Ch 7. Network Models

Seong Kyun Jung

Seoul National University

June 1, 2023

Table of Contents

Table of Contents

Nonlinearity through rectification

- Recall the equation for the firing rates in the linear model

$$\tau_r \frac{dv}{dt} = -v + Wu + Mv$$

  where u was the firing rate vector of the input layer, v the vector of the output layer.
- Negative firing rates in a linear model?
- The simplest way to fix this problem is to use rectification:

$$F(Wu + Mv) = [Wu + Mv]_+$$

  (in essence, simply overriding negative components with zeros)

Nonlinearity through rectification

- We investigate properties of the rectified continuous model

$$\tau_r \frac{dv(\theta)}{dt} = -v(\theta) + \left[h(\theta) + \frac{\lambda_1}{\pi}\int_{-\pi}^{\pi} d\theta' \cos(\theta - \theta') v(\theta')\right]_+$$

and try to compare its characteristics to the linear model

$$\tau_r \frac{dv(\theta)}{dt} = -v(\theta) + h(\theta) + \frac{\lambda_1}{\pi}\int_{-\pi}^{\pi} d\theta' \cos(\theta - \theta') v(\theta')$$

(no real reason for taking $M = \cos$ is given; probably out of expedience)

Characteristic 1 - Nonlinear Amplification

- Compared to the linear model, the amplification is smoother and smaller.
- The restriction $\lambda_1 < 1$, necessary in the linear model, is no longer necessary.

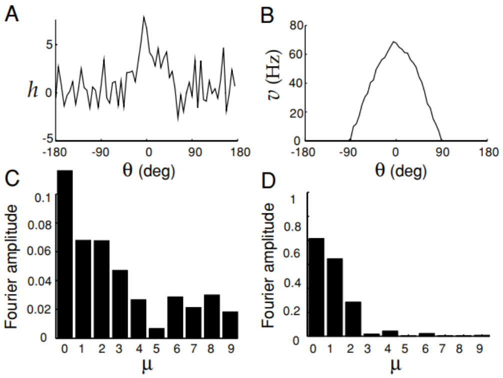

Figure 7.9 Selective amplification in a recurrent network with rectification. (A) The input $h(\theta)$ of the network plotted as a function of preferred angle. (B) The steady-state output $v(\theta)$ as a function of preferred angle. (C) Fourier transform amplitudes of the input $h(\theta)$. (D) Fourier transform amplitudes of the output $v(\theta)$. The recurrent coupling took the form of equation 7.33 with $\lambda_1 = 1.9$.

Characteristic 2 - Nonlinear Input Selection

- The output of a rectified network tends to be unimodal; thus unlike the linear network, superposition no longer holds

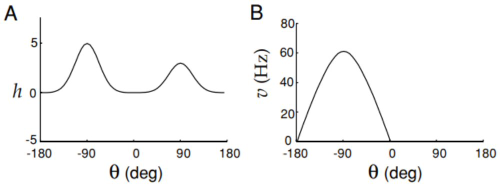

Figure 7.12 Winner-takes-all input selection by a nonlinear recurrent network. (A) The input to the network consisting of two peaks. (B) The output of the network has a single peak at the location of the higher of the two peaks of the input. The model is the same as that used in figure 7.9.

Characteristic 3 - Sustained Activity

- Similarly to the case of input integration, where a linear network showed "memory", rectified networks exhibit this behavior as well.

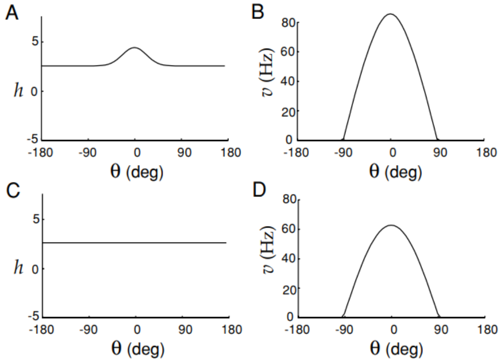

Figure 7.14 Sustained activity in a recurrent network. (A) Input to the neurons of the network consisting of localized excitation and a constant background. (B) The activity of the network neurons in response to the input of panel A. (C) Constant network input. (D) Response to the constant input of panel C when it immediately followed the input in A. The model is the same as that used in figures 7.9, 7.12, and 7.13.

Maximum Likelihood and Network Recoding

- Pouget et al. (1998): interpret the location of the characteristic pattern (i.e. the peak of the output) as the neuron's recoding of noisy stimulus
- Recurrent networks can be used to improve recoding performance of simple methods like vector decoding

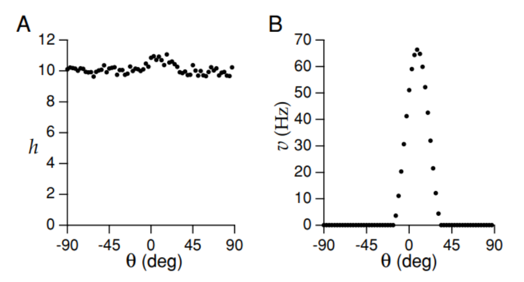

Applications - Models of the Primary Visual Cortex

- Ben-Yishai, Bar-Or, Sompolinsky (1995): simulate behavior of simple cells:

$$\tau_r \frac{dv(\theta)}{dt} = -v(\theta) + [h + \int_{\frac{\pi}{2}}^{\frac{\pi}{2}} \frac{d\theta'}{\pi}(-\lambda_0 + \lambda_1 \cos(2(\theta - \theta')))v(\theta')]_+$$

- Chance, Nelson, Abbott (1999): simulate behavior of complex cells when $\lambda_1 \approx 1$:

$$\tau_r \frac{dv(\phi)}{dt} = -v(\phi) + [h(\phi) + \frac{\lambda_1}{2\pi} \int_{-\pi}^{\pi} d\phi' v(\phi')]_+$$

Network Stability

- In the presence of constant input, what happens to the network?
- The most intuitive conjecture is that the network converges to a steady state; this is not guaranteed but does often hold.
- **Lyapunov functions** help in this analysis.

Network Stability - the linear case

- We analyze the linear model case with instantaneous adaptation; that is,

$$\tau_s \frac{dI}{dt} = -I + h + M * F(I)$$

where I is the synaptic current, $F(I)$ is the firing rate vector

- The function L satisfies $\frac{dL}{dt} < 0$ whenever $\frac{dI}{dt} \neq 0$:

$$L(\mathbf{I}) = \sum_{a=1}^{N_v} \left(\int_0^{I_a} dz_a \, z_a F'(z_a) - h_a F(I_a) - \frac{1}{2} \sum_{a'=1}^{N_v} F(I_a) M_{aa'} F(I_{a'}) \right) \quad (7.40)$$

- Since $\frac{dL}{dt}$ depends linearly on $(\frac{dI_a}{dt})^2$, if L is bounded from below, convergence to a steady state is guaranteed.

Associative Memory

- Short-term memory is associated with sustained activity of the neurons (input integration, rectified recurrent network)
- What happens in long-term memory? A similar input is given, then we "recall" something
- "In a network associative memory, recurrent weights are adjusted so that the network has a set of discrete fixed points identical to the patterns of activity that represent the stored memories" (p. 261)

Constructing an Associative Memory Network

- Memory pattern: v^m where m is the index of the memory (so v^m is an entire vector of firing rates)
- The network's goal is to modify the recurrent weights s.t. the v^m become **basins of attraction**.
- Suppose the network evolves according to the delayed evolution + nonlinear model, with the assumption $h = 0$:

$$\tau_r \frac{dv}{dt} = -v + [Mv]_+$$

- A necessary condition for recall is that the v^m must all be fixed points:

$$v^m = [Mv^m]_+, m = 1, 2, ..., N_{mem}$$

Computation of Associative Memory Network

- Big Question: How many patterns can be stored in a memory network?
- We make some simplifying assumptions:
  - Suppose the patterns consist of "active" neurons (with firing rate c) and "inactive" neurons (with firing rate 0)
  - Suppose each pattern consists of αN_v active neurons

Computation of Associative Memory Network

- We try to construct conditions where a certain number of patterns can be supported by a network.
- Suppose we have a matrix K where

$$K \cdot v^m = \lambda v^m$$

  for $m = 1, 2, ..., N_{mem}$.
- Take $M = K - \frac{n^T n}{\alpha N_v}$ where $n = (1, 1, ..., 1)$.
- A sufficient condition for

$$F(\lambda v^m - cn) = F(M \cdot v^m) = v^m$$

  is that for inactive neurons,

$$F(-c) = 0$$

  and for active neurons,

$$F(c(\lambda - 1)) = c$$

Computation of Associative Memory Network

- It is relatively easy to find an activation function F that satisfy the two conditions; so it remains to build a matrix K
- The matrix

$$K = \frac{\lambda}{c^2 \alpha N_v (1 - \alpha)} \Sigma_{n=1}^{N_{mem}} v^n (v^n - \alpha cn)$$

  satisfies $K \cdot v^m \approx \lambda v^m$, using the fact that the active neurons in different memories are uncorrelated.
- Thus in this simplified model, $N_{mem} \approx N_v$.

Computation of Associative Memory Network

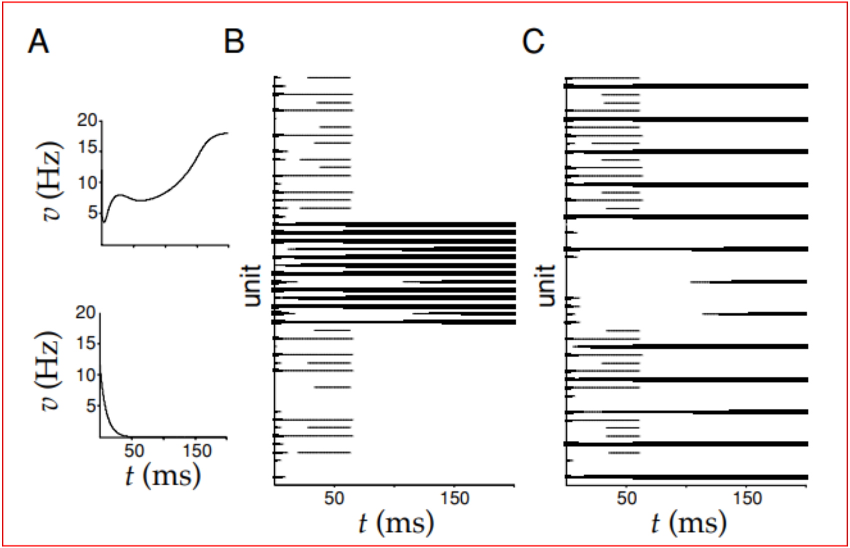

Table of Contents

Table of Contents

7.5. Excitatory-Inhibitory Networks

- Until now, the matrix M describing the recurrent weights was assumed to be symmetric.
- We now relax this assumption; now, the matrix is no longer symmetric.
- Simplifying into a model where all excitatory neurons and inhibitory neurons are homogeneous, we have two equations which describe the firing rates:

$$\tau_E \frac{dv_E}{dt} = -v_E + [M_{EE} v_E + M_{EI} v_I - \gamma_E]_+$$

$$\tau_I \frac{dv_I}{dt} = -v_I + [M_{II} v_I + M_{IE} v_E - \gamma_I]_+$$

Phase-Plane Methods and Stability Analysis

- A big question is whether such models still exhibit constant steady-state behavior.
- In the simplified case, a phase plane describes the evolution of the two parameters v_E, v_I and their steady-state behavior.

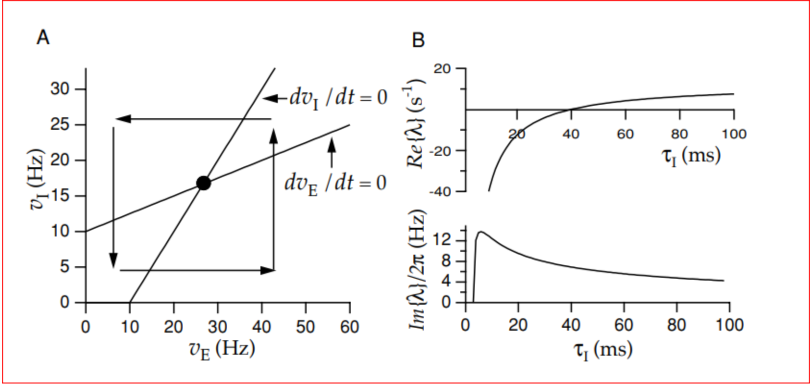

Phase-Plane Methods and Stability Analysis

- Since the system is a system of first-order linear ODEs near the fixed point, the coefficients given by the matrix

$$\begin{pmatrix} \frac{M_{EE}-1}{\tau_E} & \frac{M_{EI}}{\tau_E} \\ \frac{M_{IE}}{\tau_I} & \frac{(M_{II}-1)}{\tau_I} \end{pmatrix}$$

  determine the stability of the system.
- In particular, the real parts of the eigenvalues of this matrix determine exponential attraction or repulsion around the fixed point.

Phase-Plane Methods and Stability Analysis

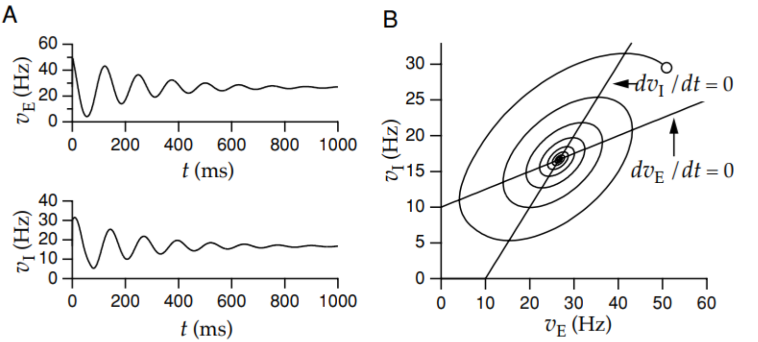

Figure 7.18 Activity of the excitatory-inhibitory firing-rate model when the fixed point is stable. (A) The excitatory and inhibitory firing rates settle to the fixed point over time. (B) The phase-plane trajectory is a counterclockwise spiral collapsing to the fixed point. The open circle marks the initial values $v_E(0)$ and $v_I(0)$. For this example, $\tau_I = 30$ ms.

Phase-Plane Methods and Stability Analysis

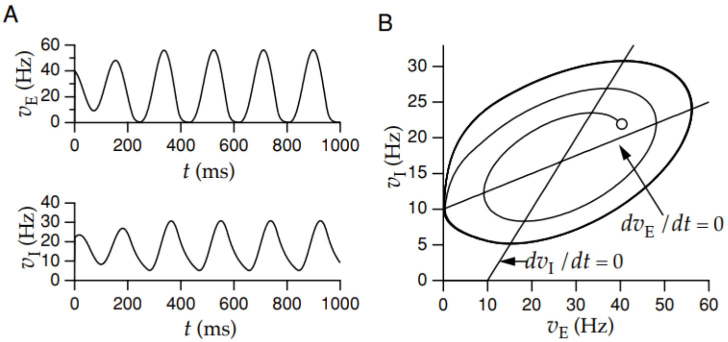

Figure 7.19 Activity of the excitatory-inhibitory firing-rate model when the fixed point is unstable. (A) The excitatory and inhibitory firing rates settle into periodic oscillations. (B) The phase-plane trajectory is a counterclockwise spiral that joins the limit cycle, which is the closed orbit. The open circle marks the initial values $v_E(0)$ and $v_I(0)$. For this example, $\tau_I = 50$ ms.

Phase-Plane Methods and Stability Analysis

- Note that without rectification, there would be no limit cycle in Fig 7.19.
- The changes in the parameters causes the behavior of the system to change (Hopf bifurcation); in the Figures this was caused by τ_I.

Table of Contents

Applications of Oscillatory Networks

- Li, Hopfield (1989): modeled the olfactory bulb's oscillatory behavior using an excitatory-inhibitory network with $M_{EE} = M_{II} = 0$
- Li, Dayan (1999): used excitatory-inhibitory dynamics to construct recurrent networks which :
  - allows for more amplification
  - does not show tuned (eg. angle-dependent) output with untuned (eg. angle-independent) input

Applications of Oscillatory Networks

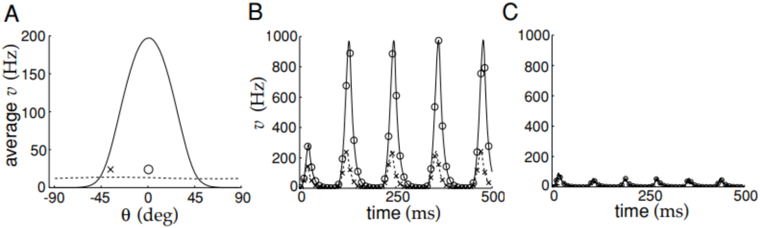

Figure 7.23 Selective amplification in an excitatory-inhibitory network. (A) Time-averaged response of the network to a tuned input with $\Theta = 0^\circ$ (solid curve) and to an untuned input (dashed curve). Symbols "o" and "x" mark the 0° and -37° points seen in B and C. (B) Activities over time of neurons with preferred angles of $\theta = 0^\circ$ (solid curve) and $\theta = -37^\circ$ (dashed curve) in response to a modulated input with $\Theta = 0^\circ$. (C) Activities of the same units shown in B to a constant input. The lines lie on top of each other, showing that the two units respond identically. The parameters are $\tau_E = \tau_I = 10$ ms, $h_I = 0$, $M_{EI} = -\delta(\theta - \theta')/\rho_\theta$, $M_{EE} = (1/\pi\rho_\theta)[5.9 + 7.8\cos(2(\theta - \theta'))]_+$, $M_{IE} = 13.3/\pi\rho_\theta$, and $M_{II} = 0$. (After Li and Dayan, 1999.)

Table of Contents

Method 1. The Boltzmann Machine

- We consider a model where a neuron takes one of two states (fire or not fire), which is determined stochastically.
- Suppose the state of unit a is determined by

$$I_a(t) = h_a(t) + \Sigma_{a'=1}^{N_v} M_{aa'} v_{a'}(t)$$

$$P[v_a(t + \Delta t) = 1] = \frac{1}{1 + exp(-I_a)}$$

Note that F is the sigmoid function, and that the decision-making process is stochastic.

Solving the Boltzmann Machine

- The equation describes a **Markov chain**, where the state at $t + \Delta t$ depends only on the state of t, and not more previous states.
- (Glauber dynamics) v "converges" to the probability distribution

$$P[v] = \frac{exp(-E(v))}{Z}, Z = \Sigma_v exp(-E(v))$$

where

$$E(v) = -h \cdot v - \frac{1}{2} v^T M v$$

Method 2. Mean-field Approximation

- The mean-field approximation approximates the Boltzmann machine
- Instead of I depending on v, thus making I stochastic, we can think of the case where I is deterministic by

$$\frac{dI}{dt} = -I + h + M \cdot F(I)$$

with $F =$(sigmoid), and stochasticity is introduced by

$$P[v_a = 1] = F(I_a)$$

- The probability distribution for the entire vector v is

$$Q[v] = \Pi_{a=1}^{N_v} F(I_a)^{v_a} * (1 - F(I_a))^{1-v_a}$$

Relationship between the two methods

- It can be shown that the Lyapunov function of eq 7.40 can be expressed as
$$L(I) = D_{KL}(Q, P) + K$$
where K is constant and D_{KL} is the Kullback-Leibler divergence.
- Thus, the dynamics of eq 7.39 can be described as trying to minimize the difference between Q and P.
- The Boltzmann machine allows us to determine how changing M affects the distribution of the output. Thus, it will continue to be used in later chapters (eg. how do neurons "learn"?)

Thank You