CNS Assignment: Feedforward and Recurrent Mechanisms for Orientation Selectivity in Visual Cortex

November 2025

Submission Details

The deadline for this assignment is 21st November at 12pm (standard late penalties apply). The assignment should be submitted via the course Learn page in pdf format, using the naming convention B0123456.pdf (i.e., your exam number). You should include the code you used to produce any figures in the assignment as an appendix (the code will not be marked, but may be checked where results are wrong and/or for plagiarism detection purposes, do not include any code in your main report). Your submission should be no longer than **6 pages** (excluding references, front matter, and appended code). To ensure fair and consistent marking, any content that extends beyond the 6 page limit will not be awarded any marks.

Plagiarism and Academic Misconduct

You may discuss this assignment with others, but direct copying of code, text, or mathematical results is not allowed. Any text that is directly copied from other sources should be placed in quotation marks with the source cited at the end of the sentence. Text which is written in your own words but based on material from other sources also requires a citation.

Use of AI tools is not permitted for this assignment.

Please remember the good scholarly practice requirements of the University regarding work for credit. You can find guidance at the School page https://informatics.ed.ac.uk/taught-students/all-students/your-studies/academic-misconduct. This also has links to the relevant University pages.

Marks for Style of Presentation (10 marks)

You should present your findings in the style of a scientific report. By this we mean that the document should include text broken into appropriate sections (e.g., by question number and topic) with equations (where relevant), figures (with axes clearly labeled, font size similar to the font in the main text, figure captions to explain their content, and numbering/lettering of figures and subpanels in order to refer to them from the main text), and references (if relevant). You should consider how to best communicate your findings in a clear and concise way to the reader. This includes designing figures appropriately, for example by making use of subplots and plotting multiple things on the same graph to save space or for aid of visual comparison, setting the axis limits appropriately,

using figure legends and colour schemes to delineate different things plotted on the same graph, etc. You are encouraged to convey your findings clearly and succinctly - longer submissions are not necessarily better. To take these factors into account, 10 marks have been allocated for the quality of presentation with regards to the overall layout and style, the quality of figures, and the scientific clarity and precision of written answers.

Introduction

Primary visual cortex (V1) is the first area in the visual system to generate selectivity for the orientation of a visual stimulus [1]. The emergence of orientation selectivity in V1 from non-orientation-tuned lateral geniculate nucleus (LGN) inputs may be viewed as a simple form of computation, and it has been proposed that an understanding of the underyling circuit mechansisms could shed light on similar computations occurring throughout the cortex, forming a kind of "canonical cortical computation" [2,3]. However, despite 60 years of intense experimental investigation, the circuit mechansisms underlying the emergence of orientation selectivity in V1 are still debated [2]. In particular, although it is is now generally agreed that orientation tuning is generated via a combination of feedforward [1] and recurrent [4,5] processes, the precise role of feedforward input, recurrent connectivity, and neuronal nonlinearities in shaping orientation-tuned responses in V1 is still a contentious issue [2].

A central thread of this debate involves the question of recurrent sharpening and amplification of orientation tuning curves by local synapses between V1 neurons [6]. Although sharpening/amplification of tuning curves has long been hypothesised as a mechanism to boost the signal to noise ratio in network responses and enhance information transmission [7,8], a number of studies have called this assumption into question [9-12].

In this assignment you will investigate how the orientation of a visual stimulus s is encoded in the responses of a simple network model of V1, and how this encoding is shaped by the recurrent connectivity of the network. To do so, you will simulate the network in different connectivity regimes (reflecting distinct hypotheses for the architecture of V1) and apply decoders to infer the stimulus value s from the simulated network responses. You will see how the decoder performance depends on the type of decoder used, the time window that neural activity is decoded from, and the connectivity between neurons in the model. You will then consider the possible implications of these models for the relationship between sensory coding and cortical connectivity.

References:

[1] Hubel DH, Wiesel TN. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J Physiol. 1962

[2] Priebe NJ. Mechanisms of Orientation Selectivity in the Primary Visual Cortex. Annu Rev Vis Sci. 2016

[3] Carandini M, Heeger DJ. Normalization as a canonical neural computation. Nat Rev Neurosci. 2011

[4] Ben-Yishai R, Bar-Or R, Sompolinsky H. Theory of orientation tuning in visual cortex. PNAS 1995

[5] Sompolinsky H, Shapley R. New perspectives on the mechanisms for orientation selectivity, Current Opinion in Neurobiology 1997

[6] Harris K., Mrsic-Flogel T. Cortical connectivity and sensory coding. Nature 2013

[7] Douglas RJ, Koch C, Mahowald M, Martin KA, Suarez HH. Recurrent excitation in neocortical circuits. Science. 1995
[8] Lien AD, Scanziani M. Tuned thalamic excitation is amplified by visual cortical circuits. Nat Neurosci. 2013
[9] Seriès P, Latham PE, Pouget A. Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. Nat Neurosci. 2004
[10] Beck J, Bejjanki VR, Pouget A. Insights from a simple expression for linear fisher information in a recurrently connected population of spiking neurons. Neural Comput. 2011
[11] Chettih SN, Harvey CD. Single-neuron perturbations reveal feature-specific competition in V1. Nature. 2019
[12] Chadwick A, Khan AG, Poort J, Blot A, Hofer SB, Mrsic-Flogel TD, Sahani M. Learning shapes cortical dynamics to enhance integration of relevant sensory input. Neuron. 2023

Ring Network Model

We will consider a network comprised of N orientation-tuned (or direction-tuned) V1 neurons receiving feedforward input from LGN and interacting via local synaptic connections. We model the average feedforward input to neuron i for the stimulus orientation s as $u_i(s) = u_0 + u_1 \cos(s - s_i)$ and the connection weight from neuron j to neuron i as $W_{ij} = W_0 + W_1 \cos(s_i - s_j)$. The preferred stimuli of the neurons are arranged evenly over the stimulus space $s \in [-\pi, \pi)$ as $s_i = 2\pi(i-1)/N - \pi$, where $i = 1, ..., N$. We model the network activity using the firing rate model

$$\tau \frac{dV_i}{dt} = -V_i + \sum_{j=1}^{N} W_{ij}\phi(V_j) + u_i(s) + \eta_i(t),$$

where $\eta_i(t)$ is Gaussian white noise representing noisy feedforward input to the network and $\phi(x) = \beta[x]_+$ is an element-wise nonlinearity mapping V_i to the instantaneous firing rate $r_i = \phi(V_i)$ (which we take to be measured in units of spikes per millisecond). This equation can be solved numerically via the Euler method using the approximation:

$$\tau \frac{V_i(t + \delta t) - V_i(t)}{\delta t} = -V_i(t) + \sum_{j=1}^{N} W_{ij}\phi(V_j(t)) + u_i(s) + \sqrt{\frac{\tau}{\delta t}}\sigma_i \times q_i(t),$$

where $q_i(t) \sim N(0, 1)$ is a normally distributed random variable drawn independently for each neuron at each simulation timestep and σ_i determines the magnitude of noise for neuron i. Finally, to model spiking activity of this network we use an inhomogeneous Poisson process. When implementing the above Euler solver, we therefore add a Poisson step $n_i(t) = \text{Poisson}(r_i(t)\delta t)$ which captures the number of spikes n_i emitted by neuron i at simulation timestep t (i.e., in a window of duration δt).

Simulation Protocol and Parameter Settings

For the following questions, simulate the above model for the stimulus $s = 0$ with $N = 100$ neurons, with a timestep of $\delta t = 1$ ms and simulation time of 500 ms, starting from initial condition $V_i(0) = 0$. Use the following parameters: $\tau = 50$ ms, $\beta = 0.1$, $\sigma_i = 0.1$, $u_0 = u_1 = 0.5$. You should repeat your simulations for each for each of the following weight settings:

1. Hubel and Wiesel regime (feedforward only): $W_0 = W_1 = 0$
2. Uniform inhibition regime: $W_0 = -4$, $W_1 = 0$
3. Strongly recurrent regime: $W_0 = -10$, $W_1 = 11$

You should perform your simulations directly in units of milliseconds. For numerical efficiency, it is recommended to implement the above equations in terms of matrix and vector operations rather than using `for` loops over neurons.

Visualisation of Network Responses in Different Regimes (40 marks)

1. Create example heatmaps (on single trials) for the evolution of $\mathbf{V}(t)$ and $\mathbf{r}(t)$ and raster plots of the spiking activity for each of the three regimes. Compare these to heatmaps of the corresponding input $\mathbf{u}(t) + \sqrt{\frac{\tau}{\delta t}}\sigma\mathbf{q}(t)$. Discuss how the different connectivity regimes differ in their response properties and how these response properties are shaped by 1) the recurrent connectivity 2) the threshold nonlinearity. [20 marks]
2. Simulate a large enough number of trials (with different random seeds for the noise term) to compute the population tuning curves $\mathbf{f}(s = 0, t) = \langle \mathbf{n}(t)/\delta t \rangle$ at each time from stimulus onset t (you may find it necessary to average spikes over some time window $[t - \Delta t/2, t + \Delta t/2]$ to reduce noise). Plot these tuning curves for each regime at a set of representative time points t. How do these tuning curves evolve over time in each regime? Comment on why these population tuning curves (for a single stimulus $s = 0$) are sufficient to determine the single-neuron tuning curves $f_i(s)$ without needing to simulate responses to each stimulus s in this model. [10 marks]
3. Compute the noise correlations $\Sigma_{ij} = \langle (\bar{n}_i - \langle \bar{n}_i \rangle)(\bar{n}_j - \langle \bar{n}_j \rangle) \rangle$ between neurons using spike counts $\bar{n}_i = \sum_{t \geq 400ms} n_i(t)$ summed over the last 100 ms of the simulation. Present the correlation matrices Σ as heatmaps for each regime. Comment on what you observe and how the form of the noise correlations reflects the network dynamics and response properties. Based on what you learned about noise correlations in lectures, how would you expect these to influence neural coding for stimulus orientation in each network? [10 marks]

Performance of Decoders (20 marks)

Assume that you record the spiking activity from the above network but do not have access to the underlying model. You may assume that you have estimated the tuning curves of each neuron and identified their preferred stimuli s_i (which you may take as their true values). Where you are asked to decode over a particular time window, you may discard any such time windows in which no neuron spikes.

4. Apply the winner take all and population vector decoders to the population response within the time window from stimulus onset ($t = 0$) up to each time t, i.e. using the cumulative spike counts $\sum_{t' \leq t} n_i(t')$. Compute and plot the root mean squared error $\sqrt{MSE(s = 0)}$ of

each decoder as a function of t (plot this in degrees rather than radians). Comment on what you observe. [10 marks]

5. Repeat the decoding analysis of question 4, but this time applying the two decoders to spiking activity within the time windows $t' \in [t_0, t_0 + \Delta t]$ of width $\Delta t = 25ms$ and with centres $t_c = t_0 + \Delta t/2$ covering the full 500 ms response period. Plot the root mean squared error of the population vector and winner take all decoders as a function of t_c. Comment on how the results compare to those obtained using the cumulative decoder above. [10 marks]

Interpretation and Implications (30 marks)

6. You are told that, if you wanted to decode the stimulus directly from the network input $\mathbf{u}(s) + \boldsymbol{\eta}(t)$, the optimal decoder at any single time point t is the population vector decoder. Moreover, you are told that the optimal decoder of the network input over the full time window $t \in [0, 500]$ is to first average (or sum) the input over time and then apply the population vector decoder to this average input. Together, this involves optimally weighting and combining the inputs both over neurons and over time. Can you use this to explain the performance of the different decoders (population vector vs WTA and instantaneous vs cumulative) applied to the output of the three networks? [Hint: consider the network-decoder cascade as a decoder of the input itself, and think about how it relates to the optimal decoder.] In your answer, discuss the sources of suboptimality of each network/decoder combination relative to the optimal decoder of network input. [20 marks]

7. Compare the findings of this project to the predictions derived in lectures for independent Poisson neurons and correlated Gaussian neurons. In particular, based on your findings, is it necessarily true that networks with sharper tuning curves and/or weaker correlations have higher information (i.e., better decoding performance)? In lectures the predictions were derived from simple tuning curve + noise encoding models, while in this assignment neural coding was studied in a recurrent network model. Based on what you have found in this project, discuss why these two approaches can lead to different conclusions, and the advantages and limitations of each approach. [10 marks]