



Networks of Neurons (2 of 2)


Angus Chadwick

School of Informatics, University of Edinburgh, UK

Computational Neuroscience (Lecture 12, 2023/2024)

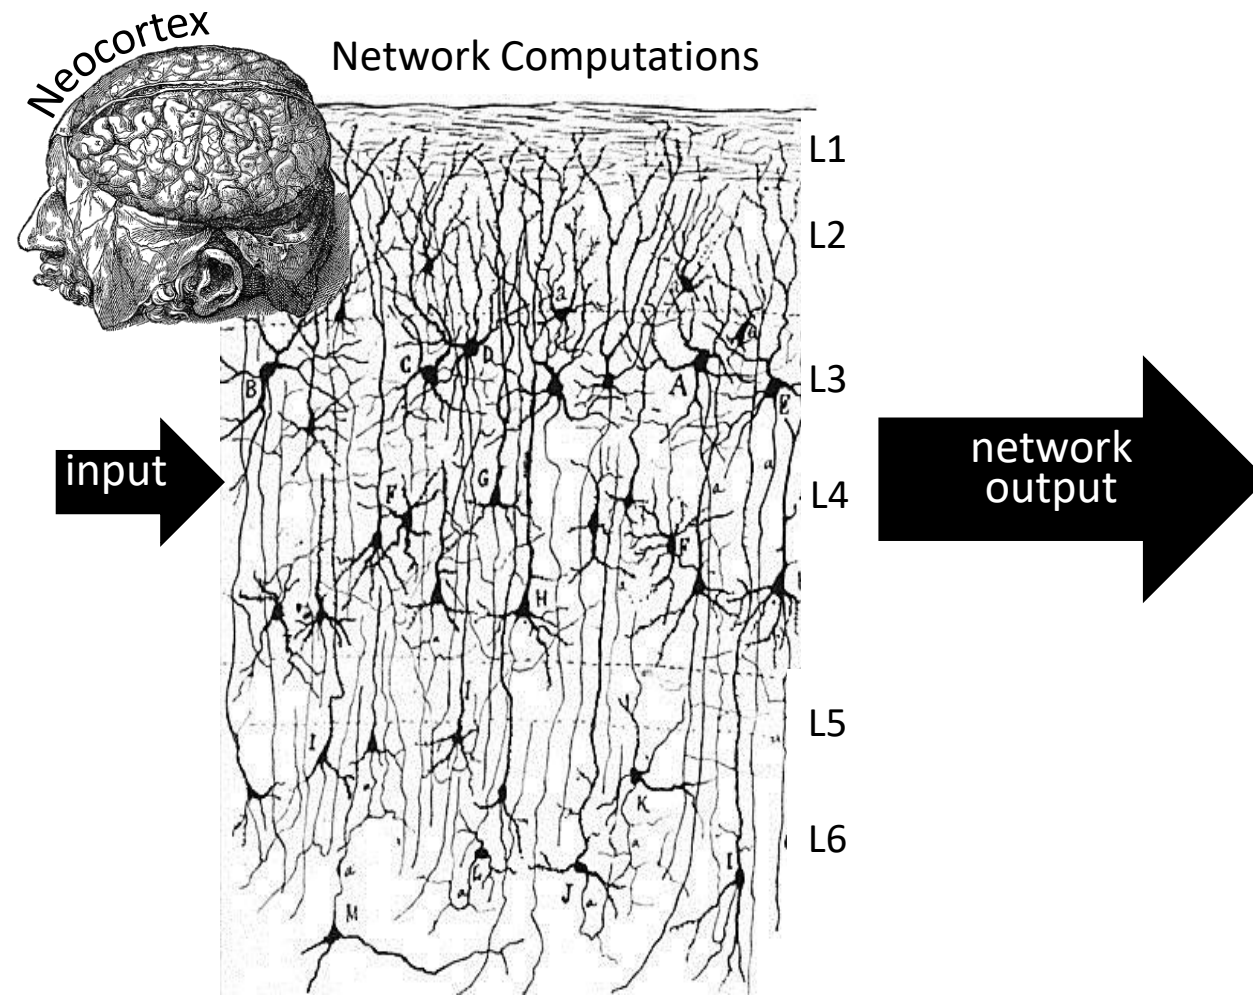
Outline of Lecture

- Cortical networks
- Models for orientation tuning (the ring network)
- State space perspective
- Linear Stability Analysis
- Stability of E-I networks


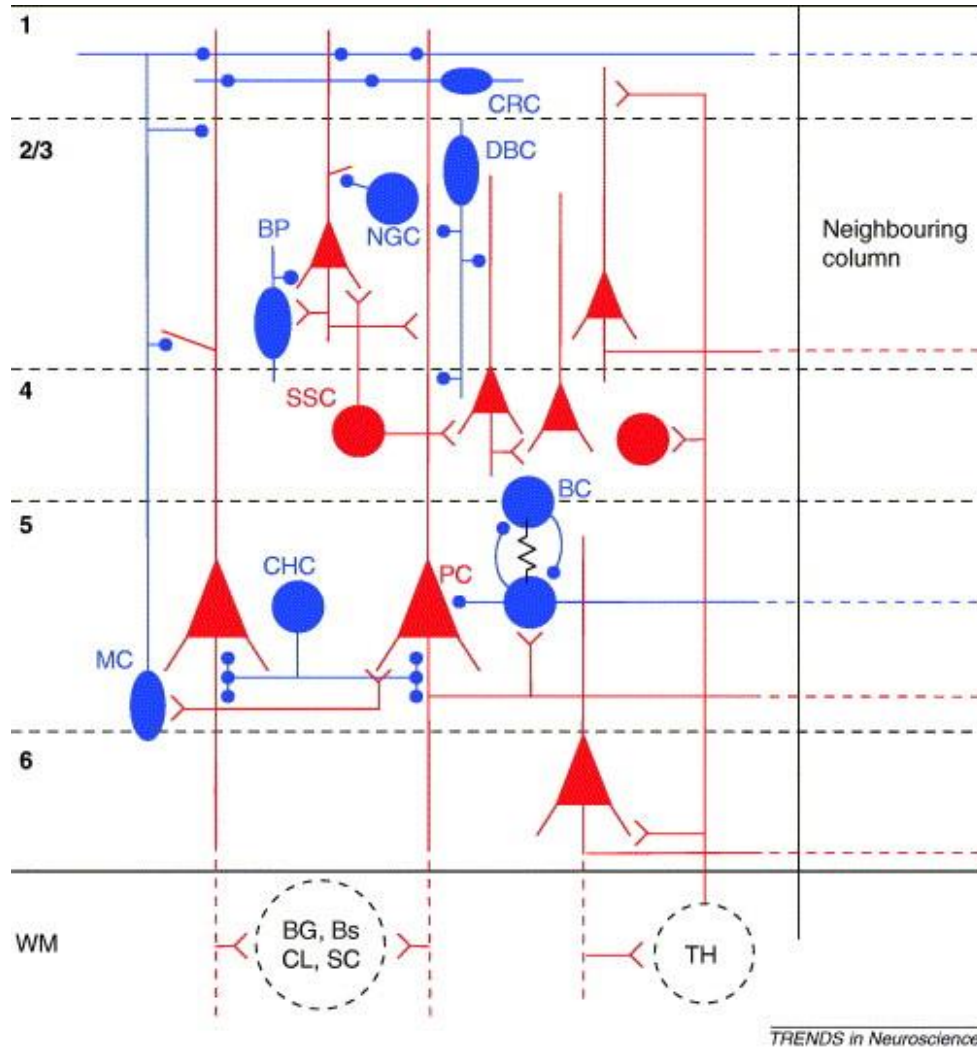


Drawings of cortical layers by Ramon y Cajal (Golgi [left, middle] and Nissl staining methods)

Cortical Computations



Is There a “Canonical” Cortical Computation?





Random vs Structured Networks

- Connections in cortex look fairly random, but also exhibit structure
- For example, connections depend on cell type (excitatory vs inhibitory) and on stimulus preference
- How do networks with random vs structured connections behave?
- Random networks are already interesting: chaos, E-I balance, reservoir computing, etc.
- Structured connectivity causes the network to respond in specific ways to specific input patterns, which can be the basis for useful computations


Connections in Visual Cortex

- Connections from a local region of tree shrew V1 (black dots)
- Locally unspecific (random?), but long range connections are selective to orientation preference



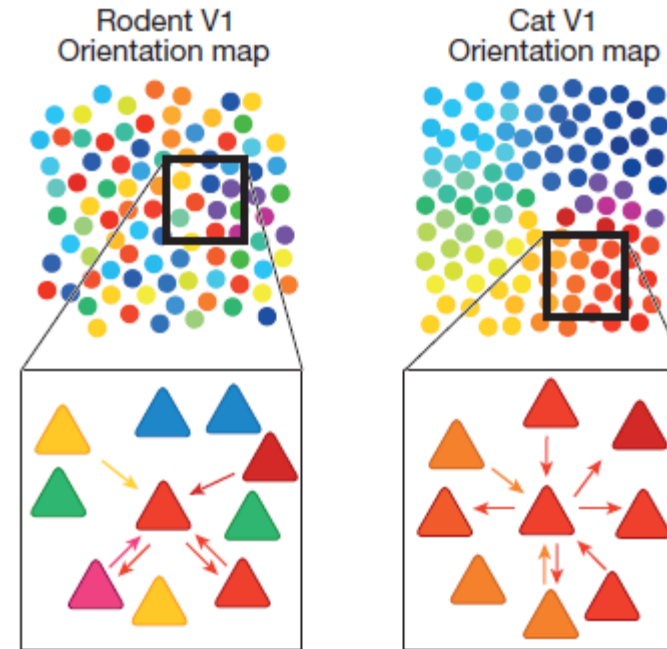
Connections in Visual Cortex

- Neurons with similar preferred orientations have stronger connections.



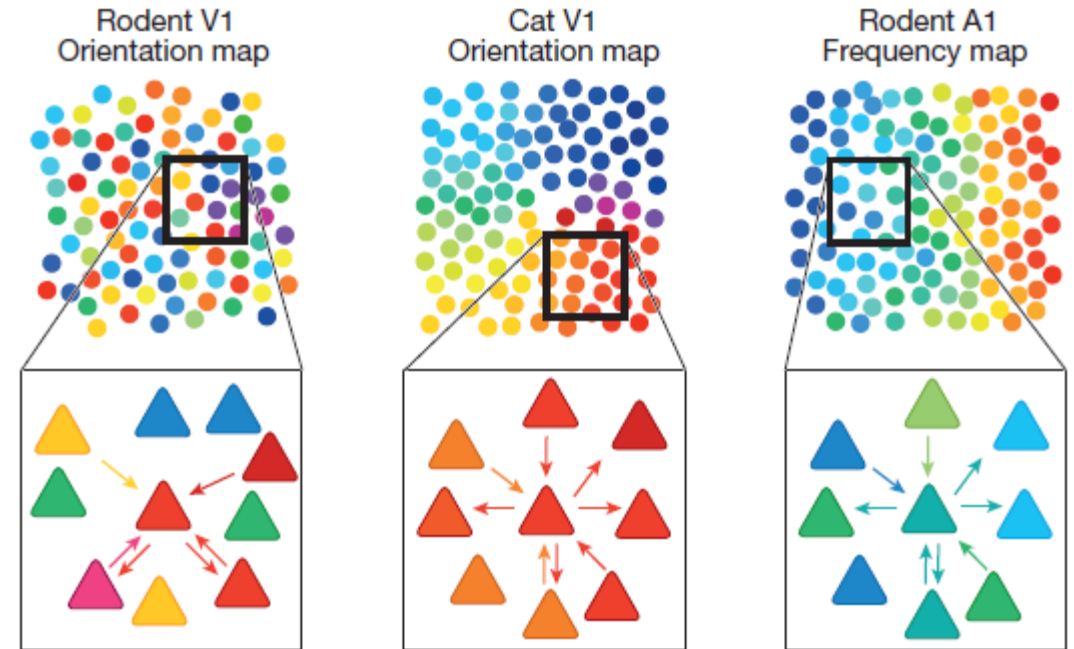
Connections in Visual Cortex

- Neurons with similar preferred orientations have stronger connections.
- This seems to be true for both topographic and salt and pepper maps



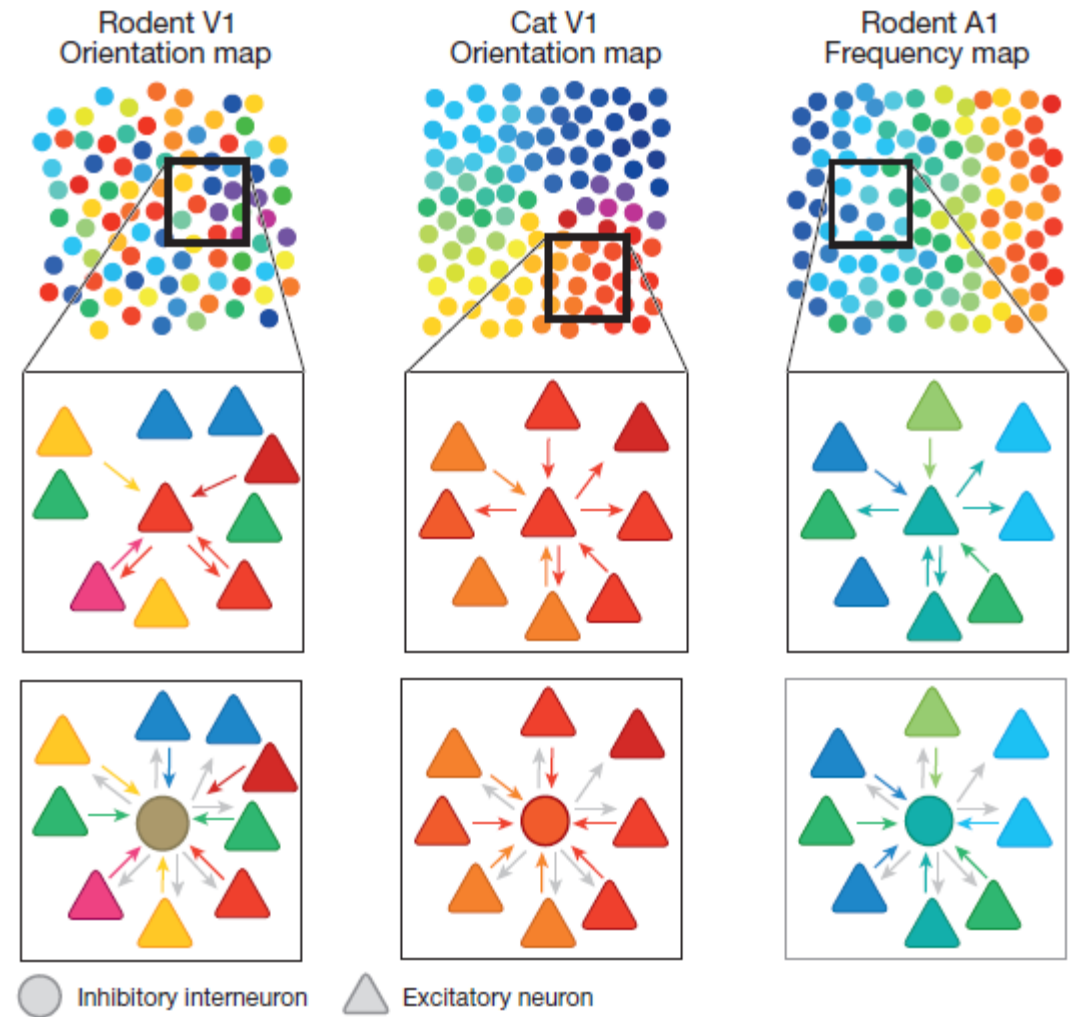
Connections in Visual Cortex

- Neurons with similar preferred orientations have stronger connections.
- This seems to be true for both topographic and salt and pepper maps



Connections in Visual Cortex

- Neurons with similar preferred orientations have stronger connections.
- This seems to be true for both topographic and salt and pepper maps
- Inhibitory neurons seem to have much broader/less selective connections in local circuit (but some debate)
- How do these local V1 connections influence orientation tuning?

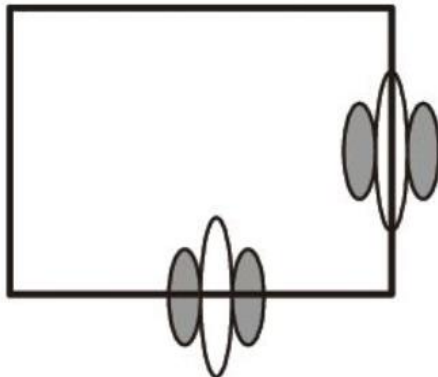



Feedforward Models for Orientation Tuning

- Hubel and Wiesel proposed an elegant conceptual model for the emergence of orientation tuning: summation of spatially shifted on-off receptive fields
- A **weighted sum of thalamic (LGN) feedforward inputs** can produce an elongated receptive field which responds selectively to stimulus orientation

Orientation selectivity


No stimulus in receptive field, no response



Non-preferred stimulus, no response

Preferred stimulus, large response

Simple cell sums LGN inputs




Recurrent Models for Orientation Tuning: The Ring Network

- An alternative hypothesis: connectivity between V1 neurons **selectively amplifies** feedforward inputs that are only **weakly tuned** to orientation
- Continuous network of V1 neurons labelled by preferred orientation (on a **ring**)

$$\tau \frac{dr(\theta)}{dt} = -r(\theta) + \left[\int_{-\pi/2}^{\pi/2} W(\theta, \theta') r(\theta') d\theta' + u(\theta - \theta_s) \right]_+$$

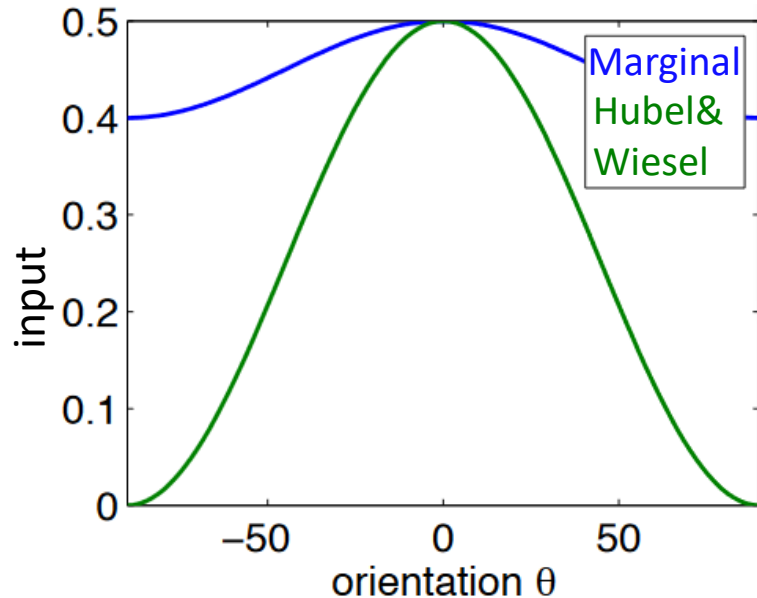

$$W(\theta, \theta') = A \cos(2(\theta - \theta')) + B$$

$$u(\theta - \theta_s) = C \cos(2(\theta - \theta_s)) + D$$



- Recurrent connectivity W depends only on **difference in preferred orientations**
- Inputs u depend only on **difference between preferred orientation and stimulus orientation**
- The model has *rotational symmetry* (this enables us to solve it analytically)

The Ben-Yishai Recurrent Ring Model: Two Regimes



- Varying inputs and connectivity yields either:
 - **Hubel and Wiesel regime** (strongly tuned inputs and weak/no recurrent connections)
 - a novel **“marginal” regime** (strongly tuned recurrent weights and weakly tuned inputs)

$$W(\theta, \theta') = A \cos(2(\theta - \theta')) + B$$

$$u(\theta - \theta_s) = C \cos(2(\theta - \theta_s)) + D$$

Tuning Curves in the Two Regimes

- Both regimes generate orientation tuning curves
- In the Hubel and Wiesel regime, responses are determined only by feedforward input and threshold nonlinearity
- In the marginal regime:
 - shape of tuning curves becomes **independent of the shape of inputs**
 - weakly tuned inputs are **selectively amplified** into strongly tuned outputs by recurrent connectivity
 - a tuned response bump can spontaneously form even when feedforward inputs have no tuning at all! (called “spontaneous symmetry breaking” in physics)

Dynamics of the Ben-Yishai Model – Stimulus Rotation

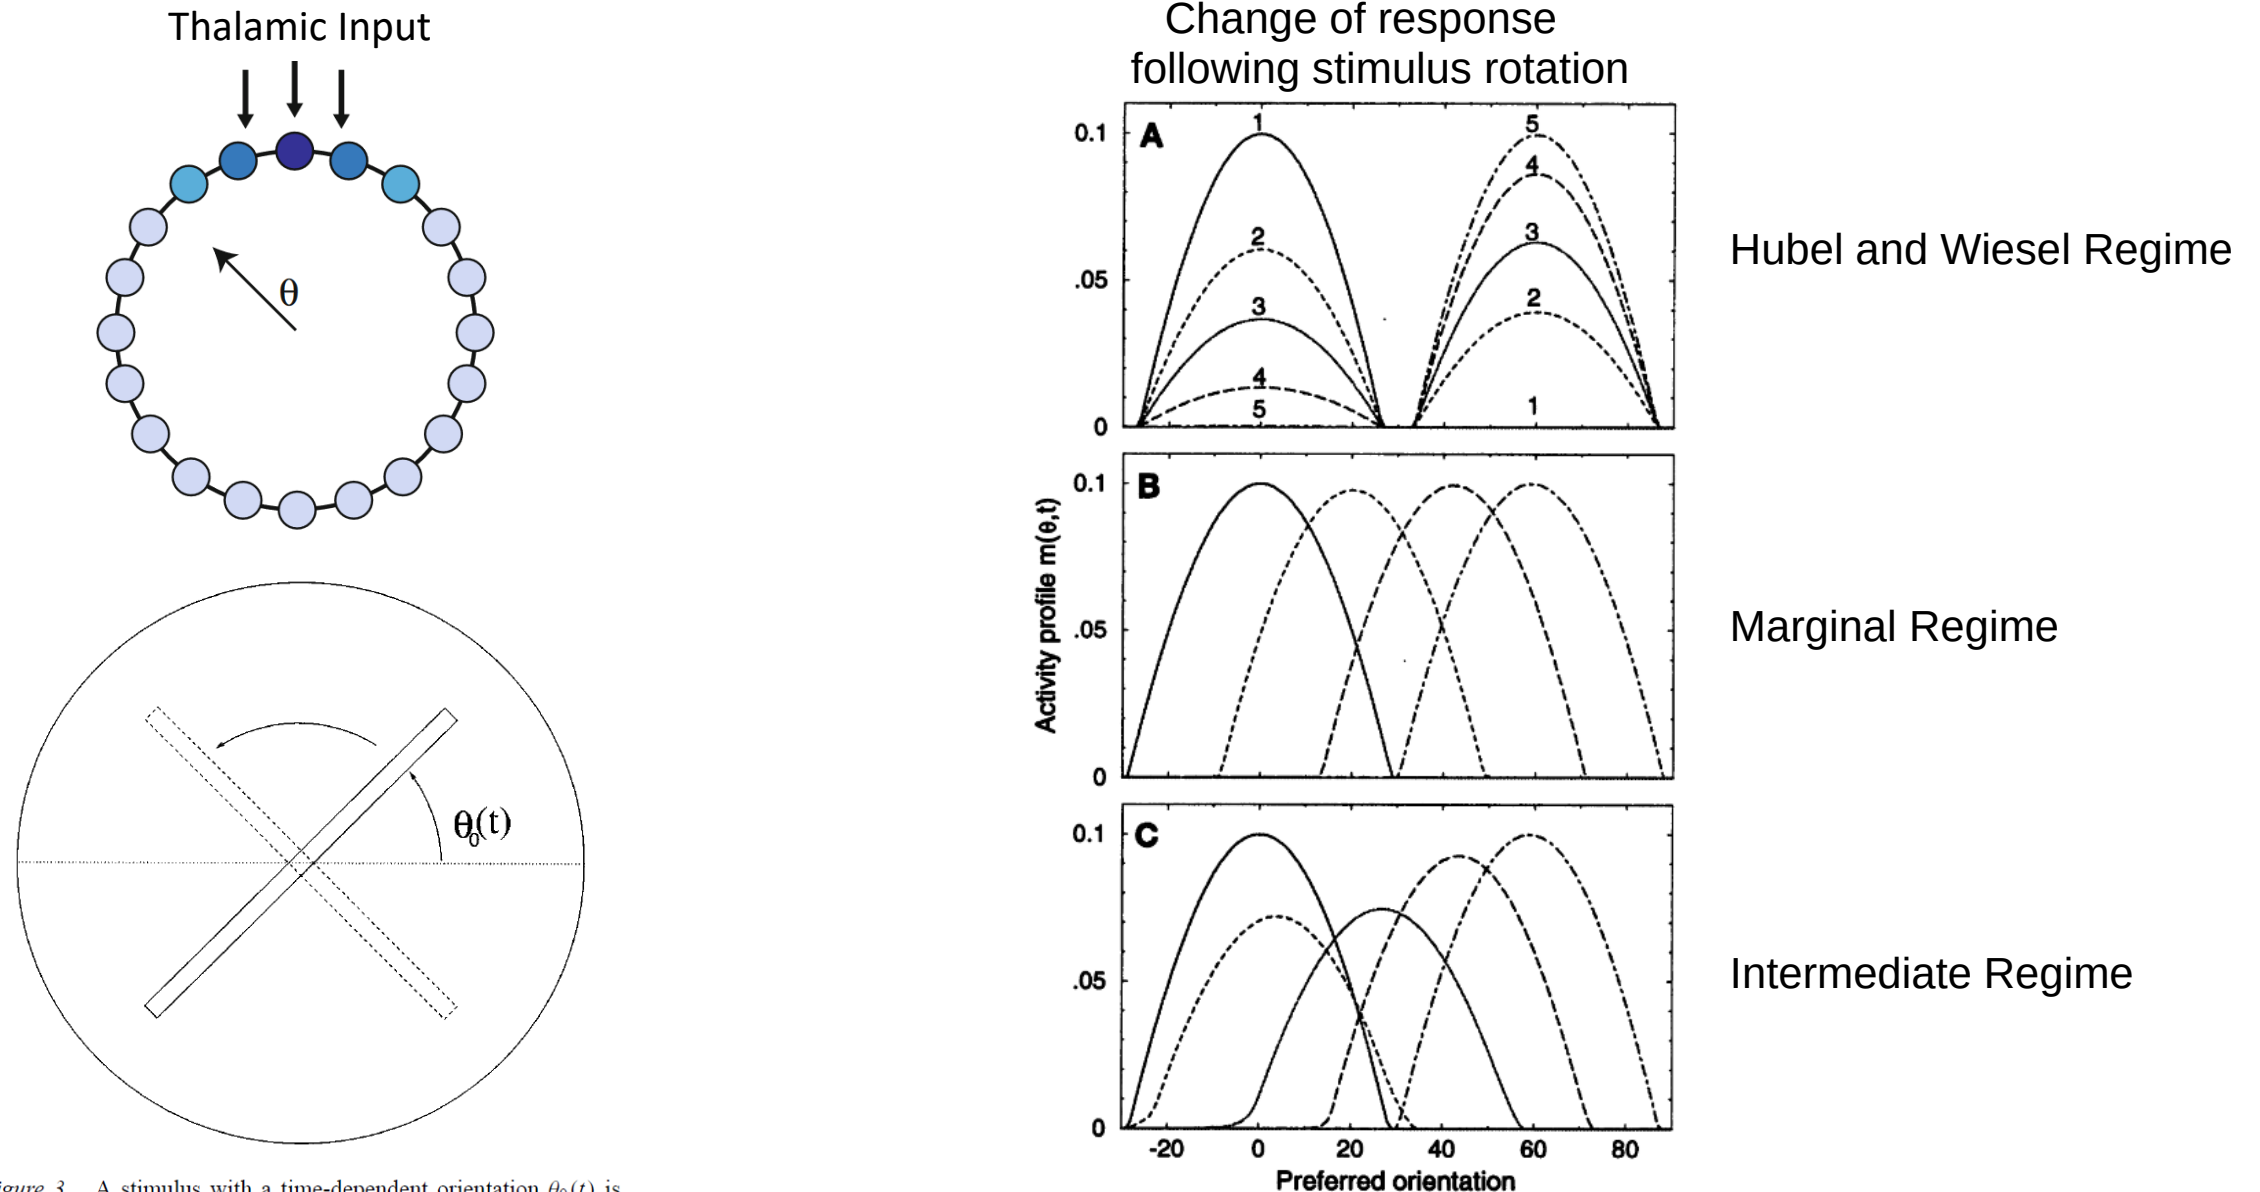


Figure 3. A stimulus with a time-dependent orientation $\theta_0(t)$ is presented in a common receptive field.

Contrast-Invariant Orientation Tuning

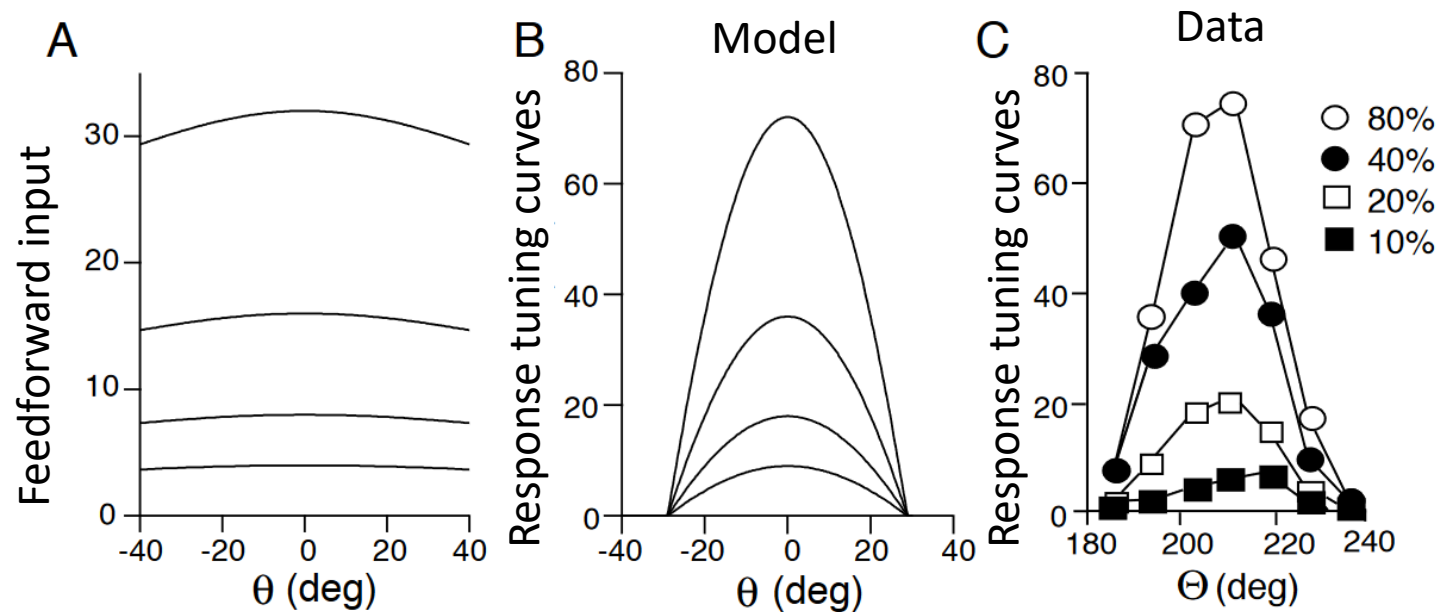



Figure 7.10: The effect of contrast on orientation tuning. A) The feedforward input as a function of preferred orientation. The four curves, from top to bottom, correspond to contrasts of 80%, 40%, 20%, and 10%. B) The output firing rates in response to different levels of contrast as a function of orientation preference. These are also the response tuning curves of a single neuron with preferred orientation zero. As in A, the four curves, from top to bottom, correspond to contrasts of 80%, 40%, 20%, and 10%. The recurrent model had $\lambda_0 = 7.3$, $\lambda_1 = 11$, $A = 40$ Hz, and $\epsilon = 0.1$. C) Tuning curves measure experimentally at four contrast levels as indicated in the legend. (C adapted from Sompolinsky and Shapley, 1997; based on data from Sclar and Freeman, 1982.)

Which Regime is V1 in?

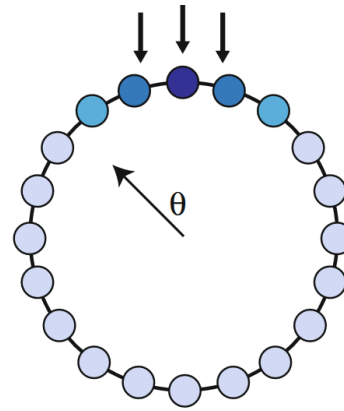
- There has been much debate about the contribution of feedforward and recurrent mechanisms to orientation tuning
- We now know that feedforward input has roughly the same degree of orientation tuning as recurrent input (ruling out a full-on marginal regime) [*Lien and Scanziani, 2013*]
- But recurrent connections do amplify tuning curves, and may be important for many things, such as: contrast-invariant tuning curves, attentional modulation, interactions between simultaneously presented stimuli, etc.
- Spontaneous bump formation is more relevant to theories of working memory and navigation than to V1 orientation tuning

Ring Networks Beyond the Visual Cortex

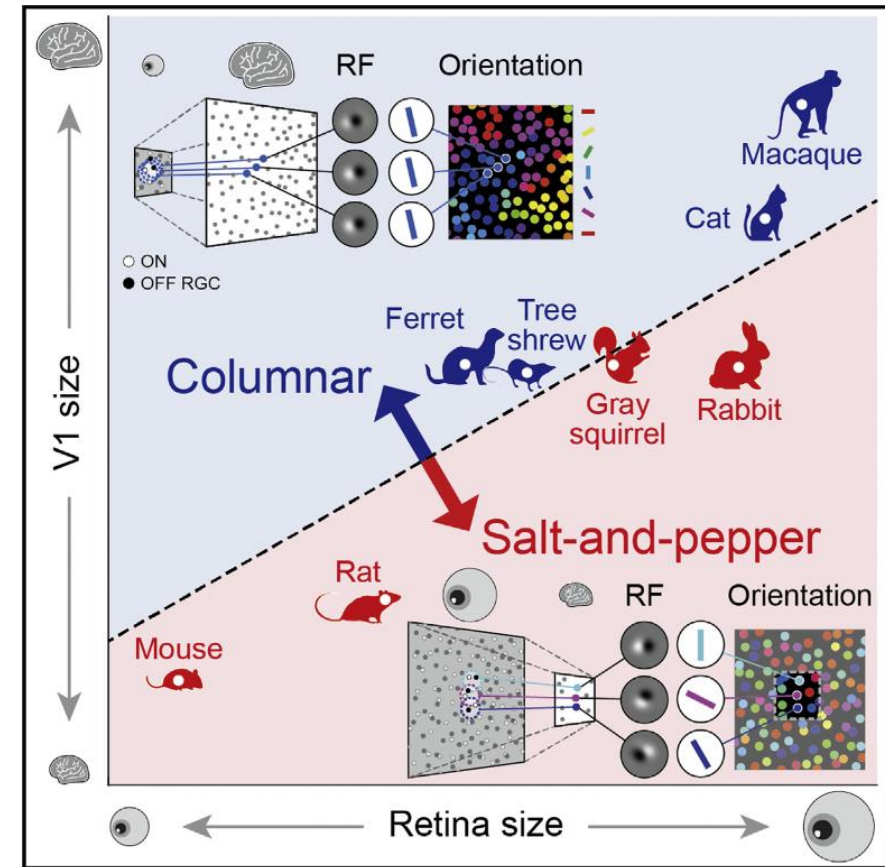

Mouse V1



A Conceptual Model for Rodent V1
(functionally arranged on a ring)



Neurons are *Physically* Arranged on a Ring
in the Fly Brain!



NEUROSCIENCE

Ring attractor dynamics in the *Drosophila* central brain

Sung Soo Kim,^{*} Hervé Rouault,^{*} Shaul Druckmann,[†] Vivek Jayaraman[†]

Ring attractors are a class of recurrent networks hypothesized to underlie the representation of heading direction. Such network structures, schematized as a ring of neurons whose connectivity depends on their heading preferences, can sustain a bump-like activity pattern whose location can be updated by continuous shifts along either turn direction. We recently reported that a population of fly neurons represents the animal's heading via bump-like activity dynamics. We combined two-photon calcium imaging in head-fixed flying flies with optogenetics to overwrite the existing population representation with an artificial one, which was then maintained by the circuit with naturalistic dynamics. A network with local excitation and global inhibition enforces this unique and persistent heading representation. Ring attractor networks have long been invoked in theoretical work; our study provides physiological evidence of their existence and functional architecture.

A Ring in the Fly Brain

Schematic of fly brain ring






Image of real fly brain ring

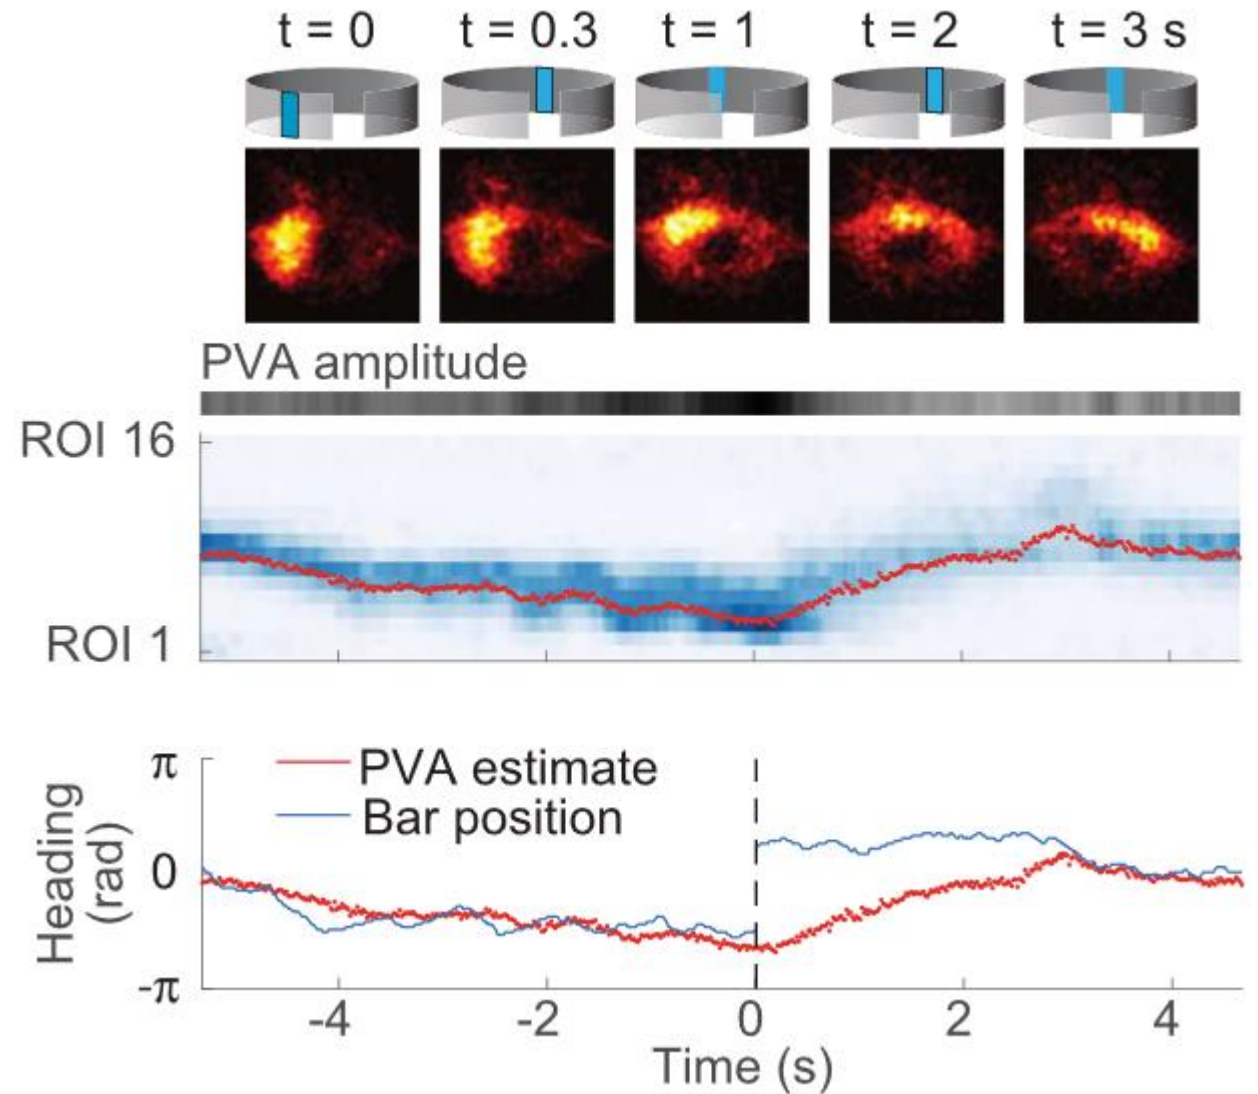


A Ring in the Fly Brain

Fly in virtual reality



Activity of fly ring in virtual reality



Summary: The Ring Network

- The ring network assumes that connectivity depends only on difference in preferred orientation (angle around ring)
- This is consistent with connectivity in V1:
 - excitatory neurons connect preferentially to neighbours with similar orientation preference to their own
 - inhibitory neurons seem to connect randomly connect to their neighbours (but debated...)
 - this would imply local excitation and long range inhibition around the ring
- The ring network can be in different regimes (Hubel and Wiesel and “marginal”)
- The ring model has since found applications in many domains outside of V1 (e.g., navigation, working memory, etc.)

Dynamics in State Space

It is useful to rewrite the network dynamics in vector notation:

$$\tau_m \frac{dr_i}{dt} = -r_i + \phi \left(\sum_j w_{ij} r_j + I_{ext,i}(t) \right) \quad \tau_m \frac{d\mathbf{r}}{dt} = -\mathbf{r} + \phi (W\mathbf{r} + \mathbf{I}_{ext}(t))$$


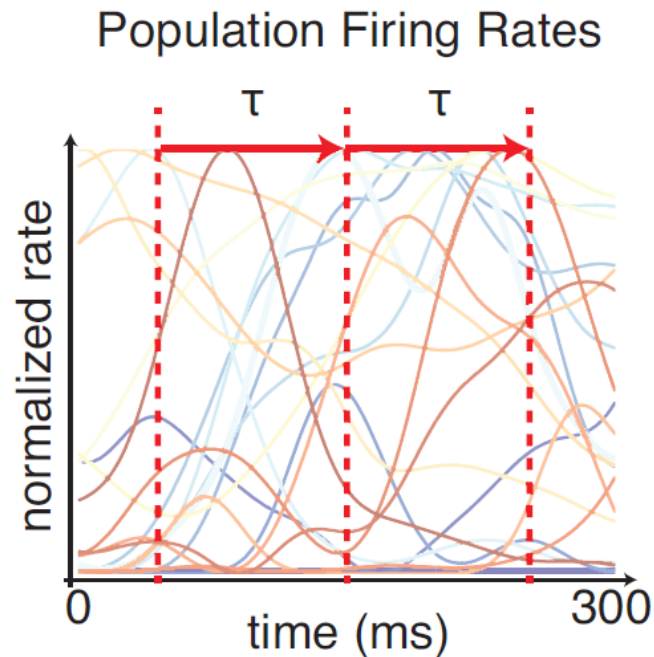
Dynamics in State Space

It is useful to rewrite the network dynamics in vector notation:

$$\tau_m \frac{dr_i}{dt} = -r_i + \phi \left(\sum_j w_{ij} r_j + I_{ext,i}(t) \right) \quad \tau_m \frac{d\mathbf{r}}{dt} = -\mathbf{r} + \phi (W\mathbf{r} + \mathbf{I}_{ext}(t))$$

This allows us to visualise the network dynamics in **state space**:

- A pattern of firing rates \mathbf{r} defines a point in an N -dimensional space.
- Network dynamics generate trajectories $\mathbf{r}(t)$ in N -dimensional space.



Computation Through Dynamics



Available online at www.sciencedirect.com

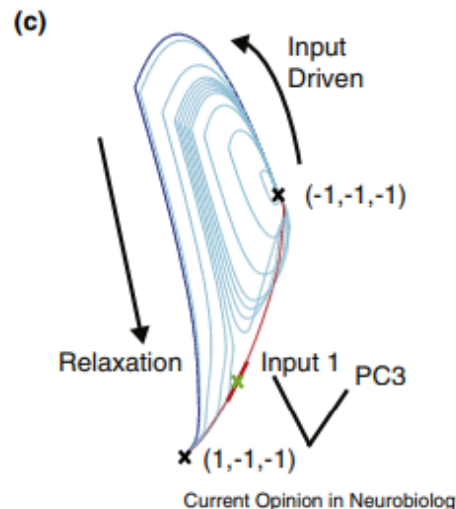


ScienceDirect

Current Opinion in
Neurobiology

Annual Review of Neuroscience

Computation Through Neural
Population Dynamics

Neural circuits as computational dynamical systems



Saurabh Vyas,^{1,3} Matthew D. Golub,^{2,3}
David Sussillo,^{2,3,4} and Krishna V. Shenoy^{1,2,3,5}

Dynamics in state space can be a
substrate for computations.

Figure shows a toy example (3-bit
memory) – the network
remembers the state of three
inputs (+1 or -1), storing them in
8 fixed points in its state space.




Fixed Points

For constant input I_{ext} , the fixed points of the network are defined as:

$$\tau_m \frac{d\mathbf{r}}{dt} = -\mathbf{r} + \phi(W\mathbf{r} + \mathbf{I}_{ext}) \quad \frac{d\mathbf{r}}{dt} = 0 \implies \mathbf{r} = \phi(W\mathbf{r} + \mathbf{I}_{ext})$$

Fixed points form the basis of various computations in recurrent networks.

For example, a fixed point may represent a decision or a memory stored in the network. A line of fixed points may store a continuous valued variable (e.g. the spatial location of an object).



Linear Dynamical Systems

A common choice of transfer function is threshold-linear:

$$\phi(\mathbf{x}) = \beta [\mathbf{x} - \theta]_+ \quad \tau_m \frac{d\mathbf{r}}{dt} = -\mathbf{r} + \beta [W\mathbf{r} + \mathbf{I}_{ext}(t) - \theta]_+$$

Linear Dynamical Systems

A common choice of transfer function is threshold-linear:

$$\phi(\mathbf{x}) = \beta [\mathbf{x} - \theta]_+ \quad \tau_m \frac{d\mathbf{r}}{dt} = -\mathbf{r} + \beta [W\mathbf{r} + \mathbf{I}_{ext}(t) - \theta]_+$$

If all neurons are above the threshold θ , this reduces to a purely linear model:

$$\phi(\mathbf{x}) = \beta(\mathbf{x} - \theta) \quad \tau_m \frac{d\mathbf{r}}{dt} = -\mathbf{r} + \beta (W\mathbf{r} + \mathbf{I}_{ext}(t) - \theta)$$

Linear Dynamical Systems

A common choice of transfer function is threshold-linear:

$$\phi(\mathbf{x}) = \beta [\mathbf{x} - \theta]_+ \quad \tau_m \frac{d\mathbf{r}}{dt} = -\mathbf{r} + \beta [W\mathbf{r} + \mathbf{I}_{ext}(t) - \theta]_+$$

If all neurons are above the threshold θ , this reduces to a purely linear model:

$$\phi(\mathbf{x}) = \beta(\mathbf{x} - \theta) \quad \tau_m \frac{d\mathbf{r}}{dt} = -\mathbf{r} + \beta (W\mathbf{r} + \mathbf{I}_{ext}(t) - \theta)$$

Rearranging terms gives a simpler form:

$$\frac{d\mathbf{r}}{dt} = A\mathbf{r} + \mathbf{u}(t) \quad A = (\beta W - I)/\tau_m$$
$$\mathbf{u}(t) = \beta(\mathbf{I}_{ext}(t) - \theta)/\tau_m$$

Such linear dynamical systems are a popular model due to their analytical tractability.

Linear Dynamical Systems: Fixed Points and Stability

For constant input \mathbf{u} , we can find the fixed points of an LDS analytically:

$$\left. \frac{d\mathbf{r}}{dt} \right|_{\mathbf{r}=\mathbf{r}^*} = 0 \implies \mathbf{r}^* = -A^{-1}\mathbf{u}$$

Linear Dynamical Systems: Fixed Points and Stability

For constant input \mathbf{u} , we can find the fixed points of an LDS analytically:

$$\left. \frac{d\mathbf{r}}{dt} \right|_{\mathbf{r}=\mathbf{r}^*} = 0 \implies \mathbf{r}^* = -A^{-1}\mathbf{u}$$

If the network is initialised at $\mathbf{r}(0)$, activity evolves over time as:

$$\mathbf{r}(t) - \mathbf{r}^* = \sum_i c_i \mathbf{v}_i e^{\lambda_i t}$$

Where \mathbf{V}_i are the eigenvectors of A , λ_i are the eigenvalues, and C_i are constants related to the initial condition $\mathbf{r}(0)$.

Linear Dynamical Systems: Fixed Points and Stability

We can now perform a **stability analysis** on the LDS: a fixed point is *stable* if network activity is attracted towards the fixed point, and *unstable* if activity is repelled away from the fixed point.

We have the solution:
$$\mathbf{r}(t) - \mathbf{r}^* = \sum_i c_i \mathbf{v}_i e^{\lambda_i t}$$


Thus, if all **eigenvalues** of A have negative real part, the fixed point is *stable*; if any eigenvalue has positive real part, the fixed point is *unstable*.

Networks of neurons in the brain ought to be stable. For example, epileptic seizures may arise due to unstable network dynamics.

Linear Dynamical Systems: Fixed Points and Stability

For 2D systems, dynamics around fixed point can be visualised as a phase plane. Stability depends on the trace and determinant of the dynamics matrix A .

Poincaré Diagram: Classification of Phase Portraits in the $(\det A, \text{Tr } A)$ -plane



$$\text{Trace}(A) = \lambda_1 + \lambda_2$$

$$\text{Det}(A) = \lambda_1 \lambda_2$$

When do both eigenvalues have negative real part?


- $\text{Det}(A) < 0$ if each has a different sign (saddle point)
- $\text{Det}(A) > 0$ if both have same sign
- If $\text{Det}(A) > 0$ and $\text{Trace}(A) < 0$, both are negative

Stability of Linear E-I Networks

Consider a linear network comprised of an E population and an I population:

$$A = (\beta W - I) / \tau_m = \begin{bmatrix} (\beta W_{EE} - 1) / \tau_m & -\beta W_{EI} / \tau_m \\ \beta W_{IE} / \tau_m & (-\beta W_{II} - 1) / \tau_m \end{bmatrix}$$

What are the conditions for the network to be stable?



Stability of Linear E-I Networks

Consider a linear network comprised of an E population and an I population:

$$A = (\beta W - I) / \tau_m = \begin{bmatrix} (\beta W_{EE} - 1) / \tau_m & -\beta W_{EI} / \tau_m \\ \beta W_{IE} / \tau_m & (-\beta W_{II} - 1) / \tau_m \end{bmatrix}$$


What are the conditions for the network to be stable?

We need both eigenvalues to have negative real part...

As we just saw, this means that we need:

$$\text{Trace}(A) = A_{EE} + A_{II} < 0$$

$$\text{Det}(A) = A_{EE}A_{II} - A_{EI}A_{IE} > 0$$



Stability of Linear E-I Networks

To determine conditions for stability, we need to do some algebra:

$$A = (\beta W - I) / \tau_m = \begin{bmatrix} (\beta W_{EE} - 1) / \tau_m & -\beta W_{EI} / \tau_m \\ \beta W_{IE} / \tau_m & (-\beta W_{II} - 1) / \tau_m \end{bmatrix}$$

$$\text{Trace}(A) = A_{EE} + A_{II} < 0$$

$$\text{Det}(A) = A_{EE}A_{II} - A_{EI}A_{IE} > 0$$

Stability of Linear E-I Networks

To determine conditions for stability, we need to do some algebra:

$$A = (\beta W - I) / \tau_m = \begin{bmatrix} (\beta W_{EE} - 1) / \tau_m & -\beta W_{EI} / \tau_m \\ \beta W_{IE} / \tau_m & (-\beta W_{II} - 1) / \tau_m \end{bmatrix}$$

$$\text{Trace}(A) = A_{EE} + A_{II} < 0 \implies W_{EE} < W_{II} + 2/\beta$$

$$\text{Det}(A) = A_{EE}A_{II} - A_{EI}A_{IE} > 0$$

Stability of Linear E-I Networks

To determine conditions for stability, we need to do some algebra:

$$A = (\beta W - I) / \tau_m = \begin{bmatrix} (\beta W_{EE} - 1) / \tau_m & -\beta W_{EI} / \tau_m \\ \beta W_{IE} / \tau_m & (-\beta W_{II} - 1) / \tau_m \end{bmatrix}$$

$$\text{Trace}(A) = A_{EE} + A_{II} < 0 \implies W_{EE} < W_{II} + 2/\beta$$

$$\text{Det}(A) = A_{EE}A_{II} - A_{EI}A_{IE} > 0 \implies W_{EI}W_{IE} > (W_{EE} - 1/\beta)(W_{II} + 1/\beta)$$

Stability of Linear E-I Networks

To determine conditions for stability, we need to do some algebra:

$$A = (\beta W - I) / \tau_m = \begin{bmatrix} (\beta W_{EE} - 1) / \tau_m & -\beta W_{EI} / \tau_m \\ \beta W_{IE} / \tau_m & (-\beta W_{II} - 1) / \tau_m \end{bmatrix}$$

$$\text{Trace}(A) = A_{EE} + A_{II} < 0 \implies W_{EE} < W_{II} + 2/\beta$$

$$\text{Det}(A) = A_{EE}A_{II} - A_{EI}A_{IE} > 0 \implies W_{EI}W_{IE} > (W_{EE} - 1/\beta)(W_{II} + 1/\beta)$$

Thus, stability requires: 1) E-E loops are sufficiently weak 2) E-I-E loops are strong enough to counteract E-E loops. But (surprisingly) also depends on strength of I-I loops...

Some aspects are intuitive (role of E-E and E-I-E) while others are less intuitive (role of I-I)

Summary: Linear Networks and E-I Stability

- We can approximate nonlinear networks as linear ones
- This is especially valid for threshold-linear neurons, but also works for other nonlinearities (via linearisation of nonlinear system)
- Linearisation allows us to perform a stability analysis around a fixed point
- Stability analysis of linear E-I model revealed conditions for network stability – inhibition must be strong enough to cancel runaway excitation
- A similar analyses can be applied to analyse paradoxical inhibition in inhibitory-stabilised networks

Linear Networks: Advantages and Limitations

- Linear networks have many advantages: we can solve them analytically and we can fit them to data to infer dynamics from neural recordings
- Real neurons in the brain are nonlinear – what do linear networks miss?
- Linear networks *can't* exhibit: chaos, multiple separate fixed points
- Linear networks *can* generate a line of fixed points, but this requires an eigenvalue with real part exactly zero – nonlinear networks don't require such fine tuning
- Many important insights have been generated using linear networks, but they are too limited for some applications

Summary of Neural Networks (Lecture 2)

- Cortical connectivity is structured, not random
- Structured connectivity can be useful for computations
- The ring network incorporates properties of connectivity from visual cortex
- Feedforward and recurrent mechanisms for orientation tuning can both emerge in the ring model
- State space analysis allows tools from dynamical systems and linear algebra to be applied to understand neural circuit computation
- Linear stability analysis reveals the conditions under which inhibition can stabilise E-I networks