




**ITqb nova**

# The basal ganglia learn a temporal basis for the organization of behavior

Dissociable features of striatal dynamics bias the what, when and how of interval timing

**Filipe Miguel Vieira Sobreira Rodrigues**



Dissertation presented to obtain the **Ph.D degree** in  
Neuroscience

Oeiras, January, 2024



# **The basal ganglia learn a temporal basis for the organization of behavior**

## **Dissociable features of striatal dynamics bias the what, when and how of interval timing**

Copyright © Filipe Miguel Vieira Sobreira Rodrigues, Instituto de Tecnologia Química e Biológica António Xavier, NOVA University Lisbon.

The Instituto de Tecnologia Química e Biológica António Xavier and the NOVA University Lisbon have the right, perpetual and without geographical boundaries, to file and publish this dissertation through printed copies reproduced on paper or on digital form, or by any other means known or that may be invented, and to disseminate through scientific repositories and admit its copying and distribution for non-commercial, educational or research purposes, as long as credit is given to the author and editor.



*Para a minha mãe.*



## ACKNOWLEDGEMENTS

I am two days away from having to submit this, ten days from having something I would be happy with, and five days past trusting what I'm writing<sup>1</sup>. Still, I worry that if I do not come up from the trenches of the discussion (or down from the clouds of the introduction) now I will risk incurring an even larger debt to the one I already feel I have with all the people that got me here. With me having the emotional maturity of a tangerine, what follows is sure to fall short of even making a dent on that feeling of indebtedness, but it is a genuine attempt.

Starting with Joe. I am eternally and immeasurably grateful for the opportunity you gave that fresh-out-of-college little amalgam of confusion with not much else going for him besides a quirky application letter. Joining your lab, the work and people there, have all shaped my thinking and being beyond recognition. You in particular, have taught me more than I am sure I am even aware of. The "brain" will never not be be this impossibly daunting thing, about which more than half the time I struggle even formulating questions, let alone answers. But the clarity of your synthesis gives me hope that I might one day get some semblance of understanding myself.

Second (but really first on many accounts), Tiago. You are the best Gru a minion could ask for<sup>2</sup>. Inside the lab, you showed me all the ropes, and then had the rare humility to let me show you some (flimsy twines) back. If I am anything approaching a scientist today, I became one under your watch. I hope I never fail to do justice to the deep care and respect for behavior you instilled in me day in and day out. Outside of the lab, you became an instant

---

<sup>1</sup>Story of my life.

<sup>2</sup>This one was Joana's idea:)

friend for life. You are one of these borderline-annoying people who you meet for five minutes and realise they have it all figured out. It is hard to think of a person with their priorities more in order – scratch the borderline, it is annoying. But it is also inspiring. Thank you.

Bruno. You are one of those people I feel an instant connection with on the basis of how far from having stuff figured out we both are. It is relieving, relatable and just plain nice (take a note Tiago). Mind you, you are also extremely annoying in the impossible-to-live-up-to standards you set. But also so ridiculously talented and hardworking that jealousy gives way to admiration before it has the slightest chance to fester. I look forward to all the great things that await you.

Asma. Your humanity is one of the things I remember best from my interviewing in the lab. After a full day of intimidating conversations I was feeling about as out of place as one can, but after a low-stakes chat with you about your work (in a thankfully dimly lit meeting room) that all went away somehow. You were and continue to be an inspiration in your appreciation for the details that matter. Getting to learn from you has been a great privilege, and I hope it keeps going for many years.

Mauro – another super human being. Whatever sits opposite of tangerines in the spectrum of emotional wisdom, you're it. I am really thankful for all our interactions over the years. You left a big cross-legged-shaped not-lunch-having good-enough hole in the open lab, but I am glad to know that you are doing well and are excited about what you are doing again. Murphy sure had it personal with you, but you showed him. Suerte!

Marga. It makes me really happy to have been a part of your journey. Your stay in the lab was one of the best times of my PhD. And, although I did not have nearly as much of an influence as you think in helping you get to where you are today, I am super proud of whatever *little*<sup>3</sup> influence I may have had. Now, if, maybe, you could please, if it is not too much to ask, possibly, when you have time, start owning you to your independent worth as a scientist, that'd be great. There would be no cooling story without you, I hope you at least know that.

Teresa. You are more than a friend at this point. You have crossed that threshold where it even feels weird saying nice things about you (see my brothers below). So I just won't. Bah. OK. I feel really lucky to know that

---

<sup>3</sup>As little as ~6% by Tiago's estimates.

I always have you in my corner<sup>4</sup>. There. Something nice. Now back to bickering please. (I'm also always in yours).

Margarida. You are pathologically genuine, it is maybe a problem, and I wish more people had it. You are also incredibly sharp, I think, you only say every fifth word so I have to infer a lot, but from what I gather you are brilliant. Thank you for all the transparency, support, lalalas and discussions, it was always nice knowing you were just one desk away.

Renato. Our plans sure took some unexpected turns (someone should do something with those expensive air bearings at some point). But alas, *time heals all wounds* (or whatever platitude you think fits best here – smokes pipe). I am really glad that you have landed on your feet, and hope that things go as smoothly as possible for you from hereon out – you are really close! Thank you for all the science talk, it was always nice, but especially now that I am apparently into microbiology?!

Simon, you showed up just as my jadedness was starting to peak. All the incredible people that had first welcomed me into the lab had moved on, but this dinosaur wasn't (isn't) going anywhere. *Why bother even learning this new master's student's name, three months and he'll be on his way too* – I am really glad I bothered (even though you didn't "Sy-mun"). You're a cool guy, talented, generous and wise way beyond your years<sup>5</sup>. Also, the chocolate thing. Merci.

To all other Paton lab members, past and present, my INDP classmates, thesis committee members, the hardware platform, and the countless other amazing people at the Champalimaud that helped me one way or the other, I am so sorry that I don't have a paragraph in me for each and every one of you. It really says more about me (and my current state) than it does about you. I hope that that will have changed in the meantime if I ever accomplish anything worth an acknowledgements section again. For now just know that I appreciate and admire you all.

Joana. I saved you for last because I just know I am going to fall short. I am sorry in advance. Thanking you just feels weird. Should I thank myself next? *We did this*. Your name should be on the cover. Actually it is! (I keep forgetting that's my name now). Perfect:) Thank you for being my best friend, partner, and other half. I am sorry that you have become such a big part of who I am that you sometimes get the appreciation I usually reserve

---

<sup>4</sup>Case in point, it's 00:34 AM – you're giving me feedback on my discussion:)

<sup>5</sup>You have no sense of style though.

for myself. Forget about this thesis, thank you for every single great thing in my life.

Sofs. I don't know why you would be reading this in 2029 (or whenever it is that kids learn to read, English, don't feel bad if it's later I just made up a number). Did mom put you up to this? Anyway, all these other people (stop reading now if you're one of them) are great and all, I love them. But. You are on a league of your own. You hacked me the moment you were born, introduced to this new kind of love, one that fills me with an unwavering sense of purpose and dread. It's weird, scary, and the best feeling in the world. Thank you.

Pedro and Tiago. I considered creating an "anti-acknowledgements" section just for the two of you. You did not help me at all during any of it. Are you even going to read this? You guys suck. And smell. And I love you<sup>6</sup> (I don't really, but mom will probably read this and you know this mooshy stuff makes her happy).

João and Vasco, the brothers I wish I had. And have really. Thank you for the constant friendship, guidance and laughter. Whether you are aware of it or not you each played a big part in making this happen. You are not really role models because by now I've learned that the ways in which you two are special are not even accessible to me, but that's ok, I'm just thankful that I get to share in your virtue. Never change.

Guida, Teresa (Ticha? sogra? temos de resolver isto..) e tia Mariana. Não sei bem como expressar a dívida e gratidão que sinto. Obrigado por serem a nossa rede de apoio, sem a vossa ajuda a Sofia não teria nada para ler em 2029.

Avó Beatriz e avô Horácio. Obrigado por todo o apoio de que nem memória tenho. Ao avô em particular, sinto que tenho qualquer coisa por agradecer que não consigo articular.

Pai, obrigado por me incutires o teu sentido de brio.

Mãe, não há palavras. Não vale grande coisa, mas esta tese é para ti.

---

<sup>6</sup>Prove that it was me who wrote this. You can't.

”  
*“Time is the primordial context. For mobile organisms, evolution has selected for neurobehavioral mechanisms that anticipate predictable events rather than those that react to them after the fact.”*

— **John Gibbon**, *Toward a neurobiology of temporal cognition: advances and challenges* (1997)  
(Psychophysicist)



## RESUMO

Os gânglios da base são um conjunto evolutivamente antigo de estruturas subcorticais que se pensa terem dotado os primeiros vertebrados com a capacidade de explorar contingências temporalmente alargadas a serviço da tomada de decisões e do controlo motor. Estas funções dependem criticamente do tempo e de informação sensoriomotora, que podem ambas ser extraídas do estado em evolução de populações de neurónios na principal estrutura de entrada dos gânglios da base, o corpo estriado. No entanto, debate-se se a atividade estriatal subjaz a processos de decisão latentes e dinâmicos ou à cinemática do movimento observável. Aqui, caracterizamos e comparamos os impactos da temperatura e da intensidade do estímulo em múltiplos aspectos das respostas da população estriatal e no comportamento de ratos coletados em várias tarefas de temporização. O arrefecimento e as intensidades mais baixas causaram dilatação; o aquecimento e as intensidades mais altas, contração, tanto da atividade neural quanto dos padrões de julgamento no tempo. No entanto, a temperatura não afectou de forma semelhante a cinemática do movimento, produzindo, em vez disso, um efeito não monotónico no vigor do movimento, que foi acompanhado por um efeito não monotónico nos níveis basais de atividade estriatal. Ambos os efeitos da temperatura replicaram a variabilidade endógena na dinâmica estriatal durante o comportamento. Além disso, a intensidade do estímulo afectou adicionalmente a amplitude das actividades estriatais. Por último, os perfis de resposta no estriado foram melhor descritos como um contínuo, argumentando contra hipóteses onde classes funcionais de neurónios de rampa sustentam selectivamente o processamento temporal. Esses dados

forneem evidências convincentes de que o curso temporal da atividade neural em populações heterogéneas de neurónios estriatais dita a velocidade de um processo latente que é usado para orientar escolhas, mas não o controlo motor contínuo. Mais amplamente, representam evidências adicionais importantes de que a dinâmica populacional estriatal pode fornecer uma base temporal que sustenta a organização do comportamento ao mapear informações temporais para ação.

**Palavras-chave:** tomada de decisões, controlo motor, temporização, corpo estriado, gânglios da base, dinâmicas de população, electrofisiologia, temperatura

## ABSTRACT

The basal ganglia are an evolutionarily ancient set of subcortical structures that are collectively thought to have endowed early vertebrates with the ability to exploit temporally extended contingencies in service of decision-making and motor control. These functions are critically dependent on timing and sensorimotor information, which can both be extracted from the evolving state of neural populations in the major input structure of the basal ganglia, the striatum. Yet, it is debated whether striatal activity underlies latent, dynamic decision processes or the kinematics of overt movement. Here, we characterized and compared the impacts of temperature and stimulus intensity on multiple aspects of striatal population responses and the behavior of rats, collected across multiple interval timing tasks. Cooling and lower intensities caused dilation; warming and higher intensities, contraction, of both neural activity and patterns of judgment in time. However, striatal temperature changes did not similarly affect movement kinematics, instead producing a non-monotonic effect on overall vigor that was accompanied by a non-monotonic effect on baseline levels of activity. Both effects of temperature mimicked endogenous variability in striatal dynamics during behavior. Furthermore, stimulus intensity additionally affected the amplitude of striatal activities. And lastly, response profiles in the striatum were best described as a continuum, arguing against hypotheses where functional clusters of ramping neurons selectively underpin temporal processing. These data provide compelling evidence that the time course of evolving neural activity in heterogeneous populations of striatal neurons dictates the speed of a latent process that is used to guide choices, but not continuous

motor control. More broadly, they suggest that striatal population dynamics may provide basal ganglia circuits with a flexible time basis to support the organization of behavior by allowing for an arbitrary mapping of temporal information onto action.

**Keywords:** decision-making, motor control, interval timing, striatum, basal ganglia, population dynamics, electrophysiology, temperature

# CONTENTS

|                                                                                                                                           |             |
|-------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| <b>List of Figures</b>                                                                                                                    | <b>xix</b>  |
| <b>Acronyms</b>                                                                                                                           | <b>xlvi</b> |
| <b>1 General introduction</b>                                                                                                             | <b>1</b>    |
| 1.1 The organization of behavior . . . . .                                                                                                | 1           |
| 1.1.1 Time and hierarchy . . . . .                                                                                                        | 2           |
| 1.1.2 Heterarchy . . . . .                                                                                                                | 7           |
| 1.1.3 Flexibility . . . . .                                                                                                               | 10          |
| 1.1.4 Topography . . . . .                                                                                                                | 15          |
| 1.1.5 Prediction . . . . .                                                                                                                | 19          |
| 1.2 Inlook . . . . .                                                                                                                      | 23          |
| <b>2 Manipulating striatal temperature during interval categorization</b>                                                                 | <b>25</b>   |
| 2.1 Introduction . . . . .                                                                                                                | 25          |
| 2.2 Results . . . . .                                                                                                                     | 28          |
| 2.2.1 Striatal temperature modified neural population speed<br>monotonically and baseline activity levels nonmono-<br>tonically . . . . . | 28          |
| 2.2.2 Striatal population speed predicts temporal judgments<br>across tasks with different immobility requirements                        | 38          |
| 2.2.3 Striatal temperature produced monotonic effects on<br>temporal judgments . . . . .                                                  | 45          |
| 2.2.4 Striatal temperature produced nonmonotonic effects<br>on movement kinematics . . . . .                                              | 48          |

|          |                                                                                                                                                                                                                      |           |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| 2.3      | Discussion . . . . .                                                                                                                                                                                                 | 55        |
| 2.4      | Methods . . . . .                                                                                                                                                                                                    | 61        |
| 2.4.1    | Subjects . . . . .                                                                                                                                                                                                   | 61        |
| 2.4.2    | Behavioral apparatus . . . . .                                                                                                                                                                                       | 62        |
| 2.4.3    | Behavioral training . . . . .                                                                                                                                                                                        | 62        |
| 2.4.4    | Thermoelectric device . . . . .                                                                                                                                                                                      | 63        |
| 2.4.5    | Surgical procedures . . . . .                                                                                                                                                                                        | 66        |
| 2.4.6    | Temperature manipulation protocol . . . . .                                                                                                                                                                          | 69        |
| 2.4.7    | Implant placement confirmation . . . . .                                                                                                                                                                             | 70        |
| 2.4.8    | Data analysis . . . . .                                                                                                                                                                                              | 70        |
| 2.4.9    | Psychophysical data analysis . . . . .                                                                                                                                                                               | 70        |
| 2.4.10   | Electrophysiological data analysis . . . . .                                                                                                                                                                         | 72        |
| 2.4.11   | Continuous behavioral data analysis . . . . .                                                                                                                                                                        | 76        |
| 2.4.12   | Statistics and reproducibility . . . . .                                                                                                                                                                             | 77        |
| <b>3</b> | <b>Manipulating stimulus intensity during delayed interval discrimination</b>                                                                                                                                        | <b>79</b> |
| 3.1      | Introduction . . . . .                                                                                                                                                                                               | 80        |
| 3.2      | Results . . . . .                                                                                                                                                                                                    | 81        |
| 3.2.1    | Rats did not discriminate vibration durations as idealized and reported . . . . .                                                                                                                                    | 81        |
| 3.2.2    | The second stimulus dominated the delayed duration comparison . . . . .                                                                                                                                              | 84        |
| 3.2.3    | Multiple features of DLS dynamics were systematically affected by the intensity of the second stimulus . . . . .                                                                                                     | 89        |
| 3.2.4    | DLS neural responses displayed negligible sensitivity to the intensity of the first stimulus . . . . .                                                                                                               | 94        |
| 3.2.5    | The speed of DLS dynamics covaried with duration judgments . . . . .                                                                                                                                                 | 96        |
| 3.2.6    | The <i>relative</i> proportion of DLS neurons modulated by different task variables matched their relative impact on performance . . . . .                                                                           | 101       |
| 3.2.7    | Differences in time decodability from putative clusters of ramping and non-ramping DLS neurons were biased by confounds unrelated to those category labels that were introduced by categorization criteria . . . . . | 107       |

|          |                                                                                                                                  |            |
|----------|----------------------------------------------------------------------------------------------------------------------------------|------------|
| 3.2.8    | No evidence for the existence of ramping and non-ramping clusters in DLS responses recorded during the second stimulus . . . . . | 113        |
| 3.3      | Discussion . . . . .                                                                                                             | 116        |
| 3.4      | Methods . . . . .                                                                                                                | 117        |
| 3.4.1    | Behavioral analyzes . . . . .                                                                                                    | 117        |
| 3.4.2    | Single neuron analyzes . . . . .                                                                                                 | 118        |
| 3.4.3    | Population analyzes . . . . .                                                                                                    | 120        |
| 3.4.4    | Simulations . . . . .                                                                                                            | 122        |
| 3.4.5    | Statistics . . . . .                                                                                                             | 124        |
| <b>4</b> | <b>General discussion</b>                                                                                                        | <b>127</b> |
| 4.1      | Scaling mechanisms . . . . .                                                                                                     | 129        |
| 4.2      | Discrete vs. continuous control . . . . .                                                                                        | 137        |
| 4.3      | Hysteresis . . . . .                                                                                                             | 143        |
| 4.4      | Outlook . . . . .                                                                                                                | 147        |
|          | <b>Bibliography</b>                                                                                                              | <b>153</b> |



## LIST OF FIGURES

- 2.1 Schematic of our implantable thermoelectric device (TED) targeting the dorsal striatum (DS). The thick horizontal line splits the diagram into two differently scaled subregions (see 5 mm scale bars). Top region is subdivided in the middle between a front view (left half, with white strokes and gray fills) and a cutaway through the center of the implant (right half). A coronal slice of a rat's brain is shown in gray for reference (+0.84 mm AP, adapted from Paxinos and Watson, 2013). . . . . 29
- 2.2 Temperature manipulation protocol. Top: Target setpoint of the TED device over a representative session. Bottom: Representative segment of thermistor readings (black trace) at the lower plate of the TED module shown in Fig. 2.1. Gradient-colored trace indicates the temperature estimated from in vivo calibration experiments (see Fig. 2.3). . . . . 30

|     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |    |
|-----|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 2.3 | Spatiotemporal characterization of our TED. (A) Schematic of the preparation in which we set our TED to one of several manipulation temperatures ( $T = \{5, 15, 20, 25, 30, 45\}^{\circ}\text{C}$ ) while measuring temperature at its lower plate and probe tip simultaneously. (B) Temperature measured at the TED plate and probe tip thermistors. (C) Temperature traces measured at the probe tip thermistor during manipulation blocks aligned to block transitions. Solid lines represent model fits. (D) Schematic of the preparation in which we set our TED to one of several manipulation temperatures ( $T = \{15, 25, 42, 45\}^{\circ}\text{C}$ ) while measuring temperature at its lower plate a movable temperature probe simultaneously. (E) Decay parameters for models fit to manipulation temperatures (as shown in (C)) across the 2 experiments (manipulation temperatures that are common to both experiments are connected with solid colored lines). (F) Same as (E), but for the gain / offset parameter across all model fits. . . . . | 31 |
| 2.4 | (A) Schematic of the preparation used to elicit, record and manipulate striatal activity dynamics under different temperatures. (B) ventrobasal complex (VB) thalamic stimulation protocol overlaid with five single-trial examples of evoked striatal activity recorded under the control temperature condition. Blue ticks depict light pulses. (C) Normalized spike density functions (SDFs) of recorded striatal neurons ( $N = 335$ , before enforcing the minimum firing rate selection criterion, see Methods), under the control temperature. Units were ordered by their angular position in the subspace defined by the first two principal components (PCs), describing dynamics between 150 ms and 750 ms from stimulation onset. . . . .                                                                                                                                                                                                                                                                                                              | 32 |
| 2.5 | Activity of three putative striatal units (arrowheads in Fig. 2.4C.) aligned to the onset of VB stimulation. Top: SDFs split by temperature (mean $\pm$ standard error of the mean (SEM)). Bottom: Spike raster plots, wherein trials are ordered chronologically. The superimposed gradient-colored trace indicates temperature setpoint over the course of the recording session. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 33 |

2.6 Temporal scaling as opposed to shifting provided a better account of temperature’s effect on neural activity. (A) Top: Simulated spike density functions exhibiting bidirectional and dose-dependent temporal scaling with temperature. Middle: Templates built by warping a control spike density function (thicker black line) in time by scale factors ranging from 0.625 (maximum contraction) to 1.75 (maximum dilation). Note that when applying this method to data, this control response is not the same as the one shown in the top panel, as the two are built using two non-overlapping random sets of control trials. Bottom: Thick lines represent the coefficient of determination ( $R^2$ ) for all scaled templates in the middle panel regressed against each of the target spike density functions shown at the top. We computed this objective function for each neuron-temperature condition pair and took its global maximum as the corresponding temporal scaling factor, highlighted here by the larger markers. Thinner lines and smaller markers depict  $R^2$  values for a similar regression procedure applied to a series of shifted, as opposed to scaled, templates. (B) Same as (A), except that for artificially temporally shifted responses relative to control (top), temporally shifted templates (middle), and their regression outcomes (bottom). The thinner lines and smaller markers respectively represent the  $R^2$  curves and maxima resulting from regressing the scaled templates from A (middle) against the shifted targets in B (top). Conversely, the result of regressing shifted templates against scaled targets is plotted in the same manner in (A, bottom). . . . . 34

2.7 (A) To assess whether the effects of temperature on individual striatal responses were better accounted for by temporal scaling or shifting, we built two separate spiking models in which we either injected one effect or the other. Briefly, we modeled 500 control firing rate functions as Gaussian bumps defined over 1.5 s with means spanning the interval from 150 ms to 750 ms (Fig. 1e) and a standard deviation of 50 ms. The amplitudes of the resulting probability density functions were rescaled so that their distribution of mean firing rates matched that of striatal data. Next, we created one additional rate function per neuron per manipulation condition by either shifting or scaling its control response in time. Again, the distribution of generative temporal scaling factors and shifts used was informed by the empirical distributions of these metrics extracted from striatal data. We then generated 150 spike trains of each condition per neuron by sampling spike times from inhomogeneous Poisson point processes with the aforementioned condition-specific responses as their time-dependent rate parameters. From this point on, we proceeded to analyze the resulting surrogate spike data in the exact same way we did for the striatal data, by first averaging trials within condition, generating libraries of templates and then computing temporal scaling factors and shifts. Finally, for each 'neuron'-condition pair within each model, we stored the  $R^2$  values corresponding to the best-matching scaled and shifted templates and subtracted the former from the latter to build the distributions shown here at the top (scaling model) and middle (shifting model) panels. Thick solid sigmoidal lines represent the CDFs of each condition's  $R^2$  difference. Thin vertical black lines denote control mean differences. Small horizontal colored lines link the respective means of the corresponding manipulation and control distributions. (B) Same as (A), but for striatal data. . . 35

- 2.8 Temperature rescaled striatal population activity in time leading to graded monotonic modulation of decoded time estimates. **(A)** Distribution of percentage change in temporal scale relative to control dilation ( $N = 268$ , after enforcing the minimum firing rate selection criterion, see Methods). Histograms (top) and corresponding cumulative density functions (CDF, bottom) of dilation conditioned on manipulation temperature (same color scheme as in Fig. 1). **(B)** Left: Distributions of neuronal response dilations as a function of striatal temperature change ( $N = 175$ , after excluding unreliable scaling factor estimates, see Methods, one-sample two-tailed t-tests, all significant  $|t_s(174)| = [4.70, 9.37]$ ,  $P = [3.84e-17, 5.28e-6]$ , non-significant  $t(174) = 1.16$ ,  $P = 0.25$ ; repeated measures ANOVA followed by post-hoc contrasts with Tukey correction for multiple comparisons,  $F(3, 696) = 51.29$ ,  $P = 5.88e-30$ ; smallest  $q_s(348) = 2.96$ ,  $P = 0.016$ ). Markers represent temperature-split response dilations for each cell. Boxplots show population medians (horizontal thick lines) and interquartile range (IQR) (colored bars). Right: Distribution of stretch in neuronal responses (one-sample two-tailed t-test,  $t(174) = -10.15$ ,  $P = 2.70e-19$ ). **(C)** Average control population activity projected onto its first 3 PCs (solid black line), and temperature-split averages (colored markers) projected onto that control trajectory at the time point indicated by the last arrowhead in (E) and (F). Ghosted versions of these visualizations show their 2D projections onto all possible combinations of PCs 1, 2 and 3. The dashed black arrow indicates the direction of time. . . . . 36
- 2.9 Distributions of changes in baseline firing rate as a function of striatal temperature change (one-sample two-tailed t-test all significant  $|t_s(267)| = [5.24, 15.01]$ ,  $P = [3.38e-37, 3.31e-7]$ , non significant  $t(267) = 0.18$ ,  $P = 0.86$ ; repeated measures ANOVA followed by post-hoc contrasts with Tukey correction for multiple comparisons,  $F(3, 1068) = 72.29$ ,  $P = 1.44e-42$ ; smallest  $t_s(534) = 4.02$ ,  $P = 3.40e-4$ ). Markers represent temperature-split change in baseline firing rate for each cell. Boxplots show population medians (horizontal thick lines) and IQR (colored bars). . . . . 37

|      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |    |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 2.10 | Decoded posterior probability of time given the state of striatal populations at four arbitrary time points, averaged separately across trials for each temperature condition (mean $\pm$ SEM). . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 37 |
| 2.11 | Difference between trial-averaged decoded estimates for manipulation (left: extreme cooling; middle: mild cooling; right: warming) and control conditions, shown for the time window highlighted in Fig. 1E. Grayscale arrowheads indicate time points used in 2.10. The identity line is shown in dashed black. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 38 |
| 2.12 | Schematic of behavioral setup (left) and event diagram illustrating a correct trial (right) in both the no-fixation (orange) and fixation (blue-green) versions of the interval categorization task. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 39 |
| 2.13 | (A) Psychometric performance in the no-fixation version of the task (N = 3 rats, mean $\pm$ SEM). (B) Same as (A), for the fixation task variant (N = 3 rats). . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 40 |
| 2.14 | (A) Normalized SDFs of all striatal neurons (N = 441) recorded during presentations of the longest stimulus in our set (2.4 s) in the no-fixation task variant. Units are ordered by their angular position in the subspace defined by the first two PCs describing firing dynamics. (B) Same as (A), for the fixation task variant (N = 483 neurons). . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 40 |
| 2.15 | Left: Schematic of hypothetical relationship between trial to trial variability in the speed of striatal dynamics, blue indicating slower, and red indicating faster, and the resulting distributions of internal time estimates at each interval offset. Assuming a fixed internal decision boundary, long or short judgments are colored as black or gray portions of the distributions. Right: Schematic depicting hypothetical, graded shifts in average internal timing speeds when conditioning on choice for the two intervals nearest to the decision boundary. Specifically, short (long) stimuli incorrectly categorized as long (short) should reflect pronounced biases towards faster (slower) than usual dynamics, and correctly categorized short (long) stimuli should reflect milder biases towards relatively slower (faster) dynamics. . . . . | 42 |

|      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |    |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 2.16 | (A) Left: Distributions of neuronal response dilations for each near-boundary stimulus-choice pair introduced in Fig. 2.15. Markers represent response dilations for each cell and stimulus-choice condition. Boxplots show population medians (horizontal thick lines) and IQR (colored bars). Right: Distribution of stretch in neuronal responses (one-sample two-tailed t-test, $t(440) = -7.21$ , $P = 2.45e-12$ ). (B) Same as (A), for the fixation task variant (one-sample two-tailed t-test, $t(482) = -5.28$ , $P = 1.93e-7$ ). . . . .                                                                                                                                                                                                                                                                               | 42 |
| 2.17 | (A) Average population activity from 2.14 projected onto its first three PCs (solid black line), and averages (colored markers) from each stimulus-choice pair projected onto that reference trajectory at the time of stimulus offset for the short near-boundary stimulus (1.38 s). Ghosted versions of these visualizations show their 2D projections onto all possible combinations of PCs 1, 2 and 3. The dashed black arrow indicates the direction of time. (B) Same as (A), for the fixation task variant. . . . .                                                                                                                                                                                                                                                                                                       | 43 |
| 2.18 | (A) Average decoded <i>maximum a posteriori</i> (MAP) estimates for each stimulus-choice condition defined in 2.15 for the no-fixation task version. Colored solid lines and patches represent medians and IQR across 512 concatenations (see Methods). (B) Same as (A) for the fixation task variant. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 43 |
| 2.19 | (A) Performance in the no fixation version of the interval categorization task conditioned on striatal population speed ( $N = 3$ rats, see Methods): fast (red), typical (black), slow (dark blue) or slowest (light blue). Psychometric curves fit to cross-animal averages of psychophysical data split by neural speed, respectively shown as solid lines and markers of matching color (mean $\pm$ SEM). Bottom-right inset: Differences in proportion of long choices from each atypical speed condition to the typical speed condition (mean $\pm$ propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each speed condition's psychometric fit. Solid black lines represent MAP point estimates. (B) Same as (A), but for the fixation task variant ( $N = 3$ rats). . . . . | 44 |

2.20 Performance in both task variants was qualitatively similar before and after TED implantation. **(A)** Performance of rats trained on the no-fixation variant of the interval categorization task and implanted with the initial version of our custom TED targeting the DS (N = 4) on the last day of training before surgery. Gray lines are psychometric fits to individual animals. The black line is a fit to the average across animals. The underlying choice data is shown for the cohort average (mean  $\pm$  SEM). **(B)** Same as (A), but for rats trained on the fixation task variant and implanted with the initial (N = 1) or final version of our custom TED targeting the DS (N = 6). **(C)** Schematized timeline for a typical experimental rat. Elongated bars illustrate long periods of time, whereas squares represent individual daily sessions, which, with the exception of surgeries, lasted for 2 hours. **(D-E)** Same as **(A-B)**, but for the last post-surgery training session before the first temperature manipulation session. . . . . 46

2.21 Performance of rats implanted with the final version of our custom TED targeting the DS in the fixation version of the interval categorization task conditioned on striatal temperature (N = 6 rats): warm (red), control (black), cool (teal) or coolest (cyan). Psychometric curves fit to cross-animal averages of psychophysical data split by temperature, respectively shown as solid lines and markers of matching color (mean  $\pm$  SEM). Bottom-right inset: Differences in proportion of long choices from each experimental condition to the control condition (mean  $\pm$  propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each temperature condition’s psychometric fit. Solid black lines represent MAP point estimates. Within-animal data were pooled across the first two sessions of temperature manipulations, with the first one contributing the most (see Methods). 47

2.22 **(A)** Performance of rats implanted with the initial version of our custom TED targeting the DS in the no fixation version of the interval categorization task conditioned on striatal temperature (N = 4 rats): control (black), cool (teal). Psychometric curves fit to cross-animal averages of psychophysical data split by temperature, respectively shown as solid lines and markers of matching color (mean  $\pm$  SEM). Bottom-right inset: Differences in proportion of long choices from the cool to the control condition (mean  $\pm$  propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each temperature condition's psychometric fit. Solid black lines represent MAP point estimates. **(B)** Same as (A), but for the one rat trained in the fixation task variant and implanted with the initial version of our TED. . . 48

2.23 Effect of temperature on psychophysical thresholds. **(A)** Markers represent MAP estimates and transparent patches the corresponding 95% confidence intervals of threshold parameters fit to individual animals' performance on control (vertical axis) versus manipulation blocks (horizontal axis). Each animal contributes one data point of each color. The identity line is plotted as a diagonal line, with its thick solid portion highlighting the region of the main axes that is shown in the bottom-right inset axes. Top-right inset: Distribution of threshold differences between manipulation and control conditions (mean  $\pm$  SEM). Bottom-right inset: Same as main axes, but with data pooled from two pilot experiments using a single cooling dose in either the fixation (N = 1) or the no-fixation (N = 4) task variants. **(B)** Distributions of threshold dilation as a function of induced temperature changes (one-sample two-tailed t-tests, all significant  $|t_s(5)| = [3.25, 5.17]$ ,  $P = [4.01e-3, 2.31e-2]$ ; repeated measures ANOVA followed by post-hoc contrasts with Tukey correction for multiple comparisons,  $F(2, 15) = 17.30$ ,  $P = 1.22e-4$ ; smallest  $q_s(10) = 4.45$ ,  $P = 1.52e-2$ ). Markers represent individual animals. Boxplots show animal means (horizontal thick lines) and SEM (colored bars). 49

2.24 Manipulating primary motor cortex (M1) temperature did not produce discernible effects on timing judgments. **(A)** Average categorization performance in the fixation task version at the onset of M1 temperature manipulations. Psychometric functions fit to cross-animal averages ( $N = 4$ ) of temperature-split psychophysical data, respectively shown as solid lines and markers of matching color (mean  $\pm$  SEM). Bottom right inset: Average differences in proportion of long choices from each manipulation condition to control (mean  $\pm$  propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each condition's psychometric fit. Solid black lines represent the MAP point estimates implicit in the fits shown in the main axes. **(B)** Animal-split categorization behavior. Markers represent MAP estimates and transparent patches the corresponding 95% confidence intervals of threshold parameters fit to individual animals' performance on control (vertical axis) versus manipulation blocks (horizontal axis). Inset: Distribution of threshold differences between manipulation and control conditions. Markers represent individual animal differences (mean  $\pm$  SEM). **(C)** Left: Distributions of threshold dilation as a function of induced temperature changes in M1 (all non-significant  $|t_s(3)| = [0.93, 2.90]$ ,  $P = [0.06, 0.42]$ ; repeated measures ANOVA,  $F(2, 9) = 1.93$ ,  $P = 0.20$ ). Markers represent individual animals. Boxplots show animal means (horizontal thick lines) and SEM (colored bars). Right: Distribution of threshold stretch split by manipulation target (DS:  $N = 6$ ; M1:  $N = 4$ , see Methods). Markers represent individual animals, and their size and color denote bootstrapped significance (see Methods, one-sample two-tailed  $t$ -tests,  $t_{DS}(5) = -4.37$ ,  $P = 7.19\text{e-}3$ ;  $t_{M1}(3) = -1.79$ ,  $P = 0.17$ ; two-sample two-tailed  $t$ -test,  $t(8) = -2.62$ ,  $P = 3.14\text{e-}2$ ). Boxplots show animal means (horizontal thick lines) and SEM (colored bars).

|      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |    |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 2.25 | Threshold dynamics aligned to and across block transitions early during DS and M1 temperature manipulations (left and right, respectively). Condition-split cross-animal average thresholds (mean $\pm$ SEM) were computed using trials that fell into a sliding window lasting 90 s (half the block duration) that was swept from the preceding to the succeeding control blocks in increments of 9 s. Each marker corresponds to one sweep, and its color shading denotes the fraction of that sweep's window that was inside a control block (with black markers corresponding to 100% control trials), and by extension its complement that fell in a manipulation block (with pure manipulation colors corresponding to 100% manipulation trials). . . . . | 51 |
| 2.26 | (A) Illustration of three trial epochs: stimulus offset (left), reaction (middle) and choice (right). (B) 2D trajectories of implant position aligned to reaction from a representative example DS-implanted animal. Dashed (solid) lines correspond to short (long) choices with individual control trials ghosted in the background, and condition-split averages on top (mean $\pm$ SEM). Filled and open markers show average position at reaction and choice, respectively. . . . .                                                                                                                                                                                                                                                                        | 51 |
| 2.27 | (A) 2D trajectories of implant position during stimulus presentation from a representative example DS-implanted animal. Dashed (solid) lines correspond to short (long) choices with individual control trials ghosted in the background, and condition-split averages on top (mean $\pm$ SEM). (B) Same as (A), but split into x (top) and y (bottom) coordinates and plotted all the way up to reaction and choice. Filled and open markers show average position at reaction and choice, respectively. (C) Same data as in (A) and the stimulus presentation portion of (B), combined into an overall speed metric (see Methods). . . . .                                                                                                                    | 52 |

2.28 **(A)** 2D trajectories of implant position aligned to reaction from a representative example DS-implanted animal. Dashed (solid) lines correspond to short (long) choices with individual control trials ghosted in the background, and condition-split averages on top (mean  $\pm$  SEM). Filled and open markers show average position at reaction and choice, respectively. **(B)** Same as (A), split into x (top) and y (bottom) coordinates. **(C)** Same as (B), combined into an overall speed metric. . . . . 52

2.29 **(A)** Effect of temperature on the temporal scale of speed profiles recorded during the stimulus period. Left: Implant speed profile dilation for DS animals (one-sample two-tailed t-tests, non-significant  $t_s(5) = [0.44, 2.02]$ ,  $P = [0.10, 0.68]$ ; repeated measures ANOVA,  $F(2, 15) = 2.07$ ,  $P = 0.16$ ). Right: Implant speed stretch for DS ( $N = 6$ ) and M1 ( $N = 4$ ) animals (one-sample two-tailed t-tests,  $t_{DS}(5) = -2.02$ ,  $P = 0.10$ ,  $t_{M1}(3) = -0.83$ ,  $P = 0.47$ ; two-sample two-tailed t-test,  $t(8) = -0.65$ ,  $P = 0.53$ ). **(B)** Same as (A) for dilation and stretch of speed profiles recorded between reaction and choice. Left: Implant speed profile dilation for DS animals (one-sample two-tailed t-tests,  $t(5) = [0.11, 0.87]$ ,  $P = [0.43, 0.92]$ ; repeated measures ANOVA,  $F(2, 15) = 0.11$ ,  $P = 0.89$ ). Right: Implant speed stretch for DS ( $N = 6$ ) and M1 ( $N = 4$ ) animals (one-sample two-tailed t-tests,  $t_{DS}(5) = -0.07$ ,  $P = 0.95$ ;  $t_{M1}(3) = -0.76$ ,  $P = 0.51$ ; two-sample two-tailed t-test,  $t(8) = 0.56$ ,  $P = 0.59$ ). Markers represent individual animals, and their size and color denote bootstrapped significance. . . . . 53

2.30 Left: Same as in (E, left) for average implant speed dilation of DS animals (one-sample two-tailed t-tests, significant  $t_s(5) = [-3.64; -2.57]$ ,  $P = [0.05; 0.02]$ , non-significant  $t(5) = 0.57$ ,  $P = 0.59$ ; repeated measures ANOVA followed by post-hoc contrasts with Tukey correction for multiple comparisons,  $F(2, 15) = 4.57$ ,  $P = 0.03$ ; significant  $q_s(10) = 4.03$ ,  $P = 0.02$ ; largest non-significant  $q_s(10) = 2.21$ ,  $P = 0.16$ ). Right: Model comparison between two one-parameter models of temperature's effect on average implant speed dilation in DS animals (one-sample two-tailed t-tests,  $t_{\beta+\Delta T}(5) = -1.19$ ,  $P = 0.29$ ;  $t_{\beta+|\Delta T|}(5) = 3.21$ ,  $P = 0.02$ ; paired-sample two-tailed t-test,  $t(5) = -2.72$ ,  $P = 0.04$ ). The goodness of fit of the monotonic (left) and the non-monotonic (right) model, as measured by its coefficient of determination ( $R^2$ ), is shown for individual animals (open markers) and the population (mean  $\pm$  SEM). The  $\beta$  coefficient was constrained to only have positive (negative) values in the monotonic (non-monotonic) model. . . . . 54

2.31 (A) Left: median reaction time dilation for DS animals (one-sample two-tailed t-tests,  $t(5) = [1.31, 2.01]$ ,  $P = [0.10, 0.25]$ ; repeated measures ANOVA,  $F(2, 15) = 1.43$ ,  $P = 0.27$ ). Right: median reaction time stretch for DS ( $N = 6$ ) and M1 ( $N = 4$ ) animals (one-sample two-tailed t-tests,  $t_{DS}(5) = -1.84$ ,  $P = 0.12$ ;  $t_{M1}(3) = -0.52$ ,  $P = 0.64$ ; two-sample two-tailed t-test,  $t(8) = -0.94$ ,  $P = 0.37$ ). (B) Same as (A), but for median movement times (left: one-sample two-tailed t-tests,  $t(5) = [0.28, 1.04]$ ,  $P = [0.35, 0.79]$ ; repeated measures ANOVA,  $F(2, 15) = 0.60$ ,  $P = 0.56$ ; right: one-sample two-tailed t-tests  $t_{DS}(5) = -1.01$ ,  $P = 0.36$ ;  $t_{M1}(3) = -1.04$ ,  $P = 0.38$ ; two-sample two-tailed t-test,  $t(8) = -0.30$ ,  $P = 0.77$ ). . . . . 55

|      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |    |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 2.32 | Baseline DS firing rate correlates with reaction but not movement times in the fixation version of the interval categorization task. <b>(A)</b> Average TED implant speeds aligned to stimulus onset recorded during control blocks in the fixation version of the interval categorization task (DS cohort, N = 6 animals). Solid black lines and shaded gray patches represent individual animal medians and IQR ranges, respectively. <b>(B)</b> Distribution of correlation coefficients between baseline firing rates of individual striatal neurons (N = 483) and subsequent reaction (top; one-sample two-tailed t-test, $t(482) = -8.56$ , $P = 1.55e-16$ ) or movement times (bottom; one-sample two-tailed t-test, $t(482) = 0.15$ , $P = 0.88$ ) for the DS-recorded animals trained on the fixation variant of the interval categorization task (N = 3). . . . . | 56 |
| 3.1  | <b>(A)</b> Schematic of the behavioral apparatus with a rat depicted at the initiation/stimulus delivery port. The mapping of choice spout (left or right) to choice category (duration of the second stimulus (T2) > duration of the first stimulus (T1) or T2 < T1) was different from the depicted one for one for the dorsolateral striatum (DLS)-implanted rats. <b>(B)</b> Event diagram of an example correct trial in the delayed duration comparison task. . . . .                                                                                                                                                                                                                                                                                                                                                                                                 | 82 |
| 3.2  | <b>(A)</b> Proportion of T2 > T1 choices for all stimulus duration pairs pooled across the three DLS-implanted rats. Dotted lines above, below and at the identity line represent normalized time difference (NTD) levels of +0.35, -0.35 and 0, respectively. <b>(B)</b> Same as (A), but for stimulus intensity and normalized intensity difference (NID). The black arrow highlights the intensity pair that was undisclosed in both the original study and its erratum (intensity of the first stimulus (I1) = 64; intensity of the second stimulus (I2) = 64 mm s <sup>-1</sup> ). . . . .                                                                                                                                                                                                                                                                             | 83 |

3.3 Distribution of stimulus features in the delayed duration comparison task. **(A)** Main axes: frequency of stimulus duration pairs (T1 and T2) relative to the most common pair (T1 = 161; T2 = 334 ms). Left: Distribution of sampled T2 durations for individual DLS-implanted animals (in gray) and pooling across them (black). **(B)** Same as (A), but for stimulus intensity. The black arrow in the main axes highlights the intensity pair that was undisclosed in both the original study and its erratum (I1 = 64; I2 = 64 mm s<sup>-1</sup>). 83

3.4 Performance of simulated agents strictly driven by NTD and biased by NID. **(A)** Proportion of T2 > T1 choices for all duration pairs pooled across three simulated agents. Dotted lines above, below and at the identity line represent NTD levels of +0.35, -0.35 and 0, respectively. **(B)** Proportion of T2 > T1 choices as a function of NTD and grouped by NTD level (left) or stimulus pair (right). Gray markers and lines refer to individual agents, and black and white ones to cross-agent averages for each NTD level. Dashed lines on the right connect the two stimulus pairs with smaller absolute duration differences ( $|T2 - T1|$ ). **(C)** Proportion of T2 > T1 choices conditioned on all possible values of the numerator (main axes) and denominator (inset) terms in the definition of NTD. **(D-F)** Same as (A-C), but for stimulus intensity and NID. **(G)** Same data as in (A), but overlaying the performance rule that generated them. For each duration pair, a line was drawn between its T1 and T2 coordinates and the horizontally shifted point for which generative performance was closest to simulated performance on that pair. **(H)** Similar to (G), but for stimulus intensity. **(I)** Coefficients of a logistic regression model fit to simulated choice data using all four stimulus features (T1, T2, I1 and I2) as regressors. . . . . 85

|     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |    |
|-----|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 3.5 | (A) Proportion of $T2 > T1$ choices as a function of NTD and grouped by NTD level (left) or stimulus duration (right). Gray markers and lines refer to individual animals, and black and white ones to cross-animal averages for each NTD level. Dashed lines on the right connect the two stimulus pairs with smaller absolute duration differences ( $ T2 - T1 $ ). (B) Proportion of $T2 > T1$ choices conditioned on all possible values of the numerator (main axes) and denominator (inset) terms in the NTD definition. (C-D) Same as (A-B), but for stimulus intensity and NID. The black arrows highlight the intensity pair that was undisclosed in both the original study and its erratum ( $I1 = 64$ ; $I2 = 64 \text{ mm s}^{-1}$ ).                  | 86 |
| 3.6 | (A) Depiction of how contraction bias may have affected estimates of $T1$ in the delayed duration comparison. Same data and analysis as in Fig. 1C overlaying an hypothetical performance rule where the probability of reporting $T2 > T1$ is a monotonic (sigmoidal) function of effective NTD. For each duration pair, a line was drawn between its $T1$ and $T2$ coordinates and the horizontally shifted point for which hypothetical performance was closest to rats' performance on that pair. The vertical solid line depicts the expected value of the $T1$ distribution, and the black arrow indicates the one stimulus pair the authors singled-out as potentially affected by contraction bias. (B) Similar to (A), but for stimulus intensity. . . . . | 87 |
| 3.7 | Coefficients of a logistic regression model fit to choice data using all four stimulus features ( $T1$ , $T2$ , $I1$ and $I2$ ) as regressors. Gray markers refer to fits for individual animals, and black ones to a fit pooling data across all three DLS-implanted animals. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 88 |
| 3.8 | (A) Proportion of $T2 > T1$ choices as a function of $T2$ . Gray markers refer to individual animals, and black ones data pooled across all three DLS-implanted animals. (B) Pooled discrimination performance as a function of $T2$ and conditioned on $I2$ . Main axes: psychometric functions fit to $I2$ -split choices. Marker size represents the relative number of trials per $T2$ value normalized within each $I2$ condition. Inset: differences between choice proportions in extreme $I2$ (high: $I2 = 117 \text{ mm s}^{-1}$ ; low: $I2 = 34 \text{ mm s}^{-1}$ ) and intermediate $I2$ trials ( $I2 = 64 \text{ mm s}^{-1}$ , mean $\pm$ propagated SEM). (C) Same as (B), but conditioned on $I1$ . . . . .                                          | 89 |

3.9 Multiple parameters of the psychometric function were affected by the intensity of the second stimulus, but not that of the first. **(A)** Schematic illustrating the parameterization of the psychometric function used throughout (adapted from H. H. Schütt et al., 2016). **(B)** Model comparison between psychometric fits to choices of individual rats in middle-I2 ( $I2 = 64 \text{ mm s}^{-1}$ ) and extreme-I2 trials (high:  $I2 = 117 \text{ mm s}^{-1}$ ; low:  $I2 = 34 \text{ mm s}^{-1}$ ), measured as the difference between each model's Bayesian information criterion (BIC) score and that of the rat-specific model with the lowest BIC across low and high intensities. Extreme-I2 models are labeled according to which parameters of the psychometric function were free and which were set to those of the corresponding middle-I2 fits (e.g., "none" means that all parameters were free, and "m" means that only the threshold parameter was fixed). Small and large markers correspond to individual animals and cross-animal averages (mean  $\pm$  SEM), respectively. The dashed black line corresponds to the I2-agnostic cross-animal average, according to which models were sorted along the  $x$ -axis. **(C)** Same as (B), but for I1. **(D)** Point estimates for each of the four parameters of the psychometric function introduced in (A) estimated using only middle-intensity ( $x$ -axis) or extreme-intensity trials ( $y$ -axis). Small circles correspond to individual animals, and crosses to cross-animal averages (mean  $\pm$  SEM). . . . . 90

3.10 **(A)** Average normalized population activity conditioned on I2 and aligned to the onset of second stimulus (S2) (mean  $\pm$  SEM across neurons). **(B)** Responses of four example DLS neurons during S2. For each neuron, spike density functions (mean  $\pm$  SEM across trials) are conditioned on I2 and plotted above its spike raster, where trials are first sorted by I2, and then by T2 within each I2 level. . . . . 91

|      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |    |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 3.11 | (A) I2-split population activity projected onto each of the first three PCs computed using I2-agnostic population activity aligned to the onset of S2 (padded with 500 ms before S2 onset). (B) Population activity conditioned on I2 projected onto a 2D section of the subspace defined by the three PCs shown in (A). The 3-axis gizmo at the center depicts the projections of the axes defined by the first three PCs onto the chosen 2D section, the dashed black arrow indicates the direction of time, and the gray patches group population states recorded at each possible stimulus offset across I2. . . . .                     | 92 |
| 3.12 | (A) Neurometric curves trained using S2-aligned population activity from middle-I2 trials conditioned on I2 (main axes), and differences between the proportion of $T2 > T1$ neural "reports" in high ( $I2 = 117 \text{ mm s}^{-1}$ ) and low ( $I2 = 34 \text{ mm s}^{-1}$ ) intensity trials compared to middle intensity ( $I1 = 64 \text{ mm s}^{-1}$ ) trials (inset, mean $\pm$ propagated SEM). (B) Time estimates decoded from DLS activity around S2 presentation conditioned on I2. Main axes: mean posterior distributions across runs. Inset: maximum a posteriori (MAP) point estimates averaged across runs (median and IQR). | 93 |
| 3.13 | (A) Average normalized population activity conditioned on I1 and aligned to the onset of S2 (mean $\pm$ SEM across neurons). (B) Responses of four example DLS neurons during S2. For each neuron, spike density functions (mean $\pm$ SEM across trials) are conditioned on I1 and plotted above its spike raster, where trials are first sorted by I1, and then by T2 within each I1 level. . . .                                                                                                                                                                                                                                          | 95 |
| 3.14 | (A) I1-split population activity projected onto each of the first three PCs computed using I1-agnostic population activity aligned to the onset of S2 (padded with 500 ms before S2 onset). (B) Population activity conditioned on I1 projected onto a 2D section of the subspace defined by the three PCs shown in (A). The 3-axis gizmo at the center depicts the projections of the axes defined by the first three PCs onto the chosen 2D section, the dashed black arrow indicates the direction of time, and the gray patches group population states recorded at each possible stimulus offset across I1. . . . .                     | 95 |

3.15 (A) Neurometric curves trained using S2-aligned population activity from middle-I1 trials conditioned on I1 (main axes), and differences between the proportion of T2 > T1 neural "reports" in high ( $I1 = 117 \text{ mm s}^{-1}$ ) and low ( $I1 = 34 \text{ mm s}^{-1}$ ) intensity trials compared to middle intensity ( $I1 = 64 \text{ mm s}^{-1}$ ) trials (inset, mean  $\pm$  propagated SEM). (B) Time estimates decoded from DLS activity around S2 presentation conditioned on I1. Main axes: mean posterior distributions across runs. Inset: maximum a posteriori (MAP) point estimates averaged across runs (median and IQR). 96

3.16 DLS dynamics recorded during the first stimulus were negligibly affected by the intensity of the first stimulus. (A) Average normalized population activity conditioned on I1 and aligned to the onset of first stimulus (S1) (mean  $\pm$  SEM across neurons). (B) Responses of four example DLS neurons during S1. For each neuron, spike density functions (mean  $\pm$  SEM across trials) are conditioned on I1 and plotted above its spike raster, where trials are first sorted by I1, and then by T1 within each I1 level. (C) I1-split population activity projected onto each of the first three PCs computed using I1-agnostic population activity aligned to the onset of S1 (padded with 500 ms before S1 onset). (D) Population activity conditioned on I1 projected onto a 2D section of the subspace defined by the three PCs shown in (C). The 3-axis gizmo at the center depicts the projections of the axes defined by the first three PCs onto the chosen 2D section, the dashed black arrow indicates the direction of time, and the gray patches group population states recorded at each possible stimulus offset across I1. (E) Neurometric curves trained using S1-aligned population activity from middle-I1 trials conditioned on I1 (main axes), and differences between the proportion of T2 > T1 neural "reports" in high ( $I1 = 117 \text{ mm s}^{-1}$ ) and low ( $I1 = 34 \text{ mm s}^{-1}$ ) intensity trials compared to middle intensity ( $I1 = 64 \text{ mm s}^{-1}$ ) trials (inset, mean  $\pm$  propagated SEM). (F) Time estimates decoded from DLS activity around S1 presentation conditioned on I1. Main axes: mean posterior distributions across runs. Inset: maximum a posteriori (MAP) point estimates averaged across runs (median and IQR). 97

|      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |     |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 3.17 | Systematic biases in timekeeping across stimulus pairs are obscured (revealed) by conditioning trials on the outcome (category) of choices. <b>(A)</b> Hypothetical distributions of internal estimates of elapsed time as a function of experimental time since the onset of the first (left) and second (right) stimulus, plotted both continuously (colored gradient in the background) and discretely at the set of possible stimulus offset times (white distributions in the foreground). <b>(B)</b> Joint distributions of internal estimates of T1 and T2 for two pairs of stimuli with their respective marginals projected on each axis. <b>(C)</b> Joint and marginal distributions of T1 and T2 estimates for a stimulus pair for which T2 was shorter (left) or larger (right) than T1, conditioned on whether a direct comparison of those estimates would produce correct or incorrect choices. <b>(D)</b> Same as (C) but for all stimulus pairs. <b>(E-F)</b> Same as (C-D) but conditioned on the category of hypothetical choices strictly informed by a comparison of T1 and T2 estimates. . . . | 99  |
| 3.18 | <b>(A)</b> Average normalized population activity conditioned on choice and aligned to the onset of S2 (mean $\pm$ SEM across neurons). <b>(B)</b> Responses of four example DLS neurons during S2. For each neuron, spike density functions (mean $\pm$ SEM across trials) are conditioned on choice and plotted above its spike raster, where trials are first sorted by choice, and then by T2 within each choice category. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 100 |
| 3.19 | <b>(A)</b> Choice-split population activity projected onto each of the first three PCs computed using choice-agnostic population activity aligned to the onset of S2 (padded with 500 ms before S2 onset). <b>(B)</b> Population activity conditioned on choice projected onto a 2D section of the subspace defined by the three PCs shown in (A). The 3-axis gizmo at the center depicts the projections of the axes defined by the first three PCs onto the chosen 2D section, the dashed black arrow indicates the direction of time, and the gray patches group population states recorded at each possible stimulus offset across choice categories. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 100 |

|      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |     |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 3.20 | (A) Neurometric curves trained using S2-aligned population activity from choice-agnostic trials conditioned on choice. (B) Time estimates decoded from DLS activity around S2 presentation conditioned on choice. Main axes: mean posterior distributions across runs. Inset: maximum a posteriori (MAP) point estimates averaged across runs (median and IQR). . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 101 |
| 3.21 | Top: proportion of DLS neurons that were either up- or down-modulated by several experimental and behavioral factors across multiple task epochs. Signed modulation was assessed by quantifying the proportion of significant positive and negative coefficients in multivariable linear regression models fit to spike counts recorded in a 334-ms window aligned to the onset and/or offset of multiple task events. The "~" in ~reaction and ~choice refers to the approximate nature of the alignment for these epochs (334 ms and 668 ms after go cue for reaction and choice, respectively), which unlike all the others, lacked trial-to-trial annotated event times. The factor/regressor cross-correlation matrix is shown as an inset in the top left. Bar colors are factor-specific and respect all previously introduced color schemes (where applicable), and bar opacity refers to the significance level used to assess modulation (1% and 5% for opaque and transparent colors, respectively). Bottom: same as above, but highlighting modulation by I1 and I2. . . . . | 103 |
| 3.22 | (A) Total proportion of DLS neurons that were up- or down-modulated by I1 (in gray) and I2 (in black) across the different task epochs. (B) Same as (A), but for stimulus durations. . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 104 |
| 3.23 | Opposing effects of the duration of the spike integration window used in Fig. 3.21 on the average number of recorded spikes and the number of trials contributing to that estimate during S2 presentation epochs. Top: average number of spike counts across trials as a function of spike integration window for the intermediate I2 condition. Bottom: number of trials with recorded spike counts at each time point relative to the onset (left) or offset (right) of the second stimulus, split by I2 condition (median and IQR across neurons). . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 105 |

3.24 (A) I2-conditioned generative spike density functions used in simulated positive controls for the same four example DLS neurons shown in previous figures, aligned to the onset (left) or offset (right) of the second stimulus. (B) False negative (top) and false positive (bottom) rates computed by calculating the proportion of I2-modulated neurons during the S2 presentation epochs averaged over two separate sets of 25 runs corresponding to positive and negative controls. For each run 100% of the simulated neuronal responses were either positively modulated by I2 (when computing false negative rates) or not modulated by I2 (when computing false positive rates). . . . . 106

3.25 Bottom: proportion of DLS neurons labeled as "ramping"(up and down, where applicable) and "non-ramping"using the classification and grouping criteria from Toso and colleagues. Task epochs within which classification criteria were first enforced by the authors are washed out, and stimulus epochs used to subsequently assess decoding performance for each cluster are vivid. "Ramps" in each stimulus period (S1 and S2) do not preserve their initial direction qualifier because they were defined as the union of up- and down-ramping labels across the respective onset and offset epochs (an S2 "ramp" might have ramped up during the S2 onset epoch and down during S2 offset epoch). Top: cluster-split distributions of temporal tuning across DLS responses recorded during 1-s windows centered on each task event (washed out), or aligned to the onset of each stimulus period (vivid). Statistical departure from unimodality was assessed within cluster using Hartigan's dip test, and differences between clusters were assessed using two-sample Kolmogorov-Smirnov tests with Bonferroni correction for multiple comparisons (not significant (NS):  $p>0.05$ , \*:  $p<0.05$ , \*\*:  $p<0.01$ ). . . . . 109

|      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |     |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 3.26 | Bottom: proportion of DLS neurons labeled as "ramping"(up and down, where applicable) and "non-ramping"using the classification and grouping criteria from Toso and colleagues. Task epochs within which classification criteria were first enforced by the authors are washed out, and stimulus epochs used to subsequently assess decoding performance for each cluster are vivid. "Ramps" in each stimulus period (S1 and S2) do not preserve their initial direction qualifier because they were defined as the union of up- and down-ramping labels across the respective onset and offset epochs (an S2 "ramp" might have ramped up during the S2 onset epoch and down during S2 offset epoch). Top: cluster-split distributions of temporal tuning across DLS responses recorded during 1-s windows centered on each task event (washed out), or aligned to the onset of each stimulus period (vivid). Statistical departure from unimodality was assessed within cluster using Hartigan's dip test, and differences between clusters were assessed using two-sample Kolmogorov-Smirnov tests with Bonferroni correction for multiple comparisons (NS: $p > 0.05$ , *: $p < 0.05$ , **: $p < 0.01$ ). . . . . | 110 |
| 3.27 | Specification of the model used to simulate Gaussian bump responses (top) and generative spike rates in ten example trials of two pairs of simulated neurons that varied either in their stereotypy (bottom left) or gain (bottom right) parameters. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 111 |
| 3.28 | Distributions of temporal tuning (top), stereotypy (middle) and gain (bottom) parameters agnostic to (in white) or split by whether the corresponding simulated neurons were labeled as "ramps" or "non-ramps" using the authors' criteria. Horizontal gray and black bars above temporal tuning distributions at the top represent the epochs in which Toso et al. enforced their clustering criteria and evaluated decoding performance, respectively. Statistical departure from unimodality was assessed for the temporal tuning of "ramps" using Hartigan's dip test (NS: $p > 0.05$ , *: $p < 0.05$ , **: $p < 0.01$ ). . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 111 |

|      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |     |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 3.29 | (A) Main axes: decoded time estimates from two equally sized subpopulations ( $N = 150$ ) of Gaussian bumps varying along their stereotypy, gain, and temporal tuning parameters, and conditioned on whether they were labeled as "ramping" or "non-ramping" (mean posterior mean $\pm$ mean posterior standard deviation (SD) across 100 simulations, each with 100 trials). Insets: decoded time estimates from a single simulated trial (posterior mean $\pm$ posterior SD) illustrating the decoding accuracy (top) and precision (bottom) metrics used in (B). (B) Comparison of decoding accuracy (top) and precision (bottom) from "ramping" and "non-ramping" neurons, for simulations from four differently parameterized models (left) and DLS data recorded during the two stimulus periods (right, median and IQR across simulations). Model variants were defined by whether their stereotypy and/or gain parameters were fixed ( $\eta = \eta_0$ and/or $\gamma = \gamma_0$ ) or sampled ( $\eta \sim$ and/or $\gamma \sim$ ) from the corresponding distributions shown in white in 3.28. Temporal tuning parameters ( $\mu$ ) always varied across simulated neurons and were drawn from the same uniform distribution for all model variants. Statistical differences across cluster and epoch distributions were assessed using two-sample Kolmogorov-Smirnov tests with Bonferroni correction for multiple comparisons (NS: $p > 0.05$ , *: $p < 0.05$ , **: $p < 0.01$ ). . . . | 112 |
| 3.30 | (A) Normalized spike density functions of all DLS neurons recorded during the presentation of the second stimulus. Neurons are ordered by polar angle in the subspace defined by the first two PCs describing firing dynamics during S2. (B) Same as (A), but for a hybrid simulation wherein neurons labeled as "ramps" using the authors' criteria were replaced by simulated ramps (evenly split between up- and down-ramping). . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 114 |
| 3.31 | (A) Joint distribution of PC coefficients conditioned on the authors' cluster identities. (B) Same as (A), but for a hybrid simulation wherein neurons labeled as "ramps" using the authors' criteria were replaced by simulated ramps (evenly split between up- and down-ramping). . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 115 |

3.32 (A) Distributions of clusterability metrics - polar angle of PC coefficients shown in 3.31 (top) and Pearson's correlation coefficient between the responses shown in 3.30 and time (bottom) - either agnostic to (in white) or split by cluster identity as assigned by the authors' classification method. Statistical departure from unimodality was assessed using Hartigan's dip test (NS:  $p > 0.05$ , \*:  $p < 0.05$ , \*\*:  $p < 0.01$ ). (B) Same as (A), but for a hybrid simulation wherein neurons labeled as "ramps" using the authors' criteria were replaced by simulated ramps (evenly split between up- and down-ramping). . . . . 115



|     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |     |
|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 4.2 | Boundary manipulation protocol. The thin horizontal dotted lines represent the four categorical boundaries rats experienced in these sessions (N = 5). Boundary changes followed the same rules as the temperature manipulation experiments: a control-manipulation-control 3-min block design with boundaries drawn at random and without replacement from the set $B = 1.32, 1.5, 1.68, 1.85$ s until exhaustion (at which point the set was replenished and the sampling process resumed). . . . .                                                                                                                                                                                                                                                                                                                                                      | 149 |
| 4.3 | Rats adapt to a dynamic decision boundary. <b>(A)</b> Performance in the first two sessions of boundary manipulations conditioned on categorical boundary (N = 5 rats): 1.32 (yellow), 1.5 (black), 1.68 (dark blue) or 1.85 (light blue) s. Psychometric curves fit to cross-animal averages of psychophysical data split by boundary, respectively shown as solid lines and markers of matching color (mean $\pm$ SEM). Bottom-right inset: Differences in proportion of long choices from each experimental condition to the control condition (mean $\pm$ propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each boundary condition's psychometric fit. Solid black lines represent MAP point estimates. <b>(B)</b> Same as (A), but for the last two sessions of boundary manipulations. . . . .       | 149 |
| 4.4 | <b>(A)</b> Distributions of percentage change in threshold relative to control (dilation) as a function of which categorical boundary was in effect for the first two boundary manipulation sessions. Markers represent individual threshold dilations, linked within animals by thin solid black lines. Boxplots show animal means (horizontal black lines) and SEM (colored bars). <b>(B)</b> Same as (A) but for the last two manipulation sessions (N = 5; one-sample two-tailed t-tests, significant $t_s(4) = [-7.21; 2.89]$ , $P = [0.04; 2e3]$ , non-significant $t(4) = 2.61$ , $P = 0.06$ ; repeated measures ANOVA followed by post-hoc contrasts with Tukey correction for multiple comparisons, $F(2, 12) = 15.64$ , $P = 5e-4$ ; smallest significant $q_s(8) = 5.18$ , $P = 0.01$ , non-significant $q_s(8) = 2.43$ , $P = 0.15$ ). . . . . | 150 |

4.5 Trial history affects categorization performance. Left: psychometric curves fit to cross-animal data pools either agnostic to (dashed black line) or conditioned on (solid colored lines) the duration of the previous trial and it having been correctly categorized ( $N = 22$  rats). Bottom-right inset: Differences in proportion of long choices from each trial history condition to the history-agnostic condition (mean  $\pm$  propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each trial history condition's psychometric fit. Solid black lines represent MAP point estimates. Right: same as before, but looking at choices following incorrect choices. . . . . 150

4.6 Coefficients of a logistic regression model fit to choice data using categorical regressors for multiple task variables on current and past trials. White markers refer to single animals, and black or colored (in the case of those that directly map onto the conditions introduced in Fig. 4.5) ones to averages across animals.  $\beta_0$  – intercept; B – broken fixation; C – correct choice; E – incorrect choice; S – short choice; L – long choice;  $t$  – index of the current trial;  $[0.60, \dots, 2.40]$  – stimulus duration (s). . . . . 151

## ACRONYMS

|                          |                                                                                                                                                              |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b><i>C. elegans</i></b> | <i>Caenorhabditis elegans</i> (pp. 5–16, 21, 22)                                                                                                             |
| <b><i>E. coli</i></b>    | <i>Escherichia coli</i> (pp. 2–10)                                                                                                                           |
| <b>ARS</b>               | area-restricted search (pp. 12, 16)                                                                                                                          |
| <b>BG</b>                | basal ganglia (pp. 17, 19–22, 26–28, 48, 56, 57, 59, 127, 128, 137, 139, 143)                                                                                |
| <b>BIC</b>               | Bayesian information criterion (pp. xxxv, 90, 118)                                                                                                           |
| <b>ChR-2</b>             | Channelrhodopsin-2 (p. 30)                                                                                                                                   |
| <b>CNS</b>               | central nervous system (p. 16)                                                                                                                               |
| <b>CPG</b>               | central pattern generator (p. 17)                                                                                                                            |
| <b>CR</b>                | conditioned response (p. 14)                                                                                                                                 |
| <b>CS</b>                | conditioned stimulus (p. 14)                                                                                                                                 |
| <b>DA</b>                | dopamine (pp. 11–14, 20–22, 129, 130, 133, 135, 146, 147)                                                                                                    |
| <b>DLS</b>               | dorsolateral striatum (pp. xxxii–xlii, 79–83, 88, 89, 91–98, 100, 101, 103–107, 109, 110, 112–114, 118, 120, 122, 124, 128–131, 133, 134, 136–143, 146, 147) |
| <b>DMS</b>               | dorsomedial striatum (pp. 131, 133, 134, 142)                                                                                                                |
| <b>DS</b>                | dorsal striatum (pp. xix, xxvi–xxxii, xliv, 29, 30, 38, 45–48, 50–56, 64, 66–68, 72, 128, 129, 138, 144, 146, 148)                                           |
| <b>FI</b>                | fixed interval (pp. 129, 136)                                                                                                                                |

|             |                                                                                                                                                 |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>FSI</b>  | fast-spiking interneuron ( <i>p.</i> 132)                                                                                                       |
| <b>GABA</b> | $\gamma$ -aminobutyric acid ( <i>pp.</i> 19, 20)                                                                                                |
| <b>GFP</b>  | green fluorescent protein ( <i>p.</i> 70)                                                                                                       |
| <b>GLM</b>  | generalized linear model ( <i>p.</i> 81)                                                                                                        |
| <b>GPe</b>  | globus pallidus externus ( <i>pp.</i> 19–22)                                                                                                    |
| <b>GPI</b>  | globus pallidus internus ( <i>pp.</i> 19–21)                                                                                                    |
| <b>I1</b>   | intensity of the first stimulus ( <i>pp.</i> xxxii–xxxvii, xxxix, 82–90, 93–98, 102–104, 106, 118–120, 123)                                     |
| <b>I2</b>   | intensity of the second stimulus ( <i>pp.</i> xxxii–xxxvi, xxxix, xl, 82–94, 98, 101–107, 118–120, 123)                                         |
| <b>IPI</b>  | inter-press interval ( <i>pp.</i> 131, 140–143)                                                                                                 |
| <b>IQR</b>  | interquartile range ( <i>pp.</i> xxiii, xxv, xxxii, xxxvi, xxxvii, xxxix, xlii, 36, 37, 42, 43, 56, 93, 96, 97, 101, 105, 112)                  |
| <b>ISI</b>  | inter stimulus interval ( <i>pp.</i> 81, 94, 102, 106)                                                                                          |
| <b>ITOI</b> | inter-trial-onset interval ( <i>p.</i> 63)                                                                                                      |
| <b>LDA</b>  | linear discriminant analysis ( <i>p.</i> 120)                                                                                                   |
| <b>LED</b>  | light-emitting diode ( <i>pp.</i> 62, 67, 68)                                                                                                   |
| <b>M1</b>   | primary motor cortex ( <i>pp.</i> xxviii–xxxi, 45, 50, 51, 53, 55, 64, 68, 128, 131)                                                            |
| <b>MAP</b>  | <i>maximum a posteriori</i> ( <i>pp.</i> xxv–xxviii, xliv–xlvi, 43, 44, 47–50, 71, 148–150)                                                     |
| <b>MFC</b>  | medial frontal cortex ( <i>pp.</i> 133, 134)                                                                                                    |
| <b>MLR</b>  | mesenchepalic locomotor region ( <i>pp.</i> 16, 17, 22)                                                                                         |
| <b>MSN</b>  | medium spiny neuron ( <i>p.</i> 102)                                                                                                            |
| <b>NID</b>  | normalized intensity difference ( <i>pp.</i> xxxii–xxxiv, 82–86, 122, 123)                                                                      |
| <b>NS</b>   | not significant ( <i>pp.</i> xl–xlili, 77, 109–112, 115)                                                                                        |
| <b>NTD</b>  | normalized time difference ( <i>pp.</i> xxxii–xxxiv, 82–87, 122, 123)                                                                           |
| <b>PBS</b>  | phosphate buffered saline ( <i>p.</i> 70)                                                                                                       |
| <b>PC</b>   | principal component ( <i>pp.</i> xx, xxiii–xxv, xxxvi–xxxviii, xlii, xlili, 32, 33, 36, 40, 43, 74, 91, 92, 94, 95, 97, 100, 113–115, 120, 140) |

|             |                                                                                                                                             |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>PCA</b>  | principal component analysis ( <i>pp.</i> 73, 91, 98, 120)                                                                                  |
| <b>PCB</b>  | printed circuit board ( <i>pp.</i> 62, 64, 65)                                                                                              |
| <b>PFA</b>  | paraformaldehyde ( <i>p.</i> 70)                                                                                                            |
| <b>PID</b>  | proportional-integral-derivative ( <i>pp.</i> 64, 65)                                                                                       |
| <b>RL</b>   | reinforcement learning ( <i>p.</i> 26)                                                                                                      |
| <b>RNA</b>  | ribonucleic acid ( <i>p.</i> 1)                                                                                                             |
| <b>RPE</b>  | reward prediction error ( <i>p.</i> 22)                                                                                                     |
| <b>S1</b>   | first stimulus ( <i>pp.</i> xxxvii, xl, xli, 81, 94, 96–98, 102, 106–110, 118, 119, 130, 140)                                               |
| <b>S2</b>   | second stimulus ( <i>pp.</i> xxxv–xlii, 81, 91–96, 98, 100–102, 104–110, 113, 114, 118, 119, 123)                                           |
| <b>SC</b>   | superior colliculus ( <i>p.</i> 17)                                                                                                         |
| <b>SD</b>   | standard deviation ( <i>pp.</i> xlii, 112)                                                                                                  |
| <b>SDF</b>  | spike density function ( <i>pp.</i> xx, xxiv, 32, 33, 40, 118, 120, 123)                                                                    |
| <b>SEM</b>  | standard error of the mean ( <i>pp.</i> xx, xxiv–xxxi, xxxiv–xxxviii, xliv–xlvi, 33, 37, 40, 44, 46–52, 54, 89–91, 93, 95–97, 100, 148–150) |
| <b>SNc</b>  | substantia nigra pars compacta ( <i>pp.</i> 20, 129, 135, 147)                                                                              |
| <b>SNr</b>  | substantia nigra pars reticulata ( <i>pp.</i> 19–22)                                                                                        |
| <b>SPN</b>  | spiny projection neuron ( <i>pp.</i> 19–22, 129, 130, 132, 133, 135, 136, 139, 143)                                                         |
| <b>STDP</b> | spike-timing-dependent plasticity ( <i>p.</i> 132)                                                                                          |
| <b>STN</b>  | subthalamic nucleus ( <i>p.</i> 19)                                                                                                         |
| <b>T1</b>   | duration of the first stimulus ( <i>pp.</i> xxxii–xxxiv, xxxvi–xxxviii, 82–89, 92, 93, 96–99, 102, 106, 117, 119, 120, 122, 123)            |
| <b>T2</b>   | duration of the second stimulus ( <i>pp.</i> xxxii–xxxviii, 82–89, 91–93, 95–100, 102, 107, 117–120, 122, 123)                              |
| <b>TEC</b>  | thermoelectric cooling device ( <i>pp.</i> 63–65)                                                                                           |
| <b>TED</b>  | thermoelectric device ( <i>pp.</i> xix, xx, xxvi, xxvii, xxxii, 28–31, 45–49, 56, 61, 63, 65–69)                                            |
| <b>TWR</b>  | tap withdrawal response ( <i>pp.</i> 13, 15)                                                                                                |
| <b>US</b>   | unconditioned stimulus ( <i>p.</i> 14)                                                                                                      |

|            |                                                                      |
|------------|----------------------------------------------------------------------|
| <b>V1</b>  | primary visual cortex ( <i>pp.</i> 81, 142)                          |
| <b>VB</b>  | ventrobasal complex ( <i>pp.</i> xx, 30, 32, 33, 58, 66, 67, 72, 73) |
| <b>VPL</b> | ventral posterior lateral nucleus ( <i>pp.</i> 130, 140)             |
| <b>VTa</b> | ventral tegmental area ( <i>p.</i> 20)                               |
| <b>ZIP</b> | Zeta inhibitory peptide ( <i>pp.</i> 131, 133)                       |

## GENERAL INTRODUCTION

However it was that life on Earth began (~3.6 billion years ago), whether in a warm little pond<sup>1</sup> or in a world of ribonucleic acid (RNA), at some smeared point(s) in evolutionary time things must have gone from: non-cellular to cellular; prokaryote to eukaryote; autotroph to heterotroph; non-predator to predator; non-sexual to sexual; small to large; unicellular to multicellular; multicellular to tissuegrade; sessile to motile; non-neural to neural (in no strict order, delineation or direction; Bengston, 2002).

Each of these transitions belies some of the most profound and unsolved scientific questions in biology. Neuroscience in particular, is by and large preoccupied with brains and behaviors, were it not its ultimate goal to understand how the former give rise to the latter, but getting there will no doubt require distilling the general principles that govern each, and since both are the staggered product of evolution, one might as well start at the start.

### 1.1 The organization of behavior

“Let there be behavior”, and there was behavior – no. As with most dichotomies, distinctions get less different the closer we get to their defining boundaries, especially when ambiguous definitions are involved (Calhoun & El Hady, 2023). Yet, it might still be fruitful to exaggerate near inflection

---

<sup>1</sup>In a 1871 letter to Joseph Hooker, Charles Darwin famously speculated that “[...] if (and oh what a big if) we could conceive in some warm little pond with all sorts of ammonia and phosphoric salts, light, heat, electricity etcetera present, that a protein compound was chemically formed, ready to undergo still more complex changes [...]” (Griffiths, 2007).

points and consider *that* original behavior. One sensible proposal follows from the premise of *homeostasis* (Cannon, 1932), and posits that both physiology and behavior branched out from the same basic negative feedback mechanisms of metabolism, respectively leveraging internal and external dynamics on opposite sides of the cell membrane such that the system as a whole, organism plus environment, flowed towards desirable states (Ashby, 1960; Cisek, 2019). Through this lens, *taxis* navigation stands out as a plausible root for the taxonomy of behavior.

Chemo-, photo-, thermo-, and aerotaxis respectively (and allegedly) enabled those early single-celled “behavers” to navigate chemical, light, temperature, and oxygen-level gradients towards benefit and away from harm. How they did it is not known, but if extant chemotactic organisms provide any clue, then not only was *time* crucial to behavior since its conception, but so too was *hierarchy*.

### 1.1.1 Time and hierarchy

The flagellated bacteria *Escherichia coli* (*E. coli*) are unicellular prokaryotes. They are devoid of neurons and muscles (tautologically). Yet, through a network of receptor proteins acting on signaling proteins acting on protein effectors, they are able to chemotax their way towards a range of attractants and away from a range of repellents (Adler, 1966), making them a model organism for the study of primitive behavior.

As a consequence of their microscopic size (a few microns long), *E. coli* function essentially as point sensors (Macnab & Koshland, 1972), unable to measure gradients by comparing spatial concentration differences directly like some of their (larger) eukaryote counterparts (Swaney et al., 2010). Instead, they employ temporal gradient sensing by moving over distances several times their own length and comparing current and past chemical environments (Hazelbauer, 2012; Webre et al., 2003). In the absence of gradients, this results in random walks (reminiscent of Brownian motion), consisting of relatively long (~1 s) straight *runs* and short (~0.1 s) reorienting *tumbles* (H. Berg & Brown, 1972). Runs are the product of concerted *counterclockwise* rotation of (on average four) rotary motor proteins (spinning at roughly one revolution per 0.01 s), which causes the flagella – long, thin, helical, semi-rigid filaments – to which these are attached to bundle

up and propel the cell forward, in a corkscrew<sup>2</sup> fashion (H. C. Berg, 2003; Vladimirov & Sourjik, 2009). Reversing the direction of rotation of one or more individual flagellar motors to *clockwise* disrupts this bundle and causes the cell to stop and tumble erratically, and when those motors return to anticlockwise rotation the cell runs off in a new random direction (Wadhams and Armitage, 2004).

This ability to move is not inherently beneficial but, critically, *E. coli*'s random walks become biased in the presence of attractant or repellent gradients (H. Berg & Brown, 1972). It is this bias that enables chemotaxis, and it results from the chemosensory system directly impinging on the motor system by reducing (increasing) the probability of tumbles relative to runs for swims along (un)favorable directions (Segall et al., 1986). Specifically, the likelihood of motors spinning clockwise (and causing a tumble) is enhanced when a free-floating signaling protein binds to the cytoplasmic end of the flagellar motor. This binding is in turn contingent on that signaling protein first being phosphorylated by a kinase that is itself regulated by the activity of transmembrane chemoreceptors – it is inhibited (excited) by increases (decreases) in the concentration of attractants and vice versa for repellents (Cluzel et al., 2000). Lastly, a constitutively active phosphatase guarantees rapid turnover ( $\sim 0.1$  s) of the phosphorylated signalling protein back to its unphosphorylated state, so that tumbling can cease (Segall et al., 1986). In a nutshell, improvements in chemical environment lead to less kinase activity which leads to less phosphorylated signalling protein around which leads to less probable motor switching which leads to less tumbling. And the opposite applies when things get worse.

But how is the temporal concentration gradient itself –  $\Delta C/\Delta t$  – computed in the first place? When we say that *E. coli* cells “compare” present and “remembered” past chemical “concentrations” across “time” and then “decide” whether to run or tumble, we run the risk of inadvertently inviting literal interpretations of these symbolic metaphors that would have us look for direct molecular correlates of the intervening quantities and operations. Yet, as we just saw for the mechanism governing the switch between counterclockwise and clockwise rotation of flagellar motors, there is no strict

---

<sup>2</sup>It is worth emphasizing that intuitions about movement through water do not carry over well from the macroscopic to the microscopic world. Due to their size and range of travelling speeds, *E. coli* experience very high viscosity and almost no inertia, so when they stop, they stop immediately, and they do not coast or displace liquid (Purcell, 1976).

requirement for any particular instantiation of explicit representations. The same seems true for gradient sensing. Rather than time being readily decodable from the state of some dedicated molecule, its encoding is implicit in the dynamics of the system as it *adapts* to stimuli in the environment. Concretely, in *E. coli* there exist two separate enzymes that counteract each other to respectively methylate and demethylate the cytoplasmic domain of chemoreceptors. Methylation, which happens at a constitutive rate, affects receptors in two distinct ways: it decreases their ability to suppress the kinase that ultimately triggers tumbles; it decreases their affinity to ligands. The net effect is that any level of attractant, however high, is only efficient at inhibiting tumbles for so long – methylation will eventually catch up, leading to higher kinase activity and by extension, higher probability of tumbling. In this way the system can be said to adapt to positive stimuli. Demethylation has the opposite effect, but is contingent on the kinase activity that arises from decrements in attractant binding, thus allowing for adaptation to negative stimuli (Borkovich et al., 1992). Put simply, if things keep changing coherently on the sensory side, they tend not to change on the motor side. No explicit records of time or concentration need be invoked. It seems that more often than not biological representations are *pragmatic* in this way, presumably because they evolved to produce effective control over and under a wide range of conditions and constraints, and not faithful descriptions of the environment (Cisek & Kalaska, 2010).

Abstracting mechanistic detail away, bacterial chemotaxis can be summarized as a time-dependent computation informing a switch between two relatively *discrete* behavioral motifs, each arising from the relatively *continuous* movement of lower-level actuators. At both levels there is a time-base to speak of, and these are (conceivably) dissociable. One could imagine manipulating the temporal bases upon which the delayed concentration comparison relies and affecting behavior at the higher and more discrete level, without directly affecting it at the lower and more continuous one. And conversely, one could in principle modify the rotation speed of flagellar motors and thus affect the kinematics of runs and tumbles without directly affecting their relative expression. This hierarchical framing of behavioral control and conjecture about the partial dissociability of the many timescales over which it unfolds will be a recurring theme throughout this thesis.

Speaking of which – neuroscience.

Taxis navigation in general, and chemotaxis in particular, are thought to

have evolved multiple times independently, within and across the domains of life (Jékely, 2011). In fact, the existence of chemotactic single-celled eukaryotes is often invoked as a likely prerequisite for the very formation of the first neural networks, seeing as it is one of the identified mechanisms for neuronal migration and axonal guidance during development (Tessier-Lavigne & Goodman, 1996). However, its first implementation *with* those nervous systems likely did not occur until early bilaterians (~600 million years ago), for which the nematode *Caenorhabditis elegans* (*C. elegans*), with its 302 neurons and 132 muscle cells<sup>3</sup> (Cook et al., 2019; White, 1986), can be considered a model organism (Bennett, 2021). These soil-dwelling roundworms are also able to chemotax towards numerous volatile and water-soluble compounds, including byproducts of bacterial metabolism (Bargmann et al., 1993; Bargmann & Horvitz, 1991; S. Ward, 1973). Strikingly (and somewhat ironically), the neurally-mediated chemotaxis that allows modern bacteri-ovores like *C. elegans* to feed on *E. coli* is in many ways similar to the bacterial flavor of this behavior.

One notable parallel is in the gradient computation itself – despite being big enough (~1-mm long) to in theory enable spatial gradient sensing, the chemosensory organ arrangement of *C. elegans* is such<sup>4</sup> that they still need to rely on single-point comparisons through time in order to assess gradients (S. Ward, 1973; S. Ward et al., 1975). Akin to *E. coli*, the system’s chemosensory response to fluctuating stimuli is largely linear and can be captured by a biphasic temporal filter: the more recent past (up to ~10 s ago) is given a large weight; the more distant past (from ~10 to ~60 s ago) is given a small weight of opposite sign, and sensory neurons respond to the difference (Kato et al., 2014). The result is a form of adaptation that equates to a running comparison extending tens of seconds into the past.

Another striking similarity lies in the hierarchical organization of *C. elegans* locomotive behavior, which coarsely clusters into extended forward *runs* (lasting tens of seconds to minutes) and comparatively brief reorienting episodes called *pirouettes* – pauses followed by reversals followed by  $\Omega^5$  turns

---

<sup>3</sup>These numbers refer to the adult hermaphrodite, which represents the overwhelming majority of *C. elegans* individuals in laboratory settings. Adult males have 385 neurons and 155 muscles (Cook et al., 2019).

<sup>4</sup>*C. elegans* crawls on its side by alternately contracting body wall muscles along the dorsoventral axis, and their chemosensory neurons are ~8  $\mu\text{m}$  apart along the mediolateral axis (orthogonal to the plane of motion; S. Ward et al., 1975).

<sup>5</sup>Sharp turning events with head-to-tail bends (named after the etchings they leave on agar once performed; Croll, 1975).

(in the seconds range). In place of the syllabic rotating flagella of *E. coli*, each of these stereotyped behaviors subsumes coordinated sub-second body bends in *C. elegans*. Forward runs are initiated at the head and propagate *posteriorly*, whereas the reversals that initiate pirouettes originate at the tail and propagate *anteriorly* (Croll, 1975; Kaplan et al., 2020). Furthermore, dwell times in both locomotory states are exponentially distributed, indicating that *C. elegans* locomotion is also fundamentally stochastic (Zhao et al., 2003). Lastly, just as with *E. coli* (barring a different gymnastics analogy), pirouettes are more likely to occur when a worm heads down (up) a gradient of attractant (repellent) and less likely to occur when a worm ascends (descends) it (Pierce-Shimomura et al., 1999).

At a low level, the neural circuit that implements chemotaxis (and the other kinds of taxis performed by *C. elegans*) can be drafted as two (mostly) feedforward chains with four stages: sensory neurons<sup>6</sup>; sensory interneurons; premotor interneurons; and motor neurons (Faumont et al., 2012). To a first approximation, sensory neurons are modality-specific (taste, smell, touch, etc.) and split into ON cells or OFF cells, respectively excited by increases or decreases in stimulus strength (Chalasani et al., 2007; Suzuki et al., 2008; Thiele et al., 2009). Multiple sensory neurons converge on individual sensory interneurons of two major functional types, defined by whether they use that cross-modal information to promote runs and suppress pirouettes or the other way round (Gray et al., 2005). Both run and pirouette sensory interneurons project mainly to premotor interneurons, which again tend to fall into two broad classes. One pool that mostly participates in forward locomotion and another that is recruited during reversals (the initial phase of a pirouette), but the two are heavily interconnected (presumably through inhibition; Cook et al., 2019; White, 1986). The motor neurons that ultimately generate undulatory thrust are distributed along the length of the animal and innervate local patches of dorsal and ventral musculature in the body wall. They are connected to upstream premotor interneurons via both electrical and chemical synapses and tend to inherit their functional role, thus establishing a *balancing* act that ultimately determines the direction of motion (Kawano et al., 2011). Positive and negative (chemo)taxis to a wide range of (chemical) stimuli then follows naturally from this push-pull arrangement.

---

<sup>6</sup>Some sensory neurons are coupled by gap junctions to hub-like interneurons that may serve the purpose of reporting an averaged, multimodal assessment of the local environment.

At a high level, this too can all be reduced to a sensorimotor loop, wherein temporal differentiation of sensory inputs biases the selection between two competing motor programs, that mostly differ in the direction with which (roughly) the same sequence of motor primitives unfolds. And again, we can posit that the time bases that inform behavior at both levels are to some extent dissociable. This tight algorithmic correspondence in spite of the massive differences in implementation<sup>7</sup> is a testament to the generality of *time* and *hierarchy* as organizing principles of behavior (Tinbergen, 1951), with time taking primacy as both a scaffold for decision and a determinant of hierarchical depth. Normative arguments for why that is may be founded on the intertwined notions of modularity, conservation, efficiency, balance, and pragmatism, all bound by “the arrow of time” (Popper, 1956).

### 1.1.2 Heterarchy

So far I have mostly focused on the similarities between *E. coli* and *C. elegans* chemotaxis, but as one might expect given the rifts in scale and complexity across the two, there are many key differences.

For one, unlike *E. coli*, *C. elegans* complement their biased random walk strategy (also referred to as *klinokinesis* in *C. elegans* circles) with *klinotaxis* (Fraenkel & Gunn, 1940), wherein heading is intermittently adjusted (every second or so) towards where the gradient is steepest (Ino & Yoshida, 2009). This steering takes place mainly during forward runs and is thought to result from a process of active sensing in which olfactory or gustatory signals are temporally coupled with the sinusoidal head sweeps that accompany undulatory movement (Izquierdo & Lockery, 2010; S. Ward, 1973).

The two strategies differ fundamentally in at least three regards. First, klinokinesis manifests as a bias in the selection between two distinct motor programs (Pierce-Shimomura et al., 1999), whereas klinotaxis consists of a small but steady steering bias (Ino & Yoshida, 2009; S. Ward, 1973). Second, the relevant sensory information for klinokinesis is temporally integrated over the course of tens of seconds to minutes (Pierce-Shimomura et al., 1999). In contrast, the time constants associated with klinotaxis are constrained by

---

<sup>7</sup>At the risk of flogging the metaphor to death, I will point out that the protein-based implementation of *E. coli* chemotaxis is also amenable to summary as a cascade with four levels: chemoreceptors (of which there are five types; Grebe and Stock, 1998), kinases (which integrate across receptor “modalities”), signaling protein (a common currency for control), and rotary motors (bistable effectors). One key aspect does seem exclusively neural though – pathway splitting (Gjorgjieva et al., 2014).

the  $\sim 1$ -s period of head oscillations (Ino & Yoshida, 2009; S. Ward, 1973). And third, although gradients are necessarily computed along the direction of travel in klinokinesis (in a trivial sense), the computation itself is agnostic to the current somatosensory and motor states of the animal. Klinotaxis on the other hand is inherently spatiotemporal, gradients are computed over time as worms swing their heads perpendicularly to the forward direction and thus depend on knowledge of ongoing head kinematics. And the asymmetries carry over to the implementation, with the two strategies involving partially dissociable circuits with different associated timescales (Ino & Yoshida, 2009; Liu et al., 2018), despite overlapping on the input and output sides (same sensory neurons and same muscles).

Even disregarding size differences for a moment, it is hard to imagine how *E. coli* might achieve something like klinotaxis. The lack of any direct coupling between specific chemoreceptors and specific rotary motors simply precludes any possibility for steering during runs. This is because in *E. coli*, relative changes in binding of an attractant or repellent to any receptor along the cell membrane are equally poised to affect the concentration of that single signaling protein that is then free to diffuse in the cytoplasm and bind to flagellar motors indiscriminately. It was presumably only with the advent of specialized neurons and muscles ( $\sim 700$  million years ago), and the wiring between sensory inputs and motor outputs that came with them (Arendt et al., 2016), that steering became possible (a point I will return to in section 1.1.4). And more importantly, the possibility of simultaneous coexpression of multiple policies aimed at achieving similar goals with the same peripherals (as seen in *C. elegans* chemotaxis) was likely predicated on those nervous systems first becoming sufficiently deep such that they could then become sufficiently parallel, with multiple routes between sensors and effectors.

Critically, these routes are not one-way streets where information flows in as afference and out as efference. They are loops. Sensory stimuli elicit actions with sensory consequences that are correlated with those actions. These reafferent contingencies can be leveraged directly (as done by *E. coli* and *C. elegans* in klinokinesis), or anticipated given copies of the very motor commands that led to them – *efference copy* (Jékely et al., 2021). Indeed, in *C. elegans* klinotaxis, a dedicated subset of pirouette sensory interneurons subtracts oscillatory sensory reafferent signals from the corollary discharges of dorsal and ventral motor neurons in the forward locomotion pool, and

then projects the residuals back onto those same motor neurons in the form of inhibition. If a gradient is detected perpendicularly to the forward direction, this produces an asymmetry that steers the worm away from or towards its source (depending on modality and valence). If there is no gradient detected, either because there is none to begin with or because the worm is already aligned with it, then there is no asymmetry in inhibition onto dorsal and ventral forward motor neurons and hence, no steering (Liu et al., 2018).

Altogether, this is one example of multiplicity in behavioral control. There exist many others, some of which I have mentioned in passing. For example, that *E. coli* have more than one type of chemoreceptor, and that in *C. elegans* there are multiple modalities on top of that, already speaks to a kind of distributed control. But the klinokinesis and klinotaxis example is special in that it obviates the need for crosstalk between, and feedback within, the two main pathways of the taxis hierarchy in *C. elegans*. Strictly speaking, this *recurrence* renders the resulting organizational structure non-hierarchical, or at the very least non-canonical (Dawkins, 1976). As far as I can tell, there is no universally agreed upon term for something that has the general appearance of an overlapping branching hierarchy but where elements in lower levels can have causal influence on those above, though one proposal has been *heterarchy* (McCulloch, 1945) – a system where executive control is distributed and the modules interact not strictly in series like master and slave, but more like a community of experts in parallel (Merel et al., 2019; Winston, 1972). Another option is to relax the implicit assumption of static direction of flow between levels in a hierarchy and embrace “*changing boss-ships as an interesting complicating factor*” (Dawkins, 1976), such that what has causal influence on what can rapidly switch depending on whether information or control are being considered. Aspects of both the deposing heterarchical view and the reconciling hierarchical one resonate with many successful ideas in systems neuroscience<sup>8</sup> and so, with questionable propriety, I am coopting heterarchy to loosely mean both.

---

<sup>8</sup>If I am allowed a millennial moment (on top of a verging-on-pedantic moment), I might also make a meta point and add that it resonates with my struggles in writing this introduction – for the one part there has to be a unidirectional flow of information (you will not read this in reverse), but for the other I have to (or ought to) be mindful of what comes next and exert control to restrain that flow, so that the whole thing does come across disjointed.

### 1.1.3 Flexibility

One of the primary functions of behavior, and by extension, of nervous systems, is to maximize evolutionary fitness in complex and unpredictable environments. Doing so hinges on the ability to weigh competing inputs across multiple sensory streams in a way that is both aligned with current physiological needs and reflective of past experience.

Basic forms of sensory integration are at play even in bacteria. For example, when *E. coli* are exposed to composite gradients of a mixture of ligands to which two of their chemoreceptor-types are tuned, chemotactic decisions depend on relative concentrations as well as on the expression ratio of the recruited receptors (Adler & Tso, 1974; Grebe & Stock, 1998; Kalinin et al., 2010). However, the coupling between specific chemicals and behavioral responses is fixed – the molecules that *E. coli* deems attractive or repulsive are physiologically determined to be so, they are directly beneficial or detrimental<sup>9</sup>. In *C. elegans*, this mapping is less stringent, with some stimuli (e.g., odors, tastes and temperatures) functioning as cues – means to ends rather than ends in themselves. It is also not as static, meaning that how a particular stimulus informs behavior is determined by a combination of innate, contextual and experience-dependent factors (Ghosh et al., 2016; Ishihara et al., 2002).

Indeed, in an assay where in order to reach the source of an appetitive gradient worms have to first cross a barrier of a harmful chemical, the probability of crossing increases monotonically with the concentration of attractant, and decreases with the concentration of repellent (Ishihara et al., 2002). Notably, this multisensory summation process is not rigid. It is modulated by hunger state, with the proportion of animals crossing the barrier increasing as a function of starvation (Ghosh et al., 2016). The circuit implicated in this cross-modal and context-dependent integration mainly consists of two sensory neurons converging onto a single sensory interneuron of the pirouette-promoting type. The projections from the attractant- and repellent-sensitive neurons are respectively inhibitory and excitatory, leading to fewer (more) pirouettes when aligned with a relatively higher (lower) gradient in attractant compared to that of the repellent. The context-dependency then comes in due to this interneuron also secreting a

---

<sup>9</sup>Nicely illustrated in chemotaxis to sugar, with bacteria partially depleting the gradient as they climb it (Adler et al., 1973).

*neuromodulator*<sup>10</sup>, one that covaries with satiation level and exerts a top-down facilitating influence on the sensory neuron of the aversive modality (Ghosh et al., 2016). This results in a sensitivity to increments in repellent that is inversely proportional to the animals' hunger, fostering a progressively higher tolerance to threat (i.e., likelier barrier crossings) as worms become increasingly starved. Thus, *C. elegans* is capable of using information about internal states to bias decisions that entail conflicting sensory signals.

There is one other neuromodulator of the biogenic monoamine family that has been implicated in this and many other behaviors in *C. elegans* – dopamine (DA). In total, there are eight<sup>11</sup> DA-synthesizing mechanosensory neurons in the head of *C. elegans* that also corelease glutamate, and these act both through classic synapses (429 in total; White, 1986) as well as extrasynaptically by releasing DA into the worm's body cavity (Sulston et al., 1975). Consistent with a highly conserved evolutionary role, many of the mammalian enzymes involved in DA synthesis, vesicle loading and reuptake find homologs in *C. elegans* and, importantly, the homology extends to receptors – *C. elegans* also expresses D1- and D2-like dopamine receptors, and these have also been linked to antagonistic effects on both neuronal excitability and behavior (I will elaborate on these in section 1.1.5; Vidal-Gadea and Pierce-Shimomura, 2012). Case in point, in the threat-reward paradigm described above, *C. elegans* mutants lacking D1-like receptors display an increased tendency to cross the aversive barrier as compared to wild-types, and the opposite is true for D2-like receptor mutants (D. Wang et al., 2014).

Dopamine signaling in *C. elegans* has additionally been tied to various aspects of another class of complex decisions that is universally crucial for survival and reproduction – *foraging* for food. Though chemotaxis is a reliable strategy when it comes to locating resources on a small scale, it can prove insufficient in ecological niches like those occupied by *C. elegans* in the wild<sup>12</sup>, in which resource distributions are dynamic and prone to “clumping”. Strikingly, when entering a given bacterial lawn (a patch of food), worms

---

<sup>10</sup>Tyramine, which together with octopamine, dopamine and serotonin makes up the four biogenic monoamines that act as neuromodulators in *C. elegans* (Chase & Koelle, 2007).

<sup>11</sup>Four pairs of bilaterally symmetric dopaminergic neurons in the head of the hermaphrodite, the male has an additional six in the tail (totaling 14; Sulston et al., 1975).

<sup>12</sup>*C. elegans* is a so-called fruit nematode, which forages for blooms of bacteria in compost heaps containing decaying vegetable material (e.g., rotting fruits, flowers, stems, Kiontke et al., 2011).

will slow their locomotion rate in a DA- and D2-receptor-dependent manner, seeing as mutants lacking the ability to either synthesize DA or express D2-like receptors fail to exhibit this slowing response (Chase et al., 2004; Sawin et al., 2000). Furthermore, exogenous exposure of *C. elegans* to DA across a range of concentrations also slows down movement, and also in a way that depends on DA receptor type (Schafer & Kenyon, 1995): D1-like receptor knock-out mutants exhibit a lower threshold for this DA-induced immobility; D2-like mutants, a heightened resistance to it (Chase et al., 2004). Altogether, this bidirectional neuromodulatory regulation of locomotion speed may be an adaptive mechanism by which *C. elegans* calibrates the amount of time spent in any given food patch, thus fine-tuning *exploitation* (Charnov, 1978).

After exhausting a particular patch (or being experimentally removed from it), off-food worms will search the immediate area for additional food before expanding their search to include more distant, surrounding areas (T. Hills et al., 2004). The initial local search is characterized by a high frequency of pirouettes, which then decreases abruptly ~15 minutes after food was last encountered (Calhoun et al., 2014). This neatly exemplifies area-restricted search (ARS), a ubiquitous foraging strategy found in all major eumetazoan phyla (e.g., cnidarians, arthropods, chordates, molluscs and nematodes), and an optimal search algorithm under the kinds of resource distributions that follow from the tendency of living things to grow, reproduce, and disperse in spatially autocorrelated ways (T. T. Hills, 2006). Dopamine, through its effect on downstream premotor interneurons, onto which glutamatergic inputs also converge (White, 1986), seems to also be a key component in regulating ARS, specifically the transitions between exploitative and explorative states (Calhoun et al., 2014; T. Hills et al., 2004). Mutant worms unable to synthesize DA fail to dwell upon removal from food and instead start to roam immediately, an effect that can be rescued by exogenous application of DA (T. Hills et al., 2004). Crucially, this is contingent on normal glutamatergic signalling, as mutations in genes specific to the expression of glutamate receptors in premotor interneurons (postsynaptic to DA), are sufficient to eliminate ARS (even in the presence of exogenous DA; T. Hills et al., 2004; A. Maricq et al., 1995). (Expect to see this motif – DA modulating behavior via its effect on the efficiency of glutamatergic synapses – resurfacing). A model consistent with all of these observations suggests that while on food, DA resulting from the detection of that food acts to slow down movement so that worms do not wander off. If they do wander off, then lingering

DA acts to increase the probability of them wandering back in through its modification of glutamatergic synapses. If however, a patch is fully depleted, then following a period of intensive local search to ensure that that is really the case, DAergic activity decreases, the frequency of pirouettes diminishes, and worms switch to a more expansive global search (Chase & Koelle, 2007).

In a broader sense, examples like these, of modulation of behavior by latent states that echo recent history, can be framed as forms of short-term behavioral plasticity. To longer-term forms, we usually reserve words like *learning*, and *memory* (though the distinction seems fuzzy at best). Remarkably, several forms of learning, both nonassociative and associative, have been demonstrated across the spectrum of *C. elegans* behavior (Ardiel & Rankin, 2010). One of the first and arguably best-studied instances is the (nonassociative) *habituation* of the tap withdrawal response (TWR)<sup>13</sup>, an escape reflex whereby worms perform pirouettes in response to a nonlocalized mechanical stimulus (a tap to the plate containing the worm<sup>14</sup>). In brief, repeated tapping results in a reduction of both amplitude and frequency of the TWR (Rankin et al., 1990). As seen in other organisms, habituation, and spontaneous recovery from it (i.e., forgetting), depend critically on the schedule of stimulus delivery – shorter (longer) inter-tap-intervals lead to faster (slower) and deeper (shallower) habituation, as well as faster (slower) recovery. Additionally, learning rate depends on the presence of food – when tested off of food, worms habituate faster than when tested on it (Kindt et al., 2007). And lastly, the mechanosensory habituation of pirouettes in the context of the TWR does not generalize to nonmechanically-induced pirouettes outside of it (Wicks & Rankin, 1997), suggesting that plasticity lies somewhere upstream of the pirouette-promoting premotor interneurons. Indeed, the neural underpinnings of habituation and retention of the TWR memory (lasting up to ~24 h post-training) that have been identified so far imply a combination of: DA- and D1-receptor-dependent attenuation of excitability in mechanosensory neurons (Kindt et al., 2007; Sanyal et al., 2004; Suzuki et al., 2003); depotentiation of glutamatergic synapses onto premotor interneurons (Rose et al., 2003), presumably via DA (T. Hills et al., 2004). (Again *that* motif). Looks to be that non-associative learning joins the

---

<sup>13</sup>Analogous to the famous habituation of the gill-withdrawal reflex in *Aplysia californica* (Pinsker et al., 1970).

<sup>14</sup>Following up on anecdotal reports of *C. elegans* showing decreasing responses to repeated touches (Croll, 1975).

non-exhaustive list of non-exclusive roles for DA in *C. elegans* covered so far: reward signalling, motor control and foraging.

Onto associative learning then – *C. elegans* is also capable of simple forms of classical conditioning (Ardiel & Rankin, 2010), specifically (and vitally) those of the *simultaneous* kind, in which the presentations of the conditioned stimulus (CS) and the unconditioned stimulus (US) coincide in time (Pavlov, 1927). If this condition is met, worms will learn that a given CS (taste, smell, temperature or oxygen-level) predicts a given US (food, no food or aversive stimulus), as reported by a given conditioned response (CR) (chemotaxis, thermotaxis or aerotaxis). In the case of choice tests between diffusive gradients of two chemicals<sup>15</sup> for which naïve worms display no initial preference, preferences can be induced reliably and bidirectionally by pairing each chemical with either food or a repellent during an initial conditioning phase (Wen et al., 1997). Moreover, *C. elegans* can dramatically overwrite innate preferences for particular compounds if these are first paired with no access to food or with aversive stimuli for extended periods of time (in the hours range) prior to testing (Hukema et al., 2008; Ishihara et al., 2002; Saeki et al., 2001). And crucially, the plasticity underlying the switch from migration towards to migration away from what initially was a chemoattractant requires both dopamine and glutamate (Hukema et al., 2008). (Again). Despite an unimpaired ability to chemotax, *C. elegans* mutants lacking an enzyme necessary for DA synthesis fail to acquire the CR unless preincubated with exogenous DA; mutant strains with defects in glutamatergic signaling show general deficits in chemotaxis and still approach the CS after conditioning (Hukema et al., 2008). It seems that DA's part in learning is as ancient as DA itself<sup>16</sup>.

The remarkable capacity of *C. elegans* to flexibly adapt default preferences in light of recent (and not so recent) experience as well as current homeostatic goals begs the question: if steering, sensorimotor integration, foraging, learning, and memory are no challenge, what is it that these worms, and bilaterians in general, can not do? In keeping with the theme so far, of trying to understand the basic building blocks of behavior and its neural bases

---

<sup>15</sup>Thematic harmony with chapter 2 would have me explore conditioning with thermotaxis, but (sadly) that would require providing additional context and it seems that the chemotaxis case is the most thoroughly worked out (Ardiel & Rankin, 2010).

<sup>16</sup>It bears noting that there exist multiple DA systems in modern vertebrates, which may or may not have all stemmed from a root common to the DA system found in nematodes (Yamamoto & Vernier, 2011).

though phylogenetic refinement (Cisek, 2019), I will next be exploring the anatomical and behavioral innovations of chordates and vertebrates, and inch my way up to mammals.

### 1.1.4 Topography

*Imagine a vast sheet of paper on which straight Lines, Triangles, Squares, Pentagons, Hexagons, and other figures, instead of remaining fixed in their places, move freely about, on or in the surface, but without the power of rising above or sinking below it, very much like shadows – only hard and with luminous edges – and you will then have a pretty correct notion of my country and countrymen. [...] Nothing was visible, nor could be visible, to us, except Straight Lines.*

Abbott, 1884, *Flatland: A romance of many dimensions*

The quote above may as well have been “uttered” by our last common ancestors with *C. elegans*, who likely lived fairly flat and benign lives on, or just below, the microbial rugs of the Ediacaran ocean floors – biomats (Seilacher, 1999). Though technically predators (by most definitions), they were grazers in essence, chemotaxing from lawn to lawn while avoiding (mostly) environmental hazards. This would change at the turn of the period (~530 million years ago), when a nuanced interplay between increases in water oxygenation and a surge in macroscopic predation ignited the so-called Cambrian explosion<sup>17</sup> – a burst of diversification that ultimately led to the divergence of most modern eumetazoan phyla, including chordates and eventually, vertebrates (Bengston, 2002; Fox, 2016).

The shift from the “peaceful Garden of Ediacara to the Cambrian bloodbath of predator eating predator” (Jacobs, 2012) meant hard times for basal bilaterians bearing any semblance to *C. elegans*, with their sluggish pace and lack of discernible defenses beyond close-range withdrawal reflexes (e.g., the TWR). This, together with the surplus of energy that the elevated oxygen levels afforded, presumably spurred a positive feedback loop of sensorimotor innovations leading to more and more sophisticated pursuit and evasive behaviors – especially in the lineage of vertebrates (Cisek, 2021). Indeed,

---

<sup>17</sup>Though the term “explosion” undoubtedly applies in the scope of geological time, it is a bit of a misnomer when it comes to evolution, a process of gradual accumulation over successive generations that explores variation in heritable traits.

fossil records from the early Cambrian show that worm-like animals started to burrow several centimeters into the sediments beneath the biomat around this time, which may have provided refuge as well as access to previously untapped ecological niches. It is also plausible that those same selective pressures drove some animals in the opposite direction, up into the water column above the seabed (Fox, 2016). In either case, the departure from a two-dimensional life of tranquility to a three-dimensional one of hostility is sure to have placed increasing demands on fast maneuverability and spatial sensing abilities. This required physical adaptations of sense organs themselves, which went from something like the (in practice) single-point chemo-, mechano- and photosensory neurons of *C. elegans* (Bargmann & Horvitz, 1991; Sawin et al., 2000; A. Ward et al., 2008) to the laterally symmetric pairs of sensor arrays found in most modern vertebrates.

But first, there were chordates, for whom the closest living proxy is held to be the amphioxus (*Branchiostoma floridae*; Holland, 2014) – a small (~3-cm long) fish-like filter-feeder that, though able to swim, mostly lives half-buried in the sand, oblique to the sea bottoms. Its entire central nervous system (CNS) consists of a nerve cord with a slight swelling at the anterior end – the cerebral or sensory vesicle. Despite its simplicity, this vesicle is considered a precursor to the vertebrate brain, as it contains homologs of many of its key structures: a diencephalon; hints of a telenchepalon; a small mesencephalon, and a rhombencephalon (Holland, 2015). The former is essentially a rudimentary hypothalamus connected to a central non-image-forming eye, and is thought to be mainly involved in circadian rhythms and homeostatic regulation, as well as control of foraging behaviors – specifically the shift from filter-feeding patterns of local exploitation to those producing more global exploration. As with *C. elegans*, these ARS-like foraging decisions seem critically dependent on dopamine and its interaction with D1- and D2-like receptors (T. T. Hills, 2006; Zieger et al., 2017). These tend to be mostly expressed at glutamatergic synapses in a probable homolog of the vertebrate mesencephalic locomotor region (MLR), which is also heavily implicated in generating escape responses (Lacalli & Candiani, 2017).

When disturbed, amphioxi quickly leave their burrow, swim a short distance, and then rapidly burrow again, tail first into the sand (D’aniello et al., 2023). The initial escape is prompted by a reduction in light detected at a single frontal eye patch (posited to have evolved from the rostral *C. elegans*-like photoreceptors of early bilaterians; A. Ward et al., 2008), resulting

in *bilateral* excitatory input to a putative precursor of the vertebrate optic tectum (Lacalli, 2018) – the superior colliculus (SC) in mammals. These prototectal regions then project *ipsilaterally* to the MLR homolog, which in turn drives central pattern generators (CPGs) in the spinal cord that implement fast undulatory swimming (Cisek, 2019; Lacalli & Candiani, 2017). Crucially, the frontal eye of the amphioxus faces inward, posteriorly, with an opaque concave pigment cup anterior to it<sup>18</sup> (Lacalli, 2018), such that a shadow coming from anywhere but above will trigger an undirected escape. This configuration seems particularly well suited to evade predators venturing up from the biomats below.

Early vertebrates would then elaborate on these two foundational chordate systems for behavioral control, namely the hypothalamic and the prototectal circuits for foraging and escape (Cisek, 2021). Starting with the latter, the frontal eye patch is thought to have split into two and migrated to the sides of the head, eventually becoming the retinas of the lateral eyes (Lacalli, 2018). In the oldest group of now-living vertebrates – lamprey (*Petromyzon marinus*), a ~60-cm long predatory jawless fish – the temporal and nasal parts of each retina project *contralaterally* to the rostral and caudal portions of each tectum, which in turn project to *contralateral* and *ipsilateral* premotor locomotor centres in the brainstem, respectively (Suryanarayana et al., 2021). The result is a tectal co-registration of a retinotopic map of the surrounding space with an egocentric motor map downstream. A *topographic* Braitenberg-vehicle-like configuration that enables rapid oriented escape (approach) from (towards) potential predators (prey) in the local environment (Braitenberg, 1984; Isa et al., 2021).

The other major innovation of early vertebrates was the expansion of one of two main segments of the hypothalamus, which “mushroomed” out into what would one day become the telencephalon (Benito-Gutiérrez et al., 2021; Puelles & Rubenstein, 2015). In lamprey, this consists of an outer *pallium* (Latin for mantle) and an inner *subpallium* (as in under the mantle), which are respectively seen as precursors to the mammalian cortical mantle (both allo and neo) and the subcortical basal ganglia (BG). The key insight here, by Paul Cisek, goes back to the operational definitions of physiology and behavior adopted in this introduction (see 1.1), and is that functionally, this telencephalic outgrowth of the peduncular segment of the hypothalamus is

---

<sup>18</sup>Unsurprisingly, there are plenty of videos on the internet of amphioxi swimming face-first into the ground.

trying to achieve the very same thing as its diencephalic terminal counterpart – homeostasis (Cannon, 1932; Puelles & Rubenstein, 2015). It is just that where the latter might promote vasoconstriction and shivering to deal with a drop in body temperature, the former would bias ancient locomotor circuits towards shelter; where the latter might adjust metabolic rates in response to a growing hunger, the former would promote hunting. One deals in physiological control on the inside, the other in behavioral control on the outside. Much of the history of vertebrate evolution is then the history of how that control stretched further and further out into the environment (both in time and space), hand in hand with the expansion, subdivision and specialization of the pallial and subpallial systems of the telencephalon<sup>19</sup> (Cisek, 2021).

Notably, a common design pattern arises both within and across the cortical and subcortical structures that subsequently developed. From the early sensory areas of mammals where visuotopic; tonotopic; odotopic; gustotopic and somatotopic maps can be found, to the mototopic maps in later motor areas, and all of the more associative maps in between, topographical organization seems ubiquitous (Kaas, 1997). One reason for this topographical heterarchy might be that this effectively groups neurons on the basis of how frequently they interact, thus minimizing requirements for long, slow, and metabolically costly connections. This reasoning extends to the bridging of congruent and adjoining maps, which are also the ones that are likely to interact the most. A second reason may have to do with how topography lends itself to other, more local circuit motifs (e.g., lateral inhibition) from which efficient and predictive coding schemes may then naturally arise. Crucially, the benefit of codes that reduce redundancy is not just to minimize unnecessary neural activities. They effectively structure information such that, on the one hand, an internal model of the environment causing the past sensory inputs is built up, while on the other, the current sensory state is represented concisely, thus greatly simplifying the task of detecting meaningful contingencies in the environment (Barlow, 1961; R. P. Rao & Ballard, 1999) – two functions that one can loosely ascribe to cortical and basal ganglia circuits.

---

<sup>19</sup>The elephant in the room being the brain system of which the largest in the animal kingdom is ironically found in the elephant – the cerebellum. This system is all but lacking in the jawless lamprey, but is thought to have sprouted later from the dorsal hindbrain of jawed fish, as an adaptive filter for reafference (Bodznick et al., 1999).

### 1.1.5 Prediction

In isolation, the retinotectospinal circuits that are present in lamprey are capable of supporting spatial navigation in a stimulus-response manner – particular visual features, such as a looming stimulus, elicit particular responses, such as an oriented escape. However, implicit in the detection of a loom is, once again, a kind of temporal comparison, one that exploits the fact that when something approaches fast, its retinal fingerprint expands fast. Tacitly, the associated time horizon is small. A lamprey relying on these circuits alone may flee from a predator in one moment, but be otherwise unfazed by its lingering presence shortly after it has left the retina. One possible advantage provided on the one hand by the development of intratectal processing layers (Isa et al., 2021), but most noticeably by that of complementary retinotelencephalic circuits (Cisek, 2021), may have been that these endowed early vertebrates with the ability to exploit temporal regularities that were less local in time, by simple virtue of providing a deeper path for neural activity to loose itself in (Buonomano & Merzenich, 1995; Buonomano, 2003). This would have permitted inferring temporal contingencies on much longer timescales than what was previously possible, thereby freeing learning from the immediacy of the environment (Lau et al., 2017; Pezzulo et al., 2021).

Such a time-horizon-extending recurrent organization can already be seen in the lamprey subpallium, which contains many of the defining circuit-level features of the mammalian BG, including a striatum-like main input structure with two populations of (inhibitory) spiny projection neurons (SPNs) that use  $\gamma$ -aminobutyric acid (GABA) as a neurotransmitter. One subpopulation predominantly expresses D1-type dopamine receptors and projects directly to homologs of the globus pallidus internus (GPi) and substantia nigra pars reticulata (SNr), while the other mainly expresses D2 receptors and projects to those output nuclei indirectly via homologs of the globus pallidus externus (GPe) and subthalamic nucleus (STN) (Stephenson-Jones et al., 2012; Stephenson-Jones et al., 2011). STN projections are glutamatergic (excitatory), whereas GPi, SNr and most GPe projection neurons

are GABAergic<sup>20</sup>. Together, the *striatonigral* and *striatopallidal* circuits respectively make up the *direct* and *indirect* pathways of the BG, which in mammals are widely believed to respectively promote and suppress actions (Alexander et al., 1990; Cruz et al., 2022; Kravitz et al., 2010; Tai et al., 2012) via their opposing effects on BG output nuclei. Specifically, (in)direct pathway activity selectively (dis)inhibits GABAergic projection neurons in the GPi and SNr (Freeze et al., 2013), and thus (increases) decreases their tonic inhibition onto downstream thalamic and brainstem nuclei (Deniau & Chevalier, 1985; Roseberry et al., 2016). Given that this general blueprint also exists in the lamprey subpallium, antagonistic effects of the two pathways on its behavior seem all but inevitable, but (as far as I know) are yet to be demonstrated.

There exists at least one additional subdivision all across the striatum of lamprey that is highly conserved across vertebrates, one that creates a three-dimensional mosaic of histochemically defined compartments – *striosomes* (or patches) – within the broader striatal *matrix* (Graybiel & Ragsdale, 1978; Stephenson-Jones et al., 2013). Striatal projection neurons expressing D1- and D2-like dopamine receptors are uniformly distributed over the two divisions, which mostly differ in their downstream projection targets. Striosome SPNs project preferentially to (and thus inhibit) habenula-projecting neurons in the GPe and dopamine neurons in the SNc/VTA homolog, whereas matrix SPNs follow the canonical direct and indirect pathways. Because of this difference in projection patterns the two are posited to be differentially recruited in predicting when rewards will come, and what actions might increase later predictions (Joel et al., 2002).

Regardless of pathway and compartment, vertebrate SPNs in the striatum: rest at hyperpolarized membrane potentials owing to inward rectifying potassium channels (Ericsson et al., 2011; Uchimura & North, 1989); receive converging glutamatergic inputs from functional modules within the pallium and thalamus (Foster et al., 2021; Hunnicutt et al., 2016; Ocaña et al., 2015; Suryanarayana et al., 2020), and are densely innervated by broad DAergic arborizations coming from the SNc and VTA (or the posterior tuberculum in lamprey; Gerfen and Surmeier, 2011; Grillner and Robertson, 2016; Perez-Fernandez et al., 2014). These in turn get their input from the tectum, pallium

---

<sup>20</sup>A subset of tonically active glutamatergic GPe neurons projects specifically to a homolog of the lateral habenula, which in turn exerts an inhibitory influence on DAergic neurons in the posterior tuberculum – the lamprey version of the substantia nigra pars compacta (SNc) and ventral tegmental area (VTA) (Stephenson-Jones et al., 2013; von Twickel et al., 2019).

and striosome SPNs in the striatum. Moreover, in lamprey as in mammals, dopamine release has opposite effects on the excitability of the two SPN subtypes, increasing it for those expressing D1 receptors, and decreasing it for D2-receptor expressing neurons (Ericsson et al., 2013; Surmeier et al., 2007). And much like *C. elegans* and mammals, a subset of DA neurons in the lamprey also corelease glutamate (Sulzer et al., 1998; von Twickel et al., 2019; White, 1986), but whether or not phasic DA there also leads to long term synaptic potentiation (depression) at glutamatergic corticostriatal and thalamostriatal synapses onto D1 (D2) SPNs is not yet established (Iino et al., 2020; Kreitzer & Malenka, 2007; Reynolds et al., 2001; Shen et al., 2008). Demonstrating that this modulation of heterosynaptic plasticity by DA is already present at the (known) root of vertebrate phylogeny would have tremendous implications for the study of BG function, seeing as it is presumably through this mechanism that sparse reward signals teach striatal circuits about the arbitrary states and actions that led to those rewards (Xiong et al., 2015; Znamenskiy & Zador, 2013). Lastly, at the output nuclei of both lamprey and mammals, GPi and SNr neurons are tonically active at rest and project to motor regions in the thalamus and brainstem, including the tectum, which crucially, projects back to the striatum via the thalamus (Foster et al., 2021; McElvain et al., 2021; Stephenson-Jones et al., 2013).

One key difference in *topography* between the lamprey subpallium and the mammalian BG is that in the striatum of the former there is no strict subdivision between dorsal and ventral subdomains, instead all of it may be more akin to the dorsal striatum of mammals (Suryanarayana et al., 2021). A key *topological* difference, seems to be that in the lamprey there are no thalamocortical projections “closing” the cortico-BG-thalamo-cortico loops. The same seems generally true in sauropsids, though reciprocal palliothalamic projections are present in birds (presumably via homoplasy; Butler, 2008). Another disparity is the fact that GPe and GPi homologs do not form separate nuclei in lamprey, but are instead intermingled in the pallidum (Suryanarayana et al., 2021). It is also the case that this distinction is less clear in rodents than in primates, and that the GPi gets proportionally larger from monkeys, to rats, to mice (Steiner & Tseng, 2016). Together with the fact that unlike the SNr, the GPi does not target midbrain nor hindbrain locomotor regions (in lamprey, rodents and primates; Steiner and Tseng, 2016), all of these parallels and discrepancies speak to a primordial function of the thalamo-BG-thalamo loops in biasing downstream brainstem

structures, primarily or even exclusively through the SNr. The possibility to communicate those control signals back up to the pallium via reciprocal corticothalamic projections may have only developed later, possibly side by side with the expansion of the cerebellum and its projections to the thalamus, and both may have then been progressively elaborated on from mammals to primates (Cisek, 2021; Kuramoto et al., 2011; Montgomery, 2023).

The topographical, reentrant and opponent circuit architecture of the vertebrate subpallium appears as though tailored to enable an arbitrary mapping of states onto actions – a policy – that then manifests through selective disinhibition of phylogenetically older motor centres (Mink, 1996), where preexisting mappings are (comparatively) hardwired (Isa et al., 2021). In contrast to the ancient role of tonic hypothalamic DA in signalling average reward over time and biasing foraging decisions (as seen with food intake in *C. elegans*; T. T. Hills, 2006), phasic midbrain DA in vertebrates instead seems to signal unexpected deviations from that cached value – a reward prediction error (RPE) (Schultz et al., 1997). In principle, the striosome-DA and striosome-GPe-habenula pathways could provide a substrate for computing something akin to the time derivative of value that features in the computation of RPEs (Sutton & Barto, 2018). Such a teaching signal could then be used to reinforce adaptive policies via its asymmetric effect on matrix SPNs in the direct and indirect pathways (Joel et al., 2002), in light of temporally extended state information conveyed initially by BG-thalamus loops, and later by those involving cortex (Cisek, 2021). But crucially, the reigns of the actual moment-to-moment control remain largely with downstream systems such as the tectum and MLR. The pallium simply built on their egocentric, temporally myopic and essentially unimodal representations to catalyze learning. And the subpallium leveraged that to better predict the multi-term consequences of actions and chunk them into sequences (Graybiel, 1998). Adaptive behavior is so because it is predictive. The vertebrate telencephalon took this old idea and ran with it by expanding on its temporal horizon.

## 1.2 Inlook

This was quite the detour<sup>21</sup> to get to rats in IKEA® buckets judging intervals of time. Of all of the ideas introduced in this chapter, which I strongly believe to be of general relevance to the work presented in this thesis, the ones that are of more particular interest for framing its core results are that: behavior has never not been informed by time, spatiotemporal contingencies are arguably why it evolved in the first place, complex nervous systems “merely” made it more far- and multisighted; behavior is inherently heterarchical, and the result of control systems that likewise, are as nested, parallel, and recurrent as they are topographically organized.

To this evolutionary and conceptual backdrop, I will subsequently explore how the basal ganglia, a uniquely vertebrate solution, participates in interval timing, a uniquely vertebrate behavior. First, by using temperature to explicitly test the dissociable contributions of variability in the temporal scale, and baseline levels of neural activity in the rat dorsal striatum onto discrete decisions about time, and the continuous movements used to report them. Second, by leveraging inherent biases in interval discrimination due to stimulus intensity to further argue that multiple features in signals of the sensorimotor striatum point to them encoding the temporal bases that more generally, inform action selection in the context of reinforcement learning. And last, by trying to square these and other unpublished yet related results with seemingly compatible and conflicting ideas in the field about how endogenous variation in different aspects of striatal dynamics might be flexibly controlled to guide timing behavior in all of its scales.

---

<sup>21</sup> Honestly, I think it may have gotten away from me at times, but I hope it was not entirely useless for you (it certainly was not for me, even if not all of the utility found its way onto the page, or sticks with me once the dust settles).



## MANIPULATING STRIATAL TEMPERATURE DURING INTERVAL CATEGORIZATION

The basal ganglia are thought to contribute to decision-making and motor control. These functions are critically dependent on timing information, which can be extracted from the evolving state of neural populations in their main input structure, the striatum. However, it is debated whether striatal activity underlies latent, dynamic decision processes or kinematics of overt movement. Here, we measured the impact of temperature on striatal population activity and the behavior of rats, and compared the observed effects with neural activity and behavior collected in multiple versions of a temporal categorization task. Cooling caused dilation, and warming contraction, of both neural activity and patterns of judgment in time, mimicking endogenous decision-related variability in striatal activity. However, temperature did not similarly affect movement kinematics. These data provide compelling evidence that the timecourse of evolving striatal activity dictates the speed of a latent process that is used to guide choices, but not continuous motor control. More broadly, they establish temporal scaling of population activity as a likely neural basis for variability in timing behavior.

### 2.1 Introduction

Much of behavior is dependent on time. Humans and other animals must extract temporal structure from the environment to learn to anticipate events, to understand relationships between actions and consequences, and estimate

time, implicitly or explicitly, to plan and properly sequence and coordinate action. For tasks as varied as waiting at a stoplight to a hummingbird foraging for nectar, time is fundamental.

Timing mechanisms appear to be distributed across the nervous system, reflecting the importance of time information for much of brain function (Paton & Buonomano, 2018). However, one common requirement among diverse functions is the need to create an index, ordering and spacing information along the temporal dimension such that useful relations can be extracted and outputs appropriately coordinated. On the scale of seconds to minutes at which much of behavior unfolds, neuronal population dynamics represent a candidate means of both encoding and generating temporal patterns. Artificial neural network models have explored evolving population activity as a basis for timing sensory events (Buonomano & Merzenich, 1995) and movements (J. Wang et al., 2018). And correlations between behavior and the time course of population activity have lent some support to the hypothesis that time-varying patterns of activity within a population perform temporal computations (Bakhurin et al., 2017; Jazayeri & Shadlen, 2015; J. Kim et al., 2013; Merchant & Averbeck, 2017). However, correlations do not imply causation. The critical prediction of these “population clock” hypotheses is that interventions capable of slowing or speeding the temporal evolution of activity should cause a corresponding dilation or contraction of the temporal functions performed using that activity.

One brain system where time information appears to be critical is the basal ganglia (BG), an evolutionarily ancient set of brain structures thought to contribute to appropriate action selection based on experience. A dominant view holds that the BG embed core features of reinforcement learning (RL) algorithms (Doya, 1999). In mammals, inputs from a diverse set of territories in cortex, thalamus, and limbic brain structures convey information about the state of the world that converges with dense dopaminergic innervation in the major input area of the BG, the striatum. The input from dopamine neurons is thought to teach striatal circuits about the value of taking particular actions in a given state, information that can ultimately be conveyed to downstream brainstem and thalamo-cortical motor circuits to bias selection or otherwise specify features of actions (Alexander et al., 1990). To accomplish such a function, the BG would need access to information about ordering and spacing along the temporal dimension, either implicitly or explicitly, both to extract meaningful relations between the

environment, actions and outcomes that drive learning (Gallistel & Gibbon, 2000), and to coordinate the production of actions in time (Graybiel, 1998; Hidalgo-Balbuena et al., 2019). Interestingly, data from people with BG disorders (Freeman et al., 1996; Malapani et al., 1998) and human fMRI (Coull et al., 2004; S. M. Rao et al., 2001) have consistently identified the BG as being involved in timing behavior. In addition, lesions and pharmacological manipulations of the striatum can cause deficits in temporal estimation and reproduction (Meck, 2006; Mello et al., 2015). Lastly, recordings from striatal populations have demonstrated that time information can be decoded from neural activity, and this information correlates with variability in timing behavior (Bakhrin et al., 2017; Gouvêa et al., 2015; D. Z. Jin et al., 2009; J. Kim et al., 2018; Matell et al., 2003). Specifically, the state of striatal population activity continuously changes along reproducible trajectories during behavioral tasks that require time estimation, advancing more quickly when animals report long judgments, and more slowly when they report short judgments (Bakhrin et al., 2017; Gouvêa et al., 2015; Mello et al., 2015; J. Wang et al., 2018).

To test whether variability in the speed of BG population dynamics merely correlates with or directly regulates timing function, we sought to experimentally manipulate dynamics as animals reported temporal judgments. Interestingly, despite being composed of elements with differing temperature dependencies (e.g., ion channel conductances and synaptic transmission) (Robertson & Money, 2012), neural circuits in at least some systems can produce patterns of activity that systematically slow down or speed up with decreasing or increasing temperature (Hamaguchi et al., 2016; Tang et al., 2010). For this reason, temperature manipulations offer a potential method to test hypotheses regarding the relationship between the speed of neural dynamics and function (Banerjee et al., 2021). Indeed temperature manipulations in the zebra finch have been used to identify area HVC (proper name) as a locus within the song production circuit that contributes to the temporal patterning of bird song (Long & Fee, 2008). Similar temperature manipulations have identified a subregion of human motor cortex that regulates the speed of speech (Long et al., 2016), and a region of rodent medial frontal cortex that controls the timing of a delayed movement (Xu et al., 2014). However, temperature can have distinct effects on neural activity depending on cellular and circuit-level characteristics of the area in question (Long & Fee, 2008; Owen et al., 2019).

Here, we used a custom thermoelectric device (TED) to systematically vary the temperature of striatal tissue, both warming and cooling relative to a baseline condition. We found that temperature affected overall activity levels non-monotonically, with both warming and cooling relative to a control temperature producing lower baseline firing rates overall. In contrast, temperature manipulations caused monotonic and graded changes in the temporal scaling of neural activity, mimicking decision-related variability in activity observed during multiple versions of a temporal judgment task. Temperature manipulations also caused bidirectional and graded changes in animals' timing judgments, mirroring the temperature-dependent modification of temporal scaling of neural activity as well as the observed relationship between temporal scaling of activity and animals' judgments. Strikingly, these results were not accompanied by similar effects of temperature on movement execution. Instead, though more modest than the effects of temperature on animals' judgments, the pattern of average speed of animals' movements was a non-monotonic function of temperature, similar to the observed non-monotonic effect of temperature on baseline firing rate.

Together, these findings imply that distinct aspects of behavior may be controlled by temporal evolution of population activity and the overall levels of activity. Such distinctions suggest that continuously evolving patterns of activity in the striatum support the placement of discrete behavioral transitions in time, as required for action selection and decision-making, but do not provide moment by moment motor command-related signals required for continuous control. Instead of the BG providing high dimensional continuous control signals, overall level of activity appears to provide a low dimensional gain signal applied to motor programs executed elsewhere, consistent with the role of BG circuits in modulating movement vigor.

## 2.2 Results

### 2.2.1 Striatal temperature modified neural population speed monotonically and baseline activity levels nonmonotonically

Temperature has been shown in some systems to alter the speed of neural population activity while maintaining its general pattern (Hamaguchi et al., 2016; Tang et al., 2010). Thus, we first examined the effect of temperature on

striatal neural activity.

First, we developed a TED (Aronov & Fee, 2011) (Fig. 2.1) based on the Peltier effect and used it to achieve closed-loop control over the temperature of silver metal probes implanted in brain tissue.




Figure 2.1: Schematic of our implantable TED targeting the DS. The thick horizontal line splits the diagram into two differently scaled subregions (see 5 mm scale bars). Top region is subdivided in the middle between a front view (left half, with white strokes and gray fills) and a cutaway through the center of the implant (right half). A coronal slice of a rat's brain is shown in gray for reference (+0.84 mm AP, adapted from Paxinos and Watson, 2013).

To characterize the spatiotemporal profile of temperature changes in the brain, we measured temperature at different distances from the tip of a probe implanted in dorsal striatum (DS, Fig. 2.3), setting our TED to either a control level approximating normal body temperature, a warm condition or two levels of cooling (Fig. 2.2). We applied temperature manipulations using a block design of control-manipulation-control with transitions occurring at regular three-minute intervals (Fig. 2.2).

Manipulation temperatures were drawn at random and without replacement from the aforementioned set until its exhaustion, at which point the set was replenished and the sampling process resumed. We found that temperature near the tips of the probes tracked block changes, reaching asymptote within  $\sim 60$  s of transitions (Fig. 2.2, Fig. 2.3B, C), and that temperature changes fell off to minimal levels within 6 mm of the probe tip (Fig. 2.3E, F).

This *in vivo* characterization of the implant confirmed that temperature manipulations were largely localized to striatal tissue and that manipulation

## CHAPTER 2. MANIPULATING STRIATAL TEMPERATURE DURING INTERVAL CATEGORIZATION




Figure 2.2: Temperature manipulation protocol. Top: Target setpoint of the TED device over a representative session. Bottom: Representative segment of thermistor readings (black trace) at the lower plate of the TED module shown in Fig. 2.1. Gradient-colored trace indicates the temperature estimated from in vivo calibration experiments (see Fig. 2.3).

blocks of three-minute duration would allow for assessing effects of striatal temperature on neural activity and behavior. Next we characterized the features of neural activity that were affected by temperature, adapting a paradigm for optogenetically inducing patterns of striatal activity under anesthesia (Hidalgo-Balbuena et al., 2019).

This approach allowed us to generate large numbers of highly reproducible bouts of population activity free from neural variability directly related to ongoing behavior, and was thus well-suited for studying the direct effects of temperature on striatal activity. Briefly, we used a viral strategy to express Channelrhodopsin-2 (ChR-2) in the VB (Fig. 2.4A), a somatosensory thalamic area that projects directly to the striatum, and indirectly may influence striatal activity through thalamocortical projections (see Supplementary Information Text). Three weeks post-infection, under urethane anesthesia, we implanted an optic fiber over VB, a single insulated silver probe of our TED into DS and an adjacent Neuropixels (Jun et al., 2017) silicon probe (Fig. 2.4A). We then delivered 50-ms trains of blue light pulses at 100 Hz once every 1.5 s and recorded DS neural activity (Fig. 2.4B, C). Stimulation caused a brief volley of striatal activity, followed shortly thereafter by reproducible patterns of firing across the population over hundreds of milliseconds after




Figure 2.3: Spatiotemporal characterization of our TED. **(A)** Schematic of the preparation in which we set our TED to one of several manipulation temperatures ( $T = \{5, 15, 20, 25, 30, 45\}$  °C) while measuring temperature at its lower plate and probe tip simultaneously. **(B)** Temperature measured at the TED plate and probe tip thermistors. **(C)** Temperature traces measured at the probe tip thermistor during manipulation blocks aligned to block transitions. Solid lines represent model fits. **(D)** Schematic of the preparation in which we set our TED to one of several manipulation temperatures ( $T = \{15, 25, 42, 45\}$  °C) while measuring temperature at its lower plate a moveable temperature probe simultaneously. **(E)** Decay parameters for models fit to manipulation temperatures (as shown in (C)) across the 2 experiments (manipulation temperatures that are common to both experiments are connected with solid colored lines). **(F)** Same as (E), but for the gain / offset parameter across all model fits.

## CHAPTER 2. MANIPULATING STRIATAL TEMPERATURE DURING INTERVAL CATEGORIZATION




Figure 2.4: (A) Schematic of the preparation used to elicit, record and manipulate striatal activity dynamics under different temperatures. (B) VB thalamic stimulation protocol overlaid with five single-trial examples of evoked striatal activity recorded under the control temperature condition. Blue ticks depict light pulses. (C) Normalized SDFs of recorded striatal neurons ( $N = 335$ , before enforcing the minimum firing rate selection criterion, see Methods), under the control temperature. Units were ordered by their angular position in the subspace defined by the first two PCs, describing dynamics between 150 ms and 750 ms from stimulation onset.

the last light pulse had been delivered (Fig. 2.4C). For assessing the effect of temperature on neural responses we focused on this longer lasting activity because, unlike the initial volley, later responses reflected autonomous dynamics of the system as it settles as opposed to the initial direct response to stimulation.

While the general patterning of individual-neuron responses over time was maintained across different temperatures, the time course of this pattern systematically varied depending on temperature, advancing more slowly the colder the temperature, and more quickly when temperature was raised above baseline (Fig. 2.5).

To quantify temperature-dependent warping in the time course of neural responses, we computed a scaling factor for each neuron in each of the four temperature conditions (see Methods, Fig. 2.6A).

Temporal scaling as opposed to shifting of responses provided a significantly better explanation of the effect of temperature on firing rates across the population (Fig. 2.6, 2.7).

Across all recorded neurons, distributions of response dilation were ordered inversely as a function of temperature (Fig. 2.8A, B), with the majority of cells exhibiting time-contracted firing profiles (dilation  $< 0\%$ ) under warming, and time-dilated firing profiles (dilation  $> 0\%$ ) under cooling conditions.




Figure 2.5: Activity of three putative striatal units (arrowheads in Fig. 2.4C.) aligned to the onset of VB stimulation. Top: SDFs split by temperature (mean  $\pm$  SEM). Bottom: Spike raster plots, wherein trials are ordered chronologically. The superimposed gradient-colored trace indicates temperature setpoint over the course of the recording session.

At a population level this led to systematic differences in the speed with which population state advanced along its typical trajectory in PC state space (Fig. 2.8C).

Temperature also modified baseline firing rate of striatal neurons, however, this effect was distinct from the effect on temporal scaling. Whereas temperature produced a monotonic effect on neural temporal scaling across the sampled temperatures (Fig. 2.5, Fig. 2.8), baseline firing rates varied non-monotonically as a function of temperature, with both warming and cooling of striatal tissue relative to a physiologically normal control value resulting in lower firing rates (Fig. 2.9).

Previous studies have shown that estimates of elapsed time decoded from striatal populations can predict timing judgments (Gouvêa et al., 2015; J. Kim et al., 2018). To assess whether temperature effects resembled endogenous, decision-related variability in population activity during behavior and its impact on the readout of decision variables, we decoded elapsed time from the population under different temperatures. Briefly, we first characterized the “typical” temporal profiles of striatal responses using a subset of control trials, and then applied a probabilistic decoding approach to estimate elapsed time based only on the observed state of the recorded population in remaining




Figure 2.6: Temporal scaling as opposed to shifting provided a better account of temperature’s effect on neural activity. (A) Top: Simulated spike density functions exhibiting bidirectional and dose-dependent temporal scaling with temperature. Middle: Templates built by warping a control spike density function (thicker black line) in time by scale factors ranging from 0.625 (maximum contraction) to 1.75 (maximum dilation). Note that when applying this method to data, this control response is not the same as the one shown in the top panel, as the two are built using two non-overlapping random sets of control trials. Bottom: Thick lines represent the coefficient of determination ( $R^2$ ) for all scaled templates in the middle panel regressed against each of the target spike density functions shown at the top. We computed this objective function for each neuron-temperature condition pair and took its global maximum as the corresponding temporal scaling factor, highlighted here by the larger markers. Thinner lines and smaller markers depict  $R^2$  values for a similar regression procedure applied to a series of shifted, as opposed to scaled, templates. (B) Same as (A), except that for artificially temporally shifted responses relative to control (top), temporally shifted templates (middle), and their regression outcomes (bottom). The thinner lines and smaller markers respectively represent the  $R^2$  curves and maxima resulting from regressing the scaled templates from A (middle) against the shifted targets in B (top). Conversely, the result of regressing shifted templates against scaled targets is plotted in the same manner in (A, bottom).




Figure 2.7: **(A)** To assess whether the effects of temperature on individual striatal responses were better accounted for by temporal scaling or shifting, we built two separate spiking models in which we either injected one effect or the other. Briefly, we modeled 500 control firing rate functions as Gaussian bumps defined over 1.5 s with means spanning the interval from 150 ms to 750 ms (Fig. 1e) and a standard deviation of 50 ms. The amplitudes of the resulting probability density functions were rescaled so that their distribution of mean firing rates matched that of striatal data. Next, we created one additional rate function per neuron per manipulation condition by either shifting or scaling its control response in time. Again, the distribution of generative temporal scaling factors and shifts used was informed by the empirical distributions of these metrics extracted from striatal data. We then generated 150 spike trains of each condition per neuron by sampling spike times from inhomogeneous Poisson point processes with the aforementioned condition-specific responses as their time-dependent rate parameters. From this point on, we proceeded to analyze the resulting surrogate spike data in the exact same way we did for the striatal data, by first averaging trials within condition, generating libraries of templates and then computing temporal scaling factors and shifts. Finally, for each ‘neuron’-condition pair within each model, we stored the  $R^2$  values corresponding to the best-matching scaled and shifted templates and subtracted the former from the latter to build the distributions shown here at the top (scaling model) and middle (shifting model) panels. Thick solid sigmoidal lines represent the CDFs of each condition’s  $R^2$  difference. Thin vertical black lines denote control mean differences. Small horizontal colored lines link the respective means of the corresponding manipulation and control distributions. **(B)** Same as (A), but for striatal data.

## CHAPTER 2. MANIPULATING STRIATAL TEMPERATURE DURING INTERVAL CATEGORIZATION




Figure 2.8: Temperature rescaled striatal population activity in time leading to graded monotonic modulation of decoded time estimates. **(A)** Distribution of percentage change in temporal scale relative to control dilation ( $N = 268$ , after enforcing the minimum firing rate selection criterion, see Methods). Histograms (top) and corresponding cumulative density functions (CDF, bottom) of dilation conditioned on manipulation temperature (same color scheme as in Fig. 1). **(B)** Left: Distributions of neuronal response dilations as a function of striatal temperature change ( $N = 175$ , after excluding unreliable scaling factor estimates, see Methods, one-sample two-tailed  $t$ -tests, all significant  $|t_s(174)| = [4.70, 9.37]$ ,  $P = [3.84e-17, 5.28e-6]$ , non-significant  $t(174) = 1.16$ ,  $P = 0.25$ ; repeated measures ANOVA followed by post-hoc contrasts with Tukey correction for multiple comparisons,  $F(3, 696) = 51.29$ ,  $P = 5.88e-30$ ; smallest  $q_s(348) = 2.96$ ,  $P = 0.016$ ). Markers represent temperature-split response dilations for each cell. Boxplots show population medians (horizontal thick lines) and IQR (colored bars). Right: Distribution of stretch in neuronal responses (one-sample two-tailed  $t$ -test,  $t(174) = -10.15$ ,  $P = 2.70e-19$ ). **(C)** Average control population activity projected onto its first 3 PCs (solid black line), and temperature-split averages (colored markers) projected onto that control trajectory at the time point indicated by the last arrowhead in (E) and (F). Ghosted versions of these visualizations show their 2D projections onto all possible combinations of PCs 1, 2 and 3. The dashed black arrow indicates the direction of time.

trials (see Methods). Estimates of elapsed time derived from ongoing population activity systematically led ahead and lagged behind true time during warming and cooling blocks, respectively. This can be observed by sampling the output of the decoder at discrete delays from stimulation onset (Fig. 2.10), and more continuously by subtracting decoder output during control blocks from that during the three different manipulation conditions (Fig. 2.11). Relative to control, cooler temperatures gradually shifted decoded estimates of time earlier, and the warmer temperature shifted decoded estimates of time later.

These data demonstrate that, under anesthesia, the temporal scaling of




Figure 2.9: Distributions of changes in baseline firing rate as a function of striatal temperature change (one-sample two-tailed  $t$ -test all significant  $|t_s(267)| = [5.24, 15.01]$ ,  $P = [3.38e-37, 3.31e-7]$ , non significant  $t(267) = 0.18$ ,  $P = 0.86$ ; repeated measures ANOVA followed by post-hoc contrasts with Tukey correction for multiple comparisons,  $F(3, 1068) = 72.29$ ,  $P = 1.44e-42$ ; smallest  $t_s(534) = 4.02$ ,  $P = 3.40e-4$ ). Markers represent temperature-split change in baseline firing rate for each cell. Boxplots show population medians (horizontal thick lines) and IQR (colored bars).




Figure 2.10: Decoded posterior probability of time given the state of striatal populations at four arbitrary time points, averaged separately across trials for each temperature condition (mean  $\pm$  SEM).




Figure 2.11: Difference between trial-averaged decoded estimates for manipulation (left: extreme cooling; middle: mild cooling; right: warming) and control conditions, shown for the time window highlighted in Fig. 1E. Grayscale arrowheads indicate time points used in 2.10. The identity line is shown in dashed black.

neural response profiles by temperature can reproduce features of endogenous variability in the time course of population activity and time encoding shown in previous studies to correlate with timing judgments.

### 2.2.2 Striatal population speed predicts temporal judgments across tasks with different immobility requirements

A number of studies have observed that ongoing behavior may correlate with timing performance (De Kock et al., 2021; Gouvêa et al., 2014; Kawai et al., 2015; Killeen & Fetterman, 1988; Machado et al., 2009; Safaie et al., 2020). This raises the possibility that previously observed correlations between the speed of neural population activity and timing judgments might reflect the encoding of kinematic features by striatal populations (Dhawale et al., 2021; Rueda-Orozco & Robbe, 2015) combined with the adoption of embodied strategies for timing (Killeen & Fetterman, 1988). Such a scenario may indicate a more indirect role for striatal population dynamics in the decision process. To examine this possibility, we studied DS population dynamics during behavior in two versions of a temporal categorization task (Fig. 2.12) that differed in the degree to which animals were free to move around the behavioral box during interval estimation.

We trained rats, in both versions of the task, to report intervals of time as either shorter or longer than a 1.5-s category boundary (Fig. 2.12). Briefly, rats were placed in a rectangular behavioral box with three nose ports positioned at head level. Trials began with an auditory tone triggered by the subjects'




Figure 2.12: Schematic of behavioral setup (left) and event diagram illustrating a correct trial (right) in both the no-fixation (orange) and fixation (blue-green) versions of the interval categorization task.

entry into the central “initiation” nose port. After an interval of silence during which subjects were either free to move about the box (no-fixation) (Gouvêa et al., 2015) or required to maintain their snout positioned in the central port (fixation) (Cruz et al., 2022), a brief second tone was delivered. This second tone acted both as stimulus offset, defining the duration of the interval subjects were asked to judge, and a “go” cue, freeing subjects to report their choice at one of two equidistant ports on either side of the initiation port. Choices reported at one of the lateral nose ports following short stimuli (<1.5 s) and at the opposite lateral nose port after long stimuli (>1.5 s) were defined as “correct” and resulted in delivery of a water reward at the choice port. “Incorrect” and premature choices (in the no-fixation version of the task) or premature departures from the central port (in the fixation version), were punished with an error tone and a time penalty added to the inter-trial-onset interval.

Accuracy of timing judgments (Fig. 2.13) and several key features of the recorded striatal population activity in relation to behavior were similar in the two paradigms.

First, in both task variants, striatal population state continuously evolved through a set of non-repeating states during interval presentation (Fig. 2.14).

Furthermore, the time course of this evolution was systematically related to animals’ judgments, a relationship we examined using several approaches that facilitated comparison with the effect of temperature on neural activity described above.

Given previous observations that the time course of striatal responses




Figure 2.13: (A) Psychometric performance in the no-fixation version of the task ( $N = 3$  rats, mean  $\pm$  SEM). (B) Same as (A), for the fixation task variant ( $N = 3$  rats).




Figure 2.14: (A) Normalized SDFs of all striatal neurons ( $N = 441$ ) recorded during presentations of the longest stimulus in our set (2.4 s) in the no-fixation task variant. Units are ordered by their angular position in the subspace defined by the first two PCs describing firing dynamics. (B) Same as (A), for the fixation task variant ( $N = 483$  neurons).

covaries with duration judgments, conditioning trials of a given interval on animals' judgments should result in distributions of neural population responses with systematically different temporal scaling (Fig. 2.15). To test this prediction, we focused on intervals where long and short judgments were most balanced, corresponding to the two intervals nearest to 1.5 s. Because rats correctly categorized the majority of trials, even for these "difficult" stimuli, we expect the majority of any hypothetical distribution of population speeds across trials to result in population states at interval offset that lie on the "correct" side of the animal's internal representation of the decision boundary. Thus, conditioning neural data on whether the animal judged 1.38 s as "long" or "short" should create two distributions of trials with respect to population speed: one with an average speed that is slightly slower than the overall average in the case of correct, "short" judgments (Fig. 2.15, right panel, top left), and another with an average that is shifted to an even larger extent above the overall average toward faster speeds in the case of incorrect, "long" judgments (Fig. 2.15, right panel, top right). Conversely, conditioning neural data on judgment of the 1.62-s interval should create one distribution of trials with an average speed that is slightly faster than the overall average in the case of correct, "long" judgments (Fig. 2.15, right panel, bottom right), and another with an average that is shifted to an even larger extent below the overall average toward slower speeds in the case of incorrect, "short" judgments (Fig. 2.15, right panel, bottom left).

Using the temporal scaling metric applied to recordings under anesthesia above, we found that response dilation of individual neurons systematically varied in accordance with these predictions (Fig. 2.16). As with manipulation of temperature, this led to systematic differences in the speed with which population state advanced along its typical trajectory in neural state space (Fig. 2.17), and systematic differences in the speed of decoded time estimates derived from the population (Fig. 2.18).

Lastly, to gain a sense of how population speed might relate to choices across a broader range of stimuli, instead of conditioning neural activity on behavior, we conditioned behavior on the speed with which neural population activity evolved along its typical trajectory during the interval period in the two tasks. In both behavioral scenarios, rats were systematically biased toward making "short" judgments the more slowly the population state progressed on single trials, particularly for stimuli closer to the 1.5-s decision boundary, leading to a monotonic relationship between population

## CHAPTER 2. MANIPULATING STRIATAL TEMPERATURE DURING INTERVAL CATEGORIZATION




Figure 2.15: Left: Schematic of hypothetical relationship between trial to trial variability in the speed of striatal dynamics, blue indicating slower, and red indicating faster, and the resulting distributions of internal time estimates at each interval offset. Assuming a fixed internal decision boundary, long or short judgments are colored as black or gray portions of the distributions. Right: Schematic depicting hypothetical, graded shifts in average internal timing speeds when conditioning on choice for the two intervals nearest to the decision boundary. Specifically, short (long) stimuli incorrectly categorized as long (short) should reflect pronounced biases towards faster (slower) than usual dynamics, and correctly categorized short (long) stimuli should reflect milder biases towards relatively slower (faster) dynamics.




Figure 2.16: (A) Left: Distributions of neuronal response dilations for each near-boundary stimulus-choice pair introduced in Fig. 2.15. Markers represent response dilations for each cell and stimulus-choice condition. Boxplots show population medians (horizontal thick lines) and IQR (colored bars). Right: Distribution of stretch in neuronal responses (one-sample two-tailed t-test,  $t(440) = -7.21$ ,  $P = 2.45e-12$ ). (B) Same as (A), for the fixation task variant (one-sample two-tailed t-test,  $t(482) = -5.28$ ,  $P = 1.93e-7$ ).




Figure 2.17: **(A)** Average population activity from 2.14 projected onto its first three PCs (solid black line), and averages (colored markers) from each stimulus-choice pair projected onto that reference trajectory at the time of stimulus offset for the short near-boundary stimulus (1.38 s). Ghosted versions of these visualizations show their 2D projections onto all possible combinations of PCs 1, 2 and 3. The dashed black arrow indicates the direction of time. **(B)** Same as **(A)**, for the fixation task variant.




Figure 2.18: **(A)** Average decoded MAP estimates for each stimulus-choice condition defined in 2.15 for the no-fixation task version. Colored solid lines and patches represent medians and IQR across 512 concatenations (see Methods). **(B)** Same as **(A)** for the fixation task variant.

## CHAPTER 2. MANIPULATING STRIATAL TEMPERATURE DURING INTERVAL CATEGORIZATION

speed and psychometric threshold (Fig. 2.19) – the stimulus duration at which the sigmoid crosses 50% probability of the animal making either a long or short choice.




Figure 2.19: (A) Performance in the no fixation version of the interval categorization task conditioned on striatal population speed ( $N = 3$  rats, see Methods): fast (red), typical (black), slow (dark blue) or slowest (light blue). Psychometric curves fit to cross-animal averages of psychophysical data split by neural speed, respectively shown as solid lines and markers of matching color (mean  $\pm$  SEM). Bottom-right inset: Differences in proportion of long choices from each atypical speed condition to the typical speed condition (mean  $\pm$  propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each speed condition’s psychometric fit. Solid black lines represent MAP point estimates. (B) Same as (A), but for the fixation task variant ( $N = 3$  rats).

Thus, variability in timing judgments was associated with the same features of neural activity that were modified by temperature in the first set of experiments. Furthermore, this was true whether or not subjects were required to maintain their position in the initiation port during the interval stimuli, suggesting that decision-related variability in the time course of striatal responses did not result from encoding of detailed kinematic information alone. Next, having established that temperature provides a means of manipulating decision-related variability in striatal activity, we tested whether and how striatal temperature modified subjects’ decisions, movement parameters, or both.

### 2.2.3 Striatal temperature produced monotonic effects on temporal judgments

We implanted two additional groups of rats previously trained in either the version of the interval categorization task: not requiring ( $N = 4$ ) or requiring fixation ( $N = 7$ ), with a TED targeting its probe tips to DS. Critically, the no-fixation cohort and one rat in the fixation cohort were implanted with a predecessor of the TED we described above, which was unsuited for prolonged extreme cooling and incapable of warming. As such, for those animals temperature manipulations were restricted to a single mild cooling level.

In advance of temperature manipulations, and for both task variants, rats' performance was virtually perfect for easy stimuli, progressively more variable as stimuli approached the categorical boundary and well described by a sigmoid psychometric function with a threshold close to the experimentally imposed decision boundary of 1.5 s (Fig. 2.20).

Strikingly, at the onset of temperature manipulations, all eleven subjects across the two task versions displayed a systematic shift towards short judgments when the striatum was cooled (Fig. 2.21, 2.22, see Supplementary Information Text). Furthermore, subjects implanted with the latest TED version (Fig. 2.1) exhibited bidirectional and monotonic changes in their categorization behavior as a function of temperature: rats were more likely to report short judgments during cooling blocks, and long judgments during warm blocks (Fig. 2.21, 2.23), particularly for intervals nearer to the 1.5-s categorical boundary (Fig. 2.21). Importantly, the larger the magnitude of the cooling manipulation, the larger the change in choice behavior. The systematic changes in the subjects' judgments caused by temperature were most reliably captured by shifts in the threshold parameter of the psychometric function (Fig. 2.21, top left inset).

Thresholds tracked differences between control and manipulation temperatures in both sign and magnitude for all individual animals (Fig. 2.23).

Additionally, the circuit mechanism underlying the behavioral effects of temperature did not seem to involve overlying primary motor cortex (M1), through which the insulated portion of the probes passed, because direct manipulation of M1 temperature in an additional set of four rats performing the fixation version of the task produced significantly smaller effects on choice behavior (Fig. 2.24).

## CHAPTER 2. MANIPULATING STRIATAL TEMPERATURE DURING INTERVAL CATEGORIZATION




Figure 2.20: Performance in both task variants was qualitatively similar before and after TED implantation. (A) Performance of rats trained on the no-fixation variant of the interval categorization task and implanted with the initial version of our custom TED targeting the DS ( $N = 4$ ) on the last day of training before surgery. Gray lines are psychometric fits to individual animals. The black line is a fit to the average across animals. The underlying choice data is shown for the cohort average (mean  $\pm$  SEM). (B) Same as (A), but for rats trained on the fixation task variant and implanted with the initial ( $N = 1$ ) or final version of our custom TED targeting the DS ( $N = 6$ ). (C) Schematized timeline for a typical experimental rat. Elongated bars illustrate long periods of time, whereas squares represent individual daily sessions, which, with the exception of surgeries, lasted for 2 hours. (D-E) Same as (A-B), but for the last post-surgery training session before the first temperature manipulation session.




Figure 2.21: Performance of rats implanted with the final version of our custom TED targeting the DS in the fixation version of the interval categorization task conditioned on striatal temperature ( $N = 6$  rats): warm (red), control (black), cool (teal) or coolest (cyan). Psychometric curves fit to cross-animal averages of psychophysical data split by temperature, respectively shown as solid lines and markers of matching color (mean  $\pm$  SEM). Bottom-right inset: Differences in proportion of long choices from each experimental condition to the control condition (mean  $\pm$  propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each temperature condition's psychometric fit. Solid black lines represent MAP point estimates. Within-animal data were pooled across the first two sessions of temperature manipulations, with the first one contributing the most (see Methods).

Moreover, the time course of any trend of an effect appeared delayed relative to that of the effect on the DS cohort (Fig. 2.25), consistent with volume conduction of cortical temperature manipulations down to the striatum (Fig. 2.3).

Thus, striatal temperature manipulations, shown above to produce changes in the population response under anesthesia that mimic decision-related temporal scaling of firing patterns (Fig. 2.5, 2.8, 2.10, 2.11), caused highly reproducible, parametric variation in a decision variable used by rats to guide duration judgments during the task. As in the recordings during behavior (Fig. 2.13, 2.14, 2.16, 2.17, 2.18, 2.19), this effect was robust to differences in task design with respect to the degree that animals were free to move during interval presentation (Fig. 2.12). These data suggest that the systematic changes in timing judgments induced by temperature in the striatum are thus due to the effect of temperature on the temporal scaling of neuronal responses within populations of neurons that contribute to the




Figure 2.22: (A) Performance of rats implanted with the initial version of our custom TED targeting the DS in the no fixation version of the interval categorization task conditioned on striatal temperature ( $N = 4$  rats): control (black), cool (teal). Psychometric curves fit to cross-animal averages of psychophysical data split by temperature, respectively shown as solid lines and markers of matching color (mean  $\pm$  SEM). Bottom-right inset: Differences in proportion of long choices from the cool to the control condition (mean  $\pm$  propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each temperature condition's psychometric fit. Solid black lines represent MAP point estimates. (B) Same as (A), but for the one rat trained in the fixation task variant and implanted with the initial version of our TED.

decision.

#### 2.2.4 Striatal temperature produced nonmonotonic effects on movement kinematics

What can these results tell us about the nature of the functional involvement of striatal circuits in behavior? One of the most commonly observed effects of dysfunction in BG circuits is a change in the speed of movement, such as bradykinesia in the case of Parkinson's disease. Indeed, inhibiting direct pathway striatal projection neurons has been shown to slow movement without affecting other aspects of its execution (Cruz et al., 2022; Panigrahi et al., 2015). Since the anesthetized recordings described above revealed qualitatively distinct effects of striatal temperature on population speed (monotonic), and baseline activity levels (non-monotonic), we wondered whether there were any correlates of either or both effects to be found in measures of animals' movements. We focused on fixation sessions, where




Figure 2.23: Effect of temperature on psychophysical thresholds. (A) Markers represent MAP estimates and transparent patches the corresponding 95% confidence intervals of threshold parameters fit to individual animals' performance on control (vertical axis) versus manipulation blocks (horizontal axis). Each animal contributes one data point of each color. The identity line is plotted as a diagonal line, with its thick solid portion highlighting the region of the main axes that is shown in the bottom-right inset axes. Top-right inset: Distribution of threshold differences between manipulation and control conditions (mean  $\pm$  SEM). Bottom-right inset: Same as main axes, but with data pooled from two pilot experiments using a single cooling dose in either the fixation ( $N = 1$ ) or the no-fixation ( $N = 4$ ) task variants. (B) Distributions of threshold dilation as a function of induced temperature changes (one-sample two-tailed  $t$ -tests, all significant  $|t_s(5)| = [3.25, 5.17]$ ,  $P = [4.01e-3, 2.31e-2]$ ; repeated measures ANOVA followed by post-hoc contrasts with Tukey correction for multiple comparisons,  $F(2, 15) = 17.30$ ,  $P = 1.22e-4$ ; smallest  $q_s(10) = 4.45$ ,  $P = 1.52e-2$ ). Markers represent individual animals. Boxplots show animal means (horizontal thick lines) and SEM (colored bars).

both warming and cooling were applied, as bidirectional manipulation of temperature was observed under anesthesia to differentiate temperature's effects on temporal scaling and baseline firing rate. We first tracked the position of the TED in video using a markerless tracking algorithm (Mathis et al., 2018) (Fig. 2.26).

We then asked how temperature affected position over time either during the interval period (Fig. 2.27), or during choice movements (Fig. 2.28). We did not observe any significant monotonic effect of temperature in any case. Rats were positioned similarly during the interval period (Fig. 2.27A), and followed similar paths from the fixation port to the choice ports across the different temperatures (Fig. 2.28).

## CHAPTER 2. MANIPULATING STRIATAL TEMPERATURE DURING INTERVAL CATEGORIZATION




Figure 2.24: Manipulating M1 temperature did not produce discernible effects on timing judgments. **(A)** Average categorization performance in the fixation task version at the onset of M1 temperature manipulations. Psychometric functions fit to cross-animal averages ( $N = 4$ ) of temperature-split psychophysical data, respectively shown as solid lines and markers of matching color (mean  $\pm$  SEM). Bottom right inset: Average differences in proportion of long choices from each manipulation condition to control (mean  $\pm$  propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each condition's psychometric fit. Solid black lines represent the MAP point estimates implicit in the fits shown in the main axes. **(B)** Animal-split categorization behavior. Markers represent MAP estimates and transparent patches the corresponding 95% confidence intervals of threshold parameters fit to individual animals' performance on control (vertical axis) versus manipulation blocks (horizontal axis). Inset: Distribution of threshold differences between manipulation and control conditions. Markers represent individual animal differences (mean  $\pm$  SEM). **(C)** Left: Distributions of threshold dilation as a function of induced temperature changes in M1 (all non-significant  $|t_s(3)| = [0.93, 2.90]$ ,  $P = [0.06, 0.42]$ ; repeated measures ANOVA,  $F(2, 9) = 1.93$ ,  $P = 0.20$ ). Markers represent individual animals. Boxplots show animal means (horizontal thick lines) and SEM (colored bars). Right: Distribution of threshold stretch split by manipulation target (DS:  $N = 6$ ; M1:  $N = 4$ , see Methods). Markers represent individual animals, and their size and color denote bootstrapped significance (see Methods, one-sample two-tailed  $t$ -tests,  $t_{DS}(5) = -4.37$ ,  $P = 7.19\text{e-}3$ ;  $t_{M1}(3) = -1.79$ ,  $P = 0.17$ ; two-sample two-tailed  $t$ -test,  $t(8) = -2.62$ ,  $P = 3.14\text{e-}2$ ). Boxplots show animal means (horizontal thick lines) and SEM (colored bars).




Figure 2.25: Threshold dynamics aligned to and across block transitions early during DS and M1 temperature manipulations (left and right, respectively). Condition-split cross-animal average thresholds (mean  $\pm$  SEM) were computed using trials that fell into a sliding window lasting 90 s (half the block duration) that was swept from the preceding to the succeeding control blocks in increments of 9 s. Each marker corresponds to one sweep, and its color shading denotes the fraction of that sweep's window that was inside a control block (with black markers corresponding to 100% control trials), and by extension its complement that fell in a manipulation block (with pure manipulation colors corresponding to 100% manipulation trials).




Figure 2.26: (A) Illustration of three trial epochs: stimulus offset (left), reaction (middle) and choice (right). (B) 2D trajectories of implant position aligned to reaction from a representative example DS-implemented animal. Dashed (solid) lines correspond to short (long) choices with individual control trials ghosted in the background, and condition-split averages on top (mean  $\pm$  SEM). Filled and open markers show average position at reaction and choice, respectively.

## CHAPTER 2. MANIPULATING STRIATAL TEMPERATURE DURING INTERVAL CATEGORIZATION




Figure 2.27: (A) 2D trajectories of implant position during stimulus presentation from a representative example DS-implanted animal. Dashed (solid) lines correspond to short (long) choices with individual control trials ghosted in the background, and condition-split averages on top (mean  $\pm$  SEM). (B) Same as (A), but split into x (top) and y (bottom) coordinates and plotted all the way up to reaction and choice. Filled and open markers show average position at reaction and choice, respectively. (C) Same data as in (A) and the stimulus presentation portion of (B), combined into an overall speed metric (see Methods).




Figure 2.28: (A) 2D trajectories of implant position aligned to reaction from a representative example DS-implanted animal. Dashed (solid) lines correspond to short (long) choices with individual control trials ghosted in the background, and condition-split averages on top (mean  $\pm$  SEM). Filled and open markers show average position at reaction and choice, respectively. (B) Same as (A), split into x (top) and y (bottom) coordinates. (C) Same as (B), combined into an overall speed metric.

In addition, the speed profiles with which rats executed their choices were not scaled monotonically in time as a function of temperature (Fig. 2.29).




Figure 2.29: (A) Effect of temperature on the temporal scale of speed profiles recorded during the stimulus period. Left: Implant speed profile dilation for DS animals (one-sample two-tailed t-tests, non-significant  $t_5(5) = [0.44, 2.02]$ ,  $P = [0.10, 0.68]$ ; repeated measures ANOVA,  $F(2, 15) = 2.07$ ,  $P = 0.16$ ). Right: Implant speed stretch for DS ( $N = 6$ ) and M1 ( $N = 4$ ) animals (one-sample two-tailed t-tests,  $t_{DS}(5) = -2.02$ ,  $P = 0.10$ ,  $t_{M1}(3) = -0.83$ ,  $P = 0.47$ ; two-sample two-tailed t-test,  $t(8) = -0.65$ ,  $P = 0.53$ ). (B) Same as (A) for dilation and stretch of speed profiles recorded between reaction and choice. Left: Implant speed profile dilation for DS animals (one-sample two-tailed t-tests,  $t(5) = [0.11, 0.87]$ ,  $P = [0.43, 0.92]$ ; repeated measures ANOVA,  $F(2, 15) = 0.11$ ,  $P = 0.89$ ). Right: Implant speed stretch for DS ( $N = 6$ ) and M1 ( $N = 4$ ) animals (one-sample two-tailed t-tests,  $t_{DS}(5) = -0.07$ ,  $P = 0.95$ ;  $t_{M1}(3) = -0.76$ ,  $P = 0.51$ ; two-sample two-tailed t-test,  $t(8) = 0.56$ ,  $P = 0.59$ ). Markers represent individual animals, and their size and color denote bootstrapped significance.

However, we did observe a modest non-monotonic effect in average movement speeds as a function of temperature (Fig. 2.30, left), as evidenced by absolute temperature change being a better predictor of average speed over animals than signed temperature change (Fig. 2.30, right). This was in stark contrast to the clearly monotonic effect of temperature on decisions (Fig. 2.23) and temporal scaling of population activity (Fig. 2.8).

Next we analyzed animals' time to react to stimulus offset and the time taken to move to the choice port. We did not observe significant differences in either reaction times or movement times depending on temperature across rats (Fig. 2.31), however, we again observed a trend wherein both warming and cooling appeared to delay movement initiation following stimulus offset

## CHAPTER 2. MANIPULATING STRIATAL TEMPERATURE DURING INTERVAL CATEGORIZATION




Figure 2.30: Left: Same as in (E, left) for average implant speed dilation of DS animals (one-sample two-tailed t-tests, significant  $t_s(5) = [-3.64; -2.57]$ ,  $P = [0.05; 0.02]$ , non-significant  $t(5) = 0.57$ ,  $P = 0.59$ ; repeated measures ANOVA followed by post-hoc contrasts with Tukey correction for multiple comparisons,  $F(2, 15) = 4.57$ ,  $P = 0.03$ ; significant  $q_s(10) = 4.03$ ,  $P = 0.02$ ; largest non-significant  $q_s(10) = 2.21$ ,  $P = 0.16$ ). Right: Model comparison between two one-parameter models of temperature's effect on average implant speed dilation in DS animals (one-sample two-tailed t-tests,  $t_{\beta+\Delta T}(5) = -1.19$ ,  $P = 0.29$ ;  $t_{\beta+|\Delta T|}(5) = 3.21$ ,  $P = 0.02$ ; paired-sample two-tailed t-test,  $t(5) = -2.72$ ,  $P = 0.04$ ). The goodness of fit of the monotonic (left) and the non-monotonic (right) model, as measured by its coefficient of determination ( $R^2$ ), is shown for individual animals (open markers) and the population (mean  $\pm$  SEM). The  $\beta$  coefficient was constrained to only have positive (negative) values in the monotonic (non-monotonic) model.

(Fig. 2.31A), similar to the observed effects of temperature on both average movement speed (Fig. 2.30) and baseline striatal firing under anesthesia (Fig. 2.9). Anecdotally, the animal with the largest effect of temperature on psychometric threshold was also the animal with the strongest non-monotonic relationship between reaction times and temperature.

Lastly, endogenous variability in baseline firing rates recorded in the fixation version of the interval categorization task also covaried with reaction times (Fig. 2.32).

These observations suggest that the temporal evolution of striatal population state controls the time course of decision variables used by rats to determine *when* to do *what* action, but that this feature of neural activity is not what determines the moment by moment kinematics of movement execution. Instead, we find evidence that the striatum may provide a lower dimensional gain signal coded in the firing rate of neurons that controls




Figure 2.31: **(A)** Left: median reaction time dilation for DS animals (one-sample two-tailed t-tests,  $t(5) = [1.31, 2.01]$ ,  $P = [0.10, 0.25]$ ; repeated measures ANOVA,  $F(2, 15) = 1.43$ ,  $P = 0.27$ ). Right: median reaction time stretch for DS ( $N = 6$ ) and M1 ( $N = 4$ ) animals (one-sample two-tailed t-tests,  $t_{DS}(5) = -1.84$ ,  $P = 0.12$ ;  $t_{M1}(3) = -0.52$ ,  $P = 0.64$ ; two-sample two-tailed t-test,  $t(8) = -0.94$ ,  $P = 0.37$ ). **(B)** Same as (A), but for median movement times (left: one-sample two-tailed t-tests,  $t(5) = [0.28, 1.04]$ ,  $P = [0.35, 0.79]$ ; repeated measures ANOVA,  $F(2, 15) = 0.60$ ,  $P = 0.56$ ; right: one-sample two-tailed t-tests  $t_{DS}(5) = -1.01$ ,  $P = 0.36$ ;  $t_{M1}(3) = -1.04$ ,  $P = 0.38$ ; two-sample two-tailed t-test,  $t(8) = -0.30$ ,  $P = 0.77$ ).

overall vigor (Cruz et al., 2022; Panigrahi et al., 2015; Turner & Desmurget, 2010).

## 2.3 Discussion

Previous work has demonstrated that the speed of neural population activity along reproducible trajectories in neural space can correlate with variability in the timing of actions (Crowe et al., 2014; Murakami et al., 2014; J. Wang et al., 2018) and time-dependent decisions (Gouvêa et al., 2015; J. Kim et al., 2013; Leon & Shadlen, 2003). Indeed, there is growing evidence that this may represent a general principle for temporal processing in the brain (Paton & Buonomano, 2018). Here we provide additional causal evidence for this principle, by showing that experimental temperature manipulations in the striatum can be used to slow down or speed up both patterns of neural activity under anesthesia and the temporal evolution of latent decision variables that rats use to guide behavior, specifically their categorical judgments of duration. These results appeared to reflect a direct impact of temperature




Figure 2.32: Baseline DS firing rate correlates with reaction but not movement times in the fixation version of the interval categorization task. **(A)** Average TED implant speeds aligned to stimulus onset recorded during control blocks in the fixation version of the interval categorization task (DS cohort,  $N = 6$  animals). Solid black lines and shaded gray patches represent individual animal medians and IQR ranges, respectively. **(B)** Distribution of correlation coefficients between baseline firing rates of individual striatal neurons ( $N = 483$ ) and subsequent reaction times (top; one-sample two-tailed  $t$ -test,  $t(482) = -8.56$ ,  $P = 1.55e-16$ ) or movement times (bottom; one-sample two-tailed  $t$ -test,  $t(482) = 0.15$ ,  $P = 0.88$ ) for the DS-recorded animals trained on the fixation variant of the interval categorization task ( $N = 3$ ).

on decision-related (or action selection-like) and not moment by moment movement execution-related (or continuous control-like) processes, because the effects of temperature on timing judgments and movement kinematics were qualitatively distinct. It has long been appreciated in fields as diverse as neuroscience, robotics, and artificial intelligence that control of movement is likely facilitated by a hierarchy of control mechanisms (Merel et al., 2019). Even relatively simple nervous systems must both switch between distinct motor pattern generators, as well as continuously control their output. The observation that an intervention capable of modulating striatal population speed differentially impacts the evolution of decision variables, and those required for the continuous control of movement, is consistent with the proposal that the BG act as a mid-level controller, important for selecting among (Doya, 1999), linking (Graybiel, 1998), or modulating different actions, but provides evidence against their issuing moment by moment commands required for movement execution. One caveat of this interpretation is that any potential movements occurring during the interval period, when animals' judgments were presumably forming, were below our threshold for

detection. Thus, we analyzed the kinematics of choice movements, which necessarily occurred after the temporal judgment had concluded. It is thus possible that temperature produced similar effects on judgments, and the kinematics of putative, undetected movements during the decision process, and that its effects on movement kinematics changed qualitatively when it came to executing choice movements a few hundred milliseconds later. We consider this possibility less likely for both parsimony, and because of existing data that would appear inconsistent with the BG providing detailed moment by moment kinematic information for use in behavioral control. For example, dorsal striatal lesions have been shown to disrupt the sequencing of behavioral motifs, without affecting their execution (Markowitz et al., 2018). Even data demonstrating that striatal lesions can impact detailed kinematics of learned stereotyped movements can be reconciled with a more specific role for the BG in discrete control, as learned movements appear to involve the concatenation of motor elements, or sequencing, implying that discrete transitions must be appropriately placed within the overall time course of the behavior (Dhawale et al., 2021). Observations that kinematic information can be decoded from the striatum during tasks that involve timing (Dhawale et al., 2021; Rueda-Orozco & Robbe, 2015) are consistent with this view as well, because BG circuits would likely benefit from, while potentially not depending exclusively on, access to continuous information about the state of commands that are sent to effector systems to guide when to produce particular actions, even if they are not responsible for sending detailed kinematic control signals to downstream circuits. Our results are reminiscent of the observation that cooling an orofacial region of motor cortex in singing mice slows certain aspects of the song timing while leaving individual note durations unchanged (Okobi et al., 2019). Interestingly, tonic inhibition of the direct feedforward pathway of the BG at its initiation point in the striatum can produce a slowing of movement (Panigrahi et al., 2015), suggesting that the influence of the BG on control parameters related to vigor may not be through dynamics in the higher dimensional space of population firing, but rather through a low dimensional, gain-like, modulation of motor programs that are largely implemented by circuitry elsewhere. Consistent with this we found that endogenous variability in the baseline firing rate of striatal neurons covaried with reaction times, and that the non-monotonic effect of striatal temperature on this feature resembled a pattern in rats' behavior where both warmer and cooler temperatures produced delayed and slower

choice movements.

During behavior, animals continuously interact with the environment, and this interaction drives significant amounts of neural activity across many brain areas. Pinpointing features of neural activity that underlie latent processes, such as aspects of cognition, in the face of behavior-driven neural activity is a difficult problem requiring multiple approaches. One approach involves studying neural activity across behavioral tasks that vary along dimensions that are orthogonal to the process of interest. For example, signals related to a decision and not its report should be invariant to movement conditions of the task (Parker & Newsome, 1998). In the current study, we found that variability in the speed of striatal population responses correlated with subjects' temporal judgments across two task variants that differed in the degree that subjects were free to move about during stimulus presentation. This consistency argues against a scenario where movement is the principal driver of the correlation between neural variability and temporal judgment. A priori, it is impossible to determine whether such correlations reflect a causal relationship from that neural activity to behavior, an unobserved source of neural activity elsewhere that directly causes both striatal neural activity and behavior, or potentially behavior that causes striatal neural activity through reentrant sensory input. Previous analysis of high speed video during the no-fixation version of the temporal judgment task we study here demonstrated that decision-related information appeared in striatal activity several hundred milliseconds before it appeared in outward behavior, arguing against decision signals being driven by the sensory consequences of behavior alone (Gouvêa et al., 2015). Nonetheless, the existence of such varied possibilities underlies the need for so-called causal interventions to help determine whether neural activity is a cause, corollary or a consequence of behavior. Here, to examine the impact of a temperature intervention on population dynamics, we studied dynamics elicited by optogenetic activation of VB thalamus under anesthesia, and not the endogenous patterns of activity observed during behavior. This approach has the drawback of not allowing for simultaneous assessment of the impact of temperature on neural activity and behavior, as well as neural activity being generated via circuit mechanisms that are distinct from those at play during behavior. However, it possesses advantages as well. Specifically, behavioral sources of neural variability are removed altogether, allowing for a much more direct assessment of temperature's effects on features of neural population dynamics that

were independently seen to correlate with timing judgments in multiple sets of experiments. In terms of delineating the nature of the causal influences between the various potential factors at play, technically this independence strengthens our confidence that the causal chain runs from temperature, to population dynamics, to animals' temporal judgments (Pearl, 2009).

What circuit mechanisms might give rise to task-relevant striatal population activity studied here? Previous views have suggested that the BG largely inherit their patterns of activity from inputs, often hypothesized to originate in cortex (Matell & Meck, 2004; J. Wang et al., 2018). However, if time-varying striatal activity underlying temporal judgment were simply inherited from some other brain area, striatal temperature manipulations would be expected to produce a minimal shift and not a significant rescaling of behavior (Long & Fee, 2008). As in the vocal control circuit of songbirds, our data thus appear to be inconsistent with a mechanism where the relevant dynamics are simply inherited by the brain area targeted for temperature manipulations, in our case the dorsal striatum. In songbirds, there is evidence that a combination of local circuit mechanisms in pallial area HVC (Egger et al., 2020) and a larger reentrant circuit involving HVC and multiple other brain areas are involved in generating the temporally patterned activity underlying song timing (Hamaguchi et al., 2016; Long & Fee, 2008). A similar scenario may underlie the mechanisms for generating temporally patterned striatal activity involved in temporal judgments. First, while most network modeling efforts that use neural dynamics for computation have relied at least in part on recurrent excitation, recent work suggests that it may be possible for a largely inhibitory, striatum-like circuit to produce complex spatio-temporal dynamics given sustained excitatory input (Murray & Escola, 2017). The striatum may also represent one stage in a larger reentrant circuit involving multiple brain systems, where the larger circuit contributes to generation of dynamic patterns of activity that govern the evolution of decision variables. In this view, delays or advances induced by cooling or warming would accumulate with each cycle through the circuit, resulting in temporal rescaling with temperature. Such a circuit could in principle involve cortex, BG structures and thalamus (Alexander et al., 1990), or subcortical areas such as downstream BG structures, superior colliculus and thalamus (McHaffie et al., 2005). However, our data suggest that any reentrant circuit mechanism involving cortex did not involve primary motor cortex, as temperature manipulations there had negligible effects on choice

behavior, consistent with previous studies demonstrating that manipulating motor cortex does not affect expert behavior in motor timing tasks (Kawai et al., 2015; Xu et al., 2014). However, orbito-frontal and medial frontal cortical areas have been shown to encode temporal information during both motor timing and temporal judgment tasks, albeit less accurately than the striatum (Bakhrin et al., 2017; J. Kim et al., 2013), and cooling of medial frontal cortical structures has been shown to delay movements (Xu et al., 2014), potentially indicating involvement of frontal cortical structures. In addition, the activity of midbrain dopamine neurons correlates with and can directly cause changes in timing judgments (Soares et al., 2016), suggesting that dopaminergic neuromodulation may additionally tune the time course of network activity through its action on striatal circuits.

It has been proposed that timing processes are distributed in the brain (Paton & Buonomano, 2018), and that networks of neurons implicitly possess a rich capacity to act as timekeeping mechanisms through the time-varying patterns of activity they tend to produce during behavior, sometimes termed a “population clock” (Buonomano & Merzenich, 1995). While the data presented here strongly support this hypothesis, in principle the kinds of computations performed by earlier more algorithmic, information processing accounts (Gibbon, 1977) of timing might well be embedded in the type of population activity we describe here. This possibility is reflected in a growing belief that the brain performs many of its computations through dynamics (Vyas et al., 2020).

Our percepts, thoughts, and actions are continuously intertwined and regulated in time, and understanding the neural basis of temporal processing has been argued to be a necessary prerequisite for general models of cognition (Finnerty et al., 2015). Yet understanding how the brain appropriately orders and spaces information along the temporal dimension to encode signals that are functions of time has been an enduring challenge for neuroscience. Here we provide compelling evidence that the time course of activity in populations of striatal neurons directly influences the time course of a timing process used to guide decision-making. The data not only imply a causal link between temporal scaling of a population response and a temporal basis for computation in the brain, but argue that, in the hierarchy of behavioral control, striatal dynamics may act at an interface between cognition and motor function to help guide whether and when to produce what actions, but not to provide the moment by moment, continuously evolving

motor commands required for movement execution. Understanding the precise circuit mechanisms responsible for establishing and modulating the timescale of neural activity in these circuits, and which specific computations this activity subserves, represent important future directions toward understanding how the brain flexibly operates on internally computed information to produce adaptive and intelligent behavior.

## 2.4 Methods

All experimental procedures were in accordance with the European Union Directive 2010/63/EU and approved by the Champalimaud Foundation Animal Welfare Body (protocol number: 2017/013) and the Portuguese Veterinary General Board (project approval: 0421/000/000/2018).

### 2.4.1 Subjects

A total of 32 adult Long-Evans hooded rats (*Rattus norvegicus*) between the ages of 6 and 24 months were used in this study. Two rats were used in an acute experiment aimed at characterizing the spatiotemporal profile of our temperature manipulation. Another 4 animals were used in an acute experiment combining electrophysiological recordings, temperature manipulation and optogenetic stimulation. 26 wild-type males were trained in the interval categorization task (across the fixation and no-fixation variants), of which 15 were chronically implanted with a custom TED that allowed for temperature manipulation experiments, 6 were implanted with 32-channel tungsten microwire moveable array bundles (Innovative Neurophysiology): 3 unilaterally (previously published data (Gouvêa et al., 2015)) and 3 bilaterally. Five were used in behavioral manipulation experiments. Prior to surgery (where applicable), animals were kept in pairs in transparent cages with HEPA (High-Efficiency Particulate Air) filters on a 12-hour light-dark cycle (with lights ON at 8 am), at 21 °C and relative humidity of 50%. All experimental procedures were performed during the light phase of the cycle. Animals used in behavioral experiments had *ad libitum* access to food and were water-deprived.

### 2.4.2 Behavioral apparatus

The behavioral apparatus consisted of a 36 cm tall, 22.5 cm wide, and 35 cm long plastic storage box (TROFAST, Ikea) with three floor-level custom nose ports, a speaker (LS00532, Pro Signal) nearing the top of the opposite wall and a custom-made lid that provided uniform lighting and allowed for overhead video recordings (Flea3 FL3-U3-13S2, Point Grey Research Inc.) through an aperture. Each cylinder-shaped nose port was made up of 3D printed components housing a white light emitting diode (light-emitting diode (LED)), an infrared (IR) emitter-sensor pair that enabled the detection of port entries and exits and the accompanying printed circuit board (printed circuit board (PCB)) (Champalimaud Foundation Scientific Hardware Platform). Additionally, the two lateral ports (positioned symmetrically around the central one) were each equipped with a metallic spout connected to a 20  $\mu$ L water syringe via a solenoid valve (LHDA1231215H, Lee Company). All sensors, actuators and peripherals were respectively monitored, controlled and kept in the same temporal reference frame using a custom finite state machine implemented by a microcontroller I/O board (Arduino Mega 2560, Arduino) and an interfacing PCB (Champalimaud Foundation Scientific Hardware Platform). Finally, detected port events and other task-relevant behavioral data were timestamped, serially communicated to a Windows 10 desktop computer and stored as a parseable text file using a custom Python script. Video was acquired at 60 FPS with a resolution of 1280 x 960 pixels in 8-bit grayscale using Bonsai (Lopes et al., 2015).

### 2.4.3 Behavioral training

Leading up to the experimental sessions reported in this paper, animals were first trained in two hour-long daily sessions five times a week in various “tasks” of increasing complexity. During this stage, termed *Poking101*, rats were progressively introduced to the following rules: (un)lit ports are (un)responsive, meaning that nose-poking into a lit port will cause it to turn off and trigger some task event, whereas doing so at an unlit port is inconsequential; entering a lit lateral port results in a reward delivery of 25  $\mu$ L of water paired with a brief secondary reinforcer (auditory tone, 1750 Hz, 150 ms); entering the central port when it is illuminated initiates a trial and can lead to both lateral ports lighting up. This is contingent on the animal’s snout continuing to interrupt the IR beam at the center port for the

entirety of a “fixation delay” (fixation version), or the animal not making a premature entry at either lateral port during the “withhold delay” (no-fixation version). Whichever the task variant, this imposed delay starts off at 0 s and is adaptively marched up towards 3 s (within and across sessions) and consists of an interval of silence demarcated by two brief auditory tones (7500 Hz, 150 ms). Failure to withhold premature departures from the central port (fixation version) or choices (no-fixation version) causes the current trial to be aborted, eliciting an error tone (150 ms of white noise) and adding a timeout of 15 s to the already ticking 9-s inter-trial-onset interval (ITOI). Once animals were able to reliably maintain fixation at the central port (fixation version) or defer choices (no-fixation version) for 3 s, training on the interval categorization task began (Gouvêa et al., 2015; Gouvêa et al., 2014). In it, instead of waiting for a fixed amount of time and collecting a reward at either lateral port once it elapsed, rats were asked to wait for a variable delay on each trial and to then categorize it as either shorter or longer than a 1.5-s boundary. “Short” judgments were registered at one of the lateral nose ports and “long” judgments at the opposite one. Rewards were contingent on stimulus and judgment categories being the same. When this was not the case, an error tone (150 ms of white noise) was played and a time penalty of 10 s was added to the ITOI. Pairs of stimuli symmetric about the categorical boundary were gradually introduced (from easiest to hardest) into the discrete sampling set that animals experienced, until reaching  $I = 0.6, 1.05, 1.38, 1.62, 1.95, 2.4$  s (fixation version) or  $I = 0.6, 1.05, 1.26, 1.38, 1.62, 1.74, 1.95, 2.4$  s (no-fixation version). A correction-loop procedure was used such that, following 3 incorrect categorizations of any given stimulus, only that stimulus was presented to the animal until its associated error count dropped below 3. This training mechanism was disabled during manipulation and recording sessions. It took  $\sim 3$  months for rats to reach asymptotic performance.

#### 2.4.4 Thermoelectric device

##### 2.4.4.1 Design

We used a custom-made implantable TED (weighing  $\sim 30$  g) based on the Peltier effect to manipulate temperature in neural tissue. The implant consisted of a heat dissipation module, a thermoelectric cooling (thermoelectric cooling device (TEC)) module (01801-9A30-12CN, Custom Thermoelectrics),

a 10 k $\Omega$  thermistor (9707204, Farnell) and two 15-mm long sharpened silver probes. These were insulated down to, but excluding, the tips with a thin layer of PTFE low density thread seal tape (00686081520745, Gasoila) and polyimide tubing. The main distinguishing factors between the implant's initial prototype (used in the single cooling dose pilot experiments) and its final version (used on the DS and M1 bidirectional manipulation cohorts performing the fixation task variant), were that the former was constructed with a passive aluminum heatsink (ICKS25X25X18,5, Fischer Elektronik) and 0.5-mm thick probes insulated with 1-mm wide polyimide tubing (95820-11, Cole-Parmer), whereas the latter had active heat dissipation via a water block (WBA-1.00-0.49-AL-01, Custom Thermoelectrics) and 1-mm thick probes insulated with 2-mm wide polyimide tubing (95820-13, Cole-Parmer). This water block was used in tandem with a peristaltic pump (200-SMA-150-050, Williamson), male and female Luer adapters (WZ-45504-00, Cole-Palmer) and the required interfacing tubing (WZ-06407-71, Cole-Palmer), allowing for a continuous flow ( $\sim 15 \text{ mL min}^{-1}$ ) of room temperature water through the water block's inner chambers. The TEC's upper plate was glued to the bottom of the heatsink using thermal glue (TBS20S, TBS), which was also used to secure the thermistor at the center of this module's lower plate. Finally, the two sharpened silver probes were soldered onto the TEC's lower plate (one on each side of the thermistor) using a mixture of lead (419424, Farnell) and silver solder (SDR-9703-030, Custom Thermoelectrics), at a distance of 5 mm from each other. This inter-probe spacing corresponds to two times the ML stereotaxic coordinate of all our DS-targeted implants. Lastly, an RJ45 (85513-5014, Molex) connector was added on top of the heat sink and a custom 3D-printed spacer was mounted on its bottom, both secured using epoxy resin (2022-1, Araldite).

### 2.4.4.2 Closed-loop control

The implant was plugged into a custom-made PCB (developed by the Champalimaud Foundation Scientific Hardware Platform and available upon request) via an ethernet cable. This PCB implemented a proportional-integrative-derivative (proportional-integral-derivative (PID)) controller that was designed to bring the implant's thermistor measurement to any experimenter-defined target temperature (within the TEC's range of operation). Briefly, the thermistor readout was continuously compared to the current temperature

setpoint in order to compute an absolute error term (proportional channel), a cumulative error (integrative channel) and an instantaneous change in error (derivative channel). These three error terms were then linearly combined, with weights set by the resistive and capacitive components of the hardware modules that implemented them, and used to modulate the control current driving the TEC. This negative feedback mechanism was optimized so that the target temperature could be reached with negligible delays, steady-state errors and over/undershoots. The resulting closed loop control allowed for stable, safe and transient temperature manipulations, as it required less user intervention, monitorization and arbitration than open loop alternatives. The PID's setpoint was communicated through a serial communication pin from an additional Arduino Mega 2560 board that implemented the temperature manipulation protocol *per se*, meaning it controlled both when to transition into a new block and which temperature to transition to. All block types lasted for 3 minutes, except for control ones in our single cooling dose experiment, which were twice as long to accommodate slower heat dissipation due to this initial experiment's characteristic passive heatsink. In all cases, block transition times and target temperatures were respectively signaled via a brief digital pulse and an additional serial communication port to the task-implementing Arduino board. Both the PID-implementing PCB and the block-controlling Arduino were connected to a computer running Windows 10, where a LabView-based (National Instruments) graphical user interface (TEC visualizer, Champalimaud Foundation Scientific Hardware Platform) enabled online visualization and saving of digitized thermistor temperature measurements (sampled at 100 Hz). Lastly, to prevent irreversible tissue damage in the eventuality of a partial compromise of the closed loop system leading to its "opening", an additional failsafe mechanism was implemented in the PCB's firmware, ensuring that the TED was automatically disabled if the registered thermistor temperature ever dipped below 0 °C or rose above 55 °C.

#### **2.4.4.3 Calibration**

A calibration curve between different set temperatures at the lower plate of the TED module and temperature measurements at the tip of the silver probes was derived from an acute preparation with an anesthetized rat. Lower plate temperature was set to each value in  $T = 5, 15, 20, 25, 30, 45$  °C,

in blocks of 4 minutes, always preceded and followed by a control block of the same duration ( $T = 36^\circ\text{C}$ ). Temperature at one of the tips of the implant's probe was measured by a second thermistor glued along the probe axis to the polyimide insulation layer. In a separate acute experiment, we positioned a thermistor probe angled at  $30^\circ$  at different distances to the implant's tip ( $D = 0.3, 0.5, 0.9, 2, 5.75\text{ mm}$ ), and for each of them repeated the aforementioned blocked calibration procedure (but with manipulation temperatures drawn from the reduced set  $T = 15, 25, 42, 45^\circ\text{C}$ ). Lastly, all implants were tested individually post-assembly to ensure their respective TED modules were functioning steadily and properly calibrated, using the same 4-min block protocol in warmed agarose gel (1.5%), which has similar thermal properties to brain tissue (S. Kim et al., 2007).

## **2.4.5 Surgical procedures**

### **2.4.5.1 Acute temperature measurements and calibration**

Rats were anesthetized with 2.0-4.5% isoflurane. Animals' body temperature was continuously monitored and maintained at  $35^\circ\text{C}$  by a rectal probe connected to a closed-loop heating system (FHC, <https://www.fh-co.com>). After being anesthetized and before making the first incision, we administered dexamethasone ( $2\text{ mg kg}^{-1}$ ), carprofen ( $5\text{ mg kg}^{-1}$ ) and a saline solution of atropine ( $0.05\text{ mg kg}^{-1}$ ) subcutaneously (SC). We stereotaxically targeted the DS unilaterally ( $+0.84\text{ mm AP}$ ,  $+2.5\text{ mm ML}$  from Bregma (Paxinos & Watson, 2013). Immediately following these procedures, animals were perfused for histological confirmation of the measurements' location.

### **2.4.5.2 Viral injections**

Following the same procedure for induction and maintenance of anesthesia used for acute temperature measurements and calibration, we stereotaxically targeted the VB of the thalamus for viral delivery ( $-2.3\text{ mm AP}$ ,  $\pm 2.8\text{ mm ML}$ ,  $6.6\text{ mm DV}$  from Bregma (Paxinos & Watson, 2013). We injected 300 nL of rAAV5-CamKII-hChR2(H134R)-EYFP (titer  $\sim 10^{12}\text{ GC\%}$ ; University of Pennsylvania Vector Core) using an automated microprocessor controlled microinjection pipette with micropipettes pulled from borosilicate capillaries (Nanoject II, Drummond Scientific). Injections were performed at  $0.2\text{ Hz}$  with  $2.3\text{ nL}$  injection volumes per pulse. For all injections, the micropipette

was kept at the injection site 10 minutes before withdrawal. Craniotomies were then covered with Kwik-Cast (WPI) and the skin was closed with sutures (Vicryl, Ethicon Inc.). Animals were allowed to fully recover on a warming pad and returned to the home cage when fully alert. During the 3 days following surgery, animals were given carprofen ( $5 \text{ mg kg}^{-1}$ , SC).

#### **2.4.5.3 Acute optogenetic stimulation, extracellular recordings and temperature manipulation**

Following 3-6 weeks for viral expression, 4 rats were anesthetized with two doses of urethane, the first at  $0.7 \text{ g kg}^{-1}$  of body weight and the second at  $0.35 \text{ g kg}^{-1}$  20 minutes after. Additionally, we administered dexamethasone ( $2 \text{ mg kg}^{-1}$ ), carprofen ( $5 \text{ mg kg}^{-1}$ ) and a saline solution of atropine ( $0.05 \text{ mg kg}^{-1}$ ) subcutaneously. Animals were then kept with isoflurane at 0.5-1% until at least 30 minutes before electrophysiological recordings began. Animals' body temperature was continuously monitored and maintained at  $35^\circ\text{C}$  by a rectal probe connected to a closed-loop heating system (FHC, <https://www.fh-co.com>) throughout the experiments. We opened a large rectangular craniotomy over the left hemisphere (4 mm AP by 3 mm ML from Bregma (Paxinos & Watson, 2013)), centered in the same target location as the chronic implants. A  $300 \mu\text{m}$  diameter and 0.37 NA optic fiber (Doric) was targeted to VB ( $-2.3 \text{ mm AP}$ ,  $\pm 2.8 \text{ mm ML}$ ,  $6.2 \text{ mm DV}$  from Bregma (Paxinos & Watson, 2013)), inserted at a  $39^\circ$  angle and secured with blue light cured self adhesive resin cement (RelyX™ Unicem 2 Self-Adhesive Resin Cement, 3M). A small silver ground wire was inserted under the skull of the opposite hemisphere. A TED similar to the one used for chronic implants (with a single silver probe at a  $90^\circ$  angle relative to the heat sink, to accommodate the geometrical demands of the experimental preparation) was lowered to the same DS target location ( $+0.84 \text{ mm AP}$ ,  $-2.5 \text{ mm ML}$ ,  $4 \text{ mm DV}$  from Bregma (Paxinos & Watson, 2013)). This modified device was calibrated and behaved similarly to the ones used for chronic manipulations. Finally, a Neuropixels probe (Phase 3A Option 3, IMEC (Jun et al., 2017)) was placed caudally relative to the temperature probe, and slowly lowered to target ( $5$  to  $6.5 \text{ mm DV}$ ) and allowed to stabilize in the tissue for at least 30 minutes before starting recordings and stimulation protocols. Seldomly, and for longer recording protocols, an additional dose of urethane was necessary to maintain anesthesia ( $0.2 \text{ g/Kg}$ ). A  $473 \text{ nm}$  LED source (Doric) was

connected to the implanted optical fiber using a patch cord (400  $\mu\text{m}$  core, 0.48 NA) and set to 3.5-5.5 mW at the end of the fiber and controlled using a dedicated arduino that was also responsible for switching the block temperature identity through a serial communication with the TED controller. Each stimulation trial consisted of a single train of 5, 1-ms long, pulses at 100 Hz (each train lasting 50 ms in total). Each trial was separated by a period of 1.5 seconds. Electrophysiological and peripheral synchronization (LED and temperature probe) data were simultaneously acquired using SpikeGLX software (<https://billkarsh.github.io/SpikeGLX/>) at 30kHz. Local-field potential gain and action potential gain were set at 250 and 500, respectively, and split at 300 Hz. Each block at a specific temperature lasted 3 minutes. Temperature identities were drawn, without replacement, from the available set of 3 temperatures and were always intercalated with a control block. This protocol was repeated twice for a total of 2 blocks for each manipulation condition. Immediately following these procedures, animals were perfused for histological confirmation of the measurements' location.

### 2.4.5.4 Chronic TED implantation

Rats ( $N = 15$ ; 5 in the single cooling dose pilot experiment, 6 and 4 in the bidirectional DS and M1 temperature manipulation experiments, respectively) underwent surgery around 3 months after they started training. During the implantation of the TED rats were anesthetized with 2.0-4.5% isoflurane. Animals' body temperature was continuously monitored and maintained at 35 °C by a rectal probe connected to a closed-loop heating system (FHC, <https://www.fh-co.com>). After being anesthetized and before making the first incision, we administered dexamethasone (2 mg kg<sup>-1</sup>), carprofen (5 mg kg<sup>-1</sup>) and a saline solution of atropine (0.05 mg kg<sup>-1</sup>) subcutaneously. We stereotactically targeted the DS bilaterally (+0.84 mm AP,  $\pm 2.5$  mm ML from Bregma (Paxinos & Watson, 2013)). Two craniotomies and durotomies matching the diameter of the silver probes were made. 5 support screws were placed: 1 in the occipital plate, 2 posterior and 2 anterior to the location of the craniotomies. The cranial bone was covered with self-curing dental adhesive resin cement (Super-Bond, C&B) to improve adherence to the dental acrylic used to secure the implant. The TED was then slowly lowered perpendicular to the brain surface to a depth of 4 mm from cortical surface. The craniotomies were covered with Kwik-Cast (WPI) and the implant was

fitted into place and secured with several layers of dental acrylic (the first of which mixed with gentamicin). The procedure ended with suturing (Vicryl, Ethicon Inc.) the skin anterior and posterior to the implant. Animals were allowed to fully recover on a warming pad and returned to the home cage once fully alert. Animals were then individually housed to minimize implant damage. During the 3 days following surgery, animals were injected once a day with carprofen ( $5 \text{ mg kg}^{-1}$ , SC). Animals were allowed to recover for a week after the surgery with food and water *ad libitum*.

#### **2.4.5.5 Chronic electrophysiology implantation**

These procedures followed a protocol similar to the one used for the implantation of the chronic TED implants. For a detailed description, see (Gouvêa et al., 2015).

#### **2.4.6 Temperature manipulation protocol**

Following one week of recovery from surgery, all animals were again water-deprived and gradually resumed behavioral training. Once they were performing at approximately pre-surgical levels in their assigned version of the interval categorization task, they were subjected to temperature manipulation sessions. These 2-h sessions consisted of 6- or 3-minute control blocks (for the pilot and bidirectional temperature manipulation experiments, respectively), in which the TED was set to body temperature ( $\sim 36^\circ\text{C}$ ), always interleaved with 3-minute manipulation blocks, in which the TED was either set to  $25^\circ\text{C}$  in the single cooling dose pilot experiments, or one of 3 manipulation doses ( $15$ ,  $25$  and  $42^\circ\text{C}$ ) in all bidirectional manipulation experiments. Manipulation temperatures were drawn at random and without replacement from the aforementioned set until its exhaustion, at which point the set was replenished and the sampling process resumed. Sessions invariably started and ended with a control block and animals were not explicitly cued to block transitions. Manipulation sessions were interleaved with washout sessions, in which the TED's controller was disabled, and correction-loop training was reinstated.

### **2.4.7 Implant placement confirmation**

Rats were sacrificed with transcardiac perfusion with phosphate-buffered saline (phosphate buffered saline (PBS)), followed by 4% (wt/vol) paraformaldehyde (PFA). Following perfusion, brains were left in 4% PFA for 24 h and then moved to a 30% sucrose solution (wt/vol) in PBS for 2 to 3 days. For chronic electrophysiology, single-cooling dose and acute experiments, a vibratome was used to section the brain into 50  $\mu\text{m}$  coronal or 40  $\mu\text{m}$  sagittal slices, respectively. Coronal slices were stained with NISSL and sagittal slices series were alternated with NISSL or immunostained with a primary antibody against green fluorescent protein (GFP) (A-6455, Invitrogen) and a secondary antibody conjugated with AlexaFluor 488 (ab150077), and finally, incubated in DAPI. Images were acquired with a stereoscope (Lumar V12, Zeiss) or a slide scanner (Axio Scan Z1, Zeiss). For animals subjected to bidirectional chronic temperature manipulations, a 1 T MR scanner (ICON, Brucker) was used to collect MRI data. A T2-weighted structural image of the brains was collected using a Rapid Imaging with Refocused Echoes (RARE) pulse sequence. The sequence used had a repetition time of 2800 ms, echo time of 90 ms and a RARE factor of 12. The field of view was set to  $28 \times 15 \times 20 \text{ mm}^3$ , the spatial resolution of the images was  $150 \times 150 \times 150 \mu\text{m}^3$  or  $80 \times 80 \times 80 \mu\text{m}^3$  and a matrix of  $187 \times 100 \times 133$  voxels was acquired after 8 averages during a 7-hour scanning.

### **2.4.8 Data analysis**

Unless otherwise stated, all data were analyzed using custom MATLAB 2019b (<https://www.mathworks.com>) scripts.

### **2.4.9 Psychophysical data analysis**

#### **2.4.9.1 Preprocessing**

Trials with reaction times greater than 1 s, or movement times greater than 2 s were labelled as outliers and excluded from all reported analyses. This resulted in less than 5% of all trials being removed. In order to make balanced comparisons across animals and temperature conditions, data from the first two manipulation sessions of every chronically implanted animal were pooled together chronologically up to the point where there were 10 trials per stimulus for each manipulation condition and 40 trials per stimulus for

the control temperature condition. The same pooling procedure was applied in reverse for the last two temperature and boundary manipulation sessions (see Supplementary Information).

#### 2.4.9.2 Psychometric function

We used the *Psignifit* (H. Schütt et al., 2015) toolbox to fit the following 4-parameter psychometric function to all interval categorization data:

$$\Psi(x; m, w(\alpha), \lambda, \gamma) = \gamma + (1 - \lambda - \gamma) \cdot S(x; m, w(\alpha))$$

$$S_{\logistic}(x; m, w(\alpha)) = \frac{1}{1 + \exp^{-2 \ln(1/\alpha - 1)(x - m)/w}}, \alpha = 0.05$$

In this parameterization, a strictly monotonic sigmoid function  $S$  from the stimulus level  $x$  onto the unit interval  $[0, 1]$ , is specified by  $m = S^{-1}(0.5)$  and  $w = S^{-1}(1 - \alpha) - S^{-1}(\alpha)$ , namely the threshold and width parameters. This is independent of the choice of  $S$ , which in our case is the logistic function. The hyper-parameter  $\alpha$ , which sets the span of  $w$  along the vertical axis, was set to 0.05. To account for stimulus-independent choices,  $S$  is scaled by two additional free parameters,  $\lambda$  and  $\gamma$ , which respectively control the upper and lower asymptotes of the psychometric function  $\Psi$ . The and parameters were fixed across temperatures at values found through fitting the corresponding control temperature data.

#### 2.4.9.3 Dilation and stretch metrics

We adopted the dilation and stretch definitions from (Long & Fee, 2008). Briefly, dilation ( $D$ ) of any scalar metric  $x$  (e.g., threshold MAP), was calculated as the percent difference from unity in the ratio of a given temperature's estimate over that of control.

$$D = \left( \frac{x_{\Delta T}}{\Delta T = 0} - 1 \right) \times 100$$

Stretch ( $S$ ) was defined as the slope coefficient in a least squares linear regression using dilation as the response variable and the magnitude of our temperature manipulation (induced temperature differences around the implant's tip) as the sole predictor.

$$D = S \times \Delta T + \varepsilon$$

## 2.4.10 Electrophysiological data analysis

Unless otherwise stated, what follows applies to all three electrophysiology data sets analyzed in this paper: acute DS recordings (N = 335 neurons, across 4 animals), chronic DS recordings during the no-fixation (N = 441 neurons, across 3 animals) (Gouvêa et al., 2015), and the fixation (N = 483, across 3 animals) versions of the interval categorization task. We use the generic terms condition to refer to striatal temperature in the acute data set, and to stimulus-choice pairs in the two chronic data sets. The reference condition corresponds either to control temperature (acute) or to correct categorizations of the 2.4-s stimulus (chronic). Lastly, trial onset refers to the onset of VB stimulation (acute) or stimulus onset (chronic).

### 2.4.10.1 Preprocessing

In the case of the acute recordings, we focused on contiguous periods of stable activity (without abrupt changes in firing profile/rate), which in practice meant trimming the very beginning and/or end sessions where needed. For the three animals we recorded from chronically and unilaterally in the DS during performance in the no-fixation required version of the interval categorization task, preprocessing was done as described when this data set was first published (Gouvêa et al., 2015). For the remaining animals, we used a semi-automated offline approach to detect and sort recorded action potentials into well-isolated units and clusters of multi-unit activity. Detection, sorting and inference of the relative depth of each unit, were done using *KiloSort2* ([github.com/MouseLand/Kilosort2](https://github.com/MouseLand/Kilosort2)), whereas curation of the resulting clusters was performed using *Phy* ([github.com/cortex-lab/phy](https://github.com/cortex-lab/phy)). Prior to any of the analyses shown in the main figures, we further selected validated units with an intersectional approach that used firing rate and recording depth in the case of the acutely recorded data. Briefly, in order to survive this selection step, units had to have a mean firing rate of 0.5 Hz or higher and in the case of the acute recordings, have been recorded at a contact that was later inferred to be in the DS, by comparing its position along the Neuropixels probe to a dip in the distribution of recorded cell depths – likely corresponding to a characteristic “silence” when transitioning from gray (cortex) to white (corpus callosum) to gray (striatum) matter.

### 2.4.10.2 Single-neuron responses

Spike density functions were built on a trial-by-trial basis by first counting spike times in 2-ms bins and then convolving the resulting histograms with a causal kernel specified by a gamma distribution with shape and scale parameters of  $k = 2$  and  $\theta = 75$  ms, respectively. Baseline firing rates in the acute recordings were computed in a 500 ms window preceding VB stimulation. To compute temporal scaling factors for each unit-condition pair, we first upsampled reference spike density functions by a factor of 10 and then warped them in time using 1000 scale factors, linearly spaced in the range  $[0.625, 1.75]$ . Both the upsampling and time-warping steps were performed using linear interpolation. Next, we linearly regressed all time-warped templates against the spike density function of each condition and stored the corresponding coefficient of determination ( $R^2$ ). The scale factor that maximized this template-matching metric, is what we operationally defined as the temporal scaling factor for that unit-condition pair. In the case of control scaling factors, we split data into two random non-overlapping sets of trials and arbitrated which one was used to construct templates and which one was the target. To account for an edge artifact where the scaling factors of a small percentage of neurons were estimated to be at the lower and upper bounds of our scaling range ( $[0.625, 1.75]$ ), these neurons were excluded from the analysis in Fig. 2B. Regarding response dilation and stretch, we used the same definitions from the psychophysical analysis section, except that scaling factors ( $f$ ) replaced ratios of temperature over control estimates when calculating dilation:

$$D = (f - 1) \times 100$$

### 2.4.10.3 Low-dimensional representations of population state

We used principal component analysis (PCA) to enable visualization of striatal population trajectories in representative 3-dimensional subspaces. Briefly, we first averaged spike density functions across trials of a given condition and concatenated them into  $N \times (T \times K)$  matrices, where  $N$  is the number of neurons for that data set,  $T$  the number of 2-ms time bins relative to trial onset, and  $K$  the number of experimental conditions. After normalizing these data to have zero mean and unit standard deviation along the temporal dimension, we used trial-averages corresponding to the reference condition

to find the three orthogonal directions that maximally captured variance in said data. Lastly, each condition’s trial-averaged population activity was then projected onto the subspace defined by these PCs, with reference trajectories plotted for all time points, and non-reference trajectories plotted only for arbitrary time points and projected onto the corresponding reference trajectory so as to prevent visual clutter.

#### **2.4.10.4 Psychometric curves split by population state at stimulus offset**

This analysis was adapted from its original introduction in (Gouvêa et al., 2015), and was only applied to the two chronic data sets in the current paper. Briefly, for all individual trials in each session with 5 or more simultaneously recorded units, we projected population activity at stimulus offset onto the median trajectory traversed by that striatal ensemble during the entire stimulus presentation period. We then normalized these projections by the length of that median trajectory. In addition to this temporal scaling metric, we computed an outlieriness metric as the average point-by-point minimum distance between each trial’s trajectory and its session’s median trajectory (the 5% most extreme trials in these “outlieriness” distributions were removed from the subsequent analyses). After pooling all normalized projections over all sessions and animals, we then partitioned, for each stimulus, the resulting distributions into  $Q$  groups. For ease of comparison,  $Q$  was chosen to match the number of conditions in the temperature manipulation experiments during behavior in the single cooling dose pilot ( $Q = K = 2$ ) and the bidirectional temperature manipulation ( $Q = K = 4$ ) experiments. Psychometric curves were then fit to trials from each group as described above.

#### **2.4.10.5 Decoding time from ongoing population activity**

We used a naïve (flat prior) Bayes decoder to continuously compute probability distributions over elapsed time using multi-session concatenations of putative striatal population activity aligned to trial onset.

Briefly, we:

1. Discretized time  $t$  into  $B$  2-ms bins, such that  $b \in [1, B]$  and  $t_b \in [0, T]$  ms, with  $T = 1500$  ms for the acute recordings and  $T = 2400$  ms for the chronic recordings;

2. Fit an encoding model to each neuron  $n \in [1, N]$  at each point in experimental time  $t_b$ :

$$p(r_n|t_b)$$

Which we determined empirically by querying cell-specific single-trial spike density functions  $r_n$  at the time interval  $[t_b, t_{b+1}]$ , and smoothing the resulting rate histograms with a Gaussian kernel ( $\mu = 0, \sigma = 10$ ). This was done using a subset of trials making up half of all reference trials;

3. Made conditional independence assumptions about neurons, regardless of whether or not they were recorded simultaneously:

$$p(t_b|R) \propto \prod_{n=1}^N p(r_n|t_b)$$

Where  $R = (r_1, r_2, \dots, r_N)$  is a novel, to-be-decoded instance of concatenated population activity recorded at a known condition and point in time since trial onset;

4. Used Bayes' rule in creating a decoding model that linearly combined all individual neuron encoding models with a uniform prior over decoded time  $p(t_b) = 1/B$ :

$$p(t_b|R) = \frac{p(t_b) \prod_{n=1}^N p(r_n|t_b)}{p(R)}$$

Where the probability  $p(R)$  for the population vector  $R$  to occur does not have to be estimated explicitly, as it indirectly follows from normalizing the posterior distribution  $p(t_b|R)$  such that it sums to 1 over all possible decoder outcomes, i.e., elapsed time as decoded from striatal ongoing activity.

Once the time-dependence in the responses of striatal cells recorded during a set of training trials is known, this Bayesian approach directly addresses the inverse problem: given the firing rates of the same cells, now recorded during previously unseen test trials, how likely is it for any and all  $b$  units of time to have passed.

## 2.4.11 Continuous behavioral data analysis

### 2.4.11.1 Preprocessing

Full session 2-h videos recorded during the fixation version of the task were first cut into 15-s long clips (one per trial) aligned on stimulus onset ([−5, 10] s). Offline tracking of the position of several implant features was performed using *DeepLabCut* (Mathis et al., 2018). Contiguous low confidence position estimates (likelihood < 0.85) of 6 samples or less were linearly interpolated and tracking time series were subsequently smoothed with a 6-sample median filter. In order to standardize units of distance across animals and sessions, we computed session-wise background frames by taking the median pixel intensities from a random sample of 3000 frames from each session’s full video. Next, we arbitrarily picked one background frame from a particular session of a particular animal as a reference, and computed the affine transformations that would best align each session’s background frame to that reference. These transforms were then applied to both the  $x$  and  $y$  coordinates of each tracked feature’s position and finally, converted to approximately SI units of distance by exploiting knowledge of the behavioral box’s dimensions and assuming no nonlinear optical distortions.

### 2.4.11.2 Implant speed

Instantaneous implant speed ( $s$ ) was computed by taking the one-step forward difference of the average (across three implant features) displacement ( $\Delta r$ ) with respect to sampling time ( $\Delta t$ ).

$$s = \frac{\Delta r}{\Delta t} = \frac{\sqrt{x^2 + y^2}}{\Delta t}$$

To compute speed profile dilation and stretch, we used the same template-matching approach described above for single-neuron responses, but using average condition-split speed across time instead of condition-split spike density functions. Lastly, average speed dilation was calculated in the same way as threshold dilation, with each condition’s average – the mean across single trial median speeds computed between reaction (the first detected exit time from the initiation port following stimulus offset) and 350 ms (~the time animals took to settle into a nose port) after choice (the first detected entry time at a choice port) – replacing threshold point estimates.

### 2.4.12 Statistics and reproducibility

No statistical method was used to predetermine sample size but our sample sizes are similar to previous studies recording and/or perturbing neural activity (Xu et al., 2014), (Gouvêa et al., 2015), (Mello et al., 2015), (Jurado-Parras et al., 2020). When conditions contained less than 5 animals, results were confirmed to be significant within each animal. Data exclusion criteria are specified in the relevant methods sections. Trained animals were randomly assigned to the various experimental groups. Temperature manipulation sessions consisted of control blocks interleaved with block-wise randomized manipulation doses. In both variants of the interval categorization task, stimuli were drawn at random from trial to trial. Data collection and analysis were not performed blind to the conditions of the experiments.

Data distribution was assumed to be normal but this was not formally tested. Unless otherwise stated, we used one-sample two-tailed t-tests whenever assessing the statistical significance of shifts in distributions, which we visually afforded with vertical solid black lines connecting the distribution's mean to zero. When examining differences across distributions, we used either two-sample two-tailed t-tests when comparing striatal and motor-cortical stretch distributions, paired-sample two-tailed t-tests when comparing models of dilation as a function of temperature change, or repeated measures ANOVA followed by post-hoc contrasts with Tukey correction for multiple comparisons when comparing dilation distributions across conditions but within cohorts. We visually afforded these two- and paired-sample tests with horizontal solid black lines connecting the two underlying distributions, offset vertically for clarity. In all cases, we denote test outcomes near the respective visual affordance with the following notation: \*,  $P < 0.05$ ; \*\*,  $P < 0.01$ ; NS, not significant. In the case of single animal stretch estimates, we assessed their statistical significance at a 5% level by bootstrapping. Specifically, we computed these point estimates for 1000 random samples per manipulation condition, constructed by sampling equal numbers of trials with replacement from the control condition while preserving stimulus identity. For each iteration, we then performed linear regression on bootstrapped dilations and stored the respective slope coefficient as that iteration's stretch. Bootstrapped significance was consistently denoted by larger dark-filled markers, as opposed to smaller white ones.



## MANIPULATING STIMULUS INTENSITY DURING DELAYED INTERVAL DISCRIMINATION

To behave adaptively, the brain must register temporal structure in the environment and use it to organize behavior. The DLS integrates sensorimotor input, and is necessary for accurate timing and structuring behavior in general. However, if DLS provides the basis for mapping temporal features in the environment to behavior, its activity should predict variation in that mapping. A reanalysis of DLS population activity in rats comparing the duration of two sequentially presented vibratory stimuli revealed a striking correspondence between neural activity and behavior. Varying vibration intensity of the second stimulus induced systematic biases in temporal judgments, and corresponding biases in multiple features of DLS activity during stimulus presentation, including population coding of time. In contrast, the same intensity manipulations applied to the first stimulus affected neither behavior nor neural activity. Furthermore, neuronal response profiles were best described as a continuum, arguing against hypotheses where categories of responses, e.g., ramping activity, selectively underpin temporal processing. These data represent important additional evidence that striatal population dynamics support the organization of behavior by mapping temporal information to action.

### 3.1 Introduction

Evidence for the involvement of striatum in interval timing comes from a variety of sources, including disease states (Malapani et al., 1998; Rowe et al., 2010), lesions (Meck, 2006; Wolff et al., 2022), neuroimaging (Coull et al., 2004; S. M. Rao et al., 2001), neurophysiology (Gouvêa et al., 2015; D. Z. Jin et al., 2009; J. Kim et al., 2018; Matell et al., 2003; Mello et al., 2015; Monteiro et al., 2023), pharmacology (Gouvêa et al., 2015; A. V. Maricq & Church, 1983; Mello et al., 2015), optogenetics (Cruz et al., 2022) and temperature manipulations (Monteiro et al., 2023). A recent study challenged this body of work by concluding that though time information is inherent to neural activity in the DLS, it is otherwise decoupled from animals' estimates of elapsed time, as inferred from their duration judgments (Toso, Reinartz, et al., 2021). The data used to support this claim was recorded during two variants of a delayed comparison task. Briefly, rats were presented with two sequential vibratory stimuli that varied in both duration and intensity, and trained to either compare the duration and ignore the intensity (*delayed duration comparison*), or compare the intensity and ignore the duration (*delayed intensity comparison*) of the two stimuli. The authors' central conclusion that striatal activity was decoupled from temporal judgment was based on four purported features of DLS activity:

1. though rats' judgments about the duration of vibration stimuli were reliably biased by their intensity, activity in the DLS was not similarly modulated;
2. the DLS did not encode stimulus duration differently on correct and incorrect trials;
3. elapsed time was decodable from DLS activity in all epochs of the delayed duration comparison task, and not just during stimulus periods;
4. in rats trained in the delayed intensity comparison task variant, stimulus duration was encoded by DLS even though it was experimentally irrelevant.

Here, we focus on the delayed duration comparison task to demonstrate that 1 and 2 are false – DLS activity was unambiguously modulated by stimulus intensity and predictive of animals' duration judgments. This was readily

apparent in the population mean activity when conditioned on intensity or choice and corroborated by multiple quantitative analyses, including state space visualization, the construction of neurometric curves, probabilistic population decoding of elapsed time and finally a generalized linear model (GLM) that seeks to partition the relative contribution of different task variables to explainable variance in neural activity. We also argue that 3 and 4 logically do not constitute support for the claim that striatal activity is not used to guide judgment of duration. By analogy, would the response of a primary visual cortex (V1) neuron to a visual cue incidentally presented within its receptive field, but without behavioral importance for some task at hand lead one to conclude that V1 does not support visual behavior? Just as we do not go intermittently blind in the absence of an explicitly visual task at hand, we would not expect timing mechanisms to shut off whenever we are not tasked with judging the passage of time.

In addition, we detail the several data analysis problems that lead us to question the central conclusions of the original study (Toso, Reinartz, et al., 2021) and in its erratum (Toso et al., 2023). Upon addressing these issues, we find the familiar correlation between speed of striatal dynamics and judgments about duration, as well as a striking agreement between the degree to which different task variables impacted choices and were encoded by DLS neurons.

## 3.2 Results

### 3.2.1 Rats did not discriminate vibration durations as idealized and reported

Toso et al. performed chronic extracellular recordings from the left DLS of three rats trained in a delayed duration comparison task (Fig. 3.1). Briefly, rats were rewarded for correctly reporting which of two vibratory stimuli delivered to the right vibrissae was of a longer duration. Each trial consisted of a series of delays: a post-initiation delay of 700 ms; the S1 period; an inter stimulus interval (ISI) lasting 2 s; the S2 period; and a 500-ms post-stimulus delay (Fig. 3.1B). Animals were required to maintain their snout at the stimulus delivery port from trial initiation up until an auditory “go” cue was presented at the offset of the post-stimulus delay, at which point they were free to report their choice at one of the two lateral spouts (Fig. 3.1A).

### CHAPTER 3. MANIPULATING STIMULUS INTENSITY DURING DELAYED INTERVAL DISCRIMINATION

Importantly, in addition to the two vibrations varying along the task-relevant axis of duration, they varied along the task-irrelevant axis of intensity. Thus, individual stimulus pairs were defined on the basis of their duration ( $T1$  and  $T2$ ) and intensity ( $I1$  and  $I2$ ), and chosen such that the sets of all possible durations  $T = \{112, 161, 232, 334, 481, 694, 1000\}$  ms and intensities  $I = \{34, 64, 117\}$  mm s<sup>-1</sup> were the same for both stimuli (Hernández et al., 1997).




Figure 3.1: (A) Schematic of the behavioral apparatus with a rat depicted at the initiation/stimulus delivery port. The mapping of choice spout (left or right) to choice category ( $T2 > T1$  or  $T2 < T1$ ) was different from the depicted one for one for the DLS-implanted rats. (B) Event diagram of an example correct trial in the delayed duration comparison task.

Figure 1 shows all duration (Fig. 3.2A) and intensity pairs (Fig. 3.2B), as well as the corresponding performance pooled across the three DLS-implanted rats. Rather than annotating and/or using a diverging 3-hue color map to indicate  $P(T2 > T1)$ , as done here (Fig. 3.2) and in previous studies employing comparison tasks (Akrami et al., 2018; Fassihi et al., 2014; Romo et al., 1999), the authors denoted performance using only color in a sequential 2-hue color map. This makes it difficult to detect differences across stimulus pairs and, together with the authors' undeclared omission of the most prevalent intensity pair in the data ( $I1 = 64$ ;  $I2 = 64$  mm s<sup>-1</sup>, Fig. 3.3), creates the impression that performance is well summarized as primarily driven by normalized time difference (NTD, Fig. 3.5A, left):

$$NTD = \frac{T2 - T1}{T1 + T2}$$

and biased by normalized intensity difference (NID, Fig. 3.5C, left):

$$NID = \frac{I2 - I1}{I1 + I2}$$




Figure 3.2: (A) Proportion of  $T_2 > T_1$  choices for all stimulus duration pairs pooled across the three DLS-implanted rats. Dotted lines above, below and at the identity line represent NTD levels of +0.35, -0.35 and 0, respectively. (B) Same as (A), but for stimulus intensity and NID. The black arrow highlights the intensity pair that was undisclosed in both the original study and its erratum ( $I_1 = 64$ ;  $I_2 = 64$  mm.s<sup>-1</sup>).




Figure 3.3: Distribution of stimulus features in the delayed duration comparison task. (A) Main axes: frequency of stimulus duration pairs ( $T_1$  and  $T_2$ ) relative to the most common pair ( $T_1 = 161$ ;  $T_2 = 334$  ms). Left: Distribution of sampled  $T_2$  durations for individual DLS-implanted animals (in gray) and pooling across them (black). (B) Same as (A), but for stimulus intensity. The black arrow in the main axes highlights the intensity pair that was undisclosed in both the original study and its erratum ( $I_1 = 64$ ;  $I_2 = 64$  mm.s<sup>-1</sup>).

If animals' choices were indeed independent of  $T1$  and  $T2$  within each NTD level, it follows that the probability of reporting  $T2 > T1$  as a function of all possible values of  $T2 - T1$  (the numerator in the definition of NTD) should produce a step-like function (Fig. 3.4).

This was not the case in the data (Fig. 3.5A, B). The same rationale and prediction applies to NID and intensity difference (Fig. 3.4), and again the data did not conform (Fig. 3.5C, D). For both duration and intensity we show that  $P(T2 > T1)$  was a non-monotonic function of stimulus difference (Fig. 3.5B, D). Moreover, if animals' decisions were truly governed by NTD and NID, then the denominators in their respective definitions, the aggregate stimulus duration ( $T1 + T2$ ) and intensity ( $I1 + I2$ ), should act only as normalization terms and as such not affect performance (Fig. 3.4). Instead, we found that rats' propensity to judge  $T2$  as larger than  $T1$  increased monotonically as a function of both these quantities (insets in Fig. 3.5B, D).

Together, these observations reveal the inadequacy of NTD and NID as criteria for grouping trials when summarizing rats' performance in this task and warrant further characterization of the behavior before searching for neural correlates of animals' judgments.

### 3.2.2 The second stimulus dominated the delayed duration comparison

We have shown that NTD and NID do not adequately capture performance in this task, and that the probability of rats reporting  $T2$  as larger than  $T1$  did not steadily increase as a function of the signed difference between these durations. Though this means that rats were not doing a one-to-one comparison of nominal stimulus durations, it is entirely consistent with them doing an asymmetric comparison due to what is commonly referred to as *contraction bias* or the *central tendency effect* – the tendency for perceptual judgments to gravitate towards the center of the distribution of previous stimuli. This behavioral regularity has been observed repeatedly across animal models and stimulus modalities (Akrami et al., 2018; Fassihi et al., 2014; Hollingworth, 1910; Jazayeri & Shadlen, 2010; Romo et al., 1999; Urai et al., 2017), and is proposed to arise from a Bayesian computation (Ashourian & Loewenstein, 2011; Xiang et al., 2021) wherein noisy representations of stimulus magnitudes are combined with prior information about the corresponding distribution of magnitudes (although see Boboeva et al.,




Figure 3.4: Performance of simulated agents strictly driven by NTD and biased by NID. **(A)** Proportion of T2 > T1 choices for all duration pairs pooled across three simulated agents. Dotted lines above, below and at the identity line represent NTD levels of +0.35, -0.35 and 0, respectively. **(B)** Proportion of T2 > T1 choices as a function of NTD and grouped by NTD level (left) or stimulus pair (right). Gray markers and lines refer to individual agents, and black and white ones to cross-agent averages for each NTD level. Dashed lines on the right connect the two stimulus pairs with smaller absolute duration differences ( $|T2 - T1|$ ). **(C)** Proportion of T2 > T1 choices conditioned on all possible values of the numerator (main axes) and denominator (inset) terms in the definition of NTD. **(D-F)** Same as (A-C), but for stimulus intensity and NID. **(G)** Same data as in (A), but overlaying the performance rule that generated them. For each duration pair, a line was drawn between its T1 and T2 coordinates and the horizontally shifted point for which generative performance was closest to simulated performance on that pair. **(H)** Similar to (G), but for stimulus intensity. **(I)** Coefficients of a logistic regression model fit to simulated choice data using all four stimulus features (T1, T2, I1 and I2) as regressors.

### CHAPTER 3. MANIPULATING STIMULUS INTENSITY DURING DELAYED INTERVAL DISCRIMINATION




Figure 3.5: (A) Proportion of  $T_2 > T_1$  choices as a function of NTD and grouped by NTD level (left) or stimulus duration (right). Gray markers and lines refer to individual animals, and black and white ones to cross-animal averages for each NTD level. Dashed lines on the right connect the two stimulus pairs with smaller absolute duration differences ( $|T_2 - T_1|$ ). (B) Proportion of  $T_2 > T_1$  choices conditioned on all possible values of the numerator (main axes) and denominator (inset) terms in the NTD definition. (C-D) Same as (A-B), but for stimulus intensity and NID. The black arrows highlight the intensity pair that was undisclosed in both the original study and its erratum ( $I_1 = 64$ ;  $I_2 = 64 \text{ mm s}^{-1}$ ).

2023 for an alternative non-Bayesian explanation). Indeed we find that rats' performance was well accounted for by assuming that the effective magnitudes of T1 and I1 drifted towards the expected value of the respective sampling distributions during the inter-stimulus delay (Fig. 3.6). The authors acknowledge that this phenomenon may have led to animals' high error rate for the stimulus pair with the shortest possible T1 (T1 = 112; T2 = 232 ms, Fig. 3.6A), but in doing so underplay its ubiquity by suggesting that there is something unique about this particular stimulus duration pair.




Figure 3.6: **(A)** Depiction of how contraction bias may have affected estimates of T1 in the delayed duration comparison. Same data and analysis as in Fig. 1C overlaying an hypothetical performance rule where the probability of reporting  $T_2 > T_1$  is a monotonic (sigmoidal) function of effective NTD. For each duration pair, a line was drawn between its T1 and T2 coordinates and the horizontally shifted point for which hypothetical performance was closest to rats' performance on that pair. The vertical solid line depicts the expected value of the T1 distribution, and the black arrow indicates the one stimulus pair the authors singled-out as potentially affected by contraction bias. **(B)** Similar to (A), but for stimulus intensity.

To quantify the asymmetry in the contribution of duration and intensity of the two stimuli onto animals' judgments, we fit a logistic regression model using all four stimulus features (T1, T2, I1 and I2) as regressors and choice category ( $T_2 > T_1$ ) as the response variable (Fig. 3.7). Consistent with a strong contraction bias, T2 dominated the comparison, with an associated regression weight more than five times that of T1. Moreover, I1 had virtually no bearing on rats' decisions, meaning that the intensity bias reported by the authors (Fig. 3.5C, left) was driven by I2 alone.

Accordingly, plotting discrimination performance as a function of T2




Figure 3.7: Coefficients of a logistic regression model fit to choice data using all four stimulus features (T1, T2, I1 and I2) as regressors. Gray markers refer to fits for individual animals, and black ones to a fit pooling data across all three DLS-implemented animals.

produced monotonic functions (Fig. 3.8A), that systematically differed if conditioned on I2 (Fig. 3.8B) but not when conditioned on I1 (Fig. 3.8C). Lastly, fitting intensity-split choices with 4-parameter psychometric functions followed by model comparison confirmed the lack of detectable changes on parameters of psychophysical interest due to I1 and indicated that the effect of I2 could not be captured by any single psychometric parameter on its own (Fig. 3.8B, C, 3.9).

Beyond providing a more accurate summary of a rich behavior, these results inform subsequent analyzes of neural activity and constrain their interpretations. For example, lack of modulation of DLS neural activity by I1 takes on a different meaning after establishing that I1 did not meaningfully impact animals' reports (Fig. 3.7, 3.8C, 3.9) – the main behavioral covariate of internal estimates of elapsed time provided by the experimental paradigm. Likewise, the striking difference in the degree to which T1 and T2 predicted choices might warrant a similar asymmetry in emphasis on neural activity aligned to the first and second stimuli (Fig. 3.7). Additionally, comparing neural activity recorded during nominally correct and incorrect trials is not readily interpretable due to the non-monotonic relationship between performance and experimental difficulty (Fig. 3.5B). Lastly, the fact that that extreme-I2 conditions did not simply produce laterally shifted, temporally or vertically rescaled versions of the intermediate-I2 psychometric curve (Fig. 3.8B, 3.9) argues against intensity affecting choices exclusively through




Figure 3.8: (A) Proportion of  $T_2 > T_1$  choices as a function of  $T_2$ . Gray markers refer to individual animals, and black ones data pooled across all three DLS-implanted animals. (B) Pooled discrimination performance as a function of  $T_2$  and conditioned on  $I_2$ . Main axes: psychometric functions fit to  $I_2$ -split choices. Marker size represents the relative number of trials per  $T_2$  value normalized within each  $I_2$  condition. Inset: differences between choice proportions in extreme  $I_2$  (high:  $I_2 = 117 \text{ mm.s}^{-1}$ ; low:  $I_2 = 34 \text{ mm.s}^{-1}$ ) and intermediate  $I_2$  trials ( $I_2 = 64 \text{ mm.s}^{-1}$ , mean  $\pm$  propagated SEM). (C) Same as (B), but conditioned on  $I_1$ .

any *one* property of the underlying decision variable (e.g., its temporal scale).

For these reasons, we focused most of the analyzes below on DLS activity recorded during the presentation of the second stimulus, refrained from using experimentally-defined correctness as a contrast when assessing if DLS responses covaried with duration judgments, and did not share in the authors' assumption that temporal scaling of population dynamics is the only possible neural correlate of the observed  $I_2$ -induced behavioral bias.

### 3.2.3 Multiple features of DLS dynamics were systematically affected by the intensity of the second stimulus

Having established that rats' behavior can be better described as categorization of  $T_2$  biased by  $I_2$  than discrimination of the two stimulus durations biased by their intensities, we next sought to characterize the impact of  $I_2$  on DLS activity at the single-cell and population levels during the second stimulus period. In their original publication, the authors reported that “None of the neurons recorded [in the DLS] was found to encode  $I$  [intensity] by its firing rate” (Toso, Reinartz, et al., 2021), and this was later amended to “2% of neurons encoded stimulus intensity during stimulus 2 presentation” in their erratum (Toso et al., 2023). Viewing neural activity aligned to the onset of




Figure 3.9: Multiple parameters of the psychometric function were affected by the intensity of the second stimulus, but not that of the first. **(A)** Schematic illustrating the parameterization of the psychometric function used throughout (adapted from H. H. Schütt et al., 2016). **(B)** Model comparison between psychometric fits to choices of individual rats in middle- $I_2$  ( $I_2 = 64 \text{ mm s}^{-1}$ ) and extreme- $I_2$  trials (high:  $I_2 = 117 \text{ mm s}^{-1}$ ; low:  $I_2 = 34 \text{ mm s}^{-1}$ ), measured as the difference between each model's BIC score and that of the rat-specific model with the lowest BIC across low and high intensities. Extreme- $I_2$  models are labeled according to which parameters of the psychometric function were free and which were set to those of the corresponding middle- $I_2$  fits (e.g., “none” means that all parameters were free, and “ $m$ ” means that only the threshold parameter was fixed). Small and large markers correspond to individual animals and cross-animal averages (mean  $\pm$  SEM), respectively. The dashed black line corresponds to the  $I_2$ -agnostic cross-animal average, according to which models were sorted along the  $x$ -axis. **(C)** Same as (B), but for  $I_1$ . **(D)** Point estimates for each of the four parameters of the psychometric function introduced in (A) estimated using only middle-intensity ( $x$ -axis) or extreme-intensity trials ( $y$ -axis). Small circles correspond to individual animals, and crosses to cross-animal averages (mean  $\pm$  SEM).

the second stimulus and conditioned on its intensity, both as population averages (Fig. 3.10A) and example neuronal responses (Fig. 3.10B), suggests a much higher-than-reported prevalence and strength of I2 modulation in the DLS. This is true in spite of cross-neuron averages being, in general, poor indicators of overall modulation, owing to their inability to distinguish between lack of modulation and symmetry in up- and down-modulation across the population. We will return later in the manuscript to this issue to explain why the authors reported so few neurons correlated with intensity.




Figure 3.10: **(A)** Average normalized population activity conditioned on I<sub>2</sub> and aligned to the onset of S<sub>2</sub> (mean  $\pm$  SEM across neurons). **(B)** Responses of four example DLS neurons during S<sub>2</sub>. For each neuron, spike density functions (mean  $\pm$  SEM across trials) are conditioned on I<sub>2</sub> and plotted above its spike raster, where trials are first sorted by I<sub>2</sub>, and then by T<sub>2</sub> within each I<sub>2</sub> level.

To further examine how I<sub>2</sub> affected the time evolution of population activity of the DLS during the S<sub>2</sub> period, we turned to dimensionality reduction using PCA. Briefly, we computed PCs using unconditioned S<sub>2</sub>-aligned average activity, and then projected average activity conditioned on I<sub>2</sub> onto each of the first three PCs (Fig. 3.11A). This confirmed that a substantial amount of variance in DLS responses recorded during the presentation of the second stimulus is indeed explained by I<sub>2</sub>, and at least partially along the same intrinsic dimensions of neural space as T<sub>2</sub>-driven variance. Interestingly, I<sub>2</sub> seemed to affect multiple features of DLS activity during this period, scaling it both in time and amplitude. This is already visible at the level of single-neuron responses (Fig. 3.10B) and PC projections (Fig. 3.11A), but is most evident when looking at neural trajectories in PC

space (Fig. 3.11B), which seem to systematically vary with I2 in both speed and eccentricity. This dual effect of I2 on the temporal scale and overall gain of neural activity is reminiscent of its effect on multiple parameters of the psychometric function (Fig. 3.12A, 3.9), suggesting that the DLS may contribute to rats' decisions in this task through dissociable features of its population dynamics.




Figure 3.11: (A) I2-split population activity projected onto each of the first three PCs computed using I2-agnostic population activity aligned to the onset of S<sub>2</sub> (padded with 500 ms before S<sub>2</sub> onset). (B) Population activity conditioned on I2 projected onto a 2D section of the subspace defined by the three PCs shown in (A). The 3-axis gizmo at the center depicts the projections of the axes defined by the first three PCs onto the chosen 2D section, the dashed black arrow indicates the direction of time, and the gray patches group population states recorded at each possible stimulus offset across I2.

Next, to characterize the contributions of intensity modulation of striatal activity onto the pattern of rats' choices conditional on I2 and across T2, we performed cross-validated decoding of stimulus discrimination categories using neural population activity recorded at the end of S<sub>2</sub> during middle-intensity trials. Specifically, we used half of all intermediate-I2 ( $I2 = 64 \text{ mm s}^{-1}$ ) trials to train a linear classifier to discriminate the durations of the two stimuli (T1 and T2) based only on population spike counts recorded in the last 112 ms (duration of the shortest possible T2) before S<sub>2</sub> offset. Neural "judgments" in all remaining trials (including all low- and high-intensity trials), were then simply the classes ( $T2 > T1$  or  $T2 < T1$ ) predicted by the classifier using those trials' population state just prior to stimulus offset. The resulting neurometric curves split by I2 (Fig. 3.12A) provided

a striking qualitative match to their psychometric counterparts (Fig. 3.8B), thus confirming that DLS population state nearing the offset of the second stimulus covaried with intensity-biased judgments for all possible stimulus durations.

Lastly, to characterize temporal scaling of population activity in DLS due to I2-driven modulation, we decoded elapsed time continuously from the population state during S2 presentation using a cross-validated naïve Bayes decoder. To do so, we first characterized the “typical” temporal profiles of striatal responses using a subset of all middle-intensity trials, and then applied a probabilistic decoding approach to estimate elapsed time based only on the observed state of the recorded population in remaining trials. Estimates of elapsed time derived from ongoing population activity systematically led ahead and lagged behind true time during high- and low-I2 trials, respectively (Fig. 3.12B).




Figure 3.12: (A) Neurometric curves trained using S2-aligned population activity from middle-I2 trials conditioned on I2 (main axes), and differences between the proportion of  $T_2 > T_1$  neural “reports” in high ( $I_2 = 117 \text{ mm s}^{-1}$ ) and low ( $I_2 = 34 \text{ mm s}^{-1}$ ) intensity trials compared to middle intensity ( $I_1 = 64 \text{ mm s}^{-1}$ ) trials (inset, mean  $\pm$  propagated SEM). (B) Time estimates decoded from DLS activity around S2 presentation conditioned on I2. Main axes: mean posterior distributions across runs. Inset: maximum a posteriori (MAP) point estimates averaged across runs (median and IQR).

This decoding approach is something the authors tried as well, though with a different outcome and subsequent interpretation. The discrepancy results from a combination of the low-level implementation errors detailed in the supplementary text and a more conceptual error specific to the design of this analysis – the authors balanced trials for choice on a per stimulus basis

when training their decoder, which amounts to training on flat psychometric curves, effectively baking into the decoder an insensitivity to features that correlate with choice (such as intensity).

Altogether, our findings demonstrate that DLS dynamics and animals' duration judgments were similarly modulated by the intensity of the second stimulus.

### **3.2.4 DLS neural responses displayed negligible sensitivity to the intensity of the first stimulus**

We observed that the intensity of the first stimulus had no systematic net effect on animals' duration judgments (Fig. 3.7, 3.8C, 3.9). As such, the possibility that this stimulus feature affected DLS activity past S1 and the ISI and into the S2 period would be hard to reconcile with the DLS being a substrate for said I1-independent decisions. To address this, and for ease of comparison, we applied the exact same treatment used to characterize the impact of I2 on S2-aligned DLS activity in the previous section (Fig. 3.10, 3.11, 3.12) to I1-conditioned activity recorded during S2. We did not observe any orderly relationship between I1 and S2-aligned: population averages (Fig. 3.13A); single-neuron responses (Fig. 3.13B, same examples from Fig. 3.10B); PC projections (Fig. 3.14A); PC trajectories (Fig. 3.14B); neurometric curves (Fig. 3.15A); and decoded time estimates (Fig. 3.15B).

Furthermore, I1 modulation of DLS dynamics during the S1 period was also virtually non-existent, if not entirely absent (Fig. 3.16), thus standing in stark contrast to the effect of I2 on S2-aligned activity (Fig. 3.10, 3.11, 3.12).

Since there was nothing experimentally different about the delivery of the two stimuli, this striking asymmetry argues against the DLS faithfully encoding all aspects of an animal's immediate sensory state, and in favor of its activity reflecting a learned mapping between states and actions. It should be noted that the authors reported this uneven effect of intensity on DLS activity during the two stimuli in their erratum (Toso et al., 2023), though their qualitatively correct observation was unsatisfactorily accounted for in the absence of the accompanying characterization of the asymmetric impact of I1 and I2 on behavior.

Taken together, these data speak to a remarkable indifference of DLS responses to the intensity of the first stimulus relative to that of the second.




Figure 3.13: **(A)** Average normalized population activity conditioned on I1 and aligned to the onset of S2 (mean  $\pm$  SEM across neurons). **(B)** Responses of four example DLS neurons during S2. For each neuron, spike density functions (mean  $\pm$  SEM across trials) are conditioned on I1 and plotted above its spike raster, where trials are first sorted by I1, and then by T2 within each I1 level.




Figure 3.14: **(A)** I1-split population activity projected onto each of the first three PCs computed using I1-agnostic population activity aligned to the onset of S2 (padded with 500 ms before S2 onset). **(B)** Population activity conditioned on I1 projected onto a 2D section of the subspace defined by the three PCs shown in (A). The 3-axis gizmo at the center depicts the projections of the axes defined by the first three PCs onto the chosen 2D section, the dashed black arrow indicates the direction of time, and the gray patches group population states recorded at each possible stimulus offset across I1.




Figure 3.15: (A) Neurometric curves trained using S2-aligned population activity from middle-I1 trials conditioned on I1 (main axes), and differences between the proportion of  $T_2 > T_1$  neural "reports" in high ( $I_1 = 117 \text{ mm s}^{-1}$ ) and low ( $I_1 = 34 \text{ mm s}^{-1}$ ) intensity trials compared to middle intensity ( $I_1 = 64 \text{ mm s}^{-1}$ ) trials (inset, mean  $\pm$  propagated SEM). (B) Time estimates decoded from DLS activity around  $S_2$  presentation conditioned on I1. Main axes: mean posterior distributions across runs. Inset: maximum a posteriori (MAP) point estimates averaged across runs (median and IQR).

### 3.2.5 The speed of DLS dynamics covaried with duration judgments

Systematic correlations between the speed of striatal dynamics and interval timing have been observed across laboratories, animal models, interval timing paradigms and striatal domains (Gouvêa et al., 2015; J. Kim et al., 2018; Mello et al., 2015; Monteiro et al., 2023; J. Wang et al., 2018). Adding to the evidence provided by these correlative studies, our group has recently published evidence in support of this link being causal, by using temperature to slow down and speed up population dynamics in the dorsal striatum of rats and observing corresponding biases towards short and long duration judgments (Monteiro et al., 2023). Perhaps then the most surprising of the central claims of Toso and colleagues was that there was no relationship between DLS activity and animals' duration discrimination. Importantly, this was inferred from the overlap in time estimates decoded from S1- and S2-aligned DLS activity recorded during correct and incorrect trials. This analysis is problematic not just because nominal correctness can not be used to meaningfully dichotomize rats' behavior (Fig. 3.5) and because rats mostly used S2 to guide choice (Fig. 3.7), but more importantly because, even




Figure 3.16: DLS dynamics recorded during the first stimulus were negligibly affected by the intensity of the first stimulus. **(A)** Average normalized population activity conditioned on I1 and aligned to the onset of S1 (mean  $\pm$  SEM across neurons). **(B)** Responses of four example DLS neurons during S1. For each neuron, spike density functions (mean  $\pm$  SEM across trials) are conditioned on I1 and plotted above its spike raster, where trials are first sorted by I1, and then by T1 within each I1 level. **(C)** I1-split population activity projected onto each of the first three PCs computed using I1-agnostic population activity aligned to the onset of S1 (padded with 500 ms before S1 onset). **(D)** Population activity conditioned on I1 projected onto a 2D section of the subspace defined by the three PCs shown in (C). The 3-axis gizmo at the center depicts the projections of the axes defined by the first three PCs onto the chosen 2D section, the dashed black arrow indicates the direction of time, and the gray patches group population states recorded at each possible stimulus offset across I1. **(E)** Neurometric curves trained using S1-aligned population activity from middle-I1 trials conditioned on I1 (main axes), and differences between the proportion of T2 > T1 neural "reports" in high (I1 = 117 mm s<sup>-1</sup>) and low (I1 = 34 mm s<sup>-1</sup>) intensity trials compared to middle intensity (I1 = 64 mm s<sup>-1</sup>) trials (inset, mean  $\pm$  propagated SEM). **(F)** Time estimates decoded from DLS activity around S1 presentation conditioned on I1. Main axes: mean posterior distributions across runs. Inset: maximum a posteriori (MAP) point estimates averaged across runs (median and IQR).

assuming rats performed as idealized, of how correct and incorrect trials map onto over- and underestimation of duration changes across time within each stimulus period. Incorrectly discriminating stimulus pairs for which T2 is longer than T1 presumably corresponds to faster-than-usual timekeeping during S1 and/or slower-than-usual timekeeping during S2, whereas the opposite is true for errors when T2 is shorter than T1 (Fig. 3.17). The same logic holds in reverse for correct trials. Therefore, pooling decoder output conditioned on trial outcome indiscriminately across stimulus pairs would obscure these signatures within each category for all but extreme T1 and T2 values (equal to or longer than 694 ms), since these durations signal which choice is correct unambiguously and do not overlap in time with any other (ambiguous) durations. However, since the authors had to exclude these trials due to data scarcity, their analysis was incapable of detecting even this remnant of the correlation between time as encoded by the DLS and rats' choices.

To circumvent this issue, we instead conditioned neural activity on whether rats judged T2 to be greater than T1, which should partition trials into two groups with consistent opposite underlying biases in timekeeping during the presentation of S2. We thus subjected S2-aligned DLS activity conditioned on choice to the same battery of analyses we used for I2 (Fig. 3.10, 3.11, 3.12) and I1 (Fig. 3.13, 3.14, 3.15, 3.16). Consistent with previous findings, we saw that DLS responses covaried with animals' judgments. Interestingly, choice modulation, though present, was perhaps not as evident as that driven by I2 when looking at population mean activity (Fig. 3.18A) or responses from the individual neurons we have been showcasing throughout (Fig. 3.18B). The reason for this became apparent with PCA (Fig. 3.19), which revealed that unlike intensity-, decision-related variability seems to express itself almost exclusively through the speed of DLS dynamics. This was corroborated by both the horizontal shift in neurometric curves (Fig. 3.20A) and the steady divergence between decoded time estimates conditioned on animals' reports (Fig. 3.20B), which can be understood as the accumulating effects of slow-biased (fast-biased) neural trajectories preceding  $T2 < T1$  ( $T2 > T1$ ) reports resulting in population states at S2 offset that increasingly lag behind (lead ahead) their  $T2 > T1$  ( $T2 < T1$ ) counterparts.

Together these results demonstrate that judgments about the duration of the second stimulus could be reliably predicted by the speed with which patterns of DLS activity evolved during its presentation.




Figure 3.17: Systematic biases in timekeeping across stimulus pairs are obscured (revealed) by conditioning trials on the outcome (category) of choices. (A) Hypothetical distributions of internal estimates of elapsed time as a function of experimental time since the onset of the first (left) and second (right) stimulus, plotted both continuously (colored gradient in the background) and discretely at the set of possible stimulus offset times (white distributions in the foreground). (B) Joint distributions of internal estimates of T1 and T2 for two pairs of stimuli with their respective marginals projected on each axis. (C) Joint and marginal distributions of T1 and T2 estimates for a stimulus pair for which T2 was shorter (left) or larger (right) than T1, conditioned on whether a direct comparison of those estimates would produce correct or incorrect choices. (D) Same as (C) but for all stimulus pairs. (E-F) Same as (C-D) but conditioned on the category of hypothetical choices strictly informed by a comparison of T1 and T2 estimates.




Figure 3.18: **(A)** Average normalized population activity conditioned on choice and aligned to the onset of S2 (mean  $\pm$  SEM across neurons). **(B)** Responses of four example DLS neurons during S2. For each neuron, spike density functions (mean  $\pm$  SEM across trials) are conditioned on choice and plotted above its spike raster, where trials are first sorted by choice, and then by T2 within each choice category.




Figure 3.19: **(A)** Choice-split population activity projected onto each of the first three PCs computed using choice-agnostic population activity aligned to the onset of S2 (padded with 500 ms before S2 onset). **(B)** Population activity conditioned on choice projected onto a 2D section of the subspace defined by the three PCs shown in (A). The 3-axis gizmo at the center depicts the projections of the axes defined by the first three PCs onto the chosen 2D section, the dashed black arrow indicates the direction of time, and the gray patches group population states recorded at each possible stimulus offset across choice categories.




Figure 3.20: (A) Neurometric curves trained using  $S_2$ -aligned population activity from choice-agnostic trials conditioned on choice. (B) Time estimates decoded from DLS activity around  $S_2$  presentation conditioned on choice. Main axes: mean posterior distributions across runs. Inset: maximum a posteriori (MAP) point estimates averaged across runs (median and IQR).

### 3.2.6 The *relative* proportion of DLS neurons modulated by different task variables matched their relative impact on performance

A surprising difference between our analyses and those presented in the original paper and its correction is in the number of neurons identified as modulated by stimulus intensity. We found this discrepancy to be driven by a combination of issues related to the particular statistical criteria the authors enforced.

It is increasingly appreciated that single neuron activity is often modulated by multiple factors (Kobak et al., 2016; Mante et al., 2013; Rigotti et al., 2013), indicating that to fully understand a neuron’s response properties, both sound experimental design that unconfounds task variables and analysis methods that quantify their relative contributions are necessary. In this section we provide a summary for how different task variables affected the activity of DLS neurons at different moments in the task, ultimately more focused on I2 during  $S_2$  to explore a technical issue regarding the operational definition of intensity-selective neurons used by Toso and colleagues – for a neuron to be considered as modulated by intensity, spike counts aligned to the onset of either stimulus were required to be significantly correlated with intensity (of that stimulus) for at least five consecutive and non-overlapping

100-ms bins. We find that these criteria are statistically underpowered to detect cells modulated by intensity in light of the typical firing statistics of striatal neurons (Steiner & Tseng, 2016) and the total amount of data available. The nature of this task (and most other sensory timing tasks) is such that the number of trials recorded during the presentation of a given stimulus decreases at discrete boundaries defined by the set of its possible durations. This means that with the authors' approach, statistical power drops sharply with each subsequent regression, to the point where only trials corresponding to the two longer stimulus durations (694 and 1000 ms) can contribute past the fourth bin-wise regression. Yet, all regressions were treated equally in how they contributed to a neuron being labeled as modulated by intensity. Furthermore, considering the characteristically low firing rates of medium spiny neurons (MSNs, the vast majority of striatal neurons, Steiner and Tseng, 2016), the response variable – spike counts recorded within 100 ms – is effectively varying between 0 and 2, thus impairing statistical power even further.

To try and navigate these constraints while facilitating a direct comparison with the authors' approach, we fit separate linear regression models to each neuron's spike counts aligned to multiple task events (trial initiation, S1, ISI, S2 and "go" cue), and calculated the proportions of neurons that were significantly up- and down-modulated by the different regressors (Fig. 3.21). Crucially, instead of regressing spike counts against I1 or I2 separately, we used a multivariable design matrix that included not just I1 and I2, but also T1 and T2 as well as choices and rewards in the current and previous trials. This was done so as to account for any systematic (Fig. 3.2) as well as spurious correlations between predictors, and thus better parse variance that reflected direct modulation by each stimulus feature independently of one another as well as of premotor-, motor-, and reward-related signals.

As mentioned above, in the specific case of stimulus epochs, the fact that most neurons were only recorded for a few trials for any given combination of vibration durations and intensities means that as time relative to the onset (offset) of the stimulus increases (decreases), the number of trials that continue to meet the selection criteria up to (from) that point decreases. This leads to a trade off in the selection of the *spike integration window* for this analysis. Longer (shorter) windows will lead to a larger (smaller) range in the support of the distribution of spike counts over trials (Fig. 3.23 up), but at the cost of fewer (more) trials contributing samples to that distribution




Figure 3.21: Top: proportion of DLS neurons that were either up- or down-modulated by several experimental and behavioral factors across multiple task epochs. Signed modulation was assessed by quantifying the proportion of significant positive and negative coefficients in multivariable linear regression models fit to spike counts recorded in a 334-ms window aligned to the onset and/or offset of multiple task events. The "~" in ~reaction and ~choice refers to the approximate nature of the alignment for these epochs (334 ms and 668 ms after go cue for reaction and choice, respectively), which unlike all the others, lacked trial-to-trial annotated event times. The factor/regressor cross-correlation matrix is shown as an inset in the top left. Bar colors are factor-specific and respect all previously introduced color schemes (where applicable), and bar opacity refers to the significance level used to assess modulation (1% and 5% for opaque and transparent colors, respectively). Bottom: same as above, but highlighting modulation by I1 and I2.

### CHAPTER 3. MANIPULATING STIMULUS INTENSITY DURING DELAYED INTERVAL DISCRIMINATION




Figure 3.22: (A) Total proportion of DLS neurons that were up- or down-modulated by I1 (in gray) and I2 (in black) across the different task epochs. (B) Same as (A), but for stimulus durations.

(Fig. 3.23 bottom).

To both inform the choice of this parameter and create negative and positive controls for this analysis, we generated surrogate spike data with similar firing statistics to those found in the real data. Briefly, for each neuron we first computed the cross-trial average spike density function aligned to the onset of S2 using only intermediate-I2 trials ( $I2 = 64 \text{ mm s}^{-1}$ ). This average response was then either used directly as the generative spike rate function for all trials in the case of negative controls, or first multiplied by a gain parameter proportional to I2 in the case of positive controls (Fig. 3.24A). The resulting rate functions were then used to sample spike times from inhomogeneous Poisson point processes on a trial-by-trial basis. Running our regression analysis with these surrogate data allowed us to estimate false positive (type I errors) and negative rates (type II errors) for a range of spike integration windows, which suggested 334 ms as the optimal time window given the previously mentioned constraints imposed by data scarcity (Fig. 3.23). Even with this optimized window, the estimated false negative




Figure 3.23: Oposing effects of the duration of the spike integration window used in Fig. 3.21 on the average number of recorded spikes and the number of trials contributing to that estimate during S2 presentation epochs. Top: average number of spike counts across trials as a function of spike integration window for the intermediate I2 condition. Bottom: number of trials with recorded spike counts at each time point relative to the onset (left) or offset (right) of the second stimulus, split by I2 condition (median and IQR across neurons).

rate was close to 70% (80%) at a significance threshold of 5% (1%) (Fig. 3.24B), meaning that our detected 9.7% (4.7%) of I2-modulated neurons (Fig. 3.21B, left) may actually reflect that as much as 32.3% (23.5%) of DLS neurons were modulated by the intensity of the second stimulus during its presentation.

Granted, these error rates are intimately tied to the nature and magnitude of intensity modulation we simulated in our positive controls, which, though tuned to approximately match effect sizes seen in the data (compare Fig. 3.10B and Fig. 3.24A), were not explored systematically. Still, it is worth stressing that for simplicity, in our positive controls intensity had a multiplicative effect on the amplitude, and not the temporal scale, of generative spike rates. Modulation through temporal scaling would arguably be even harder to detect, due to the need for its effects to accumulate over time, and how that interacts with the number of contributing trials decreasing as a function of elapsed time since stimulus onset (Fig. 3.23 bottom). Case in point, the proportion of neurons modulated by the upcoming choice, which we have shown to manifest almost entirely in the speed of population dynamics (Fig. 3.19, 3.20), seem comparatively low during S2 (Fig. 3.21A). This means that the estimated false negative rates may well be underestimating the true rate of type II errors in our quantification. For all of these reasons, instead




Figure 3.24: (A) I2-conditioned generative spike density functions used in simulated positive controls for the same four example DLS neurons shown in previous figures, aligned to the onset (left) or offset (right) of the second stimulus. (B) False negative (top) and false positive (bottom) rates computed by calculating the proportion of I2-modulated neurons during the S2 presentation epochs averaged over two separate sets of 25 runs corresponding to positive and negative controls. For each run 100% of the simulated neuronal responses were either positively modulated by I2 (when computing false negative rates) or not modulated by I2 (when computing false positive rates).

of reporting an “exact” fraction of intensity-modulated neurons, that would be inextricably linked to arbitrary parameters in its calculation (e.g., a spike integration window lasting  $x$  ms and a significance level of  $y$ ), cell type-specific firing statistics, and what features of their responses were affected, we emphasize that it was definitely not zero, nor was it two percent.

Regarding modulation by other task variables, such as I1, choice and reward, we again find that the authors’ estimates were systematically below ours (even before accounting for false negative rates, which would have to be estimated separately on an epoch-by-epoch basis), but would once more caution against statements about the absolute fraction of neurons displaying sensitivity to a particular task factor. However, *relative* statements strike us as potentially meaningful, depending on the question at hand. For example, intensity modulation was much more apparent during and after the second stimulus than during and after the first, where it was comparatively absent (Fig. 3.21).

In yet another example, DLS activity strongly reflected the duration of S1 early in the ISI period, but the proportion of neurons carrying information about which T1 had been presented seemed to steadily decrease across

subsequent epochs, to the point where it was roughly between one third and one sixth (depending on significance thresholds) the proportion of neurons modulated by T2 right after the offset of S2 (Fig. 3.21A, 3.22B). This ratio of neural modulation is reminiscent of the ratio of contribution of these variables onto rats' decisions (Fig. 3.7), and is consistent with previous proposals that contraction bias arises not at the encoding stage (during S1), but gradually up to the decision stage (presumably during or shortly after S2, Ashourian and Loewenstein, 2011).

Together, these results reiterate a fundamental aspect of the data – I2 modulation was prevalent among DLS neurons – and moreover, they underscore the accordance between the extent to which different stimulus features affected temporal judgments and modulated neural activity across the DLS.

### **3.2.7 Differences in time decodability from putative clusters of ramping and non-ramping DLS neurons were biased by confounds unrelated to those category labels that were introduced by categorization criteria**

Two of the seven main figures in Toso, Reinartz, et al., 2021 were devoted to the issue of *ramping* vs. *non-ramping* neurons in the DLS, and how time estimates decoded from ramps were more reliable than those decoded from non-ramps. While we understand in theory the interest in ramping profiles due to long-existing computational accounts of timing as resulting from an integration process (Gibbon, 1991; Treisman, 1963), their existence as a distinct functional class within neural populations is by no means axiomatic. Yet, like others before them (Donnelly et al., 2015; Emmons et al., 2017), the authors assumed these categories exist and applied several ad-hoc criteria to classify neurons into the two groups without first providing evidence for the clusterability of their activity profiles. Additionally, though the differential time encoding of putative ramps and non-ramps was ultimately only assessed during stimulus presentation periods, these were not the epochs used in the clustering steps. Instead, the authors first enforced their classification criteria in two 1-s windows centered on the onset and offset of each stimulus, and later labeled the union of up- and down-ramping neurons identified in those stimulus-flanking windows as stimulus period ramps (and all remaining neurons as non ramps, Fig. 3.25 bottom). In other words, a neuron was considered to ramp during the stimulus period if it

ramped (up- or down-) from 500 ms before to 500 ms after stimulus onset and/or offset, leading to an odd set of possibilities: a neuron ramping exclusively before or exclusively after stimulus presentation will be labeled a stimulus *period* ramp, regardless of its response profile during stimulus presentation; a neuron with a unimodal response profile (a “bump”) peaking half way through a given stimulus period will be considered a ramp (since its activity increases going forward from stimulus onset and decreases going backward from its offset). Evidence for this last scenario at play can be seen in the contrasting shapes of the distributions of temporal tuning (see methods) of ramps during the task epochs in which these were first identified – bimodal – and during stimulus epochs – unimodal (Fig. 3.25 top). In all cases, these were statistically different from the strictly unimodal temporal tuning distributions of non-ramps, even during stimulus epochs (Fig. 3.25 top).

In addition to the issue of *when* classification criteria were applied, we found the criteria themselves problematic considering that the purpose of this analysis was to assess differences in time decodability from activities in the two resulting classes. First, to be classified as a ramp, a neuron’s firing rate was required to be stereotyped over trials. Were this criterion used to select cells for further analysis in general, we see no issue. However, non-stereotypical cells were instead placed in the non-ramping category. Unsurprisingly, this produced a marked asymmetry in the stereotypy coefficients (see methods) of ramps and non-ramps during both S1 and S2 (Fig. 3.26 left). It then becomes circular that decoding performance, which hinges on the reproducibility of neural responses across trials, was better with ramping than with non-ramping neurons. Second, as a proxy for a monotonicity requirement, the average firing rate of ramping neurons was required to be significantly correlated with elapsed time, and with a linear regression slope that crossed an arbitrary threshold in *absolute* value. This last detail is again problematic, as it selects for neurons with high dynamic range in their firing rate over those with a relatively lower range (Fig. 3.26 right). Intuitively, this translates into systematically poorer temporal informativeness in non-ramps (just as an image with an 8-bit color depth can resolve fewer colors than a 24-bit one). Additionally, because firing rate ranges correlate with overall levels of firing rate (data not shown), lower ranges also translate into higher variability in spike timing, which further hinders (benefits) precision when decoding time from the responses of non-ramps (ramps).




Figure 3.25: Bottom: proportion of DLS neurons labeled as "ramping" (up and down, where applicable) and "non-ramping" using the classification and grouping criteria from Toso and colleagues. Task epochs within which classification criteria were first enforced by the authors are washed out, and stimulus epochs used to subsequently assess decoding performance for each cluster are vivid. "Ramps" in each stimulus period (S1 and S2) do not preserve their initial direction qualifier because they were defined as the union of up- and down-ramping labels across the respective onset and offset epochs (an S2 "ramp" might have ramped up during the S2 onset epoch and down during S2 offset epoch). Top: cluster-split distributions of temporal tuning across DLS responses recorded during 1-s windows centered on each task event (washed out), or aligned to the onset of each stimulus period (vivid). Statistical departure from unimodality was assessed within cluster using Hartigan's dip test, and differences between clusters were assessed using two-sample Kolmogorov-Smirnov tests with Bonferroni correction for multiple comparisons (NS:  $p > 0.05$ , \*:  $p < 0.05$ , \*\*:  $p < 0.01$ ).

To validate these intuitions, we rigorously applied the authors' operational definitions of ramping and non-ramping neurons to a population of simulated Gaussian bumps with variable temporal tuning (when they peak on average), stereotypy (how much peak times vary across trials), and gain (how high are those peaks relative to baseline), each controlled by independently sampled parameters specific to each simulated neuron (Fig. 3.27).

As expected, simulated bumps labeled as ramps were systematically associated with higher stereotypy and gain parameters, as well as with distributions in temporal tuning parameters that were bimodal in flanking epochs (in which classification occurred) and unimodal in the epoch of interest (in which decoding performance was evaluated, Fig. 3.28). Despite this difference in shape, and similarly to what was observed in DLS responses, temporal tuning distributions during the decoding epoch were different for




Figure 3.26: Bottom: proportion of DLS neurons labeled as "ramping" (up and down, where applicable) and "non-ramping" using the classification and grouping criteria from Toso and colleagues. Task epochs within which classification criteria were first enforced by the authors are washed out, and stimulus epochs used to subsequently assess decoding performance for each cluster are vivid. "Ramps" in each stimulus period (S1 and S2) do not preserve their initial direction qualifier because they were defined as the union of up- and down-ramping labels across the respective onset and offset epochs (an S2 "ramp" might have ramped up during the S2 onset epoch and down during S2 offset epoch). Top: cluster-split distributions of temporal tuning across DLS responses recorded during 1-s windows centered on each task event (washed out), or aligned to the onset of each stimulus period (vivid). Statistical departure from unimodality was assessed within cluster using Hartigan's dip test, and differences between clusters were assessed using two-sample Kolmogorov-Smirnov tests with Bonferroni correction for multiple comparisons (NS:  $p > 0.05$ , \*:  $p < 0.05$ , \*\*:  $p < 0.01$ ).

the two groups (Fig. 3.25, 3.28).

Altogether, these selection biases greatly impaired the accuracy and precision of time decoding in the non-ramping category relative to the ramping one (Fig. 3.29A, B, left), though they were each sufficient individually to produce significant changes in decoding performance (Fig. 3.29B, left). This exercise demonstrates how the authors' results, which we also replicate here (Fig. 3.29B, right), derive from artifacts induced by their selection criteria unrelated to the original objective of identifying neurons with different temporal response profiles. Put differently, the criteria pre-determined the result. Incidentally, it bears noting that in line with the asymmetric contribution of S1 and S2 onto rats' decisions, decoding performance (agnostic to cluster identities) was significantly better during S2 than S1 (Fig. 3.29B, right), likely




Figure 3.27: Specification of the model used to simulate Gaussian bump responses (top) and generative spike rates in ten example trials of two pairs of simulated neurons that varied either in their stereotypy (bottom left) or gain (bottom right) parameters.




Figure 3.28: Distributions of temporal tuning (top), stereotypy (middle) and gain (bottom) parameters agnostic to (in white) or split by whether the corresponding simulated neurons were labeled as "ramps" or "non-ramps" using the authors' criteria. Horizontal gray and black bars above temporal tuning distributions at the top represent the epochs in which Toso et al. enforced their clustering criteria and evaluated decoding performance, respectively. Statistical departure from unimodality was assessed for the temporal tuning of "ramps" using Hartigan's dip test (NS:  $p > 0.05$ , \*:  $p < 0.05$ , \*\*:  $p < 0.01$ ).

due to (at least in part) more stereotypical DLS responses during the second stimulus relative to the first (Fig. 3.26 left).




Figure 3.29: (A) Main axes: decoded time estimates from two equally sized subpopulations ( $N = 150$ ) of Gaussian bumps varying along their stereotypy, gain, and temporal tuning parameters, and conditioned on whether they were labeled as "ramping" or "non-ramping" (mean posterior mean  $\pm$  mean posterior SD across 100 simulations, each with 100 trials). Insets: decoded time estimates from a single simulated trial (posterior mean  $\pm$  posterior SD) illustrating the decoding accuracy (top) and precision (bottom) metrics used in (B). (B) Comparison of decoding accuracy (top) and precision (bottom) from "ramping" and "non-ramping" neurons, for simulations from four differently parameterized models (left) and DLS data recorded during the two stimulus periods (right, median and IQR across simulations). Model variants were defined by whether their stereotypy and/or gain parameters were fixed ( $\eta = \eta_0$  and/or  $\gamma = \gamma_0$ ) or sampled ( $\eta \sim$  and/or  $\gamma \sim$ ) from the corresponding distributions shown in white in 3.28. Temporal tuning parameters ( $\mu$ ) always varied across simulated neurons and were drawn from the same uniform distribution for all model variants. Statistical differences across cluster and epoch distributions were assessed using two-sample Kolmogorov-Smirnov tests with Bonferroni correction for multiple comparisons (NS:  $p > 0.05$ , \*:  $p < 0.05$ , \*\*:  $p < 0.01$ ).

In summary, these results demonstrate that the authors' assertion regarding the relatively superior decodability of elapsed time from ramping as compared to non-ramping DLS responses is unfounded.

### 3.2.8 No evidence for the existence of ramping and non-ramping clusters in DLS responses recorded during the second stimulus

We show above how the authors' categorization criteria for labeling a neuron as ramping or non-ramping was problematic for the validity of claims regarding relative accuracy of time decoding in the two categories of cells. However, we wondered whether there was any structure in the data that would suggest such categories in the first place. If the data is deemed to not possess sufficient structure to be meaningfully partitioned, then clustering becomes arbitrary, potentially misleading and arguably unwarranted (Adolfsson et al., 2019). We thus sought to assess in an unbiased manner whether there is evidence for categories of DLS response profiles in the data, and, if so, whether those categories correspond to ramping and non-ramping patterns of activity. To do this, we adopted the *clusterability via multimodality approach*, which formalizes the intuitive notion that if clusters do exist, distributions along features that define them must be multimodal (Adolfsson et al., 2019).

To evaluate our ability to detect ramps should these exist, we analyzed the clusterability of two data sets. A *real* one, consisting of DLS activity recorded during the presentation of the second stimulus, and a *hybrid* one that consisted of those very same S2-aligned real responses but where the neurons originally labeled as ramps (using the criteria from Toso et al.) were replaced with actual simulated ramps (evenly split between up- and down-ramping, see methods). As a starting point, in both cases we plotted normalized spike density functions for all neurons ordered according to their temporal profile of activity (Fig. 3.30).

Specifically, responses were sorted by the angle of their first two PC coefficients in polar coordinates. We did not visually detect discontinuities in firing profiles themselves, which appeared more like a continuum. Although this was true even in the hybrid case, where we know there were ramps, methods used to order cells are known to produce visually misleading effects. Further inspection of the joint distribution of the PC coefficients used to sort those neural response rasters revealed that indeed simulated ramps tended to cluster in two poles opposite to each other, corresponding to up- and down-ramping neurons (Fig. 3.31B). The same phenotype was not observed in the case of ramps inferred by the authors among strictly real DLS activities, which instead seemed much more spread out in PC space, if not




Figure 3.30: (A) Normalized spike density functions of all DLS neurons recorded during the presentation of the second stimulus. Neurons are ordered by polar angle in the subspace defined by the first two PCs describing firing dynamics during  $S_2$ . (B) Same as (A), but for a hybrid simulation wherein neurons labeled as "ramps" using the authors' criteria were replaced by simulated ramps (evenly split between up- and down-ramping).

for displaying a slight tendency towards higher eccentricities (Fig. 3.31A). This bias along the radial polar coordinate was to be expected given the previously described bias in firing rate ranges induced by the clustering method (Fig. 3.26 right).

To quantify these observations and formally assess the existence of ramping and non-ramping clusters, we selected two features that satisfied the recently proposed *efficiency*, *algorithm independence* and *effectiveness* requirements for well-suited measures of clusterability (Adolfsson et al., 2019): 1) polar angle in PC coefficient space, and 2) Pearson's correlation coefficient between response profiles and elapsed time since the onset of  $S_2$ . For both measures, we found that their distributions were unimodal in the case of real data (Fig. 3.32A), and multimodal in the case of hybrid data (Fig. 3.32B). Furthermore, real data distributions restricted to putative ramps remained unimodal, indicating that the authors' classification approach fell short of even dichotomizing a continuum (Fig. 3.32A).

These results argue strongly against DLS activity profiles naturally clustering into ramping and non-ramping subpopulations, and instead favor a more holistic view of the heterogeneous dynamics underlying time encoding




Figure 3.31: (A) Joint distribution of PC coefficients conditioned on the authors' cluster identities. (B) Same as (A), but for a hybrid simulation wherein neurons labeled as "ramps" using the authors' criteria were replaced by simulated ramps (evenly split between up- and down-ramping).




Figure 3.32: (A) Distributions of clusterability metrics - polar angle of PC coefficients shown in 3.31 (top) and Pearson's correlation coefficient between the responses shown in 3.30 and time (bottom) - either agnostic to (in white) or split by cluster identity as assigned by the authors' classification method. Statistical departure from unimodality was assessed using Hartigan's dip test (NS:  $p > 0.05$ , \*:  $p < 0.05$ , \*\*:  $p < 0.01$ ). (B) Same as (A), but for a hybrid simulation wherein neurons labeled as "ramps" using the authors' criteria were replaced by simulated ramps (evenly split between up- and down-ramping).

by populations of striatal neurons (J. Wang et al., 2018).

### 3.3 Discussion

We have shown that the activity of neural populations in dorsal striatum correlates with animals' judgments of stimulus duration, consistent with a growing number of studies that have focused on a variety of different behavioral tasks (Bakhurin et al., 2017; Gouvêa et al., 2015; D. Z. Jin et al., 2009; J. Kim et al., 2018; Matell et al., 2003; Mello et al., 2015; Monteiro et al., 2023; J. Wang et al., 2018). Despite our contrasting conclusions to those of the original authors, we find the data set to be a valuable addition to this body of work, in that the behavioral task leverages an innate tendency for duration judgments to be biased by stimulus intensity to study potential neural correlates of timing behavior. Prior work employing motor timing tasks has involved training animals to respond at different times, depending on the identity of an explicit cue or recent trial history with respect to temporal availability of rewards (Malapani & Fairhurst, 2002; Mello et al., 2015). Temporal judgment tasks, or what are sometimes referred to as sensory timing tasks, often leverage trial by trial variability in judgments applied to particular durations to ask similar questions (Paton & Buonomano, 2018). However, the behavioral task used here takes advantage of the fact that stimuli of differing intensity are subject to systematically modified temporal estimation processes to take experimental control over internal mechanisms related to the judgment of duration (Matthews et al., 2011; Pardo-Vazquez et al., 2019; Toso, Fassihi, et al., 2021; Toso, Reinartz, et al., 2021). The fact that modification of temporal judgments via this apparently innate mechanism caused similar changes in striatal activity to those observed when timing behavior is modified using other means further strengthens the hypothesis that patterns of striatal activity indeed underlie animals' temporal judgments.

How does this relate to the notion of temporal "perception" on which the original study focused to a large degree? We remain agnostic to the animal's subjective experience, and thus prefer the term "judgment" which we operationally define as the animal's report. To perform this task, animals must transform information regarding the duration of vibratory sensory stimuli into a binary report, or choice, but also obey other requirements of the task, such as initiating a trial and maintaining their presence in a nose port for an extended period (Cruz et al., 2022). All of these elements of the behavior

involve the mapping together of sensory or internal variables with particular actions, and implicitly require a time-base as part of that process. In other words, during training animals learn a policy for behavioral control that must rest on a temporal basis. A wealth of data implicate the striatum in learning policies for behavioral control that maximize expected reward (Barnes et al., 2005; Cruz et al., 2022; Lau & Glimcher, 2008; Lauwereyns et al., 2002; Reynolds et al., 2001; Samejima et al., 2005; Schultz et al., 1997; Tai et al., 2012; Xiong et al., 2015). Thus, we tend to view striatal correlates of timing behavior as reflecting how neural dynamics provide a temporal basis for constructing intermediate representations required for behavioral control policies, such as the value of state-action combinations, which may or may not impinge on the animals' subjective experience of duration. That said, the animal's pattern of choices is the only window we have as external observers into any potential subjective percept of interval duration that may exist. The correspondence between timing behavior and neural activity in the striatum, reproducibly observed across many studies with differing behavioral task designs, suggests that striatal activity may at least reflect such a putative subjective perception of duration, independent of its ultimate origin.

## 3.4 Methods

### 3.4.1 Behavioral analyzes

#### 3.4.1.1 Psychometric functions

We used the *Psignifit* toolbox (H. H. Schütt et al., 2016) to fit the following 4-parameter psychometric function to all duration discrimination data:

$$\Psi(x; m, w(\alpha), \lambda, \gamma) = \gamma + (1 - \lambda - \gamma) \cdot S(x; m, w(\alpha))$$

In this parameterization, a strictly monotonic sigmoid function  $S$  from the stimulus level  $x$  (corresponding to either T1 or T2) onto the unit interval  $[0, 1]$  is specified by  $m = S^{-1}(0.5)$  and  $w = S^{-1}(1 - \alpha) - S^{-1}(\alpha)$ , namely the *threshold* and *width* parameters. This is independent of the choice of  $S$ , which in our case was the reversed cumulative Gumbel distribution. The hyper-parameter  $\alpha$ , which sets the span of  $w$  along the vertical axis, was set to 0.05. To account for stimulus-independent choices,  $S$  is scaled by two additional free parameters,  $\lambda$  and  $\gamma$ , which respectively control the upper and lower asymptotes of the psychometric function  $\Psi$ .

### 3.4.1.2 Psychometric model comparison

To assess whether and how stimulus intensity impacted specific parameters of the psychometric function described above, we first built a reference psychometric curve for each stimulus (S1 and S2), by fitting all of its four parameters to rats' choices as a function of stimulus duration using only middle-intensity trials. Then, for each of the remaining intensity conditions (high and low), we fit a battery of models that varied combinatorially in which parameters were let free and which were fixed to those of the corresponding middle-intensity reference model. Lastly, we computed the BIC for each tested model and each tested intensity condition, and took relatively lower (higher) BIC scores as evidence for (against) the corresponding models.

$$\text{BIC} = k \ln(n) - 2 \ln(\hat{L})$$

Where  $n$  is the number of observations (trials),  $k$  is the model-specific number of free parameters and  $\hat{L}$  is the corresponding maximum likelihood.

## 3.4.2 Single neuron analyzes

### 3.4.2.1 Spike density functions

Spike density functions (SDFs) were built on a trial-by-trial basis by taking the provided histograms of spike times (computed by the original authors using a 1-ms bin) and convolving them with a causal kernel specified by a gamma distribution with shape and scale parameters of  $k = 2$  and  $\theta = 50$  ms, respectively. Normalized SDFs were computed as z-scores of stimulus-aligned SDFs averaged across all possible stimulus durations.

### 3.4.2.2 Neuron selection

In order to be considered for analyzes, individual DLS neurons were required to have displayed an average firing rate during the presentation of the second stimulus of at least 1 Hz (554/602 neurons), and to have been recorded for at least one trial of all possible combinations of T2 and I1 (567/602 neurons), and T2 and I2 (588/602 neurons). An additional set of 42 neurons were deemed unstable by visual inspection of spike rasters and SDFs aligned to several task events. Altogether, these criteria resulted in a selection of 485 out of the original 602 neurons recorded in the delayed duration comparison

task variant. Note that unlike the authors, we used the same selection criteria for all reported analyzes.

#### **3.4.2.3 Modulation assessment**

In order to quantify the proportion of neurons that were modulated by several behavioral (e.g., choice on the current and previous trials) and experimental factors (e.g., I1, T1, I2 and T2 on the current trial) across the different task epochs, we used these as regressors in multivariable linear regression to fit each neuron's spike counts recorded in a 334-ms time window aligned to multiple task events (e.g., trial initiation, onset and offset of S1 and S2). For every neuron and every task epoch, the p-values for all factor coefficients were corrected for multiple comparisons using the Benjamini-Hochberg procedure. Lastly, we quantified the proportion of significant coefficients at two significance levels ( $\alpha = 0.01$  and  $\alpha = 0.05$ ) for each factor across epochs.

#### **3.4.2.4 Temporal tuning**

The temporal tuning of each neuron was estimated separately for each epoch in which the authors enforced their criteria for classifying neurons as ramping or non-ramping (1-s windows centered on the onset and offset of S1 and S2), and stimulus epoch (1-s windows aligned to the onset of S1 and S2), as the center of mass of its baseline-subtracted cross-trial average spike density function.

#### **3.4.2.5 Stereotypy coefficient**

For each neuron and within each stimulus epoch, we calculated the Pearson correlation coefficient between two cross-trial average spike density functions, each computed using randomly selected but complementary halves of all its single-trial responses recorded during stimulus presentation. We repeated this procedure 100 times and took the average correlation over permutations to be the stereotypy coefficient for that neuron and stimulus.

#### **3.4.2.6 Firing rate range**

A stimulus-specific firing rate range was defined as the difference between the maximum and minimum of each neuron's cross-trial average spike density function during that stimulus.

### 3.4.3 Population analyzes

#### 3.4.3.1 Low dimensional representations of population state

We used PCA to enable visualization of stimulus-aligned population responses in a representative subspace conditioned on task variables of interest. Briefly, we first averaged stimulus-aligned SDFs across all trials for a given neuron and concatenated all the resulting average responses into an  $N \times T$  matrix, where  $N$  is the total number of neurons and  $T$  is the number of 1-ms time bins. After normalizing these reference data to have zero mean and unit standard deviation along the temporal dimension, we used PCA to find the three orthogonal directions that captured the most variance in stimulus-aligned DLS dynamics. Lastly, when considering a given contrast (e.g., I2), each condition's trial-averaged population activity was first z-scored using the means and standard deviations estimated from the aforementioned condition-agnostic reference data and then projected onto each of the first three PCs separately, or onto a 2D subspace consisting of an informative section through the space defined by the first three PCs.

#### 3.4.3.2 Neurometric curves

We used linear discriminant analysis (LDA) to find the hyperplane that best separated population state vectors recorded just before the offset of the second stimulus depending on whether T2 was larger or smaller than T1. Briefly, in order to combine trial-by-trial responses of units that were largely recorded separately, rather than simultaneously, pseudo-trial population state vectors were first constructed by concatenating random samples (drawn uniformly and without replacement) from each neuron's distribution of spike counts in the last 112 ms preceding stimulus offset. This was done separately (but with cross-validation) for each training and test conditions. The former consisted of all possible combinations of T1 and T2 wherein both I1 and I2 were equal to  $64 \text{ mm s}^{-1}$  (middle intensity), and the latter varied depending on which contrast was used: for I2, test conditions were defined as all possible combinations of T1, T2 and I2 for which I1 =  $64 \text{ mm s}^{-1}$ ; for I1, combinations of T1, T2, and I1 for which I2 =  $64 \text{ mm s}^{-1}$  were used; and for choice, test pseudo-trials were combinations of T1 and T2 for which I1 = I2 =  $64 \text{ mm s}^{-1}$ . We then used LDA to train a classifier to distinguish training pseudo-trials according to what choice they corresponded to, and then used it to “judge”

test pseudo-trials. Lastly, for each contrast condition, we fit these “neural judgments” with a 4-parameter psychometric function, analogously to what we did for behavioral ones (see 3.4.1.1 above).

### 3.4.3.3 Time decoding

We used a naïve (flat prior) Bayes decoder to compute probability distributions over elapsed time using a repeated 2-fold cross-validation scheme with condition-averaged population responses over 100 repetitions.

Briefly, for each repetition we:

1. Discretized time  $t$  into  $B$  1-ms bins, such that  $b = 1, \dots, B$  and  $t_b \in [-240, 481]$  ms relative to the onset of the stimulus of interest;
2. Computed each neuron’s average spike count  $n_i$  through time  $t$  across trials for the training condition and all test conditions using 2-fold cross-validation:

$$n_i \forall i = 1, \dots, N$$

3. Computed each neuron’s tuning curve as a function of time  $r_i(t)$  by convolving its average spike counts through time recorded during the training condition trials with a Gaussian kernel  $g(t, \mu, \sigma)$  with mean  $\mu = 0$  and standard deviation  $\sigma = 50$  ms:

$$r_i(t) = (n_i * g)(t)$$

$$g(t, \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{t-\mu}{\sigma}\right)^2}$$

4. Assumed a Poisson spike generation model such that:

$$p(n_i | t_b) = \frac{r_i^{n_i}}{n_i!} e^{-r_i}$$

5. Made conditional independence assumptions about neurons, regardless of whether or not they were recorded simultaneously:

$$p(t_b | \mathbf{n}) \propto \prod_{i=1}^N p(n_i | t_b)$$

Where  $\mathbf{n} = (n_1, n_2, \dots, n_N)$  is a to-be-decoded test instance of concatenated population activity recorded at a known condition and point in time since the onset of the stimulus of interest;

6. Used Bayes' rule to create a decoding model that linearly combined all individual neuron encoding models with a uniform prior over decoded time  $p(t_b) = 1/B \forall b = 1, \dots, B$ :

$$p(t_b|\mathbf{n}) = \frac{p(t_b) \prod_{i=1}^N p(n_i|t_b)}{p(\mathbf{n})}$$

Where the probability  $p(\mathbf{n})$  for the population vector  $\mathbf{n}$  to occur does not have to be estimated explicitly, as it indirectly follows from normalizing the posterior distribution  $p(t_b|\mathbf{n})$  such that it sums to 1 over all possible decoder outcomes, i.e., elapsed time as decoded from DLS ongoing activity;

7. Computed point estimates of decoded time ( $\mu_b$ ), accuracy ( $\langle \varepsilon \rangle$ ) and precision ( $\langle \sigma \rangle$ ):

$$\begin{aligned} \mu_b &= \sum_{b=1}^B p(t_b|\mathbf{n}) \cdot t_b \\ \langle \varepsilon \rangle &= \frac{1}{B} \sum_{b=1}^B \varepsilon_b = \frac{1}{B} \sum_{b=1}^B |t_b - \mu_b| \\ \langle \sigma \rangle &= \sum_{b=1}^B \sigma_b = \sum_{b=1}^B \sqrt{\sum_{j=1}^B p(t_b|\mathbf{n}) \cdot (t_j - \mu_b)^2} \end{aligned}$$

### 3.4.4 Simulations

#### 3.4.4.1 Binary choices

To simulate agents driven by nominal NTD and NID, we sampled choices on a trial-by-trial basis from Bernoulli processes according to:

$$p(T2 > T1)_k = \frac{1}{1 + e^{-(\beta_T \cdot \text{NTD}_k + \beta_I \cdot \text{NID}_k)}}$$

Where  $p(T2 > T1)_k$ ,  $\text{NTD}_k$  and  $\text{NID}_k$  are the Bernoulli process parameter, normalized time difference and normalized intensity difference on trial  $k$ , respectively. The coefficients  $\beta_T$  and  $\beta_I$  were chosen such that the overall proportion of simulated choices as a function of, and grouped by NTD and NID (Fig. 3.4B, E left) were similar to those observed in rats (Fig. 3.5A, C left).

$$\beta_T = 3, \beta_I = 1$$

Lastly, to ensure that stimulus and trial statistics were matched between simulated and real data, the trial-wise T1, T2, I1 and I2 values (that determined  $\text{NTD}_k$  and  $\text{NID}_k$ ) used in simulations were dictated by the real data.

#### 3.4.4.2 Scalar time estimates

To illustrate how averaging decoded time estimates conditioned on the outcome of choices would mask underlying systematic biases in timekeeping, we simulated marginal distributions over internal time estimates for each stimulus duration assuming scalar timing with a Weber fraction  $w = 0.2$ . Specifically, internal time distributions were computed by convolving Dirac delta distributions shifted by each possible stimulus duration with Gaussian kernels of increasing standard deviation ( $\sigma = w \times t$ ). Lastly, when computing joint distributions of T1 and T2 estimates across stimulus pairs, these were assumed to be independent.

#### 3.4.4.3 Spikes times

In three separate instances, we resorted to analyzing spike data simulated from inhomogeneous Poisson point processes on a trial by trial basis (i.e., with a time-varying generative spike rate  $\lambda_{n,k}(t)$  for neuron  $n$  on trial  $k$ ), in a manner that matched (either roughly or exactly depending on the instance) the statistics of real data collected by Toso et al. in terms of numbers of neurons, numbers and types of trials per neuron, and firing statistics. SDFs were estimated from simulated spike trains exactly as described for real data in the spike density functions subsection of the single neuron analyzes section above.

When estimating false positive and false negative rates of our regression-based method for identifying I2-selective neurons during S2 presentation, we generated two sets of surrogate data: a positive and a negative control, with 100% and 0% of neurons modulated by I2, respectively. In the negative control case, for each simulated neuron and on each simulated trial,  $\lambda_{n,k}(t)$  was set to the (real) mean S2-aligned SDF of that neuron averaged over middle-I2 trials ( $\text{I2} = 64 \text{ mm s}^{-1}$ ). For the positive control, the same was true except that generative rates were multiplied by a gain factor proportional to each trial's I2 ( $\lambda_k(t) = \ln(\text{I2}_k)/\ln(64)$ ) before sampling spike times.

When demonstrating that the authors' operational definitions of ramping and non-ramping neurons predetermined them finding that time was better decoded from the activity of ramps than non-ramps, we simulated populations of Gaussian bumps with three independent parameters: a temporal tuning parameter controlling the average peak time of the simulated bump ( $\mu$ ); a stereotypy parameter controlling trial-to-trial variability in that peak time ( $\eta$ ); and a gain parameter controlling peak amplitude relative to baseline ( $\gamma$ ).

$$\lambda_{n,k}(t) \propto \gamma_n \cdot \mathcal{N}(v_k \sim \mathcal{N}(\mu_n, \eta_n), \sigma_n)$$

To assess if variability in each parameter on its own was sufficient to drive performance differences in time decoding between ramps and non-ramps, we simulated four different populations of our model bumps that differed with respect to which parameters were held constant and which were sampled from arbitrary distributions (informed by the biases observed in real data):

1.  $\mu \sim \mathcal{U}[-500, 1500] \text{ ms}, \eta = 0 \text{ ms}, \gamma = 3 \text{ s}^{-1}$
2.  $\mu \sim \mathcal{U}[-500, 1500] \text{ ms}, \eta \sim \text{Exp}(250) \text{ ms}, \gamma = 3 \text{ s}^{-1}$
3.  $\mu \sim \mathcal{U}[-500, 1500] \text{ ms}, \eta = 0 \text{ ms}, \gamma \sim \text{Exp}(3) \text{ s}^{-1}$
4.  $\mu \sim \mathcal{U}[-500, 1500] \text{ ms}, \eta \sim \text{Exp}(250) \text{ ms}, \gamma \sim \text{Exp}(3) \text{ s}^{-1}$

Lastly, when generating ground truth data to validate our clusterability analysis, we adopted a hybrid approach wherein the subset of DLS neurons that had been labeled as ramps using the authors' classification criteria were replaced with simulated ramps proper, equally split among up- and down-ramping ones. Importantly, simulated ramps inherited the trial (number and type) and firing (baseline and range) statistics of the real neurons they replaced.

### 3.4.5 Statistics

Departures from unimodality were assessed using Hartigan's dip test (Hartigan & Hartigan, 1985), and differences between distributions were assessed using two-sample Kolmogorov-Smirnov tests with Bonferroni correction for multiple comparisons (where applicable).

In all cases, test outcomes were denoted as: NS,  $p \geq 0.05$ ; \*,  $p < 0.05$ ; \*\*,  $p < 0.01$ , written in the same color as that of the corresponding distribution

for unimodality tests, and in black near the horizontal solid black lines plotted above the underlying distributions for two-sample tests.



## GENERAL DISCUSSION

Timing in the range of seconds to minutes is essential to most non-trivial forms of learning (Paton & Buonomano, 2018), and seems to be an ability that is unique among vertebrates. The BG, an ancient set of subcortical structures with a highly conserved circuit architecture (Grillner & Robertson, 2016), is one of the defining systems of the vertebrate condition (Cisek, 2021). Its main input structure – the striatum – receives converging inputs from the entire cortical mantle and diverging inputs from both the thalamus and dopaminergic neurons in the midbrain (Foster et al., 2021; Hunnicutt et al., 2016; Matsuda et al., 2009). It has been implicated in interval timing by multiple lines of evidence (Gouvêa et al., 2015; D. Z. Jin et al., 2009; J. Kim et al., 2018; Malapani et al., 1998; A. V. Maricq & Church, 1983; Meck, 2006; Mello et al., 2015), but its specific role in supporting this and other forms of adaptive behavior is still highly contended. Some camps claim that there exist distinct functional classes among striatal neurons on the basis of their firing profiles, and that ramping neurons in particular are of crucial importance to keeping track of time (Donnelly et al., 2015; Emmons et al., 2017). Others are more agnostic to the functional regimes of individual neurons and instead propose that timekeeping in the striatum is implicit in the dynamics of its heterogeneous populations of neurons (Bakhurin et al., 2017; J. Wang et al., 2018). And some even question whether striatal activity selectively underpins, or merely covaries with, reported estimates of elapsed time (Jurado-Parras et al., 2020; Toso, Reinartz, et al., 2021). More broadly, regarding the role of the BG, many groups favor the view that it is mainly involved in learning how to map state information onto discrete actions that

are then largely implemented elsewhere in the brain (Graybiel, 1998; Mink, 1996), while others argue that the BG directly learns a more fine-grained policy for continuous control (Desmurget & Turner, 2010; Park et al., 2020). Here, I hope to have helped constrain these debates by characterizing how different features of population dynamics in the dorsal striatum impact interval timing in the rat.

First, we used temperature to manipulate neural activity in the DS during an interval categorization task (Gouvêa et al., 2014). Our results provide compelling evidence for the causal nature of the link between the speed of striatal dynamics and judgments about duration. Then, we showed that variability in this feature of neural activity does not seem to control the moment by moment details of movement execution, despite its dramatic effect on rats' categorical decisions. Instead the data suggest that the striatum impacts overall vigor through a lower-dimensional gain signal expressed in baseline levels of activity. And lastly, we observe that applying the same temperature manipulation to one of the main sources of input to the dorsal striatum – M1 – does not produce any discernible effects on expert performance (Monteiro et al., 2023).

Second, we reanalyzed data from the DLS recorded during a delayed duration discrimination with respect to two successive vibratory stimuli delivered to the whiskers (Toso, Reinartz, et al., 2021, 2023). We find a strong agreement between the degree to which different task variables impacted performance and neural activity. Specifically, vibration intensity affected both the temporal scale (i.e., speed) and amplitude scale of DLS responses, while intensity-independent variability in choices mainly covaried with the speed of DLS dynamics. Lastly, we emphasize that the heterogeneity of response profiles across the DLS does not seem to comport with a dichotomy of ramping and non-ramping neurons (Rodrigues et al., 2023).

Both sets of experiments can be said to bias duration estimates by manipulating intrinsic dynamics in the dorsal striatum, thereby revealing the features of neural activity that might covary with those decisions endogenously. In reduced laboratory tasks with static rules such as the ones leveraged here, endogenous variability in neural dynamics might seem suboptimal, a nuisance, but in a natural setting with volatile contingencies, it provides a substrate for adaptation. Therefore, elucidating the mechanisms by which nervous systems flexibly adjust the temporal scale (and other features) of neural responses to accommodate non-stationarities in the statistics of the

environment is a likely prerequisite to understanding adaptive behavior.

## 4.1 Scaling mechanisms

In a previous study from the lab, dorsostriatal dynamics were shown to rescale with the changing behavioral demands of a serial fixed interval (FI) reward schedule (Mello et al., 2015). Specifically, when trained to lever-press after a given FI had elapsed, rats reliably pressed in anticipation of interval offset, and DS responses tiled the time between consecutive rewards reproducibly and unambiguously. Furthermore, as the FI changed across blocks of trials, rats' behavior also changed, with press onset times shifting earlier and later as a function of FI. Likewise, firing patterns in the DS stretched and contracted to meet the current schedule, by up to a factor of five and in as little as three trials. Such rapid and drastic adaptation implies similarly fast and substantial changes in the circuit properties that control striatal speed.

In a separate study employing a similar interval categorization task to the no-fixation variant described in chapter 2, bulk fluorescence imaging in the mouse SNc revealed that DAergic neurons encoded information about trial-to-trial variability in time estimates – more (less) SNc activity at interval offset was systematically associated with a higher propensity for short (long) judgments. Congruently, activation or inhibition of these neurons during stimulus presentation was sufficient to respectively slow down or speed up estimates of elapsed time, as inferred from the thresholds of the corresponding psychometric curves. Given that these effects were transient, evident on stimulated trials but not on those that immediately followed, they are more likely to be explained by direct consequences of SNc DA on the excitability of its downstream targets than by long-term plasticity<sup>1</sup> changes in those same target regions (Soares et al., 2016).

The SNc projects primarily to the DLS, where DA-dependent effects on excitability depend critically on what dopamine receptors striatal SPNs express. This antagonistic arrangement may at first seem at odds with the reported coherent effects of SNc DA on duration judgments and (presumably) on the speed of DLS dynamics (Soares et al., 2016). However, in another set of experiments with mice now performing the fixation required version of the

---

<sup>1</sup>The time constants of DA-dependent spine enlargement at the dendrites of SPNs are in the minute-range (Iino et al., 2020).

interval categorization task, our group noted an asymmetric contribution of indirect and direct pathway DLS SPNs during time estimation. Monitoring bulk activity in each pathway during stimulus presentation revealed that the need to actively suppress tempting actions (the short action early on and the long one later in the stimulus) results in DLS dynamics being disproportionately driven by indirect pathway SPNs. Accordingly, inhibiting D2- but not D1-receptor expressing DLS SPNs during the to-be-judged intervals disrupted fixation and “released” the action that would have been most likely should the interval have ended then (Cruz et al., 2022). This unbalanced recruitment of the indirect pathway suggests that DA’s effect on timing may have been mediated by a net reduction in the excitability of the SPNs that participate in DLS dynamics during stimulus presentation.

Dopamine alone can not sustain dynamics in the strictly inhibitory circuit that is the striatum, especially in light of the hyperpolarized resting membrane potentials of SPNs (Uchimura & North, 1989). Outside of the lab, but in work that directly inspired our acute temperature manipulation experiments, Hidalgo-Balbuena et al. showed that in anesthetized rats receiving a train of mechanical stimuli to the forelimb, DLS activity was primarily driven by inputs coming from S1 in cortex and the ventral posterior lateral nucleus (VPL) of the thalamus. Interestingly, time relative to the onset of the stimulation was more readily decodable from DLS responses than from those found upstream in S1, wherein each stimulus in the train seemed to reset neural dynamics. Next, in a separate group of rats overtrained on a temporal production task that involved running on a treadmill (and thus rhythmic mechanical stimulation of the limbs), the authors leveraged the fact that all animals developed a qualitatively similar stereotyped motor sequence to probe the relative contributions of DLS, S1 and VPL onto various aspects of that behavior. A fundamental observation was that no combination of lesions, pharmacological, or optogenetic manipulations in the three circuit nodes under consideration had any detectable impact on rats’ ability to control ongoing movement kinematics. This was in spite of some interventions having pronounced effects on the temporal patterning of stereotyped motor sequences. The most striking of which were those produced by bilateral *activation* and *inhibition* of VPL terminals in the DLS, which induced *over* and *underproduction* of the target interval, respectively (Hidalgo-Balbuena et al., 2019). Gradation in the level of tonic excitatory input from the thalamus thus joins DA-dependent modification of SPN excitability as a potential

circuit-level mechanism for controlling the temporal scale of neural activity in the DLS.

These results relate to a set of seminal studies probing the neural bases of both acquisition and execution of the idiosyncratic and stereotyped motor sequences that rats inevitably converge on in a different temporal production task. Originally, Kawai et al. demonstrated that an intact M1 is critical for rats to learn to double-press a lever with a inter-press interval (IPI) of ~700 ms; but otherwise unnecessary<sup>2</sup> once animals become proficient in their personalized behavioral “dances” (Kawai et al., 2015). More recently, the Ölveczky lab has built on this work by showing that the DLS, but not the dorsomedial striatum (DMS), is required for both learning and executing these complex motor skills (Dhawale et al., 2021; Wolff et al., 2022). But more importantly to the present matter, they provide compelling evidence for the case that learned stereotyped motor programs are encoded in the DLS, and that unlike corticostriatal, thalamostriatal synapses play a big role in that encoding. Specifically, Wolff and colleagues first showed that silencing DLS-projecting axons coming from M1 had no effect on expert performance, whereas the same manipulation at the axon terminals from the parafascicular and the rostral intralaminar nuclei of the thalamus resulted in a marked *underproduction* of the target IPI. Lastly, injecting a Zeta inhibitory peptide (ZIP)<sup>3</sup> in the DLS of expert rats reverted performance to levels similar to those observed in naïve animals (Wolff et al., 2022). Though the authors’ did not share in this perspective, one way to interpret the thalamus-DLS silencing experiment in light of the previous results from the Rueda-Orozco group (Hidalgo-Balbuena et al., 2019), is to construe silencing as an extreme form of inhibition that results in DLS (and by extension behavioral) dynamics unfolding exceedingly fast.

Previous theoretical work has tried to tackle precisely these aspects of striatal function. How can an all-inhibitory circuit like the striatum generate reproducible patterns of sequential activity at arbitrary speeds? And how can it do so in a manner that is so critically dependent on cortex early in training, but not so after that learning has occurred? To recapitulate these properties in a simulated striatum-like network model, Murray et

---

<sup>2</sup>Anecdotally, this was not true for *the* rat that adopted a “dance” with precise digit control (i.e., dexterity; Kawai et al., 2015).

<sup>3</sup>ZIP blocks an enzyme that is necessary for the maintenance of potentiated glutamatergic synapses (Sacktor, 2011).

al. focused on a somewhat underappreciated aspect of striatal circuitry – its recurrent and plastic inhibition (Wilson & Groves, 1980). They leveraged two forms of SPN-SPN plasticity for which there is empirical support: long-term anti-Hebbian<sup>4</sup> learning (Rueda-Orozco et al., 2009) and short-term synaptic plasticity (Tecuapetla et al., 2007). The former is a type of spike-timing-dependent plasticity (STDP) that leads to depotentiation when a post-synaptic neuron fires (reliably) shortly after the pre-synaptic one. This allows arbitrary sequences of patterned excitatory input (presumed to come from cortical afferents) to be internalized in the recurrent connections of SPNs, thus carving a path of least resistance through the network. Short-term synaptic depression on the other hand, results in fleeting “winner-take-all” dynamics (which have also been reported in the striatum; Wickens et al., 1991) in the presence of shared background excitation – “winning” SPNs become less efficient at inhibiting others the longer that they have been active, eventually allowing those that are least inhibited to become the next “winners”. This mechanism makes it possible for learned sparse activity sequences in the striatum to unfold at arbitrary speeds depending on the level of excitatory tonic input (Murray & Escola, 2017). Though the authors did not commit to thalamus as *the* source of that common drive in this study, they certainly suggested it as likely candidate. However, one crucial discrepancy between the model and existing data is that the relationship between excitatory tonus and temporal scale in the former is opposite to that between thalamic drive and timekeeping in the latter (Hidalgo-Balbuena et al., 2019; Murray & Escola, 2017; Wolff et al., 2022). One (admittedly ad hoc) way to reconcile the two is to posit a sign-flip in between thalamic axon terminals and SPN dendrites, by for instance invoking inhibitory interneurons. Indeed, there is some experimental evidence showing that thalamic afferents target fast-spiking interneurons (FSIs) disproportionately when compared to other cell-types within the striatum (including SPNs; Johansson and Silberberg, 2020). Alternatively, temporal rescaling could instead be mediated by dopamine, possibly though a concerted effect on the time constants of synaptic depression, or on the excitability of the intervening SPNs. Both of these possibilities were also entertained by authors (Murray & Escola, 2017).

However, in a follow-up paper, the same group now reconsidered their previously core idea of sequences materializing exclusively at the level of

---

<sup>4</sup>Fire together, *do not* wire together (Hebb, 1949).

SPN-SPN synapses. This was in large part due to the ZIP and thalamus-DLS silencing experiments from the Ölveczky lab having highlighted the importance of thalamostriatal synapses for skill learning and retention (Wolff et al., 2022), but in no small part because of the well-documented role of DA in learning, which had been all but neglected in their original model. In their revised formulation, Murray et al. propose a two-pathway circuit architecture, wherein fast, DA-dependent, reinforcement learning drives plasticity at corticostriatal synapses, while Hebbian learning modifies thalamostriatal synapses independently of DA and on a much slower timescale (Murray & Escola, 2020). This setup leads to two key covert processes that unfold with repeated practice: *input alignment*, by which the effects of thalamic inputs on downstream striatal populations become increasingly similar to those of cortical inputs (i.e., aligned); and *control transfer*, via which striatal activity goes from being predominantly driven by the cortical pathway to being mostly driven by the thalamic one. This new model indeed recapitulated the need for thalamostriatal synapses in sequence execution after learning while at the same time providing an explanation for why continued practice can have effects on performance even after a behavior has been learned, but at the cost of a mechanism in which to flexibly control its execution speed. One speculative way of combining the strengths of both approaches would be to assert that while the cortical pathway has the reigns, associative plasticity in the thalamic pathway actively biases what initially was a broadcast permissive excitatory tonus – without which cortical inputs would be insufficient to drive striatal units – such that it becomes stronger for those SPNs that intervene in any given sequence. This would have the added benefit of capturing the one result from the Wolff et al. data that the two-pathway model can not explain – the thalamic pathway is necessary not just for expert execution but also for initial learning, as evidenced by rats’ inability to relearn the task so long as DLS-projecting thalamic axons are silenced (Wolff et al., 2022).

Tying many of these ideas together (some of which into knots), Wang et al. manipulated and recorded from (separately): the primate medial frontal cortex (MFC); downstream caudate (homolog of the rat DMS) nucleus, and upstream thalamic regions, always in the context of a two-interval, two-effector temporal production task. Briefly, on each trial monkeys were explicitly cued as to which effector they should use (eye or hand) and which interval they should produce (800 or 1500 ms), with all four possible combinations randomly interleaved (J. Wang et al., 2018). Pharmacological

inactivations in each brain area during this task led to graded performance deficits, with thalamic infusions having the most pronounced effects; followed by the considerably smaller effect sizes in MFC and lastly, caudate. This gradation seems to be in accordance with the observation that in a comparable motor timing task in the rat, the DLS is disproportionately recruited over the DMS (Dhawale et al., 2021), which leads me to wonder if perturbing the putamen (homolog of the rat DLS) in the monkey would not have led to much sharper drops in temporal production than those observed when the caudate was targeted. Still, the authors reported that in all three of the brain regions they monitored in this task, neural responses covaried with monkeys' motor timing behavior, both within and across target intervals. Remarkably though, whereas in the MFC and caudate that covariability mainly expressed itself in the temporal scale of neural responses, in the thalamus it did so chiefly in a much lower-dimensional, offset-like feature of neural activity.

This nicely adheres to the idea that excitatory tonus provided by the thalamus might play a part in setting the speed with which recurrent dynamics unfold in downstream target regions. To investigate these and other potential network-level mechanisms underlying flexible control of production speed, Wang et al. trained the output of a recurrent neural network to breach a fixed threshold at a desired time, which was specified by the magnitude of a contextual excitatory input. After training, the activity of model units emulated many of the key features of MFC and caudate neurons, namely the heterogeneity of their response profiles and their temporal rescaling as a function of production interval (i.e., input tonus). Within the network, input and recurrent dynamics played complementary roles, with the former shifting initial and terminal fixed points along a dimension that was orthogonal to the space wherein the trajectories arising from the latter then unfolded. Accordingly, different production intervals were systematically associated with changes in the level (temporal scale) of activity in the input (recurrent) subspace, which was governed by the network's input (recurrent) weights: longer target intervals  $\rightarrow$  higher tonic inputs  $\rightarrow$  slower recurrent dynamics; conversely, shorter  $\rightarrow$  lower  $\rightarrow$  faster. This thus predicted that nonscaling components even within the dynamics of MFC and caudate responses should also covary with production intervals, as was later borne out by the data. Last but not least, to gain insight into exactly how higher and lower tonic inputs respectively give rise to a general slowing down or

speeding up of recurrent dynamics, the authors dissected a reduced model with only two mutually-inhibiting units that shared background excitation. This exercise led to a fundamental insight: stronger inputs drive individual units towards where their nonlinear input-output functions start to saturate, where the slopes of those activation functions are shallower – shallower slopes correspond to smaller derivatives and thus longer time constants. It is then through these single-neuron nonlinearities that an external tonic input can exert control over how fast or slow dynamics unfold within the recipient network (J. Wang et al., 2018).

*Trying* to put all of it together, it may be that though tonic external drive and dopamine function as two independent circuit-level knobs for adjusting the temporal scale of previously learned patterns of neural activity in the dorsal striatum, they ultimately act via the same mechanism at the cell level – external tonus sets the sensitivity of striatal neurons to recurrent inputs by indexing their saturating nonlinearities, and DA stretches or squashes those transfer functions directly. In the case of a population of SPNs that is heavily biased towards D2-receptor expressing neurons (Cruz et al., 2022), the net result is that more DA will lead to fainter activation functions, which for any given input level will then lead to shallower slopes and by extension, slower dynamics. Conversely, less DA; amplified nonlinearities; steeper slopes; faster dynamics (Soares et al., 2016).

Perhaps surprisingly then, in a time production task without strict immobility requirements, wherein mice lick throughout the production interval, *more* DA seems to predict *shorter* time estimates and *less* of it correlates with *slower* timekeeping. Accordingly, *activating* and *inhibiting* SNc neurons shifted production intervals *earlier* and *later*, respectively (Hamilos et al., 2021). These results are diametrically opposite to those observed by our group (Soares et al., 2016). Yet, the two are entirely reconcilable by positing that the change from active suppression in one task to continuous movement in the other might be accompanied by a qualitative change in the sampling bias of SPNs that take part in the dynamics that are then modulated by DA – from mostly D2- (Cruz et al., 2022) to mostly D1-receptor expressing SPNs (Hamilos et al., 2021). Beyond these two studies, assessing the role of DA in timekeeping in light of the behavioral demands of the task at hand and how that might shape the DA-receptor landscape in downstream striatal regions, may provide some relief to the many apparently disparate accounts that have been reported (Mikhael & Gershman, 2019). Ultimately,

it just makes sense that the message should lie at the intersection between neuromodulator and receptor.

The synthesis provided by Wang and colleagues also offers a way to avoid the inhibitory interneuron-requiring kludge I had proposed above to “fix” the discrepancy in the direction of how changes in thalamic input level translate to changes in striatal speed, as seen in data (Hidalgo-Balbuena et al., 2019; Wolff et al., 2022) and in my informal concocting of the two models proposed by Murray and Escola (Murray & Escola, 2017, 2020). Additionally, it makes a number of testable predictions concerning the role of the thalamus in the behavioral tasks that we have studied in the lab. Chief among them is that recording from thalamic inputs to the DLS in either the serial FI or the interval categorization tasks should reveal that activity there is low-dimensional and proportional to animals’ estimates about duration. Furthermore, if recorded simultaneously, DLS dynamics should be systematically (slower) faster when thalamic tonus is (higher) lower (Hidalgo-Balbuena et al., 2019; Murray & Escola, 2017; J. Wang et al., 2018; Wolff et al., 2022). And manipulating those thalamic afferents should bias the temporal scale of striatal activity and duration judgments accordingly (Hidalgo-Balbuena et al., 2019).

The effects of our temperature manipulations seem less amenable to summary via this cell-intrinsic mechanism alone, seeing as if anything, temperature has been shown to affect SPN excitability in the reverse direction to what one might intuit – warming (cooling) elicits a graded outward (inward) potassium current which results in membrane hyperpolarization (depolarization). In other words, higher and lower temperatures lead to lower and higher excitability, respectively (Owen et al., 2019). On its own this might predict that the impact of temperature on temporal scaling would be opposite to what we observed, with warming leading to dilation and cooling to contraction of striatal responses. However, the experiments from Owen et al. were performed in slice preparations (Owen et al., 2019), which do not recapitulate many of the cell-extrinsic processes that temperature is also known to affect, including synaptic transmission and axonal conduction (Banerjee et al., 2021). The fact that we see clear effects of temperature on temporal rescaling (Fig. 2.5, 2.8, 2.10, 2.11) speaks to a net effect that overturns this cell-intrinsic, potassium-channel-mediated partial effect. In this sense, though it certainly proved a powerful method for establishing the speed of population dynamics in the striatum as a contributing cause of variability in duration

judgments, the exact circuit-level mechanisms by which it did so may not all find parallels in what goes on endogenously.

The whisker vibration intensity manipulation leveraged by the Diamond lab on the other hand (Fassihi et al., 2014; Toso, Fassihi, et al., 2021; Toso, Reinartz, et al., 2021), is almost certain to be exploiting cell- and network-level mechanisms that are all instantiated in the natural setting. Given the speed-control mechanisms we have been considering, the fact that in their data there are clear monotonic effects of intensity on both the speed of DLS dynamics and rats' discrimination judgments predicts that there should be an inverse relationship between the activity of DLS-projecting thalamic neurons and vibration intensity. Indeed, in anesthetized rats receiving repeated whisker stimulation, neurons in the parafascicular and medial posterior nuclei of the thalamus that project to whisker-responsive regions of the DLS show precisely this phenotype – declines in activity as a function of stimulation frequency (Alloway et al., 2014). Whether or not this is also the case in awake behaving animals remains to be seen, but if so it would provide extra support to the idea that tonic input from the thalamus sets the pace of striatal dynamics, possibly via its interaction with the activation nonlinearities (J. Wang et al., 2018) or dynamic synaptic properties (Murray & Escola, 2017) therein.

## 4.2 Discrete vs. continuous control

One of the biggest points of contention regarding BG function has to do with its role in motor control. In the specific case of the dorsal striatum, it is often debated whether signals there reflect learned transitions between the behavioral syllables that make up procedural memories (Hidalgo-Balbuena et al., 2019; Markowitz et al., 2018; Martiros et al., 2018) or the moment-by-moment motor commands required to express them (Dhawale et al., 2021; Jurado-Parras et al., 2020; Mizes et al., 2023). There is seemingly no shortage of data to invoke in support of one or the other (Mink, 2018; Park et al., 2020), but more often than not it feels like the distinction hinges on loose definitions more so than on fundamental differences. A heterarchical framing of behavior like the one introduced in chapter 1 might provide firmer grounds on which to then lay out concrete predictions about the roles of areas *A*, *B* and *C* in relation to hierarchs *x*, *y*, and *z*, or levels *i*, *ii* and *iii*. Though none of the experiments I have directly analyzed in this thesis were

designed with the primary goal of settling this debate, they still provide valuable data points that ought to constrain it, and might even offer some *partial* reconciliation, with vigor as broker. By emphasizing the existence of multiple features of striatal dynamics and their dissociable contributions onto different aspects of rats' behavior, I hope to counter lines of reasoning that would preemptively pit these two sets of hypotheses as intrinsically incompatible before distilling their arguments into testable statements about data. Full disclosure though, I am partial to this discussion.

Our bidirectional temperature manipulations experiments had two main effects on neural activity in the dorsal striatum of anesthetized rats: a *monotonic* one on the temporal scale of its dynamics (Fig. 2.8); a *non-monotonic* one on the baseline level of activity there (Fig. 2.9), which were then putatively matched to two main effects of striatal temperature on timing behavior: an unambiguous and *monotonic* one on rats' duration judgments (Fig. 2.23); a modest and *non-monotonic* one on the speed of choice movements (Fig. 2.30). Even in the absence of the latter, in which case there would be no significant behavioral correlate of the nonlinear effect of temperature on striatal base firing, the obvious lack of a monotonic one (Fig. 2.29) already speaks lengths to the division of labor among brain systems that ultimately determines *when* to do *what*, and *how*. Still, the comparatively smaller effect of striatal temperature on movement kinematics was corroborated by a similarly non-monotonic trend in *reaction*, but not *movement* time profiles (Fig. 2.31). Together, these results generated the prediction that endogenous variability in striatal baseline firing might correlate with the delay with which choice movements were initiated (i.e., reaction time), but not their duration (i.e., movement time). Indeed, when we went back to look for it in our chronic recordings of striatal activity during interval categorization requiring fixation, we found a significant anti-correlation between overall levels of DS activity before trial initiation and the delay with which rats later reacted to stimulus offset, but none with how long they then took to report their choices (Fig. 2.32). Unbeknownst to us at the time of writing the temperature manipulation manuscript<sup>5</sup>, experiments in which the putamen (the DLS homolog in primates) of baboons (*Papio papio*) was reversibly cooled during a reactive reaching task produced analogous results: reaction times increased with cooling; movement times, reaching accuracy and precision all went

---

<sup>5</sup>Shame on us.

unaffected by it (Beaubaton et al., 1980). It is possible that this deferment of movement initiation is primarily mediated by temperature's effect on the firing of direct pathway striatal projection neurons, seeing as inhibition of D1- but not D2-type SPNs in the DLS also delayed choice movements in mice performing our interval categorization task (Cruz et al., 2022).

What is it that is so fundamentally different between reaction and movement times? I suspect that the specificity of our temperature effects, and similar ones reported in both rats (Xu et al., 2014) and primates (Beaubaton et al., 1980), is more related to the *transition* aspect of going from not moving to moving than to reactive or execution components per se. It just so happens that in our tasks, the movements that rats and mice use to report their judgments about duration are ballistic, monosyllabic (Markowitz et al., 2018). Because of this, the striatum, and the BG more generally, while continuously involved in biasing those judgments during stimulus presentation, need only participate once in the capacity of gating them – at the moment the choice action is “released” (Mink, 1996). In more multisyllabic behaviors, one might expect striatal temperature manipulations to also dilate and contract movement duration, though the general tendency of animals and BG circuits to “chunk” well-learned complex spatiotemporal routines into supraordinate syllables in themselves might speak to the contrary (Graybiel, 1998; X. Jin & Costa, 2010). Indeed, in a task wherein rats were trained to press two levers in an arbitrary sequence of three lever-pressing actions, SPNs in the DLS seemed to disproportionately bracket the onset and offset of the reinforced sequence. Interestingly, this heightened modulation was contingent on a particular press sequence having been reinforced, seeing as incorrectly performed sequences did not elicit these flanking responses from putative SPNs (Martiros et al., 2018).

These observations resonate with some of our reanalyses of the delayed interval discrimination data from Toso et al. (Toso, Reinartz, et al., 2021). To recap, we observed a systematic relationship between the intensity of the *second* stimulus in the comparison and the accompanying neural activity in sensorimotor striatum. Specifically, vibration intensity during that stimulus period affected the speed and gain of DLS responses monotonically (Fig. 3.11), to the point where clear intensity-dependent up-modulation is visible at the level population averages (Fig. 3.10). This stood in stark contrast to the utter lack of modulation of DLS activity by the intensity of the *first* stimulus during its presentation (Fig. 3.16). Since the two stimuli were indistinguishable

from an experimental standpoint, both in terms of sensory inputs and immobility requirements, this striking asymmetry would seem to argue against the DLS leveraging immediate sensorimotor states to control ongoing movement kinematics. Instead, the fact that a similarly striking imbalance could be readily observed in the relative contribution of the two stimuli onto animals' decisions (Fig. 3.7, 3.8), seems to favor the hypothesis that DLS activity informs a learned mapping between states and actions. This point is reiterated by the perturbation experiments from Hidalgo-Balbuena et al. – in a task requiring fine-grained control over position, speed and acceleration, manipulating/lesioning the DLS and its afferents (S1 in cortex or VPL in the thalamus) affected neither of these kinematic parameters, despite clearly impacting the temporal structure of behavior (Hidalgo-Balbuena et al., 2019). And the same was observed in a separate study using the exact same task when lesions were restricted to just the DLS (Jurado-Parras et al., 2020).

All of these results seem to directly contradict recent accounts of DLS as crucial for generating the detailed movement patterns that underlying the motor skills that rats acquire in the previously mentioned IPI production task (Dhawale et al., 2021), as well as in a yet unmentioned “piano-playing” task (Mizes et al., 2023). In the context of the former, Dhawale et al. reported no enhanced tuning of DLS neurons to the boundaries of learned motor sequences, nor to transitions in the elemental motifs that might compose them. Though the authors attributed the discrepancy with previous findings (X. Jin & Costa, 2010; Martiros et al., 2018) to a categorical difference between “species-typical” behaviors (e.g., pressing a lever  $x$  times) and “learned skills” (e.g., idiosyncratic stereotyped routines during interval timing), we see it possible that the reported relative lack of bracketing or sequencing responses instead reflects differences in data analysis. Specifically, by grounding their argument primarily on the profile of temporally warped neural activity during the IPI, averaged across DLS neurons and animals, the authors incur the risk of masking existing nonuniformities in the temporal tuning of those responses, if: there is substantial trial-to-trial variability in production intervals (as there seems to be); up- and down-modulation are similarly common at syllable transitions (as was shown to be the case in spontaneously moving mice; Markowitz et al., 2018); transitions happen at different points within the distinctive sequences of each rat (were they not idiosyncratic). Indeed, when the authors instead plot DLS activities aligned to the onset of the first press from a single rat as neural trajectories in PC space, a clear “kink” can

be observed around (at least) the first lever press (Dhawale et al., 2021). This is suggestive of fast, transient modulation explaining a substantial amount of variance in the data. The lack of an equally clear kink near the moment of the second press may reflect a smearing of time-locked phasic responses due to cross-trial variability in the produced IPIs, or it may be that the “chunked” motor sequence includes subsequent movements leading up to the reward magazine, or that there is some asymmetric recruitment of the DLS in initiating versus terminating learnt action sequences, or even that, as the authors concluded, there really is none at all. Replotting neural trajectories for the remaining individual rats, and separately for the first and the second lever presses, as well as for the moment of reward delivery, might help to better arbitrate between these possibilities.

Next, Dhawale and colleagues leveraged the fact that rats tended to alternate between slight variations of their private stereotyped routines – sequence modes – to ask whether there were any signs of enriched modulation at the choice-points between them. The authors concluded that indeed there were, but that importantly, they *followed* mode transitions rather than *leading* them, thus implying that the DLS could not have possibly contributed to the decision to execute one over the other (Dhawale et al., 2021). However, I would once more like to nitpick and consider an alternative explanation. Starting with the very notion of *modes*, which are put forth as these deliberate twists on a common behavioral backbone aimed at the same goal (as illustrated by the analogy with different types of tennis serves). After inspecting the example video<sup>6</sup> that the authors supplied in their supplementary information, my admittedly subjective and speculative take is that these variants might instead reflect unanticipated deviations from the initially intended motor program, which then have to be corrected. If this is accurate (and representative), one way to reconcile the DLS not leading peak trajectory discriminability would be to recast modes as unplanned yet frequent motor errors that are subsequently corrected, in which case it would make more sense for the DLS to lag (and select how to get back on track after an error) and not lead (and select the erroneous motif in the first place). Indeed, the kink in neural space following peak movement trajectory discriminability in

---

<sup>6</sup>In supplementary video 1, mode 2 looks like the rat’s left forelimb slips from holding the lever and gets momentarily stuck underneath it during an attempt to resume the exact the same sequence that is observed in mode 1 (Dhawale et al., 2021).

the putatively erroneous mode is compatible with a discrete reactive selection that then brings the two modes back together. In this view, instead of a choice-point between existing modes representing deliberately arbitrating between a flat and a slice serve, it would be more like adapting to a poorly thrown ball mid-serve.

Last, but most crucially, the authors claimed that DLS activity encodes task-related movement kinematics. As a statement about the data, we see no problem with this. That movement kinematics would be partially decodable from the motor commands that generated them, be them discrete or continuous, is not surprising. However, as a statement about *representation* and *control*, we again tend to disagree. To support their claim, Dhawale et al. used two complementary analyses that respectively focused on the *encoding* and *decoding* of kinematic parameters (i.e., horizontal and vertical position, velocity and acceleration of the head and forelimbs), by and from the activity of individual neurons in the DLS, as compared to those in the DMS. The authors took the fact that the fine-grained details of the learned motor patterns could (not) be predicted by, and were (not) predictive of, the activity of simultaneously recorded DLS (DMS) units, as evidence for the DLS controlling those details (Dhawale et al., 2021). Setting aside the issue of using a deep neural network to decode kinematics from neural activity and then implying a direct link (i.e., linear) between the two, the very premise of basing conclusions about the function of the DLS in motor control on across-region comparisons regarding *one* single hypothesis, rather than within-region comparisons of *multiple* hypotheses, is problematic. The activity of neurons in V1 is likely to correlate with many aspects of ongoing movement kinematics, possibly more so than neurons with more temporally extended tuning like those found in the DMS, but that on its own does not argue for V1 controlling them. Interestingly, in their preprint version of this paper, the authors did contemplate alternative hypotheses regarding the representational nature of DLS responses. There, remarkably, they found that a simpler model using sparse Gaussian bumps tiling the IPI as regressors, outperformed the more complex kinematics model (despite the former having nearly half the number of bases of the latter). This was later deemphasized in light of a small performance boost in a combined model that used the two sets of predictors (Dhawale et al., 2019, 2021), but as we see it, parsimony would rule for a discrete sequencing role for the DLS over and above a continuous control one. Concretely, I would predict that if the

DLS were cooled or warmed during the IPI task, inter-syllable delays would respectively dilate and contract (and by extension, so would IPIs), but the individual syllables within the accompanying motor routines would not. Given that the Ölveczky lab has reached out to us regarding the technical details of how we performed our temperature manipulations (and that we have promptly shared everything we have), I am excited by the prospect of not having to wait long to be proven wrong.

In a more recent study, Mizes et al. doubled down on this interpretation of BG and more precisely, of DLS function, using a combination of electrophysiological recordings and lesions during an elegant discrete motor sequence production assay, in which rats are rewarded for pressing three “piano” keys in an arbitrary yet prescribed order. Here too the authors report no evidence for population activity in the DLS preferentially marking the start and/or stop of automatic motor sequences; nor for it delineating individual key presses and orienting body movements within those sequences (Mizes et al., 2023). But again, these conclusions rely mostly on time-warped neural activity averaged across neurons and combined across rats, together with the assumption that motor sequences begin and end with the first and third presses, respectively. In work employing a closely-resembling task, unwarped responses of tentative SPNs in the DLS seemed to clearly demarcate the onset and offset of rewarded three-press behavioral sequences, but were not tightly locked to the first and last presses in those sequences (Martiros et al., 2018). Indeed, even in the Mizes data, if one extrapolates from the unwarped responses of example neurons and kinks in neural trajectories, it starts to look like DLS activity might be reflecting transitions between individual behavioral elements. Moreover, forelimb trajectories after bilateral lesions to the DLS look qualitatively similar to those reported by Hidalgo-Balbuena, wherein the temporal patterning of motor sequences was globally affected, but the kinematics seemed to be locally preserved.

### 4.3 Hysteresis

The observed monotonic relationship between striatal temperature and the speed of striatal dynamics arises from a combination of cell- and circuit-level temperature dependencies, not all of which need be coherent. Though many of these may not be readily accessible to nervous systems as means with which to regulate the temporal scale of neural activity, temperature is

likely to put the system in a state wherein it then deploys those endogenous mechanisms to compensate for the exogenous perturbation.

While the effects of temperature on timing judgments reported here were initially robust (Fig. 2.21, 2.23, 2.25), they steadily decreased over the course of manipulation sessions (Fig. 4.1). In spite of us trying to minimize opportunities for adaptation to temperature changes by applying them in blocks of trials lasting only three minutes, interspersed with control blocks (Fig. 2.2), and alternating manipulation sessions with “washout” ones where temperature was not manipulated (Fig. 2.20C), we wondered whether rats might still have learned to counter the effects of temperature. To explore this possibility, we performed an additional set of behavioral experiments in which the categorical boundary between intervals rewarded for short and long judgments was shifted with the same blockwise structure (Fig. 4.2) as temperature had been (Fig. 2.2). This manipulation was devised to create a similar scenario to that of striatal temperature manipulations in terms of unsignalled and surprising feedback. Indeed, animals developed the ability to shift their decision thresholds on a block-by-block manner over a small number of sessions, mirroring the observed diminishing effects of temperature on temporal judgments (Fig. 4.3, 4.4). Thus, while temperature manipulations might also produce physiological adaptations that render neural systems more robust to future temperature variations, we found that the changing impact of temperature on timing judgments was consistent with rats learning to compensate for decision variables whose features might be changing<sup>7</sup>. Although we did not record DS responses during either the temperature nor the boundary manipulation experiments, it is plausible that adaptation in both cases also involved a flexible rescaling of striatal dynamics, possibly via one of the mechanisms explored in the previous section. Alternatively, adaptation could reflect modification of some other feature of neural activity, in the DS and/or elsewhere in the brain.

This remarkable ability of rats to lean on recent trial history to flexibly adjust their behavior to volatile contingencies had us questioning whether something similar may be at play even when those task rules are static. To explore this idea, we leveraged training data from 22 rats, each with

---

<sup>7</sup>More generally, these data emphasize the need to continuously evaluate the behavioral impact of manipulating neural systems from the moment those manipulations begin, as opposed to evaluating end-point or average effects alone, seeing as the adaptation ability of animals may overcome what are initially significant effects on performance.

~20 sessions, totalling ~150000 trials of asymptotic performance in the fixation variant of the interval categorization task. We started by conditioning psychometric curves on whether the previous stimulus had been correctly categorized or not, and on what its duration had been (Fig. 4.5). This analysis revealed that rats use a “win-stay” strategy that is weighted by decision uncertainty, with correct choices based on ambiguous (unambiguous) stimuli leading to larger (smaller) biases in the following trial. After incorrect decisions, rats also show hints of a “lose-switch” strategy, but whether or not those biases are graded as a function of decision confidence is less clear. This asymmetry is likely due to the considerable difference in the amounts of data that are available to us for correct and incorrect trials, together with the higher probability of selecting for disengagement when conditioning on errors. To understand if the picture painted by psychometric curves fit irrespective of animal identity was representative of individual rats, we next fit a multivariable logistic regression model to choice data from each rat separately, using categorical predictors for multiple task variables on current and past trials. Unsurprisingly, the predictors associated with the current stimulus duration consistently accrued the largest coefficients, but those reflecting the interaction of correct choices and stimulus duration on the previous trial also attained significant weights across all animals (Fig. 4.6). Furthermore, they reiterated our initial observation, that history biases following rewarded trials are proportional to decision uncertainty.

A similar history- and confidence-dependency had been previously reported in rats performing odor identification and categorization (Mendonça et al., 2020), and was theorized to reflect optimal stimulus-response learning in dynamic environments (Drugowitsch et al., 2019). In this way, apparently suboptimal behavior arises from optimal learning algorithms being tested outside of the context in which they evolved (Mendonça et al., 2020). Critically, this work relaxes the common assumption that all choice variability is due to stochastic stimulus (or timing) noise, and makes the additional assumption that a volatile subjective category boundary is adjusted on a trial by trial basis. This does not imply an explicit representation of said dynamic categorical boundary against which equally explicit single-trial stimulus percepts are then compared. An equivalent albeit more parsimonious account would be to posit that the plastic boundary is implicit in the values of the available actions, which can themselves be approximated using

reproducible neural activities (such as those found in the striatum) as temporal bases. However, the mechanisms through which these adjustments are realized in the brain are far from being well understood.

Using our interval categorization task as an example, one possible starting point would be to ask if the systematic biases we have observed in the speed of striatal dynamics (Gouvêa et al., 2015; Monteiro et al., 2023) correlate with the expected biases in performance given recent trial history. Though we have not yet addressed this question thoroughly, preliminary analyses and modeling efforts building on the work of Asma Motiwala (Cruz et al., 2022; Gouvêa et al., 2015; Motiwala et al., 2022) seem to indicate that this is not the case – history-dependent biases do not appear to manifest in the temporal scale of striatal responses (data not shown). If this turns out to stand the test of more systematic inquiry, it suggests that hysteresis effects are either not mediated by the DS, or implemented via some other aspect of neural activity there. Given the observation that gain modulation of DLS responses by vibration intensity also correlated with intensity biases in interval timing behavior (Fig. 3.8, 3.7; Rodrigues et al., 2023), one possibility we are excited about has to do with precisely this feature – *gain* – and how it may be asymmetrically modified by the phasic release of DA in the DS, a signal that is known to scale with outcome uncertainty (Fiorillo et al., 2003; Soares et al., 2016). Specifically, in light of prominent views of the basal ganglia as a policy learning system (Doya, 1999; Lau et al., 2017; Mink, 1996), activities recorded in the striatum are likely to reflect something like the multiplexing of action- and time-related information, a sort of decomposed policy for behavioral control that mirrors both a temporal basis set provided by neural population dynamics and motor information subject to reward-driven learning (Gershman et al., 2014; Ludvig et al., 2008; Motiwala et al., 2022). One could then imagine a scenario wherein the bases that inform the value of the short action are disproportionately amplified following a correct short choice when compared to those used to approximate the value function of the long action (which may even be synergistically dimmed). And vice versa for correct long judgments. This would produce a kind of lopsided warping in the geometry of striatal trajectories in neural state space conditioned on trial history, in stark contrast to: the speed changes induced by temporal rescaling; the offsets induced by tonic inputs and baseline activity levels. Furthermore, it seems plausible that this gain amplification should be leveled by the informativeness of

reward, which from an information theoretic viewpoint equates to how unpredictable it was, which in turn tends to follow from the interaction of stimulus difficulty and decision uncertainty/confidence (Drugowitsch et al., 2019; Mendonça et al., 2020). Whether this *asymmetric* (for actions and rewards on either side of the boundary) and *graded* (by the magnitude of reward prediction errors) modulation of striatal gain is mediated by DA's effect on plasticity or excitability, we do not yet know, but further exploring existing photometry data in the lab (Cruz et al., 2022; Soares et al., 2016), by conditioning SNc and (in)direct pathway DLS activity on hemisphere, choice, and history, may provide some clues.

## 4.4 Outlook

The brain is a self-perpetuating, heterarchical, plastic and open dynamical system shaped by evolution and experience in a volatile world. It possesses intrinsic dynamics, and at the same time is subject to a continuous stream of inputs from the environment (many of which reafferent). Both of these factors – intrinsic dynamics and inputs from the environment – give rise to observable patterns of neural activity and adaptive behavior. A mechanistic understanding of how variability in either of the former explains variability in either of the latter is a fundamental goal of systems neuroscience. One of those systems – the basal ganglia – is a defining trait of the vertebrate brain, and one that is crucial for the equally-defining ability of vertebrates to bridge between temporally distant events. Its main input structure – the striatum – sits in a privileged anatomical and functional position to learn a time basis for value function approximation in service of biasing and gating decisions during interval timing (in particular) and reinforcement learning (more broadly), but only through its interaction with preexisting systems for behavioral control.




Figure 4.1: The effects of temperature diminished over manipulation sessions. (A) Threshold dilation across the first three temperature manipulation sessions for the DS cohort trained in the fixation version of the task ( $N = 6$ ). Small markers and thin lines linking them refer to threshold dilations for individual animals. Larger markers correspond to cross-animal averages (mean  $\pm$  SEM), and their facecolor being anything other than black indicates that the underlying dilation distribution was significantly shifted from zero ( $p < 0.05$ , one-sample two-tailed  $t$ -test). The gradient bar and arrow symbolize the uneven contribution of the last two sessions to the data pool shown in (B) and (C), with the last session contributing the most (see Methods). (B) Performance in the last two sessions of temperature manipulations conditioned on striatal temperature ( $N = 6$  rats): warm (red), control (black), cool (teal) or coolest (cyan). Psychometric curves fit to cross-animal averages of psychophysical data split by temperature, respectively shown as solid lines and markers of matching color (mean  $\pm$  SEM). Bottom-right inset: Differences in proportion of long choices from each experimental condition to the control condition (mean  $\pm$  propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each temperature condition's psychometric fit. Solid black lines represent MAP point estimates. (C) Threshold dynamics aligned to and across block transitions in the last two sessions of DS temperature manipulations (left and right, respectively). Condition-split cross-animal average thresholds (mean  $\pm$  SEM) were computed using trials that fell into a sliding window lasting 90 s (half the block duration) that was swept from the preceding to the succeeding control blocks in increments of 9 s. Each marker corresponds to one sweep, and its color shading denotes the fraction of that sweep's window that was inside a control block (with black markers corresponding to 100% control trials), and by extension its complement that fell in a manipulation block (with pure manipulation colors corresponding to 100% manipulation trials).




Figure 4.2: Boundary manipulation protocol. The thin horizontal dotted lines represent the four categorical boundaries rats experienced in these sessions ( $N = 5$ ). Boundary changes followed the same rules as the temperature manipulation experiments: a control-manipulation-control 3-min block design with boundaries drawn at random and without replacement from the set  $B = 1.32, 1.5, 1.68, 1.85$  s until exhaustion (at which point the set was replenished and the sampling process resumed).




Figure 4.3: Rats adapt to a dynamic decision boundary. **(A)** Performance in the first two sessions of boundary manipulations conditioned on categorical boundary ( $N = 5$  rats): 1.32 (yellow), 1.5 (black), 1.68 (dark blue) or 1.85 (light blue) s. Psychometric curves fit to cross-animal averages of psychophysical data split by boundary, respectively shown as solid lines and markers of matching color (mean  $\pm$  SEM). Bottom-right inset: Differences in proportion of long choices from each experimental condition to the control condition (mean  $\pm$  propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each boundary condition's psychometric fit. Solid black lines represent MAP point estimates. **(B)** Same as **(A)**, but for the last two sessions of boundary manipulations.




Figure 4.4: (A) Distributions of percentage change in threshold relative to control (dilation) as a function of which categorical boundary was in effect for the first two boundary manipulation sessions. Markers represent individual threshold dilations, linked within animals by thin solid black lines. Boxplots show animal means (horizontal black lines) and SEM (colored bars). (B) Same as (A) but for the last two manipulation sessions ( $N = 5$ ; one-sample two-tailed  $t$ -tests, significant  $t_s(4) = [-7.21; 2.89]$ ,  $P = [0.04; 2e3]$ , non-significant  $t(4) = 2.61$ ,  $P = 0.06$ ; repeated measures ANOVA followed by post-hoc contrasts with Tukey correction for multiple comparisons,  $F(2, 12) = 15.64$ ,  $P = 5e-4$ ; smallest significant  $q_s(8) = 5.18$ ,  $P = 0.01$ , non-significant  $q_s(8) = 2.43$ ,  $P = 0.15$ ).




Figure 4.5: Trial history affects categorization performance. Left: psychometric curves fit to cross-animal data pools either agnostic to (dashed black line) or conditioned on (solid colored lines) the duration of the previous trial and it having been correctly categorized ( $N = 22$  rats). Bottom-right inset: Differences in proportion of long choices from each trial history condition to the history-agnostic condition (mean  $\pm$  propagated SEM). Top-left inset: Marginal posterior distributions of the threshold parameter for each trial history condition's psychometric fit. Solid black lines represent MAP point estimates. Right: same as before, but looking at choices following incorrect choices.




Figure 4.6: Coefficients of a logistic regression model fit to choice data using categorical regressors for multiple task variables on current and past trials. White markers refer to single animals, and black or colored (in the case of those that directly map onto the conditions introduced in Fig. 4.5) ones to averages across animals.  $\beta_0$  – intercept; B – broken fixation; C – correct choice; E – incorrect choice; S – short choice; L – long choice; t – index of the current trial; [0.60, ..., 2.40] – stimulus duration (s).



## BIBLIOGRAPHY

- Abbott, E. A. (1884). *Flatland: A romance of many dimensions* (Vol. s12-III). <https://doi.org/10.1093/nq/s12-III.68.299-k>
- Adler, J., Hazelbauer, G. L., & Dahl, M. M. (1973). Chemotaxis toward sugars in *Escherichia coli*. *Journal of Bacteriology*, 115(3), 824–847. <https://doi.org/10.1128/jb.115.3.824-847.1973>
- Adler, J. (1966). Chemotaxis in Bacteria. *Science*, 153(3737), 708–716.
- Adler, J., & Tso, W.-W. (1974). "Decision"-Making in Bacteria: Chemotactic Response of *Escherichia coli* to Conflicting Stimuli. *Science*, 184, 8–10.
- Adolfsson, A., Ackerman, M., & Brownstein, N. C. (2019). To cluster, or not to cluster: An analysis of clusterability methods. *Pattern Recognit.*, 88, 13–26.
- Akrami, A., Kopec, C. D., Diamond, M. E., & Brody, C. D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature*, 554(7692), 368–372.
- Alexander, G. E., Crutcher, M. D., & DeLong, M. R. (1990). Basal ganglia-thalamocortical circuits: Parallel substrates for motor, oculomotor, "prefrontal" and "limbic" functions. *Prog. Brain Res.*, 85, 119–146.
- Alloway, K. D., Smith, J. B., & Watson, G. D. (2014). Thalamostriatal projections from the medial posterior and parafascicular nuclei have distinct topographic and physiologic properties. *Journal of Neurophysiology*, 111(1), 36–50. <https://doi.org/10.1152/jn.00399.2013>
- Ardiel, E. L., & Rankin, C. H. (2010). An elegant mind: Learning and memory in *Caenorhabditis elegans*. *Handbook of Behavioral Neuroscience*,

## BIBLIOGRAPHY

---

- 22(604), 91–111. <https://doi.org/10.1016/B978-0-12-415823-8.00009-5>
- Arendt, D., Tosches, M. A., & Marlow, H. (2016). From nerve net to nerve ring, nerve cord and brain-evolution of the nervous system. *Nature Reviews Neuroscience*, 17(1), 61–72. <https://doi.org/10.1038/nrn.2015.15>
- Aronov, D., & Fee, M. S. (2011). Analyzing the dynamics of brain circuits with temperature: Design and implementation of a miniature thermoelectric device. *J. Neurosci. Methods*, 197(1), 32–47.
- Ashby, W. R. (1960). *Design for a brain*.
- Ashourian, P., & Loewenstein, Y. (2011). Bayesian inference underlies the contraction bias in delayed comparison tasks.
- Bakhurin, K. I., Goudar, V., Shobe, J. L., Claar, L. D., Buonomano, D. V., & Masmanidis, S. C. (2017). Differential encoding of time by prefrontal and striatal network dynamics. *J. Neurosci.*, 37(4), 854–870.
- Banerjee, A., Egger, R., & Long, M. A. (2021). Using focal cooling to link neural dynamics and behavior. *Neuron*, 109(16), 2508–2518.
- Bargmann, C. I., Hartwig, E., & Horvitz, H. R. (1993). Odorant-selective genes and neurons mediate olfaction in *C. elegans*. *Cell*, 74(3), 515–527. [https://doi.org/10.1016/0092-8674\(93\)80053-H](https://doi.org/10.1016/0092-8674(93)80053-H)
- Bargmann, C. I., & Horvitz, H. R. (1991). Chemosensory neurons with overlapping functions direct chemotaxis to multiple chemicals in *C. elegans*. *Neuron*, 7(5), 729–742. [https://doi.org/10.1016/0896-6273\(91\)90276-6](https://doi.org/10.1016/0896-6273(91)90276-6)
- Barlow, H. B. (1961). Possible Principles Underlying the Transformations of Sensory Messages. In *Sensory communication* (pp. 216–234). <https://doi.org/10.7551/mitpress/9780262518420.003.0013>
- Barnes, T. D., Kubota, Y., Hu, D., Jin, D. Z., & Graybiel, A. M. (2005). Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature*, 437(7062), 1158–1161. <https://doi.org/10.1038/nature04053>
- Beaubaton, D., Amato, G., Trouche, E., & Legallet, E. (1980). Effects of putamen cooling on the latency, speed and accuracy of a pointing movement in the baboon. *Brain Research*, 196(2), 572–576. [https://doi.org/10.1016/0006-8993\(80\)90427-8](https://doi.org/10.1016/0006-8993(80)90427-8)
- Bengston, S. (2002). Origins and early evolution of predation. *Palaeontological society papers*, 57(3), 457–468. <https://doi.org/10.1111/pala.12105>

- Benito-Gutiérrez, È., Gattoni, G., Stemmer, M., Rohr, S. D., Schuhmacher, L. N., Tang, J., Marconi, A., Jékely, G., & Arendt, D. (2021). The dorsoanterior brain of adult amphioxus shares similarities in expression profile and neuronal composition with the vertebrate telencephalon. *BMC Biology*, 19(1), 1–19. <https://doi.org/10.1186/s12915-021-01045-w>
- Bennett, M. S. (2021). Five Breakthroughs: A First Approximation of Brain Evolution From Early Bilaterians to Humans. *Frontiers in Neuroanatomy*, 15(August), 1–34. <https://doi.org/10.3389/fnana.2021.693346>
- Berg, H., & Brown, D. (1972). Chemotaxis in *Escherichia coli* analysed by three-dimensional tracking. *Nature*, 239(672), 2327–2332.
- Berg, H. C. (2003). The rotary motor of bacterial flagella. *Annual Review of Biochemistry*, 72, 19–54. <https://doi.org/10.1146/annurev.biochem.72.121801.161737>
- Boboeva, V., Pezzotta, A., Clopath, C., & Akrami, A. (2023). *From recency to central tendency biases in working memory: A unifying network model*.
- Bodznick, D., Montgomery, J. C., & Carey, M. (1999). Adaptive mechanisms in the elasmobranch hindbrain. *Journal of Experimental Biology*, 202(10), 1357–1364. <https://doi.org/10.1242/jeb.202.10.1357>
- Borkovich, K. A., Alex, L. A., & Simon, M. I. (1992). Attenuation of sensory receptor signaling by covalent modification. *Proceedings of the National Academy of Sciences of the United States of America*, 89(15), 6756–6760. <https://doi.org/10.1073/pnas.89.15.6756>
- Braitenberg, V. (1984). *Vehicles: experiments in synthetic psychology*.
- Buonomano, D. V., & Merzenich, M. M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. *Science*, 267(5200), 1028–1030.
- Buonomano, D. V. (2003). Timing of neural responses in cortical organotypic slices. *Proceedings of the National Academy of Sciences of the United States of America*, 100(8), 4897–4902. <https://doi.org/10.1073/pnas.0736909100>
- Butler, A. B. (2008). Evolution of the thalamus: A morphological and functional review. *Thalamus and Related Systems*, 4(1), 35–58. <https://doi.org/10.1017/S1472928808000356>

## BIBLIOGRAPHY

---

- Calhoun, A. J., Chalasani, S. H., & Sharpee, T. O. (2014). Maximally informative foraging by *Caenorhabditis elegans*. *eLife*, 3, 1–13. <https://doi.org/10.7554/elife.04220>
- Calhoun, A. J., & El Hady, A. (2023). Everyone knows what behavior is but they just don't agree on it. *iScience*, 26(11), 108210. <https://doi.org/10.1016/j.isci.2023.108210>
- Cannon, W. B. (1932). The wisdom of the body.
- Chalasani, S. H., Chronis, N., Tsunozaki, M., Gray, J. M., Ramot, D., Goodman, M. B., & Bargmann, C. I. (2007). Dissecting a circuit for olfactory behaviour in *Caenorhabditis elegans*. *Nature*, 450(7166), 63–70. <https://doi.org/10.1038/nature06292>
- Charnov, E. (1978). Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, 752(4), 739–752.
- Chase, D. L., & Koelle, M. R. (2007). Biogenic amine neurotransmitters in *C. elegans*. *WormBook : the online review of C. elegans biology*, 1–15. <https://doi.org/10.1895/wormbook.1.132.1>
- Chase, D. L., Pepper, J. S., & Koelle, M. R. (2004). Mechanism of extrasynaptic dopamine signaling in *Caenorhabditis elegans*. *Nature Neuroscience*, 7(10), 1096–1103. <https://doi.org/10.1038/nn1316>
- Cisek, P. (2019). Resynthesizing behavior through phylogenetic refinement. *Attention, Perception, and Psychophysics*, 81(7), 2265–2287. <https://doi.org/10.3758/s13414-019-01760-1>
- Cisek, P. (2021). Evolution of behavioural control from chordates to primates. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 1–31.
- Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual Review of Neuroscience*, 33, 269–298. <https://doi.org/10.1146/annurev.neuro.051508.135409>
- Cluzel, P., Surette, M., & Leibler, S. (2000). An ultrasensitive bacterial motor revealed by monitoring signaling proteins in single cells. *Science*, 287(5458), 1652–1655. <https://doi.org/10.1126/science.287.5458.1652>
- Cook, S. J., Jarrell, T. A., Brittin, C. A., Wang, Y., Bloniarz, A. E., Yakovlev, M. A., Nguyen, K. C., Tang, L. T., Bayer, E. A., Duerr, J. S., Bülow, H. E., Hobert, O., Hall, D. H., & Emmons, S. W. (2019). Whole-animal connectomes of both *Caenorhabditis elegans* sexes. *Nature*, 571(7763), 63–71. <https://doi.org/10.1038/s41586-019-1352-7>

- Coull, J. T., Vidal, F., Nazarian, B., & Macar, F. (2004). Functional anatomy of the attentional modulation of time estimation. *Science*, 303(5663), 1506–1508.
- Croll, N. A. (1975). Components and patterns in the behaviour of the nematode *Caenorhabditis elegans*. *Journal of Zoology*, 176(2), 159–176. <https://doi.org/10.1111/j.1469-7998.1975.tb03191.x>
- Crowe, D. A., Zarco, W., Bartolo, R., & Merchant, H. (2014). Dynamic representation of the temporal and sequential structure of rhythmic movements in the primate medial premotor cortex. *J. Neurosci.*, 34(36), 11972–11983.
- Cruz, B. F., Guiomar, G., Soares, S., Motiwala, A., Machens, C. K., & Paton, J. J. (2022). Action suppression reveals opponent parallel control via striatal circuits. *Nature*, 607(7919), 521–526.
- D’aniello, S., Bertrand, S., & Escriva, H. (2023). The Natural History of Model Organisms: *Amphioxus* as a model to study the evolution of development in chordates. *eLife*, 12, 1–15. <https://doi.org/10.7554/ELIFE.87028>
- Dawkins, R. (1976). *Growing points in ethology*.
- De Kock, R., Gladhill, K. A., Ali, M. N., Joiner, W. M., & Wiener, M. (2021). How movements shape the perception of time. *Trends Cogn. Sci.*, 25(11), 950–963.
- Deniau, J. M., & Chevalier, G. (1985). Disinhibition as a basic process in the expression of striatal functions. *Brain Research*, 13(7), 277–280. [https://doi.org/10.1016/0166-2236\(90\)90109-N](https://doi.org/10.1016/0166-2236(90)90109-N)
- Desmurget, M., & Turner, R. S. (2010). Motor Sequences and the Basal Ganglia: Kinematics, Not Habits. *Journal of Neuroscience*, 30(22), 7685–7690. <https://doi.org/10.1523/JNEUROSCI.0163-10.2010>
- Dhawale, A. K., Wolff, S. B. E., Ko, R., & Ölveczky, B. P. (2019). The basal ganglia can control learned motor sequences independently of motor cortex. *bioRxiv*, 827261. [https://www.biorxiv.org/content/10.1101/827261](https://www.biorxiv.org/content/10.1101/827261v1%0Ahttps://www.biorxiv.org/content/10.1101/827261v1.abstract)  
<https://www.biorxiv.org/content/10.1101/827261v1.abstract>
- Dhawale, A. K., Wolff, S. B. E., Ko, R., & Ölveczky, B. P. (2021). The basal ganglia control the detailed kinematics of learned motor skills. *Nat. Neurosci.*, 24(9), 1256–1269.
- Donnelly, N. A., Paulsen, O., Robbins, T. W., & Dalley, J. W. (2015). Ramping single unit activity in the medial prefrontal cortex and ventral

- striatum reflects the onset of waiting but not imminent impulsive actions. *Eur. J. Neurosci.*, 41(12), 1524–1537.
- Doya, K. (1999). What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Netw.*, 12(7), 961–974.
- Drugowitsch, J., Mendonça, A. G., Mainen, Z. F., & Pouget, A. (2019). Learning optimal decisions with confidence. *Proceedings of the National Academy of Sciences of the United States of America*, 116(49), 24872–24880. <https://doi.org/10.1073/pnas.1906787116>
- Egger, R., Tupikov, Y., Elmaleh, M., Katlowitz, K. A., Benezra, S. E., Picardo, M. A., Moll, F., Kornfeld, J., Jin, D. Z., & Long, M. A. (2020). Local axonal conduction shapes the spatiotemporal properties of neural sequences. *Cell*, 183(2), 537–548.e12.
- Emmons, E. B., De Corte, B. J., Kim, Y., Parker, K. L., Matell, M. S., & Narayanan, N. S. (2017). Rodent medial frontal control of temporal processing in the dorsomedial striatum. *J. Neurosci.*, 37(36), 8718–8733.
- Ericsson, J., Silberberg, G., Robertson, B., Wikström, M. A., & Grillner, S. (2011). Striatal cellular properties conserved from lampreys to mammals. *Journal of Physiology*, 589(12), 2979–2992. <https://doi.org/10.1113/jphysiol.2011.209643>
- Ericsson, J., Stephenson-Jones, M., Pérez-Fernández, J., Robertson, B., Silberberg, G., & Grillner, S. (2013). Dopamine differentially modulates the excitability of striatal neurons of the direct and indirect pathways in lamprey. *Journal of Neuroscience*, 33(18), 8045–8054. <https://doi.org/10.1523/JNEUROSCI.5881-12.2013>
- Fassihi, A., Akrami, A., Esmaeili, V., & Diamond, M. E. (2014). Tactile perception and working memory in rats and humans. *Proc. Natl. Acad. Sci. U. S. A.*, 111(6), 2331–2336.
- Faumont, S., Lindsay, T. H., & Lockery, S. R. (2012). Neuronal microcircuits for decision making in *C. elegans*. *Current Opinion in Neurobiology*, 22(4), 580–591. <https://doi.org/10.1016/j.conb.2012.05.005>
- Finnerty, G. T., Shadlen, M. N., Jazayeri, M., Nobre, A. C., & Buonomano, D. V. (2015). Time in cortical circuits. *J. Neurosci.*, 35(41), 13912–13916.
- Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete Coding of Reward Dopamine Neurons. *Science*, 299(March), 1898–1902. <https://doi.org/10.1126/science.1077349>

- Foster, N. N., Barry, J., Korobkova, L., Garcia, L., Gao, L., Becerra, M., Sherafat, Y., Peng, B., Li, X., Choi, J. H., Gou, L., Zingg, B., Azam, S., Lo, D., Khanjani, N., Zhang, B., Stanis, J., Bowman, I., Cotter, K., ... Dong, H. W. (2021). The mouse cortico–basal ganglia–thalamic network. *Nature*, 598(7879), 188–194. <https://doi.org/10.1038/s41586-021-03993-3>
- Fox, D. (2016). What sparked the Cambrian explosion? *Nature*, 530(7590), 268–270. <https://doi.org/10.1038/530268a>
- Fraenkel, G., & Gunn, D. (1940). The Orientation of Animals.
- Freeman, J. S., Cody, F. W., O’Boyle, D. J., Craufurd, D., Neary, D., & Snowden, J. S. (1996). Abnormalities of motor timing in huntington’s disease. *Parkinsonism Relat. Disord.*, 2(2), 81–93.
- Freeze, B. S., Kravitz, A. V., Hammack, N., Berke, J. D., & Kreitzer, A. C. (2013). Control of Basal Ganglia Output by Direct and Indirect Pathway Projection Neurons. *Journal of Neuroscience*, 33(47), 18531–18539. <https://doi.org/10.1523/JNEUROSCI.1278-13.2013>
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychol. Rev.*, 107(2), 289–344.
- Gerfen, C. R., & Surmeier, D. J. (2011). Modulation of striatal projection systems by dopamine. *Annual Review of Neuroscience*, 34, 441–66. <https://doi.org/10.1146/annurev-neuro-061010-113641>
- Gershman, S. J., Moustafa, A. A., & Ludvig, E. A. (2014). Time representation in reinforcement learning models of the basal ganglia. *Frontiers in Computational Neuroscience*, 7(January), 194. <https://doi.org/10.3389/fncom.2013.00194>
- Ghosh, D. D., Sanders, T., Hong, S., McCurdy, L. Y., Chase, D. L., Cohen, N., Koelle, M. R., & Nitabach, M. N. (2016). Neural Architecture of Hunger-Dependent Multisensory Decision Making in *C. elegans*. *Neuron*, 92(5), 1049–1062. <https://doi.org/10.1016/j.neuron.2016.10.030>
- Gibbon, J. (1977). Scalar expectancy theory and weber’s law in animal timing. *Psychol. Rev.*, 84(3), 279.
- Gibbon, J. (1991). Origins of scalar timing. *Learn. Motiv.*, 22(1), 3–38.
- Gjorgjieva, J., Sompolinsky, H., & Meister, M. (2014). Benefits of pathway splitting in sensory coding. *Journal of Neuroscience*, 34(36), 12127–12144. <https://doi.org/10.1523/JNEUROSCI.1032-14.2014>

## BIBLIOGRAPHY

---

- Gouvêa, T. S., Monteiro, T., Motiwala, A., Soares, S., Machens, C., & Paton, J. J. (2015). Striatal dynamics explain duration judgments. *Elife*, 4.
- Gouvêa, T. S., Monteiro, T., Soares, S., Atallah, B. V., & Paton, J. J. (2014). Ongoing behavior predicts perceptual report of interval duration. *Front. Neurobot.*, 8, 10.
- Gray, J. M., Hill, J. J., & Bargmann, C. I. (2005). A circuit for navigation in *Caenorhabditis elegans*. *Proceedings of the National Academy of Sciences of the United States of America*, 102(9), 3184–3191. <https://doi.org/10.1073/pnas.0409009101>
- Graybiel, A. M. (1998). The basal ganglia and chunking of action repertoires. *Neurobiol. Learn. Mem.*, 70(1-2), 119–136.
- Graybiel, A. M., & Ragsdale, C. W. (1978). Histochemically distinct compartments in the striatum of human, monkey, and cat demonstrated by acetylthiocholinesterase staining. *Proceedings of the National Academy of Sciences of the United States of America*, 75(11), 5723–5726. <https://doi.org/10.1073/pnas.75.11.5723>
- Grebe, T. W., & Stock, J. (1998). Bacterial chemotaxis: The five sensors of a bacterium. *Current Biology*, 8(5), 154–157. [https://doi.org/10.1016/s0960-9822\(98\)00098-0](https://doi.org/10.1016/s0960-9822(98)00098-0)
- Griffiths, G. (2007). Cell evolution and the problem of membrane topology. *Nature Reviews Molecular Cell Biology*, 8(12), 1018–1024. <https://doi.org/10.1038/nrm2287>
- Grillner, S., & Robertson, B. (2016). The Basal Ganglia Over 500 Million Years. *Current Biology*, 26(20), R1088–R1100. <https://doi.org/10.1016/j.cub.2016.06.041>
- Hamaguchi, K., Tanaka, M., & Mooney, R. (2016). A distributed recurrent network contributes to temporally precise vocalizations. *Neuron*, 91(3), 680–693.
- Hamilos, A. E., Spedicato, G., Hong, Y., Sun, F., Li, Y., & Assad, J. A. (2021). Slowly evolving dopaminergic activity modulates the moment-to-moment probability of reward-related self-timed movements. *eLife*, 10, 1–38. <https://doi.org/10.7554/eLife.62583>
- Hartigan, J. A., & Hartigan, P. M. (1985). The Dip Test of Unimodality. *The Annals of Statistics*, 13(1), 70–84. <https://doi.org/10.1214/aos/1176346577>

- Hazelbauer, G. L. (2012). Bacterial chemotaxis: The early years of molecular studies. *Annual Review of Microbiology*, 66, 285–303. <https://doi.org/10.1146/annurev-micro-092611-150120>
- Hebb, D. O. (1949). The Organization of Behavior; A Neuropsychological Theory. *The American Journal of Psychology*, 63(4), 633. <https://doi.org/10.2307/1418888>
- Hernández, A., Salinas, E., Garcia, R., & Romo, R. (1997). Discrimination in the sense of flutter: New psychophysical measurements in monkeys. *J. Neurosci.*, 17(16), 6391–6400.
- Hidalgo-Balbuena, A. E., Luma, A. Y., Pimentel-Farfan, A. K., Peña-Rangel, T., & Rueda-Orozco, P. E. (2019). Sensory representations in the striatum provide a temporal reference for learning and executing motor habits. *Nat. Commun.*, 10(1), 1–15.
- Hills, T., Brockie, P. J., & Maricq, A. V. (2004). Dopamine and Glutamate Control Area-Restricted Search Behavior in *Caenorhabditis elegans*. *Journal of Neuroscience*, 24(5), 1217–1225. <https://doi.org/10.1523/JNEUROSCI.1569-03.2004>
- Hills, T. T. (2006). Animal foraging and the evolution of goal-directed cognition. *Cognitive Science*, 30(1), 3–41. [https://doi.org/10.1207/s15516709cog0000\\_50](https://doi.org/10.1207/s15516709cog0000_50)
- Holland, L. Z. (2014). Genomics, evolution and development of amphioxus and tunicates: The Goldilocks principle. *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution*, 324(4), 342–352. <https://doi.org/10.1002/jez.b.22569>
- Holland, L. Z. (2015). The origin and evolution of chordate nervous systems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1684). <https://doi.org/10.1098/rstb.2015.0048>
- Hollingworth, H. L. (1910). The central tendency of judgment. *The Journal of Philosophy, Psychology and Scientific Methods*, 7(17), 461–469.
- Hukema, R. K., Rademakers, S., & Jansen, G. (2008). Gustatory plasticity in *C. elegans* involves integration of negative cues and NaCl taste mediated by serotonin, dopamine, and glutamate. *Learning and Memory*, 15(11), 829–836. <https://doi.org/10.1101/lm.994408>
- Hunnicutt, B. J., Jongbloets, B. C., Birdsong, W. T., Gertz, K. J., Zhong, H., & Mao, T. (2016). A comprehensive excitatory input map of the striatum reveals novel functional organization. *eLife*, 5, e19103. <https://doi.org/10.7554/eLife.19103>

## BIBLIOGRAPHY

---

- Iino, Y., Sawada, T., Yamaguchi, K., Tajiri, M., Ishii, S., Kasai, H., & Yagishita, S. (2020). Dopamine D2 receptors in discrimination learning and spine enlargement. *Nature*, (June 2019), 1–6. <https://doi.org/10.1038/s41586-020-2115-1>
- Iino, Y., & Yoshida, K. (2009). Parallel use of two behavioral mechanisms for chemotaxis in *Caenorhabditis elegans*. *Journal of Neuroscience*, 29(17), 5370–5380. <https://doi.org/10.1523/JNEUROSCI.3633-08.2009>
- Isa, T., Marquez-Legorreta, E., Grillner, S., & Scott, E. K. (2021). The tectum/superior colliculus as the vertebrate solution for spatial sensory integration and action. *Current Biology*, 31(11), R741–R762. <https://doi.org/10.1016/j.cub.2021.04.001>
- Ishihara, T., Iino, Y., Mohri, A., Mori, I., Gengyo-Ando, K., Mitani, S., & Katsura, I. (2002). HEN-1, a secretory protein with an LDL receptor motif, regulates sensory integration and learning in *Caenorhabditis elegans*. *Cell*, 109(5), 639–649. [https://doi.org/10.1016/S0092-8674\(02\)00748-1](https://doi.org/10.1016/S0092-8674(02)00748-1)
- Izquierdo, E. J., & Lockery, S. R. (2010). Evolution and analysis of minimal neural circuits for klinotaxis in *Caenorhabditis elegans*. *Journal of Neuroscience*, 30(39), 12908–12917. <https://doi.org/10.1523/JNEUROSCI.2606-10.2010>
- Jacobs, L. F. (2012). From chemotaxis to the cognitive map: The function of olfaction. *Proceedings of the National Academy of Sciences of the United States of America*, 109(SUPPL.1), 10693–10700. <https://doi.org/10.1073/pnas.1201880109>
- Jazayeri, M., & Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nat. Neurosci.*, 13(8), 1020–1026.
- Jazayeri, M., & Shadlen, M. N. (2015). A neural mechanism for sensing and reproducing a time interval. *Curr. Biol.*, 25(20), 2599–2609.
- Jékely, G. (2011). Origin and early evolution of neural circuits for the control of ciliary locomotion. *Proceedings of the Royal Society B: Biological Sciences*, 278(1707), 914–922. <https://doi.org/10.1098/rspb.2010.2027>
- Jékely, G., Godfrey-Smith, P., & Keijzer, F. (2021). Reafference and the origin of the self in early nervous system evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 376(1821), 20190764.

- Jin, D. Z., Fujii, N., & Graybiel, A. M. (2009). Neural representation of time in cortico-basal ganglia circuits. *Proc. Natl. Acad. Sci. U. S. A.*, 106(45), 19156–19161.
- Jin, X., & Costa, R. M. (2010). Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature*, 466(7305), 457–462. <https://doi.org/10.1038/nature09263>
- Joel, D., Niv, Y., & Ruppín, E. (2002). Actor-critic models of the basal ganglia: new anatomical and computational perspectives. *Neural Networks*, 15(4-6), 535–547. [https://doi.org/10.1016/S0893-6080\(02\)00047-3](https://doi.org/10.1016/S0893-6080(02)00047-3)
- Johansson, Y., & Silberberg, G. (2020). The Functional Organization of Cortical and Thalamic Inputs onto Five Types of Striatal Neurons Is Determined by Source and Target Cell Identities. *Cell Reports*, 30(4), 1178–1194.e3. <https://doi.org/10.1016/j.celrep.2019.12.095>
- Jun, J. J., Steinmetz, N. A., Siegle, J. H., Denman, D. J., Bauza, M., Barbarits, B., Lee, A. K., Anastassiou, C. A., Andrei, A., Aydın, Ç., Barbic, M., Blanche, T. J., Bonin, V., Couto, J., Dutta, B., Gratiy, S. L., Gutnisky, D. A., Häusser, M., Karsh, B., . . . Harris, T. D. (2017). Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551(7679), 232–236.
- Jurado-Parras, M.-T., Sfaie, M., Sarno, S., Louis, J., Karoutchi, C., Berret, B., & Robbe, D. (2020). The dorsal striatum energizes motor routines. *Curr. Biol.*, 30(22), 4362–4372.e6.
- Kaas, J. H. (1997). Topographic maps are fundamental to sensory processing. *Brain Research Bulletin*, 44(2), 107–112. [https://doi.org/10.1016/S0361-9230\(97\)00094-4](https://doi.org/10.1016/S0361-9230(97)00094-4)
- Kalinin, Y., Neumann, S., Sourjik, V., & Wu, M. (2010). Responses of *Escherichia coli* bacteria to two opposing chemoattractant gradients depend on the chemoreceptor ratio. *Journal of Bacteriology*, 192(7), 1796–1800. <https://doi.org/10.1128/JB.01507-09>
- Kaplan, H. S., Salazar Thula, O., Khoss, N., & Zimmer, M. (2020). Nested Neuronal Dynamics Orchestrate a Behavioral Hierarchy across Timescales. *Neuron*, 105(3), 562–576.e9. <https://doi.org/10.1016/j.neuron.2019.10.037>
- Kato, S., Xu, Y., Cho, C. E., Abbott, L. F., & Bargmann, C. I. (2014). Temporal Responses of *C.elegans* Chemosensory Neurons Are Preserved in Behavioral Dynamics. *Neuron*, 81(3), 616–628. <https://doi.org/10.1016/j.neuron.2013.11.020>

## BIBLIOGRAPHY

---

- Kawai, R., Markman, T., Poddar, R., Ko, R., Fantana, A. L., Dhawale, A. K., Kampff, A. R., & Ölveczky, B. P. (2015). Motor cortex is required for learning but not for executing a motor skill. *Neuron*, 86(3), 800–812.
- Kawano, T., Po, M. D., Gao, S., Leung, G., Ryu, W. S., & Zhen, M. (2011). An imbalancing act: Gap junctions reduce the backward motor circuit activity to bias *C. elegans* for forward locomotion. *Neuron*, 72(4), 572–586. <https://doi.org/10.1016/j.neuron.2011.09.005>
- Killeen, P. R., & Fetterman, J. G. (1988). A behavioral theory of timing. *Psychol. Rev.*, 95(2), 274–295.
- Kim, J., Ghim, J.-W., Lee, J. H., & Jung, M. W. (2013). Neural correlates of interval timing in rodent prefrontal cortex. *J. Neurosci.*, 33(34), 13834–13847.
- Kim, J., Kim, D., & Jung, M. W. (2018). Distinct dynamics of striatal and prefrontal neural activity during temporal discrimination. *Front. Integr. Neurosci.*, 12, 34.
- Kim, S., Tathireddy, P., Normann, R. A., & Solzbacher, F. (2007). Thermal impact of an active 3-D microelectrode array implanted in the brain. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 15(4), 493–501.
- Kindt, K. S., Quast, K. B., Giles, A. C., De, S., Hendrey, D., Nicastro, I., Rankin, C. H. H., & Schafer, W. R. (2007). Dopamine Mediates Context-Dependent Modulation of Sensory Plasticity in *C. elegans*. *Neuron*, 55(4), 662–676. <https://doi.org/10.1016/j.neuron.2007.07.023>
- Kiontke, K. C., Félix, M. A., Ailion, M., Rockman, M. V., Braendle, C., Pénigault, J. B., & Fitch, D. H. (2011). A phylogeny and molecular barcodes for *Caenorhabditis*, with numerous new species from rotting fruits. *BMC Evolutionary Biology*, 11(1). <https://doi.org/10.1186/1471-2148-11-339>
- Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C. E., Kepecs, A., Mainen, Z. F., Qi, X.-L., Romo, R., Uchida, N., & Machens, C. K. (2016). Demixed principal component analysis of neural population data. *Elife*, 5.
- Kravitz, A. V., Freeze, B. S., Parker, P. R., Kay, K., Thwin, M. T., Deisseroth, K., & Kreitzer, A. C. (2010). Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry. *Nature*, 466(7306), 622–626. <https://doi.org/10.1038/nature09159>

- Kreitzer, A. C., & Malenka, R. C. (2007). Endocannabinoid-mediated rescue of striatal LTD and motor deficits in Parkinson's disease models. *Nature*, 445(7128), 643–647. <https://doi.org/10.1038/nature05506>
- Kuramoto, E., Fujiyama, F., Nakamura, K. C., Tanaka, Y., Hioki, H., & Kaneko, T. (2011). Complementary distribution of glutamatergic cerebellar and GABAergic basal ganglia afferents to the rat motor thalamic nuclei. *European Journal of Neuroscience*, 33(1), 95–109. <https://doi.org/10.1111/j.1460-9568.2010.07481.x>
- Lacalli, T. (2018). Amphioxus, motion detection, and the evolutionary origin of the vertebrate retinotectal map. *EvoDevo*, 9(1), 4–8. <https://doi.org/10.1186/s13227-018-0093-2>
- Lacalli, T., & Candiani, S. (2017). Locomotory control in amphioxus larvae: New insights from neurotransmitter data. *EvoDevo*, 8(1), 1–8. <https://doi.org/10.1186/s13227-017-0067-9>
- Lau, B., & Glimcher, P. W. (2008). Value representations in the primate striatum during matching behavior. *Neuron*, 58(3), 451–463.
- Lau, B., Monteiro, T., & Paton, J. J. (2017). The many worlds hypothesis of dopamine reward prediction error: implications of a parallel circuit architecture in the basal ganglia. *Current Opinion in Neurobiology*, 46, 241–247. <https://doi.org/10.1016/j.conb.2017.08.015>
- Lauwereyns, J., Watanabe, K., Coe, B., & Hikosaka, O. (2002). A neural correlate of response bias in monkey caudate nucleus. *Nature*, 418(JULY), 413–417. <https://doi.org/10.1038/nature00844.1>
- 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.
- Liu, H., Yang, W., Wu, T., Duan, F., Soucy, E., Jin, X., & Zhang, Y. (2018). Cholinergic Sensorimotor Integration Regulates Olfactory Steering. *Neuron*, 97(2), 390–405.e3. <https://doi.org/10.1016/j.neuron.2017.12.003>
- Long, M. A., & Fee, M. S. (2008). Using temperature to analyse temporal dynamics in the songbird motor pathway. *Nature*, 456(7219), 189–194.
- Long, M. A., Katlowitz, K. A., Svirsky, M. A., Clary, R. C., Byun, T. M., Majaj, N., Oya, H., Howard, M. A., 3rd, & Greenlee, J. D. W. (2016). Functional segregation of cortical regions underlying speech timing and articulation. *Neuron*, 89(6), 1187–1193.

## BIBLIOGRAPHY

---

- Lopes, G., Bonacchi, N., Frazão, J., Neto, J. P., Atallah, B. V., Soares, S., Moreira, L., Matias, S., Itskov, P. M., Correia, P. A., Medina, R. E., Calcaterra, L., Dreosti, E., Paton, J. J., & Kampff, A. R. (2015). Bonsai: An event-based framework for processing and controlling data streams. *Front. Neuroinform.*, 9, 7.
- Lourenço, J. M. (2021). *The NOVAthesis L<sup>A</sup>T<sub>E</sub>X Template User's Manual*. NOVA University Lisbon. <https://github.com/joaomlourenco/novathesis/raw/main/template.pdf>
- 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. <https://doi.org/10.1162/neco.2008.11-07-654>
- Machado, A., Malheiro, M. T., & Erilagen, W. (2009). Learning to time: A perspective. *J. Exp. Anal. Behav.*, 92(3), 423–458.
- Macnab, R. M., & Koshland, D. E. (1972). The gradient-sensing mechanism in bacterial chemotaxis. *Proceedings of the National Academy of Sciences of the United States of America*, 69(9), 2509–2512. <https://doi.org/10.1073/pnas.69.9.2509>
- Malapani, C., Rakitin, B., Levy, R., Meck, W. H., Deweer, B., Dubois, B., & Gibbon, J. (1998). Coupled temporal memories in parkinson's disease: A dopamine-related dysfunction. *J. Cogn. Neurosci.*, 10(3), 316–331.
- Malapani, C., & Fairhurst, S. (2002). Scalar timing in animals and humans. *Learn. Motiv.*, 33(1), 156–176.
- Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474), 78–84.
- 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.
- Maricq, A., Peckol, E., Driscoll, M., & Bargmann, C. I. (1995). Mechanosensory signalling in *C. elegans* mediated by the GLR-1 glutamate receptor. *Nature*, 378(7), 603–605.
- Markowitz, J. E., Gillis, W. F., Beron, C. C., Neufeld, S. Q., Robertson, K., Bhagat, N. D., Peterson, R. E., Peterson, E., Hyun, M., Linderman, S. W., Sabatini, B. L., & Datta, S. R. (2018). The striatum organizes

- 3D behavior via Moment-to-Moment action selection. *Cell*, 174(1), 44–58.e17.
- Martiros, N., Burgess, A. A., & Graybiel, A. M. (2018). Inversely Active Striatal Projection Neurons and Interneurons Selectively Delimit Useful Behavioral Sequences. *Current Biology*, 28(4), 560–573.e5. <https://doi.org/10.1016/j.cub.2018.01.031>
- Matell, M. S., & Meck, W. H. (2004). Cortico-striatal circuits and interval timing: Coincidence detection of oscillatory processes. *Brain Res. Cogn. Brain Res.*, 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. *Behav. Neurosci.*, 117(4), 760–773.
- Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., & Bethge, M. (2018). DeepLabCut: Markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.*, 21(9), 1281–1289.
- Matsuda, W., Furuta, T., Nakamura, K. C., Hioki, H., Fujiyama, F., Arai, R., & Kaneko, T. (2009). Single nigrostriatal dopaminergic neurons form widely spread and highly dense axonal arborizations in the neostriatum. *Journal of Neuroscience*, 29(2), 444–453. <https://doi.org/10.1523/JNEUROSCI.4029-08.2009>
- Matthews, W. J., Stewart, N., & Wearden, J. H. (2011). Stimulus Intensity and the Perception of Duration. *Journal of Experimental Psychology: Human Perception and Performance*, 37(1), 303–313. <https://doi.org/10.1037/a0019961>
- McCulloch, W. S. (1945). A heterarchy of values determined by the topology of nervous nets. *Bulletin of Mathematical Biophysics*, 89–93.
- McElvain, L. E., Chen, Y., Moore, J. D., Brigidi, G. S., Bloodgood, B. L., Lim, B. K., Costa, R. M., & Kleinfeld, D. (2021). Specific populations of basal ganglia output neurons target distinct brain stem areas while collateralizing throughout the diencephalon. *Neuron*, 109(10), 1721–1738.e4. <https://doi.org/10.1016/j.neuron.2021.03.017>
- McHaffie, J. G., Stanford, T. R., Stein, B. E., Coizet, V., & Redgrave, P. (2005). Subcortical loops through the basal ganglia. *Trends Neurosci.*, 28(8), 401–407.

## BIBLIOGRAPHY

---

- Meck, W. H. (2006). Neuroanatomical localization of an internal clock: A functional link between mesolimbic, nigrostriatal, and mesocortical dopaminergic systems. *Brain Res.*, 1109(1), 93–107.
- Mello, G. B. M., Soares, S., & Paton, J. J. (2015). A scalable population code for time in the striatum. *Curr. Biol.*, 25(9), 1113–1122.
- Mendonça, A. G., Drugowitsch, J., Vicente, M. I., DeWitt, E. E., Pouget, A., & Mainen, Z. F. (2020). The impact of learning on perceptual decisions and its implication for speed-accuracy tradeoffs. *Nature Communications*, 11(1), 1–15. <https://doi.org/10.1038/s41467-020-16196-7>
- Merchant, H., & Averbeck, B. B. (2017). The computational and neural basis of rhythmic timing in medial premotor cortex. *J. Neurosci.*, 37(17), 4552–4564.
- Merel, J., Botvinick, M., & Wayne, G. (2019). Hierarchical motor control in mammals and machines. *Nat. Commun.*, 10(1), 5489.
- Mikhael, J. G., & Gershman, S. J. (2019). Adapting the flow of time with dopamine. *J. Neurophysiol.*, 121(5), 1748–1760.
- Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, 50(4), 381–425. [https://doi.org/10.1016/S0301-0082\(96\)00042-1](https://doi.org/10.1016/S0301-0082(96)00042-1)
- Mink, J. W. (2018). Basal ganglia mechanisms in action selection, plasticity, and dystonia. *European Journal of Paediatric Neurology*, 22(2), 225–229. <https://doi.org/10.1016/j.ejpn.2018.01.005>
- Mizes, K. G., Lindsey, J., Escola, G. S., & Ölveczky, B. P. (2023). Dissociating the contributions of sensorimotor striatum to automatic and visually guided motor sequences. *Nature Neuroscience*, 26. <https://doi.org/10.1038/s41593-023-01431-3>
- Monteiro, T., Rodrigues, F. S., Pexirra, M., Cruz, B. F., Gonçalves, A. I., Rueda-Orozco, P. E., & Paton, J. J. (2023). Using temperature to analyze the neural basis of a time-based decision. *Nat. Neurosci.*
- Montgomery, J. C. (2023). Roles for cerebellum and subsumption architecture in central pattern generation. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, (100). <https://doi.org/10.1007/s00359-023-01634-w>
- Motiwala, A., Soares, S., Atallah, B. V., Paton, J. J., & Machens, C. K. (2022). Efficient coding of cognitive variables underlies dopamine response

- and choice behavior. *Nature Neuroscience*, 25(6), 738–748. <https://doi.org/10.1038/s41593-022-01085-7>
- Murakami, M., Vicente, M. I., Costa, G. M., & Mainen, Z. F. (2014). Neural antecedents of self-initiated actions in secondary motor cortex. *Nat. Neurosci.*, 17(11), 1574–1582.
- Murray, J. M., & Escola, G. S. (2017). Learning multiple variable-speed sequences in striatum via cortical tutoring. *Elife*, 6.
- Murray, J. M., & Escola, G. S. (2020). Remembrance of things practiced with fast and slow learning in cortical and subcortical pathways. *Nature Communications*, 11(1), 1–12. <https://doi.org/10.1038/s41467-020-19788-5>
- Ocaña, F. M., Suryanarayana, S. M., Saitoh, K., Kardamakis, A. A., Capantini, L., Robertson, B., & Grillner, S. (2015). The lamprey pallium provides a blueprint of the mammalian motor projections from cortex. *Current Biology*, 25(4), 413–423. <https://doi.org/10.1016/j.cub.2014.12.013>
- Okobi, D. E., Jr, Banerjee, A., Matheson, A. M. M., Phelps, S. M., & Long, M. A. (2019). Motor cortical control of vocal interaction in neotropical singing mice. *Science*, 363(6430), 983–988.
- Owen, S. F., Liu, M. H., & Kreitzer, A. C. (2019). Thermal constraints on in vivo optogenetic manipulations. *Nat. Neurosci.*, 22(7), 1061–1065.
- Panigrahi, B., Martin, K. A., Li, Y., Graves, A. R., Vollmer, A., Olson, L., Mensh, B. D., Karpova, A. Y., & Dudman, J. T. (2015). Dopamine is required for the neural representation and control of movement vigor. *Cell*, 162(6), 1418–1430.
- Pardo-Vazquez, J. L., Castiñeiras-de Saa, J. R., Valente, M., Damião, I., Costa, T., Vicente, M. I., Mendonça, A. G., Mainen, Z. F., & Renart, A. (2019). The mechanistic foundation of Weber’s law. *Nature Neuroscience*, 22(9), 1493–1502. <https://doi.org/10.1038/s41593-019-0439-7>
- Park, J., Coddington, L. T., & Dudman, J. T. (2020). Basal Ganglia Circuits for Action Specification. *Annual Review of Neuroscience*, 43, 485–507. <https://doi.org/10.1146/annurev-neuro-070918-050452>
- Parker, A. J., & Newsome, W. T. (1998). Sense and the single neuron: Probing the physiology of perception. *Annu. Rev. Neurosci.*, 21, 227–277.
- Paton, J. J., & Buonomano, D. V. (2018). The neural basis of timing: Distributed mechanisms for diverse functions. *Neuron*, 98(4), 687–705.
- Pavlov, I. P. (1927). Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex.

## BIBLIOGRAPHY

---

- Paxinos, G., & Watson, C. (2013). *The rat brain in stereotaxic coordinates: Hard cover edition*. Academic Press.
- Pearl, J. (2009). *Causality*. Cambridge University Press.
- Perez-Fernandez, J., Stephenson-Jones, M., Suryanarayana, S. M., Robertson, B., & Grillner, S. (2014). Evolutionarily Conserved Organization of the Dopaminergic System in Lamprey: SNc/VTA Afferent and Efferent Connectivity and D2 Receptor Expression. *The Journal of Comparative Neurology*.
- Pezzulo, G., Parr, T., & Friston, K. (2021). The evolution of brain architectures for predictive coding and active inference. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 377(1844). <https://doi.org/10.1098/rstb.2020.0531>
- Pierce-Shimomura, J. T., Thomas, M. M., & Lockery, S. R. (1999). The Fundamental Role of Pirouettes in *Caenorhabditis elegans* Chemotaxis. *Journal of Neuroscience*, 4.
- Pinsker, H., Kupfermann, I., Castellucci, V., & Kandel, E. (1970). Habituation and dishabituation of the gill-withdrawal reflex in *Aplysia*. *Science*, 167(3926), 1740–1742. <https://doi.org/10.1126/science.167.3926.1740>
- Popper, K. (1956). The Arrow of Time. *Nature*, (1949), 1956.
- Puelles, L., & Rubenstein, J. L. (2015). A new scenario of hypothalamic organization: Rationale of new hypotheses introduced in the updated prosomeric model. *Frontiers in Neuroanatomy*, 9(MAR), 1–23. <https://doi.org/10.3389/fnana.2015.00027>
- Purcell, E. M. (1976). Life at low Reynolds number. <https://doi.org/10.4310/MRL.2010.v17.n4.a10>
- Rankin, C. H., Beck, C. D., & Chiba, C. M. (1990). *Caenorhabditis elegans*: A new model system for the study of learning and memory. *Behavioural Brain Research*, 37(1), 89–92. [https://doi.org/10.1016/0166-4328\(90\)90074-O](https://doi.org/10.1016/0166-4328(90)90074-O)
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87. <https://doi.org/10.1038/4580>
- Rao, S. M., Mayer, A. R., & Harrington, D. L. (2001). The evolution of brain activation during temporal processing. *Nat. Neurosci.*, 4(3), 317–323.

- Reynolds, J. N., Hyland, B. I., & Wickens, J. R. (2001). A cellular mechanism of reward-related learning. *Nature*, 413(6851), 67–70.
- Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451), 585–590.
- Robertson, R. M., & Money, T. G. A. (2012). Temperature and neuronal circuit function: Compensation, tuning and tolerance. *Curr. Opin. Neurobiol.*, 22(4), 724–734.
- Rodrigues, F. S., Monteiro, T., Motiwala, A., & Paton, J. J. (2023). The dorsolateral striatum encodes a temporal basis for the organization of behavior. *bioRxiv*. <https://doi.org/10.1101/2023.11.13.566826>
- 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.
- Rose, J. K., Kaun, K. R., Chen, S. H., & Rankin, C. H. (2003). GLR-1, a Non-NMDA Glutamate Receptor Homolog, Is Critical for Long-Term Memory in *Caenorhabditis elegans*. *Journal of Neuroscience*, 23(29), 9595–9599. <https://doi.org/10.1523/jneurosci.23-29-09595.2003>
- Roseberry, T. K., Lee, A. M., Lalive, A. L., Wilbrecht, L., Bonci, A., & Kreitzer, A. C. (2016). Cell-Type-Specific Control of Brainstem Locomotor Circuits by Basal Ganglia. *Cell*, 164(3), 526–537. <https://doi.org/10.1016/j.cell.2015.12.037>
- Rowe, K. C., Paulsen, J. S., Langbehn, D. R., Duff, K., Beglinger, L. J., Wang, C., O'Rourke, J. J. F., Stout, J. C., & Moser, D. J. (2010). Self-paced timing detects and tracks change in prodromal huntington disease. *Neuropsychology*, 24(4), 435–442.
- Rueda-Orozco, P. E., Mendoza, E., Hernandez, R., Aceves, J. J., Ibanez-Sandoval, O., Galarraaga, E., & Bargas, J. (2009). Diversity in long-term synaptic plasticity at inhibitory synapses of striatal spiny neurons. *Learning and Memory*, 16(8), 474–478. <https://doi.org/10.1101/lm.1439909>
- Rueda-Orozco, P. E., & Robbe, D. (2015). The striatum multiplexes contextual and kinematic information to constrain motor habits execution. *Nat. Neurosci.*, 18(3), 453–460.
- Sacktor, T. C. (2011). How does PKM $\zeta$  maintain long-term memory? *Nature Reviews Neuroscience*. <http://dx.doi.org/10.1038/nrn2949>

- Saeki, S., Yamamoto, M., & Iino, Y. (2001). Plasticity of chemotaxis revealed by paired presentation of a chemoattractant and starvation in the nematode *Caenorhabditis elegans*. *Journal of Experimental Biology*, 204(10), 1757–1764. <https://doi.org/10.1242/jeb.204.10.1757>
- Safaie, M., Jurado-Parras, M.-T., Sarno, S., Louis, J., Karoutchi, C., Petit, L. F., Pasquet, M. O., Eloy, C., & Robbe, D. (2020). Turning the body into a clock: Accurate timing is facilitated by simple stereotyped interactions with the environment. *Proceedings of the National Academy of Sciences*, 117(23), 13084–13093.
- Samejima, K., Ueda, Y., Doya, K., & Kimura, M. (2005). Representation of action-specific reward values in the striatum. *Science*, 310(5752), 1337–1340.
- Sanyal, S., Wintle, R. F., Kindt, K. S., Nuttley, W. M., Arvan, R., Fitzmaurice, P., Bigras, E., Merzs, D. C., Hébert, T. E., Van Der Kooy, D., Schafer, W. R., Culotti, J. G., & Van Tol, H. H. (2004). Dopamine modulates the plasticity of mechanosensory responses in *Caenorhabditis elegans*. *EMBO Journal*, 23(2), 473–482. <https://doi.org/10.1038/sj.emboj.7600057>
- Sawin, E. R., Ranganathan, R., & Horvitz, H. R. (2000). *C. elegans* locomotory rate is modulated by the environment through a dopaminergic pathway and by experience through a serotonergic pathway. *Neuron*, 26(3), 619–631. [https://doi.org/10.1016/S0896-6273\(00\)81199-X](https://doi.org/10.1016/S0896-6273(00)81199-X)
- Schafer, W. R., & Kenyon, C. J. (1995). A calcium-channel homologue required for adaptation to dopamine and serotonin in *Caenorhabditis elegans*. <https://doi.org/10.1038/375073a0>
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599.
- Schütt, H., Harmeling, S., Macke, J., & Wichmann, F. (2015). Psignifit 4: Pain-free bayesian inference for psychometric functions.
- Schütt, H. H., Harmeling, S., Macke, J. H., & Wichmann, F. A. (2016). Painfree and accurate bayesian estimation of psychometric functions for (potentially) overdispersed data. *Vision Res.*, 122, 105–123.
- Segall, J. E., Block, S. M., & Berg, H. C. (1986). Temporal comparisons in bacterial chemotaxis. *Proceedings of the National Academy of Sciences of the United States of America*, 83(23), 8987–8991. <https://doi.org/10.1073/pnas.83.23.8987>

- Seilacher, A. (1999). Biomat-related lifestyles in the Precambrian. *Palaios*, 14(1), 86–93. <https://doi.org/10.2307/3515363>
- Shen, W., Flajolet, M., Greengard, P., & Surmeier, D. J. (2008). Dichotomous Control Dopaminergic of Striatal Synaptic Plasticity. *Science*, 321(5890), 848–851. <https://doi.org/10.1126/science.1160575>.Dichotomous
- Soares, S., Atallah, B. V., & Paton, J. J. (2016). Midbrain dopamine neurons control judgment of time. *Science*, 354(6317), 1273–1277.
- Steiner, H., & Tseng, K. Y. (2016). *Handbook of basal ganglia structure and function*. Academic Press.
- Stephenson-Jones, M., Ericsson, J., Robertson, B., & Grillner, S. (2012). Evolution of the basal ganglia: Dual-output pathways conserved throughout vertebrate phylogeny. *Journal of Comparative Neurology*, 520(13), 2957–2973. <https://doi.org/10.1002/cne.23087>
- Stephenson-Jones, M., Kardamakis, A. A., Robertson, B., & Grillner, S. (2013). Independent circuits in the basal ganglia for the evaluation and selection of actions. *Proceedings of the National Academy of Sciences of the United States of America*, 110(38). <https://doi.org/10.1073/pnas.1314815110>
- Stephenson-Jones, M., Samuelsson, E., Ericsson, J., Robertson, B., & Grillner, S. (2011). Evolutionary conservation of the basal ganglia as a common vertebrate mechanism for action selection. *Current Biology*, 21(13), 1081–1091. <https://doi.org/10.1016/j.cub.2011.05.001>
- Sulston, J., Dew, M., & Brenner, S. (1975). Dopaminergic Neurons in the Nematode *Caenorhabditis elegans*. *Biological Reviews*, 69(3), 345–417.
- Sulzer, D., Joyce, M. P., Lin, L., Geldwert, D., Haber, S. N., Hattori, T., & Rayport, S. (1998). Dopamine neurons make glutamatergic synapses in vitro. *Journal of Neuroscience*, 18(12), 4588–4602. <https://doi.org/10.1523/jneurosci.18-12-04588.1998>
- Surmeier, D. J., Ding, J., Day, M., Wang, Z., & Shen, W. (2007). D1 and D2 dopamine-receptor modulation of striatal glutamatergic signaling in striatal medium spiny neurons. *Trends in Neurosciences*, 30(5), 228–235. <https://doi.org/10.1016/j.tins.2007.03.008>
- Suryanarayana, S. M., Pérez-Fernández, J., Robertson, B., & Grillner, S. (2020). The evolutionary origin of visual and somatosensory representation

- in the vertebrate pallium. <https://doi.org/10.1038/s41559-020-1137-2>
- Suryanarayana, S. M., Robertson, B., & Grillner, S. (2021). The neural bases of vertebrate motor behaviour through the lens of evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 377(1844). <https://doi.org/10.1098/rstb.2020.0521>
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (Second). The MIT Press. <http://incompleteideas.net/book/the-book-2nd.html>
- Suzuki, H., Kerr, R., Bianchi, L., Frøkjær-Jensen, C., Slone, D., Xue, J., Gerstbrein, B., Driscoll, M., & Schafer, W. R. (2003). In vivo imaging of *C. elegans* mechanosensory neurons demonstrates a specific role for the MEC-4 channel in the process of gentle touch sensation. *Neuron*, 39(6), 1005–1017. <https://doi.org/10.1016/j.neuron.2003.08.015>
- Suzuki, H., Thiele, T. R., Faumont, S., Ezcurra, M., Lockery, S. R., & Schafer, W. R. (2008). Functional asymmetry in *Caenorhabditis elegans* taste neurons and its computational role in chemotaxis. *Nature*, 454(7200), 114–117. <https://doi.org/10.1038/nature06927>
- Swaney, K. F., Huang, C. H., & Devreotes, P. N. (2010). Eukaryotic chemotaxis: A network of signaling pathways controls motility, directional sensing, and polarity. *Annual Review of Biophysics*, 39(1), 265–289. <https://doi.org/10.1146/annurev.biophys.093008.131228>
- Tai, L.-H., Lee, A. M., Benavidez, N., Bonci, A., & Willbrecht, L. (2012). Transient stimulation of distinct subpopulations of striatal neurons mimics changes in action value. *Nat. Neurosci.*, 15(9), 1281–1289.
- Tang, L. S., Goeritz, M. L., Caplan, J. S., Taylor, A. L., Fisek, M., & Marder, E. (2010). Precise temperature compensation of phase in a rhythmic motor pattern.
- Tecuapetla, F., Carrillo-Reid, L., Bargas, J., & Galarraga, E. (2007). Dopaminergic modulation of short-term synaptic plasticity at striatal inhibitory synapses. *Proceedings of the National Academy of Sciences*, 104(24), 10258–63. <https://doi.org/10.1073/pnas.0703813104>
- Tessier-Lavigne, M., & Goodman, C. S. (1996). The Molecular Biology of Axon Guidance Published by : American Association for the Advancement of Science The Molecular Biology of Axon Guidance. *Science*, 274(5290), 1123–1133.

- Thiele, T. R., Faumont, S., & Lockery, S. R. (2009). The neural network for chemotaxis to tastants in *Caenorhabditis elegans* is specialized for temporal differentiation. *Journal of Neuroscience*, 29(38), 11904–11911. <https://doi.org/10.1523/JNEUROSCI.0594-09.2009>
- Tinbergen, N. (1951). The study of instinct.
- Toso, A., Fassihi, A., Paz, L., Pulecchi, F., & Diamond, M. E. (2021). A sensory integration account for time perception. *PLoS Comput. Biol.*, 17(1), e1008668.
- Toso, A., Reinartz, S., Pulecchi, F., & Diamond, M. E. (2021). Time coding in rat dorsolateral striatum. *Neuron*, 109(22), 3663–3673.e6.
- Toso, A., Reinartz, S., Pulecchi, F., & Diamond, M. E. (2023). Erratum: Time coding in rat dorsolateral striatum. *Neuron*, 111(4), 585–594. <https://doi.org/10.1016/j.neuron.2023.01.012>
- Treisman, M. (1963). Temporal discrimination and the indifference interval: Implications for a model of the “internal clock”. *Psychological Monographs: General and Applied*, 77(13), 1–31.
- Turner, R. S., & Desmurget, M. (2010). Basal ganglia contributions to motor control: A vigorous tutor. *Curr. Opin. Neurobiol.*, 20(6), 704–716.
- Uchimura, N., & North, R. A. (1989). in Rat Nucleus Accumbens Neurons. *Journal of Neurophysiology*, 62(6), 1280–1286.
- Urai, A. E., Braun, A., & Donner, T. H. (2017). Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nat. Commun.*, 8, 14637.
- Vidal-Gadea, A. G., & Pierce-Shimomura, J. T. (2012). Conserved role of dopamine in the modulation of behavior. *Communicative and Integrative Biology*, 5(5), 440–447. <https://doi.org/10.4161/cib.20978>
- Vladimirov, N., & Sourjik, V. (2009). Chemotaxis: How bacteria use memory. *Biological Chemistry*, 390(11), 1097–1104. <https://doi.org/10.1515/BC.2009.130>
- von Twickel, A., Kowatschew, D., Saltürk, M., Schauer, M., Robertson, B., Korsching, S., Walkowiak, W., Grillner, S., & Pérez-Fernández, J. (2019). Individual Dopaminergic Neurons of Lamprey SNc/VTA Project to Both the Striatum and Optic Tectum but Restrict Co-release of Glutamate to Striatum Only. *Current Biology*, 29(4), 677–685.e6. <https://doi.org/10.1016/j.cub.2019.01.004>

## BIBLIOGRAPHY

---

- Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation through neural population dynamics. *Annu. Rev. Neurosci.*, 43, 249–275.
- Wadhams, G. H., & Armitage, J. P. (2004). Making sense of it all: Bacterial chemotaxis. *Nature Reviews Molecular Cell Biology*, 5(12), 1024–1037. <https://doi.org/10.1038/nrm1524>
- Wang, D., Yu, Y., Li, Y., Wang, Y., & Wang, D. (2014). Dopamine receptors antagonistically regulate behavioral choice between conflicting alternatives in *C. elegans* e115985. *PLoS ONE*, 9(12), 1–18. <https://doi.org/10.1371/journal.pone.0115985>
- Wang, J., Narain, D., Hosseini, E. A., & Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. *Nat. Neurosci.*
- Ward, A., Liu, J., Feng, Z., & Xu, X. Z. (2008). Light-sensitive neurons and channels mediate phototaxis in *C. elegans*. *Nature Neuroscience*, 11(8), 916–922. <https://doi.org/10.1038/nn.2155>
- Ward, S. (1973). Chemotaxis by the nematode *Caenorhabditis elegans*: identification of attractants and analysis of the response by use of mutants. *Proceedings of the National Academy of Sciences of the United States of America*, 70(3), 817–821. <https://doi.org/10.1073/pnas.70.3.817>
- Ward, S., Thomson, N., White, J. G., & Brenner, S. (1975). Electron microscopical reconstruction of the anterior sensory anatomy of the nematode *caenorhabditis elegans*. *Journal of Comparative Neurology*, 160(3), 313–337. <https://doi.org/10.1002/cne.901600305>
- Webre, D. J., Wolanin, P. M., & Stock, J. B. (2003). Bacterial chemotaxis. *Current Biology*, 13(2), 501–517. <https://doi.org/10.1146/annurev.ph.44.030182.002441>
- Wen, J. Y., Kumar, N., Morrison, G., Rambaldini, G., Runciman, S., Rousseau, J., & Van Der Kooy, D. (1997). Mutations that prevent associative learning in *C. elegans*. *Behavioral Neuroscience*, 111(2), 354–368. <https://doi.org/10.1037/0735-7044.111.2.354>
- White, J. G. (1986). The structure of the nervous system of the nematode *Caenorhabditis Elegans*. *Philosophical Transactions of the Royal Society*, 340, 1–340.

- Wickens, J. R., Alexander, M. E., & Miller, R. (1991). Two dynamic modes of striatal function under dopaminergic-cholinergic control: Simulation and analysis of a model. *Synapse*, 8(1), 1–12. <https://doi.org/10.1002/syn.890080102>
- Wicks, S. R., & Rankin, C. H. (1997). Effects of tap withdrawal response habituation on other withdrawal behaviors: the localization of habituation in the nematode *Caenorhabditis elegans*. *American Psychological Association*.
- Wilson, C. J., & Groves, P. M. (1980). Fine Structure and Synaptic Connections of the Common Spiny Neuron of the Rat Neostriatum: A Study Employing Intracellular Injection of Horseradish. *The Journal of Comparative Neurology*, 615(April 1979), 1–17. [papers3://publication/uuid/D25323A7-4661-4530-A95F-81C7FBBF46838](https://pubmed.ncbi.nlm.nih.gov/6154661/)
- Winston, P. (1972). The M.I.T. Robot. In *Machine Intelligence*, 7, ed. B. Meltzer D. Michie.
- Wolff, S. B. E., Ko, R., & Ölveczky, B. P. (2022). Distinct roles for motor cortical and thalamic inputs to striatum during motor skill learning and execution. *Sci Adv*, 8(8), eabk0231.
- Xiang, Y., Graeber, T., Enke, B., & Gershman, S. J. (2021). Confidence and central tendency in perceptual judgment. *Atten. Percept. Psychophys.*, 83(7), 3024–3034.
- Xiong, Q., Znamenskiy, P., & Zador, A. M. (2015). Selective corticostriatal plasticity during acquisition of an auditory discrimination task. *Nature*, 521(7552), 348–351. <https://doi.org/10.1038/nature14225>
- Xu, M., Zhang, S.-Y., Dan, Y., & Poo, M.-M. (2014). Representation of interval timing by temporally scalable firing patterns in rat prefrontal cortex. *Proc. Natl. Acad. Sci. U. S. A.*, 111(1), 480–485.
- Yamamoto, K., & Vernier, P. (2011). The evolution of dopamine systems in chordates. *Frontiers in Neuroanatomy*, 5(MARCH), 1. <https://doi.org/10.3389/fnana.2011.00021>
- Zhao, B., Khare, P., Feldman, L., & Dent, J. A. (2003). Reversal frequency in *Caenorhabditis elegans* represents an integrated response to the state of the animal and its environment. *Journal of Neuroscience*, 23(12), 5319–5328. <https://doi.org/10.1523/jneurosci.23-12-05319.2003>
- Zieger, E., Lacalli, T. C., Pestarino, M., Schubert, M., & Candiani, S. (2017). The origin of dopaminergic systems in chordate brains: Insights

## BIBLIOGRAPHY

---

- from amphioxus. *International Journal of Developmental Biology*, 61(10-12), 749–761. <https://doi.org/10.1387/ijdb.170153sc>
- Znamenskiy, P., & Zador, A. M. (2013). Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. *Nature*, 497(7450), 482–485. <https://doi.org/10.1038/nature12077>



ITqb nova



Oeiras, January, 2024

The basal ganglia learn a temporal basis for the organization of behavior

Filipe S. Rodrigues

