

Student number:

First two letters of last name:

Mark:

Welcome to the final assignment! Building on the previous assignments which explored representation and computation, this assignment focuses on the neural basis of spatial navigation, and its intimate relationship to memory. To do so, we will explore an **actual data set** from one of the most studied brain structures in neuroscience: the hippocampus. As you'll recall, this is the structure where long-term potentiation (LTP) was first discovered (Bliss and Lømo, 1973), the structure that famous patient H.M. had removed to cure his epilepsy – only to find he had lost the ability to form new episodic memories or find his way around – and the structure where O'Keefe and Dostrovsky (1971) discovered place cells.

This assignment will **not** use Nengo¹. Instead, the data set is stored in the form of two video files. Two ways of accessing the videos are available:

- YouTube: [video 1](#) and [video 2](#) (also linked from LEARN). Remember that you can hit the fullscreen button and select HD (720p) quality for best results. In addition, you have the option of using YouTube's [HTML5 trial site](#), which gives you a playback speed control (but only if you use certain browsers).
- Downloading the video files to your computer from LEARN, so you can use a video player of your choice. If you do this, please note that the video was encoded² with XVID, so if you have an older video player you may need to download the [codec](#).

The videos have sound, but the ability to hear this is not required to complete the assignment. It may help you understand what is going on more easily, however. (Plus, for what it's worth, I think it's cool.)

The submission instructions are the same as for the previous assignment. Just to be clear, they are:

- Enter your answers into the form fields in this file, remembering to save your work. You can ignore messages about "Data collection workflow" or "Submitting your form".
- Make sure the file you submit actually has your answers!
- Only PDF submissions to the designated LEARN dropbox can be accepted. Please name your submission 12345678.pdf, replacing 12345678 with your student number.
- Explain your reasoning, but do not exceed the space available in the form fields *without using the scroll-bar*. I will not read whatever can only be seen by scrolling. Filling the entire box is not required to obtain full credit.
- Late submissions will incur a 30 percentage point penalty per 24 hours.
- The number of points (out of 100) is indicated before each question. As before, the number of points should not be taken as indication of how many distinct steps there are to the answer or how long it should be.
- I regret that your submission cannot be returned to you; however, the concepts from the assignment will be discussed in class the week following the submission deadline.

¹A Nengo simulation of place cells is available [here](#) if you are interested.

²If this word made you think of neuroscience, give yourself a mental bonus point.

Getting to know the data

Open the **first video** using your preferred method. It will probably play automatically, but pause and rewind it for now. Remember to adjust the quality and size settings so you have a clear, large view.

You should see a plethora of cryptic panels, as in the figure below:

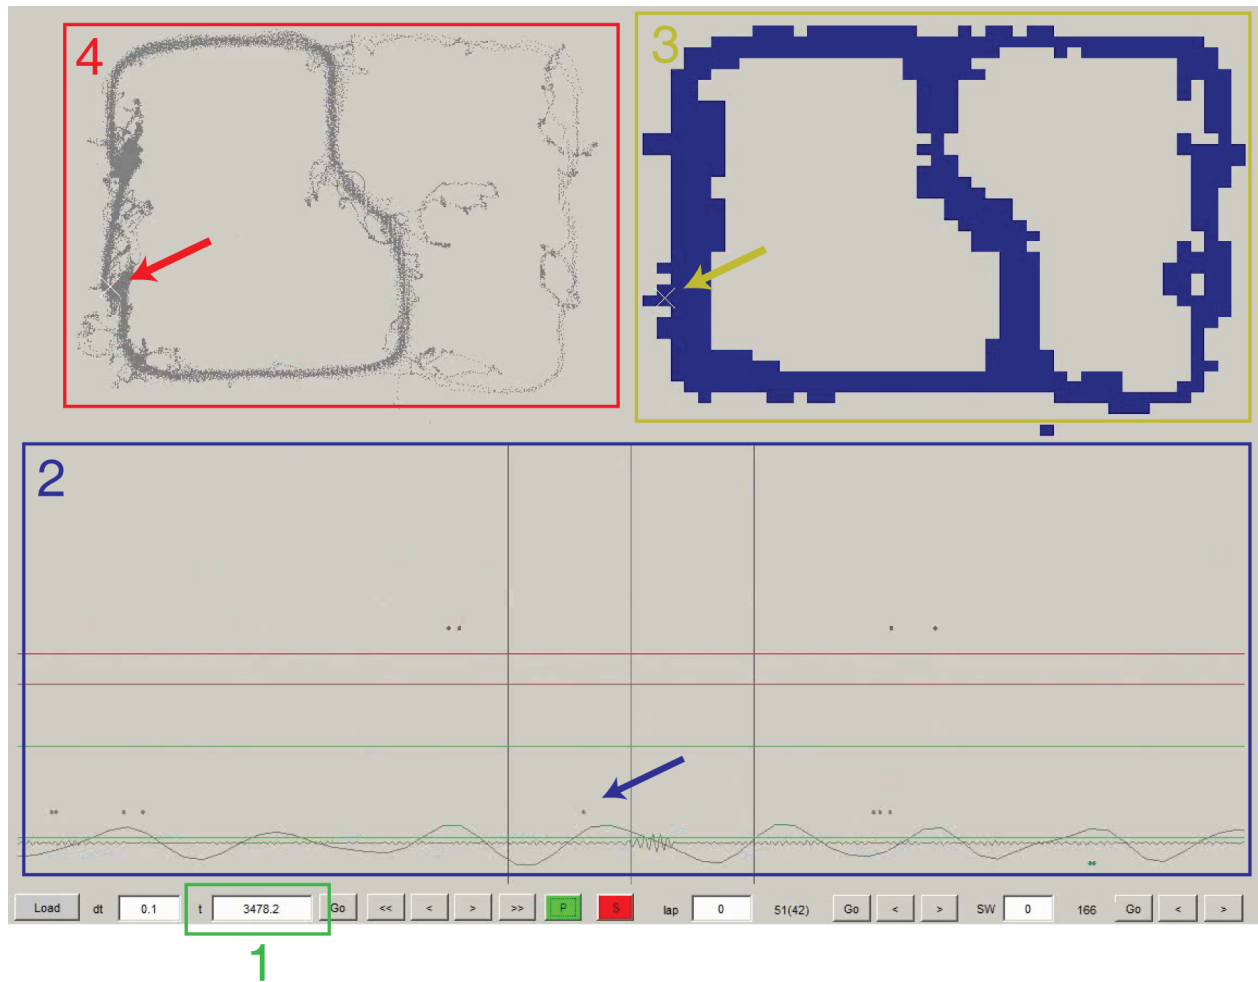


Figure 1: Annotated screenshot of the first frame of the first video. See the text for an explanation of the different panels and features highlighted here. If your screen does not look like this (minus the boxes and arrows), make sure you loaded the first, not the second, video.

Box 1 (green, bottom left) tells you the current experiment time, in seconds. **Box 2** (blue) is a rasterplot, showing spikes recorded from approximately 50 hippocampal place cells as a rat runs a spatial navigation task. As usual, each row corresponds to one neuron. Here, different neurons are assigned different colours to help you distinguish them³. The activity of these neurons is decoded

³Some additional information you won't need is also shown here, including local field potentials at the bottom of the plot (one filtered in the theta frequency band, 6-10Hz, and another in the ripple band, 150-200Hz), and some horizontal

to generate the position estimate in **Box 3**, using the method explained in lecture notes 10.1 (dark blue indicates low probability, red indicates high probability).

Box 4 shows you a number of things: first, the current location of the rat on the maze. This is indicated by the faint white x ; make sure you can find it (the red arrow points to it; it will be easier to see once things start moving). In this experiment, the rat will be running around on a track or “maze”: his path over the whole session is indicated in gray. At this time, the rat is on the bottom left section of the track.

The path is actually made up of many tiny dots, because every 1/30th of a second, a video camera takes a sample of the rat’s current location. Notice how there seem to be many “left” laps but only few “right” laps; the gray path on the right is much thinner than on the left. This is because in this particular session, only “left” laps were rewarded! There are actually two reward sites, which appear as solid gray blobs on the left side. This occurs because the rat tends to spend a lot of time there. As you will see, the rat only runs the maze in one direction (counterclockwise for left laps, and clockwise for right laps).

Box 4 also shows you the preferred firing locations (i.e. the location of the place field(s)) of those cells that are currently active. Right now, there is only one cell firing one spike in the current time window. This current time window (200 ms) is indicated by the vertical lines on the rasterplot; the arrow points to the one spike that is currently in it. This neuron has a preferred firing field in the middle of the maze – if you look carefully, you can spot a faint circle to indicate it. One such circle will be drawn for each spike in the rasterplot current time window.

If you have sound enabled, make sure it is **not** set to maximum volume! Then hit Play.

Over about 12 seconds, the rat will run a lap. You can see the white x , indicating its location, moving around the maze in both **Box 4** and the decoded location plot. If you have sound enabled, you will also hear a noisy progression from a low rumble to some intermittent high notes⁴. You should also see many spikes fly by in the rasterplot. The video runs in real time (you can verify this by looking at the experiment time, **Box 1**).

Watch the experiment for a few more laps. 3 should be enough. Pause the video.

Question 1 (10 marks):

score

Above, I claimed that the neurons in the rasterplot were “place cells”. Based on the data provided in the video, how can you verify this claim? Pick a single neuron in the rasterplot and describe what you would look for, in the video, to convince yourself or someone else that it is probably a place cell⁵.

lines of no importance.

⁴Each neuron is assigned a key on a piano keyboard according to where on the maze it has its place field. Cells with place fields in the lower left have low tones; higher and higher tones are assigned going counterclockwise along the maze. I did this to make it easier to perceive patterns in the data; the real hippocampus is not believed to implement such audio features.

⁵Note: of course, you can imagine more definitive experiments to test this, such as turning the lights off, or rotating some cues in the room. For this question, however, you have to work with the data you are given.

Question 2 (10 marks):

score

Recall that in Assignment 1, which row a neuron appeared, in the rasterplot, was unrelated to its preferred stimulus. In this rasterplot, do you find any relationship between which row a neuron appears on, and its preferred stimulus (location)? Explain your reasoning.

Rewind the video to the beginning, again, and play. As the rat runs, now look at the decoding plot in the top right of the screen.

Recall that to make this type of plot, the spiking activity from all the neurons is combined to estimate the location of the rat, based on the neural activity alone.

Question 3 (10 marks):

score

What about this decoded plot tells you that the recorded neurons contain information about the rat's location?

You probably noticed that the decoded estimate does not always match the actual location of the animal perfectly. Examples of this often occur when the animal is resting at the reward sites.

Question 4 (10 marks):

score

Look at the rest period that occurs between 33 and 37 seconds in the video. Why does the decoding not match the animal's location?

An example of a frequent such “error” is that when the rat reaches the first turn on the central stem of the maze, the neural activity often decodes to the second turn (the choice point at the top of the maze). An example of this occurs at experiment time 3529.2 (in **Box 1**).

Question 5 (10 marks):

score

What does such a systematic error tell you about the tuning curves of the relevant neurons? (Recall that the decoding method uses the tuning curves, and the spikes in a given time window, to estimate the rat’s position.)

As the animal runs, it does so at a certain speed that can be estimated from the change in location of the white x in a given unit of time. Likewise, we can estimate the speed at which the *neural representation* moves.

Question 6 (10 marks):

score

Explain how you can estimate approximately, at a glance, the speed of the neural representation of location from the rasterplot. To be clear, you may want to contrast what the rasterplot would look like for a slowly moving representation, and for a rapidly moving one.

The above phenomena illustrate that hippocampal neurons, as a population, can provide the animal with a representation of its location on the maze. However, there are also systematic deviations from this representation. Are such deviations just random noise? Or is there something deeper going on? We will explore this next, using the **second video**.

Pull up the second video on your screen. Watch it for a while, just past experiment time 4771.2 (30 seconds into the video is enough).

Question 7 (10 marks):

score

Inspect the event that occurred at experiment time 4771.2. Describe what happened, choosing your words carefully for informativeness and precision.

Question 8 (10 marks):

score

How does the speed at which the neural representation moves in the previous question compare to the update speed during normal running?

Question 9 (10 marks):

score

Relate your observation(s) in Questions 7 and/or 8 to a real-life phenomenon you have experienced. Make sure to explain what specific aspect or property of your previous observation(s) you are using, and how it maps onto your real-life example.

Now inspect the event at experiment time 4910.7.

Question 10 (10 marks):

score

You'll notice that fewer neurons participate in this event than in the previous one, as well as some other small differences. There is however a striking difference that does not seem compatible with the term "replay". What is it? Why is this property incompatible with a simple concept of replay? Assume that the rat runs in one direction only.

There are many other interesting phenomena that can be explored with these data. In case you have noticed something of interest, I have provided a bonus opportunity below. If I like your answer, you'll get one bonus mark. If your insight is so novel no one has noticed it before, you are welcome to analyze it further in the van der Meer lab (see the note below)!

Bonus question (1 mark):

score

Identify and describe a phenomenon or event of interest, in either of the videos, not yet discussed in this assignment.

Postscript: None of the information below is relevant to completing the assignment, but in case you are curious: I recorded the data used to make these videos when I was a post-doc in [A. David Redish's lab](#) at the University of Minnesota. I used 8 "tetrodes" (32 channels) to record from the dorsal CA1 pyramidal cell layer in the hippocampus. These data feature in a number of publications, including [van der Meer & Redish \(2011a\)](#) *Journal of Neuroscience*, [van der Meer et al. \(2010\)](#) *Neuron*, and [Gupta et al. \(2010\)](#) *Neuron*. In the [van der Meer lab](#) at Waterloo, we currently have a 96-channel recording system set up to record neurons from 24 tetrodes simultaneously. I thank A. David Redish for his permission to use these data for this course. The videos were created using custom MATLAB scripts and the freely available VirtualDub video editing software.