Lecture 3: Recurrent networks

Jochen Braun

Otto-von-Guericke-Universität Magdeburg,
Cognitive Biology Group

Theoretical Neuroscience II, SS 2020

Lecture 3: Recurrent networks

Think recurrent networks as infinite cascade of feedforward networks. How can we analyze behaviour and convergence? **Linear approximation** *provides mathematical tools: eigenvectors of connectivity matrix.*
Linear solution: *Obtain time-development of activity* $\boldsymbol{v}(t)$, *given input* $\boldsymbol{h}$ *and initial condition* $\boldsymbol{v}(0)$, *as linear combination of eigenvectors* $\boldsymbol{e}_\mu$, *with time-varying coefficients* $c_\mu(t)$. *Eigenvalues* λ_μ *determine amplification: development in direction* $\boldsymbol{e}_\mu$ *is stable if* $0 < \lambda_\mu < 1$, *unstable if* $1 < \lambda_\mu$. **Linear network can** *selectively amplify dominant eigenvector* $\boldsymbol{e}_1$, *or linear combinations of degenerate eigenvectors* $\boldsymbol{e}_{1,2}$.
Non-linear networks *are (slightly) more biological, but must be analyzed in simulation. They can amplify selectively, disproportionately amplify largest input ('winner-take-all'), and multiplicatively scale inputs ('gain modulation').*

Organisation of lecture

- **1. Recurrent networks**
- **2. Solution of linear approximation**
- **3. Capabilities of *linear* recurrent networks**
- **4. Capabilities of *non-linear* recurrent networks**

1. Recurrent networks, further simplified

Consider a two-layer network with input units $[u_1, u_2, \ldots]$ driving output units $[v_1, v_2, \ldots]$, which are also recurrently connected among themselves. The dynamic equations of output activity v_m is

$$\tau \frac{dv_m}{dt} = -v_m + F\left(\sum_n W_{mn}\, u_n + \sum_{m'} M_{mm'}\, v_{m'} \right)$$

where W_{mn} is the synaptic weight of the feedforward connection from u_n to v_m, and $M_{mm'}$ is the synaptic weight of the recurrent connection from $v_{m'}$ to v_m.

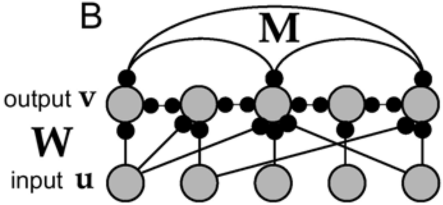

Dynamic equation of recurrent network

As always, it is convenient to collect the dynamic equations of all output units into a single vector equation

$$\tau \frac{d\boldsymbol{v}}{dt} = -\boldsymbol{v} + F\left(\boldsymbol{W} \cdot \boldsymbol{u} + \boldsymbol{M} \cdot \boldsymbol{v}\right)$$

where $\boldsymbol{W}$ and $\boldsymbol{M}$ are the feedforward and recurrent synaptic weight matrices, respectively. This is the dynamic equation of the entire network.

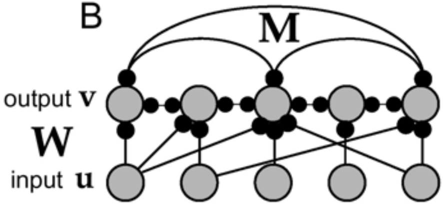

Combine feedforward inputs

To simplify even further, we can combine all feedforward inputs into a vector

$$\tau \frac{d\boldsymbol{v}}{dt} = -\boldsymbol{v} + F\left(\boldsymbol{h} + \boldsymbol{M} \cdot \boldsymbol{v}\right), \qquad \boldsymbol{h} = \boldsymbol{W} \cdot \boldsymbol{u}$$

where $\boldsymbol{h}$ is the (combined) feedforward inputs and $\boldsymbol{M}$ is the recurrent synaptic weight matrix.

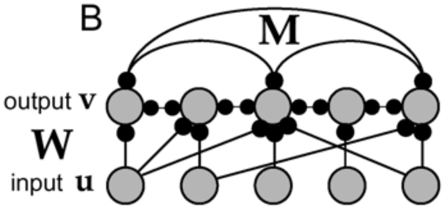

Steady-states in nature

In a steady-state, the rate of increments is exactly balanced by the rate of decrements.

In a neural network, the excitatory synaptic inputs are exactly balanced by inhibitory synaptic inputs.

Steady-state condition

A steady-state, each component v_o of the rate vector $\boldsymbol{v}_{ss}$ must satisfy

$$0 = \underbrace{-v_o}_{relaxation} + \underbrace{F\left(h_o + \sum_{o'} M_{oo'} \cdot v_{o'}\right)}_{excitatory\ and\ inhibitory\ inputs}$$

In other words, for each population o, the total effect of excitatory and inhibitory inputs must exactly balance the effect of relaxation, so that the time-derivative ${}^{dv_o}/_{dt}$ vanishes.

Determining the existence and location of such vectors requires more advanced methods, involving eigenvectors and eigenvalues of recurrent connectivity $\boldsymbol{M}$.

Linear approximation

To obtain a general idea of the dynamic possibilities, we consider a *drastically simplified* system, in which the transfer function has been replaced by the identity function (this is not biological!):

$$\boldsymbol{v}_{out} = \boldsymbol{F}\left(\boldsymbol{v}_{in}\right) \qquad \rightarrow \qquad \boldsymbol{v}_{out} = \boldsymbol{v}_{in}$$

The dynamic equation then simplifies to:

$$\tau \frac{d\boldsymbol{v}}{dt} = -\boldsymbol{v} + \boldsymbol{F}\left(\boldsymbol{h} + \boldsymbol{M} \cdot \boldsymbol{v}\right) \qquad \rightarrow \qquad \tau \frac{d\boldsymbol{v}}{dt} = -\boldsymbol{v} + \boldsymbol{h} + \boldsymbol{M} \cdot \boldsymbol{v}$$

For this simplification, we seek the time-evolution of the output vector $\boldsymbol{v}(t)$, given a constant input vector $\boldsymbol{h}$ and initial output vector $\boldsymbol{v}(0)$.

Example

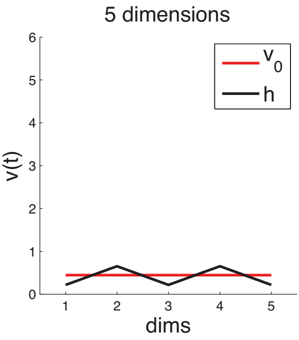

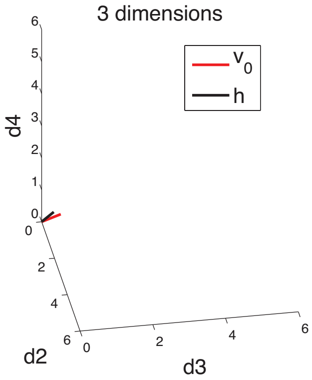

$$\boldsymbol{M} = \begin{pmatrix} 0.1 & 0 & 0 & 0 & 0 \\ 0 & 0.3 & 0.28 & 0.2 & 0 \\ 0 & 0.28 & 0.5 & 0.28 & 0 \\ 0 & 0.2 & 0.28 & 0.3 & 0 \\ 0 & 0 & 0 & 0 & 0.1 \end{pmatrix}, \quad \boldsymbol{h} = \begin{pmatrix} 0.2 \\ 0.6 \\ 0.2 \\ 0.6 \\ 0.2 \end{pmatrix}, \quad \boldsymbol{v}_0 = \begin{pmatrix} 0.4 \\ 0.4 \\ 0.4 \\ 0.4 \\ 0.4 \end{pmatrix}$$

Note that v_1 and v_5 don't communicate with v_2, v_3, and v_4.

Example, ctd

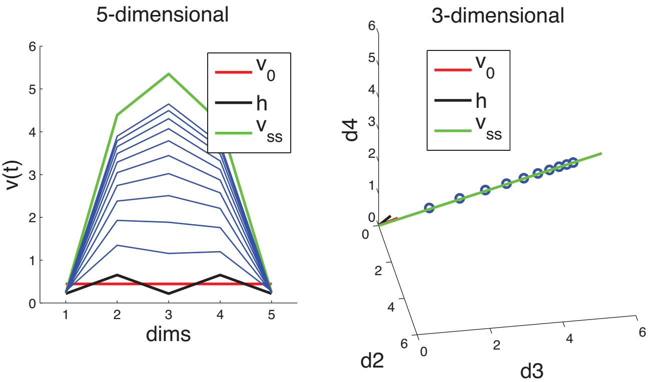

$$M = \begin{pmatrix} 0.1 & 0 & 0 & 0 & 0 \\ 0 & 0.3 & 0.28 & 0.2 & 0 \\ 0 & 0.28 & 0.5 & 0.28 & 0 \\ 0 & 0.2 & 0.28 & 0.3 & 0 \\ 0 & 0 & 0 & 0 & 0.1 \end{pmatrix}, \quad \boldsymbol{h} = \begin{pmatrix} 0.2 \\ 0.6 \\ 0.2 \\ 0.6 \\ 0.2 \end{pmatrix}, \quad \boldsymbol{v}_0 = \begin{pmatrix} 0.4 \\ 0.4 \\ 0.4 \\ 0.4 \\ 0.4 \end{pmatrix}$$

Note that v_1 and v_5 don't communicate with v_2, v_3, and v_4

Summary so far

- Recurrent networks have richer dynamics than feedforward networks, but are more difficult to analyze.
- In *linearized* networks, the activation function $F(\boldsymbol{M} \cdot \boldsymbol{u} + \boldsymbol{h})$ is replaced by $\boldsymbol{M} \cdot \boldsymbol{u} + \boldsymbol{h}$.
- Linearized networks capture many qualitative aspects of the non-linear network.
- Outlook: the analysis of linearised networks will revolve around the eigenvectors/eigenvalues of $\boldsymbol{M}$.

2. Eigenvectors and eigenvalues

We will use of the *eigenvectors* and *eigenvalues* of the connectivity matrix ('eigen' is German for 'one's own'). In a system with N units, the connectivity matrix is of size $N \times N$ and has N eigenvectors $\boldsymbol{e}_\mu$ with associated eigenvalues λ_μ:

$$\boldsymbol{M} \cdot \boldsymbol{e}_\mu = \lambda_\mu \, \boldsymbol{e}_\mu, \qquad \mu = 1, 2, \ldots, N$$

Matrix scales eigenvectors proportionally (by eigenvalue), but rotates (and scales) all other vectors.

Eigenvectors have unit length. Eigenvectors of symmetric matrices (biologically plausible!) are orthogonal:

$$\boldsymbol{e}_\mu \cdot \boldsymbol{e}_\nu = \delta_{\mu\nu} = \begin{cases} 1 & \text{if} \quad \mu = \nu \\ 0 & \text{if} \quad \mu \neq \nu \end{cases}$$

Greek letters μ and ν index eigenvectors, roman letters i, j, n etc index network populations.

Example

Use Matlab-expression *[e,lambda] = eig(M)* to obtain eigenvectors and eigenvalues of $\boldsymbol{M}$.

$$\boldsymbol{M} = \begin{pmatrix} 0.1 & 0 & 0 & 0 & 0 \\ 0 & 0.3 & 0.28 & 0.2 & 0 \\ 0 & 0.28 & 0.5 & 0.28 & 0 \\ 0 & 0.2 & 0.28 & 0.3 & 0 \\ 0 & 0 & 0 & 0 & 0.1 \end{pmatrix}$$

$$\boldsymbol{e}_1 = \begin{pmatrix} 0 \\ 0.5 \\ 0.7 \\ 0.5 \\ 0 \end{pmatrix}, \quad \lambda_1 = 0.9, \qquad \boldsymbol{e}_2 = \begin{pmatrix} 0 \\ 0.5 \\ -0.7 \\ 0.5 \\ 0 \end{pmatrix}, \quad \lambda_2 = 0.1, \qquad \boldsymbol{e}_3 = \begin{pmatrix} 0 \\ 0.7 \\ 0 \\ -0.7 \\ 0 \end{pmatrix}, \quad \lambda_3 = 0.1, \ldots$$

Orthonormal basis

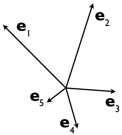

Given the orthogonal coordinate system (or basis) provided by the normalized eigenvectors, we can write any N-dimensional vector as a linear expansion of eigenvectors

$$\boldsymbol{x} = \sum_{\mu=1}^{N} a_{\mu} \, \boldsymbol{e}_{\mu}$$

where a_{μ} are linear coefficients.

Solution as a linear expansion of eigenvectors

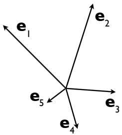

Similarly, we can write the time-varying, N-dimensional vector $\boldsymbol{v}(t)$, which is solution of our dynamic equation, as a linear expansion of eigenvectors

$$\boldsymbol{v}(t) = \sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu$$

where $c_\mu(t)$ are time-dependent, linear coefficients.

Equations for coefficients (board)

Substituting the linear expansion in

$$\tau \frac{d\boldsymbol{v}}{dt} = -\boldsymbol{v} + \boldsymbol{h} + \boldsymbol{M} \cdot \boldsymbol{v}, \qquad \boldsymbol{v}(t) = \sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu$$

leads to

$$\tau \sum_{\mu=1}^{N} \frac{dc_\mu(t)}{dt}\, \boldsymbol{e}_\mu = -\sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu + \boldsymbol{h} + \sum_{\mu=1}^{N} c_\mu(t)\, \lambda_\mu \cdot \boldsymbol{e}_\mu$$

and

$$\tau \frac{dc_\nu(t)}{dt} = -(1-\lambda_\nu)\, c_\nu(t) + \boldsymbol{h} \cdot \boldsymbol{e}_\nu \qquad \forall \nu$$

If the time-dependent coefficients $c_\nu(t)$ satisfy these equations for all ν, then the time-dependent vector $\boldsymbol{v}(t)$ satisfies the dynamic equation of the network.

Details

$$\tau \frac{d\boldsymbol{v}}{dt} = -\boldsymbol{v} + \boldsymbol{h} + \boldsymbol{M} \cdot \boldsymbol{v}, \qquad \boldsymbol{v}(t) = \sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu$$

$$\tau \frac{d}{dt} \sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu = -\sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu + \boldsymbol{h} + \boldsymbol{M} \cdot \left(\sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu \right)$$

$$\tau \sum_{\mu=1}^{N} \frac{dc_\mu(t)}{dt}\, \boldsymbol{e}_\mu = -\sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu + \boldsymbol{h} + \sum_{\mu=1}^{N} c_\mu(t)\, \underbrace{\boldsymbol{M} \cdot \boldsymbol{e}_\mu}$$

Using the eigenvector property $\boldsymbol{M} \cdot \boldsymbol{e}_\mu = \lambda_\mu\, \boldsymbol{e}_\mu$, we get

$$\tau \sum_{\mu=1}^{N} \frac{dc_\mu(t)}{dt}\, \boldsymbol{e}_\mu = -\sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu + \boldsymbol{h} + \sum_{\mu=1}^{N} c_\mu(t)\, \lambda_\mu\, \boldsymbol{e}_\mu$$

Taking the dot-product with an eigenvector $\boldsymbol{e}_\nu$

$$\tau \sum_{\mu=1}^{N} \frac{dc_\mu(t)}{dt}\, \boldsymbol{e}_\mu = -\sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu + \boldsymbol{h} + \sum_{\mu=1}^{N} c_\mu(t)\, \lambda_\mu\, \boldsymbol{e}_\mu \quad \Big| \cdot \boldsymbol{e}_\nu$$

and using the orthogonality of eigenvectors $\boldsymbol{e}_\mu \cdot \boldsymbol{e}_\nu = \delta_{\mu\nu}$, the sums disappear, leaving only

$$\tau \frac{dc_\nu(t)}{dt}\, \boldsymbol{e}_\nu \cdot \boldsymbol{e}_\nu = -c_\nu(t)\, \boldsymbol{e}_\nu \cdot \boldsymbol{e}_\nu + \boldsymbol{h} \cdot \boldsymbol{e}_\nu + c_\nu(t)\, \lambda_\nu\, \boldsymbol{e}_\nu \cdot \boldsymbol{e}_\nu$$

$$\tau \frac{dc_\nu(t)}{dt} = -(1 - \lambda_\nu)\, c_\nu(t) + \boldsymbol{h} \cdot \boldsymbol{e}_\nu \qquad \forall \nu$$

Time-dependent coefficients

For constant input $\boldsymbol{h}$, the dynamic equation

$$\tau \frac{dc_\mu(t)}{dt} = -(1-\lambda_\mu)\, c_\mu(t) + \boldsymbol{h}\cdot\boldsymbol{e}_\mu$$

has the solution

$$c_\mu(t) = \frac{\boldsymbol{h}\cdot\boldsymbol{e}_\mu}{1-\lambda_\mu}\left[1-\exp\left(-\frac{1-\lambda_\mu}{\tau}t\right)\right] + \boldsymbol{v}_0\cdot\boldsymbol{e}_\mu\,\exp\left(-\frac{1-\lambda_\mu}{\tau}t\right)$$

If $\lambda_\mu > 1$, the exponential term is *positive* and coefficient $c_\mu(t)$ *grows* exponentially with time. If $\lambda_\mu < 1$, the exponential term is *negative* and $c_\mu(t)$ approaches asymptotic value.

Note that for $t = 0$, we obtain

$$c_\mu(0) = \boldsymbol{v}_0\cdot\boldsymbol{e}_\mu \quad \Leftrightarrow \quad \boldsymbol{v}_0 = \sum_{\mu=1}^{N} c_\mu(0)\,\boldsymbol{e}_\mu$$

Case $0 < \lambda_\mu < 1$

For eigenvectors with eigenvalues $0 < \lambda_\mu < 1$, coefficients relax exponentially from an initial value $\boldsymbol{v}_0 \cdot \boldsymbol{e}_\mu$ to a steady-state value $\frac{\boldsymbol{h} \cdot \boldsymbol{e}_\mu}{1-\lambda_\mu}$:

$$\lim_{t \to 0} \left[1 - \exp\left(- \frac{1-\lambda_\mu}{\tau} t \right) \right] = 0, \qquad \lim_{t \to 0} \exp\left(- \frac{1-\lambda_\mu}{\tau} t \right) = 1$$

$$\lim_{t \to \infty} \left[1 - \exp\left(- \frac{1-\lambda_\mu}{\tau} t \right) \right] = 1, \qquad \lim_{t \to \infty} \exp\left(- \frac{1-\lambda_\mu}{\tau} t \right) = 0$$

$$\lim_{t \to \infty} c_\mu(t) = \frac{\boldsymbol{h} \cdot \boldsymbol{e}_\mu}{1 - \lambda_\mu}, \qquad \lim_{t \to 0} c_\mu(t) = \boldsymbol{v}_0 \cdot \boldsymbol{e}_\mu$$

The factor $1/_{1-\lambda_\mu}$ determines how much the input projection onto $\boldsymbol{e}_\nu$ is *amplified* by the recurrent dynamics. The closer λ_μ to one, the larger the amplification.

Example, again

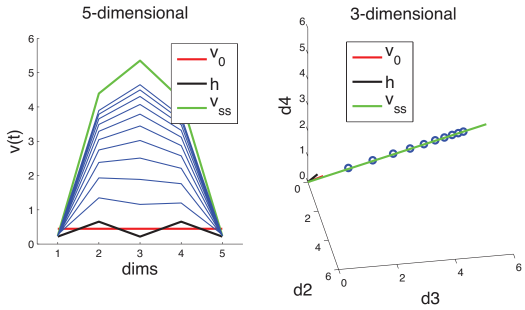

$$\boldsymbol{h} = \begin{pmatrix} 0.2 \\ 0.6 \\ 0.2 \\ 0.6 \\ 0.2 \end{pmatrix}, \quad \boldsymbol{e}_1 = \begin{pmatrix} 0 \\ 0.5 \\ 0.7 \\ 0.5 \\ 0 \end{pmatrix}, \quad \boldsymbol{e}_2 = \begin{pmatrix} 0 \\ 0.5 \\ -0.7 \\ 0.5 \\ 0 \end{pmatrix}, \quad \lambda_1 = 0.9, \qquad \boldsymbol{v}_{ss} = \lim_{t\to\infty} c_1(t)\, \boldsymbol{e}_1$$

$$\lim_{t\to\infty} c_1(t) = \frac{\boldsymbol{h}\cdot\boldsymbol{e}_1}{1-\lambda_1} = 10\, \boldsymbol{h}\cdot\boldsymbol{e}_1 = 10\,(0.3+0.14+0.3) = 7.4$$

$$\lim_{t\to\infty} c_2(t) = \frac{\boldsymbol{h}\cdot\boldsymbol{e}_2}{1-\lambda_2} = 1.1\, \boldsymbol{h}\cdot\boldsymbol{e}_2 = 1.1\,(0.3-0.14+0.3) = 0.51$$

Example, ctd

Changing to eigenvector coordinates, we see that most of the growth occurs in $c_1(t)$, that is, along the the dominant eigenvector $\boldsymbol{e}_1$:

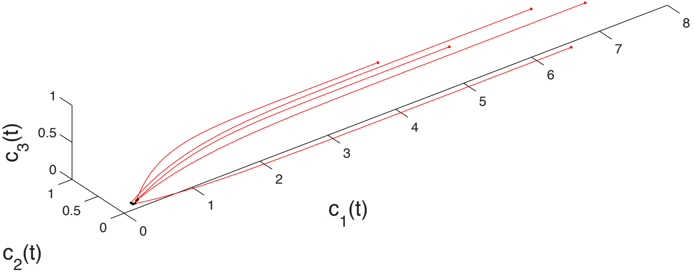

$$\boldsymbol{e}_1 = \begin{pmatrix} 0 \\ 0.5 \\ 0.7 \\ 0.5 \\ 0 \end{pmatrix}, \quad \lambda_1 = 0.9, \qquad \boldsymbol{e}_2 = \begin{pmatrix} 0 \\ 0.5 \\ -0.7 \\ 0.5 \\ 0 \end{pmatrix}, \quad \lambda_2 = 0.1, \qquad \boldsymbol{e}_3 = \begin{pmatrix} 0 \\ 0.7 \\ 0 \\ -0.7 \\ 0 \end{pmatrix}, \quad \lambda_3 = 0.1, \ldots$$

Case $1 < \lambda_\mu$

For eigenvectors with eigenvalues $1 < \lambda_\mu$, the exponents $\frac{\lambda_\mu - 1}{\tau}$ are positive and the exponential terms

$$\exp\left(-\frac{1-\lambda_\mu}{\tau}\,t\right)$$

grow without bound, so that the recurrent dynamics is unstable.

The larger λ_μ, the faster this exponential growth, so that the activity vector will eventually be dominated by the dominant eigenvector:

$$\boldsymbol{v}(t) = \sum_{\mu=1}^{N} c_\mu(t)\,\boldsymbol{e}_\mu \approx c_\nu(t)\,\boldsymbol{e}_\nu, \qquad \lambda_\nu = \max_\mu \lambda_\mu$$

Summary

- The recurrent dynamics amplifies input components aligned with eigenvectors of the connectivity matrix.

$$\boldsymbol{h} \cdot \boldsymbol{e}_\mu$$

- The amplification factor $\frac{1}{1-\lambda_\mu}$ depends on the eigenvalue associated with each eigenvector.
- When $0 < \lambda_u < 1$, amplification is finite and the network is stable, i.e., converges to a steady-state.
- When $\lambda_u > 1$ for one or more μ, amplification is infinite and the network is unstable.

- The case $\lambda_\mu = 1$ is special and will be discussed separately.

3. Capabilities

⋆

linear

⋆

recurrent networks

We consider linear networks with symmetric connectivity $\boldsymbol{M}$

$$\tau \frac{d\boldsymbol{v}}{dt} = -\boldsymbol{v} + \boldsymbol{h} + \boldsymbol{M} \cdot \boldsymbol{v},$$

with eigenvectors and eigenvalues

$$\boldsymbol{M} \cdot \boldsymbol{e}_\mu = \lambda_\mu \, \boldsymbol{e}_\mu,$$

and their expanded solution

$$\boldsymbol{v}(t) = \sum_{\mu=1}^{N} c_\mu(t) \, \boldsymbol{e}_\mu, \qquad \boldsymbol{e}_\mu \cdot \boldsymbol{e}_\nu = \delta_{\mu\nu}$$

We have shown above that the coefficients satisfy

$$\tau \frac{dc(t)}{dt} = -(1-\lambda) \, c(t) + \boldsymbol{h} \cdot \boldsymbol{e}, \quad c(0) = \boldsymbol{v}_0 \cdot \boldsymbol{e}, \quad c(\infty) = \frac{\boldsymbol{h} \cdot \boldsymbol{e}}{1-\lambda}$$

Selective amplification

Suppose that one eigenvalue (which we call λ_1) of the recurrent weight matrix is close to unity, and all others are close to zero. Then the eigenvector expansion will be dominated by the first eigenvector:

$$\boldsymbol{v}(t) = \sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu \approx c_1(t)\, \boldsymbol{e}_1$$

and for times $t >> \tau$, the steady-state activity will be

$$\boldsymbol{v}_{ss} \approx \frac{\boldsymbol{h} \cdot \boldsymbol{e}_1}{1 - \lambda}\, \boldsymbol{e}_1$$

The steady state response is dominated by the projection of the input onto the dominant eigenvector. This projection is amplified by a (potentially large) factor $\frac{1}{1-\lambda_1}$.

$$\boldsymbol{v}_{ss} \approx \frac{\boldsymbol{h} \cdot \boldsymbol{e}_1}{1 - \lambda} \boldsymbol{e}_1$$

Input vector $\boldsymbol{h}$, eigenvector $\boldsymbol{e}_1$, normal vector $\boldsymbol{o}$, length of dot product $(\boldsymbol{h} \cdot \boldsymbol{e}_1)$, vector of this length in direction of eigenvector $(\boldsymbol{h} \cdot \boldsymbol{e}_1)\boldsymbol{e}_1$:

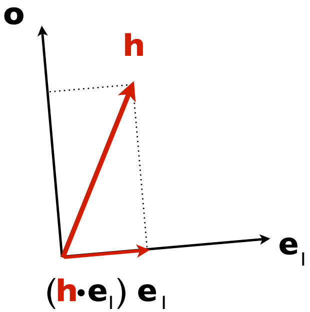

Example

Matrix, eigenvalues, and development in eigenvector coordinates:

$$M = \begin{bmatrix} 0.63 & 0.27 & 0 & 0 & 0 \\ 0.53 & 0.37 & 0 & 0 & 0 \\ 0 & 0 & 0.1 & 0 & 0 \\ 0 & 0 & 0 & 0.1 & 0 \\ 0 & 0 & 0 & 0 & 0.1 \end{bmatrix}, \qquad \lambda_1 = 0.9,\ \lambda_{\mu\neq 1} = 0.1$$

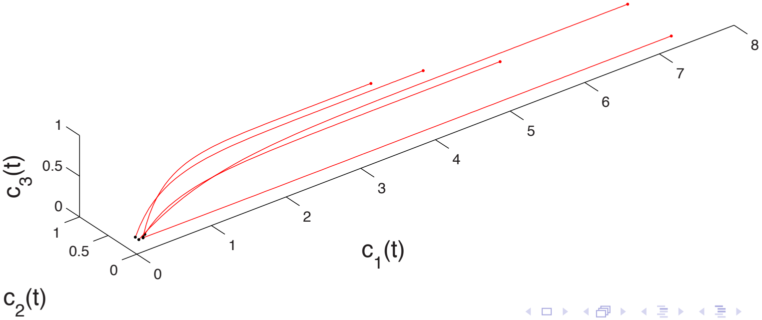

Summary

- When a single eigenvector dominates connectivity, steady-state activity assumes a stereotypical shape (set by the dominant eigenvector):

$$\boldsymbol{v}_{ss} \approx \boldsymbol{e}_1$$

- The steady-state amplitude depends on the projection (dot product) of input and eigenvalue

$$\boldsymbol{v}_{ss} \approx (\boldsymbol{e}_1 \cdot \boldsymbol{h})\, \boldsymbol{e}_1$$

- The amplification of this projection depends on the eigenvalue λ

$$\boldsymbol{v}_{ss} \approx \frac{\boldsymbol{e}_1 \cdot \boldsymbol{h}}{1 - \lambda}\, \boldsymbol{e}_1$$

- One particular shape is *selectively amplified* (to the extent that it is present in the input).

Selective amplification with degenerate eigenvectors

Suppose that two eigenvectors have the same eigenvalues $\lambda_1 = \lambda_2$ close to but less than one. Then the time-dependent activity is

$$\boldsymbol{v}(t) = \sum_{\mu=1}^{N} c_\mu(t)\, \boldsymbol{e}_\mu \approx c_1(t)\, \boldsymbol{e}_1 + c_2(t)\, \boldsymbol{e}_2$$

and the steady-state activity is

$$\boldsymbol{v}_{ss} \approx \frac{(\boldsymbol{h} \cdot \boldsymbol{e}_1)\, \boldsymbol{e}_1 + (\boldsymbol{h} \cdot \boldsymbol{e}_2)\, \boldsymbol{e}_2}{1 - \lambda} = \alpha\, \boldsymbol{e}_1 + \beta\, \boldsymbol{e}_2$$

Thus, the network amplifies the projection of the input onto the plane defined by $\boldsymbol{e}_1$ and $\boldsymbol{e}_2$.

$$\boldsymbol{v}_{ss} \approx \frac{(\boldsymbol{h} \cdot \boldsymbol{e}_1)\, \boldsymbol{e}_1 + (\boldsymbol{h} \cdot \boldsymbol{e}_2)\, \boldsymbol{e}_2}{1 - \lambda} = \alpha\, \boldsymbol{e}_1 + \beta\, \boldsymbol{e}_2$$

Input vector $\boldsymbol{h}$, eigenvectors $\boldsymbol{e}_{1,2}$, normal vector $\boldsymbol{o}$, length of dot product $(\boldsymbol{h} \cdot \boldsymbol{e}_{1,2})$, vectors aligned with eigenvectors $(\boldsymbol{h} \cdot \boldsymbol{e}_{1,2})\boldsymbol{e}_{1,2}$:

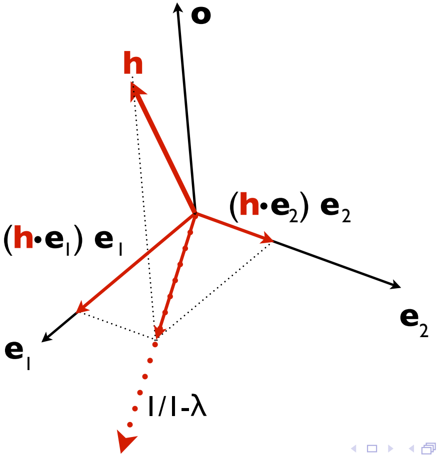

Connectivity with degenerate eigenvectors

Recurrent network with $N = 20$ units. Left: 20×20 connectivity matrix (heat-map). Right: 20 eigenvalues, with two degenerate eigenvalues just below unity.

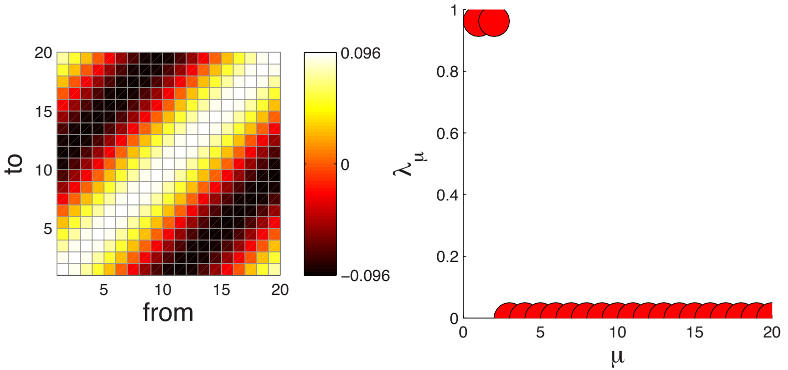

Amplification of different input vectors h:

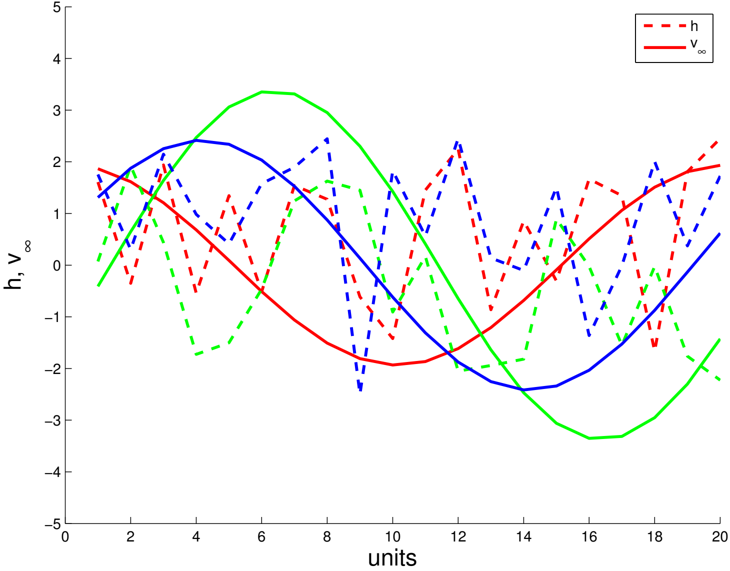

Summary

- When two degenerate eigenvectors dominate connectivity, steady-state activity conforms to a family of stereotypical shape (linear combinations of the dominant eigenvectors):

$$\boldsymbol{v}_{ss} \propto \; \alpha \, \boldsymbol{e}_1 + \beta \, \boldsymbol{e}_2$$

- The projection of the input onto dominant eigenvectors selects the *shape* of the steady-state:

$$\boldsymbol{v}_{ss} \propto \; (\boldsymbol{e}_1 \cdot \boldsymbol{h}) \; \boldsymbol{e}_1 + (\boldsymbol{e}_2 \cdot \boldsymbol{h}) \; \boldsymbol{e}_2$$

- The degenerate eigenvalue λ determines the amplification (scale) of the steady-state:

$$\boldsymbol{v}_{ss} \simeq \frac{(\boldsymbol{e}_1 \cdot \boldsymbol{h}) \; \boldsymbol{e}_1 + (\boldsymbol{e}_2 \cdot \boldsymbol{h}) \; \boldsymbol{e}_2}{1 - \lambda}$$

- One particular member of a family of shapes is *selectively amplified* (to the extent that it is present in the input).

Input integration (case $\lambda_\mu = 1$)(optional)

Suppose a recurrent weight matrix has one eigenvalue $\lambda_1 = 1$ and all other eigenvalues $\lambda_{\mu\neq 1} < 1$. In this case, the general dynamic equation for coefficients

$$\tau \frac{dc_\mu(t)}{dt} = -(1-\lambda_\mu)\, c_\mu(t) + \boldsymbol{h} \cdot \boldsymbol{e}_\mu$$

simplifies for $\mu = 1$ to

$$\tau \frac{dc_1(t)}{dt} = \boldsymbol{h} \cdot \boldsymbol{e}_1$$

or, for a time-dependent input $\boldsymbol{h}(t)$

$$c_1(t) = c_1(0) + \frac{1}{\tau} \int_0^t \boldsymbol{e}_1 \cdot \boldsymbol{h}(t')\, dt'$$

Such a network integrates phasic inputs over time and maintains a memory of the integral value:

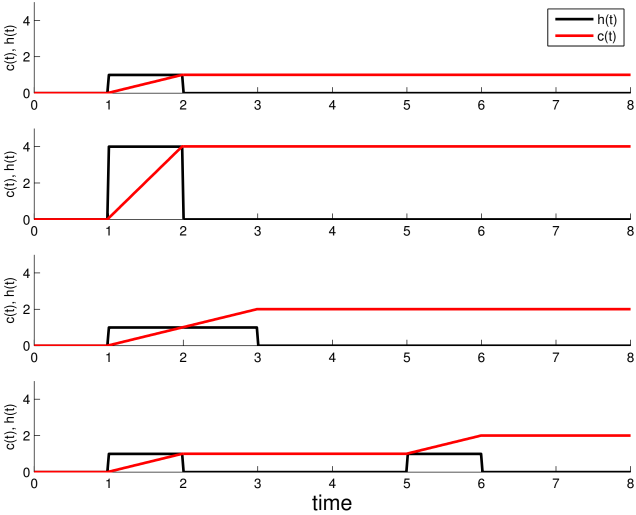

Fig. 7.7 of D&A: Cartoon of burst and integrator neurons involved in maintaining horizontal eye position.

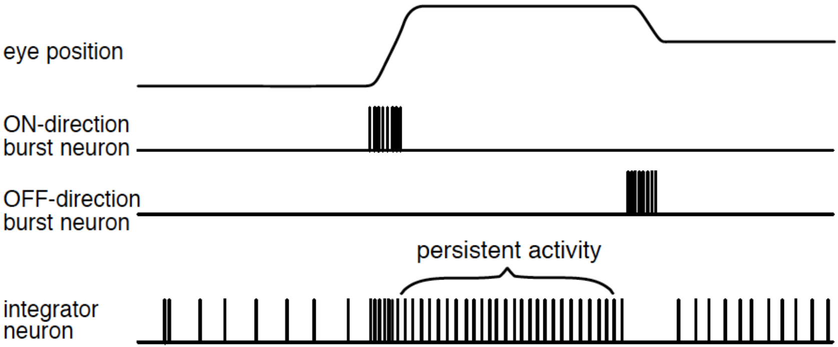

Points to note

- When a single eigenvector dominates connectivity, steady-state activity amplifies one particular shape (the dominant eigenvector), if it is contained in the input.
- With two degenerate eigenvectors, steady-state activity amplifies one particular member of a family of shapes (linear combinations of dominant eigenvectors), which best matches the input.
- When the dominant eigenvalue is unity, the network integrates phasic inputs over time and maintains a memory of the integral value.

4.Capabilities of

⋆

nonlinear

⋆

recurrent networks

Linear model do not adequately describe networks of biological neurons. We now consider some nonlinear networks in simulation, gaining biological realism but sacrificing analytical tractability. Similar to linear networks, the behaviour of nonlinear networks is also governed by the connectivity matrix $\boldsymbol{M}$ (and its eigenvector and eigenvalues).

The nonlinearity we consider is a simple rectification with threshold

$$\tau \frac{d\boldsymbol{v}}{dt} = -\boldsymbol{v} + \boldsymbol{F}\left(\boldsymbol{h} + \boldsymbol{M} \cdot \boldsymbol{v}\right) = -\boldsymbol{v} + \left[\boldsymbol{h} + \boldsymbol{M} \cdot \boldsymbol{v} - \gamma\right]_+$$

where γ is a threshold vector (often chosen to be zero).

Selective amplification

Recurrent network with $N = 20$ units. Left: 20×20 connectivity matrix (heat-map). Right: 20 eigenvalues, with two degenerate dominant eigenvalues.

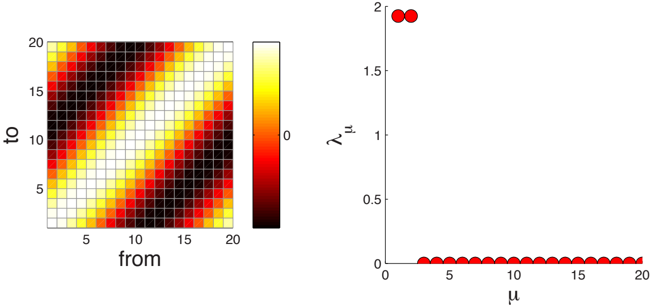

Note: stability of non-linear network does not depend simply on size of eigenvalues.

Amplification of example input (four time steps):

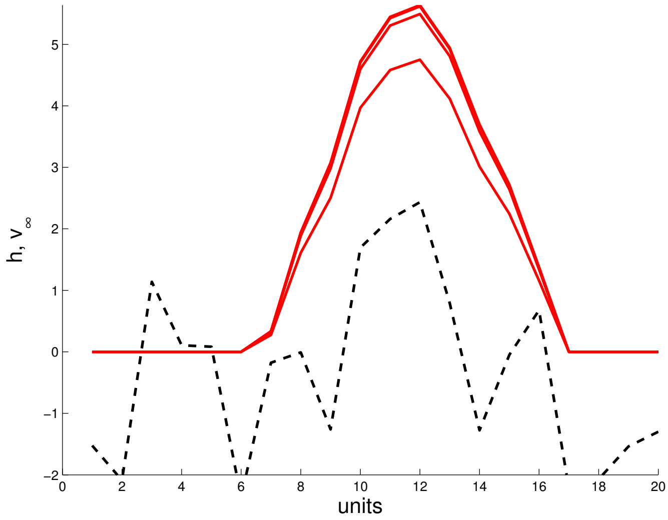

Amplification of example input (four time steps):

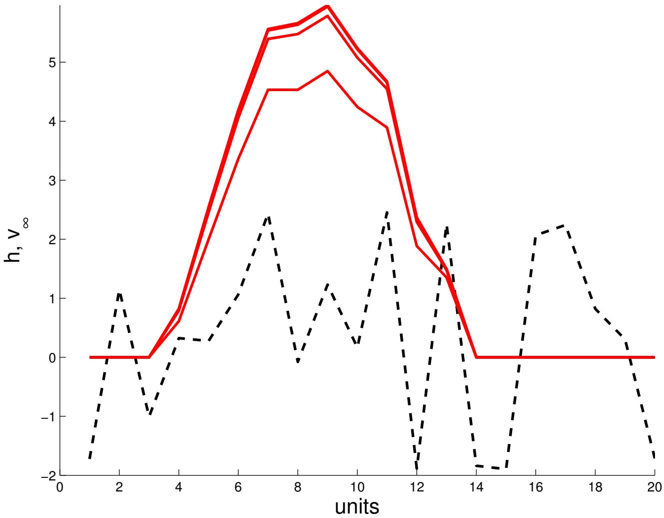

Summary selective amplification

- When two degenerate eigenvector dominate connectivity, steady-state activity assumes a stereotypical shape
- This shape is the rectified version of a linear combination of the dominant eigenvectors.
- The steady-state amplitude depends on the projection (dot product) of input onto the plane of eigenvectors
- The network selectively amplifies the linear combination that best matches the input.

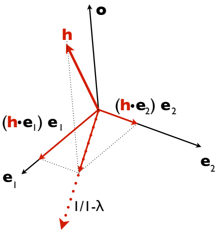

Winner-take all

The same network can perform a winner-take all computation, amplifying the largest of several inputs:

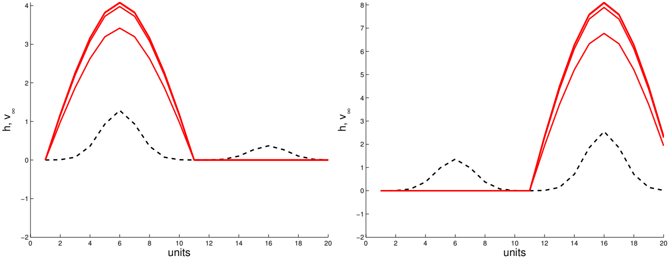

Summary winner-take-all

- The network selectively amplifies the linear combination that best matches the input.
- Due to the non-linearity, this does not reflect all inputs components proportionately.
- Larger input components are amplified disproportionately.
- This results in a *winner-take-all* functionality.

Multiplicative gain modulation

An additive offset to the input translates into a multiplicative scaling of the output. This is because the shape of the response profile is fixed by the connectivity, so that additional input cannot broaden the base of the output.

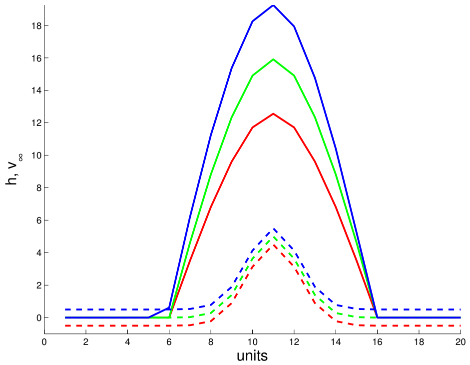

Input: dashed curves. Steady-states: solid curves.

Summary gain modulation

- Gain modulation is the multiplicative scaling of a population response by an additional input.
- It is widely observed in sensory populations (e.g., transformation of EYE to HEAD coordinates).
- In a recurrent network, an additive input results in a multiplicative scaling.
- The multiplication is performed by the network, not by individual neurons.

Sustained activity (optional)

When a selective input is replaced by a uniform input, the network retains a memory of the earlier input. Here, input is shown on the left, network response on the right. The earlier input/response are shown in black, the subsequent input/response in red.

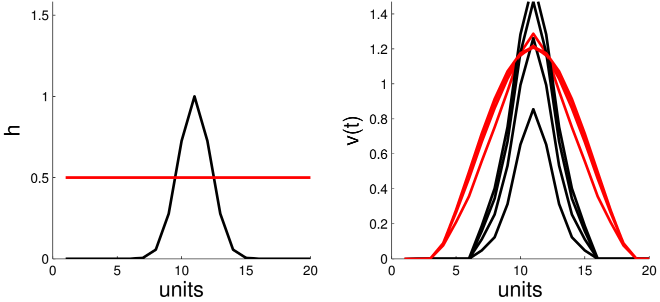

Summary sustained activity

- In the special case of an eigenvalue near unity, the network integrates input over time and retains a memory of the integral value.

- The pattern of population activity is set by the dominant eigenvector.

- The amplitude of population activity represents a time-integral of phasic inputs (inasmuch as they align with the eigenvector).

Overall summary

- **Why is the qualitative behaviour of recurrent networks – both linear and non-linear – shaped by eigenvectors of the connectivity matrix?**
- **Recurrent networks retain activity patterns that propagate (pass unchanged) through their connectivity matrix.**
- **By definition, the eigenvectors of the connectivity matrix are such patterns.**

$$\lambda_i \, \boldsymbol{e}_i = \boldsymbol{M} \cdot \boldsymbol{e}_i$$

- **The 'dominant' eigenvalues are comparable to each other, but larger than all others.**
- **The recurrent dynamics is stable when *all* eigenvalues are smaller than unity.**

- **When dominant eigenvalues are just below unity, the network performs selective amplification $\times \frac{1}{1-\lambda}$.**
- **Specifically, it amplifies the linear combination of degenerate eigenvectors that best matches the input (projection of input onto degenerate plane).**
- **Today we considered examples with one or two dominant eigenvectors.**

- **For a linear recurrent network, the dynamics can be obtained analytically in terms of an eigenvalue expansion.**
- **For non-linear recurrent networks, the qualitative behaviour remains unchanged.**

- **Recurrent networks can perform 'selective amplification', 'winner-take-all', and 'multiplicative-gain-modulation'.**

Next: Associative memory