EGT2
ENGINEERING TRIPOS PART IIA

Thursday 2 May 2024 9:30 to 11:10

Module 3G3 – CRIB

INTRODUCTION TO NEUROSCIENCE

Answer not more than ***three*** *questions.*

All questions carry the same number of marks.

The ***approximate*** *percentage of marks allocated to each part of a question is indicated in the right margin.*

Write your candidate number ***not*** *your name on the cover sheet.*

STATIONERY REQUIREMENTS
Single-sided script paper

SPECIAL REQUIREMENTS TO BE SUPPLIED FOR THIS EXAM
CUED approved calculator allowed
Engineering Data Book

10 minutes reading time is allowed for this paper at the start of the exam.

You may not start to read the questions printed on the subsequent pages of this question paper until instructed to do so.

You may not remove any stationery from the Examination Room.

1 A neuroscientist performs an *in vitro* study of a specific type of cortical neuron. They inject a current step, $I_{\text{ext}}(t)$, into a neuron of this type, and simultaneously record the cell's membrane potential, $V(t)$. The resulting voltage time course is shown in Fig. 1.

(a) Briefly explain why, during the initial period when $I_{\text{ext}} = 0$, the recorded membrane potential is not zero; what biophysical factors determine the non-zero value of this resting potential? [10%]

Answer: The membrane potential at rest is determined by the interplay of several types of ionic currents, the magnitude of each of which is determined by (i) the (potentially voltage-dependent) conductance or permeability of the membrane to that ionic species, and (ii) the ratio of ionic concentrations inside and outside the cell. For a typical cortical neuron, the relative balance of the main ionic currents (potassium, sodium and chloride) leads to an equilibrium potential of about −70 mV.

(b) Explain why the electrophysiological behaviour of this neuron, as probed by this simple step current injection, cannot be fully captured by the standard Hodgkin-Huxley (HH) model presented in lectures. [15%]

Answer: The Hodgkin-Huxley model does not have any long time scales that would explain the slow adaptation of this neuron (on the order of a couple of 100 ms). The typical membrane time constant is ~ 20 ms, and the slowest channel (potassium) has a maximum time constant of a few ms.

(c) The scientist proposes to rescue the HH model by including an additional current, I_{A}, mediated by slow, voltage-dependent, high-threshold potassium channels. This current takes the following form:

$$I_{\text{A}}(t) = -\bar{g}_{\text{A}}\, w(t)\, \left[V(t) - E_{\text{K}^+}\right] \tag{1}$$

$$\text{with} \quad \tau_{\text{A}} \frac{dw}{dt} = -w(t) + w_\infty(V(t)) \tag{2}$$

$$\text{and} \quad w_\infty(V) = \frac{1}{1 + \exp\left(-\frac{(V+35)}{10}\right)} \quad (V \text{ measured in mV}) \tag{3}$$

It is estimated that $\tau_{\text{A}} = 200$ ms.

(i) State the biophysical interpretations of $\bar{g}_{\text{A}}$, E_{K^+} and $w(t)$. [15%]

Answer: $\bar{g}_{\text{A}}$, measured in mS (or mS/cm^2 if all other membrane-related quantities are also measured per unit of membrane area) represents the peak conductance for this type of slow potassium channels, i.e. the conductance that would results from all channels being open ($w(t) = 1$). E_{K^+} is the reversal

potential for potassium ions (typically ~ -85mV) and $w(t) \in [0 : 1]$ represents the momentary fraction of open channels of that type.

(ii) Sketch the temporal evolution that you expect $w(t)$ to exhibit during a voltage clamp experiment whereby the membrane voltage, having been at rest for a long time, is suddenly clamped at 0 mV for 500 ms. [10%]

Answer: See the left panel in the figure below: at rest, $w(t) = w_\infty(-70\text{mV}) \approx 0$, and it then ramps up exponentially with timescale $\tau_A = 200$ ms towards $w_\infty(0 \text{ mV}) \approx 1$ during the voltage clamp.

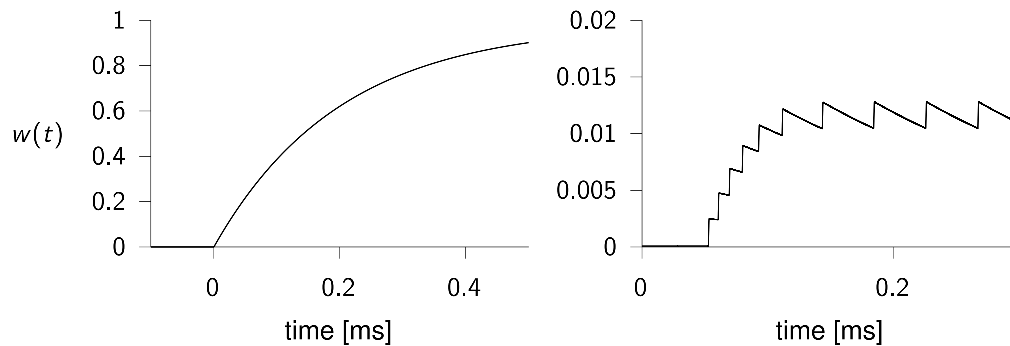

(iii) Sketch the temporal evolution that you expect $w(t)$ to exhibit during the 150 ms-long current injection experiment of Fig. 1, assuming that $V(t)$ in this model is indeed as shown on that figure. You may approximate each action potential by a square pulse of duration 0.5 ms. [15%]

Answer: See the right panel in the figure above. During each (very narrow) action potential, $w(t)$ ramps up exponentially towards $w_\infty(\text{high voltage}) \approx 1$, such that by the end of the action potential, it has effectively increased by $\approx 1 - e^{-0.5/200} = 0.0025$. It then decays slowly towards $w_\infty(\text{resting potential}) \approx 0$ on a timescale of 200 ms; therefore, when action potentials are fired in close succession, $w(t)$ has time to accumulate. At some point, action potentials become spaced enough that $w(t)$ has time to decay following each AP by just as much as it grew during the AP. This occurs as the neuron settles into a regime of regular firing.

(iv) Based on your answer to part (c)(iii), explain how the inclusion of $I_A(t)$ makes the HH model a better model for this neuron. [15%]

Answer: Each time the cell fires an AP, the slow potassium conductance increases, thereby slightly hyperpolarising the cell for several hundred ms. This accumulates over time (until a steady-state is reached) and overall slows down firing if the externally injected current does not compensate for that growing source of hyperpolarising current.

(v) Suppose that, during the injection of a steady current of some value $I_{\text{ext}} = a$, the neuron settles into a regime of regular firing, and the conductance due to the

 (TURN OVER

slow potassium channels never exceeds 2% of its theoretical maximum. Derive an upper bound for the steady-state firing rate of the neuron. You may use the same square pulse approximation to the AP as in part (c)(iii), and indeed your may find your answer to that question a helpful starting point. [20%]

Answer: The question states that $w(t) \leq 0.02$; to derive an *upper bound* for the neuron's firing rate, we can assume that, when the neuron starts firing at steady intervals Δ, $w(t)$ reaches exactly 0.02 by the end of each AP (anything lower than 0.02 would imply a lower firing rate). During the interval Δ (measured in ms), $w(t)$ decreases to a value of $0.02e^{-\Delta/200}$. At steady state, this value must also be equal to the value of $w(t)$ just before each AP, i.e. $0.02 - (1 - e^{-0.5/200})$ (c.f. answer to question (c)(iii)). Thus, Δ satisfies $e^{-\Delta/200} = 1 - \frac{1-e^{-0.5/200}}{0.02} \approx 1 - \frac{0.0025}{0.02} = 0.875$. This yields $\Delta = 200 \log(1/0.875) \approx 26.7$ ms, implying a firing rate of less than 38 Hz.

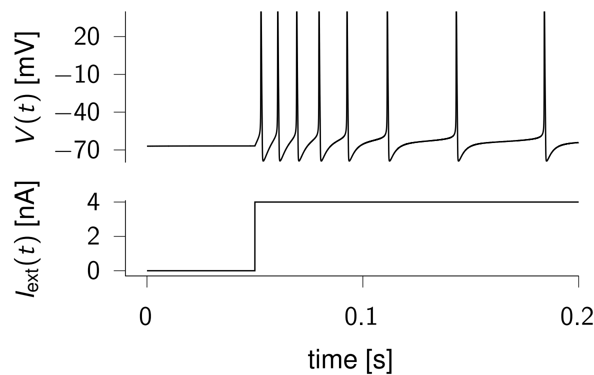

Fig. 1

2 (a) Briefly describe the intracellular and extracellular neuronal recording techniques. State two advantages and two limitations of extracellular recordings, relative to intracellular recordings. [15%]

Answer: For intracellular recordings, an electrode penetrates the cell's membrane, and its potential is compared to a reference electrode outside. For extracellular recordings, the electrode is kept outside the cell – yet small action potential-related signals can still be picked up. Advantages over intracellular recordings: any two of the following: much easier technique, doesn't damage the cell and therefore enables long recordings, can be scaled up to large arrays of electrodes. Limitations: cannot resolve sub-threshold membrane potential fluctuations, signal mixes up APs from multiple cells requiring spike sorting.

(b) Urn 'A' contains 6 yellow balls and 3 green balls. Urn 'B' contains 3 yellow and 3 green balls. An urn is chosen at random, and a ball drawn at random from that urn. Which of the following two questions constitutes the best illustration of the Bayesian perception framework, and why? Note: you are NOT required to solve these problems!

(i) "What is the probability that the ball is red?"

(ii) "The ball is red. What is the probability that it came from Urn 'A'?"

[15%]

Answer: The second problem is analogous to the way perception is formalised as an inference problem (or "inverse probability problem") in the Bayesian framework. The brain receives sensory evidence (here represented by the colour of the ball that was drawn), which is an incomplete reflection of the state of the world (here, the urn that was used) over which there therefore remains some uncertainty. Perception is then construed as the process through which the brain forms a "posterior belief" about the state of the world ("what is the probability that Urn A was used?"). This is done by combining a likelihood function ("with what probability would each of these two urns yield a red ball?") with a prior distribution ("what is the probability of drawing from Urn A in the first place?").

(c) This question is about a variant of the motion discrimination task discussed in lectures in the context of decision making. A monkey is trained to make a perceptual judgement about the net direction of motion (left or right) in a dynamic random-dot display exhibiting a certain degree of motion coherence c, presented for some limited duration T. After a short post-stimulus delay, the monkey has to indicate its direction choice by making an eye movement to one of the two direction-choice targets, and is subsequently rewarded for a correct decision with R mL of fruit juice. On a random half of the trials, the monkey is also given the option to instead choose a smaller ($0.8R$) but certain reward by making a saccade to a third target. This "sure target" is shown during

the delay period, at least 500 ms after the random-dot motion is extinguished. During motion viewing, the monkey does not know whether the sure-bet option will arise.

(i) Assume that, in each trial, the monkey does not only estimate the most likely motion direction given the sensory evidence but also the probability p of being correct about it ("confidence"), and that its decision maximises the expected reward given p. How large must the monkey's estimate of p be, for it to forgo the sure-bet option? [15%]

Answer: If the monkey believes in the most probable direction with probability p, it can expect a reward of pR if is forgoes the sure-bet option, versus $0.8R$ if it doesn't. Thus, a rational monkey would forgo the sure-bet option whenever $p > 0.8$.

(ii) For $c = 5\%$ and $c = 50\%$, *qualitatively* sketch the frequency (fraction of trials) with which you would expect to see the monkey choose the sure-bet option, as a function of T. Provide brief justifications. [15%]

Answer: Both curves should decrease as a function of T. The curve for $c = 5\%$ should lie above the curve for $c = 50\%$. That is because task difficulty decreases with both T and c, such that for larger T or larger c the monkey will be more certain about the motion direction and therefore have a lower probability of opting for the sure bet.

(iii) Consider 100 trials with constant c and T. State two factors that could cause variability in the monkey's internal confidence estimate across those trials. [15%]

Answer: Even for fixed c and T, the time course of the random dot display can be more or less informative about the motion direction, depending on the specific patterns of random incoherent motion sampled in the particular trial. This would lead to more or less uncertainty. In addition, even if that noise pattern was always the same, several sources of noise in the neural processes that underlie the monkey's estimate of confidence could lead to variability.

(iv) For any fixed c and T, the monkey's left-vs-right decisions are found to be better on average when it is offered the sure-bet option but opts out, compared to when it is not offered the choice. Does this speaks in favour of, or against, the assumptions of part (c)(i)? Why? [15%]

Answer: The fact that the monkey is more likely to choose the sure bet when the task is hard (c.f. answer to (c)(ii)) doesn't in itself speak for or against the assumption of (c)(i). After all, the monkey could simply estimate task difficulty and opt for the sure bet whenever it deems the trial too hard, but irrespective of the specific uncertainty it holds in that particular trial. However, the observation made here that, for equal average task difficulty (i.e. for trials with fixed c and T), the monkey does better

when it chooses to forgo the sure-bet option relative to when it is not given the choice, suggests that it indeed decides based on its internal sense of uncertainty (which will fluctuate across trials even for fixed c and T, as discussed in (c)(iii)).

(v) Based on what you know about the behaviour of certain LIP neurons in the standard version of this task, explain why, in the variant discussed here, it is reasonable here to expect the same LIP neurons to be somewhat predictive of whether or not the monkey will choose the sure-bet option (in trials where it is offered the choice). [10%]

Answer: As discussed in lectures, the activity of many LIP neurons appear to reflect the log odds ratio of the two directions conditioned on the history of sensory observations, as formalised algorithmically by the sequential likelihood ratio test. The value of this log likelihood ratio at the end of the stimulus presentation should not only dictate the most likely motion direction (hence the response of the monkey, if it has to decide) but also its subjective confidence which in turn should determine whether or not to accept the sure-bet option. Thus, the activity of LIP neurons should predict sure vs. high-stakes bet choices in trials where the monkey is given that choice.

3 This question is about synaptic transmission in the frog's nerve-muscle synapse (the so-called neuromuscular junction, NMJ). The NMJ works very much like an ordinary synapse between two neurons, and it has the following specifics:

- The NMJ uses acetylcholine (ACh) as the neurotransmitter, and mainly ionotropic receptors (so-called nicotinic ACh receptors) that are permeable to Na^+, K^+, and Ca^{2+} ions.
- Postsynaptic potentials in the NMJ are excitatory and they are called "end-plate potentials" (EPP).

(a) Describe the main events in the NMJ following the generation of an action potential in the presynaptic (motor) neuron that lead to the generation of an EPP in the muscle. [5%]

Answer:

(i) Presynaptic action potential propagates along the axon of the presynaptic cell.

(ii) Presynaptic action potential arrives at the presynaptic terminal (the synaptic bouton), depolarising the membrane there.

(iii) Voltage-dependent Ca^{2+} channels open in the presynaptic membrane.

(iv) Intracellular Ca^{2+} concentration increases at the presynaptic terminal.

(v) ACh-containing presynaptic vesicles fuse with the presynaptic membrane.

(vi) The ACh content of presynaptic vesicles is released into the synaptic cleft.

(vii) ACh molecules diffuse across the synaptic cleft.

(viii) When reaching the postsynaptic membrane, ACh molecules bind to postsynaptic (nicotinic) ACh receptors.

(ix) Postsynaptic ACh receptors open ion channels in the postsynaptic membrane.

(x) Na^+ and Ca^{2+} ions enter and K^+ ions leave the postsynaptic muscle spindle through these ion channels, so that their intracellular concentration increases and decreases, respectively.

(xi) On balance, there are more positive charges entering than ions leaving the muscle spindle, which therefore becomes depolarised.

(b) ACh molecules in the synaptic cleft are normally broken down by an enzyme called acetylcholinesterase. Prostigmine is a drug that blocks acetylcholinesterase. Describe how the shape of EPPs changes qualitatively when prostigmine is applied in the NMJ. [10%]

Answer: It becomes larger and longer because more ACh molecules bind to ACh receptors, and they bind for a longer time because they only disappear after a longer time from the synaptic cleft.

(c) Some snake venoms contain α-neurotoxins, which are nicotinic ACh receptor antagonists. Describe how EPPs change when an α-neurotoxin is applied to the NMJ. [5%]

Answer: They are diminished or completely abolished, as ACh cannot bind to the ACh receptor.

(d) In the NMJ, so-called spontaneous miniature EPPs (mEPPs) can be recorded without stimulation, in the absence of any action potential being generated in the presynaptic neuron. Fig. 3A shows the distribution of the amplitude of these mEPPs. The histogram shows the empirical distribution, the solid line shows a smooth distribution fit to the data. For answering the questions below, you may find the following additional information useful:

- The time course of mEPPs is indistinguishable from that of ordinary (action potential-induced) EPPs, but the amplitude of EPPs is much larger (around 70 mV).
- Prostigmine and α-neurotoxins have the same effect on mEPPs as on EPPs.
- The frequency, but not the amplitude, of mEPPs can be increased by a small depolarisation of the presynaptic terminal (which still does not lead to the generation of action potentials).
- The voltage response to the elementary current through a single ACh receptor-channel is approximately 0.3 μV.

Given these facts, answer the following questions

(i) Describe how the mechanism responsible for generating mEPPs is different from what you described in response to question part (a). [15%]

Answer:

A. Presynaptic action potential *does not* propagate along the axon of the presynaptic cell.

B. Presynaptic action potential *does not* arrive at the presynaptic terminal (the synaptic bouton), depolarising the membrane there.

C. Voltage-dependent Ca^{2+} channels *do not* open in the presynaptic membrane.

D. Intracellular Ca^{2+} concentration *does not* increase considerably at the presynaptic terminal.

E. ACh-containing presynaptic vesicles fuse *spontaneously* with the presynaptic membrane.

The rest is the same.

(ii) Explain with reasons why mEPPs are much smaller than EPPs and what determines the spread in their amplitude as shown in Fig. 3A. [10%]

Answer: For EPPs, typically several vesicles dock and fuse at around the same time, while for EPPs,

it is only one vesicle at a time. The spread is mainly due to variability in the transmitter content of vesicles, and in the number of ACh receptors binding ACh postsynaptically.

(e) Fig. 3B shows the amplitude distribution of ordinary, action potential-induced EPPs in the NMJ when intracellular Ca^{2+} concentration in the motor neuron is lowered. The histogram shows the empirical distribution, the solid line shows a fit to the part of the data above 0. For answering the questions below, note the following observations about this distribution (you may find it useful to consider these in comparison to Fig. 3A):

- The distribution has multiple peaks.
- The first peak is at 0 mV and it is very narrow (dark grey).
- Successive peaks are at integral multiples of cca. 0.4 mV and they widen progressively.

Given these observations, and your knowledge of the mechanism of synaptic transmission, answer the following questions:

(i) What happens on those occasions that result in a 0 mV EPP? [5%]

Answer: Synaptic failure: accidentally, no vesicles are released.

(ii) Why does synaptic transmission appear to be "quantal": why are EPP amplitudes (approximately) integral multiples of the same value? [20%]

Answer: This is caused by neurotransmitter molecules being released in vesicles. Each vesicle contributes a quantum of EPP, so the total amplitude of the EPP is mainly determined by the number of vesicles released in response to an action potential.

(iii) What causes the spread in EPP amplitudes around the quantal values (the peaks of the distribution), and why do the peaks become wider? [10%]

Answer: The same factors that determine spread in the size of mEPPs (see above): variability in the transmitter content of vesicles, and in the number of ACh receptors binding ACh postsynaptically. As the total EPP is the sum of postsynaptic responses to individual vesicles released presynaptically, its variance scales with the number of vesicles released – hence the widening.

(iv) How does this distribution change if the intracellular Ca^{2+} concentration in the motor neuron is increased? [20%]

Answer: Presynaptic Ca^{2+} concentration controls the number of vesicles released but not their neurotransmitter content, and so when it is increased, the probability of failures (the amplitude of the first peak of the histogram at 0 mV) decreases, and the probability of larger number of vesicles released

 (cont.

(the amplitude of subsequent peaks) increases.

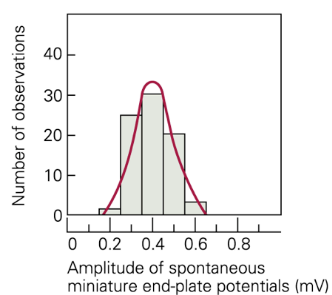

Fig. 3A

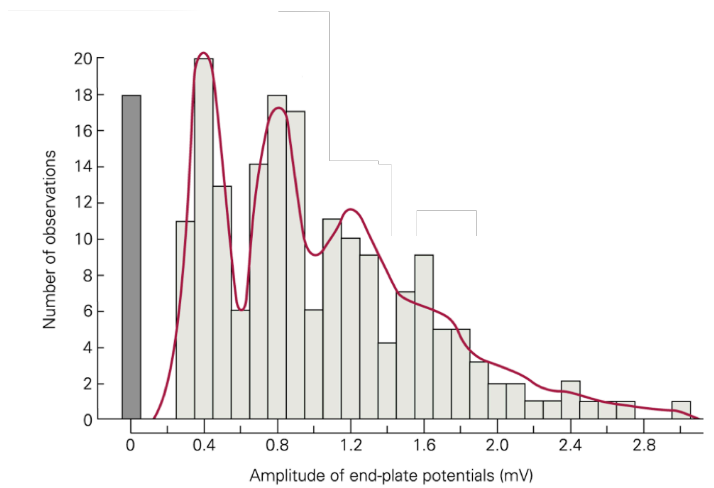

Fig. 3B

4 In the lectures, we discussed an experiment in which the contributions of the hippocampus to spatial learning have been studied at different stages of learning. Answer the following questions regarding this experiment.

(a) What species and experimental apparatus was used? [10%]

Answer: Rats and a plus maze. (Morris water maze is also an acceptable apparatus as long as answers to subsequent sub-questions correctly describe how to distinguish different behavioural strategies at different stages of learning, and the role of the hippocampus in them.)

(b) What was the behavioural task animals had to perform? What was the protocol for training and testing animals? [10%]

Answer: Animals were started from one arm of the maze and had to find a reward in one of the adjacent arms in the maze. They were trained by always using the same combination of starting and rewarded arm (for example, they were always started from the South arm, with the reward being in the West arm of the maze). They were tested by starting from the arm that is opposite from the starting arm used during training (in the example, the North arm).

(c) What kind of behavioural strategies might animals use to solve this task, and how can we tell which strategy they are using by observing their behaviour? [10%]

Answer: They can use a "place strategy" that is based on remembering the allocentric location of the rewarded arm, such that they approach that arm from whichever arm they are started from (e.g. go to the West arm), i.e. both during training and testing. Alternatively, they can use a "response" strategy, such that they always execute the same egocentric actions (take the left arm when you reach the centre of the maze), which would result in they choosing the opposite arm during testing than during training.

(d) What brain regions were blocked in this experiment? [5%]

Answer: The hippocampus and the striatum (caudate nucleus).

(e) What pharmacological manipulations were used for perturbing the operation of these brain regions, when were these used in the experiment, and what is the mechanism of their action? [10%]

Answer: Lidocaine injections were used to reversibly block activity in either of the given brain regions

(hippocampus or striatum) immediately before testing the animals after 8 or 16 days of training. Lidocaine is a Na^+ channel blocker, so it blocks the generation and propagation of action potentials, and as a result silences the given brain area.

(f) What was the control pharmacological manipulation used, when and where was it used, and what does it control for? [5%]

Answer: In a control group of animals, saline was injected at the same times as in the experimental group with lidocaine injections and to the same brain regions. Saline does not perturb neural activity but acts as a control for the unspecific effects of surgery and injection.

(g) How many experimental groups were used in the experiment, and approximately how many animals were used in each group? [10%]

Answer: There was a different experimental group for each combination of lidocaine vs. saline injection, hippocampal vs. striatal injection, and short (8-day) vs. long (16-day) training, which gives us $2 \times 2 \times 2 = 8$ experimental groups in total. The number of animals in each group was 12-14 (about ten).

(h) What was the pattern of results observed in the experiment for each experimental group? [10%]

Answer: The following table shows the dominant behavioural strategy for each experimental group.

injection	where	when	strategy
saline	striatum	8 days	place
saline	hippocampus	8 days	place
lidocaine	striatum	8 days	place
lidocaine	hippocampus	8 days	place/response
saline	striatum	16 days	response
saline	hippocampus	16 days	response
lidocaine	striatum	16 days	place
lidocaine	hippocampus	16 days	response

(i) What was the interpretations of these results in terms of how behaviour changes over the course of spatial learning, and how different brain regions contribute to navigation? [10%]

Answer: The place strategy takes relatively little training to acquire, while it takes longer training for the response strategy to be acquired. Once the response strategy has been acquired, animals (with intact brains) will rely on that over the place strategy. Thus, as a default, animals rely on the place strategy early in learning, and on the response strategy late in learning. In turn, the place strategy relies on the hippocampus,

 (TURN OVER

and the response strategy relies on the striatum. Animals can revert to the other strategy when their default strategy becomes unavailable (e.g. due to lesioning of the brain area on which the default strategy relies) as long as the other strategy has already been acquired. (Thus, animals can revert to the place strategy when their striatum is blocked after long training, but their behaviour is mostly random when their hippocampus is blocked after short training due to the lack of a fully acquired response strategy at that point.)

(j) A strategy that was not explicitly considered in the original experiment is that animals might use distinctive visual cues to guide them to the reward. How would you distinguish this strategy from those you described in your answer to question (c)? [20%]

Answer: Distinguishing from response strategy: same approach as used in the original experiment to distinguish place and response strategies. During training, start animals always from the South arm and reward them in the West arm, but test them (in extinction) by starting them from the North arm. If they choose the East arm, then they are using the response strategy, whereas if they choose the West arm, they are using the visual (or place) strategy.
Distinguishing from place strategy: switch lights off during testing. If animals choose East and West arm with roughly equal probability, they are using a visual strategy, if they keep choosing the West arm, they are using the place strategy.

END OF PAPER

Numerical answers

Q1(c)(v): 38 Hz.
Q2(c)(i): $p > 0.8$.