

# Homework 4

## Systems & Theoretical Neuroscience [SWC & Gatsby]

Due: Monday Dec. 11th

### 1 Optimal control

For this question you will build a system that controls a two-joint robot arm, such as the one shown in the diagram below. The motor commands that you can send to the arm are  $\phi_1$  and  $\phi_2$ , the angle that each respective joint makes with the x axis. However, your control system is noisy, such that the final angle of each joint is given by  $\phi_a + \epsilon_a$ , where  $\epsilon_a$  is Gaussian with mean zero and variance 0.1 radians.

The goal of your control system is to move the joints such that the tip of your robot finger (the arrowhead in figure 1) is in the green reward region, which will win you a reward of 10 points. If instead you end up in the red penalty region, you will incur a penalty of -50 points. Your goal here is to infer which motor commands to send in order to optimise your expected reward.




Figure 1: The robot arm control problem. The targets' sizes are  $s$  by  $s$ , with  $s=0.5$ . The distance between the target and the base of the arm  $d = 1$ , the length of the lower arm piece  $l_1 = .75$  and the upper arm piece  $l_2 = 1.25$ . The angles  $\phi_1$  and  $\phi_2$  take values between  $[-\pi, \pi]$  radians.

- The problem described here is a typical problem in optimal control theory. Wolpert and Gharamani [WG00] discuss what optimal optimal control theory means for the nervous system, and how the brain uses internal models to achieve optimal control. Read the paper. What is a forward model for? And what are inverse models for? How are both useful in optimal motor control?
- Write down an equation for how the end position of the arm  $(x_2, y_2)$  (the arrowhead in the figure) depends on the control system's command  $\phi_1, \phi_2$  and the noise  $\epsilon_1, \epsilon_2$ .

- c) Write down an equation for how the reward depends on the end position  $(x_2, y_2)$ .

Together, the function for the movement endpoint and the reward function you have just written down constitute a forward model mapping from internal signals to expected arm position and returns. Simply put, a forward model takes a motor command as input and predicts the predicted bodily state as output. Forward models are useful in optimal control because they allow you to predict your expected sensory feedback while planning a motor action.

Now, let's implement our forward model in actual code. We have prepared a Jupyter Notebook with some plotting functions to help you get started.

- d) In the Jupyter Notebook, write a function `get_reward`, and a function `forward_model`, following the answers you gave in b) and c). Run the cell below. Does the robot arm behave as you expected? What happens if you make the variance higher?

As an important complement to forward prediction models, control systems make use of inverse models that tell us: given the expected returns, what is the optimal motor command to send? Inverse models use the desired position as input, and estimate the necessary motor commands to achieve this.

- e) Work out the expected reward for all combinations of  $\phi_1, \phi_2$  on the interval  $[-\pi, \pi]$  radians by sampling in Python. Plot the results as a heat map (with `plt.contourf()`). What is the optimal combination of angles, given your current noise levels? What happens if your variance gets larger?
- f) (BONUS) Find the analytical solution for the combination  $\phi_1, \phi_2$  that gives the highest expected reward.
- g) Repeat the last step with the penalty being -100. How did your optimal motor command change?
- h) How would you use such forward models for planning movements? Why would you want to use an inverse model instead? Would you still need a forward model if you had a correct inverse model?

## References

- [WG00] Daniel M Wolpert and Zoubin Ghahramani. "Computational principles of movement neuroscience". In: *Nature neuroscience* 3 (2000), pp. 1212–1217.

## 2 Forward models in electric fish

The mormyrid family of weakly electric fish is able to detect and capture prey in murky water through active electrical sensing of its environment. Using a specialized organ in its tail, the mormyrid generates an electrical field around it by emitting electrical organ discharges (EODs). Through electrical sensing receptors, the mormyrid is then able to detect changes in the surrounding electric field to infer the movement of objects in its surroundings.



Figure 2: an electric fish - *c. phantasticus*, aka the elephantnose fish

- What computation does the nervous system need to carry out in order to detect prey using *self-generated* electric fields? Suggest what inputs are required for this, and a mechanism for doing this.
- Does your suggested approach work while the fish is moving? If not, suggest some possible improvements to your mechanism to enable it to work during self-motion.

The electrosensory lobe (ELL) in the mormyrid electric fish brain has been identified as a key structure for electrolocation. A simplified circuit diagram is shown in figure 3.




Figure 3: A cartoon of the ELL circuit. Sensory input arrives to the green cell in the ELL via the orange axon. Other internal signals arrive to the blue and gray cells in the *eminentia granularis posterior* (EGp) via the dark blue fibers coming in from the top. These blue cells then provide input to the green ELL cell via its branching dendrites depicted. For the purposes of this homework, one can ignore the grey cells.

- Based on its function and anatomy, identify an analogous brain structure in mammals. Can you correspond particular cell types in the mormyrid ELL (i.e. different colored cells in figure 3) to those in the mammalian brain? (Hint: focus on the blue and green cells - what internal signals would you expect the dark blue fibers to be carrying?)

The further investigate the role of the ELL in electrolocation, you design an experiment to record ELL cell responses to stimuli in the environment. You fix your pet fish, Eric (Figure 2), in agar and

record from his ELL while manually stimulating a central nucleus that generates EOD commands, which result in EODs being emitted from the electric organ. After several trials, you introduce an electrosensory stimulus for 9 minutes. You then remove this stimulus and continue stimulating and recording for another set of trials. You obtain the raster plot in figure 4:

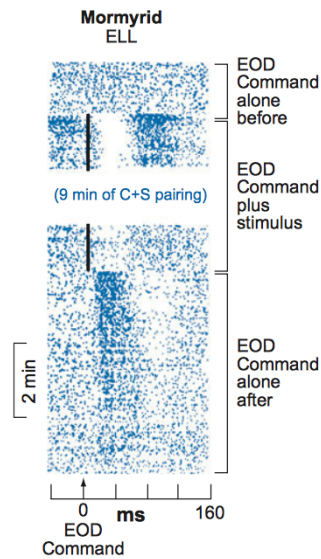


Figure 4: Each row corresponds to an individual trial, each blue dot to a spike. The EOD command signal occurs at 0 ms and the stimulus onset time is given by the black vertical bar.

- d) In the first set of trials when the EOD command is emitted but no stimulus is present, there is no response measured in ELL and the firing remains at baseline. Why does this make sense?
- e) Describe what happens throughout the EOD + stimulus (C + S) pairing and after the stimulus has been turned off.
- f) What does this data imply about the function of the electrosensory lobe and electric organ? What role does it seem to be playing in the computation you described in part (a)?
- g) How might these responses be achieved? Relate your proposed mechanism to the wiring diagram in figure 3. Which cell in figure 3 do you think we are recording from?
- h) Suggest a mechanism by which the cell's responses would change in this way over the course of the experiment. (Hint: think of a plasticity rule for the synapse from blue to green cells in figure 3 - what would Hebbian plasticity do?)

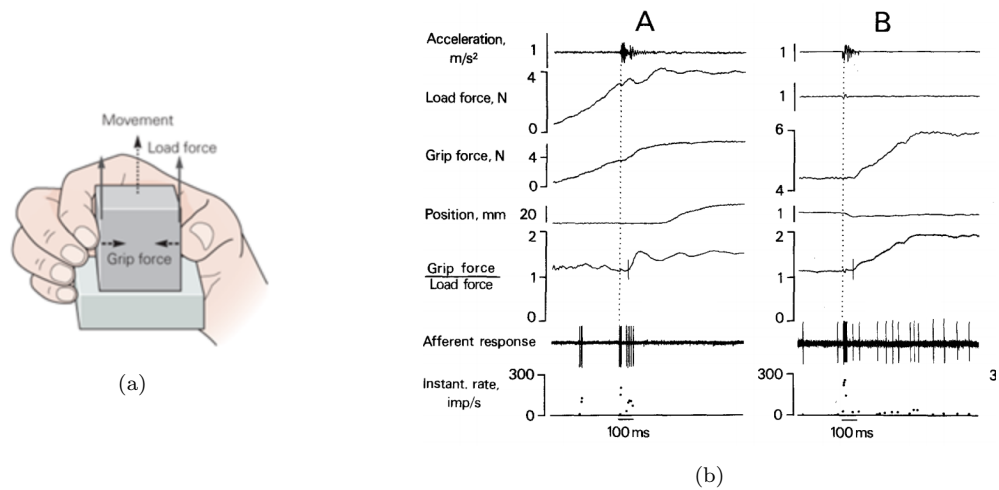


Figure 5: Data from two typical trials A and B. Top row shows the acceleration of the object during the lift, as measured by an accelerometer attached to it. The vertical dotted line indicates the onset of a slip event. Second and third rows show the load and grip force, as defined in figure 5a. Fourth row shows the position of the object, showing that the slip event in trial A occurred in the actions leading to a lift, whereas the slip event in trial B occurred during the hold. The fifth row shows the ratio of grip force to lift force, which is the critical control variable for avoiding an actual slip of the object. Last row shows recordings from a tactile mechanoreceptor cell afferent fiber.

### 3 Feedback control in human precision grip


Consider the experiment illustrated in figure 5a, in which a subject is asked to lift an object. Not shown is the machinery built into this innocent-looking cube so that the experimenter can precisely measure the grip force and load force produced by the subject, as well as the resulting acceleration of the object during the lift.

In the first such experiment, the subject is simply asked to lift the object and hold it for 5 seconds. The data from two typical trials are plotted in figure 5, along with simultaneous recordings from the afferent fiber of a tactile mechanoreceptor cell with its receptive field located under the region of skin in contact with the object.

Note the sudden oscillations in the acceleration of the object at certain times during the lift. Such “slip events“ are generally imperceptible by the subject.

- a) About how long does it take for the subject to respond to a slip event? Is this enough time for the brain to respond to a sensory stimulus? Speculate where the circuits mediating this response might be. Where are the likely targets of the tactile afferent fibers?
- b) Explain how and why the behavioral responses in trial A and B differ. (Hint: think about *how* the grip force ratio is increased in each case)

These data lead you to hypothesize that the tactile afferent fiber responses are driving the resulting motor behavior through some kind of reflex loop. You thus decide to test this hypothesis by directly stimulating the tactile receptors with an electrode to verify if they can directly drive a motor response. While the subject held the object in the air, you measured the grip force after



(a) Representative grip force profiles for a subject of each group. Vertical dotted line with black arrow on top indicates time of stimulation, second vertical dotted line indicates onset of force response.

(b) Representative slip events and corresponding force responses during voluntary slow release of the object.

Figure 6

electrical stimulation of the skin area in contact with the object. In one group of subjects, you tell them you will electrically stimulating their finger before doing so. In a second group, you stimulate their finger unexpectedly. The resulting grip force responses are plotted in figure 6a.

- c) How do the grip force responses differ in each group?
- d) Speculate on how such specific contextual modulation could be implemented within the short response time. Consider both the cases described in part (b) and part (c). What kind of information is needed by the circuit to produce these different responses?

All the observations until now suggest that these behavioral responses are the result of some kind of involuntary reflex. You thus decide to test whether it could interfere with a voluntary movement by having subjects voluntarily try to produce the opposite behavior. You instruct the subjects this time to lift the object and, during the hold, slowly release it. Naturally, as the hand gradually loosens its grip on the object, slip events start to occur. You find that these do in fact trigger force responses that interfere with the instructed movement (figure 6b).

- e) How do these force response profiles relate to the ones observed earlier? (Hint: are they more like the force response profiles in figure 5 and 6a bottom, or like those in figure 6a top?) What does this tell you about the circuit(s) driving these responses?
- f) Contrast this kind of control to that explored in question 1. When and how might one be more useful than the other? Are both necessary?