

Assignment 3

Theoretical Neuroscience

TAs:

Kirsty McNaught (kirsty.mcnaught.16@ucl.ac.uk)

Lea Duncker (lea.duncker.11@ucl.ac.uk)

Jorge Menendez (jorge.menendez.15@ucl.ac.uk)

Due 8 March, 2018

1. Continuous time Hopfield networks

Consider a continuous time Hopfield network,

$$\frac{dx_i}{dt} = \phi(h_i) - x_i \quad (1)$$

where ϕ is the gain function (taken to be non-negative, more or less sigmoidal, and saturating), N is the number of neurons, and h_i is the synaptic drive,

$$h_i \equiv \sum_{j=1}^N J_{ij} x_j. \quad (2)$$

We'll let

$$J_{ij} = \frac{1}{Nf(1-f)} \sum_{\mu=1}^p \xi_i^\mu (\xi_j^\mu - f) \quad (3)$$

where the ξ_i^μ are random binary vectors, a fraction f of which are 1,

$$\xi_i^\mu = \begin{cases} 1 & \text{probability } f \\ 0 & \text{probability } 1 - f. \end{cases} \quad (4)$$

There are several differences between this formulation and the one we used in class: the x_i are continuous rather than discrete; the gain function is smooth and non-negative (the latter ensuring that the x_i will be non-negative); the elements of the patterns are 0 and 1 rather than -1 and 1 ; and the probability of 1 is f rather than $1/2$. However, the analysis is nearly identical.

As usual, the goal is to find the equilibria. With this formulation, the equilibria aren't necessarily all that close to the patterns, ξ_i^μ . However, we still expect the equilibria to be at least related to the patterns. With that in mind, we define the overlaps, denoted m_μ , via

$$m_\mu = \frac{1}{Nf(1-f)} \sum_i (\xi_i^\mu - f)x_i. \quad (5)$$

If $x_i = \xi_i^\mu$, m_μ will be close to 1, whereas if x_i is independent of ξ_i^μ , m_μ will be close to zero. At an equilibrium, we expect one of the m_μ to be large and the rest to be small.

(a) For this question, we'll let J_{ij} be symmetric, but otherwise arbitrary. Define the "energy" E via

$$E \equiv -\frac{1}{2} \sum_{ij} x_i J_{ij} x_j + \sum_i \psi(h_i) \quad (6)$$

where ψ is obeys

$$\frac{d\psi(x)}{dx} = \phi(x). \quad (7)$$

Show that

$$\frac{dE}{dt} = - \sum_{ij} \frac{dx_i}{dt} J_{ij} \frac{dx_j}{dt}. \quad (8)$$

Thus, if J_{ij} is symmetric and positive definite (consistent with Eq. (3) if $f = 0$), then E is a non-decreasing function of time. I tried, but could not find, a Lyapunov function when J_{ij} is symmetric but not positive definite. That does not mean one does not exist. For extra credit, find one!

For the rest of the questions, use the connection matrix given in Eq. (3).

(b) Show that

$$h_i = \sum_{\mu} \xi_i^{\mu} m_{\mu}. \quad (9)$$

(c) Show that the m_{μ} obey the equation

$$\frac{dm_{\nu}}{dt} = \frac{1}{N} \sum_i \frac{\xi_i^{\nu}}{f} \phi\left(m_{\nu} + \sum_{\mu \neq \nu} \xi_i^{\mu} m_{\mu}\right) - \frac{1}{N} \sum_i \frac{1 - \xi_i^{\nu}}{1 - f} \phi\left(\sum_{\mu \neq \nu} \xi_i^{\mu} m_{\mu}\right) - m_{\nu}. \quad (10)$$

(d) Define

$$\zeta_i \equiv \sum_{\mu \neq \nu} \xi_i^{\mu} m_{\mu}. \quad (11)$$

We're going to treat ζ_i as a random variable with respect to index, i . Because ξ_i^{ν} and ξ_i^{μ} are uncorrelated, it follows that ξ_i^{ν} is independent of ζ_i . Thus, in the large N limit, the equation for the m_{μ} becomes

$$\frac{dm_{\nu}}{dt} = \Phi(m_{\mu}) - m_{\mu} \quad (12)$$

where

$$\Phi(m) \equiv \int d\zeta p(\zeta) [\phi(m + \zeta) - \phi(\zeta)]. \quad (13)$$

Note that $\Phi(m)$ is just a smoothed, and offset, version of $\phi(m)$.

Equation (12) has an equilibrium at $m_{\nu} = 0$. Assuming $p(\zeta)$ does not change with time, is it possible for this equilibrium to be stable, and still have a stable equilibrium at $m_{\mu} > 0$?

- (e) Again assuming $p(\zeta)$ is constant, show that if $\Phi'(0) > 0$, then we're guaranteed to have a stable equilibrium with $m_{\mu} > 0$.
- (f) Would the answers to the above two questions change if we dropped the (unrealistic) assumption that $p(\zeta)$ is constant?
- (g) Assume that x_i is independent of ξ_j^{μ} when $\mu \neq \nu$. Show that ζ_i is a zero mean Gaussian random variable with variance, denoted σ^2 , given by

$$\sigma^2 = \frac{p-1}{Nf(1-f)} \langle x_i^2 \rangle. \quad (14)$$

This would suggest that small f is bad, since it increases the variance. However, if $x_i \propto \xi_i^{\nu}$, then $\langle x_i^2 \rangle \propto f$, and $\sigma^2 \propto 1/(1-f)$.

- (h) This isn't a question, but there are a couple things to notice. For the system to have a “memory” – a fixed point for which m_μ is $\mathcal{O}(1)$ – the smoothed gain function, $\Phi(m)$, must be sufficiently steep. Thus, σ can't be too big (because the larger σ is, the more the gain function is smoothed; see Eq. (13)). Given Eq. (14), for small f the variance should scale as p/N , which would mean that the capacity shouldn't depend much on f (at least when f is small). However, I told you in class (and in the Hopfield writeup) that capacity scales as $1/f$. I have never been able to find a simple explanation for the $1/f$ scaling.

2. Coupled line attractor

Consider a coupled network of N neurons whose units evolve according to

$$\frac{dr_i}{dt} = \phi\left(\sum_j W_{i-j}r_j + h_i\right) - r_i \quad (15a)$$

$$\tau \frac{dh_i}{dt} = g(t) \sum_j A_{i-j}r_j - h_i. \quad (15b)$$

We'll take W to be symmetric: $W_{i-j} = W_{j-i}$. Assume that when $h_i = 0$, Eq. (15a) has a stable equilibrium given by $f(\theta_i - \theta)$,

$$f(\theta_i - \theta) = \phi\left(\sum_j W_{i-j}f(\theta_j - \theta)\right) \quad (16)$$

where the θ_i are equally spaced. Assume that this equation is satisfied for all θ , making it a true line attractor.

- (a) In the limit that $g(t)$ is infinitesimally small, show that the position on the line attractor, θ , evolves according to

$$\tau \frac{d^2\theta}{dt^2} + \frac{d\theta}{dt} = g(t) \sum_{ij} v_{0i}^\dagger(\theta) \phi'_i A_{i-j} f(\theta_j - \theta) \quad (17)$$

where

$$\phi'_i \equiv \phi'\left(\sum_j W_{i-j}f(\theta_j - \theta)\right), \quad (18)$$

$\mathbf{v}_0^\dagger(\theta)$ is the adjoint eigenvalue of the linearized dynamics whose eigenvalue is 0,

$$\sum_j v_{0j}^\dagger(\theta) \phi'_j W_{j-i} = v_{0i}^\dagger(\theta) \quad (19)$$

and it's normalized so that

$$\sum_i v_{0i}^\dagger(\theta) f'(\theta_i - \theta) = 1. \quad (20)$$

- (b) Recall that the adjoint eigenvector is related to $f(\theta_i - \theta)$ via

$$v_{0i}^\dagger(\theta) = \frac{f'(\theta_i - \theta) / \phi'_i(\theta)}{Z} \quad (21)$$

where

$$Z \equiv \sum_i \frac{f'(\theta_i - \theta)^2}{\phi'_i(\theta)}. \quad (22)$$

Consequently, θ evolves according to

$$\tau \frac{d^2\theta}{dt^2} + \frac{d\theta}{dt} = \frac{g(t)}{Z} \sum_{ij} f'(\theta_i - \theta) A_{i-j} f(\theta_j - \theta). \quad (23)$$

Show that in the large N limit, Z is independent of θ .

- (c) Show that in the large N limit, the right hand side of Eq. (23) becomes independent of θ . Show also that if A_{i-j} is even ($A_{i-j} = A_{j-i}$), the right hand side is zero.

3. Networks with time-varying dynamics

Consider a network of N neurons that evolves according to

$$\frac{dx_i}{dt} = \phi \left(\sum_j W_{ij} x_j + \sum_\mu J_{i\mu} z_\mu + \sum_\mu C_{i\mu} u_\mu(t) \right) - x_i \quad (24)$$

where $u_\mu(t)$ is a control signal, ϕ is the gain function (as usual, it's more or less sigmoidal), and \mathbf{z} is related to \mathbf{x} via

$$z_\mu = \sum_j A_{\mu j} x_j. \quad (25)$$

In this setting the dimensionality of both \mathbf{z} and \mathbf{u} is typically much less than N , but that's not necessary for the questions.

- (a) Show that z_μ evolves according to

$$\frac{dz_\mu}{dt} = \sum_i A_{\mu i} \phi \left(\sum_j W_{ij} x_j + \sum_\nu J_{i\nu} z_\nu + \sum_\nu C_{i\nu} u_\nu(t) \right) - z_\mu \quad (26)$$

Thus, if $W_{ij} = 0$,

$$\frac{dz_\mu}{dt} = f_\mu(\mathbf{z}, \mathbf{u}(t)) \quad (27)$$

where the function f_μ is given by a neural network with one hidden layer.

- (b) Assume the goal of the network is to produce as output the function $z_\mu^*(t)$. Show that under the learning rule

$$\Delta A_{\mu i} = -\eta (z_\mu^*(t) - z_\mu(t)) x_i(t), \quad (28)$$

the instantaneous error, $(z_\mu^*(t) - z_\mu(t))^2$, decreases. Assume that η , the learning rate, is small. Is there any guarantee that the total error, which is the time average of $(z_\mu^*(t) - z_\mu(t))^2$, will decrease?