Version ML/3 – CRIB

EGT2
ENGINEERING TRIPOS PART IIA

Monday 10 May 2021 09:00 to 10:40

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 and at the top of each answer sheet.*

STATIONERY REQUIREMENTS
Write on single-sided paper.

SPECIAL REQUIREMENTS TO BE SUPPLIED FOR THIS EXAM
CUED approved calculator allowed.
You are allowed access to the electronic version of the Engineering Data Books.

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

The time taken for scanning/uploading answers is 15 minutes.

Your script is to be uploaded as a single consolidated pdf containing all answers.

1 (a) Suppose a neuron is well described by an *integrate and fire model* with a fixed firing threshold, V_{thr}, and with sub-threshold membrane potential, V, modelled as a single passive compartment with capacitance C and Ohmic membrane conductances, g_i, with reversal potentials, E_i:

$$C\,\frac{dV}{dt} = \sum_i g_i\ (E_i - V)$$

(i) Write down expressions for the resting potential, V_{m}, of the neuron and the membrane time constant, τ. [5%]

Answer: Setting the membrane equation to steady state and solving gives:

$$V_m = \frac{\sum_i g_i E_i}{\sum_i g_i}.$$

Rearranging the membrane equation in the form $\tau\dot{V} = V_m - V$ gives:

$$\tau = \frac{C}{\sum_i g_i}.$$

(ii) An experimentalist applies a train of depolarising impulses to the neuron that cause the membrane potential to increase by an amplitude S, as shown in Fig. 1. Each pulse is separated by a time interval, T, equal to the membrane time constant, τ. Assuming the neuron remains below firing threshold, V_{thr}, derive an expression for the peak membrane potential, V_{p}. What is the minimum impulse amplitude, $S = S_0$, that will eventually cause the neuron to reach the threshold? [20%]

Answer: Solving $\tau\dot{V} = V_m - V$, with no input we have

$$V(t) = V(0)\exp(-t/\tau) + V_m.$$

If the membrane receives a pulse of magnitude S every interval $T = \tau$ then after N pulses we have, by linearity:

$$\begin{aligned} V(\tau N) - V_m &= S + S\exp(-\tau/\tau) + S\exp(-2\tau/\tau) + \ldots + S\exp(-N\tau/\tau) \\ &= S(1 + \mathrm{e}^{-1} + \mathrm{e}^{-2} + \ldots + \mathrm{e}^{-N}), \end{aligned}$$

a geometric series. Therefore as $N \to \infty$, $V(\tau N) \to V_p = V_m + \frac{S}{1-\mathrm{e}^{-1}}$.
Thus for the neuron to eventually fire, $S_0 > (V_{\text{thr}} - V_m)(1 - \mathrm{e}^{-1})$.

(iii) The experimentalist next applies a drug that introduces an additional conductance, g_{m}, with reversal potential E_{m}. Suppose g_{m} is equal to the pre-drug membrane conductance, $\sum_i g_i$. Assuming the impulses are kept at the same amplitude, $S = S_0$, derive an expression for the minimum value of E_{m} that will allow the neuron to fire. Write your expression in terms of V_{thr} and V_{m}. [20%]

Answer: Using the expression for V_m in part (i), in the presence of the drug the neuron will settle to a new equilibrium potential V_m^* where

$$V_m^* = \frac{g_m E_m + \sum_i g_i E_i}{g_m + \sum_i g_i} = \frac{V_m + E_m}{2}$$

(since $g_m = \sum_i g_i$.) Also, from (i) the new membrane time constant will be $\tau' = \tau/2$. Therefore, with the input stimuli kept as in (ii), the maximum deflection will be:

$$\begin{aligned} V(\tau N) - V_m^* &= S + S\exp(-2\tau/\tau) + S\exp(-4\tau/\tau) + ... + S\exp(-2N\tau/\tau) \\ &= S(1 + \mathrm{e}^{-2} + \mathrm{e}^{-4} + ... + \mathrm{e}^{-2N}) \\ &\to \frac{S}{1 - \mathrm{e}^{-2}}. \end{aligned}$$

Substituting $S = (V_{\text{thr}} - V_m)(1 - \mathrm{e}^{-1})$ we see that the neuron will fire if:

$$V_{\text{thr}} - V_m^* < V(\tau N) - V_m^* = \frac{(V_{\text{thr}} - V_m)(1 - \mathrm{e}^{-1})}{1 - \mathrm{e}^{-2}} = \frac{(V_{\text{thr}} - V_m)}{1 + \mathrm{e}^{-1}}$$

Rearranging gives the condition for firing in terms of E_m:

$$E_m > \frac{1 - \mathrm{e}^{-1}}{1 + \mathrm{e}^{-1}} V_m + \frac{2\mathrm{e}^{-1}}{1 + \mathrm{e}^{-1}} V_{\text{thr}}.$$

Addendum: since the question for this part does not explicitly state that the inter-pulse intervals (IPI) are kept fixed to their value in part ii (while it explicitly says their amplitude is kept fixed), we decided to mark as correct the answers that assumed the new IPI is equal to the new membrane time-constant (i.e. those that assumed the "rule" for setting IPI, rather than its value, is held fixed). In this case the voltage decay between the pulses will remain equal to e^{-1}, and thus the maximum deflection will now be given by:

$$\begin{aligned} V(\tau N) - V_m^* &= S(1 + \mathrm{e}^{-1} + \mathrm{e}^{-2} + ... + \mathrm{e}^{-N}) \\ &\to \frac{S}{1 - \mathrm{e}^{-1}}. \end{aligned}$$

Substituting $S = (V_{\text{thr}} - V_m)(1 - \mathrm{e}^{-1})$ we see that the neuron will fire if:

$$V_{\text{thr}} - V_m^* < V(\tau N) - V_m^* = V_{\text{thr}} - V_m$$

Rearranging then yields the condition:

$$E_m > V_m.$$

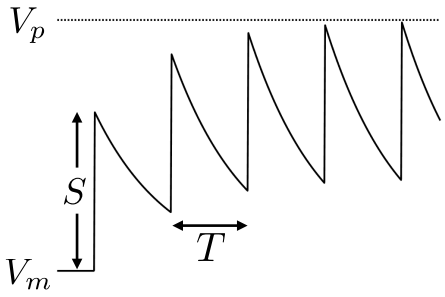

Fig. 1

(b) Explain why perception can be thought of as probabilistic inference, and why it relies on prior knowledge and prior expectations. [15%]

Answer: Sensory inputs are almost always ambiguous and noisy, so behaviourally relevant state variables of the world cannot be inferred from such inputs with full certainty. Probabilistic inference is the self-consistent framework for inference under uncertainty, and thus provides a good theoretical framework for understanding perception. To infer the variables of interest as optimally as possible (e.g. to partially resolve the ambiguities in sensory inputs or to partially denoise them), the brain can rely on prior knowledge of the world and of the regularities governing its state variables, gained through individual experience or during evolution.

(c) The left panel of Fig. 2 depicts a checkerboard, with dark and light squares, and a cylinder casting a shadow on the checkerboard. Square *A* is commonly perceived darker than square *B*, even though the pixels in these two regions of the image have exactly the same physical colour. (To check that squares *A* and *B* indeed have the same physical colour, the right panel of Fig. 2 shows the same image as the left panel but with two superimposed vertical bars that have uniform colours.)

Explain the following aspects of this visual illusion using the probabilistic inference model of perception. (Think of the pixel colours as the sensory inputs/variables, s.)

(i) Specify the following components (by reference to features of the objects or object-parts depicted in the left panel of Fig. 2, or to the rules and regularities governing them) and briefly justify your answer in each case:

A. the variable(s) inferred by perception; [10%]

Answer: The intrinsic surface colour (or reflectance) of the different objects and object parts, and the checks in particular, as intrinsic colors tend be behaviourally relevant and more invariant; more so than the pixel colours carried by sensory input s.

B. a variable that plays a causal role in determining s, but is (partially) discarded or compensated for by perception; [10%]

Answer: Variations in illumination, in particular due to the shadow cast by the cylinder. Because these variables are typically less relevant to our behavioural needs and goals.

C. key pieces of prior knowledge/expectations. [10%]

Answer: (1) The squares in the checkerboard are one of two types: dark (D) vs. light (L). (2) The diagonally connected dark squares all have the same intrinsic surface colour; same for the light checks. (3) Objects cast shadows which tend to have smooth borders. (4) A surface of given intrinsic colour appears darker under a shadow than outside of it.

(ii) Describe how the components you described above come together to explain the illusion. [10%]

Answer: The brain is primarily interested in inferring the intrinsic colours of the squares, x, and in the process of doing so it partially discards or compensates for the illuminant/shadow, which together with x determines the sensory input s. Apparent colour (s) can be used to locally decide if a square is darker or brighter than its neighbours. Then given the prior expectation that all dark squares are of the same intrinsic colour, the visual system infers that square A is dark ($x_A = D$) while square B is bright ($x_B = L$). Thus A is perceived as darker than B.

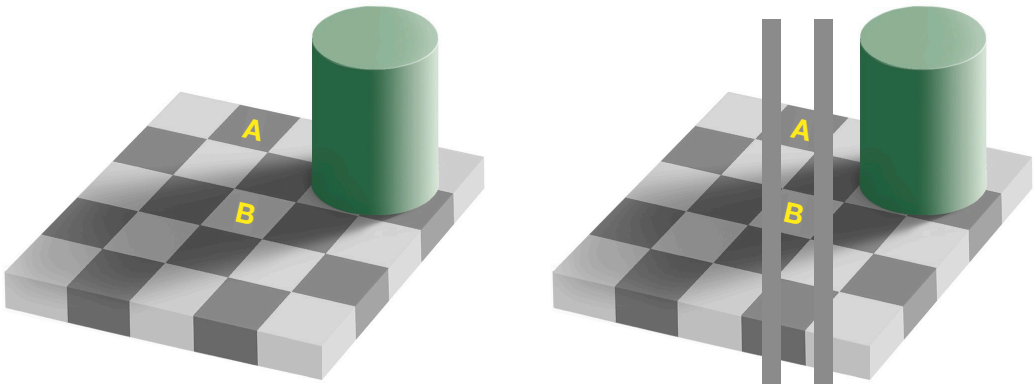

Fig. 2

2 (a) Do conductances that depolarise the neuronal membrane potential always increase the likelihood of an action potential? Explain your reasoning. [10%]

Answer: No. A conductance with reversal potential above resting potential can shunt stimulus current if the conductance is large enough.

(b) This question is about the drift-diffusion model of decision making, whereby a decision variable, x_t, starts at $x_0 = 0$ and evolves according to the stochastic differential equation

$$dx_t = \mu\, dt + dW_t$$

until it reaches one of the two bounds at $\pm B$. (Note that here we have fixed the coefficient of dW_t to 1, but allow B to be a free model parameter.) Explain, with reasons, how increasing B affects the decision maker's speed and accuracy, and thus the speed-accuracy trade-off. [20%]

Answer: The higher B, the later x_t will reach one or the other bound, on average, as it has to diffuse or drift a longer distance to reach the bounds. Thus increasing B increases the average decision time and lowers speed. Since the average decision time increases, more evidence is accumulated and thus the decisions become more accurate. In sum, increasing B favors accuracy over speed, in the speed-accuracy trade-off.

(c) Crickets have two "cerci" (Fig. 3, left, rear appendages in black), each covered with hairs that are deflected by air currents, thus providing information about the wind direction θ (relative to heading direction). For low air velocity, only four sensory neurons relay information about θ to the brain. Specifically, the firing rate f_i of the i^{th} neuron is a half-wave rectified cosine function of θ (Fig. 3, right), $f_i(\theta) = 40\,[\cos(\theta - \theta_i)]_+$, where θ_i is the neuron's preferred wind direction and $[x]_+$ equals x if $x > 0$ and 0 otherwise.

(i) Show that $f_i(\theta) = 40\,[\vec{v} \cdot \vec{c}_i]_+$, where $\vec{v}$ is the normalised air velocity vector and $\vec{c}_i$ is the normalised preferred direction vector of neuron i (Fig. 3, left). [10%]

Answer: This follows directly from the fact that the dot product between two unit-norm vectors (here, $\vec{v}$ and $\vec{c}_i$) is equal to the cosine of the angle between the two vectors. Here, the angle between $\vec{v}$ and $\vec{c}_i$ is equal to $\theta - \theta_i$, hence $40\,[\vec{v} \cdot \vec{c}_i]_+ = 40\,[\cos(\theta - \theta_i)]_+$.

(ii) Assume the wind direction is constant. Give an expression for the reconstruction (or "decoding") of $\vec{v}$ from the firing rates of these four neurons, assuming a sufficiently long time window is used to obtain near-perfect rate estimates. Explain why you think this system uses four neurons. [30%]

Answer: It is easy to see that if $\vec{v}$ is in the quadrant defined by $\vec{c}_1$ and $\vec{c}_2$, then $\vec{v} = (\vec{v}\cdot\vec{c}_1)\vec{c}_1 + (\vec{v}\cdot\vec{c}_2)\vec{c}_2 = (r_1/40)\vec{c}_1 + (r2/40)\vec{c}_2$ where r_i is the firing rate response of neuron i. If $\vec{v}$ is in the quadrant defined by $\vec{c}_2$ and $\vec{c}_3$, then $r_1 = 0$ and we need to add a contribution from the third neuron (and it does not hurt to keep the r_1 term, since it is zero): $\vec{v} = (r_1/40)\vec{c}_1 + (r2/40)\vec{c}_2 + (r_3/40)\vec{c}_3$. Repeating the argument for the other two quadrant, we arrive at a general formula that works for any $\vec{v}$: $\vec{v} = \frac{1}{40}\sum_i r_i\vec{c}_i$; in other words, the wind direction can be reconstructed as a weighted sum of the four neurons' preferred directions, weighted by the corresponding neuron's firing rate. We must rely on four neurons, because firing rates cannot be negative, so e.g. a negative contribution of $\vec{c}_1$ to the current wind direction must be encoded by an opponent neuron (neuron 3).

(iii) Now, assume that firing rate estimates are noisy, e.g. based on a finite time window. What mathematical approach would you use to formalise optimal decoding of the wind direction? [30%]

Answer: There are several valid ways of formalising optimal decoding based on noisy firing rate responses. One close to what we have discussed in lectures is to choose an estimate $\hat{\theta}$ that minimises the expected squared distance $\langle(\hat{\theta} - \theta)^2\rangle_{p(\theta)}$, where $p(\theta)$ is some prior over wind direction (e.g. isotropic). This estimate would be the mean of the posterior distribution $p(\theta|r_1,\ldots,r_4) \propto p(r_1,\ldots,r_4|\theta)p(\theta) = p(\theta)\prod_i p(r_i|\theta)$. Here, the likelihood term $p(r_i|\theta)$ describes the noise in the neural responses (e.g. Poisson distribution with a mean determined by the tuning curves).

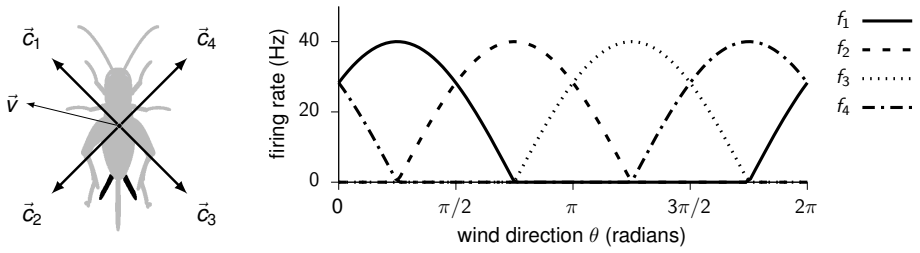

Fig. 3

3 (a) Write short notes on the following:

(i) Dale's principle; [5%]

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

(ii) ionotropic receptor; [5%]

Answer: It is a class of neurotransmitter receptors that are also ion channels at the same time, so that ligand binding leads to the opening of the channel in the same molecule.

(iii) spaced training; [5%]

Answer: It is a kind of training in which training events (or sessions), e.g. of CS-US pairs, are presented with a large time interval between them, such as weeks. It typically leads to a slower build-up but longer-term retention of training effects associated with structural changes, e.g. the creation of new or the elimination of existing synapses.

(iv) conditioned stimulus; [5%]

Answer: A stimulus to which an untrained animal does not give any specific response, but after training (systematically pairing it with an unconditioned stimulus, US) it does evoke a response (without the US).

(v) Mg^{2+} block. [5%]

Answer: The NMDA receptor's extracellular mouth is blocked by a Mg^{2+} ion below a threshold membrane potential value (around −40 mV) such the receptor remains closed unless the postsynaptic cell is depolarised above this threshold. This means that the opening of the NMDA receptor on the postsynaptic cell requires a coincidence of presynaptic activation (glutamate release by the presynaptic cell, sensed by glutamate binding to the receptor) and postsynaptic activation (postsynaptic membrane potential exceeding the threshold). This is critical for some forms of associative learning, e.g. in the *Aplysia*.

(b) This question is about the role of hippocampal learning in spatial navigation.

(i) Describe an experimental approach for studying how navigation depends on information stored in the hippocampus depending on the amount of experience. In your answer, describe the following:

A. the apparatus; [5%]

Answer: The plus maze experiment is a classical example for this. As the name suggests, this uses a plus-shaped maze as an apparatus.

B. the behavioural paradigm for training and testing; [10%]

Answer: During training the animal (rat) is always started from the same (e.g. the 'south') arm of the maze and is rewarded (by food or drink) by going to another fixed (e.g. the 'west') arm of the maze. In test trials (a single test trial per animal), the animal is started from a novel (e.g. the 'north') arm of the maze, such that the originally trained arm is at a different relative location to its starting location than during training.

C. the key behavioural measure of learning; [10%]

Answer: Performance in test trials is measured by whether the animal chooses the originally trained arm (i.e. its behaviour is driven by the spatial location of that arm), called a 'place' choice, or the arm that is in the same relative location to the new starting location as the originally trained arm was to the starting location used during training, and thus requires the same sequence of motor actions to be reached, a 'response' choice.

D. the pharmacological interventions that need to be used in the experiment: the drug(s), the time at which they are applied, and the brain structure(s) to which they are applied. [10%]

Answer: Reversible inactivation of either the hippocampus or, as a control, the striatum (caudate nucleus) by lidocaine (local anaesthetic, voltage dependent sodium channel blocker which does prevents action potential firing) prior to testing can be used to test the respective contributions of these brain areas to navigation. As a control, saline needs to be injected to the same brain areas.

(ii) Describe the results you expect from the experiment you specified above, and what these results tell us about the kind of information that is and is not stored in the hippocampus, and when that information is used for navigation. [20%]

Answer: Basic finding with control animals: after short training 'place' choices dominate, after extended training 'response' choices dominate. These results are abolished, or even reverted, when hippocampus is inactivated after short, but not extended, training, and when striatum is inactivated after extended, but not short, training. This shows that hippocampal information is predominantly used for navigation after short, but not extended training, and that this information is fundamentally spatial. In contrast, information stored in the striatum is used for navigation after extended, but not short training, and this information is fundamentally response-based.

(iii) Explain how you would extend the experimental approach above to show that

the hippocampal information used for navigation is stored by LTP. In your answer, include any changes to the behavioural paradigm or any additional drugs that you may use, as well as the results you would expect from the experiment and the interpretation of those results. Make sure that your proposed experiment controls for potential effects of LTP on retrieval, and the effects of hippocampal LTP on storing other kinds of information. [20%]

Answer: Inject a selective NMDA receptor antagonist (e.g. AP5), or as a control saline, to the hippocampus chronically, during training, or as a control after training. After short training, the dominance of place choices should be abolished, even without lidocaine injection. After long training, there should be no such effect, i.e. the dominance of response choices should remain intact. Saline injection should have no effect in either case. There should also be no effect if NMDA receptor antagonist is injected after training. This means that storage, but not retrieval, of place information in the hippocampus is LTP-dependent.

4 (a) This question is about classical conditioning in the *Aplysia*.

(i) Name three different neurotransmitters that are involved in the induction of classical conditioning. For each neurotransmitter, explain which cell releases it, on which cell it exerts its effects, the kind of molecules the target cell uses to bind it, and the effects binding has on the target cell. [30%]

Answer:

•Serotonin (5-HT) [10%] is released by the facilitating interneuron due to the activation of the US-pathway (from the tail) and acts on membrane-bound metabotropic receptors on the sensory neuron which are coupled (via a G-protein) to the adenylyl cyclase enzime. This enzime also acts as a coincidence detector between the CS and the US because CS activation (siphon stimulation) causes an action potential in the sensory neuron which leads to Ca^{2+} influx, and it is the increased Ca^{2+} concentration due to this influx that boosts the effects of serotonin receptor-activation on the enzyme. The adenylyl cyclase produces cyclic AMP (cAMP) which in turn activates three different pathways modifying the future behaviour of the cell such that 1. voltage-dependent K^+ current is decreased (longer action potential), 2. more vesicles are in the active zone (more transmitter released), 3. voltage-dependent Ca^{2+} increased (more transmitter released).

•Glutamate (Glu) [10%] is released by the sensory neuron and binds to membrane-bound glutamate receptors on the motor neuron, including NMDA receptors. These glutamate receptors depolarise the motor neuron and thus lead to a response (gill withdrawal). In addition, NMDA receptors act as coincidence detectors of presynaptic activation (caused by the CS) and postsynaptic activation (which could be due to direct activation from the US, as the US-pathway synapses onto the motor neuron). The mechanism of this is described in the answer to the previous question. NMDA receptor activation leads to Ca^{2+} influx, which activates a number of biochemical pathways in the motor neuron including an increase in postsynaptic non-NMDA glutamate receptor-dependent responses and the production of retrograde messengers (see below).

•Nitrogen monoxide (NO) [10%] or carbon monoxide (NO) are released as retrograde messengers by the motor neuron to the sensory neuron. Their release is one of the events triggered by NMDA receptor activation (described above). They do not bind to membrane-bound receptors in the sensory cell but instead bind directly to enzymes in its cytoplasma and result in enhanced transmitter release.

(ii) Describe the differences between the training protocols used to induce (short-term) sensitisation and conditioning, and the resulting differences in the behavioural effects of these two forms of learning. In addition, explain the mechanisms by which differences in training lead to the differences in behaviour. [20%]

Answer:

•Differences in training protocols [5%]. During sensitisation, the "CS" (harmless stimulus) is only used for testing but it is never used in training. Only the "US" (noxious stimulus) is used during training. During conditioning, the CS is consistently paired with the US during training. This means that the CS must precede the US within a short time interval (< 1 s).

•Differences in behaviour [5%]. Following sensitisation, *any* initially harmless stimulus evokes a response, while after conditioning only the CS that was paired with the US evokes a response, i.e. the learning is *associative*: the US and the designated CS have been associated. Typically, the effects of conditioning are also greater and last longer than those of sensitisation.

•Mechanisms linking differences in training to differences in behaviour [10%]. During sensitisation, it is only the adenylyl cyclase in the sensory neuron that is activated, and only via the serotonin-receptor pathway. That primes the sensory neuron so that action potentials generated by later CS (or "CS") stimulation lead to greater responses in this cell (see above), but it is only this cell that changes its behaviour and these changes do not typically last very long (at least for short-term sensitisation). In contrast, during conditioning, the adenylyl cyclase in the sensory neuron is already primed by stimulation by the CS before the serotonin-receptor pathway is activated and so its activated for a longer time. In addition, the functioning of the motor neuron is also altered byNMDA receptor activation (see above). Note that both the strong activation of the adenylyl cyclase in the sensory neuron and NMDA receptor neuron activation in the motor neuron are dependent on the *coincidence* of the CS and the US, such that the CS must precede the US. This is because only in that way is the adenylyl cyclase primed already (by the CS-evoked action potential in the sensory neuron) by the time serotonin is released onto the sensory neuron by the facilitatory interneuron (due to the US), and only then is there still glutamate (released by the CS-activated sensory neuron) binding to the NMDA receptor at the time the motor neuron is depolarised (by the US). Thus these effects will be specific to the CS-pathway. Other sensory pathways transmitting other harmless stimuli may be transiently sensitised but these effects will diminish much faster than those of conditioning, by which time only the conditioned pathway will lead to responses, leading to specificity of the association.

(b) The activity of GABAergic neurons in the ventral tegmental area has been shown to strongly correlate with reward prediction.

(i) Describe how you expect the activity of such a neuron to vary over time during the course of a classical conditioning trial (using a single CS) before and after training. [20%]

Answer: Before learning: it should be zero (i.e. at baseline) throughout the trial. After learning: it should be zero beofre the CS, elevated between the CS and the US, and then decay back to zero following the US.

(ii) Based on their main neurotransmitter, what effect do you expect these neurons to have on dopaminergic neurons in the VTA, and how would you interpret this effect in the context of the Rescorla-Wagner learning theory of dopaminergic activity? [30%]

Answer: GABA is an inhibitory neurotransmitter, and so these neurons should inhibit dopaminergic neurons. This is consistent with the activity of dopaminergic neurons signalling reward prediction errors, according to the Rescorla-Wagner theory, because reward prediction errors can be computed by subtracting reward prediction from actual reward, which thus needs reward prediction to be taken

into account with a negative sign. Inhibitory neurons conveying a reward prediction signal are ideally suited for this, as inhibition can effectively substract from the postsynaptic neuron's activity.

END OF PAPER

THIS PAGE IS BLANK

Assessor's comments:

Q1 Parts *c* and *b* (in that order) were answered well by most candidates who attempted this question – though some who answered *b*, more or less stated the question and provided the formula for Bayes' rule and meaning of likelihood and prior, and missed the main point of the question. Candidates did worse on part *a*, and in particular only one candidate got full mark on part *iii*.

Q2 Almost all candidates got parts *b* and *c(i)* right, and more than half got part *a* right. The hardest part was *c(ii)*, but even here most had indicated enough understanding to get at least half the mark. Most who got full mark did not do it the way itÕs written in the crib and did not provide that expression (instead they used arc-cosine to get an expression for the angle itself, as opposed to the unit vector indicating the wind direction).

Q3 Part *a* was generally well solved, though many candidates failed to note for spaced training that it typically results in structural changes, or described it as relevant only to habituation. Part *b* was more challenging, though still based closely on lecture material. Instead of the plus maze (as described in lectures and in the crib), many candidates chose the Morris Water maze (MWM) as the basic apparatus, with which it is more challenging to study the required question experimentally – especially with its standard form using a starting position that always changes. Many candidates also failed to address how hippocampal contributions to navigation depend on the amount of experience, or used lidocaine during training rather than testing, or (in part *b(iii)*) failed to control for the effects of drugs (AP5) during testing.

Q4 A popular question. Part *a(i)* was generally well solved, though many candidates failed mention NO (or CO) as a neurotransmitter and included GABA instead which does not have a specific role in this context. Part *a(ii)* was also generally well solved, though few students noted both major differences in the effects of sensitisation and conditioning: that the latter is input specific and also stronger (longer lasting). Part *b* was challenging, requiring substantial generalisation from lecture material, and very few candidates could solve it (near) perfectly. While almost all candidates realised that a GABAergic neuron is inhibitory, many mistook this for receiving (rather than "sending") signals with a negative sign. Many candidates did not quite understand the distinction between reward prediction error (signalled by dopamine) and reward prediction (signalled by the cells in question),

and almost all (even those who included the correct equation for the Rescorla-Wagner, or the temporal difference learning rule) failed to notice that computing the former requires the latter with a negative sign, thus making the neurons in question ideally placed to convey this signal to dopaminergic neurons.