Encoding priors in the brain: A reinforcement learning model for mouse decision making

Sanjukta Krishnagopal, Peter Latham

December 14, 2021

Abstract

In two-alternative forced choice tasks, prior knowledge can improve performance, especially when operating near the psychophysical threshold. For instance, if subjects know that one choice is much more likely than the other, they can make that choice when evidence is weak. A common hypothesis for these kinds of tasks is that the prior is stored in neural activity. Here we propose a different hypothesis: the prior is stored in synaptic strengths. We study the International Brain Laboratory task, in which a grating appears on either the right or left side of a screen, and a mouse has to move a wheel to bring the grating to the center. The grating is often low in contrast which makes the task relatively difficult, and the prior probability that the grating appears on the right is either 80% or 20%, in (unsignaled) blocks of about 50 trials. We model this as a reinforcement learning task, using a feedforward neural network to map states to actions, and adjust the weights of the network to maximize reward, learning via policy gradient. Our model uses an internal state that stores an estimate of the grating and confidence, and follows Bayesian updates, and can switch between engaged and disengaged states to mimic animal behavior. This model reproduces the main experimental finding – that the psychometric curve with respect to contrast shifts after a block switch in about 10 trials. Also, as seen in the experiments, in our model the difference in neuronal activity in the right and left blocks is small – it is virtually impossible to decode block structure from activity on single trials if noise is about 2%. The hypothesis that priors are stored in weights is difficult to test, but the technology to do so should be available in the not so distant future.

1 Introduction

Animals live in a dynamic environment where they constantly make decisions. Presumably, they make these decisions through some mechanism of weighing their options and making a decision that they believe is correlated with a positive reward. Decision making is ubiquitous in the animal kingdom and has been extensively studied in fields including mathematics, psychology, machine learning and neuroscience [1, 2, 3]. However, many of these works make overly simplistic assumptions about information processing models in animals, partly because this is indeed quite difficult to model and has limited experimental support [4]. Understanding how neural systems work together to make decisions, and adapt their behavior in response to feedback is an important open question. Explaining decision making in animals may involve a model of sensory processing, past experience integration, fatigue tracking, aggregation of environmental stimulus, internal and external contexts etc., all of which play a role in decision making. In particular, how animals store priors and exploit priors to make decisions when evidence is weak is a topic of growing interest [5].

In recent years, there have been significant advancements in the tools available to record neural data with unprecedented accuracy [6]. Consequently, there has been a growth in the standardization of experiments and availability of data on animal behavior and decision making tasks. Specifically, International Brain Lab [7] is a global neuroscience collaboration consisting of 22 labs that aggregate standardized data from mice in order to investigate complex adaptive behavior. This opens up possibilities to test theoretical models and validate them against experimental observations. Understanding how an animal makes decisions, specifically what policy it uses, would be very insightful in interpreting animal behavior. Several models [8, 9, 10] have made discoveries on learning and decision making in mice using this dataset.

The field of machine learning, particularly reinforcement learning (RL) [11], has made tremendous advances in recent years, often rivaling human-like performance in decision-making tasks [12, 13]. Indeed, learning in neural pathways is thought to be a consequence of change in synaptic weights through interaction with a dynamic environment with the goal of optimizing some long term objective, which is the fundamental principle of RL [14]. The most general RL framework consists of an agent (animal) interacting with an environment through an action, and receiving a reward as a consequence of their chosen action. The action transitions the agent into a new state

(where a state can simply be the new position of the animal). In this work, we use a type of policy gradient reinforcement learning [15, 16], which aims to maximize the 'expected' reward under a certain policy. Policy gradient algorithms typically learn a parametrized policy (a distribution of probabilities over action space given a state). They have been successfully used in several machine learning tasks [17] as well as explainable neuroscience models for inferring animal learning rules directly from data [10].

In this work, we consider a perceptual decision-making task executed by head-fixed mice [7]. The stimulus is a grating that appears on either the right or the left of the screen with varying contrast (ranging from 0 to 1). Turning the wheel moves the stimulus, and the mouse is rewarded with water if it brings the stimulus to the center of the visual field in less than 60s [18]. The stimulus is presented in 'blocks'; in a block, the probability that the stimulus appears on the right is either 80% or 20%. Block changes are unsignaled. A key experimental feature of mouse behavior is that the mouse learns the prior, in the sense that they are biased towards the block direction, i.e., the mice tend to pick the side that appeared more frequently in recent trials. However, despite behavioral evidence the mouse learn biases in the stimulus direction, it is difficult to decode what block the animal is in from neural activity.. This is an important question in understanding how brain activity evolves in decision making tasking, and important to capture in realistic models.

Contrary to machine learning tasks, animals often have intrinsic states that lead organisms to engage in exploration, play, and other biological behavior [19]. Experimental evidence suggests that attention mechanisms play a key role in shaping ongoing learning in animals [20]. In this work, we use RL to infer the internal model of the mouse through their behavioral data alone. In particular, we introduce two key features to the conventional RL framework: (1) an *internal* state of the agent in the RL model, which stores (a) an estimate of the mouse's perception of the contrast, and (b) confidence in its perception. Contrary to conventional RL models, where the state typically represents the external position, the addition of an internal state is intuitive from a biological perspective. (2) Our agent operates in two attentional modes: attentive (where the mouse updates it internal state through feedback from the environment), and inattentive (no update of the estimate of stimulus, and confidence decreases monotonically). The existence of two states is natural and consistent with experimental results, with the inattentive state encompassing factors such as boredom, fatigue etc.

We find that an RL model that incorporates these features demonstrates

activity patterns largely similar to the mice. This model captures the block-dependent behavior, and, as in experiments, it is hard to decode block from neural activity. This unexpected result matches experimental evidence and provides insight into how priors are encoded by the animal, predicting that priors can also be stored in weights.

2 Methods

2.1 Experiment

We consider a perceptual decision-making task executed by mice in the IBL collaboration (22 theoretical and computational neuroscience labs)[21, 22]. The mouse experiments are carried out through a standardized process in 11 labs. The head-fixed mouse is presented with a visual stimulus and is trained to select one of two choices by turning a steering wheel placed under it. The stimulus is a grating that appears on either the right or the left of the screen with varying contrast (ranging from 0 to 1). Turning the wheel moves the stimulus, and the mouse is rewarded with water if it brings the stimulus to the center of the visual field in less than 60s [18]. The stimulus is presented in 'blocks'; in a block, the probability that the stimulus appears on the right is either 80% or 20%. Block changes are unsignaled. The stimulus side is a Bernoulli process with the parameter 'p' denoting the probability that the stimulus is on the right.

Detailed experimental information and analysis of behavioral features aggregated over several mice (and several trials) is presented in (cite). The key features found in experiments are listed below:

- Accuracy is proportional to contrast, i.e., mouse tend to be more accurate when the presented grid has high contrast.
- Movements elicited by high contrast stimuli typically had shorter trial time and higher peak velocity.
- The mouse learns block biases, i.e., demonstrates bias towards the block direction.
- Block biases are not easily decoded from the neural activity of individual neurons.

2.2 Model

We study a generic two-alternative forced choice (2AFC) task, commonly used for studying decision-making behavior, and use an RL framework to model it. We use a reinforcement learning (RL) framework in which states are mapped to actions via a feedforward network with a single hidden layer; the hidden layer consists of 64 fully connected units implementing ReLU nonlinearities. Time is discretized, and at each time point the input to the network is the state (the mouse's estimate of the contrast, its confidence in that estimate, an attentional variable), and the output is the action (do nothing, or move the wheel left or right).

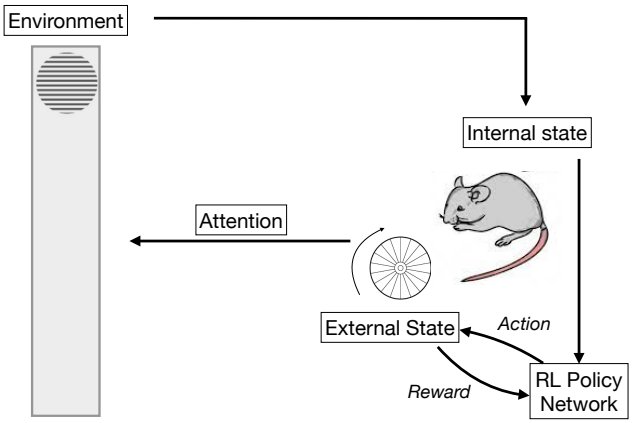

Figure 1: An illustration outlining the model.

Unlike machines, animals have attention spans that can be affected by a variety of factors. In experiments, a lack of attention can stem from distracted, fatigue, boredom etc., and a results in a period of inactivity. It then becomes important to model attention in order to fully explain the mouse behavior. We find, in our work, that attention is a critical component of our model in replicating mouse behavior. In our model, at every time point t, the mouse has two possIBLe attention states: attentive $\alpha_t = 1$, and inattentive $\alpha_t = 0$.

The mouse's estimate of contrast is determined as follows. The actual contrast of the grating, denoted c, lies between -1 and 1, with sign indicating side (negative corresponds to left, positive to right), and magnitude indicating difficulty ($c = 0$ means no grating, so the mouse has to guess); typical

psychometric curves (probability of choosing right versus contrast) are shown in Fig. 5. On each timestep the mouse receives noisy evidence about the true contrast c. That evidence, denoted κ_t, is drawn from a Gaussian distribution centered around the true contrast: $\kappa_t \sim \mathcal{N}(c, \beta^2)$. The mouse's internal state, denoted $s_t \equiv [\mu_t, \sigma_t^2]$, is the estimate for the mean and variance of the contrast given past (and current) input. With a small twist: to model the fact that the mouse often does not pay attention, we introduce an attentional state, α_t, with $\alpha_t = 1$ corresponding to engaged and $\alpha_t = 0$ to disengaged. Taking into account attention, and assuming a Bayesian update when the mouse is engaged and added noise otherwise, we have

$$\mu_{t+1} = \mu_t + \alpha_t \frac{\beta^2(\kappa_t - \mu_t)}{\beta^2 + \sigma_t^2} + (1 - \alpha_t)\zeta \tag{1a}$$

$$\sigma_{t+1}^2 = \sigma_t - \alpha_t \frac{\sigma_t^4}{\beta^2 + \sigma_t^2} + (1 - \alpha_t)\xi\,. \tag{1b}$$

where ζ and ξ are Gaussian noise that add 10% uncertainty to both mean and variance.

Every trial commences with the agent (mouse) at the center of the screen and a stimulus in right/left with a certain. The trial ends when the external position of the agent matches one of the decision boundaries, and the mouse receives a reward R_s. This reward is positive if it reaches the side of the stimulus, and a negative if it gets to the wrong decision boundary. Additionally a negative binary reward $R_{T'}$ is introduced to disincentivize very long trials. Analogous to the experiments, if trials do not receive a result within a threshold time T' they incur a negative reward of $R_{T'} = -\omega$, else $R_{T'} = 0$İn addition each gets a small negative reward $R_t = -rho$ for each time point that passes without reaching the target, analogous to accumulated fatigue. The total reward under a trial is given by $R = R_s + R_{T'} + TR_t$, where T is total trial time. In addition, the wheel moves according to $x_{t+1} = x_t + a_t\alpha_t$. The reward is then used to update the policy through policy gradient reinforcement learning. The goal is to learn a policy such that the expected reward is maximized.

The internal state updates through interactions with the environment and a policy is learned on this state, whereas taking an action modifies only the external state, and the reward is obtained when the external state, i.e., the physical position of the agent aligns with the target. Crucially, the action is *predicated on* the internal state only (through a policy), whereas it *affects* only the external state update. While this is unintuitive for machine learning

applications, it's a natural aspect of modeling animal behavior. An illustration of the process presented in Fig. 1

2.3 Policy gradient algorithm

The goal of the agent is to maximize the objective function: accumulated reward in every trial [15]. In other words, the agent determines the optimal action a_t given a state s_t. This is given by the *policy*. Consider a policy $\pi_\theta(a_t|s_t) = P(a_t|s_t, \theta)$ parametrized by the weights θ. We aim to maximize the expected reward under the policy

$$J(\pi_\theta) = \mathbb{E}_{\tau\sim\pi_\theta} R(\tau) \tag{2}$$

where τ denotes a trajectory chosen under the policy, and $R(\tau)$ is the total reward under that trajectory. One can fine tune a vector of parameters θ in order to learn a policy that optimizes the objective. The standard approach in machine learning is to optimize the policy through gradient ascent as follows:

$$\theta_{k+1} = \theta_k + \alpha \nabla_\theta J(\pi_\theta) \tag{3}$$

However, the gradient of an objective which contains the expectation ($\nabla_\theta \mathbb{E}_{\tau\sim\pi_\theta} R(\tau)$ using Eq. 2) is nontrivial. However, we can simplify it as follows:

$$\begin{aligned} \nabla_\theta \mathbb{E}_{\tau\sim\pi_\theta} R(\tau) &= \nabla_\theta \int \pi_\theta(\tau) R(\tau) d\tau \\ &= \int \nabla_\theta \pi_\theta(\tau) R(\tau) d\tau \\ &= \int \pi_\theta(\tau) \nabla_\theta \log \pi_\theta(\tau) R(\tau) d\tau \end{aligned} \tag{4}$$

Putting this together,

$$\nabla \mathbb{E}_{\tau\sim\pi}[R(\tau)] = \mathbb{E}_{\tau\sim\pi}[R(\tau) \nabla \log \pi_\theta(\tau)] \tag{5}$$

Now, using the definition:

$$\pi_\theta(\tau) = \Pi_{t=1}^{T} \pi_\theta(a_t|s_t) P(s_{t+1}, r_{t+1}|s_t, a, t) \tag{6}$$

Taking its log:

$$\log \pi_\theta(\tau) = \sum_{t=1}^{T} \log \pi_\theta(a_t|s_t) + \sum_{t=1}^{T} \log P(s_{t+1}, r_{t+1}|s_t, a, t) \tag{7}$$

The second term is independent of the parameters θ, hence disappears upon taking gradient w.r.t θ.

$$\nabla \mathbb{E}_{\tau \sim \pi}[R(\tau)] = \mathbb{E}_{\tau \sim \pi}\left[R(\tau) \sum_{t=1}^{T} \nabla \log \pi_\theta(a_t|s_t)\right] \tag{8}$$

Using the above, one can calculate gradients without knowing a model for the environment, i.e., without $p(s_{t+1}|s_t, a_t)$ (model-free approach) which is often difficult to estimate. The expectation in Eq. 8 is approximated through sampling several trajectories. The log likelihoods are simply the outputs of the policy network (softmax at the output layer yields $p(a_t|s_t)$).

Although the gradient of the parametrized policy does not depend on the reward, $R(\tau)$ adds a lot of variance over large samples (cite). However, replacing trajectory reward $R(\tau)$ with discounted reward $G(\tau)$ (cite Reinforce) in Eq. 5 reduces the variance in sampled trajectories and allows for online weight updates. For further discussion on policy gradient approaches see (cite). Our approach is outlined below:

1. Initialize the policy parameters (neural network weights) θ of at random.
2. Generate a trajectory τ on policy: The neural network takes a state as input and outputs a distribution over actions $s_1, a_1, s_2, a_2 \ldots s_T, a_T$.
3. Learning: aggregate reward, state and action vector for the entire trial $t = 1 \ldots T$, update policy parameters as follows:

   $$\theta \rightarrow \theta + \alpha G_t(\tau) \nabla_\theta \log \pi_\theta(A_t|S_t) \tag{9}$$

   where α is the learning rate, γ is the discount factor that discounts future rewards G_t is the accumulated reward such that rewards at t' forward steps from the current time are discounted by $\gamma^{t'}$
4. Repeat steps 2 and 3 until convergence.

3 Results

3.1 Model trajectories and internal state

We use a policy gradient model to approximate mouse behavior. The policy is parametrized by a 2 layer neural network with 64 nodes in the hidden layer

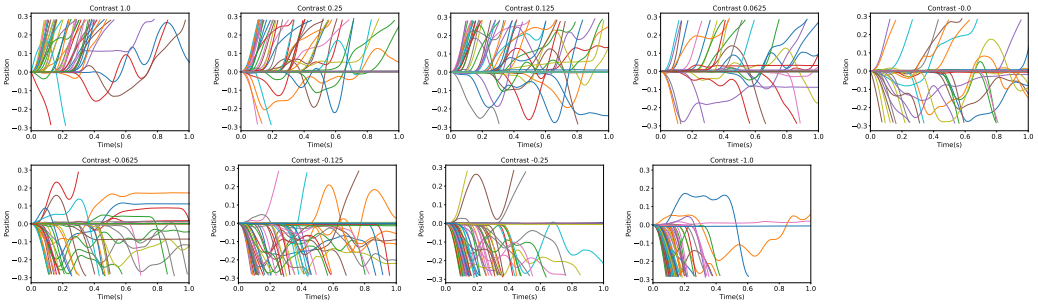

Figure 2: Mouse trajectories over several trials of a single mouse (KS003) sorted by contrast of stimulus, plotted upto 1 sec.

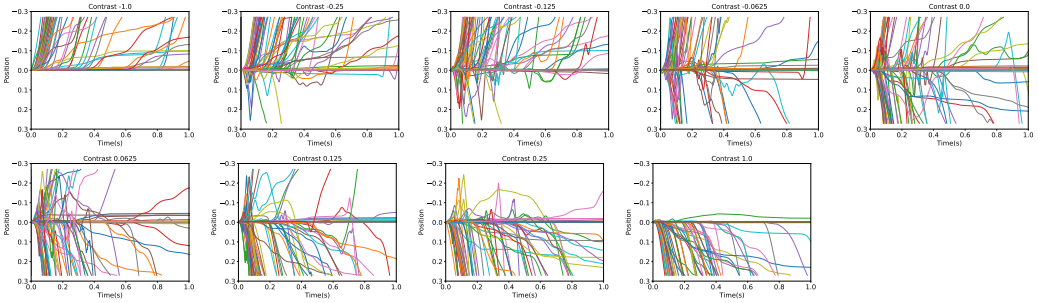

Figure 3: RL model trajectories over several epochs sorted by contrast of stimulus. Trajectories plotted after first 1000 epochs of learning time. learning rate $\alpha = 0.01$, discount factor $\gamma = 0.99$, $R = 10$, policy network has one hidden layer with 64 neurons.

and a tanh nonlinearity. The output layer is passed to a softmax layer to generate action probabilities. The learning rate is 0.001, and the discount parameter γ is set to 0.99. Fig. 2 plots trajectories over several trials of a single trained mouse. Notable features are (1) the performance is best for high absolute values of contrast. (2) A high proportion of trajectories wait a small amount of time to accumulate some evidence, and then rotate the wheel decisively towards one direction. The wait time is inversely proportional to absolute value of contrast. (c) Some trajectories make their first wheel movement very late (>1sec). (d) a small number of trajectories make changes in direction of velocity after accumulating additional evidence.

We attempt to capture these features in our model of mouse decision making. Fig. 3 shows equivalent RL-trajectories across various values of contrast of the stimulus. RL-trajectories are generated after 100 epochs of training, each trajectory is one epoch. The time scales are mapped onto real time corresponding to mouse movement by setting the fastest time to target to be equivalent. Reiterating, modeling attention is an essential component in modeling wheel movements, and attention times are statistically matched to the mouse. The RL generated trajectories do indeed capture the features of the mouse trajectories with some differences that we would like to address: (a) Trajectory curves aren't smooth: unlike the wheel movement task where velocity is naturally continuous, RL wheel movement is a discrete state variable, hence changes in direction are sharp. (b) The time scales of time to target are spread differently: since the notion of real time in RL is somewhat artificial, the mapping of timescales is somewhat artificial. Hence, the trajectories may seem like stretched out versions of the mouse trajectories. Now that we have addressed the main differences, it is natural to see that the fundamental features of mouse decision making are in fact captured by the RL model!

3.2 Statistical features of mice behavior

Mice behavior tends to have certain distinct statistical features that tend to be common across mice from different labs and trials (cite IBL paper):

- Trial accuracy (likelihood of mice succeeding at the task) is highest for trials with high contrast, and least at zero contrast
- Similarly, trial time is on average longest for zero accuracy and is inversely proportional to absolute value of contrast

- The mouse learn block structure, and it is known, in experiments, that decoding block structure from neural activity alone is difficult.
- Mice on very long trials can get distracted and lose track of the goal.

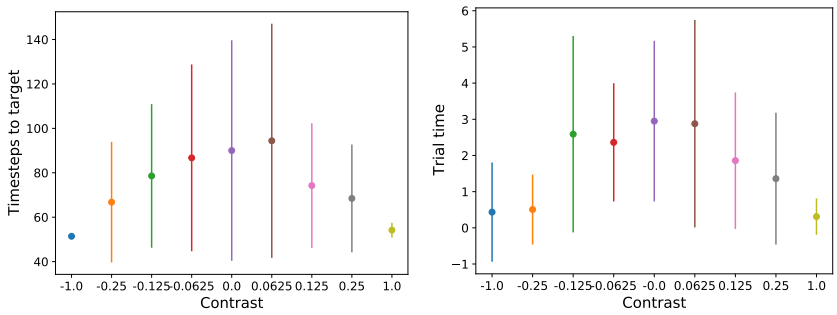

Figure 4: Trial time for (left) mouse, and (right) RL model as a function of contrast for the blocks.

Indeed one may naturally expect an animal to have the above features. However, a fully trained naive ML algorithm would not see the third feature. In the absence of realistic modeling of internal states, RL would significantly outperform biological agents like the mice, especially in a simple task like this. It then becomes essential to make unconventional and biologically motivated choices of internal states that may explain how a machine learning algorithm can emulate animal-like behavior. Here we use epochs in the context of the RL model to be analogous to trial in the context of mouse experiments, and timestep in RL to be analogous to real time in the experiment.

Fig. 4 shows the average trial time as a function of contrast, There is a correlation between trial time and absolute value of contrast, similar to animal behavior, despite the variance being relatively high.

3.3 Learning priors

Animals learn priors when exposed to biased data. The IBL dataset consists of block data - data recorded from mice exposed to stimulus in blocks with biased stimulus (80% right, 20%left or vice-versa). Experimental evidence [22] suggests that mice in blocks learn the prior, i.e., show preference for right when they are in the right dominant block. Fig. 5 shows that our model does indeed capture biases: the psychometric curves are shifted toward the block

prior. The psychometric curves shift because the network continues to learn through the blocks. Not surprisingly, then, higher learning rate (larger η) leads to larger shifts, although there is a fair amount of noise.

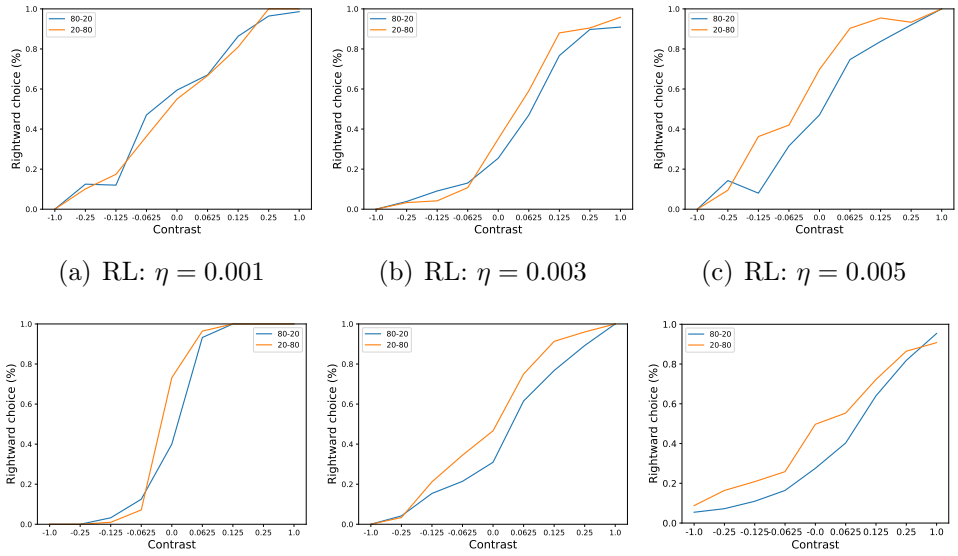

(a) RL: $\eta = 0.001$ (b) RL: $\eta = 0.003$ (c) RL: $\eta = 0.005$

(d) RL: $\eta = 0.01$ (e) RL: $\eta = 0.006$ (f) Agg - 40 mice

Figure 5: (a-e) RL psychometric curves (percent rightward choices) for different learning rates η. Comparison of (e) RL and (f) aggregate over 40 mouse. Blocks switch every 50 epochs in the RL model. $\beta = 0.5$

3.4 Neural evidence for block structure

Figs. 6 show the (a) activity of hidden neurons, and (b) their fractional difference for both blocks. Their are similar within a noise of 2%, making decoding of which block the signal came from extremely difficult. Thus, while it is possIBLe to learn something about block structure from activity, it is hard. Fig. 6c shows that the network learns the block in about 10 trials, consistent with what is seen in mice.

4 Discussion

In this work we study mouse behavior, specifically decision making in response to a stimulus, through a reinforcement learning model. Specifically, we

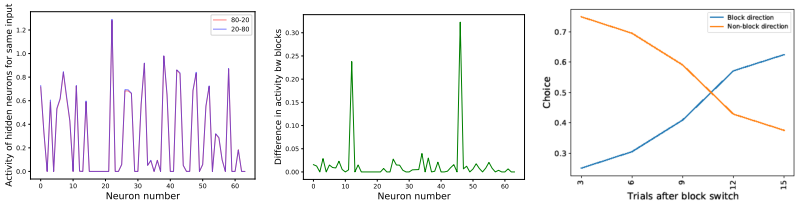

(a) (b) (c)

Figure 6: (a) Activity of 64 hidden neurons when presented with the same zero contrast input for for both blocks - the two curves overlap almost exactly. (b) Difference between neuron activity between blocks as a fraction of average activity across blocks for each neuron. (c) Choice of direction (left vs right) representing timescale of learning of new prior as a function of time after block switch on zero contrast trials.

study the IBL task where the mouse has to move a wheel in response to a grating stimulus of varying contrast that appears in blocks. However, animal behavior is often a result of a combination of factors that conventional goal-driven RL fails to capture. We attempt to capture some of these nuances through introducing an internal state, that consists of an evidence integrated approximation of the contrast, and a confidence measure of how certain the agent is in this approximation. The internal state follows Bayesian updates. The learned policy is a mapping from the internal state to the action, whereas the action modifies only the external state (wheel position). Additionally, our agent can switch between engaged and disengaged states to mimic animal behavior. Additionally, the model learns several of the statistical properties of mouse behavior including accuracy and trial time as a function of contrast. The stimulus is displayed in blocks, i.e., with probability 80% or 20% of being on the right side of the screen. This block structure introduces a bias, and in cases with zero or weak evidence, the mouse tends to move the wheel in the block direction, shifting the psychometric curves, indicating that the mouse encodes this prior. Our RL model reproduces the main experimental finding – that the psychometric curve with respect to contrast shifts after a block switch in about 10 trials. Also, as seen in the experiments, in our model the difference in neuronal activity between the right and left blocks is small making it very difficult to decode block structure from activity on single trials if noise is about 2%. Hence, here we hypothesis that priors can also be stored in the weights on the network, contrary to conventional wisdom, but in agreement with experimental observation. The hypothesis that priors are

stored in weights is difficult to test, but the technology to do so should be available in the not so distant future.

While this provides new insight into internal models that may explain animal decision making, there are a few caveats. As with all reinforcement learning methods, there are a large number of parameters to be selected and models are somewhat sensitive to parameter selection. Some of our parameters were chosen to match the experiments, and other were fine-tuned to obtain optimal results. Additionally, we want to emphasize that policy gradient is not the only RL algorithm one may use for this task, and other policy-based RL algorithms such as actor-critic, although requiring additional training, may be equally effective if endowed with the right internal model. Our choice of the internal model was motivated by biological plausibility and experimental data. The last caveat is that introducing noise is critical to modeling biological systems, and decoding priors from noisy activity is incredibly difficult. Hence, in this work, through our RL model, we present the novel hypothesis that priors can be stored in the weights.

5 Data Availability

Data is available from IBL [7]. All code is available on github at $github.com/chimeraki/RL_ibl$.

References

[1] Sabrina M Tom, Craig R Fox, Christopher Trepel, and Russell A Poldrack. The neural basis of loss aversion in decision-making under risk. *Science*, 315(5811):515–518, 2007.

[2] Joshua I Gold and Michael N Shadlen. The neural basis of decision making. *Annu. Rev. Neurosci.*, 30:535–574, 2007.

[3] Ward Edwards. The theory of decision making. *Psychological bulletin*, 51(4):380, 1954.

[4] Imran Noorani and RHS Carpenter. Full reaction time distributions reveal the complexity of neural decision-making. *European Journal of Neuroscience*, 33(11):1948–1951, 2011.

[5] Elisabetta Versace, Antone Martinho-Truswell, Alex Kacelnik, and Giorgio Vallortigara. Priors in animal and artificial intelligence: where does learning begin? *Trends in cognitive sciences*, 22(11):963–965, 2018.

[6] Liam Paninski and John P Cunningham. Neural data science: accelerating the experiment-analysis-theory cycle in large-scale neuroscience. *Current opinion in neurobiology*, 50:232–241, 2018.

[7] The International Brain Laboratory, Valeria Aguillon-Rodriguez, Dora Angelaki, Hannah Bayer, Niccolo Bonacchi, Matteo Carandini, Fanny Cazettes, Gaelle Chapuis, Anne K Churchland, Yang Dan, et al. Standardized and reproducible measurement of decision-making in mice. *eLife*, 10, 2021.

[8] Nicholas A Roy, Ji Hyun Bak, Athena Akrami, Carlos D Brody, Jonathan W Pillow, et al. Extracting the dynamics of behavior in decision-making experiments. *bioRxiv*, 2020.

[9] Rylan Schaeffer, Mikail Khona, Leenoy Meshulam, Ila Rani Fiete, et al. Reverse-engineering recurrent neural network solutions to a hierarchical inference task for mice. *bioRxiv*, 2020.

[10] Zoe Ashwood, Nicholas A Roy, Ji Hyun Bak, and Jonathan W Pillow. Inferring learning rules from animal decision-making. *Advances in Neural Information Processing Systems*, 33, 2020.

[11] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[12] Peter Dayan and Nathaniel D Daw. Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience*, 8(4):429–453, 2008.

[13] Aurelio Cortese, Hakwan Lau, and Mitsuo Kawato. Unconscious reinforcement learning of hidden brain states supported by confidence. *Nature communications*, 11(1):1–14, 2020.

[14] Yael Niv. Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53(3):139–154, 2009.

[15] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pages 387–395. PMLR, 2014.

[16] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.

[17] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.

[18] Christopher P Burgess, Armin Lak, Nicholas A Steinmetz, Peter Zatka-Haas, Charu Bai Reddy, Elina AK Jacobs, Jennifer F Linden, Joseph J Paton, Adam Ranson, Sylvia Schröder, et al. High-yield methods for accurate two-alternative visual psychophysics in head-fixed mice. *Cell reports*, 20(10):2513–2524, 2017.

[19] Nuttapong Chentanez, Andrew Barto, and Satinder Singh. Intrinsically motivated reinforcement learning. *Advances in neural information processing systems*, 17, 2004.

[20] Angela Radulescu, Yael Niv, and Ian Ballard. Holistic reinforcement learning: the role of structure and attention. *Trends in cognitive sciences*, 23(4):278–292, 2019.

[21] Valeria Aguillon-Rodriguez, Dora Angelaki, Hannah Bayer, Niccolo Bonacchi, Matteo Carandini, Fanny Cazettes, Gaelle Chapuis, Anne K Churchland, Yang Dan, Eric Dewitt, et al. Standardized and reproducible measurement of decision-making in mice. *Elife*, 10:e63711, 2021.

[22] Larry F Abbott, Dora E Angelaki, Matteo Carandini, Anne K Churchland, Yang Dan, Peter Dayan, Sophie Deneve, Ila Fiete, Surya Ganguli, Kenneth D Harris, et al. An international laboratory for systems and computational neuroscience. *Neuron*, 96(6):1213–1218, 2017.