Assignment 5

G4360 Introduction to Theoretical Neuroscience

Two types of Ring networks: Bump attractors or an SSN
Once again, the things you have to do are in red.

Bump attractor: We will construct discrete dynamics on a grid along a ring from a continuous model on a ring. In the continuous model (following Goldberg, Rokni & Sompolinsky, Neuron 42:489-500 (2004)), θ is a continuous variable from 0 to 2π, $r(\theta)$ is the response of the unit preferring orientation $\theta/2$ (preferred orientation runs around the ring from 0^o to 180^o), and $\mathbf{h}(\theta)$ is its input, $\mathbf{W}(\theta - \theta')$ is the connection between the units preferring θ and θ', and v_{th} is a threshold for firing:

$$\tau \frac{dr(\theta)}{dt} = -r(\theta) + \left[\int_0^{2\pi} \frac{d\theta'}{2\pi} W(\theta - \theta') r(\theta') + h(\theta) - v_{th} \right]_+ \tag{1}$$

Here, $[x]_+$ is rectification: $= x$ if $x > 0$, $= 0$ otherwise.
We move this to a grid of 180 grid positions around the ring, separated by $\Delta\theta = 2\pi/180$, starting from $\theta = 0$. The grid positions are θ_i, $i = 1, \ldots, 180$. There is a single unit at each position, which projects both positive and negative synapses. It has firing rate $r_i = r(\theta_i)$, and $\mathbf{r}$ is the resulting vector of rates, and similarly for h_i and $\mathbf{h}$; $\mathbf{v}_{th}$ is the vector all of whose elements are v_{th}; and $W_{ij} = W(\theta_i - \theta_j)\frac{\Delta\theta}{2\pi}$ and $\mathbf{W}$ the resulting matrix, giving dynamics

$$\tau \frac{d\mathbf{r}}{dt} = -\mathbf{r} + [\mathbf{W}\mathbf{r} + \mathbf{h} - \mathbf{v}_{th}]_+ \tag{2}$$

We define $\mathbf{W}$ and $\mathbf{h}$ from

$$W(\theta) = W_0 + 2W_1 \cos(\theta) \tag{3}$$
$$h(\theta) = h_0 + 2h_1 \cos(\theta - \theta_h) \tag{4}$$

Choose $1 < W_1 < 2$ and $W_0 + W_1 < 2$ with $0 < W_0 < 1$ (the dynamics lose stability if $W_1 > 2$ or $W_0 > 1$ or, approximately, $W_0 + W_1 > 2$; for $W_1 < 1$, there is no bump solution). τ is arbitrary, it just sets the time units; you could set it to 10 msec, or just to 1.
To compute the dynamics, you could use simple Euler, found by replacing $\frac{d\mathbf{r}}{dt}$ with $\frac{\mathbf{r}(t+\Delta t) - \mathbf{r}(t)}{\Delta t}$: $\mathbf{r}(t + \Delta t) = (1 - \frac{\Delta t}{\tau})\mathbf{r}(t) + \frac{\Delta t}{\tau}[\mathbf{W}\mathbf{r}(t) + \mathbf{h} - \mathbf{v}_{th}]_+$ Using $\Delta t/\tau = 0.1$ should be sufficient (you could check that cutting Δt in half doesn't noticeably change the outcome, which is a decent check that your time step is small enough). You could alternatively use a better method if you prefer.

a. First consider a uniform input, $h_1 = 0$. Verify that for $h_0 < v_{th}$, even if you start with a random initial condition of positive activations, the dynamics will decay to $\mathbf{r} = 0$. Simulate for a couple of values of $h_0 > v_{th}$, say $h_0 = v_{th} + 1$ and $h_0 = v_{th} + 10$. Verify that if your initial rates are not all exactly equal but have any nonzero noise (positive noise if the initial rates are otherwise zero), no matter how small, the dynamics will evolve to a bump solution (they will probably evolve to a bump solution even for an initial condition $\mathbf{r} = 0$, due to numerical noise in the simulation). There is "dynamical symmetry breaking": the dynamics and the input are circularly symmetric, but the circularly symmetric activity pattern (the uniform pattern) is unstable to any small perturbation, and the bump solution, which breaks the circular symmetry by choosing a particular location on the circle, is stable. For noisy initial conditions the bump should appear at a random location (probably selected by where some weighted sum over the initial noise in a local region is largest, but likely to appear random to you), with a common shape and height for a given h_0. How do the shape and height change for the different values of h_0?

b. Analytically, the bump activity should reach 0 at an angle ψ from the bump center, where $2W_1G_1(\psi) = 1$ and $G_1(\psi) = \frac{1}{2\pi}\left(\psi - \frac{\sin(2\psi)}{2}\right)$, with $0 < \psi < \pi$ (this is the analytic solution for continuous θ; might be slightly changed by going to a discrete grid). Does this appear to agree with your simulations? (I will place in the course directory a file, ring-model.pdf, that gives the analytics for those who are interested.)

c. Now add a weak tuned input h_1, say $h_1 = 0.1(h_0 - v_{th})$. Does this choose the bump location? Does the bump appear to be otherwise similar or identical?

d. Finally, simulate with the same parameters except $0 < W_1 < 1$. Now you should find that the uniform solution is stable, and there is no bump solution to a uniform input. What steady state do you arrive at for a non-uniform input (nonzero h_1), and how does it compare to the bump solution for $W_1 > 1$?

Additional things you might try (optional): explore the dynamics of the bumps in one or both of two ways:

- For the case with $h_1 = 0$: Add time-varying noise to the simulation, say adding some small i.i.d. noise to each h_i at each timestep, drawing the noise anew at each time step. You should find that the steady-state bump will drift in location, roughly as a random walk meaning the distance the bump travels over some time will grow as the squareroot of the time;

- For a case with $h_1 > 0$: After the steady state is reached in response to a tuned stimulus centered at θ_i, instantaneously turn that stimulus off and turn on another tuned stimulus of the same strength at a different location. How does the bump move from one location to the other – does one bump shrink while the other grows, or does the bump rotate from one position to the other? Does this depend on whether the 2nd bump is relatively near to or far from the first? How long does the change take?

Stabilized supralinear network (SSN): We'll use the same grid of 180 positions on a ring, but now there is an E and an I cell at each position. We'll consider the ring to span 180^o, representing a preferred orientation, so the grid points have spacing 1^o. We use a power-law input/output function. We use connectivity with no "Mexican hat"; we take the four connectivity functions ($E \to E$, $E \to I$, $I \to E$, $I \to I$) to have the same width, differing only in their strengths. We define these functions on the grid: the connection between the unit of type Y (E or I) at position θ_j to the unit of type X at position θ_i is

$$W_{ij}^{XY} = J^{XY} e^{-\frac{\Theta(\theta_i,\theta_j)^2}{2\sigma_W^2}} \tag{5}$$

Here, $\Theta()$ is the shortest distance around a circle defined by

$$\Theta(\theta_i, \theta_j) = \text{Min}\left(|\theta_i - \theta_j|, 180^o - |\theta_i - \theta_j|\right) \tag{6}$$

For parameters, use $J^{EE} = 0.044$, $J^{IE} = 0.042$, $J^{EI} = 0.023$, $J^{II} = 0.018$, $\sigma_W = 32^o$ (this is all following Rubin, Van Hooser and Miller, Neuron, 2015; see also Ahmadian, Rubin & Miller, Neural Computation, 2013).

We take $\mathbf{r} = \begin{pmatrix} \mathbf{r}_E \\ \mathbf{r}_I \end{pmatrix}$, where $\mathbf{r}_E$ and $\mathbf{r}_I$ are the firing rates of the E and I cells respectively, both ordered in the same way around the ring (*e.g.*, from 1^o to 180^o). Our dynamical equations are

$$\mathbf{T}\tau_E \frac{d\mathbf{r}}{dt} = -\mathbf{r} + k(\mathbf{W}\mathbf{r} + \mathbf{h})_+^n \tag{7}$$

where $(\mathbf{x})_+^n$ is applied element by element and, for a given element x_i, $(x_i)_+^n = x_i^n$ if $x_i > 0$; $= 0$, otherwise. We'll take $k = 0.04$ and $n = 2$. Take $\tau_E = 20ms$ and take $\mathbf{T}$ to be a diagonal matrix with entries 1 for the E cells and 1/2 for the I cells, *i.e.* $\tau_I = 10ms$. (The faster τ_I may not be necessary but helps to ensure stability).

For an input $\mathbf{h}$ of a stimulus with direction θ_0, the input to both the E and the I units at θ_i is $h_i = ce^{-\frac{\Theta(\theta_i,\theta_0)^2}{2\sigma_h^2}}$. Here, c is a constant (c for 'contrast') that you will vary to vary the strength of the stimulus. Take $\sigma_h = 30^o$.

a. First, for a single stimulus of orientation of your choice θ_0, simulate the response, starting from an initial condition $\mathbf{r} = 0$, for $c = \{1.25, 2.5, 5, 10, 20, 40\}$. Again, use first-order Euler, a time step of $1ms$ should be fine. For each c, simulate until a steady state is reached by some criterion (change per timestep gets sufficiently small).

For the steady state, for the E unit and the I unit at the stimulus center, plot, as a function of c:

- Their firing rate;
- Their feedforward input, their net recurrent input ($E - I$, where E is the recurrent excitatory input and I is the recurrent inhibitory input, taken to have a positive sign), and their total input (feedforward + net recurrent).
- The percent of the unit's input that is feedforward or is recurrent, counting recurrent input now as $E + I$ and total input as $FF + E + I$
- For the recurrent input, the percent of it that is excitatory: $\frac{E}{E+I}$

You should see: saturation of excitatory firing rates; a transition from a feedforward-dominated regime for weak input, to a recurrent-dominated regime for stronger input; that for stronger input, the recurrent input largely cancels or 'balances' the feedforward input; and that the recurrent input becomes more inhibition-dominated for stronger stimuli.

b. Now consider adding a 2nd stimulus 90^o away from the first (*i.e.*, on the opposite side of the ring). By symmetry, that stimulus by itself should produce a response exactly like the response to the θ_0 stimulus, except shifted by 90^o. So you don't need to simulate response to that stimulus alone; but simulate response to the two stimuli shown at the same time, again for the given values of c (same c for both stimuli). You know by symmetry that the responses must be identical at each stimulus center. So, choosing the units at one of the stimulus centers, for the E and for the I units, plot the ratio of their steady-state response when both stimuli are shown together, to the sum of their steady-state responses to the two stimuli when each stimulus is shown alone; that is, the ratio of the actual response to the response you would get if responses to the two stimuli sum linearly. You should find that this ratio is > 1, representing supralinear summation, for weaker inputs but < 1, representing sublinear summation, for stronger inputs.

c. For at least some, if not all, of the c values, you probably want to plot, with preferred orientation from 0^o to 180^o on the x axis, a curve of E unit responses vs. position x on the ring, as follows: the sum of the two responses to each stimulus shown alone; and the response to the two stimuli shown together. In another graph, plot the same for the I unit. This will allow you to directly see the supralinear and sublinear summation.

Other things you might want to try (optional):

- Give a uniform input of varying strengths to the network; do you ever see non-uniform solutions emerge? (you shouldn't)
- Consider adding the two stimuli with different c values; you should see the emergence of "winner-take-all" behavior, where the greater the difference between the c values for the two different stimuli, the more the response to the weaker stimulus is suppressed (relative to its response if shown alone with that c value) and the more the response to the stronger stimulus approaches the response if it were shown by itself.