Assignment Network Plasticity, notes to answers

Neural Computation 2016-2017. Mark van Rossum

20th December 2016

Model and setup

In this assignment we consider a highly abstracted version of a piece of sensory cortex. We consider N neurons connected via plastic synapses to $M = 50$ inputs. For convenience, the inputs are arranged on a circle $\frac{2\pi}{M}, 2\frac{2\pi}{M}, \ldots, 2\pi$.

The input is structured as follows: The inputs are binary (0/1). On a given timestep a random unit on the circle is chosen, and this unit as well as the next $d-1$ units in the clockwise direction are turned on. In other words, the stimulus is a randomly placed arc of fixed length. Every timestep a different stimulus is presented. (check for yourself your simulation of this is working correctly).

Initially, the neural activity of neuron j, y_j, is given by $y_j = [\sum_i w_{ji} x_i]_+$. The synapses, described by a $N \times M$ matrix, are every time-step updated according to the rule $w_{ji} \to w_{ji} + \epsilon x_i (y_j^2 - y_j/10)$. After this update the weights are multiplicatively normalized such that for each neuron j, $\sum_i w_{ij}^2 = 1$. Initialize the weights with uniform random numbers.

Take the learning rate $\epsilon = 0.02$, and stimulus width $d = 5$.

In principle we could track the system after every time-step, but to reduce computation, we let the system run in chunks of 100 time-steps. We call such a chunk an 'episode'.

Questions

Question 1 (5 points) First we study independent neurons. You can set $N = 1$, or run multiple neurons in parallel. Plot the receptive field (i.e. the weight vector) after a few episodes. Describe it's shape. What happens if not all stimuli are equally likely?

- *You should find that the neuron develops a single bell-shaped receptive field with some fringes on the edge. If you simulate longer you find that these fringes typically disappear.*
- *There are many ways to make some stimuli more common. You can for instance alternate between 2 locations. You should find that the receptive field will be locate at either position. Thus the learning is competitive: only a few locations 'win'.*
- *The shape is due to the fact that also input close to the center of the receptive field can lead to post activity, but one that is generally less, so the weight increase will be smaller. If you fix the stimulus at a single location, you find a square-shaped receptive field.*

Question 2 (10 points) Track the receptive field across episodes. What happens if you vary ϵ? Also compare your findings to the case of classical Hebbian learning in which case we have $w_{ji} \to w_{ji} + \epsilon x_i y_j$ (and still normalize as above). Explain the difference and justify your claims.

- *You should see faster moving receptive field with both larger* ϵ.
- *With Hebbian learning, the RF moves around more. It might also have multiple peaks or fringes*
- *A nice demonstration that one student did is to track the RF development of parallel independent neurons receiving identical input. For Hebbian learning, the RFs converge quickly to be the same.*
- *The Hebbian learning is less stable because it is less competitive. Ignoring the normalization for now, the weight update is always positive. In contrast, in the original learning rule, small weights are punished.*
- *Note that the receptive field can actually not jump, but always moves continuously.*

Because the inputs are placed on a ring, we can describe the weight vector of each neuron by a population vector (see lecture notes). We are interested in its angle, called ϕ.

Question 3 (10 points) Describe how you calculated the angle. Plot the angle across episodes. The auto-covariance of the angle is $\langle(\phi(i)-\bar{\phi})(\phi(i+k)-\bar{\phi})\rangle$ where $\phi(i)$ is the angle after episode i, $\bar{\phi}$ is the average angle, and the average is over trials or episodes. Plot the auto-covariance as a function of episode (k) for both cases of the previous question. Describe in interpret the various aspects of the plot (peak, limits, etc.).

- *You can calculate the angle as* $\arctan(\sum_i w_i \cos\psi_i, \sum_i w_i \sin\psi_i)$ *where* $\psi_i = 2\pi i/M$ *runs between 0 and* 2π*. (Note, it is not necessary, nor recommended to work in degrees).*
- *The autocorrelation is a symmetric peak. Decaying to zero for large k, because for long enough time the RF will be uncorrelated. There might be some damped oscillations in the curve. I think these are numerical artefacts, arising from not long enough simulations; there is nothing in the simulations that would cause oscillations.*
- *Note that in the limit of large k, the angles should be uncorrelated,* $\langle\phi(i)\phi(i+k)\rangle = \langle\phi(i)\rangle\langle\phi(i+k)\rangle$*. For* $k=0$, $\langle\phi(i)\phi(i)\rangle - \bar{\phi}^2 = var(\phi(i))$. $\bar{\phi}=0$ *and, for a long enough simulation* ϕ *follows a uniform distribution between* $-\pi$ *and* π, *var*$=\pi^2/3 \approx 3.3$*. In particular for the more stable quadratic learning rule, you might not see those limits because you would need to simulate much longer to sample all angles.*
- *With standard Hebb learning the autocorrelation decays much faster than the quadratic one.*

Next, we set the number of neurons to $N = 5$ and let them interact. The interactions are all-to-all and all the same strength, but self-interactions are excluded.

For the dynamics we use $y_j = [\sum_i w_{ji} x_i + q \sum_{j' \neq j} \sum_i w_{j'i} x_i]_+$. Thus when q is positive, the interaction is excitatory.

Question 4 (5 points) The above equation for y_j is an approximation. Assuming firing rate dynamics as in the lecture notes and assuming that the neural dynamics are much faster than the duration of the stimulus, what would be a more correct equation for the value of y_j reached at the end of each stimulus?

- *Consider one of the neurons. The used equation assumes that the activity of the other neurons is described by* $\sum w_{j'i} x_i$. *It thereby ignores the indirect effect of the lateral connections on the other neurons.*
- *To be more correct, the activities have to be solved iteratively. Assuming the rate dynamics from the lecture notes: One could simulate* $\tau dy_j(t)/dt = -y_j(t) + h_j$, *with* $h_j = [\sum_i w_{ji} x_i + q \sum_{j' \neq j} y_{j'}]_+$ *until steady state is reached.*
- *If the rectification can be ignored, you could write the steady state directly:* $y_j^{ss} = (I - Q)^{-1} (\sum_i w_{ji} x_i)$, *where* $Q_{ij} = (1 - \delta_{i,j}) q$. *See Dayan and Abbot book for more about this case.*

Question 5 (10 points) Study the angles that develop for both positive and negative values of q. Also plot the covariance function of the angles between pairs of neurons. Interpret the results.

- *For positive (and zero) q, the interactions are excitatory and the neurons will share the same RF.*
- *For negative q the RFs spread out to equally cover the stimulus space.*
- *Can't make q too negative ($\lesssim -0.4$) as then the inhibition is too strong and there is never activity.*