Homework 2
Systems & Theoretical Neuroscience [SWC and Gatsby]

Due: Wed, 1st November

1 Vision

1.1

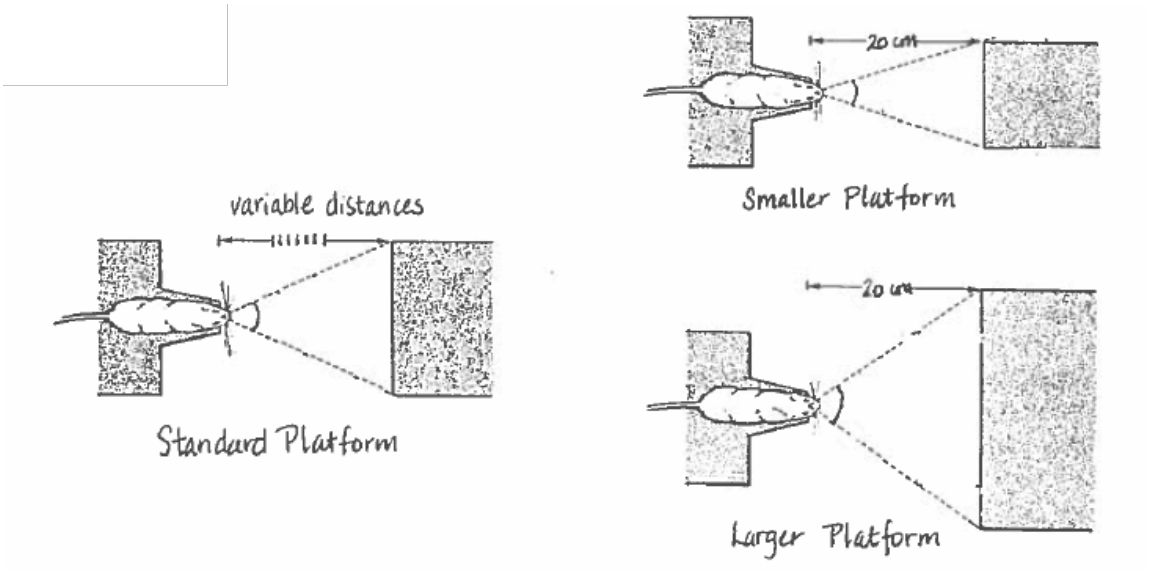

Figure 1

Your lab partner is interested in studying visuomotor cueing and has trained a fleet of 10 gerbils to jump from a take-off platform to a standard-size landing platform, over a range of distances, from 10 to 35 cm (see figure 1). He finds the gerbils to be quite adept and efficient at this task, and coudl typically gauge their jump to be exactly the right distance, not under-jumping (in which case they would fall into a tub of water), and not over-jumping (in which case they would land on a heating pad that was just a tad on the hot side) either.

After the training is complete, he conducts his experiment by testing the gerbils with the standard landing platform placed at various distances, interspersed at random with a landing platform that

Table 1: Experimental results from the gerbil jumping task.

platforms:	Standard						Small	Large
gaps (cm):	10	15	20	25	30	35	20	20
total trials:	10	10	10	10	10	10	10	10
over jumps:	1	0	1	0	0	0	10	0
under jumps:	0	0	0	0	1	1	0	10

is smaller than the standard one and one that is larger. His results are shown in table 1.

a) Judging from these results, what do you think is the relevant cue the gerbils are using to gauge their jumping distance?

Table 2: Results for the jumping task in gerbils with lesioned striate cortex.

platforms:	Standard						Small	Large
gaps (cm):	10	15	20	25	30	35	20	20
total trials:	10	10	10	10	10	10	10	10
over jumps:	0	1	0	0	0	0	10	0
under jumps:	0	0	0	0	0	1	0	10

Your lab partner is discussing his results enthusiastically with you, when, as a joke, you remark, 'I wonder if you could get blind gerbils to jump as well'. Your lab partner, in a sadistic fit, decides that this is a wonderful idea, and embarks upon a second experiment in which he removes the striate cortex from each of his gerbils, and after a recovery period, tests them again using the same paradigm as above. The results from his second experiment are shown in table 2.

b) The results in table 2 are quite surprising. Even though the 'blind' gerbils are a bit reluctant to jump, when persuaded to jump they seem to do just as well as before. Since their primary visual area is ablated, what other visual pathways are still intact, and might be subserving this jump calculation?

You have become very interested in these results and proceed to search the existing literature for similar experiments. You find a similar one involving rats, which are close cousins of the gerbils. In this experiment, one group of rats was trained to consider two landing platforms of the same size, one placed closer and the other placed farther away, with the one farther away baited with a favourite rat food. These rats quickly learned to always jump to the farther landing platform. When the striate cortices of these rats were ablated, however, none of the rats were able to perform this task anymore.

c) Assuming that both this reported result and your lab partner's result are significant and believable, what is a difference between the tasks in the two experiments that could explain the discrepancy between the results? What does this suggest to you about cortical vs subcortical visual functions?

d) Would you expect humans to show this type of blindsight behaviour, as exhibited by your lab partner's gerbils? If so, to what degree as compared to gerbils; same, more or less? And why?

1.2

Visual processing at the retinal level relies heavily on lateral interactions. In the outer plexiform layer, horizontal cells provide antagonistic inputs from adjacent regions of the visual field, allowing for border and contrast enhancement. At the inner plexiform level, the cells that mediate lateral interactions are the amacrine cells. The following study is an investigation into these second level lateral interactions of the retina. We record intracellulary from the retinal cells of a mudpuppy, *Necturus maculosus* (a real life animal, no joke) to find out.

First, the responses of horizontal cells and amacrine cells to a bar of light are recorded.

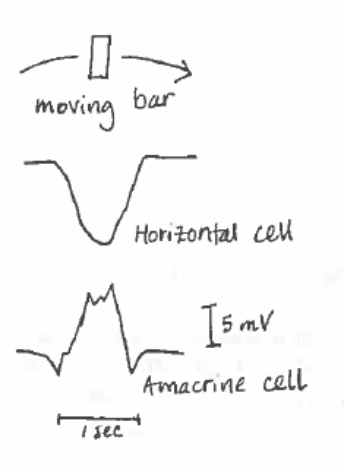

Figure 2

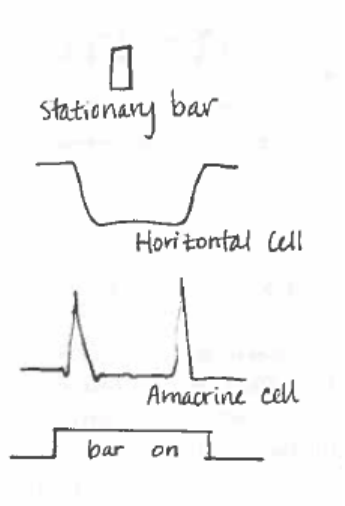

Figure 3

In figure 2, the bar is moved slowly across the receptive field of each cell. In figure 3, a stationary bar is presented in the receptive field centre for 1 second.

a) Based on these data, describe qualitatively the important differences in receptive field properties for these two cells.

Now we get fancy and use a more complicated stimulus configuration to investigate the effects of stationary vs moving peripheral stimuli on a central cell's response to a central stimulus. The set-up includes a four-vaned windmill pattern, each vane equivalent to the bar of light used above, that can be stationary or spinning, along with a central spot that can be flashed on and off (see figure 4). Recordings are made under 3 conditions from cells with receptive fields positioned directly beneath the central spot: central spot flash alone, central spot flash with stationary windmill and central spot flash with spinning windmill.

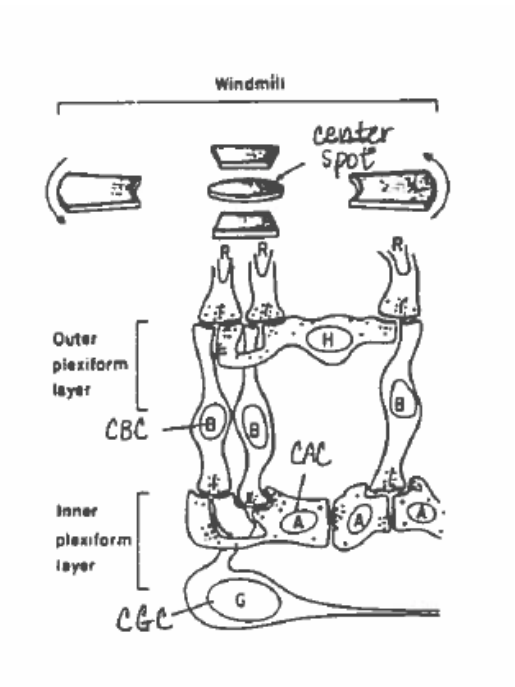

Figure 4

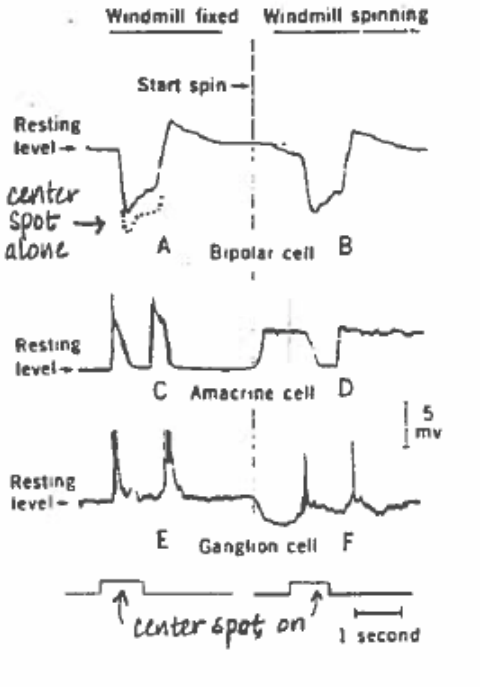

Figure 5

The central bipolar cell's response to the flash alone is indicated by the dotted curve in A of figure 5. In the presence of a stationary windmill pattern, the response to the central spot is somewhat reduced, as indicated by the solid curve in A. When the windmill pattern is rotating instead of stationary (figure 5B), the observed effect is no different from that of the stationary windmill.

b) To what cell type would you attribute the change from the dotted curve to the solid curve in A? Why do you think there is essentially no difference between the solid curves in A and B?

In trace C, the recording from the amacrine cell under the central spot (CAC), there is no significant difference between its response to the central flash with windill stationary from its response to the central flash alone (dotted and solid traces are superimposed). However, when the windmill starts spinning, the response becomes more interesting: the cell is depolarised to a new sustained level and the response to the centre spot is reversed in sign!

c) Which cells are most directly responsible for the depolarising effect of the spinning windmill on the central amacrine cell, and why is this depolarising effect sustained instead of transient?

In traces E and F, you now record from a ganglion cell, which happens to be an ON/OFF type, firing both at the onset and the offset of the central spot. Such cells are much more common in lower vertebrates than in the mammalian brain.

d) In this ganglion cell (CGC), when the windmill started spinning, there is a sustained hyperpolarisation, and the response to the central flash is reduced. How might you explain these effects

e) (Just for fun) Considering that the diet of the mudpuppy consists mainly of small bugs that are similar to the size of the central spot in this set-up, what do you think would be the real life, day-to-day significance of having cells with these particular connections and response characteristics?

2 Audition

2.1

Jamie, your pet adult barn owl, has been trained since early adulthood to sit quietly in a darkened, anechoic (echo-free) room and to orient his head accurately in response to bursts of sound. In response to the stimuli, which come from a speaker which can be positioned at any azimuth or elevation, he turns his head to face the speaker (remember, since it's dark, he can't actually see the speaker). Using an appropriate laser and head mounted mirror, you can accurately monitor the direction Jamie is facing. (Note that owls cannot move their eyes).

To test the role of abnormal binaural cues and their effect on sound localisation, you insert a dense foam plug into Jamie's right ear. Now Jamie, who used to localise almost perfectly in the dark makes consistent localisation errors.

a) What sort of errors would you expect Jamie to make? Explain why the foam plug would cause these sorts of errors.

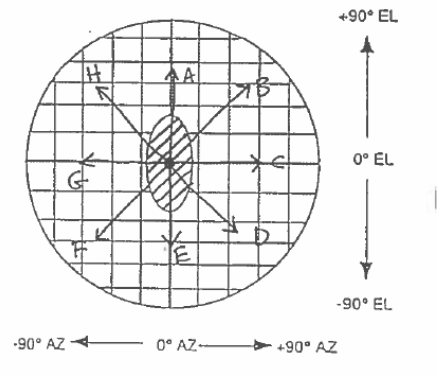

Figure 6

The owl's inferior colliculus (IC), a midbrain auditory structure contains space-specific neurons. That is, under normal, non-plugged conditions, a neuron in this region will respond to auditory stimuli from a restricted spatial location, as shown in figure 6.

b) Keeping in mind your answer for part (a), choose the change in response you would expect to observe while recording from such a neuron with the ear-plug in place (one of arrows A-H).

In order to detect where sounds come from, barn owls use the *interaural time difference* (ITD): the difference between the arrival times of a sound in the left and right ear.

c) Explain how the ITD could be a good measure for localising the origin of a sound, and draw a neural circuit that could implement this system.

Another measure that Jamie uses for localising sounds is the *interaural intensity difference* (IID): the difference in sound amplitude between his two ears.

d) Explain how the IID could be used for localisation, and draw a neural circuit that could implement this system.

e) Jamie, like all barn owls, has differently sized and placed ears: his left ear is in fact pointing slightly down, while his right ear is pointing up. Explain why this might be evolutionarily advantageous for an owl, in terms of sound localisation.

3 Olfaction

3.1

a) Give a brief (bullet point) comparison of visual and olfactory systems from sensory transduction to perception.

Your PI is curious to know how mitral cells in the olfactory bulb respond to a range of molecules that are structurally very similar to each other. To investigate this, you design an experiment where you present mice with a set of structurally related molecules while recording from single units in a single glomerulus. Figure 7 shows the spike trains you obtain. The black trace at the bottom indicates the respiratory cycle of the animal, and the black bar indicates the moment of odour exposure.

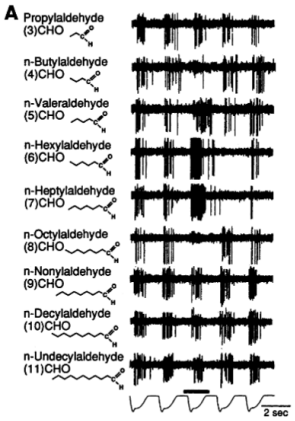

Figure 7: Single unit recordings of a mitral cell to a range of molecules. The black trace at the bottom shows the (artificial) respiration cycle.

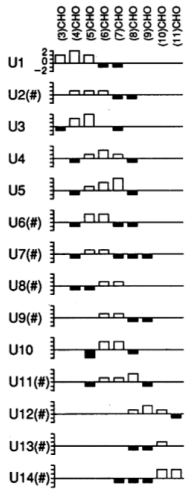

Figure 8: Firing rates in response to different molecules (horizontal axis) in different glomeruli (vertical axis)

b) Describe what you see in the spike trains. Do you see a general pattern, and how do the responses to different (but related) molecules differ?

You then repeat this experiment for a range of glomeruli, and compute the firing rates during the odour exposure. The results are plotted in figure 8. On the basis of figure 8, the researchers performing this study concluded that the olfactory bulb might be performing a kind of *contrast enhancement*.

c) Explain how the data shown in figure 8 indicate a contrast enhancement effect.

d) Draw a neural circuit that could achieve contrast enhancement. What is the most important feature of your circuit?

In a next experiment, you want to know how the concentration of an odour affects mitral cell responses in the olfactory bulb. To achieve this, you load the olfactory epithelium with a calcium sensitive dye in a mouse that expresses a different dye in mitral cells. This gives you a readout of both the input (ORNs containing Fura-2) and output (mitral cells expressing ArcLight) of the olfactory bulb (figure 9a). You expose your mouse to an odour at a variety of concentrations and get the data displayed in figure 9c.

e) Describe and explain the observed effect, what it might be good for, and what further experiments you could do to better understand the circuitry involved in implementing it.

   For additional (but highly recommended) reading, see: `https://www.ncbi.nlm.nih.gov/pubmed/20060600`

f) Design an experiment that would allow you to determine how the olfactory cortex performs template matching in order to identify/recognise familiar and non-familiar odorants.

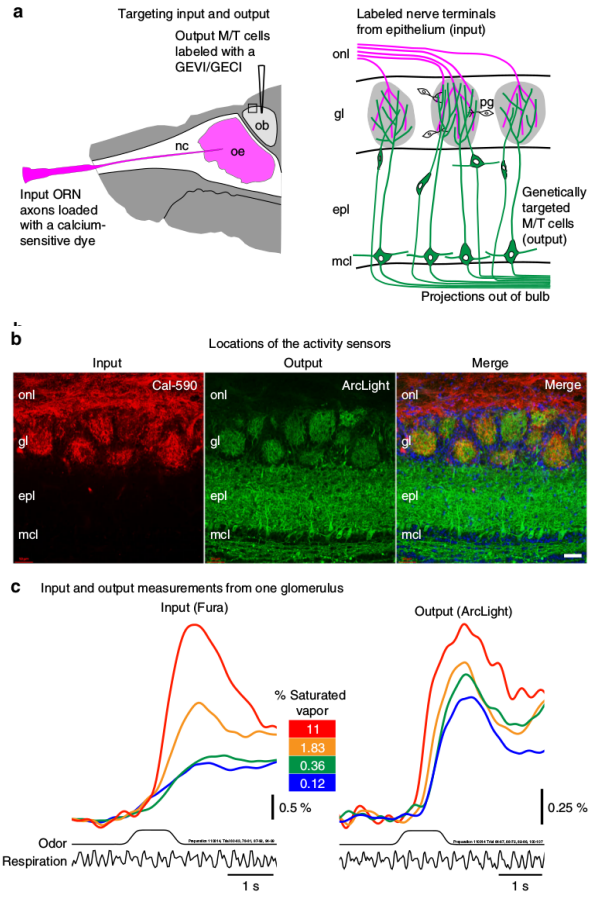

Figure 9: Responses of olfactory receptor neurons (left) and mitral cells (right)

3.2

Odors in the world are made up of different concentrations of different airborne molecules. In this question we model odors as N-dimensional, real valued variables $\mathbf{x} \in \mathbb{R}^N$. Each component x_i is a single odorant molecule. For instance, x_1 might be octanol, x_2 might be benzaldehyde, and so on. An 'odor' such as 'banana' is a particular combination of odorants, i.e. a particular value of the vector $\mathbf{x}$.

In the olfactory bulb of the mouse, olfactory receptor axons and mitral cell dendrites come together in spherical bundles called *glomeruli*. Each olfactory receptor neuron expresses a single type of olfactory receptor, and (miraculously) these converge by receptor type onto the glomeruli. The activation of a glomerulus reflects the total activation of all the ORNs that express its corresponding receptor. Each ORN responds in varying amounts to a selection of possible odorant molecules.

Mice have around 1800 glomeruli, but there are many more possible odors in the world. The goal of the olfactory bulb is to infer the true odors in the world based on these lower dimensional observations.

We model the glomeruli responses $\mathbf{y} \in \mathbb{R}^M$ (where $M << N$) as a linear combination of odorants, i.e. $\mathbf{y} = \mathbf{A}\mathbf{x} + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_y^2)$.

a) What is the size/shape of this $\mathbf{A}$ matrix? How can you interpret the rows or columns?

b) Write down the likelihood $p(\mathbf{y} \mid \mathbf{x})$.

We assume the neural circuit is computing the *maximum a posteriori* estimate of $\mathbf{x}$, i.e. maximising the log of the posterior $p(\mathbf{x} \mid \mathbf{y})$ with respect to $\mathbf{x}$. We assume the prior probability of odorants is a zero-mean gaussian prior distribution with independent and constant variance in each dimension:

$$p(\mathbf{x}) = \frac{1}{Z} \exp\left\{ -\frac{\mathbf{x}^T \mathbf{x}}{2\sigma_x^2} \right\} \tag{1}$$

where Z is a normalising constant and σ_x is a scalar.

c) Write down an expression for the log posterior.

One way a neural circuit can solve an optimisation is to set up time dynamics that perform gradient ascent, i.e. if we're maximising $L(\mathbf{x})$ with respect to $\mathbf{x}$, we update our estimate of $\mathbf{x}$ at every time step with:

$$\mathbf{x}_t = \mathbf{x}_{t-1} + \frac{1}{\tau_2} \frac{\partial L}{\partial \mathbf{x}} \tag{2}$$

The continuous time dynamical system that corresponds to this is given by:

$$\tau \frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = \frac{\partial L}{\partial \mathbf{x}} \tag{3}$$

where τ_1, τ_2 are time constants that control the speed at which we update $\mathbf{x}$.

Typically a dynamics equation for neurons $\mathbf{x}$ will take the form:

$$\tau_1 \frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = -\mathbf{x} + f(\mathbf{x}, \mathbf{y}, \text{any other neurons}) \tag{4}$$

where $-\mathbf{x}$ is a leak term that drives the neuron to zero in the absence of inputs, and $f(...)$ is some function of inhibitory and/or excitatory input from itself and other neurons in the network.

d) Compute a dynamics equation in the form of Equation 4 which performs gradient ascent on your log posterior in given in c).

e) Many parts of the brain are believed to perform predictive coding - that is, use a forward model to make a prediction of some observation and compare it to the incoming data. We can reformulate the MAP optimisation in this way by substituting an auxiliary variable $\mathbf{r} = \mathbf{y} - \mathbf{A}\mathbf{x}$ into our posterior and applying a Lagrange multiplier to enforce the constraint $\mathbf{r} = \mathbf{y} - \mathbf{A}\mathbf{x}$.

Note: If you are unable to complete this part, please get the solution from a colleague and use it to complete the rest of the question.

i) Write down the full Lagrangian for this problem, which is a function of $\mathbf{x}$, $\mathbf{r}$ and $\boldsymbol{\lambda}$ (your lagrange multiplier).

ii) Differentiate with respect to $\mathbf{r}$ and set to zero.

iii) Substitute your result from e)ii) into your Lagrangian to get a new Lagrangian which is a function of $\mathbf{x}$ and $\boldsymbol{\lambda}$ only.

iv) Differentiate to find update equations in both $\mathbf{x}$ and $\boldsymbol{\lambda}$. They should both be in the form given in Equation 4. *Hint: think carefully about the direction in which you should be updating* $\boldsymbol{\lambda}$.

f) In parts d) and e), you computed two different sets of dynamics which achieve the same goal. In Figure 10 we show the two circuits implied by each solution. Arrows represent excitatory connections and circles are inhibitory connections. Label the connections with their corresponding weight matrices according to your update equations. Can you make plausible suggestions as to type of olfactory cell each layer of neurons corresponds to? Give reasons for your choice.

Figure 10: Two possible olfactory circuits.
Model from part d) Model from part e)

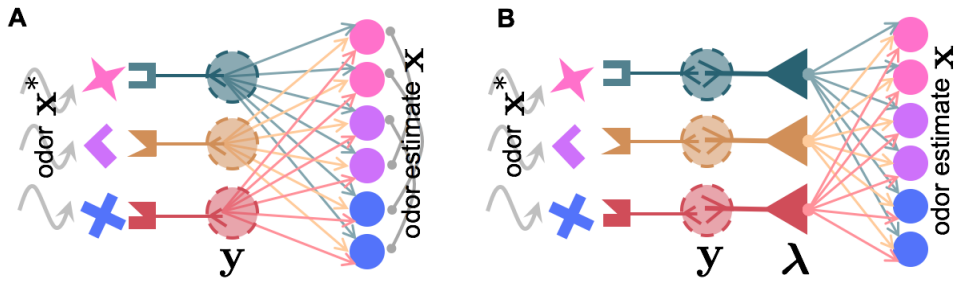

g) You would like to determine which of the above circuits is closest anatomically to that observed in the mouse olfactory bulb. How would you determine this experimentally? Do you think either of the proposed circuits is correct? If so, which one? If not, give reasons for why they are a bad fit.

h) We used a zero-mean gaussian prior $p(\mathbf{x})$ to describe the prior likelihood of different odorants in the world. Do you think this was reasonable? Can you suggest additional constraints that would be make it more realistic? How would you learn the parameters describing your prior?