



Networks of Neurons (1 of 2)


Angus Chadwick

School of Informatics, University of Edinburgh, UK

Computational Neuroscience (Lecture 11, 2025/2026)

Outline of Lecture


- Introduction to networks
- Spiking and firing rate network models
- Feedforward and recurrent networks
- Excitatory-inhibitory networks



The “rainbow” (an image of a real network)

Outline of Lecture

- **Introduction to networks**
- Spiking and firing rate network models
- Feedforward and recurrent networks
- Excitatory-inhibitory networks




The “brainbow” (an image of a real network)


Networks of Neurons

- Neurons in the brain connect to form *networks*
- Networks of neurons *transform* external inputs into a pattern of spiking output
- This input-output transformation may be viewed as a *computation*
- These computations are thought to be the neural basis for *perception, decision making, memory, motor control*, etc.
- In this lecture we focus on models of *biological* networks of neurons (in contrast to artificial neural networks)

Networks of the Brain



Networks Dynamics and Computation



What do we want to know? (1 of 2)


- *What* are the properties of neural networks? (experimental/descriptive modelling question)
- *How* do networks generate their activity patterns and associated computations? (mechanistic question)
- *Why* are networks the way they are? What are problems do they solve and what constraints do they solve them under? (“teleological” question, ultimately appeals to evolution)

Marr's 3 Levels of Analysis

We are mainly taking a “bottom-up” approach to understanding networks...

But remember Marr's three levels
- we saw examples of “top-down” approaches with e.g. sparse coding, predictive coding, etc.

"...trying to understand perception by studying only neurons is like trying to understand bird flight by studying only feathers. It cannot be done" - David Marr (1982/2010, p. 27)



What do we want to know? (2 of 2)


- How do the *connectivity* and *single-neuron properties* relate to the *global properties* (dynamics, input-output transformation, etc.)?
- How do networks *learn*? What are the learning rules? What are the initial conditions?
- What aspects of biological networks can be distilled into *artificial applications* (AI, neural hardware, etc.)? What aspects can we ignore?
- When do networks go wrong? How does this relate to *disease*?

What Tools do we Use?

- **Experiment:** measurements of connectivity, neuronal activity, etc.
- **Mathematical analysis:** on simplified or reduced models (e.g., firing rate models, mean field theory)
- **Numerical simulation:** often of more biologically-plausible models
- **Machine learning:** e.g., by training an artificial neural network to solve a task and studying the learned network solution
- **Neural data analysis tools:** e.g., inference of network dynamics from recordings of neural activity

Outline of Lecture

- Introduction to networks
- **Spiking and firing rate network models**
- Feedforward and recurrent networks
- Dynamics of recurrent networks
- Excitatory-inhibitory networks



The “brainbow” (an image of a real network)

Spiking Network Models: An Example

- Consider a leaky integrate and fire neuron receiving inputs through a set of synapses:

Membrane potential dynamics:

$$\tau_m \frac{dV}{dt} = -(V - E_m) - I_{syn}(t) / g_m$$

Spike-reset rule:

$$\text{if } V(t) \geq V_{\text{threshold}}, \text{ set } V(t) = V_{\text{reset}}$$

Synaptic input current:

$$I_{syn}(t) = \sum_{j=1}^{N_{\text{synapse}}} g_{syn,j}(t)(V(t) - E_{syn,j})$$

Spiking Network Models: An Example

- Consider a leaky integrate and fire neuron receiving inputs through a set of synapses:

Membrane potential dynamics:
$$\tau_m \frac{dV}{dt} = -(V - E_m) - I_{syn}(t)/g_m$$

Spike-reset rule:
$$\text{if } V(t) \geq V_{\text{threshold}}, \text{ set } V(t) = V_{\text{reset}}$$

Synaptic input current:
$$I_{syn}(t) = \sum_{j=1}^{N_{\text{synapse}}} g_{syn,j}(t)(V(t) - E_{syn,j})$$

- Now consider multiple such neurons coupled together via these synapses:

Membrane potential of neuron i :
$$\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) - \sum_{j \neq i} g_{syn,ij}(t)(V_i - E_{syn,ij})/g_{m,i}$$

Spiking Network Models

- A spiking network consists of:
 - 1) A **neuron model** (e.g., Hodgkin-Huxley, leaky integrate and fire, etc.)
 - 2) A **synapse model** (e.g., delta synapse, exponential synapse, etc.)
 - 3) A set of **parameters** for the synaptic and membrane conductances, etc.
- Mathematically, a spiking network model is a set of **coupled differential equations**
- For example, a network of Hodgkin-Huxley neurons has $4N$ neuron equations, and $O(N^2)$ synapse equations (depending on synapse model, could be more...)
- We typically make simplifications to reduce the number/complexity of equations

A Simple Spiking Network






Figure 5.20 Two synaptically coupled integrate-and-fire neurons. (A) Excitatory synapses ($E_s = 0$ mV) produce an alternating, out-of-phase pattern of firing. (B) Inhibitory synapses ($E_s = -80$ mV) produce synchronous firing. Both model neurons have $E_L = -70$ mV, $V_{th} = -54$ mV, $V_{reset} = -80$ mV, $\tau_m = 20$ ms, $r_m \bar{g}_s = 0.05$, $P_{max} = 1$, $R_m I_e = 25$ mV, and $\tau_s = 10$ ms.

A More Complex Spiking Network



spike raster



neuron trace

[Izhikevitch, 2003]

Spiking Network Models – Current-Based Synapses

- The leaky integrate and fire model with **conductance-based synapses** is:

$$\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) - \sum_{j \neq i} g_{syn,ij}(t)(V_i - E_{syn,ij}) / g_{m,i}$$

Spiking Network Models – Current-Based Synapses

- The leaky integrate and fire model with **conductance-based synapses** is:

$$\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) - \sum_{j \neq i} g_{syn,ij}(t)(V_i - E_{syn,ij}) / g_{m,i}$$

- We often simplify these synaptic currents by **neglecting their voltage-dependence**:

$$\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) + \sum_{j \neq i} w_{ij} g_{syn,ij}(t) \quad w_{ij} = -(\langle V_i \rangle - E_{syn,ij}) / g_{m,i}$$

Spiking Network Models – Current-Based Synapses

- The leaky integrate and fire model with **conductance-based synapses** is:

$$\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) - \sum_{j \neq i} g_{syn,ij}(t)(V_i - E_{syn,ij}) / g_{m,i}$$


- We often simplify these synaptic currents by **neglecting their voltage-dependence**:

$$\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) + \sum_{j \neq i} w_{ij} g_{syn,ij}(t) \quad w_{ij} = -(\langle V_i \rangle - E_{syn,ij}) / g_{m,i}$$

- This approximation assumes that $(V - E_{syn})$ is roughly constant (E_{syn} is synaptic reversal potential: -80 mV for inhibitory and 0 mV for excitatory)
- These **current-based** synapse models are easier to work with analytically, and form the basis of further simplified firing-rate network models

Dales Law

- Synapses are either excitatory or inhibitory (E or I)
- For any neuron, all of its outgoing synapses are the same type (E or I)
- This suggests an important simplification - we can write the synaptic conductance as:



conductance of synapse from j to i

fixed weight of j to i
(number of ion channels etc.)

timecourse of all outgoing synapses from neuron j

$$g_{syn,ij}(t) = \bar{g}_{syn,ij} g_{syn,j}(t)$$

- This reduces the number of synaptic conductance equations from N^2 to N
- This makes quite strong assumptions – in reality, half of synapses may fail on any given spike, synapses may release different amounts of neurotransmitter, etc.

Firing Rate Network Models

Firing rate models abstract the spiking of a neuron into a firing rate (spikes/second)

We can replace our spiking network models with firing rate models using **two key approximations**:

- 1) We replace spikes with the **time- or neuron-averaged firing rates**
- 2) We model the inputs as an average current

Spiking Network [current based synapses]

$$\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) + \sum_j w_{ij} g_{syn,ij}(t) + I_{ext,i}(t) / g_{m,i}$$

Firing Rate Network Models

Firing rate models abstract the spiking of a neuron into a firing rate (spikes/second)

We can replace our spiking network models with firing rate models using **two key approximations**:

- 1) We replace spikes with the **time- or neuron-averaged firing rates**
- 2) We model the inputs as an average current

Spiking Network [current based synapses]

$$\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) + \sum_j w_{ij} g_{syn,ij}(t) + I_{ext,i}(t) / g_{m,i}$$

Rate Approximation $g_{syn,ij}(t) \rightarrow \langle g_{syn,ij}(t) \rangle \approx \bar{g}_{syn,ij} r_j(t) \approx \bar{g}_{syn,ij} \phi(V_j)$

firing rate transfer function



Firing Rate Network Models

Firing rate models abstract the spiking of a neuron into a firing rate (spikes/second)

We can replace our spiking network models with firing rate models using **two key approximations**:

- 1) We replace spikes with the **time- or neuron-averaged firing rates**
- 2) We model the inputs as an average current

Spiking Network [current based synapses]

$$\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) + \sum_j w_{ij} g_{syn,ij}(t) + I_{ext,i}(t) / g_{m,i}$$

firing rate **transfer function**

Rate Approximation $g_{syn,ij}(t) \rightarrow \langle g_{syn,ij}(t) \rangle \approx \bar{g}_{syn,ij} r_j(t) \approx \bar{g}_{syn,ij} \phi(V_j)$


Firing Rate Network $\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) + \sum_j \bar{w}_{ij} \phi(V_j) + I_{ext,i}(t) / g_{m,i}$

Firing Rate Network Models

What do we mean by firing rate? What are we averaging over?


Two interpretations: average over neurons or time (or both). Usually we don't specify which.

Time-Average



Neuron-Average

rate = average over pool of equivalent neurons
(several neurons, single run)



activity

$$A = \frac{1}{\Delta t} \frac{n_{\text{act}}(t; t + \Delta t)}{N}$$

Firing Rate Network Models

- Advantages of firing rate network models:

- Numerically efficient
- Analytically tractable
- Can train to optimise cost functions (e.g., train RNNs via backprop)
- Simpler to visualise/understand

- Disadvantages:

- Lose information about spike timing
- Assumptions required for averaging may be unrealistic
- Sometimes a firing rate network behaves very differently when converted to a spiking network (e.g., may become unstable, or require fine-tuning of parameters)

Firing Rate Models: Assumptions

- When is it reasonable to approximate a neuron's spike pattern in terms of its rate?
- We are ignoring spike patterns, e.g. synchrony, sequences, refractory periods, etc.
- Firing rate models make sense for neurons that fire asynchronously and irregularly (e.g., Poisson firing)
- This is often a reasonable approximation for neurons in cortex, but some brain regions or brain states have reliable temporal spike patterns that are ignored by firing rate models

Two Firing Rate Network Models

- There are two common variants of firing rate network models in the field
- We discussed the following model earlier:

$$\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) + \sum_j \bar{w}_{ij} \phi(V_j) + I_{ext,i}(t) / g_{m,i} \quad \begin{array}{l} V_i = \text{membrane potential} \\ \phi = \text{“transfer function”} \\ \quad \text{(the f-I curve)} \end{array}$$

- This would usually be rewritten in a simpler form:

$$\tau_m \frac{dV_i}{dt} = -(V_i - E_m) + \sum_j w_{ij} \phi(V_j) + u_i(t)$$

Two Firing Rate Network Models

- There are two common variants of firing rate network models in the field
- We discussed the following model earlier:

$$\tau_{m,i} \frac{dV_i}{dt} = -(V_i - E_{m,i}) + \sum_j \bar{w}_{ij} \phi(V_j) + I_{ext,i}(t) / g_{m,i} \quad \begin{array}{l} V_i = \text{membrane potential} \\ \phi = \text{"transfer function"} \\ \quad \text{(the f-I curve)} \end{array}$$

- This would usually be rewritten in a simpler form:



$$\tau_m \frac{dV_i}{dt} = -(V_i - E_m) + \sum_j w_{ij} \phi(V_j) + u_i(t)$$

- A very common alternative equation for a firing rate network is:

$$\tau_m \frac{dr_i}{dt} = -r_i + \phi \left(\sum_j w_{ij} r_j + u_i(t) \right) \quad \begin{array}{l} r_i = \text{"firing rate"} \\ \quad \text{(spikes per second)} \end{array}$$

Transfer Functions are f-I Curves

- Many choices of transfer function are available: rectified-linear, rectified-power law, sigmoid, rectified-tanh etc.
- Often a threshold-linear function is used, and sometimes even a purely linear function.
- Neurons can't have negative firing rates, and their firing rates saturate for very strong input currents, so threshold-sigmoid transfer functions may be more realistic.




Summary: Network Models

- Spiking network models couple together neurons via synapses to capture spiking patterns of the network
- Firing rate models abstract spike patterns into a single firing rate for each neuron
- Firing rate models can be justified using various approximations, and are much simpler and easier to work with, but less realistic
- Usually, our goal is to create the simplest model that captures some phenomenon, not the most detailed model. Unless we are interested in spike timing, firing rate models are the default option

Outline of Lecture


- Introduction to networks
- Spiking and firing rate network models
- **Feedforward and recurrent networks**
- Excitatory-inhibitory networks




The “rainbow” (an image of a real network)

Computations in Feedforward Networks

- Feedforward networks are those with no loops
- Feedforward networks are simple to analyse and can capture important properties of sensory systems (as we saw with convolutional neural networks)
- But they omit the recurrent and top-down (feedback) connections found in the brain, and are therefore unrealistic biologically




Feedforward Networks: Synfire Chains



- A synfire chain is a feedforward spiking network comprising multiple layers
- Spiking propagates in volleys from one layer to the next
- We can ask questions such as: will spiking propagate indefinitely or eventually die out? How does this depend on connectivity?
- This can offer a useful model of sensory processing (first ~ 100 ms after stimulus onset).

Feedforward vs Recurrent Processing

Box 2. The feedforward sweep



REVIEW

Initial feedforward sweep is like a synfire chain. We can already recognise and categorise objects with this.

Later, recurrent and top-down processes take place, which are necessary for higher level “conscious” perception

The distinct modes of vision offered by feedforward and recurrent processing

Victor A.F. Lamme and Pieter R. Roelfsema

An analysis of response latencies shows that when an image is presented to the visual system, neuronal activity is rapidly routed to a large number of visual areas. However, the activity of cortical neurons is not determined by this feedforward sweep alone. Horizontal connections within areas, and higher areas providing feedback, result in dynamic changes in tuning. The differences between feedforward and recurrent processing could prove pivotal in understanding the distinctions between attentive and pre-attentive vision as well as between conscious and unconscious vision. The feedforward sweep rapidly groups feature constellations that are hardwired in the visual brain, yet is probably incapable of yielding visual awareness; in many cases, recurrent processing is necessary before the features of an object are attentively grouped and the stimulus can enter consciousness.


Trends Neurosci. (2000) 23, 571–579

Recurrent Networks: Fixed Points, Oscillations and Chaos


- Recurrently connected networks of neurons admit many kinds of **emergent collective behaviour** depending on their connectivity and transfer functions
- We can use **dynamical systems theory** to study these different kinds of behaviour
- Can exhibit fixed points (attractors), oscillations (limit cycles) or chaotic dynamics
- These dynamics confer different **computational capabilities** to the network
- Attractors may be useful for **memory**, oscillations may be useful for **neuronal communication**, chaos may be useful for high-capacity **input discrimination**

Network Dynamics: Fixed Points, Oscillations and Chaos


One Fixed Point (Point Attractor)




Multiple Discrete Fixed Points




A Continuum of Fixed Points (Line Attractor)



Stable Oscillation (Limit Cycle)



Chaos (Lorenz Attractor)




Summary: Feedforward and Recurrent Networks

- Two models of feedforward networks: deep neural networks and synfire chains
- Offer a good model for the multi-layer propagation of sensory activity following stimulus onset
- Recurrent networks are far richer due to their dynamics
- Dynamics can be used for computation, e.g. fixed points for memory, line attractors for evidence integration, etc.
- We will see examples of such networks later in the course
(ring network, Hopfield network, etc.)

Outline of Lecture

- Introduction to networks
- Spiking and firing rate network models
- Feedforward and recurrent networks
- **Excitatory-inhibitory networks**



The “rainbow” (an image of a real network)

E-I Firing Rate Networks

- Some models consider a **single population** of cells whose outgoing synaptic weights may be positive or negative (**violating Dale's principle**)
- Other models include **separate E and I populations**:

$$\tau_E \frac{dr_i^E}{dt} = -r_i^E + \phi\left(\sum_j W_{ij}^{EE} r_j^E - \sum_j W_{ij}^{EI} r_j^I + u_i^E\right)$$

$$\tau_I \frac{dr_i^I}{dt} = -r_i^I + \phi\left(\sum_j W_{ij}^{IE} r_j^E - \sum_j W_{ij}^{II} r_j^I + u_i^I\right)$$

- For firing rate models, it is often assumed that we are averaging over both E and I neurons, so don't need to treat E and I separately
- For spiking models, E and I neurons are usually treated separately


Balanced Excitatory-Inhibitory Network Models

Chaos in Neuronal Networks with Balanced Excitatory and Inhibitory Activity

C. van Vreeswijk and H. Sompolinsky

Neurons in the cortex of behaving animals show temporally irregular spiking patterns. The origin of this irregularity and its implications for neural processing are unknown. The hypothesis that the temporal variability in the firing of a neuron results from an approximate balance between its excitatory and inhibitory inputs was investigated theoretically. Such a balance emerges naturally in large networks of excitatory and inhibitory neuronal populations that are sparsely connected by relatively strong synapses. The resulting state is characterized by strongly chaotic dynamics, even when the external inputs to the network are constant in time. Such a network exhibits a linear response, despite the highly nonlinear dynamics of single neurons, and reacts to changing external stimuli on time scales much smaller than the integration time constant of a single neuron.

—————




Balanced Excitatory-Inhibitory Network Models

- The theory of balanced network dynamics **explains the irregular asynchronous regime** of cortex
- Network generates this regime even when inputs are constant!
- Network has sparse but strong synapses. Take limit of many neurons and solve analytically (uses tools from statistical physics)
- Behaviour of network:
 - excitatory and inhibitory inputs to each neuron are **strong** and **balanced**, and almost cancel one another
 - Fluctuations in external input are rapidly cancelled by recurrent inhibition
 - Spiking is generated by **fluctuations** in E-I balance

Evidence for E-I Balance

- E-I balance has been confirmed in multiple experiments. A successful theoretical prediction! Below shows how E and I inputs to a neuron track each other in time.



Behavioral/Systems/Cognitive

Neocortical Network Activity *In Vivo* Is Generated through a Dynamic Balance of Excitation and Inhibition

Bilal Haider, Alvaro Duque, Andrea R. Hasenstaub, and David A. McCormick


Department of Neurobiology, Kavli Institute for Neuroscience, Yale University School of Medicine, New Haven, Connecticut 06510

Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities

Michael Okun & Ilan Lampl

Other Dynamical Regimes

- Other dynamical regimes are possible, and have been studied theoretically




Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons

NICOLAS BRUNEL

LPS, Ecole Normale Supérieure, 24 rue Lhomond, 75231 Paris Cedex 05, France

brunel@lps.ens.fr

Received December 23, 1998; Revised June 9, 1999; Accepted June 9, 1999






Asynchronous Activity and E-I Balance

- Two states in cortex:

- 1) balanced regime, characterised by irregular, asynchronous, Poisson spiking
- 2) synchronous regime, characterised by global, synchronous fluctuations in network activity

Cortical state is modulated by attention, alertness, etc.



Summary: E-I Networks

- Networks in the brain are made of excitatory and inhibitory neurons
- When coupled together, these can generate interesting dynamical behaviour, including:
 - Balanced network dynamics (with asynchronous irregular firing)
 - Synchronised dynamics (with large collective fluctuations)

Summary of Neural Networks (Lecture 1)

- Networks can be used to perform computations
- Two main classes of network – feedforward and recurrent
- Two main strategies for modelling biological networks – spiking and firing rate
- Recurrent networks exhibit rich and complex dynamics
- During alert brain states, cortical networks are characterised by E-I balance and asynchronous activity

Bibliography

- Lecture notes Ch 10
- Dayan and Abbott Ch 7 [network dynamics, firing rate vs spiking networks]
- Lamme and Roelfsema (2000) [feedforward vs recurrent processing]
- Van Vreeswijk and Sompolinsky (1996) [balanced network dynamics]
- Harris and Thiele (2011) [two cortical states]