Version GH/3 – CRIB

EGT2
ENGINEERING TRIPOS PART IIA

Tuesday 23 April 2019 14:00 to 15:40

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.*

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.

1 (a) In a few sentences, explain what a spike train is, and describe the classical problem of "elucidating the neural code for sensation". [20%]

Answer: A spike train is a sequence of action potentials emitted by a neuron over a certain time window, e.g. in response to a sensory stimulus. This sequence can have zero, one, or more spikes, and is typically discretized using some finite time resolution (typically 1 to 10 ms) and represented as a long binary word. Spike trains are very high-dimensional mathematical objects. Elucidating the neural code for sensation means understanding two aspects:

- what information about the stimulus is encoded in the spike train (e.g. luminance, intensity of pressure on the skin, ...), and
- how that information is represented in the spike train, i.e. which features of the spike train carries that information (usual suspects include the total number of spikes fired, the latency of the first spike, the first interspike-interval, ...).

(b) Consider a cylindrical neurite of radius R, with axial resistivity r_a (in $\Omega \cdot \mathrm{mm}$), unit membrane resistance r_m (in $\Omega \cdot \mathrm{mm}^2$), and unit membrane capacitance c_m (in $\mathrm{nF}/\mathrm{mm}^2$).

(i) Explain, with reason, whether the following statement is true or false: "A neurite of radius $2R$ has the same membrane time constant as a similar neurite of radius R". [15%]

Answer: True. The membrane time constant is the product $r_m c_m$ which does not depend on the radius of the neurite.

(ii) Explain, with reason, whether the following statement is true or false: "The total axial resistance to current flow is the same for a portion of neurite of length L and radius $2R$ as it is for a portion of length $L/2$ and radius R (of an otherwise similar neurite)". [15%]

Answer: False. The axial resistance to current flow is inversely proportional to the number of charge carriers in a cross section, which in turn scales with R^2. On the other hand, the longitudinal resistance is proportional to the length L of the portion of neurite considered. Thus, the resistance to current flow is lower (by a factor of two) for a length L and radius $2R$ than for a length $L/2$ and radius R.

(iii) Write down a formula for the *length constant* of the neurite. How does it depend on the ion channel density in the membrane? [20%]

Answer: The length constant, λ, is given by:

$$\lambda = \sqrt{\frac{R r_m}{2 r_a}}$$

where R is the radius of the neurite [mm], r_m is the unit membrane resistivity [$\Omega \cdot \mathrm{mm}^2$], r_a is the unit axial resistivity [$\Omega \cdot \mathrm{mm}$]. The membrane resistivity is inversely proportional to the channel density,

ρ, therefore the length constant scales as

$$\lambda \propto \sqrt{1/\rho}$$

(iv) Consider a brief current pulse injected somewhere along the neurite. Explain, with the aid of a sketch, how the propagation of this input signal along the neurite depends on the length constant. You may assume the neurite has infinite length. [20%]

Answer: For a transient pulse, the spatial profile of the voltage along the neurite begins as a very sharp peak at the injection site, and broadens over time with a decaying amplitude (right). At any time t, the spatial profile is bell-shaped with lengthscale $2\lambda\sqrt{t/\tau_m}$ where λ is the length constant of the neurite, and $\tau_m = r_m c_m$ is the membrane time constant. Therefore, the length constant determines how far the signal will have propagated after a given time t/τ_m.

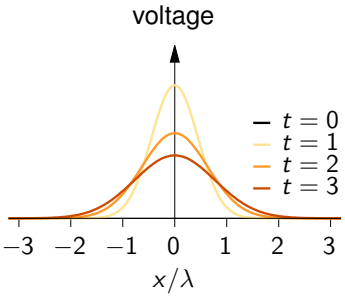

(v) Describe three ways that signal propagation can be improved in a neurite, along with any disadvantages or biologically relevant constraints for each. [10%]

Answer: Signal attenuation can be reduced by either (1) increasing the neurite diameter, R; (2) reducing the channel density; or (3) using active channels (e.g. action potential). Option (1) is subject to space requirements and/or metabolic load due to large cell volume; option (2) may make the neuron vulnerable to ion channel noise; option (3) may make it difficult to sum input linearly (for example).

2 (a) Write short notes on the following:

(i) The physical and biological features from which uncertainty in sensory perception arises; [15%]

Answer: The sensory transduction process is always noisy; for example, photoreceptors exhibit spontaneous voltage fluctuations due to various sources of intrinsic noise or shot noise. Moreoever, even if sensors were noiseless, the sensory representation of the world would remain fundamentally ambiguous due to the physical properties of the world and of our sensors; for example, we only have two eyes and the front of a object almost always obstructs its back, meaning that the full 3D extent of an object present in the visual scene can never be fully and exactly reconstructed.

(ii) The properties of mechanoreceptors that determine the spatial resolution of stimulus encoding. [15%]

Answer: The sizes of their receptive fields; their noise levels; and the spatial density of receptors.

(b) A subject engages in an auditory detection task. Brief flashes of light mark the beginning and end of each trial. In each trial, a pure-frequency tone with intensity $I > 0$ is presented with 50% probability; this divides trials into "ON trials" (sound) and "OFF trials" (no sound).

The subject's brain gathers noisy sensory evidence summarized by a number s (arbitrary units), assumed to be drawn in each trial from a normal distribution. This distribution has a fixed unit variance $\sigma^2 = 1$, and a mean μ directly proportional to I in ON trials and equal to zero in OFF trials. The subject then reports hearing a sound if s is greater than the subject's "internal decision threshold" θ, and not hearing a sound otherwise.

(i) Give expressions for the hit rate (the probability to report hearing a sound in an ON trial) and false alarm rate (the probability to report hearing a sound in an OFF trial), as functions of μ and θ. You may use the standard normal cumulative density function defined as

$$\Phi(y) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{y} e^{-x^2/2} dx.$$

[15%]

Answer: The hit rate is $P(s > \theta)$ under $s \sim \mathcal{N}(\mu, 1)$, and is therefore equal to $1 - \Phi(\theta - \mu)$. Similarly, the false alarm rate is $P(s > \theta)$ under $s \sim \mathcal{N}(0, 1)$, which equals $1 - \Phi(\theta)$.

(ii) Sketch the corresponding Receiver Operating Characteristic (ROC) curve *qualitatively*, for $\mu = 1$ and for $\mu = 2$, and provide a brief interpretation. [20%]

Answer:

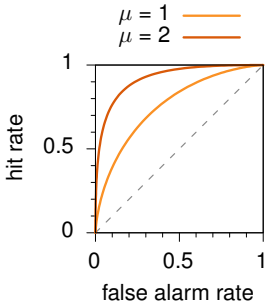

The ROC curve is defined as the hit rate shown as a function of the false alarm rate as the subject varies their internal decision threshold θ. See the required plots on the left. These curves always begin at (0,0) and end at (1,1) as the threshold goes from $+\infty$ to $-\infty$. As the sound intensity (in ON-trials) increases (i.e. μ increases), the curves pass nearer the ideal point (0,1): it becomes easier and easier to place the threshold in such a way as to achieve a good hit rate while keeping the false-alarm rate low.

(iii) The experimenter now pays the subject one pound for each hit, but deducts from their payment L pounds for each false alarm. In order to maximize their expected payment, where should the subject place their internal decision threshold θ? Derive an expression as a function of μ and L. [20%]

Answer: The expected return is given by $\mathscr{L}(\theta) = 1 - \Phi(\theta - \mu) - L(1 - \Phi(\theta))$. Differentiating w.r.t. θ gives (up to a multiplicative constant)

$$\frac{d\mathscr{L}}{d\theta} \propto -\exp(-(\theta-\mu)^2/2) - L\exp(-\theta^2/2)$$

Setting this expression to 0 and solving for θ gives the optimal threshold

$$\theta = \frac{\mu}{2} + \frac{\log L}{\mu}.$$

(iv) How would you build on the above experiment to investigate whether humans make optimal decisions? [15%]

Answer: One could run the experiment in multiple sessions, with a different value of the sound intensity I in each session. By measuring hit and false alarm rates, one can then infer the subject's internal threshold, and compare it to the optimal threshold. Because certain modelling assumptions have gone into deriving expressions for the hit and false-alarm rates (e.g. normal distribution, unit variance), one can only make such comparisons on a qualitative level. As an example of a qualitative feature, we may note that for $L > 1$ the optimal threshold has a minimum as a function of μ; this could be checked for. Moreover, this minimum scales as $\sqrt{\log(L)}$. This dependence could be checked e.g. by testing different subjects with different values of L.

3 This question is about evidence accumulation. A monkey performs a classical random dot motion discrimination task. In each trial, a random dot field moves with some coherence level c on a display, in direction $m \in \{L;R\}$ (Left or Right) chosen randomly with equal probability. The monkey must infer m and report its choice as soon as possible by making a saccade in the corresponding direction, thereby ending the trial, collecting a reward (if correct), and proceeding to the next trial.

During the experiment, a scientist records the firing activity of a single neuron in area MT of this monkey. Time in each trial is discretized in successive, adjacent bins of 50 ms duration indexed by an integer variable t ($t = 1,2,\ldots$). The response x_t of this neuron in each time bin is binarized such that $x_t = 0$ if the neuron is silent in that bin, and $x_t = 1$ if at least one action potential is produced. The recorded neuron is found to fire stochastically, such that for a trial with direction m, its response x_t appears to be drawn independently in each time bin from a Bernoulli distribution with mean α_m.

(a) What recording technique is this scientist most likely to be using, and why? [10%]

Answer: In-vivo recordings of action potentials in awake animals are most often performed extracellularly (e.g. via metal microelectrodes or glass micropipettes), especially when there is no need to record subthreshold voltage fluctuations (as is apparently the case here).

(b) Another area of the monkey's brain observes the responses of the recorded neuron as they unfold during each trial, and represents the cumulated log-likelihood ratio

$$D_t = \log \frac{p(x_1, x_2, \ldots, x_t | m = R)}{p(x_1, x_2, \ldots, x_t | m = L)}$$

in each time bin of each trial. Derive a simple update rule for D_t as a function of D_{t-1}, x_t, α_L and α_R. [Hint: the probability mass function for a Bernoulli variable x of mean α can be conveniently written as $p(x) = \alpha^x (1-\alpha)^{1-x}$.] [30%]

Answer:

$$\begin{aligned}
D_t &= \log \frac{\prod_{i=1}^{t} \alpha_R^{x_i} (1-\alpha_R)^{1-x_i}}{\prod_{i=1}^{t} \alpha_L^{x_i} (1-\alpha_L)^{1-x_i}} \\
&= \log \prod_{i=1}^{t} \left(\frac{\alpha_R}{\alpha_L}\right)^{x_i} \left(\frac{1-\alpha_R}{1-\alpha_L}\right)^{1-x_i} \\
&= \log \prod_{i=1}^{t-1} \left(\frac{\alpha_R}{\alpha_L}\right)^{x_i} \left(\frac{1-\alpha_R}{1-\alpha_L}\right)^{1-x_i} + \log \left(\frac{\alpha_R}{\alpha_L}\right)^{x_t} \left(\frac{1-\alpha_R}{1-\alpha_L}\right)^{1-x_t} \\
&= D_{t-1} + x_t \log \frac{\alpha_R}{\alpha_L} + (1-x_t) \log \frac{1-\alpha_R}{1-\alpha_L}
\end{aligned}$$

(c) The monkey reports rightward motion as soon as D_t exceeds some threshold θ, and leftward motion as soon as D_t falls below $-\theta$. Assume $\alpha_L = 0.2$ and $\theta = 3$. Sketch the timecourse of D_t for two example correct "right trials" ($m = R$), one with $\alpha_R = 0.4$ and the other with $\alpha_R = 0.8$. Explain the differences between the two. [20%]

Answer: When the neuron fires in time bin t ($x_t = 1$), D_t experiences a positive jump by $\log(\alpha_R/\alpha_L)$. When the neuron is silent, the jump is negative, of size $\log((1-\alpha_R)/(1-\alpha_L))$. For $\alpha_R = 0.4$, we get slighly more negative jumps than positive jumps in D_t, but the positive jumps are (more than twice) larger. For $\alpha_R = 0.8$, negative jumps are exactly as large as positive jumps, but positive jumps are much more likely, such that the threshold is reached much sooner on average.

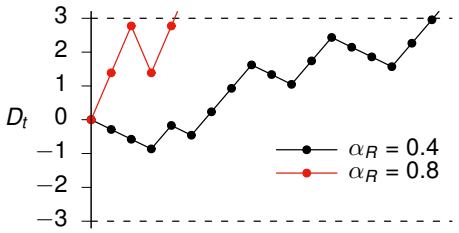

(d) The scientist now changes the frequency of occurence of left and right trials, such that $p(m = R) = 0.6$ (and therefore $p(m = L) = 0.4$). How would you alter the sequential update rule for D_t to take into account this information? [20%]

Answer: The log posterior ratio would then need to be used instead of the log likelihood ratio. Non-equal prior probabilities for $m = L$ and $m = R$ would result in an additive offset of D_t equal $\log(0.6/0.4)$, which is the right initial condition to use for D_0 before the trial even begins.

(e) Finally, the scientist searches for neurons that encode the decision variable D_t. Suggest one brain area where such neurons are likely to be found, and give two properties of firing responses that you would expect from such "decision neurons" but *not* from sensory neurons such as the one recorded by the scientist in area MT. [20%]

Answer: Such decision neurons are usually found in area LIP. In contrast to MT neurons that tend to display noisy but steady activity during the trial, LIP neurons have activity that ramps up (or down) from trial onset, at a rate that typically increases with the motion coherence. Moreover, aligning their activity to the decision time shows that LIP neurons whose activity ramps up during e.g. right trials consistently reach the same "threshold" value just before the eye saccade is made.

4 (a) Describe the evolutionary advantage of learning. Include the following in your answer:

- The relationship between genotype and phenotype, and how learning can affect this relationship.
- The circumstances under which learning can be beneficial, and the way in which learning can increase fitness under these circumstances.

You may illustrate your answer with simple diagrams. [40%]

Answer:

•As the diagram below illustrates, in general, the genotype (ie. the genes of an organism) determine the phenotype, which in turn determines the fitness of the organism in a given environment. The nervous system and the behaviour it produces are part of the phenotype and a major determinant of its fitness. Behaviour in particular changes the environment itself which in turn provides feedback to the nervous system in the form of sensory inputs. Learning is a process of reconfiguring the nervous system based on experience, and thus changing the behaviours it emits, and ultimately the fitness the organism achieves.

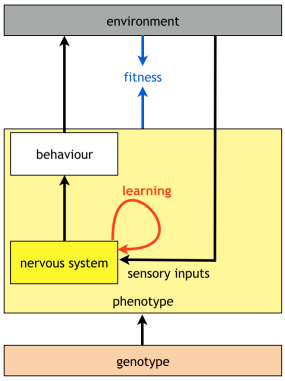

•When the environment, and thus the fitness landscape it induces, is non-stationary or the fitness landscape is very 'spiky' (see diagrams below), changes in the genome may have too low a probability to find high fitness regions of the fitness landscape. Learning helps by acting on faster time scales than genetic changes, thus mitigating the effects of environmental non-stationarities, and providing mechanisms by which organisms can adaptively change their phenotype rather than the 'random search' of genetic modifications.

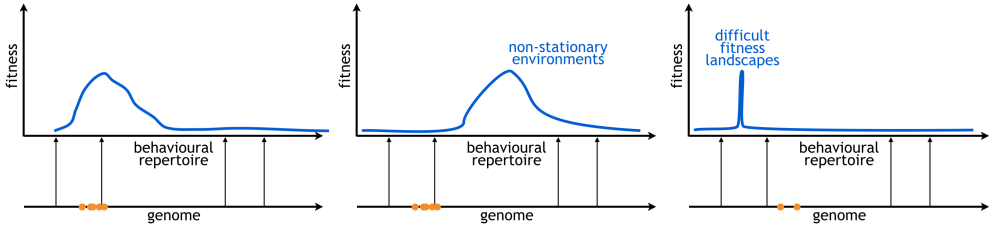

(b) Imagine that we engineer a new receptor channel that has a reversal potential which is equal to the resting potential of the cell. Sketch the postsynaptic potential trace when

(cont.

this receptor is activated at a time when the postsynaptic cell is at rest. In addition, explain the effect of receptor activation on the postsynaptic cell's behaviour, with special regard to whether this effect is excitatory, inhibitory, or neither. [30%]

Answer: The sketch should show that the postsynaptic potential remains constant (ie. a flat line) as a function of time at the resting membrane potential (about -70 mV). The postsynaptic potential does not change upon receptor activation because the reversal potential is the same as the membrane potential and so the driving force is zero. This however does *not* mean that there is no effect on the postsynaptic cell. The main effect is an increase in the net membrane conductance (or decrease in membrane resistance), ie. a shunt, which makes it harder for the cell to reach the threshold for emitting action potentials. This is simply due to Ohm's law: the decreased net membrane resistance means that more current needs to be passed via excitatory receptor channels to achieve the same amount of depolarisation that is necessary for reaching the firing threshold. This effect is inhibitory (shunting inhibition).

(c) In contrast to the interpretation of dopamine signals based on prediction errors discussed during lectures, another theory posits that dopamine signals are related to 'surprise'. According to this theory, dopaminergic cells increase their firing rates whenever something salient but unexpected happens. In this context, a stimulus is said to be salient if it is rewarding itself, or if the animal has associated it with reward.

(i) Explain how this alternative theory would account for the pattern of dopamine signals before and after training in a classical conditioning paradigm. [15%]

Answer: Before conditioning, both the CS and the US are unexpected, but the CS is not salient (because it is not associated with reward yet), while the US is salient, so the theory predicts increased dopaminergic activity at the time of the US but not at the time of the CS. After conditioning, the animal can predict the US based on the CS, which means that the arrival of the CS itself is still unexpected but it has become salient, while the US is still salient but not unexpected any more. Therefore, the theory predicts increased dopaminergic activity at the time of the CS but not at the time of the US. This is precisely the basic experimental observation: dopamine signals at the time of the US before training, and at the time of the CS after training.

(ii) Devise an experimental paradigm that can adjudicate between the surprise- and prediction error-based accounts of the dopamine response. Explain, with reasons, which theory is favoured by the available experimental data discussed during lectures. [15%]

Answer: Several different answers are possible. A main difference between the two theories is that prediction errors are signed quantities, whereas surprise is an unsigned quantity. Thus, for example,

one could simply use ‘catch’ trials after training in a classical conditioning paradigm to distinguish which quantity is encoded in dopamine signals. In these catch trials, the US is omitted. This omission of the US is salient and unexpected, therefore the surprise theory predicts increased dopaminergic activity. In contrast, omission of expected reward results in a negative prediction error, and thus the prediction-error theory predicts a decrease in dopaminergic activity. It is the latter that is empirically observed: dopaminergic activity drops below baseline at the time of the omitted US. Therefore, (at least this kind of) experimental data favour prediction error over surprise as the theoretical account of dopamine signals.

END OF PAPER