

The effective use of the time domain in neural network models

A.Ravishankar Rao, ravirao@us.ibm.com
Guillermo Cecchi, gcecchi@us.ibm.com

Computational Biology Center
IBM Research
Yorktown Heights, New York, USA

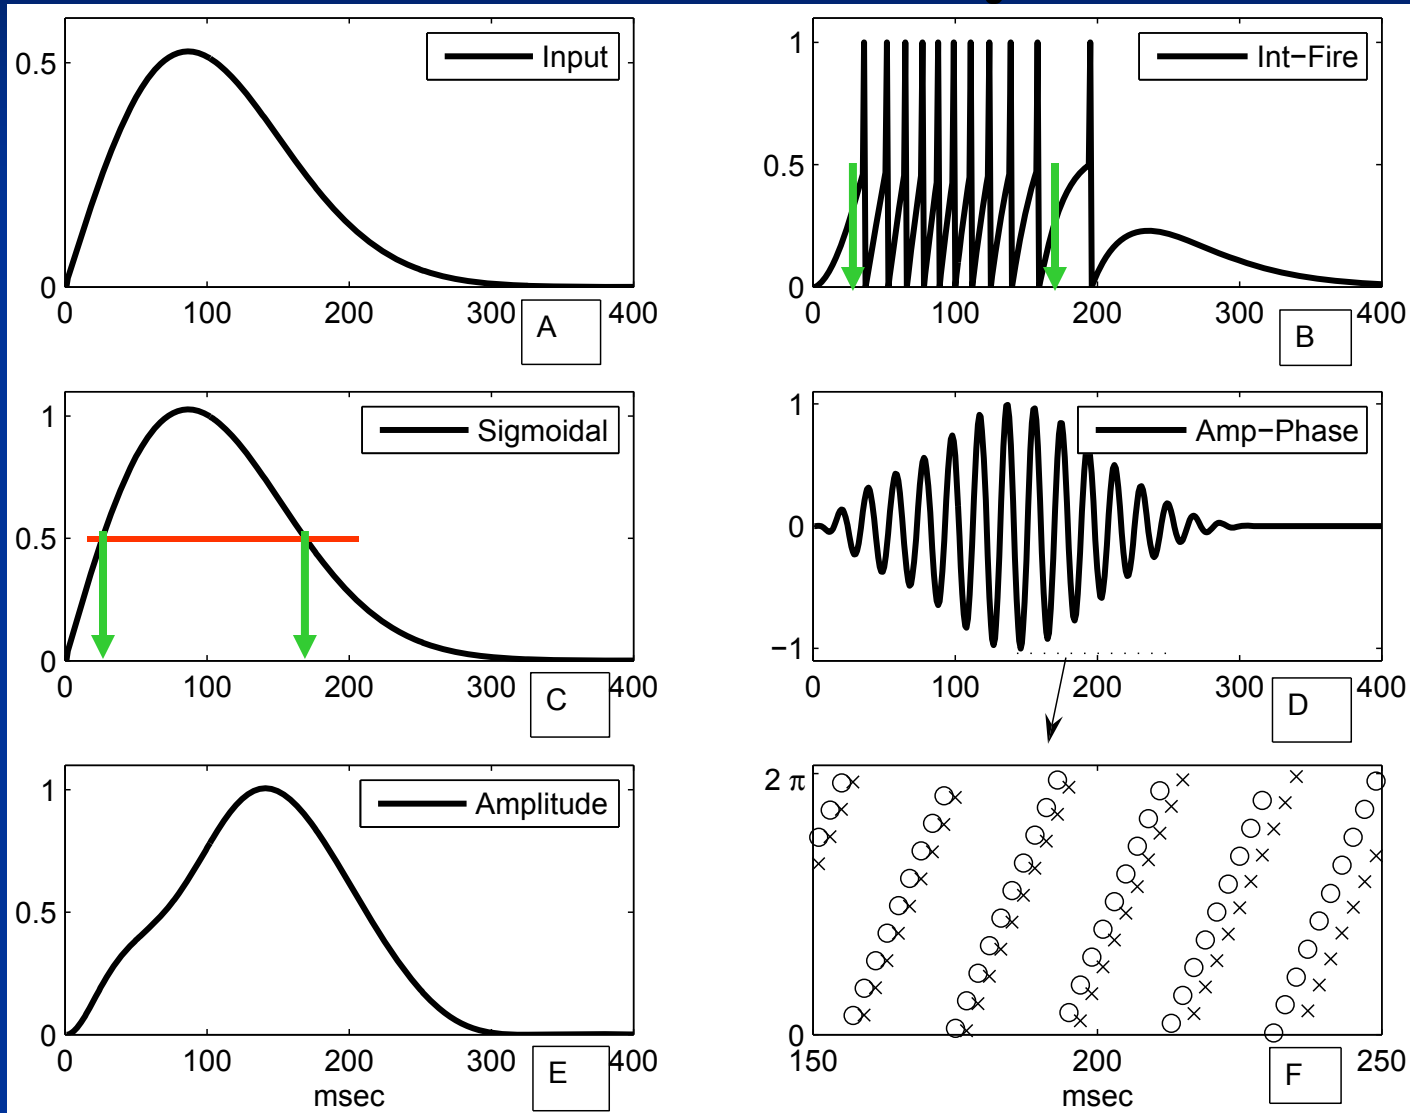
Outline

- Background/Motivation
- Brief review of 1D/2D systems of differential equations
 - Qualitative modeling
- Models of temporal neural behavior
 - Wilson-Cowan cortical dynamics
 - Wilson-Cowan oscillator
 - FitzHugh-Nagumo neural oscillator
- Networks of oscillators
 - Kuramoto model and synchronization
 - Network topology and synchronization
 - Relevance: the binding problem, and a learning network that addresses this problem
- Stimulus representation through temporal encoding
 - Winnerless networks
- Challenges in applying the theory
 - Metrics for neural spike trains
 - Discrepancies between the models and actual neurophysiological data
- Conclusions
 - A rich area for future research
 - Opportunities for interactions with other disciplines, including neuroscience

Motivation: temporal phenomena in biological systems

- Rhythms
 - Circadian rhythms
 - Heartbeats
 - Central pattern generators for locomotion
- Neuroscience
 - Spikes
 - Oscillations observed in EEG recordings
 - Binding Problem
 - Representation of stimuli (odors): Winnerless networks
- E.g. recent work in neuroscience indicates increasing importance of timing
 - "Driving fast-spiking cells induces gamma rhythm and controls sensory responses." J. A. Cardin, et al. Nature, 2009.
 - "Millisecond-Timescale Optical Control of Neural Dynamics in the Nonhuman Primate Brain." X. Han et al. Neuron, 2009.
 - "Timing, Timing, Timing: Fast Decoding of Object Information from Intracranial Field Potentials in Human Visual Cortex," H. Liu et al, Neuron, 2009.
- Nevertheless, a significant amount of research in the neural networks field does not take the temporal dimension into account effectively

Comparison of different models of neural activity



1D and 2D differential equations



$$\frac{dx}{dt} = \lambda - \gamma x$$

★ Solution:

$$x(t) = x_0 e^{-\gamma t} + \frac{\lambda}{\gamma} (1 - e^{-\gamma t})$$

★ Initial condition determines evolution of the system

★ In two dimensions, we consider $\dot{X} = AX$

★ The initial values, and eigenvalues of A determine the evolution of the system

★ $X(t) = \alpha_1 e^{\lambda_1 t} c_1 + \alpha_2 e^{\lambda_2 t} c_2$ where λ are eigenvalues, and c are eigenvectors, and α is determined by initial conditions


★ Solving differential equations using Eulers method:

★ Taylor series approximation

$$x(t_0 + \Delta t) = x(t_0) + \Delta t \frac{dx}{dt} \Big|_{t_0}$$

★ For better accuracy, use methods like Runge-Kutta

$$\lambda = 3, \gamma = 1$$



Qualitative theory of differential equations: use signs

$$\dot{V} = F(V)$$

At equilibrium

$$F(V) = 0$$



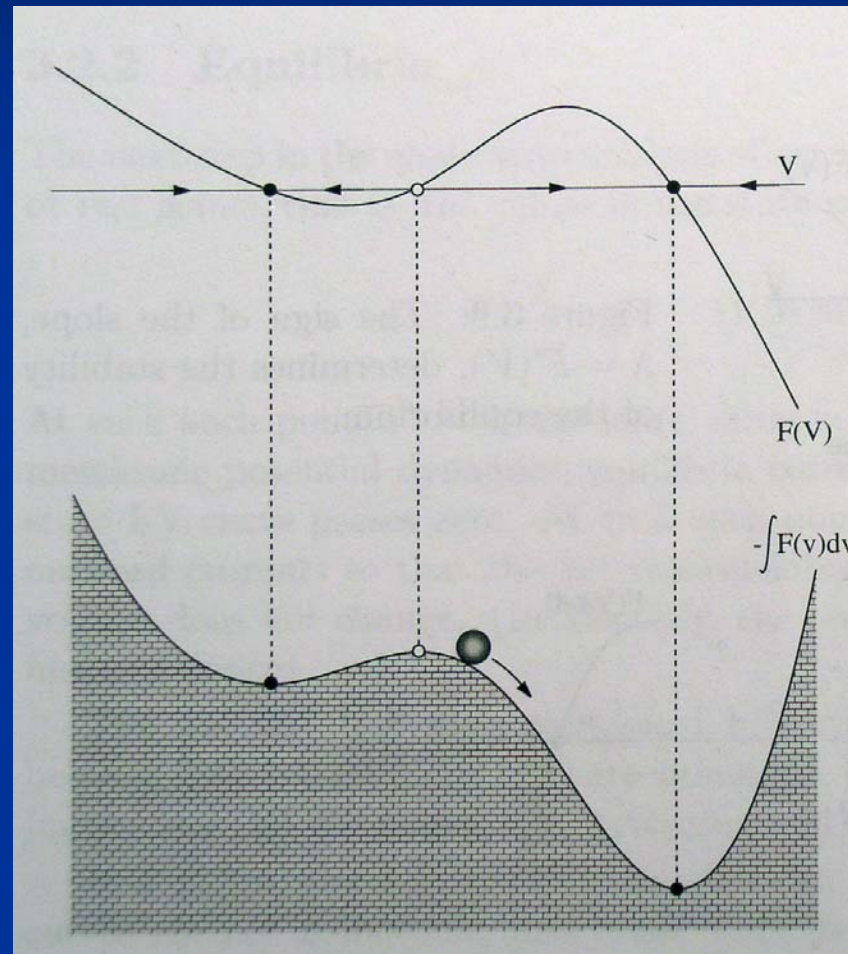



Figure 3.9: The sign of the slope, $\lambda = F'(V)$, determines the stability of the equilibrium.



Mechanistic interpretation: energy landscape






Attraction domains



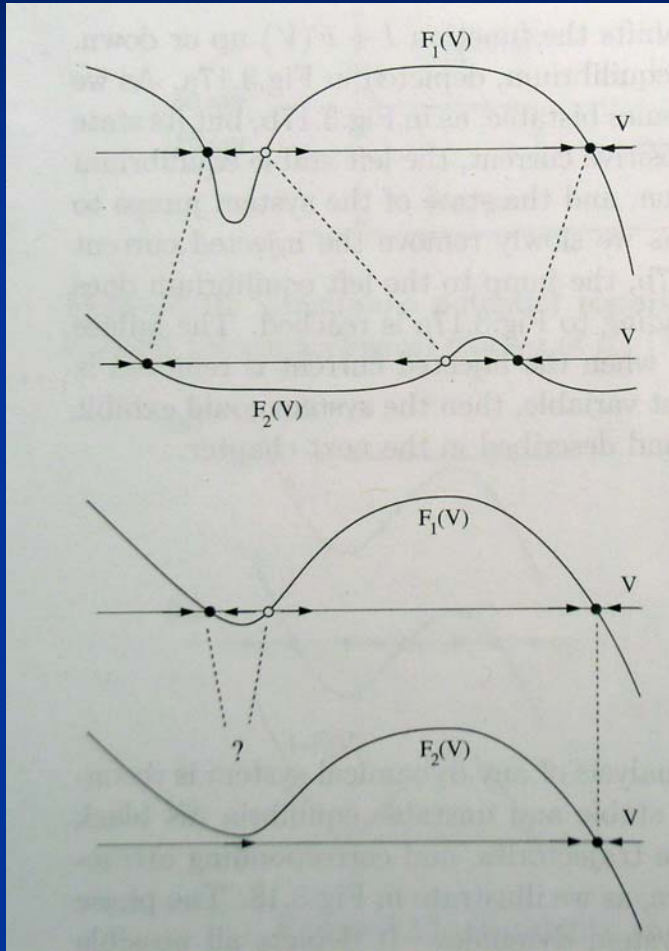
Two attraction domains are separated by an unstable equilibrium

Phase portrait of a 1D system

$$\dot{V} = F(V)$$



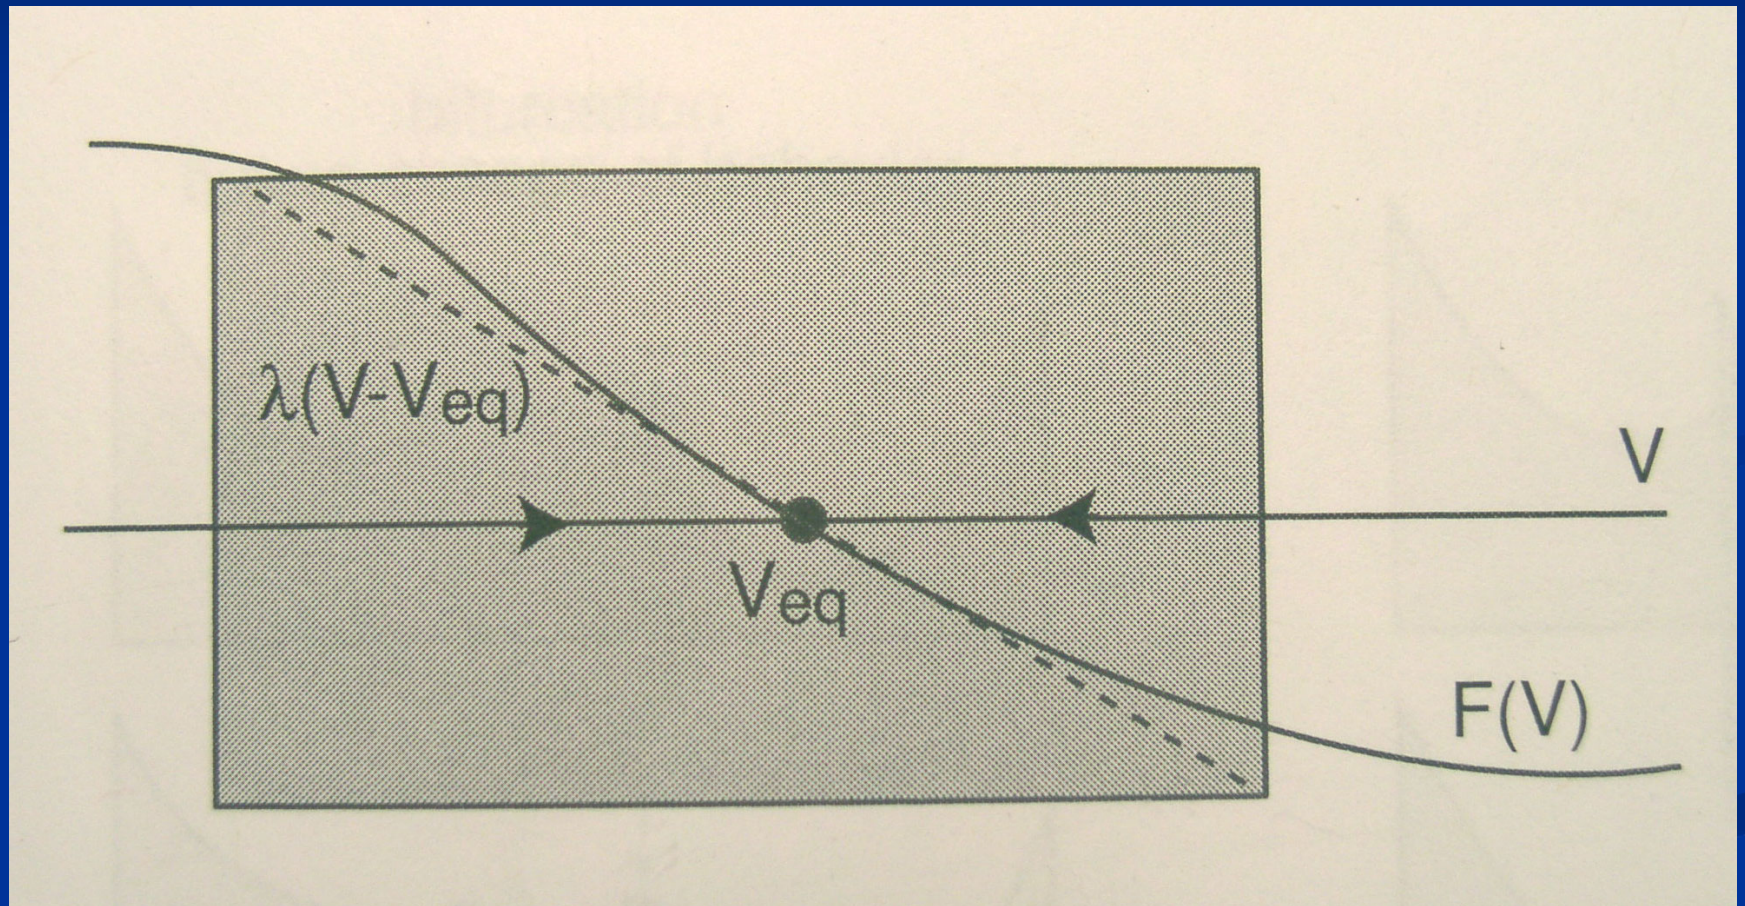
Depiction of equilibrium states, trajectories and attraction domains



These two systems are qualitatively identical

These two systems are qualitatively different

Hartman-Grobman theorem: linearization



Can replace a more complex system with a simpler one with identical phase portrait

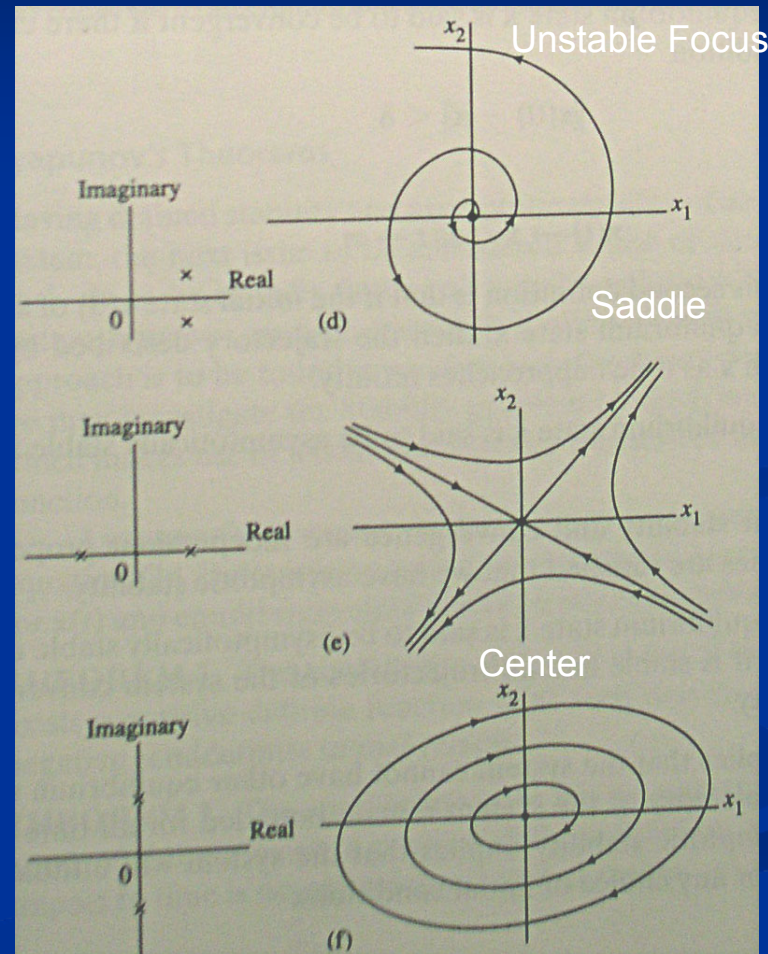
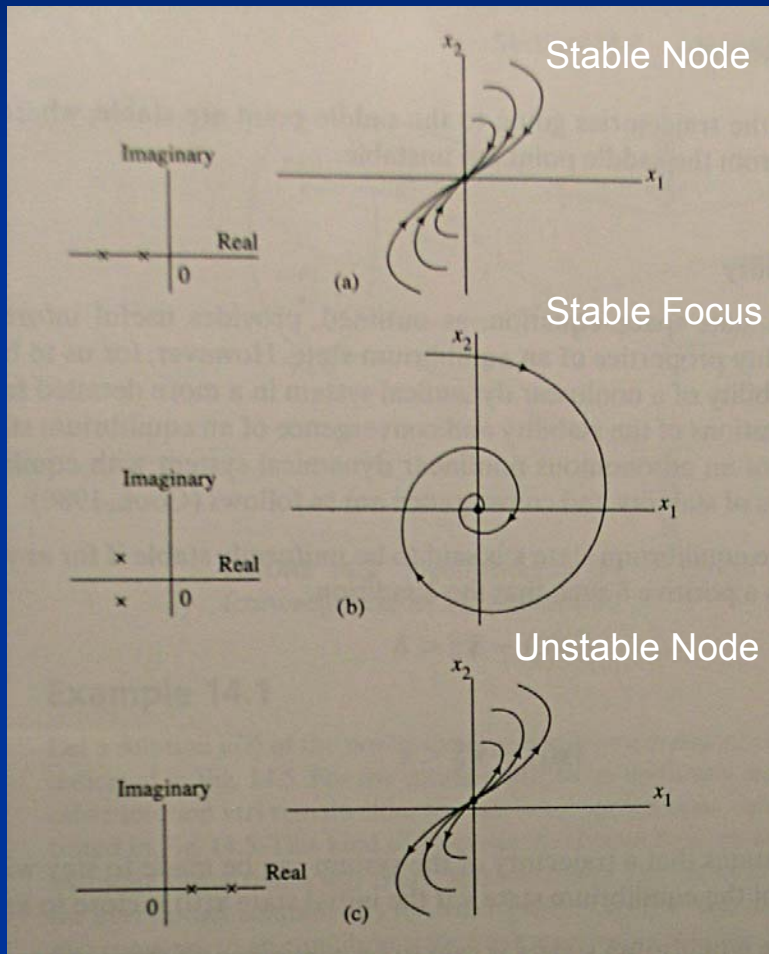
Two-dimensional phase portraits

TABLE 14.1 Classification of the Equilibrium State of a Second-Order System

Type of Equilibrium State $\bar{\mathbf{x}}$	Eigenvalues of the Jacobian Matrix \mathbf{A}
Stable node	Real and negative
Stable focus	Complex conjugate with negative real parts
Unstable node	Real and positive
Unstable focus	Complex conjugate with positive real parts
Saddle point	Real with opposite signs
Center	Conjugate purely imaginary

$$\dot{X} = AX$$

Equivalence classes for 2D phase portraits

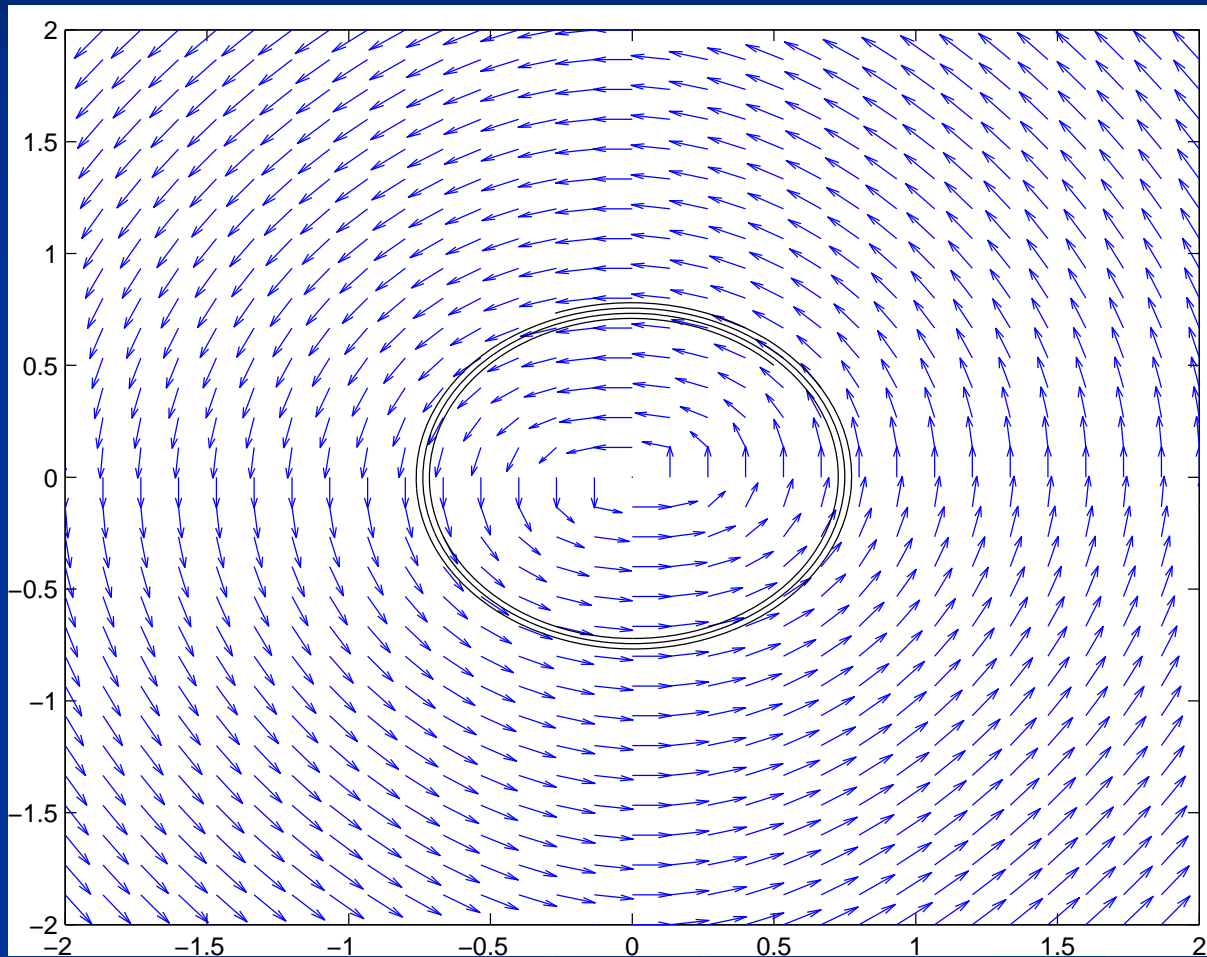


Powerful descriptive tool

Oscillations

- Oscillations occur in a linear system when there is one pair of eigenvalues that is purely imaginary.

Phase portrait when $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$



Example: Creating 2D phase portraits in MATLAB


- Governing equation is $\dot{X} = AX$
- A is a 2×2 matrix
- X consists of a 2d coordinate $[x_1 \ x_2]$
- First select range for x values: $x_1 = -2:0.1:2$
- $x_2 = -2:0.1:2$
- Use `meshgrid` to create a grid of 2D points
- `[X1, X2] = meshgrid(x1,x2);`
- `[m,n] = size(X1);` (should be 41×41)

Creating 2D phase portraits

- `reshape(X1, 1, m*n)` gives a 1x1681 matrix.
- `X = [reshape(X1, 1, m*n); reshape(X2, 1, m*n)]` gives a 2x1681 matrix
- Compute `Xdot` at each grid point: `Xdot = AX`
- `V1 = reshape(Xdot(1,:), m, n); V2 = reshape(Xdot(2,:), m, n);`
- Normalize velocity vectors:
`V1 = V1./sqrt(V1.^2 + V2.^2 + 0.00001);`
`V2 = V2./sqrt(V1.^2 + V2.^2 + 0.00001)`
- `quiver(X1, X2, V1, V2); axes([-2 2 -2 2]);`

Example: tracing the trajectory of a solution

- Select the matrix A
- Select a starting point: say $x_1 = 0.5$, $x_2 = 0.5$
- Select an integration step, say $dt = 0.01$
- Select number of steps, say 2000
- $x_1 = \text{zeros}(2000, 1)$; $x_2 = \text{zeros}(2000, 1)$;
- $x_1(1) = 0.5$; $x_2(1) = 0.5$;
- for $i = 2:2000$
 - $x_1(i) = x_1(i-1) + dt * (A(1,1)*x_1(i-1) + A(1,2)*x_2(i-1))$;
 - $x_2(i) = x_2(i-1) + dt * (A(2,1)*x_1(i-1) + A(2,2)*x_2(i-1))$;
 - end;
- Overlay with phase portrait plot: hold on;
- $\text{plot}(x_1, x_2, 'k')$;
- This plots one trajectory



Euler integration

Linear stability analysis

Given the following nonlinear system

$$\frac{dX}{dt} = F(X)$$

Compute an equilibrium point, x_{eq}

$$F(x_{eq}) = 0$$


Compute the Jacobian matrix A at x_{eq}

$$A = \begin{bmatrix} \frac{\partial F_1}{\partial x_1} & \frac{\partial F_1}{\partial x_2} & \cdots \\ \frac{\partial F_2}{\partial x_1} & \frac{\partial F_2}{\partial x_2} & \cdots \\ \vdots & \vdots & \frac{\partial F_N}{\partial x_N} \end{bmatrix}$$

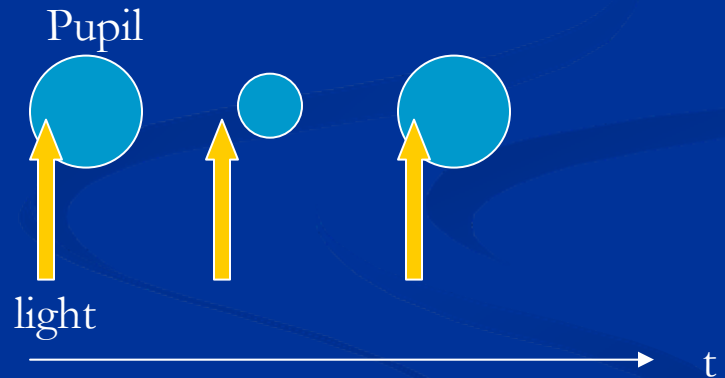
If all eigenvalues of the linear system have negative real parts, the nonlinear system is asymptotically stable. If at least one eigenvalue has positive real part, the system is unstable

Simple ways of generating oscillations


- Integrate and fire neural model with constant input



- Feedback with delays



- Mutual inhibition with fatigue



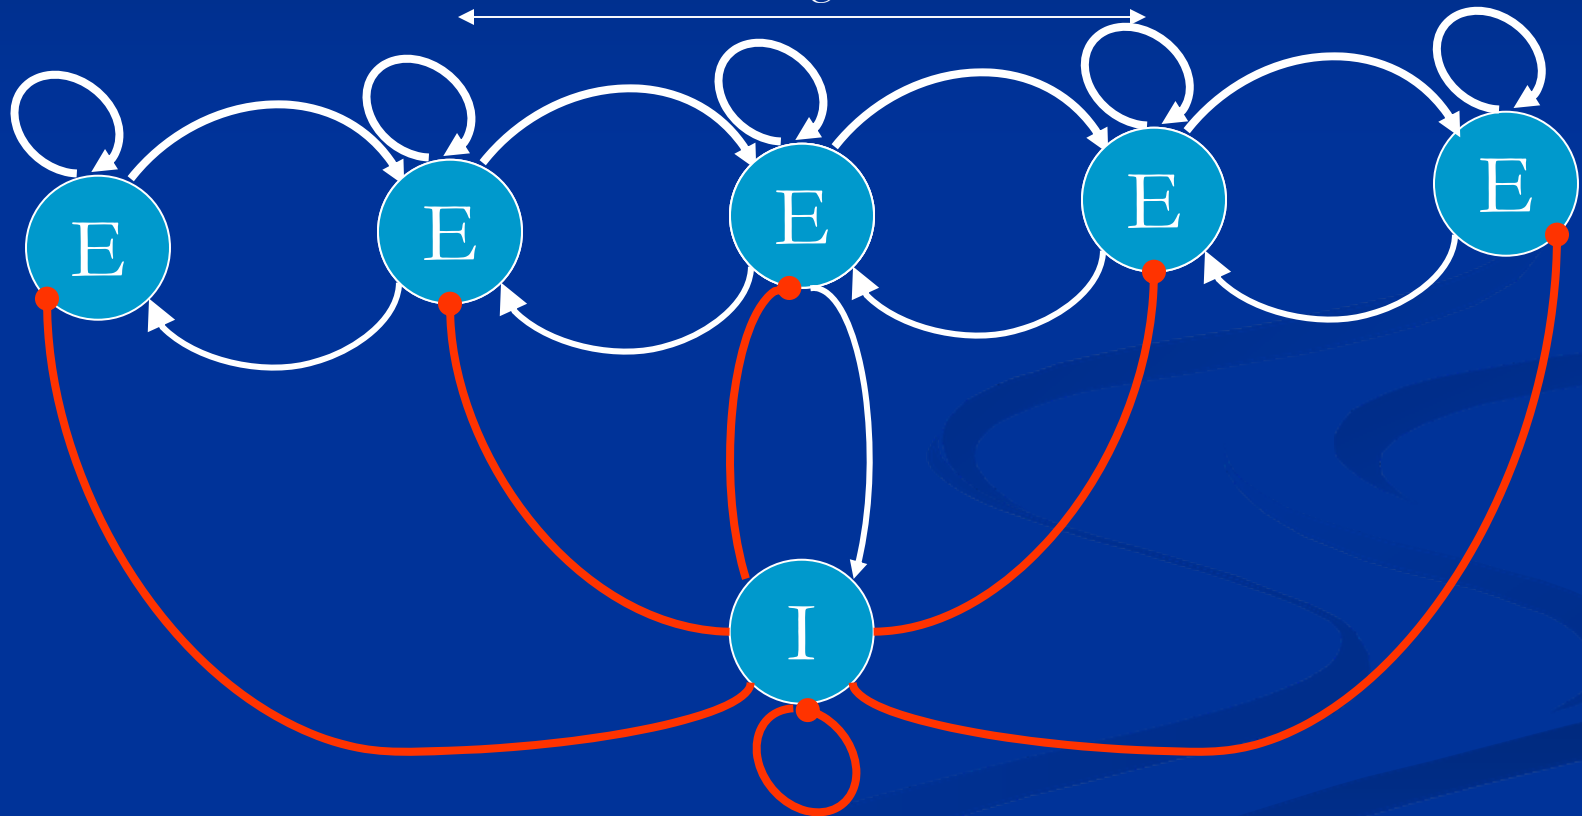
Wilson-Cowan cortical dynamics

Anatomical considerations:
All types of connections exist

model

Cortical Position

Short range excitation



Long range inhibitory feedback

(Only one inhibitory neuron is shown for clarity)

Wilson-Cowan cortical dynamics model

- ★ $E(t)$ = proportion of excitatory cells firing per unit time
- ★ $I(t)$ = proportion of inhibitory cells firing per unit time
- ★ Let r be the refractory period. Some cells cannot fire during this period.
- ★ Then at time $t + \tau$, proportion of sensitive E cells, $E_s = 1 - r_e E$
- ★ Let c_e = average number of excitatory synapses per E cell
- ★ Let g_e = average number of inhibitory synapses per E cell
- ★ Let $P(t)$ be external excitation for E cells, $Q(t)$ that for I cells
- ★ Net input excitation for E cells is $c_e E - g_e I + P(t)$
- ★ Net response is $S_e[c_e E - g_e I + P(t)]$ for E cells and $S_i[c_i E - g_i I + Q(t)]$ where S is a sigmoid
- ★ Probability of cell being sensitive is independent of its input excitation, so $E(t + \tau) = (1 - r_e E)S_e[c_e E - g_e I + P(t)]$
- ★ Use Taylor series expansion: $E(t + \tau) = E(t) + (dE/dt)\tau$

Wilson-Cowan cortical dynamics model



$$\tau \frac{dE}{dt} = -E + (1 - r_e E) S_e [c_e E - g_e I + P(t)]$$

$$\tau \frac{dI}{dt} = -I + (1 - r_i E) S_i [c_i E - g_i I + Q(t)]$$

★ Introducing spatial coordinates x , a spatial weight distribution w , and setting $r = 0$ we obtain



$$\tau \frac{dE}{dt} = -E(x) + S_e \left[\sum_x w_{EE} E(x) - \sum_x w_{IE} I(x) + P(x) \right]$$

$$\tau \frac{dI}{dt} = -I(x) + S_i \left[\sum_x w_{EI} E(x) - \sum_x w_{II} I(x) + Q(x) \right]$$

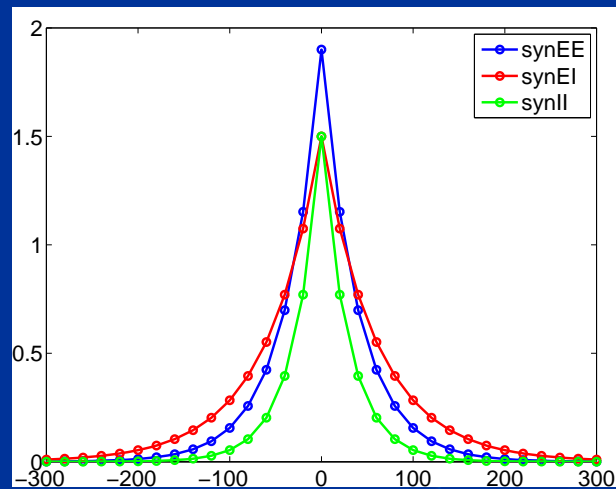
★ These pairs of equations exhibit the same dynamics

★ The choice of model parameters affects the behavior of the network, as we will see

Wilson-Cowan cortical dynamics

WCcortexSTM.m model

- ★ EE and IN are 1D arrays of neural units.
- ★ Sigmoidal function is $S(x) = 100x^2 / (\theta^2 + x^2)$
- ★ synEE, synEI, synII are spatial patterns of synaptic weights.
- ★ Weights vary as $w_{ij}(\Delta x) = b_{ij} \exp(-\Delta x / \sigma_{ij})$.
- ★ $\sigma_{EE} = 40, \sigma_{EI} = 60, \sigma_{II} = 30$.
- ★ $b_{EE} = 1.95, b_{EI} = 1.4, b_{II} = 2.2$
- ★ Parameters are chosen to prohibit the formation of spatially uniform states.
- ★ Use Euler method for integration. Time step $\Delta t = 0.5 \text{ msec}$



Monotonically decreasing distance function

Wilson-Cowan cortical dynamics model

WCcortexSTM.m

Inner loop of computation

```
EEresp = NeuralConv(synEE, EE) - NeuralConv(synEI, IN) + P;  
EEresp = (EEresp.*(EEresp > 0)).^2;  
INresp = NeuralConv(synEI, EE) - NeuralConv(synII, IN) + Q;  
INresp = (INresp.*(INresp > 0)).^2;  
EE = EE + (DeiT/DT)*(-EE + 100*EEresp./(20^2 + EEresp));  
IN = IN + (DeiT/DT)*(-IN + 100*INresp./(40^2 + INresp));
```

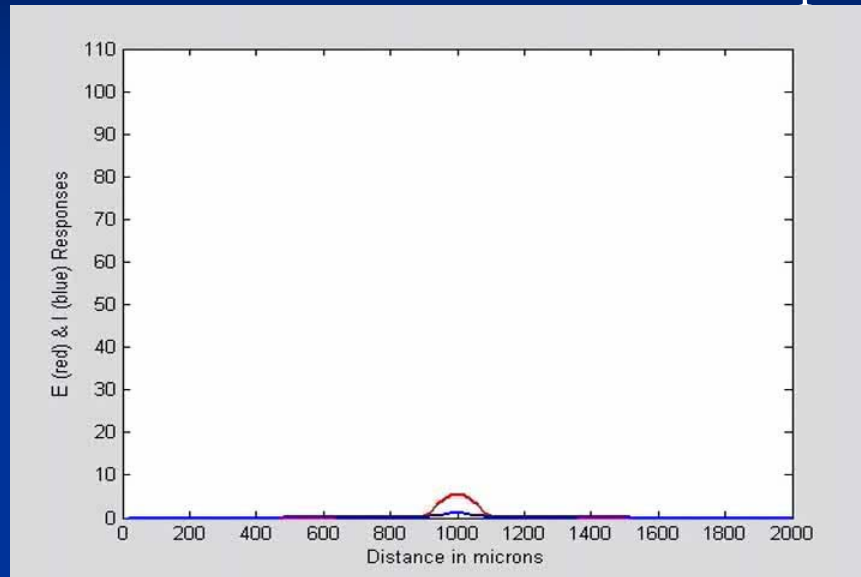
P is the stimulus. Here, P is 1.0 for 10 ms.

Q = 0


Wilson-Cowan cortical dynamics

WCcortexSTM.m **model**

Cortex stimulated with a narrow pulse, width=100 microns, for 10ms



Movie of spatial cortical response



Temporal response of central neural unit at 1000 microns

- Network reaches an asymptotically stable state that retains information about the stimulus location.
- This represents a short term memory mode.
- Short-range recurrent excitation stabilizes a narrow pulse of activity
- Longer range inhibition prevents excitation from spreading


Wilson-Cowan cortical dynamics model

WCcortexSTM.m

Cortex stimulated with a narrow pulse, width=100 microns, for 10ms

Movie of phase portrait plot


WCCortexSTM_movie_small_stimulus_phase_plot_WMV V9.wmv




Wilson-Cowan cortical dynamics model

WCcortexSTM.m

Cortex stimulated with a wider pulse, width=500 microns for 10ms



Movie of spatial cortical response



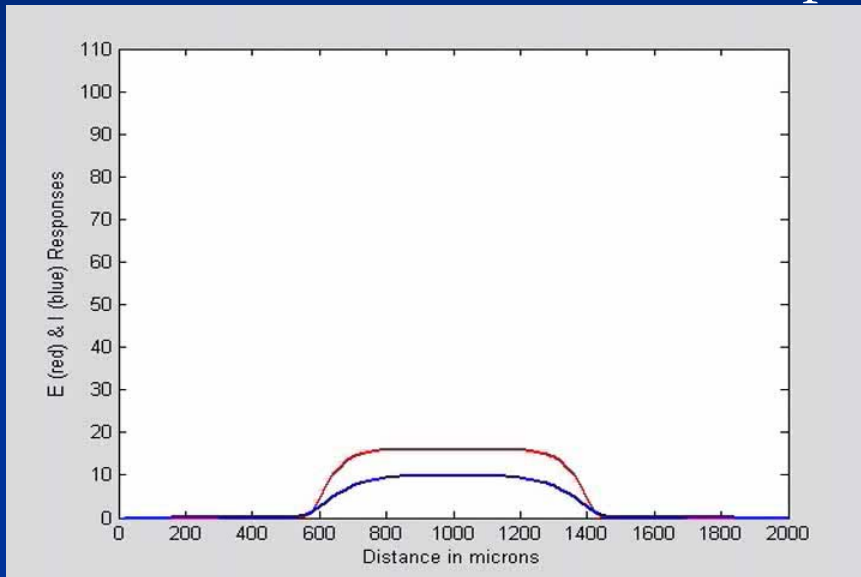
Temporal response of central neural unit at
1000 microns

- Network reaches an asymptotically stable state that retains information about the stimulus location.
- This represents a short term memory mode.
- Short-range recurrent excitation stabilizes a narrow pulse of activity
- Longer range inhibition prevents excitation from spreading


Wilson-Cowan cortical dynamics model

WCcortexSTM.m

Cortex stimulated with a wider pulse, width=500 microns for 10ms



Movie of spatial cortical response



Temporal response of central neural unit at
1000 microns

- Network reaches an asymptotically stable state that retains information about the stimulus location.
- This represents a short term memory mode.
- Short-range recurrent excitation stabilizes a narrow pulse of activity
- Longer range inhibition prevents excitation from spreading

Wilson-Cowan cortical dynamics

WCcortexOSC.m **model**


- ★ All parameters are the same as before, except that
- ★ $b_{EE} = 1.9, b_{EI} = 1.5, b_{II} = 1.5$
- ★ Parameters are chosen to prohibit the formation of spatially uniform states.

Wilson-Cowan cortical dynamics model

WCcortexOSC.m

Cortex stimulated with a narrow pulse, width=100 microns, for 10ms

WCCortexOSC_movie_small_stimulus_WMV V9.wmv



Movie of spatial cortical response

Temporal response of central neural unit at
1000 microns


- Connectivity parameters have been changed
- Spatially localized oscillations, about 20 Hz

Wilson-Cowan cortical dynamics model

WCcortexOSC.m

Cortex stimulated with a wider pulse, width=400 microns for 10ms

WCCortexOSC_movie_large_stimulus_WMV V9.wmv



Temporal response of central neural unit at
1000 microns

Movie of spatial cortical response

- More spatially complex oscillations are produced.
- There is a spatially adjacent pair of synchronized oscillations.

Wilson-Cowan cortical dynamics


WCcortexOSC.m

model

Cortex stimulated with a wider pulse, width=400 microns for 10ms

Movie of phase portrait plot


WCCortexOSC_movie_large_stimulus_phase_plot_WMV V9.wmv




Wilson-Cowan cortical dynamics model

Several interesting phenomena are exhibited through this model

- An Active Transient Mode
 - Gives rise to a dramatic response amplification
 - This transient amplification may be a mechanism for detection of weak stimuli
- Generation of standing waves
- Explanation of visual hallucinations by extending model to 2D
- Image processing tasks are possible with a 2D model
 - Mantere et al, *Journal of Mathematical Imaging and Vision* 2, 251-259 (1992)



Input



Cortical representation

Linear vs. nonlinear oscillations


- In linear systems, the only possible oscillations involve sines and cosines
- If initial conditions are changed, amplitude of oscillation changes
- Noise in biological systems can play havoc with such oscillators
- Clearly unsuitable for control of vital functions, e.g. breathing, heartbeats, locomotion
- This implies the use of non-linear oscillators
- Example

$$\frac{dr}{dt} = 1 - r^2$$

$$\frac{d\theta}{dt} = \omega$$


- Robust wrt initial conditions, noise, numerical imprecision
- Stable limit cycle:
 - when $r > 1$, dr/dt is negative, which reduces r
 - when $r < 1$, dr/dt is positive, which increases r

Linear Oscillator




- Initial condition 1
- Initial condition 2

Non-linear Oscillator



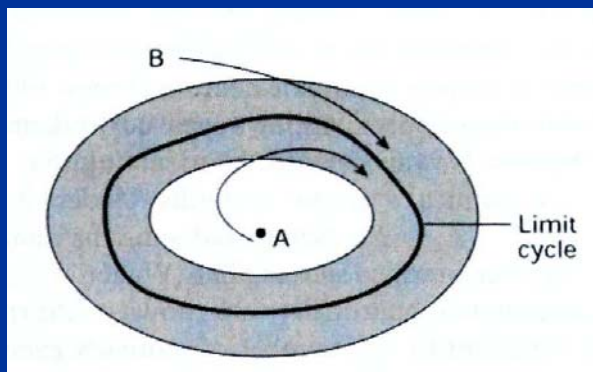
Nonlinear oscillation: Limit cycles

- A trajectory $X(t)$ of a dynamical system is an **oscillation** if $X(T+t) = X(t)$ for some unique $T > 0$ for all t
- T is the period of oscillation
- Uniqueness of T implies equilibrium points are excluded
- For linear systems there are infinitely many periodic solutions within any small neighborhood of a given oscillation.
- Behavior exhibited by nonlinear systems: limit cycles
- Def. All trajectories in a sufficiently small region of a limit cycle are spirals.
- Limit cycle can be asymptotically stable, or unstable



Nonlinear oscillations: Poincare-Bendixon Theorem


- Given an annular region in a constant-coefficient 2D system satisfying (1) annulus contains no equilibrium points (2) all trajectories crossing the boundary of the annulus enter it.
- Then, the annulus must contain at least one asymptotically stable limit cycle



Region A contains a steady state (unstable)
Trajectories enter the annulus from Regions B and A.

The annulus must contain a limit cycle

Wilson-Cowan network oscillator




- Localized (non-spatial) version of the original Wilson-Cowan (1972) equations
- Simplest demonstration of those equations showing a limit cycle
- Assume all E neurons receive identical stimuli and have identical synaptic strengths
- Reduces to a 2 neuron model, where S is the sigmoid function

$$\frac{dE}{dt} = 0.2[-E + S(1.6E - I + K)]$$


$$\frac{dI}{dt} = 0.1[-I + S(1.5E)]$$

Wilson-Cowan network oscillator

WCoscillator.m



Time course



Phase portrait

Movie: [WCoscillator_movie_1_WMV V9.wmv](#)

- Input, $K=20$
- Equilibrium: Unstable spiral point
- Trajectories that enter the box $[0 \ 100 \ 0 \ 100]$ always stay within it
- Using Poincare-Bendixon theorem, this implies existence of a limit cycle

Robustness: All trajectories approach the limit cycle asymptotically

FitzHugh-Nagumo model

- Simplest equations for neural spike generation
- Represent a 2D simplification of the original 4D Hodgkin-Huxley neural equations
- V = voltage across the axon membrane
- I_{inp} = input current
- R = recovery variable (outward K^+ current)


$$\frac{dV}{dt} = 10\left[V - \frac{V^3}{3} - R + I_{inp}\right]$$

$$\frac{dR}{dt} = 0.8[-R + 1.25V + 1.5]$$

- Time constant for V is 0.1ms, and for R is 1.25ms
 - Activation process is much faster than the recovery process
 - Fast vs. slow

FitzHugh-Nagumo model


FitzHugh.m



Time course

- Input, $I_{inp}=1.5$ generates action potentials
- Subthreshold input, $I_{inp}=0.5$
 - generates no action potential
- Equilibrium: Unstable spiral point
- Trajectories enter the dotted green box as shown
 - E.g. at #1, choose sufficiently large value of V
 - The line to the left intersects trajectories moving down, left
- Using Poincare-Bendixon theorem, this implies existence of a limit cycle

$$dv/dt = 0$$



Phase portrait

■ FitzHugh_movie_WMV V9.wmv

Robustness: All trajectories approach the limit cycle asymptotically

Network of interconnected oscillators: synchronization

- Synchronization is observed in natural phenomena
 - Networks of heart pacemaker cells
 - Synchronously flashing fireflies
- Sync.avi from
 - <http://go.owu.edu/~physics/StudentResearch/2005/BryanDaniels/index.html>

Strogatz, "From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators, Physica D, 2000

Winfree-Kuramoto Model

Based on Winfree's initial work, Kuramoto showed that for weakly coupled, near identical limit-cycle oscillators:

$$\dot{\theta}_i = \bar{\omega}_i + \sum_{j=1}^N \Gamma_{ij}(\theta_j - \theta_i)$$

The simplest case, the Kuramoto model, corresponds to

$$\Gamma_{ij}(\theta_j - \theta_i) = \frac{K}{N} \sin(\theta_j - \theta_i)$$

where $K > 0$ is the coupling strength. Given the rotational symmetry of the equations, one can write

$$\dot{\theta}_i = \omega_i + \frac{K}{N} \sum_{j=1}^N \sin(\theta_j - \theta_i)$$


where $\omega_i = \Omega - \bar{\omega}_i$ is the deviation with respect to the ensemble mean frequency.

Strogatz, "From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators, Physica D, 2000

Winfree-Kuramoto Model

By introducing an *order parameter*

$$r e^{i\psi} = \frac{1}{N} \sum_{n=1}^N e^{i\theta_n}$$



it becomes evident that this model is based on a *mean field* approximation:


$$\dot{\theta}_i = \omega_i + K r \sin(\psi - \theta_i)$$

That is, each oscillator is coupled to the mean field quantities r and ψ , as the *ensemble* amplitude and phase.

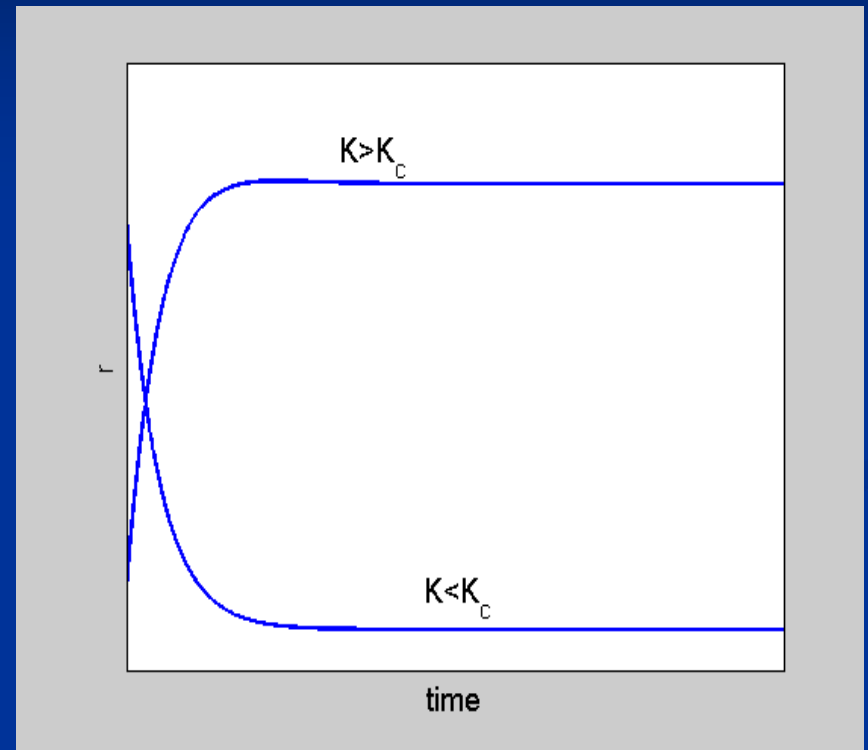
The ensemble synchronization depends on the value of K , and the spread of natural frequencies. For a given spread, there is a *phase transition* as a function of K .

Strogatz, "From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators, Physica D, 2000

Winfree-Kuramoto Model



K =Coupling strength

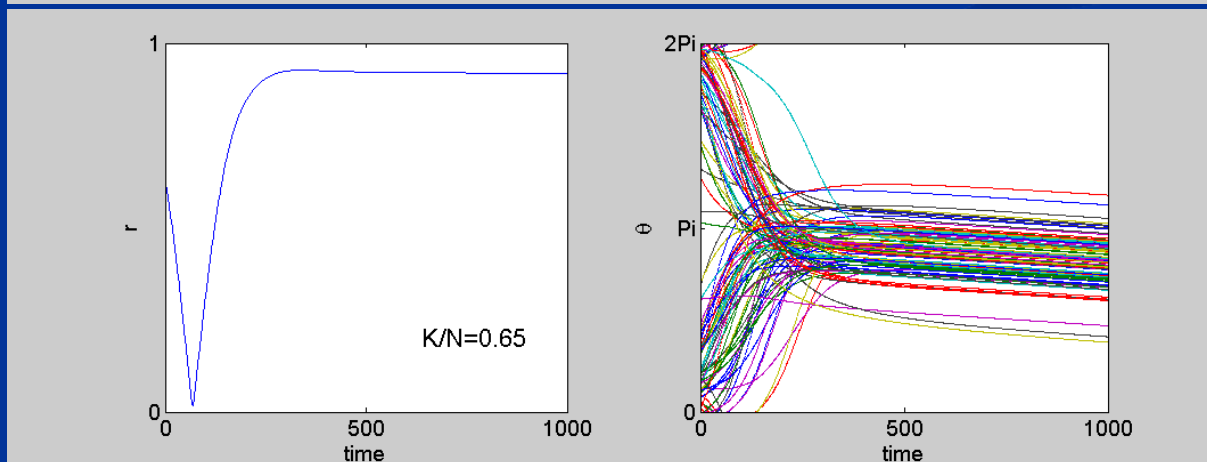



Phase transition: un-synchronized behavior to global synchronization as a function of the coupling strength

Winfree-Kuramoto Model

K =Coupling strength
 N =number of nodes
in the network

$K/N=0.05$
 $N=100$



$K/N=0.65$
 $N=100$

Synchronization and Topology

The synchronization of an array of oscillators depends heavily on the network *topology*.




Fig. 14.1. The three basic classes of complex networks. (a) ER random networks, (b) small-world networks (SWNs), (c) scale-free networks (SFNs)

Small-world: structured, but short

Regimes


regular



small-world



random



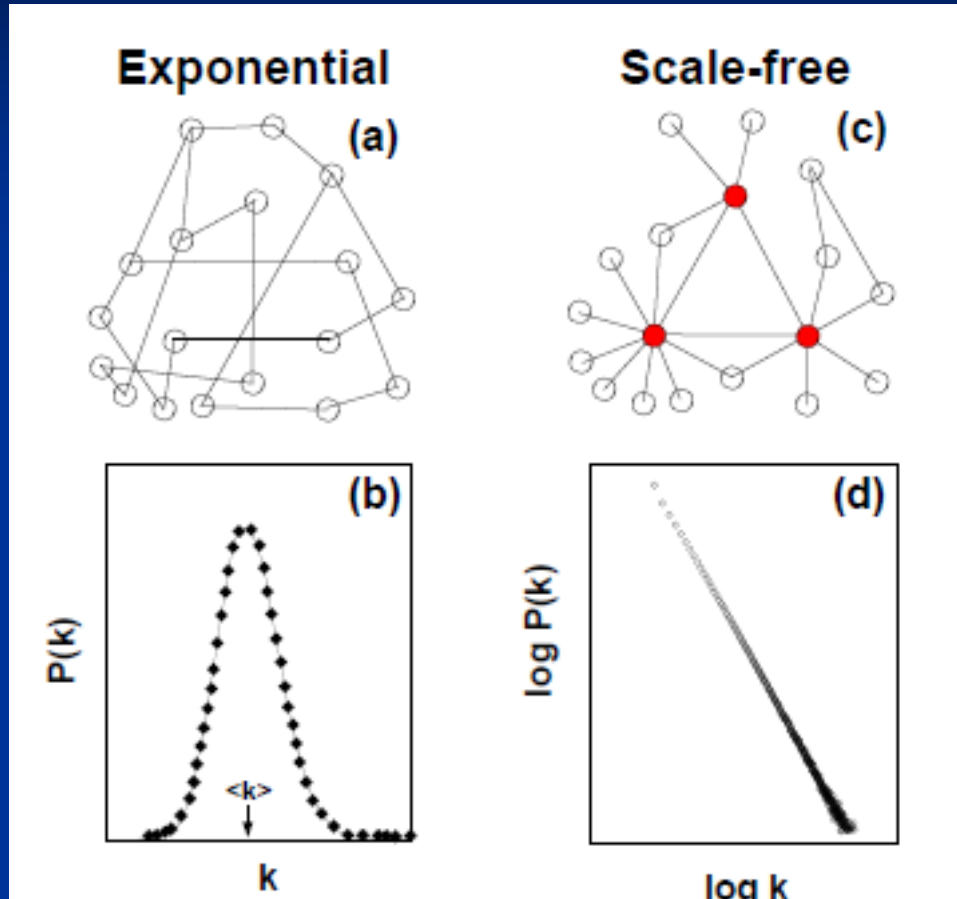
$L(p)$ = characteristic path length (shortest path length between 2 vertices, averaged).

$C(p)$ = clustering coefficient (measures local Density of triangulation).

“shortcut” parameter →

Watts & Strogatz, 1998

Scale-free: a few hubs



degrees scale as

$$d(k) \sim k^{-\alpha}$$

Random (Erdős):
All nodes are equal
Mean degree well defined

Scale-free:
Some nodes are hubs
Mean degree not defined

E.g. Map of roads

E.g. Map of airline routes

Barabasi, 2000

Synchronization: small-world

Assume a systems of oscillators:

$$\dot{x}_i = F(x_i) - d \sum_j W_{ij} x_j$$

Such that $\sum_j W_{ij} = 0 \quad \forall i$, and F determines the autonomous dynamics. The coupling strength is d . If all oscillators are identical, the completely synchronized state, $x_i(t) = s(t) \quad \forall i$, $\dot{s} = F(s)$, is a solution of the above equation. Let X be the mean field oscillation

$$X = \frac{1}{N} \sum_{i=1}^N x_i$$

We can measure the phase organization with two variables:

$$Var(X)$$

and

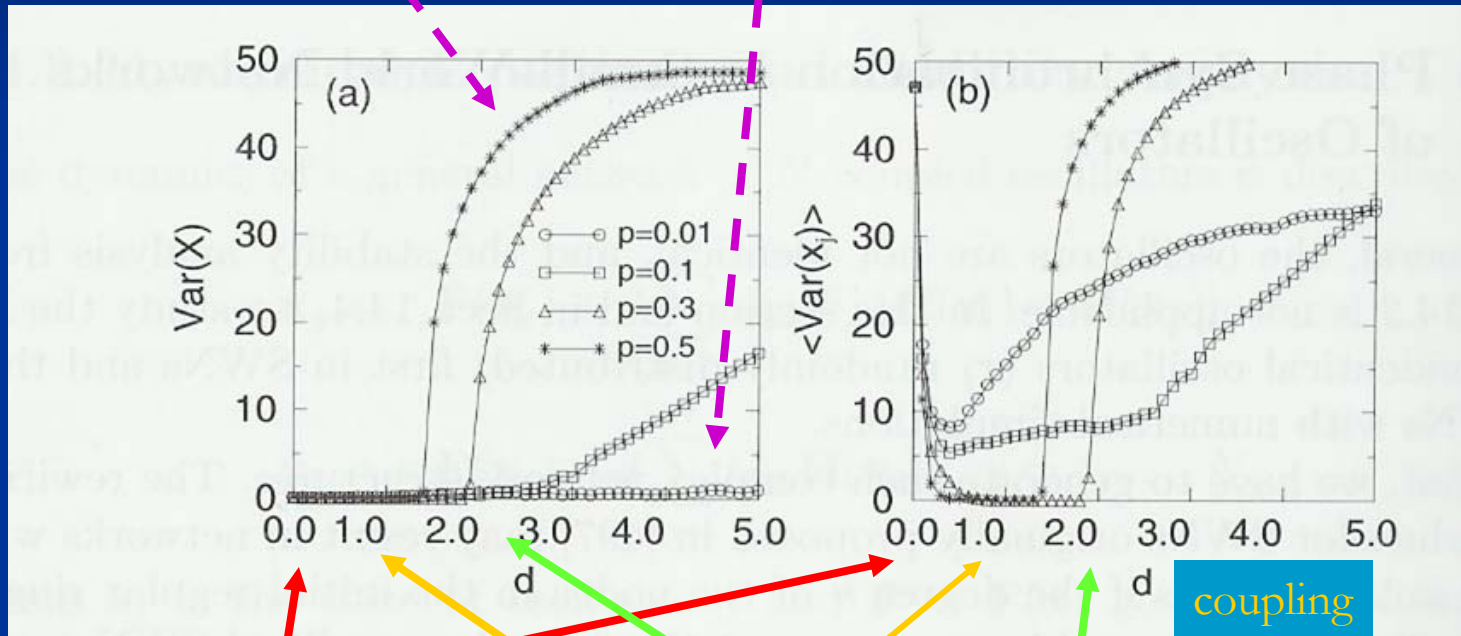
$$\langle Var(x_i) \rangle$$

The small-world shortcut parameter p determines the degree of synchronization, as a function of the coupling parameter d : intuitively, *for the same total number of links*, the shortcuts spread synchrony around.

Synchronization: small-world

more shortcuts

less shortcuts




Individual units oscillate, uncoupled

Individual oscillations are killed, without synchrony

Global synchrony emerges


Synchronization: small-world

Similar behavior for $\text{Var}(X)$ as a function of the shortcut parameter p , for the same coupling constant.



← Disordered regime ↔ Oscillator death ↔ Global synchrony →

Synchronization: scale-free



Mean degree = 10

Networks with more heterogeneous degrees generate weaker collective synchronization

Synchronization for different scaling constants: $\alpha = 3$ (\pm), $\alpha = 4$ (\square) and $\alpha = \infty$ (\triangle). Solid line: analytic result for $\alpha = \infty$ (i.e., homogenous network) and $N \rightarrow \infty$


Synchronization: scale-free

Order parameter $s_j = \langle \sin(\Delta\phi_j) \rangle^2 + \langle \cos(\Delta\phi_j) \rangle^2$

$\Delta\phi_j$ is the phase difference w.r.t. the mean field

Order parameter

Distance to the mean




There is a clear dependency between the *degree* of a node and its *synchronizability*: for the hubs, the order parameter approaches 1

The distance to the mean $\Delta X(k) = Var(\bar{X} - x(k))$ is an inverse power-law in k . The larger k is, the smaller is the distance $\Delta X(k)$


Relevance: the binding problem, and a learning network that addresses it

- A central problem: the binding of distributed representations
- Known as the binding problem (or Rosenblatt's Superposition Catastrophe)




Specific Solutions and Design Inspiration

- The binding problem (or Rosenblatt's Superposition Catastrophe)



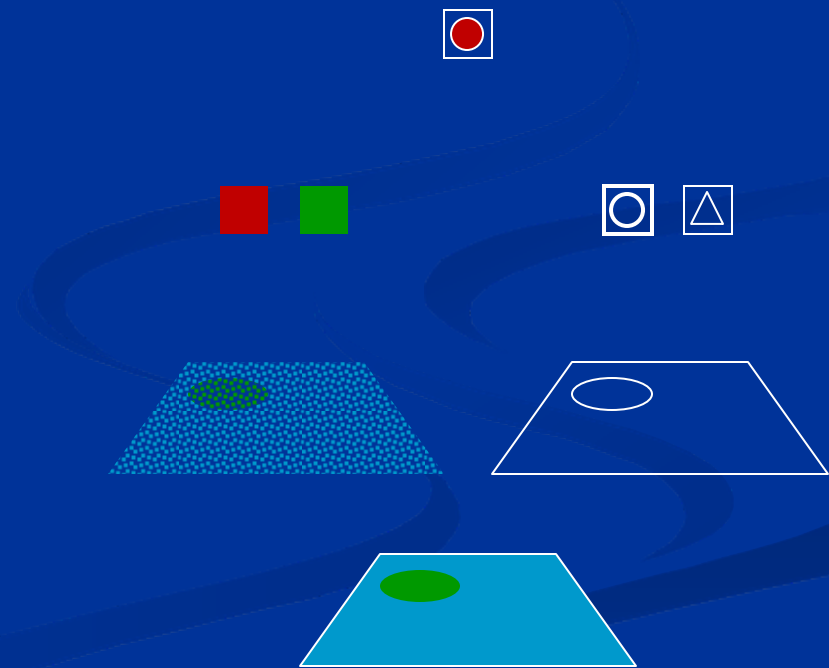
Specific Solutions and Design Inspiration

- The binding problem (or Rosenblatt's Superposition Catastrophe)




Specific Solutions and Design Inspiration

- The binding problem (or Rosenblatt's Superposition Catastrophe)




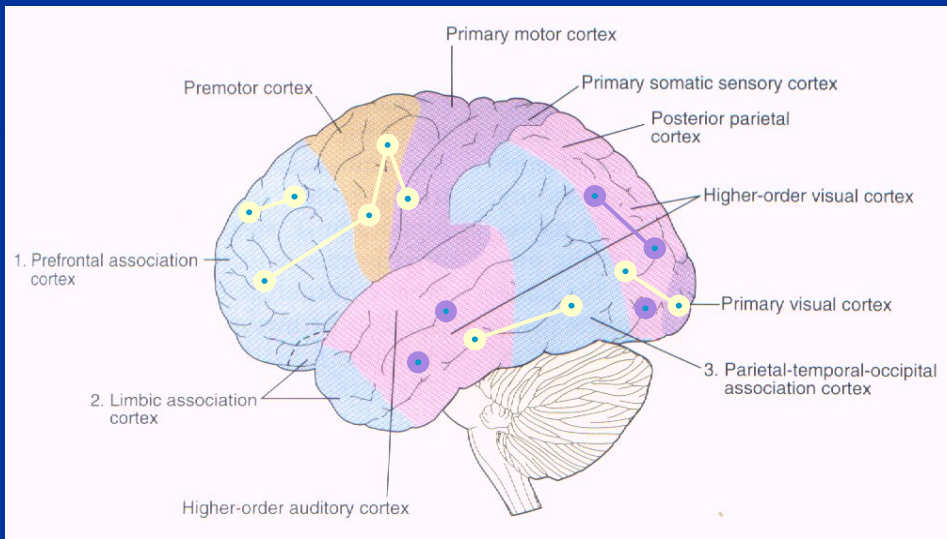
Specific Solutions and Design Inspiration

- The binding problem (or Rosenblatt's Superposition Catastrophe)
 - Brain networks decompose visual information into maps of different attributes
 - Brain networks recombine attributes
 - Information about relationships between attributes is lost
- Neuroscience observation
 - Synchronous neural activity reflecting global visual input properties (Gray & Singer)
- Inspiration
 - Modulate classification information with timing information containing the lost relationship information



Synchronous Networks


- Allows coordination among weakly or indirectly connected units



Networks of oscillatory units

- Observations by Altmann, Eckhorn and Singer (1986), Gray and Singer (1987) on synchronous oscillations in the cortex
- Von der Malsburg first proposed a network of oscillatory units to solve the “cocktail party problem” (1986)
- Cheng, Wang and Liu: applications to image segmentation (2000).

Network connectivity



Feedforward

Lateral

Feedback

Objective Function

$$E = \langle \mathbf{y}W\mathbf{x} - \frac{1}{2}\mathbf{y}^2 - \frac{1}{2} \sum_n \mathbf{W}_n^2 + \frac{1}{2} \lambda S(\mathbf{y}) \rangle_{\mathcal{E}}$$

$$S(\mathbf{y}) = N (\langle y_n^2 \rangle_{\mathcal{N}} - \langle y_n \rangle_{\mathcal{N}}^2) = \sum_{n=1}^N y_n^2 - \frac{1}{N} \left(\sum_{n=1}^N y_n \right)^2$$

- Reward faithfulness of representation (alignment between y and $W\mathbf{x}$)
- Constrain y and W
- Reward sparseness

Objective function

- Imposing synaptic normalization

$$E = \langle \mathbf{y}W\mathbf{x}^T + \frac{1}{2}\lambda S(\mathbf{y}) - \frac{1}{2}\mathbf{y}^2 \rangle_{\mathcal{E}}$$

Extension to oscillatory units

$$p_n = x_n e^{i\phi_n}, \quad q_n = y_n e^{i\theta_n}$$

Use complex extension of energy

$$\mathcal{C}(E) = \mathbf{q}W\bar{\mathbf{p}} + \frac{1}{2}\lambda S(\mathbf{q}) - \frac{1}{2}\mathbf{q}\bar{\mathbf{q}}$$

$$E_s = E + \beta \operatorname{Re}[\mathcal{C}(E)]$$

Variance

$$S(\mathbf{q}) = \sum_{n=1}^N q_n \overline{q_n} - \frac{1}{N} \left(\sum_{n=1}^N q_n \right) \left(\sum_{n=1}^N \overline{q_n} \right)$$

Gradient ascent

$$E_S = \left\langle \sum_{n,m} y_n W_{nm} x_m (1 + \beta \cos \Psi_{nm}) - \alpha \sum_n y_n^2 (1 + \beta) - \gamma \sum_{n \neq m} y_n y_m (1 + \beta \cos \Phi_{nm}) \right\rangle \varepsilon$$

$$\Psi_{nm} = \theta_n - \phi_m$$

$$\Phi_{nm} = \theta_n - \theta_m$$

Upper unit phase Lower unit phase

$$\Delta y_n \sim \frac{\partial E_S}{\partial y_n}$$

$$\Delta \theta_n \sim \frac{1}{y_n} \frac{\partial E_S}{\partial \theta_n}$$

Gradient ascent: update equations

$$\Delta y_n \sim \sum_j W_{nj} x_j [1 + \cos(\phi_j - \theta_n)] - \alpha y_n$$

$$- \gamma \sum_k y_k [1 + \beta \cos(\theta_k - \theta_n)]$$

$$\Delta \theta_n \sim \beta \sum_j W_{nj} x_j \sin(\phi_j - \theta_n)$$



$$- \beta \gamma \sum_k y_k \sin(\theta_k - \theta_n)$$

$$\Delta \phi_n \sim \sum_j W_{jn} y_j \sin(\theta_j - \phi_n)$$

$$\Delta W_{ij} \sim y_i x_j [1 + \beta \cos(\phi_j - \theta_i)]$$

Modified Hebbian update

Network behavior: inputs







Amplitude response over 10 oscillation cycles

Hebbian learning rule is applied to weights after the network settles (~250 iterations)

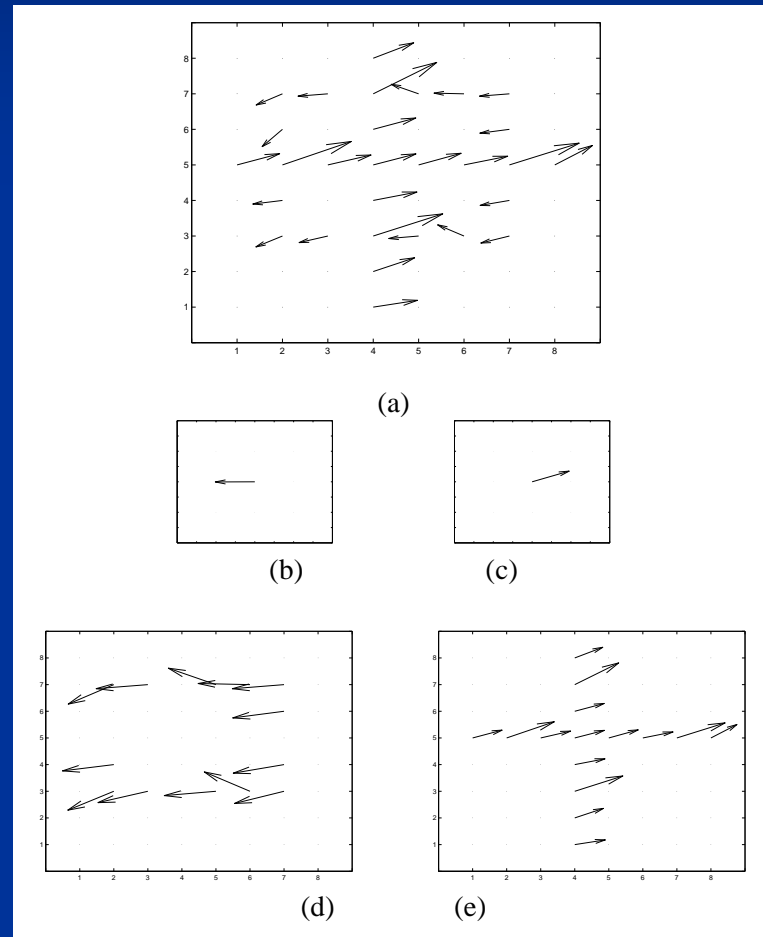
This procedure is repeated 1000 times.

Network behavior after training: Superposition of inputs

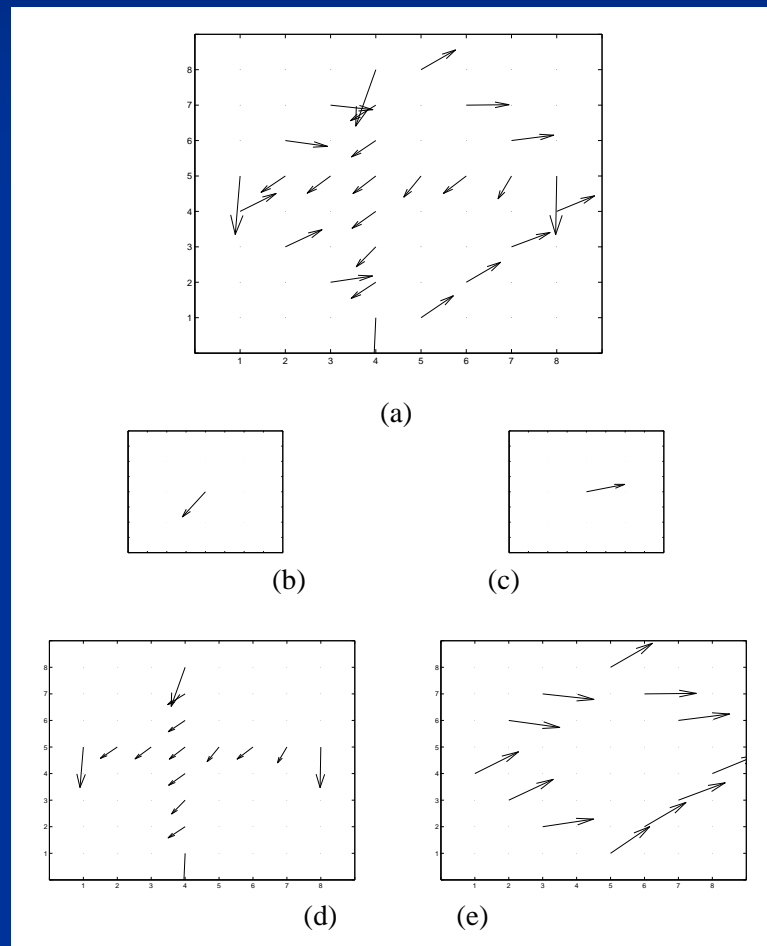



Blue: upper layer
Red: lower layer

Phase behavior



Phase behavior




Stimulus representation through temporal coding: Winnerless Competition

- Computation is spatio-temporal
- The system never reaches a fixed point
- However, responses to different stimuli are repeatable

Background: attractor auto-associative networks

- Computation is based on **fixed points**
- Dynamics is a nuisance: the answer is obtained **when** the system reaches a stable point of the dynamics

Hopfield Model



Non-linear (Boolean) Units

Hopfield, 1982, 1984

Hopfield Model

Providing that $T_{ij} = T_{ji}$, it is possible to construct an energy function:

$$E = -\frac{1}{2} \sum_{i,j} T_{ij} V_i V_j$$

That is minimized by the evolution of the network:

$$\Delta E = -\Delta V_i \sum T_{ij} V_j < 0$$

Being a quadratic form, it will reach a (local) minimum. Similarly, a Hebbian learning rule can be derived by prescribing minimization over weight updates:

$$\Delta T_{ij} \sim V_i V_j \quad \Rightarrow \quad \Delta E = -\Delta T_{ij} V_i V_j < 0$$

But synaptic normalization is required during learning so that $|T|$ stays bounded.

Hopfield Model


But the storage capacity is severely limited; for a network with N neurons, the number of patterns that can be stored without errors is:

$$M = \frac{N}{2 \log N}$$



If a small amount of noise is accepted, the storage capacity improves to:

$$M = 0.14N$$

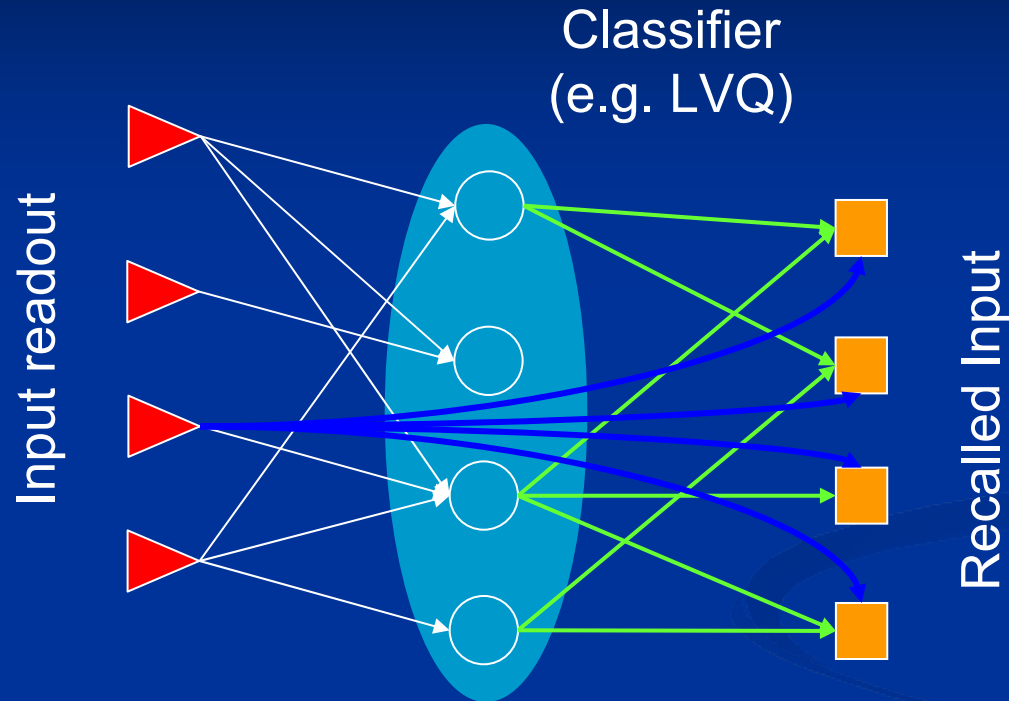
In fact, there is a *phase transition* at this threshold.



Training set






Hopfield and classification



The codebook vectors $\{\vec{W}_i\}$ are the stored patterns, and can be used with symmetric connections, or learned in a third layer, to reconstitute the memory, regardless of the network implementation of the classifier.


Attractor vs Winnerless Competition



Linear stability analysis of a fixed point

Global stability

Winnerless Competition



$$\tau_1 \frac{dx_i(t)}{dt} = f[x_i(t)] - y_i(t) - z_i(t)[x_i(t) - \nu] + 0.35 + S_i,$$

$$\frac{dy_i(t)}{dt} = x_i(t) - by_i(t) + a,$$

$$\tau_2 \frac{dz_i(t)}{dt} = \sum_j g_{ji} G[x_j(t)] - z_i(t).$$


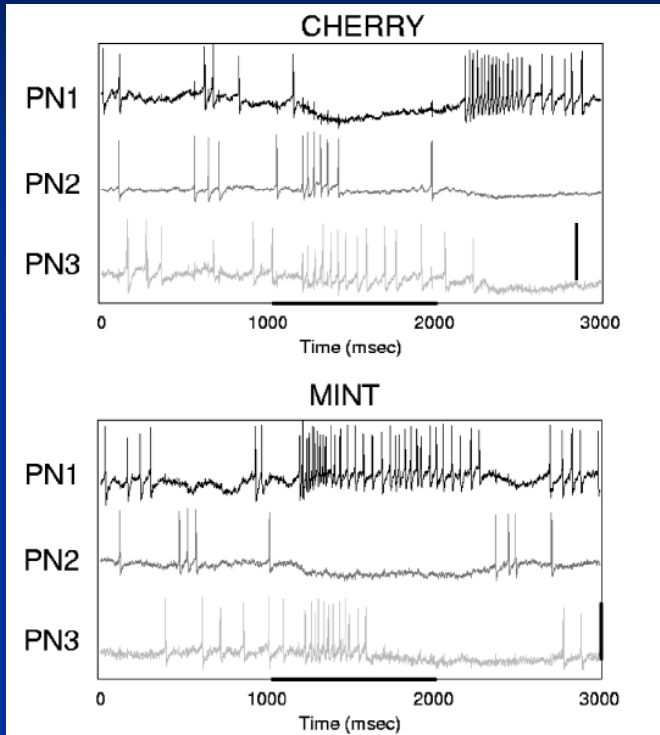
FitzHugh-Nagumo (reduced Hodgkin-Huxley):
Other non-linearities (e.g. Integrate-and-Fire).
 $g_{ij} \neq g_{ji}$ implies no energy functional

Abarbanel, 2001

Winnerless Competition

Odor response (insect)


Winnerless model



Abarbanel, 2001

Laurent, 2001

Activity proceeds along heteroclinic orbits
(those connecting saddle points)



Introduction to temporal analysis of neural recordings: spike metrics

- What is the right way to measure “distances” between spike trains?
- Neuroscientists often rely on just counting the total number of spikes
- Precise temporal information is known to be relevant in some cases
- Echo-location
- Auditory processing
- How to include time in the distance?
- A very popular method was introduced by Victor & Purpura (1997)
- Spike train metrics

Spike Count

Commonly used, but ...

spikes



All these trains, from the same or different units, cannot be distinguished


Victor-Purpura Metric

- Adapted from for genetic sequence alignment (Sellers, 1984)
- Accounts for the total number of spikes
- Considers temporal locality of spikes
- Operations to measure the distance between spike trains as **costs**
 - Spike count: **insert/delete** cost
 - Spike distortion: **time shift** cost

Victor-Purpura Metric


- Insert/Delete: the cost to pay for having more or less spikes
 - Cost = 1
- Shifting spikes for the same neuron: the cost to pay for moving the same spike back or forth in time
 - Cost = $q |\Delta t|$
 - q is a “temporal resolution” parameter
 - $q=0$ is equivalent to spike count
 - Increasing q is increasing sensitivity
 - Obtain distance, $D(q)$
- Shifting spikes across neurons:
 - Multi-neuronal recording is considered to be a sequence of labeled events
 - Add a rule that sets the cost of changing the labels: obtain $D(q,k)$
 - k is the cost to pay for shifting in time the same spike in different neurons
 - $k = 0$: origin of the spike is irrelevant
 - If $k < 2$ spikes can be switched between neurons
 - If $k > 2$ they can't, as it is cheaper to insert and delete a spike in the same neuron
- The Sellers algorithm requires only one forward sweep

Victor-Purpura Metric



A diagram of a sequence of elementary steps that transforms spike train A into spike train B. Each rectangle represents one spike, and the line that they rest on denotes time. Each elementary step is one of three types: deletion of a spike (deleted spike shown in red), insertion of a spike (inserted spike shown in green), or shifting a spike in time (blue arrows).


Victor-Purpura Metric



Cost: 1

~~Delete~~


Victor-Purpura Metric



Cost: $q|\Delta t|$

Shift within Neurons

Victor-Purpura Metric



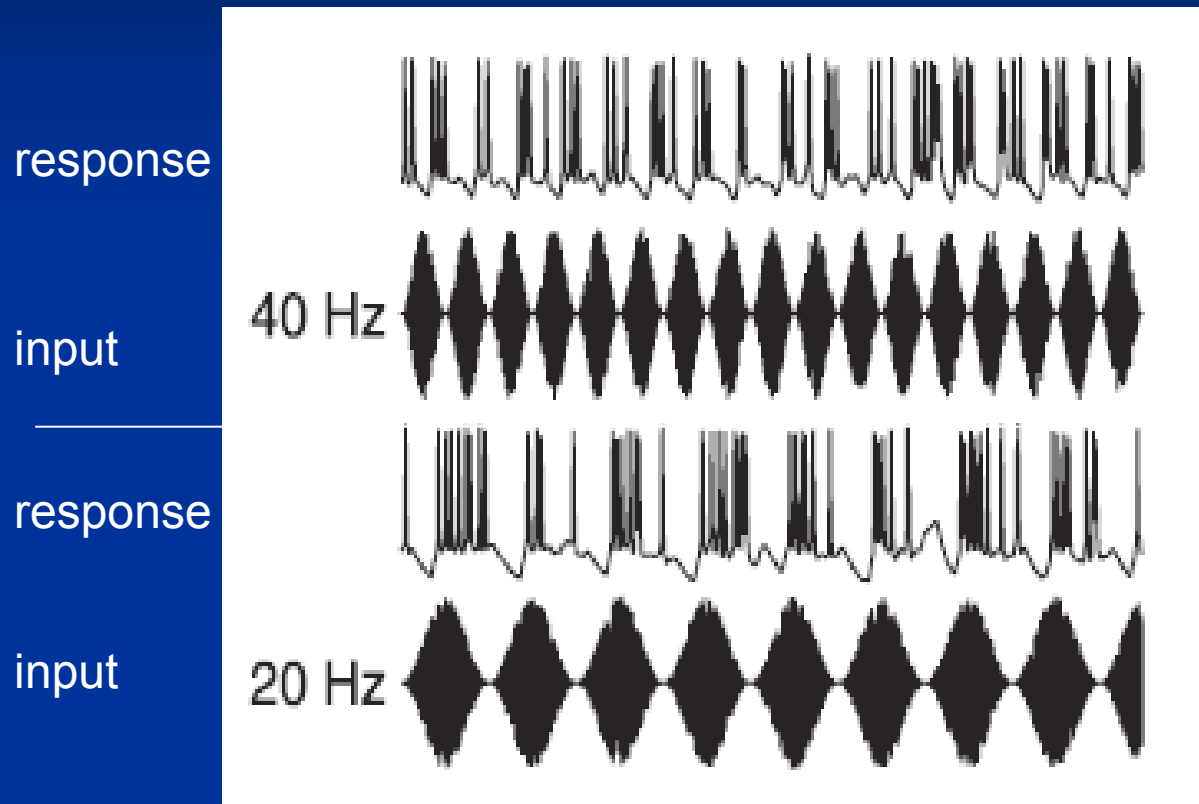
Cost: k

Shift Between Neurons

Application

- Discrimination of amplitude-modulated auditory stimuli in locusts (Wohlgemuth & Ronacher, 2007)
- Classification of neural responses based on spike metrics
 - Provides information about
 - The necessary evaluation time window that a neuron uses
 - The optimal temporal resolution of processing (q parameter)
 - Yields clues about coding principles
 - Helps us understand respective contributions of spike counts vs. spike timing

Application



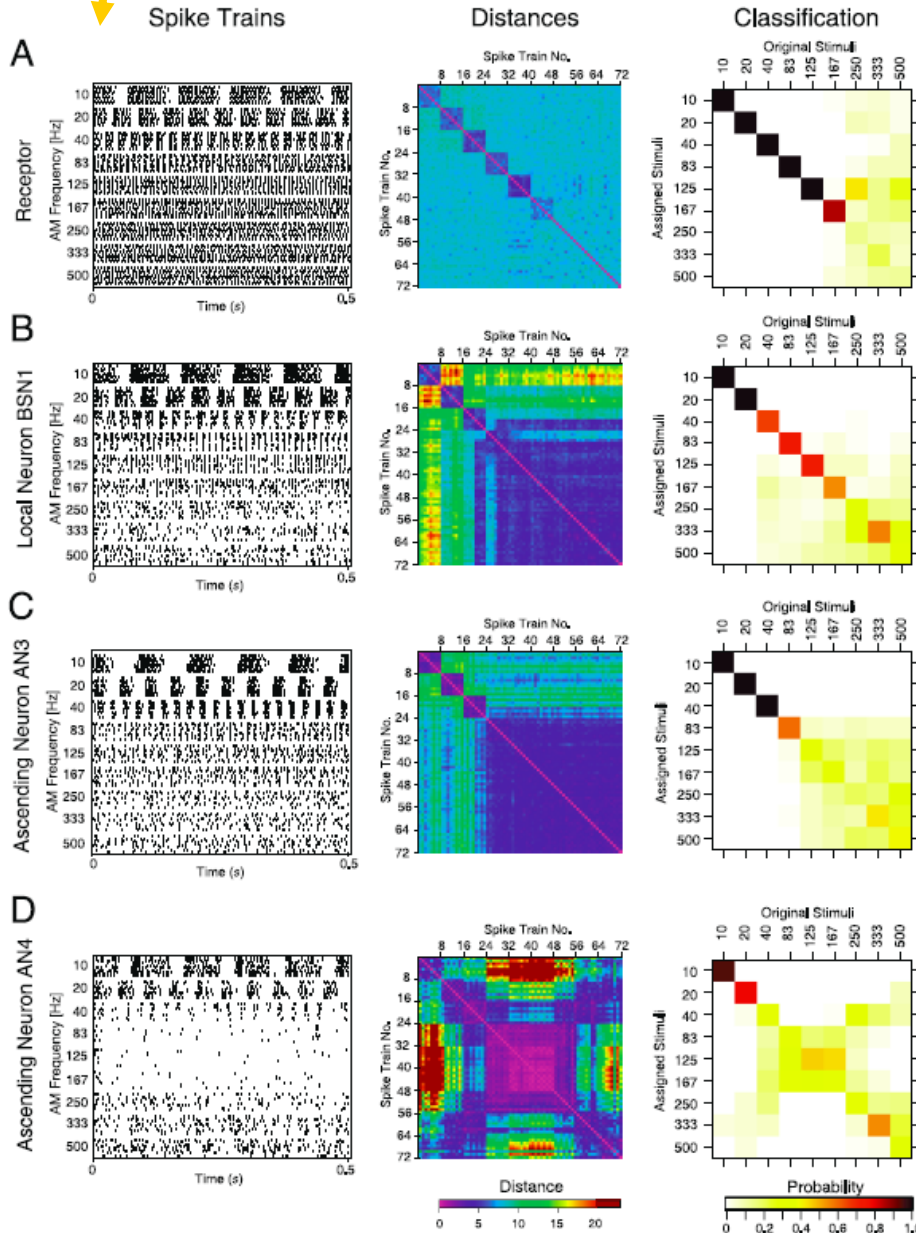
Three classes
of neurons

Auditory receptor
neurons

Local neurons

Ascending
neurons

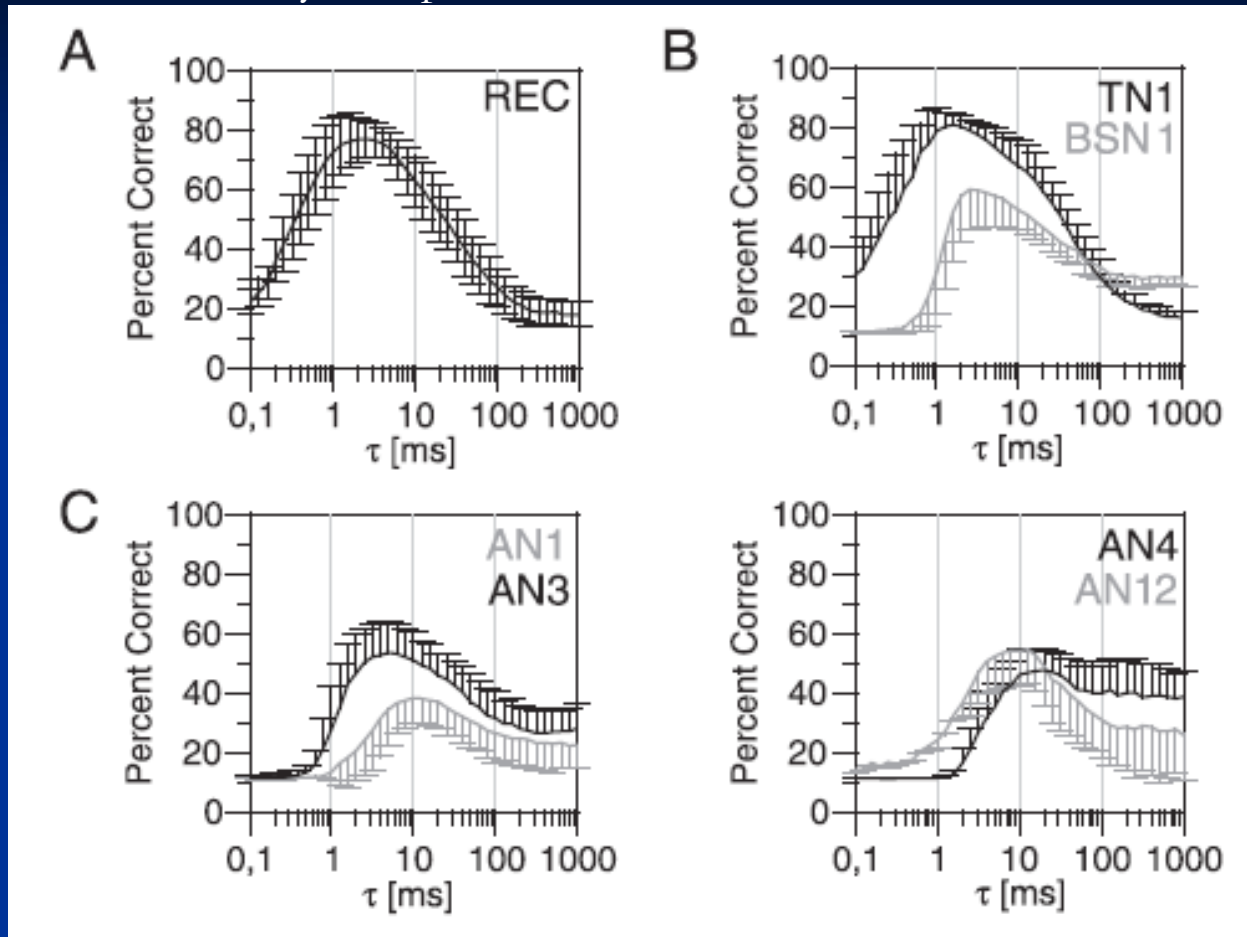
Spike trains for
different stimuli



Classification
based on
spike metric

Wohlgemuth
& Ronacher,
2007

Rec = Auditory Receptor neuron



TN = local neurons with primary-like responses
BSN = bisgemental neuron (interneuron)

AN = ascending interneuron

- Classification accuracy as a function of temporal resolution, i.e. the q parameter

Discrepancies between models and actual neurophysiological data

- Models are idealized
- Neural data are noisy
 - E.g. there may be gaps in a neuron's firing (cycle skipping)
 - How can models be built to accommodate such imprecision?
 - See “Model this! Seven empirical phenomena missing in the models of cortical oscillatory dynamics,” Danko Nikolic, IJCNN 2009.

Relevance of this research to our daily lives

- NY Times article, 5/5/09 “Ear Plugs to Lasers: The Science of Concentration”
- Based on "Driving fast-spiking cells induces gamma rhythm and controls sensory responses." J. A. Cardin, et al. Nature, 2009
- Interplay between bottom-up and top-down processes:
 - Neurons in pre-frontal cortex (planning center) oscillate in unison
 - pre-frontal cortex can override sensory input
 - However, significant concentration is required to override a strong input like a TV commercial
 - Limit multitasking; start your day with the most important task; meditate; build a stimulus shelter
 - Implications for our own lives, and for education

Conclusions

- We have reviewed a rich area of research and exploration
 - Tried to achieve a balance between breadth & depth
- Several questions still open for investigation
 - Using oscillatory networks on more realistic input data, scaling up the network sizes
 - Combining different sensory modalities
 - Closing the perception-action loop
- Opportunities for interaction with neuroscience
 - Analysis of spike train recordings
 - Developing explanatory and predictive models
- Thanks for your attention. Feedback is always welcome!