SISSA
Scuola
Internazionale
Superiore di
Studi Avanzati

Neuroscience Area – PhD course in
Cognitive Neuroscience

Tactile perception and memory in rats and humans:

task and species differences within a general framework

Candidate:
Davide Giana

Advisor:
Mathew E. Diamond

Academic Year 2022-2023

TABLE OF CONTENTS

ABSTRACT

Perceptual memories are the storage of our experiences; they are the basis for understanding the external world and guiding our decisions. Despite fast-paced research in the field, behavioural and cognitive constructs tend to be custom built around the investigators' preferred task, and general principles across tasks seem to be missing.

To address these issues, alongside psychophysical experiments, I aimed to build a computational model comprised of interconnected functional units, each performing a specific task-independent operation; the interaction between units and the readout of the system is controlled in a top-down mechanism, depending on goal-oriented requirements. I also aimed to relate this model to neuronal activity in two cortical regions, which are believed to be part of a network involved in perceptual decision making.

Two groups of subjects, rats and humans, performed each of two different tasks requiring the elaboration of vibrotactile stimuli: (I) a categorization task, where a single stimulus must be judged ("strong" or "weak") according to an implicit boundary, and (II) a delayed comparison task, where two stimuli are delivered in each trial, the first of which must be stored in short-term memory to be compared to the second.

The results show that several aspects of trial history, such as recent stimuli and recent choice outcomes, factor into the choice of the current trial. Neuronal recordings from prefrontal and posterior parietal cortices, performed in rats executing both tasks, revealed that cortical activity accounted for decisional variables and past trials' stimuli and choices, correlating with task performance and trial history, as emerged from behavioural results. These finding extend a previous model, thus helping to account for perceptual memory in a task-general manner.

GENERAL INTRODUCTION

The perception of stimuli is far from a passive report of the firing of sensory neurons: it is an active process, that includes prediction as well as faithful representation (Bubic et al., 2010; Summerfield et al., 2006; Weilnhammer et al., 2017). The capacity for predictive coding allows the brain to take advantage of the environment statistical regularities, both to reduce the computational load on perception, and to reduce the noisiness of stimuli and memories.

The past is a reasonable prediction of the future. In sensory environments, statistics may change in the order of seconds or minutes. Therefore, an intuitively sound idea is that, for sensory systems, the recent past provides a context against within which information must be culled from ongoing events. Typical hallmarks of how the environment statistics affect perception are the stimulus history effects: the perception of current stimuli is influenced by the stimuli received in the past. For instance, receiving the same stimulus twice can lead to a different percept, depending on the events that led up to that stimulus instance (Cicchini et al., 2018; Jou et al., 2004; Kiyonaga et al., 2017). These effects can be either attractive, making the current percept more similar to past ones, or repulsive, magnifying the distance between current and past percept. Whether these effects are sensory, perceptual, decisional, or a combination of all is still a hotly debated topic in the scientific community.

The aim of the present study is to investigate general principles of the traces of past stimuli stored in the brain – what we call perceptual memory – and to provide a framework that is, as much as possible, task- and species-independent. To do so, I explored the behaviour of rats and humans in two task designs previously established in the Tactile Perception and Learning Lab. I combined the working memory (WM) task developed in (Fassihi et al., 2014) and the reference memory (RM) task developed in (Hachen et al., 2021) by having rats learn to perform both tasks. The two tasks are operationally similar, are performed in the same apparatus, and require the elaboration of the same vibrotactile stimuli, delivered to the whiskers. However, the tasks involve different cognitive resources: crucially, the delayed comparison task (working memory task, WM) requires the mnemonic retention of the first

stimulus of each trial, until the second is delivered, to allow the comparison of the two stimuli; the categorisation task (reference memory task, RM), instead, requires the immediate categorisation of a stimulus against an implicit boundary, with no explicit retention required.

Having subjects learn and perform both tasks allowed me to investigate whether and how stimulus history effects are modulated by the different cognitive resources required; also, I investigated whether the magnitude of such effects in one task predicted the magnitude in the other, hinting at shared, perhaps hardwired, physiological mechanisms.

To test the species-specific component of the process, I also had human subjects perform the same tasks, and assessed whether their behaviour mirrored that of rodents. Both rats and humans showed remarkable similarity in WM, with stimulus history affecting them with the same dynamics and temporal profile, the main difference being that rats showed greater magnitude of effect. In other words, human participants were better able to “isolate” the current trial. In RM, rats and humans also performed the task in the same way; differently from WM, though, stimulus history was masked by choice history, which acted in the direction opposite to stimulus history. Thus, while determining whether stimulus history correlated between the two tasks proved inconclusive, a case can be made about the interaction of stimulus and choice history as task- and modality-dependent.

In order to examine the neuronal correlates of these effects, I implanted five rats with microelectrode arrays in two cortical areas, prefrontal cortex (PFC) and posterior parietal cortex (PPC), areas that have been implied in literature to be involved in the mediation of attention, stimulus memory and trial history in tasks with similar design to the ones described in this work. I recorded the neuronal activity from PFC and PPC during the execution of WM and RM tasks and examined how the activity related to both current stimuli and trial history, in order to find shared mechanisms between tasks and individuals. Both PFC and PPC appeared to be involved in WM and RM. PFC was more involved than PPC in encoding the tasks’ decision variables, and its activity was predictive of subjects’ performance. At the same time, both PFC and PPC were involved in coding for stimulus and choice history, correlating with the magnitude of the effects during different sessions.

Finally, I further developed the behavioural model first presented in Hachen et al., 2021, originally conceived to produce a computational framework for RM, to account for the WM task as well, in both rats and humans. Extending the model allowed to explain the contraction bias observed in WM within the same framework established previously for RM. Thus, my

extension of the model, was able to reconcile two apparently opposite stimulus history effects through task-independent interactions, in both rats and humans.

Overall, these results show that stimulus history in a delayed comparison task, both in a short and long timescale, act in the same fashion in rats and humans, hinting at a conserved mechanism. Neuronal correlates of decision variables and both stimulus and choice history are distributed across PFC and PPC, providing further evidence that these cortical areas act as part of a single network mediating present and past behaviour-relevant variables. Moreover, the effects described were merged seamlessly in the previous model, thus providing a single, task-independent framework to account for stimulus history in tactile perception modality.

CHAPTER 1

BEHAVIOURAL ANALYSIS OF STIMULUS AND CHOICE HISTORY AND THEIR NEURAL CORRELATES

INTRODUCTION

In rats, tactile perception is a well-studied sensory modality, with very clear anatomical pathways from sensory receptor neurons to primary somatosensory cortex (S1) (Diamond & Toso, 2023). In rats, vibrissal touch is a critical modality for actively exploring the world; as such, a large portion of the cerebral cortex involves sensorimotor processing, and within the primary somatosensory cortex (S1), a large portion devoted to vibrissal inputs (vibrissal somatosensory cortex, vS1) (Diamond & Toso, 2023).

In previous publications from the Tactile Perception and Learning Lab, rats have been shown to be able to perform well in both a delayed comparison design (in this thesis referred to as Working Memory, WM) (Fassihi et al., 2014) and a single stimulus categorisation design (in this thesis referred as Reference Memory, RM) (Hachen et al., 2021), on a par with human subjects. Building upon the expertise in shaping behaviour through operant conditioning, we trained them to perform both RM and WM, in separated sessions, to compare the extent of history effects, in the same subject, across tasks.

History effects include the influence of recent stimuli and choices on current decisions, even when those stimuli and choices no longer carry information about what choice should be made next. Often called serial dependence, this effect has been widely studied in different modalities (Barbosa & Stein, 2020; Fischer & Whitney, 2014; Fritsche et al., 2017; Hachen et al., 2021; Liberman et al., 2014). History has been reported to act either in an attractive or repulsive way, depending on the modality, task, and time-scale involved (Bliss et al., 2017; Fornaciai & Park, 2018; Hachen et al., 2021; Manassi et al., 2018); it still widely debated whether serial dependence acts at the sensory, perceptual, or decisional level, or

a combination of the three (Bosch et al., 2020; Cicchini et al., 2017; Fornaciai & Park, 2018; Fritsche et al., 2017; Kiyonaga et al., 2017). Overall, a cohesive description of serial dependence is still missing, especially one that transcends a single task design.

Together with serial dependence, the phenomenon of contraction bias, first described in (Hollingworth, 1910), is also in the library of known history effects: differently from serial dependence, it is described as an over- or underestimation of stimuli, depending on the magnitude of the stimulus relative to the range employed during the experiment; also, it accounts for how the *distribution* of past stimuli affects current decisions, instead of the single stimuli delivered in the recent past. This effect is always studied in delayed comparison designs (Akrami et al., 2018; Ashourian & Loewenstein, 2011; Fassihi et al., 2017; Tal-Perry & Yuval-Greenberg, 2022), and is often described as a regression of memory traces towards the mean of previously presented stimuli. While there have been many attempts at modelling contraction bias (Ashourian & Loewenstein, 2011; Jou et al., 2004), integrating a Bayesian statistical approach, there is still much to be understood. Moreover, a shared framework encompassing both serial dependence and contraction bias is still missing.

By assessing both qualitatively and quantitatively serial dependence and contraction bias in WM and RM, with each requiring different cognitive resources in processing the same physical features, we aimed at building a coherent description of the effects of stimulus history across two different experimental paradigms.

Besides the phenomenological description, a cohesive neurophysiological underpinning of history effects is also lacking. While many cortical areas and networks have been implied in stimulus memory and history processing, we chose to focus on two reciprocally connected areas in the rat cortex: prefrontal cortex (PFC) and posterior parietal cortex (PPC). PFC is loosely defined in rodents (Barthas & Kwan, 2017), as it is associated with many functions, among which sensory integration, motor planning and execution and attention, which are more segregated in primate cortex; PFC recordings presented in this thesis include populations from secondary motor (M2), anterior cingulate (ACC) and infralimbic (IL) cortices.

PFC has been implicated in adaptation to visual stimuli distribution (Wang et al., 2020), short-term memory retention of vibrotactile stimuli (Brody et al., 2003; Fassihi et al., 2017; Romo et al., 1999), serial dependence and contraction bias (Barbosa & Stein, 2020;

Benozzo et al., 2023; Serrano-Fernández et al., 2023). We aimed at uncovering both a stimulus representation and an activity modulation dependent on stimulus history that would be conserved across WM and RM.

PPC activity during task execution has been correlated before with multimodal sensory integration (Whitlock, 2017), behavioural flexibility, together with PFC, (Sohn et al., 2000), and trial history representation (Akrami et al., 2018). By recording PPC, together with PFC, I aimed at testing the hypothesis that these two cortical areas are part of a shared network responsible for present and past trials representation, and that this representation is predictive of stimulus history observed in behaviour.

The driving motivation behind our search for a single coherent framework arises from notions of efficiency. The sensory inputs involved in many memory tasks may be essentially equivalent. Thus, in the present case, RM and WM tactile stimuli are drawn from the same set. Yet, when the tasks are broken down into a sequence of operations, they may seem quite different. Our intuition was that the nervous system could not contain sufficient processing machinery (neuronal populations) to devote one distinct network to each possible perceptual memory task. If the brain uses some set of networks to construct and store a long-term boundary (RM) and to store the most recent stimulus in a short-term buffer (WM) there might be some dynamics characterizing the set of networks across tasks. It seemed conceivable that what distinguishes one task from another is not so much the networks' dynamics, but which representation is "read out" in order to solve the problem at hand. Our overall goal, then, is to determine whether general dynamic properties concerning the representation of the current and past stimuli, and choices, could be identified as a common property of different behaviours.

RESULTS

Task design

Each subject, rat or human, performed two different tasks, involving the processing of vibrotactile stimuli and the reporting of a decision. The first task, called Working Memory (WM), is schematically represented in Figure 1-A; it involved receiving a sequence of two stimuli in each trial, a base and a comparison, and comparing the two stimuli in order to judge which was stronger. This design allowed us to measure how the memory trace of the base stimulus was affected by stimulus history effects while being retained by the subject. The second task, called Refence Memory (RM), is schematically represented in Figure 1-B; it involved receiving a single stimulus in each trial, which the subject must then categorise as weak or strong, in relation to a fixed, absolute boundary. This design allowed us to assess how the stimulus percept was immediately affected by stimulus history. The memory of interest here is not the base stimulus, but the boundary itself. In WM, the relevant memory is ”local”, i.e.: ideally, it is within the current trial. In RM, the relevant memory is ”non-local”, i.e.: ideally, it is unaffected by the current trial. By having each subject perform sessions for both tasks, we measured whether the magnitude and direction of stimulus history effects in one task predicted this stimulus history effects in the other task, thus testing the hypothesis of a task-independent, and possibly species-independent, cognitive and physiological mechanism.

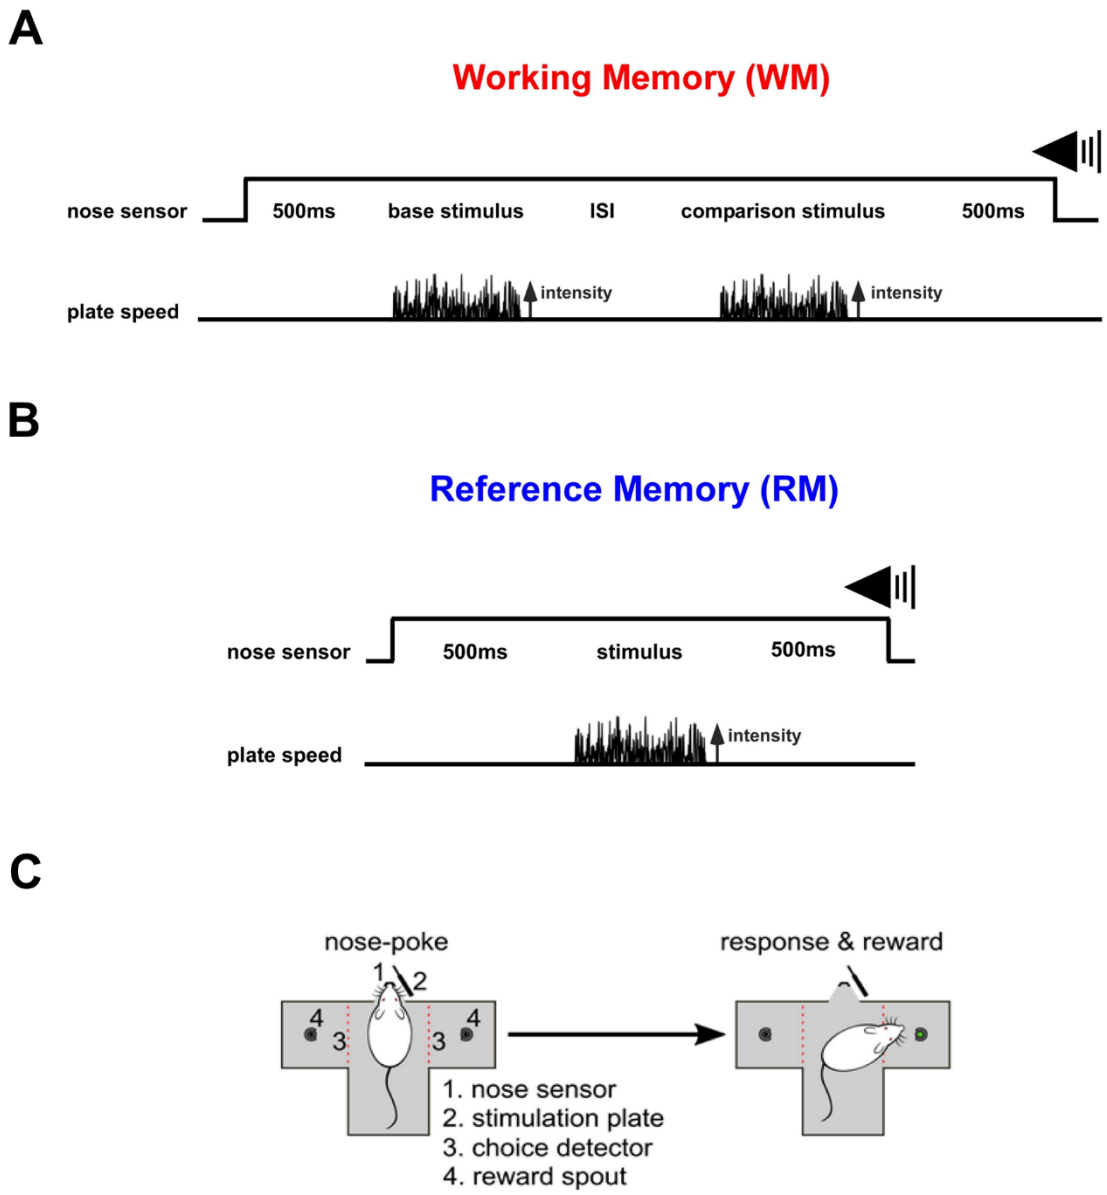

Figure 1: WM and RM designs. (A) Working memory trial design. Each trial begins with the subject contacting a dedicated sensor (nose poking for rats, finger poking for humans). After a pre-stimulus delay of 500 ms, the base stimulus is delivered, lasting 334 ms. Then, after a variable inter-stimulus interval (ISI), the comparison stimulus is delivered, lasting 334 ms. After a post-stimulus delay of 500 ms, an auditory go cue signals the subject that they can break contact and report their choice. After the outcome of that choice is revealed, the subject can initiate a new trial at any time. **(B)** Reference memory trial design. As in A, except there is no ISI and no comparison stimulus. **(C)** Rat setup schematic design.

Contraction bias

First, we examined how subjects performed in WM, and looked for hallmarks of stimulus history effects. One such effect, crucial to our modeling, is the contraction bias, first described in (Hollingworth, 1910). In general, when affected by contraction bias, subjects consistently underestimate large stimuli, and consistently overestimate small stimuli; large

and small, in this context, are relative to the distribution of stimuli presented to the subjects. The posited explanation is that the working memory of the stimulus, even as the subject prepares the response, shifts (or “contracts”) towards the center of the stimulus distribution. As such, contraction bias could be called a global effect, meaning that it becomes apparent even without examining any specific sequence of recent stimuli. A second effect, often called serial dependence, describes how stimuli received by the subject in the recent past affect the perception of, and the choices made upon, subsequent stimuli. This effect can act in two ways: either attractive, making stimulus percepts more similar to the past ones, or repulsive, with opposite effect.

We designed the stimuli pair sets to highlight contraction bias. Rats performed the WM task on the stimulus set depicted in Figure 2-A, in which each circle represents a pair, in base/comparison coordinates, along with the percentage of correct trials for each stimuli pair (the circled numbers). Each pair is characterised by a Normalised Speed Difference (NSD) value, as described in Equation 1, which assumes the Weber’s fraction as the correlate of difficulty in comparing each stimuli pair. The larger a pair’s NSD, the easier the comparison between the two stimuli, and vice versa. The pairs in the upper half of the plot, above the dashed line, share a positive NSD, as all of them are characterised by a comparison stimulus stronger than base. Conversely, pairs in the lower half plot, below the dashed line, share a negative NSD, and are characterised by base stimuli stronger than comparison. Pairs closer to the dashed line have a smaller NSD value, irrespective of its sign, and thus are harder to discriminate, as the relative difference between their comparison and base stimuli is smaller. Of the pairs composing the set, some share the same nominal difference (NSD=±0.3), and are, in theory, equally discriminable. Those are the pairs arranged in two diagonals of five pairs each, one above and one below the dashed diagonal. Supplementary Figure 1-A shows the performance for these 10 pairs, divided according to the average stimulus speed in each pair, irrespective of order. While performance was well above chance for all speeds, there was a small positive trend between speed and performance. However, this trend is too weak to account for the results shown in the behavioural section of this chapter.

The 10 diagonal pairs are arranged to ensure that the subjects attend to both stimuli, as neither contains enough information by itself to solve the task. In addition, three different sets of 7 psychometric pairs each, arranged vertically and indicated by the coloured arrows (red, orange, and yellow). The vertically arranged sets of stimulus pairs are characterized

by a different, fixed base speed in each set, while the comparison speed changes from pair to pair. In this way, each of the three sets allowed us to compute a psychometric curve over the NSD range.

The corresponding psychometric curves, computed according to equation 2 (Methods), are depicted in Figure 2-B, showing the probability of judging the comparison stimulus as stronger than base against the NSD value (N=10 rats, pooled). In the red set (base=58 mm/s, red arrow in Figure 2-A), the probability of judging the comparison as stronger was higher for all pairs, compared to the other two sets. Conversely, for the yellow set (base=108 mm/s, yellow arrow in Figure 2-A), the comparison-stronger-than-base probability was lower for all pairs. The orange set (base=80 mm/s, orange arrow in Figure 2-A) assumed an intermediate position. Figure 2-C shows that, while sensory acuity (given by curve slope, $1/\sigma$, where σ is the standard deviation of the normal distribution underlying the psychometric function) was not significantly different across the set, the difference between the curves was captured by the psychometric function inflection point (μ), which represents the NSD value at which acuity is maximised, and by lapses (γ and λ, Figure 2-D), which are classically interpreted to represent the rate of non-sensory errors committed when the NSD is very low or very high, respectively. The curves' shift described can be explained in the framework of contraction bias by considering the base stimulus intensity, relative to the centre of the stimuli distribution (Supplementary Figure 2-A). When the base was weaker than the average stimulus (red), rats were more likely to judge the base as stronger than the comparison; when the base was stronger than the average (yellow), rats were more likely to judge the base as weaker than the comparison. As such, the base stimulus memory was attracted, or contracted, towards the centre of the stimulus distribution; subjects then remembered the base stimulus as stronger (in the red set) or weaker (in the yellow set) than its real-time percept, biasing choices in opposite directions. The base stimulus in the orange set was close to the average, and produced the least biased psychometric curve; was remembered more faithfully, resulting in an intermediate effect.

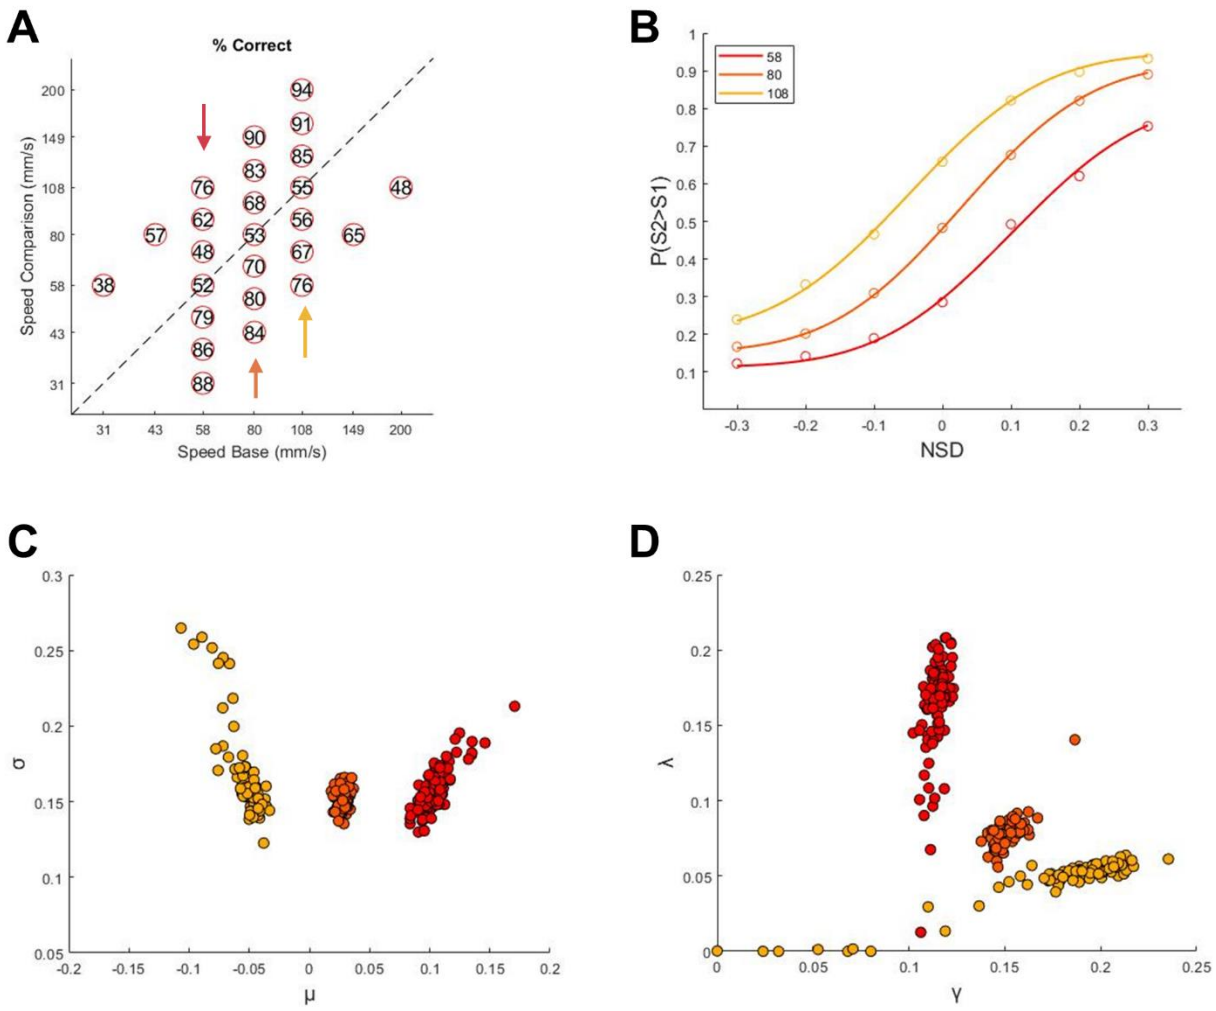

Figure 2: Working Memory, psychometric curves, rats. (A) Stimuli pairs set employed for rats, where each circle represents a base-comparison pair. Axes are stimuli nominal speeds in mm/s, in log scale. Base stimulus on x axis, comparison stimulus on y axis. The numbers within circles are the percentage of correct choices for each stimulus pair. Coloured arrows (red, orange, and yellow) indicate the three vertical psychometric sets **(B)** psychometric curves for the three sets indicated in A, with NSD on x axis and probability of reporting comparison stimulus as stronger than base on y axis. **(C)** Scatter plot of inflection point (μ) against reciprocal of acuity (σ). Each dot represents the parameters recovered from one bootstrapped iteration, separated according to the psychometric set as in B. **(D)** Scatter plot of lower lapse (γ) against upper lapse (λ). Each dot represents the parameters recovered from one bootstrapped iteration, separated according to the psychometric set as in B. N=10 rats.

For human subjects, there is no need to design the stimuli pairs set to force them to solve the task as designed, as they can be directly instructed to do so (i.e.: “hold the memory of the first stimulus on each trial”); therefore, we designed the stimulus set for the human version of WM to test contraction bias differently from that of rats. For human subjects, the stimulus set, in Figure 3-A, was comprised of two psychometric sets: a vertical one, similar to rats, in which the base stimulus had a fixed value for all pairs (80 mm/s), and a horizontal one, in which the comparison stimulus had a fixed value (also 80 mm/s). Figure 3-B shows the psychometric curves obtained from the vertical (dark green) and horizontal (light green)

sets (N=16). Interestingly, human subject exhibited results qualitatively similar to rats. In Figure 3-C and D, the light green psychometric, pertaining to the horizontal set, was characterized by the lowest acuity (largest σ) and greatest lapses (γ and λ); symmetrically, the dark green psychometric, pertaining to the vertical set, showed higher acuity and almost zero lapses.

Here, contraction bias affected performance in opposite ways in the two sets. In the horizontal set, performance was always hindered by the bias, as the memory trace of the base stimulus was always contracted towards the centre of the distribution (Supplementary figure 2-B), that is, towards the fixed comparison intensity. In this way, the contraction bias reduced the difference between base stimulus memory trace and comparison stimulus, impairing performance and leading to lower acuity. At the same time, the increased lapses are harder to describe: usually, lapses are interpreted as non-sensory errors, for example due to inattention or motor execution errors. However, the lapses shown here (Figure 2-D and 3-D) are due to a systematic contraction of memory in a specific direction, leading to errors that are not explainable via psychometric curve fitting alone. In the second chapter of this thesis, this point will be raised again, showing that a simulated lapse-less subject will produce finite lapses in a psychometric fitting, when under the effect of memory contraction.

In the vertical set, the memory trace was contracted towards the centre of the distribution (Supplementary figure 2-B), which was very close to the base stimulus of the vertical set: thus, contraction reduced the noisiness in the base stimulus memory trace, stabilising memory and improving performance.

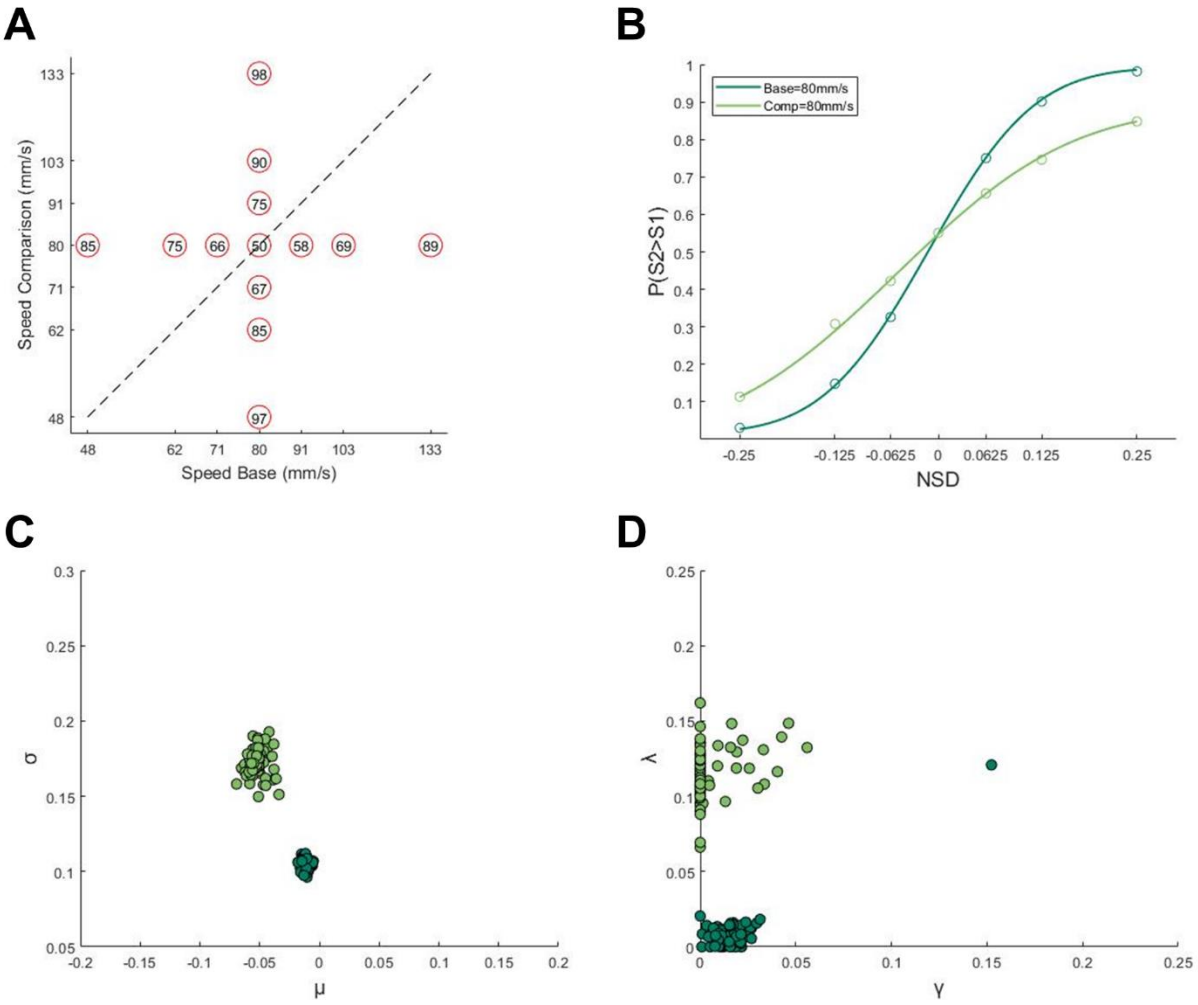

Figure 3: Working Memory in humans. (A) Stimuli pairs set employed. Each circle represents a base-comparison pair. Axes are stimuli nominal speeds in mm/s, in log scale; base stimulus on x axis, comparison stimulus on y axis. The numbers within circles are the percentage of correct choices for each stimulus pair. Coloured arrows (light green, dark green) indicate the horizontal and vertical psychometric sets, respectively **(B)** psychometric curves for the two sets indicated in A, with NSD on x axis and probability of reporting comparison stimulus as stronger than base on y axis. **(C)** Scatter plot of inflection point (μ) against reciprocal of sensitivity (σ). Each dot represents the parameters recovered from one bootstrapped iteration, separated according to the psychometric set as in B. **(D)** Scatter plot of lower lapse (γ) against upper lapse (λ). Each dot represents the parameters recovered from one bootstrapped iteration, separated according to the psychometric set as in B. N=16 humans.

Stimulus serial dependence in WM

We have shown that contraction bias is a global stimulus history effect, emerging from all trials taken together, in both rats and human subjects. We then searched for *local* effects, which would depend on the stimuli delivered in the recent past. We investigated whether and how the choice in the current trial (n) was modulated by the intensity of stimuli delivered in the preceding trial (n-1), even though those are non-informative for the current trial. First we considered the effect of the average speed (speeds of base plus comparison divided by

two) of trial n-1 (Figure 4-A); the effect of the single stimuli within trial n-1, taken separately, is more complex to interpret and will be discussed later. In rats, the slope of the linear regression is consistently negative across subjects (k = -7.9E-4, *p* = 6.9E-4). This means that, as the stimulus intensity in trial n-1 increases, the probability of judging comparison as stronger than base in trial n decreases. In Figure 5-A, human subjects showed the same trend (k=-2.7E-3, p-value=4.2E-6). Similarly, to the description of contraction bias above, these results can be interpreted as follows: the base stimulus memory of trial n is attracted to the average intensity delivered in trial n-1. Thus, if the average intensity in trial n-1 was weak, the base stimulus memory trace of trial n would become weaker as the subject awaits the comparison stimulus, and the comparison stimulus would be more likely to be judged as stronger than the base (vice versa when the average intensity in trial n-1 was strong).

Next, we examined trials further in the past. In rats (Figure 4-B), trial n-2 showed the same direction in stimulus history effect, although to a lesser degree (k = -3.5E-4, *p* = 1E-3). The same was true for humans (Figure 5-B) (k = -1.5E-3, *p* = 1.1E-3). Combining the results of n-1 and n-2, we surmised that extending the analysis to trials further in the past would show a decaying effect of past trials' stimuli. As a test of this prediction, Figure 4-C shows the linear regression coefficients for slope obtained from rats from trial n-1 to trial n-9, while Figure 5-C shows the same coefficients obtained from humans. Both groups show overall negative coefficients for all trials considered, further confirming the coherency of the effect across subjects and species. As expected, the magnitude of the effect decayed as we considered trials further in the past, until becoming insignificant. This proves a recency component in the attraction of base stimulus to stimuli presented in the past. Rats and humans, however, differed in the extent to which past stimuli affected the current trial: while rats showed significant coefficients up to trial n-7, humans retained a measurable effect only up to trial n-3. This result allows us to refine the interpretation of what the base stimulus is attracted to: in each current trial, n, the base stimulus memory trace is attracted to the intensity of stimuli presented in the most recent trial n-1, which carries the largest portion of the effect, and is also attracted, to a lesser and lesser extent, to the intensity of stimuli presented in trials n-2, n-3 etc. Because human subjects usually completed more trials per unit of time than did rats, the more restricted history effect in humans cannot be explained by positing some constant time interval (rather than some number of trials) as the determinant of past-trial effects.

Taking into account the global (contraction bias) and local (serial dependence) effects, we hypothesized that the base stimulus memory trace is attracted to a weighted, dynamically updated average of past stimuli. This dynamic intensity attractor would, on average, remain close to the centre of the overall stimulus distribution (Supplementary figure 2-A and B) and would explain the contraction bias which emerges when pooling all trials together (Figures 2 and 3). Also, how far in the past the attraction extends might vary across species and even across individuals, as shown by serial dependence results (Figures 4 and 5). In the second chapter of the thesis, we will develop these hypotheses to explain stimulus history effects within one conceptual framework, building upon recent literature (Hachen et al., 2021).

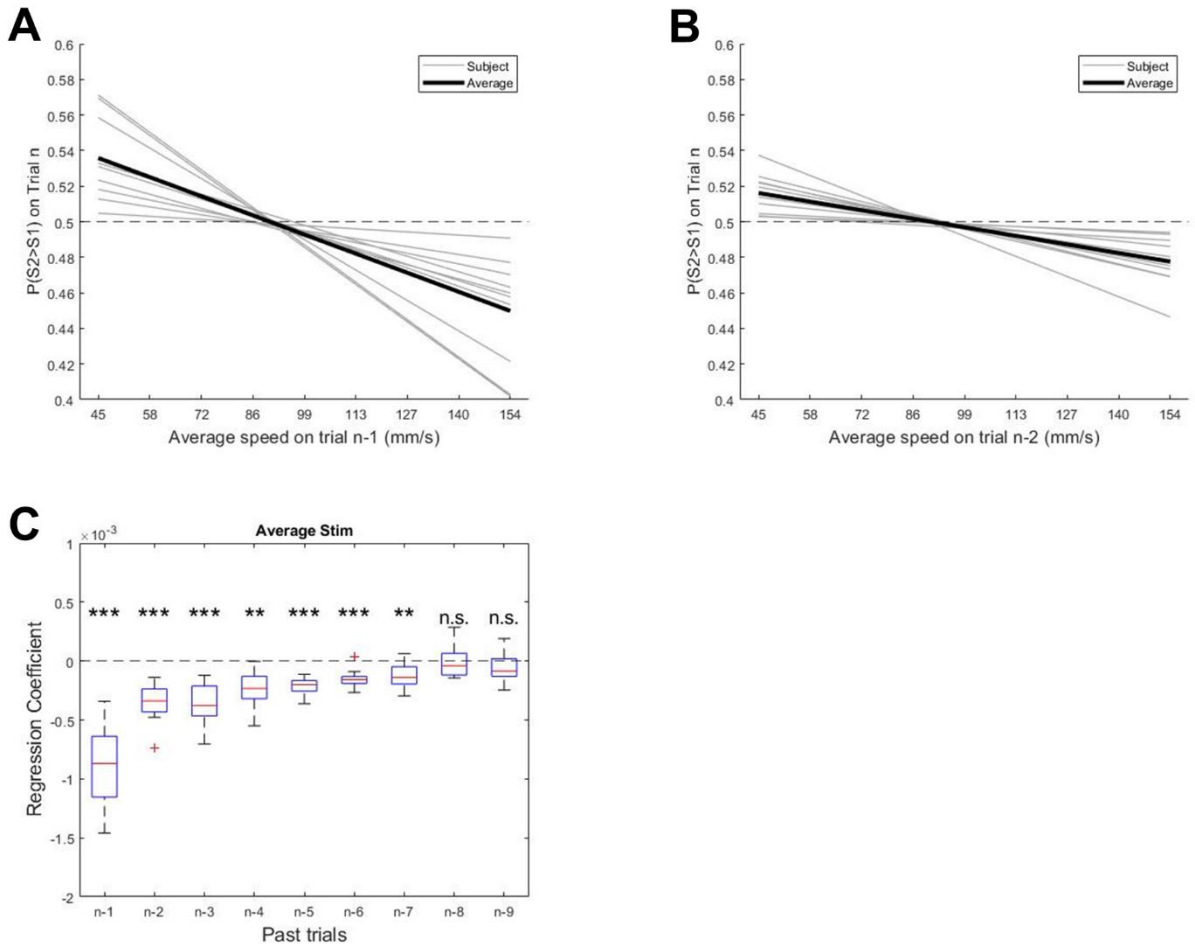

Figure 4: Working Memory history effects in rats. (A) Effect of stimuli in trial n-1 on trial n decision. Average nominal speed of base and comparison stimuli in trial n-1, in mm/s, on x axis, against probability of reporting comparison stimulus as stronger than base during trial n, on y axis. Grey lines are linear regression curves for each subject, black line is linear regression for pooled subjects. **(B)** as in A, for trial n-2, in mm/s, on x axis. **(C)** Linear regression slope coefficients. Past trials on x axis, regression slope coefficient, as in A, for past trials on y axis (Student's t-test). N=10 rats.

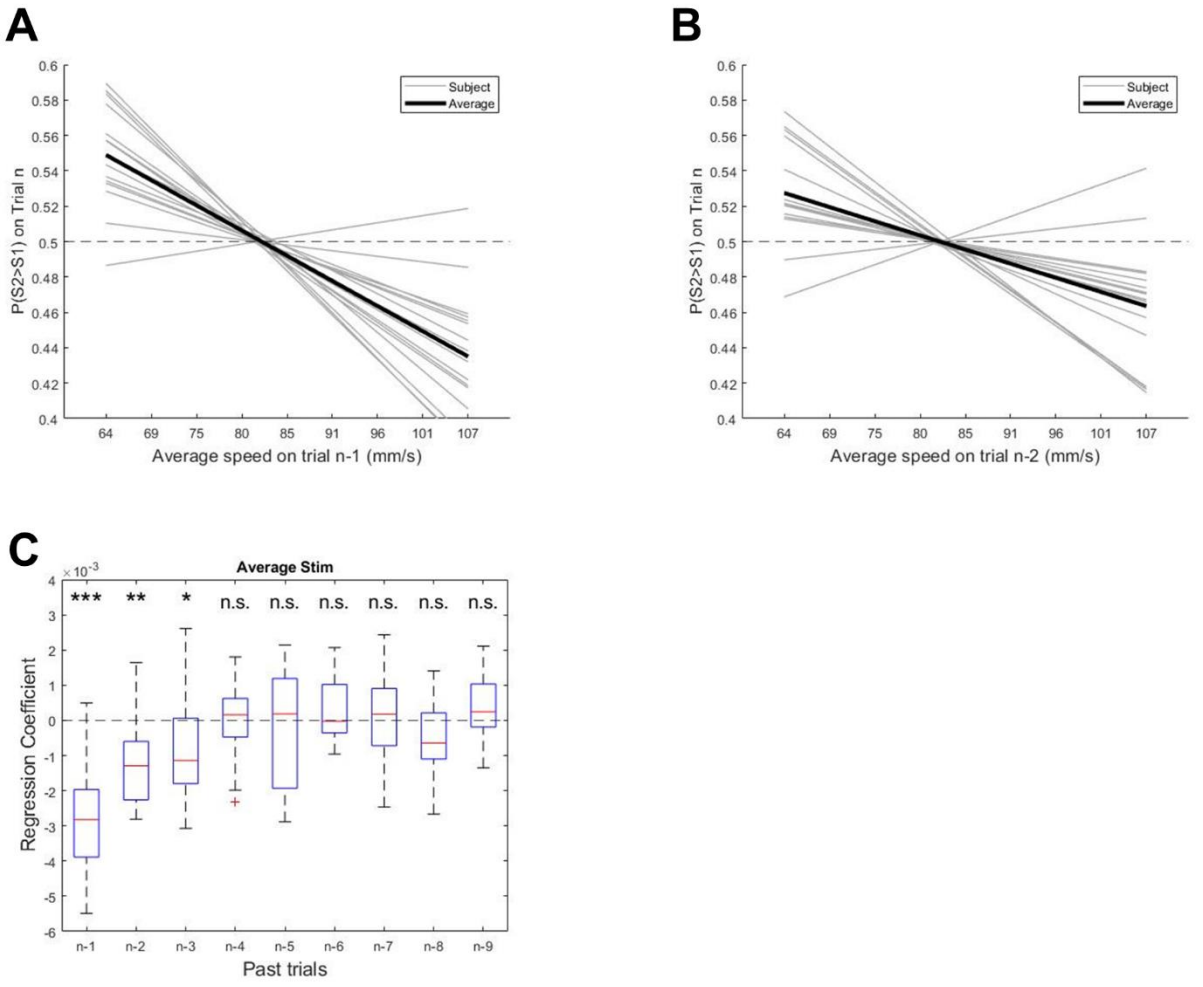

Figure 5: Working Memory history effects in humans. (A) Effect of stimuli in trial n-1 on trial n decision. Average nominal speed of base and comparison stimuli in trial n-1, in mm/s, on x axis, against probability of reporting comparison stimulus as stronger than base during trial n, on y axis. Grey lines are linear regression curves for each subject, black line is linear regression for pooled subjects. **(B)** as in A, for trial n-2, in mm/s, on x axis. **(C)** Linear regression slope coefficients. Past trials on x axis, regression slope coefficient, as in A, for past trials on y axis (Student's t-test). N=16 humans.

Choice serial dependence in WM

Until now, we only considered stimulus history. However, subjects can be affected by choice history as well, which would constitute a non-sensory bias associated with either repeating or switching choice from one trial to the next. As this effect can further interact with the outcome (rewarded or unrewarded) of previous choices, it can interact with stimulus history in a non-trivial manner, as will become apparent when discussing the RM results in the next section. To account simultaneously for stimulus and choice history, we extended the results shown in Figure 4-C and Figure 5-C, employing a Generalised Linear Model (GLM). Figure 6-A and Figure 7-A, left, show the β recovered, for rats and humans respectively, for trials from n backwards. To determine how far back the model needed to go, we measured the

progressive prediction improvement, adding one trial at a time, via Akaike Information Criterion (AIC), in Figure 6-B and Figure 7-B, for rats and humans respectively. The predictor employed for trial n was the NSD of the pair presented in the current trial; the predictors employed for past trials were the average intensity of stimuli presented in that trial. The first predictor, extracted from trial n, had a positive weight, and the largest magnitude. This means that the current-trial NSD is the most important predictor, a result that confirms that subjects were carrying out the task as expected. The coefficient for past stimuli confirmed what was shown and discussed above: past trials having significant weights, up to trial n-5 for rats and n-2 for humans. The negative weights mean that the stronger the average intensity was in a given past trial, the less likely the subject was to judge the comparison as stronger than the base in the current trial – trial n base stimulus memory was attracted towards trial n-*p*, where *p* denotes the past trial. Humans had a larger β for trial n, which is in line with the fact that they showed less contraction bias. Figure 6-A, right, shows the β pertaining to past choices, separated according to the outcome of that choice, on trial n. In rats, only the most recent trials generated a significant effect, and the effect direction was reversed depending on the outcome (rewarded, unrewarded). Correct choice in trial n-1 had positive weight, meaning that it exerted an attractive effect: rats were more likely to repeat a choice that led to a positive outcome. Conversely, an incorrect choice in trial n-1 had a larger negative weight, meaning that it exerted a repulsive effect: rats were more likely to make a choice opposite to the one that led to an unrewarded outcome. Humans (Figure 7-A, right), on the other hand, show a repulsive effect towards choice n-1 irrespectively of the outcome of that choice.

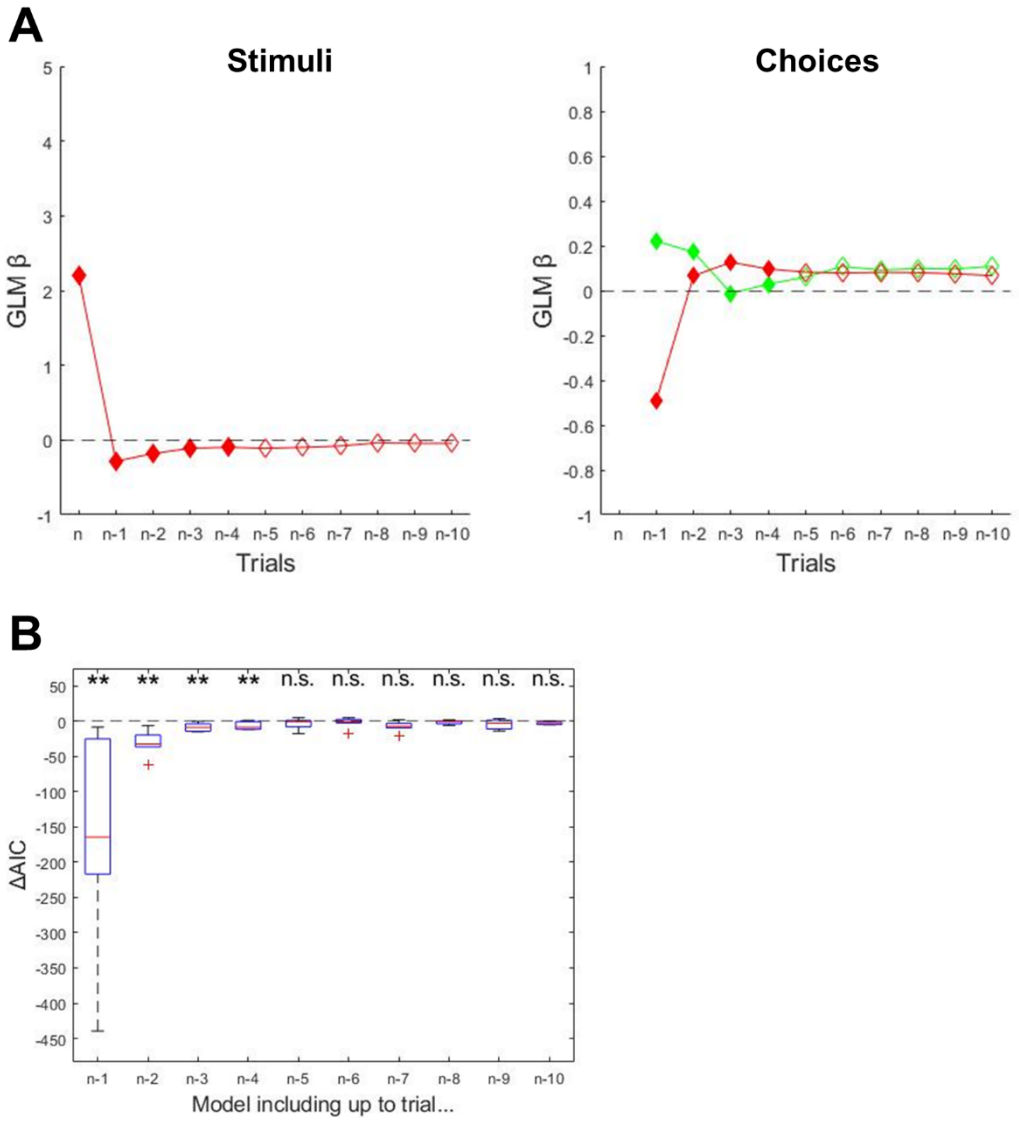

Figure 6: Working Memory Generalised Linear Model for rats. (A) Standardised weights for stimuli (left) and choices (right), rewarded choices in green, unrewarded in red. Trials on x axis. Filled diamonds denote weights for trials deemed significant in B. **(B)** Akaike's Information Criterion difference between consecutive model iterations. X axis denotes up to which trial is included in that model iteration; Y axis represents the ΔAIC between the preceding iteration and the one denoted on x axis. For n-1, the preceding iteration included only trial n (Student's t-test). N=10 rats.

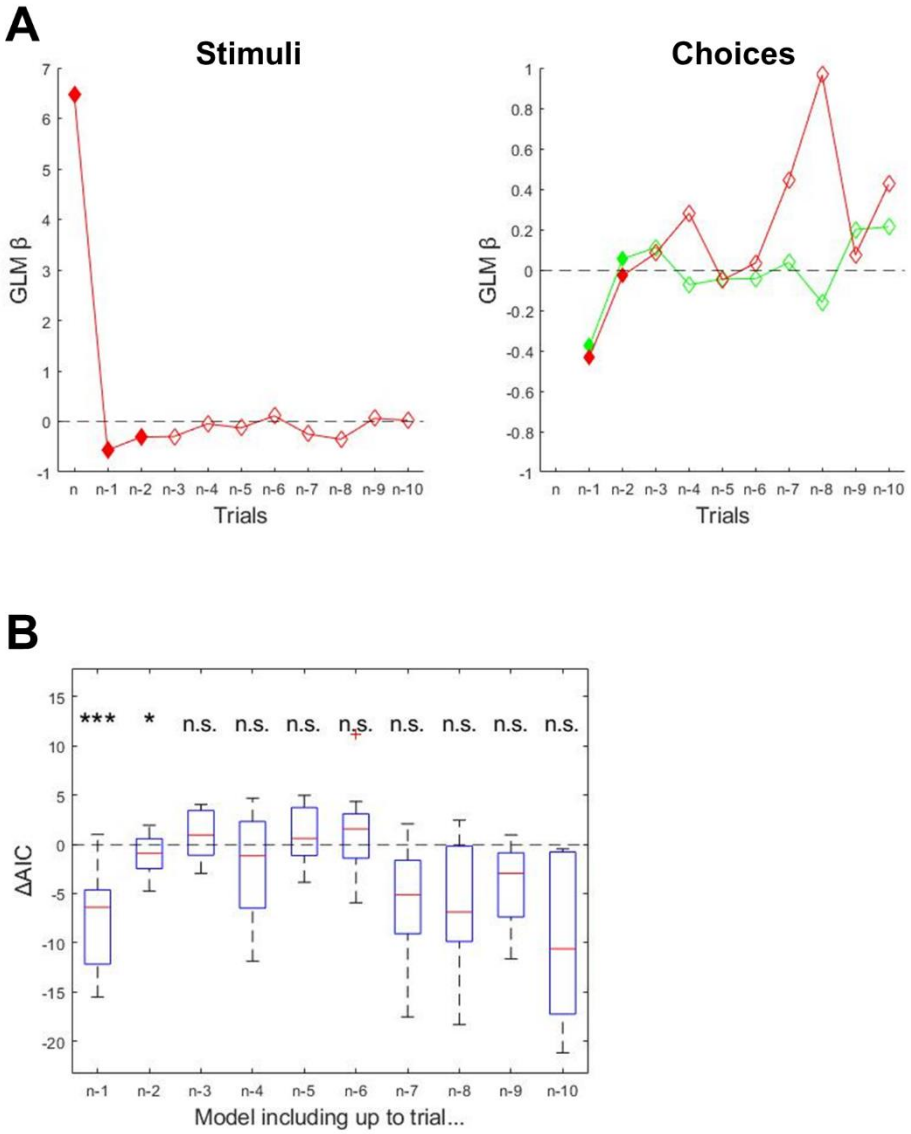

Figure 7: Working Memory Generalised Linear Model for humans. (A) Standardised weights for stimuli (left) and choices (right), rewarded choices in green, unrewarded in red. Trials on x axis. Filled diamonds denote weights for trials deemed significant in B. **(B)** Akaike's Information Criterion difference between consecutive model iterations. X axis denotes up to which trial is included in that model iteration; Y axis represents the ΔAIC between the preceding iteration and the one denoted on x axis. For n-1, the preceding iteration included only trial n (Student's t-test). N=16 humans.

Base and comparison prevalence in serial dependence

So far, we simplified the stimulus history effect by characterising each trial by the *average* intensity of the two stimuli presented in that trial. The n-1 stimulus effect can be further dissected and analysed as potentially having distinct base and comparison effects. Figure 8-A shows that, in rats (red), both stimuli shared negative coefficients, thus exerting the same sign of effect – attraction – on the base stimulus of trial n, as discussed above. However, the base stimulus effect was both larger and more consistent across subjects.

This is in contrast with the recency exhibited by the effect of past trials: the base stimulus in trial n-1 is delivered before the comparison in the same trial, so the comparison is more recent and might be expected to show the larger coefficient. In fact, this was the case in humans (green): the comparison in trial n-1 carried the larger and more consistent effect. This hints at a difference in how the two species process the stimulus history. While humans exhibit recency, wherein the comparison in trial n-1 is closer in time to trial n and exerts a stronger attraction, rats are more affected by the base stimulus – that is, the stimulus required to be held in memory. This difference might be due to the cortical networks involved. We speculate that the neuronal populations responsible in the rat brain for retaining base stimulus memory, and those responsible for causing the attraction of such memory to past stimuli, are partially or largely overlapping, possibly employing a form of multiplex coding. When the base stimulus is loaded into memory, its trace becomes the representation to which the next base stimulus is attracted. In humans, these populations might be more separated. Possibly, the encoding of stimulus history and working memory traces are anatomically segregated in humans; thus, stimulus history would not be disproportionately affected by base stimulus n-1 and, instead, the most recent stimulus (i.e.: comparison n-1) would retain prevalence in the stimulus history. If this could be proven, it would explain why rats are more affected by base stimulus than comparison in trial n-1, and vice versa for humans.

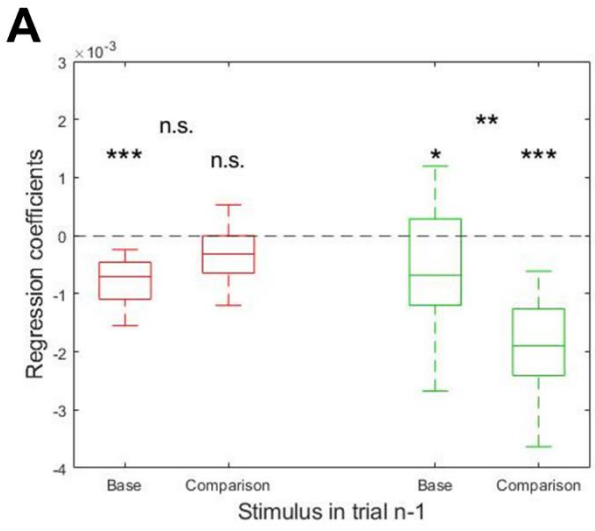

Figure 8: Regression slope coefficients for base and comparison in WM. (A) Regression slope coefficients for probability of reporting comparison stimulus as stronger than base during trial n, computed on either base or comparison in trial n-1. Rats in red, humans in green. N=10 rats, N=16 humans (Student's t-test)

Stimulus serial dependence in RM

Having established a framework for history effects in WM task, we then analysed whether history effects are comparable in the RM task. While categorisation tasks are not usually analysed in a way that might reveal contraction bias, as there is no explicit comparison and no memorisation required, stimuli must nonetheless be compared to a threshold – a boundary between categories. In such a design, the boundary is implicit and must be discovered by subjects during sessions. The human subject, in our experiments, is never given instructions like "You are about to receive the stimulus whose intensity constitutes the category boundary." Humans, like rats, must form a boundary by generating and maintaining a "phantom" perceptual memory which, by trial and error, works as the boundary for comparison on each trial. In Hachen et al., 2021, my colleagues studied in depth how past stimuli affect choices in the RM task. To summarise: past trials' stimulus intensity values exerted a repulsive effect on current choices, meaning that subjects were more likely to categorise trial n stimulus as weak after strong (n-1, and earlier) stimuli, and vice versa. This effect decayed for trials further in the past, exhibiting the same recency we showed above in the WM task. This repulsive effect was also modelled by positing that the subject's internal boundary, or criterion, is attracted towards the most recent stimulus. Since the trial n-1 stimulus exerted a greater effect when the intertrial time from n-1 to n was longer, the investigators included that the attraction of the criterion towards stimulus n-1 followed a decay function defined by a time constant, τ. We will describe this model more in depth, and further develop it, in the second chapter. In this chapter, I will describe how rat and human subjects performed the RM task and how stimulus history affected them.

The stimulus dynamic ranges employed for the two groups (rats, humans) were the same as those employed in their respective WM designs, to avoid task differences due to stimulus representation nonlinearity. Thus, the human stimulus range was narrower than the rats' range, as the dynamical range in WM was different between the two species. In both groups, the category boundary between strong and weak lay in the middle, in a linear scale, of the respective ranges. Figures 9-A and 10-A show psychometric curves for rats and humans, respectively. Human subjects performed better than rats as measured both by lower lapse rates (as was the case in WM), and by better acuity. Based on these characteristics, we expected rats to be more susceptible to stimulus history effects than humans, as a stronger serial dependence would lead to more misjudgements in the task design employed for RM. To our surprise, neither group showed a significant stimulus n-1 effect, contrary to what was

seen in WM. Figures 9-B and 10-B show nine psychometric curves, again for rats and humans, each computed according to the trial n-1 stimulus: neither rats nor humans showed a systematic separation due to stimulus n-1, as visible from the μ and σ distribution in Figures 9-C and 10-C. This result deviates from what has been described in Hachen et al., 2021. Figure 11-A shows stimulus history curves for rats, and Figure 12-A for humans, as linear regression curves, comparable to those given for WM in Figures 4-A and 5-A. In RM, however, there is only one stimulus per trial, and no within-trial stimulus averaging was needed. A negative slope in the history curve means that, after a weak stimulus in trial n-1, the subjects were more likely to judge stimulus n as strong, and vice versa. In the case of a negative slope, stimulus n-1 exerts a repulsive effect on the judgement of trial n; in the case of a positive slope, stimulus n-1 exerts an attractive effect on the judgement of trial n.

Confirming the psychometric result shown in Figures 9-C and 10-C, and differently from WM, the effect was inconsistent across subjects: while all subjects showed the same direction of effect (i.e.: slope sign) in WM curves, here subjects exhibited either repulsive or attractive stimulus n-1 effect, leading to no consistent effect, in either species. Only when considering trials beyond n-1, a significant negative coefficient could be recovered, in Figures 11-C and 12-C, for rats and humans, respectively. We surmise that this could be due to a choice attraction effect, which could partially or completely obscure the past stimuli repulsion.

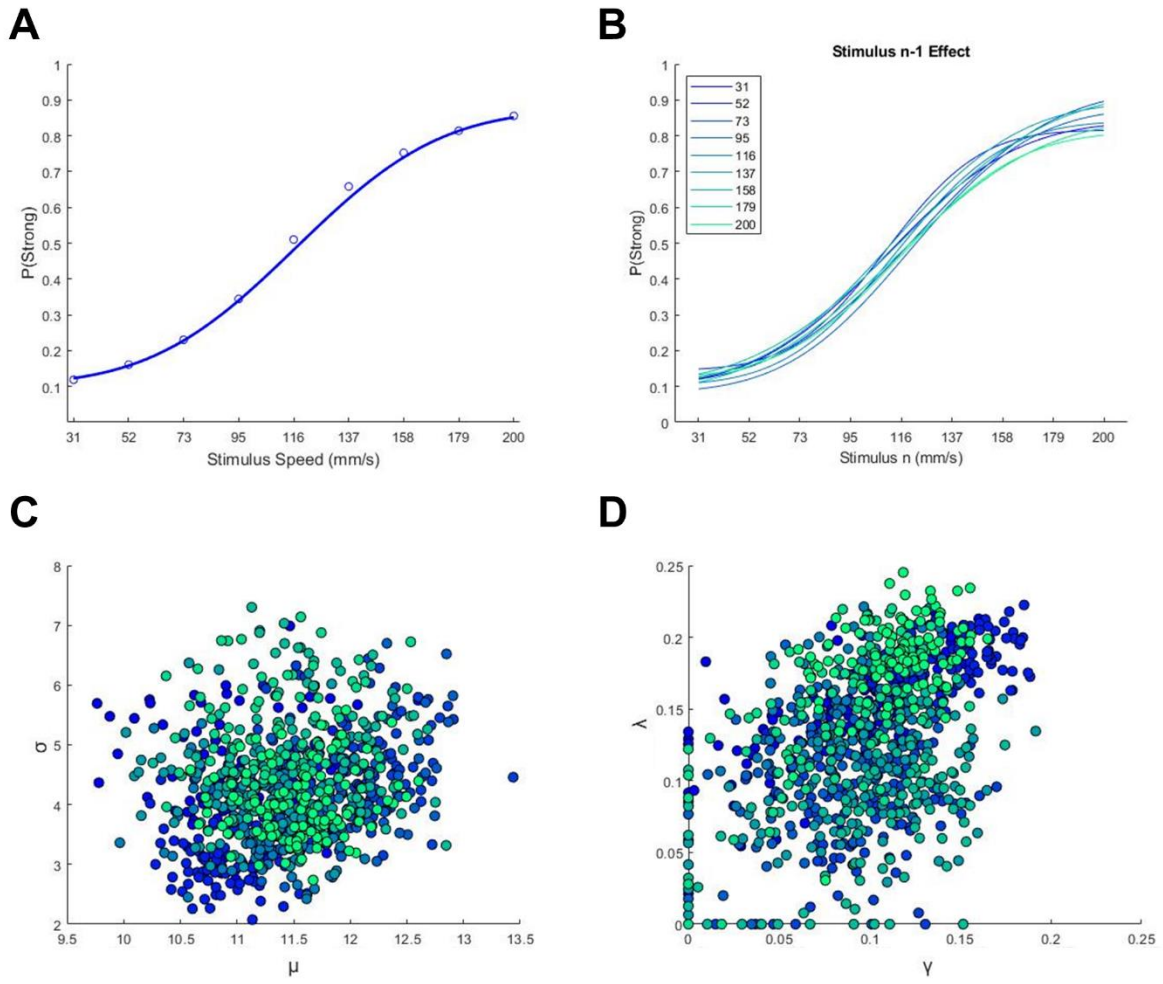

Figure 9: Reference Memory psychometric curves in rats. (A) psychometric curve for the stimulus set, with nominal stimulus speed on x axis and probability of reporting the stimulus as strong on y axis. **(B)** psychometric curves, as in A, conditioned on the stimulus presented in trial n-1. **(C)** Scatter plot of inflection point (μ) against reciprocal of sensitivity (σ). Each dot represents the parameters recovered from one bootstrapped iteration, separated according to stimulus n-1, as in B. **(D)** Scatter plot of lower lapse (γ) against upper lapse (λ). Each dot represents the parameters recovered from one bootstrapped iteration, separated according to stimulus n-1, as in B. N=10 rats.

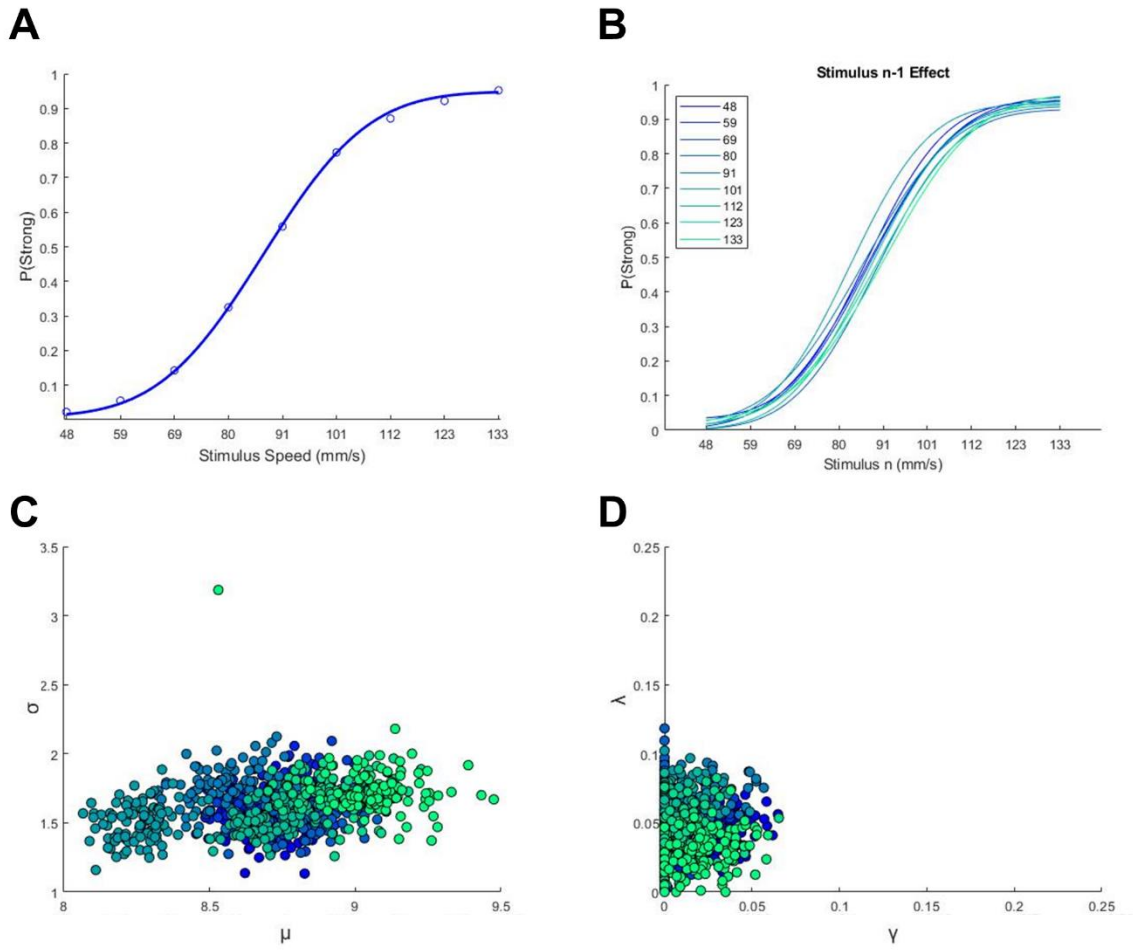

Figure 10: Reference Memory psychometric curves in humans. (A) psychometric curve for the stimulus set, with nominal stimulus speed on x axis and probability of reporting the stimulus as strong on y axis. **(B)** psychometric curves, as in A, conditioned on the stimulus presented in trial n-1 **(C)** Scatter plot of inflection point (μ) against reciprocal of sensitivity (σ). Each dot represents the parameters recovered from one bootstrapped iteration, separated according to stimulus n-1, as in B. **(D)** Scatter plot of lower lapse (γ) against upper lapse (λ). Each dot represents the parameters recovered from one bootstrapped iteration, separated according to stimulus n-1, as in B. N=16 humans.

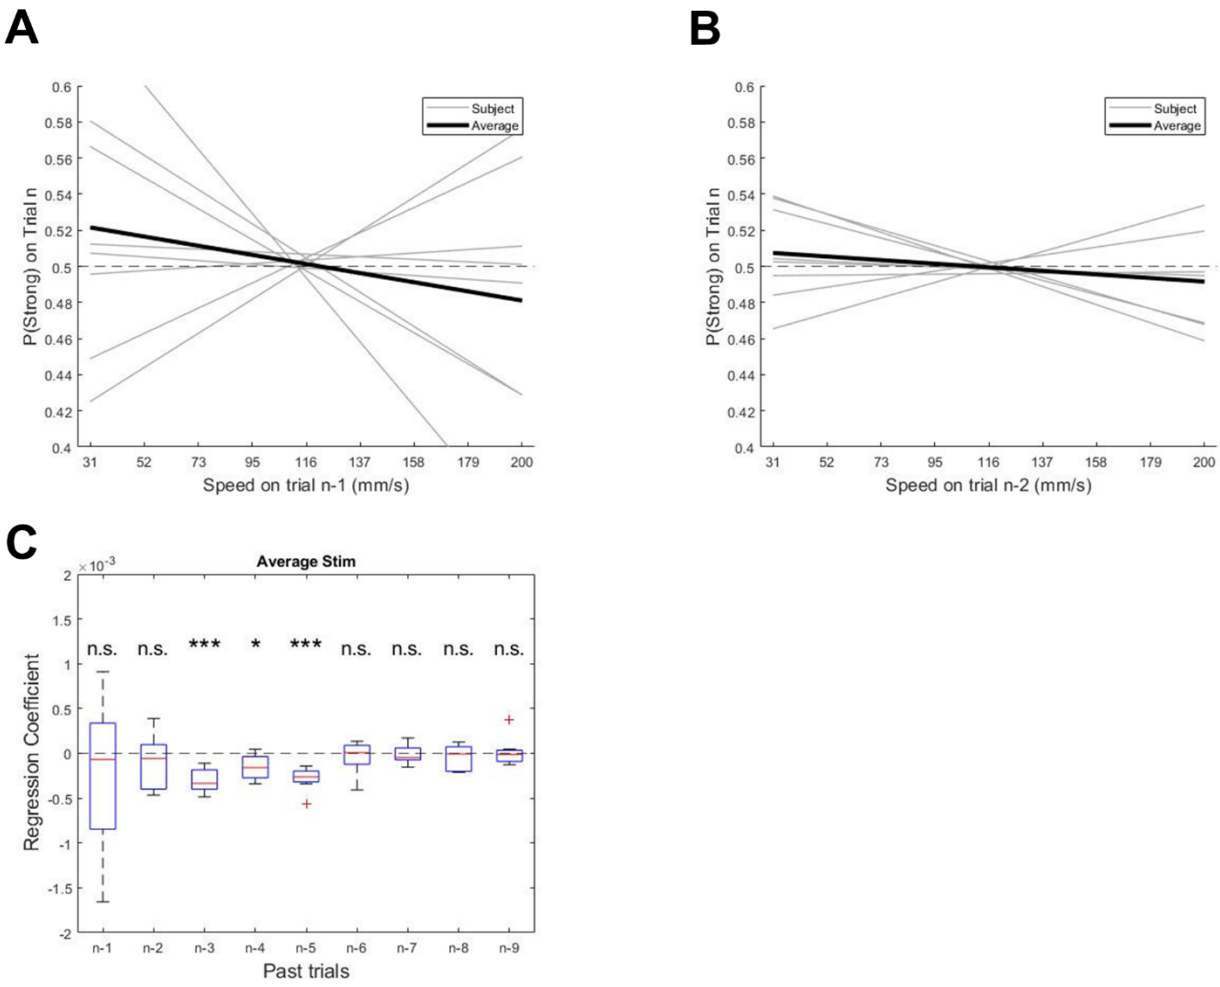

Figure 11: Reference Memory history effects in rats. (A) Effect of stimuli in trial n-1 on trial n decision. Nominal speed of stimulus in trial n-1, in mm/s, on x axis, against probability of reporting the stimulus as strong during trial n, on y axis. Grey lines are linear regression curves for each subject, black line is linear regression for pooled subjects. **(B)** as in A, for trial n-2, in mm/s, on x axis. **(C)** Linear regression slope coefficients. Past trials on x axis, regression slope coefficient, as in A, for past trials on y axis (Student's t-test). N=10 rats.

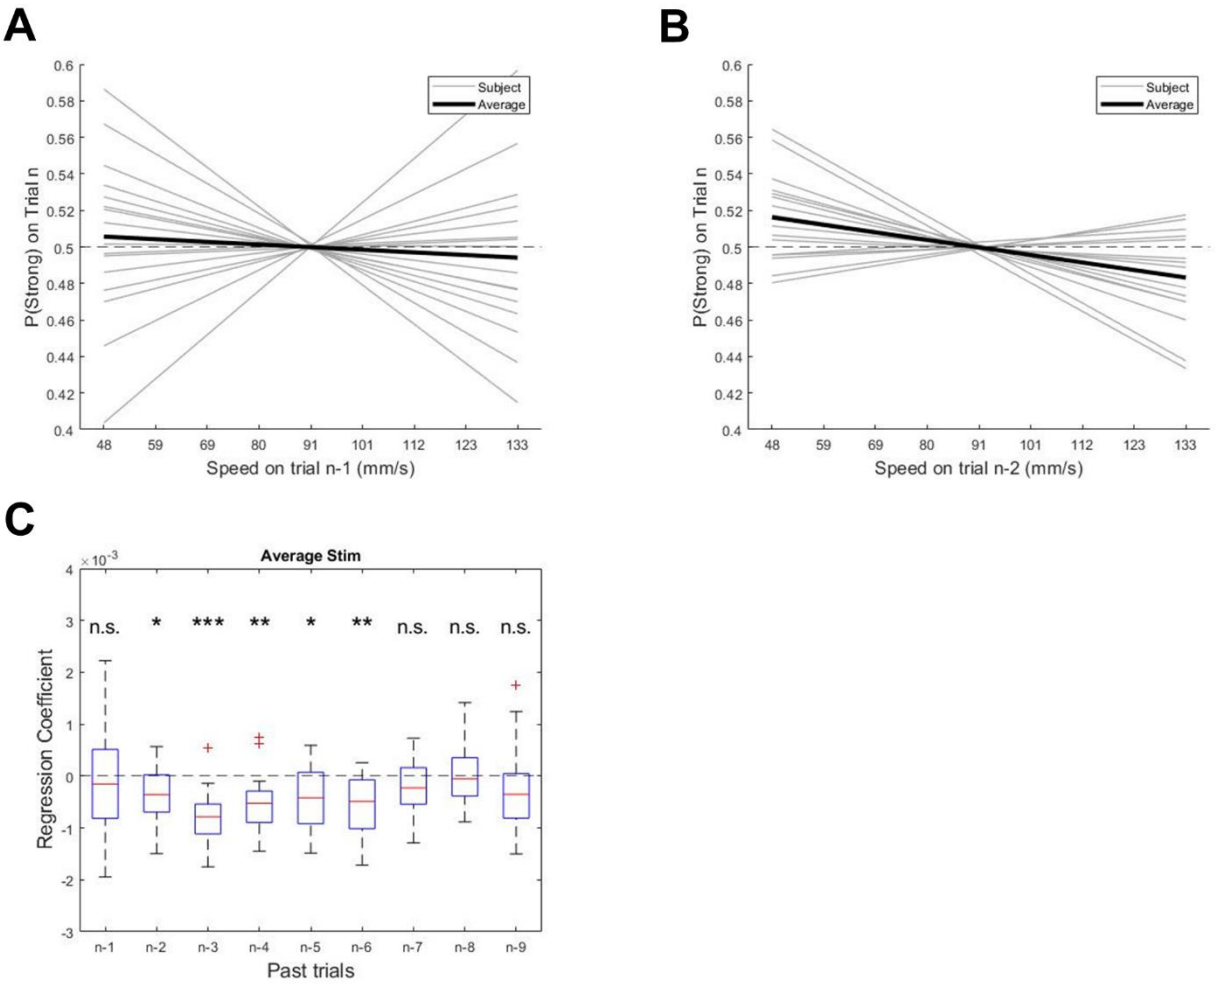

Figure 12: Reference Memory history effects in humans. (A) Effect of stimuli in trial n-1 on trial n decision. Nominal speed of stimulus in trial n-1, in mm/s, on x axis, against probability of reporting the stimulus as strong during trial n, on y axis. Grey lines are linear regression curves for each subject, black line is linear regression for pooled subjects. **(B)** as in A, for trial n-2, in mm/s, on x axis. **(C)** Linear regression slope coefficients. Past trials on x axis, regression slope coefficient, as in A, for past trials on y axis (Student's t-test). N=16 humans.

Choice serial dependence in RM

In RM, differently from WM, stimulus intensity on trial n was fully correlated to the correct choice to be made (except for trials that presented the category boundary stimulus). Thus, an attractive effect to the previous choices would run counter to the repulsion from the previous stimuli. To test for the possibility that two opposing forces largely cancelled each other, we performed a GLM analysis again, including both past stimuli and choices, as described for WM. Figures 13-A, left, and 14-A, left, for rats and humans, respectively, show the β recovered when, as predictor for trial n, we employed the stimulus intensity delivered in for selected trials. As in WM, trial n β is large and positive: this was expected, as the subjects are performing the task as intended, and thus correctly map the current trial stimulus to current trial choice. As in WM, the overall β of past stimuli are negative, meaning

they exert a repulsive effect: after a strong stimulus, subjects, whether rat or human, are more likely to report the current stimulus as weak, and vice versa. Humans show a significant effect up to trial n-3, with a profile of magnitude decay, in time, similar to that shown in WM (Figure 7-A). Rats, on the other hand, show a much shallower profile for past stimuli β, with little discernible decay; also, there is no clear progression towards non significance: Figure 13-B shows, instead, a noisy sequence of significant and non-significant trial history, differently from WM (Figure 6).

Figure 13-A, right, shows the β for past choices, separated according to the outcome. Rats showed the same opposite effects that were seen in WM: in particular, the positive β for correct choices in trial n-1 opposes the negative β for stimulus n-1, effectively cancelling each other. Humans (Figure 14-A, right), on the other hand, behaved differently than they did in WM, producing a positive β for correct choices in trial n-1, instead of negative β for both outcomes; thus, humans exhibited a strong choice repetition after a rewarded trial, masking the stimulus repulsion effect, as happened in rats.

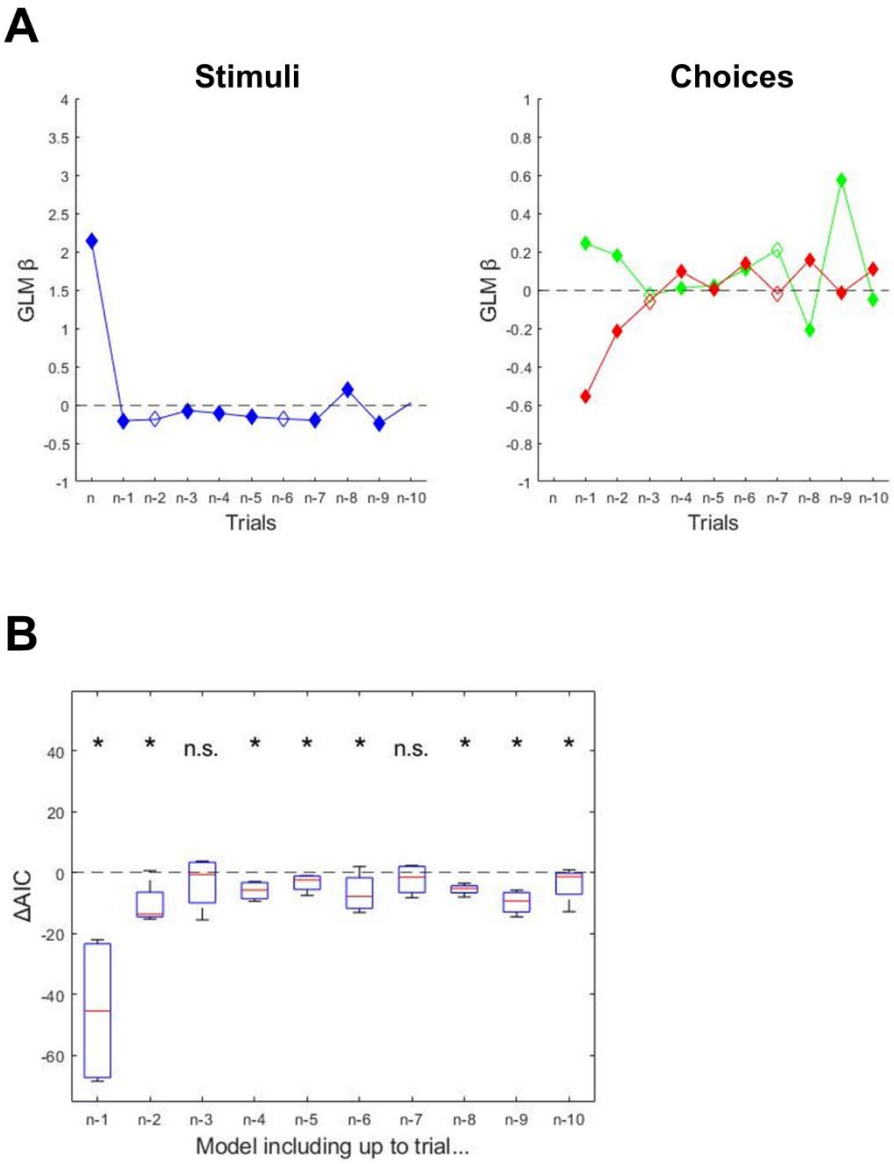

Figure 13 Reference Memory Generalised Linear Model for rats. (A) Standardised weights for stimuli (left) and choices (right), rewarded choices in green, unrewarded in red. Trials on x axis. Filled diamonds denote weights for trials deemed significant in B. **(B)** Akaike's Information Criterion difference between consecutive model iterations. X axis denotes up to which trial is included in that model iteration; y axis represents the ΔAIC between the preceding iteration and the one denoted on x axis. For n-1, the preceding iteration included only trial n (Student's t-test). N=10 rats.

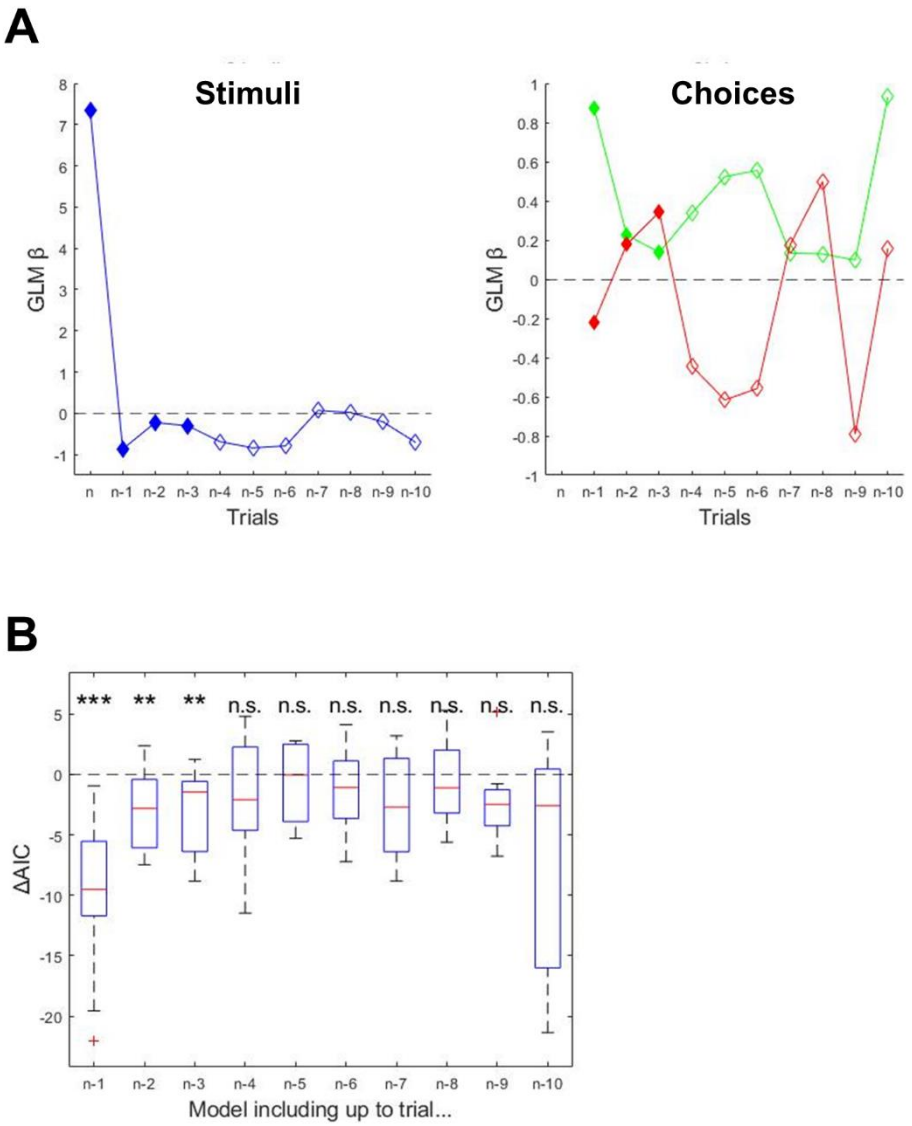

Figure 14: Reference Memory Generalised Linear Model for humans. (A) Standardised weights for stimuli (left) and choices (right), rewarded choices in green, unrewarded in red. Trials on x-axis. Filled diamonds denote weights for trials deemed significant in B. **(B)** Akaike's Information Criterion difference between consecutive model iterations. X axis denotes up to which trial is included in that model iteration; y axis represents the ΔAIC between the preceding iteration and the one denoted on x axis. For n-1, the preceding iteration included only trial n (Student's t-test). N=16 humans.

Neuronal recordings

5 rats underwent chronic electrode implant surgery to provide to allow the extracellular recording of neuronal activity, after they had successfully completed training in both tasks and their performance had been stably above 75%, in both tasks, for at least 2 months. These implants consisted of two microelectrode arrays, composed of 12 electrodes each, arranged in a 4 by 3 grid. One array was placed in prefrontal cortex (PFC) and one in posterior parietal cortex (PPC). The electrodes' tips were positioned inside the V^{th} cortical layer. The implants also contained a mechanism to allow the electrodes to be moved perpendicularly to the cortex surface for several hundred micrometres, allowing us to explore the span of the V^{th} layer. All rats were recorded during the execution of both WM and RM tasks: each session yielded one ensemble of simultaneously recorded units.

For each ensemble, we computed the firing rate of each unit in non-overlapping bins of 20ms, then smoothed across bins with a gaussian kernel with σ=100ms and a window of 500ms. The firing rates were then aligned to the beginning of each trial, identified by the nose-poking. Figure 15-A and B show the firing rate obtained from two different sample units, one each from PFC and PPC, averaging over trials where the subject reported different choices (A) or trials following different stimulus category (B).

Decoding procedure

Having detected that at least some units were sensitive to behaviour-relevant variables, we then attempted to decode those same variables, one recording session at a time. To do so, we sorted the spikes recorded during each session into single units, and grouped together the units recorded simultaneously. We will refer to these groups as ensembles: each recorded session then yielded one PFC ensemble and one PPC ensemble, whose units we analysed together. WM recordings produced 54 ensembles for PFC across all rats (units per ensemble: 8.8 ± 6) and 43 ensembles for PPC across all rats (units per ensemble: 6.3 ± 3.2). When separated according to the rats' behavioural task, RM recordings produced 37 ensembles for PFC (units per ensemble: 7.3 ± 5.1) and 29 ensembles for PPC (units per ensemble: 4.4 ± 3.7). All unit-count values are mean ± standard deviation.

Figure 15-C shows an example of decoding from one PFC ensemble during the RM task. The curves represent the decoding accuracy, expressed as distance between the decoder accuracy and the shuffled decoder accuracy, in units of standard deviation (see equation 7

and Methods). The shuffled data was produced by shuffling the trials' labels, such as stimulus, choice, etc., thus disrupting the correlation between neuronal activity and trial events. This shuffled decoder allowed us to measure the accuracy expected from "chance" decoding. The horizontal dashed line represents the threshold for statistical significance (distance = 1.96 x SD): the higher the curve extends above the dashed line, the better the decoder performed, compared to the decoder operating on shuffled data. The decoding window began one second before nose-poking, and ended one second after the go cue (which signalled to the subjects to withdraw and proceed towards their chosen reward spout; see Figure 1). As the shortest inter trial interval (ITI) duration possible was 2s, controlled via software, the portions of ITI before and after each trial never overlap.

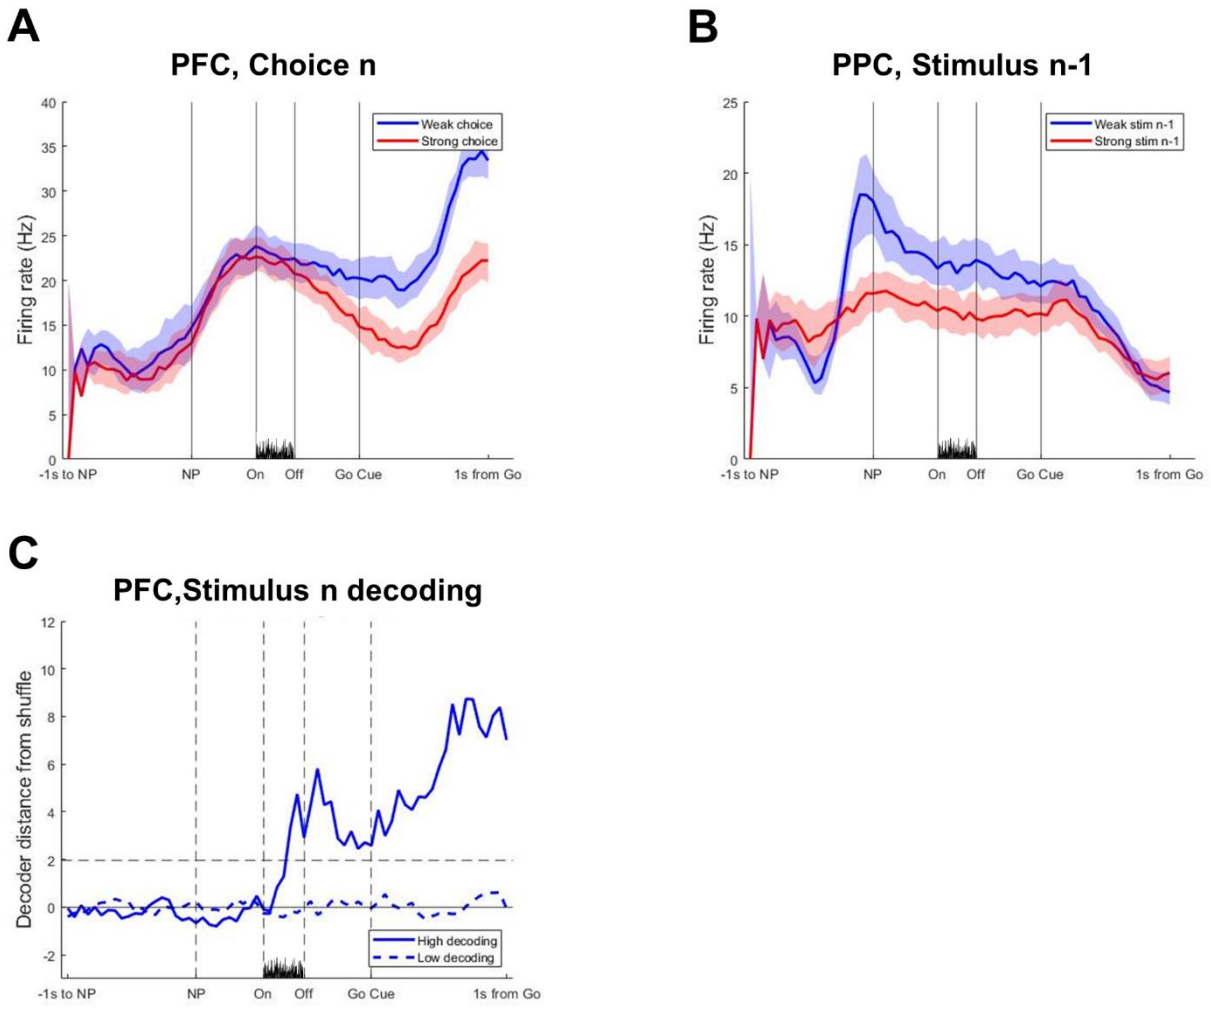

Figure 15: Sample units and decoding from recordings during RM. (A) Firing rate (Hz, y axis) of a sample unit from PFC, during RM execution. x axis is timeline, from 1s before nose-poking to 1s after the go cue. Red curve represents trials in which the rat chose strong, blue curve, weak. Shading represents the 95% c.i. for the median. **(B)** Firing rate (Hz, y axis) of a sample unit from PPC, during RM execution. x axis is timeline, identical to (A). Red curve represents trials in which stimulus n-1 belonged to strong category, blue curve represents trials in which stimulus n-1 belonged to weak category. Shading represents the 95% c.i. for the median. **(C)** Sample RM sessions decoding, from PFC, for stimulus n. x axis is timeline, as in A, y axis is distance of

decoder accuracy from shuffled decoder accuracy (see methods). Solid curve represents a high decoding session, dashed line represents a low decoding session.

Stimulus decoding in RM

Figure 16 shows how PFC (A) and PPC (B) are involved in the coding of stimuli during the RM task. A first difference between the cortical areas was the decoding of stimulus n, the stimulus of the current trial. The stimulus presentation occurred between the two dashed lines denoted by On and Off, represented by the black vibration. In PFC, it was possible to decode stimulus n already during and after stimulus presentation, with the decoder performance becoming even better after the go cue, while the subject is reporting its decision. PFC engagement during stimulus presentation and afterwards is in line with previous reports in the literature for similar task designs involving the visual modality (Wang et al., 2020). As PFC comprises a series of cortical areas whose attributed functions include motor preparation and execution, as described in the Introduction, the stronger engagement during the choice action was to be expected. In PPC, the decoder was able to identify stimulus n later in the trial, after the go cue, and with lower accuracy compared to PFC.

To explore whether such representations might have a causal role in the rats' behaviour, we proceeded to investigate whether the stimulus n decoding performance correlated with the subjects' task performance. We divided the ensembles on the basis of decoder performance: high-decoding ensembles corresponded to decoder performance in the top third of all ensembles, and low-decoding ensembles corresponded to decoder performance in the lowest third of all ensembles. As shown in Figure 16-C, high-decoding ensembles from PFC were recorded during sessions with better behavioural performance. The fact that subjects' engagement in the task correlated positively with the degree of population encoding of the stimulus is consistent with a possible causal role of PFC in the execution of this type of categorisation task (also see Wang et al., 2020.) At the same time, in panel Figure 16-E, no difference emerged from the separation according to decoding quality in PPC. Thus, we concluded that PFC might be more causally involved than was PPC in the execution of the current trial.

If not in the execution of the current trial in real time, what alternative roles might PPC participate in? PPC showed a taller and wider peak of decoding for stimulus n-1 at the beginning of trial n (Figure 16B), compared to that seen for PFC (Figure 16-A). This stimulus n-1 "reactivation" right before the beginning of the next trial (trial n) is akin to the one

described in Barbosa & Stein, 2020 in macaque spatial working memory behaviour. In this earlier work, the magnitude of reactivation, measured in decoder performance, was shown to correlate with the magnitude of stimulus after effect. We next investigated whether the magnitude of this reactivation would correlate with history effects. To do so, we divided ensembles in high- or low-decoding, according to the decoding performance for stimulus n-1, for both areas separately. Once sorted, we computed stimulus history curves for the two groups of sessions, as described for Figure 11-A in the behavioural section. The results are shown in Figure 16-G and H, for PFC and PPC, respectively. We expected that, once so divided, high-decoding sessions would show a larger stimulus history effect than the low-decoding sessions. Surprisingly, the magnitude of reactivation of stimulus n-1 in PFC yielded flat curves for both high- and low-decoding ensemble. In PPC, the high decoding sessions revealed a slight trend in the direction we expected: a negative slope in the stimulus n-1 regression curve, which arises from a stimulus repulsion. However, this trend too was non-significant.

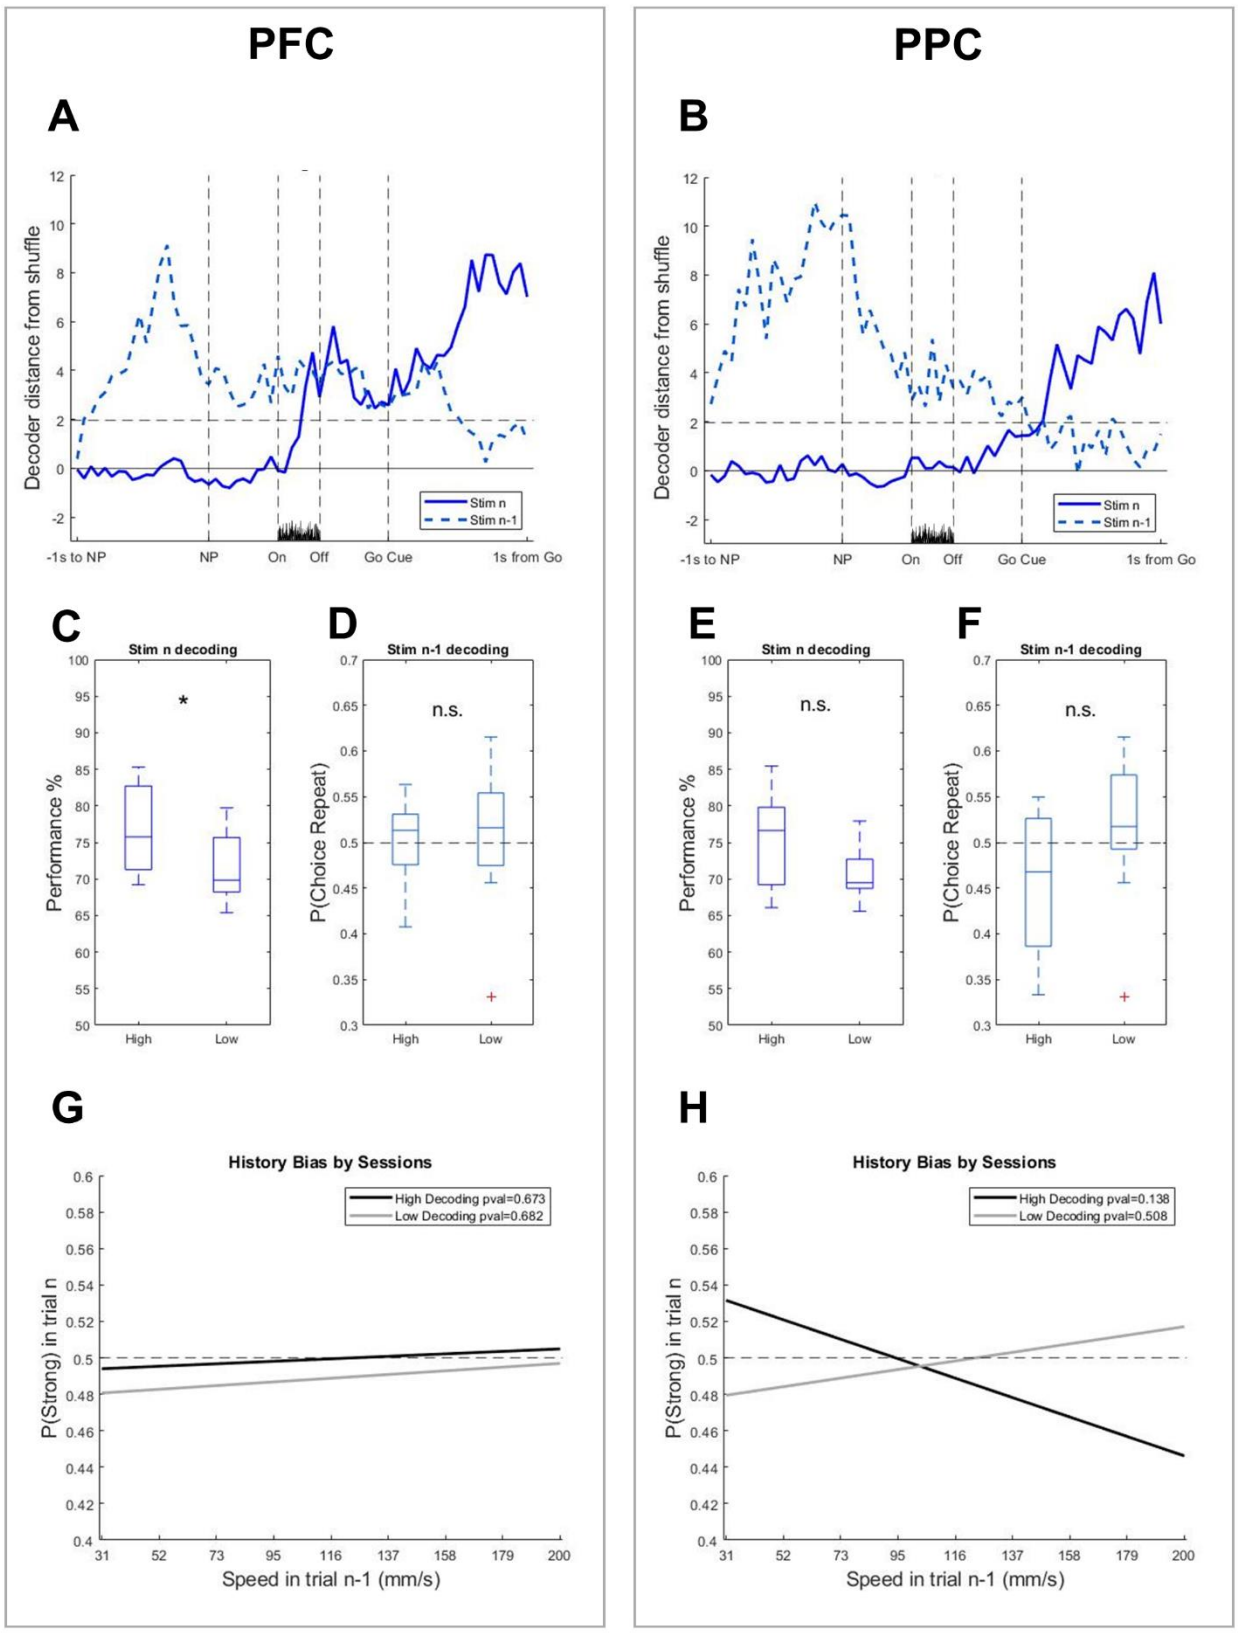

Figure 16: Stimulus decoding in RM. (A) Stimulus decoding, from PFC, for stimulus n (solid) and stimulus n-1 (dashed). Curves represent median from high decoding sessions. x axis is timeline, from 1s before nose-poking to 1s after the go cue. y axis is distance of decoder accuracy from shuffled decoder accuracy (see methods). **(B)** Same as A, for PPC. **(C)** Subjects' performance for sessions belonging to high or low stimulus n decoding ensembles, for PFC. **(D)** Subjects' choice repetition probability for sessions belonging to high or low stimulus n-1 decoding ensembles, for PFC. **(E)** Same as C, for PPC. **(F)** Same as D, for PPC. **(G)** History

curves for sessions belonging to high decoding (black) and low decoding (grey) ensembles for stimulus n-1, for PFC. x axis is stimulus n-1 nominal speed, y axis is probability of judging stimulus as strong in trial n. **(H)** Same as G, for PPC.

Choice decoding in RM

As shown in the behavioural section (Figure 13-A, left), rats exhibited a strong choice attraction after a rewarded trial. To search for a possible neural correlate of choice attraction in RM, we applied the same analysis to choices, instead of stimuli. Figure 17 shows the decoding curves for choice in PFC (A) and PPC (B). Here, the decoding performance for choice n was larger than that of stimulus n in both areas, and especially in PPC. The earliest moment at which choice n became decodable, however, still appeared in PFC, in the same window when stimulus became decodable. This is consistent with the fact that subjects were employing stimulus information to solve the task, and, in RM, each trial's stimulus determined the correct choice to be made, except when the boundary stimulus was presented. However, as seen in panel C, dividing sessions according to decoder performance for choice did not produce any difference in subjects' performance. As PFC's peak for stimulus decoding close to stimulus presentation was larger than that of choice, and vice versa for the peaks after the go cue, it is conceivable that there are either two interspersed populations in PFC, one for early stimulus coding, and one for late choice coding, or that the same population shifts from coding one to the other in different trials epochs, akin to the description given in Romo et al., 1999 for macaque vibration frequency delayed comparison. Further analysis would be required to distinguish between these, or other alternative hypotheses.

As for PPC, when sessions were divided between high- and low-decoding for choice n, high-decoding sessions show better behavioural performance. Since choice coding in PPC emerged later than in PFC, and so a causal link between PPC encoding and subjects' performance is unlikely, it is possible that the better encoding of choice found in PPC was due to the subjects' higher engagement with the task in those sessions.

As choice n-1 was proven to have a significant effect on the current trial, we also separated sessions according to the reactivation of choice n-1 decoding in both areas. However, this separation did not highlight difference in choice repetition (Figure 17-D for PFC, Figure 17-F for PPC). Also, the history curves obtained via the same separation between high-coding and low-coding for choice n-1 resemble those described for the separation according to stimulus n-1 coding: no effect in PFC, and only a trend for high coding sessions in PPC.

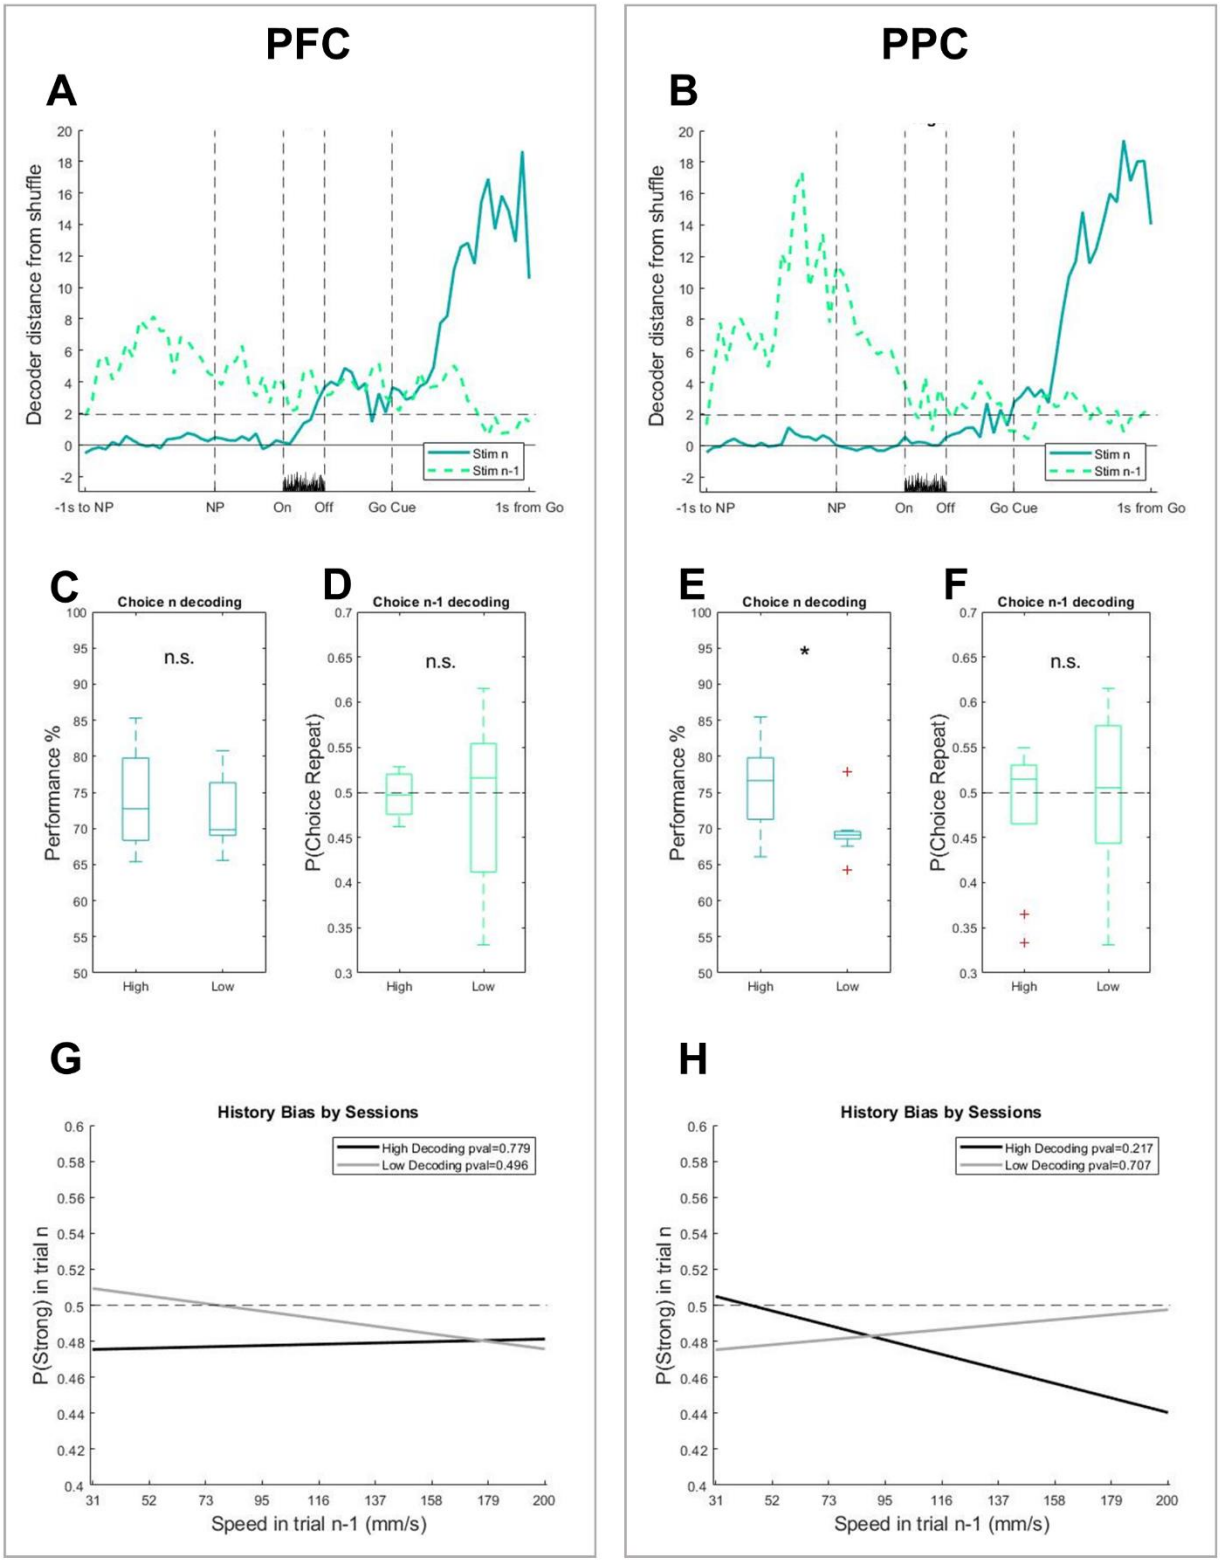

Figure 17: Choice decoding in RM. (A) Choice decoding, from PFC, for choice n (solid) and choice n-1 (dashed). Curves represent median from high decoding sessions. x axis is timeline, from 1s before nose-poking to 1s after the go cue. y axis is distance of decoder accuracy from shuffled decoder accuracy (see methods). **(B)** Same as A, for PPC. **(C)** Subjects' performance for sessions belonging to high or low choice n decoding ensembles, for PFC. **(D)** Subjects' choice repetition probability for sessions belonging to high or low choice n-1 decoding ensembles, for PFC. **(E)** Same as C, for PPC. **(F)** Same as D, for PPC. **(G)** History curves for sessions belonging to high decoding (black) and low decoding (grey) ensembles for choice n-1, for PFC. x axis is stimulus n-1 nominal speed, y axis is probability of judging stimulus as strong in trial n. **(H)** Same as G, for PPC.

Stimulus decoding in WM

We performed the same library of analyses on WM ensembles. First, we examined the decoding performance for base and comparison stimuli in both cortical areas, trying to draw a comparison with the decoding time course of the RM stimulus. We expected that PFC would display a robust decoding for both stimuli, possibly with better representation of the comparison stimulus, as it is the stimulus closer in time to choice execution. We also expected PPC to be lesser involved in the stimuli in trial n than PFC. The results are shown in Supplementary Figure 3-A and B, for PFC and PPC, respectively. PPC, indeed, proved to hold very little information about individual stimuli in trial n. At the same time, PFC too, surprisingly, resulted in a very poor decoding profile for base stimulus n, and the decoding for comparison n proved inferior to that of stimulus n in RM.

NSD decoding in WM

From the unexpectedly weak encoding of the individual stimuli of trial n, we hypothesized that the PFC-PPC network might be more involved with the tasks' decision variables, rather than the stimuli per se. The distinction would not be apparent in RM sessions, during which stimulus and decision variable coincide, but would become immediately evident in WM sessions, where an individual stimulus intensity cannot determine the correct decision. Indeed, the decoding of the NSD value defining each trial proved to be robust (Figure 18-A and B, for PFC and PPC, respectively), though to a lesser extent than that of stimulus n in RM (Figure 16-A and B). When separating ensembles according to decoder performance for NSD, high-decoding ensembles from PFC proved to come from higher performance sessions, compared to low-decoding (C). PPC, on the other hand, demonstrated no such difference (E). This result replicated that of the analogous analysis for RM (Figure 16-C and E), further reinforcing the hypothesis that PFC is actively engaged in the encoding of task-oriented decision variables, while PPC is downstream of PFC with regard to the current trial's information flow. PPC might be more involved with building a register of past-trial history.

When we computed history curves for high- and low-decoding ensembles for stimulus n-1 in RM, no differences emerged (Figure 16-G and H). For comparison n-1 coding in WM, however, both PFC and PPC presented a difference: low-decoding ensembles showed significant linear regression coefficients, similar to those shown in the behavioural results

(Figure 4-A), while the high-decoding ensembles did not. As described above in the RM results, we expected the opposite, based on literature.

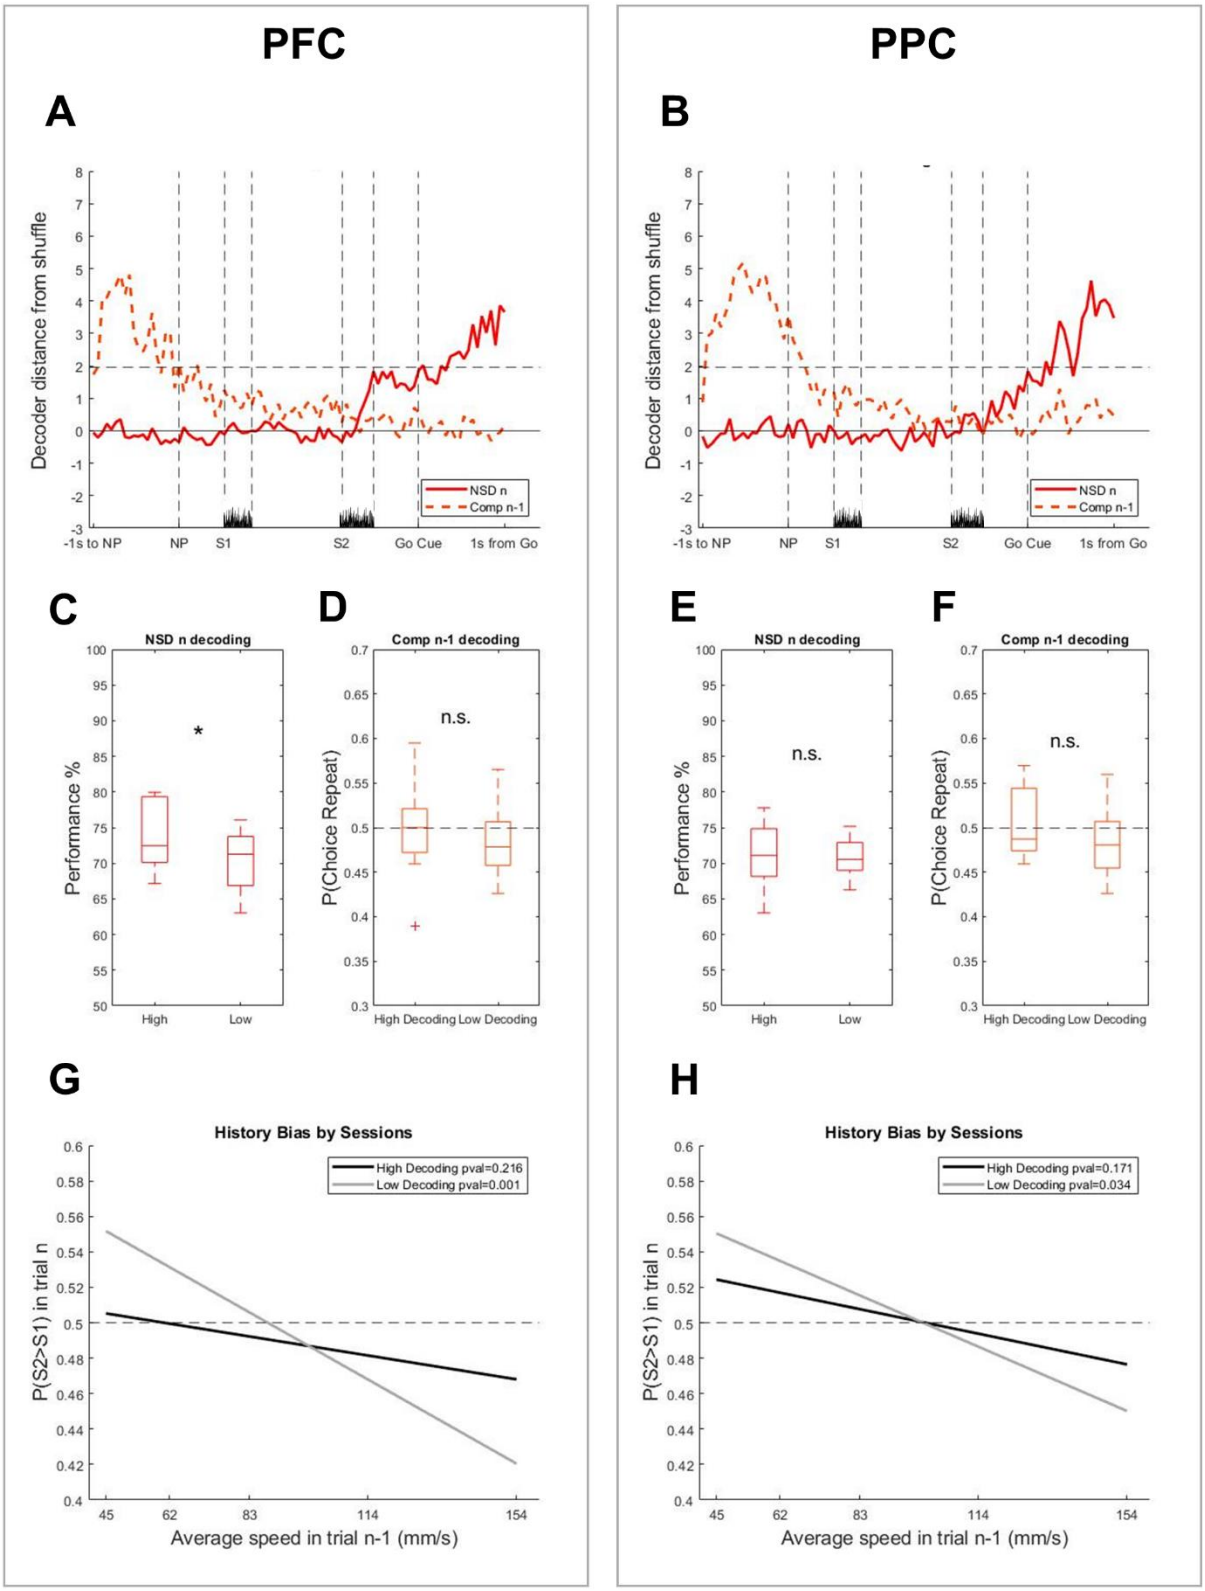

Figure 18: NSD and comparison decoding in RM. (A) Decoding, from PFC, for NSD n (solid) and comparison n-1 (dashed). Curves represent median from high decoding sessions. x axis is timeline, from 1s before nose-poking to 1s after the go cue. y axis is distance of decoder accuracy from shuffled decoder

accuracy (see methods). **(B)** Same as A, for PPC. **(C)** Subjects' performance for sessions belonging to high or low NSD n decoding ensembles, for PFC. **(D)** Subjects' choice repetition probability for sessions belonging to high or low comparison n-1 decoding ensembles, for PFC. **(E)** Same as C, for PPC. **(F)** Same as D, for PPC. **(G)** History curves for sessions belonging to high decoding (black) and low decoding (grey) ensembles for stimulus n-1, for PFC. x axis is average nominal speed of stimuli in trial n-1, y axis is probability of judging comparison as stronger than base in trial n. **(H)** Same as G, for PPC.

Choice decoding in WM

Like for RM, we extended the analysis to the decoding of choice n and choice n-1. As was the case for RM, the decoder's performance for choice n was better than that of NSD n in both cortical areas. Neither areas' decoder performance correlated with subject performance. On one hand, this result further establishes the role of PFC as encoding a pure decisional variable in the time window of choice reporting; on the other hand, PPC's encoding of choice did correlate with subjects' performance in RM (Figure 17-E), but not in WM. It is conceivable that this is due to the strong correlation between stimulus and correct choice in RM, a correlation that is absent in WM. As an alternative explanation, it is possible that PPC deals with longer scale history, which is more relevant for the execution of RM, where the categorical boundary is long history, than of WM.

Finally, the division between high- and low-decoding for choice n-1 revealed a distinction between PFC and PPC when it came to stimulus history curves: while PFC demonstrated significant curve slope (Figure 19-G) in low-decoding sessions only, as had been the case for comparison n-1 decoding, PPC's low-decoding sessions are just above the threshold for statistical significance (Figure 19-H); PPC's high-decoding sessions, instead, demonstrated significant stimulus history effect.

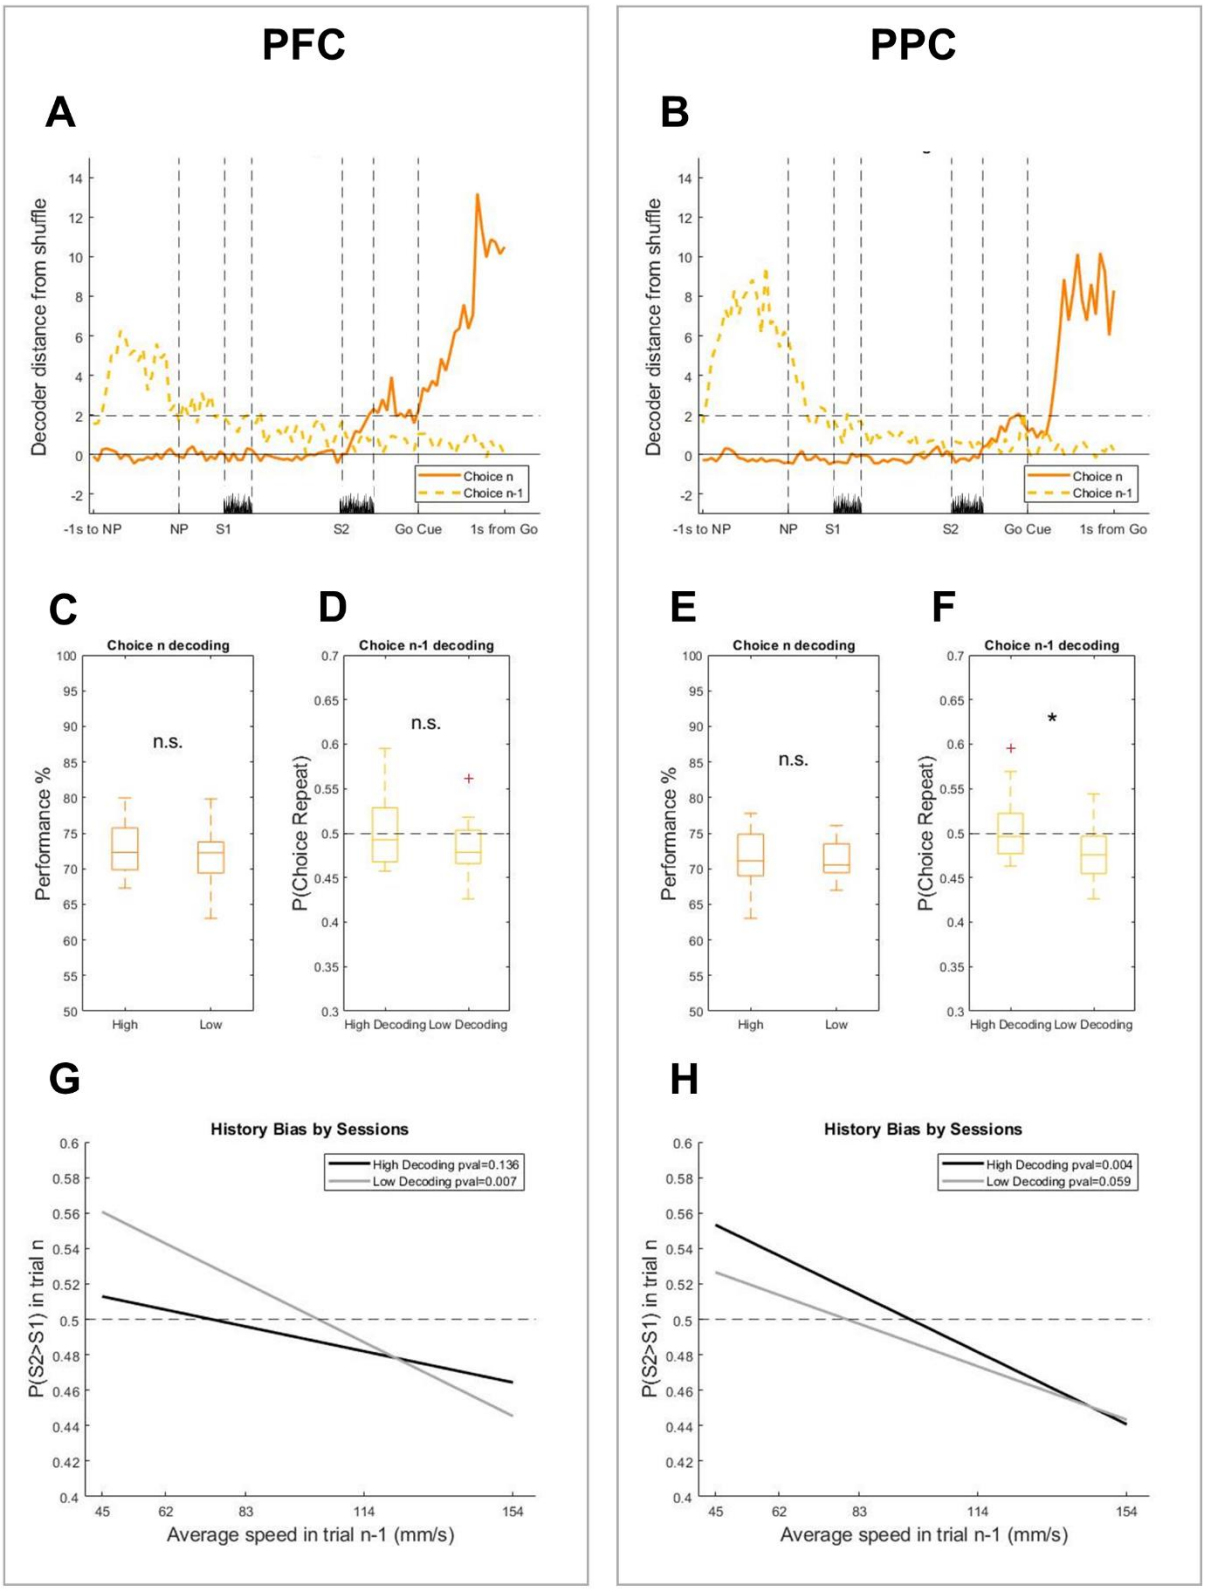

Figure 19: Choice decoding in WM. (A) Choice decoding, from PFC, for choice n (solid) and choice n-1 (dashed). Curves represent median from high decoding sessions. x axis is timeline, from 1s before nose-poking to 1s after the go cue. y axis is distance of decoder accuracy from shuffled decoder accuracy (see methods). **(B)** Same as A, for PPC. **(C)** Subjects' performance for sessions belonging to high or low choice n decoding ensembles, for PFC. **(D)** Subjects' choice repetition probability for sessions belonging to high or low choice n-1 decoding ensembles, for PFC. **(E)** Same as C, for PPC. **(F)** Same as D, for PPC. **(G)** History curves for sessions belonging to high decoding (black) and low decoding (grey) ensembles for choice n-1, for PFC. x

axis is stimulus n-1 nominal speed, y axis is probability of judging comparison as stronger than base in trial n. **(H)** Same as G, for PPC.

DISCUSSION

In the behavioural results, we showed how stimulus history acted in the context of the WM task in rats and humans. Figure 1 (rats) and Figure 4 (humans) illustrate how contraction bias acted differently according to task and experimental designs. In rats, three psychometric sets shared the same NSD values, but had different base stimuli; in this design, the psychometric curves derived from each set shifted horizontally, in a direction that would be expected if the remembered base stimulus moves towards the centre of the stimuli distribution. In other words, the memory traces of the base stimuli showed a regression to the mean of all stimuli. In humans, two psychometric sets had either fixed base (vertical set, analogous to the central rats set) or fixed comparison (horizontal). In the horizontal set, the memory traces of base stimuli, which varied from pair to pair, were contracted towards the centre of the distribution; the contraction lead to smaller effective NSDs, as base stimulus trace became more similar to comparison stimulus, thus impairing performance. In the vertical set, on the contrary, the base stimulus was already very close to the centre, so the memory traces were presumably reinforced by contraction bias. Contraction bias "anchored" the memory close to its actual perceived value, thus improving performance. These results are coherent with past descriptions of contraction bias both in the tactile (Fassihi et al., 2014, 2017; Serrano-Fernández et al., 2023) and auditory modalities (Akrami et al., 2018), and in time estimation (Tal-Perry & Yuval-Greenberg, 2022).

Rats and humans were affected by serial dependence in a similar way, as well. Figure 2 (rats) and 5 (humans) show how the effects of past stimuli extends beyond the most recent trial, n-1. While humans showed a steeper decay of regression slope coefficients, compared to rats, both species showed a reduction in slope when considering trials further in the past, similar to that shown in (Hachen et al., 2021) for the RM design. The GLM analysis of Figures 3 (rats) and 6 (humans) confirmed this temporal profile. The GLM also began to highlight the role of choice history. For both species, choice n-1 carried by far the largest weight. Humans, though, showed a *repulsive* effect for choice n-1 regardless of that choice's outcome, while rats showed opposite effects, depending on the outcome. As seen in other

experimental paradigms with rodents, rats exhibited a tendency to win-stick, and lose-switch, i.e.: repeat a choice that led to a reward, and reverse a choice that led to no reward.

These results reveal a remarkable similarity in how stimulus history affects current WM choices in both rats and humans. Namely, both species exhibited contraction bias, with base stimulus memory contracting towards the centre of the stimuli distribution; in both species past stimuli exert the same effect, although to a lesser degree in humans, and the magnitude of their effect reduces according to recency. Moreover, all these hallmarks of stimulus history effect were consistent across subjects, regardless of species. Two main differences emerge between rats and humans, beyond the effect magnitude: the first is the extent to which stimulus history extends to past trials, with rats being affected by trials much further in the past, compared to humans, who showed a shorter-lived trial history. The second is the choice outcome dependency; rats showed a tendency to repeat choices that had been rewarded, and not to repeat choices that went unrewarded, while humans showed a tendency to alternate choices from one trial to the next, irrespective of outcomes.

The results for the second task, RM, proved again similar across rats and humans, but different from both WM and previous RM literature. Figures 5 (rats) and 8 (humans) show that stimulus serial dependence in RM was much more heterogeneous across subject than it was in WM. Subjects showed no overall trend. A possible reason for this absence is shown in Figures 9 (rats) and 12 (humans): disentangling stimulus and choice serial dependence, it is possible to show that a small stimulus effect is present. However, in both groups, a strong choice n-1 attractive effect, when choice n-1 had been rewarded, opposed stimulus n-1 repulsive effect, masking it.

We formulated two hypotheses that might explain this discrepancy. The first has to do with the task design: in the present work, we employed the same dynamic stimulus range for both tasks, anchoring the RM one to WM, which we established first. It is possible that such a range was simply too large, compared to the one employed in earlier work. The large range might have masked the stimulus serial dependence to emerge: a larger range makes the task easier, with stimuli getting farther from the categorical boundary; as each stimulus became easier to categorise, choices would become less susceptible to stimulus history.

Thus, the choice attraction effect is not unexpected in the current results: it is, however, comparatively larger than the serial stimulus dependence. It is possible that, employing a narrower, and thus more difficult, stimulus range, the stimulus effect would prevail over

choice attraction. The second hypothesis is applicable to rats only. We trained rats in WM first, as it is the harder task to learn. It is possible that they then approached the RM task differently from rats who learned *only* the RM task. As stimulus history mostly affects performance negatively in WM (each trial is, ideally, independent and contraction bias tends to have a net harmful outcome), rats might have learnt to minimise its influence through the course of long training. The routine of treating each trial in isolation rather than as a continuous string might have allowed them to suppress stimulus history influences once presented with a task, RM, without an explicit memory component. However, humans surprisingly showed the same result as rats, even though half of the subjects experienced the RM task first. Moreover, the task structure provided no incentive to employ choice n-1 in solving the task, as shown in Supplementary figure 4: the structural probability of the same correct side being repeated from one trial to the next deviated slightly for 0.5 for all tasks and species; also, the deviation is mostly in the direction opposite to how the subjects employed choice n-1. For example, human subjects performing WM showed choice repulsion (Figure 7-A, left), but the actual probability of the same correct side repeating resulted slightly above 0.5, and, thus, choice attraction would have been preferable. Due to this, we would favour the first explanation.

Our strategy of having subjects perform both tasks was based on the idea that stimulus history effect magnitude in one task would predict the effect in the other task, thus providing evidence for a single phenomenological and physiological mechanism for both contraction bias and stimulus history effects, that would then be declined according to the cognitive demands of each task. Clearly, non-sensory history components need to be accounted for, too: while stimulus serial dependence and contraction bias in WM show a remarkable similarity between rats and humans, the results in RM do not clearly follow from the stimulus history effect as in WM. While rats and humans still behaved similarly in the RM task, this was due mostly to serial choice dependence. We cannot exclude that the stimulus history effect would have been similar in RM; the subjects' tendency to repeat choices, after rewarded trials, masked the stimulus repulsion effect, thus providing inconclusive evidence regarding the development of stimulus history effects in different tasks within each subject.

As for the neuronal recordings, the results show that both PFC and PPC are part of a network involved in the processing of task relevant information. Figures 15 and 17 show that PFC is more involved in the representation of current trial's decisional variables, with a direct correlation to subjects' performance. PPC, on the other hand, turned out to be more involved

in representing past trials, both with respect to stimulus and choice. This finding agrees with PPC data in an auditory delayed comparison task (Akrami et al., 2018). Due to this, and the time course of decoding curves, it is plausible that PPC is located downstream to PFC in the flow of information. However, the encoding of past stimuli and choices, in both cortical areas considered, failed to correlate systematically with either stimulus history or choice history behavioural effects, differently from what was reported in (Barbosa & Stein, 2020), where the magnitude of reactivation in macaque PFC directly correlated with serial dependence in a saccade-to-target task. While it is entirely possible that the failure to uncover this correlation depends on the task design differences, or species differences, it is also plausible that this information is encoded, in these cortical areas, with variable precision from session to session. Whether to employ it or not would then be gated by a different part of the same network, not recorded in the present work, which would act in a top-down manner on either or both PFC or PPC. Indeed, in a recent publication (Cazettes et al., 2023), the authors showed that rodent secondary motor cortex (M2), which is part of the prefrontal rodent network, can act as a reservoir of information to be employed in different viable strategies, regardless of whether that specific strategy is actively employed or not. It is not far-fetched to hypothesize the same for history effects, although they are not necessarily part of an explicit strategy employed by the animals.

While both areas showed a reactivation of past stimuli and choices, in WM the magnitude of reactivation proved to be correlated with stimulus effects. Possibly, this happened because, in RM, stimulus serial dependence was hard to detect directly, without more elaborate analysis, from the behavioural point of view. This correlation, however, proved counterintuitive, as the sessions characterized by a larger reactivation also proved to be less affected by stimulus history, and vice versa. We propose two potential hypotheses to explain this result: the first is that the reactivation is actually a counterbalancing signal, intended to reduce the bias caused by the stimulus history, and not the stimulus history signal per se. The second possibility is that the stimulus history effects are due to a flow of information from the PFC-PPC network examined here: when the magnitude of reactivation is low, the information has already been processed and delivered to another, hidden, portion of the network, responsible for the enactment of the bias due to the stimulus history; when the magnitude of reactivation is high, it means that the information is still retained in the visible portion of the network, and thus is yet to be processed, leading to no measurable stimulus history effect.

METHODS

Rat Subjects

10 Wistar male rats (Envigo) were housed in quartets and examined periodically by a veterinarian. When they reached around 400-500g they were split in pairs, due to cage dimension limits. The rats began the training protocol at the age of 8 weeks. They were motivated by water deprivation and rewarded during trials with tap water. The deprivation schedule gave access to water for 1 hour during training and for around 30 hours during weekends.

All rats were trained for intensity discrimination in two different tasks: first, they were trained to perform the delayed comparison task (WM); once proficient in it, they were introduced to the categorisation task (RM) and then alternated from one to the other on a weekly basis.

Protocols conformed to international norms and approved by the Italian Health Ministry and Committee of the International School for Advanced Studies (SISSA).

Rat Setup

The experimental setup consisted of a T-shaped Plexiglas box, measuring 25 × 25 × 38 cm. In the front wall, a hole allowed the rat to extend its head into the stimulus delivery port; here, a nose poke detected the snout by an infrared detector, while a blue LED light signalled to the rat that a new trial could begin. The setup was positioned into a Faraday cage and the experiments were run in the dark; the only light provided was a red LED and was used by the experimenter to monitor the rats during the task, using a webcam (Logitech HD Webcam C310). During the trial a shaker motor (type 4808; Bruel and Kjaer), with 12.7 mm peak-to-peak displacement, was used to generate stimuli. The motor was placed on its flank to produce motion in the horizontal dimension through a 20 × 30 mm plastic plate attached to its diaphragm. We applied a strip of double-sided adhesive tape to this plate to keep the rat's whiskers in contact and allow them to better follow the plate's movement. Two reward spouts, on 8 cm high pedestals, were place to the right and left of the nose poke; two custom-made infrared sensors, one for each spout, detected the rat's position and activated reward delivery accordingly. The water was delivered by two syringe-pumps, one for each side, controlled by a AVR32 board (National Instruments, Austin, TX). Three audio speakers were placed on three walls of the setup, one above each spout and one at the back: the central

one delivered the go cue, while the lateral ones gave a reward delivery cue, upon correct response, to act as reinforcement.

Experiments were controlled using LabVIEW software (National Instruments).

Stimulus generation

Each stimulus was a noisy vibration, obtained by stringing together randomly sampled velocities. Velocities were obtained by sampling a normal distribution with 0 mean and defined by the standard deviation σ, ranging from 3 to 20: a stimulus with low σ will feel weak, because values of velocity will be close to the mean, while a stimulus with high σ will feel strong because values of velocity will be farther from the mean; in fact, the mean stimulus speed is directly proportional to σ, and the motor amplifier gain was set so that the average stimulus speed would be 10 times the value of σ, measured in mm/s. The sequence of velocities formed one seed. There were 50 different seeds for each vibration mean speed, making the signature of a given vibration hard to recognize and making the mean speed a more salient feature.

Experimental design: delayed comparison (WM) task

Before every trial, the blue light turned on, signalling to the rat that a new trial could begin. Then, the rat had to insert its head in the hole and keep its nose in the nose poke, triggering the delivery of the stimuli. The rat had to remain in the nose poke for the entire trial, comprising: pre-stimulus delay (500ms), base stimulus (334ms), inter-stimulus delay (either 1, 2 or 4s), comparison stimulus (334ms), post-stimulus delay (500ms); after this sequence elapsed, the go cue signalled to the rat to make a choice. If the rat moved away from the nose poke before the go cue sounded, the sequence would abort, and the rat would have to wait 2s before starting a new trial. After the go cue, the rat reported its judgement based on the rule that was reinforced during training: half of the rats, randomly chosen when they arrived in the lab, were trained to go to the left spout if the base stimulus was stronger than the comparison and to the right spout if the comparison was stronger than the base, while the other half had the rule reversed.

The delayed comparison paradigm allowed the distinction of different phases and different cognitive functions and, importantly, allowed us to separate in time the perception of the

base stimulus from the decision making, involving working memory. The sequence of perceptual and decision making operations can be envisioned as:

1) Encoding the first stimulus and extracting the relevant parameter (intensity)

2) Storing the parameter value in memory and maintaining it

3) Encoding the second stimulus and extracting the relevant parameter

4) Comparing the second parameter value to the memory of the first

5) On the basis of the comparison, applying the decision rule (go right if base > comparison or, for different rats, base < comparison)

Nominal base and comparison pair difficulty does not depend on the value of the stimuli, but on the difference between them: the closer they were, the more difficult the discrimination. We expressed this difficulty using an index, the Normalized Speed Difference (NSD), computed according to Weber's Law:

$$(1)\ NSD = \frac{\sigma_{comparison} - \sigma_{base}}{\sigma_{comparison} + \sigma_{base}}$$

The value of the index expresses the relative difference between the stimuli features, and the sign expresses which one is stronger.

To generate the different stimuli pairs used to train and test rats, we used the Stimulus Generalization Matrix (SGM). In this way, in every trial, neither of the two stimuli, taken alone, contained enough information to solve the task. If either the base or the comparison was fixed across trials, the rat could solve the task by applying a threshold to the other. Instead, both stimuli varied in a similar range in the SGM, forcing the rat to compare the two to solve the task and be rewarded.

In Figure 2-A, each circle represents a stimuli pair, using the σ of the two stimuli as coordinates, in a logarithmic scale. The diagonal line, where the two stimuli are identical, represents a discrimination line: in the upper side, comparison is larger than base for all trials, and vice versa in the lower side. Moreover, the distance of each couple from the diagonal represents the NTD or NSD: the closer to the diagonal, the more difficult the trial. In the SGM, all trials are the same distance from the diagonal, so all are equally difficult.

In order to measure acuity, we added three other sets of stimuli pairs on top of the SGM. The new pairs have different grades of difficulty, meaning different NSD values, including the 0 case, when base and comparison have the same σ. Each of the new sets had fixed base σ (either 5.8, 8 or 10.8), so they appear vertically arranged on the SGM; the comparison σ varied, producing a set of 7 NSD, from -0.3 to 0.3. To avoid the rats setting a threshold on the comparison, which would happen if always presented with the same base, both SGM set and one of the vertical sets were presented during each session. To maintain the centre of the stimulus distribution fixed in the different set combinations, we manipulated the presentation probability of each pair.

Experimental design: categorisation (RM) task

The task was operated by the subject in the same manner as described for the WM task. The trial structured differed from WM, being comprised of: pre-stimulus delay (500ms), stimulus (334ms), post-stimulus delay (500ms); after this sequence elapsed, the go cue would signal the rat to make a choice. The choice reports if the animal categorised the stimulus as strong or weak, compared to an implicit boundary. The stimulus set included 9 linearly and equally spaced stimuli, employing the dynamical range set for the WM task. The 5th stimulus lied on the implicit boundary and the associated subject choice was rewarded as correct with a 50% probability, regardless of the choice.

Training stages

All rats underwent an 8-stage training protocol in order to learn to perform discrimination on both tasks.

- 1st stage: Handling. After arriving to the facility, the rats were sorted into cages, two for each cage. They had a couple of days to adjust to the new environment, then the experimenter would hold them and pet them for an hour every day, while they got used to human handling; during the last days of this stage, the water restriction schedule began and water was available only during the handling hour, so the rats would begin to associate water availability with human handling.
- 2nd stage: Nose Poking and Reward Collection. In this stage, the rat would explore the experimental apparatus and learn the basic mechanism of reward collection. The vibrating plate was present but immobile. During exploration, the rat would discover the head hole and, when the snout was detected by the nose poke sensor, an

immediate go cue would play, followed by the delivery of a small amount of water by a randomly chosen spout; at the same time, a clicking sound would be played by the corresponding speaker. After this, the blue light would turn on again and a new “trial” could begin. Rats learn quickly to poke their nose to receive a reward. At this point, the experimenter would gradually raise the waiting time before the go cue would sound: if the rat left the nose poke before the time had elapsed, no reward would be granted. When the rat was capable of waiting for two seconds, the next stage would begin.

- 3rd stage: Passive Task. In this stage, the rat would get used to the delivery of the stimuli and start learning the correlation between stimuli features and reward side. During this stage, only the relevant feature would be involved, while the other (duration for intensity rats and intensity for duration rats) would remain fixed throughout all stimuli. Now, instead of just waiting for the delay before the go cue, the rat had to attend to both the stimuli before the go cue played. Again, early withdrawal would yield no reward. After the go cue, the rat could go to either spouts, though only the correct one, according to the rat’s rule, would deliver a reward; however, during this stage, if the rat’s first choice proved wrong, it could still go to the other spout and be rewarded. The correct response depended on the comparison between the stimuli and the rule (e.g.: comparison>base go right, or vice versa) was assigned randomly to the rats and kept the same from here on. During this stage, the experimenter would manually and gradually increase the gain of the motor, until reaching the level used for the actual experimental sessions: this was done to avoid scaring the rat with too strong a stimulus, and to get them used to the vibrations.
- 4th stage: Active Task. At this stage, reward collection for the second choice was no longer allowed. The rat had to learn to actively compare the stimuli and find the correct response. We began measuring rat performance. When it was stable and solid enough, the rat would move to the next stage.
- 5th stage: Introduction of Acuity Task. At this point, we added the acuity pairs, as described above, to the SGM.
- 6th stage: Introduction of RM task. At the beginning of the working week, rats would unknowingly start a RM session instead than WM. No cue of the change was provided, except for the fact that now a single stimulus was delivered in each trial, instead of two. The side rule was kept consistent with the one taught in WM, e.g.: a rat that had been taught that if comparison is stronger than base, he should go to the

left spout, now was taught that if the stimulus is strong, he should go to the left spout. At the beginning of the first session, only the strongest and weakest stimuli would be delivered, in order to simplify the task; then, as the number of correctly categorized trials increased, the other stimuli were introduced gradually, in increasing difficulty. After a period of confusion, which usually lasted at most 3 sessions, the rats became proficient with the task.

- 7th stage: Return to WM. At the beginning of the next work week, the rats would get back to the usual WM task. Again, no cue as to the nature of the task was delivered: only the number of stimuli that were delivered per trial. We carefully monitored that rats' performance was not negatively impacted by learning the RM task. Once we were satisfied that the rats were performing at least as well as before, we moved to the last stage.
- 8th stage: Alternating tasks. At this stage, the rats alternated between one task and the other on a weekly basis.

Human subjects

16 subjects participated in both WM and RM tasks. All subjects performed two experimental sessions, with each session lasting approximately one and a half hour. The order in which each subject performed the task was randomized across subjects. All subjects were volunteers and were paid after the participation, on the basis of how well they performed. The study protocol conformed to international norms and was approved by the Ethics Committee of the International School for Advanced Studies (SISSA).

Human Setup

All experiments took place in the Tactile Perception and Learning Lab, at SISSA. The setup was identical for both experiments. The subjects sat in front of a PC screen with their right arm resting on a pillow upon the desk, with their right index fingertip triggering an identical sensor to the rats' setup to begin a trial: when doing so, their fingertip would be in contact with a plastic probe. The probe was attached to a shaker motor (type 4808; Bruel and Kjaer), producing vibrations in the horizontal direction, perpendicular to the fingertip. Subjects' responses were reported by pressing one of two buttons with the same hand. The experiment was automatized via LabVIEW (National Instruments).

Stimulus generation

In both experiments, vibrations were produced in the same manner as described in the rats' section.

Experimental design: delayed comparison (WM) task

The task was designed in the way described for rats, save for the construction of the stimuli pairs set, shown in Figure 3-A. The stimuli pairs were arranged in two sets: a vertical set, in which the base stimulus σ was fixed at 8, and a horizontal set, in which the comparison stimulus σ was fixed at 8. For both sets, NSD values ranged from -0.25 to 0.25. Stimuli pairs from both sets were randomly presented during the WM sessions.

Experimental design: categorization (RM) task

The task was designed in the way described for rats, save for the specific values of the stimulus set. The stimulus set employed the entire dynamical range established in the human version of the WM task, and was comprised of 9 linearly and equally spaced stimuli. As for rats, the 5^{th} stimulus lied on the implicit boundary between the strong and weak category, and was rewarded randomly 50% of the time.

Surgery protocol

To implant electrodes into left PFC and left PPC, both contralateral to the right-side vibrissal stimulation (n = 5 rats), rats were anesthetized with Isoflurane through a plastic snout mask. The Isoflurane anesthetic was administered via a veterinary apparatus (V-1 Tabletop, VetEquip Inc.). The concentration of the agent was variable, based on the depth of anesthesia needed and subject differences. To prepare for the operation, the rat was anesthetized with Isoflurane at 2%–2.5% mixed with oxygen (O2 98%) as calculated according to the MAC value (Minimum Alveolar Concentration: 1.35% ± 0.10%). Atropine was administered (2 mg/kg) to maintain a sustained heart rhythm and avoid mucous secretions in the respiratory tract. Once immobile, a mask was positioned on the rat's snout to provide a constant administration of Isoflurane with concentration of 2.5% and a flow rate of 1 l/min. This established a state characterized by analgesia, muscle relaxation and loss of reflexes (absence of reaction to tail and hind paw pinches and corneal air puff) and was maintained across the entire surgical procedure (opening of skin and bone). Blood oxygen

saturation was at least 98% and cardiac frequency was 350 ± 25 BPM. Once the craniotomy was completed, provided that the rat continued to show no reflexes or signs of discomfort, Isoflurane concentration was decreased to about 1.5% in 1 l/min of oxygen to prevent gas accumulation and allow a more rapid postoperative awakening. This same flow rate was maintained during the electrophysiological recordings used to verify electrode positions. At the end of the surgical procedure the isoflurane vaporizer was turned off and the rat was allowed to recover while breathing 1 l/min of oxygen. Although all inhalant anesthetics are maintained with 1.2 to 1.5 times MAC, factors that affect MAC have to be considered during the maintenance of the general anesthesia. For this reason, during the entire procedure and until the rat completely recovered from the anesthesia, its body temperature, blood oxygenation, breathing frequency and heart rate were continuously monitored. The absence of reflexes to tail and paw pinches and corneal air puff was regularly tested. In addition to the signals of the intraoperative monitoring machines, it is essential to observe physiological signs such as the color of the mucous membranes and the pupil diameter. In case of bleeding, fluids were reintroduced by subcutaneous injection to ensure proper hydration. The target cortical regions were accessed by craniotomy, using X Atlas and standard stereotaxic technique. Dura mater was removed over the entire craniotomy with a small needle. The electrodes were configured as 2 arrays of 12 electrodes (custom built, 4 rows of 3, with 200m spacing in both directions). Electrodes were sharply tapered, and shaft diameter was 75µm. Electrode impedances were between 0.5 and 1 MΩ). The electrode arrays were inserted by slowly advancing a Narishige automatic micromanipulator. After inserting the array, the remaining exposed cortex was covered with biocompatible silicone (KwikSil; World Precision Instruments). In all rats, 5 small screws were fixed in the skull as support for dental cement. All screws served as ground and reference electrodes. Before the conclusion of the operation, animals were given antibiotic (Baytril; 5 mg/kg; i.p.) and analgesic (Rimadyl; 2.5 mg/kg; i.m.). After surgery, a local antibiotic (Isaderm) was applied around the wound to help the healing. In addition, both the antibiotic and the analgesic were delivered through a water bottle for 24 h after completion of surgery. During this recovery time, rats had unlimited access to water and food.

Behavioural Analysis

The analysis described in this section was performed in the same way for rats and humans. Only sessions with a percentage of correct choices above 70% were considered. All analyses were performed using custom functions and scripts in MATLAB.

Psychometrics

The choices performed on stimuli pairs belonging to psychometric sets were fitted to Equation 2 for WM and Equation 3 for RM,

$$(2)\ P(S2 > S1) = \gamma + (1 - \gamma - \lambda)\frac{1}{2}\mathrm{erf}\left(\frac{NSD - \mu}{\sqrt{2\sigma^2}}\right)$$

$$(3)\ P(Strong) = \gamma + (1 - \gamma - \lambda)\frac{1}{2}\mathrm{erf}\left(\frac{Speed - \mu}{\sqrt{2\sigma^2}}\right)$$

where γ and λ are the lower and upper asymptotes, respectively, μ is the inflection point and σ is the sensitivity. The fit was performed employing Maximum Likelihood Estimation method. The parameters were estimated via 100 bootstrapped iterations, on a subject-by-subject basis.

History curves

History curves were estimated by fitting a linear regression on the choices in the current trials using the stimuli intensities presented in past trials as predictors. The linear regression coefficients were estimated via 100 bootstrapped iterations, employing the Least Squares method via MATLAB function *fitlm*.

GLM

The choices performed on each trial were fitted to

$$(4)\ P_k(Response) = \gamma + (1 - \gamma - \lambda)\frac{1}{e^{-X_k}}$$

where γ and λ are the lower and upper asymptotes, respectively and k=0,1,2… denotes how many trials before the current trial (N) are considered. In WM Response stands for S2>S1, and in RM Response stands for Strong, as in the psychometric fit. When considering only stimulus history, we employed:

$$(5)\ X_k = \beta_0 + \sum_{n=N}^{N-k} \beta_n S_n$$

where the predictor for trial N is either the NSD of the presented pair (WM) or the intensity of the presented stimulus (RM); the predictors for all other trials are either the average of stimuli intensities in that trial (WM) or the stimulus intensity in that trial (RM). When considering both stimuli and choices, we employed:

$$(6)\ X_k = \beta_0 + \sum_{n=N}^{N-k} \beta_n S_n + \sum_{n=N}^{N-k} (R_n C_n^R + U_n C_n^U)$$

where C_n^R represents rewarded choices trial n and could assume values of 1 (S2>S1 in WM, Strong in RM), -1 (S1>S2 in WM, Weak in RM) or 0 (choice unrewarded); C_n^U represents unrewarded choices trial n and could assume values of 1 (S2>S1 in WM, Strong in RM), -1 (S1>S2 in WM, Weak in RM) or 0 (choice rewarded); for n=N, both were always set to 0. All predictors were z-scored to make coefficients directly comparable.

The GLM analysis was repeated with increasing values of k, and the significance of coefficients for trial N-k was assessed by AIC, computed at each k^{th} step. All coefficients and AIC were estimated via 100 bootstrapped iterations, on a subject-by-subject basis.

Neuronal recording preprocessing

Neuronal recordings were spike-sorted manually via Wave Clus (Chaure et al., 2018). Single units were selected according to spike shape consistency and inter-spike interval. Two criteria had to be met for a unit to be included in the analyses. First, overall firing rate within the session had to be at least 2 Hz. Second, overall firing rate per trial for the entire session, with many different stimulus conditions intermixed across trials, did not show a significant non-zero linear correlation ($p < 0.05$) over time. The second criterion was aimed at excluding unstable recordings.

Firing rate

Firing rate was computed by binning spikes in non-overlapping 50ms time-bins, then applying a gaussian smoothing over a 500ms window with a standard deviation of 100ms

Decoding

Each recording session yielded a neuronal ensemble of simultaneously recorded units, which were analysed together. Decoding was performed by Support Vector Machine analysis, independently at each time-bin. Decoder accuracy was evaluated by partitioning the dataset according to repeated hold out cross-validation (80% training, 20% testing) for 20 iterations. For the shuffled accuracy, 20 iterations of trial labels shuffling were performed. Decoder performance was then assessed, similarly to (Barbosa & Stein, 2020), as

$$(7)\ Distance(t) = \frac{Accuracy_{Real}(t) - Accuracy_{Shuffle}(t)}{\sigma_{Shuffle}(t)}$$

where $Accuracy_{Real}$ is the average accuracy over all iterations of cross-validation for correctly ordered labels, $Accuracy_{Shuffle}$ is the average accuracy over iterations for shuffled labels, and $\sigma_{Shuffle}$ is the standard deviation over all iterations for shuffled labels.

CHAPTER 2

UNIFIED MODEL FOR TACTILE STIMULUS HISTORY IN RATS AND HUMANS

INTRODUCTION

Behavioural modelling has been a productive field of study in the last decades. These models often start as purely phenomenological, such as the reinforcement learning model (Kaelbling et al., 1995). For example, a particularly successful model, the drift diffusion model (Pedersen et al., 2017; Tavares et al., 2017) describes how evidence is integrated over time to reach a decision about a noisy stimulus.

Many theoretical models have been proven successful in predicting behaviour, and have been integrated with other descriptions of the statistical brain (Bitzer et al., 2014; Shahar et al., 2019). Even if many models dealt with abstracting behaviour through a mathematical description, their theoretical underpinnings have often provided insight into the physiological mechanisms driving behaviour. In the case of the drift diffusion model, for example, recent publications (Nunez et al., 2017; Steinemann et al., 2023) showed how the decisional variables' time course predicted by the model correlated with the activity of neuronal populations, recorded during the task.

Despite this, even very successful models are seldom extended beyond the experimental scope they are first designed for.

First presented in Hachen et al., 2021, my colleagues devised a phenomenological model, describing how, in the RM task, the history-dependent decision criterion shift they observed can be explained with a single-parameter, time-continuous model.

In the present work, we will first recount that model, and explain how we extended it to include the stimulus history effects observed in WM, in both rats and humans. We simulated

a subject acting according to the model's description, aiming to replicate the behavioural features of stimulus history observed in WM and RM. We also fitted the extended model to the behavioural data discussed in Chapter 1, and tested it against competing models in predicting the real subject's behaviour.

A neural networks model for WM and RM has also been developed in collaboration with Francesca Schönsberg, currently submitted to Cosyne 2024. This approach is different as it focuses on synaptic plasticity dynamics, and extracts the task responses without performing any explicit training on the network. Though the PhD candidate is a co-author of this second modelling project, it is not presented as part of the thesis. One long-term goal of ours is to unify the synaptic plasticity approach with the work presented below, which can be considered as belonging to mean-field approaches.

RESULTS

Model description

First, we will describe the model we further developed, starting from the description in Hachen et al., 2021, and show how tuning model parameters can reproduce the effects of stimulus history described in the behavioural results, by simulating subjects performing the two tasks under the model's assumptions.

The model structure and flow are depicted in Figure 20-A: a stimulus is delivered (either to whiskers or fingertips) and its neuronal representation is relayed to primary somatosensory cortex (S1), where the representation becomes part of a conscious percept (Diamond & Toso, 2023). As the percept is formed, it leaves a memory trace into the Short-term Buffer (STB), which acts as a temporary storage. While the trace is held in the STB, it starts to be attracted toward the representation held in the Long-term Buffer (LTB). The LTB representation, we will argue, may be approximated a sensory prior. At the same time, the sensory prior contained in the LTB is attracted, reciprocally, to the trace in the STB: each new percept provides a dynamical update to the LTB content; in this way, the LTB content is similar to a running average of past stimuli, giving more weight to recent stimuli over stimuli in the more distant past. Their reciprocal attraction is described by the Equations 8 and 9, whose solution is described in the Methods section. The model is characterized by two time constants, τ_{Long} and τ_{Short}, which describe the reciprocal decay of the buffers towards each other.

Figure 21-A shows the buffers' behaviour for two equal, or symmetrical, τ values: regardless of the two buffers' starting point at time zero, their content races converge asymptotically to the average of the two starting values, and do so with a symmetrical decay dictated by the (de facto only) time constant.

On the other hand, asymmetrical τ values lead to a qualitatively different dynamic (Figure 21-B and C): the larger τ_{Long} becomes compared to τ_{Short}, the closer the asymptote moves towards the LTB starting value, and the more similar the effective time constant becomes to τ_{Short}, as described by Equation 13. As the equations are symmetrical, the reverse is true if τ_{Short} becomes larger than τ_{Long}.

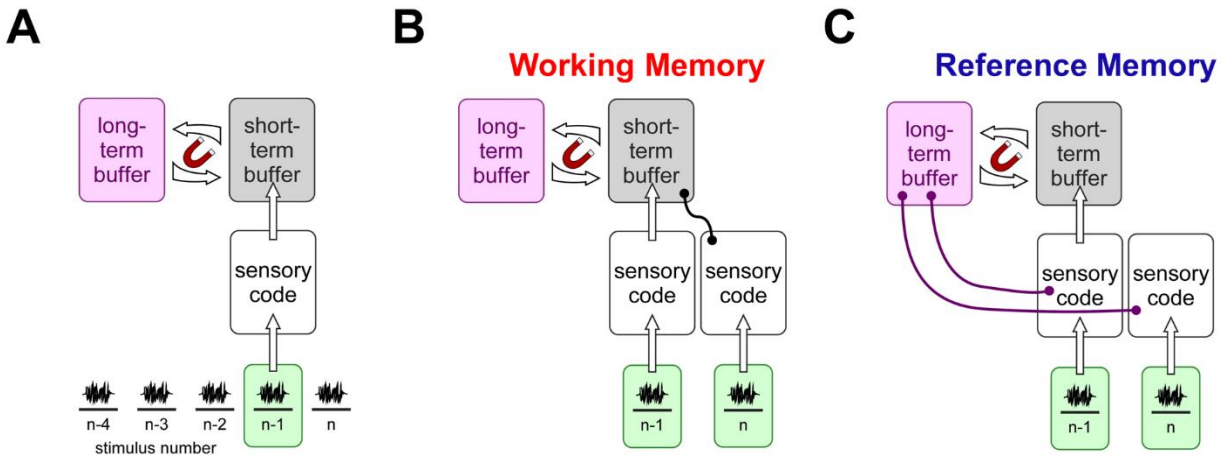

Figure 20: Schematic description of the model. (A) General description. The short-term buffer (STB) is loaded with the integrated percept of the latest delivered stimulus. Then, the STB and the long-term buffer (LTB) begin attracting each other, until a new stimulus is delivered. Then, the STB content is deleted and is ready to be loaded with the percept of the newly delivered stimulus. **(B)** Model applied to WM. The black line connects the STB to the comparison's percept, which are then compared to solve the task. **(C)** Model applied to RM. The purple line connects the LTB to the stimulus percept, which are then compared to solve the task.

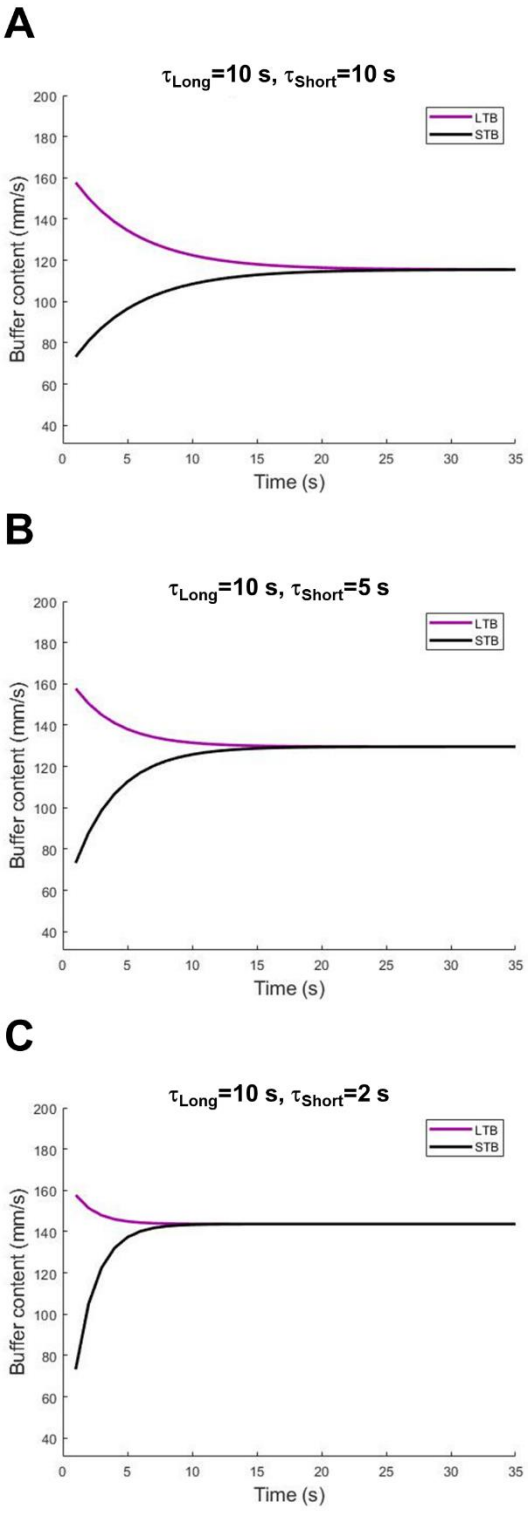

Figure 21: Short-term and long-term buffers dynamics. (A) Buffers' attraction for symmetrical time constants. x axis is time in seconds, y axis is buffer speed content. **(B)** Same as A, for asymmetrical time constants. **(C)** Same as A, for asymmetrical time constants.

The model describes how memory traces and the prior evolve over time according to the stimuli presented within the behavioural context. The STB and LTB can then be employed to describe stimulus history effects emerging in the paradigms described in the behavioural

section. In WM (Figure 20-B), the STB holds the base stimulus trace during the delay, acting as the sensory short-term memory. During the ISI, the percept trace held in the STB will be attracted to the LTB content, providing the basis for contraction bias. As the comparison stimulus is delivered and integrated into a percept, it is compared to the trace held in the STB, to make the judgment; afterwards, the comparison's percept is loaded into the STB, overwriting the previous trace, and attracts the LTB during the inter-trial interval (ITI).

Figure 22-A depicts a simulation of LTB and STB representations over some WM trials. A stimulus pair negatively affected by contraction bias is outlined in yellow; during the ISI, the STB trace (black) approaches the LTB trace (purple), moving closer to the intensity of its comparison stimulus (second black dot) and then becoming even stronger (black diamond) at comparison delivery: thus, the pair would be likely to be misjudged. The pair outlined in green, conversely, is positively affected: as the STB trace is attracted by the LTB, it is pulled further away from its comparison, magnifying the difference between them, and making the judgment easier or more likely to be correct. Last, outlined in blue, a pair in which the base stimulus is already so close to the LTB that its trace is stabilized by the buffers' interaction: thus, the base stimulus is remembered very faithfully.

In RM (Figure 20-C), each stimulus percept is compared with the LTB, which acts as the internal boundary between the weak and strong category. Then, as each percept is loaded in the STB, it attracts the LTB during the ITI, thus slowly shifting the categorical boundary and potentially leading some stimuli to shift from one category to the next. Figure 22-B depicts a simulation of LTB and STB during a RM session: in each trial, contained within two grey lines, a single stimulus is delivered, and must be categorised as either weaker or stronger than the boundary, represented by the horizontal dashed line. According to the model's description, however, each stimulus is compared, instead, to the LTB at the time of stimulus delivery. As the LTB wanders around the centre of the stimulus distribution, which is the same as the boundary, by design, most trials are categorised correctly, regardless. Sometimes, as in the trial highlighted in yellow, the LTB might wander far enough from the boundary, leading the subject to mis-categorise it. As described above, the longer τ_{Long} is, the less the LTB strays from the centre of the distribution, making this occurrence rarer.

Note that, qualitatively, in this model the STB exerts a progressive influence on the LTB through the attraction dynamics. If this attraction occurs during the intertrial interval, then the n-1 history effect – the effect of the STB on the category boundary stored in the LTB – will

increase as the ITI increases. Exactly this dynamic was shown, experimentally, in Hachen et al., 2021, lending experimental confirmation to this aspect of the model.

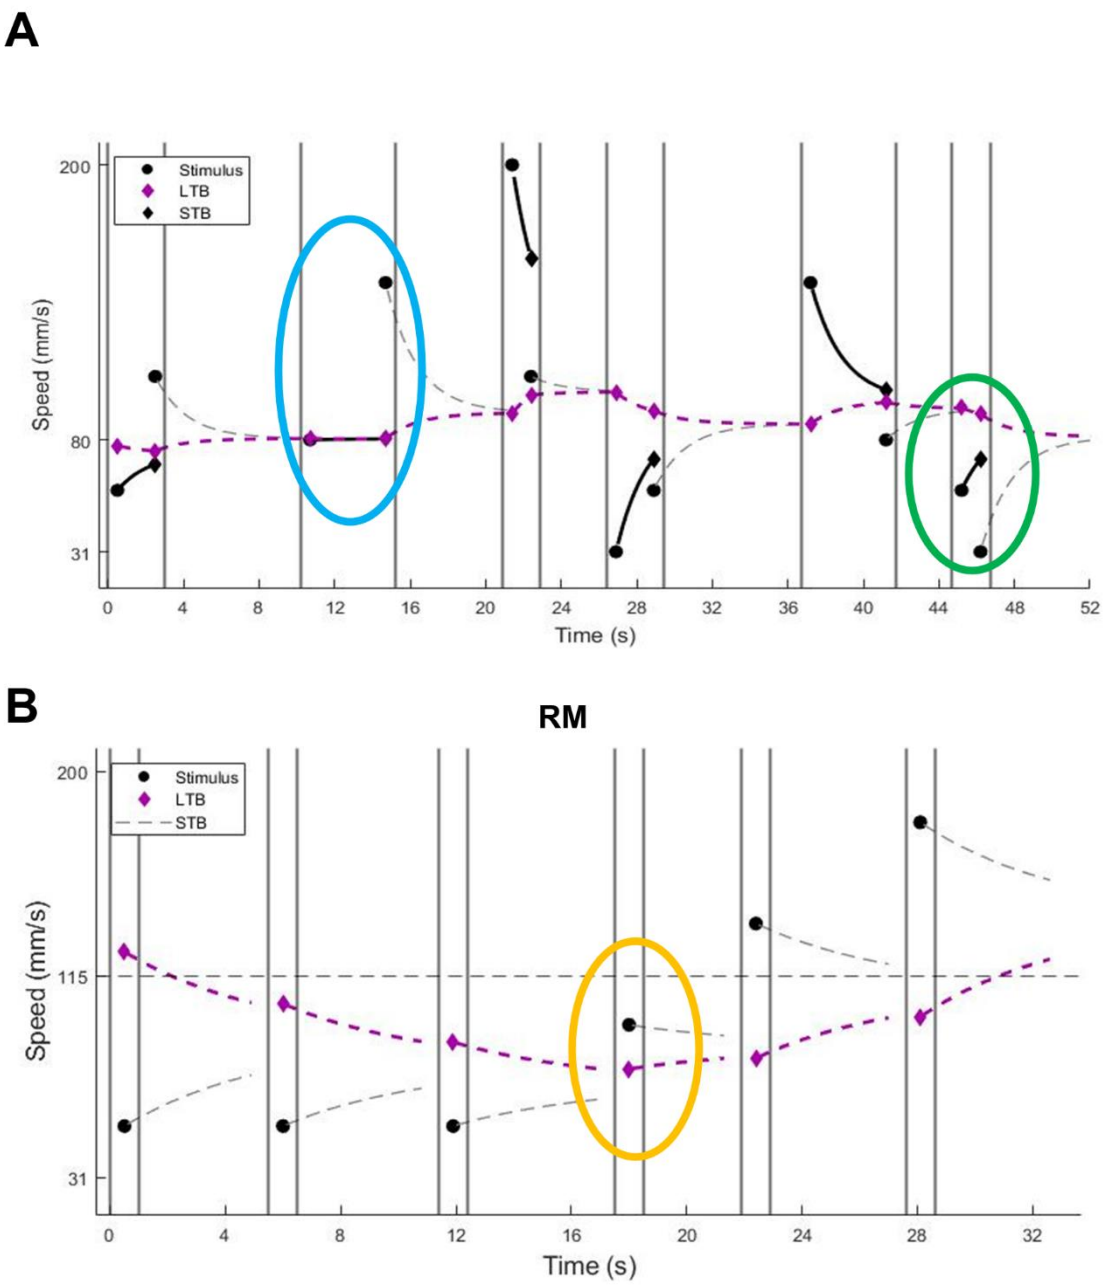

Figure 22: Buffers' timeline simulation. (A) Buffers' timeline simulation for WM. X axis is time in seconds, y axis is stimulus speed. Each trial is contained within vertical grey lines. Black dots are stimuli delivered; diamonds are buffers' contents at stimulus delivery. Solid black line is base stimulus trace in STB, dashed grey line is comparison stimulus trace in STB, dashed purple line is LTB. **(B)** Same as A, for RM.

RM simulation

To show how the model can reproduce the stimulus history effects described in the first chapter, we simulated subjects employing Equations 2 and 3 as generative models, for RM

and WM, respectively. As described in Methods, these subjects were unbiased and lapse-less, meaning that they always acted upon the decision variable (i.e.: NSD) and never chose randomly; employing the different time constants combinations, we produced different base stimulus traces at decision time to adjust the effective NSD, when simulating WM, and different internal boundaries for RM. Then, we computed history curves and fitted GLM and psychometric models on the simulated choices, in the same manner employed on the behavioural data, and analysed how different time constants affected GLM and psychometric parameters, and if our model could reproduce stimulus history effects observed in real subjects.

Figure 23 shows the simulation results applied to RM. To be as close to the description in Hachen et al., 2021, we first employed the symmetrical version of the model, where $\tau_{Long} = \tau_{Short}$. Figure 23-B shows the effect of the model on the psychometric curve computed over all trials together: μ, the inflection point and measure for bias, is unaffected by τ_{Long}, as the shift of the internal boundary in one direction on some trials, due to recent stimuli, is compensated by other trials in which the shift happened in the opposite direction; σ, the slope and measure for sensitivity, on the other hand, is negatively impacted: as τ_{Long} decreases, the sensitivity decreases as well, as the choice on each trial becomes more and more impacted by recent stimuli. Figure 24-A and B show how the same parameters, μ and σ, are affected by τ_{Long} when computing psychometrics conditioned on the stimulus delivered in trial n-1. Here, μ was more strongly affected by stimulus n-1 for shorter τ_{Long}, as the subject became more repulsed by stimulus n-1: with weaker stimuli n-1, μ would reduce, shifting the psychometric curve towards the left and increasing the overall probability of judging stimulus n as strong. On the other hand, σ was largely unaffected by stimulus n-1, depending only on τ_{Long}. Subject performance (Figure 23-C) also increase with τ_{Long}, since the negative effects described reduced as the LTB became more stable.

Figure 24-C shows stimulus history linear regression coefficients, on the left, analogous to those in Figure 11-C. The coefficients for n-1 are larger for shorter τ_{Long}, as the LTB is updated faster by the latest stimulus received. At the same time, the coefficients decay faster towards 0 for shorter τ_{Long}, as the LTB also “forgets” faster: for longer τ_{Long}, stimuli from trials further in the past still exert an effect during trial n. In Figure 24-D, the distribution of LTB values during the simulated sessions shows how longer τ_{Long} stabilises the LTB around the

centre of the stimulus range, producing a narrow distribution which reflects the smaller effect of recent stimuli, compared to shorter τ_{Long}.

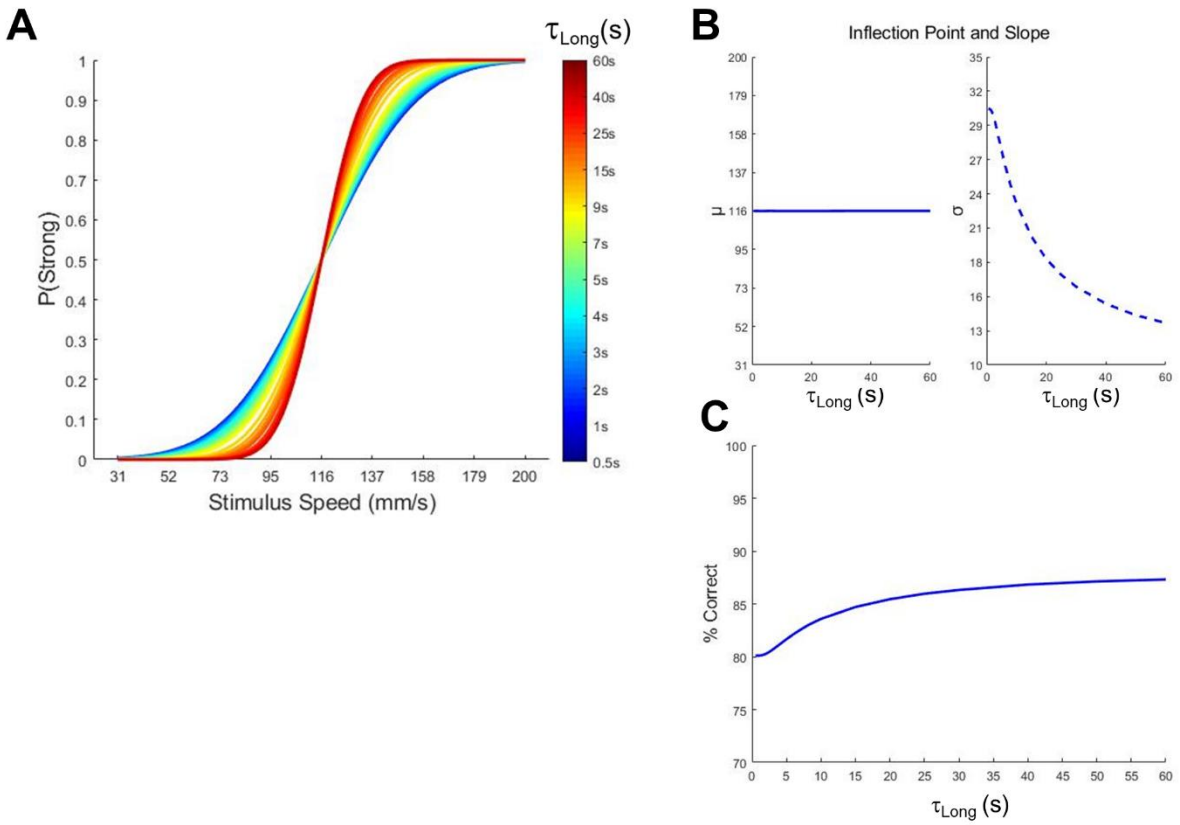

Figure 23: Psychometrics simulation, RM. (A) Psychometric curves from simulated subjects with different τ_{Long}. x axis is stimulus nominal speed, y axis is probability of judging stimulus as strong. **(B)** Inflection point (μ) and reciprocal of slope at inflection point (σ), left and right respectively. x axis is τ_{Long}. **(C)** Subject performance. x axis is τ_{Long}, y axis is percentage of correct responses.

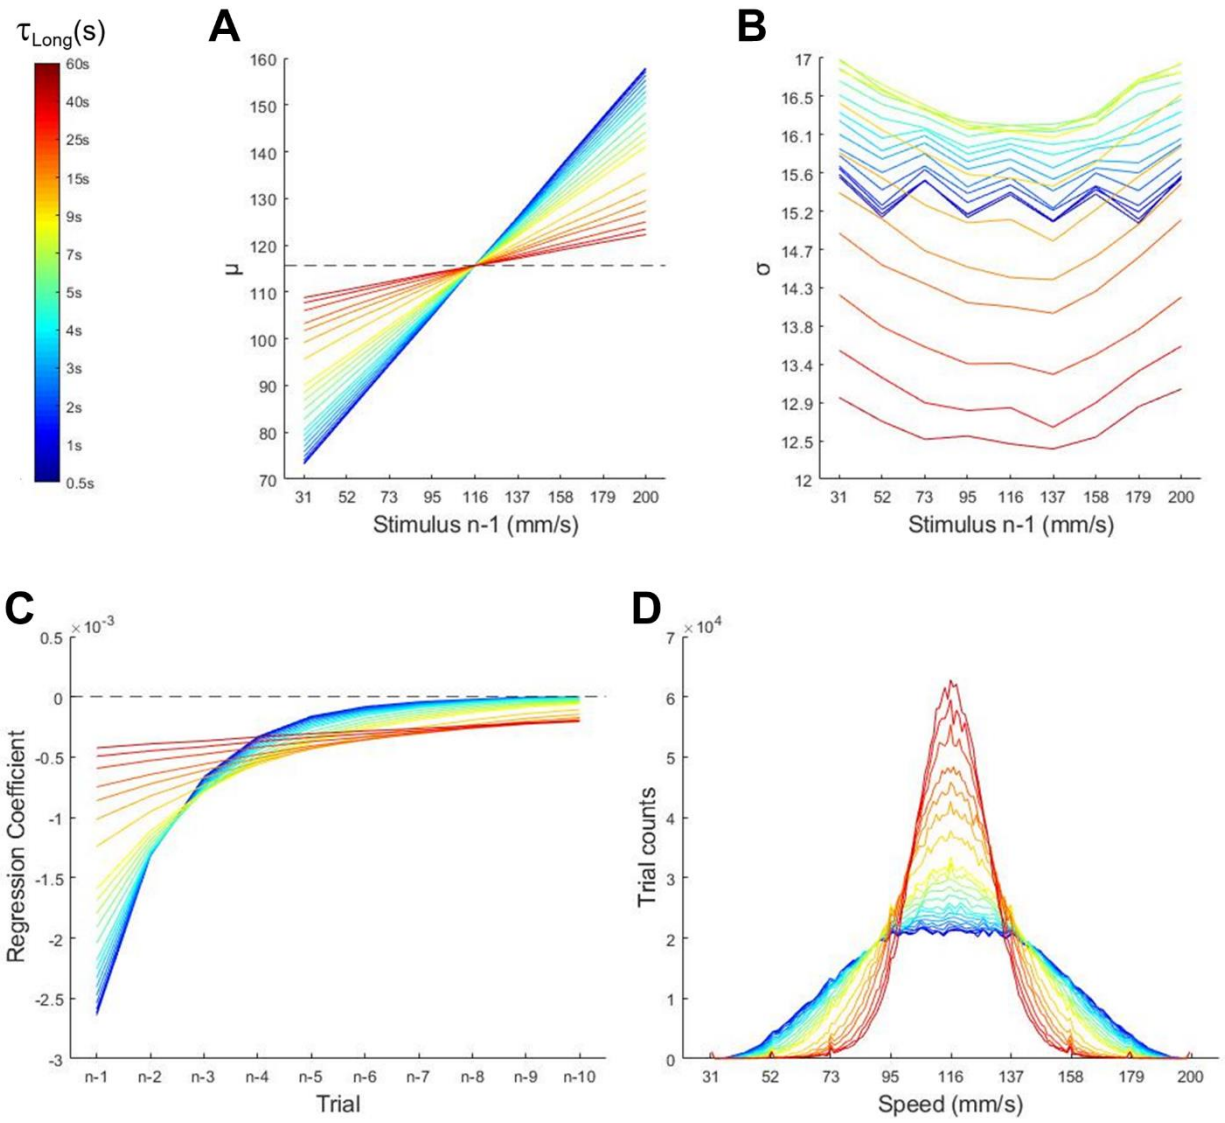

Figure 24: Stimulus history simulation, RM. (A) Inflection point (μ) for psychometric curves separated according to stimulus n-1, as in figure 7-B. x axis is stimulus n-1 nominal speed, y axis is μ. **(B)** Same as A, for σ. **(C)** Regression slope coefficients for past trials, as in figure 8-C. **(D)** LTB content distribution, for different τ_{Long}. x axis is buffer content, y axis is trial count.

WM simulation

For WM, we employed various combinations of time constants, all characterised by $\tau_{Short} <= \tau_{Long}$, as we expected the STB to be more volatile than the LTB. The justification for expecting more volatility is that the STB is reloaded and, ideally, wiped clean of memory on each trial. The functional attributes of fast reloading and scrubbing struck us, intuitively, as being properties associated with higher volatility.

Figure 25 shows the parameters recovered for the psychometric model, for various combinations of time constants values: the lapse parameters (Figure 25-D) increase rapidly

with τ_{Long} before reaching a plateau, which is higher at shorter τ_{Short}; for τ_{Short} above 5s they remain close to zero (which is the lapses values set for the generative model) for any value of τ_{Long}. The most dramatic increase is in the upper asymptote, λ, meaning that pairs for which NSD is positive (i.e.: comparison is stronger than base) are more affected: this is due to how the stimuli pairs are arranged. In the SGM, stimuli intensities are chosen to maintain NSD constant, either positive or negative, for each pair: so, the stimuli intensities become spaced exponentially, as described in Methods. In this way, while the NSD is constant, the absolute difference between stimuli in a pair is not: it is smaller for weaker pairs, and larger for stronger pairs. Thus, when the STB holds weaker stimuli traces, it needs to shift less towards the LTB to exert a noticeable effect on the NSD, since the absolute difference is smaller, compared to when the STB holds stronger stimuli traces.

The psychometric inflection point (μ) and slope (σ) are both affected as well (Figure 25-C): as the time constants increase, the slope increases (smaller σ) until it approaches the generative model value (0.1). For τ_{Short} below 5s, the slope even increases beyond the generative value: however, as already shown, this is coupled with larger lapses and a larger prediction error, so the increased slope should not be taken at face value. The effects of short time constants are overall decreasing performance, as shown in Figure 25-B, so that is not surprising: a lower slope means lower acuity, which in this case comes from misremembering the base stimulus, thus altering the decision variable NSD. Interestingly, though, even the worst possible combination, performance-wise, leads to an overall performance still above chance, meaning that, with the stimuli pair set employed, there might not be a very strong incentive to improve base stimulus retention even for subjects with a very unstable STB. As for μ, it stabilizes on slightly positive values: this reflects the fact that the centre of the stimuli distribution is slightly stronger than the median of all stimuli, due to the fact that stimuli are exponentially spaced; in this way, the STB base stimulus trace is, on average, attracted to stronger values: as a consequence, the comparison stimulus needs to be slightly stronger than the base nominal values to achieve an effective NSD that is equal to zero, thus leading to fitting positive values for μ.

Figure 26, left panels, shows regression coefficients for past trials, computed on the stimuli average for that trial, as shown previously in the behavioural results (Figure 4-C). Each panels shows curves computed from different values of τ_{Short}. As was the case in RM, Figure 24-C, shorter τ_{Short} produces larger coefficients for the most recent trials, which decay towards zero for trials further in the past. In the same way, longer τ_{Long} extend this effect

further in the past, as the LTB becomes more resistant to updating, both due to τ_{Long} itself, which stabilises the LTB value, and the asymmetrical dynamics shown in Figure 21-B and C, making the LTB asymptote closer to its initial value the shorter τ_{Short} is.

Figure 26, right panels, shows the distribution for STB and LTB, computed for the same τ_{Short} values employed for the regression coefficients. As in RM, longer τ_{Long} produces narrower distributions for the LTB, and for the STB too. However, longer τ_{short} produces STB distributions with clearer peaks centred on the values employed for the base stimuli, showing how the STB retains a more stable memory trace.

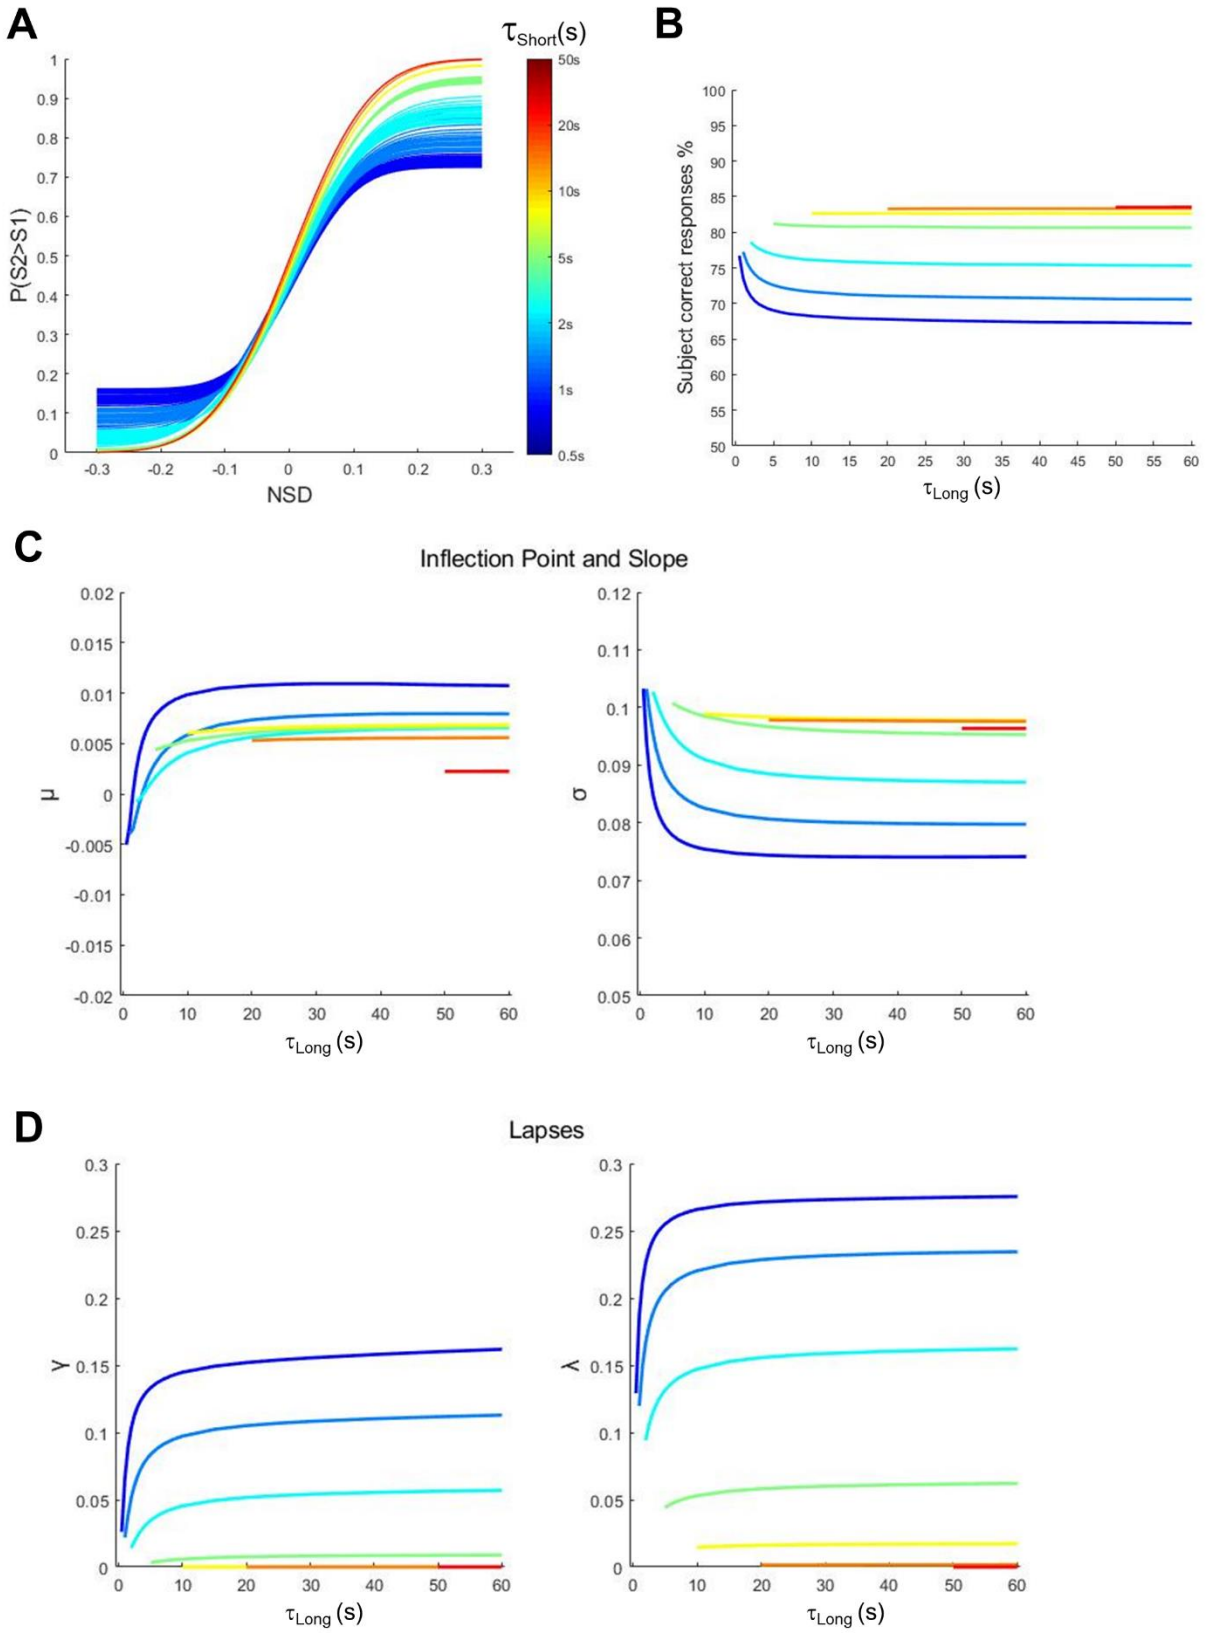

Figure 25: Psychometrics simulation, WM. (A) Psychometric curves from simulated subjects with different τ_{Short}. x axis is NSD, y axis is probability of judging comparison as stronger than base. **(B)** Subject performance as a function of time constants. x axis is τ_{Long}, y axis is percentage of correct responses. Colour denotes

different τ_{Short} **(C)** Same as figure 20-B, Colour denotes different τ_{Short}. **(D)** Same as C, for lower (γ) and upper (λ) lapses.

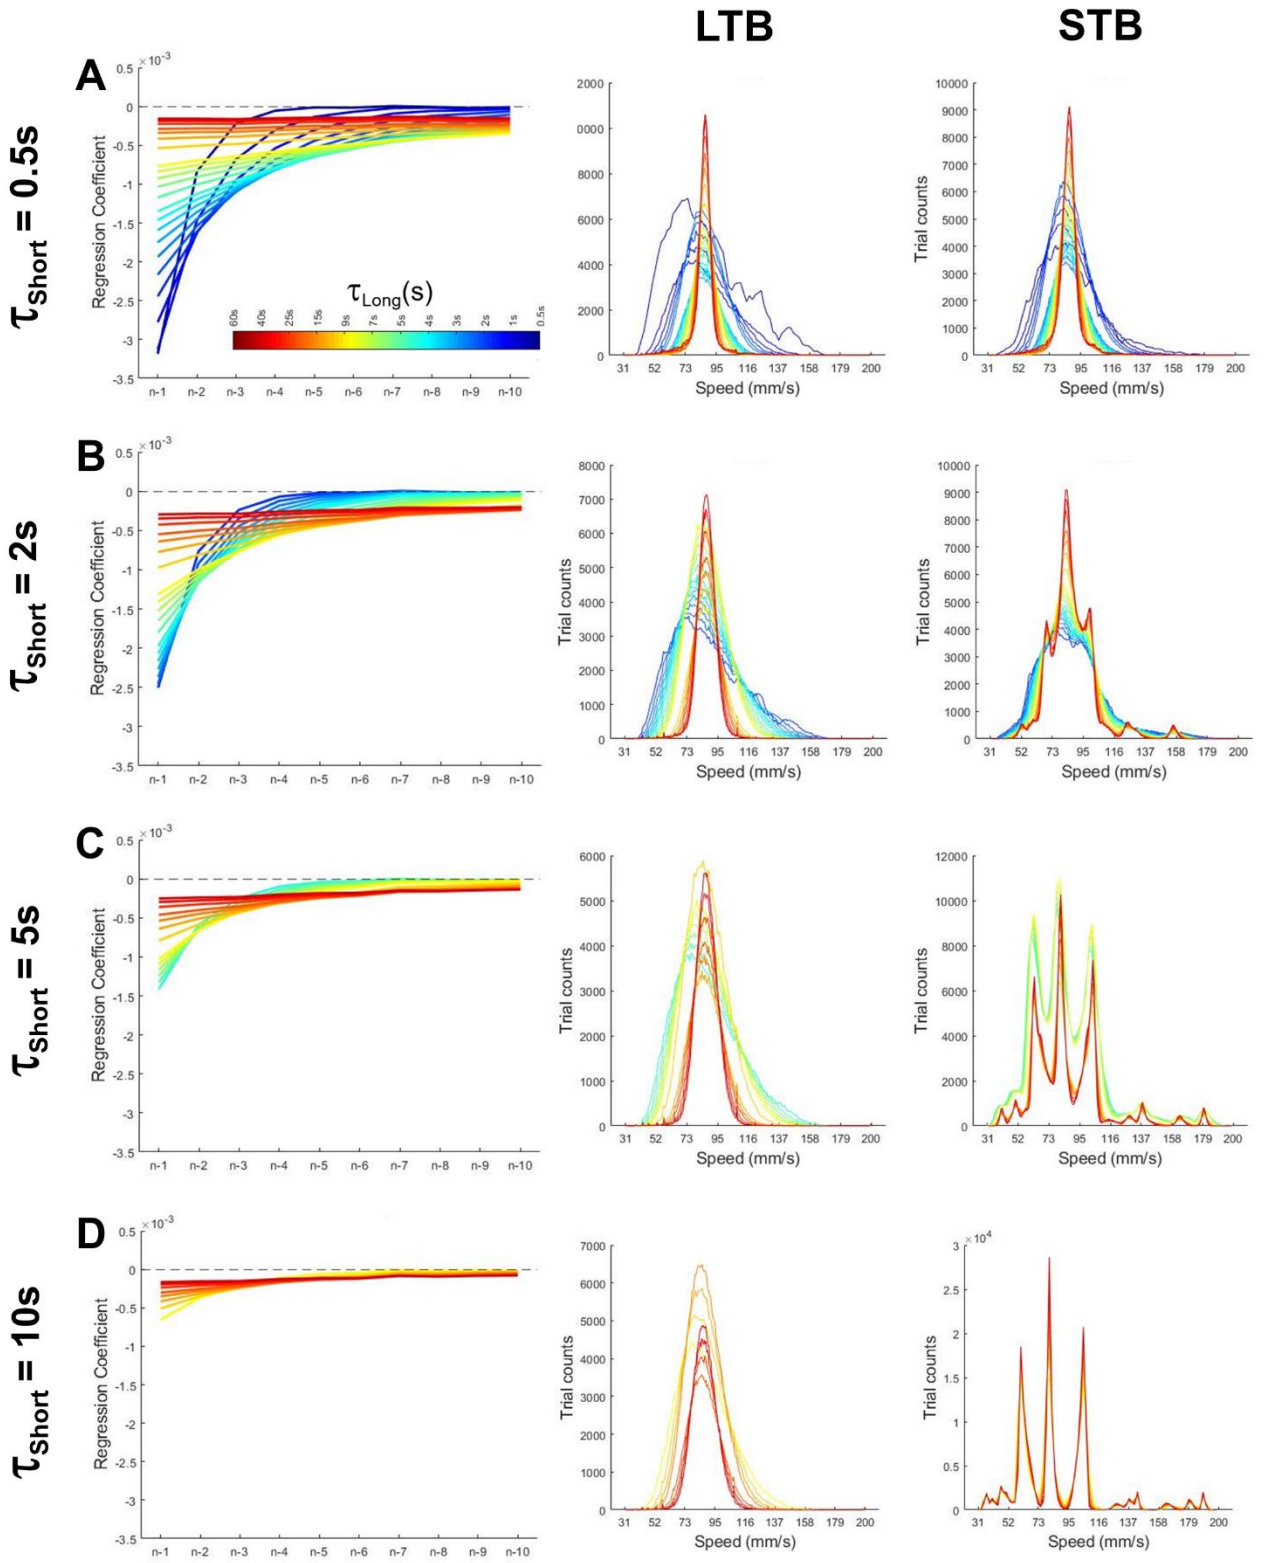

Figure 26: Stimulus history simulation, WM. (A) Left, regression slope coefficients, as in figure 2-C, for τ_{Short}=0.5s. Colour denotes τ_{Long}. Right, LTB and STB distributions for τ_{Short}=0.5s. x axis is buffer content, y axis is trial counts. Colour denotes τ_{Long}. **(B)** Same as A, for τ_{Short}=2s. **(C)** Same as A, for τ_{Short}=5s. **(D)** Same as A, for τ_{Short}=10s.

Model fitting in WM

Having established how the model behaves when applied to an ideal, simulated subject, we proceeded to fit it to the collected data. We fitted our model, to which we will refer as T (Tau), for shortness, on a subject by subject basis, recovering the fitted time constants and the predicted choices trial by trial. We then compared how well this model predicted the reals subjects' choices against the psychometric model discussed in the behavioural section and a variety of GLM models with an increasing number of predictors. These GLM models include only past trials that have been shown to be predictive in the behavioural section. The S model includes only past stimuli for RM, or the average of stimuli for WM. The C model includes only past choices. The RW model includes only past choices, separate by correct and incorrect. The SC and SRW models are the combination of the S model with either of the other two. The TC and TRW models combine, via GLM, the decision variables produced by the T model, as described above, with choice n-1 or choice n-1 separated by correct and incorrect, respectively.

For WM, Figure 27 shows the Brier Score difference between the model with lowest error, identified with a red cross, and all other models, for rats (A) and humans (B). For both groups, the T model provided the lowest prediction error, on a par with SC and SRW for humans. The middle row of Figure 27, in the same task order, shows the ΔAIC for all models, against the model with lowest AIC, again identified by a red circle: the T model was confirmed as the most economic model, combining the best predictive power with the least number of parameters needed. The bottom row of Figure 27 shows the median of time constants recovered for each subject (E, rats and F, humans). All, save one human subjects, showed a shorter τ_{Short} compared to τ_{Long}, reflecting the fact that the LTB is less volatile than the STB, which in turn tends to decay faster towards the LTB. However, the time constants values recovered from rats are far shorter than human ones, reflecting the fact that rats exhibited a much stronger dependency on past stimuli than humans.

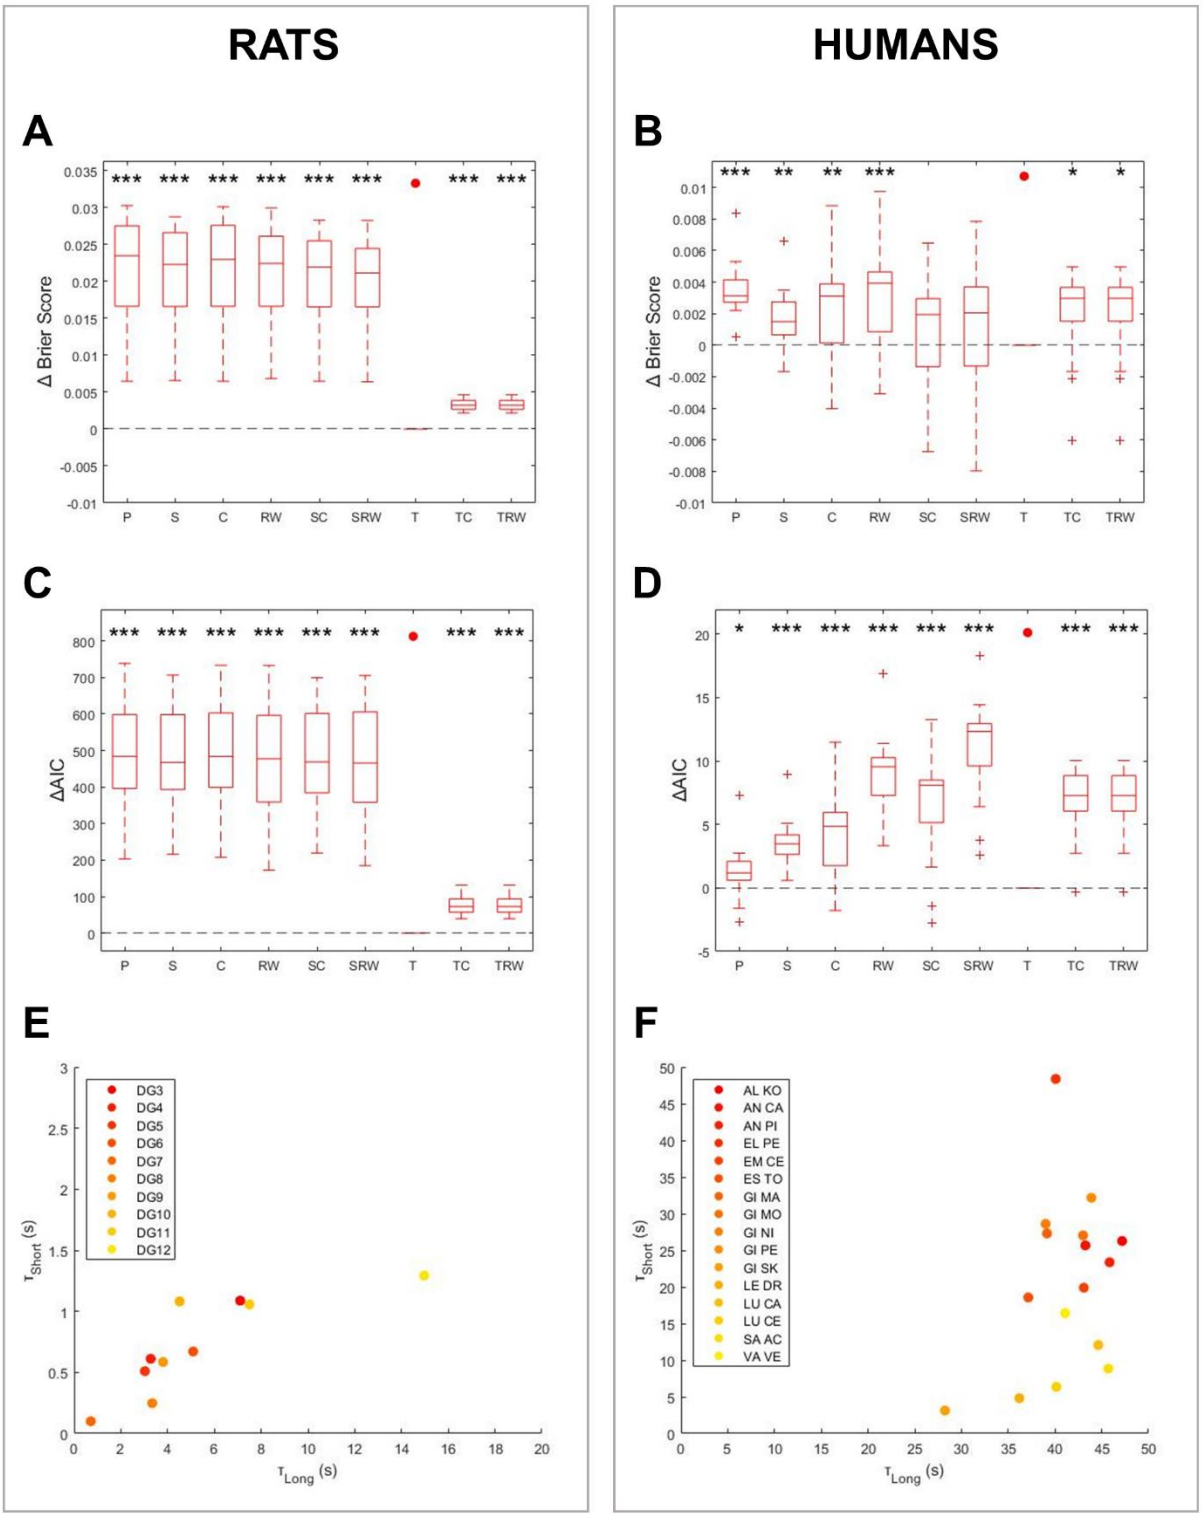

Figure 27: Model results for WM. (A) Brier Score difference for tested model, for rats, against lowest Brier Score model, marked by red circle. **(B)** Same as A, for humans. **(C)** AIC difference for tested model, for rats, against lowest AIC model, marked by red circle. **(D)** Same as C, for humans. **(E)** Time constants recovered, subject by subject, for rats. **(F)** Same as E, for humans. N=10 rats, N=16 humans (Student's t-test).

Model fitting in RM

For RM, Figure 28 shows the Brier Score difference between the model with lowest error, identified with a blue circle, and all other models, for rats (A) and humans (B). For rats, as described in the behavioural section (Figure 13-A, right), there was a more prevalent choice history effect, leading to the SRW model as the most predictive model and the T model and its variations as the least predictive. The stimulus history effect proved very weak in these subjects, overlayed with a strong choice attraction or repulsion towards previous correct or incorrect choices, respectively. For humans, the SC model provided the lowest prediction error; as described in the behavioural section (Figure 14-A, right), the stimulus history effect was overlayed with a strong choice attraction effect towards previous correct choices: as correct choices made up most of the trials, the contribution of previous incorrect choices was negligible, making them largely irrelevant for prediction purposes.

The bottom row of Figure 28, in the same group order, shows the ΔAIC for all models, against the model with lowest AIC, again identified by a blue circle: unsurprisingly, for both rats and humans the T model and its variation proved the models with largest AIC, as the subjects mostly did not conform to the model assumptions. For rats, the C model proved the most conservative model, on par with other GLM models, while the T model and its variations were significantly worse. For humans, the P model proved better than all other models, possibly in force of the human subjects performing closer to the design as intended, compared to rats.

Due to the unexpected results for the RM task, the model failed to capture the subjects' behaviour, contrarily to the success it had in a previous publication. It is my hypothesis that the choice to employ the same dynamical range established in the WM task, instead of tailoring it to RM, while reasoned, had the unforeseen consequence of inducing an unprecedented choice history effect, which effectively masked the smaller stimulus history effect. On top of that, it is possible that rats, being thoroughly trained in WM first, as it is the hardest task of the two to teach, had the unintended consequence of altering their strategy in RM, compared to naïve rats that are taught RM directly.

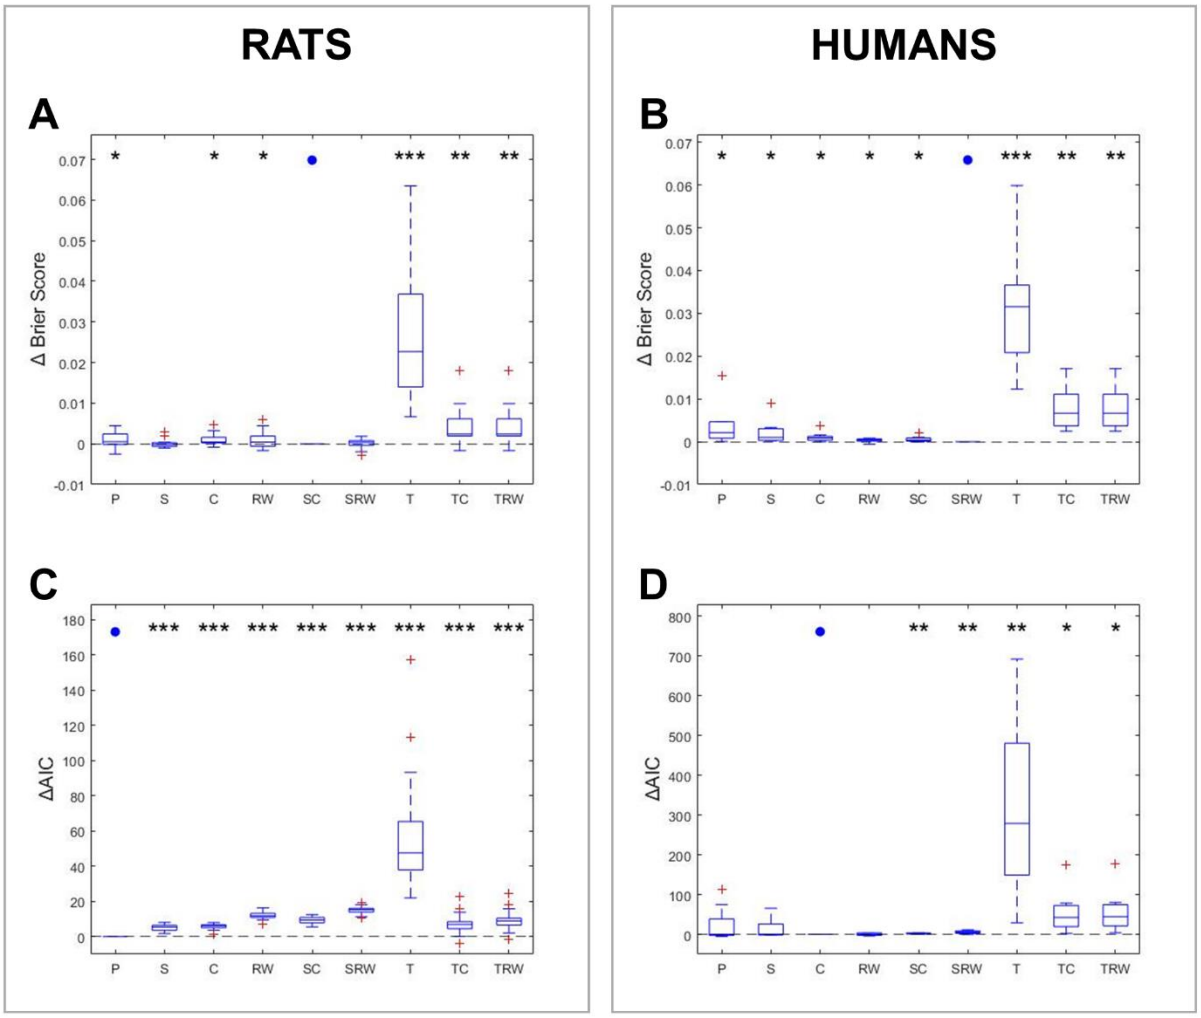

Figure 28: Model results for RM. (**A)** Brier Score difference for tested model, for rats, against lowest Brier Score model, marked by red circle. **(B)** Same as A, for humans. **(C)** AIC difference for tested model, for rats, against lowest AIC model, marked by red circle. **(D)** Same as C, for humans. **(E)** Time constants recovered, subject by subject, for rats. N=10 rats, N=16 humans (Student's t-test).

DISCUSSION

The simulations show how the extended version of the model in (Hachen et al., 2021) replicates the serial dependence observed in that work, for RM, and can adapt to account for both contraction bias and stimulus serial dependence. However, the psychometrics simulations, especially in WM, deserve some deeper discussion.

While the simulated subject always employed the same generative parameters (see Methods), regardless of time constants, the psychometric fitting (Figure 4 for RM, Figure 6 for WM) did not recover those parameters exactly. In fact, it led to much misestimation, especially in WM. For both RM and WM, σ was always overestimated. As a measure of sensitivity, σ is employed as a measure of the underlying response distribution: however, this simulation shows how it can be easily misestimated, without accounting for stimulus history.

But the WM lapses show the grossest misestimation. While the generative lapses were set to zero, the fitted lapses could be far from it, depending on the time constants combination. The lapses are classically interpreted as non-sensory mistakes: those could be due to momentary lack of attention or choices already made before stimuli presentation, for example. However, in this simulation, there are no non-sensory mistakes: the simulated subject always employs the sensory information to the best of their ability. In this sense, the simulated subject's lapses should be interpreted as history-induced, not non-sensory, lapses. If the presented model's underlying hypothesis hold, without awareness of the extent of contraction bias and stimulus history, the lapses would be interpreted as non-sensory, perhaps underestimating the subject's attention or understanding of the task's instructions.

That observation is valid only if the model described is adequately applicable. Figure 24 shows that both rats' and humans' WM stimulus history, as contraction bias and serial dependence, can be accounted for by the interaction between LTB and STB. Not only this model, once fitted, provided the best predictive power of both species' choices, but it did so with the least amount of parameters to achieve that explanatory power. The tau recovered for rats (Figure 24-E) are very short, compared to those of humans (Figure 24-F), especially with $\tau_{Short} < 2s$, which explains how rats are always strongly affected by contraction bias, with base stimulus memory traces decaying fast towards the LTB. At the same time, humans also showed values of τ_{Short} that are always shorter than τ_{Long}, except for one subject. This

reinforces the hypothesis that the information about the stimuli distribution, for which the LTB acts as a proxy, is actually more stable than the memory traces held consciously.

In RM, as was expected, the model failed to provide adequate explanatory power. As already discussed in Chapter 1, stimulus serial dependence effects were masked, in both rats and humans, by choice serial dependence, providing different results from (Hachen et al., 2021), from which the model originated. Possible explanations for this unexpected result have been discussed in Chapter 1.

Despite this, the novel extension on the previous model proved successful in capturing both contraction bias and stimulus serial dependence in WM. When applying the model to the task that most exhibited stimulus history effects, which was WM, it provided a successful fitting and captured the decision variables that the subjects employed on trial-by-trial basis, providing a dynamical description of perceptual memory. The model construction provides a phenomenological account of the reciprocal interaction between the LTB, a dynamically updated repository for the distribution of prior stimuli, and the STB, an unstable repository for working memory, which compensated for noise employing the information contained in the LTB. The model accounted for these phenomena employ only two time constants, providing time-continuous estimates of STB and LTB, ready to be employed by classical behavioural models.

Considering the results presented in (Hachen et al., 2021) and here, the interaction between LTB and STB provides adequate explanation to both stimulus serial dependence and contraction bias, in the context of two different experimental paradigms, requiring different cognitive resources, within the tactile modality. Moreover, the model was successfully applied to both rats and humans, providing evidence for conserved mechanisms underlying the aforementioned phenomena. Perhaps even more appealing, this model assumes a task-independent network of cognitive buffers, the read-out of which can be selected in a top-down, goal-directed manner: the LTB is read-out in the case of RM, the STB in the case of WM. However, both LTB and STB are present in both task designs, and interact in the same way. More research will be required to validate this model in different modalities and with different task designs.

METHODS

Model equations

We model the buffer interaction using the differential equation system:

$$\begin{cases} (8)\ \tau_{long}\dfrac{d}{dt}LTB(t) = STB(t) - LTB(t) \\ (9)\ \tau_{short}\dfrac{d}{dt}STB(t) = LTB(t) - STB(t) \end{cases}$$

which solves for LTB and STB as:

$$(10)\ LTB(t) = \frac{LTB_0 + STB_0}{2} - \frac{LTB_0 - STB_0}{2}\frac{\Delta_-}{\Delta_+} + \frac{LTB_0 - STB_0}{2}\left(\frac{\Delta_-}{\Delta_+} + 1\right)e^{-\Delta_+ t}$$

$$(11)\ STB(t) = \frac{LTB_0 + STB_0}{2} - \frac{LTB_0 - STB_0}{2}\frac{\Delta_-}{\Delta_+} + \frac{LTB_0 - STB_0}{2}\left(\frac{\Delta_-}{\Delta_+} - 1\right)e^{-\Delta_+ t}$$

Where LTB_0 and STB_0 are the respective values at t=0 and

$$(12)\ \Delta_- = \frac{1}{\tau_{Long}} - \frac{1}{\tau_{Short}}$$

$$(13)\ \Delta_+ = \frac{1}{\tau_{Long}} + \frac{1}{\tau_{Short}}$$

Simulations

STB and LTB content were computed continuously during simulated sessions for each task, with different τ_{Long} and τ_{Short} couples; their value at decision time were employed to produce simulated decisions. In WM, each trial NSD was corrected as:

$$(14)\ NSD_n = \frac{Comparison_n - STB_n}{Comparison_n + STB_n}$$

Which was then employed in equation 2. In RM, equation 3 was modified as

$$(15)\ P(Strong)_n = \gamma + (1 - \gamma - \lambda)\frac{1}{2}\mathrm{erf}\left(\frac{Speed_n - LTB_n}{\sqrt{2\sigma^2}}\right)$$

The simulations were performed with the same fixed generative parameters for both equations (μ=0, σ=0.1, γ and λ=0). The simulated choices were then fitted in the same way described in the behavioural methods to show how different time constants affect those models' parameters.

Model fitting

The computed STB and LTB were applied to psychometric equations 2 and 3 as described in the simulation results. The time constants τ_{Long} and τ_{Short} were fit via Maximum Likelihood Estimation method; the model's error and likelihood were estimated via Repeated Hold Out cross validation, using 80% of the dataset as training data and the remaining 20% as testing data. The model was fitted on a subject by subject basis for both tasks.

CONCLUSIONS

Throughout the present work, we examined how stimulus history affected judgment in two distinct perceptual memory decision-making paradigms: a categorisation (RM) and a delayed comparison (WM) task, involving vibrotactile stimuli. From the behavioural point of view, we examined two species, rats and humans, finding that they were affected by trial history in an analogous manner, exhibiting the same temporal dynamics, albeit with greater magnitude in rats.

The delayed comparison results aligned with those found in literature. On the other hand, the results in the categorisation task diverged from many works in the literature, due to choice serial dependence obfuscating stimulus serial dependence; in this condition, it was impossible to conclude whether the stimulus history effects, for each subject, correlated across tasks.

We presented an extended version of the behavioural model of the Tactile Perception and Learning Lab that succeeded in accounting for stimulus history in the delayed comparison task. Together with the original presentation of the model, we provide a single task-independent framework to describe the effects of stimulus history and account for subject variability in timescales. This framework proved successful both in rats and humans, providing evidence for a conserved physiological mechanism, despite minor species-specific differences.

The cortical activity recorded from rats showed how both PFC and PPC are part of a network responsible for coding task relevant information. Current, trial-relevant information, such as stimuli and choices are encoded along with trial history information. The encoding of the former, in PFC, correlated with subjects' attention to the task; the latter, despite not being relevant to solving the task, was proved to correlate with stimulus history in both PFC and PPC, during WM. During RM, the encoding was still present, despite no stimulus history was detectable in the subjects' behaviour. This result led us to hypothesise that a different, non-recorded region is part of the same network and acts as a gated system for the trial history information.

The results reinforce the idea of a single network of interconnected units, accounting for stimulus history in a task-independent manner, with just the read-out unit being selected in a top-down, goal-directed manner. This encourages the further study of a supra-modal perceptual network, independent of goal-directed attention and behavioural strategy.

In the future, we would extend both the phenomenological model and the extension of cortical areas involved. New task paradigms can be designed to further support the task-independency of the description presented here, and to possibly include a more organic choice history description.

SUPPLEMENTARY FIGURES

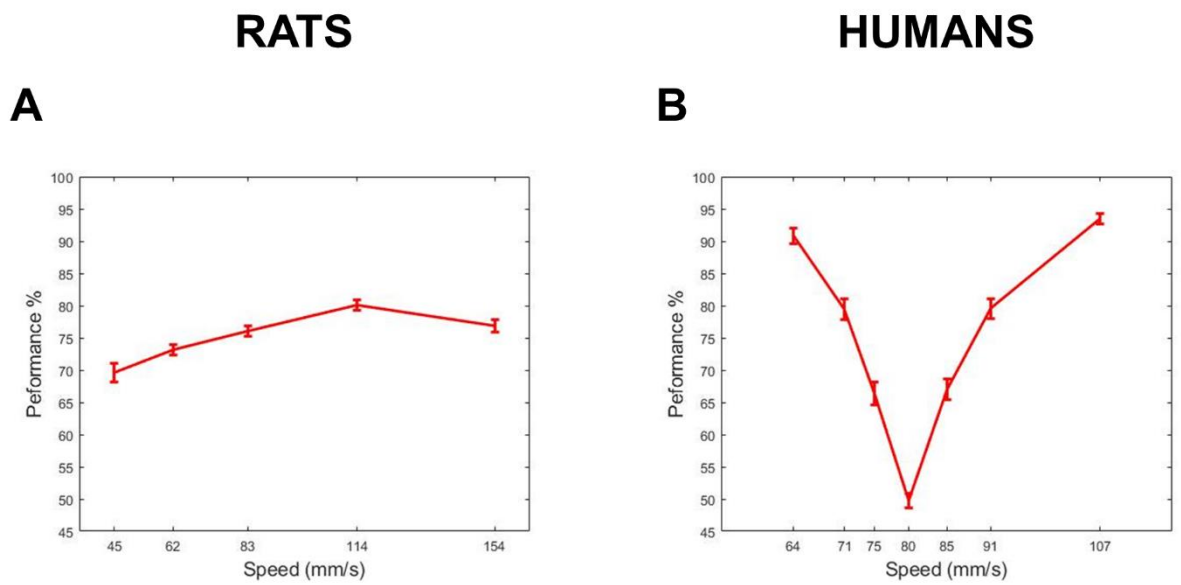

Supplementary figure 1: Performance as a function of stimulus intensity in WM. (A) Rats' performance on SGM pairs only, as a function of average stimulus intensity in each pair. x axis is average stimulus intensity in mm/s, y axis is performance as percentage of correct trials. Error bars are standard error of the mean. N=10. **(B)** Humans' performance, as a function of average stimulus intensity in each pair. x axis is average stimulus intensity in mm/s, y axis is performance as percentage of correct trials. Error bars are standard error of the mean. N=16.

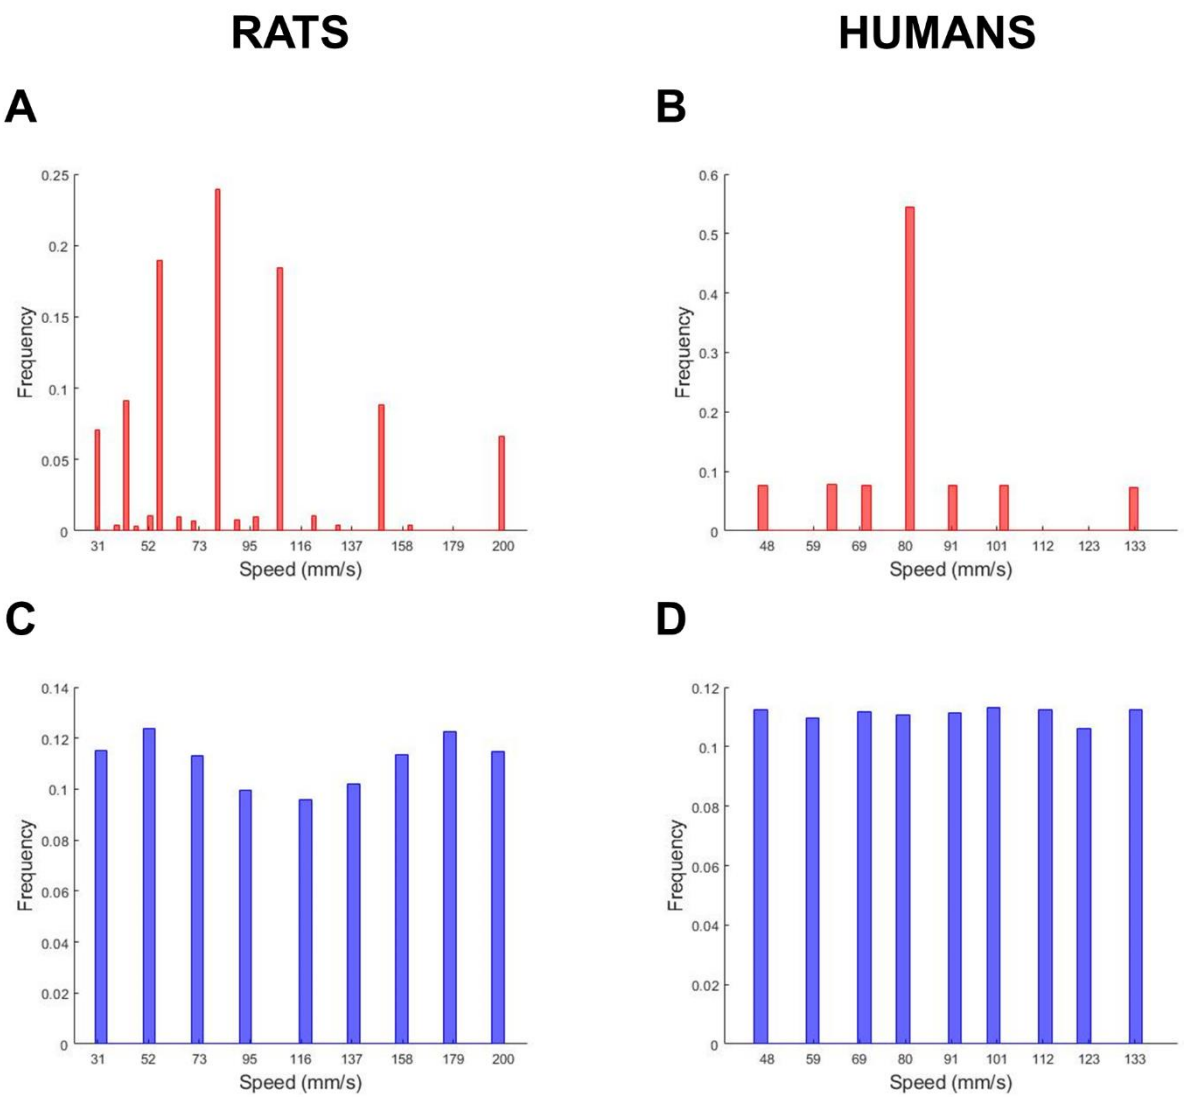

Supplementary figure 2: Stimulus distribution. (A) Stimulus distribution for WM task in rats. X axis is stimulus speed in mm/s, y axis is frequency. **(B)** As A, for humans. **(C)** As A, for RM task. **(D)** As C, for humans.

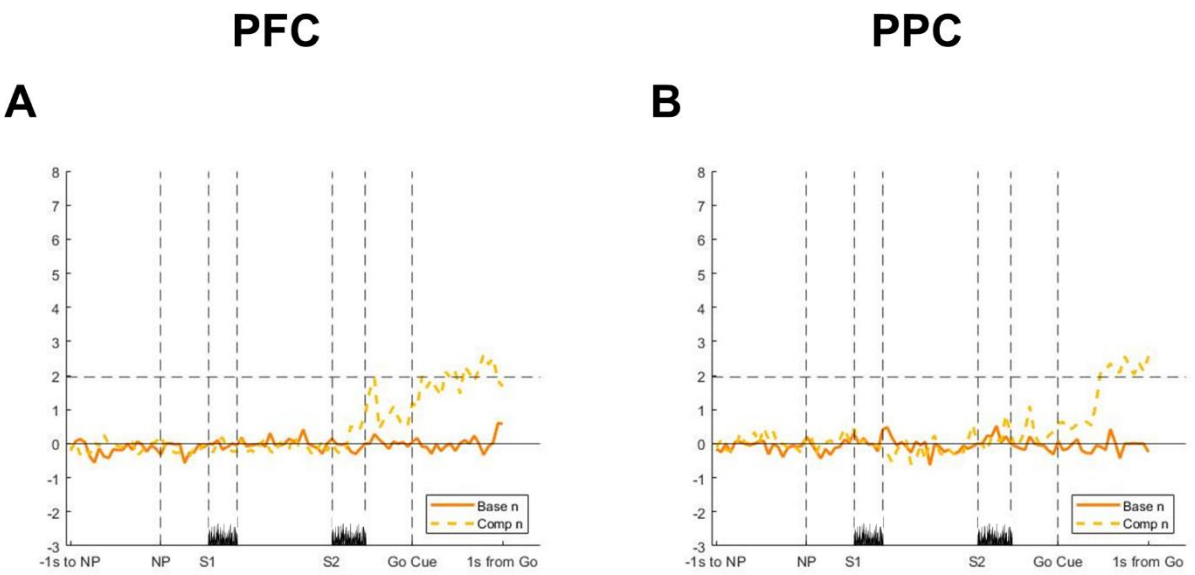

Supplementary figure 3: Base and comparison n decoding in WM. (A) Decoding, from PFC, for base stimulus n (solid) and comparison stimulus n (dashed). Curves represent median from high decoding sessions. x axis is timeline, from 1s before nose-poking to 1s after the go cue. y axis is distance of decoder accuracy from shuffled decoder accuracy (see methods). **(B)** Same as A, for PPC.

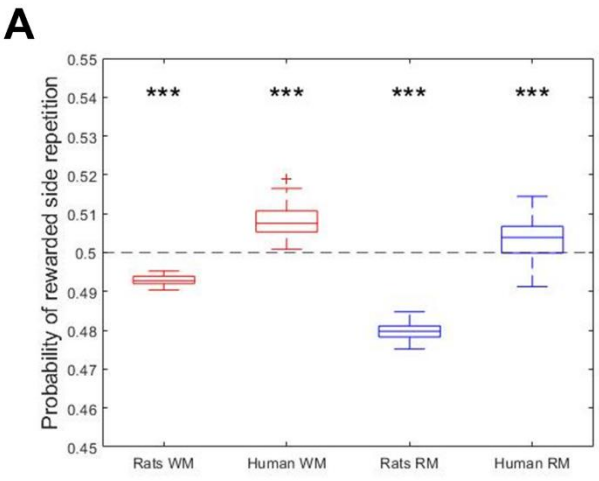

Supplementary figure 4: Rewarded side repetition. (A) Probability of the same rewarded side repeating from one trial to next, on y axis, for rats and humans and both tasks, on x axis. Dashed line marks 0.5 probability. N=10 rats, N=16 humans. (Student's t-test)

REFERENCES

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. https://doi.org/10.1038/NATURE25510

Ashourian, P., & Loewenstein, Y. (2011). Bayesian Inference Underlies the Contraction Bias in Delayed Comparison Tasks. *PLOS ONE*, *6*(5), e19551. https://doi.org/10.1371/JOURNAL.PONE.0019551

Barbosa, J., & Stein, H. ✉. (2020). Interplay between persistent activity and activity-silent dynamics in the prefrontal cortex underlies serial biases in working memory. *Nature Neuroscience*. https://doi.org/10.1038/s41593-020-0644-4

Barthas, F., & Kwan, A. C. (2017). Secondary Motor Cortex: Where 'Sensory' Meets 'Motor' in the Rodent Frontal Cortex. *Trends in Neurosciences*, *40*(3), 181–193. https://doi.org/10.1016/J.TINS.2016.11.006

Benozzo, D., Ferrucci, L., & Genovesio, A. (2023). Effects of contraction bias on the decision process in the macaque prefrontal cortex. *Cerebral Cortex*, *33*(6), 2958–2968. https://doi.org/10.1093/CERCOR/BHAC253

Bitzer, S., Park, H., Blankenburg, F., & Kiebel, S. J. (2014). Perceptual decision making: Drift-diffusion model is equivalent to a Bayesian model. *Frontiers in Human Neuroscience*, *8*(1 FEB). https://doi.org/10.3389/FNHUM.2014.00102/FULL

Bliss, D. P., Sun, J. J., & D'Esposito, M. (2017). Serial dependence is absent at the time of perception but increases in visual working memory. *Scientific Reports 2017 7:1*, *7*(1), 1–13. https://doi.org/10.1038/s41598-017-15199-7

Bosch, E., Fritsche, M., Ehinger, B. V., & de Lange, F. P. (2020). Opposite effects of choice history and evidence history resolve a paradox of sequential choice bias. *J. Vis.*, *20*(12), 9–9. https://doi.org/10.1167/jov.20.12.9

Brody, C. D., Hernández, A., Zainos, A., & Romo, R. (2003). Timing and Neural Encoding of Somatosensory Parametric Working Memory in Macaque Prefrontal Cortex. *Cerebral Cortex*, *13*(11), 1196–1207. https://doi.org/10.1093/CERCOR/BHG100

Bubic, A., Yves von Cramon, D., & Schubotz, R. I. (2010). Prediction, cognition and the brain. *Frontiers in Human Neuroscience*, *4*, 1094. https://doi.org/10.3389/FNHUM.2010.00025/BIBTEX

Cazettes, F., Mazzucato, L., Murakami, M., Morais, J. P., Augusto, E., Renart, A., & Mainen, Z. F. (2023). A reservoir of foraging decision variables in the mouse brain. *Nature Neuroscience 2023 26:5*, *26*(5), 840–849. https://doi.org/10.1038/s41593-023-01305-8

Chaure, F. J., Rey, H. G., & Quian Quiroga, R. (2018). A novel and fully automatic spike-sorting implementation with variable number of features. *Journal of Neurophysiology*, *120*(4), 1859–1871. https://doi.org/10.1152/JN.00339.2018

Cicchini, G. M., Mikellidou, K., & Burr, D. (2017). Serial dependencies act directly on perception. *Journal of Vision*, *17*(14), 6–6. https://doi.org/10.1167/17.14.6

Cicchini, G. M., Mikellidou, K., & Burr, D. C. (2018). The functional role of serial dependence. *Proceedings of the Royal Society B*, *285*(1890). https://doi.org/10.1098/RSPB.2018.1722

Diamond, M. E., & Toso, A. (2023). Tactile cognition in rodents. *Neuroscience & Biobehavioral Reviews*, *149*, 105161. https://doi.org/10.1016/J.NEUBIOREV.2023.105161

Fassihi, A., Akrami, A., Esmaeili, V., & Diamond, M. E. (2014). Tactile perception and working memory in rats and humans. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(6), 2331–2336. https://doi.org/10.1073/PNAS.1315171111/SUPPL_FILE/SM02.AVI

Fassihi, A., Akrami, A., Pulecchi, F., Schö, V., Diamond Correspondence, M. E., Schö Nfelder, V., & Diamond, M. E. (2017). Transformation of Perception from Sensory to Motor Cortex Article Transformation of Perception from Sensory to Motor Cortex. *Current Biology*, *27*, 1585-1596.e6. https://doi.org/10.1016/j.cub.2017.05.011

Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nat. Neurosci.*, *17*(5), 738–743. https://doi.org/10.1038/nn.3689

Fornaciai, M., & Park, J. (2018). Attractive Serial Dependence in the Absence of an Explicit Task. *Psychological Science*, *29*(3), 437–446. https://doi.org/10.1177/0956797617737385/ASSET/IMAGES/LARGE/10.1177_0956797617737385-FIG3.JPEG

Fritsche, M., Mostert, P., & de Lange, F. P. (2017). Opposite effects of recent history on perception and decision. *Curr. Biol.*, *27*(4), 590–595. https://doi.org/10.1016/j.cub.2017.01.006

Hachen, I., Reinartz, S., Brasselet, R., Stroligo, A., & Diamond, M. E. (2021). Dynamics of history-dependent perceptual judgment. *Nature Communications*. https://doi.org/10.1038/s41467-021-26104-2

Hollingworth, H. (1910). The Central Tendency of Judgment. *The Journal of Philosophy, Psychology and Scientific Methods*, *7*.

Jou, J., Leka, G. E., Rogers, D. M., & Matus, Y. E. (2004). Contraction bias in memorial quantifying judgment: Does it come from a stable compressed memory representation or a dynamic adaptation process? *American Journal of Psychology*, *117*(4), 543–564. https://doi.org/10.2307/4148991

Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1995). An Introduction to Reinforcement Learning. *The Biology and Technology of Intelligent Autonomous Agents*, 90–127. https://doi.org/10.1007/978-3-642-79629-6_5

Kiyonaga, A., Scimeca, J. M., Bliss, D. P., & Whitney, D. (2017). Serial Dependence across Perception, Attention, and Memory. *Trends in Cognitive Sciences*, *21*(7), 493–497. https://doi.org/10.1016/j.tics.2017.04.011

Liberman, A., Fischer, J., & Whitney, D. (2014). Serial dependence in the perception of faces. *Current Biology*, *24*(21), 2569–2574. https://doi.org/10.1016/j.cub.2014.09.025

Manassi, M., Liberman, A., Kosovicheva, A., Zhang, K., & Whitney, D. (2018). Serial dependence in position occurs at the time of perception. *Psychonomic Bulletin and Review*, *25*(6), 2245–2253. https://doi.org/10.3758/S13423-018-1454-5/FIGURES/4

Nunez, M. D., Vandekerckhove, J., & Srinivasan, R. (2017). How attention influences perceptual decision making: Single-trial EEG correlates of drift-diffusion model parameters. *Journal of Mathematical Psychology*, *76*, 117–130. https://doi.org/10.1016/J.JMP.2016.03.003

Pedersen, M. L., Frank, M. J., & Biele, G. (2017). The drift diffusion model as the choice rule in reinforcement learning. *Psychonomic Bulletin and Review*, *24*(4), 1234–1251. https://doi.org/10.3758/S13423-016-1199-Y/FIGURES/1

Romo, R., Brody, C. D., Hernández, A., & Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature 1999 399:6735*, *399*(6735), 470–473. https://doi.org/10.1038/20939

Serrano-Fernández, L., Beirán, M., Romo, R., & Parga, N. (2023). Prefrontal Cortex Neural Correlates of the Contraction Bias in Frequency Discrimination. *BioRxiv*, 2023.07.27.550794. https://doi.org/10.1101/2023.07.27.550794

Shahar, N., Hauser, T. U., Moutoussis, M., Moran, R., Keramati, M., Consortium, N. S. P. N., & Dolan, R. J. (2019). Improving the reliability of model-based decision-making estimates in the two-stage decision task with reaction-times and drift-diffusion modeling. *PLOS Computational Biology*, *15*(2), e1006803. https://doi.org/10.1371/JOURNAL.PCBI.1006803

Sohn, M. H., Ursu, S., Anderson, J. R., Stenger, V. A., & Carter, C. S. (2000). The role of prefrontal cortex and posterior parietal cortex in task switching. *Proceedings of the National Academy of Sciences*, *97*(24), 13448–13453. https://doi.org/10.1073/PNAS.240460497

Steinemann, N. A., Stine, G. M., Trautmann, E. M., Zylberberg, A., Wolpert, D. M., & Shadlen, M. N. (2023). Direct observation of the neural computations underlying a single decision. *BioRxiv*, 2022.05.02.490321. https://doi.org/10.1101/2022.05.02.490321

Summerfield, C., Egner, T., Greene, M., Koechlin, E., Mangels, J., & Hirsch, J. (2006). Predictive codes for forthcoming perception in the frontal cortex. *Science*, *314*(5803), 1311–1314. https://doi.org/10.1126/SCIENCE.1132028/SUPPL_FILE/SUMMERFIELD.SOM.PDF

Tal-Perry, N., & Yuval-Greenberg, S. (2022). Contraction bias in temporal estimation. *Cognition*, *229*, 105234. https://doi.org/10.1016/J.COGNITION.2022.105234

Tavares, G., Perona, P., & Rangel, A. (2017). The attentional Drift Diffusion Model of simple perceptual decision-making. *Frontiers in Neuroscience*, *11*(AUG), 268904. https://doi.org/10.3389/FNINS.2017.00468/BIBTEX

Wang, T. Y., Liu, J., & Yao, H. (2020). Control of adaptive action selection by secondary motor cortex during flexible visual categorization. *ELife*, *9*, 1–26. https://doi.org/10.7554/ELIFE.54474

Weilnhammer, V., Stuke, H., Hesselmann, G., Sterzer, P., & Schmack, K. (2017). A predictive coding account of bistable perception - a model-based fMRI study. *PLoS Computational Biology*, *13*(5). https://doi.org/10.1371/JOURNAL.PCBI.1005536

Whitlock, J. R. (2017). Posterior parietal cortex. *Current Biology*, *27*(14), R691–R695. https://doi.org/10.1016/J.CUB.2017.06.007

ACKNOWLEDGMENTS

Questi quattro anni sono stati un viaggio. Molto è cambiato nella mia vita da quando iniziai il dottorato. Certamente in meglio.

Grazie a Mathew, per la sua fiducia, la sua guida e le incessanti, e a volte fantasiose, discussioni scientifiche.

Grazie a Francesca, senza il cui contributo questo lavoro sarebbe assai più povero. Ma, soprattutto, grazie dell'amicizia che abbiamo costruito in questi anni.

Grazie a Annachiara, Lilith, Maria, Arianna, Marco e Fabrizio e gli altri membri passati del gruppo: avete reso, per me, il nostro laboratorio una comunità, non solo un luogo di lavoro.

Grazie a Chiara, Federico, Matteo e Giacomo. Siete una presenza gioiosa e costante nella mia vita, e questo NON è il momento più importante che condivido con voi.

Grazie ai miei genitori, Elena e Erminio, per l'amore e il supporto di tutti gli anni. Non sarò mai in grado di ripagarlo, ma sappiate che vorrei.

Grazie ai genitori "extra", Margherita e Gilberto, che ho guadagnato in età più avanzata ed estendono con la propria presenza la mia famiglia.

Grazie a Lucia, il cui amore non cessa mai di sorprendermi, e a Gaia. Voi siete, e sarete sempre, il mio progetto più grande.