Harvey Mudd College

CS 152
Neural Networks
Fall 2004

Bob Keller, Professor
keller@cs.hmc.edu
621-8483

Office Hours (1249 Olin):

- Note: 1249 is in the SW corner of Olin
- Monday, Tuesday, Wednesday 4:15-5:30
- By drop-in (as available)
- By appointment:
  - email keller@cs.hmc.edu
  - phone 621-8483

Text: “NND”

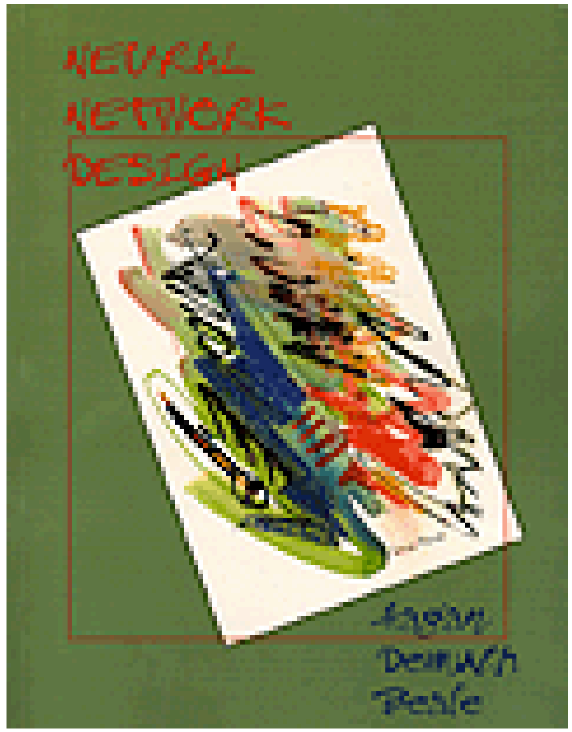

Neural Network Design by Martin T. Hagan, Howard B. Demuth, and Mark Beale, originally published by PWS Publishing Company, 1996.

Reprint available from Huntley Bookstore, or University of Colorado Bookstore at 303-492-3648.
ISBN 0-9717321-0-8.

Course Outline

Please Refer to Web Page for details:

`http://www.cs.hmc.edu/courses/current/cs152`

Neural Networks: an Eclectic Discipline

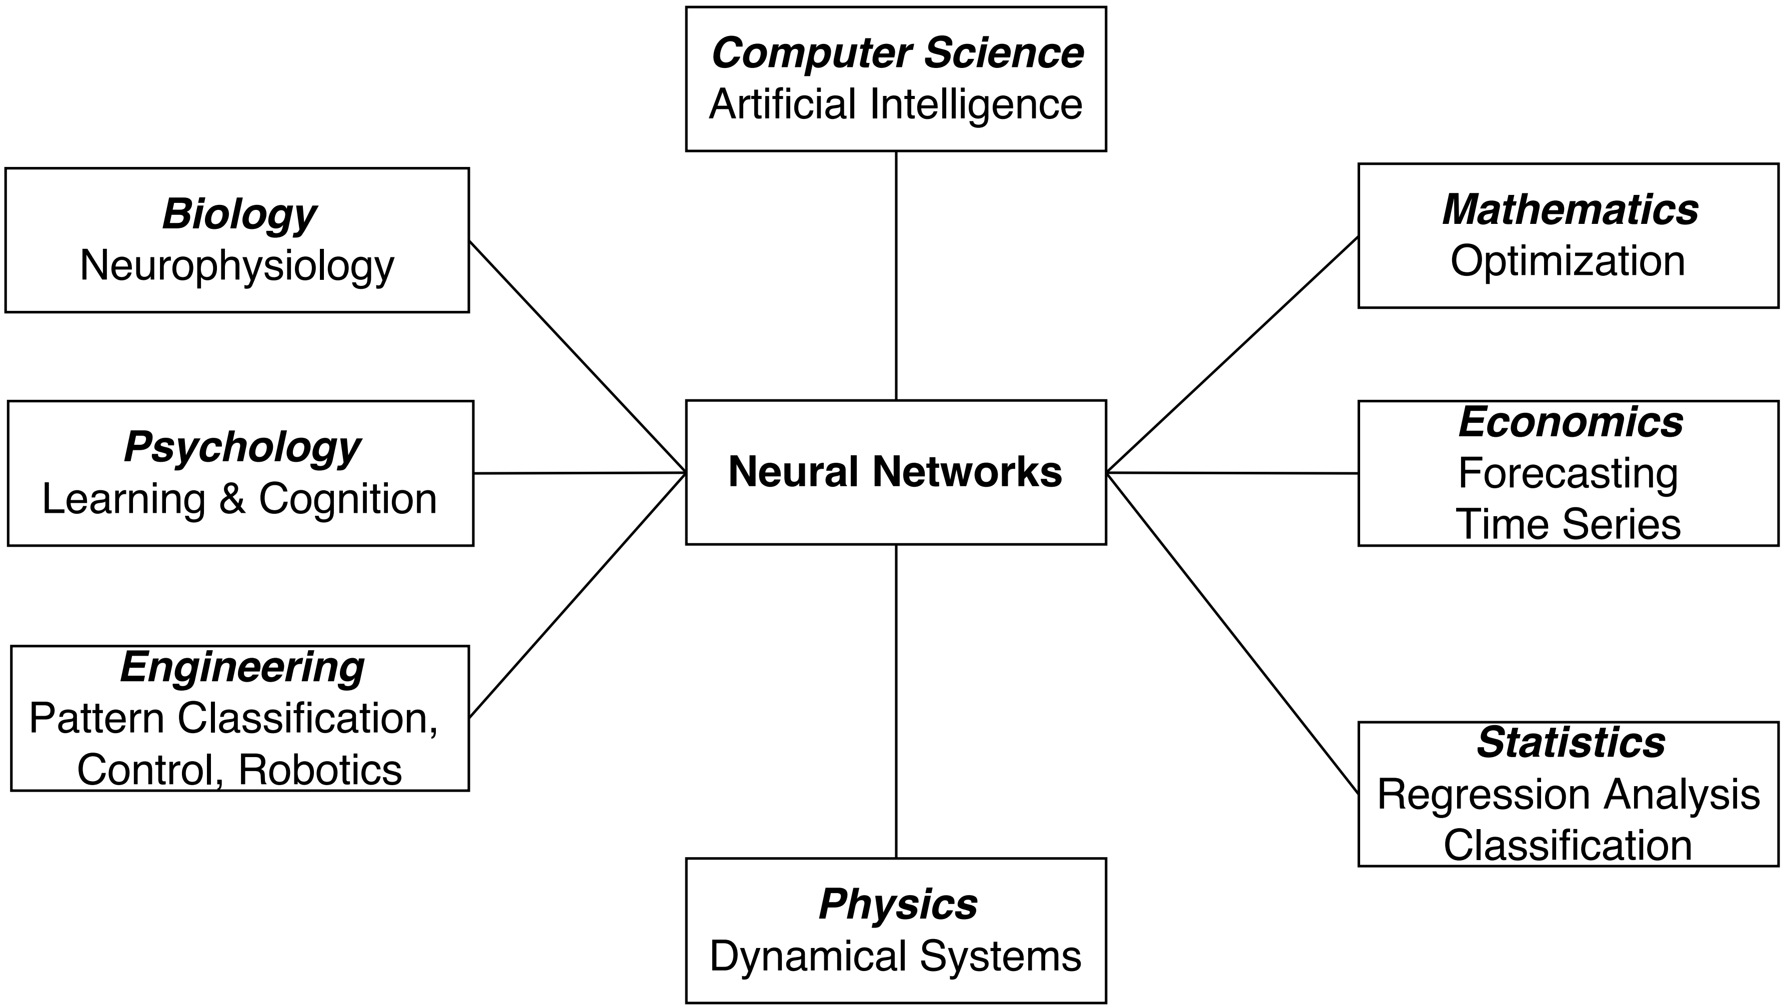

Biological Intelligence

- Intelligence, the ability to make decisions based upon input from the environment.

- Intelligence is realized by a ***network*** of ***neurons***, for example the brain and the attendant sensory and motor neurons.

"Neurons R Us"

- Not only our intelligence, but all aspects of our behavior, are the result of neural activity:
  - emotions
  - memory
  - reflexes
  - habits
  - likes and dislikes
  - addictions

Some Approaches to Artificial Intelligence

- Reverse Engineering of Biology
  - Understand real neurons well enough to model
  - ***Simulate*** neural behavior
- Artificial Neural Networks
  - Develop a parameterized model for a class of problems
  - ***Learn*** the parameters
- Simulated Evolution
  - Provide basic evolutionary mechanism for neurons
  - ***Evolve*** intelligent behavior

Fundamental Problems for a Given Neural Model

- How to ***represent*** information?
- How to characterize the ***computational capability*** of the model?
- How to achieve ***learning*** in the model?

Some Applications of Artificial Neural Networks (1 of 5)

- **Optical character recognition**
  - U.S. mail zip-code recognizer
  - Kanji: 4000 chars in 15 fonts, 99% accurate, 100k chars/sec (Sharp & Mitsubishi)
- **Communications**
  - Adaptive noise cancellation
    - Headphones
    - Conference telephones

Applications (2 of 5)

- **Process control**
  - Electric arc furnace control: 30MVA, 50kamp transformer, $2M savings
  - Steel-rolling mill controller
  - Copier uniformity control (Ricoh)
  - Anti-lock brakes, etc. (Ford)
  - Food process control (M&M)
  - Particle beam focusing (SLAC)
  - Flourescent bulb mfg. (GE)

Applications (3 of 5)

- **Financial analysis**
  - Prediction of commodities market (18% vs. 12.3% by traditional methods)
  - Mortgage risk evaluator (AVCO, Irvine)
  - Real-estate evalution (Foster Onsley Conley)
  - Portfolio management (LBS Captial)
  - Currency trading (Citibank)
- **Crime prevention**
  - Bomb sniffer (JFK airport)
  - Credit card fraud detection (Visa, etc.)

Applications (4 of 5)

- **Object classification**
  - Grading grains from video images
  - Forensics: glass classification
  - High-energy physics: particle identification
- **Warfare**
  - Missile guidance
- **Optical telescope focusing**

Applications (5 of 5)

- **Biomedical**
  - Clinical diagnoses
  - Patient mortality predictions
  - Protein structure analysis
  - Electrode placement
- **Speech recognition**
- **Game playing**
  - World backgammon champion

Some Physiological Aspects of Neurons

Reference for Biological Aspects

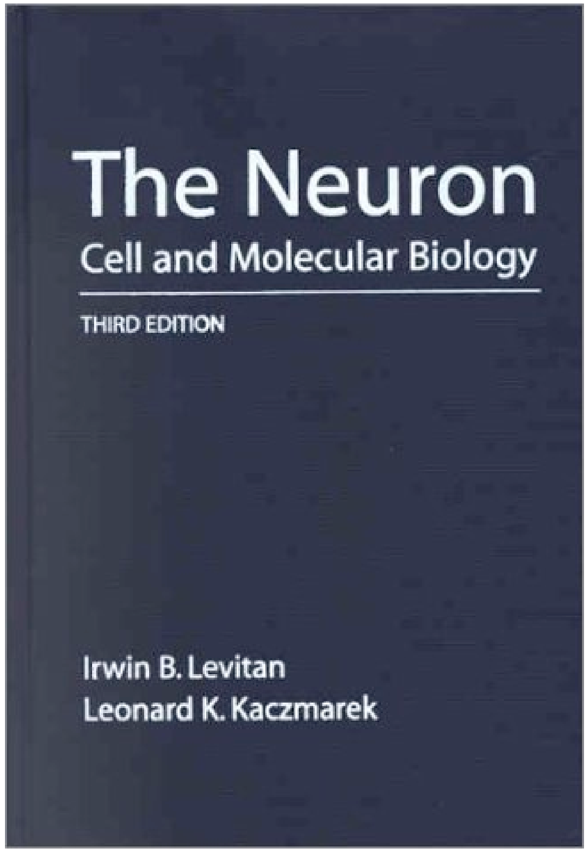

Oxford University Press.

Neuron Cell
(top half)

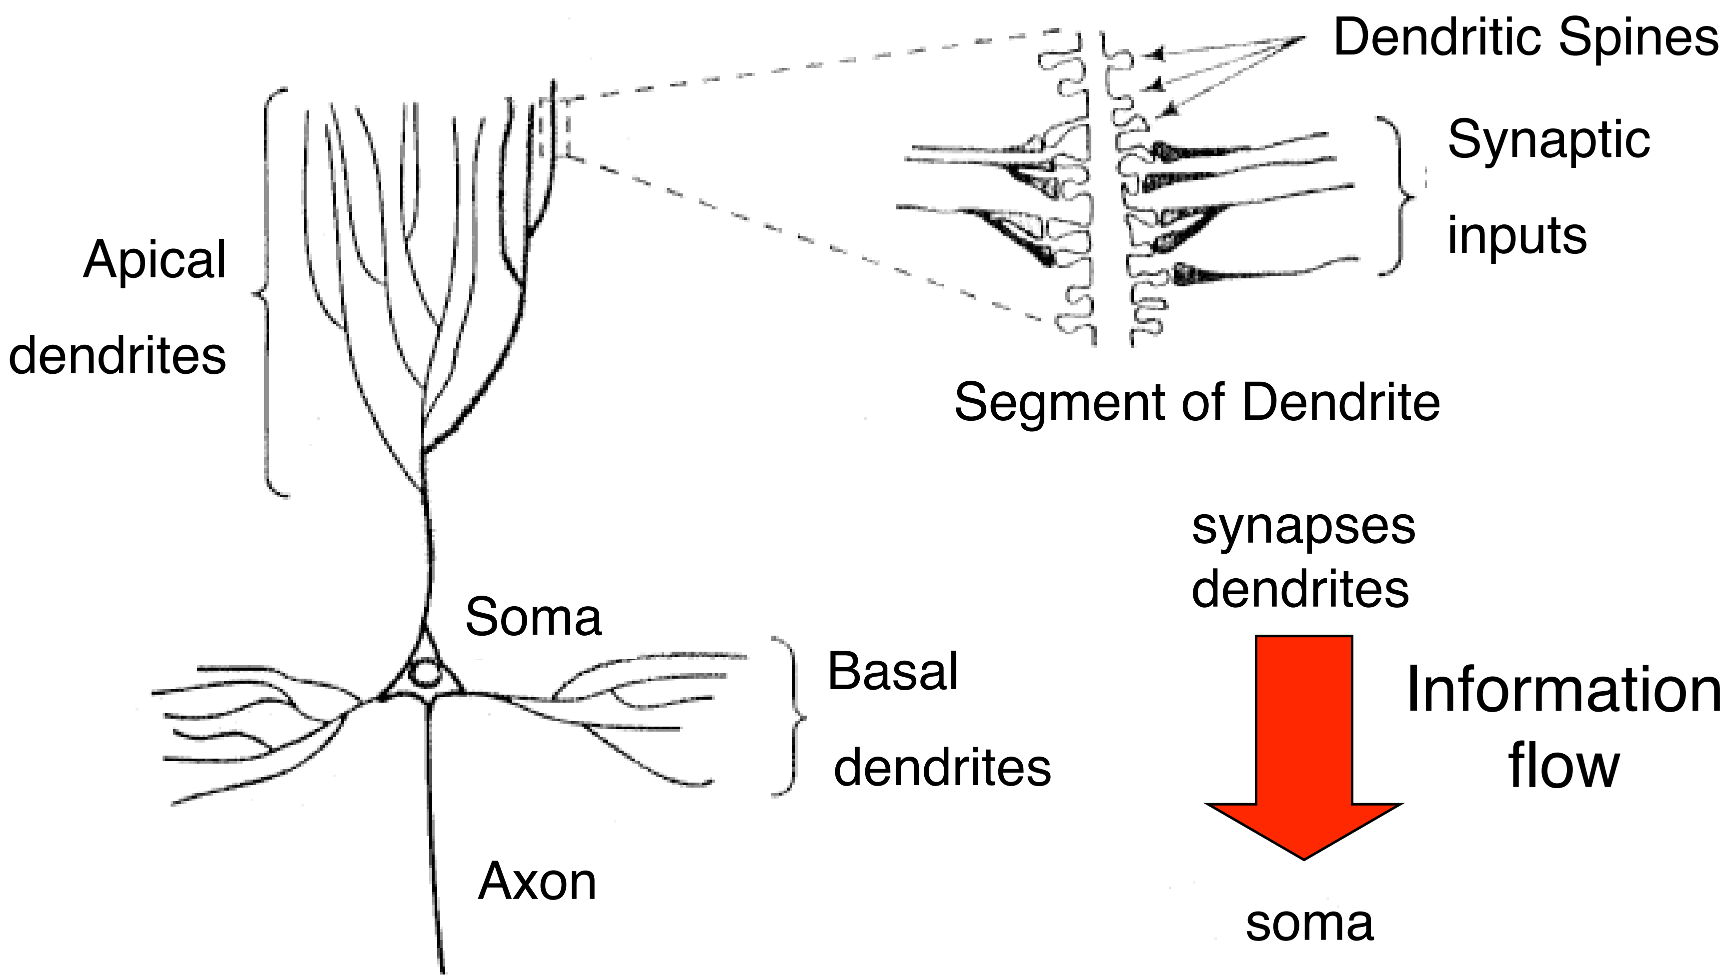

synapses
dendrites

Information
flow

soma

Neuron Cell (bottom half)

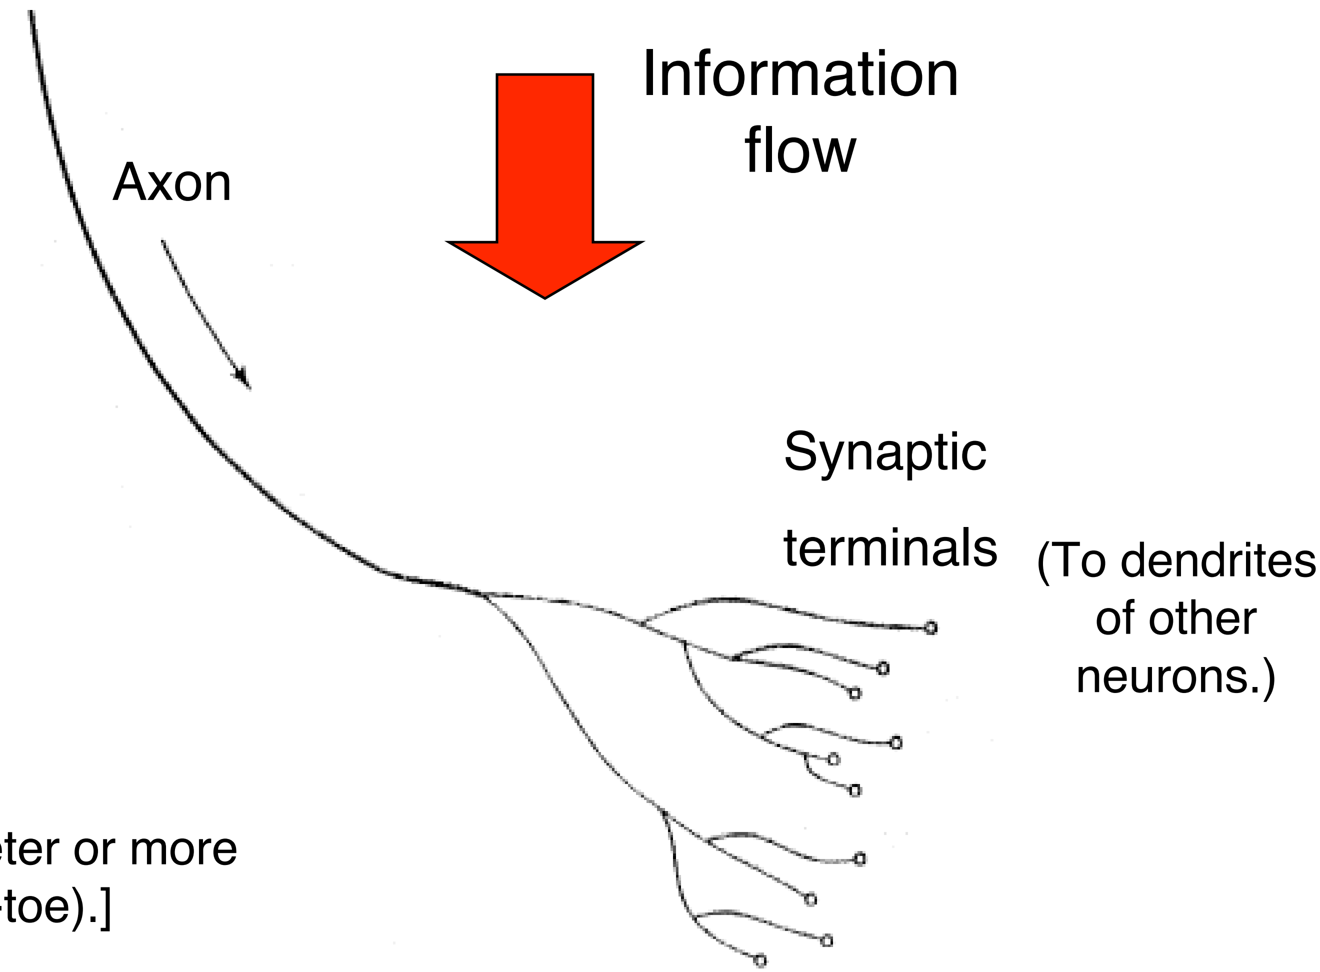

[Axon can be a meter or more long (e.g. spine-to-toe).]

Photomicrograph of one neural cell (from cerebral cortex)

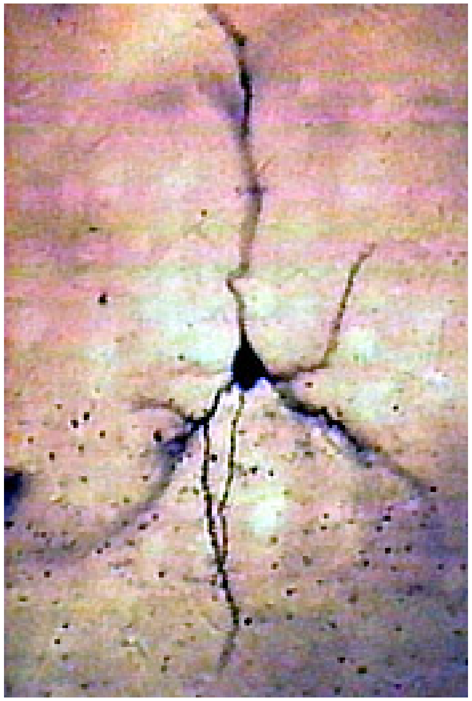

Structure of one neuron

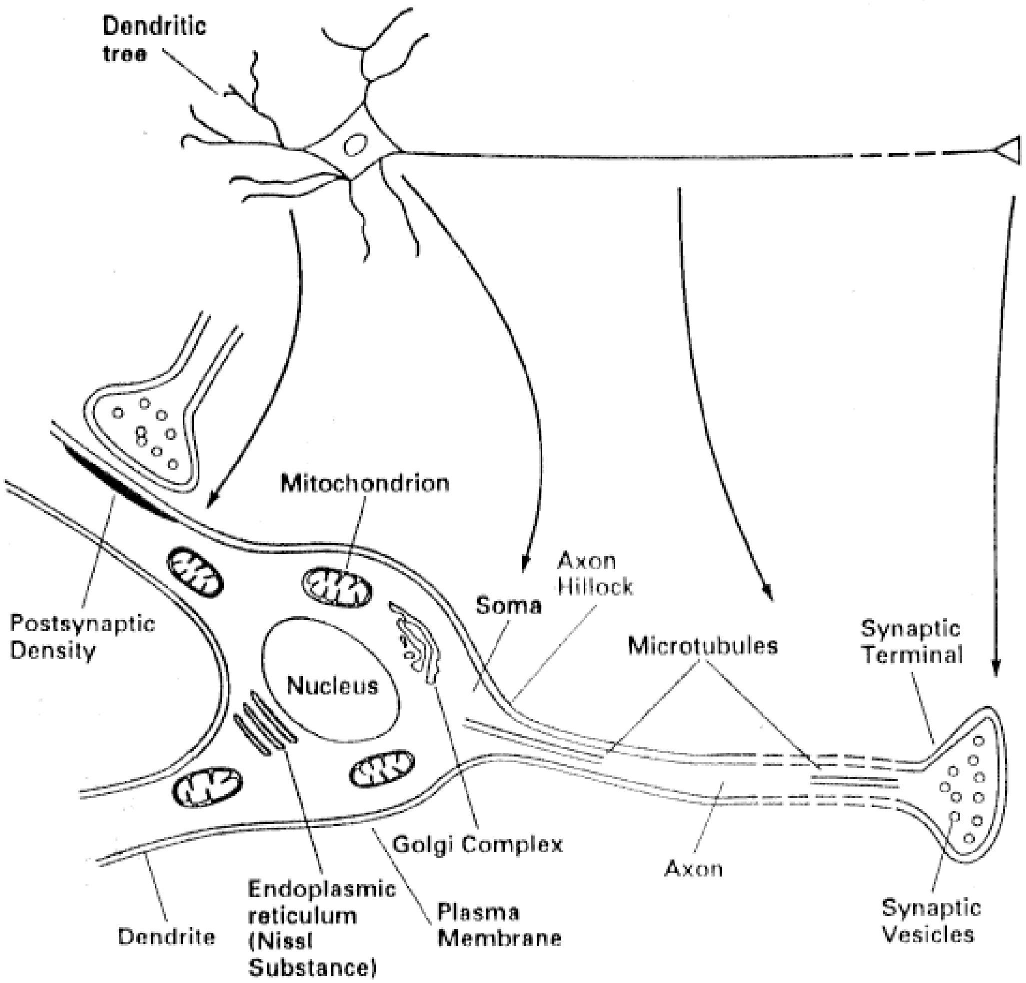

reference: Irwin B. Levitan and Leonard K. Kaczmarek, The Neuron, Oxford University Press, 1991.

Photomicrograph of network of neural cells (from the hippocampus region of the brain)

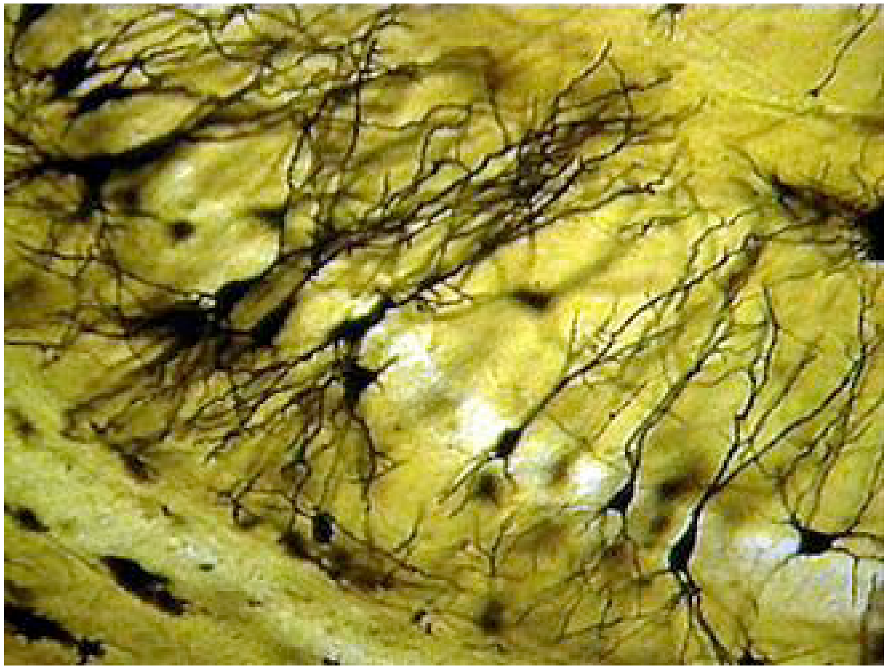

Composition of the Brain

- 10% neurons
- 90% glial (“glue”) cells

Myelin sheath around axon (consists of glial cells)

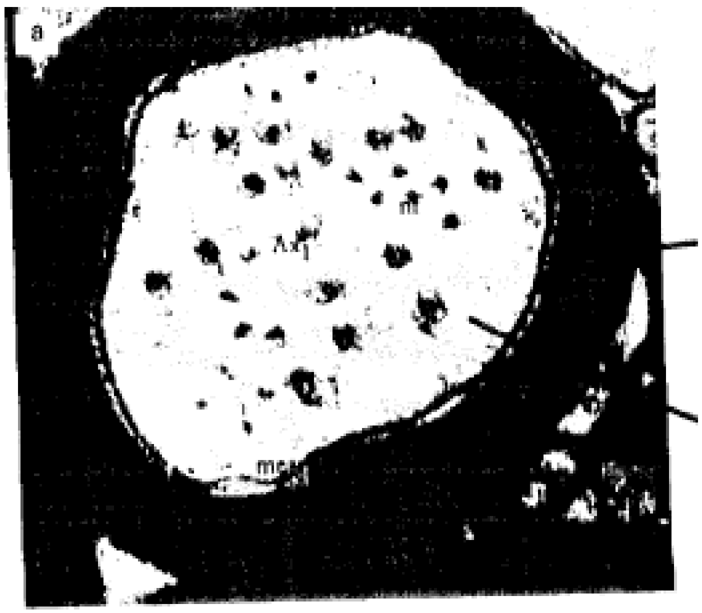

axon cross section

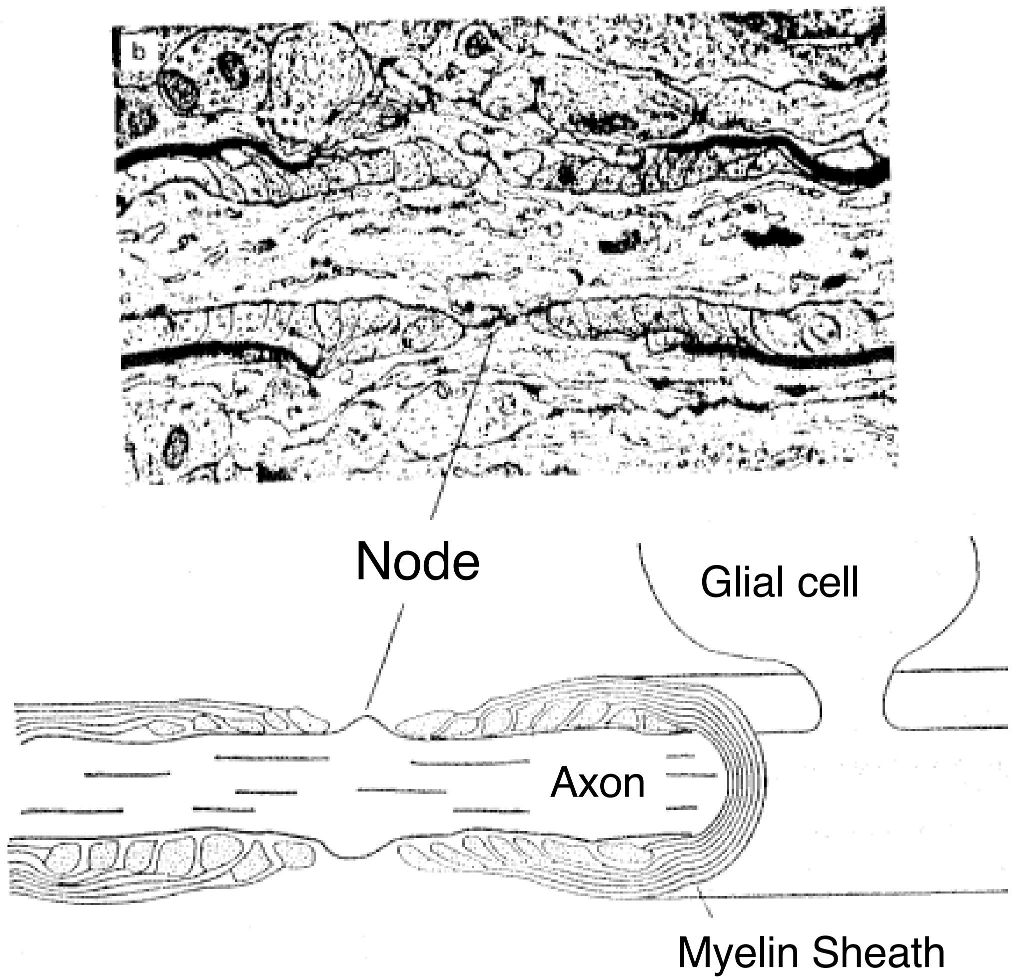

reference: Irwin B. Levitan and Leonard K. Kaczmarek, The Neuron, Oxford University Press, 1991.

Myelin Sheath (cont'd)

- Acts as insulator
- Current can flow out only at junctions (called nodes of Ranvier) to other axons
- Demyleinating diseases:
  - Myelin deficit in newborns
  - MS (multiple sclerosis)
  - ALS (amyotrophic lateral sclerosis, "Lou Gehrig's disease")

Dendrite Information Flow

- Normally dendrites receive information from synapses of other neurons

- In some cells, *both* input ***and output*** can occur through the *same* set of dendritic structures.

In addition to signal, axon carries:

- Construction material (proteins)
- Nutrients (in the form of mitochondria)
- Enzymes

Experiment determining chemical nature of neural transmission, Loewi, 1921

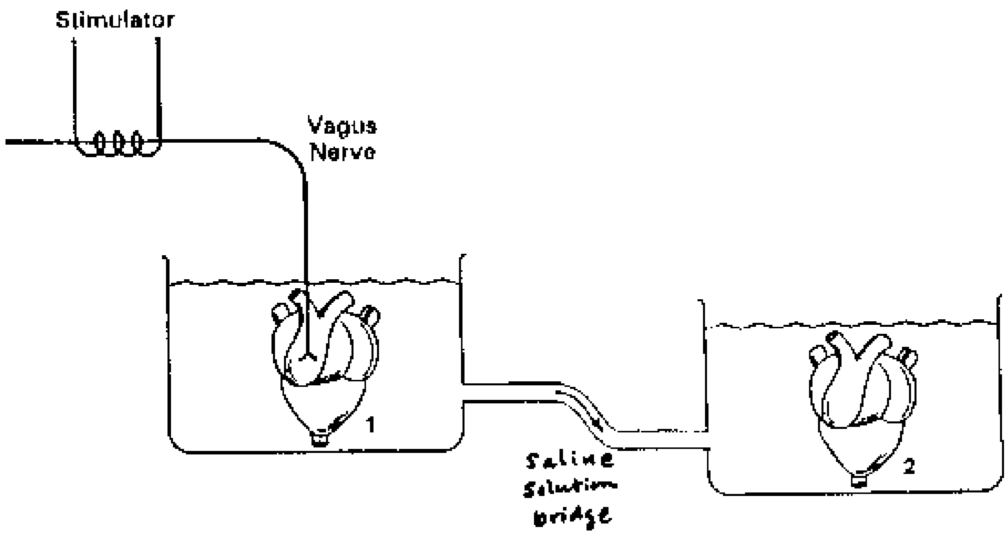

slow-down due to direct stimulation

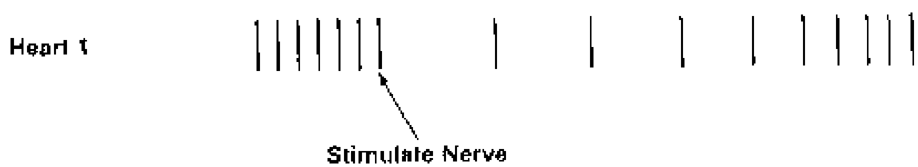

slow-down due to chemical change in solution

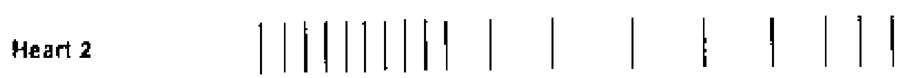

reference: Irwin B. Levitan and Leonard K. Kaczmarek, The Neuron, Oxford University Press, 1991.

Inter-Neuron Signaling

- An ionic (electro-chemical) reaction carries the signal across the gap between a synapse of one neuron and a dendrite of the next.
- The **strength** of this connection is an abstraction of the efficiency of the transfer.
- In artificial neural networks, this strength is represented by a **numeric weight**.

Chemical Synapse

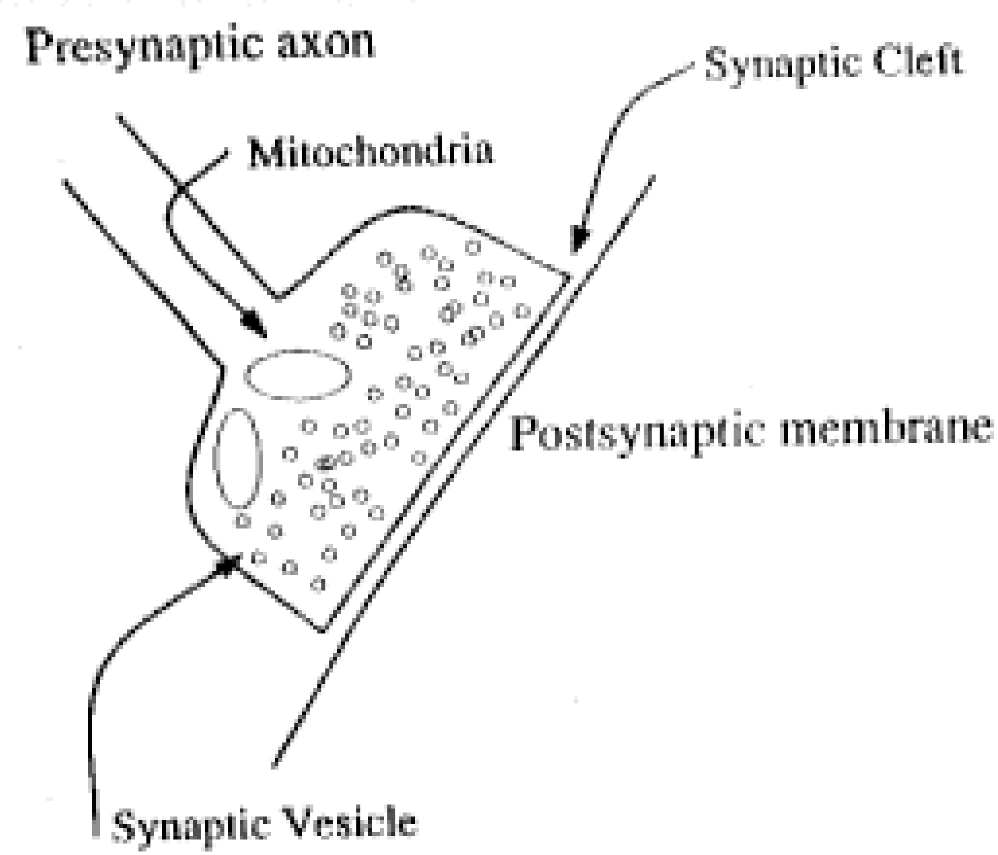

reference: James A. Anderson, An Introduction to Neural Networks, MIT Press, 1955.

Electronmicrograph of one synapse/dendrite connection

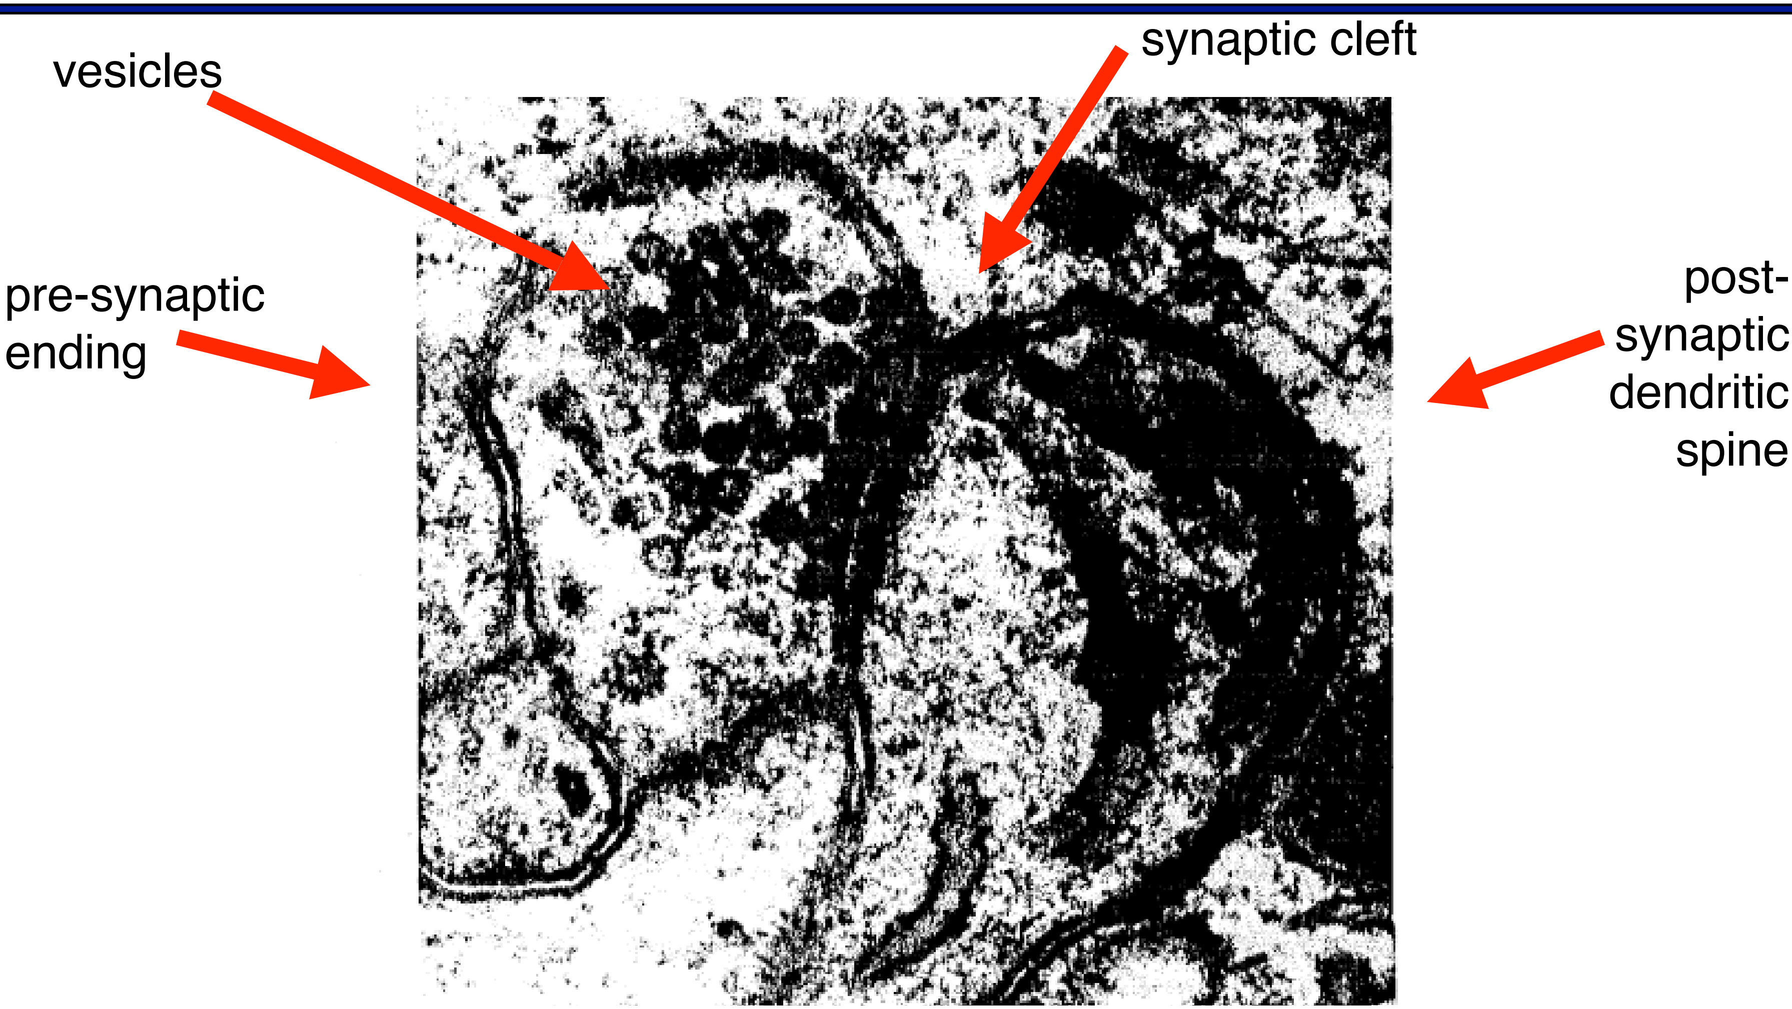

reference: Irwin B. Levitan and Leonard K. Kaczmarek, The Neuron, Oxford University Press, 1991.

Schematic sometimes used (symbolic of synaptic clefts)

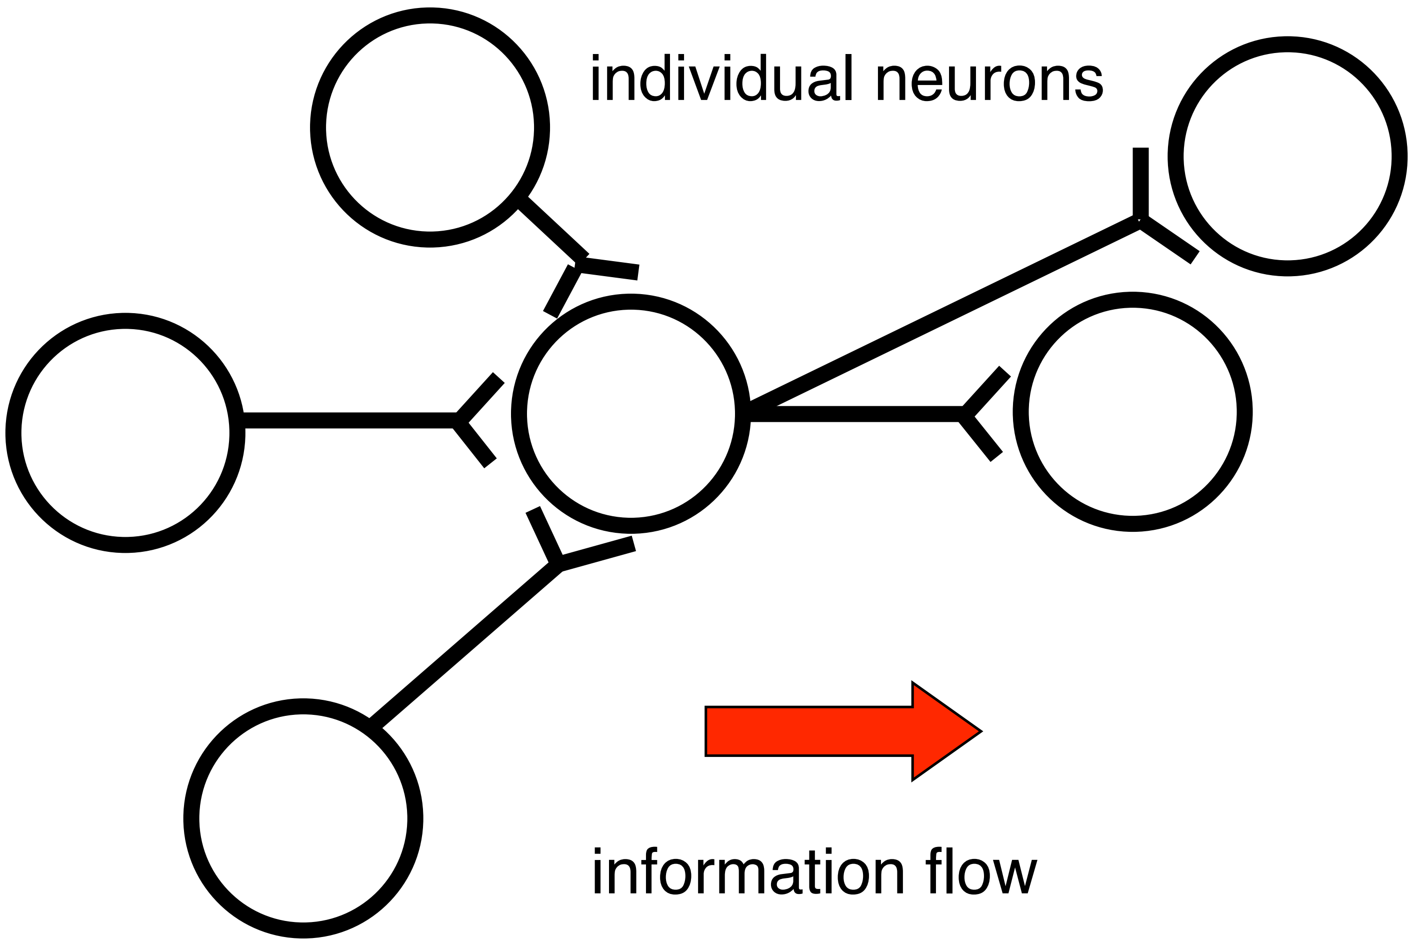

Ionic Neurotransmitter Reaction

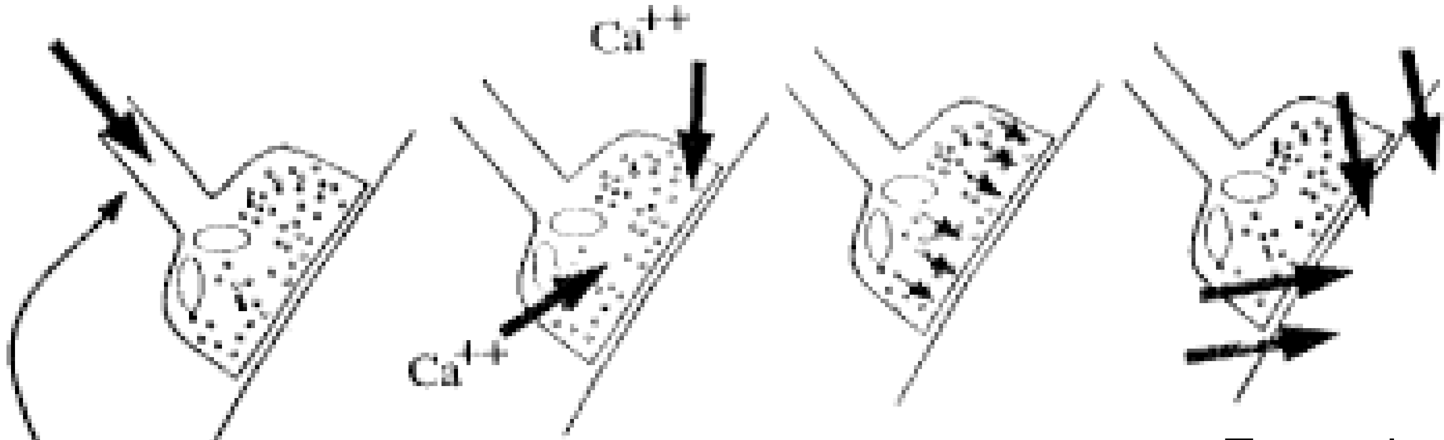

Action potential at synapes	Ca ions enter cell	Vesicles move to membrane, fuse, spill contents into cleft	Transmitter crosses cleft, causes voltage change on post-synaptic side

reference: James A. Anderson, An Introduction to Neural Networks, MIT Press, 1955.

Terminology

- **neurotransmitters**: molecules that traverse from synapse to dendrite through ion diffusion.
- **spiking**: abrupt change of output voltage
- **depolarization**: change in net input voltage toward a threshold value, at which it will “spike”
- **action potential**: the voltage change produced when the neuron spikes
- **refractory period**: period immediately after firing during which neuron is temporarily not firable

Spiking

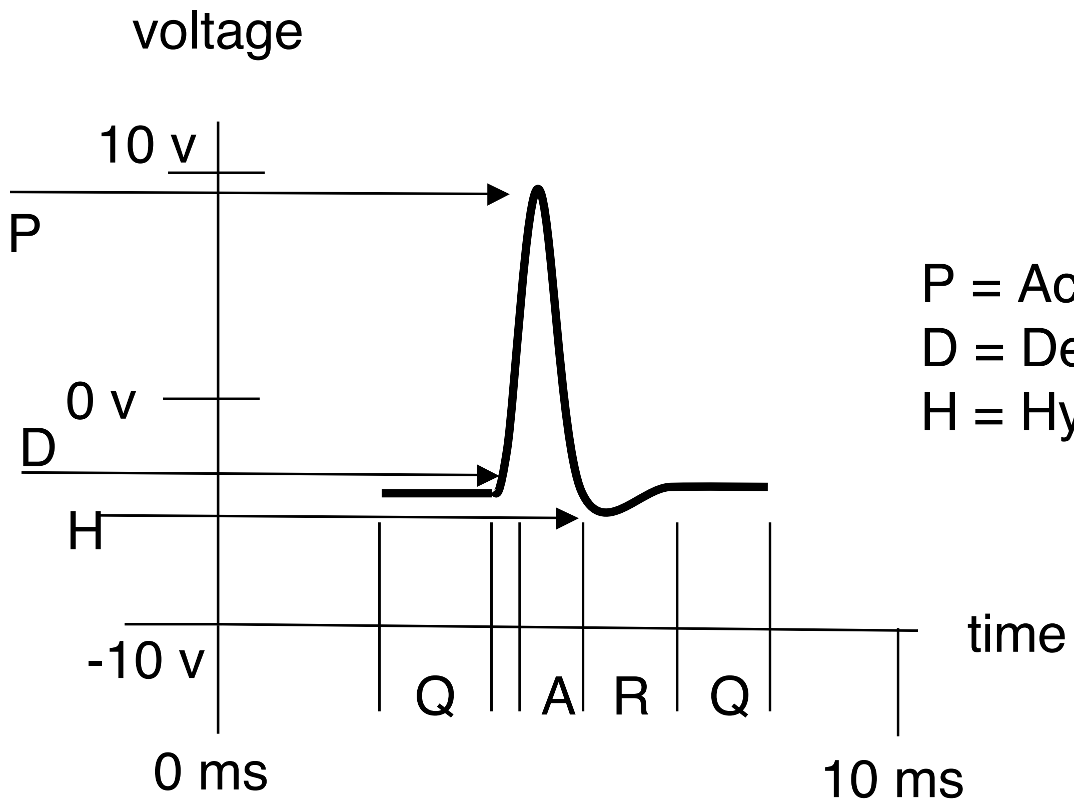

P = Action Potential
D = Depolarization Potential
H = Hyperpolarization Potential

Q = Quiescent
A = Absolute Refractory (not firable)
R = Relative Refractory (firability depends on elapsed time)

All-or-None Behavior

- The action potential is essentially binary-valued.
- The strength of the stimulus does not matter, except insofar as whether it is over or above a threshold.

Triggering phenomenon

Stimulus (summed inputs)

Response

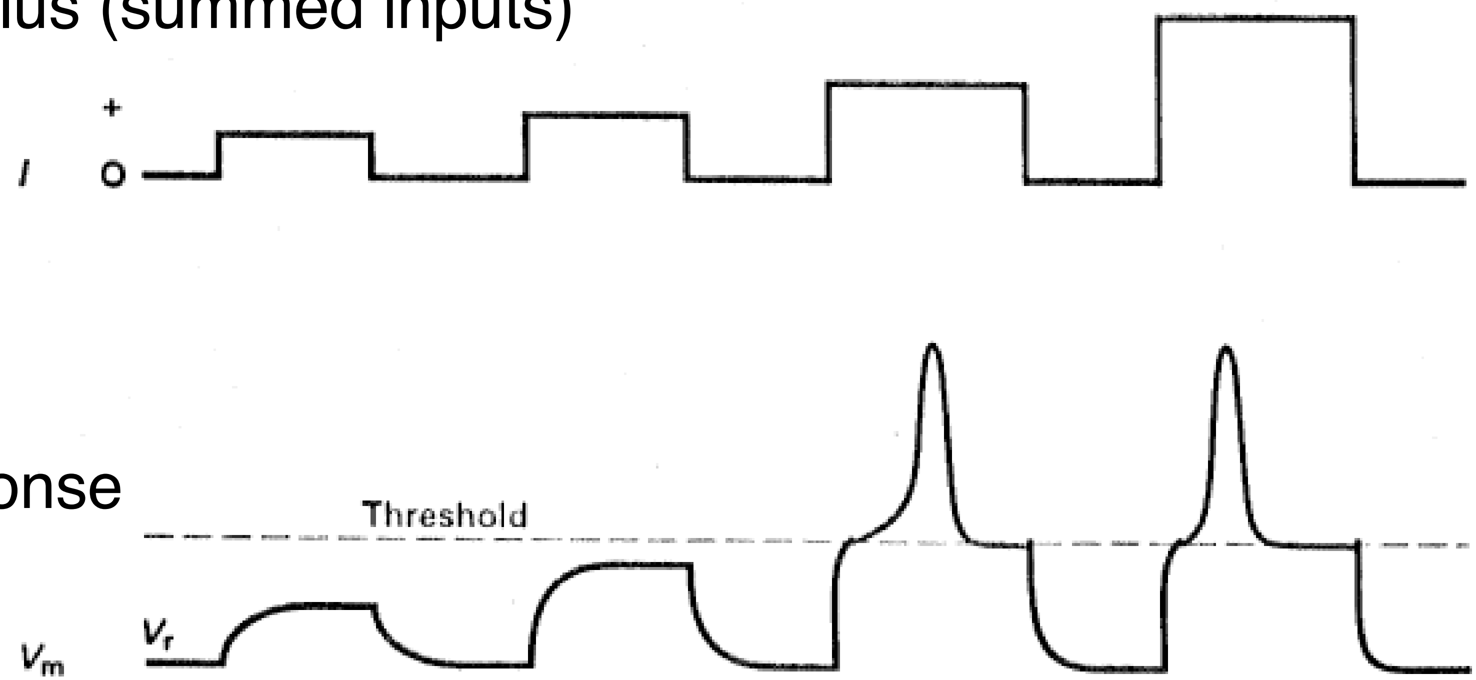

reference: Irwin B. Levitan and Leonard K. Kaczmarek, The Neuron, Oxford University Press, 1991.

Intensity

- Because of the all-or-none behavior, the neuron indicates intensity of stimulation by the **frequency** of spikes, not amplitude.

- Because of the refractory period, there is a maximum or **saturation** frequency at which the neuron can operate.

Spiking Frequency of a Neuron as a Function of (Artificial) Stimulus Magnitude

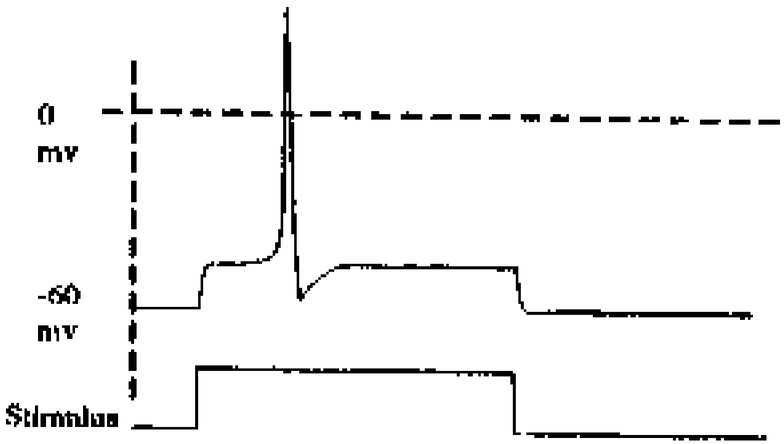

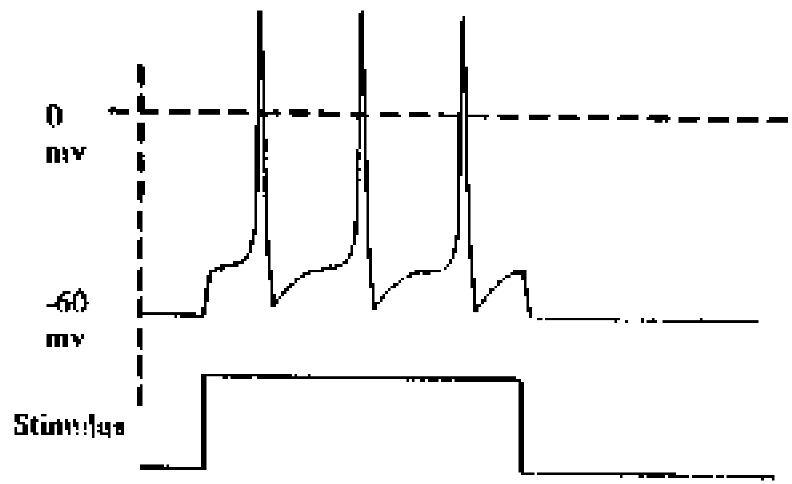

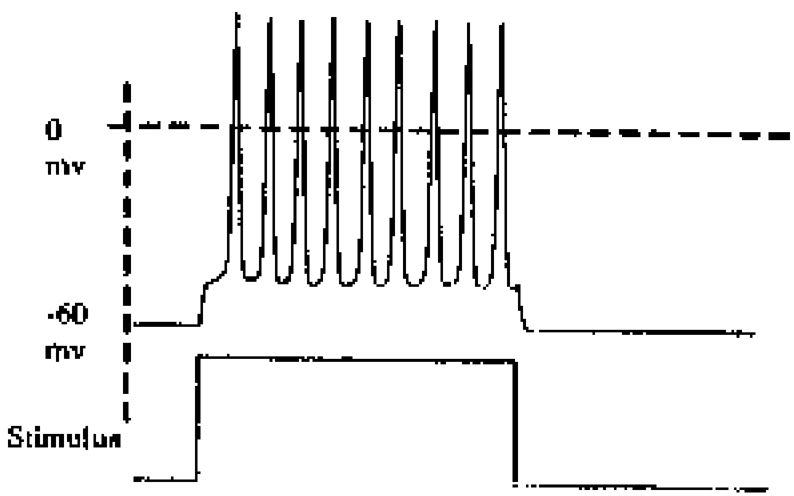

Larger stimulus higher frequency of output spiking

reference: James A. Anderson, An Introduction to Neural Networks, MIT Press, 1955.

Intensity Abstraction

- In **artificial** neural networks, we associate an output value with a neuron, which might be continuous, but have an upper limit.

- This can be viewed as a convenient abstraction of what is really the value of spiking frequency.

Sigmoid (S-shaped) Behavior

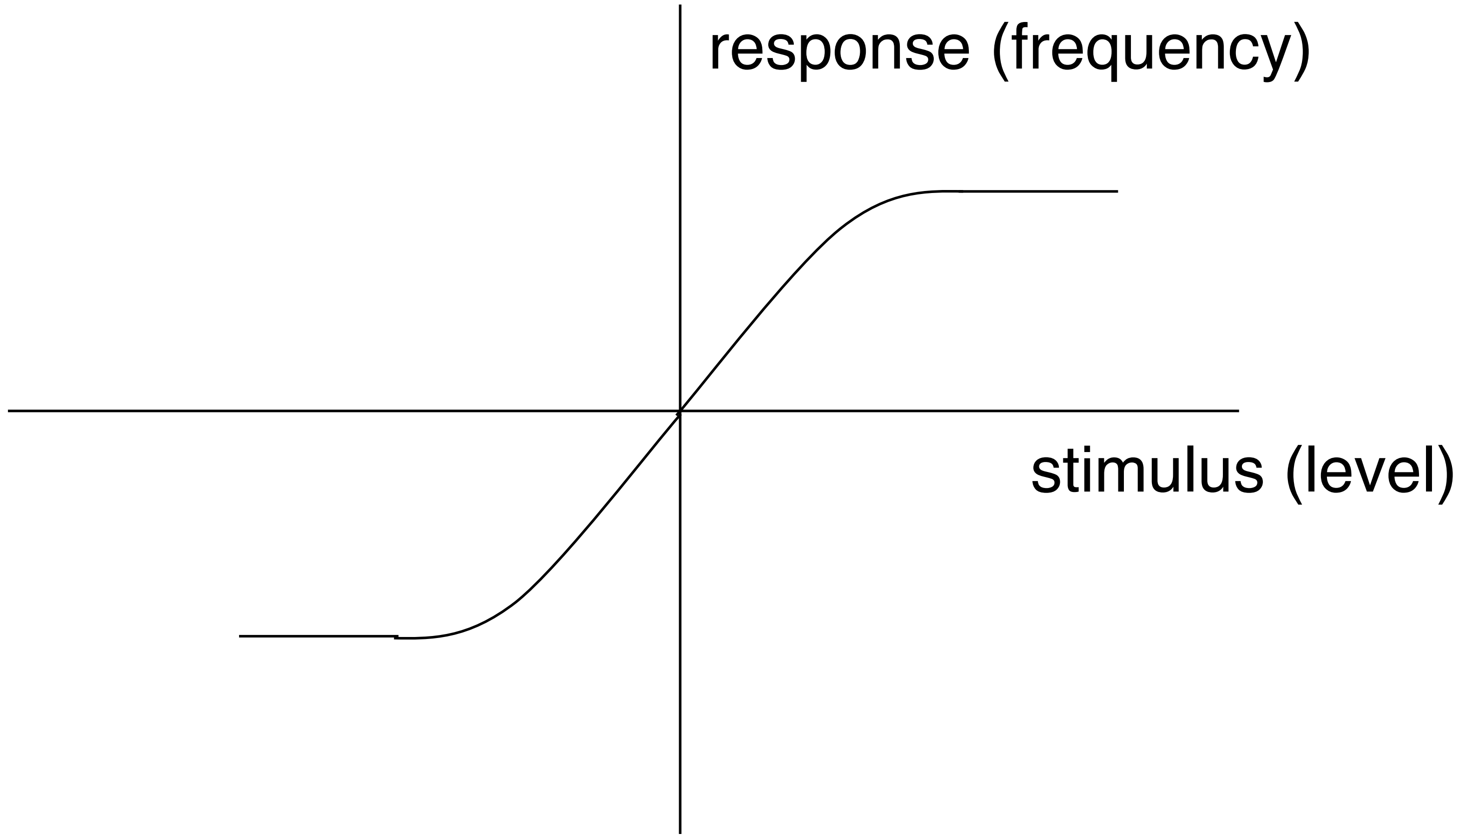

Further Abstracted Behavior

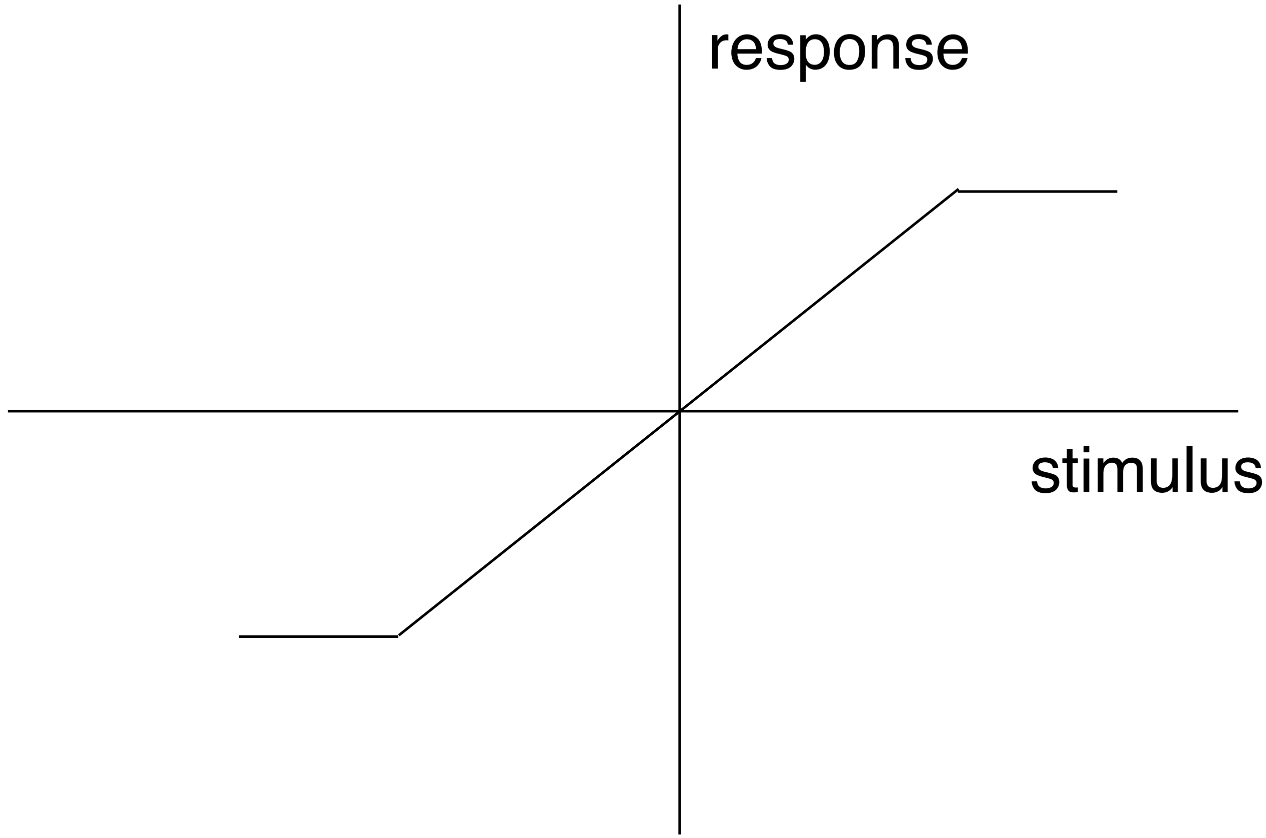

Still Further Abstracted Behavior

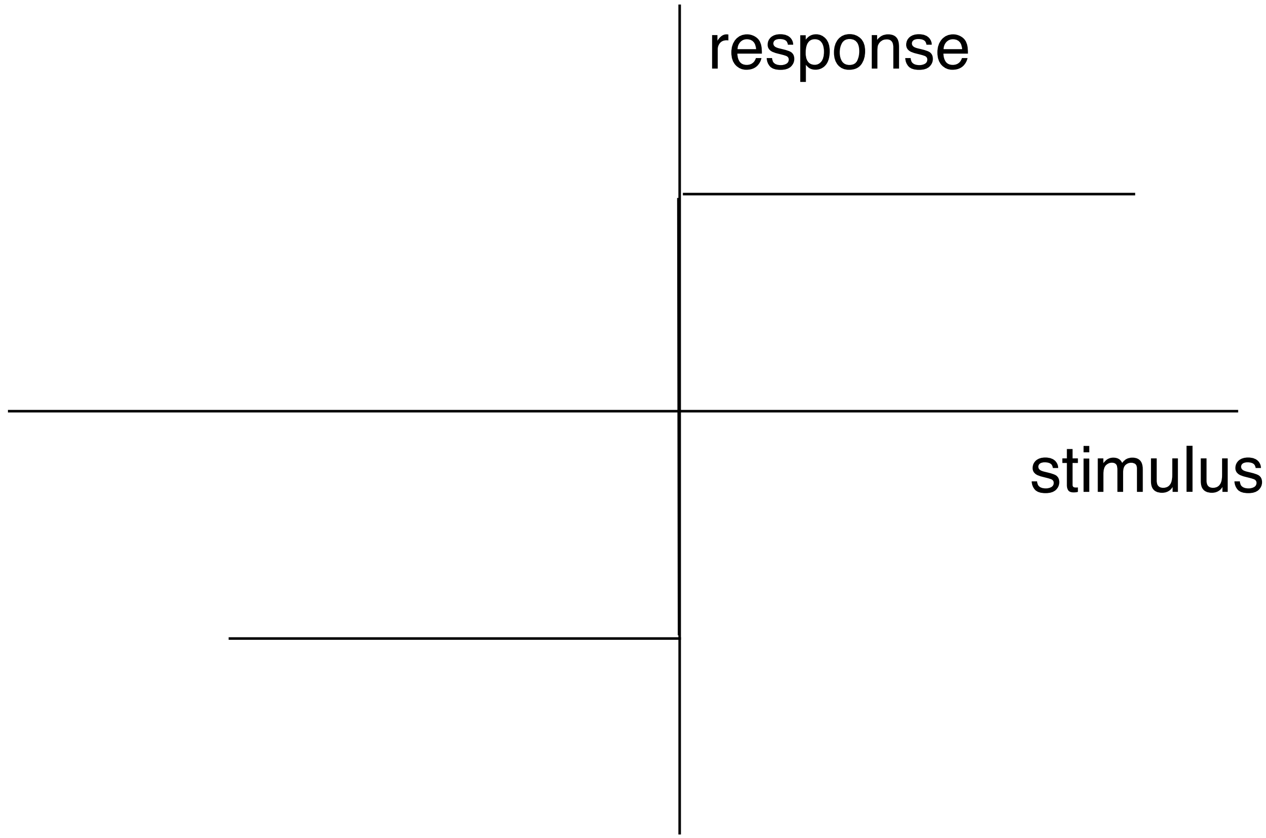

Patterns

- In addition to frequency, information can be encoded based on whether the neuron is spiking regularly (“beating”) or in bursts (“bursting”).

Various Firing Patterns

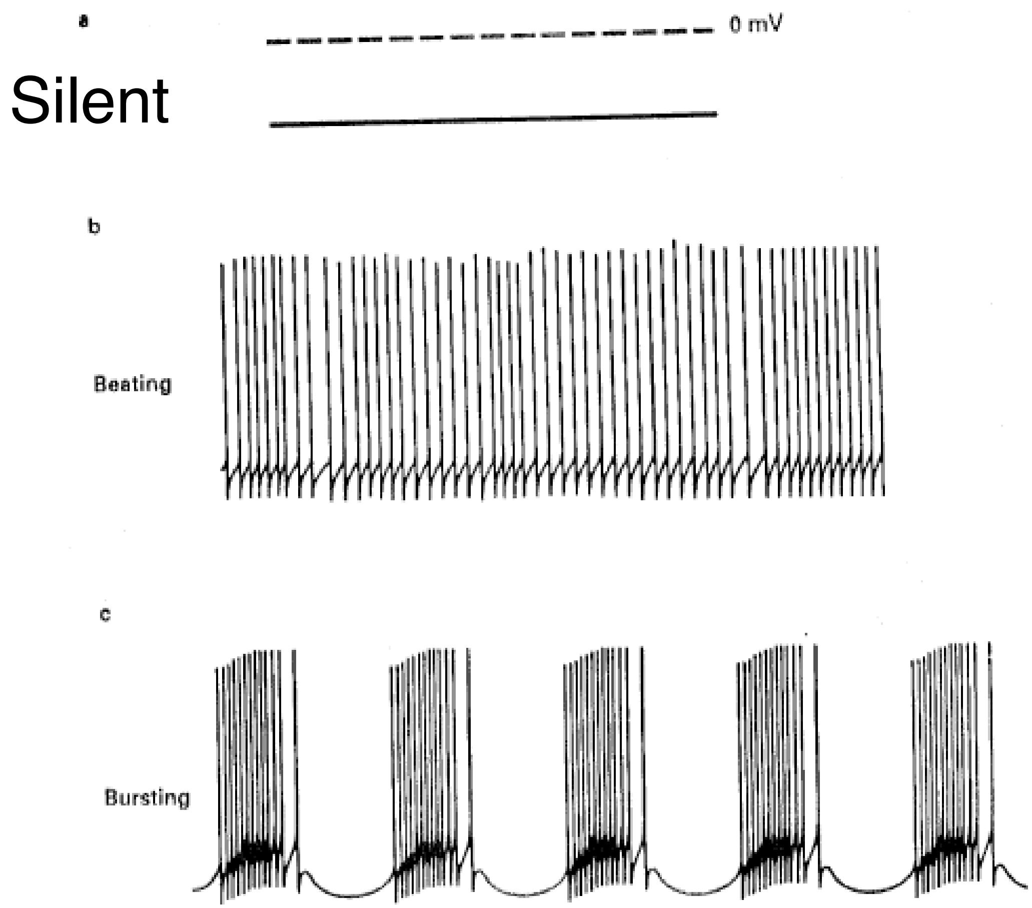

reference: Irwin B. Levitan and Leonard K. Kaczmarek, The Neuron, Oxford University Press, 1991.

Different Patterns of Responses to a Given Stimulation

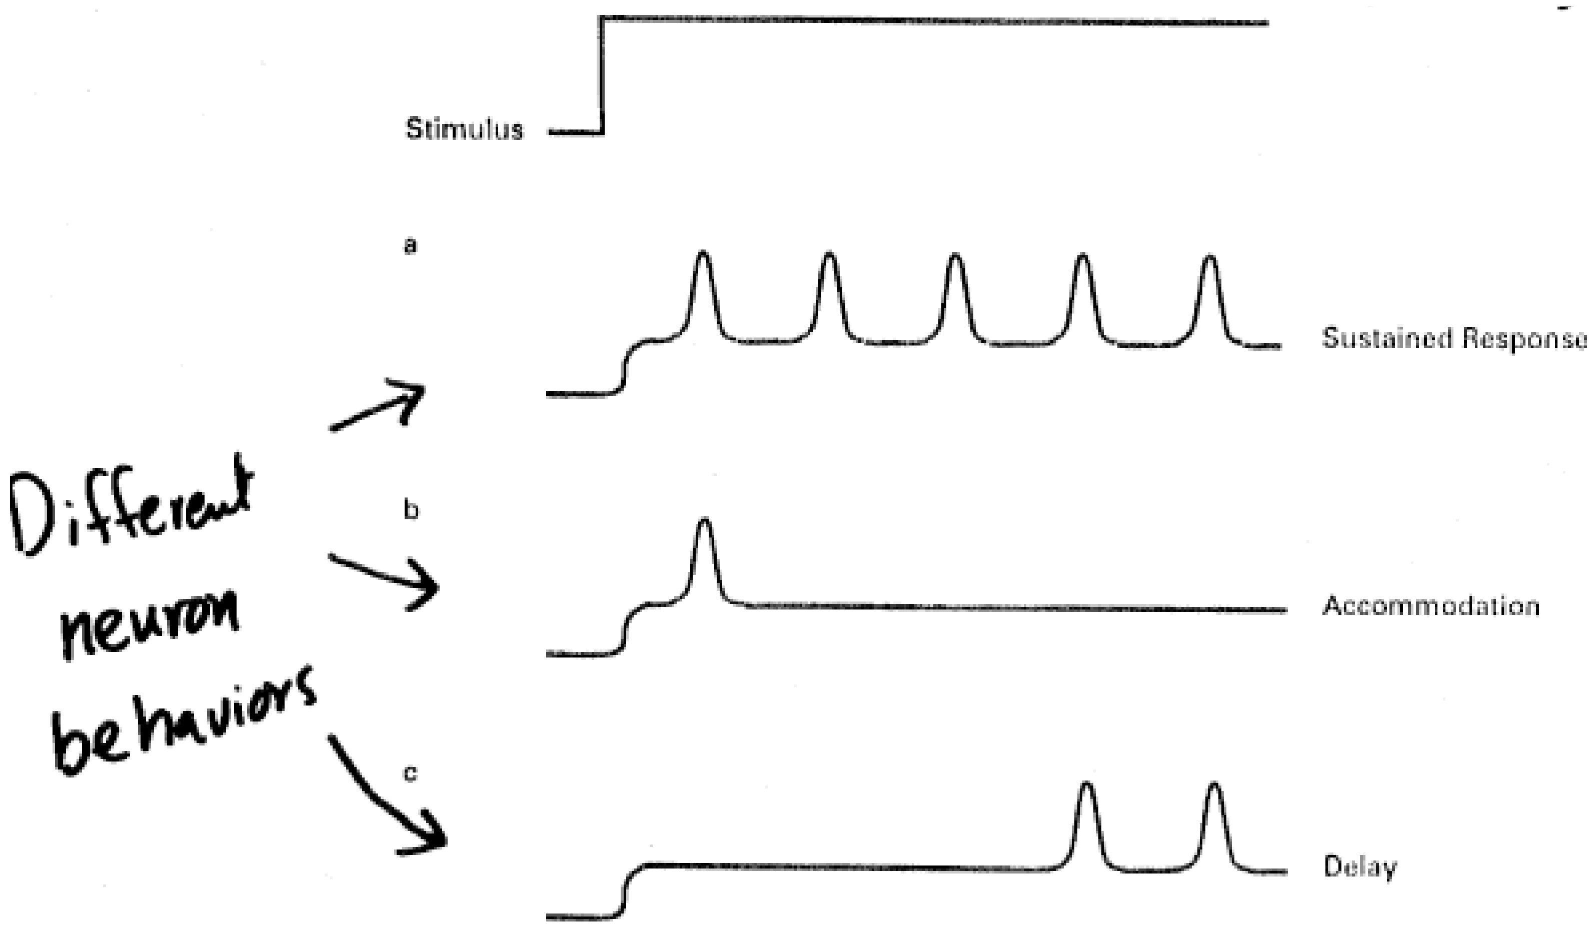

reference: Irwin B. Levitan and Leonard K. Kaczmarek, The Neuron, Oxford University Press, 1991.

Artificial Neural Nets

- Typically in ANNs, we don't try to model the different patterns. We just use a single abstract numeric value as the intensity of the stimulus or response.

Sizes, Scale

- Human estimated to have 10^{10} -10^{11} neurons.
- One neuron may connect to 10^{2} -10^{3} others.
- Therefore 10^{12} -10^{14} connections are present.

Speeds

- Switching speed ~ 1 kHz
  (1 million times slower than a computer)

- Conduction speed ~ 100 m/s
  (vs. near speed of light in a computer)

- Switching energy ~ 10^{-16} joules/op
  (vs. 10^{-5} joules/op for today's computers)

Human Nervous System

- Accounts for 1-2% of body’s weight
- Consumes ~ 25% of body’s energy

Transition to *Artificial* Neural Networks

Neural network schematic

Synaptic "weights"
(strength of connection)

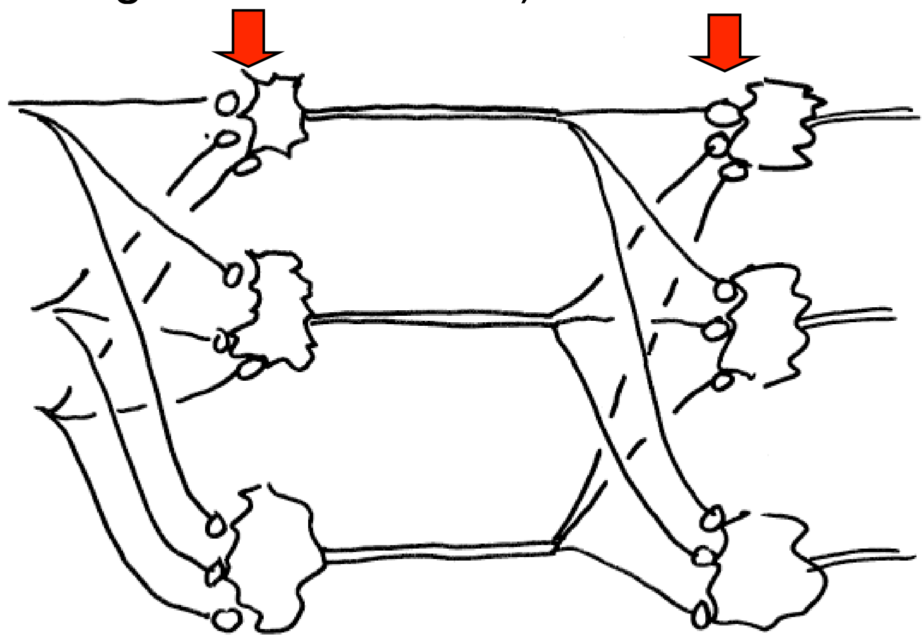

Many-to-many connections

Neural network
(further schematized)

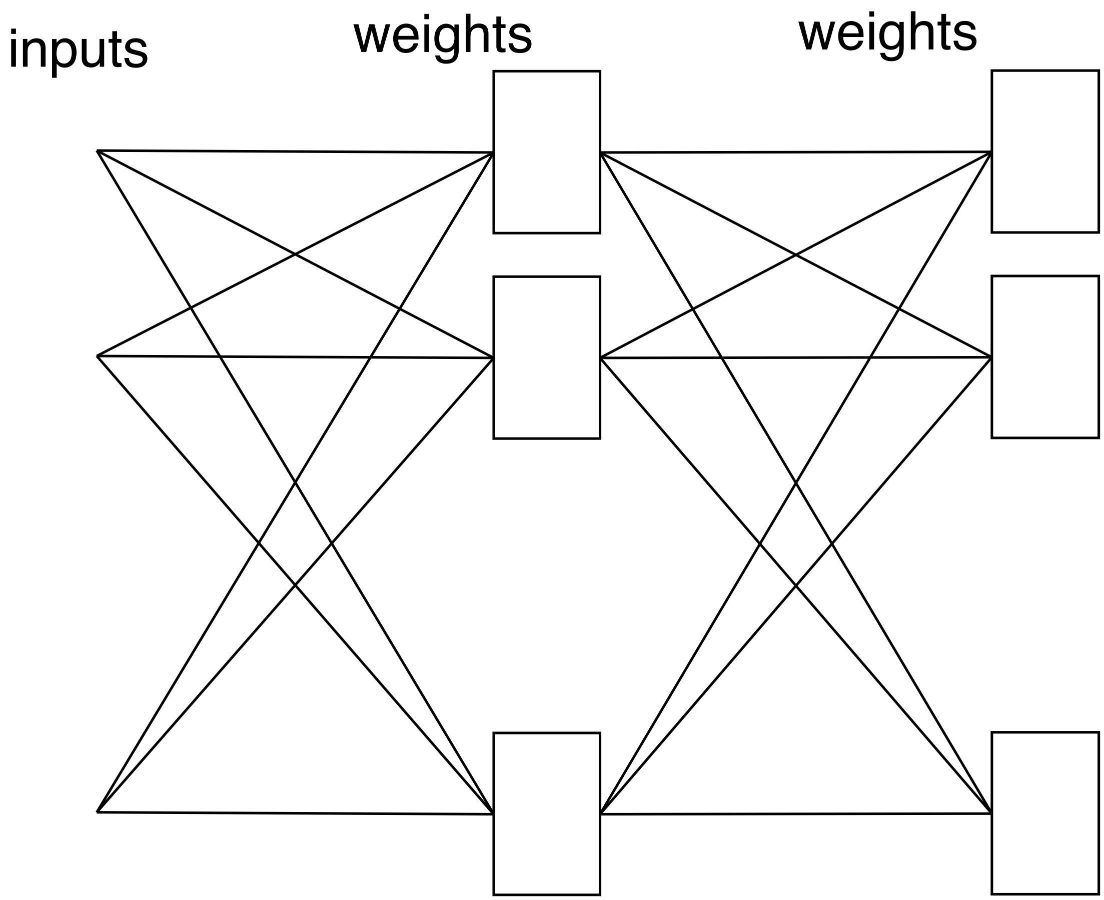

Generality

- Do we lose generality assuming a regular connection pattern?
- Do we lose generality assuming no cycles?

Abstract Functional Characteristics of Neurons

- Weighted sum multiple synaptic inputs
  - positive weight: “excitatory”
  - negative weight: “inhibitory”
- Threshold triggering phenomenon:
  - weighted sum of inputs must exceed threshold in order to cause an event.

Some NN Historical Highlights

- 1943 McCulloch and Pitts, Linear Threshold Logic Gate models
- 1949 Hebb, proposed Learning principle
- 1957 Rosenblatt's Perceptron
- 1960 Widrow & Hoff's Adaline
- 1969 Minsky & Papert (MIT), Limitations of perceptrons

Historical Highlights (cont’d)

- 1970-1980 The “neural-net winter"
- 1974 Werbos (PhD thesis, Harvard), un-noticed discovery of backpropagation: how to train multi-layer networks
- 1982 Hopfield (Princeton, then Caltech) Hopfield networks
- 1986 Rumelhart and McClelland, popularized backpropagation in multi-layer perceptrons, published "Parallel Distributed Processing"

Characteristics of Simple ANN Models

- “weight” = strength of connection
- threshold = value of weighted input below which no response is produced
- signals may be:
  - real-valued, or
  - binary-valued:
    - “unipolar” {0, 1}
    - “bipolar” {-1, 1}

McCulloch-Pitts Model, 1943

- Synchronous operation
- Binary (uni-polar) signals
- Linear threshold gates

McCulloch-Pitts Neural Model

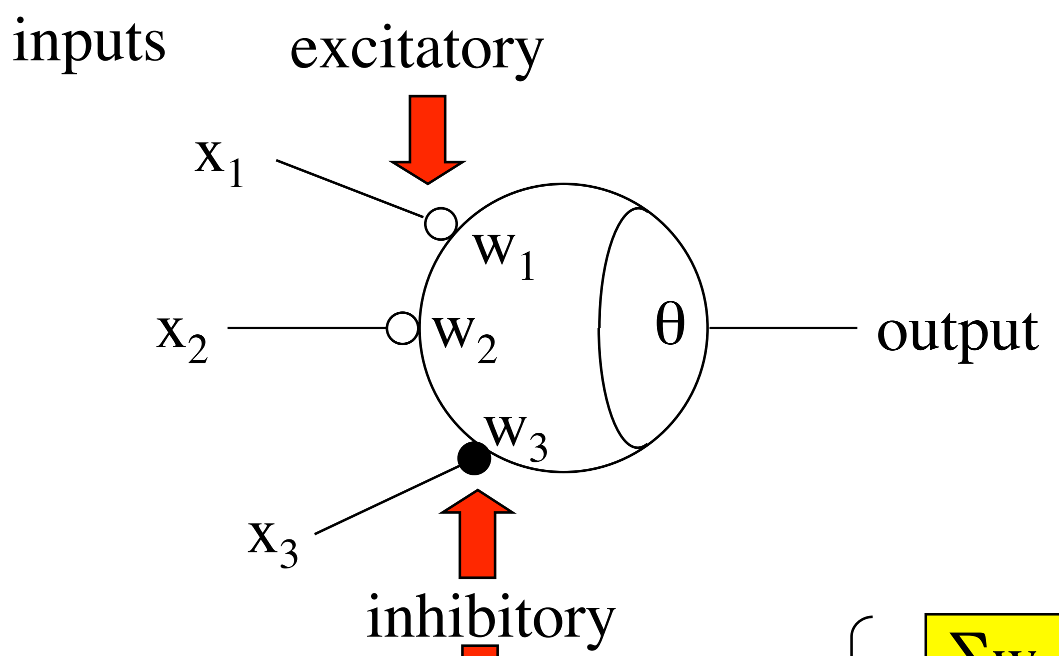

$$\text{output} = \begin{cases} 1 \text{ if } w_1 x_1 + w_2 x_2 - w_3 x_3 > \theta \\ 0 \text{ otherwise} \end{cases}$$

$\sum w_i x_i > \theta$
if we allow weights to be *signed*

How Powerful is a Network of McCulloch-Pitts Neurons?

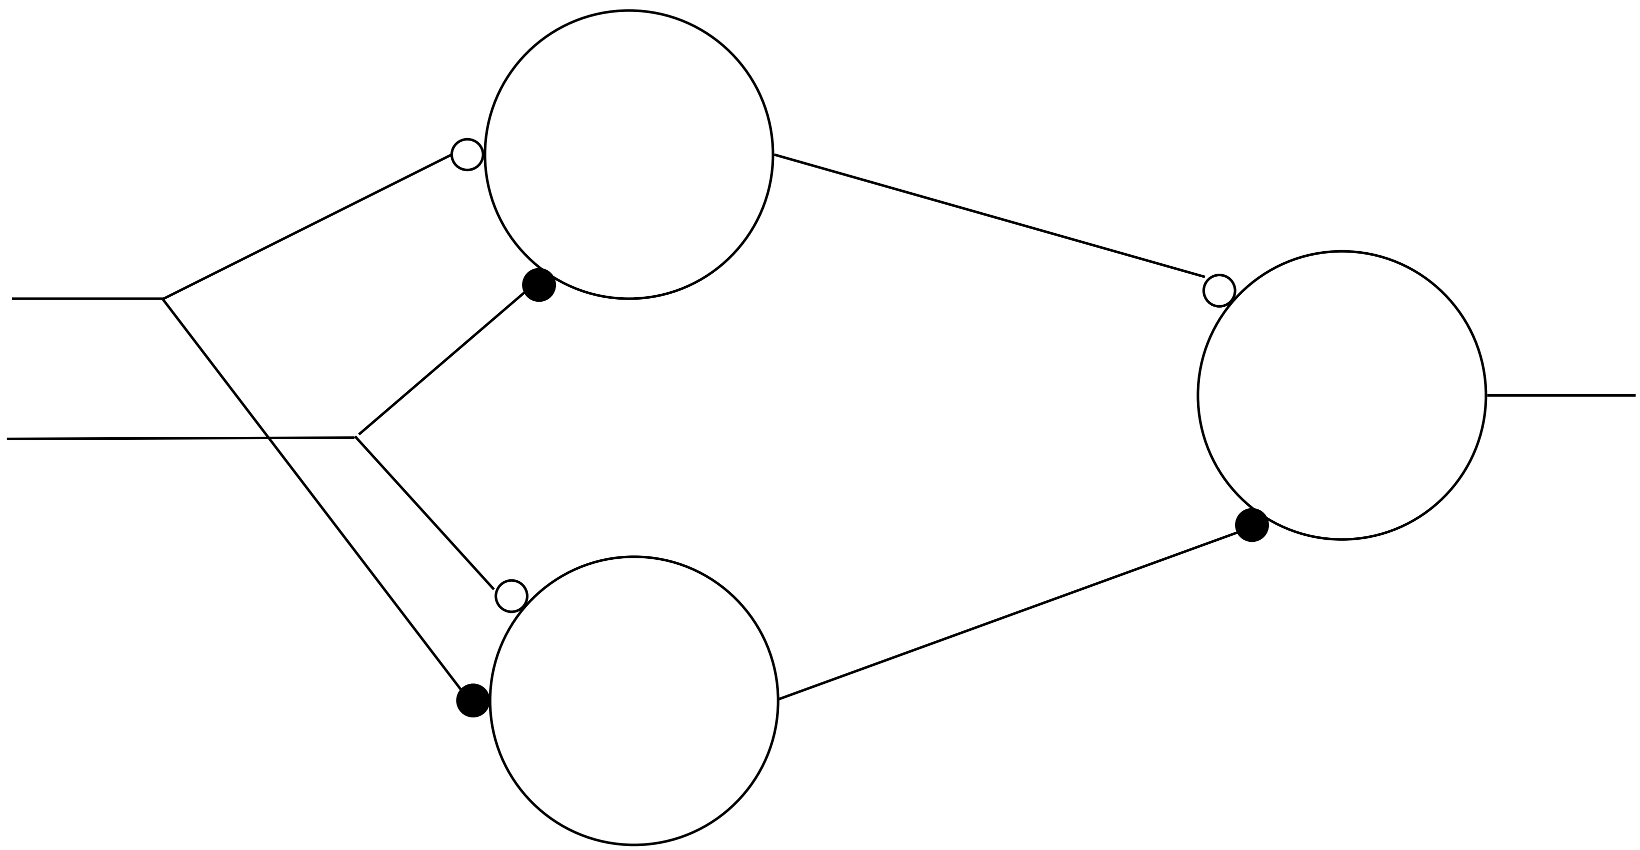

Can *any* switching function be represented?

Kleene’s paper, 1956

- “Representation of Events in Nerve Nets”
- Used McCulloch-Pitts model with possible ***feedback*** connections
- Assumed synchronous model (not realistic?)
- “Events” are essentially what we now call **regular expressions**
- Provides an exact characterization of what McCulloch-Pitts network can do

Perceptrons

Primitive Artificial Neurons

Rosenblatt's Perceptron, 1957

- Introduced the idea of **training**
- A **perceptron** is a linear threshold gate, possibly with real-number inputs, rather than just being limited to {0, 1}

Application for Perceptrons

- Classification problems:
  - Given a pre-specified set of inputs (not necessarily finite), is a given input in the set or not?
  - An input could be a retinal image

Retinal Image Classification

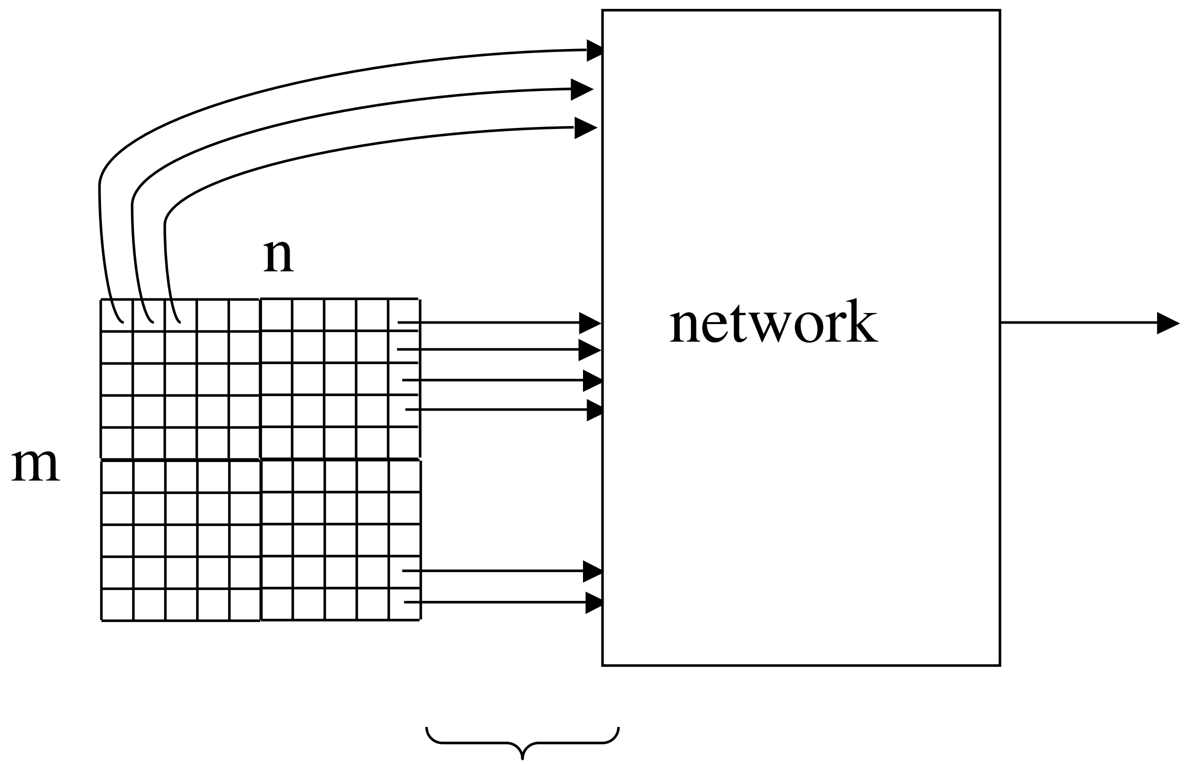

vector of m x n elements $\{x_i\}$, say gray values

Given a classification problem, try to find a perceptron to fit.

Find a vector of weights $\{w_i\}$ and a threshold θ, such that:

$$\text{output} = \begin{cases} 1 \text{ if } \sum w_i x_i > \theta \\ 0 \text{ otherwise} \end{cases}$$

$$= \begin{cases} 1 \text{ if } \{x_i\} \text{ represents a vector in the set} \\ 0 \text{ otherwise (not in the set)} \end{cases}$$

Issues

- **Existence**: Given a set, does there **exist** a perceptron that correctly classifies the set?
- **Solving**: Can the weights be found **analytically**?
- **Training**: Can the weights be found simply by presenting examples that are either in, or not in, the set?

Example

- Suppose the signals are real-valued, and the following vectors are classified as shown:
  - (3, 4) yes (in the set)
  - (6, 1) yes
  - (4, 1) no (not in the set)
  - (1, 2) no
- Is there a perceptron that classifies the set as shown, and if so, what are its weights?

Geometric Insight

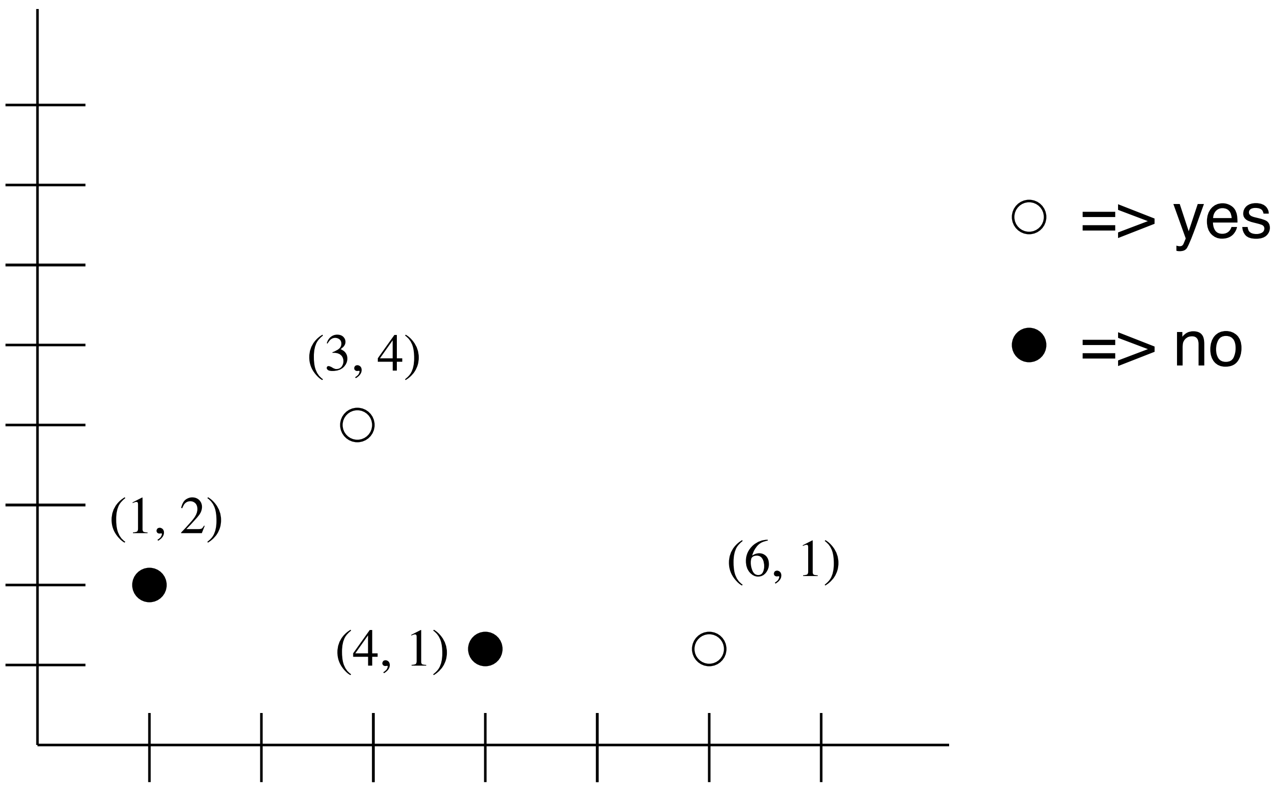

Try to Locate a Separating Line

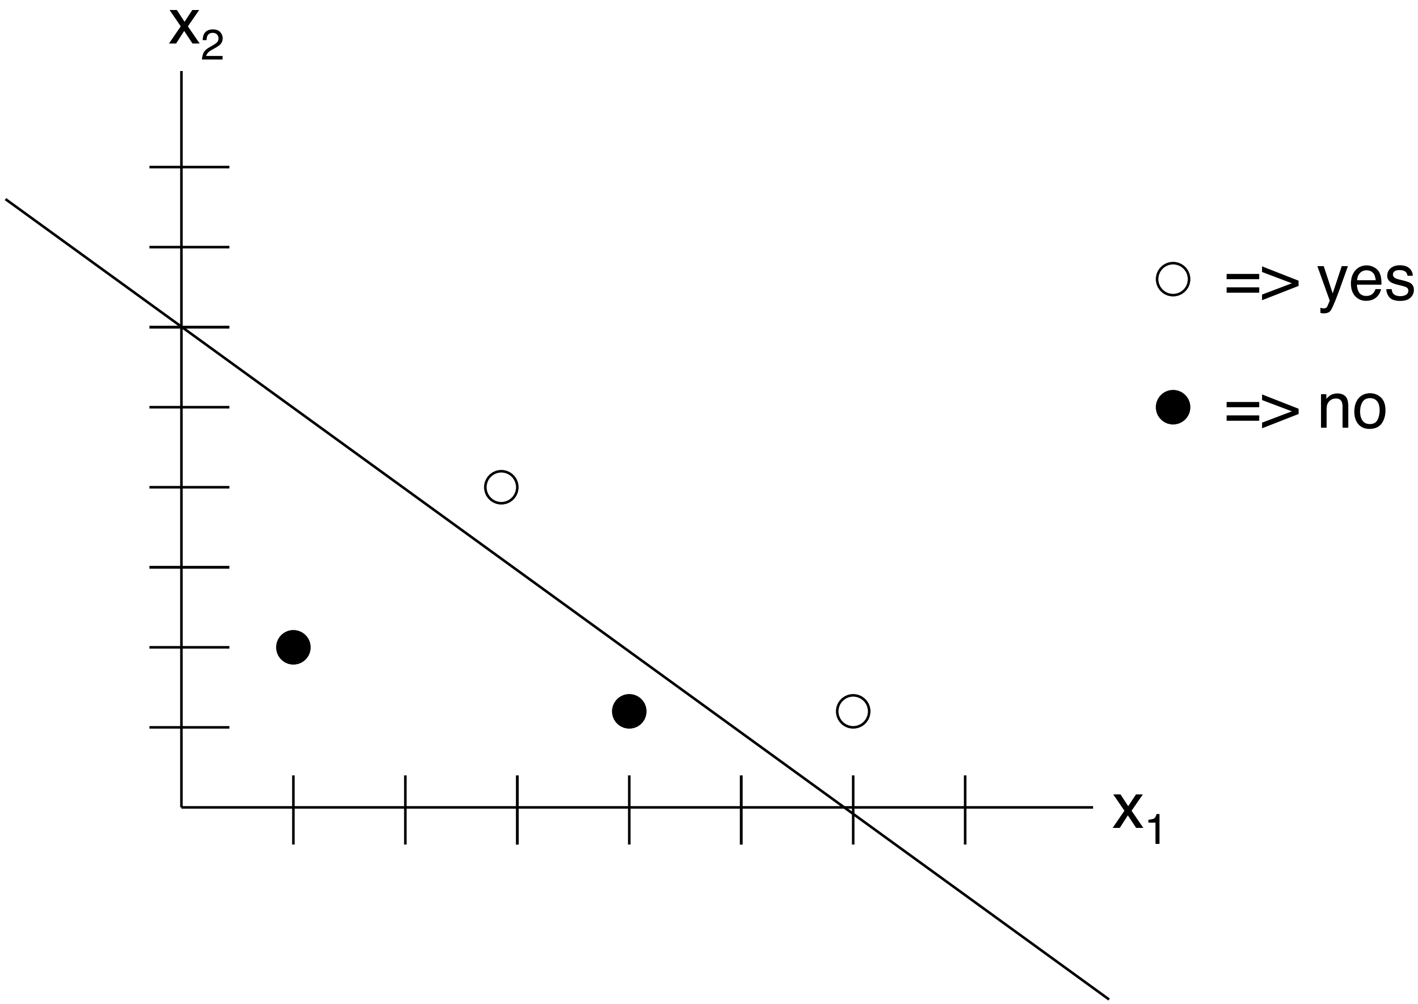

Separating Line Equation

- $x_1 + x_2 = 6$
- Points are
  - in the set if $x_1 + x_2 > 6$
  - not in the set if $x_1 + x_2 \leq 6$

Checking the points

(3, 4)	yes	$3 + 4 > 6$
(6, 1)	yes	$6 + 1 > 6$
(4, 1)	no	$4 + 1 < 6$
(1, 2)	no	$1 + 2 < 6$

General Line Equation

- $w_1 x_1 + w_2 x_2 = \theta$
- Points are
  - in the set if $w_1 x_1 + w_2 x_2 > \theta$
  - not in the set if $w_1 x_1 + w_2 x_2 \leq \theta$

- Will $\{w_i\}$ and θ always exist?

Try to
Locate a Separating Line

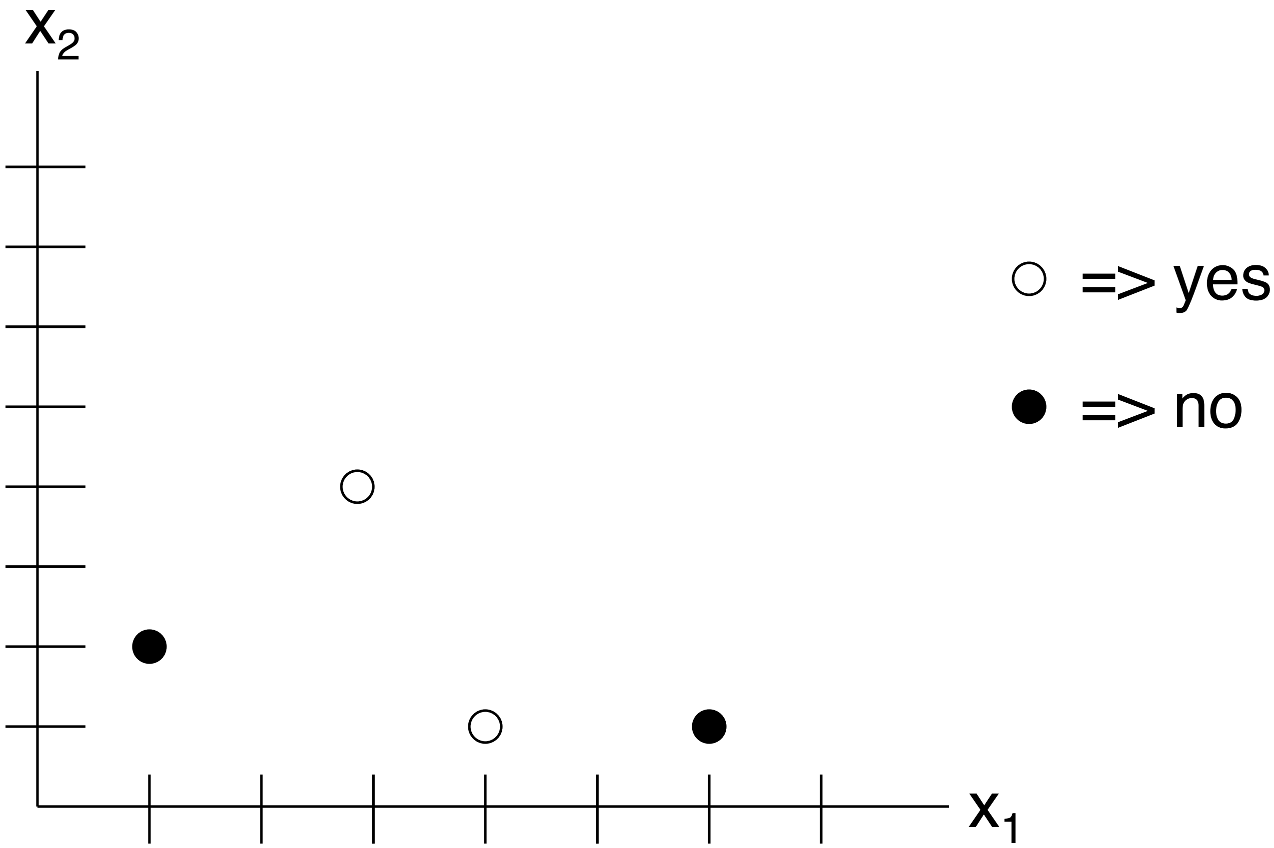

Intuitively, no line exists in this case

- Can you prove it?

Generalizing to n dimensions

- Try to find $\{w_i\}$ and θ such that $w_1 x_1 + w_2 x_2 + \ldots . + w_n x_n = \theta$ separates the points:
  - $w_1 x_1 + w_2 x_2 + \ldots . + w_n x_n > \theta$ when $(x_1, x_2, \ldots ., x_n)$ is in the set
  - $w_1 x_1 + w_2 x_2 + \ldots . + w_n x_n \leq \theta$ when $(x_1, x_2, \ldots ., x_n)$ is not in the set
  - [for now, ignore the possibility of = θ; revisit this later.]

Separating Hyperplane

- The equation
  $w_1\, x_1 + w_2\, x_2 + \ldots. + w_n\, x_n = \theta$
  defines a *hyperplane* in n-space

- If such a hyperplane exists for a classification problem, the problem is called **linearly-separable**.

Geometry

- The vector of weights (w_1, w_2 , …. w_n) is normal (**perpendicular**) to the hyperplane.

- Threshold is the distance of the hyperplane from the origin.

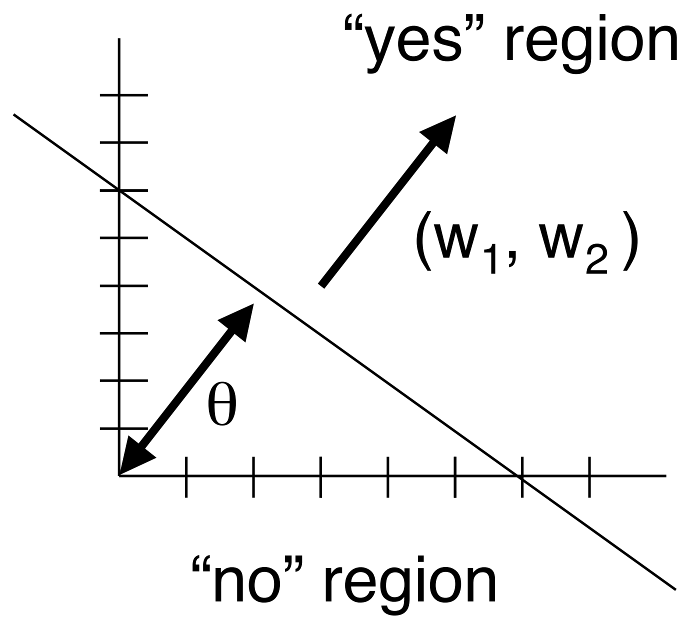

Perceptron Summary

- A perceptron can solve a classification problem iff the problem is linearly-separable.

- There are problems a single perceptron cannot solve.
- Perhaps the simplest unsolvable one is the **XOR problem**:

  yes: {(0, 1), (1, 0)}, no: {(0, 0), (1, 1)}

Some Questions to Ask

- Which switching functions are linearly separable?
- Are there any closure properties (such as closure under union or intersection) for linear separability?

Multi-level Perceptrons

- One way to alleviate the limitations of a single perceptron is to build *networks* of perceptrons.
- Another might be to allow non-linear expressions (such as squaring).
- Can the XOR problem be solved by such models?

What weights and thresholds for XOR?

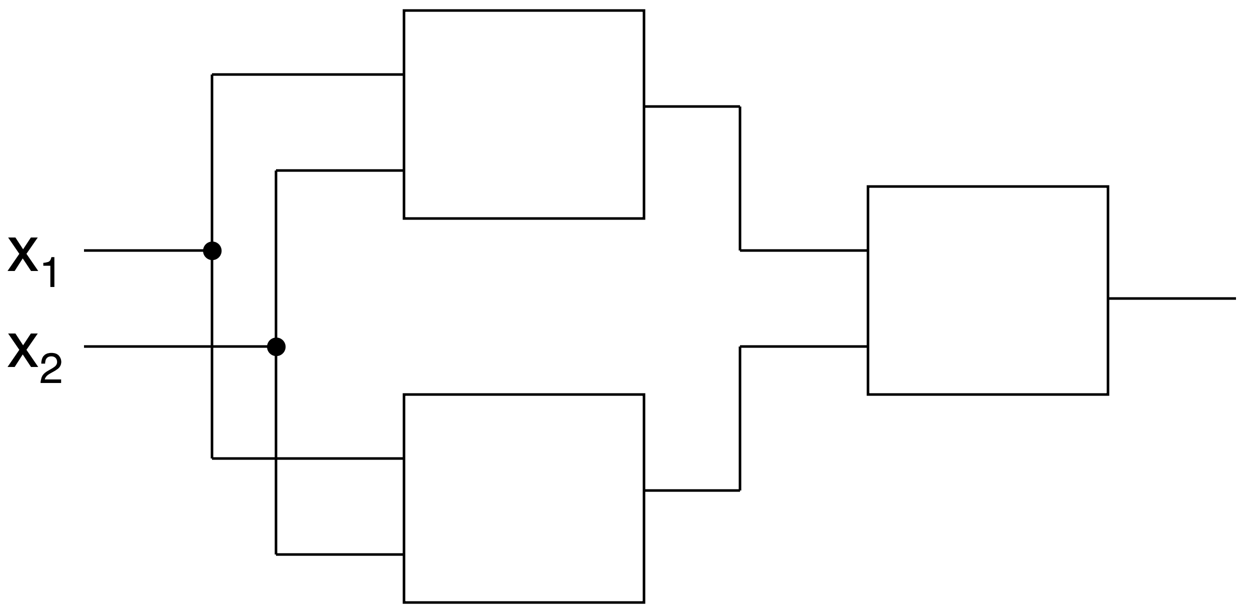

Learning

How might a neural network learn?

Hebb's Postulate, 1949

The Organization
of Behavior

A NEUROPSYCHOLOGICAL THEORY

D. O. HEBB
McGill University

1949
New York · JOHN WILEY & SONS, *Inc.*
London · CHAPMAN & HALL, *Limited*

Hebb's Postulate

A NEUROPHYSIOLOGICAL POSTULATE

Let us assume then that the persistence or repetition of a reverberatory activity (or "trace") tends to induce lasting cellular changes that add to its stability. The assumption * can be precisely stated as follows: *When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.*

Hebb’s Postulate

When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it,

some growth process or metabolic change takes place in one or both cells

such that A’s efficiency, as one of the cells firing B, is increased.

Hebb Restated (Levitan and Kaczmarek)

"When a postsynaptic neuron becomes depolarized [fires], it generates a biochemical reaction or a trophic factor that stabilizes [strengthens] the excitatory synapses that are firing at that time."

Colloquial Hebb

Neurons that fire together wire together.

Levitan and Kaczmarek

"An important aspect of [Hebb's] hypothesis is that a given presynaptic input to a cell need not, *by itself*, be of sufficient strength to induce a large depolarization in its target.

If that input is fired at the same time as a number of other inputs, and their combined action depolarizes the cell, all of these inputs will tend to be stabilized."

Levitan and Kaczmarek (cont'd)

- "If, in contrast, a given input fires asynchronously with most of the other inputs onto that cell, this input will tend to be eliminated."
- [This could be called "anti-Hebbian" learning.]

Levitan and Kaczmarek (cont'd)

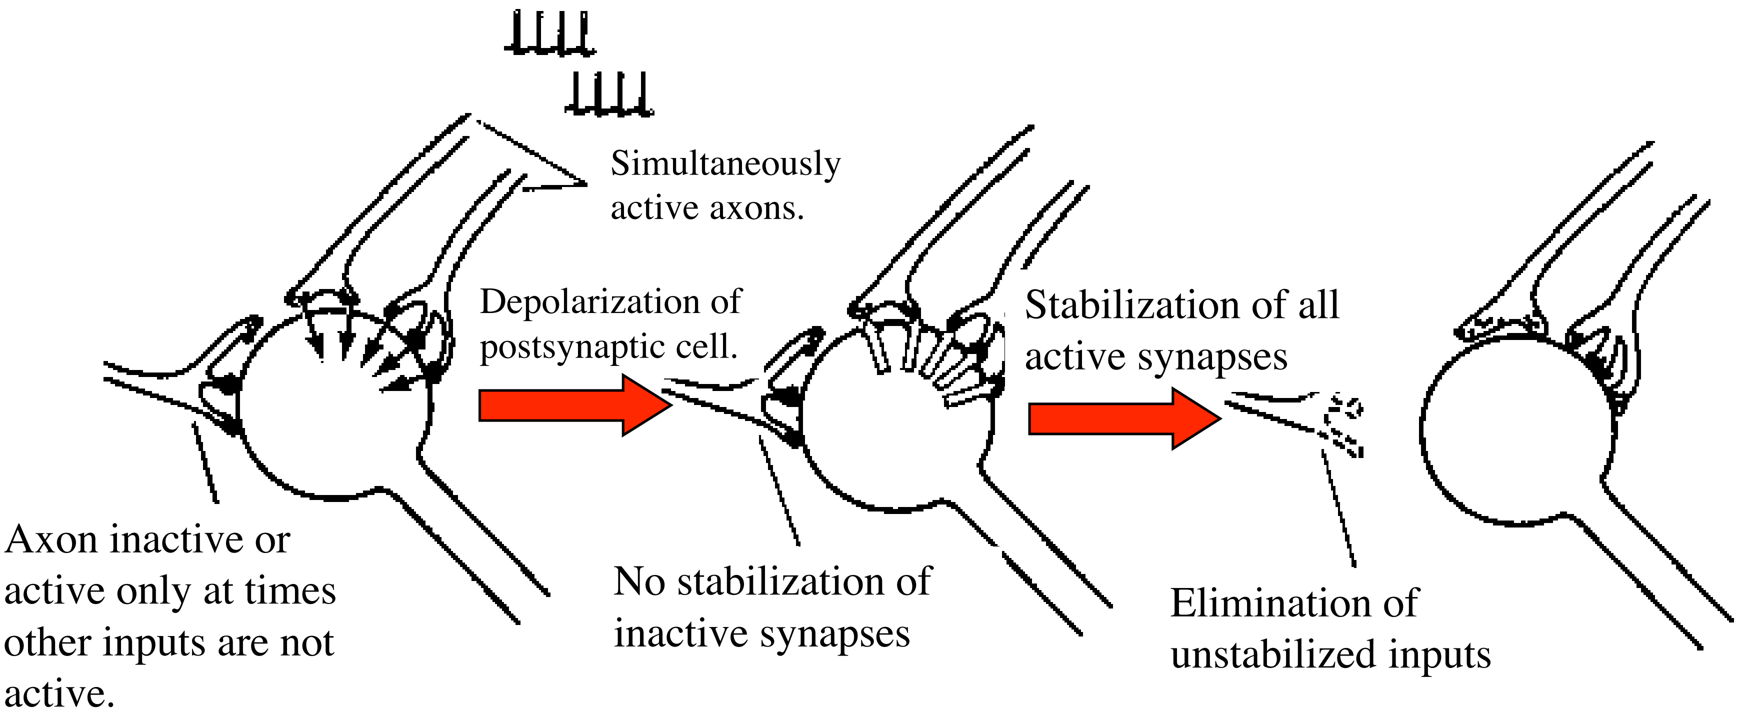

Hebb's rule. Excitatory synapses that successfully stimulate a post-synaptic neuron, or are active when the postsynaptic neuron is depolarized, are selectively stabilized.