Version GH/3 – CRIB

EGT2
ENGINEERING TRIPOS PART IIA

Monday 9 May 2022 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
Supplementary page: one extra copy of Fig. 2.

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 This question is about multisensory integration, in the context of the monkey 2AFC heading discrimination task discussed during lectures. The experimental setup is depicted in Fig. 1. A monkey sits on a motion platform which can be translated in any direction at (small) angle α relative to straight heading. A sophisticated virtual reality setup involving a 3D dot cloud allows the simulation of the visual flow consistent with any translation of the platform, whether or not this translation actually occurs physically. Here, we suppose that in any trial where the true heading is α, the visual and vestibular inputs to the monkey's brain can be statistically summarised by two quantities, $\hat{\alpha}_{\text{vis.}}$ and $\hat{\alpha}_{\text{vest.}}$, drawn independently from two normal distributions with identical means α and standard deviations $\sigma_{\text{vis.}}$ and $\sigma_{\text{vest.}}$ respectively.

(a) Explain why it is reasonable to assume that $\hat{\alpha}_{\text{vis.}}$ and $\hat{\alpha}_{\text{vest.}}$ are indepedent given α. [10%]

Answer: Sensory evidence is provided by noisy neurons in the sensory periphery (here, the inner ears and the retinas); as these neurons do not share input, and are physically separated by a long distance, the noisy processes (e.g. thermal fluctuations) that corrupt their activity for any given translation of the platform are likely to be statistically independent.

(b) Suppose the platform is physically translated, but no visual feedback is provided. Explain why $\hat{\alpha}_{\text{vest.}}$ is also called the "maximum-likelihood estimate" of the heading direction in this case. [10%]

Answer: The likelihood density $p(\hat{\alpha}_{\text{vest.}}|\alpha)$ can be viewed both as a bell-shaped function of $\hat{\alpha}_{\text{vest.}}$ that has its maximum at α, or as a (the same) bell-shaped function of α that has its maximum at $\hat{\alpha}_{\text{vest.}}$. The latter view immediately justifies calling $\hat{\alpha}_{\text{vest.}}$ the maximum likelihood estimate of the (unobserved) α.

(c) Suppose the plaform is physically translated, and consistent visual feedback is provided. Derive expressions for the multisensory maximum likelihood estimate $\hat{\alpha}_{\text{mult.}}$ and its variance $\sigma^2_{\text{mult.}}$. [30%]

Answer: By multiplying the two Gaussian likelihoods, and completing the square to form combined (unnormalized – there is no need to bother computing the normalization constant) Gaussian density over α, one readily identifies a mean given by $\hat{\alpha}_{\text{mult.}} = \hat{\alpha}_{\text{vest.}} + \frac{1}{1+(\sigma_{\text{vis.}}/\sigma_{\text{vest.}})^2}(\hat{\alpha}_{\text{vis.}} - \hat{\alpha}_{\text{vest.}})$ and variance given by $\sigma^2_{\text{mult.}} = \frac{\sigma^2_{\text{vest.}}\sigma^2_{\text{vis.}}}{\sigma^2_{\text{vest.}} + \sigma^2_{\text{vis.}}}$.

(d) Suppose you wish to show that, when combining visual and vestibular sensory cues,

monkeys form statistically optimal internal estimates of the heading direction. Describe the experimental protocol you would use, and the associated analysis you would perform. Provide reasons for your choices. [20%]

Answer: Run the task with no optic flow feedback; in each trial, the platform is translated in direction α drawn randomly from some symmetric distribution around zero, which should be wide enough for the largest (absolute) values of α to lead to near-perfect performance, but narrow enough to be able to fit the resulting psychometric curve accurately. The slope of this curve at $\alpha = 0$ is directly related to $\sigma_{\text{vest.}}$. Repeat this procedure with sole optic flow feedback and no physical translation, to estimate $\sigma_{\text{vis.}}$. Finally, run the task in the combined condition as in (c), to estimate $\sigma_{\text{mult.}}$. Basic analysis of the monkey's behaviour in the combined task can then be performed to see if the Bayesian predictions of (c) are borne out.

(e) Derive a condition on the visual and vestibular likelihood precisions, $\sigma^{-2}_{\text{vis.}}$ and $\sigma^{-2}_{\text{vest.}}$, such that they together yield the largest percentage increase in multimodal precision $\sigma^{-2}_{\text{mult.}}$ over the best of the two single-modality precisions. [20%]

Answer: This problem is more quickly solved by defining $s = \frac{1}{2}(\sigma_{\text{v}}^{-2} + \sigma_{\text{a}}^{-2})$ and $d = \frac{1}{2}(\sigma_{\text{v}}^{-2} - \sigma_{\text{a}}^{-2})$. We seek to maximise the quantity $\mathcal{L}(s, d) = 2s/\max(s + d, s - d)$ w.r.t s and d. It is easy to see (e.g. through a graph) that the minimum of the denominator w.r.t. d is attained at $d = 0$ for any setting of s, and is equal to s. Thus, the maximum of $\mathcal{L}(s, d) = 2$ for any s. This shows that any setting of the precisions in which they are equal yields the maximum percentage improvement over the single modalities.

(f) Based on your answer to the previous question, explain why a clever experimenter would in fact corrupt the visual flow feedback (e.g. by introducing noise, as done in Gu et al., *Nature Neuroscience*, 2008). [10%]

Answer: The primate visual system typically provides higher-quality evidence for self-motion than the vestibular system does. Question (e) above suggests that to get the largest integration effect, the quality of evidence should be roughly balanced over the two modalities, such that $\sigma_{\text{vest.}} \approx \sigma_{\text{vis.}}$. While $\sigma_{\text{vest.}}$ cannot easily be decreased through experimental manipulation, $\sigma_{\text{vis.}}$ can easily be increased by corrupting the optic flow feedback.

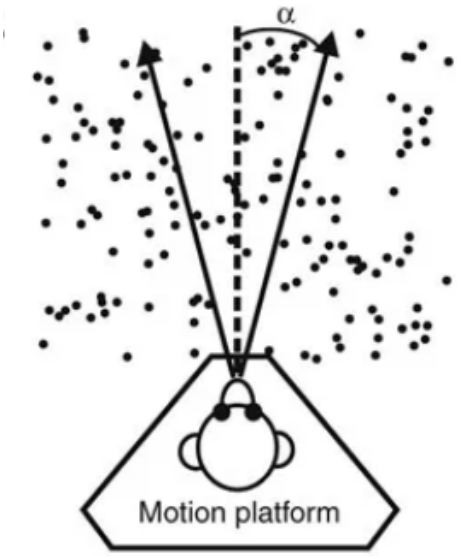

Fig. 1

2 (a) Write short notes on the following:

(i) The resting membrane potential: in particular, how it arises, and what determines its value; [15%]

Answer: Ion channels are present in the membrane and allow the movement of charges through the membrane, which otherwise acts as a capacitor. For a given ion species, there is a so-called equilibrium (Nernst) membrane potential at which ionic diffusion through the membrane (due to the concentration gradient of that ion) exactly balances the flux of ions due to the membrane potential itself. The resting membrane potential in a neuron is a weighted combination of the Nernst potentials of the various ion species for which there are ion channels. The larger the associated ionic conductance, the more that ion's Nernst potential contributes.

(ii) The action potential: in particular, what triggers it, and what terminates it; [15%]

Answer: When the membrane potential goes above a certain threshold (typically around -50 mV, sodium channels activate rapidly (more rapidly than potassium channels) and this drives the membrane potential even higher towards the Nernst potential of sodium. A millisecond or so later, the now elevated V_{m} causes sodium channels to deactivate, which allows the outward potassium current (which has also grown with V_{m}, and whose inactivation will follow more slowly) to dominate and to drive the membrane potential back to rest.

(iii) The sources of uncertainty that may limit performance in perceptual decision making: in particular, explain why making optimal decisions requires adequate treatment of such uncertainty. [15%]

Answer: Perceptual decision making is limited by two main sources of uncertainty: perceptual uncertainty due to noise in the sensory input, and uncertainty about the way internally formed decisions actually end up affecting the feedback (reward) we receive. The latter may encompass motor noise (causing errors in the behavioural report of the decision) as well as noise in the environment itself (e.g. the reward function itself may be stochastic). Making optimal decisions involves maximizing some objective function w.r.t. the decision we form internally, and using point estimates of the various perceptual and motor quantities of interest can often result in sub-optimal average return. Optimal average returns can only be obtained by accurately modelling the above-mentioned sources of uncertainty and integrating them out.

(b) A Hodgkin-Huxley model neuron is injected with a step of depolarizing current, sufficiently large to elicit firing (Fig. 2, solid black line, reproduced in each panel to facilitate comparisons in part (ii) below).

(i) Sketch the corresponding time course of the sodium and potassium gate

variables. [15%]

Answer: The three gates are unit-less variables that evolve between 0 and 1, as follows:

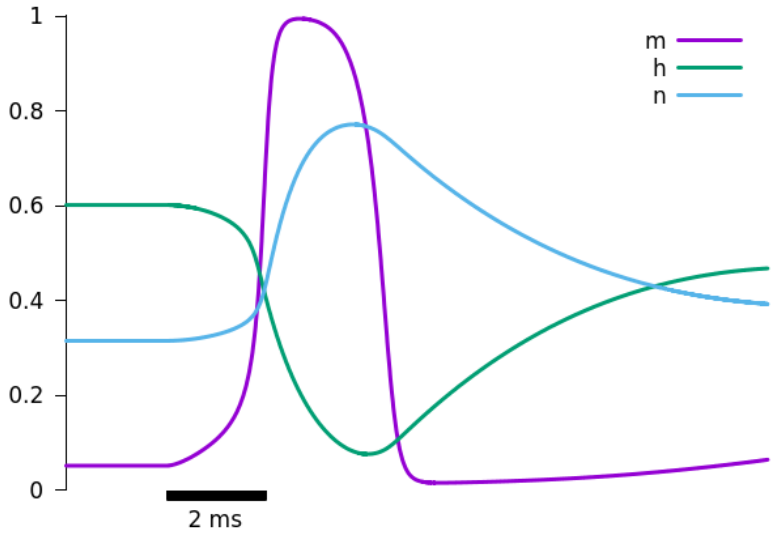

(ii) The very same current injection is performed in the following variants of the model neuron:

A: Lower reversal potential for sodium channels;

B: Larger peak potassium conductance;

C: Faster kinetics for the *m* and *h* gates;

D: Slower kinetics for the *n* gate.

The membrane potential time courses arising in these four scenarios are shown as dashed lines in the four panels of Fig. 2, in arbitrary order. On the *additional copy of Fig. 2* provided at the end of this paper, complete the legend in each panel by circling the correct scenario label (A, B, C, or D). Provide a justification for your choices. Do not forget to hand in your completed copy of Fig. 2 with your answer to this question. [40%]

Answer: **1–A**: in the HH model, the membrane potential turns back when it approaches the reversal potential of sodium channels (at that point, sodium stops flowing in). The lower peak voltage in panel 1 therefore suggests it corresponds to case A.

2–C: panel 2 exhibits faster AP upstroke; since the upstroke is primarily determined by the dynamics of the sodium channels, this suggests scenario C.

4–D: the broader AP, as well as the slower growth following after-AP hyperpolarisation, both suggest slower potassium dynamics; indeed, it is the potassium current that is responsible for hyperpolarization, and this current clearly decays slower here.

3–B: by elimination; but also, panel 3 exhibits a delayed AP and an otherwise unaltered AP height; this suggests a stronger hyperpolarizing (potassium) current overall; consistent with this, the AP

(cont.

terminates at the same time despite a delayed upstroke.

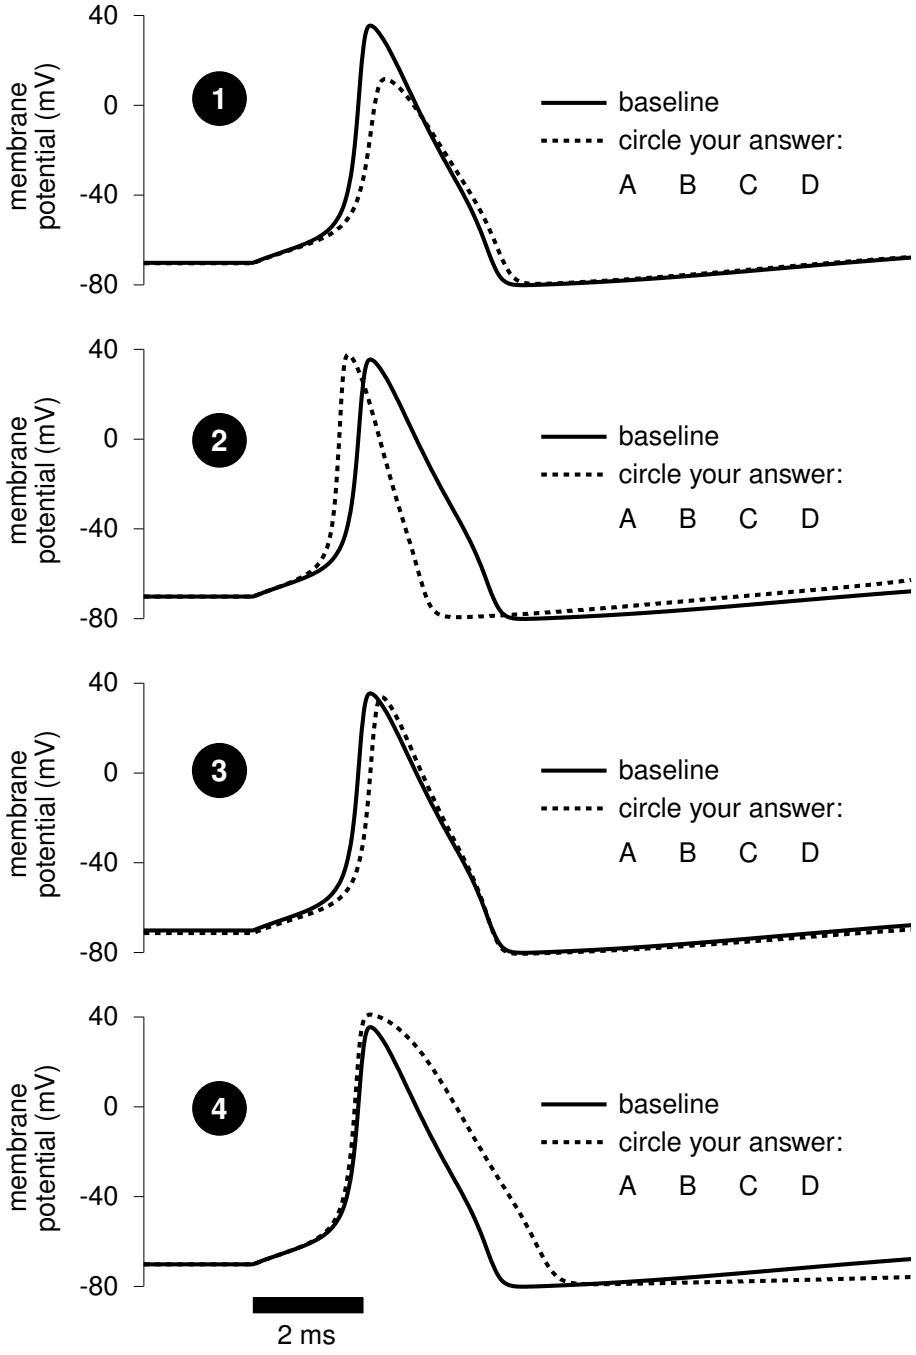

Fig. 2

3 (a) This part is about Dale's principle.

(i) What does the principle state? [10%]

Answer: It states that neurons are either excitatory or inhibitory, *i.e.* all efferent synapses of a given neuron onto other neurons have the same 'sign': they are either all excitatory or all inhibitory.

(ii) Explain how Dale's principle arises from specific regularities governing the production of neurotransmitters and their binding to receptors. [10%]

Answer: According to Dale's principle, the excitatory/inhibitory nature of a synapse is determined by (*i.e.*, depends only on) the excitatory/inhibitory type of the *pre-synaptic* neuron. However, it is ultimately the reversal potential of receptor channels in a synapse that determines the excitatory/inhibitory nature of the synapse, and receptors are expressed on the *post-synaptic* membrane. What allows Dale's principle to hold is that neurons in the brain (with minor exceptions) produce only one type of neurotransmitter *and* neurostransmitters selectively bind a specific group of receptor types, all of which have either excitatory or inhibitory effects (*i.e.*, have their reversal potentials either above or below the spiking threshold potential).

(b) Figure 3, from Bliss & Lømo 1973, shows the evolution of the population EPSP amplitude in the granule cells of the dentate gyrus area of Hippocampus, evoked by stimulation of the perforant pathway (PP) axons. The PP is stimulated at a frequency of 0.5 s^{-1}, except during the epoch marked by a black bar (bottom) during which the frequency is increased to 15 s^{-1}. Explain why high-frequency stimulation but not low-frequency stimulation leads to long-term potentiation of the EPSP amplitude. Identify the specific synaptic mechanism involved. [20%]

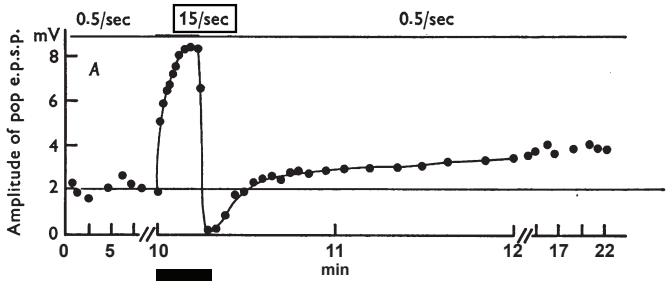

Fig. 3

Answer: The short answer is: because of the magnesium block of NMDA channels. The NMDA receptors in the post-synaptic cells (DG pyramidal neurons in this case) are activated by the glutamate released from presynaptic axonal terminals (PP axons), upon PP stimulation (which also evokes the EPSP's). However,

the Mg^{2+} ion in the NMDA channels blocks the entry of Ca^{2+}, unless the post-synaptic compartment is sufficiently depolarized. On the other hand, an individual EPSP creates a depolarization that decays to 0 on a 10-100 ms timescale. Thus when EPSP's arrive with separating intervals that are significantly longer than those timescales, the successive EPSP's do not add up. However, when EPSP's arrive in rapid succession, they do add up (as more recent EPSP's arrive before previous EPSP's have completely died) to create larger depolarizations which can reach the threshold for removal of the magnesium block of NMDA channels and allow Ca^{2+} ions to flow through glutamate-activated NMDA channels and enter the cell. The entry of this second messenger then starts the process leading to LTP. This is why only a high frequency stimulation of PP pathway can induce LTP.

(c) This part is about the Inhibitory Conditioning (IC) paradigm in reinforcement learning and the Rescorla-Wagner model.

(i) Describe the experimental protocol for the IC paradigm. [15%]

Answer: This paradigm involves an unconditioned stimulus (US) and two different conditioned stimuli (CS1 and CS2). The paradigm consists of a sequence of trials that are of two types, with trials of the two types randomly interleaved. In trials of type one, CS1 and US co-occur and CS2 does not occur, while in trials of type two, CS1 and CS2 co-occur and US does not occur.

(ii) Let u denote the binary 0/1 indicator for the unconditioned stimulus, and s_i denote the 0/1 indicator for the i^{th} conditioned stimulus (CS); there will be as many s_i as there are CS's in the IC paradigm. Write the Rescorla-Wagner update rules for all the weights, w_i, of that model. Provide simplified expressions that only involve u and w_i. [25%]

Answer: According to the Rescorla-Wagner (RW) rule, these weights are updated via $\Delta w_i = \epsilon(u - w_1 s_1 - w_2 s_2)s_i$. Since in this paradigm we always have $s_1 = 1$ and $s_2 = 1 - u$, we obtain $\Delta w_1 \propto (u - w_1 - w_2(1-u)) = u(1+w_2) - (w_1 + w_2)$ and $\Delta w_2 \propto (u - w_1 - w_2(1-u))(1-u) = -(w_1 + w_2)(1-u)$.

(iii) To which steady-state values will the weights of the Rescorla-Wagner model converge over time (assuming they *will* converge), in this paradigm? State the answer for the weights of all CS's used in the IC paradigm. [20%]

Answer: To obtain steady-state values, we have to set $\Delta w_i = 0$ (in general, on average). Using the simplified update rules derived in the previous part, we obtain the conditions $u(1+w_2) - (w_1+w_2) = 0$ and $(w_1 + w_2)(1-u) = 0$. We see that the weight values $w_1^* = -w_2^* = 1$ make both Δw_1 and Δw_2 vanish and are thus fixed points of the RW update rule. We also note that these are the only values that make the weight changes vanish independently of the trial type (i.e. independent of the value of

u). Thus if the weights do converge, they will converge to these values.

4 (a) This part is about classical conditioning of the gill-withdrawal reflex in *Aplysia*.

(i) Describe the experimental procedure for this classical conditioning. [10%]

Answer: Classical conditioning of the gill-withdrawal reflex is achieved by repeatedly delivering an innocuous stimulus (*e.g.*, gentle touching of the siphon) shortly succeeded by a noxious one (*e.g.*, electric shock to the tail).

(ii) Describe the two types of coincidence detection mechanisms involved in the induction of synaptic plasticity underlying this kind of learning. Describe each:

A. in terms of a coincidence of different stimuli and/or the behavioral response (*i.e.*, gill-withdrawal); [10%]

B. in terms of a coincidence of cellular or molecular events and variables, explaining how/why this translates to the corresponding coincidence among stimuli and/or motor response. [20%]

Answer: Two coincidences are detected: (1) coincidence of the unconditioned stimulus (US) and conditioned stimulus (CS), and (2) coincidence of the CS and the motor response (gill withdrawal). In the first case, the CS evokes the firing of action potential(s) by the pre-synaptic sensory neuron which leads to entry of Calcium ions into the axonal terminals of that neuron pre-synaptic to the motor neuron. On the other hand, the US, which follows with rapid succession, leads to release of serotonin from the interneuron synapsing onto the mentioned axonal terminals. The elevated Calcium resulting from the CS-elicited action potential modifies the cAMP-based molecular pathway that is triggered by serotonin receptors in the axonal terminal, such that more cAMP is produced when the serotonin pathway is activated by the released serotonin, leading to a stronger potentiation of the synapse. The second type of coincidence detection involves the removal of Magnesium block from the NMDA channels on the post-synaptic side of the axonal terminal (i.e. on the motor neuron): the CS leads to the release of glutamate from the sensory neuron's axonal terminal which binds to and activates the post-synaptic NMDA receptors, while the firing of the motor neuron (which is responsible for the gill withdrawal) provides the post-synaptic depolarization necessary for the removal of the magnesium ion from the NMDA receptor's ion channel. With the Magnesium ion removed, Calcium now rushes into the post-synaptic compartment through those channels, which initiates a second pathway (involving retrograde messengers) also leading to the potentiation of the synapse between sensory and motor neurons.

(b) This part is about the effect of the learning rate, ϵ, on the function and behavior of the Rescorla-Wagner (RW) model.

(i) Describe the qualitative difference in the time evolution of the RW model's

stimulus weight, w, during training in the Partial Reinforcement paradigm, for large ($\epsilon \sim 1$) *vs.* small ($\epsilon \ll 1$) learning rates. [15%]

Answer: For large ϵ, the weight fluctuates more strongly from trial to trial, as it reacts more strongly to stochastic co-occurrance (or lack thereof) of the CS and US in individual trials. For small ϵ, the the weight fluctuates much less, as it does not change appreciably in an individual trial, and fluctuates less due to effective averaging across many trials.

(ii) Explain conceptually what is learned by the model in each of the two cases in the previous part (for large learning rate, assume $\epsilon = 1$ for concreteness). [15%]

Answer: Since in the Partial Reinforcement paradigm the (single) CS is present in all trials, the RW learning rule reduces to: $\Delta w = \epsilon(u - ws)s = \epsilon(u - w)$, where u is 1 or 0 in trials with or without a US. In the $\epsilon = 1$ case, we see that if CS is followed/accompanied by the US (i.e. $u = 1$) in a trial, w jumps to (or remains at) 1, while if CS is not followed/accompanied by US ($u = 0$) it jumps to (or remains at) 0. Thus we see that, for $\epsilon = 1$, the weight "remembers" the co-occurance (or lack thereof) of CS and US in the most recent trial/experience. In the slow learning case ($\epsilon \ll 1$), by contrast, the model learns the average or statistical association between the CS and US in the past, and is not strongly affected by the most recent experience.

(iii) Explain why a large learning rate ($\epsilon \sim 1$) is typically problematic for learning in the natural environment. [15%]

Answer: Because in the natural environment, associations between stimuli (including between rewards/punishments on the one hand and stimuli that are predictive of those on the other) are typically stochastic and not deterministic. In a stochastic environment, it makes much more sense to learn *statistical* associations (between natural "CS"'s and "US"'s), based on many observations, rather than by reacting strongly to one or a few recent observations. This is what is achieved by the RW model with a small learning rate (see answer to the previous part).

(iv) Consider a modification to the RW model which has an adaptable learning rate. Suppose the environment switches from a low volatility state to a high volatility one. Should ϵ adapt to a higher or lower value to improve performance? Explain why. [15%]

Answer: In a high volatility environment, the statistics of associations between and co-occurances of different stimuli change more frequently or rapidly. Therefore, to perform well in the current state or epoch of the environment (a duration of time over which the statistics of stimulus associations is roughly constant), knowledge of statistical associations in previous epochs is not useful. It thus makes more sense for the RW to average statistics over the current epoch, but not over very long

timescales larger than the timescale over which environmental statistics change. It is thus favorable to increase the learning rate to decrease those averaging timescales, once the environment becomes more volatile and the durations of different epochs decrease.

END OF PAPER

THIS PAGE IS BLANK

Comments on Questions

[Q1]

34 candidates attempted this question, i.e. a little below chance level. The average mark was $12.4(\pm 4.0)/20$; min 6/20, max 20/20.

Part (a) was straightforward and almost universally well answered. Part (b) was meant to be relatively easy too but surprisingly only correctly answered by a small fraction of candidates! Most candidates gave very vague explanations of the sort of "vestibular feedback is the only source of information here, so $\hat{\alpha}_{\text{vest.}}$ is the ML estimate", where a proper (yet very simple) explanation would instead consider the form of the likelihood function and noting that it has a maximum at $\hat{\alpha}_{\text{vest.}}$. Part (c) was well answered by most, though some did get a little bogged down in lengthy calculations. Part (d) was only answered well by the best candidates who were able to systematically explain how behavioural data could be used to infer the variance of the internal ML estimators formed by subjects in the task. Parts (e) and (f) were correctly answered by most, but only few presented a rigorous derivation.

[Q2]

A very popular question, attempted by 46 candidates (mean 13.6/20, min 4/20, max 19/20).

Part (a) posed no particular difficulty on average, although in (a.iii) many students forgot to mention motor noise as a key source of uncertainty in decision making, and also did not give a convincing reason why making optimal decisions requires adequate treatment of the uncertainty (many answers gave only a generic argument about Bayesian inference without making it specific to decision making). Part (b.i) was generally well answered, though – strangely – many candidates decided to sketch the voltage dependence of the dyanmics of sodium and potassium gate variables, instead of their time course during the action potential as required. Part (b.ii) was reasonably well answered. Those who did answer correctly generally gave good arguments to back up their choices.

[Q3]

This unpopular question was tackled by less than 50% of the students. The distribution of marks was balanced (mean 12.4/20, min 6/20, max 20/20). The hardest part turned out to be part b. Many students had written a lot but missed the main point(s) in the answer. Many got about half of the marks in part c. Few did well in c-ii, but somehow more students nevertheless intuited the answer to part c-iii.

[Q4]

In retrospect, this question turned out to be very easy. Most of the students attempted it and the average mark was the highest among the 4 questions (mean 14.9/20, min 9/20, max 20/20). The main culprit was breaking down part b into too many sub-parts that became easier to answer.