Version GH/3 – CRIB

EGT2
ENGINEERING TRIPOS PART IIA

Tuesday 6 May 2025 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
Supplementary page: one extra copy of Fig. 2 (Question 1)

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 neuron's membrane is approximated by a sphere of radius r. The neuron's specific membrane capacitance (capacitance per unit area of membrane) is estimated to be $c_\mathrm{m} = 1\,\mu\mathrm{F}\cdot\mathrm{cm}^{-2}$. A scientist proposes the following model for how the neuron's membrane potential, $V(t)$, responds to the injection of a small current, I_ext, measured in pA:

$$C\frac{dV}{dt} = -g(V - E) + I_\mathrm{ext} \tag{1}$$

where C, g and E are constants.

(a) What does C represent, and what properties of the membrane determine its value? [10%]

Answer: Total membrane capacitance. It is the product of the specific membrane capacitance and the total membrane area ($C = 4\pi r^2 c_\mathrm{m}$).

(b) What does g represent, and what properties of the membrane determine its value? [10%]

Answer: Total membrane conductance. It is the sum of all ionic conductances in the membrane, which depends on the number of ion channels, their individual conductances when open, and (for voltage-dependent channels) their probabilities of being open at the resting potential (see below).

(c) What does E represent, and what properties of the membrane determine its value? [10%]

Answer: Resting potential. It is a weighted average of the Nernst (reversal) potentials of the various ionic species responsible for the total membrane conductance, weighted by their respective conductances (at the resting potential! thus this defines E only implicitly via self-consistency).

(d) Explain why Eq. (1) is unlikely to be a good description of the neuron's response when I_ext is large. [10%]

Answer: For large enough I_ext, the neuron should eventually fire an action potential. Equation 1 is a passive model that clearly cannot do that: it's a first-order linear ODE, such that the response to a constant large current will always be an exponential ramp towards some new steady state.

(e) The scientist injects a current step of 1 pA and records the voltage response $V(t)$ shown in Fig. 1. Using the model of Eq. (1), estimate the neuron's radius r, and explain your derivation. [30%]

(cont.

Answer: Equation 1 predicts a voltage response of the form $V(t) = E + (1 - \exp^{-t/\tau})\Delta E$, with $\Delta E = I/g$ and $\tau = C/g$. From Fig. 1, one can estimate $\tau = 20$ ms and $\Delta E = 1$ mV. Thus, $C = \tau g = \tau I/\Delta E = 20$ pF. Finally, since we also have $C = 4\pi r^2 c_m$, we can estimate $r = \sqrt{\frac{C}{4\pi c_m}} \approx 12.6$ microns.

(f) The Hodgkin-Huxley equations, as presented in lectures, are used to model the response of another neuron to a step of hyperpolarising current (Fig. 2, left; the grey box marks the stimulation period). The voltage dependence of the model's three gate variables (sodium activation m, sodium inactivation h, and potassium activation n) is shown in Fig. 2 (right, showing voltage-dependent steady-states and time constants).

(i) On the additional copy of Fig. 2 provided at the end of this paper, complete the legend by writing m, h, or n in front of the relevant line style. Do not forget to hand in your completed copy of Fig. 2 with your answer to this question. [10%]

Answer: Solid: sodium inactivation h. Long-dashed: sodium activation m. Short-dashed: potassium activation n.

(ii) Somewhat paradoxically, this hyperpolarising stimulus ends up eliciting an action potential. Based on Fig. 2, describe two potential mechanisms by which this action potential might arise (no need to explain the shape of the action potential itself). Which of these two mechanisms is likely to be the largest contributor, and why? [20%]

Answer: This phenomenon is known as 'anode break excitation'. One potential mechanism relies on the behaviour of sodium channels. Prolonged membrane hyperpolarisation to ~ -80 mV causes the h-gate to becomes nearly fully open, and the m-gate to become fully closed. After the stimulus is switched off, the membrane potential goes back to resting potential in a few ms, and the m-gate rapidly recovers its (partially) opened resting state. At that point, the slow h-gate is still more open than it normally is at rest. Thus, there is more sodium flowing into the cell, which causes further membrane depolarisation – sufficient depolarisation to trigger an action potential through the standard mechanism. A second plausible mechanism relies on potassium channels. The potassium n-gate, too, closes during stimulus-induced hyperpolarisation, and it too recovers only 'slowly' (relative to sodium activation) post-stimulus. By the time the membrane potential has gone back to rest, the outward potassium current is less than it normally is at rest – thereby also promoting further depolarisation. Given that the recovery of the n-gate is faster than that of the h-gate, the first mechanism is likely to dominate the effect.

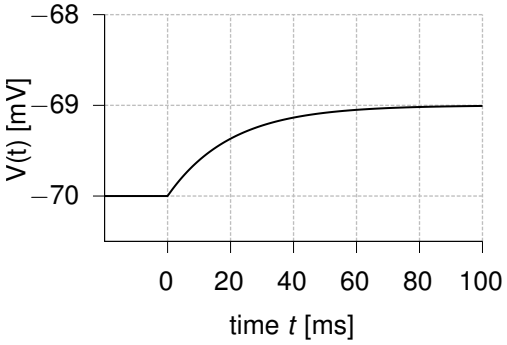

Fig. 1

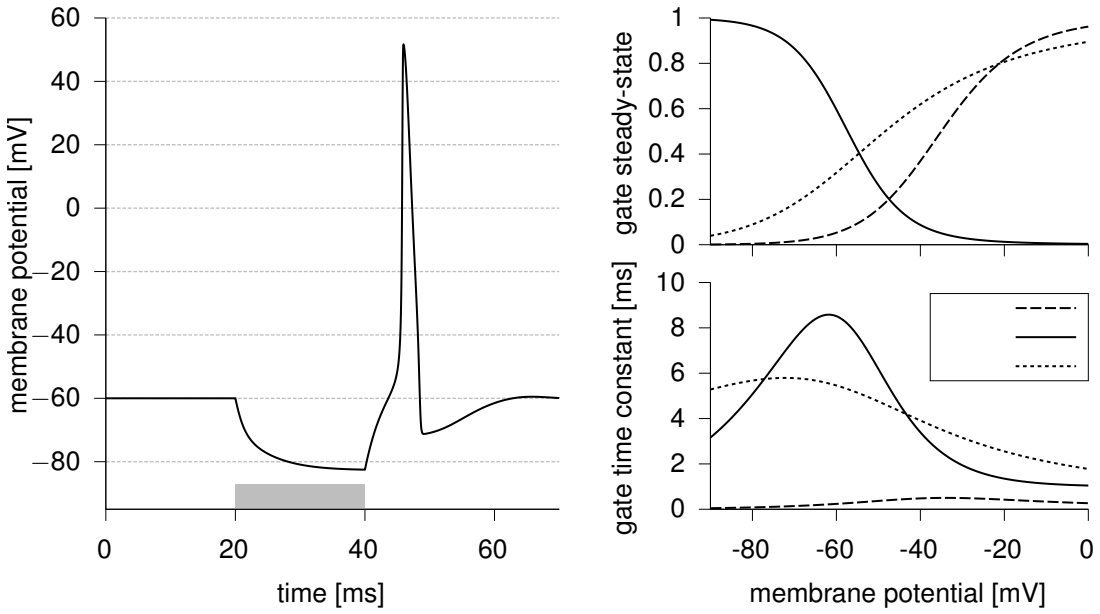

Fig. 2

 (cont.

2 In a multisensory integration experiment, a subject is asked to gaze at the center of a multimodal panel, which can deliver a brief visual and/or auditory stimulus at an arbitrary yaw location θ within the binocular visual field. Here, the yaw angle θ is measured in degrees relative to the center of the panel (with $\theta > 0$ on the right, by convention). Visual stimuli take the form of circular dot clouds ('blobs'), with the location of each dot drawn from an isotropic Gaussian distribution of variance α^2 (which the experimenter can control) around the desired location θ. Auditory stimuli are single clicks.

(a) Suppose that a visual stimulus is shown at angular position θ, with no sound. According to the statistical perception framework discussed in lectures, the visual sensory evidence can be summarised by a likelihood function $p(\text{'evidence from eyes'} \mid \theta)$. Up to a multiplicative factor, this function has the same shape as the probability density function of a Gaussian random variable with mean $\hat{\theta}_\text{v}$ ('maximum-likelihood estimate') and constant standard deviation σ_v.

(i) Does $\hat{\theta}_\text{v}$ equal θ? If yes, under what assumption? If no, why? [10%]

Answer: No, because for any given presentation of the visual stimulus at θ, multiple sources of noise in the visual system would corrupt the sensory evidence, and therefore cause the ML estimate $\hat{\theta}_\text{v}$ to randomly deviate from the true θ. Some marks can be awarded for answering "yes, on average across repeated presentations of the same stimulus".

(ii) Describe a 2AFC task that could be used to estimate the subject's σ_v without ever asking them to output precise localisation judgments. Explain what you would measure, how you would analyse the data to extract an estimate of σ_v, and what key assumption you need for this. [30%]

Answer: Instead of asking precise localisation reports, a 2AFC task can be used to estimate σ_v. In each trial, the subject is presented with two consecutive visual stimuli at two randomly chosen locations separated by a (signed) angular distance Δ. Subjects are asked to tell whether they perceived the second stimulus to be 'more to the right' compared to the first stimulus. Assuming that subjects make *comparative* judgments based on their ML estimates of the locations of the two consecutive stimuli, i.e. anwer 'yes' if $\hat{\theta}_\text{v}^{(2)} - \hat{\theta}_\text{v}^{(1)} > 0$, one can derive predictions for how subjects will respond. Specifically, for fixed Δ, $\hat{\theta}_\text{v}^{(2)} - \hat{\theta}_\text{v}^{(1)}$ will be a normal random variable with across-trial mean Δ and variance $2\sigma_\text{v}^2$. Thus, the psychometric function (proportion of trials where the subject respond 'yes', as a function of Δ) is expected to have the form $\int_0^\infty \mathcal{N}(\theta; \Delta, 2\sigma_\text{v}^2) d\theta$. A simple fit of this function to the subject's recorded behaviour will allow the extraction of σ_v.

(b) The subject is now presented with coincident visual *and* auditory stimuli, presented

simultaneously at the same location θ.

(i) Derive an expression for the maximum-likelihood estimate, $\hat{\theta}_{\mathrm{va}}$, of the stimulus location, and for the corresponding spread σ_{va}. Express your answers in terms of $\hat{\theta}_{\mathrm{v}}$, $\hat{\theta}_{\mathrm{a}}$, σ_{v} and σ_{a}, where $\hat{\theta}_{\mathrm{a}}$ and σ_{a}, defined analogously to their $(\cdot)_{\mathrm{v}}$ counterparts, are the parameters of the (Gaussian) likelihood function arising from the auditory stimulus. State and justify any key assumption(s) in your derivation. [20%]

Answer: Conditioned on a given θ_{va}, visual and auditory sensory responses can be reasonably assumed to be independent. This implies that the multimodal likelihood function is the product of the two single-modality likelihoods: $p(\text{'combined evidence'}|\theta) = p(\text{'evidence from eyes'}|\theta) \times p(\text{'evidence from ears'}|\theta) \propto \mathcal{N}(\theta; \hat{\theta}_{\mathrm{v}}, \sigma_{\mathrm{v}}^2) \times \mathcal{N}(\theta; \hat{\theta}_{\mathrm{a}}, \sigma_{\mathrm{a}}^2)$. Substituting the expression for the normal p.d.f. and completing the square gives an expression of the form $\propto \mathcal{N}(\theta, \hat{\theta}_{\mathrm{va}}, \sigma_{\mathrm{va}}^2)$, with $\hat{\theta}_{\mathrm{va}} = \hat{\theta}_{\mathrm{v}} + \frac{1}{1+\sigma_{\mathrm{a}}^2/\sigma_{\mathrm{v}}^2}(\hat{\theta}_{\mathrm{a}} - \hat{\theta}_{\mathrm{v}})$ and $\sigma_{\mathrm{va}}^2 = \frac{\sigma_{\mathrm{v}}^2\sigma_{\mathrm{a}}^2}{\sigma_{\mathrm{v}}^2+\sigma_{\mathrm{a}}^2}$.

(ii) Based on your answer to part (a)(ii), describe an experimental protocol and associated analysis that could be used to test the hypothesis that subjects combine vision and hearing optimally in their perception of stimulus locations? [30%]

Answer: One could start by measuring σ_{a} using the same 2AFC task as in part (a)(ii), now with purely auditory stimuli. Similarly, one could measure σ_{v} in the same 2AFC task, with purely visual stimuli, for various values of the blob width α (the wider these blobs, the less information they provide about the stimulus location, thus the larger σ_{v}). Finally, running the 2AFC task with combined modalities (vision+hearing) will yield estimates of σ_{va} for the various values of visual corruption (α) used, which can be compared to the predictions of part (b)(i) to test for optimal cue combination. This approach was used by Alais and Burr in their famous 2004 paper.

(c) A ventriloquist can trick their audience into thinking that their own vocalizations originate from a puppet whose mouth they move at the same time. Based on your answers to previous parts, give a possible explanation for this effect, and what it says about the relative precision of location encoding in the visual and auditory domains in humans. [10%]

Answer: A statistically optimal human who combines visual and auditory localisation cues according to the equation derived in (b)(i) should localise the ventriloquist's vocalisations somewhere between the visual (puppet's mouth) and auditory (ventriloquist's) sources, on average. Exactly where in between depends on the relative precision of visual and auditory perception. If the information (about location) provided by sounds is much weaker than that provided by vision (i.e. $\sigma_{\mathrm{a}}/\sigma_{\mathrm{v}} \gg 1$), then location perception will shift towards the puppet's mouth. That a ventriloquist can trick us suggests that this is indeed the case in humans.

3 (a) This question is about excitatory and inhibitory synaptic signalling.

(i) Describe the difference between an excitatory and an inhibitory synapse. [5%]

Answer: When activated, an excitatory synapse increases the propensity of the postsynaptic neuron to generate action potentials, an inhibitory synapse decreases it.

(ii) Describe the difference between an EPSP and an IPSP. [10%]

Answer: During an EPSP, the membrane potential becomes more depolarised (i.e. increases above the resting membrane potential), while during an IPSP, the membrane potential becomes more hyperpolarised (i.e. decreases below the resting membrane potential).

(iii) Explain with reasons if an excitatory synapse can generate an IPSP under physiological conditions. [5%]

Answer: No, it cannot. If the membrane potential is hyperpolarised below the resting membrane potential (IPSP), it will only decrease the propensity of the postsynaptic neuron to generate action potentials (inhibition). The shunting effect, characterising any (ionotropic) synaptic transmission (see below), just further adds to this inhibitory effect.

(iv) Explain with reasons if an inhibitory synapse can generate an EPSP under physiological conditions. [10%]

Answer: Yes, it can. Although depolarising the membrane potential above the resting membrane potential (EPSP) increases the propensity of the postsynaptic neuron to generate action potentials by itself (excitation), ionotropic synaptic transmission always also increases the total membrane conductance of the postsynaptic neuron. This increased total membrane conductance, or decreased membrane resistance, acts as a shunt for other (excitatory) synaptic currents and can thus decrease the propensity of the postsynaptic neuron to generate action potentials (inhibition). When the EPSP amplitude is small (below firing threshold), this shunting effect dominates, and so the synapse becomes net inhibitory.

(v) Explain with reasons if an excitatory neuron can lead to the inhibition of an other neuron. [10%]

Answer: Yes, it can, e.g. via a disynaptic pathway, such that it directly excites (synapses on) an inhibitory neuron, which in turn directly inhibits (synapses on) the target neuron.

(vi) Explain with reasons if an inhibitory neuron can lead to the excitation of an other neuron. [10%]

Answer: Yes, it can, e.g. via disynaptic disinhibition, such that it directly inhibits (synapses on) an inhibitory neuron, which in turn directly inhibits (synapses on) the target neuron.

(b) This question is about synaptic plasticity in the Aplysia gill withdrawal reflex. For each of the following pharmacological manipulations during training (but not testing), describe their effects on habituation, sensitisation, and conditioning, in terms of how the animal's gill withdrawal response (for any relevant stimulus) compares to those of naïve animals after training:

(i) Local application of lidocaine to the sensory neurons. [10%]

Answer: Lidocaine is a Na^+ channel blocker, so it prevents the sensory neurons from firing. Thus, habituation and classical conditioning are blocked as their induction requires the sensory neurons to generate action potentials, and in these paradigms responses during test will be as in naïve, untrained animals. For sensitisation, depending on the precise interpretation of the question two answers are acceptable.

1. In principle, "sensory neurons" in the question should be taken as referring to *all* sensory neurons, including those in the tail, transmitting information about the noxious stimulus (and not just those in the siphon, sensing the harmless stimulus). Under this interpretation, the correct answer for sensitisation is that it will also be blocked in the same way as habituation and classical conditioning are (described above).

2. However, it may also be acceptable to take "sensory neurons" to only refer to the sensory neurons *of the siphon* (as is explicitly suggested below, in question 3(b)(ii)). In this case, the correct answer is that the induction of sensitisation does not require the sensory neurons (in the siphon) to generate action potentials, and so it will be unaffected, such that responses to siphon stimulation in trained animals will be larger than those in naïve animals.

(ii) Local application of a potassium channel blocker to the sensory neurons (of the siphon). [20%]

Answer: Depending on the precise interpretation of the question, two answers are acceptable.

1. If only partial blocking of potassium channels is assumed, then the potassium channel blocker will make action potentials longer-lasting in the sensory neuron. However, this only happens during training, not testing. (Unlike e.g. in sensitisation or conditioning, when potassium current are decreased even during testing due to the long time constants in the 5-HT-receptor→Adenylyl cyclase→PKA pathway, i.e. the long-lasting effects of each of these receptors / enzymes.) Thus, there will be not much effect on any one of the paradigms. If anything, responses will be slightly diminished due to the depletion of presynaptic vesicles in the active zone, especially in habituation and conditioning, in which repetitive stimulation is used during training. This means that habituation (which itself is characterised by diminished responses) may be a little amplified, while conditioning (which is characterised by enhanced responses) may be a little attenuated. For conditioning this is only going to be the case for the CS+, which uses the siphon pathway (and for which responses are normally specifically enhanced), but not the CS- , which uses the mantle shelf pathway (and for which responses are not normally enhanced much), and so its specificity will also be diminished.

2. However, if perfect blocking of potassium channels is assumed, then only one action potential is generated by the sensory neuron, after which it enters depolarisation block and cannot generate any more action potentials (until the potassium channel blocker is washed out). Under this interpretation, habituation is largely blocked as it requires repeated action potential generation by the sensory neuron, classical conditioning is either blocked for the same reason, or appears as simple sensitisation (as if US alone was presented multiple times), while sensitisation is unaffected as it does not rely on action potential generation by the siphon sensory neuron (see answer to question 3(b)(i)).

(iii) Selective serotonin re-uptake inhibitors (SSRIs). [Hint: normally, serotonin molecules are actively taken up by the presynaptic cell from the synaptic cleft. This process is inhibited by SSRIs.] [20%]

Answer: Habituation is unaffected as it does not rely on serotonin signalling. So responses are diminished compared to naïve animals. Sensitisation is amplified, as SSRIs will cause more 5-HT molecules to be present in the synaptic cleft for longer, and so its postsynaptic modulatory effects on the primary (siphon→gill) pathway will be amplified and elongated. Conditioning is also amplified for the same reason. Therefore, responses to the CS+ will be larger than those in naïve animals. However, if the SSRI effect is sufficiently great, 5-HT molecules will still be present in the synaptic cleft and able to bind to postsynaptic receptors after the CS- arrives, which would also make responses to the CS- larger than in naïve animals, and thus reduce the specificity of conditioning (making it less associative, more like simple sensitisation).

4 (a) Explain what changes at the cellular level underlie the following three aspects of LTP as measured in extracellular recordings:

(i) increased amplitude of the population EPSP;

Answer: The amplitudes of EPSPs in individual cells (as could be measured intracellularly) have increased, and as they add up they create a larger extracellularly measured population EPSP. [10%]

(ii) increased amplitude of the population spike;

Answer: The amplitudes of EPSPs in individual cells (as could be measured intracellularly) have increased, and so more cells reach the firing threshold and fire action potentials in response to the stimulus. Also, because latencies have decreased (see next point), the action potentials occur more synchronously which further increases the amplitude of the population spike since. [20%]

(iii) decreased latency of the population spike.

Answer: The amplitudes of EPSPs in individual cells (as could be measured intracellularly) have increased, but this means that their initial slopes have also increased proportionally. Thus cells that for which the amplitude of EPSP exceeds the firing threshold will also reach the threshold earlier. [20%]

(b) This question is about dopaminergic cells and the temporal difference error.

(i) Describe three different experimental conditions in a simple classical conditioning experiment in which it can be shown that the activity of dopaminergic cells is consistent with a temporal difference error signal. [30%]

Answer: In a classical conditioning experiment:

A. before training, dopaminergic activity increases for some time after the presentation of the US, but not at the time of the CS;

B. after training, dopaminergic activity increases for some time after the presentation of the CS, but not at the time of the US when it is delivered;

C. after training, dopaminergic activity decreases for some time after the time when the US ought to be presented when in fact it was omitted despite the CS having been presented earlier in the trial.

(ii) Describe the activity of dopaminergic cells as a function of within-trial time in a secondary conditioning paradigm at each stage of training: before the first phase of training, after the first but before the second phase of training, and after the second phase of training. [20%]

Answer: Before the first phase of training, dopaminergic cells fire at the time of the US. After the first but before the second phase of training, dopaminergic cells fire at the time CS_1 (that was introduced during the first phase of training, and consistently precedes the US). After the second phase of training, dopaminergic cells fire at the time CS_2 (that was introduced during the second phase of

training, and consistently precedes the CS_1).

END OF PAPER

THIS PAGE IS BLANK

Version GH/3 – CRIB Candidate Number:

EGT2
ENGINEERING TRIPOS PART IIA
Tuesday 6 May 2025, Module 3G3 – CRIB, Question 1.

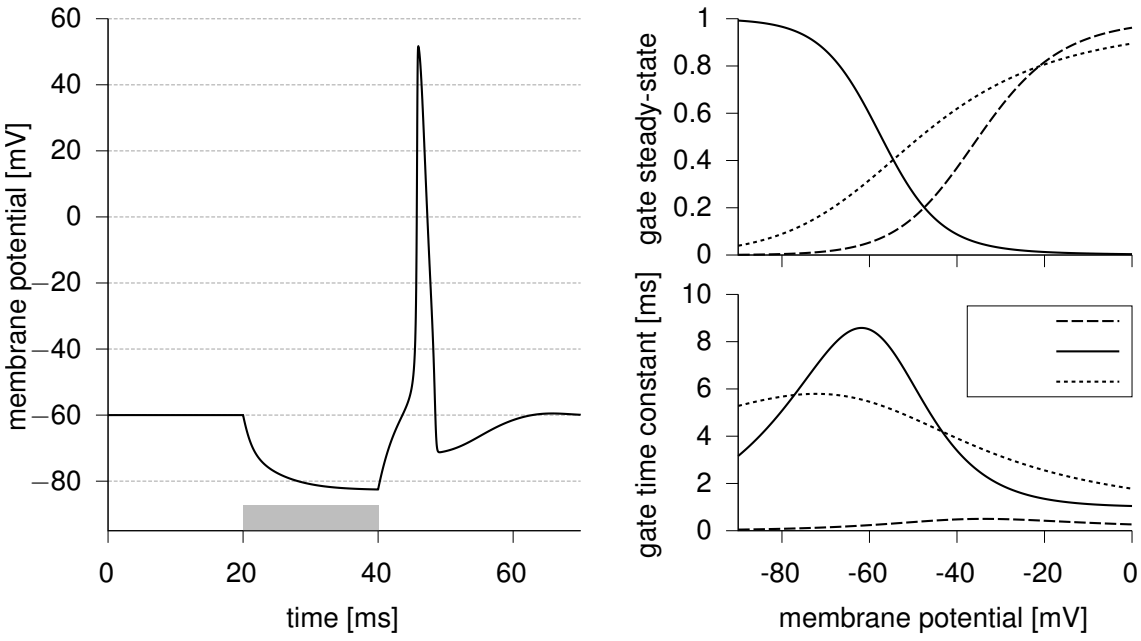

Additional copy of Fig. 2. It should be annotated as specified in Question 1, and handed in with your answers to that question.