Humans, rats and mice show species-specific adaptations to sensory statistics in categorisation behaviour

Victor Pedrosa[1,†], Elena Menichini[1,†], Quentin Pajot-Moric[1,†], Peter Vincent[1], Liang Zhou[2], Lillianne Teachen[1], Peter Latham[2], Athena Akrami[1]

[1]Sainsbury Wellcome Centre, University College London
[2]Gatsby Computational Neuroscience Unit, University College London
[†]These authors contributed equally
[*]Corresponding author: Athena Akrami, athena.akrami@ucl.ac.uk

Abstract

One key aspect of animal intelligence is the ability to find, represent and update knowledge of statistical regularities in the sensory environment, in service of adaptive behaviour. This allows animals to build appropriate priors, in order to disambiguate noisy inputs, make predictions and act more efficiently. Here, we compare human, rat, and mouse performance on a 2-alternative-forced-choice sound categorisation task. Importantly, we manipulate the stimulus statistics without affecting the overall weight of each category, and we investigate decision-making in such distinct 'statistical contexts' associated with different stimulus prior probabilities. Humans, rats, and mice develop a sensory-dependent bias that increases reward to a near-optimal level. We also explore how flexible animals are in learning different types of statistics, without any explicit cue, when switching from one prior probability to another. All species can adapt to a new underlying distribution by gradually shifting their psychometric behaviour after the switch to a new statistical context, with mice showing the slowest rate of adaptation. Despite the overall sensory-dependent bias that is similar across species, humans versus rodents show distinct trial-to-trial learning updates. Rodents show an overall win-stay choice bias, with rats incorporating decision confidence more strongly than mice. Humans, on the other hand, show a more complex learning update, and more indicative of utilising detailed sensory statistics. We propose two classes of computational models, a boundary learning (discriminative) and a stimulus-category learning (generative) model. We show how they can account for the optimal exploitation of the sensory statistics, and predict the observed species-specific trial-to-trial updates.

Keywords: Statistical learning, 2-alternative-forced-choice, Cross-species, Comparative cognition, Perceptual decision-making, Bayesian inference, Generative model, Discriminative model, Sensory statistics

Introduction

Finding and exploiting patterns in the environment is critical for living in a complex yet structured world. As sensory events unfold over time, they can reveal meaningful regularities that we can detect and learn almost automatically and incidentally. This allows us to build appropriate prior beliefs about future stimuli, which may be used to disambiguate noisy inputs, make predictions and act more efficiently (Knill and Pouget 2004; De Ridder, Vanneste, and Freeman 2014). The process by which we learn this structure from common experiences is known as statistical learning (Sherman, Graves,

and N. B. Turk-Browne 2020; Schapiro and N. Turk-Browne 2015). Extracting and exploiting statistical regularities often occurs implicitly — without instruction, and automatically (Destrebecqz and Cleeremans 2001), and is the bedrock for learning about, and responding to, a complex yet structured world. This sensitivity to environmental statistics is a pervasive property of the brain that operates over different timescales, e.g. phylogenetic (E. P. Simoncelli and Olshausen 2001), ontogenetic (Berkes et al. 2011) or sometimes it can be rapid - on the order of minutes (Aslin and Newport 2012; Schapiro and N. Turk-Browne 2015). Learning the structure of common patterns in the world allows us to adapt our behaviour and actions to favour our objectives, like collecting rewarding stimuli. For example, consider the choice of whether a piece of fruit is ready to eat. Over enough experiences of eating under-ripe fruit we learn the appropriate colour and feel after which we would choose to eat the fruit. This type of learnt categorising behaviour is a key feature of statistical learning where fundamentally continuous and uncertain sensory information is used to make a categorical decision (K. H. Britten et al. 1992; Anderson 1991). Furthermore, we are able to adapt our behaviour as the environment changes, which is essential for survival in a dynamic world (Pearson et al. 2011); fruit may ripen faster in new climates. Compounding the change in look and feel of the fruit, different varieties may be ripe at different colours when compared to the food we grew up with. In adapting to relevant structures in changing environments, animals need to integrate sensory information across time and combine that with prior information about the probability of a particular response being correct (Green and Swets 1966).

We sought to investigate this ability to detect changes, and their meaning, in the sensory environment by turning to auditory decision making. Sound categorisation is the ability of an organism to group sounds into different categories based on their acoustic features (Hausberger and Cousillas 1995; Sturdy et al. 2007; Mercado III, Orduña, and Nowak 2005; S. T. Liu et al. 2019). This behaviour is an example of a transformation from continuous data (for example frequency and pitch) to a binary decision and is essential for animals to navigate and survive in their environment, as it allows them to distinguish between sounds that are relevant to their survival and those that are not. For example, an animal needs to be able to distinguish between noises that are near (loud) or far (quiet), between the calls of its own species and those of other species, as well as between sounds that signal danger and those that indicate safety. Previous studies have examined animal performance in sound categorisation tasks to examine how brain physiology might support such a behaviour (Znamenskiy and Zador 2013; Y. Liu, Xin, and Xu 2021; Zhong et al. 2019; Mercado III, Orduña, and Nowak 2005; Hausberger and Cousillas 1995; Russ, Y.-S. Lee, and Cohen 2007; J. H. Lee et al. 2009). Typically animals are presented with stimuli chosen in some range with an arbitrary boundary placed in the middle. The animals are not explicitly shown where the boundary is (which could be done by playing some "anchor" tone) but rather must learn in which category each stimulus belongs based on receiving feedback after the chosen actions (Znamenskiy and Zador 2013; Xin et al. 2019). Some studies extend this framework by varying the probability of each category, such that one decision should become more common than the other, allowing experimenters to investigate representations of bias or prior preference for category, for instance in the visual domain (Ashwood et al. 2022; The International Brain Laboratory et al. 2021). However, the shape of the distribution of samples presented within each category is typically kept the same, even if in some studies boundary is moved to a new location (Zhong et al. 2019; Hachen et al. 2020). These experimental paradigms make it difficult to determine if the subjects are learning some true underlying statistics of the sensory environment, or instead adjust their choices purely based on reward probability of each of the two binary choices.

In this work we asked whether experimental subjects could adapt their choices, depending on the underlying distribution of sensory input, when category probability remained the same. To do so we compared human, rat, and mouse performance on a 2-alternative-forced-choice (2-AFC) sound categorisation task (Fig. **1** A). This type of 2-AFC task is an example of perceptual decision-making under uncertainty and has been widely used in rodents and humans (Otazu et al. 2009; Jaramillo and Zador 2011; Raposo et al. 2012; Kelly and O'Connell 2013; Znamenskiy and Zador 2013). Our study, however, revolved around a key manipulation neglected in previous studies: we carefully varied the stimulus statistics while maintaining the overall weights of both categories equal and studied the decisions made in distinct 'statistical contexts' associated with different stimulus prior probabilities. Theoretically, the optimal strategy for this task would be to make context-specific biased decisions (Rorie et al. 2010; Gao, Tortell, and McClelland 2011; The International

Brain Laboratory et al. 2021; Feng et al. 2009). Humans, rats, and mice develop a stimulus-dependent bias that increases reward to a near-optimal level. Moreover, we explored how flexible subjects were in learning different types of statistics, without any explicit cue, by switching statistical context from one prior sensory distribution to another. All species were able to adapt to a new underlying distribution by gradually shifting their psychometric behaviour after the switch to new statistics, with mice exhibiting the slowest rate of adaptation to the new distribution. Despite the development of similar overall stimulus-dependent biases, different species showed distinct trial-to-trial learning updates. Both our generative (stimulus-category estimation) and discriminative (boundary-estimation) models (Jebara 2012) could account for the learning and exploiting of the 'statistical context' by explaining rodent and human behaviour and predicting their trial-to-trial choice updates.

Results

Humans, rats and mice adapt to sensory statistics in a 2-AFC task

We sought to investigate the role played by sensory statistics and past experience in perception and decision-making. To this end, we implemented an auditory version of a 2-alternative-forced-choice (2-AFC) task for humans, rats and mice (Fig. **1** A). In the human version of the task, subjects were presented with a brief pure tone and were then asked to choose between two options, A or B, on the screen (Fig. **1** C). To compare results across different species, we implemented an analogous task in rats and mice.

Freely-moving rats were trained to categorise sound amplitude (Fig. **1** A and C), by poking in one of the two available water ports on the left or right side of the box (which we denote as A and B, see methods). Upon the correct choice, the reward was then delivered from either one of the ports, depending on the sound's amplitude: port A corresponded to a sound in the lower half of the range and port B corresponded to the upper half. Head-fixed mice were trained to categorise sound frequency or amplitude (Fig. **1** A and C), by licking the left or right water spout. The correct spout was determined similarly to the protocol implemented on rats and humans: low-range sounds were classified as A (spout A), whereas high-range sounds were classified as B (spout B). Error trials were followed by a 'time-out' period for rats and mice.

To investigate whether the underlying sensory statistics in these sound categorisation tasks had an impact on behaviour, we exposed subjects to one of three potential stimulus distributions: uniform, hard-A asymmetric, and hard-B asymmetric (Fig. **1** B). On uniform-distribution sessions, stimuli were randomly drawn from a uniform distribution covering the entire range of possible stimuli (Fig. **1** B, green). On hard-A asymmetric sessions, stimuli were randomly drawn from an asymmetric distribution in which stimuli belonging to category A were over-represented around the threshold (and hence were usually harder to classify), whereas stimuli belonging to category B were uniformly distributed (Fig. **1** B, blue). Conversely, on hard-B asymmetric sessions, stimuli were randomly drawn from an asymmetric distribution in which stimuli belonging to category B were over-represented around the threshold (harder to classify), whereas stimuli belonging to category A were uniformly distributed (Fig. **1** B, red). In these sound categorisation tasks, the range of frequencies presented to the subjects was specific to each species. The threshold between categories A and B, however, was always defined as the middle point of the range, and the probability of presenting a stimulus belonging to category A was the same as presenting a stimulus belonging to category B. Therefore, using these three possible distributions, we could assess whether the stimulus statistics modulate behaviour while maintaining the probability of stimulus A and stimulus B equal.

After learning the task, humans, rats and mice showed stable behaviour across sessions. To quantify their performance, we collected the responses across all subjects exposed to one particular distribution and fit their overall performance with a psychometric function. This psychometric function was defined by four parameters: the mean (the shift of the point-of-equality toward one category), the width (the inverse of the slope), and the two lapse values (error rates on easiest

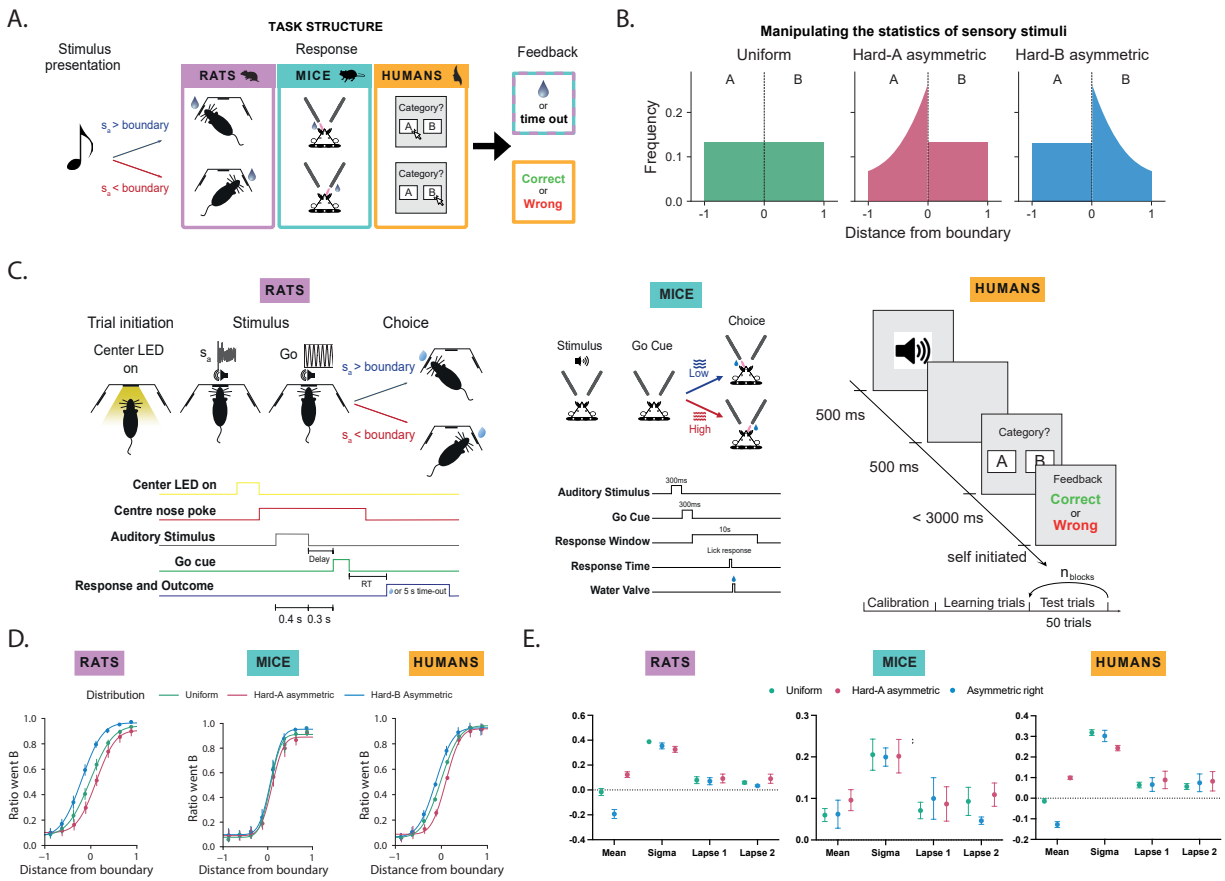

Figure 1: **Humans, rats, and mice adapt to sensory statistics in a 2-AFC task. A**: Experimental setup for each species - rats were freely moving and indicated their choice by poking in either the left or right nose port. Mice were head fixed and indicated their choice by licking either the left or right lick-spout. Humans performed experiments online and indicated their choice by clicking on a category button on the screen. Water reward was given for rodents, written feedback for humans. Rodents received a timeout after an incorrect choice **B**: Sampling distributions used in the experiments. In all three cases, the total mass on either side of the boundary is equal. **C**: Full trial structure for each species. **D**: Psychometric curves for each species (averaged across individuals) for each tested distribution, with colour codes matching those in **B**. **E**: Fit parameters (with SEM) for each species for each distribution.

stimuli), one for category A stimuli and one for category B stimuli (see methods). When exposed to uniformly distributed stimuli, humans, rats and mice exhibited a psychometric response with a mean of around zero ($mu_{humans} = -0.018$, $mu_{rats} = -0.016$, $mu_{mice} = 0.060$). Therefore, the mean of the psychometric fit for all three species coincided with the middle point of the stimulus range and the threshold. When exposed to asymmetrically distributed stimuli, however, human and rodent behaviour was biased towards the category that was over-represented around the boundary (Fig. **1** D). This bias was only observed for stimuli located around the boundary (hard stimuli) and was thus only reflected on the mean of the psychometric response (Fig. **1** E). Humans and rodents were not sensitive to changes in easy stimuli as their lapse rates remained unchanged when exposed to different distributions (Fig. **1** E). Therefore, across all three species, performance depended on the underlying sensory statistics (Fig. **1** D).

Subjects adapt to sensory statistics to optimise performance and maximise reward

We next sought to investigate whether humans and rodents would develop a better representation of the boundary position when over-exposed to hard stimuli around the boundary. To test this, we exposed humans and rats to stimuli drawn from a unimodal distribution in which stimuli around the boundary were over-represented (Fig. S1 A). All other aspects of the task were kept identical to the protocol implemented for uniform and asymmetric distributions. If the over-exposure to

hard stimuli led to a better representation of the boundary position, the psychometric fit to human and rodent performance would exhibit a sharper psychometric curve. Humans and rats, however, did not change their performance – as measured by the psychometric fit – when exposed to stimuli drawn from a unimodal distribution (Fig. S1 B). Conversely, to understand whether an under-exposure to stimuli around the boundary would affect performance, we next exposed rats and humans to stimuli drawn from a bimodal distribution in which hard stimuli were under-represented (Fig. S1 A). Again, all other details of the experimental protocols were kept unchanged. When exposed to stimuli drawn from a bimodal distribution, humans and rodents exhibited the same performance as when exposed to uniformly distributed stimuli. Therefore, the shift in the mean of the psychometric curve was only observed when the distribution of stimuli was asymmetric around the boundary and was not related to an overall over- or under-representation of stimulus samples around the boundary (Fig. S1 B), since over- or under-exposure to hard stimuli did not alter human or rat performances.

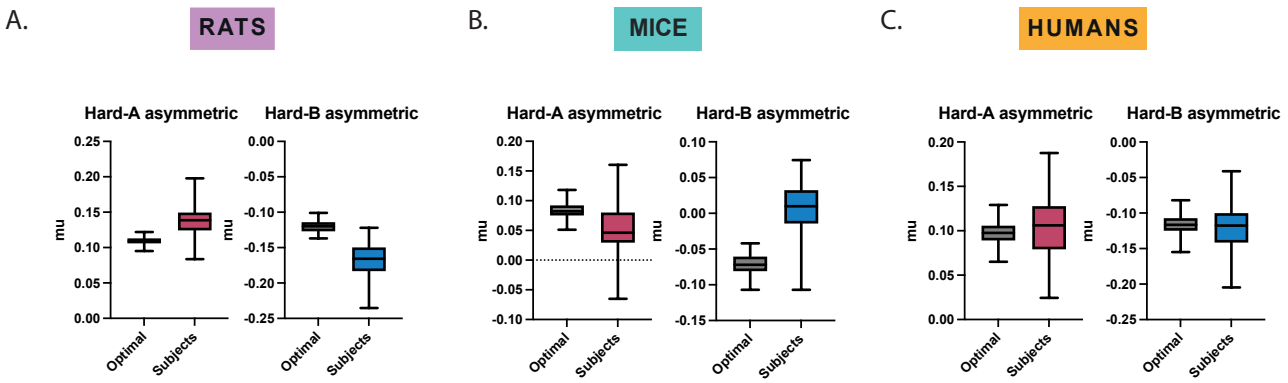

Figure 2: **Subjects adapt to sensory statistics to optimise performance and maximise reward**. The optimal mean parameter of the psychometric curves for the two skewed distributions, estimated from simulated agents with similar slopes and lapse values as rats (**A**), mice (**B**) and humans (**C**) (see Methods for further details), compared to the mean parameters calculated from data pooled across subjects.

As shown in (Fig. **1** E), humans and rodents did not change the lapses (category A and category B lapses) or the slopes of their psychometric curves (*sigma*), when exposed to different distributions. Instead, they systematically changed the mean value of their psychometrics. We predicted that rodents and humans would adapt to sensory statistics in order to optimise performance. The sensitivity of humans and rodents to sound category, as measured by the slope of the psychometric curve, is limited by sensory accuracy and internal noise. We hypothesised that, over the course of learning, subjects improved their accuracy up to a point where their sensitivity reached a plateau (Fig. S1). In order to maximise reward, however, humans and animals might be able to adapt the mean of their psychometric curve by imposing a bias for stimuli around the boundary. Since category representation was asymmetric around the boundary, subjects could improve their performance if they biased their decision towards the over-represented category for hard stimuli. To test this hypothesis, we simulated an agent that had the same psychometric parameters to our experimental subjects (rats, mice and humans, separately), except for the mean (*mu*) of the psychometric curve. The performance of the simulated agent was therefore completely determined by the parameter *mu* and the distribution of stimuli. We first calculated the value of *mu* that maximised the agent's performance on each of the asymmetric distributions. When compared to the *mu* fit to the uniform context an optimal agent should shift its psychometric curve right-ward for the hard-A asymmetric distribution (hence, positive values of *mu*) and left-ward for the hard-B asymmetric distribution (hence, negative values of *mu*). We then calculated the *mu* parameter for humans, rats and mice for each asymmetric distribution and compared it to the *mu* parameter for the uniform contexts. All species shifted their psychometric curves in the optimal direction (Fig. **2** A-C). The adaptations of subjects to sensory statistics could therefore be the result of a reward-maximisation strategy in the face of limited sensory accuracy.

Humans, rats and mice adapt to changes in sensory statistics

In our human version of the experiments, we split the experiment into 20 blocks of 50 trials. On each block, the stimulus distribution was randomly selected between hard-A or hard-B asymmetric. The random switches prevented subjects from developing an understanding of the macro-structure of the task and adapting due to this structure by developing overt biases, instead of learning from accumulated evidence from sensory exposure. To investigate whether this was the case, we analysed human behaviour by performing a block-specific psychometric fit within the experimental session. We considered two subsets of blocks: 1) blocks that were assigned a hard-A stimulus distribution and that were preceded by a hard-B block, and 2) blocks that were assigned a hard-B block and were preceded by a hard-A block. These two subsets contained the blocks that followed a switch in stimulus distribution – from hard-A to hard-B or the opposite. The psychometric fit of human performance in these blocks revealed a distribution-dependent bias as observed in fixed-distribution experiments (Fig. **3** B). This suggests that humans possess the ability to adapt to changes in stimulus distributions in a flexible and reversible manner.

We then performed switching distribution experiments in rats and mice. In these experiments, animals were split into two cohorts. For the first cohort, animals were exposed to stimuli following a hard-A asymmetric distribution throughout learning and for several sessions following performance stabilisation. These animals were then moved to a hard-B asymmetric context and kept in that context for a few sessions. For the second cohort, animals started in a hard-B asymmetric context and were later moved to a hard-A asymmetric context. We analysed these experiments by splitting the data into four groups: cohort 1 before and after the context switch, and cohort 2 before and after the switch. For both cohorts, the bias in rat psychometric response shifted from favouring one category to favouring the opposite category following the context switch (Fig. **3** B). When exposed to hard-A asymmetric stimuli, rodents exhibited a bias towards category A for stimuli around the boundary, irrespective of training history. The opposite was true for rodents exposed to the hard-B asymmetric context. In this case, animals exhibited a bias towards category B for stimuli around the boundary. Therefore, rats and mice adapted to changes in stimulus statistics. We, however, observed that the speed of this adaptation differed between rats and mice. Rats shifted the mean of their psychometric curves almost immediately, within the first session after the switch (Fig. **3** C, left) and their behaviour on the new statistical context stabilised after two sessions. Mice, however, displayed slower learning and required more sessions until the their psychometric performance plateaued on the new statistical context (Fig. **3** C, right). This difference in speed of learning was not due to the total number of trials per session (352 trials/session, on average, for mice and 216 trials/session, on average for rats).

Up to this point, when investigating human and rodent responses to asymmetrically distributed stimuli, we have imposed an asymmetric distribution in which one of the categories is uniformly distributed across the stimulus range and the other category is over-represented around the boundary. We next wondered whether the biases observed in humans and rodents were solely due to the relative difference in category representation around the boundary or whether the shape of the non-uniform category would modulate human and rodent behaviour. To address this question, we exposed humans and rats to two new stimulus distributions. Under the first one, easy-A asymmetric, the non-uniform side of the distribution over-represented easy stimuli and under-represented hard stimuli belonging to category A (Fig. S1 A). Under the second type of distribution, easy-B asymmetric, the non-uniform side of the distribution over-represented easy stimuli and under-represented hard stimuli belonging to category B (Fig. S1 A). Our hypothesis that the shift in psychometric response was a consequence of a reward-maximisation strategy suggested that subjects exposed to easy-A asymmetric distributions would demonstrate performances comparable to that of subjects exposed to hard-B asymmetric distributions – and equivalently for subjects exposed to easy-B and hard-A asymmetric distributions. In agreement with this hypothesis, human subjects exposed to easy-A asymmetric developed a bias towards category A for stimuli around the boundary whereas subjects exposed to easy-B asymmetric stimuli developed the opposite bias – towards category A (Fig. S1 B). The same behaviour was observed in rats exposed to these distributions (Fig. S1 B). Accordingly, these results suggest that the bias observed in humans and rats is due to the difference in relative probabilities of category A and category B stimuli

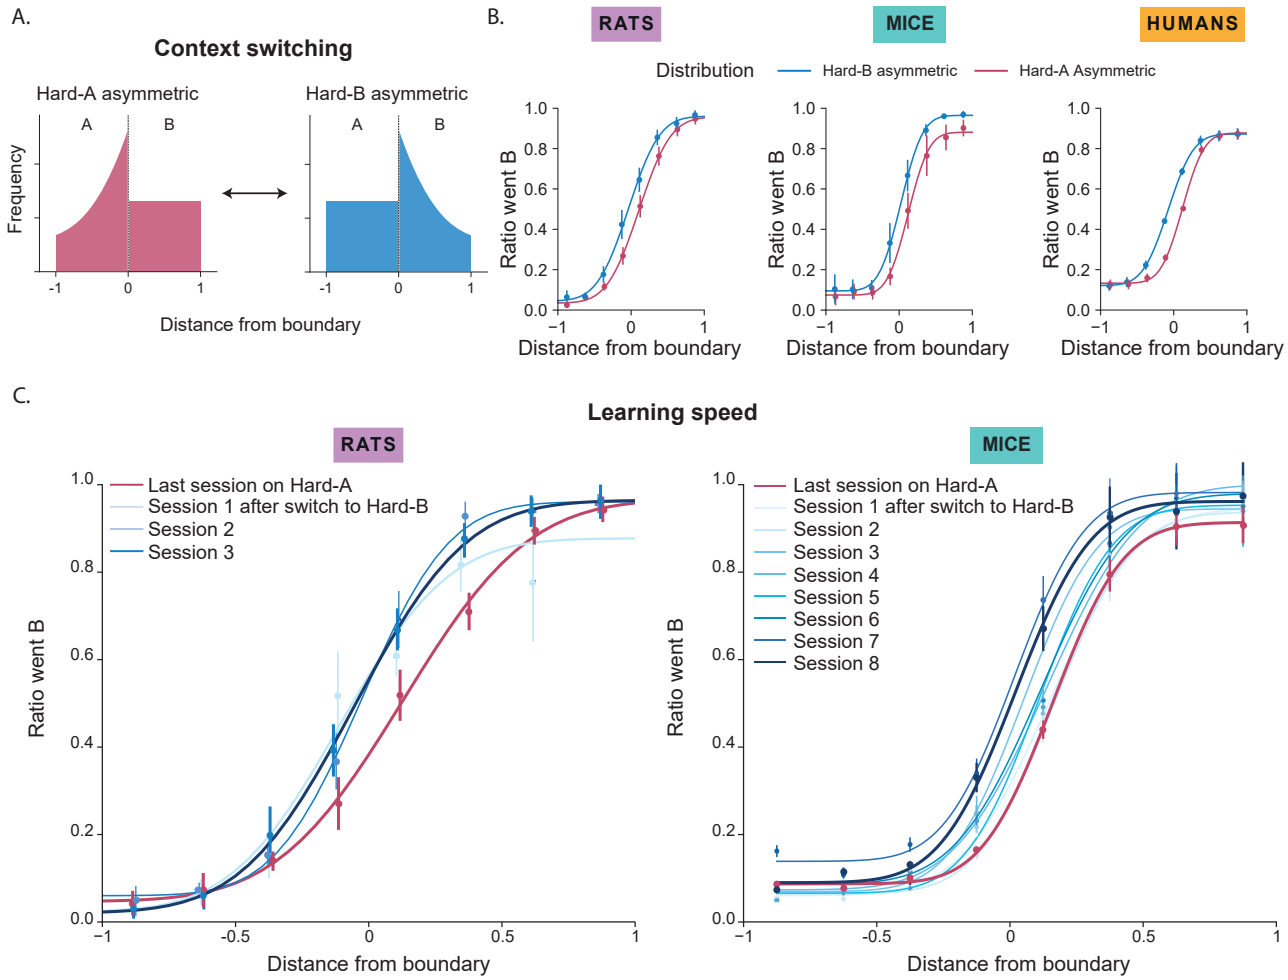

Figure 3: **Humans, rats and mice adapt to changes in sensory statistics**. **A**: Sampling distributions used in each context. **B**: Psychometric curves for the asymmetric contexts for each species. **C**: Performance across each session starting from the final session in context hard-A (red), then for each subsequent session in hard-B from light to dark blue. Psychometric curves show that rats take fewer sessions to adapt to the new statistics after a context switch, compared to mice.

around the decision boundary, rather than to the specific nature of the asymmetric distributions presented.

Boundary-estimation and Stimulus-category models to account for optimisation behaviour

Our experimental data suggest that humans and rodents exhibit similar behavioural performance and adaptation to sensory statistics. So far, we have analysed behaviour in terms of psychometric fits to overall performance. The similarities observed in our data, however, may emerge from different learning mechanisms and strategies. In order to investigate this, we implemented a computational model to perform a 2-AFC task, where our learning agent was subjected to the same manipulations as the ones imposed on human and rodent tasks. Namely, we tested the model in three different contexts, where: 1) stimuli were independently drawn from a uniform distribution, 2) stimuli were independently drawn from a hard-A asymmetric distribution, and 3) stimuli were independently drawn from a hard-B asymmetric distribution.

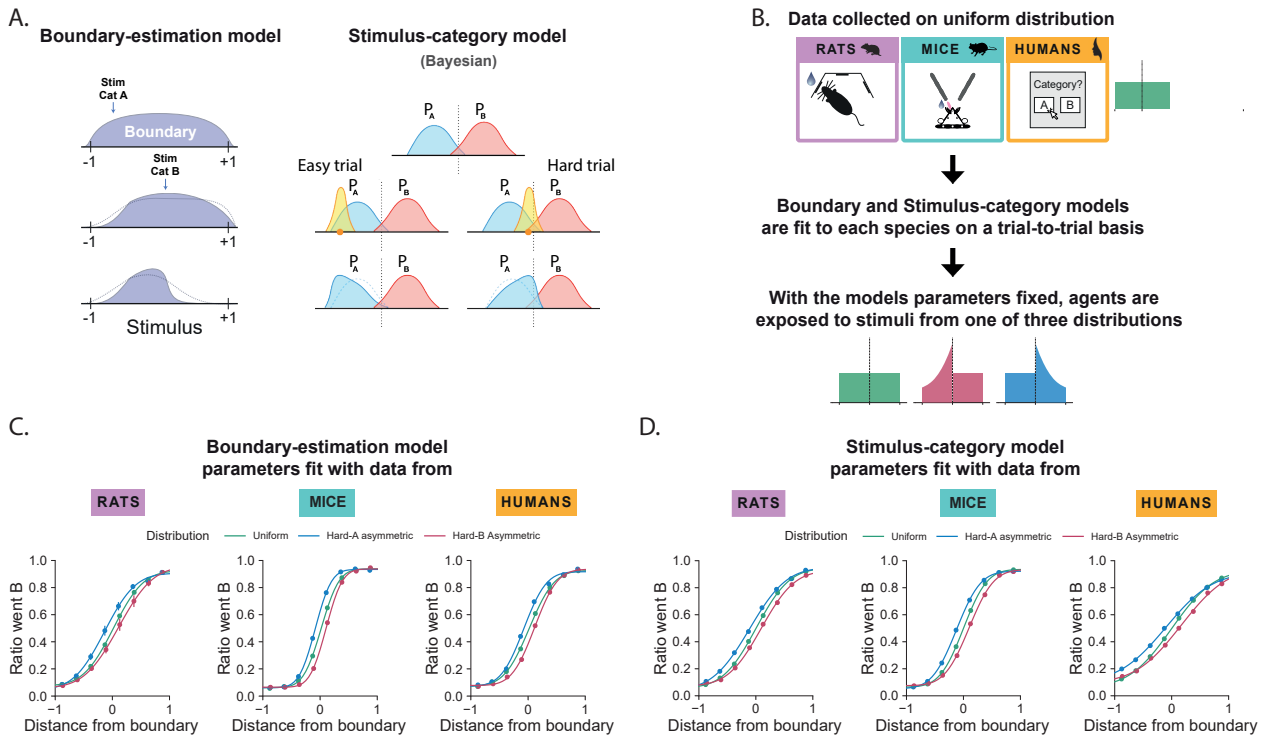

Figure 4: **Boundary-estimation and Stimulus-category models**. **A**: Schematics showing the learning process for the stimulus-category and boundary estimation models **B**: Model fitting regime - parameters fit on the initial sessions with uniform distributions are then tested on the new contexts. **C,D**: Predicted psychometric curves generated by the Boundary-estimation and Stimulus-category models (without repulsion), respectively, for each species.

We next proposed two classes of computational models (Fig. **4** A). In the first - the boundary-estimation model - the agent estimates the position of the decision boundary by building a distribution of boundary positions and updating this distribution on a trial-to-trial basis. In the second - the stimulus-category model - the agent learns a stimulus-to-category mapping by building category-specific stimulus distributions and updating these distributions on a trial-to-trial basis. This model is a form of Bayesian learner.

The boundary-estimation model utilises an agent that initially has no preconception of where the boundary may be located, but is provided with a stimulus range of -1 to 1 (Fig. **4** A). The agent's initial state is defined by a uniform distribution over this range, which represents the estimated probability of the decision boundary's position. The agent is subsequently presented with stimuli on each trial - which is perturbed by an amount determined by a noise parameter, to give rise to the perceived stimulus - and returns a binary decision (either category A or category B), with the probability of each decision being determined by the stimulus's relation to the current estimate of the boundary position. After each decision, the agent receives feedback on the correct stimulus category and updates its distribution of boundary positions accordingly. To emulate human and rodent behaviour, the model incorporates a softer update protocol controlled by two main parameters: the update precision, which reflects how accurately the stimulus position is considered when updating the model, and the learning rate, which determines the strength of the update on each trial. Additionally, the model includes a "forgetting" term, causing the boundary model to move towards a uniform distribution on each trial, and a relaxation amplitude parameter controlling the degree to which this occurs (see methods for more details). Finally, to account for lapses, in which the agent may randomly ignore the stimulus and act without reference to the boundary model, the model also includes a probability parameter for this behaviour.

The stimulus-category model operates by utilising an agent whose state is represented by two probability distributions, one

corresponding to the likelihood of stimuli belonging to category A and another corresponding to category B (Fig. **4** A). During the decision process, the agent's probability of selecting each category is proportional to the amplitude of the respective probability distribution for that category at the presented stimulus. Similar to the boundary-estimation model, this stimulus is altered by a noise parameter. Following stimulus presentation and feedback, the agent updates its model for the corresponding category by increasing the probability of stimuli surrounding the perceived stimulus for that category, while keeping the model of the opposing category unchanged. To account for lapses, similar to the boundary-estimation model, the agent may randomly select a category in some trials, with the number of such lapses determined by a predefined parameter.

To test these models, we first fit each on the trial-to-trial choice data for each species when subjects were exposed to uniformly distributed stimuli (Fig. **4** B). Using the best fit parameters, we then exposed the learning agents to stimuli drawn from one of the following three distributions: uniform, hard-A asymmetric, or hard-B asymmetric. Similar to the optimal behaviour shown by humans and rodents, when exposed to asymmetrically distributed stimuli, the agents' performance was biased toward the over-represented stimuli around the decision boundary. This bias emerged on both boundary-estimation and stimulus-category agents (Fig. **4** C and D). Therefore both computational models account for the optimal behaviour observed in humans and rodents.

Our computational model predicts divergent updating strategies for humans and rodents

Both computational models proposed here account for the shift in psychometric curves observed in humans and rodents when they are exposed to different stimulus statistics. This behaviour, however, is an average behaviour that may emerge from different trial-to-trial update strategies, when animals may systematically bias their choices, given the history of past sensory events (Akrami et al. 2018; Fassihi et al. 2014; Roy et al. 2021), their choices and outcomes (Abrahamyan et al. 2016; Busse et al. 2011), as well as their decision confidence (Lak et al. 2020). To investigate whether the boundary-estimation and the stimulus-category models implement different update strategies, we calculated the update matrix for each model under uniformly distributed stimuli. The update matrix is defined as in by the difference between the conditional psychometric function – conditioned on the stimulus presented at the previous trial – and the overall psychometric function (ibid.). This matrix is then built by averaging this difference (across subjects) for each pair of current and previous stimuli (Fig. **5**).

We first calculated the update matrix for a boundary-estimation model (Fig. **4** A). Using the best fit parameter from an agent trained on individual trials from each species, we measured the update matrix for a sub-selection of trials in which the previous trial was correctly assigned by the agent (hence, we did not directly fit the model to the update matrix). The update matrix, for agents fit to rat, mouse and human, exhibited a form of win-stay behaviour (Fig. **5** A). If the stimulus in the previous trial belonged to category A and the agent correctly categorised the stimulus, the agent was more likely to assign the stimulus in the current trial to category A again. The analogous was true for category B. This increase in the probability to repeat the agent's choice was larger if the current stimulus was close to the boundary and was also stronger if the previous stimulus was close to the boundary, similar to (ibid.).

We then similarly calculated the update matrix for the stimulus-category model (Fig. **5** B). Different from the boundary-estimation model, an agents that learns the individual category distributions manifests a trial-to-trial choice bias that is more complex than an overall 'win-stay' strategy. In particular the update matrix, for agents fit to human behaviour, reveals a 'stripe' pattern, indicating that the direction for the bias (toward or against the previous correct choice) depends on the exact interaction between the previous and current stimuli: imagine we want to probe the agent's performance in response to a 'hard' stimulus (i.e. close to the boundary) from category A. If in the previous trial, the agent has correctly categorised yet another 'hard' stimulus below the boundary as category A, it is more likely to categorise it correctly again on the current trial. However, if the previous trial was instead an 'easy' stimulus from A and was correctly categorised as A, it reduces the

chance of correctly responding A to a hard stimulus (from A) in the current trial. This pattern is less evident in the update matrices from agents fit to the rodent data (Fig. **5** B), however, qualitatively it is still present.

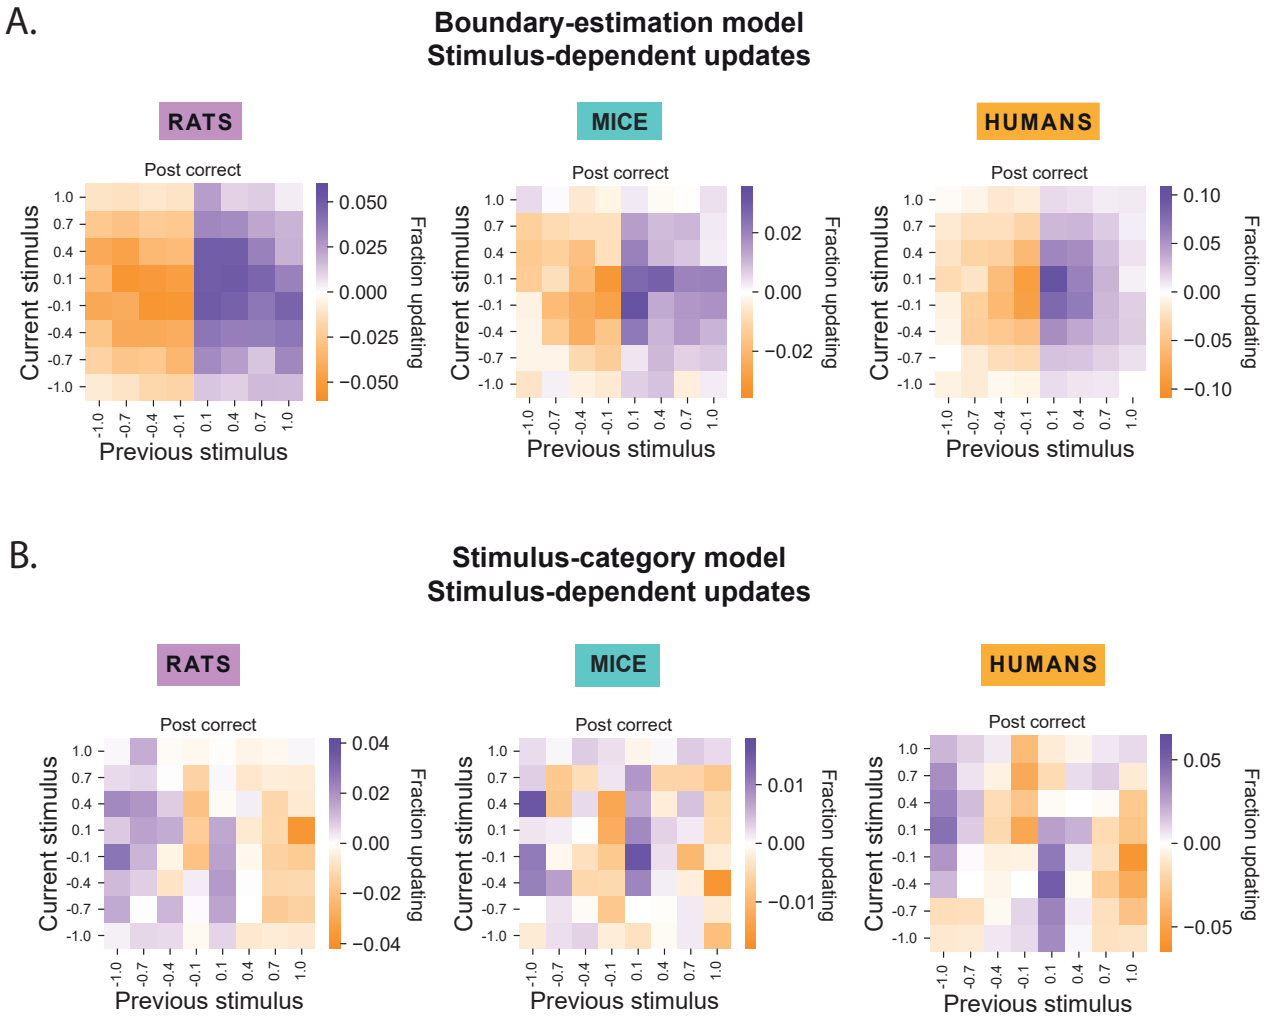

Figure 5: **Models predict divergent trial-to-trial updating strategies for humans and rodents**. Predicted update matrices for trials following correct responses for each species by the boundary-estimation model (panel **A**) and the stimulus category model (panel **B**) without repulsion. The difference between the conditional psychometric function – conditioned on the stimulus presented at the previous trial – and the overall psychometric function is first calculated. This matrix is then built by averaging this difference (across subjects) for each pair of current and previous stimuli. Only a subset of trials for which the previous trial was performed correctly is considered here. Purple colours signify a bias towards category A and orange colours signify a bias towards category B.

Humans and rodents use different updating strategies on 2AFC tasks

We next looked into how humans and rats changed and biased their performance in a trial-to-trial manner. We analysed the conditional psychometric functions dependent on the previous stimulus for all post-correct trials and compared this psychometric to their corresponding overall psychometric curves (Fig. **6**). Rats and mice tended to repeat their choice following correct ('win-stay' strategy) trials, although rats shows stronger modulation by stimulus difficultly (in both current and previous trials). Humans, instead, showed more complex behaviour, presenting the 'stripe' pattern shown for the Bayesian agent, learning the stimulus-category. To capture human behaviour, our model required local updates to each category model, whereas rat data could be captured if the agent used category-wide updates.

Recent experimental evidence and computational modelling suggest that perceptual biases may emerge from efficient

coding constraints which can lead to (repulsive) biases away from prior beliefs (Fritsche, Spaak, and De Lange 2020; Wei and Stocker 2015; Ganguli and Eero P. Simoncelli 2014). To incorporate this in our model, we implemented a repulsion term for both the boundary-estimation and the stimulus-category models (Fig. S3 A). This term shifts the stimulus presented in the current trial away from the stimulus presented in the last trial (see methods for details). The amount of repulsion is determined by a repulsion amplitude parameter. This repulsion term, together with the trial-to-trial update in each model, results in a sequence-dependent choice update.

Addition of a strong repulsion term to the boundary model could recapitulate some of non-uniformity in the trial-to-trial choice updates observed in humans and the stimulus-category model (Fig. S4.

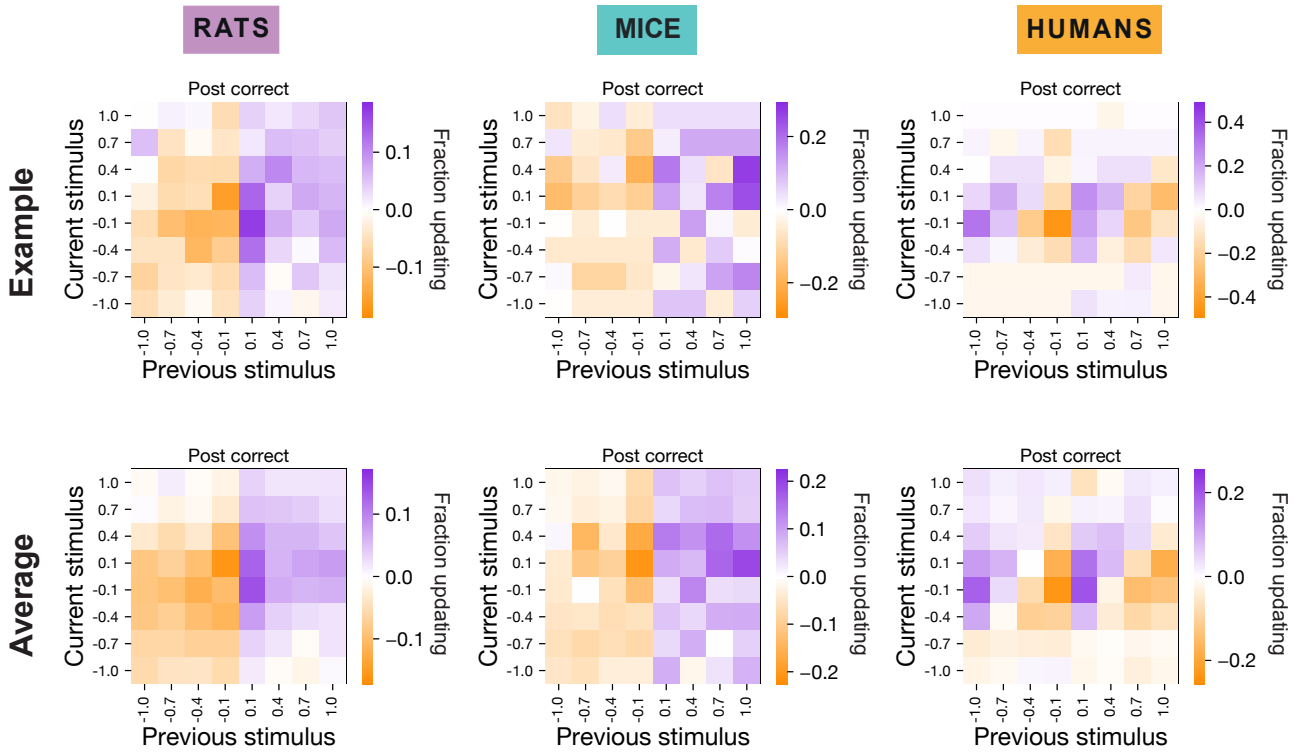

Figure 6: **Humans and rodents use different trial-to-trial updating strategies on2AFC tasks** Observed update matrices for trials following correct responses for each species (in each column) with examples of single animals shown in panel **A** and overall averages shown in panel **B**. colour codes are similar to (Fig. **5**)

Discussion

Adaptive integration of sensory and reward information

Learning the statistics of the environment is a key trait of intelligent systems for goal-directed behaviour. In this work we use a 2-AFC sound categorisation task with carefully controlled distributions of stimuli to compare the way in which humans, rats and mice change their behaviour as the statistics of the environment changes.

In principle, this kind of manipulation does not require subjects to track the change in context since the decision rule (greater or lower than some static boundary) remains the same. However, given the inherent noise in perception (Neri 2010), the optimal decision-maker on this task will make context-specific biased decisions. That is, once presented with asymmetric sampling of stimuli around the boundary, one should impose a bias towards the (locally) over-represented category for hard stimuli (Simon 1986). Subjects, in this way, could improve their overall performance and reward rate.

We found evidence for such effects in humans, rats, and mice. Indeed, when examining performance pooled across entire experimental sessions, all species showed a common effect. Previous studies have asked how animals integrate reward and sensory information dynamically in perceptual decision-making (Gao, Tortell, and McClelland 2011; Rorie et al. 2010; Feng et al. 2009). However, in those experiments, the reward probability of the binary choices was manipulated. Rorie et al showed that monkeys, when performing a classical evidence accumulation task (Kenneth H. Britten et al. 1993), will consider the reward probability of Left or Right options that are revealed to them at each trial (Feng et al. 2009; Rorie et al. 2010). Monkey's psychophysical performance was indistinguishable during the two balanced reward conditions, but imbalanced rewards (when one side provided more reward than the other, upon the correct response) biased the animals' choices strongly toward the response associated with the higher reward (Feng et al. 2009). This adaptive behaviour on perceptual decision-making tasks with unbalanced rewards was conceptualised in the framework of the leaky competing accumulator model (Usher and McClelland 2001), in which stimulus accumulators for each alternative choice accumulate noisy sensory information (Gao, Tortell, and McClelland 2011). In these models, the subjects bias their choices toward the option with larger rewards, and the extent of the bias depends on the strength of the sensory evidence. Future work is needed to see whether such a framework can explain our findings, where the reward is balanced and the sensory statistics, instead, need to be integrated to optimise final payoffs.

Trial-to-trial learning differs across species

Despite the overall similarity in biasing decisions toward the over-represented category around the boundary, the trial-to-trial strategy updates, where animals update their decision based on the history of past sensory and choice events, differed across species.

Rats and mice show an overall win-stay choice bias, with rats showing larger bias if previous trial was rewarded on hard stimuli (around the boundary), or current stimuli were harder. This suggests that rats may incorporate decision confidence more strongly than mice (Lak et al. 2020). Humans, on the other hand, show a more complex learning update which is indicative of learning detailed sensory statistics.

To understand what computational mechanisms support learning of the underlying statistics we propose two models: a Stimulus-category model and a Boundary-estimation model. In the Boundary-estimation, the agent learned only a belief about the location of the boundary, with no explicit representation of the statistics within each category. The decision was based on the location of the present stimulus with respect to the estimated boundary. The belief about the boundary was then updated based on the feedback received from this action. In the Stimulus-category model, the agent learned an explicit representation of the prior probabilities that random samples are drawn from. The decision was made based on the proportion of the amplitude of each category's probability at the present stimulus. The agent took the corresponding action assigned to the category with larger amplitude. The belief about the prior probability for that category was then updated based on the feedback received. These two models represent different classes of computation. As such, comparing their ability to fit the observed data is critical. The Bayesian model is characterised as "generative" (Jebara 2012) as it aims to explicitly learn how each category is shaped, whilst the Boundary learner is an example of a "discriminative" model as it learns to draw a boundary in the data space without explicitly learning stimulus distributions within each category. It is noteworthy to mention that the way the boundary is learned is still contingent on the distribution of stimuli.

We show that, when fit to individual trial performance on one 'statistical context', both of these two models were able to produce the observed shift in the psychometric performance when tested on a new 'statistical context'. However, these models predicted divergent updating strategies, (when the animals' choice changes depending on previous trial's stimulus and choice) (Lak et al. 2020). The boundary model, fit to individual trial performance, predicted trial-to-trial updates that only depended on the category membership of the previous trial, and updated similarly independent of where in the category the previous stimuli fell. The stimulus-category model, on the other hand, predicted updates that depended on the

exact value of the previous input within a category. In other words, these two models incorporate differing knowledge of the underlying distribution in their trial-to-trial choice updates – the Bayesian (stimulus-category) model carries details of the category shape over trials, while the Boundary model only carries the knowledge about the category membership. Only addition of a strong repulsion term to the boundary model could recapitulate some of non-uniformity in update observed in the stimulus-category model.

When comparing these models against the experimental data we found that while human data can best be explained by a generative stimulus-category model, the rodent data is best explained by a discriminative boundary-estimation model. Disentangling these classes of computational models is crucial as they may offer different predictions about the underlying brain activity.

Sound categorisation is the result of a sensorimotor mapping, where sounds belonging to one category are mapped to a specific motor act. The observed shift in the psychometric performance can be either via changes at the level of the input of the sound-to-action mapping, by modifying sensory representations (consistent with generative account), or by changing the mapping, whereby the input remains the same (consistent with discriminative account).

While our models offer a fit to the data, this does not necessarily preclude the possibility that other approaches might also be able to explain the present data. Future research will be needed to examine the full range of possible alternative models.

Methods

Human subjects

The study was approved by the University College London (UCL) Research Ethics Committee [16159/001] (London, UK). Electronic consent was obtained in accordance with the approved procedures. In all experiments participants were recruited using the online platform Prolific (*www.prolific.co* 2023). Experiments were run online using either Psychopy (Peirce et al. 2019) or Gorilla (*www.gorilla.sc* 2023) using either their APIs or custom code in JavaScript. Participants were paid a fixed amount for their time as well as a bonus payment based on their performance to motivate their continuous attention. Once a participant had completed an experiment they were excluded from any further experiments, and only participants with a high completion rate on Prolific were allowed to start the experiment. Participant demographics were matched to those of Prolific's available set of participants.

Rat subjects

All experiments were performed in accordance with the UK Home Office regulations Animal (Scientific Procedures) Act 1986 and the Animal Welfare and Ethical Review Body (AWERB). Subjects were adult male Lister Hooded rats, which were kept on a reversed 12h light-dark cycle. Rats were water-restricted and received water reward while performing the behavioural task. A total of 18 were used in the behavioural experiments.

Mouse subjects

All experiments were performed in accordance with the UK Home Office regulations Animal (Scientific Procedures) Act 1986 and the Animal Welfare and Ethical Review Body (AWERB). Two lines of mice were used in this study: CA/C57BL/6J and TL/VGAT-ChR2-EYFP. Only male mice were included in the study. The mice arrived in the experimental holding room at 6 weeks of age. Prior to the start of the experimental procedures, mice were allowed a 5-day acclimation period during which they were handled at least once a day to habituate them to the experimenter. The mice were kept on a 12-hour reverse light/dark cycle (8am to 8pm, with one hour of 'twilight' between 7am/pm and 8am/pm) during the entire experimental

period. All experiments were conducted during the 'dark' phase. The mice had free access to water and food in the form of chow, unless specified otherwise. Environmental enrichment was provided in the mouse cages as a refinement to prevent boredom and fighting between cage-mates and to provide shelter. During the experiment, to increase the motivation of the mice to perform the behavioural task, they were water deprived such that they would only receive around 1.25-1.5mL of water a day (either from performing the behavioural task or via additional supplementation provided by the experimenter). The mice's weights were monitored carefully to ensure that their weight would never fall below 75% of their baseline body weight. If a mouse's weight fell below 80% of their baseline body weight or showed signs of dehydration or discomfort, supplementary water would be provided until their weight exceeded 80% of their baseline body weight or they recovered a healthy phenotype, respectively.

Surgery

All animals were allowed one week of acclimation in the experimental unit before any surgical procedure was undertaken. Surgical procedures were never carried out in conjunction with other procedures (in our case, water restriction) and prior to surgery, mice were checked for signs of good health and a weight of at least 18g. All surgeries were performed under aseptic conditions and in accordance with the Project License and Home Office regulations. For induction and maintenance of anaesthesia, isoflurane gas was used (4% for induction, 1-2% for maintenance) as it allowed for the most reliable and direct monitoring of depth of anaesthesia. Metacam, a non-steroidal anti-inflammatory, was provided subcutaneously for preoperative analgesia (2.5µL/kg) and orally for post-operative analgesia. Eye ointment was provided at the start of the surgery to keep the animal's eyes lubricated, and the animal was placed on a heat pad to maintain its body temperature. A rectal probe was used to monitor the mouse's temperature. In most conditions, breathing rate was kept between 60-100Hz and body temperature between 36.5-37°C. To allow for the stable attachment of a stainless-steel head-bar to the animal's head, the mouse was fixed under anaesthesia in a stereotaxic frame between two ear-bars. The hair on top of the head was shaved and cleaned with an antibacterial skin cleaner to avoid any risk of infection. The skin was cut to expose the skull and cleaned of any blood present. The skull was levelled such that the difference between bregma and lambda was less than 20µm. The head-bar was placed on the skull using a custom 3D-printed head-bar holder at AP -2.95. Dental cement was used to attach the head-bar to the skull and seal the gap between the exposed skull and the skin. Once the surgery was completed, the animal was weighed again and placed in a recovery cage on a heated-rack with water, food, and a wet mash diet. A full recovery was usually observed within 30 minutes. Animals were monitored for another 3 days by recording their weight and any signs of discomfort, distress, or pain. If after three days the animal was showing all signs of a complete and successful recovery, procedures such as water restriction could be undertaken to prepare for behavioural training.

Human behaviour

In the human version of the task, on each trial, subjects were presented with a pure tone sound stimulus for a short interval and were then asked to choose between two options, A or B, by clicking on the screen. Following their choice, they were presented with written feedback informing them whether their choice was correct or wrong. The correct choice for each trial depended solely on the stimulus frequency such that low-frequency sounds belonged to category A whereas high-frequency sounds belonged to category B. The subjects were given no extra information and were only asked to pay attention to the sound and to click in one of the boxes (figure **1** panel C). The first 30 trials were defined as learning trials and determined whether or not the subject was allowed to continue the experiment. These initial 30 trials were then excluded from all the following analyses. The subjects that passed this initial phase were subject to another 1000 trials. Importantly, the sound stimulus on each trial was independent of other trials but was always drawn from the same stimulus distribution for the entire experiment – and thus, for each subject.

Rat behaviour

Rats were trained to perform an auditory two-alternative forced choice (2-AFC) task. Subjects were placed in operant training boxes, with three straight walls and one curved wall, equipped with three nose-poke ports arranged side-by-side. All ports had an LED and an infrared beam to detect rat nose poking. The side ports had speakers to deliver auditory stimuli and water spouts to deliver water reward. An LED in the centre port indicated that a trial could be freely initiated by nose-poking. The rats were trained to fixate in the centre port for up to 1.2 s, until an auditory go cue (a 6-kHz pure tone for 200 ms) signalled the end of fixation. During the fixation period, an auditory stimulus was presented for 400 ms, followed by a short delay (300 ms). At the end of the go cue, the animal could withdraw from the centre port and respond by poking in one of the two side ports to obtain water reward. A correct action was signalled by the led turning in the side port, 1.5 s clicking sound and water delivery through the spout. An incorrect action led to no water delivery and a 5 s timeout before the next trial became available. The auditory stimuli were broadband noise (2000-20,000 Hz) produced as varying sound pressure level (SPL) values sampled continuously from a fixed range of intensity between 60-95 dB. Intensity of sounds was measured and calibrated across training boxes using a sound pressure metre. The task required rats to learn to categorise the intensity of the sounds into two arbitrary categories. The decision boundary separating the two categories was set to be the median of the overall range of sound intensities. If the sound intensity was higher than the decision boundary the correct action was to poke into the left-hand nose port, whereas if the intensity was lower than the decision boundary, the correct action was to poke into the right-hand port. A separate cohort of rats was trained to use the opposite rule to map stimulus intensity and choice of side-port for counterbalancing purposes (high intensity sounds associated with poking on the right port and low intensity sounds with poking on the left port). Behavioural sessions lasted 2-4 hours. Rats performed an average of 217 valid trials in each session.

Rats were initially trained on auditory stimuli sampled from one of the three distinct distributions. Upon reaching stable performance, on separate sessions the sampling distribution was switched to the distribution with opposite asymmetry.

Mouse behaviour

Behavioural Apparatus

The behaviour was controlled using a combination of hardware and software. A HARP Sound Card (OEPS-7265) was used to control the sounds played through speakers. These speakers were designed and soldered in house. The sounds included both pure tones and white noise. A HARP Behaviour Board (OEPS-7240) was used to control the solenoid valves that enabled water release. A capacitive board (Adafruit 485-136), connected to the HARP Behaviour Board, enabled monitoring of the mice's licking. The entire head-fixation unit was essentially replicated from the IBL, with the key feature that the wheel was replaced with a custom-designed 3D printed paw support platform. The hardware for holding the water-spouts was also designed and 3D printed in-house. This allowed for a precise control of the location of the water spouts. Video data was acquired using CMOS infrared cameras. These cameras were placed in the experimental boxes to record the mice's behaviour during the task. Behavioural protocols were run using Bonsai-RX (Lopes et al. 2015), a visual reactive programming language. This allowed for precise control of the task parameters and the ability to easily modify the protocols as needed.

Behaviour

The behavioural task used in this study was a sound categorisation task. The mice were head-fixed and heard a stimulus, which could be either a pure tone or white noise, on either side of a boundary situated in the middle of the stimulus range. The mice were required to report whether that stimulus was higher or lower than the boundary by licking the right or left water spout, respectively. For pure tones, the feature the mice had to discriminate was frequency, whereas for white noise

the feature they had to discriminate was amplitude. The utilisation of both in separate cohorts of mice allowed us to verify that the effect observed was not specific to frequency and generalised across auditory modalities. The frequency range used for this experiment was 4-16kHz, and the amplitude range was 45 to 85dB. For the purposes of explanation, the lowest point of the range will be referred to as -1, the highest point of the range was referred to as 1, and the mid-point as 0. Training was performed once a day, with each session lasting between 30 minutes and 1 hour. A shaping protocol was used to teach the mice how to perform the task. The protocol was composed of the following steps:

1. In Stage 1, mice were placed in the setup and were presented with either of 2 stimuli at the extremes of the range (-1 or 1) every 10-15 seconds. Each sound was followed by the release of 5µL of water on the corresponding side (the left spout for -1, the right spout for 1). This stage typically lasted one day.
2. In Stage 2, mice were free to lick either the left or the right spout. Each lick would trigger the playing of the stimulus corresponding to the spout chosen by the mouse (-1 for the left spout, 1 for the right spout) followed by the release of 5µL of water from that spout. This stage also usually lasted one day.
3. Stage 3 was similar to Stage 2, except this time the mice had to lick one side at least 3 times in order to be able to switch to the other side. This imposed pattern was designed to prevent the development of early side-biases by the mouse. Stage 3 usually lasted 2 days.
4. In Stage 4, the mouse now had to perform the task with 2 stimuli, (-1 and 1). On any given trial, it would hear a random stimulus, followed by a 300ms gap of silence (the 'Go-Cue'), and then had a 10-second response window to lick either the left or the right spout. Correct licks were immediately rewarded with 5µL of water. Incorrect licks were punished with a 3-second 'timeout'. If the mouse did not lick within the response-window, a new trial would be presented. Trials were separated with a 2-second inter-trial-interval (ITI). Therefore, incorrect licks effectively doubled the duration of the ITI. Within 6-10 days, animals usually reached >80% performance on this stage.
5. Once threshold performance was reached on Stage 4, mice proceeded to Stage 5, where they were presented with the entire continuous range of stimuli sampled between -1 and 1. The same reward and punishment rules as in Stage 4 applied. Once mice reached >80% performance on stage 5, various distributions of stimuli were presented as described above.

On this final stage, mice performed an average of 352 valid trials in each session.

The uniform distribution was generated by sampling randomly and continuously from the entire range of stimuli (-1 to 1). Both asymmetric distributions were generated using a truncated exponential distribution, sampling either from $p(x;\lambda) = \lambda \exp^{\lambda x} + k$ between -1 and 0 for the hard-A asymmetric distribution, or from $p(x;\lambda) = \lambda \exp^{-\lambda x} + k$ between 0 and 1 for the hard-B asymmetric distribution, where k renormalises due to the truncation. The parameter λ was chosen such that sampling for the asymmetric distribution at the boundary was twice as common compared to that of the uniform distribution. Samples were generated either side of the boundary with equal probability.

Psychometric curves

Psychometric plots (as shown in Fig. **1** Fig. **3** Fig. **4**) show the probability of the subject choosing category B as a function of sensory stimulus, which is parameterised as the distance from the boundary (middle of the auditory range). The data were fit with a 4-parameter psychometric function using the *scipy.optimize.curve_fit* function from Python's *SciPy* library (Virtanen et al. 2020). This was performed by applying a non-linear least squares algorithm to find the best fit parameters that minimised the difference between the observed data and the predicted data generated by the psychometric function. We used the following equation

$$R(B) = l_A + (1 - l_A - l_B) * \Phi(s;\mu,\sigma) \tag{1}$$

where $R(B)$ is the ratio the subject chose category B for a given stimulus value s. l_A and l_B are lapse rates for categories A and B respectively. Φ is the cumulative normal distribution, parameterised by μ, the mean, and σ, the standard deviation of the underlying distribution. For these psychometric curves μ has the interpretation of $R(B)$ for a stimulus value of 0, while σ describes the slope of the psychometric, measuring the overall categorisation behaviour of the subject. The smaller σ, the steeper the slope, indicating better discrimination between category A and B.

Optimal agent

We first estimated 4 parameters of the psychometric function as described in Sect. , for each experimental subject. We then simulated an agent that has the same psychometric parameters to our experimental subjects (rats, mice and humans, separately), except for the mean of the psychometric curve. We assumed a fixed psychometric curve for this agent – unlike humans and animals who might change their psychometric response from trial to trial. We imposed the same values observed experimentally to three parameters of the psychometric curve: lapses 1 and 2 (category A and category B lapses), and the slopes of the psychometric curves (*sigma*). We assumed that these three parameters are set by the conditions of the experiment and other variables outside of the control of the subjects (internal noise and sensory accuracy). The remaining parameter, the mean of the psychometric curve (*mu*) could be adapted to improve performance. The performance of the simulated agent, therefore, was completely determined by the parameter *mu* and the distribution of stimuli. We calculated the best *mu* value that maximised the agent's performance on each distribution.

Computational models

Stimulus-category model

In the stimulus-category model, agents form two representations, one for each category. The representation is defined by a probability distribution of a stimulus s belonging to the corresponding category, $p_A(s)$ and $p_B(s)$. On each trial t, a stimulus-category model agent is presented with a stimulus s_t and outputs a binary decision, A or B. Before making the decision, we assume that the agent receives a noisy version of the stimulus by adding an extra term to s_t. This extra term, ξ_{percep}, is drawn from a normal distribution with mean zero and variance σ^2_{percep}. The noisy stimulus is then given by $\tilde{s}_t = s_t + \xi_{percep}$. Finally, this noisy stimulus is then repelled by the perceived stimulus at trial $t-1$ such that the perceived stimulus at trial t is given by $\hat{s}_t = \tilde{s}_t + \Delta_{repulsion}$, where the repulsion term is given by

$$\Delta_{repulsion} = A_{repulsion}\left(\tilde{s}_t - \hat{(s)}_{t-1}\right)\exp\left[-|\tilde{s}_t - \hat{s}_{t-1}|\right] \qquad , \tag{2}$$

where $A_{repulsion}$ is the repulsion amplitude. In simulations without repulsion $A_{repulsion}$ is set to 0. The stimulus-category model then outputs decision A with probability

$$P_A = \frac{p_A(\hat{s}_t)}{p_A(\hat{s}_t) + p_A(\hat{s}_t)} \qquad , \tag{3}$$

and decision B with probability

$$P_B = \frac{p_B(\hat{s}_t)}{p_A(\hat{s}_t) + p_A(\hat{s}_t)} \qquad . \tag{4}$$

On each trial, there is a probability p_{lapse} that the agent will choose the category randomly with probability $1/2$. Following the model decision and feedback, the agent's representation of the correct category is updated following

$$p_C^{t+1}(s) = \gamma p_C^t(s) + (1-\gamma)\, g(s|\hat{s}_t, \sigma_{update}) \qquad , \tag{5}$$

where $C \in A, B$ is the correct category at trial t, γ is a learning parameter, and $g(s|\hat{s}_t, \sigma_{update})$ is a Gaussian function centred at $\hat{s}_t$ and standard deviation σ_{update}.

Boundary-estimation model

In the boundary-estimation (BE) model, agents form a representation of the boundary by defining a probability distribution of the boundary position as a function of the stimulus value. The boundary model is thus defined by the probability distribution, $B_t(s)$, over stimulus s at trial t. In our model, we assume that this probability distribution is updated by two terms: a feedback-dependent term (or learning term), $\Delta_{learning}$, and a relaxation term, Δ_{relax}. The learning term is given by

$$\Delta_{learning}(s) = \frac{1}{1+\exp\left(-\sigma_{boundary} C\left(s-\hat{s}_t\right)\right)} \quad , \tag{6}$$

where $\sigma_{boundary}$ is related to the precision of the boundary update (large values mean precise update), C represents the stimulus category ($C = 1$ for category B and $C = -1$ for category A), and $\hat{s}_t$ is the stimulus perceived at trial t. Similarly to the stimulus-category model, this stimulus is given by a combination of noise and repulsion from the stimulus on the previous trial: $\hat{s}_t = \tilde{s}_t + \Delta^{BE}_{repulsion}$, where $\Delta^{BE}_{repulsion}$ is a repulsion term given by

$$\Delta^{BE}_{repulsion} = A^{BE}_{repulsion}\left(\tilde{s}_t - \hat{(s)}_{t-1}\right)\exp\left[-\left|\tilde{s}_t - \hat{s}_{t-1}\right|\right] \quad , \tag{7}$$

and $\hat{s}_t = s_t + \xi^{BE}_{percep}$, where s_t is the true stimulus at trial t and ξ^{BE}_{percep} is a random variable drawn from a normal distribution with mean zero and variance $\sigma^2_{percep,BE}$. $A^{BE}_{repulsion}$ is a repulsion amplitude. The post-feedback boundary model is thus given by

$$B'_{t+1}(s) = B_t(s) - \eta_{learning}\Delta_{learning} \quad , \tag{8}$$

where $\eta_{learning}$ is the learning rate for the feedback-dependent term. The relaxation term then "pushes" this post-learning boundary model back towards a uniform distribution and acts as a forgetting term. The amplitude of this relaxation term is given by the distance between the current model and the uniform distribution

$$\Delta_{relax}(s) = B'_{t+1}(s) - \frac{1}{2} \quad , \tag{9}$$

and the boundary model at trial $t+1$ is given by

$$B_{t+1}(s) = \left(B'_{t+1}(s) - \eta_{relax}\Delta_{relax}\right)/norm \quad . \tag{10}$$

Model fitting

Both models were fit to data by maximum likelihood estimation (finding the parameters θ that best explained the data, given the model). We estimated the log-likelihood of the model parameters given the data for a random selection of subjects trained on the uniform distribution. Given N responses $\{y_i\}_{i=1}^{i=N}$ from subjects and N predictions $\{\pi_i\}_{i=1}^{i=N}$ from a single simulation from a given model, the log-likelihood was defined as

$$\mathscr{L}(\theta;\{y_i\}) = \frac{-1}{N}\sum_{i=1}^{N} y_i \log(\pi_i) + (1-y_i)\log(1-\pi_i) \tag{11}$$

The optimal parameters are found via a grid search within a domain of plausible parameter values using the fmin function included in the hyperopt library (Bergstra, Yamins, and Cox 2013).

Supplementary Figures

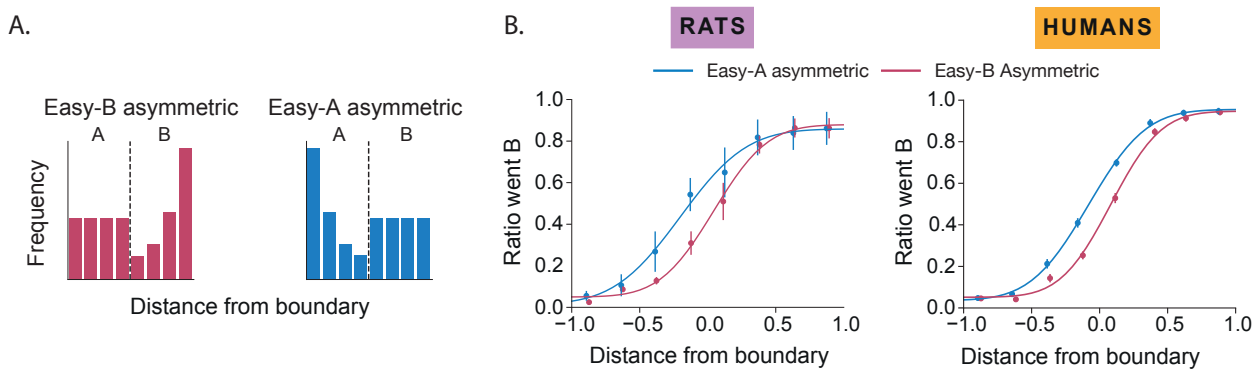

Figure S1: **Biases observed in humans and rats are solely due to the relative difference in category representation around the boundary irrespective of the shape of the non-uniform category**. **A**: Example alternative context distributions, with the respective psychometric curves (for rats and humans) shown in **B**. The shift in psychometric depends only on the local distribution of stimuli around the boundary (i.e., which category has greater mass), not the overall shape of the sampling distribution on either side of the boundary.

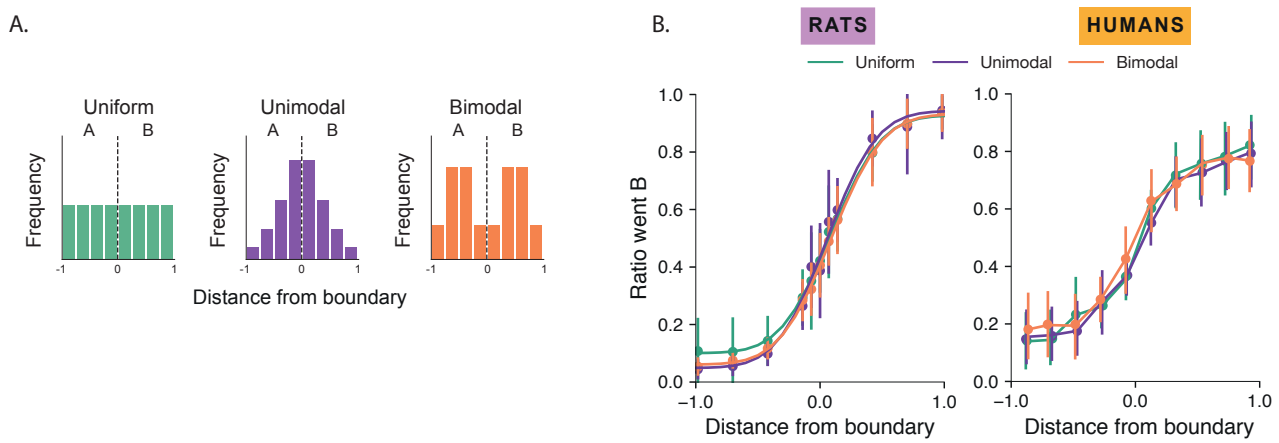

Figure S2: **Over- or under-exposure to hard stimuli did not alter human or rat performances**. Contexts with symmetric distributions of stimuli lead to oversampling of either easy (Bimodal) or difficult (unimodal, see panel **A**) simuli, but this does not vary overall performance in mice or rats.

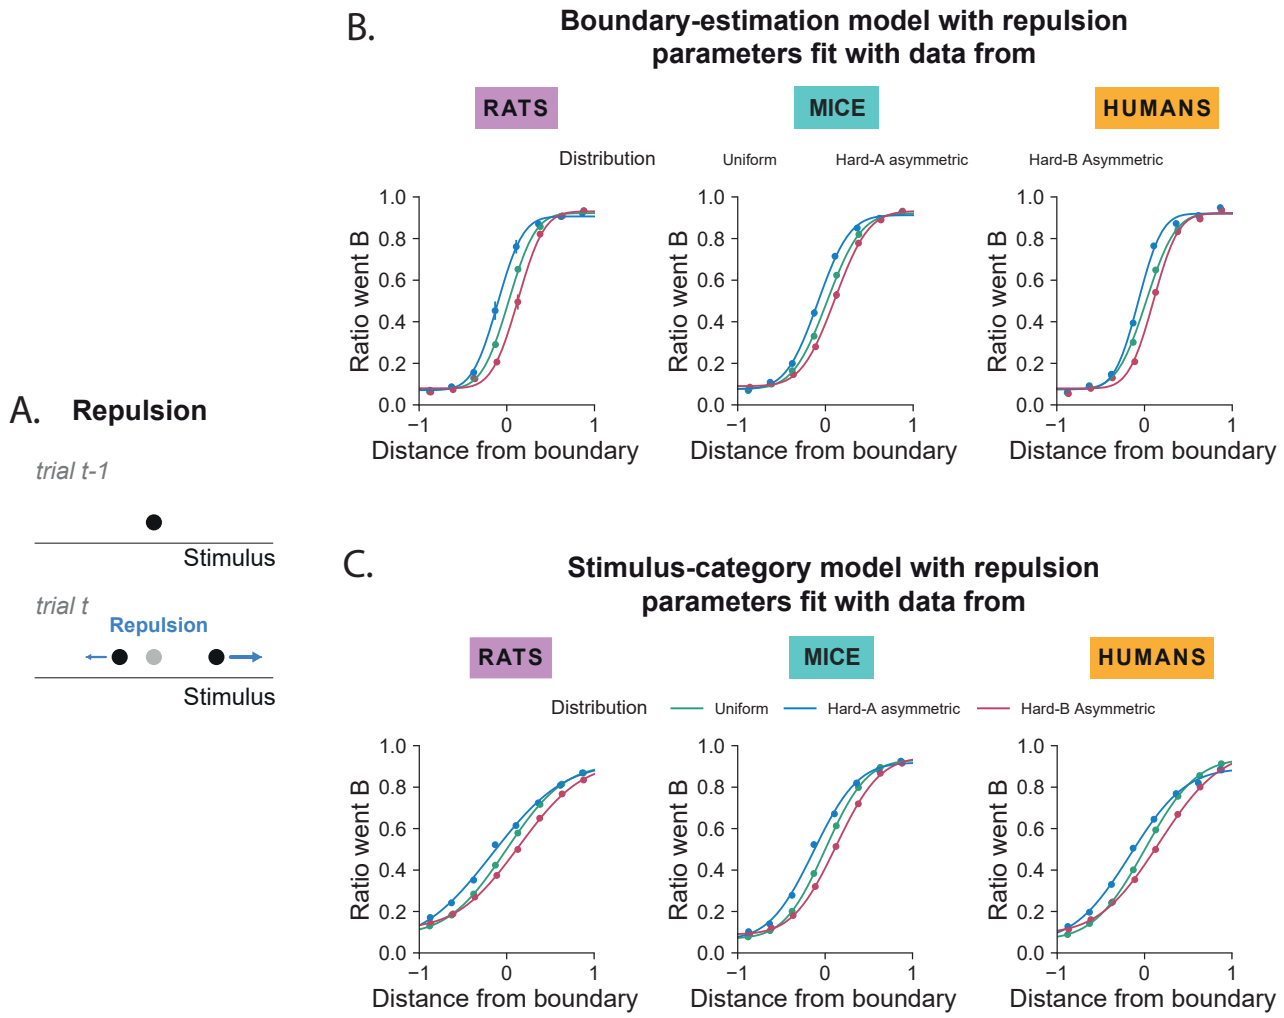

Figure S3: **With stimulus repulsion, Boundary-estimation and Stimulus-category models have similar predictions for psychometric performance**. **A**: Schematic of repulsion mechanism. Panels **B** and **C** show model predictions for psychometric curves for each species produced by the Boundary-estimation and Stimulus-category models, respectively.

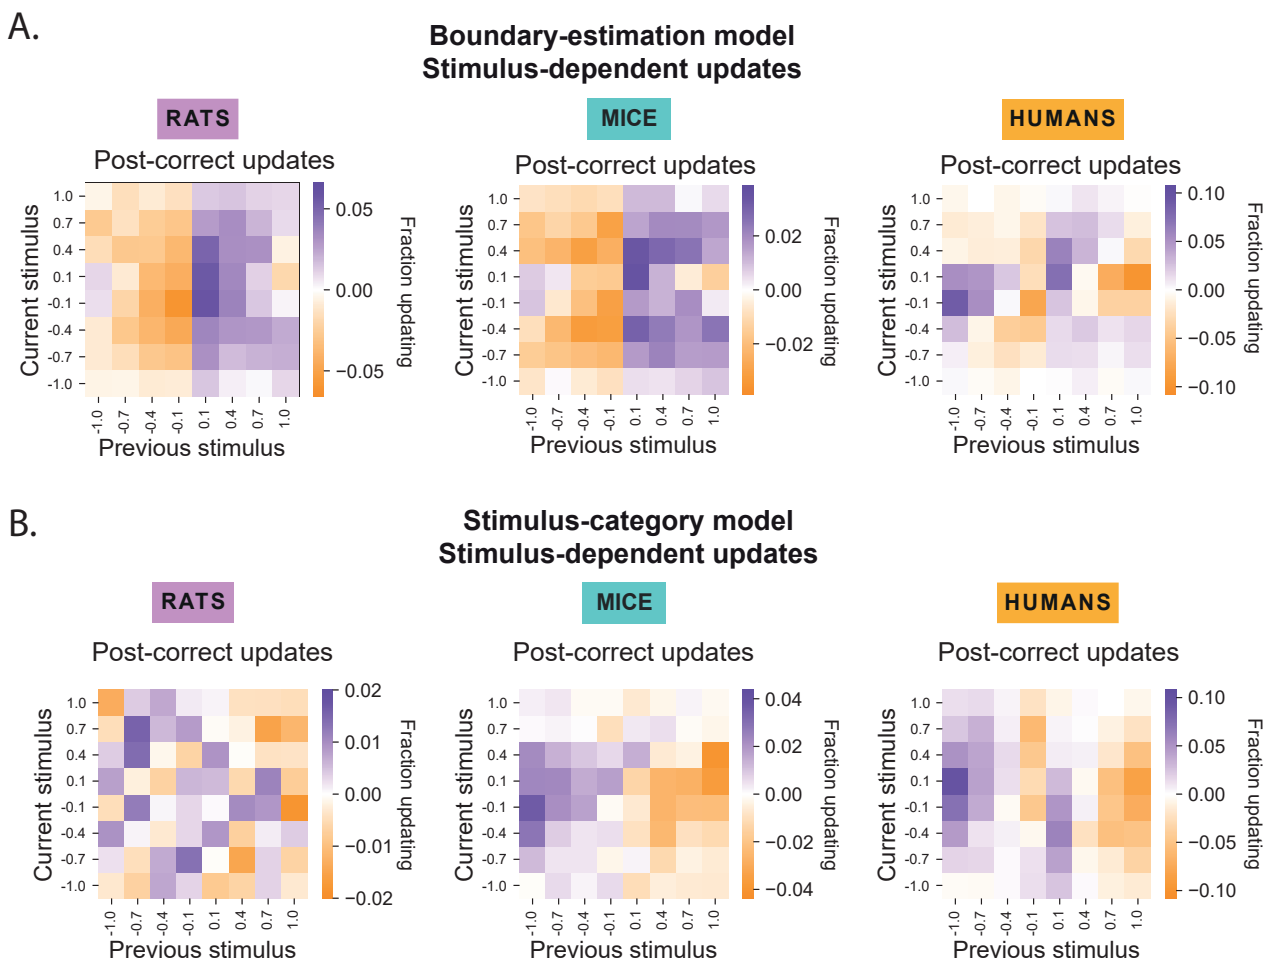

Figure S4: **With stimulus repulsion, Boundary-estimation and Stimulus-category models have divergent predictions for trial-to-trial updates** Panels **A** and **B** show model predictions for update matrices for each species produced by the Boundary-estimation and Stimulus-category models, respectively, with repulsion.

References

Abrahamyan, Arman et al. (June 2016). "Adaptable history biases in human perceptual decisions". In: *Proceedings of the National Academy of Sciences* 113.25. Publisher: Proceedings of the National Academy of Sciences, E3548–E3557. DOI: 10.1073/pnas.1518786113. URL: https://www.pnas.org/doi/full/10.1073/pnas.1518786113 (visited on 01/30/2023).

Akrami, Athena et al. (Feb. 2018). "Posterior parietal cortex represents sensory history and mediates its effects on behaviour". en. In: *Nature* 554.7692. Number: 7692 Publisher: Nature Publishing Group, pp. 368–372. ISSN: 1476-4687. DOI: 10.1038/nature25510. URL: https://www.nature.com/articles/nature25510 (visited on 01/20/2023).

Anderson, John R. (1991). "The adaptive nature of human categorization". In: *Psychological Review* 98. Place: US Publisher: American Psychological Association, pp. 409–429. ISSN: 1939-1471. DOI: 10.1037/0033-295X.98.3.409.

Ashwood, Zoe C. et al. (Feb. 2022). "Mice alternate between discrete strategies during perceptual decision-making". en. In: *Nature Neuroscience* 25.2. Number: 2 Publisher: Nature Publishing Group, pp. 201–212. ISSN: 1546-1726. DOI: 10.1038/s41593-021-01007-z. URL: https://www.nature.com/articles/s41593-021-01007-z (visited on 01/28/2023).

Aslin, Richard N. and Elissa L. Newport (June 2012). "Statistical Learning: From Acquiring Specific Items to Forming General Rules". en. In: *Current Directions in Psychological Science* 21.3. Publisher: SAGE Publications Inc, pp. 170–176.

ISSN: 0963-7214. DOI: 10.1177/0963721412436806. URL: https://doi.org/10.1177/0963721412436806 (visited on 01/19/2023).

Bergstra, James, Daniel Yamins, and David Cox (Feb. 2013). "Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures". en. In: *Proceedings of the 30th International Conference on Machine Learning*. ISSN: 1938-7228. PMLR, pp. 115–123. URL: https://proceedings.mlr.press/v28/bergstra13.html (visited on 01/30/2023).

Berkes, Pietro et al. (Jan. 2011). "Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment". eng. In: *Science (New York, N.Y.)* 331.6013, pp. 83–87. ISSN: 1095-9203. DOI: 10.1126/science.1195870.

Britten, K. H. et al. (Dec. 1992). "The analysis of visual motion: a comparison of neuronal and psychophysical performance". en. In: *Journal of Neuroscience* 12.12. Publisher: Society for Neuroscience Section: Articles, pp. 4745–4765. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.12-12-04745.1992. URL: https://www.jneurosci.org/content/12/12/4745 (visited on 01/28/2023).

Britten, Kenneth H. et al. (Nov. 1993). "Responses of neurons in macaque MT to stochastic motion signals". en. In: *Visual Neuroscience* 10.6. Publisher: Cambridge University Press, pp. 1157–1169. ISSN: 1469-8714, 0952-5238. DOI: 10.1017/S0952523800010269. URL: https://www.cambridge.org/core/journals/visual-neuroscience/article/abs/responses-of-neurons-in-macaque-mt-to-stochastic-motion-signals/C47F087B4BE2FBB6FDE7FC602BE42BDC (visited on 01/30/2023).

Busse, Laura et al. (Aug. 2011). "The Detection of Visual Contrast in the Behaving Mouse". en. In: *Journal of Neuroscience* 31.31. Publisher: Society for Neuroscience Section: Articles, pp. 11351–11361. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.6689-10.2011. URL: https://www.jneurosci.org/content/31/31/11351 (visited on 01/30/2023).

De Ridder, Dirk, Sven Vanneste, and Walter Freeman (July 2014). "The Bayesian brain: Phantom percepts resolve sensory uncertainty". en. In: *Neuroscience & Biobehavioral Reviews*. Applied Neuroscience: Models, methods, theories, reviews. A Society of Applied Neuroscience (SAN) special issue. 44, pp. 4–15. ISSN: 0149-7634. DOI: 10.1016/j.neubiorev.2012.04.001. URL: https://www.sciencedirect.com/science/article/pii/S0149763412000607 (visited on 01/30/2023).

Destrebecqz, Arnaud and Axel Cleeremans (2001). "Can sequence learning be implicit? New evidence with the process dissociation procedure". In: *Psychonomic Bulletin & Review* 8. Place: US Publisher: Psychonomic Society, pp. 343–350. ISSN: 1531-5320. DOI: 10.3758/BF03196171.

Fassihi, Arash et al. (2014). "Tactile perception and working memory in rats and humans". In: *Proceedings of the National Academy of Sciences* 111.6. Publisher: National Acad Sciences, pp. 2331–2336.

Feng, Samuel et al. (Feb. 2009). "Can Monkeys Choose Optimally When Faced with Noisy Stimuli and Unequal Rewards?" en. In: *PLOS Computational Biology* 5.2. Publisher: Public Library of Science, e1000284. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.1000284. URL: https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1000284 (visited on 01/30/2023).

Fritsche, Matthias, Eelke Spaak, and Floris P De Lange (2020). "A Bayesian and efficient observer model explains concurrent attractive and repulsive history biases in visual perception". In: *Elife* 9. Publisher: eLife Sciences Publications Limited, e55389.

Ganguli, Deep and Eero P. Simoncelli (Oct. 2014). "Efficient sensory encoding and Bayesian inference with heterogeneous neural populations". eng. In: *Neural Computation* 26.10, pp. 2103–2134. ISSN: 1530-888X. DOI: 10.1162/NECO_a_00638.

Gao, Juan, Rebecca Tortell, and James L. McClelland (Mar. 2011). "Dynamic Integration of Reward and Stimulus Information in Perceptual Decision-Making". en. In: *PLOS ONE* 6.3. Publisher: Public Library of Science, e16749. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0016749. URL: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0016749 (visited on 01/23/2023).

Green, D. M. and J. A. Swets (Jan. 1966). *Signal detection theory and psychophysics.* NTRS Author Affiliations: NTRS Document ID: 19670040997 NTRS Research Center: Legacy CDMS (CDMS). URL: https://ntrs.nasa.gov/citations/19670040997 (visited on 01/30/2023).

Hachen, I. et al. (2020). "Dynamics of history-dependent perceptual judgment". In: *bioRxiv*, p. 2020.07.12.199489. DOI: 10.1101/2020.07.12.199489.

Hausberger, Martine and Hugo Cousillas (Dec. 1995). "Categorization in birdsong: from behavioural to neuronal responses". en. In: *Behavioural Processes*. Cognition and Evolution 35.1, pp. 83–91. ISSN: 0376-6357. DOI: 10.1016/0376-6357(95)00053-4. URL: https://www.sciencedirect.com/science/article/pii/0376635795000534 (visited on 01/30/2023).

Jaramillo, Santiago and Anthony M. Zador (Feb. 2011). "The auditory cortex mediates the perceptual effects of acoustic temporal expectation". en. In: *Nature Neuroscience* 14.2. Number: 2 Publisher: Nature Publishing Group, pp. 246–251. ISSN: 1546-1726. DOI: 10.1038/nn.2688. URL: https://www.nature.com/articles/nn.2688 (visited on 01/30/2023).

Jebara, Tony (Dec. 2012). *Machine Learning: Discriminative and Generative*. en. Google-Books-ID: g5rSBwAAQBAJ. Springer Science & Business Media. ISBN: 978-1-4419-9011-2.

Kelly, S.P. and R.G. O'Connell (2013). "Internal and external influences on the rate of sensory evidence accumulation in the human brain". English. In: *Journal of Neuroscience* 33.50, pp. 19434–19441. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.3355-13.2013.

Knill, David C. and Alexandre Pouget (Dec. 2004). "The Bayesian brain: the role of uncertainty in neural coding and computation". eng. In: *Trends in Neurosciences* 27.12, pp. 712–719. ISSN: 0166-2236. DOI: 10.1016/j.tins.2004.10.007.

Lak, Armin et al. (2020). "Confidence-guided updating of choice bias during perceptual decisions is a widespread behavioral phenomenon". In: *eLife* 9, e49834. DOI: 10.7554/elife.49834.

Lee, Jung Hoon et al. (2009). "Prefrontal activity predicts monkeys' decisions during an auditory category task". In: *Frontiers in Integrative Neuroscience* 3. ISSN: 1662-5145. URL: https://www.frontiersin.org/articles/10.3389/neuro.07.016.2009 (visited on 01/30/2023).

Liu, Shi Tong et al. (Mar. 2019). "Optimal features for auditory categorization". en. In: *Nature Communications* 10.1. Number: 1 Publisher: Nature Publishing Group, p. 1302. ISSN: 2041-1723. DOI: 10.1038/s41467-019-09115-y. URL: https://www.nature.com/articles/s41467-019-09115-y (visited on 01/30/2023).

Liu, Yanhe, Yu Xin, and Ning-long Xu (2021). "A cortical circuit mechanism for structural knowledge-based flexible sensorimotor decision-making". In: *Neuron*. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2021.04.014.

Lopes, Gonçalo et al. (2015). "Bonsai: an event-based framework for processing and controlling data streams". In: *Frontiers in Neuroinformatics* 9. ISSN: 1662-5196. URL: https://www.frontiersin.org/articles/10.3389/fninf.2015.00007 (visited on 01/30/2023).

Mercado III, Eduardo, Itzel Orduña, and Jeaveen M. Nowak (2005). "Auditory Categorization of Complex Sounds by Rats (Rattus norvegicus)". In: *Journal of Comparative Psychology* 119. Place: US Publisher: American Psychological Association, pp. 90–98. ISSN: 1939-2087. DOI: 10.1037/0735-7036.119.1.90.

Neri, Peter (Dec. 2010). "How inherently noisy is human sensory processing?" en. In: *Psychonomic Bulletin & Review* 17.6, pp. 802–808. ISSN: 1531-5320. DOI: 10.3758/PBR.17.6.802. URL: https://doi.org/10.3758/PBR.17.6.802 (visited on 01/28/2023).

Otazu, Gonzalo H. et al. (May 2009). "Engaging in an auditory task suppresses responses in auditory cortex". In: *Nature neuroscience* 12.5, pp. 646–654. ISSN: 1097-6256. DOI: 10.1038/nn.2306. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4084972/ (visited on 01/30/2023).

Pearson, John M. et al. (Apr. 2011). "Posterior cingulate cortex: adapting behavior to a changing world". en. In: *Trends in Cognitive Sciences* 15.4, pp. 143–151. ISSN: 1364-6613. DOI: 10.1016/j.tics.2011.02.002. URL: https://www.sciencedirect.com/science/article/pii/S1364661311000180 (visited on 01/30/2023).

Peirce, Jonathan et al. (Feb. 2019). "PsychoPy2: Experiments in behavior made easy". en. In: *Behavior Research Methods* 51.1, pp. 195–203. ISSN: 1554-3528. DOI: 10.3758/s13428-018-01193-y. URL: https://doi.org/10.3758/s13428-018-01193-y (visited on 01/30/2023).

Raposo, David et al. (Mar. 2012). "Multisensory Decision-Making in Rats and Humans". en. In: *Journal of Neuroscience* 32.11. Publisher: Society for Neuroscience Section: Articles, pp. 3726–3735. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.4998-11.2012. URL: https://www.jneurosci.org/content/32/11/3726 (visited on 01/30/2023).

Rorie, Alan E. et al. (Feb. 2010). "Integration of Sensory and Reward Information during Perceptual Decision-Making in Lateral Intraparietal Cortex (LIP) of the Macaque Monkey". en. In: *PLOS ONE* 5.2. Publisher: Public Library of Science, e9308. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0009308. URL: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0009308 (visited on 01/30/2023).

Roy, Nicholas A. et al. (2021). "Extracting the dynamics of behavior in sensory decision-making experiments". In: *Neuron* 109.4, 597–610.e6. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2020.12.004.

Russ, Brian E., Yune-Sang Lee, and Yale E. Cohen (July 2007). "Neural and behavioral correlates of auditory categorization". eng. In: *Hearing Research* 229.1-2, pp. 204–212. ISSN: 0378-5955. DOI: 10.1016/j.heares.2006.10.010.

Schapiro, A. and N. Turk-Browne (2015). "Statistical Learning". en. In: *Brain Mapping*. Elsevier, pp. 501–506. ISBN: 978-0-12-397316-0. DOI: 10.1016/B978-0-12-397025-1.00276-1. URL: https://linkinghub.elsevier.com/retrieve/pii/B9780123970251002761 (visited on 01/19/2023).

Sherman, Brynn E., Kathryn N. Graves, and Nicholas B. Turk-Browne (Apr. 2020). "The prevalence and importance of statistical learning in human cognition and behavior". eng. In: *Current Opinion in Behavioral Sciences* 32, pp. 15–20. ISSN: 2352-1546. DOI: 10.1016/j.cobeha.2020.01.015.

Simon, Herbert A. (1986). "Rationality in Psychology and Economics". In: *The Journal of Business* 59.4. Publisher: University of Chicago Press, S209–S224. ISSN: 0021-9398. URL: https://www.jstor.org/stable/2352757 (visited on 01/28/2023).

Simoncelli, E. P. and B. A. Olshausen (2001). "Natural image statistics and neural representation". eng. In: *Annual Review of Neuroscience* 24, pp. 1193–1216. ISSN: 0147-006X. DOI: 10.1146/annurev.neuro.24.1.1193.

Sturdy, Christopher B. et al. (2007). "Auditory category perception as a natural cognitive activity in songbirds". In: *Comparative Cognition & Behavior Reviews* 2. Place: US Publisher: The Comparative Cognition Society, pp. 93–110. ISSN: 1911-4745.

The International Brain Laboratory et al. (May 2021). "Standardized and reproducible measurement of decision-making in mice". In: *eLife* 10. Ed. by Naoshige Uchida and Michael J Frank. Publisher: eLife Sciences Publications, Ltd, e63711. ISSN: 2050-084X. DOI: 10.7554/eLife.63711. URL: https://doi.org/10.7554/eLife.63711 (visited on 01/28/2023).

Usher, Marius and James L. McClelland (2001). "The time course of perceptual choice: The leaky, competing accumulator model". In: *Psychological Review* 108. Place: US Publisher: American Psychological Association, pp. 550–592. ISSN: 1939-1471. DOI: 10.1037/0033-295X.108.3.550.

Virtanen, Pauli et al. (Mar. 2020). "SciPy 1.0: fundamental algorithms for scientific computing in Python". en. In: *Nature Methods* 17.3. Number: 3 Publisher: Nature Publishing Group, pp. 261–272. ISSN: 1548-7105. DOI: 10.1038/s41592-019-0686-2. URL: https://www.nature.com/articles/s41592-019-0686-2 (visited on 01/30/2023).

Wei, Xue-Xin and Alan A. Stocker (Oct. 2015). "A Bayesian observer model constrained by efficient coding can explain 'anti-Bayesian' percepts". eng. In: *Nature Neuroscience* 18.10, pp. 1509–1517. ISSN: 1546-1726. DOI: 10.1038/nn.4105.

www.gorilla.sc (2023). en-US. URL: https://gorilla.sc/ (visited on 01/30/2023).

www.prolific.co (2023). en. URL: https://www.prolific.co/ (visited on 01/30/2023).

Xin, Yu et al. (Sept. 2019). “Sensory-to-Category Transformation via Dynamic Reorganization of Ensemble Structures in Mouse Auditory Cortex”. en. In: *Neuron* 103.5, 909–921.e6. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2019.06.004. URL: https://www.sciencedirect.com/science/article/pii/S089662731930532X (visited on 01/29/2023).

Zhong, Lin et al. (June 2019). “Causal contributions of parietal cortex to perceptual decision-making during stimulus categorization”. eng. In: *Nature Neuroscience* 22.6, pp. 963–973. ISSN: 1546-1726. DOI: 10.1038/s41593-019-0383-6.

Znamenskiy, Petr and Anthony M. Zador (May 2013). “Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination”. eng. In: *Nature* 497.7450, pp. 482–485. ISSN: 1476-4687. DOI: 10.1038/nature12077.