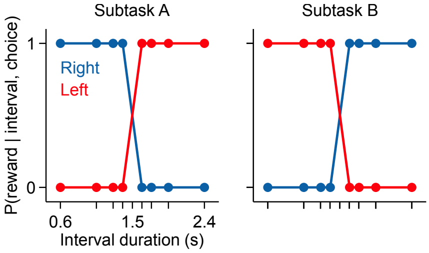

Figure 4.1. **Contingency of an interval categorization task that breaks the one-to-one mapping between categorical judgment and action.** The task is composed of two types of trials governed by the contingencies described here as subtasks A and B. Subtask A is identical to the task described in section 2.4.1, and is referred to as the *left-long* contingency. In subtask B, the mapping between reinforced action and interval category is reversed relative to subtask A; it is referred to as the *left-short* contingency.

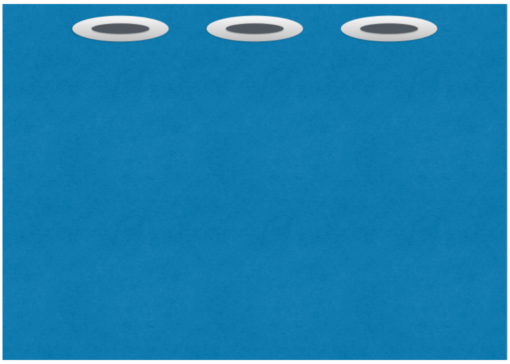

Figure 4.2. **Illustration of behavioral box used both in the original task and in the rule-switching version.** Box consists of three nose ports with infrared beams for detection of nose-poking. Center port is used for initiating trials. Choice is reported at the two laterally located ports.

4.3 Variants of the interval categorization task

Taking as starting-point the task introduced in Gouvêa et al. (2014) and described in section 2.4.1 of this monograph, we developed two different variants that implement both contingencies depicted in figure 4.1 within a single session. The first variant consists of switching between subtasks A and B across blocks of trials. The second one explores a

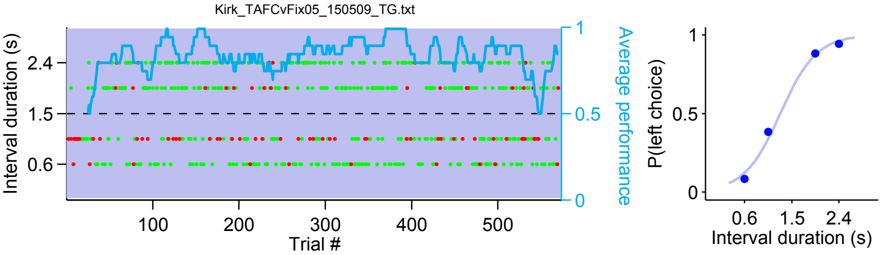

Figure 4.3. **Example session of rat performing the standard interval categorization task.** (*left*) Performance trial by trial. Dots are single trials. Position along x axis indicates trial number, while position in the y axis indicates the interval presented in that trial. Green and red dots denote correct and incorrect trials, respectively. Horizontal dashed line indicates the category boundary. Blue curve is a running average of performance (square kernel, 20 trials long). Blue background means that the *left long* contingency was employed. (*right*) Psychometric curve.

different spatial configuration that implements both subtasks under a fixed allocentric rule.

4.3.1 Rule-switching

The first task variant is a small step away from the original task: both employ the same set up (illustrated in figure 4.2), and the contingencies are identical up until an advanced stage of training. In fact, the first part of training consisted of exactly the original task (see example *left-long* session in figure 4.3). As soon as the animal was able to discriminate intervals to a level deemed satisfactory, the reverse contingency was introduced (see example *left-short* session in figure 4.4). The new contingency was then kept until good performance was reached again, and then switched back and forth in the same way for a number of sessions (figure 4.5 on page 70).

Once a subject had displayed good performance in both contingencies in separate sessions, we moved to a block-design in which both contingencies were presented within the same session. Under this arrangement, rats did successfully perform the two interval categorization subtasks with different mappings between categorical judgment and action, all within the same session (see example session in figure 4.6 on page 71).

Recording striatum activity at this stage would provide evidence about whether or not the encoding of temporal judgments reported in chapter 3 is action-specific. However, good categorization performance under both contingencies within a single session was only observed for sessions in which only easy stimuli[1] were employed. Unfortunately, performance degraded whenever difficult stimuli were introduced (for an example session, see 4.7).

While it cannot be ruled out that animals would eventually reach good categorization performance under the rule-switching task with difficult stimuli, we decided to try a different arrangement in which a given categorical judgment could be reported with either a left- or rightward action while a single, fixed rule is maintained throughout the session. That is made possible by defining the rule in terms of allocentric coordinates.

---

[1] stimulus difficulty is defined as the proximity between stimuli and the category boundary.

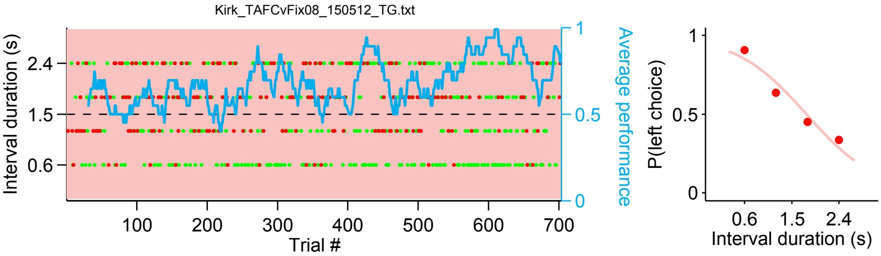

Figure 4.4. **Example session of a rat performing under the *left-short* contingency.** (*left*) Same as in figure 4.3, except that under the *left-long* contingency — as indicated by the red background. This was the third session ever performed by the subject under this contingency. (*right*) Psychometric curve.

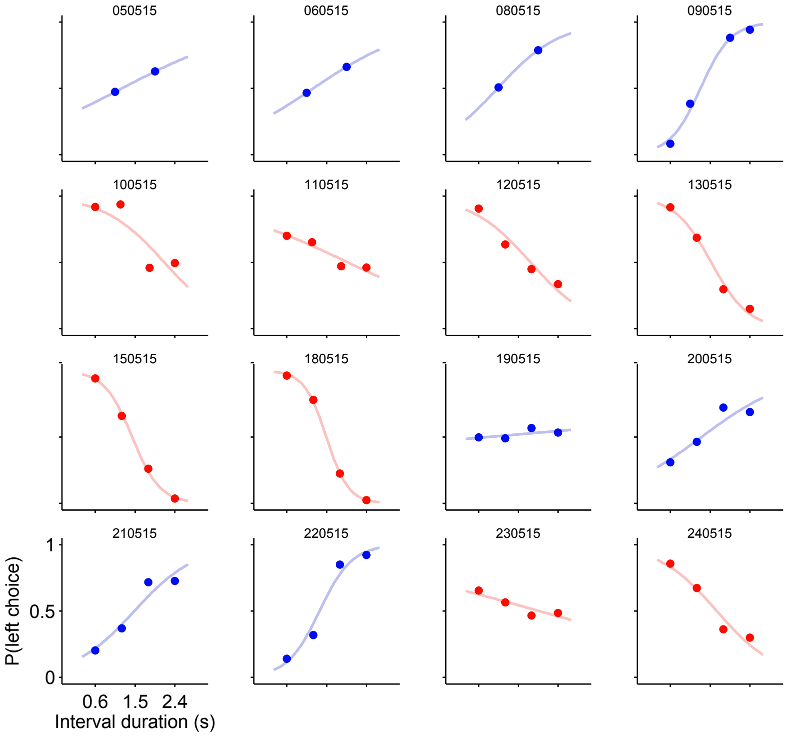

Figure 4.5. **Performance over consecutive sessions with alternating contingencies.** Psychometric curves of subject Kirk over a series of sessions. Blue curves are from *left-long* sessions, and red curves from *left-short* ones.

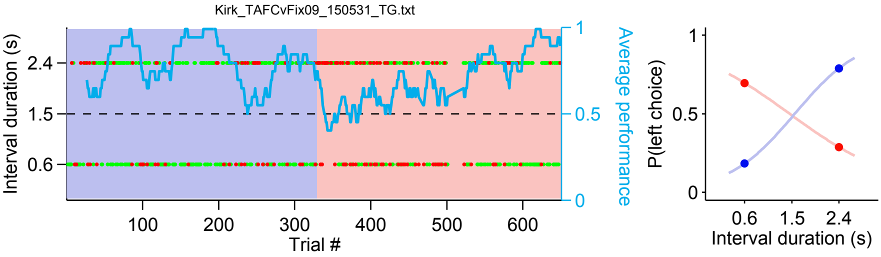

Figure 4.6. **Example session of a rat performing under both contingencies in blocks.** (*left*) Same as in figure 4.3, except both *left-long* and *left-short* contingencies are employed in different blocks — as indicated by blue and red backgrounds, respectively. Notice that all intervals are equidistant to the boundary. (*right*) Psychometric curves, following the same color code.

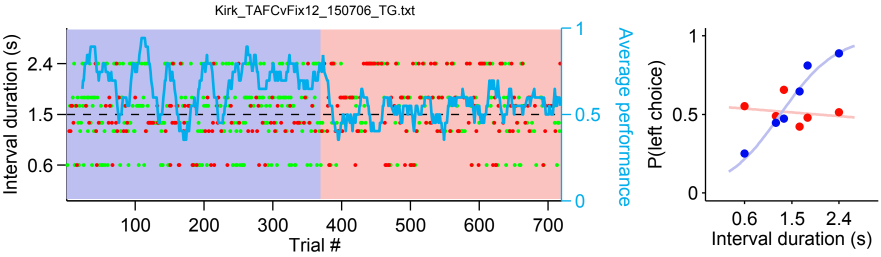

Figure 4.7. **Performance deteriorates in the rule-switching task under multiple difficulty levels.** (*left*) Example session similar to that in figure 4.6, except that stimuli of multiple difficulty levels were employed. The stimulus employed in the session depicted in figure 4.6 corresponds to the easiest of the current set. (*right*) Psychometric curves.

4.3.2 Fixed allocentric rule

While the first task variant employs rule switching to achieve performance under two different mappings between categorical judgments and actions, the second task variant achieves the same feat with a fixed rule in allocentric coordinates. The behavioral box is a modified version of that employed in the original task: the choice ports were moved to the lateral walls, and an additional initiation port was introduced opposite to the first one (see figure 4.8). Throughout a session, choosing a given choice port is only reinforced following stimuli of a single category, as if one of the choice ports were permanently labeled the *long port*, and the other one the *short port*. However, trial initiation can be elicited by one of two initiation ports positioned in opposite walls, and the action leading to a given choice port is different depending on which initiation port was used on that trial. Importantly, the active initiation port is set by the experimenter on a trial by trial basis and in a pseudo-random manner, thus ensuring the balance between different trial types.

As exemplified in figure 4.9, in the current task variant animals can perform successfully both subtasks within a single session even when near-boundary intervals are present.

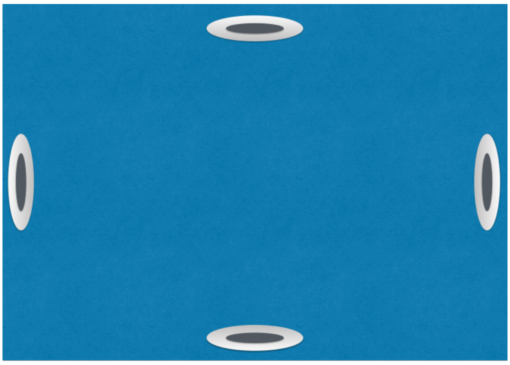

Figure 4.8. **Illustration of behavioral box used in the task with a fixed allocentric rule.** The behavioral box shown in figure 4.2 was modified to allow implementation of both subtasks with a single rule. The choice ports were moved to the lateral walls, and an additional initiation port was introduced in the wall opposite to the first initiation port.

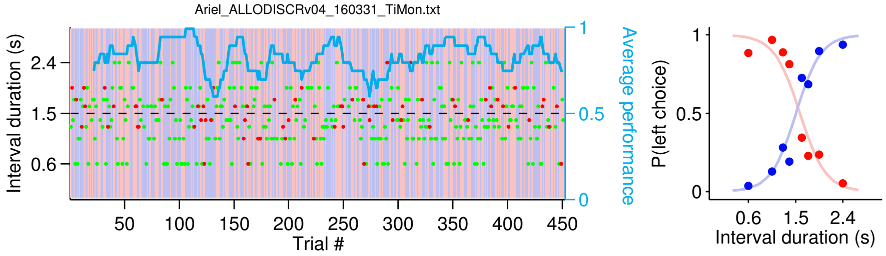

Figure 4.9. **Example session of a rat performing in the second task variant.** (*left*) Example session of a rat performing an interval categorization task in which both subtasks shown in figure 4.1 are implemented under one single rule. The rule is defined in allocentric coordinates. (*right*) Psychometric curves.

4.4 Predictions about striatal activity in the task variants given the two standing hypotheses

In chapter 2, section 2.4.1, an interval categorization task is described. In that task, judgments about the duration of a time interval are reported with one of two actions — one action for each of the two categorical judgments long/short. In chapter 3, we show that populations of striatal neurons display dynamics that contain information about how much time has passed since the onset of the interval being judged. Interestingly, striatal activity also distinguishes between subjective judgments of a same interval. However, given the one-to-one mapping between categorical judgment and actions, it is impossible to disentangle the hypotheses that I) striatal activity encodes a subjective estimate of the passage of time; or II) that striatal activity contains a time-varying encoding of specific actions.

In order to disentangle between the two scenarios, two variants of that task were developed in which, in different trials, the same two actions can be used to report opposing categorical judgments. What should striatal activity look like in one such task under each of the two hypotheses? To illustrate the neural implementation of the two hypotheses, and how the new task variants could be used to disentangle them, the following network simulations were run:

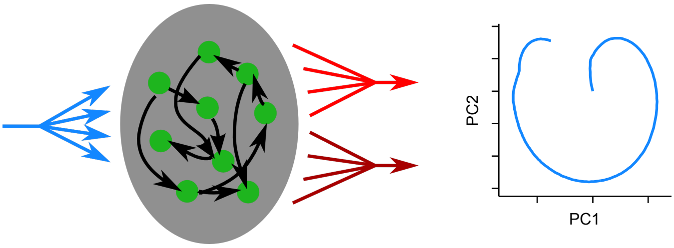

Figure 4.10. **Scheme of a network that performs two distinct interval categorization tasks with the same dynamics.** (*left*) The network is composed of 800 nonlinear units connected sparsely and recurrently. External input sets the network in motion. Activity is read-out linearly by one of two output units. (*right*) Network activity driven by the longest stimulus (see figure 4.11) depicted in principal component space.

4.4.1 Representation of time per se

The first hypothesis — that neural activity in striatum encodes a subjective estimate of the passage of time per se — is illustrated by the network depicted schematically in figure 4.10. The network, composed of nonlinear units connected sparsely and recurrently, was driven by the input represented in light blue. The input consisted of one of four temporally extended patterns: two square pulses separated in time by an interval of

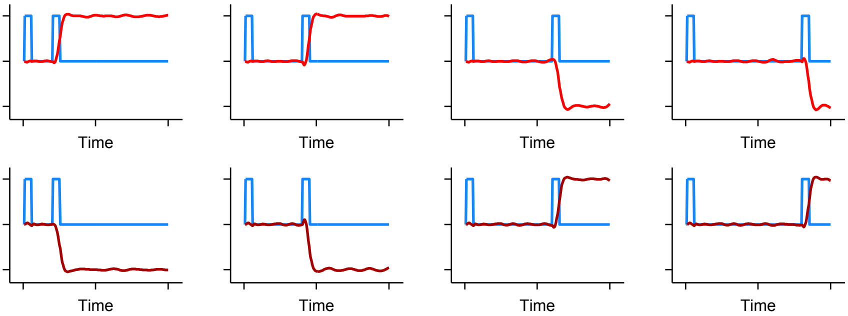

Figure 4.11. **Stimulus form and output of the network from figure 4.10.** Blue curves indicate time series applied as input to the network. Those consisted of two square pulses separated in time by one of four different intervals. Light red (top) and dark red (bottom) traces are linear projections of network activity onto read-out units trained to categorize interval duration by producing an up- or downward step. Each of the two read-out units is performing one of the subtasks described in figure 4.1.

one of four different durations (blue curves in figure 4.11). Network activity in response to the longest stimulus is depicted in principal component space in the right panel of figure 4.10.

Two linear read out units (represented in different shades of red in figures 4.10 and 4.11) were trained to produce upward or downward steps (i.e., a binary decision) in response to input stimuli. Specifically, one of the output units was trained to produce an upward step in response to the two shorter stimuli, and a downward step for the two longer stimuli (figure 4.11, top row). The other output unit was trained to produce the opposite pattern, i.e., downward step for short, and upward step for long stimuli (figure 4.11, bottom row).

In this scenario, a square pulse input sets the network activity into a dynamic trajectory. Since all four intervals begin with the same square pulse, the network will follow the same trajectory until the second pulse is presented. As previously mentioned, as long as the trajectory is stable and non-repeating, it will form the substrate for a representation of time (Buonomano, 2014). In the current implementation, time as represented by network trajectory is read by one of two distinct downstream units, each performing a different interval-action mapping. Importantly, the trajectory followed by the network on any given trial is not affected by the choice of readout unit determining behavior.

As depicted in figure 3.12 on page 54, striatal population activity expresses a non-repeating trajectory during an interval categorization task. Were the striatum to behave analogously to this network, we would expect these trajectories to remain stable when judgments about interval duration are reported with different actions.

4.4.2 Time-varying action representation

The second scenario is illustrated by the network depicted in figure 4.12. In this case, the network was driven by one of two different inputs (different shades of blue in figures 4.12 and 4.13). Inspection of network activity in reduced principal component space reveals distinct dynamic trajectories for the different inputs (right panel in figure 4.12). One single linear read out unit was trained to produce binary responses to the four intervals in a similar way as before, except that now opposing responses were produced by the same read out unit depending on which of the two inputs were used.

As in the previous implementation, this network performs two distinct timing tasks with different interval-action mappings. However, in this case, the different tasks are implemented in distinct activity patterns inside the network, and not as different readings of network activity by downstream areas. Instead, the mapping between network activity and output actions is fixed across tasks.

Were the striatum to behave analogously to this network, we would expect activity to reflect the specific actions used to report duration judgments. In particular, for two distinct timing tasks that mirror each other's interval-action mapping such as in top vs. bottom rows of figure 4.13, we should expect striatal activity to appear time-mirrored in a comparison across subtasks.

4.5 Methods

Neural network simulations

Except where noted, the description below applies to both networks presented in sections 4.4.1 and 4.4.2.

Each of the two networks was composed of non-linear neurons connected recurrently and sparsely. Specifically, $\mathbf{x}_i$, the activation levels of each neuron i, were initialized at 0 and evolved in time as governed by the following equation:

$$\tau \dot{\mathbf{x}}_i = -\mathbf{x}_i + \sum_{j=1}^{N} W_{ij}\mathbf{r}_j + \mathbf{w}_i^{in} I_t \tag{4.1}$$

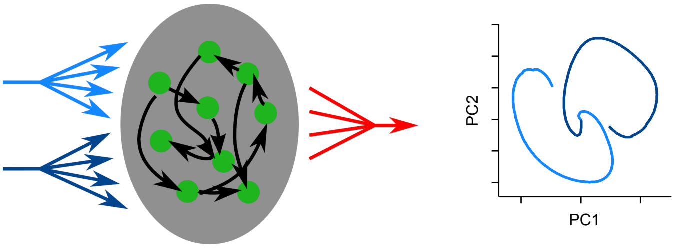

Figure 4.12. **Scheme of a network that performs two distinct interval categorization tasks with fixed activity-output mapping.** (*left*) A network similar to that in figure 4.10, except that external input can be applied through one of two different sets of input connections, and activity is read-out linearly by a single output unit. (*right*) Principal component space depiction of network activity driven by the longest stimulus (see figure 4.13) applied via each of the two sets of input connections.

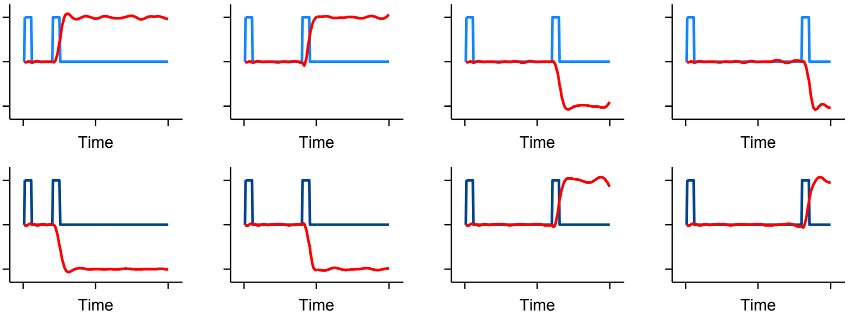

Figure 4.13. **Stimulus form and output of the network from figure 4.12.** Similarly to figure 4.11, light (top) and dark (bottom) blue curves are the time series applied as inputs to the network via each of the two sets of input connections. Only one of the inputs is active on any given trial Red traces are linear projections of network activity onto a read-out unit trained to categorize interval duration by producing an up- or downward step. The same read-out unit performs either of the two subtasks described in figure 4.1 depending on which of the inputs is active.

where $\tau = 10$, W is the connectivity matrix, $N = 800$ units, the firing rate $\mathbf{r}_j = \tanh \mathbf{x}_j$, $\mathbf{w}^{in}$ is a vector of input weights (see difference between networks below), and I_t is the input provided at time t.

The weight of the connection from neuron j onto neuron i was defined as

$$\begin{cases} W_{ij} \sim \mathcal{N}\left(0, \dfrac{g}{\sqrt{0.1N}}\right) & \text{with probability } 0.1 \\ W_{ij} = 0 & \text{with probability } 0.9 \end{cases}$$

where g is a gain factor set to 1.8. Additionally, all self connections were set to 0 (i.e., $W_{ij} = 0 \quad \forall\, i = j$).

The network in section 4.4.1 contains one single set of input connections denoted by the vector $\mathbf{w}^{in}$. Vector components describing the input connections onto each neuron i were set to independent random samples from the standard normal distribution (i.e., $\mathbf{w}_i^{in} \sim \mathcal{N}(0,1)$). The network in section 4.4.2 contains two different sets of inputs, $\mathbf{w}^{in}$ and $\mathbf{v}^{in}$. Vector $\mathbf{v}^{in}$ is generated identically but independently of vector $\mathbf{w}^{in}$. On any given trial, only one of the two input connection vectors were employed.

The input I is a 200 time step long time series composed of two square pulses (10 time steps long, amplitude = 1, baseline = 0) separated in time by one of four different intervals (see figures 4.11 and 4.13).

The network in section 4.4.2 contains one single output unit, and its output at time t was defined as

$$\mathbf{z}_t = \sum_{i=1}^{N} \mathbf{x}_i^{(t)} \mathbf{w}_i^{out} \tag{4.2}$$

where $\mathbf{x}_i^{(t)}$ is the activation level of neuron i at time t, and $\mathbf{w}^{out}$ is a read-out vector. Components of vector $\mathbf{w}^{out}$ were initialized as random samples from a Gaussian distribution with mean 0 and standard deviation $(\sqrt{N})^{-1}$, and adjusted by gradient descent on mean squared error so that the output of the network to both inputs matched the provided step-function templates. Target templates consisted of a step function with value 0 from beginning until the onset of the second input pulse, and then value 1 or -1, depending on the combination of input source and the duration of the input interpulse interval (see trained output in figure 4.13).

The network in section 4.4.1 contains a single input source and two separate output units, each of which was independently trained in a manner similar to that described in the paragraph above. Given that this network is noiseless and the input source is unique, both units were trained on the exact same network high-dimensional pattern, but with inverse target templates (see trained output in figure 4.11).

Acknowledgments

The task variant with a fixed allocentric rule emerged from discussions with Thomas Akam and Joseph Paton. Data on that variant was acquired by Tiago Monteiro, Filipe Rodrigues, and Thiago Gouvêa. All the remaining work was performed by Thiago Gouvêa.

Chapter 5

Discussion

5.1 Explaining behavior

What is the cause of intelligent behavior? If you ask a neuroscientist, they are likely to answer that it lies on the architecture of the brain (e.g., on the particular structure neural activity takes given a cognitive challenge, Douw et al., 2011). If on the contrary, the question is directed to an experimental psychologist, the answer might focus on how different behavioral responses have been reinforced in the past, and under which circumstances. These answers can be both correct without a conflict, as long as it is understood that *cause* is a word with more than one meaning. For example, under the semantic framework of the four causes proposed by Aristotle (apud Falcon, 2015), perhaps it could be said that, while neurobiology is concerned with *material* and *formal* causes of behavior (i.e., the matter it consists of, and the form assumed by matter in it, respectively), behavioral sciences address its *efficient* and *final* causes (related to its history and intent, respectively). In any case, an explanation of behavior with a satisfactory level of completeness should encompass both types of causes — a challenging, but certainly feasible endeavor.

5.1.1 Explaining behavior does not require determining the 'architecture of the mind'

Within experimental psychology, the Experimental Analysis of Behavior (EAB)[1] is one of the major paradigms. Explanations of behavior provided by EAB are elaborated by exploring regularities found between behavioral responses, the environment in which they occur, and the consequences they bring about. Importantly, this explanatory arc makes no reference to the internal structure of the behaving organism. On this regard, EAB contrasts with the major opposing paradigms, i.e., cognitive psychology and psychoanalysis, both of which evoke some sort of architecture of the mind to explain behavior. In line with the behaviorist doctrine, it can be said that, while mechanisms producing behavior are obviously a legitimate object of inquiry, formulations in terms of mental entities are bound to be underconstrained by behavioral evidence. As a consequence, the operations at the heart of the mechanism cannot be verified, and any resulting explanation of behavior is bound to be unsatisfactory. In the words of Skinner (1938, p. 20):

> Behavior has that kind of complexity or intricacy which discourages simple description and in which magical explanatory concepts flourish abundantly. Primitive systems of behavior first set the pattern by placing the behavior of man under the direction of entities beyond man himself. (...) In more advanced systems of behavior, the ultimate direction and control have been assigned to entities placed within the organism and called psychic or mental. Nothing is gained by this stratagem because most, if not all, of the determinative properties of the original behavior must be assigned to the inner entity, which becomes, as it were, an organism in its own right. (...) Some conceptions of the 'mind' and its faculties, and more recently the 'ego,' 'super-ego,' and 'id,' are examples of inner agents or organisms, designed to account for behavior, which have remained the subject of scientific investigation.

---

[1] EAB is also commonly known as Behaviorism. However, prominent behavioral scientist Skinner draws a distinction between Behaviorism and EAB. According to Skinner (1974, p. 3), "Behaviorism is not the science of human behavior; it is the philosophy of that science. Some of the questions it asks are these: Is such a science really possible? Can it account for every aspect of human behavior? What methods can it use? Are its laws as valid as those of physics and biology?"

5.1.2 Neurophysiology fills the mechanistic gap of behavioral science

By avoiding reliance on ill-defined, 'mentalist' entities, focusing instead on the relations between behavior and the environment, EAB exerts parsimony. The result is a theory that allows testable predictions about behavior, entailing applications to everyday life (Skinner, 1965). This theory, however, leaves unsolved the problem of explaining the mechanism producing behavior. EAB leaves this task for a different field of knowledge to fill — that of neurophisiology. According to Skinner (1974, p. 236–7):

> New instruments and methods will continue to be devised, and we shall eventually know much more about the *kinds* of physiological processes, chemical or electrical, which take place when a person behaves. The physiologist of the future will tell us all that can be known about what is happening inside the behaving organism. His account will be an important advance over a behavioral analysis, because the latter is necessarily "historical". (...) [The neurophysiologist] will be able to show how an organism is changed when exposed to contingencies of reinforcement and why the changed organism then behaves in a different way, possibly at a much later date. What he discovers cannot invalidate the laws of a science of behavior, but it will make the picture of human action more nearly complete.

In the decades past since Skinner wrote these words, "instruments and methods" have indeed been devised, and neurophysiology has advanced to the point of revealing "what is happening inside the behaving organism". In fact, in that same decade, neurophysiology helped settle one of the biggest debates in the history of psychology[2].

A neural representation of space

The proposed a debate opposed, supposedly, two distinct behavioral theories of spatial cognition. Tolman, Ritchie, and Kalish (1946) trained animals to navigate a maze composed by a number of corridors arranged serially and connected by right angle turns that led to a chamber containing a reinforcer. Animals were then tested on a modified version of the maze in which new paths were made available in the form

---

[2] While it is widely accepted that explaining the material mechanisms of behavior lie within the scope of neurophysiology, modern behavioral scientists and philosophers of science worry that neurophysiologists might be failing to identify the operations that underly the proposed mechanisms (Bechtel, 2005), often falling on linguistic traps such as the mereological fallacy (Bennett & Hacker, 2003; Bechtel, 2005)

of an array of parallel corridors running out from the initial location in a number of different directions. Importantly, the path running in the direction of the initial segment experienced during training was blocked. The authors found that animals would choose very often the corridor running in the direction leading to where the reward chamber was located during training.

Effectively, the experiment tested how well can animals generalize spatial trajectories. However, according to Tolman's interpretation, these results show the necessity of postulating an inner representation of space in order to explain behavior, thus ruling out EAB as an inadequate paradigm. In his words (Tolman, 1948, p. 189–92)[3]

> [T]here is a school of animal psychologists which believes that the maze behavior of rats is a matter of mere simple stimulus-response connections. (...) The rat's central nervous system, according to this view, may be likened to a complicated telephone switchboard. There are the incoming calls from sense-organs and there are the outgoing messages to muscles. (...) Learning, according to this view, consists in the respective strengthening and weakening of various of these connections; (...) Thus it is as if — although this is certainly not the way this subgroup would themselves state it — the satisfaction–receiving part of the rat telephoned back to Central and said to the girl: "Hold that connection; it was good; and see to it that you blankety-blank well use it again the next time these same stimuli come in." These theorists also assume (at least some of them do some of the time) that, if bad results — 'annoyances,' 'negative reinforcements' — follow, then this same satisfaction-and-annoyance receiving part of the rat will telephone back and say, "Break that connection and don't you dare use it next time either."

From the statement put forth by EAB that *the probability of a certain behavior being emitted is a function of the consequences it entailed in past occasions*, and presumably by assuming *behavior* can only be defined in egocentric coordinates, the author derives the implication that animals should learn exclusively a sequence of actions such as left and right turns[4]. Tolman offers the alternative hypothesis that the animal builds a *cognitive map* during learning that will allow reaching the same location following a different series of actions.

---

[3] The fragility of Tolman's interpretation of EAB is evident in his characterization of learning in structural, biological terms.

[4] In the experiment presented by Tolman et al. (1946), the behavior of pursuing the trained path is extinguished by means of blocking that passage. Spatial generalization, in this case, was likely aided by the well known effect of extinction in promoting exploration.

While Tolman's neglect of alternative behavioral explanations consistent with EAB[5] is of little relevance for the scope of the current monograph, the question about the neurophysiological mechanism by which the behavioral result comes about is very pertinent.

The behavior displayed on the test session is clearly influenced by the location rewarded in the past. If the animal shows the ability to behaviorally discriminate spatial locations, then the brain, as an essential component of the material cause of behavior, must contain distinct activity patterns that somehow correspond to the discriminated locations. Furthermore, the location-specific activity must be translated into appropriate musculoskeletal commands to implement behavior.

It was only decades after Tolman published his behavioral results that O'Keefe and Dostrovsky (1971) found a neurobiological instance of spatially-selective neural activity in the firing of hippocampal neurons (findings summarized and discussed in O'Keefe & Nadel, 1978). These findings established that the information necessary for spatial planning and generalization was present in the brain, offering for the first time the cornerstone of a legitimate mechanistic hypothesis to explain spatial learning and generalization. However, it should be noted that feature-selective activity, while intuitively useful for generating discrimination behavior, is certainly not sufficient — the ways in which this activity is read out and used by other brain areas is still not fully understood.

5.1.3 On neural representations

It is often said that feature-selective neural activity constitutes a neural *representation* of these features. It might be said, for example, that place-selective hippocampal neurons, as those presented by O'Keefe and Nadel (1978), constitute a spatial representation; or that the delay-selective striatal neurons presented in figure 3.3 on page 46 constitute a representation of elapsed time.

The concept of representation used in cognitive neuroscience is paralleled in other fields that in some way deal with organized knowledge - namely symbolic logic and Artificial Intelligence (AI).

---

[5] Namely, that the behavior learned during the training phase might be defined in allocentric coordinates, rendering the topology displayed during the test phase as conceptually surprising as that of a rat who, having learned to lever-press, alternates between left-forepaw and right-forepaw presses. From this angle, the problem at hand is that of the mechanism by which multiple topologies get grouped under the same behavioral class.

Knowledge representation in formal systems

> "[L]ogic consists of a formal or informal language together with a deductive system."
>
> (Shapiro, 2013)

In formal logic, a state of affairs can be represented by a set of symbols under a formally defined syntax. Along with this representation (or, set of statements describing the state of affairs), a deductive system is provided in the form of a set of rules for combining statements. By following those rules, new statements can be deduced. These new statements, while necessarily contained in the original representation in some way, are made explicit and evident through the deductive process.

Analogously, an AI agent stores knowledge — either hand-coded or transduced from the environment by its sensors — in the form of representations. For classical AI agents built based on the logical approach, knowledge is represented in a formal and explicit manner (Nilsson, 1991; Thomason, 2014). In fact, devising formal languages for knowledge representation — either general-purpose or application-specific — constitutes an important branch of AI (Russell & Norvig, 2003, chapters 8 and 10). With the representations established, the behavior of a classical AI agent is generated as the product of rule-based, deduction-like operations performed on them.

Distributed knowledge representation

While the drive to explain intelligent behavior lies at the heart of the study of logic and perhaps AI, the way classical AI agents operate deviates from that of biological organisms. However, an alternative paradigm within AI offers an alternative.

In contrast with the symbolic representations employed by classical AI, agents from the *connectionist* school represent behaviorally relevant variables as a distributed pattern over the activation of the component units of an Artificial Neural Network (ANN) (Rumelhart, Hinton, & McClelland, 1988). The system of rules governing how representations are transformed to generate their output (akin to a deductive system) is inspired in the architecture of the animal nervous system, and is instantiated in the connections between the units composing the ANN. Therefore, concocting an AI agent of this sort amounts to identifying (or "learning") appropriate connectivity patterns; developing methods for that concentrates a fair amount of the effort in this field (e.g., Rumelhart, Hinton, & Williams, 1986).

Within this paradigm, representations are vectors describing network state that evolve as a dynamical system governed largely by the pattern of connections between units. While the homology with the logical rules followed by classical a AI agent is not evident, the connectionist paradigm has proven very successful in a diverse set of tasks ranging from language translation (Sutskever, Vinyals, & Le, 2014) to playing video-games (Mnih et al., 2015).

Identifying representations in biological brains

It is widely accepted within the neuroscience community that the brain, as the material mediator of behavior, must contain activity that reflects the relevant environmental events that shape behavior — i.e., it must *represent* variables known to determine behavior. However, contrary to formal languages and AI agents, the brain is not human designed, and the syntax, semantics, and deductive system governing its *modus operandi* must be discovered by experimentation and inductive reasoning.

A representation is commonly understood as neural activity that I) contains information about the environment; and II) is used to generate behavior.

A common strategy to identify neural representations of sensory stimuli consists of presenting stimuli repeatedly while recording neural activity[6], often in anesthetized subjects (e.g., Hubel & Wiesel, 1962). A brain region that represents stimuli would be expected to produce different patterns that correspond to different stimuli. This strategy allows the identification of neural activity that satisfies condition I listed above.

This strategy, however, would fail to discriminate behaviorally-relevant components of stimulus-evoked neural activity from epiphenomenal agitations. This is particularly problematic in the investigation of neural representations of time: any neural pattern that changes over time in a non-repeating manner — a phenomenon expected to occur by chance in high-dimensional dynamical systems such as neural networks — would satisfy that condition (Buonomano, 2014). Yet, experiments of this sort have been used to support the claim that striatum contains a representation of time (Jin et al., 2009; Adler et al., 2012). These results, while certainly consistent with the claim, only provide weak evidence in its support.

---

[6] In a formal language, the semantics of a statement relates to its truth-condition — i.e., in what conditions it is said to be true of false. In particular for statements describing states of affairs (as opposed to tautologies), their truth-condition is given by empirical verification. Analogously, if a neural pattern in some primary sensory brain area is taken to be a representation of a sensory stimulus, it could be said that its truth-condition is given by whether the corresponding stimulus is present or absent. The correspondence between neural patterns and sensory stimuli, however, is not given and must be discovered experimentally. The experimental strategy for discovering this correspondence can, therefore, be understood as observing the neural pattern displayed by the system when the corresponding truth-condition is satisfied (i.e., when the stimulus is present).

While identifying brain signals that covary with stimuli (satisfying condition I) provides some evidence that a neural pattern constitutes a representation, the case can be made stronger by putting behavior under stimulus control and testing whether, for a fixed stimulus, fluctuations in the neural pattern covary with behavior (Parker & Newsome, 1998). Naturally, the type of neural fluctuations expected to covary with behavior are those mimicking behaviorally-relevant changes in the state of affairs — in other words, behavioral output should correlate with neural variability along stimulus dimensions.

An experimental strategy broadly employed in the study of sensory representations consists of training subjects to perform one of two behavioral responses depending on whether a unidimensional, real-valued stimulus is above or below a categorical threshold (e.g., Britten et al., 1992; Romo et al., 1999; Hanks et al., 2015). This approach is powerful because it provides comparable quantitative assessments of discrimination ability in both behavioral and neural levels (e.g., Mountcastle et al., 1969; Britten et al., 1992). Furthermore, if neural and behavioral data are obtained simultaneously (as dictated by modern standards), co-fluctuations between stimulus-evoked neural activity and stimulus-controlled behavior can be verified directly (Nienborg et al., 2012).

While this type of correlational study provides valuable evidence to test whether a certain neural activity pattern is used to generate behavior (condition II above), additional evidence can be provided by experimentally manipulating neural activity.

A complete explanation of the role played by a candidate neural representation in a given behavior should include an understanding of how this neural activity pattern interacts with that of other brain areas in the process of transforming information so as to ultimately generate (or deduce the appropriate) motor output[7]. The experimental paradigm described here — i.e., characterizing neural activity evoked by a state of affairs and observing its correlation with behavior — do not suffice to provide such a complete explanation. Instead, it seems that the prevalent experimental methods are insufficient, and that achieving that level of understanding will require technical innovation.

5.2 A representation of time in the striatum

The striatum is a forebrain nucleus that integrates the brain system known as the basal ganglia, of which it is the main input structure (Steiner & Tseng, 2010). It forms a

---

[7] Assuming the behavior in case is not covert (Skinner, 1974; Catania, 1999). In that case, perhaps "motor output" could be replaced by other long-lasting representations that, I would suppose, underlie covert behavior.

characteristic looped connectivity architecture with the thalamo-cortical system — it receives inputs from a broad set of cortical areas in a topologically organized manner, and it sends indirect projections back to those areas via the output nuclei of the basal ganglia and the thalamus (Alexander, DeLong, & Strick, 1986). Similarly to the cortical connections, loops exist with subcortical motor structures in the brainstem such as the superior colliculus (McHaffie, Stanford, Stein, Coizet, & Redgrave, 2005; Hikosaka, Takikawa, & Kawagoe, 2000).

As reviewed in section 1.4.1 on page 18, the striatum has been implicated in interval timing behavior by several lines of evidence; however, the specific role it performs — the particular operations it is reponsible for (Bechtel, 2005) — is yet unclear. One of the standing hypotheses is that it implements a neural representation of time. In support of this hypothesis, potentially time-encoding activity has indeed been observed in the striatum (Matell et al., 2003; Jin et al., 2009; Adler et al., 2012; Mello et al., 2015). Furthermore, time-encoding striatal activity has been shown to adjust to coarsely different temporal contexts (Mello et al., 2015). However, evidence that fine fluctuations in "striatal time" covary with subjective estimates of time was still missing.

5.2.1 A strengthened hypothesis

In the current monograph, the hypothesis that striatum contains a representation of time is tested with an approach similar to that described above. Time, while not constituting a sensory modality per se, is still addressable by the same methods and procedures: animals were trained on an interval discrimination task (general paradigm ans specific task described in sections 1.2.3 and 2.4.1, respectively), and striatal activity was observed and perturbed during task performance. As presented and discussed in chapter 3, the finding that striatum displays time-encoding activity during behaviorally relevant intervals was reproduced. Furthermore, the discrimination task employed revealed that the population pattern evolved faster or more slowly in statistical dependence of the subject's behavioral decision. Lastly, the ability to discriminate time intervals was impaired after transient pharmacological inactivation of the striatum, implying that this brain structure is necessary for performance of interval timing behavior.

These results strengthen the hypothesis that striatal activity constitutes a neural representation of time during performance of interval timing behavior. Alternative hypotheses, however, cannot be completely ruled out. One possibility discussed in chapter 4 is that the striatum activity observed in this task reflects an instantaneous representation of the behavioral response more likely to be emitted — i.e., a time varying representation of action values. Observing how striatal activity behaves in the task variants introduced in that chapter should disambiguate whether it represents instantaneous action-values or time per se — in other words, whether or not time is encoded in striatal population activity in an action-specific manner.

5.2.2 Origins of striatal dynamics

A complete understanding of the role played by the striatum in interval timing would require understanding how the temporal representation it contains interacts with other relevant representations in different brain areas.

The striatum receives input broadly from the cortex, in a topologically organized manner (Alexander et al., 1986; Allen Institute for Brain Science, 2015). Different cortical areas are known to display activity with temporal regularities on different timescales (Bernacchia, Seo, Lee, & Wang, 2011), and different timescales are distributed across cortex in a hierarchically organized manner, "with sensory and prefrontal areas exhibiting shorter and longer timescales, respectively" (Murray et al., 2014, p. 1). The striatal representation of time consists of a firing pattern that is temporally diverse across the population, and it is known that a gradient of cortical input — from associative to sensorimotor — exists along the striatal mediolateral axis (Allen Institute for Brain Science, 2015). Furthermore, the striatum is known to display an interesting functional organization along the mediolateral axis — while dorsolateral activity is implicated in stimulus-response associations (a.k.a. habitual behavior), the dorsomedial part is implicated in more flexible, response-outcome associations (a.k.a. goal-directed behavior) (Daw, Niv, & Dayan, 2005; Balleine, Liljeholm, & Ostlund, 2009). This functional specialization is at least partially explained by the different sources of cortical input to these areas. One interesting possibility is that also the striatal representation of time is generated by extracting temporal regularities from cortical activity in an anatomically organized manner. That possibility could be verified if cortical activity were recorded simultaneously with striatal activity during task performance — or, on a coarser way and given the knowledge about the mediolateral gradient, simply by systematically exploring this anatomical dimension. In the present study, striatal activity was probed at a dorsocentral location, and results are therefore inconclusive on this regard.

5.3 Two roles for a representation of time in producing behavior

5.3.1 Associative learning

A cornerstone of the explanation of adaptive behavior provided by EAB — or by any other major behavioral theory for that matter — is the observation that the probability of emission of a given behavioral response is modified by its consequences. Consequences that increase the probability of a behavior to reoccur are said to reinforce that behavior, while consequences that reduce that probability are said to punish that behavior (Skinner, 1974; Catania, 1999). However, in natural situations, environmental events occur in a continuous stream. Every behavior is likely to be followed by a very large number events, and identifying those that were directly caused by the emitted behavior — and therefore could be legitimately referred to as its consequences — constitutes a problem known as the *credit assignment problem.*

While the notion that intelligent behavior arises from the association of contiguous ideas has been around at least since the English empiricists (and probably much earlier; Gormezano & Kehoe, 1981), the modern *associative learning* paradigm was perhaps inaugurated by Pavlov. Since then and up until the 1960's, the intuitive idea prevailed that associations would be learned between events that occurred contiguously in time. This notion changed after a series of experiments by Rescorla (for reviews, see Rescorla, 1988a, 1988b).

In these experiments, by keeping constant within sessions the number of contiguous presentations of a neutral stimulus and a reinforcer while varying the degree to which stimulus presentation on a given trial was informative about reinforcer occurrence [i.e., the ratio between the baseline $P(\textit{reinforcer})$ and $P(\textit{reinforcer} \mid \textit{neutral stimulus})$], Rescorla showed that contiguity is not enough to establish an association; rather, associations are only learned when reinforcer is made *contingent* on stimulus presentation (Rescorla, 1966, 1967). It is as if the dogs and rats used as subjects in the experiments were employing counterfactual reasoning to infer causality between neutral stimulus (A) and the reinforcer (C): the association is only learned when, not only have C occurred when A occurred but, "[i]f A had not occurred, C would not have occurred" (Menzies, 2014).

These behavioral results are captured by an incremental learning model developed by Rescorla and Wagner (1972) in which the value of a behavioral response is adjusted to match the value of its consequences, weighted by the probability of that consequence

— in other words, associations are learned in proportion to the "surprise" elicited by consequences. Importantly, incremental learning is due to iteration made possible by the particular symbolic representation employed by the model that consists of dividing time into discrete units (i.e., trials). However, the particular choice of a trial as the basic unit of time imposes important limitations to the model, as this choice of representation naturally renders the model blind to differences in intra-trial structure.

One important effect missed by the model is seen in trace conditioning, in which the difficulty of learning an association between temporally separated stimuli depends on the inter-stimulus interval (ISI), *mutatis mutandis.* This difficulty is overcome by the Temporal-Difference (TD) method proposed by Sutton (1988), a variant of the incremental learning model by Rescorla and Wagner (1972) in which time is split in finer discrete units, dispensing with the notion of trials — a representation that feels more natural given that not every aspect of real life is organized in discrete trials. Additionally, learning in the TD method is driven by differences between reward predictions generated moment by moment, unlike in the model by Rescorla and Wagner (1972), where learning is driven exclusively by differences between predicted and *observed* consequences.

The usual choice of temporal representation used in models employing the TD method, known as Complete Serial Compound (CSC), represents time as a series of basis functions shaped as fine and non-overlapping pulses — in resemblance with a digital clock (Sutton & Barto, 1990). While this allows for great discrimination between different moments in time (e.g., an agent using a representation of this sort with sufficient resolution would perform very well on an interval discrimination task), it has the disadvantage of poor temporal generalization — as reflected in the large number of trials required for a reward signal to backpropagate to a reward-predictive stimulus. Furthermore, the high precision of this sort of representation does not appear to be biologically plausible. Alternative temporal representations for learning systems have been proposed that overcome these problems; these consist of partially overlapping basis functions that resemble smeared versions of a CSC (e.g., Grossberg & Schmajuk, 1989; Ludvig et al., 2008). Interestingly, dynamics of delay period activity displayed by neural populations in a number of brain areas resemble this pattern (Pastalkova et al., 2008; Jin et al., 2009; MacDonald et al., 2011; Adler et al., 2012; Harvey et al., 2012; Mello et al., 2015; Gouvêa et al., 2015a).

One last problem is left unsolved by this class of models: behavioral results show that the dependence of associative learning on ISI duration is actually better characterized as a dependence on the ratio between ISI and the inter-trial interval (ITI) (Terrace, Gibbon, Farrell, & Baldock, 1975; Gibbon & Balsam, 1981; Balsam & Gallistel, 2009) — in other words, learning of associations is time-scale invariant, perhaps constituting one more instance of the scalar property of interval timing (Gibbon, 1977). To capture this effect, it would be required that temporal representations scaled to adjust to temporal context. Interestingly, this property has been observed in candidate neural representations of time (Romo et al., 1999; Xu et al., 2014; Mello et al., 2015). Among formal models, one class that encodes stimulus history in the frequency domain, in a time-scale invariant manner, might be able to deal appropriately with these behavioral scenarios (Howard & Kahana, 2002; Shankar & Howard, 2012).

5.3.2 Time as a feature characterizing the environment

Besides constituting a representation useful for predicting consequences of stimuli or behavioral responses (e.g., for inferring that "*action A is consistently followed by consequence C, and consequence C rarely occurs in the absence of action A*"), temporal maps are also useful as features defining stimuli and behavioral responses themselves (e.g., the following distinct interval-timing behaviors: "*behavioral response A' consists of two key-pecks 100-ms apart, while reponse A'' consists of two key-pecks 600-ms apart*").

Interestingly, the same types of representations (e.g., the temporally smeared microstimuli representation proposed by Ludvig et al. (2008) and briefly described in section 1.3.4 on page 16) can be used both for learning behaviorally relevant consequences (when coupled with a learning strategy such as the TD method) and for producing interval timing behavior (when coupled with a Reinforcement Learning architecture such as actor-critic; Sutton & Barto, 2015). As put by Gershman et al. (2013):

> When combined with the microstimulus representation, the actor-critic architecture naturally gives rise to timing behavior: in the peak procedure, on average, responding will tend to increase toward the expected reward time and decrease thereafter. (...) Importantly, the late microstimuli are less temporally precise than the early microstimuli, in the sense that their responses are more dispersed over time. As a consequence, credit for late rewards is assigned to a larger number of microstimuli. Under the assumption that response rate is proportion to predicted value, this dispersion of credit causes the timing of actions to be more spread out around the time of reward as the length of the interval increases, one of the central empirical regularities in timing behavior [the scalar property].

An example situation in which producing appropriate behavior requires characterizing behavioral responses based on their timing is that of the bird from section 1.1.1 on page 3. In situations like that, the reward-maximizing strategy is to stay on a patch until the rate of intake drops below the average rate for the habitat, when the bird should then move to a different habitat (Charnov, 1976). A formal, algorithmic solution to this problem would therefore consist in calculating the average intake rate for a habitat, as well as local, patch-specific intake rates. Calculating rates would naturally require estimating elapsed time. Once these estimates are obtained, behavioral decisions such as that faced by the bird can be generated by a simple comparison mechanism. It has been shown in one instance that a learning strategy based on TD methods leads to suboptimal foraging behavior (Teichmann, Broom, & Alonso, 2014)). Perhaps representations of intake history that employ different time-scales in parallel (e.g., the ones proposed by Staddon & Higa, 1999; Shankar & Howard, 2012) could form the basis for a behavioral strategy that approximates the reward-maximizing one proposed by Charnov.

5.4 Formal representations, embodied cognition, and the dynamicist hypothesis

Language in cognitive neuroscience is dominated by the representationist framework — i.e., it is very common to describe the brain operations that generate behavior as representing events, and performing computations on these representations.

This view is challenged by the *embodied cognition* paradigm. According to Wilson and Foglia (2016), embodiment theories of cognition constitute a diverse set of theories that support what he calls the "embodiment thesis":

> **Embodiment Thesis**: Many features of cognition are embodied in that they are deeply dependent upon characteristics of the physical body of an agent, such that the agent's beyond-the-brain body plays a significant causal role, or a physically constitutive role, in that agent's cognitive processing.

This view is in sharp contrast with the traditional, brain-centric perspective of cognitive sciences. The contrast emerges from two points of conflict — the importance given to contributions of the body beyond the brain to cognitive function, and the non-representational/algorithmic nature of these contributions.

The latter point is only really one of conflict if representations are understood in the sense of symbols of a formal deductive/algorithmic system (see section 5.1.3 above). If instead representations are understood as high-dimensional vectors of a dynamical system (as proposed by the connectionists), the conflict disappears as both brain and body-beyond-brain are understood to implement cognitive function in the form of dynamical systems (Van Gelder, 1995, 1998). However, the question about the contribution of body-beyond-brain to cognitive function remains in contrast with current thinking in cognitive science.

On this regard, it is interesting to notice that the observations presented in chapter 2 of this monograph are consistent with the embodiment thesis. In that chapter, we show that ongoing behavior (i.e., dynamics of the body-beyond-brain) is predictive of perceptual judgments about the duration of a time interval, a phenomenon that would normally be regarded as having a high degree of "cognitive purity". As discussed in section 2.3, a plausible scenario is that the body inertia might contribute longer time scales to the dynamical system representing time and implementing the temporal decisions required by the task contingencies. Our results, however, are purely correlational, and by no means establish a causal relation between ongoing behavior and temporal judgments.

5.5 Concluding remarks

In the present work, we sought to characterize a representation of time contained in the dynamics of striatal populations of neurons.

The brain is a set of highly interconnected dynamical systems. As such, it is not surprising to find that its component networks display the kinds of temporal regularities required from a representation of time. However, while any such system can be legitimately said to carry information about time, our interest lies in characterizing the neural time-keeping mechanisms involved in producing interval timing behavior. For that purpose, we employed a psychophysical procedure that consisted of training our subjects to categorize time intervals as either long or short.

We were surprised to note that animals' ongoing behavior during the interval period was revealing of categorization biases. This result, presented in chapter 2, is consistent with an embodied cognition hypothesis of interval timing according to which animals would be using temporal regularities present in their body dynamics to derive duration judgments. Consistency, however, is a weak type of inductive link. Unfortunately, the purely correlational nature of our results prevent any strong conclusions on this regard.

We proceeded to observe striatal population activity during task performance (chapter 3). In line with observations found in primary sensory representations during similar tasks, we found that perceptual judgments covaried with fluctuations in the striatal representation of time. Despite the widespread belief that striatum is implicated in interval timing, this is, to the best of our knowledge, the first demonstration of the statistical dependence between fine fluctuations in the striatal representation of time and subjective time estimates. Importantly, this dependence could not be explained by ongoing behavior.

These observations, however, were performed under a single mapping between interval categories and behavioral responses. It is, therefore, impossible to know whether the observed signals constitute an abstract representation of time or if, instead, they are action-specific (which would not be surprising, given the well established role of striatum in motor behavior). In order to distinguish between these scenarios, we developed task variants in which the category-response mapping is not fixed (chapter 4), though striatal activity has not yet been recorded during performance of these task variants.

Taken together, our results strengthen the case that striatum contains representations that are important in implementing interval timing behavior. It remains to be clarified how does this representation come to be, as well as how is this information used by other brain systems in producing behavioral output. Experiments currently ongoing in our lab promise to shine light in some of these questions.

Appendix A

Articles published in peer-reviewed journals

A.1 Ongoing behavior predicts perceptual report of interval duration

Fac simile of article published as in Gouvêa et al., 2014.

frontiers in
NEUROROBOTICS

ORIGINAL RESEARCH ARTICLE
published: 11 March 2014
doi: 10.3389/fnbot.2014.00010

Ongoing behavior predicts perceptual report of interval duration

Thiago S. Gouvêa, Tiago Monteiro, Sofia Soares, Bassam V. Atallah and Joseph J. Paton*

Champalimaud Neuroscience Programme, Champalimaud Centre for the Unknown, Lisbon, Portugal

Edited by:
Sylvie Droit-Volet, Université Blaise Pascal, France

Reviewed by:
Bruce L. Brown, Queens College, USA
Jeremie Jozefowiez, Université Lille Nord de France, France

****Correspondence:***
Joseph J. Paton, Champalimaud Neuroscience Programme, Champalimaud Centre for the Unknown, Av. de Brasilia s/n, Doca de Pedrouços, 1400-038, Lisbon, Portugal
e-mail: joe.paton@neuro.fchampalimaud.org

The ability to estimate the passage of time is essential for adaptive behavior in complex environments. Yet, it is not known how the brain encodes time over the durations necessary to explain animal behavior. Under temporally structured reinforcement schedules, animals tend to develop temporally structured behavior, and interval timing has been suggested to be accomplished by learning sequences of behavioral states. If this is true, trial to trial fluctuations in behavioral sequences should be predictive of fluctuations in time estimation. We trained rodents in an duration categorization task while continuously monitoring their behavior with a high speed camera. Animals developed highly reproducible behavioral sequences during the interval being timed. Moreover, those sequences were often predictive of perceptual report from early in the trial, providing support to the idea that animals may use learned behavioral patterns to estimate the duration of time intervals. To better resolve the issue, we propose that continuous and simultaneous behavioral and neural monitoring will enable identification of neural activity related to time perception that is not explained by ongoing behavior.

Keywords: interval timing, time perception, decision making, perceptual decision making, choice probability, embodied cognition

1. INTRODUCTION

Animals live in naturally stochastic environments where apprehending environmental regularities is extremely important. In particular, being able to identify temporal regularities in the environment enables animals to predict future events such as the presence of mates, food or danger (Balsam and Gallistel, 2009), or to decide between alternative courses of action, e.g., deciding when to switch from exploiting a depleting food patch to exploring for new ones so as to optimize energy balance (Kacelnik and Brunner, 2002; Bateson, 2003). Behaviorally relevant temporal regularities in the environment are often on the scale of multiple seconds, therefore understanding how organisms handle time durations on this scale is extremely important to understand behavior itself.

Traditional sensory modalities such as vision, audition or tactile sensation are processed by known sensory organs and brain areas. Time perception, on the other hand, still lacks a clear and direct demonstration of how it would be implemented within the nervous system. In addition, whether the representation of temporal information is localized or distributed across different brain areas is still a matter of debate and ongoing research (Ivry and Spencer, 2004).

Neurally inspired models for interval timing include those that involve coincidence detection among oscillations of varying frequencies (Miall, 1989; Matell and Meck, 2000, 2004), integration of the noisy firing of neural populations (Simen et al., 2011) and variable firing dynamics within a population of neurons (Grossberg and Schmajuk, 1989; Buonomano and Merzenich, 1995; Meck et al., 2008; Shinomoto et al., 2011) as encoding schemes for time related information. Additionally, several abstract models of how animals track the passage of time have been proposed, many of which fall in one of two categories: accumulator models tell time by counting pulses emitted by a pacemaker and comparing it to a remembered value (Gibbon, 1977), while state based models represent time as a trajectory progressing through a sequence of states (Killeen and Fetterman, 1988; Machado, 1997; Ludvig et al., 2008). A subset of sequential state timing models posit that states reflect behavior (Killeen and Fetterman, 1988; Machado, 1997), stemming from the widely replicated observation that structured behavioral chains emerge under temporally structured reinforcement contingencies (e.g., Skinner, 1948; Hodos et al., 1962; Anderson and Shettleworth, 1977; Haight and Killeen, 1991; Machado and Keen, 2003; Balcı et al., 2008; Ölveczky, 2011; for reviews, see Staddon and Simmelhag, 1971; and Staddon, 1977).

Suggesting that interval timing is driven by behavioral state transitions implies a clear prediction that, to our knowledge, has not yet been tested: variation of behavioral chains should correlate with variations in time estimation. In the current work we continuously monitored the behavior of rodents as they categorized interval durations as longer or shorter than a learned standard interval. Idiosyncratic behavioral sequences displayed during the interval being timed were highly reproducible across trials and sessions. Moreover, the small variation present in the behavioral trajectories was often predictive of temporal judgments from very early in the trial, sometimes in advance of trial onset. These results revealed a correlation between learned motor behavior and perceptual report of elapsed time.

2. RESULTS

2.1. ANIMALS LEARNED TO CATEGORIZE TIME INTERVALS

We trained three rats and one mouse to categorize time intervals as either long or short by making left/right choices (**Figure 1A** and Materials and Methods). At each self initiated trial, two brief tones

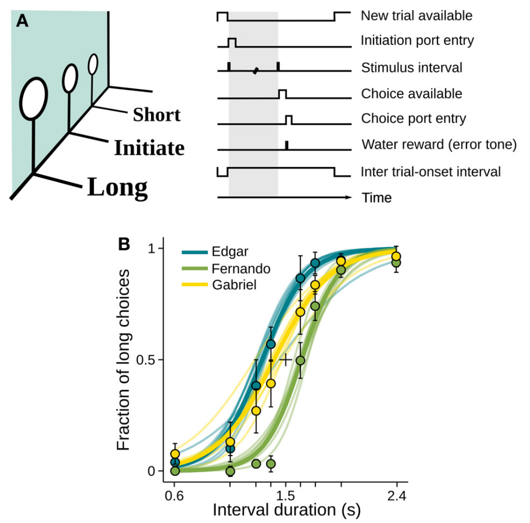

FIGURE 1 | Rats learned to categorize interval durations as short or long. (A) Animals interacted with three nose ports. Poking at the central port triggered a random stimulus interval (range: 0.6–2.4 s; top right, shaded area). A water reward became available at the left port after intervals longer than 1.5 s, and at the right port after shorter intervals. Wrong choices were punished with a brief white noise sound and a time out. **(B)** Psychometric functions show near perfect categorization of easiest stimuli (i.e., those far from the 1.5 s categorical boundary), while performance approaches chance level for intermediate durations. Circles and whiskers are mean and standard deviation across sessions. Thin lines are logistic fits to single sessions. Thick lines summarize performance of each subject. Color code identify subjects and is maintained throughout the article. $n = \{10, 8, 6\}$ sessions (rats Edgar, Fernando, and Gabriel, respectively; same for **Figures 3**, **5**, **6**); $346 \leq n \leq 558$ trials per session.

were played separated in time by an interval randomly selected from the set $I = \{0.6, 1.05, 1.26, 1.38, 1.62, 1.74, 1.95, 2.4\}$ s. Judgments about interval duration were reported at two laterally located nose ports: choosing the left side was reinforced with a drop of water after intervals longer than 1.5 s, and the right side otherwise. Incorrect choices were punished with an error tone and a time out. Animals were free to move during stimulus presentation, as long as they withheld choice until interval offset.

Each interval presentation followed by a choice constitutes one trial, and rats performed on average 456 trials per daily session (minimum = 346, standard deviation = 49.4). As revealed by their psychometric functions, animals made virtually no errors when categorizing the easiest (i.e., shortest and longest) intervals, but categorization performance declined as intervals came closer to the 1.5 s categorical boundary (**Figures 1B**, **S1A**).

2.2. ANIMALS DEVELOPED TEMPORALLY STRUCTURED BEHAVIOR

Apart from the discrete, temporally sparse behavioral measurements obtained from the nose ports, we monitored behavior continuously with a high-speed camera (rats: 120 fps; mouse: 90 fps; see Materials and Methods for details). The camera was located at the top of the behavioral box, thus only detecting motion occurring within planes parallel to the floor. Videos taken around presentation of stimulus intervals revealed highly consistent body motion patterns. To illustrate this, we overlaid video excerpts of a representative session of rat Fernando time-aligned at the onset of presentations of the longest interval. This interval duration was chosen for the analysis because it allows behavioral sequences to unfold for as long as possible before being disrupted by the interval offset tone, which acts as a go signal prompting the animal to move to a choice port. As a result of the high degree of reproducibility of the behavioral sequence, the resulting averaged images are surprisingly sharp (**Figure 2**, **Supplementary Movie 1**).

In order to quantify this effect and extend our analysis to a number of subjects and sessions, we used the aid of computer vision tracking algorithms to follow the position of each animal's head in time (see Materials and Methods for details). In agreement with the example video, head trajectories revealed body motion patterns that were very consistent across trials, as well as across sessions (**Figures 3A,B**, **S1**). Interestingly, each subject developed its own distinct trajectory. To assess the within and between subject variability in head trajectory, we computed a correlation matrix comparing all pairwise combinations of trials wherein the longest interval duration was delivered. Correlations between trajectories produced by a given subject were highly and consistently positive, whether or not they occurred in the same session. Pairs of trajectories of different subjects, on the other hand, showed near zero correlations (**Figures 3C,D**). Given the observation that trajectories are idiosyncratic and consistent from session to session, we pooled data across sessions within subjects for the following analyses.

2.3. ONGOING BEHAVIOR BEARS INFORMATION ABOUT UNFOLDING PERCEPTUAL DECISIONS

Next, we asked whether trial-to-trial variability in body trajectories carried information about the perceptual decisions being forged; were this the case, different categorizations of the same stimuli should be accompanied by distinct behavioral trajectories. For a representative session of rat Edgar, we selected all presentations of the stimulus for which choice variance (σ^2_{choice}) was highest ($I = 1.38$ s; $n = 56$ trials; $\sigma^2_{\text{choice}} = 13.3571$). The different color channels in the video were used to parse trials by choice: long choice trials were put in the red, and short choice trials in the green channel. The resulting video reveals a separation in body position from the first few hundreds of milliseconds (**Figure 4**, **Supplementary Movie 2**).

The differences in behavioral trajectories leading to different categorizations of the same stimuli imply that it should be possible to predict choice from behavioral trajectories. To quantify this effect, we employed a metric commonly used in sensory neuroscience known as *choice probability* (Britten et al., 1996; Nienborg et al., 2012). Choice probability is defined as the degree to which fluctuations of a variable during repeated presentations of a stimulus are predictive of perceptual judgments about that stimulus. This metric is commonly applied to the firing of neurons in sensory brain areas in order to estimate their involvement in the formation of percepts. We extend its use to assess whether

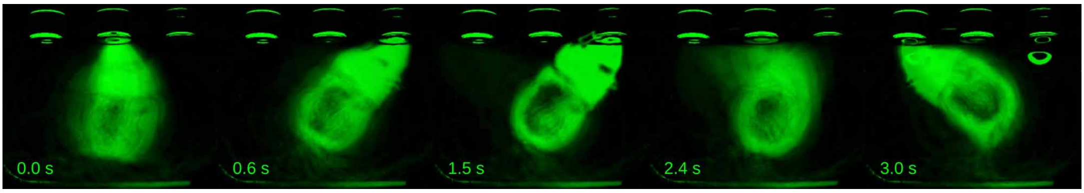

FIGURE 2 | Behavior displayed during stimulus interval is highly reproducible. A series of video frames taken from a representative session of rat Fernando at specific time points within trials were averaged across all presentations of the longest interval ($I = 2.4$ s). Times when frames were taken are indicated in seconds relative to interval onset. $n = 62$ trials.

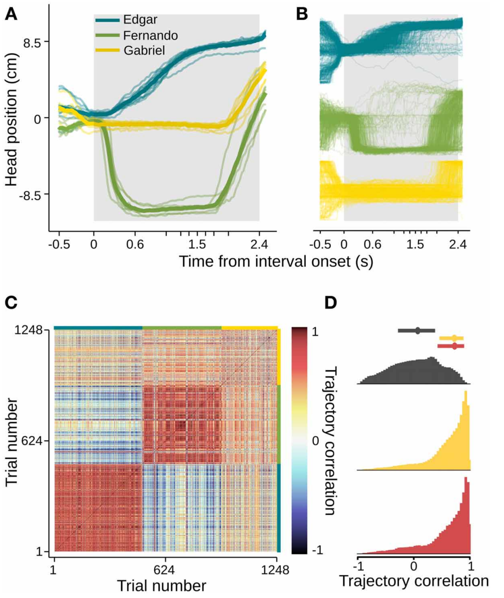

FIGURE 3 | Head trajectories are reproducible and idiosyncratic. (A) Head trajectories around presentations of longest interval. Thin lines are single session means. Thick lines are means across session means. **(B)** Trajectories at all single trials when the longest interval was presented. **(C)** Matrix of pairwise correlations between trajectories shown in **(B)**. Trials are ordered by subject (color bars framing top and right margins), then session, then trial number. **(D)** Normalized histograms of correlation coefficients. Trajectories occurring in the same session are strongly and positively correlated (bottom, red). The same is true for pairs of trajectories occurring at different sessions of a same subject (middle, yellow), and the two distributions are indistinguishable (Kolmogorov–Smirnov test, $p = 0.84$). Trajectories of different subjects, however, show near-zero correlations (top, dark gray). The latter distribution differs significantly from both within subject coefficient distributions (Kolmogorov–Smirnov tests, $p < 0.01$ in both cases). Dots and bars above histograms are medians and interquartile ranges. $n = \{10, 8, 6\}$ sessions; $n = \{491, 445, 312\}$ trials.

body trajectories carried information about unfolding perceptual judgments of time intervals.

We started by calculating choice probability from head position at individual time steps within a period extending from 0.5 s before to 2.5 s after trial initiation (**Figures 5**, **S2**, **S3**). Choice probability from head position was quantified as the area under a receiver operating characteristic (ROC) curve. This curve was calculated using distributions of head position observed in short versus long choice trials (see Materials and Methods for details). In agreement with the example video, this analysis revealed that overt behavioral sequences often allow perceptual judgment to be predicted above chance. The profile of choice probability over time differed for each individual subject (**Figure 5**). Edgar displayed a monotonically increasing profile that is significant from before stimulus onset and throughout stimulus presentation. Fernando displayed a more complex profile that was significant early in the stimulus, lost significance, and then regained small but generally significant separability from 0.7 s onward. Gabriel did not display overt head trajectories during the interval period, staying at the initiation port throughout presentation of the stimulus interval instead. However, Gabriel's choice probabilities were significant prior to trial initiation. The absence of appreciable change in Gabriel's head position during the interval period made it impossible to extract any information from this variable. However, close inspection of individual videos suggested that this rat may have produced smaller scale movements around the initiation port in the axis normal to the image plane. We were not able to quantify such movements using the current setup, and thus likely underestimated the degree to which this animal's movement may have related to choice.

Our analysis of head position represents an instantaneous behavioral analog of neuronal choice probability. While these analyses provided a measure of how well choice can be predicted from head position at each point in time, it is possible that there was information about choice contained in head position over multiple time points.

In order to evaluate the impact of head trajectories on choice, we used standard methods of clustering in high dimensional space. Briefly, trajectories were represented as sequences of head positions during a time window extending from 0.5 s before trial initiation to 0.5 s after. Since the shortest possible stimulus duration is 0.6 s, this period is common across all trials, and the animal has not yet received any information about the stimulus duration being presented. Intuitively, if behavioral trajectories are systematically related to perceptual decisions, long and short choice trials should be distributed differently in high dimensional head position space. Choice probability was estimated from single

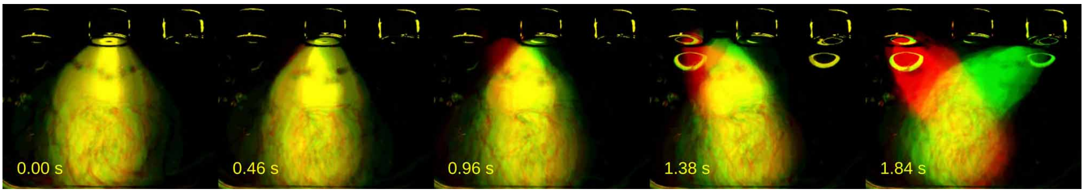

FIGURE 4 | Distinct behaviors accompany different categorizations of same stimulus. A series of video frames were taken from a representative session of rat Edgar at specific time points during presentations of a near-boundary stimulus interval ($I = 1.38$ s). Short choice trials were put on the green, and long choice trials on the red channel. Times when frames were taken are indicated in seconds relative to interval onset.

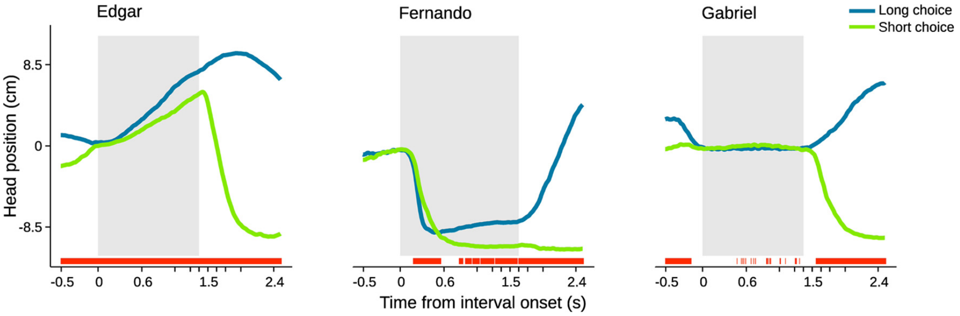

FIGURE 5 | Head trajectory is predictive of choice. Average head trajectories leading to long (blue) and short (green) categorizations of same near boundary durations. Gray shaded area indicates stimulus interval period. For each subject, the stimulus of highest choice variance across sessions was selected. Red bars indicate moments when head position is significantly predictive of choice (95% bootstrap confidence intervals on auROC curve). $n = \{678, 475, 300\}$ trials.

trial trajectories by fitting the head position sequences with multivariate Gaussian mixture models (see Materials and Methods for details). Next, we grouped trials based on choice probability in six bins with equal number of trials, and plotted the mean trajectory and psychometric curve for each bin (**Figure 6**).

All psychometric curves asymptoted at 0 and 1, but biases were strikingly different and ordered. As noted, this analysis took as input behavior occurring before stimulus identity could possibly be known by the subject. Importantly, the ordered psychometric functions do not suggest that the animal has already formed its perceptual judgment prior to stimulus presentation. Were this the case, performance should be at chance level. Rather, the animal's head trajectory exerted a bias on choice, as the difference in psychometric functions was mainly captured by the bias parameter. The same pattern would be expected if we could bin trials with respect to internal decision variables such as clock speed or decision criterion.

2.4. BEHAVIORAL TRAJECTORY IMPROVES CHOICE PREDICTION BEYOND TRIAL HISTORY

The average head trajectories during the period preceding trial initiation were strikingly related to choice probability (**Figure 6A**). This suggests that the predictive power of trajectories reflected events preceding the current trial, such as choices made and rewards received on recent trials (e.g., Sugrue et al., 2004; Lau and Glimcher, 2005). To test whether behavioral trajectories significantly improved our ability to predict choice beyond the information provided by trial history, we fit four logistic regression models to the choice data that differed in the combination of predictor variables included. We allowed different combinations of subject, current trial stimulus, recent trial history of choices, stimuli and rewards, and head trajectory to be weighed in predicting choice (**Table 1**; see Materials and Methods and **Table 2** for details).

Model no.1 captured the effect of stimuli, expected to be strong if animals learned the duration discrimination rule inherent in the task. Dummy variables standing for individual subjects were used to capture cross subject differences in psychometric functions. As expected, model no.1 predicted choice at a high success rate (**Table 1**), and was strongly significant as compared to a constant model (log likelihood ratio test, $\chi^2 = 6.01 \times 10^3$, $df = 5$, $p \ll 0.01$).

Next we assessed the contribution of behavioral trajectories during the time window used to calculate choice probability (**Figure 6**). Model no.2 maintained the predictors present in model no.1, to which it added two variables describing behavioral trajectories (i.e., projections on the first and second principal components; see Materials and Methods for details). This model showed a modest improvement in prediction success rate over model no.1. Better predictions are expected from models with

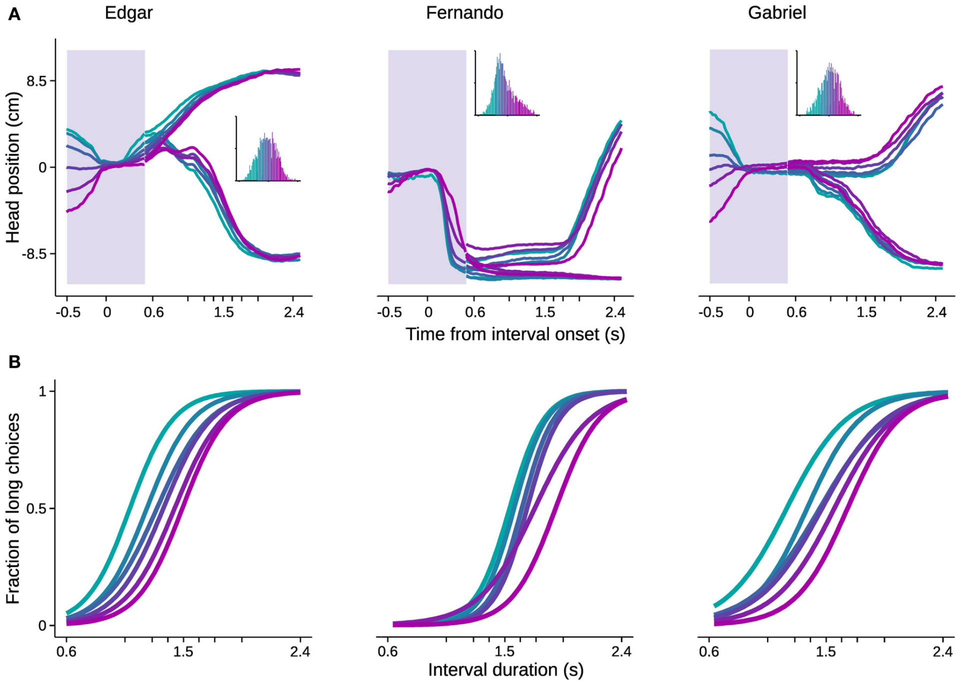

FIGURE 6 | Head trajectory reveals categorization bias. Trajectories of head position in a 1 s time window centered on interval onset were used to predict choice. The probability of trajectories conditional on choice was fitted with Gaussian mixture models. The fitted distributions were used to estimate choice probability from single trial trajectories. Trials were then binned by choice probability. **(A)** Trajectories averaged within bins are shown for the time window used for prediction (shaded area). For the remaining time, mean trajectories were further split by choice. Color code indicates bins. Inset: histogram of choice probabilities estimated from trajectories. **(B)** Psychometric curves for trials comprised in each bin. $n = \{4542, 3680, 2641\}$ trials. $n = 6$ bins.

Table 1 | Behavioral trajectories improve choice prediction beyond trial history.

Model no.	Predictor groups	Prediction success (%)	Deviance	BIC	AIC
1	Subject, interval	83.8	7103	7157.9	7115
2	Subject, interval, trajectory	84.5	6611	6711.7	6633
3	Subject, interval, trial history	84.4	6944.9	7073.1	6972.9
4	Subject, interval, trial history, trajectory	85.0	6544.5	6700.1	6578.5

Outcome of multivariate logistic regression models fit to predict choice from stimulus, trial history, and behavioral trajectory.

more free parameters. We therefore employed formal model comparison methods by calculating Bayesian and Akaike information criteria (BIC and AIC, respectively). In summary, these methods impose a cost for adding free parameters. The improvement achieved by adding extra free parameters for trajectories outweighed the costs imposed by both methods (**Table 1**). In addition, a log likelihood ratio test indicated that the improvement of model no.2 over model no.1 is highly significant ($\chi^2 = 492$, $df = 5$, $p \ll 0.01$).

In order to assess the contribution of information regarding trial history, model no.3 added to predictors in model no.1 variables describing stimulus, difficulty and reward on the preceding trial. Similar results were obtained for model no.3 as for model no.2, both compared to model no.1 (**Table 1**; log likelihood ratio test, $\chi^2 = 135$, $df = 1$, $p \ll 0.01$). A direct comparison between BIC and AIC values of models 2 and 3 indicated that trajectories were in fact slightly more informative than trial history. Furthermore, neither of the simpler models were better than model no.4, a full model incorporating variables relative to subject, stimulus interval, trajectories and trial history (**Table 1**; log likelihood ratio tests relative to model no.4; model no.2: $\chi^2 = 46$, $df = 1$, $p \ll 0.01$; model no.3: $\chi^2 = 403$, $df = 5$, $p \ll 0.01$).

3. DISCUSSION

Under environments with strong temporal regularities, animal behavior is known to become temporally structured (e.g., Skinner, 1948; Hodos et al., 1962; Anderson and Shettleworth, 1977; Haight and Killeen, 1991; Machado and Keen, 2003; Balcı

Table 2 | Specification of logistic regression models used to predict choice.

Model no.	Initial specification	Final specification
1	$S + I$	$S \times I$
2	$S + I + T_{\text{PC1}} + T_{\text{PC2}}$	$S \times I + T_{\text{PC1}} + S \times T_{\text{PC2}} + I \times T_{\text{PC2}}$
3	$S + I + R_{t-1} + d(I_{t-1}) + I_{t-1}$	$S \times I + R_{t-1}$
4	$S + I + T_{\text{PC1}} + T_{\text{PC2}} + R_{t-1} + d(I_{t-1}) + I_{t-1}$	$S \times I + T_{\text{PC1}} + S \times T_{\text{PC2}} + I \times T_{\text{PC2}} + R_{t-1}$

The four models were initially specified will all linear terms and no interactions. Models were then modified by a stepwise procedure that added interactions or removed terms so as to minimize BIC. S = subject; I = interval; T_{PC1} and T_{PC2} = trajectory of head position on a one second time window centered on interval onset, projected on the first and second principal components calculated across trials; R_{t-1} = reward on previous trial; $d(I_{t-1})$ = difficulty of the stimulus presented on previous trial, defined as the unsigned distance from the 1.5 s boundary. Models described following notation by Wilkinson and Rogers (1973).

et al., 2008; Ölveczky, 2011). Stemming from this observation, trajectories through behavioral states have been proposed to implement interval timing (Killeen and Fetterman, 1988; Machado, 1997). A prediction implied by this rationale is that trial to trial variations in the flow of behavioral sequences should correlate with trial to trial variations in temporal estimation. Consistent with this prediction, we found that behavioral trajectories differed between cases where the same intervals were categorized differently. Due to these correlated differences, choice could be significantly predicted from ongoing behavior.

While some sequential state timing models map states directly onto behaviors, many models posit a more abstract sequential state representation of time. These include neural network models that produce dynamics in response to input (Buonomano and Merzenich, 1995), successively broadening temporal basis functions for learning prediction via temporal difference or other learning rules (Grossberg and Schmajuk, 1989; Suri and Schultz, 1999; Ludvig et al., 2008), or models that time intervals through specific phase relationships amongst a diversity of oscillatory processes (Miall, 1989; Matell and Meck, 2000, 2004). Our data provides evidence in favor of state transition timing models, but it does not speak to the underlying neural mechanisms, and thus is consistent with all of the above sequential state models for timing.

Furthermore, our findings do not rule out the existence of dedicated timing mechanisms such as the pacemaker-accumulator model contained within scalar expectancy theory (Gibbon, 1977) or a sequential state timing mechanism wherein states do not map directly to behaviors. In other words the quantitative relationship between continuous behavior and subsequent choice we observed does not prove that behavior directly drives the perceptual process of time estimation. Behavior may instead simply reflect a more centrally mediated timing process. Alternatively, the rodents in our study may have indeed used behavioral sequences to estimate time in our task, with such a strategy being a useful but non-unique solution to the problem of how to estimate duration. Hence, our results alone merely suggest embodied strategies as one of a number of possible solutions to timing. Lastly, the brain is presumably responsible for the production of stereotyped motor sequences, and so in a trivial sense, the brain must be part of the system that estimates duration, even if this computation is to some degree dependent on ongoing behavior. Why might organisms include ongoing behavior in the process of timing? Limbs and muscles have mass and inertia, which may increase time constants present in movement to a degree that is harder to achieve within the nervous system alone (cf. Vogels et al., 2005). Additional experiments that manipulate the environment dynamically would help reveal whether behavior is upstream of the perceptual process of time estimation.

A number of brain structures have been implicated in interval timing. These areas include parietal cortex (Janssen and Shadlen, 2005), prefrontal cortex (Fuster, 2001; Kim et al., 2013), the basal ganglia (Maricq and Church, 1983; Matell and Meck, 2000, 2004; Fiorillo et al., 2003; Meck et al., 2008; Jin et al., 2009; Adler et al., 2012), and the cerebellum (Ivry and Keele, 1989). Interestingly from the perspective of embodied solutions to timing, these areas represent some of the most highly integrative territories in the mammalian brain, processing information concerning a broad range of sensory modalities and effector systems.

As is often the case, the full nature of interval timing likely reflects a mixture of various mechanisms, embodied and non-embodied alike (Wittmann, 2013). Animals and brains have evolved to be opportunistic, and are capable of employing a variety of strategies to solve cognitive tasks depending on the scenario they find themselves. The method of computing choice probabilities from behavior presented here represents a path forward for neuroscientists seeking to disambiguate embodied versus non-embodied components of cognitive acts such as perceptual decision making. In other decision-making contexts, it has been shown that information about an unfolding decision continually flows to the motor system, such that continuously observing behavioral output provides information about the dynamics of a decision process (Selen et al., 2012). In this way, ongoing behavior can provide a readout of the current state of decisions, which can then be compared to neural signals that are thought to be involved. We suspect that, whether studies involve human, non-human primate, rodent, or other species as experimental subjects, some neural correlates of decision variables are explained by changing motor output or sensory input resulting from ongoing behavior. We also suspect that some neural correlates of decision variables precede the emergence of decision variables in ongoing behavior, and we propose to focus on neural correlates of decision variables that meet this criterion. To do so would surely provide a better handle on the genesis of choice.

4. MATERIALS AND METHODS

All experiments were approved by the Champalimaud Foundation Bioethics Committee and the Portuguese Direção Geral Veterinária.

4.1. DURATION CATEGORIZATION TASK

Three adult male wild type Long-Evans hooded rats and one adult male PV-Cre Black 6 mouse were trained to categorize time intervals as either long or short by making left/right choices. Animals learned to trigger stimulus intervals by nosepoking at the centrally located initiation nose port when it was illuminated. Triggering a stimulus would immediately turn off the initiation port light, and cause a pair of audible tones to be played separated in time by an interval randomly selected from the set $I = \{0.6, 1.05, 1.26, 1.38, 1.62, 1.74, 1.95, 2.4\}$ s. Tones consisted of 150 ms long trains of square pulses at 7 kHz. Intervals in the set are symmetric around the 1.5 s categorical boundary, and make up four difficulty levels in geometric progression. The duration of the presented interval governed reinforcement of nosepoking at the laterally located choice ports, and responses were interpreted as the animal's perceptual judgment regarding interval duration. For intervals longer than the categorical boundary, a water reward became available for delivery upon choice of the left nose port, or at the right nose port for intervals shorter than the boundary. Incorrect responses were cued with a 150 ms long white noise sound and punished with an 11 s time out.

Animals were required to withhold poking at the choice ports during stimulus presentation, but were otherwise unrestrained. Responses occurring before interval offset were termed premature, and had the same consequences as incorrect choices (i.e., error tone and time out). Premature responses occurred in $6.7 \pm 3.3\%$ and $21.3 \pm 5.6\%$ of trials (rats and mouse, respectively; mean ± standard deviation). Stimulus intervals interrupted by premature responses were repeated in the subsequent trials, i.e., animals could not skip long intervals by making premature responses.

Nine seconds after the initiation of the previous trial, or twenty seconds following incorrect/premature responses, the initiation port would become illuminated again, indicating that a new trial could be initiated.

Sessions were selected based on categorization performance. In the selected sessions, animals correctly categorized the easiest stimuli (i.e., 0.6 and 2.4 s) at a rate of $95.8 \pm 0.03\%$, while performance reached $69.4 \pm 5.58\%$ for the hardest stimuli (i.e., 1.38 and 1.62 s), and $84.7 \pm 3.57\%$ over all eight stimuli (mean ± standard deviation). Psychometric functions are presented as logistic regressions fit to predict probability of a long choice from the duration of presented stimulus intervals. Logistic regressions were fit to single sessions. Performance of individual subjects was summarized by averaging over the parameters fit to single sessions.

4.2. BEHAVIORAL SET UP

Behavioral boxes consisted of a metal cage (mice, Island Motion, Tappan, NY, USA) or a plastic bucket (rats, IKEA, Alfragide, Portugal) containing one speaker (Cover Industrial Co., Guangdong, China) and three nose ports (Island Motion). Each nose port contained one infra-red beam/sensor pair for detecting nosepoking and one visible LED. The choice ports contained, in addition, a water tube connected to a solenoid valve for reward delivery. Valves were calibrated to deliver 25 or 5 μl of water per reward event (rats and mouse, respectively).

Except for the video camera, all sensors and effectors in the behavioral box were read and controlled by an Arduino Mega 2560 microprocessor (additional information and free software available at http://www.arduino.cc/) via a custom circuit board. The microprocessor implemented the behavioral task, and, through a serial communication port, outputted data to a desktop computer running custom software based on Python's pySerial module (freely available at http://pyserial.sourceforge.net/).

4.3. VIDEO ACQUISITION AND TRACKING

Video was acquired with a high speed camera (Flea3 FL3-U3-13S2C-CS, Point Grey Research Inc., Richmond, Canada) at 120 frames/s with a resolution of 1280 × 960 pixels in grayscale at 8 bits (rats), or 90 frames/s at 640 × 480 pixels (mouse). Video acquisition and offline tracking were performed using the in-house developed software Bonsai (freely available at http://bitbucket.org/horizongir/bonsai/downloads). To extract position of the head from the raw videos, images were background subtracted and thresholded so that the animal's body appeared as a distinct blob. For each frame, the blob's largest axis was found, and the spatial position of the axis tip corresponding to the animal's head was tracked in both x and y dimensions. All analyses were carried on position along the axis in which nose ports are lined up, while motion along the orthogonal axis was discarded.

4.4. ESTIMATING CHOICE PROBABILITY FROM ONGOING BEHAVIOR

4.4.1. Momentary head position

Choice probability given head position at a particular time, $P(C|H_t)$, was calculated by applying an ROC analysis (Green and Swets, 1966) to assess the degree of overlap between the two known distributions $P(H_t|C)$ for $C \in \{\text{long, short}\}$. The area under the ROC curve was calculated and rectified about 0.5. This number signifies the probability that an ideal observer (i.e., one with full knowledge of the distributions) would correctly categorize a new sample. Effectively, it provides an instantaneous metric of the degree to which head position is informative about unfolding perceptual decisions. A 95% confidence interval around chance level was estimated by calculating choice probability over randomly shuffled data. The procedure was repeated 100 times for each time step, and the 95th percentile was taken as the confidence interval. The analysis was applied separately for each subject and stimulus interval.

4.4.2. Behavioral trajectories

Choice probability was also estimated from behavioral trajectories. Trajectories were defined as vectors of head positions during a one second long time window centered on interval onset, and are denoted by $\mathbf{H}$. This time window was chosen because during this period subjects had no information about the stimulus being presented, thus allowing us to combine trials of different stimulus types in the same analysis. We proceeded by fitting Gaussian mixture models to the two known multivariate distributions $P(\mathbf{H}|C)$ for $C \in \{\text{long, short}\}$. On the fitting procedure, trajectories were weighted by choice variance. The rationale is that stimulus is a near-sufficient predictor of choice variability for easy stimuli, as made clear by the psychometric curves. In other words, variability emerging from other sources, such as body dynamics, would only impact decision to the extent allowed by stimulus ambiguity. Therefore, in order to capture the relevant variability, trajectories occurring in trials of a given session and stimulus type contributed to fitting with weights given by the associated binomial variance of choice. Once the distributions $P(\mathbf{H}|C)$ were fitted, and given knowledge of the marginal distributions $P(C)$ and $P(\mathbf{H})$, we could use Bayes theorem to calculate the choice probability $P(C|\mathbf{H})$ for individual trials as follows:

$$P(C|\mathbf{H}) = \frac{P(\mathbf{H}|C) \times P(C)}{P(\mathbf{H})} \qquad (1)$$

Trials were binned based on choice probability to generate the trajectories and psychometric functions presented in **Figure 6**. Similar results were observed when trajectories were fit using flat weights (data not shown). The analysis was applied separately for each subject.

4.5. GENERALIZED LINEAR MODELS

We fitted four logistic regression models to choice data. The models included different combinations of predictor variables referring to subject, stimulus, behavioral trajectory and trial history (**Table 1**). Subjects were represented by categorical variables, while stimulus duration was represented as a continuous variable. Behavioral trajectories were considered within the same time window used to calculate choice probability (i.e., a 1 s time window centered on interval onset). Given the rate at which videos were acquired, trajectory excerpts of this length are variables of 120 dimensions. To avoid adding unnecessary free parameters to the logistic models, a principal component analysis was run across the whole dataset. The first two principal components explained 89.8% of trajectory variance, and single trial projections onto these were fed to the logistic model. Trial history was represented by three variables referring to trials prior to choice: stimulus duration, stimulus difficulty (defined as the unsigned distance from the 1.5 s categorical boundary), and reward (defined as 1 for reward after long choices, 0 for no reward, and −1 for reward after short choices). Initially, models were specified with all linear terms and no interactions (**Table 2**, middle column). Models were then modified by a stepwise procedure that added interactions or removed terms so as to minimize BIC. Final model specifications are shown in the last column of **Table 2**. This analysis did not include trials in which a premature response occurred, nor their immediately succeeding trials.

ACKNOWLEDGMENTS

We thank Gonçalo Lopes for development and support of the software Bonsai used for video acquisition and tracking, Gustavo B. M. Mello for custom code used to generate supplementary movies, Matthieu Pasquet for designing the custom circuit board that interfaces computer and microcontroller, Bruno Ceña for logistic support and Asma Motiwala for comments on the manuscript.

FUNDING

Fundação para a Ciência e a Tecnologia, Champalimaud Foundation, and Bial Foundation.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: http://www.frontiersin.org/Neurorobotics/10.3389/fnbot.2014.00010/abstract

Supplementary Movie 1 | Average video of representative session by Fernando depicting trials in which stimulus interval was the longest, 2.4 s. In the depicted session, he correctly chose the long choice port for all presentations of this stimulus.

Supplementary Movie 2 | Average video of representative session by Edgar depicting trials of stimulus interval 1.38 s, the one for which choice variability was highest. Trials in which he correctly judged the interval as short are in the green color channel. Trials in which he mistakenly judged the interval as long are in the red color channel.

Figure S1 | Correlation between behavioral trajectory and temporal categorization was replicated in the mouse. (A) Psychometric function shows near perfect categorization of easiest stimuli, while performance approaches chance level for intermediate, near boundary intervals. Thin lines are logistic fits to single sessions. Thick line summarizes performance across sessions. $n = 6$ sessions. $201 \leq n \leq 521$ trials per session. **(B)** Average head trajectories around presentations of the longest interval. Thin lines are single session means. Thick line is mean of session means. $n = 109$ trials. Gray shaded area indicates stimulus interval period. **(C)** Average head trajectories leading to long (blue) and short (green) categorizations of a near boundary stimulus interval. Red bar indicates moments when head position is significantly predictive of choice (95% bootstrap confidence intervals). $n = 553$ trials. **(D)** Choice probability was estimated from head trajectory around stimulus onset time (purple shaded area), and trials were then binned by choice probability. Trajectories averaged within bins are shown for the time window used for prediction. For the remaining time, mean trajectories were further split by choice. Color code indicates bins. Inset: histogram of choice probabilities. **(E)** Psychometric curves for trials comprised in each bin. $n = 2137$ trials. $n = 6$ bins.

Figure S2 | Head trajectories leading to different categorizations of each of the eight stimulus durations. Blue and green curves depict average head trajectories leading to long and short categorizations, respectively. Stimulus period is indicated by gray shaded area. Red asterisks indicate panels presented in **Figure 5** (rats) and **Figure S1** (mouse). Mouse BII made no incorrect choices after presentations of the shortest interval.

Figure S3 | Choice probability given momentary head position. Choice probability calculated from momentary head position. Stimulus period is indicated by the gray shaded areas. Colored shaded areas indicate 95% bootstrap confidence intervals around chance. For mouse BII, analysis could not be applied to shortest interval due to absence of incorrect choices.

REFERENCES

Adler, A., Katabi, S., Finkes, I., Israel, Z., Prut, Y., and Bergman, H. (2012). Temporal convergence of dynamic cell assemblies in the striato-pallidal network. *J. Neurosci.* 32, 2473–2484. doi: 10.1523/JNEUROSCI.4830-11.2012

Anderson, M. C., and Shettleworth, S. J. (1977). Behavioral adaptation to fixed-interval and fixed-time food delivery in golden hamsters. *J. Exp. Anal. Behav.* 27, 33–49. doi: 10.1901/jeab.1977.27-33

Balcı, F., Papachristos, E., Gallistel, C., Brunner, D., Gibson, J., and Shumyatsky, G. (2008). Interval timing in genetically modified mice: a simple paradigm. *Genes Brain Behav.* 7, 373–384. doi: 10.1111/j.1601-183X.2007.00348.x

Balsam, P. D., and Gallistel, C. R. (2009). Temporal maps and informativeness in associative learning. *Trends Neurosci.* 32, 73–78. doi: 10.1016/j.tins.2008.10.004

Bateson, M. (2003). "Interval timing and optimal foraging," in *Functional and Neural Mechanisms of Interval Timing*, ed M. H. Warren (Boca Raton, FL: CRC Press), 113–141.

Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., and Movshon, J. A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis. Neurosci.* 13, 87–100. doi: 10.1017/S095252380000715X

Buonomano, D. V., and Merzenich, M. M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. *Science* 267, 1028–1028. doi: 10.1126/science.7863330

Fiorillo, C. D., Tobler, P. N., and Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299, 1898–1902. doi: 10.1126/science.1077349

Fuster, J. M. (2001). The prefrontal cortex–an update: time is of the essence. *Neuron* 30, 319–333. doi: 10.1016/S0896-6273(01)00285-9

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychol. Rev.* 84, 279–325. doi: 10.1037/0033-295X.84.3.279

Green, D. M., and Swets, J. A. (1966). *Signal Detection Theory and Psychophysics*, Vol. 1. New York, NY: Wiley.

Grossberg, S., and Schmajuk, N. A. (1989). Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Netw.* 2, 79–102. doi: 10.1016/0893-6080(89)90026-9

Haight, P. A., and Killeen, P. R. (1991). Adjunctive behavior in multiple schedules of reinforcement. *Anim. Learn. Behav.* 19, 257–263. doi: 10.3758/BF03197884

Hodos, W., Ross, G. S., and Brady, J. V. (1962). Complex response patterns during temporally spaced responding. *J. Exp. Anal. Behav.* 5, 473–479. doi: 10.1901/jeab.1962.5-473

Ivry, R. B., and Keele, S. W. (1989). Timing functions of the cerebellum. *J. Cogn. Neurosci.* 1, 136–152. doi: 10.1162/jocn.1989.1.2.136

Ivry, R. B., and Spencer, R. (2004). The neural representation of time. *Curr. Opin. Neurobiol.* 14, 225–232. doi: 10.1016/j.conb.2004.03.013

Janssen, P., and Shadlen, M. N. (2005). A representation of the hazard rate of elapsed time in macaque area lip. *Nat. Neurosci.* 8, 234–241. doi: 10.1038/nn1386

Jin, D. Z., Fujii, N., and Graybiel, A. M. (2009). Neural representation of time in cortico-basal ganglia circuits. *Proc. Natl. Acad. Sci. U.S.A.* 106, 19156–19161. doi: 10.1073/pnas.0909881106

Kacelnik, A., and Brunner, D. (2002). Timing and foraging: Gibbon's scalar expectancy theory and optimal patch exploitation. *Learn. Motiv.* 33, 177–195. doi: 10.1006/lmot.2001.1110

Killeen, P. R., and Fetterman, J. G. (1988). A behavioral theory of timing. *Psychol. Rev.* 95, 274–295. doi: 10.1037/0033-295X.95.2.274

Kim, J., Ghim, J.-W., Lee, J. H., and Jung, M. W. (2013). Neural correlates of interval timing in rodent prefrontal cortex. *J. Neurosci.* 33, 13834–13847. doi: 10.1523/JNEUROSCI.1443-13.2013

Lau, B., and Glimcher, P. W. (2005). Dynamic response-by-response models of matching behavior in rhesus monkeys. *J. Exp. Anal. Behav.* 84, 555–579. doi: 10.1901/jeab.2005.110-04

Ludvig, E. A., Sutton, R. S., and Kehoe, E. J. (2008). Stimulus representation and the timing of reward-prediction errors in models of the dopamine system. *Neural Comput.* 20, 3034–3054. doi: 10.1162/neco.2008.11-07-654

Machado, A. (1997). Learning the temporal dynamics of behavior. *Psychol. Rev.* 104, 241–265. doi: 10.1037/0033-295X.104.2.241

Machado, A., and Keen, R. (2003). Temporal discrimination in a long operant chamber. *Behav. Process.* 62, 157–182. doi: 10.1016/S0376-6357(03)00023-8

Maricq, A. V., and Church, R. M. (1983). The differential effects of haloperidol and methamphetamine on time estimation in the rat. *Psychopharmacology* 79, 10–15. doi: 10.1007/BF00433008

Matell, M. S., and Meck, W. H. (2000). Neuropsychological mechanisms of interval timing behavior. *Bioessays* 22, 94–103. doi: 10.1002/(SICI)1521-1878(200001)22:1<94::AID-BIES14>3.0.CO;2-E

Matell, M. S., and Meck, W. H. (2004). Cortico-striatal circuits and interval timing: coincidence detection of oscillatory processes. *Cogn. Brain Res.* 21, 139–170. doi: 10.1016/j.cogbrainres.2004.06.012

Meck, W. H., Penney, T. B., and Pouthas, V. (2008). Cortico-striatal representation of time in animals and humans. *Curr. Opin. Neurobiol.* 18, 145–152. doi: 10.1016/j.conb.2008.08.002

Miall, C. (1989). The storage of time intervals using oscillating neurons. *Neural Comput.* 1, 359–371. doi: 10.1162/neco.1989.1.3.359

Nienborg, H., R. Cohen, M., and Cumming, B. G. (2012). Decision-related activity in sensory neurons: correlations among neurons and with behavior. *Annu. Rev. Neurosci.* 35, 463–483. doi: 10.1146/annurev-neuro-062111-150403

Ölveczky, B. P. (2011). Motoring ahead with rodents. *Curr. Opin. Neurobiol.* 21, 571–578. doi: 10.1016/j.conb.2011.05.002

Selen, L. P., Shadlen, M. N., and Wolpert, D. M. (2012). Deliberation in the motor system: reflex gains track evolving evidence leading to a decision. *J. Neurosci.* 32, 2276–2286. doi: 10.1523/JNEUROSCI.5273-11.2012

Shinomoto, S., Omi, T., Mita, A., Mushiake, H., Shima, K., Matsuzaka, Y., et al. (2011). Deciphering elapsed time and predicting action timing from neuronal population signals. *Front. Comput. Neurosci.* 5:29. doi: 10.3389/fncom.2011.00029

Simen, P., Balcı, F., de Souza, L., Cohen, J. D., and Holmes, P. (2011). A model of interval timing by neural integration. *J. Neurosci.* 31, 9238–9253. doi: 10.1523/JNEUROSCI.3121-10.2011

Skinner, B. F. (1948). "Superstition" in the pigeon. *J. Exp. Psychol.* 38, 168–172. doi: 10.1037/h0055873

Staddon, J. E. R. (1977). "Schedule induced behavior," in *Handbook of Operant Behaviour*, eds W. K. Honig and J. E. R. Staddon (Englewood Cliffs, NJ: Prentice Hall), 125–152.

Staddon, J. E. R., and Simmelhag, V. L. (1971). The "supersitition" experiment: a reexamination of its implications for the principles of adaptive behavior. *Psychol. Rev.* 78, 3–43.

Sugrue, L. P., Corrado, G. S., and Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science* 304, 1782–1787. doi: 10.1126/science.1094765

Suri, R. E., and Schultz, W. (1999). A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. *Neuroscience* 91, 871–890. doi: 10.1016/S0306-4522(98)00697-6

Vogels, T. P., Rajan, K., and Abbott, L. (2005). Neural network dynamics. *Annu. Rev. Neurosci.* 28, 357–376. doi: 10.1146/annurev.neuro.28.061604.135637

Wilkinson, G., and Rogers, C. (1973). Symbolic description of factorial models for analysis of variance. *Appl. Stat.* 22, 392–399. doi: 10.2307/2346786

Wittmann, M. (2013). The inner sense of time: how the brain creates a representation of duration. *Nat. Rev. Neurosci.* 14, 217–223. doi: 10.1038/nrn3452

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 13 December 2013; accepted: 07 February 2014; published online: 11 March 2014.

Citation: Gouvêa TS, Monteiro T, Soares S, Atallah BV and Paton JJ (2014) Ongoing behavior predicts perceptual report of interval duration. Front. Neurorobot. ***8****:10. doi: 10.3389/fnbot.2014.00010*

This article was submitted to the journal Frontiers in Neurorobotics.

A.2 Striatal dynamics explain duration judgments

Fac simile of article published as in Gouvêa et al., 2015a.

eLIFE
elifesciences.org

RESEARCH ARTICLE

Striatal dynamics explain duration judgments

Thiago S Gouvêa†, Tiago Monteiro†, Asma Motiwala, Sofia Soares, Christian Machens, Joseph J Paton*

Champalimaud Neuroscience Programme, Champalimaud Centre for the Unknown, Lisbon, Portugal

Abstract The striatum is an input structure of the basal ganglia implicated in several time-dependent functions including reinforcement learning, decision making, and interval timing. To determine whether striatal ensembles drive subjects' judgments of duration, we manipulated and recorded from striatal neurons in rats performing a duration categorization psychophysical task. We found that the dynamics of striatal neurons predicted duration judgments, and that simultaneously recorded ensembles could judge duration as well as the animal. Furthermore, striatal neurons were necessary for duration judgments, as muscimol infusions produced a specific impairment in animals' duration sensitivity. Lastly, we show that time as encoded by striatal populations ran faster or slower when rats judged a duration as longer or shorter, respectively. These results demonstrate that the speed with which striatal population state changes supports the fundamental ability of animals to judge the passage of time.

DOI: 10.7554/eLife.11386.001

Introduction

Time, like space, is a fundamental dimension of the environment, yet how it is processed in the brain is poorly understood. A number of recent studies have identified dynamics that allow for robust representation of time by populations of neurons in multiple areas including the hippocampus (***Pastalkova et al., 2008***; ***MacDonald et al., 2011***), prefrontal (***Kim et al., 2013***; ***Xu et al., 2014***), parietal (***Leon and Shadlen, 2003***; ***Janssen and Shadlen, 2005***) and motor (***Lebedev et al., 2008***) cortices, cerebellum (***Mauk and Buonomano, 2004***), and the striatum (***Matell et al., 2003***; ***Jin et al., 2009***; ***Adler, 2012***; ***Mello et al., 2015***). However, any dynamics that result in a continuously-evolving and non-repeating population state can be used to encode time (***Buonomano, 2014***), and it is not known whether such temporal representations would inform subjects' judgments of duration or merely covary with elapsing time. The striatum, a brain structure known to be involved in reinforcement learning and decision making (***Lau and Glimcher, 2008***; ***Samejima et al., 2005***; ***Lee et al., 2015***), has been implicated in interval timing by several lines of evidence (***Hinton and Meck, 2004***; ***Harrington et al., 2009***; ***Wencil et al., 2010***; ***Malapani, 1998***; ***Meck, 2006***). However, whether dynamics in striatal activity can explain the perceptual performance of behaving subjects is unknown. To determine whether striatal ensembles drive subjects' judgments of duration, we manipulated and recorded from striatal neurons in rats performing a duration categorization psychophysical task.

Results

To measure the duration sensitivity of subjects' timing judgments, we trained rats to judge whether time intervals belonged to a long or short category (***Gouvêa et al., 2014***) (see Materials and methods; ***Figure 1a***). At each self-initiated trial, two brief tones (interval onset,

***For correspondence:** joe.paton@neuro.fchampalimaud.org

†These authors contributed equally to this work

Competing interests: The authors declare that no competing interests exist.

Funding: See page 12

Received: 03 September 2015
Accepted: 07 December 2015
Published: 07 December 2015

Reviewing editor: Timothy Behrens, Oxford University, United Kingdom

eLife digest You know someone is a good cook from their rice - grains must be well cooked, but not to the point of being mushy. Despite consistently using the same pot and stove, we, however, will sometimes overcook it. It is as if our inner sense of time itself is variable. What is it about the brain that explains this variability in time estimation and indeed our ability to estimate time in the first place?

One issue the brain must confront in order to estimate time is that individual brain cells typically fire in bursts that last for tens of milliseconds. So how does the brain use this short-lived activity to track minutes and hours? One possibility is that individual neurons in a given brain region are programmed to fire at different points in time. The overall firing pattern of a group of neurons will therefore change in a predictable way as time passes.

Gouvêa, Monteiro et al. found such predictably changing patterns of activity in the striatum of rats trained to estimate and categorize the duration of time intervals as longer or shorter than 1.5 seconds. Interestingly, when rats mistakenly categorized a short interval as a long one, population activity had travelled farther down its path than it would normally (and vice-versa for long intervals incorrectly categorized as short), suggesting that variability in subjective estimates of the passage of time might arise from variability in the speed of a changing pattern of activity across groups of neurons.

As further evidence for the involvement of the striatum, inactivating the structure impaired the rats' ability to correctly classify even the longest and shortest interval durations.

The next challenge is to determine exactly how the striatum generates these time-keeping signals, at which stage variability originates, and how the brain regions that the striatum signals to use them to control an animal's behavior.

DOI: 10.7554/eLife.11386.002

offset) were presented separated in time by an interval randomly selected from the set I = {0.6, 1.05, 1.26, 1.38, 1.62, 1.74, 1.95, 2.4} seconds. Judgments about interval duration were reported at two laterally located nose ports: choosing the left side was reinforced with water after intervals longer than 1.5 seconds (long stimuli), and the right side otherwise (short stimuli, ***Figure 1b***). Animals were required to withhold choice until interval offset. Animals made virtually no errors when categorizing the easiest (i.e. shortest and longest) intervals, but categorization performance declined as intervals approached the 1.5 second categorical boundary (***Figure 1c***).

We recorded action potentials during task performance (see Materials and methods), from populations of single striatal neurons targeting dorsal-central striatum, an area where manipulations produced timing deficits (***Meck, 2006***) (***Figure 2a***. For a reconstruction of striatal recording sites see ***Figure 2—figure supplement 1.***). We observed that striatal neurons displayed diverse firing patterns, with different units firing at different times within the interval period (***Figure 2b–d***). Can such firing patterns support duration judgments? To determine whether and the degree to which individual neurons could contribute to duration judgments, for each trial, we counted spikes in the last 500 ms of the interval period and compared spike count distributions of short vs long stimulus trials using a receiver operating characteristic (ROC) analysis (see Materials and methods). We found that the majority of neurons (~57%) preferred either short or long stimuli (***Figure 2e***; short-preferring: n = 159/433, 36.7%; long-preferring: n = 87/433, 20.1%; permutation test, $p<0.05$). As expected, short-preferring neurons displayed higher firing on average prior to the 1.5 s category boundary, after which long-preferring neurons displayed higher firing (***Figure 2f***). These averaged activity patterns resemble the likelihood of receiving reward on a moment-by-moment by basis should the animal choose short or long (compare with reward contingency in ***Figure 1b***). Such signals, previously observed in the parietal cortex of monkeys performing a similar timing task (***Leon and Shadlen, 2003***) and in the striatum in a value based decision task (***Lau and Glimcher, 2008***), are potentially useful for guiding choice. However, were animals' judgments indeed guided by such signals, it should be possible to predict choices reported later in the trial using neural activity collected during interval stimuli. Indeed, in trials wherein a near boundary interval was judged as long, firing of the short (long) preferring subpopulation dropped (rose) faster, so that the two curves crossed before the 1.5 s boundary (***Figure 2g***). Conversely, in trials wherein the same interval was judged as short,

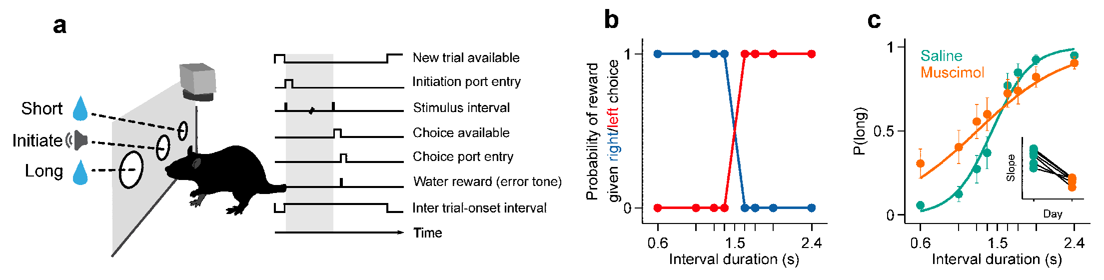

Figure 1. Rats judged interval durations as either long or short. (**a**) Rats triggered interval stimuli (i.e. two brief auditory tones separated by a silent interval of random duration) by inserting their snout into a central port. Following interval offset, animals reported their long vs short judgment at two lateral choice ports. Correct trials yielded a water reward, while incorrect or premature responses produced a white noise sound and a time out. Top view, high-speed video was acquired throughout task performance. (**b**) Reward contingency. (**c**) Averaged psychometric curves following bilateral muscimol or saline injections in dorsal striatum (mean ± standard deviation across session means, and logistic fit; n = 3 rats, 4 sessions each). Inset: slope of psychometric curves on consecutive saline and muscimol sessions. All raw data for ***Figure 1*** can be found in ***Figure 1—source data 1***.
DOI: 10.7554/eLife.11386.003

The following source data and figure supplement are available for figure 1:

Source data 1. The. txt file contains trial by trials stimulus (*Interval*), choice (*choiceLong*), animal (*Name*), treatment (*MuscimolDose*) and session (*Date*).
DOI: 10.7554/eLife.11386.004

Figure supplement 1. Histological confirmation of cannula placements.
DOI: 10.7554/eLife.11386.005

the two curves evolved more slowly so that at the time of interval offset the short preferring subpopulation was still firing at a higher level and a crossing point had not yet been reached (***Figure 2h***).

The observation of large proportions of short- and long-preferring neurons whose dynamics predicted choice is evidence that duration judgments are guided by the state of striatal populations. Might the information afforded by ensembles of striatal neurons account for the pattern of subjects' judgments across all stimuli? To test this hypothesis, we compared session to session fluctuations in behavioral performance with the separability of activity states of simultaneously recorded ensembles at the offset of short as compared to long intervals. Briefly, for each trial in a session we characterized neural population state as a vector $\mathbf{r} = (r_1, r_2,\dots, r_N)$, where r_n is the number of spikes fired by neuron $n \in [1,N]$ within the last 500 ms of the interval period. Next, for each session we found the linear discriminant that best separated population state vectors according to whether they came from a long or a short interval trial (***Figure 3a***; see Materials and methods). A threshold placed along the linear discriminant was then used as a decision rule (black line in ***Figure 3a***) to generate a 'neural duration judgment' for each trial. This procedure allowed us to obtain, for each session, a quantitative description of how well simultaneously recorded neurons could categorize stimuli, i.e., a neurometric function comparable to the behavioral psychometric function (***Figure 3b***). Consistent with duration information being encoded at the population level, we found that for sessions in which greater numbers of neurons were recorded simultaneously (i.e. upper tercile of sessions with regard to population size) psychometric and neurometric performances were similar and strongly correlated ($r^2 = 0.76$, $p<0.001$; ***Figure 3c***). These results demonstrate that a read out of stimulus category from even modestly-sized ensembles of striatal neurons was in many cases sufficient to explain the pattern of duration judgments produced by behaving subjects.

It has been previously reported that duration judgments could be predicted by animals' ongoing behavior during the interval period (***Gouvêa et al., 2014***; ***Matthews and Lerer, 1987***; ***Killeen and Fetterman, 1988***; ***Fetterman et al., 1998***; ***Machado, 1997***; ***Machado and Keen, 2003***). In addition, it is well known that striatal neurons can fire around movements (***Alexander and Crutcher, 1990***; ***Jin and Costa, 2010***). Could the categorization performance of striatal ensembles reflect activity related to movements the animal might be making during the task? To test to what degree ongoing behavior could explain the categorization performance of striatal neural activity, we applied an analogous classification analysis to video images taken of the animal just before interval offset (see Materials and methods). We found that our ability to categorize intervals using video frames was consistently poorer as compared to neural data collected at the analogous time periods during the task (***Figure 3b***, inset in ***Figure 3c***, ***Video 1***, ***Figure 3—figure supplement 1***, ***Figure 3—figure supplement 2a***). In contrast, we were able to categorize stimuli as well as the animal using video frames taken at the point when animals expressed their choice at one of the reward ports (***Figure 3—figure supplement 2b***). Furthermore, movement related responses in the striatum are known to occur both pre- and post-movement onset, much later than in other motor areas such as pre-motor and motor cortex (***Alexander and Crutcher, 1990***). Thus, if purely movement-related activity were responsible for the categorization performance of striatal ensembles, we would expect ensemble performance to display a similar time course to that of video frames. Applying the same analyses at multiple points in time ranging from 500 ms preceding to 500 ms following stimulus offset revealed a strikingly different profile of categorization performance for video frames as compared to neural ensembles (***Figure 3d–e***). Specifically, the time course of duration categorization by neural ensembles was best correlated with the duration categorization by video frames when using spikes collected between 400 ms and 200 ms preceding a reference video frame. These indicate that the categorization performance of striatal neurons was not simply related to the immediate sensorimotor state of the animal, and instead likely reflects that striatal neurons encode an internal neural representation of the state of animals' categorical decisions.

We have shown thus far that categorical information about interval duration contained in the firing of striatal populations at the time of stimulus offset can explain the precision of animals' judgments about duration. However, in the task employed here, categorical judgments must be derived from a continuously evolving decision variable that represents how much time has elapsed since the onset of the stimulus. As indicated by the diversity of firing patterns (***Figure 2d***), the state of population activity evolved continuously during interval stimuli (***Figure 3g***, ***Figure 4a***, ***Figure 3—figure supplement 3b***), a feature not captured by binary classification. Might trial to trial variations in population state predict the apparent speed of animals' internal representation of elapsed time? To test this possibility, we performed two additional analyses.

First, we projected the state of simultaneously recorded neuronal populations at stimulus offset in individual trials onto the mean trajectory within each session, noted the fraction of the mean trajectory traversed for each trial, and pooled the data for each stimulus over all sessions within a given subject. Indeed, when population state at stimulus offset had advanced relatively more or less along the mean trajectory, animals were more likely to judge intervals as long or short respectively (***Figure 3f–g***, ***Figure 3—figure supplement 3***). This effect was observed most consistently for interval stimuli that were closer to the category boundary, and thus variations in projected population state led to horizontal shifts in the psychometric curves (see Materials and methods). These data are consistent with striatal population state encoding a representation of elapsed time that is used by animals to determine categorical judgments. Indeed, such a pattern of population activity has been proposed as a suitable neural code for elapsing time (***Buonomano, 2014***; ***Buonomano and Merzenich, 1995***).

However, if such a representation encodes elapsed time, and not only subjects' judgments in this task, neural activity should continuously traverse a non-repeating trajectory in state space in a manner that predicts duration judgments during presentation of particular stimuli. Indeed, even in a low dimensional projection of population activity, we found that network state ran ahead or behind depending on whether the animal judged a near boundary stimulus as long or short (***Figure 4b–c***, ***Figure 4—figure supplement 1b,c***, ***Figure 4—figure supplement 2b,f,j***). The correspondence between population trajectory and duration judgments further suggests that striatal dynamics may form an internal representation of elapsed time that informed categorical decisions about duration. To directly test this hypothesis, we focused on stimuli near the category boundary and decoded time from the population using a naive Bayes decoder and asked whether such a representation correlated with animals' judgments, exhibiting choice probability (***Britten et al., 1996***). We found that decoded estimates of time ran faster or slower when animals judged a given stimulus as long or short, respectively (***Figure 4d–g***, ***Figure 4—figure supplement 1d-g***, cross validated naive Bayes decoder; see Materials and methods). This indicates that striatal activity provides information about elapsing time, the continuously varying decision variable necessary to inform judgments in the task. Furthermore, if this information were read out and used to guide judgments, those judgments would match those of the rats.

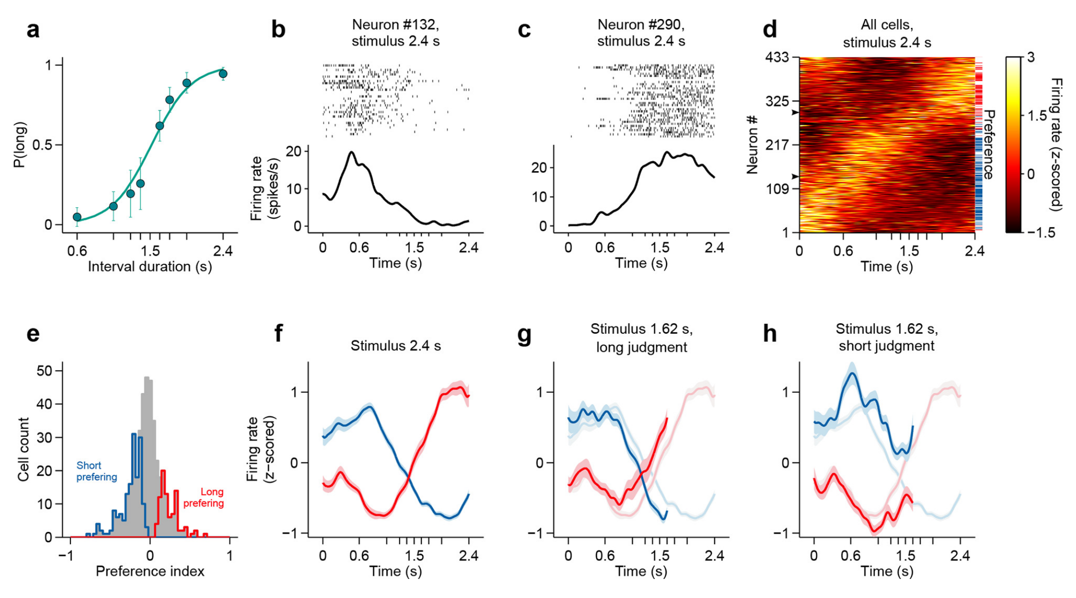

Figure 2. Dynamics of striatal subpopulations predict duration judgments. (**a**) Psychometric function for neural recording sessions (mean ± standard deviation across sessions and logistic fit, n = 37 sessions from 3 rats). (**b,c**) Raster plot and peri-stimulus time histogram (PSTH) of two example cells for trials in which the longest stimulus interval (2.4 s) was presented. Time = 0 corresponds to stimulus onset. (**d**) Normalized PSTHs of all neurons for trials in which the longest stimulus interval was presented. Arrowheads indicate cells shown in (**b,c**). Blue and red ticks indicate cells with significant short and long preferences, respectively. (**e**) Histogram of preference indices. Blue and red outlines indicate subpopulations with significant short and long preferences, respectively. (**f**) Averaged, normalized PSTH of the two subpopulations outlined in (**e**) for trials in which the longest stimulus interval was presented (mean ± SEM). (**g**) Same as in (**f**), for trials in which a near-boundary stimulus interval (1.62 s) was judged as long. For comparison, curves shown in (**f**) are reproduced as a watermark. (**h**) same as (**g**) for trials in which the stimulus was judged as short. For single subjects, see ***Figure 2—figure supplement 2*** . Behavior and neural spike count data for ***Figure 2*** and ***Figure 2—figure supplements 1 and 2*** can be found in ***Figure 2—source data 1***.

DOI: 10.7554/eLife.11386.006

The following source data and figure supplement are available for figure 2:

Source data 1. Folder with raw data for ***Figures 2–4***.

DOI: 10.7554/eLife.11386.007

Electrophysiological recordings in dorsal striatum. DOI: 10.7554/eLife.11386.008

Figure supplement 2. Dynamics of striatal subpopulations predict duration judgments.

DOI: 10.7554/eLife.11386.009

If the striatal activity we describe above directly supported task performance, manipulating the striatum should modify duration judgments. To test whether this was the case, we bilaterally injected the GABAa receptor agonist muscimol (see Materials and methods). As a result, the duration sensitivity of animals' judgments dropped significantly as compared to interleaved saline control sessions (***Figure 1c***; psychometric slope on saline sessions = [1.03 1.20] vs on muscimol sessions = [0.28 0.67]; 95% confidence intervals), yet animals otherwise performed normally. These results, by demonstrating that duration categorization in this task was dependent on a normally functioning striatum, suggest that the neural signals we observed directly supported duration judgments. However, the possibility that muscimol infusions changed other functions important for task performance such as reward processing or memory for the mapping between time and choice can not be ruled out.

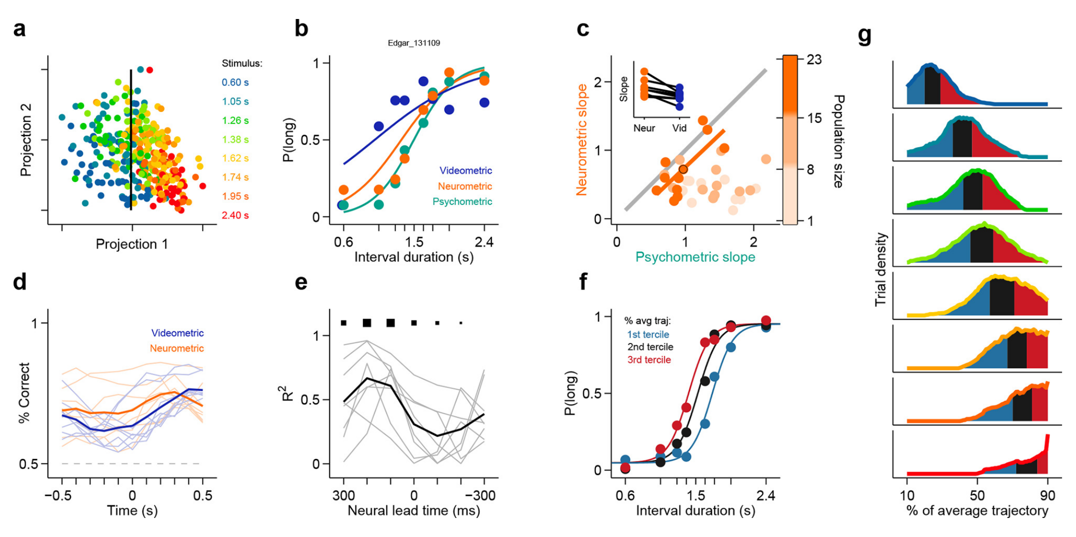

Figure 3. Simultaneously recorded population state at interval offset can explain behavioral performance. (**a**) Low dimensional representation of population state at interval offset for one example session. Black line is the decision rule (see text). (**b**) Example psychometric, neurometric and videometric curves for the same session as in (**a**). (**c**) Slopes of psychometric and neurometric curves for all sessions. Color indicates terciles of population size. Highlighted data point corresponds to the session in (a-b). Inset: regression slope of neurometric and videometric curves for sessions in the upper tercile. See ***Figure 3—figure supplement 2*** for psychometric-videometric comparison at interval offset and choice. (**d**) Performance of an ideal observer analysis in predicting stimulus category, applied to neural (orange) and video (blue) data obtained at different times relative to interval offset. Thin lines corresponds to individual sessions. Thick lines are averages. (**e**) The orange and blue curves (thin lines in panel (**d**)) for corresponding sessions were regressed against each other at different time shifts. The regression R^2 values for each session are shown in thin grey lines. The average over all sessions is shown in black. Sizes of black squares indicate the number of sessions with significant positive correlations (largest squares at 200 and 100 ms correspond to 5 sessions and smallest one at -200 ms to 1, out of a total of 8 sessions). (**f**) Psychometric curves constructed from trials separated according to whether the population state at stimulus offset had advanced more or less along the mean trajectory. Color indicates terciles shown in (**g**). (**g**) Distributions of projection on normalized mean trajectory for all trials for each stimulus are shown (stimuli color coded as in [**a**]). The equal area bins shown correspond to the groups of trials used for constructing the three psychometric curves shown in panel (**f**). Data in f-g are from rat Bertrand. See ***Figure 3—figure supplement 3*** for the remaining two subjects.

DOI: 10.7554/eLife.11386.010

The following figure supplements are available for figure 3:

Figure supplement 1. Image frames at the end of the neural analysis window do not show a consistent separation between short and long stimulus trials.

DOI: 10.7554/eLife.11386.011

Figure supplement 2. Behavior at the end of the neural analysis window did not explain the categorization performance of neural populations.

DOI: 10.7554/eLife.11386.012

Figure supplement 3. Population state at interval offset can explain behavioral performance.

DOI: 10.7554/eLife.11386.013

Discussion

Attempts to understand the neural mechanisms of time estimation have begun to focus on continuously evolving population dynamics as a general mechanism for time encoding across the brain (***MacDonald et al., 2011***; ***Mauk and Buonomano, 2004***; ***Buonomano, 2014***; ***Buonomano and Merzenich, 1995***; ***Gershman et al., 2013***). According to this view, time may be encoded by any reproducible pattern of activity across a population of neurons for as long as the pattern is continuously changing and non-repeating. However, no study to date has directly compared the speed of such "population clocks" with the duration judgments of the behaving subjects in which they are found. We show that as rats judged the duration of interval stimuli, striatal neurons displayed dynamics in firing rate that contained information about elapsed time. Furthermore, this information

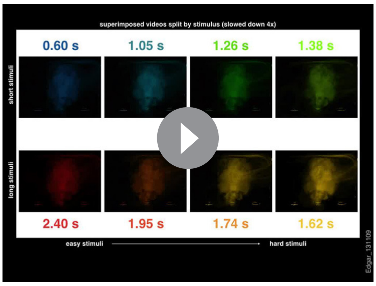

Video 1. Video clips from the entire stimulus period do not show a consistent separation between short and long stimulus trials. Superimposed video clips (thresholded and background subtracted) for each stimulus type and for the corresponding stimulus duration (same color conventions as in ***Figure 3a*** and ***Figure 3—figure supplement 1***). Red stop marks signal the end of each video and the corresponding stimulus duration.

DOI: 10.7554/eLife.11386.014

was sufficient to account for the animals' perceptual decisions, and was not accompanied by systematic differences in outwardly expressed behavior over time. Combined with the observation that striatal inactivation caused a specific decrement in timing performance, these data suggest that striatal dynamics form a central neural representation of time that guides animals' decisions about duration. Such a coding mechanism in the striatum is well situated to inform the appropriate selection of actions through downstream circuitry involving the globus pallidus, substantia nigra, and various extrinsic connections between the basal ganglia and brainstem, thalamic, and cortical motor areas (***Steiner and Tseng, 2010***). An intriguing question for future studies is how the striatal dynamics we observed during the interval discrimination task are generated. Neurons in multiple cortical layers spread across the entire cortical mantle, as well as thalamic, pallidal, and neuromodulatory populations provide input to the striatum. While time coding has been assessed in some of these populations, a careful analysis of simultaneously recorded populations that might uncover causal relationships between signals in multiple brain areas has not been carried out. Furthermore, local striatal circuitry may also play a role in shaping dynamics. However, the coding properties tested here could be tested in other brain areas where timing signals have been identified such as the hippocampus (***Pastalkova et al., 2008***; ***MacDonald et al., 2011***), medial prefrontal (***Kim et al., 2013***; ***Xu et al., 2014***), parietal (***Leon and Shadlen, 2003***; ***Janssen and Shadlen, 2005***) and motor (***Lebedev et al., 2008***) cortices, and the cerebellum (***Mauk and Buonomano, 2004***), among others. By comparing the signals recorded simultaneously in multiple brain areas during time estimation tasks, it should be possible to identify signatures of functional interaction between brain areas where they exist. Such an approach promises to elucidate where and how time information encoded at the population level is used by the brain to support the myriad time-dependent functions we and other organisms rely on for survival.

Materials and methods

Subject

Six male Long-Evans hooded rats (*Rattus norvegicus*) between the ages of 6 and 24 months were used for this study. Three rats were used for neural recordings and three rats for pharmacological manipulations. All experiments were in accordance with the European Union Directive 86/609/EEC and approved by the Portuguese Veterinary General Board (Direcção-Geral de Veterinária, project approval 014303 - 0420/000/000/2011).

Behavior

Rats were trained to perform a previously described two-alternative forced-choice timing task (***Gouvêa et al., 2014***). Briefly, animals had to categorize time intervals as either long or short by making left/right choices. For each session the animals were placed in a custom made behavioral box containing 3 nose ports and a speaker. Each trial was self-initiated by entry into the central nose port and was followed by a pair of brief auditory tones (square pulses at 7,500 Hz, 150 ms) separated by an interval selected randomly out of 8 possible durations (0.6, 1.05, 1.26, 1.38, 1.62, 1.74, 1.95 and 2.4 s). Judgments were reported at two laterally located nose ports. Left responses were reinforced with a drop of water (solenoid valves, Lee Company) after intervals longer than 1.5 seconds, and right responses otherwise. Incorrect responses were punished with a brief white noise

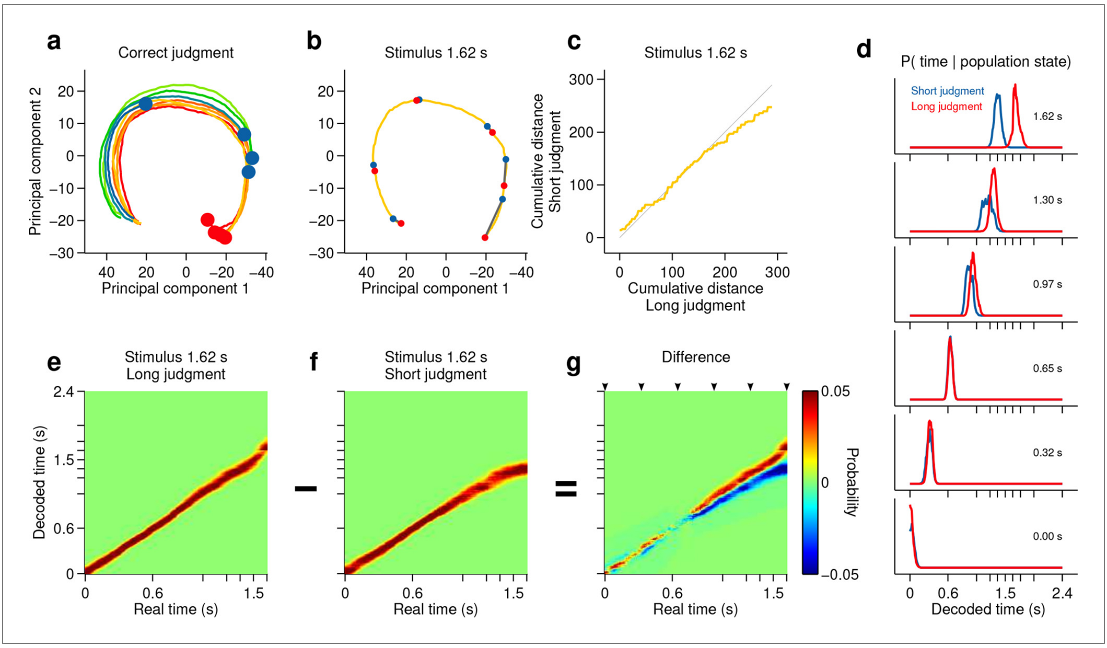

Figure 4. Smoothly changing population state encodes elapsing time in accordance with perceptual report for a long stimulus. (**a**) Low dimensional representation of population state during entire interval period of correct trials. Line colors indicate interval duration (warmer colors are longer intervals, as in ***Figure 3a***). Dots are placed at the interval offset end, and their color indicates judgment (blue: short; red: long). (**b–g**) Population state and decoded time for a single long, near boundary stimulus interval (1.62 s). (**b**) Yellow curve is same as in (**a**). Red dots are 6 time points evenly spaced between interval onset and offset. Blue dots are projections of population state during short judgment trials. Grey lines link population states at equivalent time points. (**c**) Average cumulative distance travelled in full neural space along trajectory represented in (**b**) on long versus short judgment trials. (**d**) Posterior probability of time given population state at the time points indicated in (**b**), averaged within trials of each judgment type. (**e**,**f**) Same as (**d**) for the entire interval period. (**g**) Difference between posteriors for long and short judgment trials. Arrowheads indicate same time points used in (**b**,**d**). n = 433 neurons from 3 rats. See ***Figure 4—figure supplement 1*** for a different near boundary stimulus, and ***Figure 4—figure supplement 2*** for data from individual subjects.

DOI: 10.7554/eLife.11386.015

The following figure supplements are available for figure 4:

Figure supplement 1. Smoothly changing population state encodes elapsing time in accordance with perceptual report for a short stimulus.

DOI: 10.7554/eLife.11386.016

Figure supplement 2. Single subjects show smoothly changing population states that encode elapsing time in accordance with perceptual report.

DOI: 10.7554/eLife.11386.017

burst (150 ms) and a time out (10 s). High speed video (120 fps) was collected from above during task performance. Psychometric functions were fitted using the following two-parameter logistic function

$$f(x) = \frac{1}{1 + e^{-b(x-c)}}$$

where b controls the slope and c is the inflection point of the curve.

Electrophysiology

Rats were implanted with 32-channel tungsten microwire moveable array bundles (***Figure 2—figure supplement 1a***, Innovative Neurophysiology) under isoflurane anaesthesia. All recordings (***Figure 2—figure supplement 1***) targeted dorsal striatum with coordinates centred at +0.2 mm AP and ± 3 mm ML (rat Bertrand), and +0.84 mm AP and ± 2.5 mm ML (rats Edgar and Fernando), from Bregma. Rats were given a week of post-surgical recovery and array placements were confirmed with histology (**Figure 2—figure supplement 1c**). Neural signals were recorded at 30 kHz during behavior, amplified and band-pass filtered at 250–750 Hz (Cerebus - Blackrock Microsystems). Each independent bundle was moved 50-100 μm after every recording session to ensure that independent neural populations were sampled across recording sessions. Waveforms corresponding to action potentials from single neurons were sorted offline using principal component analysis (PCA) (offline sorter, Plexon). All remaining analysis were run in Matlab (Mathworks) software. We selected all isolated units with a mean session firing rate >0.5 Hz and from sessions with >70% correct performance (averaged across all stimuli) and a minimum of 250 trials (n=433 cells, 37 recording sessions, 3 animals; rat Bertrand: 136 units, 10 sessions; rat Edgar: 163 units, 9 sessions; rat Fernando: 134 units, 18 sessions). The general result was found in all subjects. Sample size was not computed during study design. To build PSTHs, spikes were counted in 2-ms bins and convolved with a Gaussian kernel with 25-bin standard deviation. PSTHs in **Figure 2d** were ordered by angular position in the space formed by the first 2 principal components describing firing dynamics (i.e., dimensions are all time bins within interval period, samples are each neuron's mean PSTH). This method (**Geffen et al., 2009**) orders cells with respect to their dynamics while taking into consideration the full response profile over the relevant temporal window, and not just a single response feature such as peak response time.

Pharmacology

We implanted 3-mm 20-gauge stainless steel guide cannulae (Bilaney) bilaterally into the striatum of 3 rats [+0.84 mm anterior-posterior (AP), ± 2.5 mm medial-lateral (ML), from Bregma, and -3mm dorsal-ventral (DV, from cortex surface) under isoflurane anesthesia. After one week of post-surgical recovery and 4 days of training, rats were injected with either vehicle (saline, PBS 1x) or muscimol (GABA-A agonist, 100 mg/L (rats Albert and Yuri) and 300 mg/L (rat Zack), Sigma^{TM}) solutions in four alternate days. Two 1-μL syringes (Hamilton), attached to an injection pump (Harvard Apparatus) through 20-gauge internal cannulae that extended 1.5 mm bellow the guide cannulae, injected 0.6 μL of solution per site during 2.5 min. The internal cannulae were left in place for an additional 1.5 min and the rats were given a 45-min recovery period in their home-cage before starting the task. Cannula placements were confirmed by histology (**Figure 1—figure supplement 1**). The general result was found in all sessions of all subjects. Sample size was not computed during study design.

Preference index

We counted spikes during the last 500 ms of the stimulus period, and built two separate spike count distributions for short and long judgment trials. Next, we used a ROC analysis to measure the separation between distributions (95% bootstrap confidence interval, 1000 iterations). We then transformed the area under the ROC curve ($auROC \in [0,1]$) into a preference index ($k = 2*auROC - 1$; $k \in [-1,1]$). We adopted the convention that neurons with positive preference indices fired preferentially for long stimuli (**Figure 2e**).

Low dimensional representations of population state

We refer to the vector describing instantaneous firing rates (measured within 500-ms wide, 10-ms apart, overlapping time bins) across a population of neurons as the population state. The population state vector is a high dimensional variable (i.e., it has as many dimensions as neurons). With the purpose of visualizing population state in 2d plots, we employed standard dimensionality reduction techniques. In **Figure 3a**, we chose to represent in the abscissa a direction that emphasizes the separability between short and long stimulus trials (i.e., the direction that maximizes variance between groups while minimizing variance within groups; Fisher's linear discriminant; see below), and in the ordinate the axis of maximal variance that is also orthogonal to the abscissa (i.e. first principal component calculated in the null space of the linear discriminant). In **Figure 4a–b**, population state is represented in the space formed by the first 2 principal components describing population state, calculated during presentation of the interval for which choice variance is maximal (i.e. dimensions are neurons, samples are averaged spike counts for the time bins within that interval).

Neurometric curves

For each trial in a session we characterized neural population state as a vector **r** = (r_1, r_2,... , r_N), where r_n is the number of spikes fired by neuron $n \in [1,N]$ within the last 500 ms of the interval period in that trial. Next, for all trials but one from each session (*training set*; leave-one-out cross-validation procedure), we found the linear discriminant that best separated population state vectors according to whether they came from long or short interval trials (Fisher's linear discriminant analysis, LDA). The linear discriminant is given by

$$\mathbf{w} = argmax\frac{\mathbf{w}^T \mathrm{S}_B \mathbf{w}}{\mathbf{w}^T \mathrm{S}_W \mathbf{w}} = S_W^{-1}(\mu_1 - \mu_2)$$

where **w** is the vector of coefficients for the linear discriminant, S_B is the between class covariance, S_W the within class covariance and μ_1 and μ_2 are the means of all points in class 1 and class 2 respectively. A threshold placed along the linear discriminant was then used as a decision rule applied to neural data from the remaining trial (*test set*). ***Figure 3a*** depicts population vectors from an example session (projection 1: linear discriminant, no cross-validation; projection 2: first principal component of the orthogonal subspace; black line: decision rule). We iterated over this procedure until all trials had been tested, thus obtaining for each trial a 'neural duration judgment'. In analogy with behavioral judgments, we used two parameter logistic fits to obtain a quantitative description of the performance of simultaneously recorded neurons in categorizing stimuli -the neurometric function (***Figure 3b***, orange curve).

Videometric curves

Full session videos (256x192 pixels resolution) were cut into 3-s long clips with Bonsai (***Lopes et al., 2015***). Individual frames from approximately 75 ms before interval offset were used for this analysis (***Figure 3—figure supplement 1***). This buffer was added to ensure that all frames used preceded stimulus offset. Images were first represented as vectors composed of individual pixel luminance values. Given that image sequences tend to lie on curved low dimensional manifolds in pixel space (***Pless, 2003***), any slight differences in behavioral state reflected in images collected at the offset of short and long interval categories are not necessarily expected to be linearly separable. Thus, we employed isomap (***Tenenbaum and Silva, 2000***), a non-linear dimensionality reduction method, to obtain an information rich yet low dimensional representation of animals' ongoing behavior. This approach has the advantage over tracking methods that it does not make assumptions as to what part of the animals' movements might provide information about stimulus category. The neighborhood size, used to compute the shortest paths between data points, was set to 25 frames to minimize, on average, the dimensionality at which the reconstruction error elbow occurred. In analogy with the neurometric curves, for each stimulus type, we then trained a linear discriminant (leave-one-out cross-validation procedure) to classify frames into those that were recorded during trials where a 'short' or 'long' stimulus interval was presented. The classification was performed in the reduced space determined by isomap. As a positive control for the method, we repeated the same analysis for frames captured at the moment animals expressed their judgment by inserting their snout at one of the two choice ports (***Figure 3—figure supplement 2***). Here, the neighborhood size was chosen to be the minimum for which all frames (from a single session) could be included in a single embedding. This analysis was done for all usable videos (8 out of 11) of sessions in the upper tercile with regard to population size.

Time course of classification performance from neural and video data

To compare how the decoding performance using neural and video data evolved over time, the classification analyses described in Neurometric curves and Videometric curves was performed every 100 ms within a one second window centered around stimulus offset. Video frames at the each time point and neural data in a 200 ms time bin terminating at each time point were used for the analysis. This generated neural and video classification curves that described the ability of simultaneously recorded neural ensembles and video frames to correctly classify interval stimuli as long or short (***Figure 3d***). To determine the relative timing of classification ability in neural ensembles and behavior, we regressed the neural classification curve against the video classification curve for shifts ranging from −300 ms to 300 ms in 100 ms steps (***Figure 3e***).

Psychometric curves split by population state at interval offset

We projected neural activity (of populations composed of simultaneously recorded neurons) on individual trials in high dimensional neural space onto the mean trajectory of those neurons during the delay period for correct trials. Neural activity was defined as the vector of firing rates of the population obtained by convolving spike trains using a causal kernel given by gamma density function with parameters θ = 100 ms and k = 2. We normalized these projections by the length of the mean trajectory of that group of neurons for the longest interval. Pooling normalized projections over all sessions for each animal, we plotted, for each stimulus, distributions of normalized projections at interval offset. To test whether distance traversed along the mean trajectory is predictive of animals' perceptual report, we separated the distribution of pooled projections for each stimulus into 3 bins. Psychometric curves were constructed using trials from each bin. To quantify the key differences between each of these psychometric curves, we performed model comparison using the following 4 parameter logistic function

$$f(x) = d + \frac{a-d}{1+e^{-b(x-c)}}$$

where b controls the slope, c is the inflection point and a and d are the maximum and minimum values of the curve respectively. For two of three animals (Bertrand and Edgar), the model that best accounted for the differences between the three curves (based on Bayesian Information Criterion (BIC) scores) was one with only horizontal shifts between the curves. In the third animal (Fernando), the model that best fit the data was one in which the fit to the three curves differed in both horizontal shift and slope.

A trial's projection on the mean trajectory can be interpreted as a method for decoding time from neural state. Hence, trials that are outliers in the distribution of projections on the mean could potentially correspond to trials where the animal was disengaged. To remove such trials we defined a fraction (60%) of normalized trajectory around the mode of the distribution of pooled projections for each stimulus and excluded trials with projections outside this window.

Population decoder

We decoded elapsed time from striatal population activity using a cross validated, flat prior naive Bayes decoder. For each neuron $n \in [1,N]$, spike counts r_n were observed in 500-ms wide, 10-ms apart, overlapping time bins within the interval period (time referring to the right edge of the bin). For a given r_n, the probability that the current time is t was estimated as the likelihood of observing r_n spikes at time t:

$$P(t|r_n) \propto P(r_n|t)$$

To obtain the likelihood term $P(r_n \mid t)$, we estimated the joint distribution $P(r_n, t)$ by computing, for each time bin, a weighted histogram of spike counts across all correct trials. For trials in which stimulus interval i was presented, spike counts contributed to the histogram with weight w_i defined as the normalized choice variance associated with that interval,

$$w_i = P(C_S|i)P(C_L|i)$$

where C_S and C_L indicate short and long choices respectively. As a result, near boundary interval trials had a greater influence on the estimate of the joint distribution. Histograms were then smoothed using local linear regression (lowess) and normalized to unit area. When decoding from correct trials, leave-one-out cross validation was implemented by computing the joint distribution from all correct trials but the decoded one; incorrect trials were decoded using an estimate of the joint distribution computed from all correct trials. Multi-session population state vectors $\mathbf{r} = (r_1, r_2, \ldots, r_N)$ were obtained by concatenating together data from trials of same stimulus and choice type. By assuming statistical independence between spike counts of different neurons in $\mathbf{r}$, we could compute population estimates of t as the product of single neuron estimates:

$$P(t|\mathbf{r}) \propto \prod_{n=1}^{N} P(r_n|t)$$

Data presented is the average over 100 random concatenations.

Acknowledgements

We would like to thank Bassam Atallah, Brian Lau and Masayoshi Murakami for reading earlier versions of the manuscript, Gustavo Mello for help with surgeries, the Champalimaud Research Vivarium and Histology platforms, Serkan Sülün for pre-processing of video data, João Frazão and Gonçalo Lopes for help with ***Video 1***, João Semedo for technical discussions regarding Isomap, and Bruno Ceña for logistic support.

Additional information

Funding

Funder	Grant reference number	Author
Champalimaud Foundation		Thiago S Gouvêa Tiago Monteiro Asma Motiwala Sofia Soares Christian Machens Joseph J Paton
Simons Foundation	SCGB #325476	Tiago Monteiro Joseph J Paton
Bial Foundation	BIAL-BIC-188/12	Thiago S Gouvêa Tiago Monteiro Sofia Soares Joseph J Paton
Portuguese Foundation for Science and Technology		Thiago S Gouvêa Asma Motiwala Sofia Soares

The funders had no role in study design, data collection and interpretation, or the decision to submit the work for publication.

Author contributions

TSG, TM, JJP, Conception and design, Acquisition of data, Analysis and interpretation of data, Drafting or revising the article; AM, SS, CM, Analysis and interpretation of data, Drafting or revising the article

Author ORCIDs

Thiago S Gouvêa, http://orcid.org/0000-0002-0727-5838
Tiago Monteiro, http://orcid.org/0000-0002-2836-8961
Asma Motiwala, http://orcid.org/0000-0002-7693-2731
Sofia Soares, http://orcid.org/0000-0002-4594-0202
Christian Machens, http://orcid.org/0000-0003-1717-1562
Joseph J Paton, http://orcid.org/0000-0002-7693-2731

Ethics

Animal experimentation: All experiments were in accordance with the European Union Directive 86/609/EEC and approved by the Portuguese Veterinary General Board (Direcção-Geral de Veterinária, project approval 014303 - 0420/000/000/2011)

References

Adler A, Katabi S, Finkes I, Israel Z, Prut Y, Bergman H. 2012. Temporal convergence of dynamic cell assemblies in the striato-pallidal network. *The Journal of Neuroscience* **32**:2473–2484. doi: 10.1523/JNEUROSCI.4830-11.2012

Alexander GE, Crutcher MD. 1990. Neural representations of the target (goal) of visually guided arm movements in three motor areas of the monkey. *Journal of Neurophysiology* **64**:164–178.

Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA. 1996. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience* **13**:87–100. doi: 10.1017/S095252380000715X

Buonomano D, Merzenich M. 1995. Temporal information transformed into a spatial code by a neural network with realistic properties. *Science* **267**:1028–1030. doi: 10.1126/science.7863330

Buonomano DV. 2014. Neural Dynamics Based Timing in the Subsecond to Seconds Range. *Neurobiology of Interval Timing*. New York, NY: Springer New York101–117. doi: 10.1007/978-1-4939-1782-2_6

Gregor Fetterman J, Killeen PR, Hall S. 1998. Watching the clock. *Behavioural Processes* **44**:211–224. doi: 10.1016/S0376-6357(98)00050-3

Geffen MN, Broome BM, Laurent G, Meister M. 2009. Neural encoding of rapidly fluctuating odors. *Neuron* **61**: 570–586. doi: 10.1016/j.neuron.2009.01.021

Gershman SJ, Moustafa AA, Ludvig EA. 2013. Time representation in reinforcement learning models of the basal ganglia. *Frontiers in Computational Neuroscience* **7**.

Gouvêa TS, Monteiro T, Soares S, Atallah BV, Paton JJ. 2014. Ongoing behavior predicts perceptual report of interval duration. *Frontiers in Neurorobotics* **8**. doi: 10.3389/fnbot.2014.00010

Harrington DL, Zimbelman JL, Hinton SC, Rao SM. 2009. Neural modulation of temporal encoding, maintenance, and decision processes. *Cerebral Cortex* **bhp194**.

Hinton SC, Meck WH. 2004. Frontal-striatal circuitry activated by human peak-interval timing in the supra-seconds range. *Brain Research. Cognitive Brain Research* **21**:171–182. doi: 10.1016/j.cogbrainres.2004.08.005

Janssen P, Shadlen MN. 2005. A representation of the hazard rate of elapsed time in macaque area LIP. *Nature Neuroscience* **8**:234–241. doi: 10.1038/nn1386

Jin DZ, Fujii N, Graybiel AM. 2009. Neural representation of time in cortico-basal ganglia circuits. *Proceedings of the National Academy of Sciences* **106**:19156–19161. doi: 10.1073/pnas.0909881106

Jin X, Costa RM. 2010. Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature* **466**: 457–462. doi: 10.1038/nature09263

Killeen PR, Fetterman JG. 1988. A behavioral theory of timing. *Psychological Review* **95**:274–295. doi: 10.1037/0033-295X.95.2.274

Kim J, Ghim JW, Lee JH, Jung MW. 2013. Neural correlates of interval timing in rodent prefrontal cortex. *The Journal of Neuroscience* **33**:13834–13847. doi: 10.1523/JNEUROSCI.1443-13.2013

Lau B, Glimcher PW. 2008. Value representations in the primate striatum during matching behavior. *Neuron* **58**: 451–463. doi: 10.1016/j.neuron.2008.02.021

Lebedev MA, O'Doherty JE, Nicolelis MA. 2008. Decoding of temporal intervals from cortical ensemble activity. *Journal of Neurophysiology* **99**:166–186. doi: 10.1152/jn.00734.2007

Lee AM, Tai L-H, Zador A, Wilbrecht L. 2015. Between the primate and 'reptilian' brain: rodent models demonstrate the role of corticostriatal circuits in decision making. *Neuroscience* **296**:66–74. doi: 10.1016/j.neuroscience.2014.12.042

Leon MI, Shadlen MN. 2003. Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron* **38**:317–327. doi: 10.1016/S0896-6273(03)00185-5

Lopes Gonçalo, Bonacchi Niccolò, Frazão João, Neto JP, Atallah BV, Soares S, Moreira Luis, Matias S, Itskov PM, Correia Patrícia A, Medina RE, Calcaterra L, Dreosti E, Paton JJ, Kampff AR. 2015. Bonsai: an event-based framework for processing and controlling data streams. *Frontiers in Neuroinformatics* **9**. doi: 10.3389/fninf.2015.00007

MacDonald CJ, Lepage KQ, Eden UT, Eichenbaum H. 2011. Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron* **71**:737–749. doi: 10.1016/j.neuron.2011.07.012

Machado A. 1997. Learning the temporal dynamics of behavior. *Psychological Review* **104**:241–265. doi: 10.1037/0033-295X.104.2.241

Machado A, Keen R. 2003. Temporal discrimination in a long operant chamber. *Behavioural Processes* **62**:157–182. doi: 10.1016/S0376-6357(03)00023-8

Malapani C, Rakitin B, Levy R, Meck WH, Deweer B, Dubois B, Gibbon J. 1998. Coupled temporal memories in parkinson's disease: a dopamine-related dysfunction. *Journal of Cognitive Neuroscience* **10**:316–331. doi: 10.1162/089892998562762

Matell MS, Meck WH, Nicolelis MAL. 2003. Interval timing and the encoding of signal duration by ensembles of cortical and striatal neurons. *Behavioral Neuroscience* **117**:760–773. doi: 10.1037/0735-7044.117.4.760

Matthews TJ, Lerer BE. 1987. Behavior patterns in pigeons during autoshaping with an incremental conditioned stimulus. *Animal Learning & Behavior* **15**:69–75. doi: 10.3758/BF03204906

Mauk MD, Buonomano DV. 2004. The neural basis of temporal processing. *Annual Review of Neuroscience* **27**: 307–340. doi: 10.1146/annurev.neuro.27.070203.144247

Meck WH. 2006. Neuroanatomical localization of an internal clock: a functional link between mesolimbic, nigrostriatal, and mesocortical dopaminergic systems. *Brain Research* **1109**:93–107. doi: 10.1016/j.brainres.2006.06.031

Mello GBM, Soares S, Paton JJ. 2015. A scalable population code for time in the striatum. *Current Biology* **25**: 1113–1122. doi: 10.1016/j.cub.2015.02.036

Pastalkova E, Itskov V, Amarasingham A, Buzsáki G. 2008. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**:1322–1327. doi: 10.1126/science.1159775

Pless R. 2003. Image spaces and video trajectories: using isomap to explore video sequences. **3**:1433–1440 .

Samejima K, Ueda Y, Doya K, Kimura M. 2005. Representation of action-specific reward values in the striatum. *Science* **310**:1337–1340. doi: 10.1126/science.1115270

Steiner H, Tseng KY. 2010. *Handbook of Basal Ganglia Structure and Function: A Decade of Progress*. 24. Academic Press.

Tenenbaum JB, de Silva V, Langford JC. 2000. A global geometric framework for nonlinear dimensionality reduction. *Science* **290**:2319–2323. doi: 10.1126/science.290.5500.2319

Wencil EB, Coslett HB, Aguirre GK, Chatterjee A. 2010. Carving the clock at its component joints: neural bases for interval timing. *Journal of Neurophysiology* **104**:160–168. doi: 10.1152/jn.00029.2009

Xu M, Zhang S-y, Dan Y, Poo M-m. 2014. Representation of interval timing by temporally scalable firing patterns in rat prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America* **111**: 480–485. doi: 10.1073/pnas.1321314111

Appendix B

PhD in brief

B.1 Annual activities report

Fac simile of annual activities report presented to the academic services of Instituto de Tecnologia Química e Biológica / Universidade Nova de Lisboa.

Annual activities report

Thiago Santos Gouvêa

International Neuroscience Doctoral Programme
Champalimaud Centre for the Unknown, Lisbon, Portugal

February 17, 2016

2009/2010 I joined the International Neuroscience Doctoral Programme (INDP) at Instituto Gulbenkian de Ciência in September 2009. In the first year, I completed the INDP compulsory courses on biology and neuroscience, and I attended the Federation of European Neuroscience Societies (FENS) Forum 2010 in Amsterdam. At the end of the period I joined the group of Dr. Zachary F. Mainen to study the role of the basal ganglia in perceptual decision making. The project proposed at this time was later presented at the Sociedade Portuguesa de Neurociências [1].

2010/2011 I attended the 2010 Autumn School in Cognitive Neuroscience in Oxford, UK, a school featuring leading specialists in basal ganglia research. I also attended the Computational and Systems Neuroscience (Cosyne) 2011, the main conference worldwide in my field. During this year spent in the Mainen lab I learned to train animals in psychophysical tasks, and conducted pilot experiments involving pharmacological manipulation of the dopaminergic system. I presented these results later as a poster at Cosyne 2012 [2].

2011/2012 Building on the experience studying perceptual decision making acquired during the year in Dr. Mainen's lab, I joined the group of Dr. Joseph J. Paton to study the representation of subjective time in the dynamics of neural populations in the striatum. Under Dr. Paton's supervision, I designed a psychophysical task to study time perception in rats. This work was among the first to be carried at the recently inaugurated Champalimaud Centre for the Unknown (CCU), and involved designing and assembling the necessary hardware (behavioral chamber with its sensors and effectors, break out board interfacing the chamber with a computer), writing software for control of the behavioral chamber (Labview, Python, Arduino) and data analysis (Matlab), and optimizing the task contingencies that would successfully shape animals' behavior. The project was presented to the Champalimaud neuroscience community as a talk in the Champalimaud Internal Seminar Series (CISS) [3]. The behavioral results yielded a poster presented at the FENS Forum 2012 [4] as part of a satellite symposium focused on neural network dynamics. Around the same time, my labmate Gustavo Mello and I performed surgeries to implant electrode arrays on animals trained in the task, and I started collecting neurophysiology data. The first neural results were presented in preliminary form at Society for Neuroscience (SfN) 2012, the world's largest Neuroscience conference [5]. The data I generated in the Mainen lab formed the basis of another work presented at SfN 2012 [6]. In parallel, I helped establish the science communication platform *Ar – Respire Connosco* (Air – Breathe With Us) (`http://ar.neuro.fchampalimaud.org/`). My contributions included participation in planning, organizing and hosting interactive science talks for non-specialists, as well as promoting the platform at SfN 2012 [7].

2012/2013 During this year, I and my lab colleague and collaborator Tiago Monteiro acquired the bulk of our high speed video and neurophysiology data. My main focus during this year was on developing the analytical skills required by these rich datasets, as well as drafting our first manuscript. I presented our preliminary neurophysiology results as a talk in the seminar series of the Cognition and Complex Systems unit at *Universidade Federal do ABC* in São Paulo, Brazil [8]. I was one of the INDP delegates at the 1st European Neuroscience Conference by Doctoral Students (ENCODS) where, besides presenting my work [9], I and other INDP colleagues met with the conference founders. As a consequence, the 2nd edition of this FENS-sponsored conference took place in Sesimbra, Portugal, in April 2015 — though I did not directly participate in the organization of ENCODS 2. I also presented our behavioral and neurophysiology results at a prestigious conference on neural dynamics and learning [10].

2013/2014 This academic year begun with another oral presentation to the Champalimaud neuroscience community [11]. Additionally, I presented our work in poster format at two conferences taking place at CCU [12, 13], as well as at Cosyne 2014 [14]. The first article resulting from my PhD work was submitted, accepted and published in an open-access, peer-reviewed journal [15], and I started drafting our second manuscript. My lab colleague Sofia Soares managed to train mice, a genetically tractable species, to perform our psychophysical task, opening valuable experimental possibilities for the project. This branch of the project was presented by Sofia as a poster at SfN 2013 [16]. Asma Motiwala, a member of Dr. Christian Machens's group and a theorist who works closely with us, presented analysis methods she developed for our dataset as a poster at an international conference [17]. I joined a group of artists and neuroscientists at the project *Raízes da Curiosidade* (Roots of Curiosity), a co-production of Centro Cultural de Belém (CCB) and the Champalimaud Foundation. The aim of this project was to explore the common ground between art and science, and perhaps inaugurate a new, third ground. Concretely, during this year I participated in two artistic residencies, and had several creative meetings with my artist partner Catarina Vasconcelos.

2014/2015 As in the previous academic year, this one begun with an oral presentation to the Champalimaud neuroscience community [18]. After one last artistic residency, the *Raízes da Curiosidade* project culminated in a staged show presented at CCB. Additionally, we partook in a conference on art and science in CCU, organized a series of workshops for children and families in CCB's educational department *Fábrica das Artes*, wrote a book about the project [19], and featured on a documentary film released together with the book. My lab colleague Tiago Monteiro and I attended to a FENS-sponsored winter school featuring leading decision making specialists where we presented our projects — my main project and a spin off led by him — as posters [20, 21]. An abstract based on our work ranked among the top 5.4% submissions to Cosyne 2015, thus giving me the opportunity to present it as a plenary talk [22]. We deposited a draft of our second paper in an open preprint repository [23] at the same time as we submitted it to a peer-reviewed journal. The article was accepted and published in a prominent open access journal a few months later [24]. The mouse branch of the project, led by my colleague Sofia Soares, yielded a poster at SfN 2015 [25], an oral presentation at Cosyne 2016 [26], and a manuscript (in preparation). I developed a variant of our psychophysical timing task that will allow disambiguation between action specific time-varying neural signals and a more abstract neural representation of the passage of time. Neural recordings will be carried by my lab collaborators. Also during this year I was interviewed for postdoctoral positions at several labs in Germany and the United States. I have accepted an offer from Dr. Adam Kepecs, and I will join his group at the Cold Spring Harbor Laboratory soon after my thesis defense. Finally, I am currently finishing to write my thesis, and I expect to defend early in May 2016.

References

[1] Thiago S. Gouvêa, E. E. J. DeWitt, Joseph J. Paton, and Z. F. Mainen. Motivational modulation of perceptual decisions. In *Sociedade Portuguesa de Neurociências*, Lisbon, Portugal, May 2011.

[2] Thiago S. Gouvêa, Gil Costa, Eric DeWitt, and Zachary Mainen. Trial-by-trial perceptual learning during odor category decisions: Value, uncertainty and dopamine. In *Cosyne*, Salt Lake City, UT, February 2012.

[3] T. S. Gouvêa. Labels and stamps make a learning agent. In *Champalimaud Internal Seminar Series*, Lisbon, Portugal, March 2012.

[4] Thiago Gouvêa and Joseph J. Paton. Temporal judgment: an embodied strategy? In *Dynamics of Memory: What's the Evidence? (FENS 2012 Satellite)*, Barcelona, Spain, July 2012.

[5] T. S. Gouvêa and J. J. Paton. Temporal judgment in the rat: Behavior and model. In *Society for Neuroscience*, New Orleans, LA, October 2012.

[6] E. J. DeWitt, A. G. Mendonça, T. S. Gouvêa, A. Kepecs, and Z. F. Mainen. Trial-by-trial perceptual learning during odor category decisions: Reinforcement learning, uncertainty and modeling choice. In *Society for Neuroscience*, New Orleans, LA, October 2012.

[7] T. S. Gouvêa, B. A. C. Afonso, C. Afonso, N Bonnachi, W. Brendel, V. M. Corrales, P. A. Correia, G. M. P. Costa, F. A. Dehmelt, E. E. J. DeWitt, A. R. Fonseca, A. F. Hobbiss, S. Lackner, G. C. Lopes, T. G. Marques, S. P. S. Matias, A. G. Mendonça, S. V. Meyler, C. Monroy, C. E. L. Ramos, Rennie S. M., S. L. S. Soares, R. Venturini, A. Vicente, E. Vinnik, A. R. Kampff, and Z. F. Mainen. Lisbon breathes science: Public, students, and leading scientists brought together in ar. In *Society for Neuroscience*, New Orleans, LA, October 2012.

[8] T. S. Gouvêa. On embodied cognition, the basal ganglia, and neural codes for time. In *Ciclo de Seminários em Cognição (UFABC - Núcleo de Cognição e Sistemas Complexos)*, Santo André, SP, Brazil, April 2013.

[9] T. S. Gouvêa, T. Monteiro, S. L. S. Soares, and J. J. Paton. On embodied cognition, the basal ganglia, and neural codes for time. In *ENCODS - European Neuroscience Conference by Doctoral Students*, Bordeaux, France, April 2013.

[10] T. S. Gouvêa, T. Monteiro, S. L. S. Soares, and J. J. Paton. On embodied cognition, the basal ganglia, and neural codes for time. In *Temporal Dynamics in Learning: Networks and Neural Data (Janelia Conferences)*, Howard Hughes Medical Institute, Janelia Farm Research Campus, Ashburn, VA, May 2013.

[11] T. S. Gouvêa. Time perception reflects body dynamics. In *Champalimaud Internal Seminar Series*, Lisbon, Portugal, September 2013.

[12] T. S. Gouvêa and J. J. Paton. Time perception reflects body dynamics. In *Champalimaud Neuroscience Symposium*, Lisbon, Portugal, September 2013.

[13] T. S. Gouvêa, A. Motiwala, S. L. Soares, B. V. Atallah, T. Monteiro, and J. J. Paton. Time perception reflects body dynamics. In *10^o Congresso Nacional de Etologia*, Lisbon, Portugal, October 2013.

[14] Thiago S. Gouvêa, T. Monteiro, A. Motiwala, and J. J. Paton. Continuous psychometric-neurometric comparison in a perceptual decision making task. In *Cosyne*, Salt Lake City, UT, February 2014.

[15] Thiago S Gouvêa, Tiago Monteiro, Sofia Soares, Bassam V Atallah, and Joseph J Paton. Ongoing behavior predicts perceptual report of interval duration. *Frontiers in Neurorobotics*, 8, 2014.

[16] Sofia Soares, B. V. Atallah, Thiago S. Gouvêa, Tiago Monteiro, and Joseph J. Paton. Mice perform action sequences that predict choice during temporal discrimination. In *Society for Neuroscience*, San Diego, CA, November 2013.

[17] A. Motiwala, T. Gouvêa, T. Monteiro, J. Paton, and C. Machens. Quantifying continuous behavior to account for neural variability. In *9th International Conference on Methods and Techniques in Behavioral Research*, Wageningen, The Netherlands, August 2014.

[18] T. S. Gouvêa. Striatal dynamics represent perceived time. In *Champalimaud Internal Seminar Series*, Lisbon, Portugal, September 2014.

[19] Thiago Gouvêa. In Madalena Wallenstein, editor, *Raízes da Curiosidade - Tempo de Ciência e Arte*, number II in Transversalidades. Centro Cultural de Belém, Lisbon, Portugal, November 2015.

[20] Thiago S Gouvêa, Tiago Monteiro, Asma Motiwala, Sofia Soares, Christian Machens, and Joseph J Paton. Striatal dynamics represent perceived time. In *FENS-Hertie Winter School — The Neuroscience of Decision Making*, Obergurgl, Austria, January 2015.

[21] Tiago Monteiro, Gonçalo Lopes, Thiago S Gouvêa, Sofia Soares, and Joseph J Paton. Implementation of a closed-loop 2 alternative forced choice timing task. In *FENS-Hertie Winter School — The Neuroscience of Decision Making*, Obergurgl, Austria, January 2015.

[22] Thiago S Gouvêa, Tiago Monteiro, Asma Motiwala, Sofia Soares, Christian Machens, and Joseph J Paton. Striatal dynamics explain duration judgments. In *Cosyne*, Salt Lake City, UT, March 2015.

[23] Thiago S Gouvêa, Tiago Monteiro, Asma Motiwala, Sofia Soares, Christian K Machens, and Joseph J Paton. Striatal dynamics explain duration judgments. *bioRxiv*, page 020883, 2015.

[24] Thiago S Gouvêa, Tiago Monteiro, Asma Motiwala, Sofia Soares, Christian Machens, and Joseph J Paton. Striatal dynamics explain duration judgments. *eLife*, page e11386, 2015.

[25] Sofia Soares, B. V. Atallah, Thiago S. Gouvêa, Tiago Monteiro, and Joseph J. Paton. A new role for midbrain dopaminergic neurons during temporal judgments. In *Society for Neuroscience*, Chicago, IL, November 2015.

[26] Sofia Soares, Bassam Atallah, Alessandro Braga, Thiago Gouvêa, Tiago Monteiro, and Joe Paton. Midbrain dopamine neurons directly modulate duration judgments. In *Cosyne*, Salt Lake City, UT, February 2016.

References

Adler, A., Katabi, S., Finkes, I., Israel, Z., Prut, Y., & Bergman, H. (2012). Temporal convergence of dynamic cell assemblies in the striato-pallidal network. *The Journal of Neuroscience*, *32*(7), 2473–2484.

Alexander, G. E., & Crutcher, M. D. (1990). Neural representations of the target (goal) of visually guided arm movements in three motor areas of the monkey. *Journal of Neurophysiology*, *64*(1), 164–178.

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual review of neuroscience*, *9*(1), 357–381.

Allen Institute for Brain Science. (2015). *Allen mouse brain connectivity atlas.* Available from `http://connectivity.brain-map.org`.

Anderson, M. C., & Shettleworth, S. J. (1977). Behavioral adaptation to fixed-interval and fixed-time food delivery in golden hamsters. *Journal of the Experimental Analysis of Behavior*, *27*(1), 33–49. Retrieved from `http://www.ncbi.nlm.nih.gov/pubmed/1333550http://www.pubmedcentral.gov/articlerender.fcgi?artid=1333550` doi: 10.1901/jeab.1977.27-33

Artieda, J., Pastor, M. A., Lacruz, F., & Obeso, J. A. (1992). Temporal discrimination is abnormal in parkinson's disease. *Brain*, *115*(1), 199–210.

Balcı, F., Papachristos, E., Gallistel, C., Brunner, D., Gibson, J., & Shumyatsky, G. (2008). Interval timing in genetically modified mice: a simple paradigm. *Genes, Brain and Behavior*, *7*(3), 373–384.

Balleine, B. W., Liljeholm, M., & Ostlund, S. B. (2009). The integrative function of the basal ganglia in instrumental conditioning. *Behavioural brain research*, *199*(1), 43–52.

Balsam, P. D., & Gallistel, C. R. (2009). Temporal maps and informativeness in associative learning. *Trends in neurosciences*, *32*(2), 73–8. Retrieved from `http://www.ncbi.nlm.nih.gov/pubmed/19136158` doi: 10.1016/j.tins.2008.10.004

Bateson, M. (2003). Interval timing and optimal foraging. *Functional and neural* *mechanisms of interval timing*, 113–141.

Bechtel, W. (2005). The challenge of characterizing operations in the mechanisms underlying behavior. *Journal of the experimental analysis of behavior*, *84*(3), 313–325.

Beck, J. M., Kanitscheider, I., & Pouget, A. (2012). A normative theory of weber's law. In *Cosyne.* Salt Lake City, UT.

Bennett, M. R., & Hacker, P. M. S. (2003). *Philosophical foundations of neuroscience.* Blackwell Publishing.

Bernacchia, A., Seo, H., Lee, D., & Wang, X.-J. (2011). A reservoir of time constants for memory traces in cortical neurons. *Nature neuroscience*, *14*(3), 366–372.

Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., & Movshon, J. A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual neuroscience*, *13*(1), 87–100. Retrieved from `http://www.ncbi.nlm.nih.gov/pubmed/8730992`

Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *12*(12), 4745–65. Retrieved from `http://www.ncbi.nlm.nih.gov/pubmed/1464765`

Buonomano, D. V. (2014). Neural dynamics based timing in the subsecond to seconds range. In *Neurobiology of interval timing* (pp. 101–117). Springer.

Buonomano, D. V., & Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, *10*(2), 113–125.

Buonomano, D. V., & Mauk, M. D. (1994). Neural network model of the cerebellum: temporal discrimination and the timing of motor responses. *Neural Computation*, *6*(1), 38–55.

Buonomano, D. V., & Merzenich, M. M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. *Science*, 1028–1028.

Catania, A. C. (1999). Aprendizagem: comportamento, linguagem e cognição.(trad. dg souza). *Porto Alegre: ArtMed.*

Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical population biology*, *9*(2), 129–136.

Church, R. M., & Deluty, M. Z. (1977). Bisection of temporal intervals. *Journal of Experimental Psychology: Animal Behavior Processes*, *3*(3), 216.

Creelman, C. D. (1962). Human discrimination of auditory duration. *The Journal of the Acoustical Society of America*, *34*(5), 582–593.

Crozier, W., & Stier, T. (1925). Critical thermal increments for rhythmic respiratory movements of insects. *The Journal of general physiology*, *7*(3), 429–447.

Cumming, W., & Schoenfeld, W. (1958). Behavior under extended exposure to a high-value fixed interval reinforcement schedule. *Journal of the experimental analysis of behavior*, *1*(3), 245–263.

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, *8*(12), 1704–1711.

Dieguez, M. (n.d.). *http://www.historiadordofutebol.com.br/antigo.php?secao=bebeto.*

Douw, L., Schoonheim, M., Landi, D., Van der Meer, M., Geurts, J., Reijneveld, J., ... Stam, C. (2011). Cognition is related to resting-state small-world network topology: an magnetoencephalographic study. *Neuroscience*, *175*, 169–177.

Falcon, A. (2015). Aristotle on causality. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Spring 2015 ed.). `http://plato.stanford.edu/archives/spr2015/entries/aristotle-causality/`.

Ferster, C. B., & Skinner, B. F. (1957). Schedules of reinforcement.

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*, *299*(5614), 1898–1902.

Freedman, D. J., & Assad, J. A. (2006). Experience-dependent representation of visual categories in parietal cortex. *Nature*, *443*(7107), 85–88.

Fuster, J. M. (2001). The prefrontal cortex—an update: time is of the essence. *Neuron*, *30*(2), 319–333.

Geffen, M. N., Broome, B. M., Laurent, G., & Meister, M. (2009). Neural encoding of rapidly fluctuating odors. *Neuron*, *61*(4), 570–586.

Gershman, S. J., Moustafa, A. A., & Ludvig, E. A. (2013). Time representation in reinforcement learning models of the basal ganglia. *Frontiers in computational neuroscience*, *7*.

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, *84*(3), 279–325. Retrieved from `http://content.apa.org/journals/rev/84/3/279` doi: 10.1037/0033-295X.84.3.279

Gibbon, J., & Balsam, P. (1981). Spreading association in time. *Autoshaping and conditioning theory*, 219–253.

Gibbon, J., & Church, R. M. (1984). Sources of variance in an information processing theory of timing. In H. L. Roitblat, H. S. Terrace, & T. G. Bever (Eds.), (Vol. Animal Cognition, chap. 26). Psychology Press.

Goel, A., & Buonomano, D. V. (2014). Timing as an intrinsic property of neural networks: evidence from in vivo and in vitro experiments. *Philosophical transactions of the Royal Society B: Biological sciences*, *369*(1637), 20120460.

Gormezano, I., & Kehoe, E. J. (1981). Classical conditioning and the law of contiguity. *Advances in analysis of behavior*, *2*, 1–45.

Goudar, V., & Buonomano, D. (2015). Discrimination and production of spatiotemporal patterns with a single recurrent neural network. In *Cosyne.* Salt Lake City, UT.

Gouvêa, T. S., Monteiro, T., Motiwala, A., Soares, S., Machens, C. K., & Paton, J. J. (2015a). Striatal dynamics explain duration judgments. *eLife*, e11386.

Gouvêa, T. S., Monteiro, T., Motiwala, A., Soares, S., Machens, C. K., & Paton, J. J. (2015b). Striatal dynamics explain duration judgments. *bioRxiv*, 020883.

Gouvêa, T. S., Monteiro, T., Soares, S., Atallah, B. V., & Paton, J. J. (2014). Ongoing behavior predicts perceptual report of interval duration. *Frontiers in Neurorobotics*, *8*.

Green, D. M., Swets, J. A., et al. (1966). *Signal detection theory and psychophysics* (Vol. 1). Wiley New York.

Grondin, S. (2010). Timing and time perception: a review of recent behavioral and neuroscience findings and theoretical directions. *Attention, Perception, & Psychophysics*, *72*(3), 561–582.

Grossberg, S., & Schmajuk, N. A. (1989). Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Networks*, *2*(2), 79–102. Retrieved from `http://linkinghub.elsevier.com/retrieve/pii/0893608089900269` doi: 10.1016/0893-6080(89)90026-9

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806.

Haight, P. A., & Killeen, P. R. (1991). Adjunctive behavior in multiple schedules of reinforcement. *Animal Learning & Behavior*, *19*(3), 257–263. Retrieved from `http://www.springerlink.com/index/10.3758/BF03197884` doi: 10.3758/BF03197884

Hanks, T. D., Kopec, C. D., Brunton, B. W., Duan, C. A., Erlich, J. C., & Brody, C. D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature*, *520*(7546), 220–223.

Harrington, D. L., Zimbelman, J. L., Hinton, S. C., & Rao, S. M. (2009). Neural modulation of temporal encoding, maintenance, and decision processes. *Cerebral Cortex*, bhp194.

Harvey, C. D., Coen, P., & Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*, *484*(7392), 62–68.

Hikosaka, O., Takikawa, Y., & Kawagoe, R. (2000). Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiological reviews*, *80*(3), 953–978.

Hinton, S. C., & Meck, W. H. (2004). Frontal–striatal circuitry activated by human peak-interval timing in the supra-seconds range. *Cognitive Brain Research*, *21*(2), 171–182.

Hoagland, H. (1933). The physiological control of judgments of duration: Evidence for a chemical clock. *The Journal of General Psychology*, *9*(2), 267–287.

Hodos, W., Ross, G. S., & Brady, J. V. (1962). Complex response patterns during temporally spaced responding. *Journal of the experimental analysis of behavior*, *5*(4), 473–479.

Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, *46*(3), 269–299.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, *160*(1), 106–154.

Ivry, R. B., & Keele, S. W. (1989). Timing functions of the cerebellum. *Journal of Cognitive Neuroscience*, *1*(2), 136–152.

Janssen, P., & Shadlen, M. N. (2005). A representation of the hazard rate of elapsed time in macaque area lip. *Nature neuroscience*, *8*(2), 234–241.

Jazayeri, M., & Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nature neuroscience*, *13*(8), 1020–1026.

Jazayeri, M., & Shadlen, M. N. (2015). A neural mechanism for sensing and reproducing a time interval. *Current Biology*, *25*(20), 2599–2609.

Jin, D. Z., Fujii, N., & Graybiel, A. M. (2009). Neural representation of time in cortico-basal ganglia circuits. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(45), 19156–61. Retrieved from `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2776432&tool=pmcentrez&rendertype=abstract` doi: 10.1073/pnas.0909881106

Jin, X., & Costa, R. M. (2010). Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature*, *466*(7305), 457–462.

Kacelnik, A., & Brunner, D. (2002). Timing and Foraging: Gibbon's Scalar Expectancy Theory and Optimal Patch Exploitation. *Learning and Motivation*, *33*(1), 177–195. Retrieved from `http://linkinghub.elsevier.com/retrieve/pii/S0023969001911105` doi: 10.1006/lmot.2001.1110

Kandel, E. R., Schwartz, J. H., & Jessell, T. M. (Eds.). (2000). *Principles of neural science.* McGraw-Hill.

Karmarkar, U. R., & Buonomano, D. V. (2007). Timing in the absence of clocks: encoding time in neural network states. *Neuron*, *53*(3), 427–438.

Killeen, P. R., & Fetterman, J. G. (1988). A behavioral theory of timing. *Psychological review*, *95*(2), 274–95. Retrieved from `http://www.ncbi.nlm.nih.gov/pubmed/3375401`

Kim, J., Ghim, J.-W., Lee, J. H., & Jung, M. W. (2013). Neural correlates of interval timing in rodent prefrontal cortex. *The Journal of Neuroscience*, *33*(34), 13834–13847.

Kim, J., Jung, A. H., Byun, J., Jo, S., & Jung, M. W. (2009). Inactivation of medial prefrontal cortex impairs time interval discrimination in rats. *Frontiers in behavioral neuroscience*, *3*.

Knudsen, E. B., Powers, M. E., & Moxon, K. A. (2014). Dissociating movement from movement timing in the rat primary motor cortex. *The Journal of Neuroscience*, *34*(47), 15576–15586.

Lau, B., & Glimcher, P. W. (2005). Dynamic response-by-response models of matching behavior in rhesus monkeys. *Journal of the Experimental Analysis of Behavior*.

Lau, B., & Glimcher, P. W. (2008). Value representations in the primate striatum during matching behavior. *Neuron*, *58*(3), 451–463.

Lebedev, M. A., O'Doherty, J. E., & Nicolelis, M. A. (2008). Decoding of temporal intervals from cortical ensemble activity. *J. Neurophysiol.*, *99*(1), 166–186.

Lee, A., Tai, L.-H., Zador, A., & Wilbrecht, L. (2015). Between the primate and 'reptilian' brain: rodent models demonstrate the role of corticostriatal circuits in decision making. *Neuroscience*, *296*, 66–74.

Leon, M. I., & Shadlen, M. N. (2003). Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron*, *38*(2), 317–327.

Lopes, G., Bonacchi, N., Frazão, J., Neto, J. P., Atallah, B. V., Soares, S., ... others (2015). Bonsai: An event-based framework for processing and controlling data streams. *Frontiers in neuroinformatics*, *9*.

Ludvig, E. A., Sutton, R. S., & Kehoe, E. J. (2008). Stimulus representation and the timing of reward-prediction errors in models of the dopamine system. *Neural computation*, *20*(12), 3034–54. Retrieved from `http://www.ncbi.nlm.nih.gov/pubmed/18624657` doi: 10.1162/neco.2008.11-07-654

Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural computation*, *14*(11), 2531–2560.

MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron*, *71*(4), 737–749.

Machado, A. (1997). Learning the temporal dynamics of behavior. *Psychological Review*, *104*(2), 241–265. Retrieved from `http://www.ncbi.nlm.nih.gov/pubmed/9127582http://doi.apa.org/getdoi.cfm?doi=10.1037/0033-295X.104.2.241` doi: 10.1037/0033-295X.104.2.241

Machado, A., & Keen, R. (2003). Temporal discrimination in a long operant chamber. *Behavioural Processes*, *62*(1-3), 157–182. Retrieved from `http://`

`linkinghub.elsevier.com/retrieve/pii/S0376635703000238` doi: 10.1016/S0376-6357(03)00023-8

Machens, C. K., Romo, R., & Brody, C. D. (2010). Functional, but not anatomical, separation of what and when in prefrontal cortex. *The Journal of Neuroscience*, *30*(1), 350–360.

Malapani, C., Rakitin, B., Levy, R. S., Meck, W. H., Deweer, B., Dubois, B., & Gibbon, J. (1998). Coupled temporal memories in parkinson's disease: a dopamine-related dysfunction. *Cognitive Neuroscience, Journal of*, *10*(3), 316–331.

Maricq, A. V., & Church, R. M. (1983). The differential effects of haloperidol and methamphetamine on time estimation in the rat. *Psychopharmacology*, *79*(1), 10–15. Retrieved from `http://link.springer.com/10.1007/BF00433008` doi: 10.1007/BF00433008

Matell, M. S., & Meck, W. H. (2000). Neuropsychological mechanisms of interval timing behavior. *Bioessays*, *22*(1), 94–103.

Matell, M. S., & Meck, W. H. (2004). Cortico-striatal circuits and interval timing: coincidence detection of oscillatory processes. *Cognitive brain research*, *21*(2), 139–170.

Matell, M. S., Meck, W. H., & Nicolelis, M. A. L. (2003). Interval timing and the encoding of signal duration by ensembles of cortical and striatal neurons. *Behavioral Neuroscience*, *117*, 760-773.

Matthews, T. J., & Lerer, B. E. (1987). Behavior patterns in pigeons during autoshaping with an incremental conditioned stimulus. *Animal Learning & Behavior*, *15*(1), 69–75.

Mauk, M. D., & Buonomano, D. V. (2004). The neural basis of temporal processing. *Annu. Rev. Neurosci.*, *27*, 307–340.

McHaffie, J. G., Stanford, T. R., Stein, B. E., Coizet, V., & Redgrave, P. (2005). Subcortical loops through the basal ganglia. *Trends in neurosciences*, *28*(8), 401–407.

Meck, W. H. (2006). Neuroanatomical localization of an internal clock: a functional link between mesolimbic, nigrostriatal, and mesocortical dopaminergic systems. *Brain research*, *1109*(1), 93–107.

Meck, W. H., & Buhushi, C. V. (2010). Encyclopedia of psychopharmacology. In I. P. Stolerman (Ed.), (pp. 977–977). Berlin, Heidelberg: Springer Berlin Heidelberg. Retrieved from `http://dx.doi.org/10.1007/978-3-540-68706-1_1658` doi: 10.1007/978-3-540-68706-1_1658

Meck, W. H., Penney, T. B., & Pouthas, V. (2008). Cortico-striatal representation of time in animals and humans. *Current opinion in neurobiology*, *18*(2), 145–152.

Mello, G. B., Soares, S., & Paton, J. J. (2015). A scalable population code for time in the striatum. *Current Biology*, *25*(9), 1113–1122.

Menzies, P. (2014). Counterfactual theories of causation. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Spring 2014 ed.). `http://plato.stanford.edu/archives/spr2014/entries/causation-counterfactual/`.

Merchant, H., & de Lafuente, V. (2014). *Neurobiology of interval timing* (Vol. 829). Springer.

Miall, C. (1989). The storage of time intervals using oscillating neurons. *Neural Computation*, *1*(3), 359–371.

Michon, J. A. (1967). *Timing in temporal tracking.* Institute for Perception RVO-TNO Soesterberg, The Netherlands.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... others (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533.

Mountcastle, V. B., Talbot, W. H., Sakata, H., & Hyvarinen, J. (1969). Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys: Neuronal periodicity and frequency discrimination. *Journal of Neurophysiology.*

Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., ... others (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature neuroscience*, *17*(12), 1661–1663.

Nienborg, H., R. Cohen, M., & Cumming, B. G. (2012). Decision-related activity in sensory neurons: correlations among neurons and with behavior. *Annual review of neuroscience*, *35*, 463–483.

Nilsson, N. J. (1991). Logic and artificial intelligence. *Artificial Intelligence*, *47*(1-3), 31–56.

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain research*, *34*(1), 171–175.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map.* Clarendon Press Oxford.

Ölveczky, B. P. (2011). Motoring ahead with rodents. *Current opinion in neurobiology*, *21*(4), 571–8. Retrieved from `http://www.ncbi.nlm.nih.gov/pubmed/21628098` doi: 10.1016/j.conb.2011.05.002

Palnitkar, R. M., & Cannady, J. (2004). A review of adaptive neural networks. In *Southeastcon, 2004. proceedings. ieee* (pp. 38–47).

Parker, A. J., & Newsome, W. T. (1998). Sense and the single neuron: probing the physiology of perception. *Annual review of neuroscience*, *21*(1), 227–277.

Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, *321*(5894), 1322–1327.

Pastor, M., Artieda, J., Jahanshahi, M., & Obeso, J. (1992). Time estimation and reproduction is abnormal in parkinson's disease. *Brain*, *115*(1), 211–225.

Platt, J. R., & Davis, E. R. (1983). Bisection of temporal intervals by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *9*(2), 160.

Pless, R. (2003). Image spaces and video trajectories: Using isomap to explore video sequences. In *Iccv* (Vol. 3, pp. 1433–1440).

Rescorla, R. A. (1966). Predictability and number of pairings in pavlovian fear conditioning. *Psychonomic Science*, *4*(11), 383–384.

Rescorla, R. A. (1967). Pavlovian conditioning and its proper control procedures. *Psychological review*, *74*(1), 71.

Rescorla, R. A. (1988a). Behavioral studies of pavlovian conditioning. *Annual review of neuroscience*, *11*(1), 329–352.

Rescorla, R. A. (1988b). Pavlovian conditioning: It's not what you think it is. *American Psychologist*, *43*(3), 151.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning ii: Current research and theory* (pp. 64–99). Appleton-Century-Crofts.

Rice, F. (1923). A theory of chemical reactivity. *Journal of the American Chemical Society*, *45*(12), 2808–2820.

Romo, R., Brody, C. D., Hernández, A., & Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature*, *399*(6735), 470–473.

Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, *65*(6), 386.

Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1988). A general framework for parallel distributed processing. In D. E. Rumelhart, J. L. McClelland, P. R. Group, et al. (Eds.), (Vol. Parallel distributed processing). IEEE.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*, 533–536.

Russell, S. J., & Norvig, P. (2003). *Artificial intelligence: a modern approach (international edition)* (2nd ed.). Prentice Hall.

Salzman, C. D., Britten, K. H., & Newsome, W. T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, *346*(6280), 174–177.

Samejima, K., Ueda, Y., Doya, K., & Kimura, M. (2005). Representation of action-specific reward values in the striatum. *Science*, *310*(5752), 1337–1340.

Schneider, B. A. (1969). A two-state analysis of fixed-interval responding in the pigeon. *Journal of the Experimental Analysis of Behavior*.

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of neurology, neurosurgery, and psychiatry*, *20*(1), 11.

Selen, L. P., Shadlen, M. N., & Wolpert, D. M. (2012). Deliberation in the motor system: Reflex gains track evolving evidence leading to a decision. *The Journal of Neuroscience*, *32*(7), 2276–2286.

Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: seeing and deciding. *Proceedings of the national academy of sciences*, *93*(2), 628–633.

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *Journal of neurophysiology*, *86*(4), 1916–1936.

Shankar, K. H., & Howard, M. W. (2012). A scale-invariant internal representation of time. *Neural Computation*, *24*(1), 134–193.

Shapiro, S. (2013). Classical logic. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Winter 2013 ed.). `http://plato.stanford.edu/archives/win2013/entries/logic-classical/`.

Shinomoto, S., Omi, T., Mita, A., Mushiake, H., Shima, K., Matsuzaka, Y., & Tanji, J. (2011). Deciphering elapsed time and predicting action timing from neuronal population signals. *Frontiers in computational neuroscience*, *5*.

Simen, P., Balcı, F., de Souza, L., Cohen, J. D., & Holmes, P. (2011). A model of interval timing by neural integration. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *31*(25), 9238–53. Retrieved from `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3142662&tool=pmcentrez&rendertype=abstract` doi: 10.1523/JNEUROSCI.3121-10.2011

Skinner, B. F. (1938). *The behavior of organisms: an experimental analysis.* Retrieved from `http://www.bfskinner.org/product/the-behavior-of-organisms-e-book/`.

Skinner, B. F. (1948). 'Superstition' in the pigeon. *Journal of Experimental Psychology*, *38*(2), 168–172. Retrieved from `http://content.apa.org/journals/xge/38/2/168` doi: 10.1037/h0055873

Skinner, B. F. (1965). *Science and human behavior.* Simon and Schuster.

Skinner, B. F. (1974). *About behaviorism* (Kindle ed.). Retrieved from `www.amazon.com.br`.

Staddon, J., & Higa, J. (1999). Time and memory: Towards a pacemaker-free theory of interval timing. *Journal of the Experimental Analysis of Behavior*, *71*, 215–251.

Steiner, H., & Tseng, K. Y. (2010). *Handbook of basal ganglia structure and function: a decade of progress* (Vol. 24). Academic Press.

Stubbs, A. (1968). The discrimination of stimulus duration by pigeons. *Journal of the Experimental Analysis of Behavior*, *11*(3), 223.

Stubbs, D. A. (1976). Scaling of stimulus duration by pigeons. *Journal of the Experimental Analysis of Behavior*, *26*(1), 15.

Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *science*, *304*(5678), 1782–1787.

Suri, R. E., & Schultz, W. (1999). A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. *Neuroscience*, *91*(3), 871–890.

Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems 27* (pp. 3104–3112). Curran Associates, Inc. Retrieved from `http://papers.nips.cc/paper/5346-sequence-to-sequence-learning-with-neural-networks.pdf`

Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine learning*, *3*(1), 9–44.

Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of pavlovian reinforcement. In M. Gabriel & M. J (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497–537). MIT Press.

Sutton, R. S., & Barto, A. G. (2015). *Reinforcement learning: An introduction* (2nd (in progress, retrieved from `https://www.dropbox.com/s/b3psxv2r0ccmf80/book2015oct.pdf`) ed.). MIT press Cambridge.

Teichmann, J., Broom, M., & Alonso, E. (2014). The application of temporal difference learning in optimal diet models. *Journal of theoretical biology*, *340*, 11–16.

Tenenbaum, J. B., De Silva, V., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *science*, *290*(5500), 2319–2323.

Terrace, H., Gibbon, J., Farrell, L., & Baldock, M. (1975). Temporal factors influencing the acquisition and maintenance of an autoshaped keypeck. *Animal Learning & Behavior*, *3*(1), 53–62.

Thomason, R. (2014). Logic and artificial intelligence. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Winter 2014 ed.). `http://plato.stanford.edu/archives/win2014/entries/logic-ai/`.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological review*, *55*(4), 189.

Tolman, E. C., Ritchie, B. F., & Kalish, D. (1946). Studies in spatial learning. i. orientation and the short-cut. *Journal of experimental psychology*, *36*(1), 13.

Tran, N., & Fiete, I. (2015). An efficient grid cell decoder: are super-polynomial codes neurally plausible? In *Cosyne.* Salt Lake City, UT.

Treisman, M. (1963). Temporal discrimination and the indifference interval: Implications for a model of the" internal clock". *Psychological Monographs: General and Applied*, *77*(13), 1.

Uchida, N., & Mainen, Z. F. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nature neuroscience*, *6*(11), 1224–1229.

Van Gelder, T. (1995). What might cognition be, if not computation? *The Journal of Philosophy*, 345–381.

Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and brain sciences*, *21*(05), 615–628.

Vogels, T. P., Rajan, K., & Abbott, L. (2005). Neural network dynamics. *Annu. Rev. Neurosci.*, *28*, 357–376.

Wang, X.-J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends in neurosciences*, *24*(8), 455–463.

Wencil, E. B., Coslett, H. B., Aguirre, G. K., & Chatterjee, A. (2010). Carving the clock at its component joints: neural bases for interval timing. *Journal of neurophysiology*, *104*(1), 160–168.

Wilkinson, G., & Rogers, C. (1973). Symbolic description of factorial models for analysis of variance. *Applied Statistics*, 392–399.

Wilson, R. A., & Foglia, L. (2016). Embodied cognition. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Spring 2016 ed.). `http://plato.stanford.edu/archives/spr2016/entries/embodied-cognition/`.

Wing, A. M., & Kristofferson, A. (1973). The timing of interresponse intervals. *Perception & Psychophysics*, *13*(3), 455–460.

Wittmann, M. (2013). The inner sense of time: how the brain creates a representation of duration. *Nature Reviews Neuroscience*, *14*(3), 217–223.

World Trade Organization. (1995). *Agreement on trade-related aspects of intellectual property rights.* Retrieved from `https://www.wto.org/english/docs_e/legal_e/27-trips.pdf`

Xu, M., Zhang, S.-y., Dan, Y., & Poo, M.-m. (2014). Representation of interval timing by temporally scalable firing patterns in rat prefrontal cortex. *Proceedings of the National Academy of Sciences*, *111*(1), 480–485.

Znamenskiy, P., & Zador, A. M. (2013). Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. *Nature*, *497*(7450), 482–5. Retrieved from `http://www.ncbi.nlm.nih.gov/pubmed/23636333` doi: 10.1038/nature12077

ITQB-UNL | Av. da República, 2780-157 Oeiras, Portugal
Tel (+351) 214 469 100 | Fax (+351) 214 411 277

www.itqb.unl.pt