COMP 40260

Connectionism (Artificial Neural Networks) and Dynamical Systems

doing it numerically.....

Course overview

- 12 weeks, 2hr lecture (Mon, B1.08) plus 2 hr practical; (Monday Afternoon, 2-4, B1.08).
- Textbook: Elman, J. E. et al, "Rethinking Innateness: A Connectionist Perspective on Development", MIT Press, 4^{th} ed, 1999
- Theoretical **and** hands-on practical course
- Course website: http://cogsci.ucd.ie/Connectionism

Software

- Text book uses the free tLearn programme
- Stability issues, this software is dead.
- We will use some customized software suitable for learning, but not large scale simulations: Basic Prop (basicprop.wordpress.com)
- No programming experience presumed

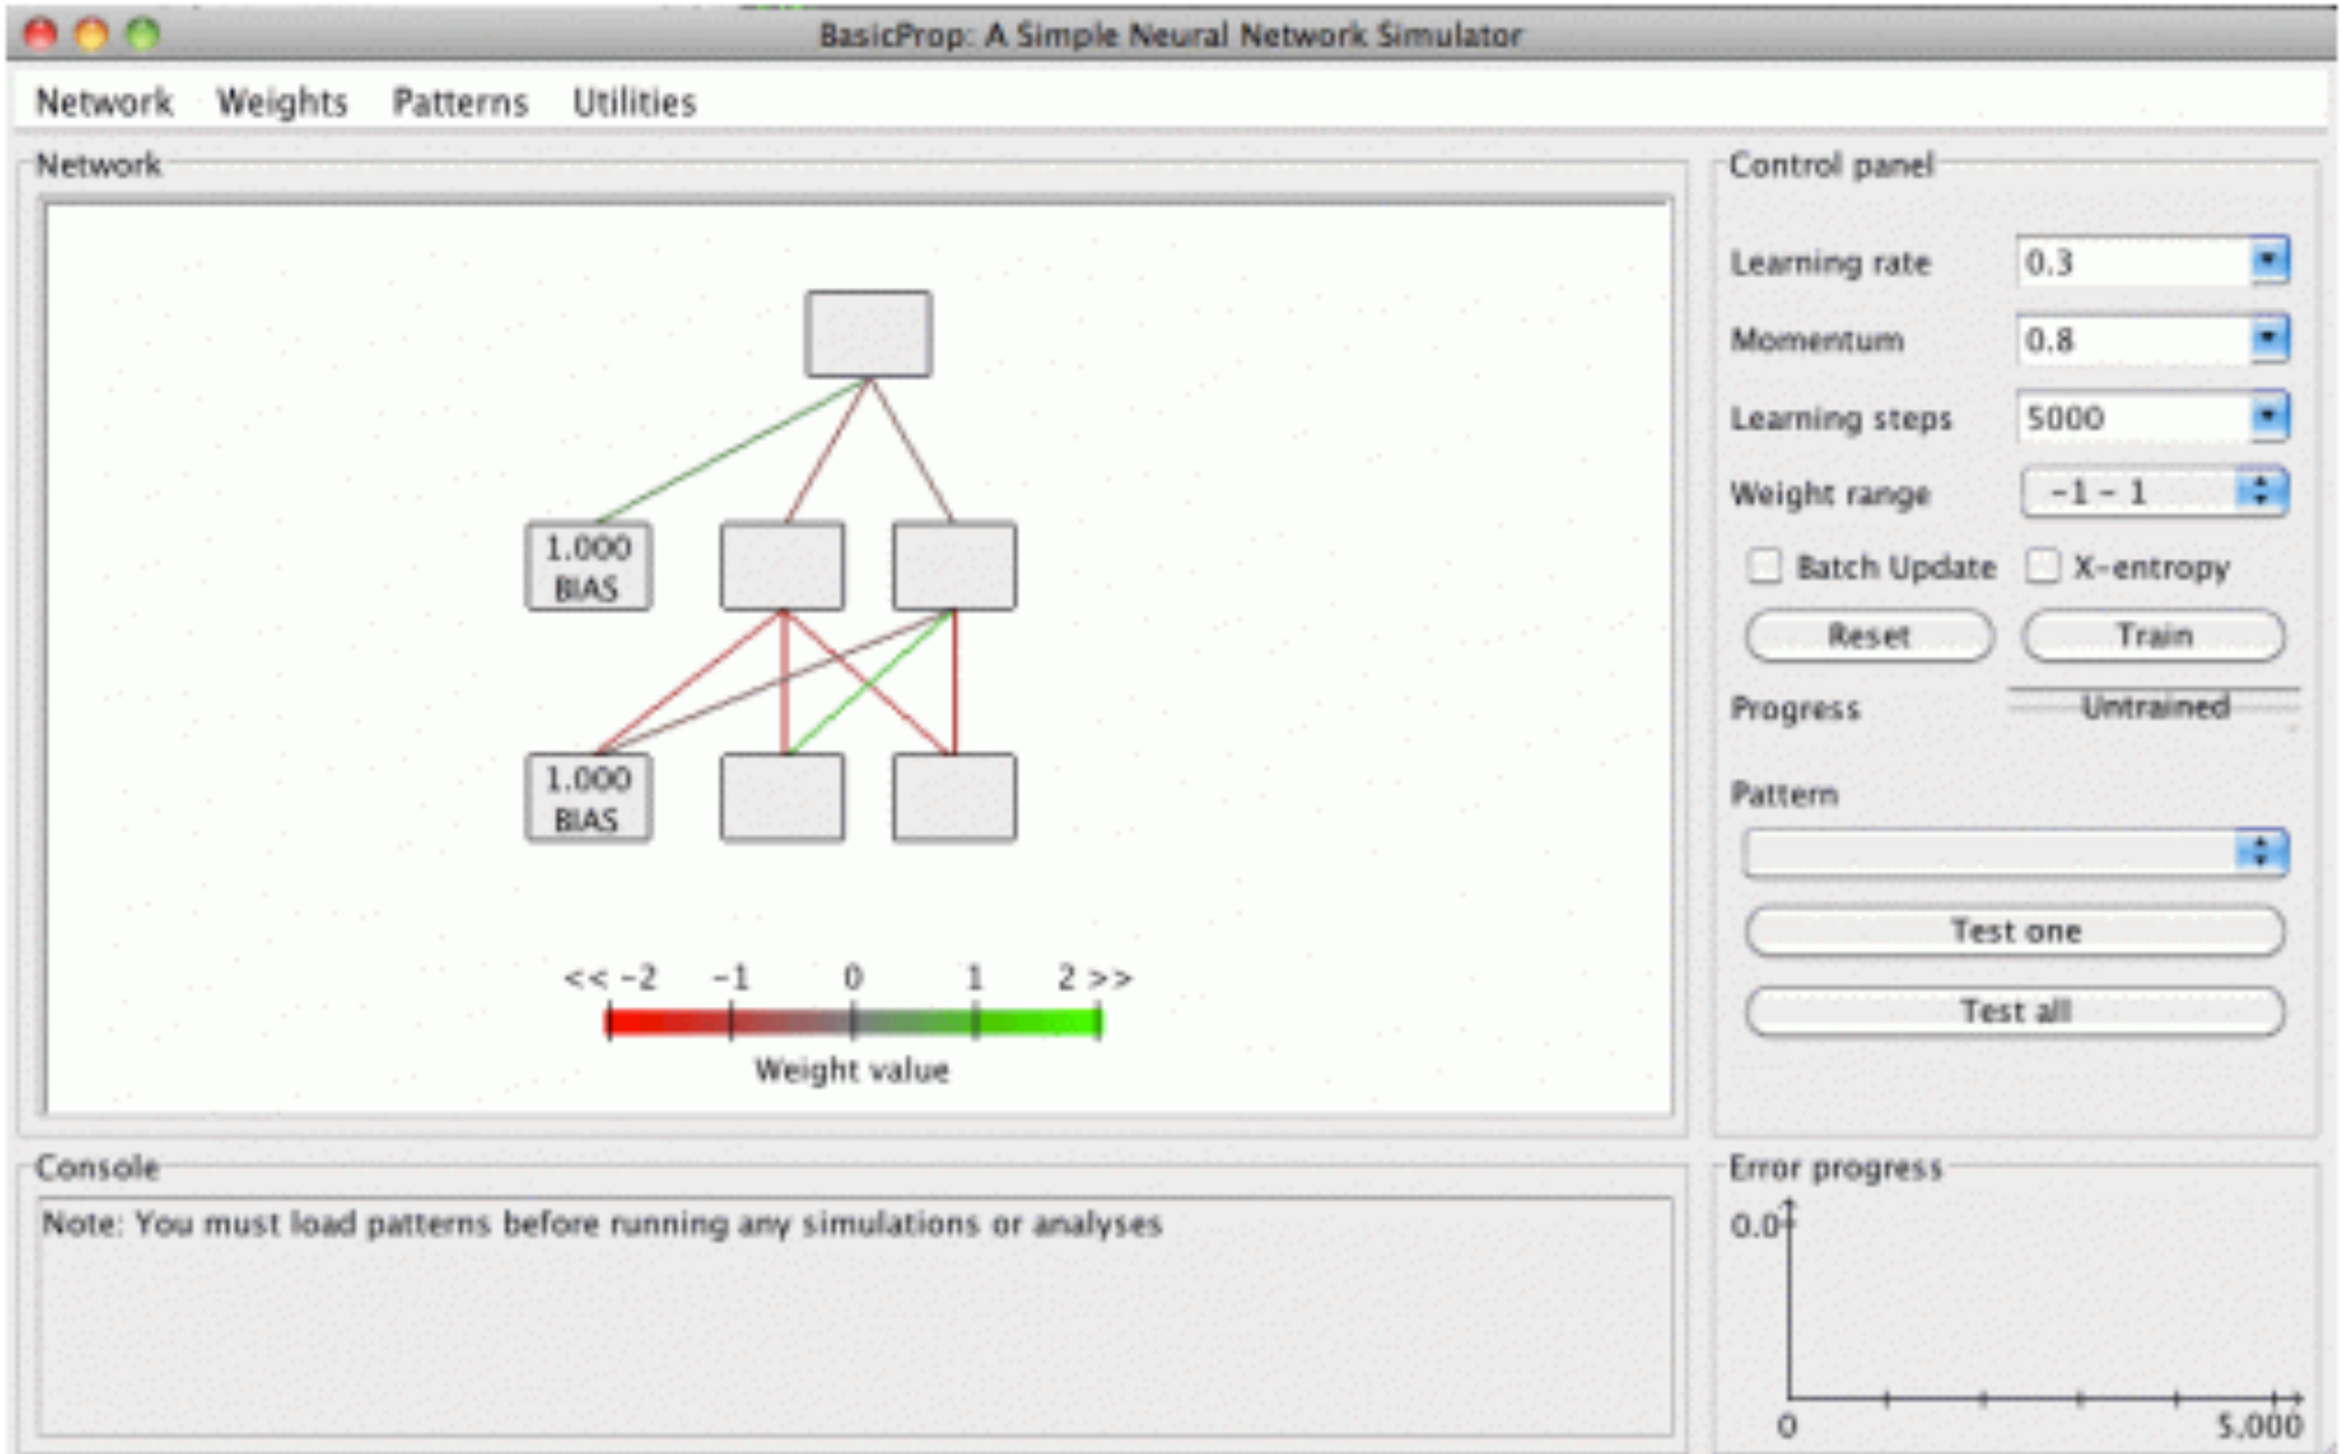
BasicProp: A Simple Neural Network Simulator
Network
Weights
Patterns
Utilities
Network
1.000
BIAS
1.000
BIAS
<< -2
-1
0
1
2 >>
Weight value
Control panel
Learning rate
0.3
Momentum
0.8
Learning steps
5000
Weight range
-1 - 1
Batch Update
X-entropy
Reset
Train
Progress
Untrained
Pattern
Test one
Test all
Console
Note: You must load patterns before running any simulations or analyses
Error progress
0.0
0
5,000

Exercises etc

- Evaluation by 3 written exercises and 1 short essay
- Due dates are on the webpage
- Essay due shortly after exams are over.

Part 1: Why Connectionism?

- Neural Networks (ANNs, RNNs)
- (Multi-layered) Perceptrons
- Connectionist Networks
- Parallel Distributed Processing (PDP)

Numerical computational models, with simple processing units, and inter-unit connections.

The great debate

- Innate knowledge vs Inductive learning
- Rationalist vs Empiricist
- Nativist vs Tabula Rasa (nature/nurture)
- Extreme points on a continuum
- Interaction of maturational factors (genetically controlled) and the Environment

Nature versus nurture is a lie. Music is not melody versus rhythm, wine is not grapes versus alcohol and we are not environment versus genes. We are their sum, their product and their expression. They dance together and we are their performance, but neither is an adversary.

Vaughan Bell, Mind Hacks

Reason 1: We need a formal theory of learning

What would a formal theory of learning look like?

Things to consider:

- What are the "atomistic" elements to your theory?
- How do those stand in relation to the world?
- What are your underlying assumptions about the nature of learning?
- Are you modelling humans? brains? The world?
- Given your elements, what processes can you model?
- How do you evaluate your success?

Some areas in which our views have evolved greatly:

1. Biology

- Genetics is a fast moving field.
- In about 1980, the accepted view was that genes were blueprints for the adult phenotype
- This placed responsibility on the genes for all “innate” features of the mature organism, and evoked a clean distinction between organism and its world
- This is not longer a plausible view
- Genes are switches that play a role in mediating the relationship between organism and world
- Many world and organism involving features affect the operation of genes

Some areas in which our views have evolved greatly:

2. Neuroscience

- Neural plasticity, in which structure and function of the CNS are altered as a result of experience, is now seen as a central process in all learning
- Neural plasticity is a life long phenomenon, not confined to immature organisms.
- Imaging techniques (eeg, fMRI, PET, MEG) have allowed for observation of some kinds of ongoing brain activity
- Underlying view of what the functional elements of the nervous system are, and what functions they might instantiate, are changing rapidly, and will continue to change

Some areas in which our views have evolved greatly:

3. Data and Machine Learning

- **Data science** has provided *vastly* greater data sets than were ever available before
- Along with these, **machine learning** techniques have exploded, and been refined on a huge variety of tasks
- Most machine learning seeks to sort, categorise, or find structure *automatically* in large data sets
- Many ML approaches support continuous online learning, allowing systems to be continually presented with new data.

Classical models of serial processing in digital computers with rule-like control structures are not much like brains.

Reason 2: Need a brain-like modelling paradigm

even if we are unsure what real brains are *doing!*

The Neuron Doctrine

Camillo Golgi

Nobel Prize in Physiology or Medicine, 1906

First shared nobel prize in medicine and physiology

Ramon Y Cajal

Golgi & Cajal did not agree, and they gave acceptance speeches criticizing each other!

Golgi had invented a staining method that allowed the intricate branching structure of cortical neurons to be seen for the first time

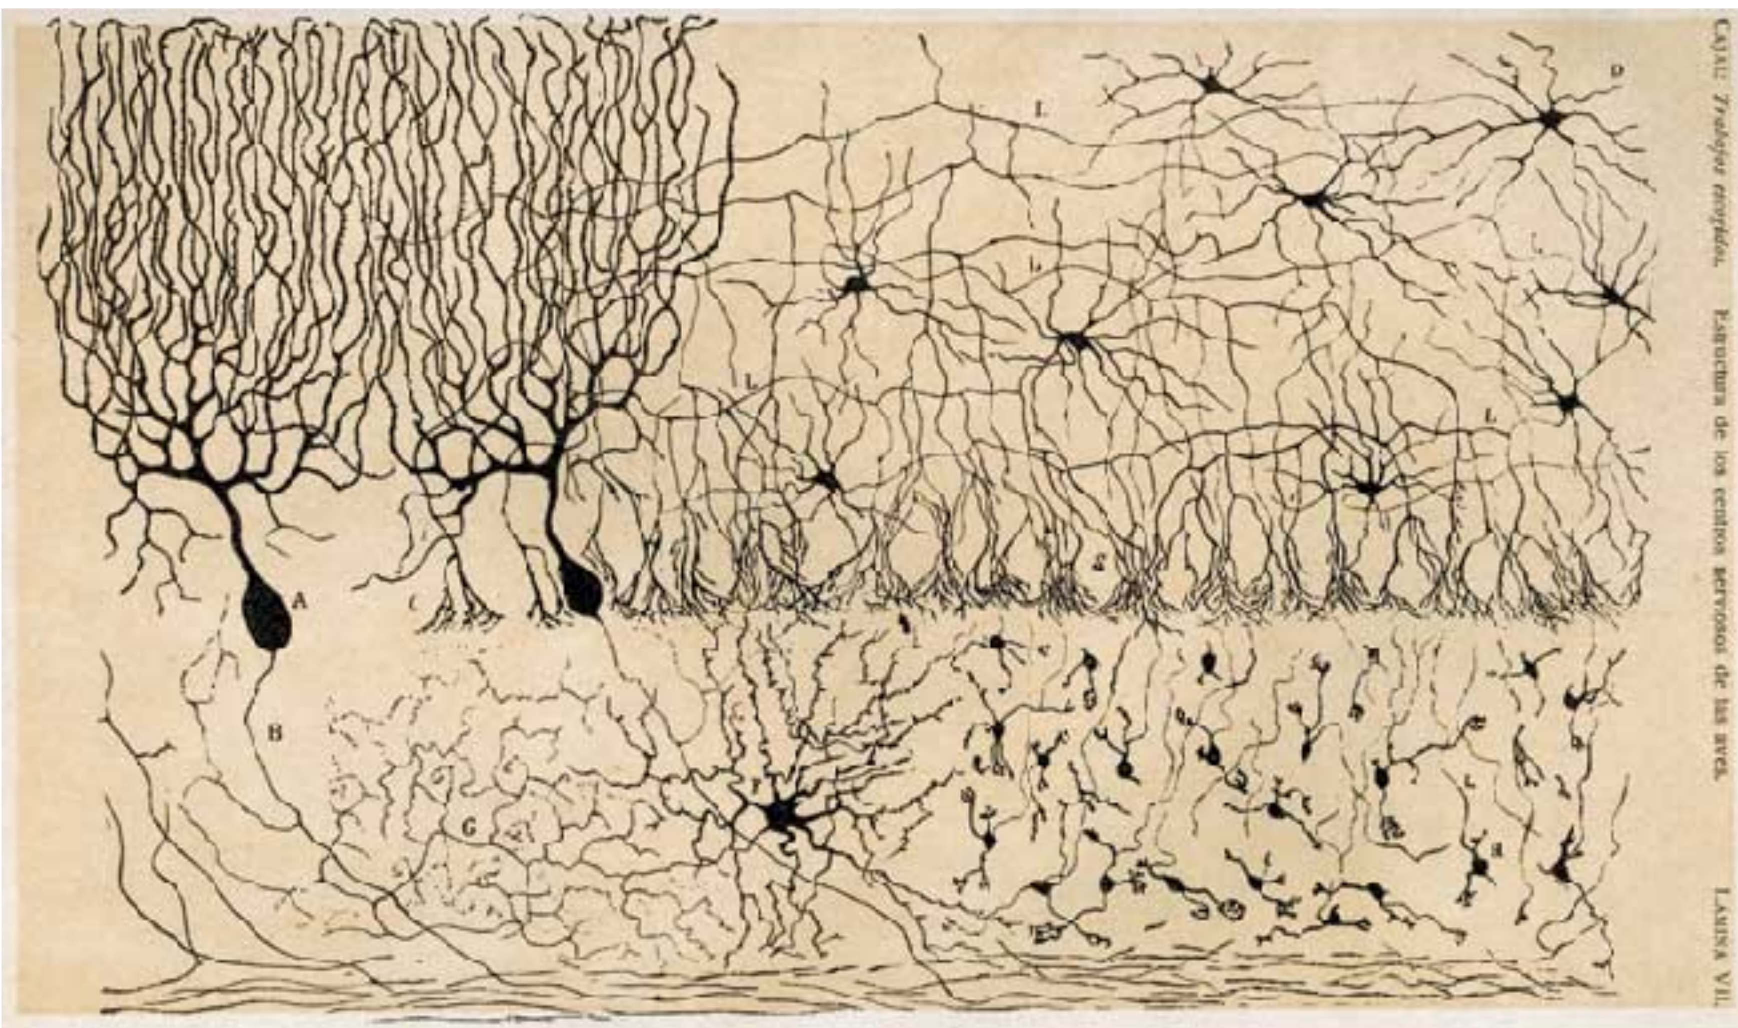

Drawing by Cajal

For Golgi, the basic functional unit of the neural system was a network of neurons, with intertwined axons. Dendrites were for nutrition only

For Cajal, the basic functional unit was the individual neuron. Individual elements did not fuse together. Dendrites provided input, axons produced outputs.

Over time, Cajal's view became the orthodoxy. It remains a central tenet of modern neuroscience . . .

. . . BUT it has gradually come under threat.

Many exceptions to the neuron doctrine

- Some spikes propagate backwards
- Some inter-neuron connections are electrical
- Glial cells (surrounding neurons) also engage in some kind of information processing
- Synapses are vastly more complicated than previously realised
- Some groups of neurons act as functional units

- Some groups of neurons act as functional units

The starting point for ANN theory lies in the early application of the concepts of information theory (Claude Shannon, Bell Labs)

(1943)

Warren S. McCulloch and Walter Pitts

A logical calculus of the ideas immanent in nervous activity

Bulletin of Mathematical Biophysics 5:115–133

Because of the "all-or-none" character of nervous activity, neural events and the relations among them can be treated by means of propositional logic. It is found that the behavior of every net can be described in these terms, with the addition of more complicated logical means for nets containing circles; and that for any logical expression satisfying certain conditions, one can find a net behaving in the fashion it describes. It is shown that many particular choices among possible neurophysiological assumptions are equivalent, in the sense that for every net behaving under one assumption, there exists another net which behaves under the other and gives the same results, although perhaps not in the same time. Various applications of the calculus are discussed.

Warren McCulloch:
Psychiatrist,
Experimental
Epistemologist,
Poet,
Militarist,
Theological Engineer

(see additional reading for more)

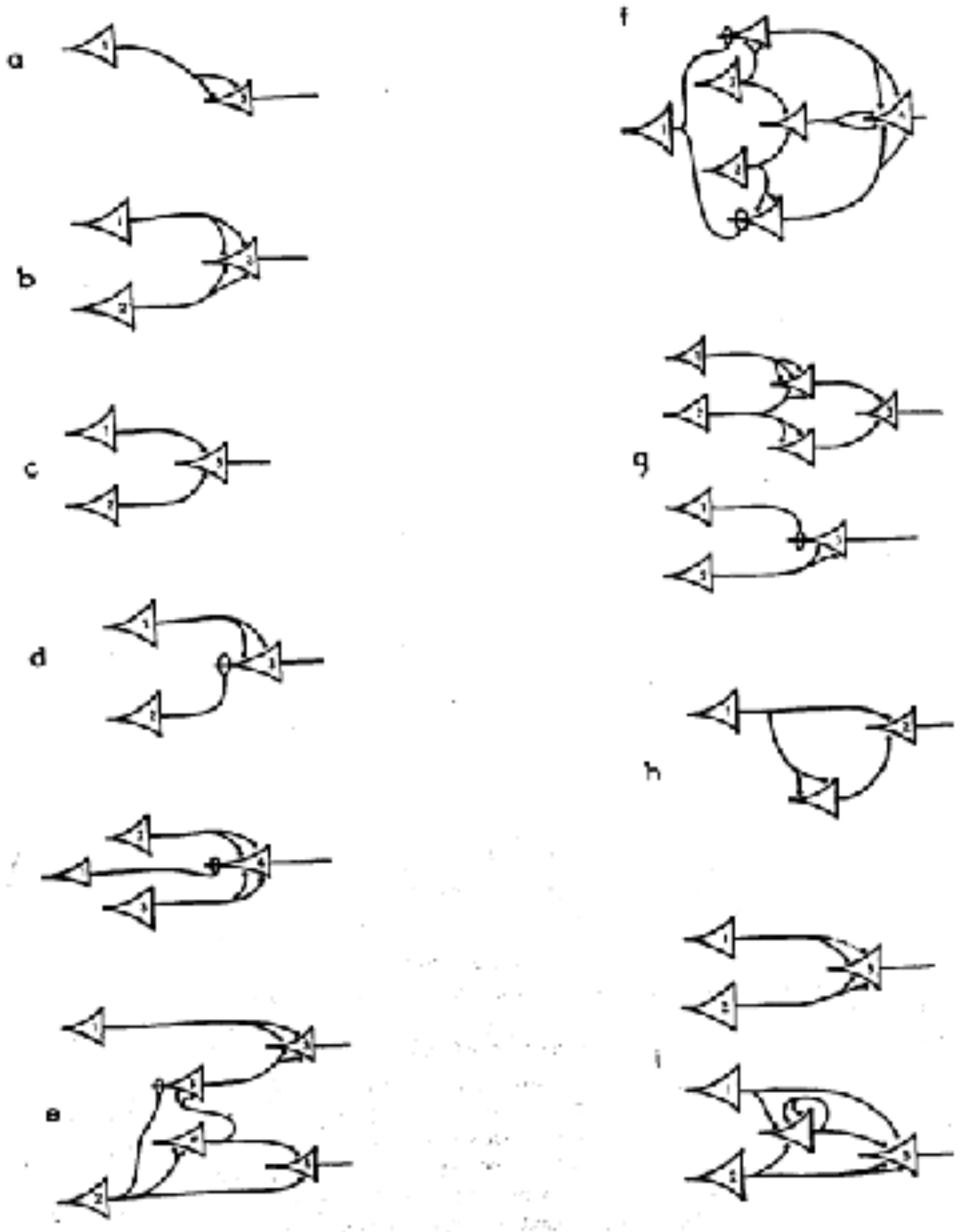

Figure 1

Hypothetical computation of logical functions by simple processing units

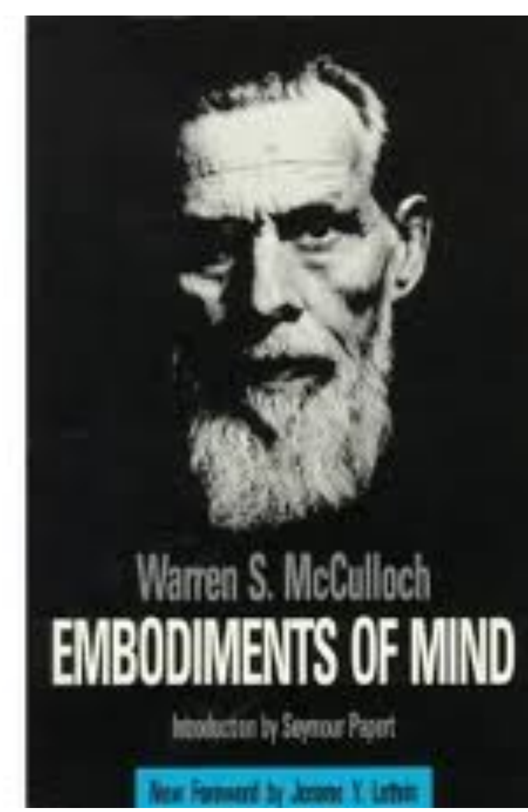

Heinz von Foerster's Principle of Undifferentiated Encoding:

"The response of a nerve cell does *not* encode the physical nature of the agents that caused its response. Encoded is only `how much' at this point on my body, but not `what'." (1973)

For Cajal and for McCulloch, the individual neuron is the computational atom. To understand brain activity, one tries to understand computations done by individual cells

In contemporary neuroscience, individual cells are frequently not assumed to be the basic computational unit. Brain activity is viewed at the level of *transient aggregates of neurons*, generating *dynamic local field potentials.* Those neurons who are grouped together in the context of one task, may disband and form other groupings in other contexts.

Interpreting brain activity is something of a black art.

How do we model the interface between an organism and its environment?

Sharp split between subject and world

Co-specification of subject and world

The sensorimotor interface with the world is very high dimensional….. and very important

Reason 3: Rich modelling of coupling relations between organism and world

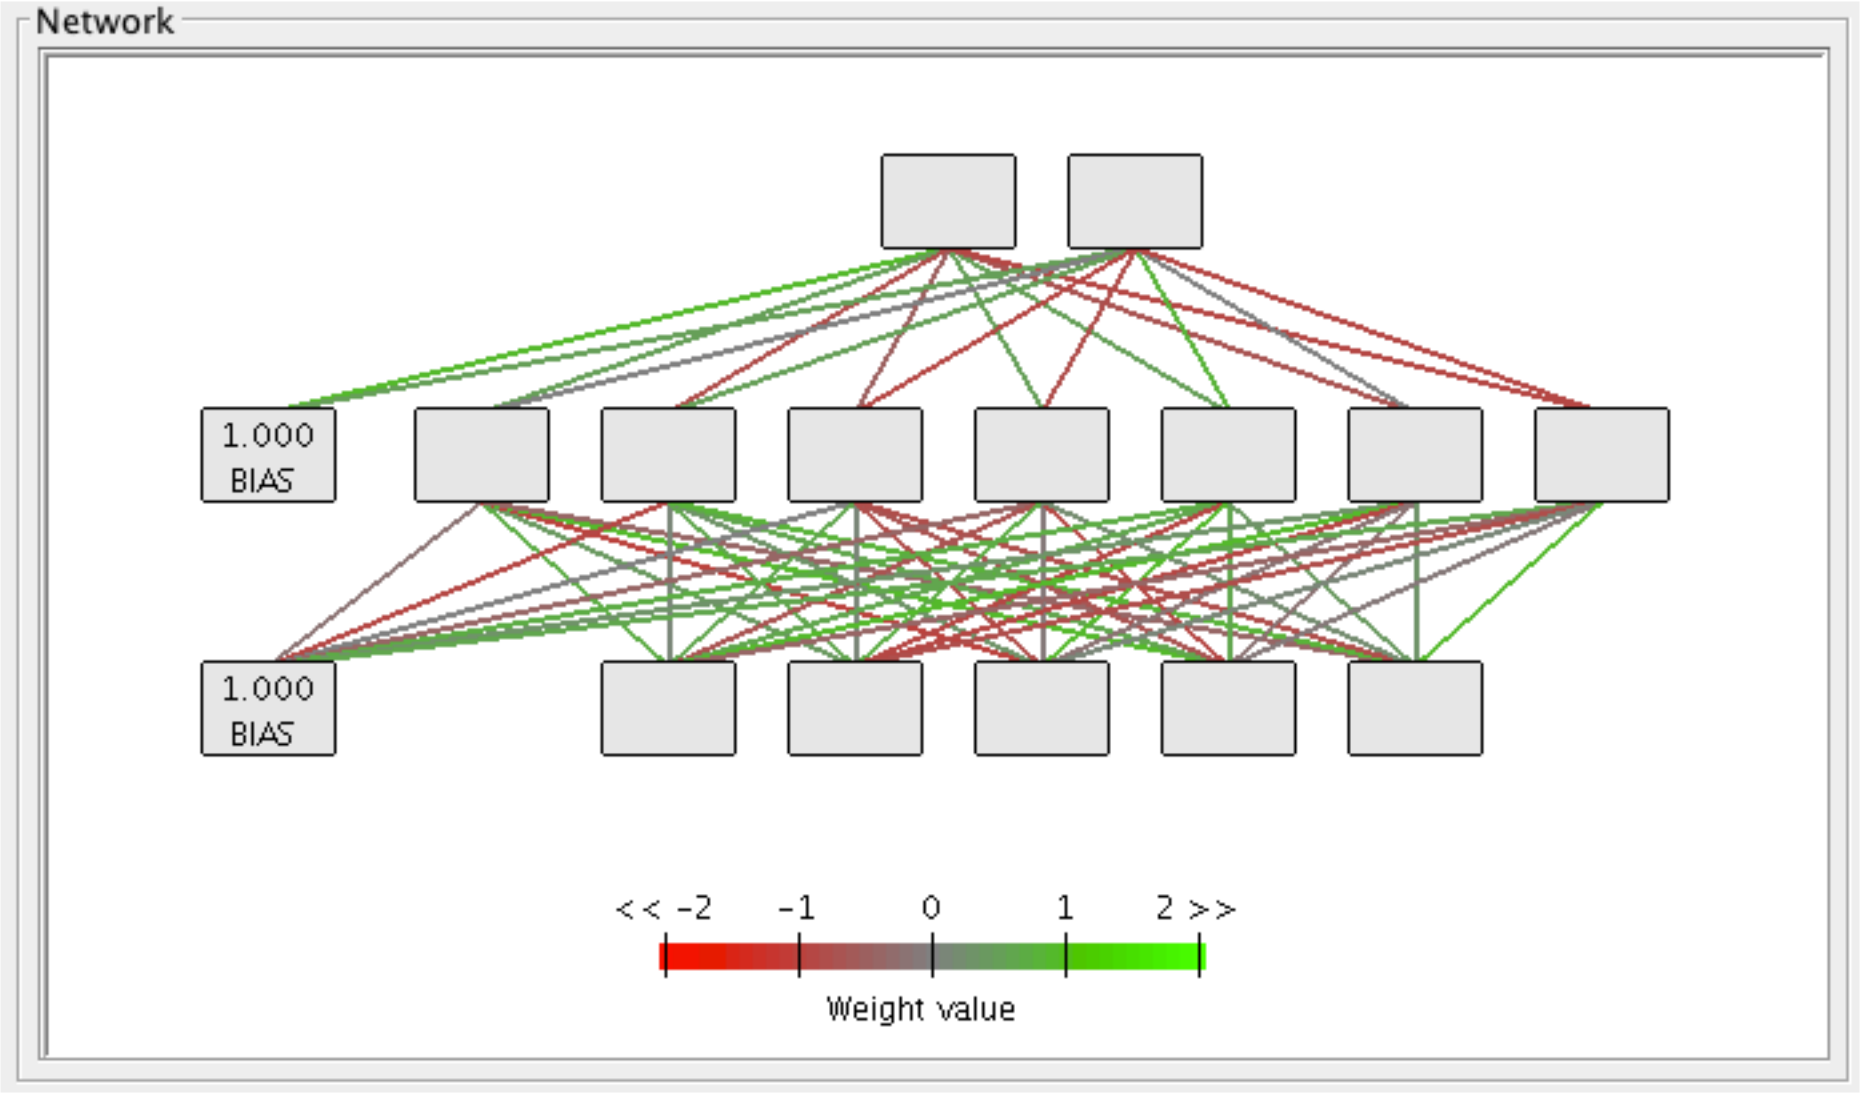

I connected 1 adruino mega, 3 servos, 1 analog infrared receiver, 3 contacts and 2 encoders from an old printer. I use a genetic algorithm that trains the 5-7-2 neural network (5 inputs, 7 neurons in the hidden layer and 2 in output). The 5 inputs are the 5 areas that the IR scans using the third servo, and the 2 outputs are the speed of each wheel servo.

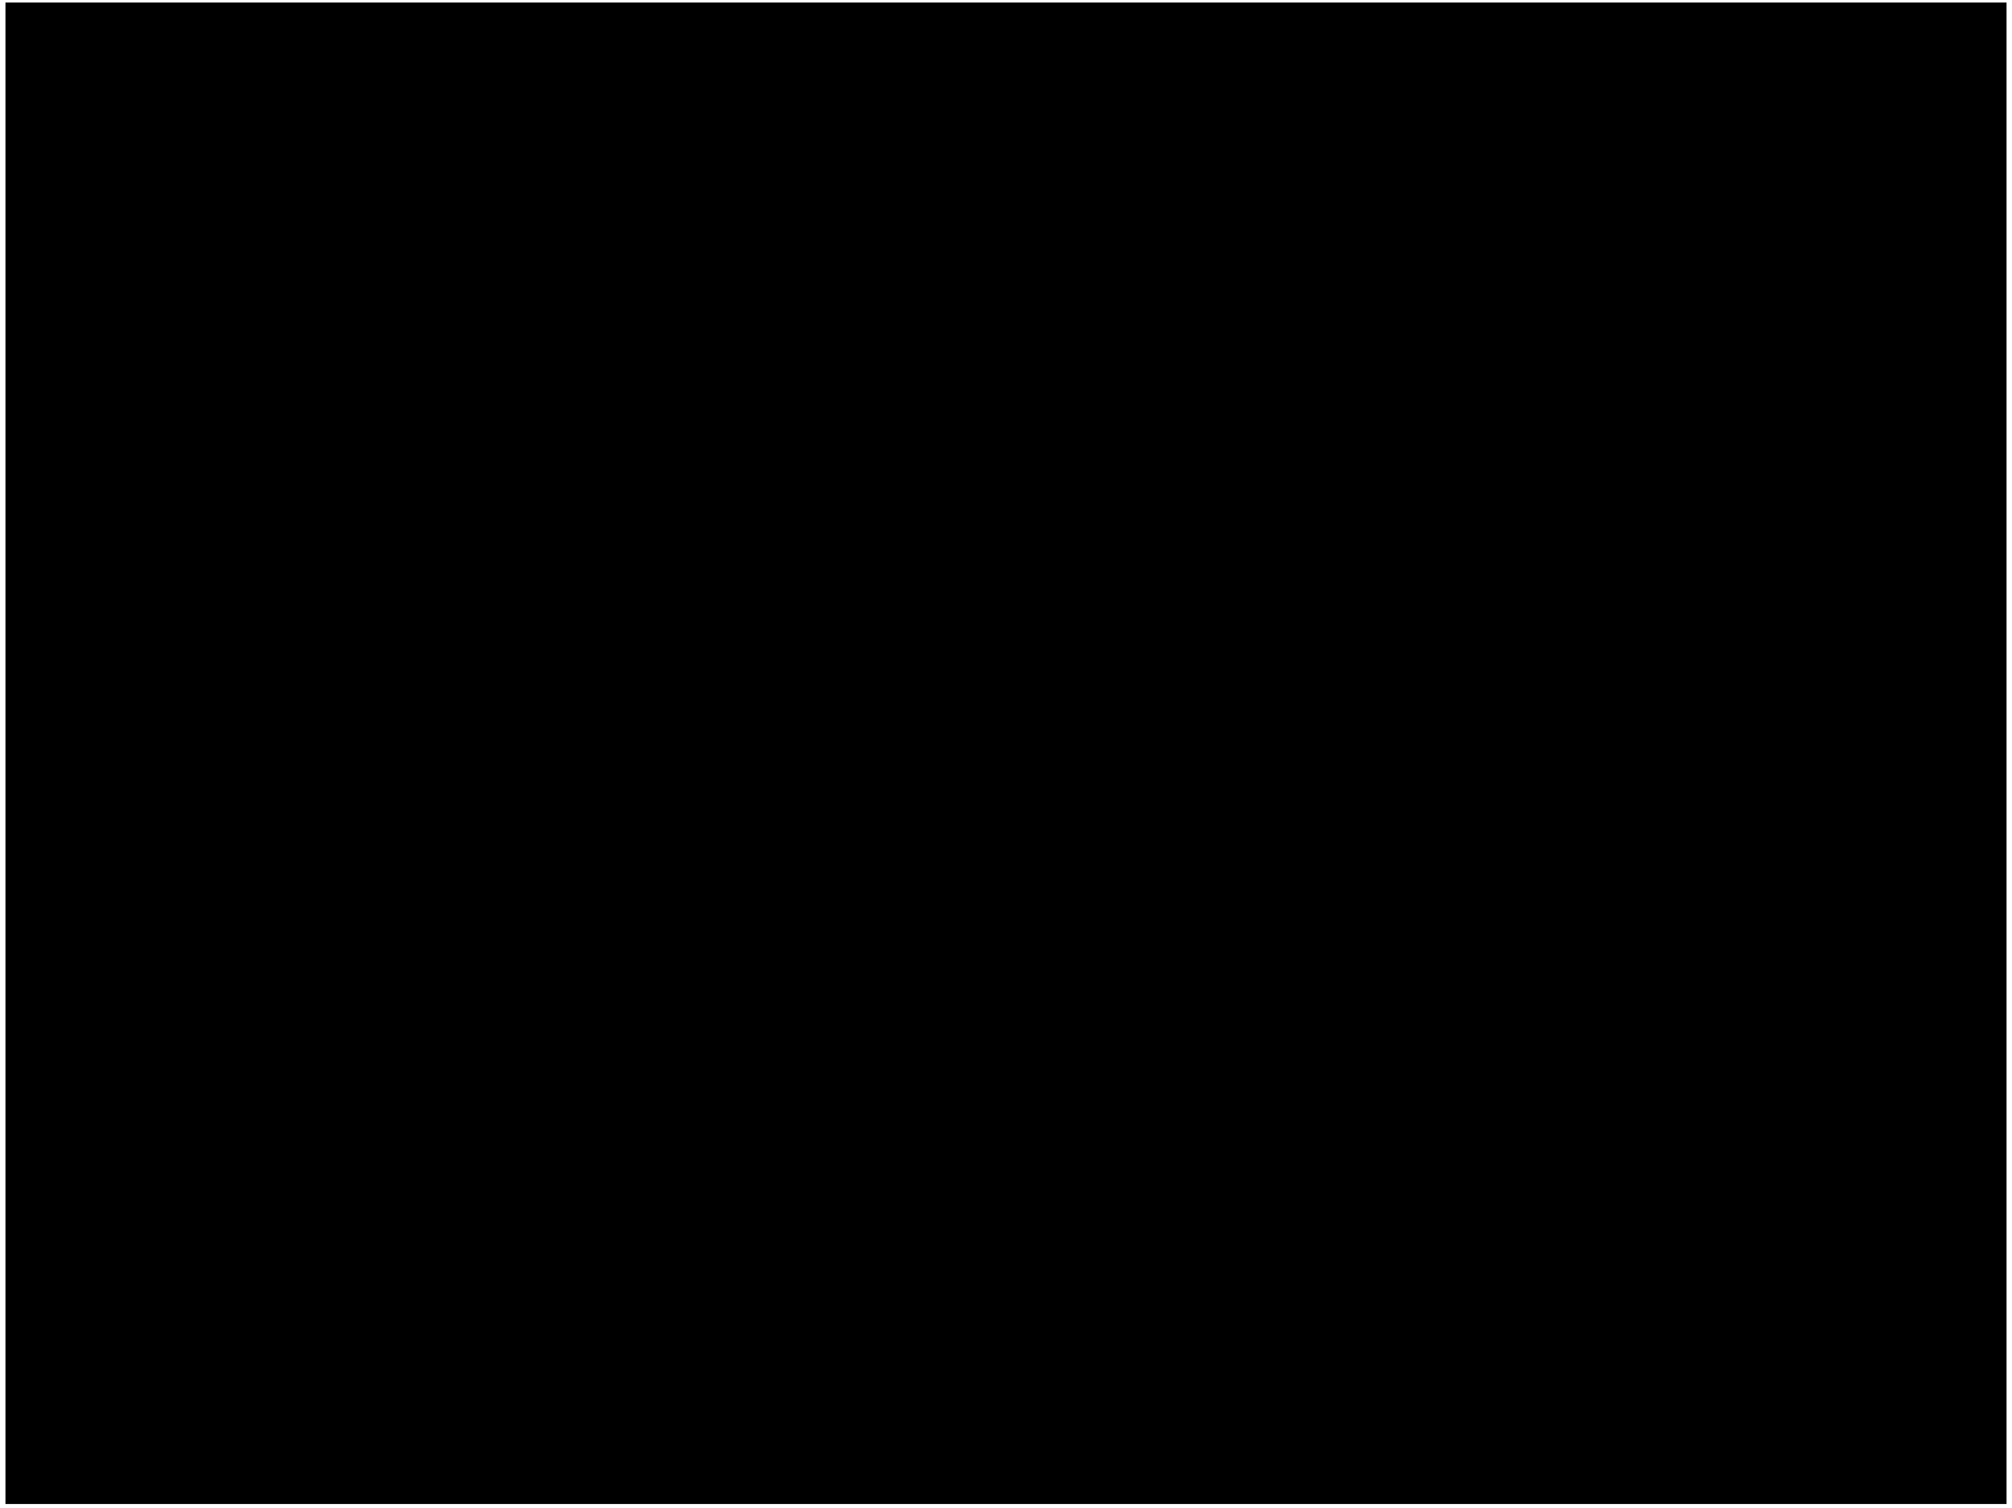

Observe start and at 6 minutes later.

Reason 4: Connectionist models are continuous with dynamical models

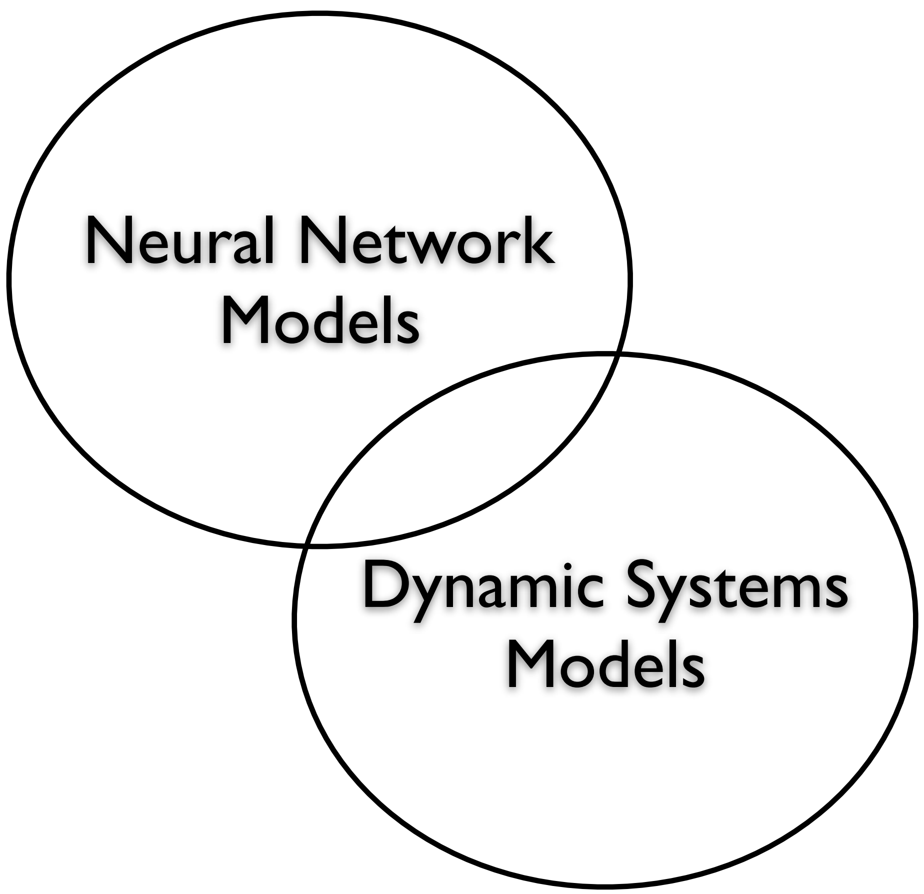

...and others

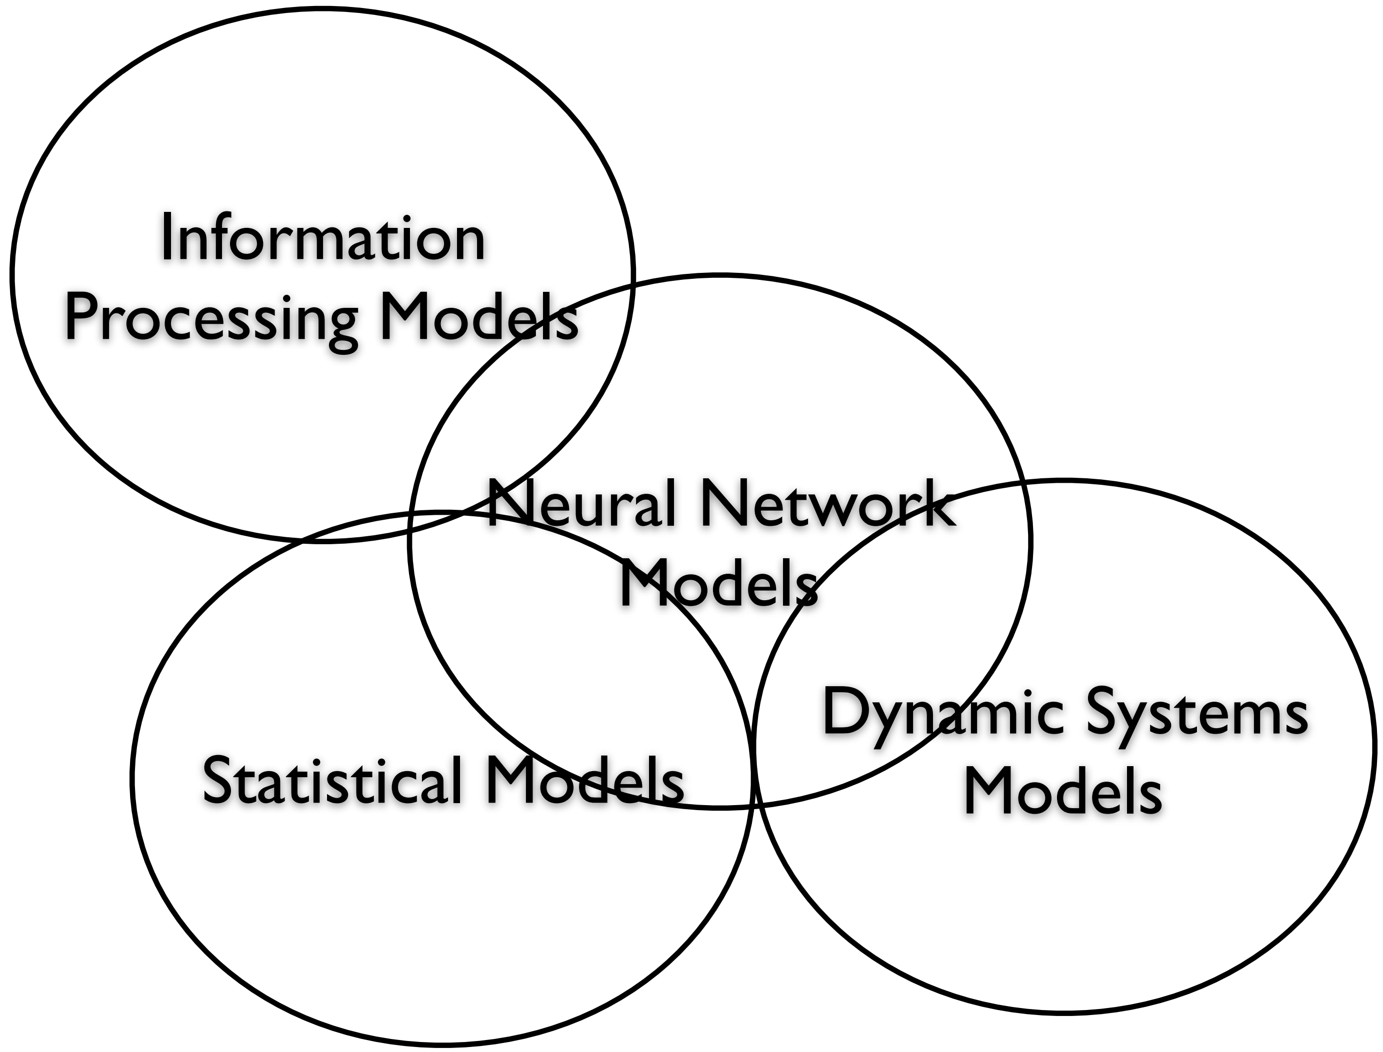
Information
Processing Models
Neural Network
Models
Statistical Models
Dynamic Systems
Models

We will distinguish between Feed Forward Networks and Recurrent Neural Networks

Output

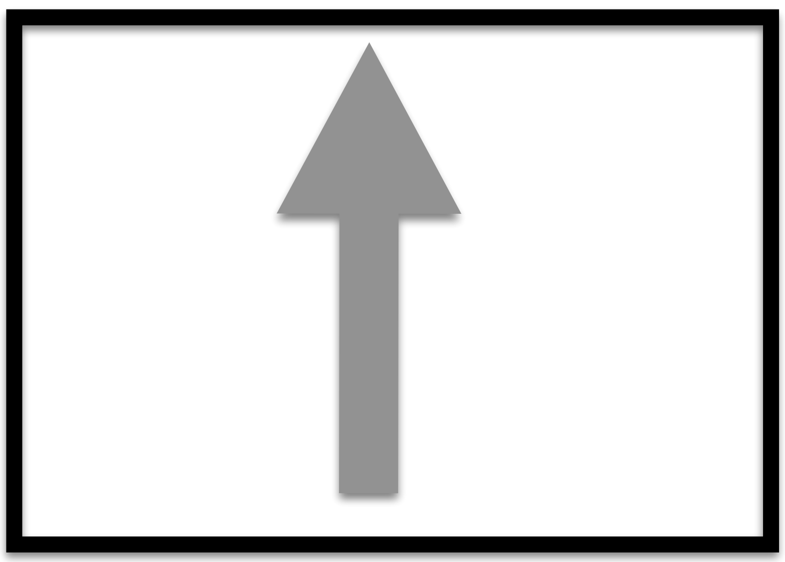

Input

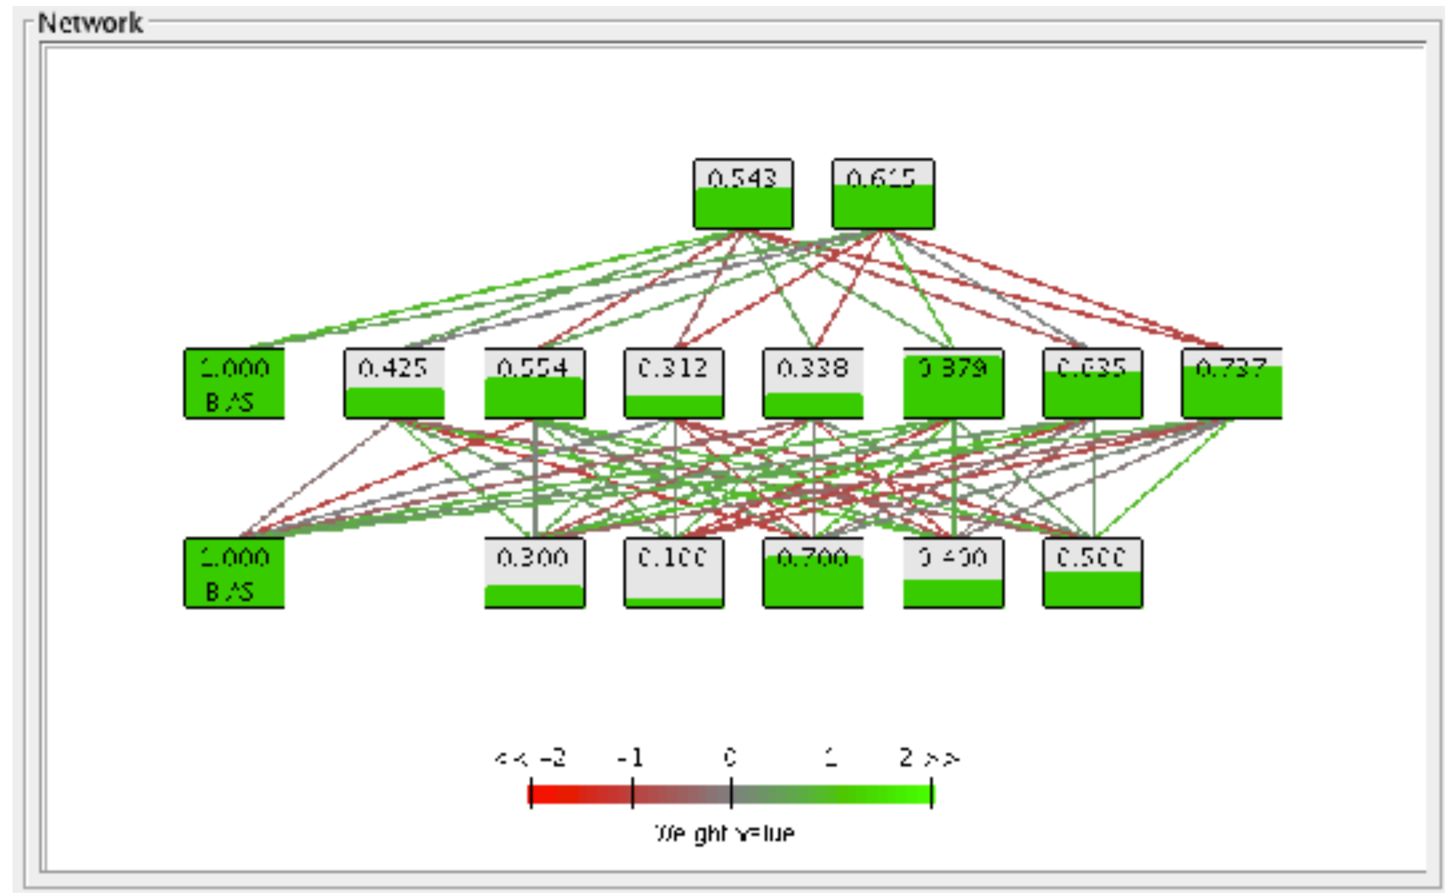

FFNs generate one input from one output

In an RNN, links form cycles, so that activity at any moment in time is a result, not only of current input, but also of past values.

This requires a different mathematical formalism: Dynamic Systems Theory.

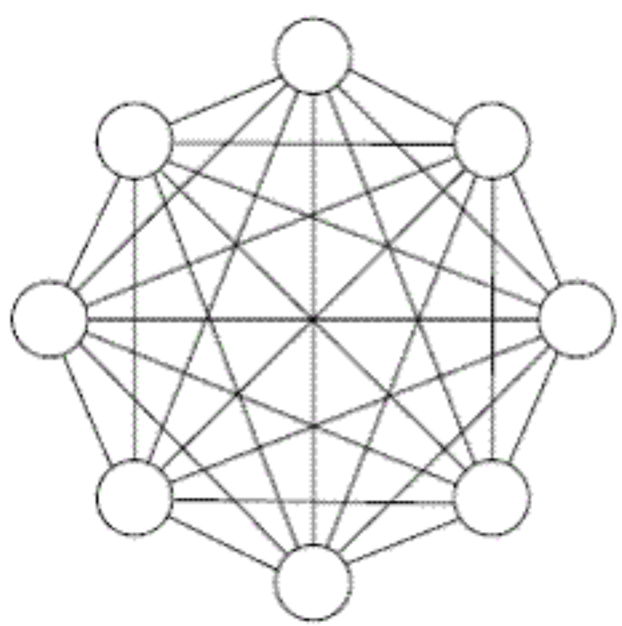

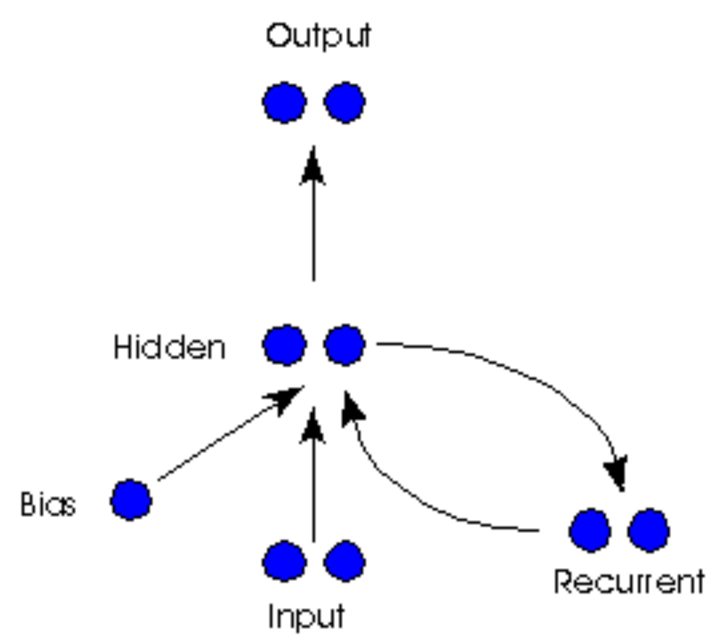

3 large sets of issues that will shape our approach to *any* neural network:

1. What do the numbers mean (to us) - this is a question of ***representation***

2. What is the ***architecture*** of the network

3. What is the approach in the network to ***time***

Representation Issues

What do the numbers in the network mean *to us?*

How do inputs relate to the real world?

How do outputs relate to the system being modelled?

How do we view the changing numbers within the system?

Architecture Issues

How many units (input, hidden, output)?

How are they connected?

How are the parameters that determine changes in these numbers determined?

What kind of units are used?

Which parameters affect individual units, and which affect the whole network?

Issues related to time

Is time-in-the-world captured in the input or output representations?

How are inputs sequenced during training?

Does the model have a meaningful relation to it's physical environment? (e.g. the robot)

ANNs and RNNs treat time in completely different ways

"Nice" properties of (some) connectionist models....

- Relatable to neural structures and processes
- Parallel computation
- Distributed representation
- Learning based on experience
- Plausible learning mechanisms (some)
- Graceful degradation
- Implementation of highly complex functions, e.g mapping from image input to text output.

Neural Inspiration: Care Required!

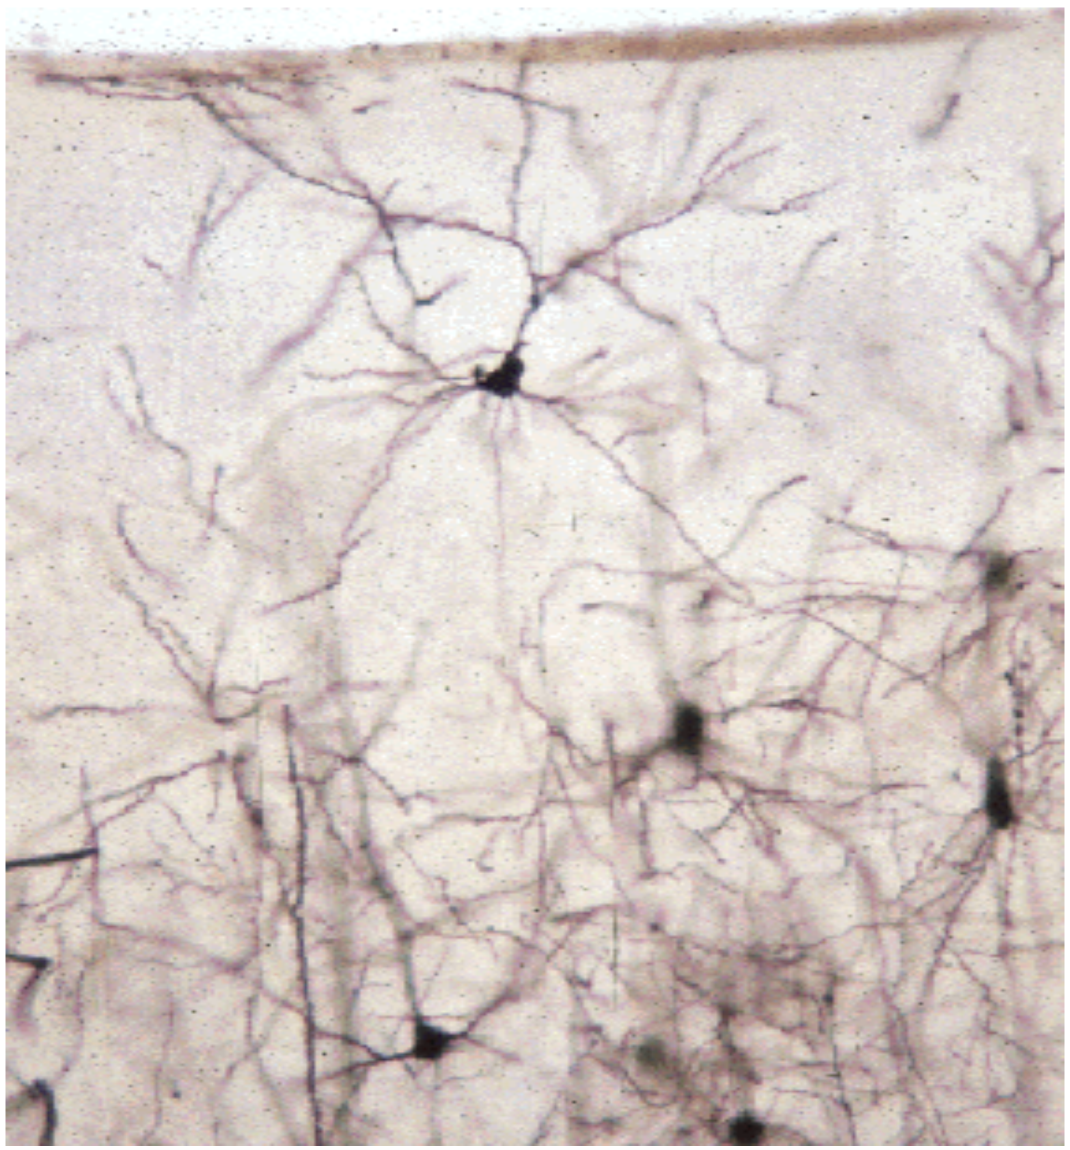

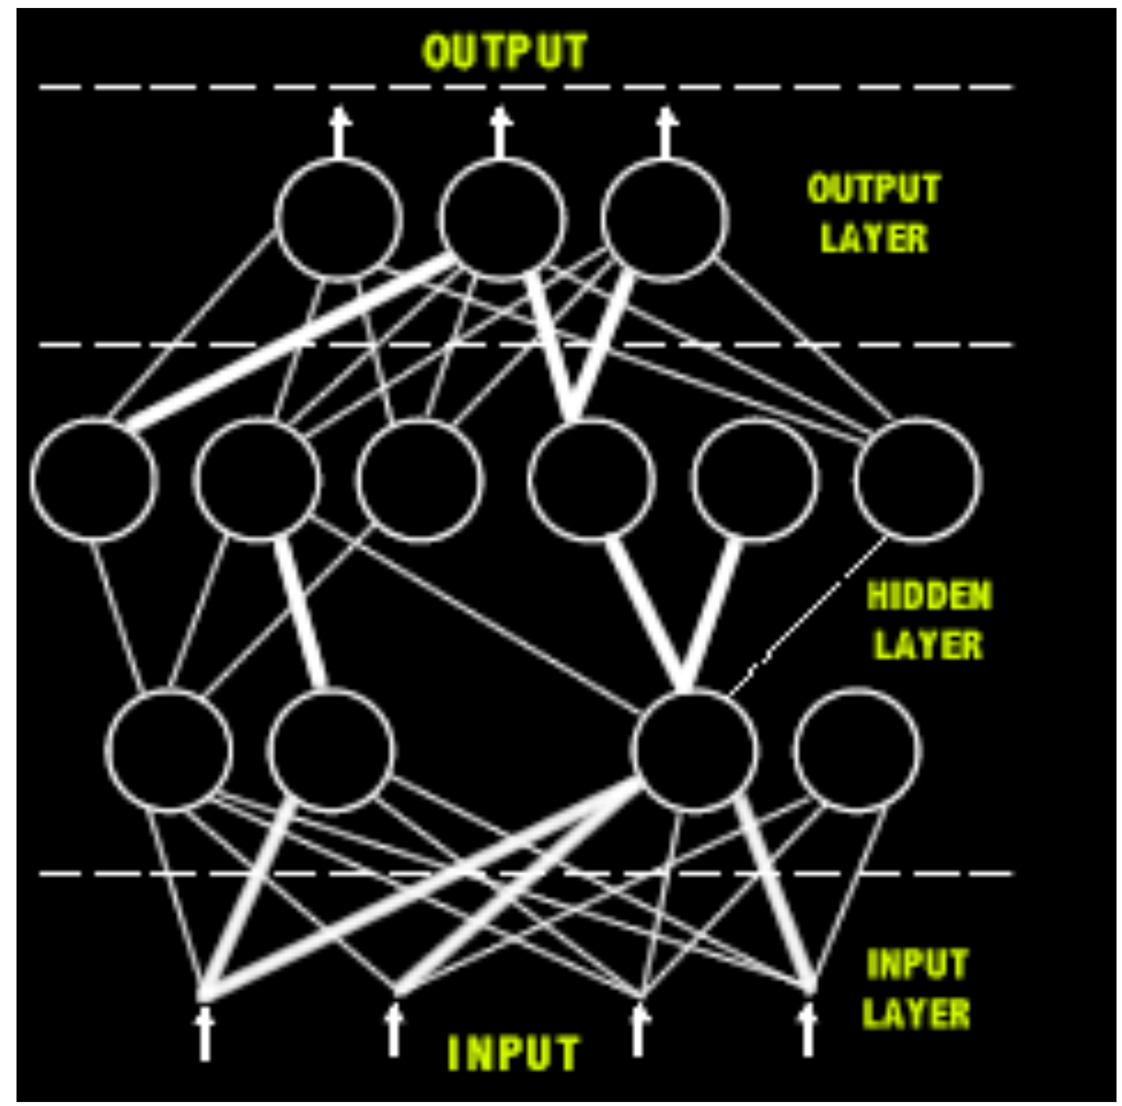

(Real) Neural Networks in the brain

- The brain is not homogeneous
- Cortex, Midbrain, Cerebellum, Brainstem
- . . . regions . . . areas
- Transient local neural assemblies

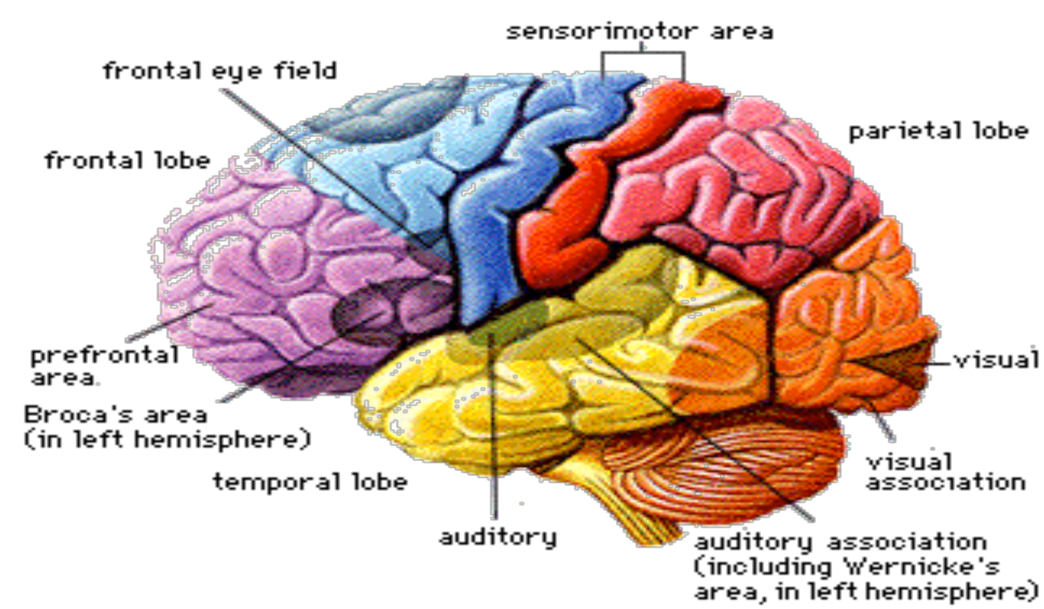

Connections

- Complex pattern of projections within and between areas
- Feedforward (sensory to central?)
- Feedback (recurrence)

Diffusion Tensor Imaging/ Diffusion Spectrum Imaging

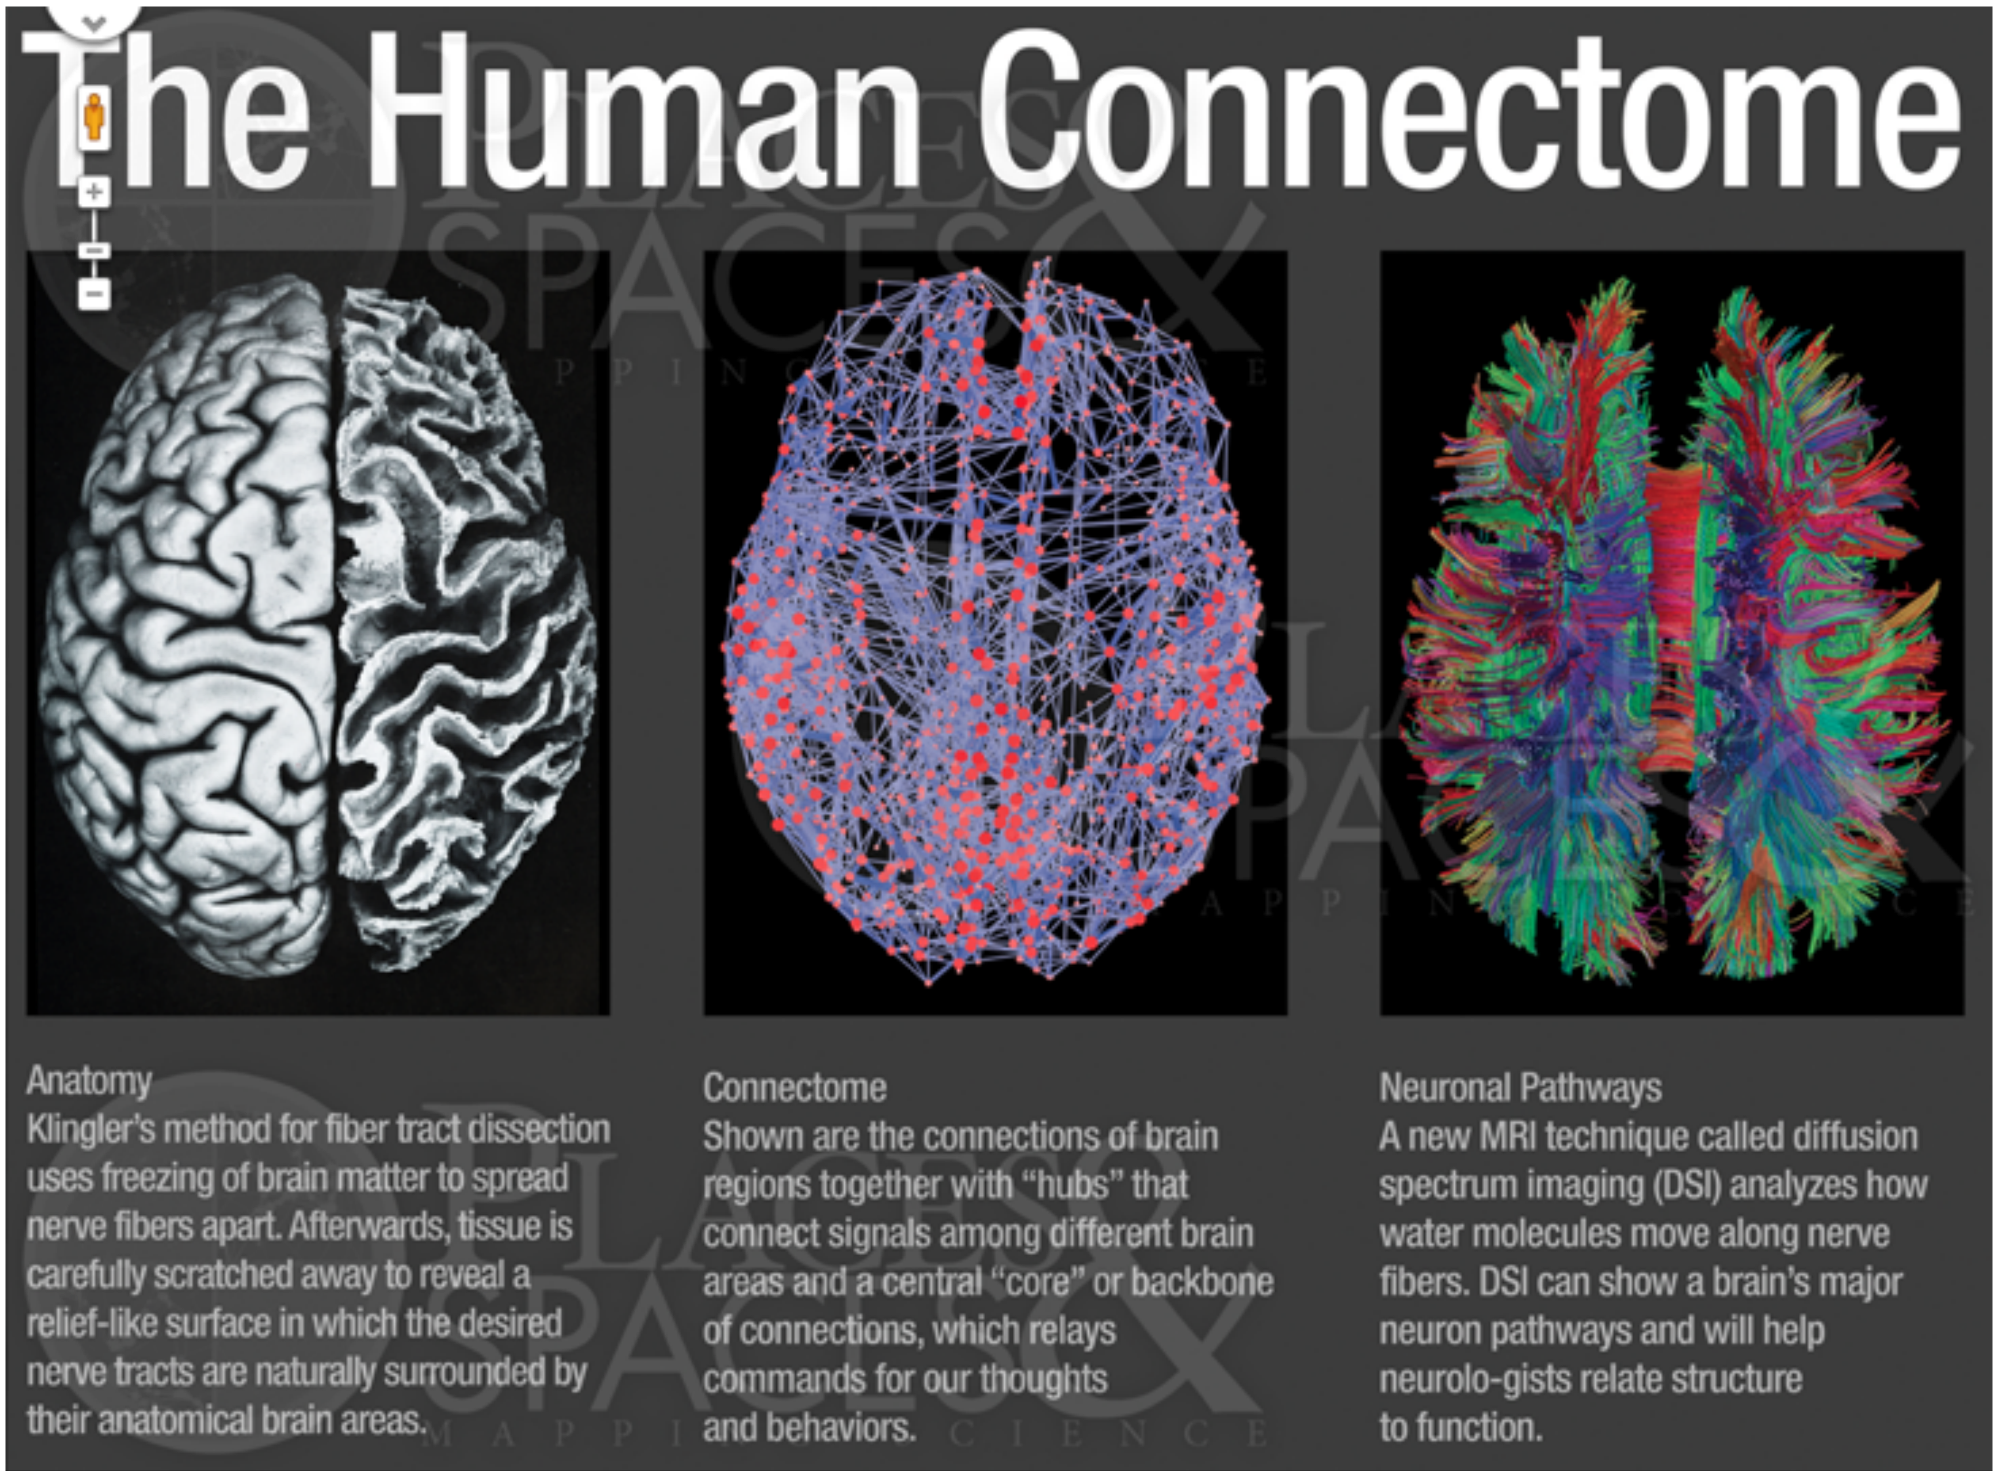

The nervous system, seen through the lens of the neuron doctrine

Neurons

- Basic computational unit is the Neuron
  - Dendrites (inputs)
  - Soma (cell body)
  - Axon (output)

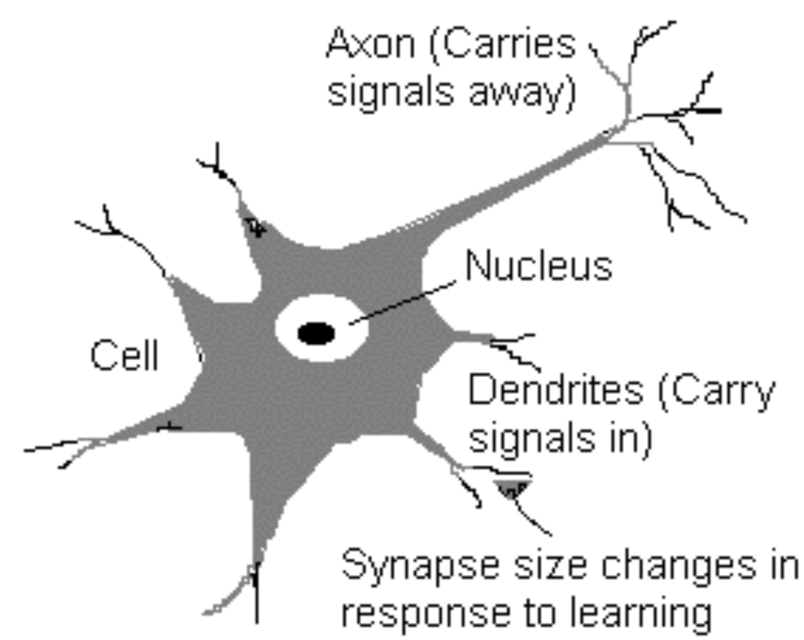

Modus operandi of a neuron

- Input (through dendrites...) from (thousands of) other neurons
- Inputs sum until a critical level (threshold) is reached
- At threshold, a spike is generated and propagates through the axon to downstream neurons (depolarization)
- Neuron enters an unresponsive refractory period
- Typical firing rates: 100 Hz (compare computer: 1,000,000,000 Hz)

Synapse

- Axons almost, but not quite, touch their target (dendrites etc)
- Neurotransmitters effect the transmission from one cell to the next through a narrow space
- This link is the synapse
- The synapse is the seat of
- learning (LTP)
- Learning uses locally available information

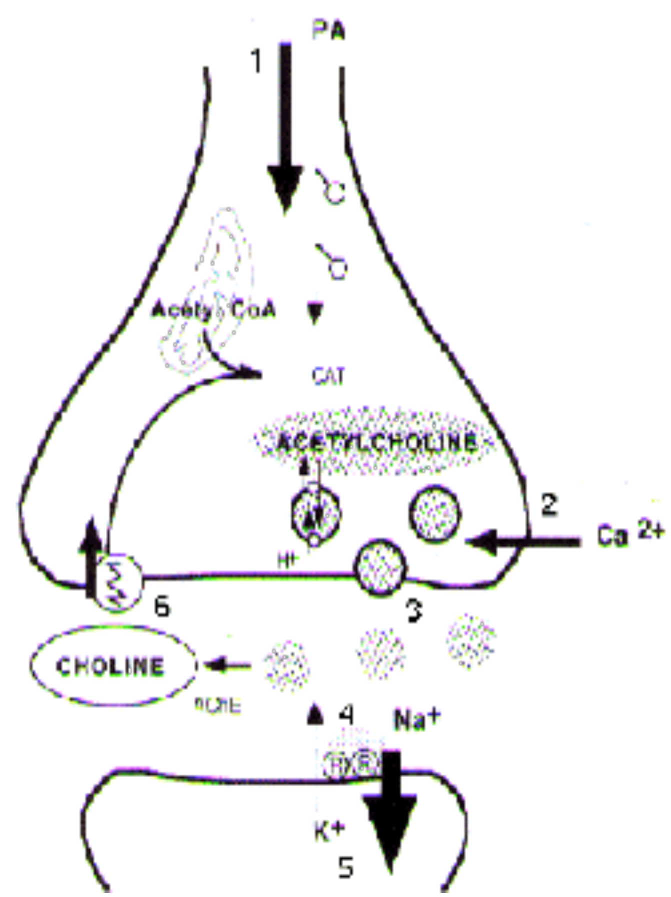

Neurons are not the only players! The tripartite synapse shows that neurons and glia cells interact, forming local dynamical systems with feedback loops.

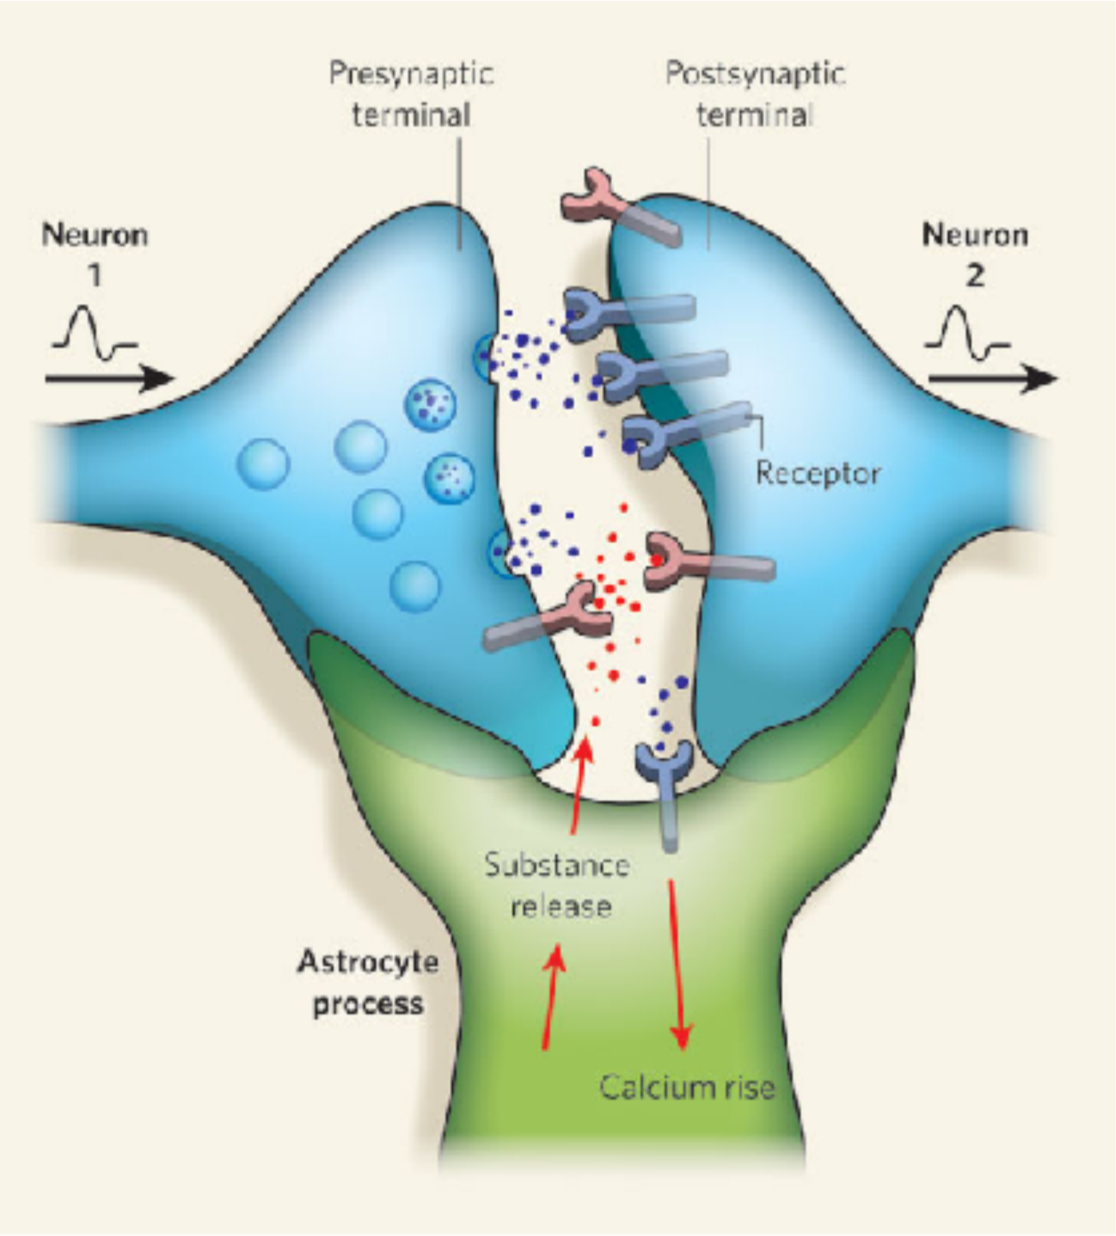

Long-term potentiation

- One way brains learn is by modification of synapses as a result of experience
- One mechanism is LTP:
  - An enduring (>1 hour) increase in synaptic efficiency that results from high frequency stimulation of an afferent (input) pathway
- Hebb's postulate (1949):
  - When an axon of cell A...excites cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells so that A's efficiency as one of the cells firing B is increased
- Bliss and Lomo discovered LTP in hippocampus in 1973

LTP continued

- Synapses become more or less important over time (plasticity)
- LTP is based on experience
- LTP is based only on local information (Hebb's postulate)

The McCulloch-Pitts Model Neuron

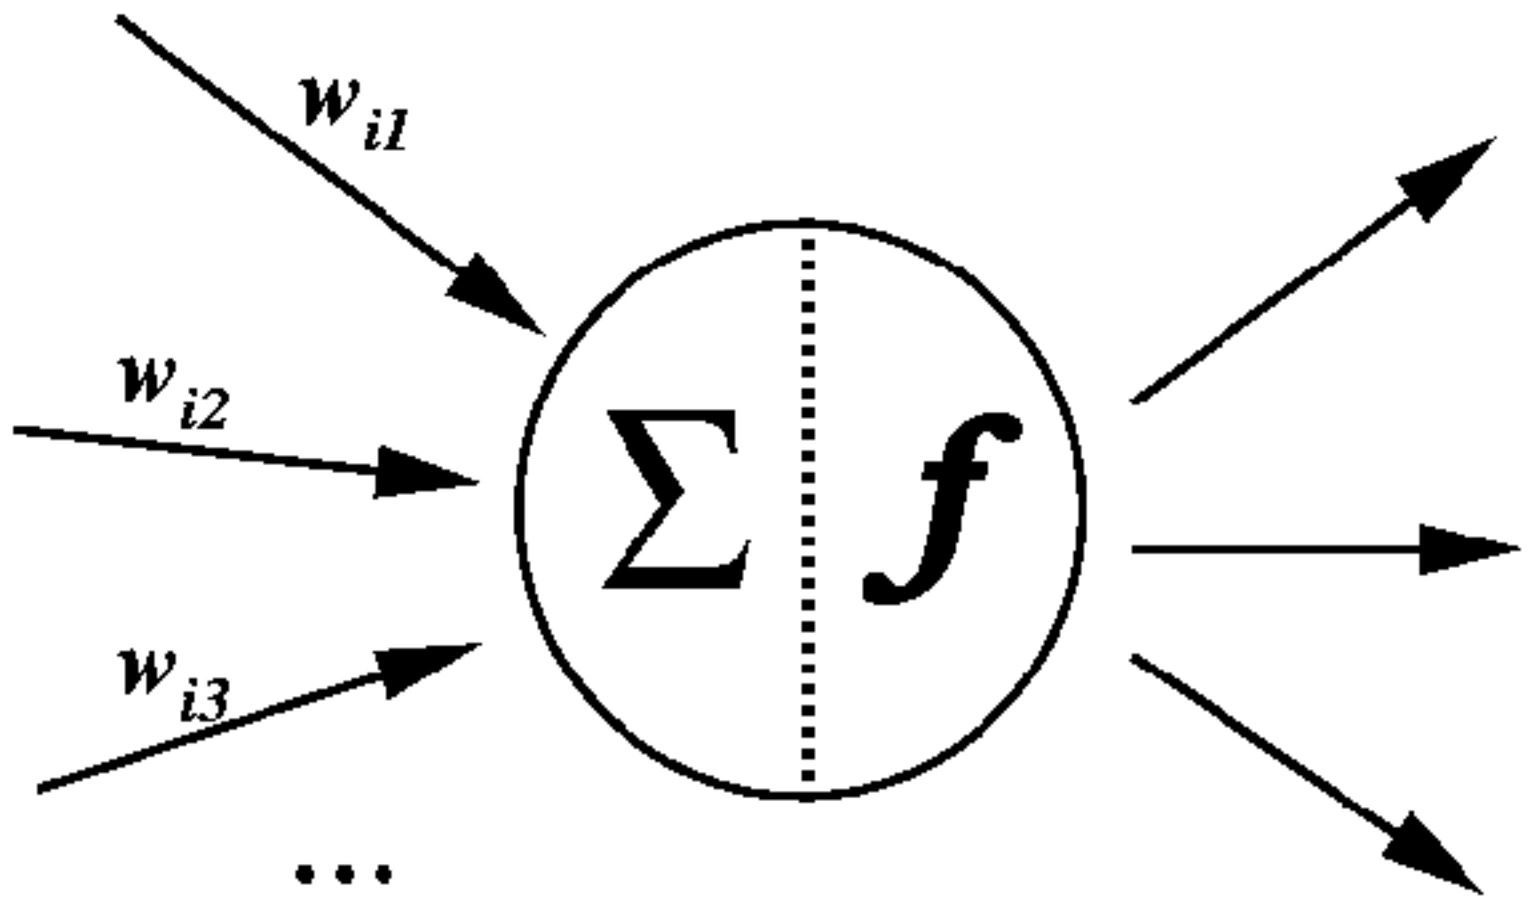

$$y_i = f(net_i)$$

Simple Artificial Neuron

- Neuron (biology) -> unit or node (model)

net input is the weighted sum of all contributing values

$$\mathrm{net}_i = \sum_j w_{ij} y_j$$

Unit activation is some function of net_i.

If *f* is the identity function, the unit is a 'linear' unit, and its activation (output) is the same thing as its net input.

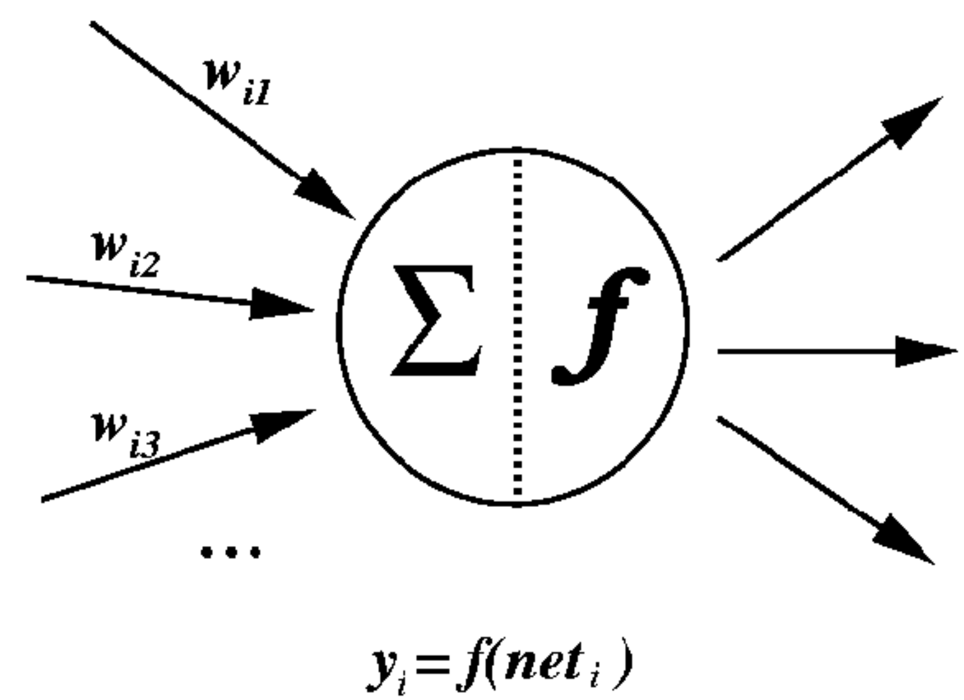

Inputs and Outputs

Individual units in a model network receive input, and compute their corresponding output. This output is called the *unit activation.*

This is different from the input to the network as a whole, which leads to the network as a whole producing output.

The output of the network as a whole is simply the activations of the output units

Activation function (sigmoid)

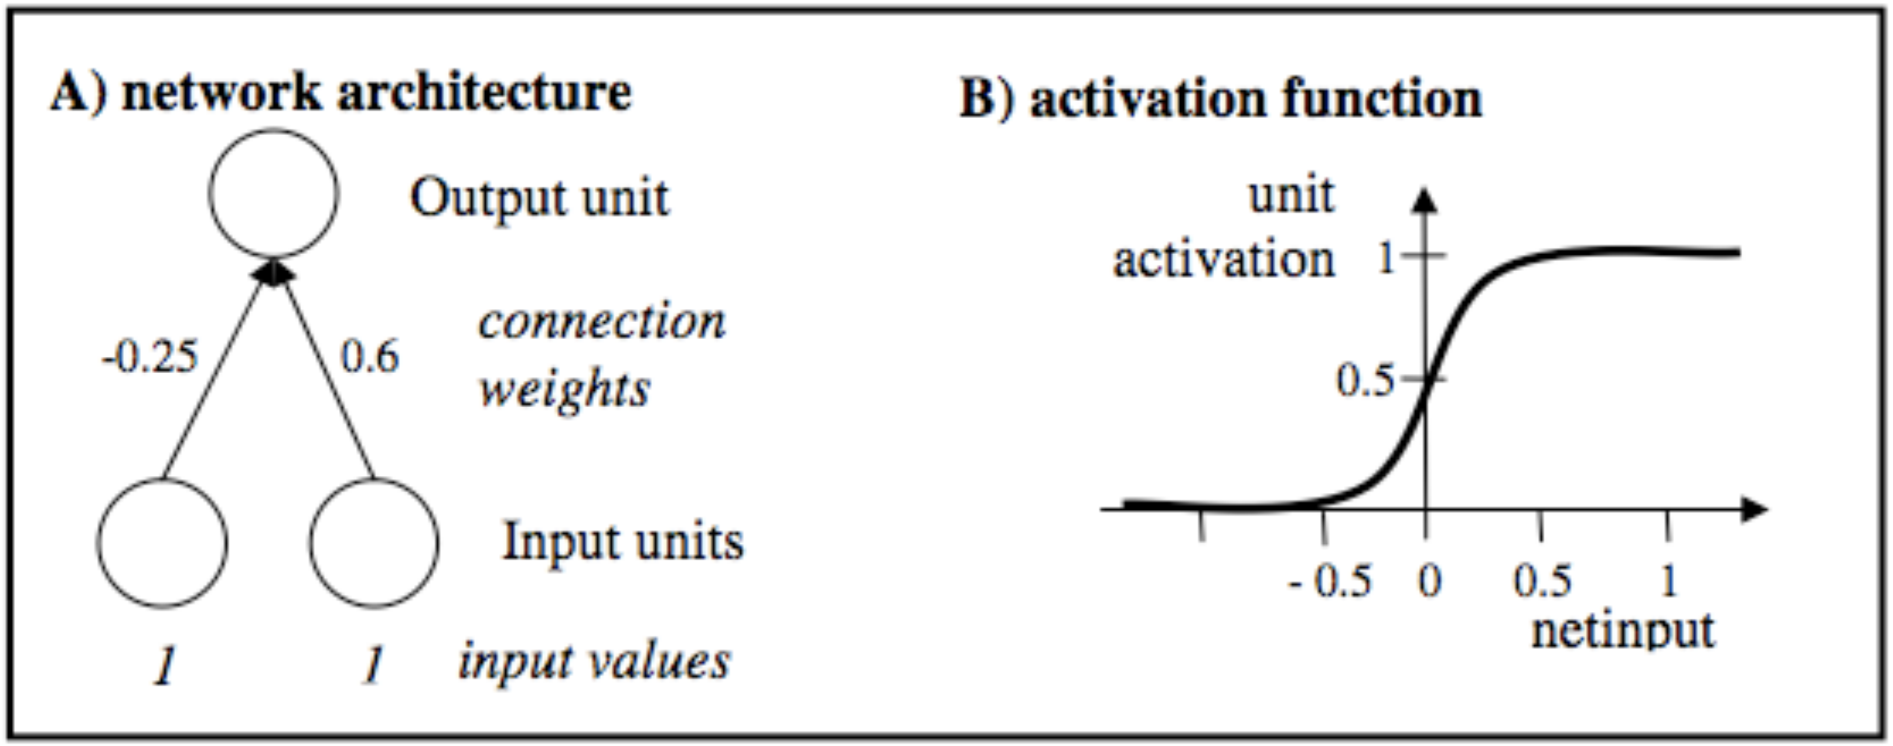

$y=1/(1+e^{-x})$

The range of possible unit activations is from 0 to 1.

An alternative function, too infrequently used, is the *tanh* function, with a range of -1 to 1

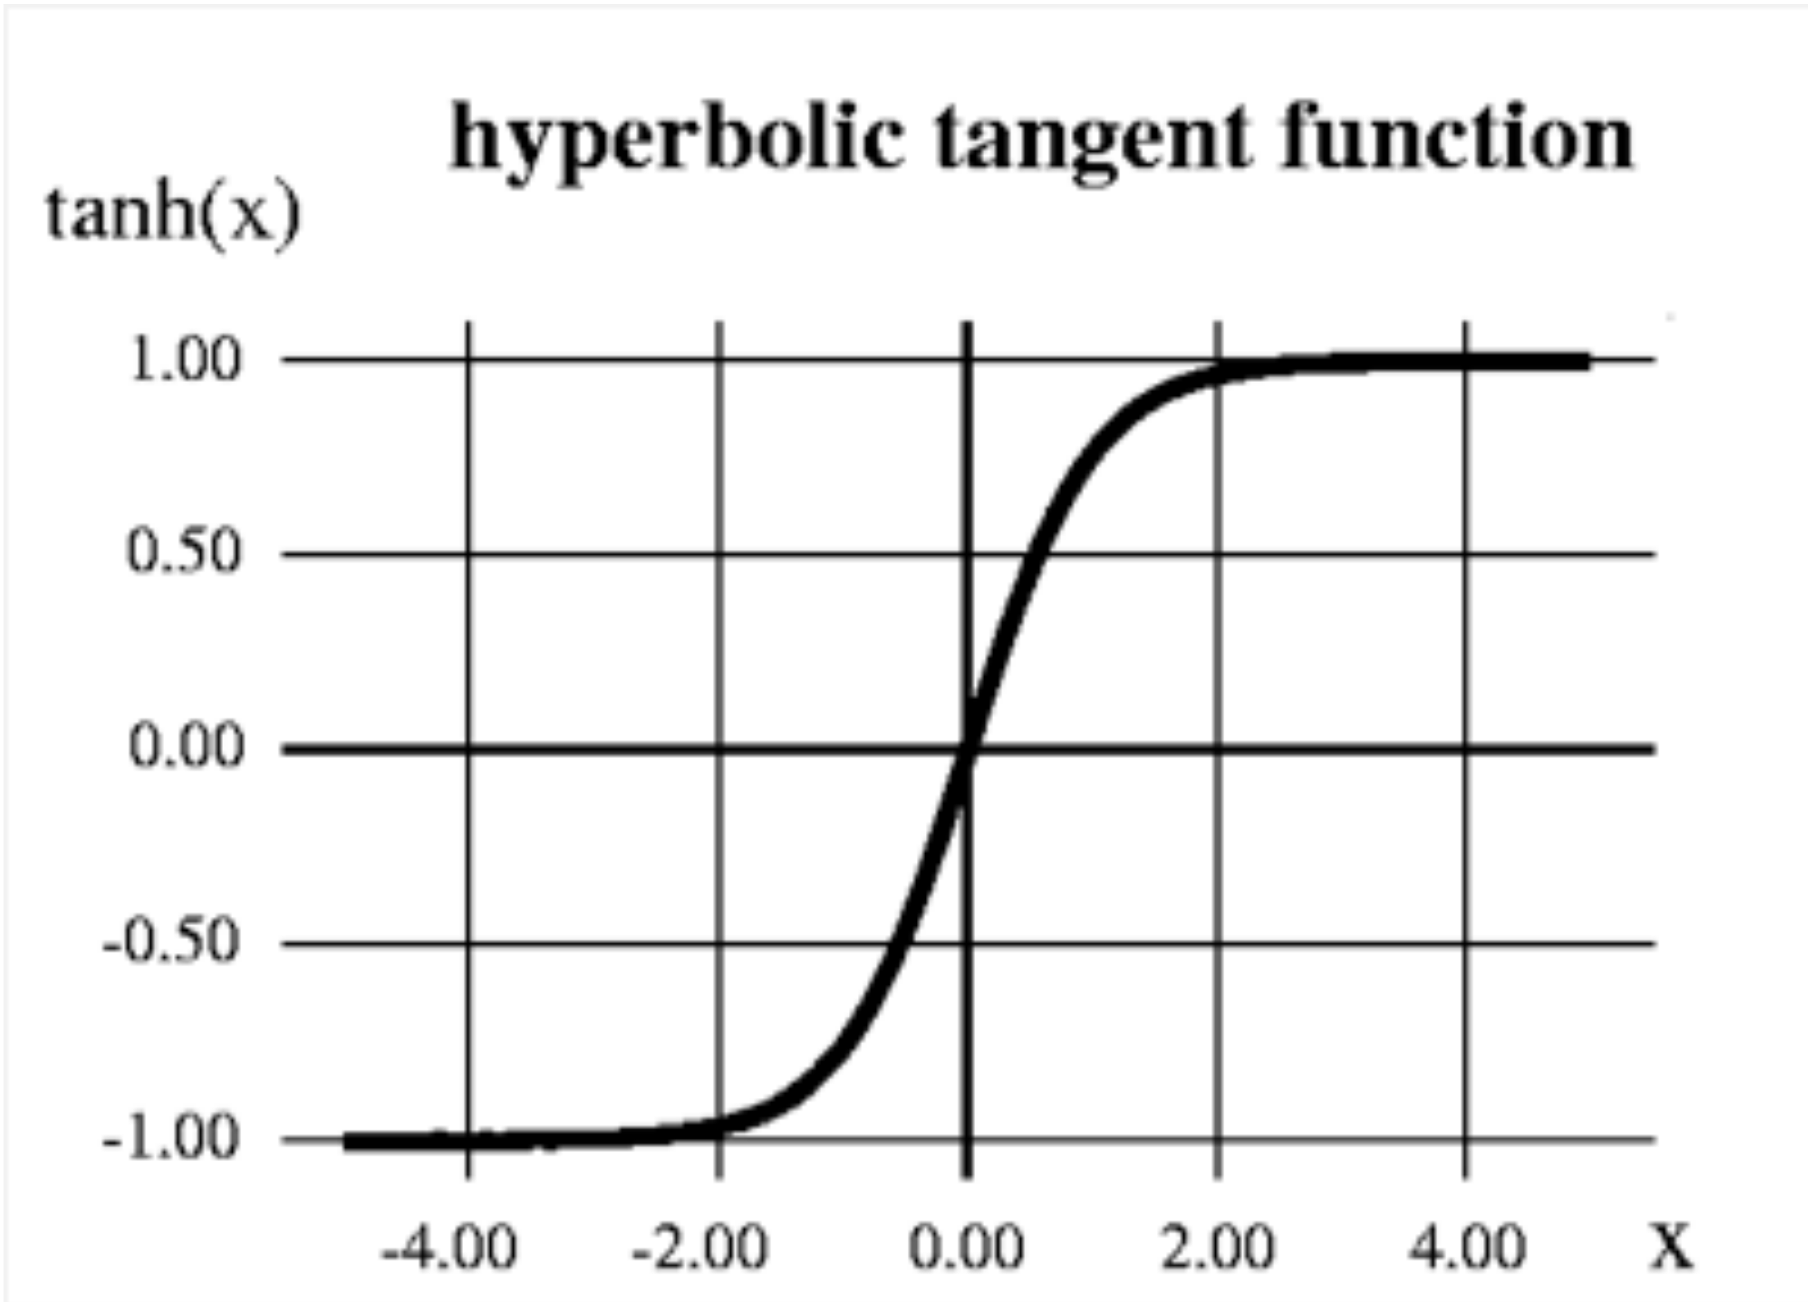

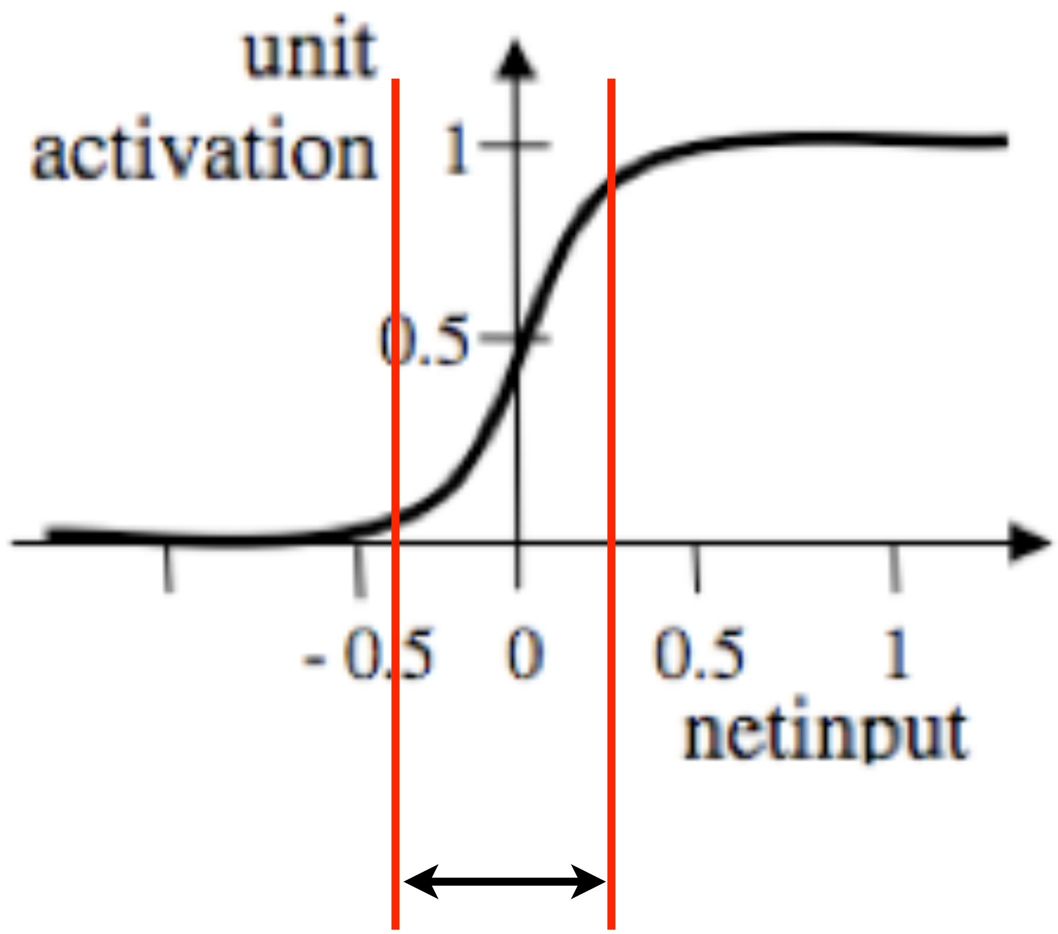

Unit responds sensitively to different inputs in this range only

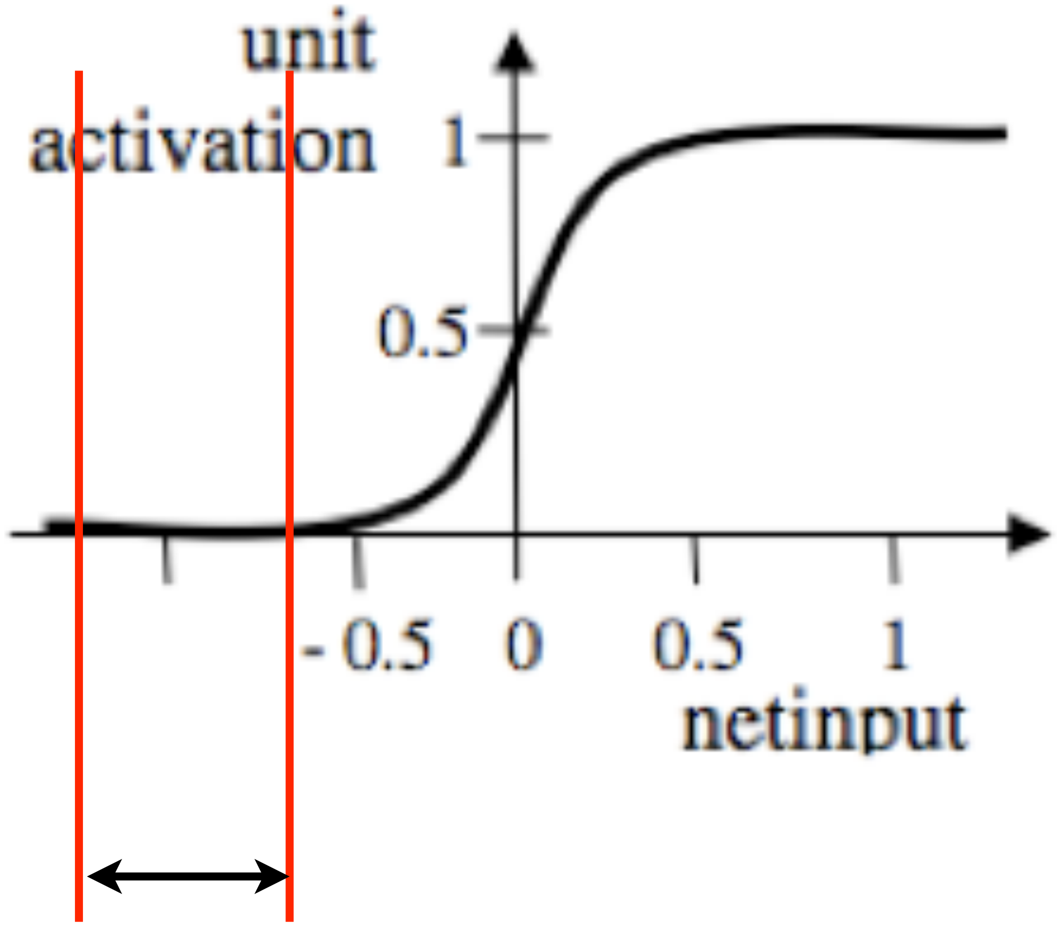

If input range lies elsewhere, we need a (trainable) means to move this range of differential sensitivity

Add a “bias” input: a constant value added to the net input.

$$\text{net}_i = \sum_j w_{ij} y_j + bias$$

OR

Add an extra unit, with constant activation of 1.0

This is a bias unit, and is conventionally part of the input to all non-input units.

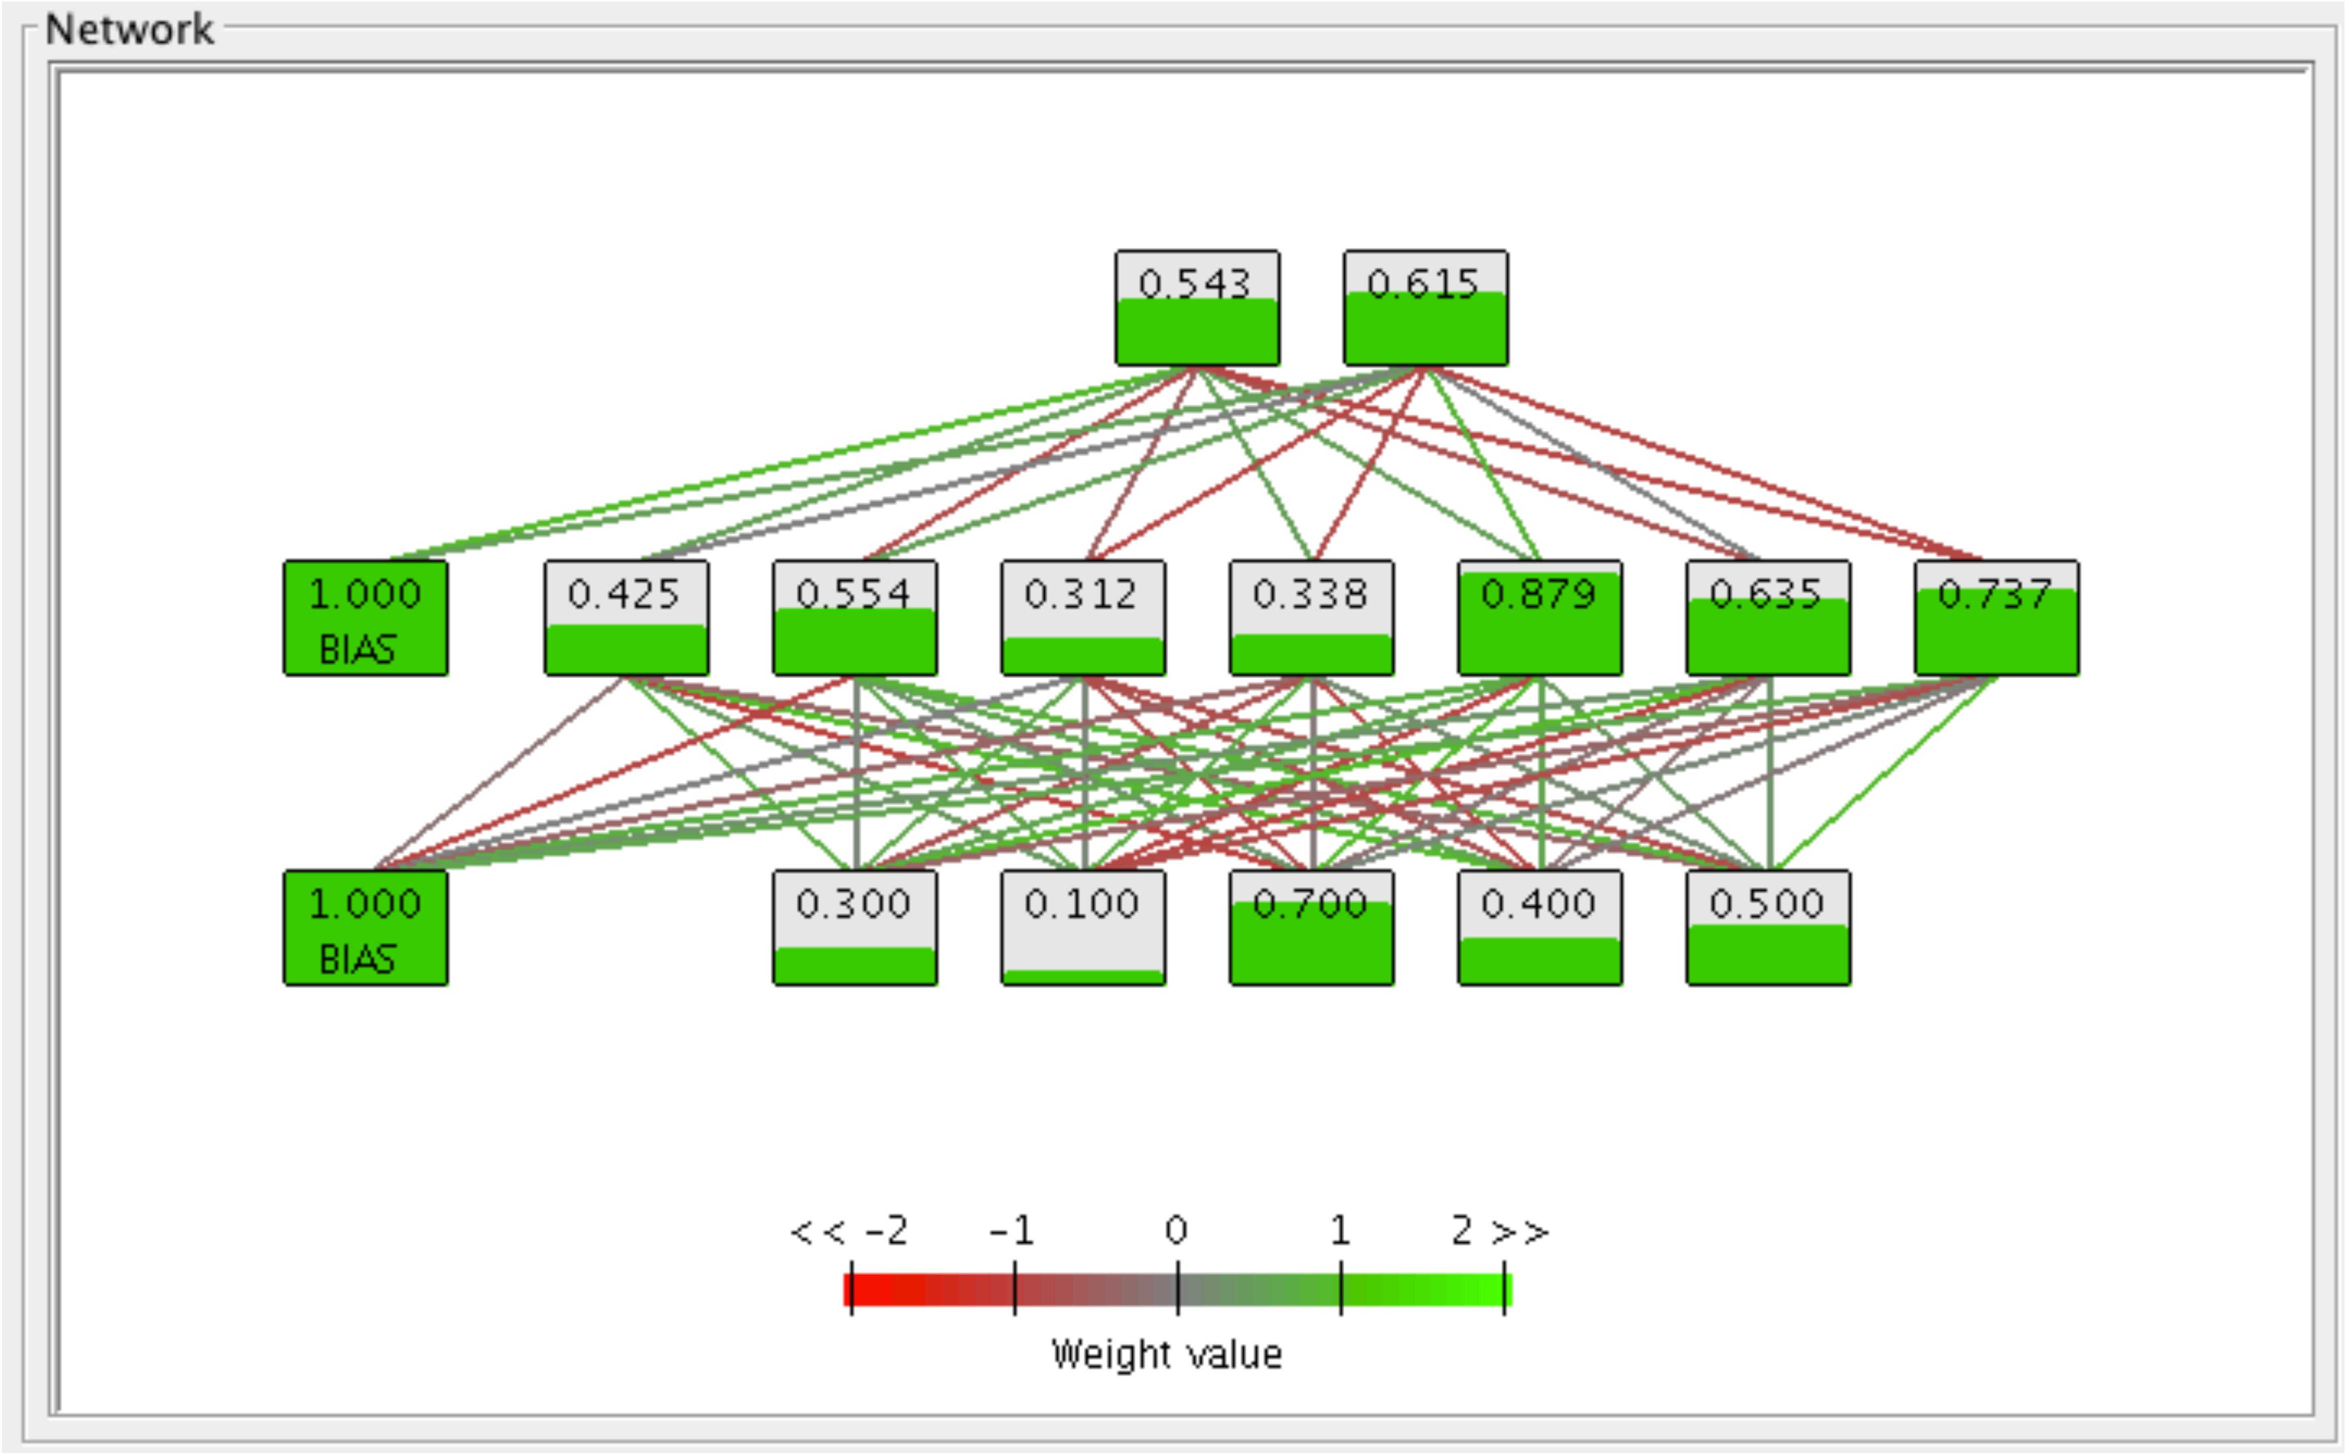
Network
0.543
0.615
1.000
BIAS
0.425
0.554
0.312
0.338
0.879
0.635
0.737
1.000
BIAS
0.300
0.100
0.700
0.400
0.500
<< -2
-1
0
1
2 >>
Weight value

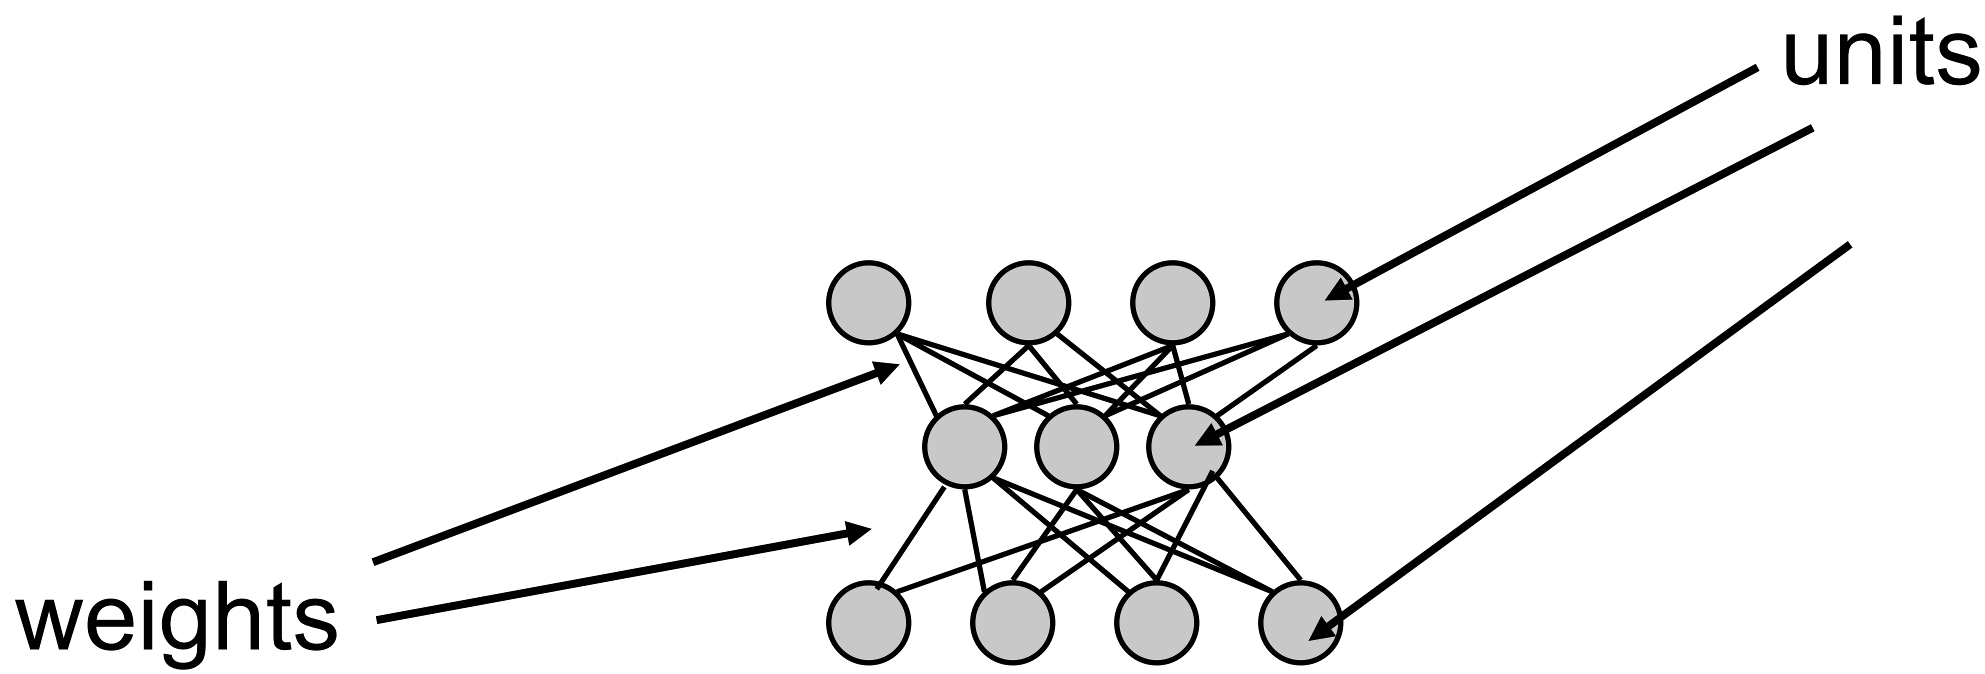
units
weights

Mapping from input to output

input layer

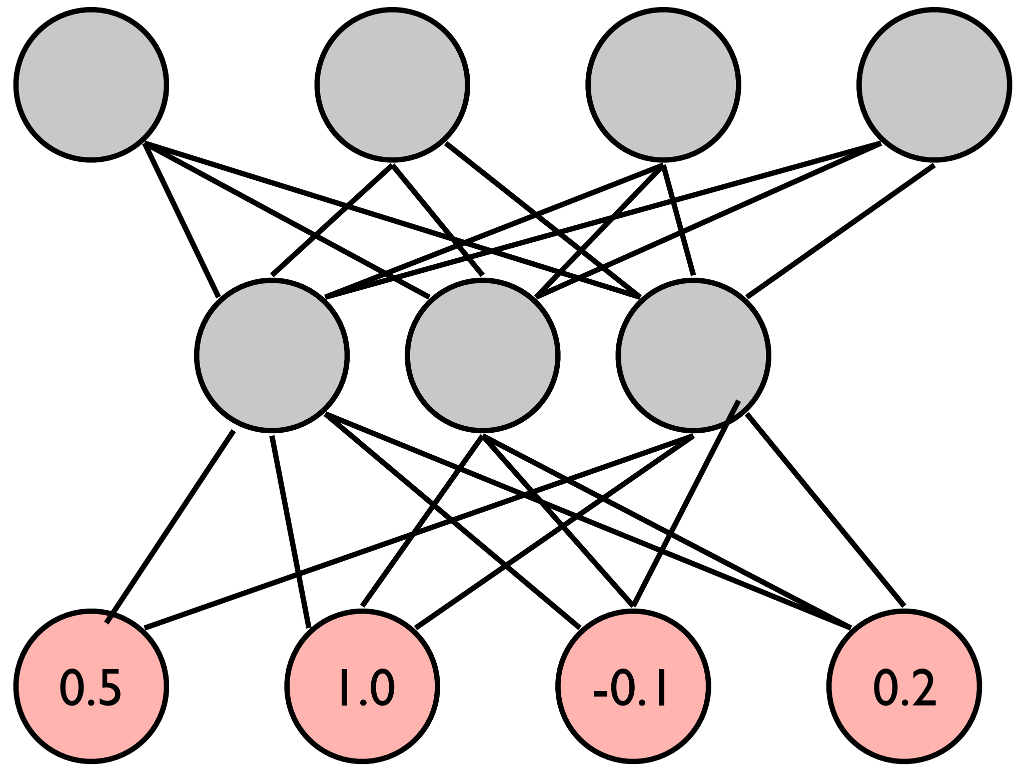

Input pattern: **<**0.5, 1.0,-0.1,0.2**>**

Mapping from input to output

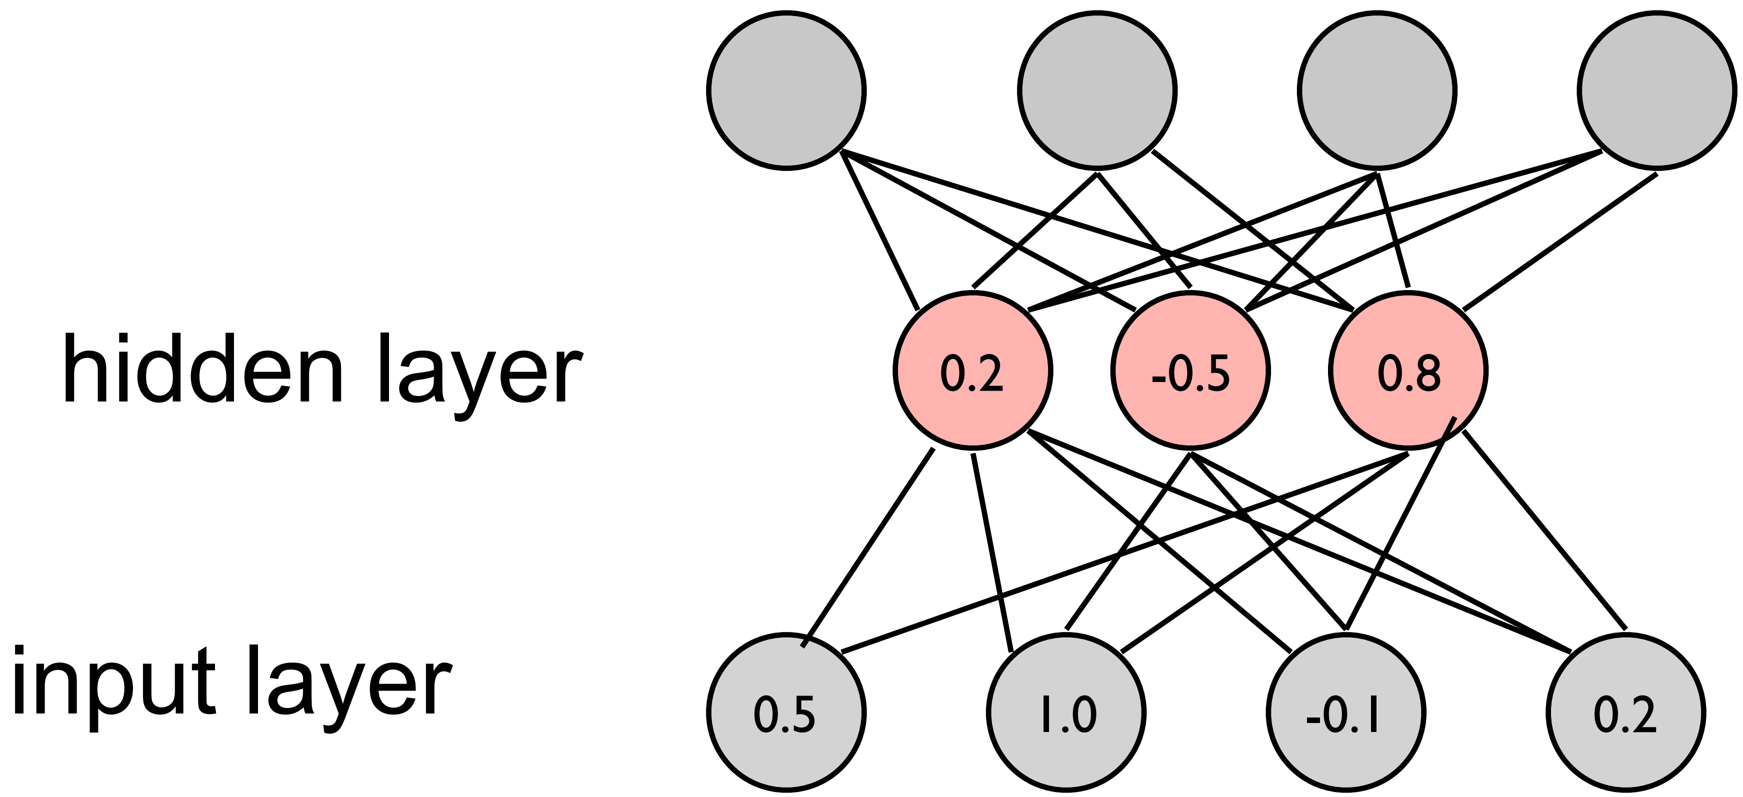

Input pattern: **<**0.5, 1.0,-0.1,0.2**>**

Mapping from input to output

Output pattern: **<**-0.9, 0.2,-0.1,0.7**>**

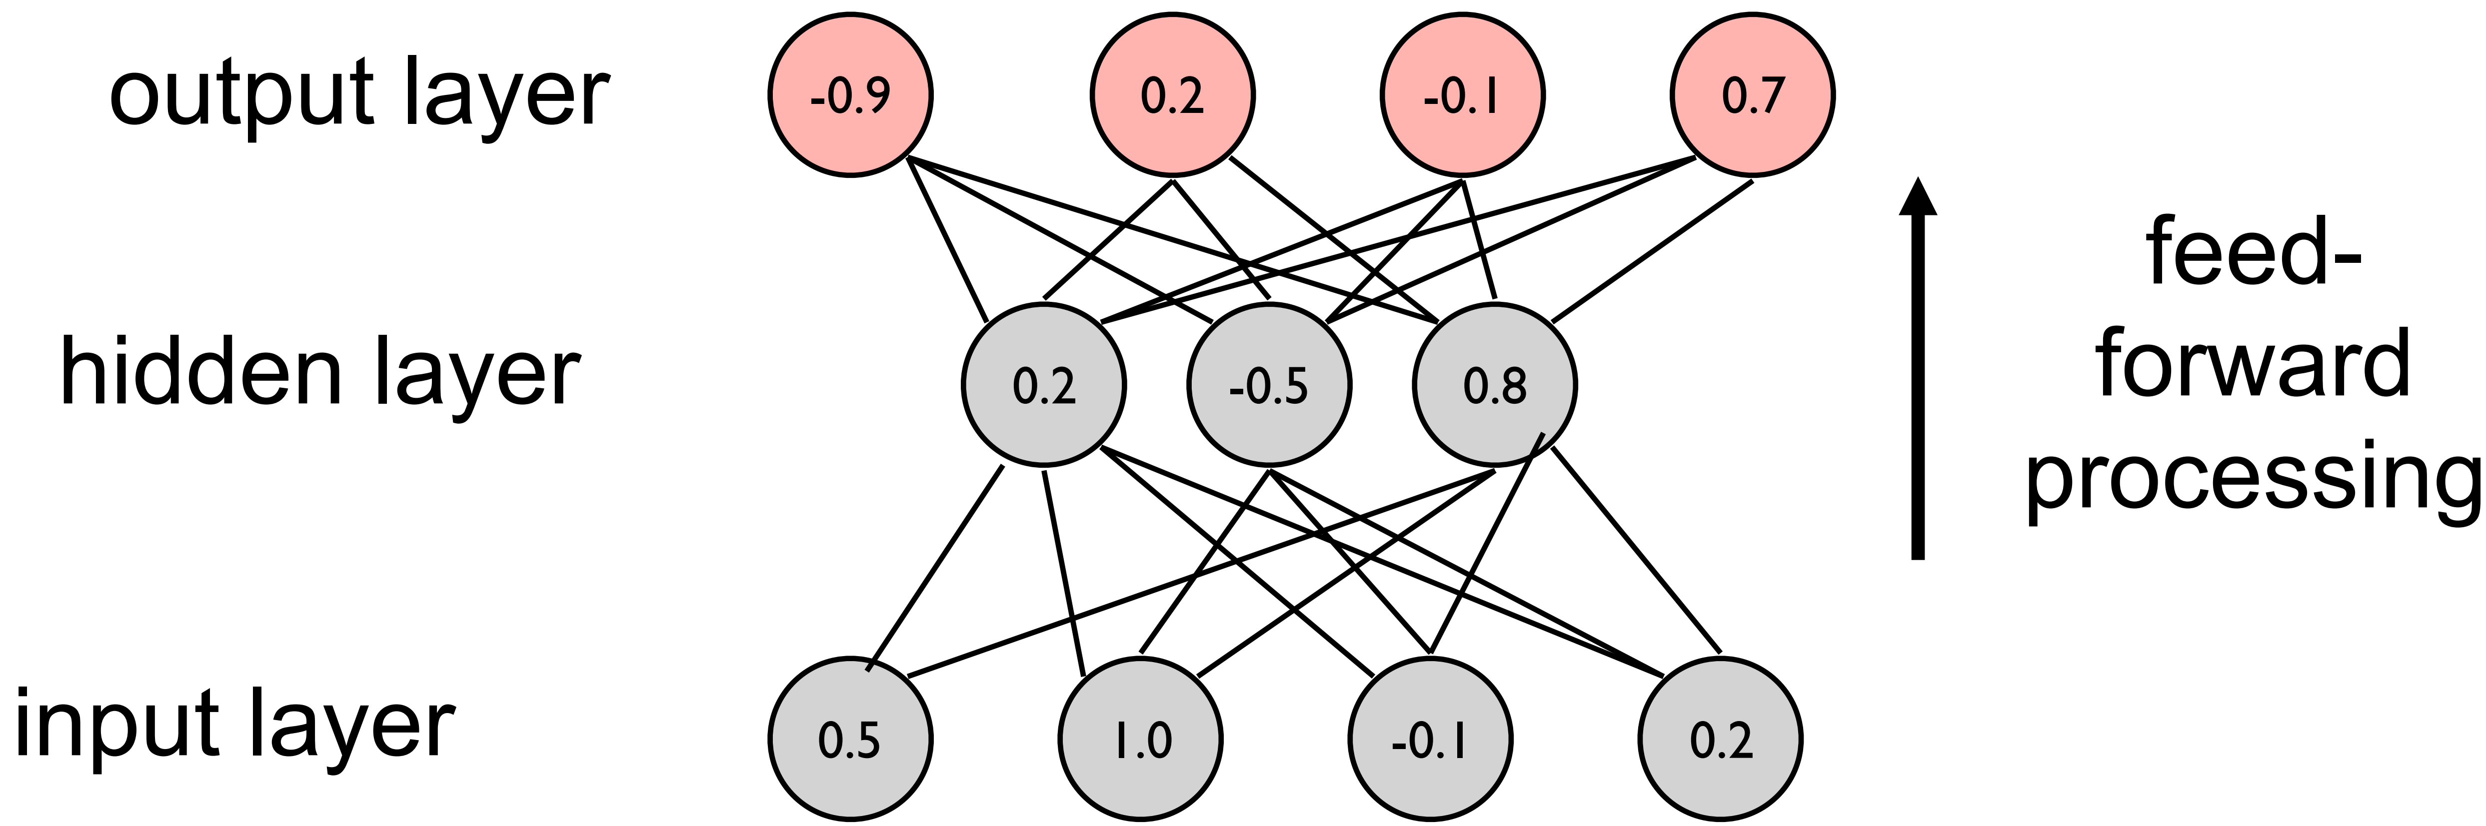

Input pattern: **<**0.5, 1.0,-0.1,0.2**>**

Other architectures........

Recurrent Network

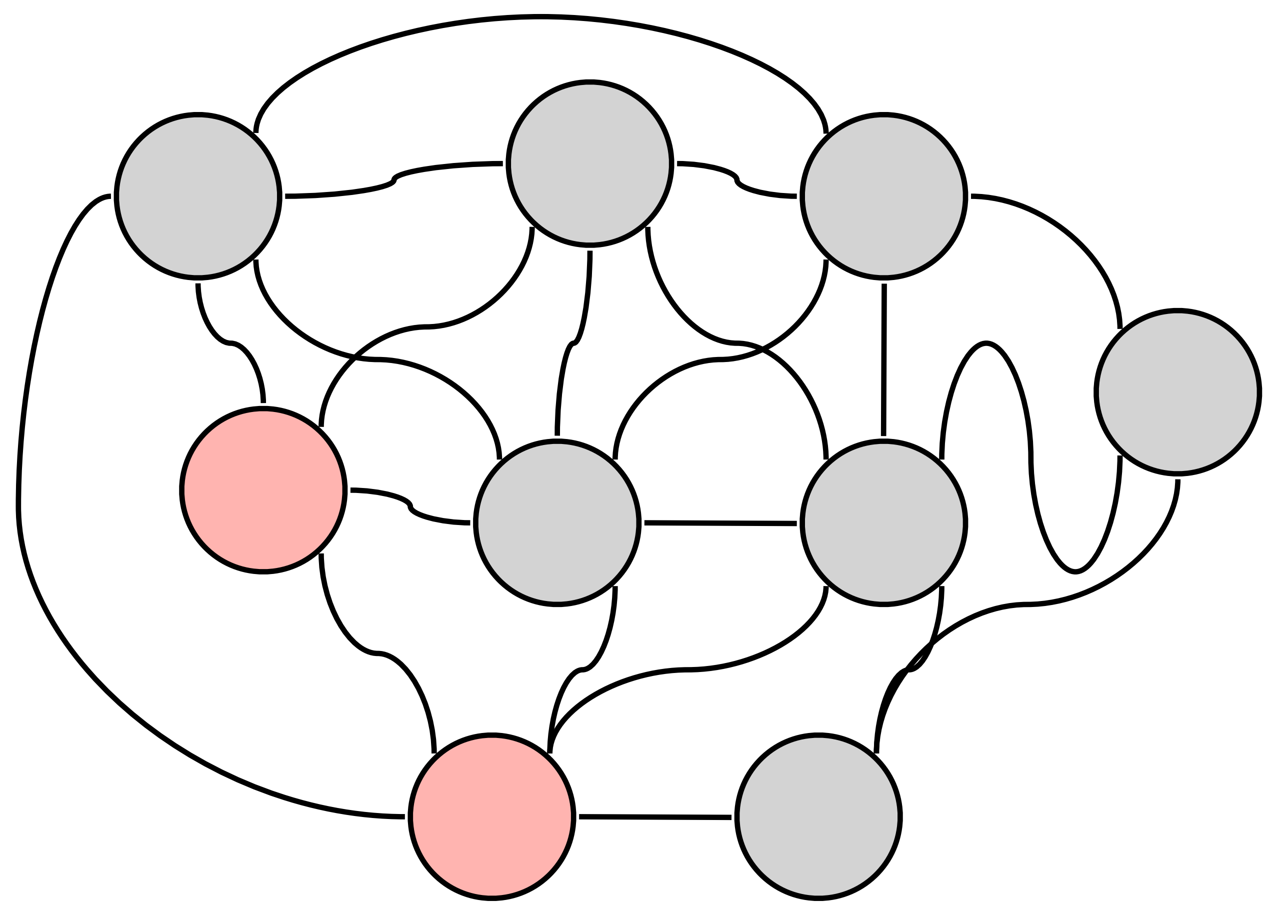

Other architectures........

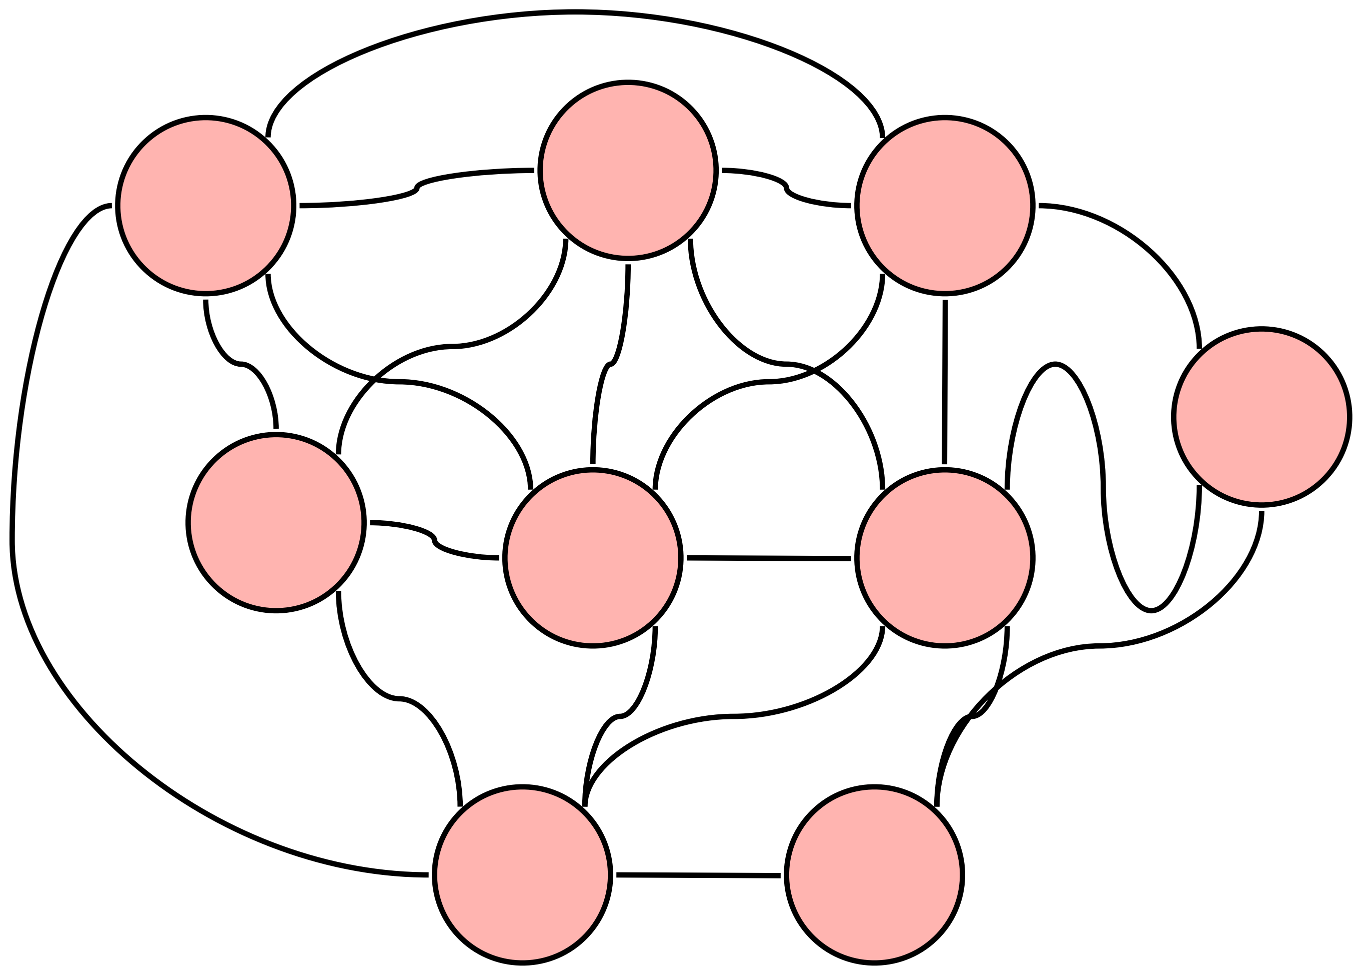

Other architectures........

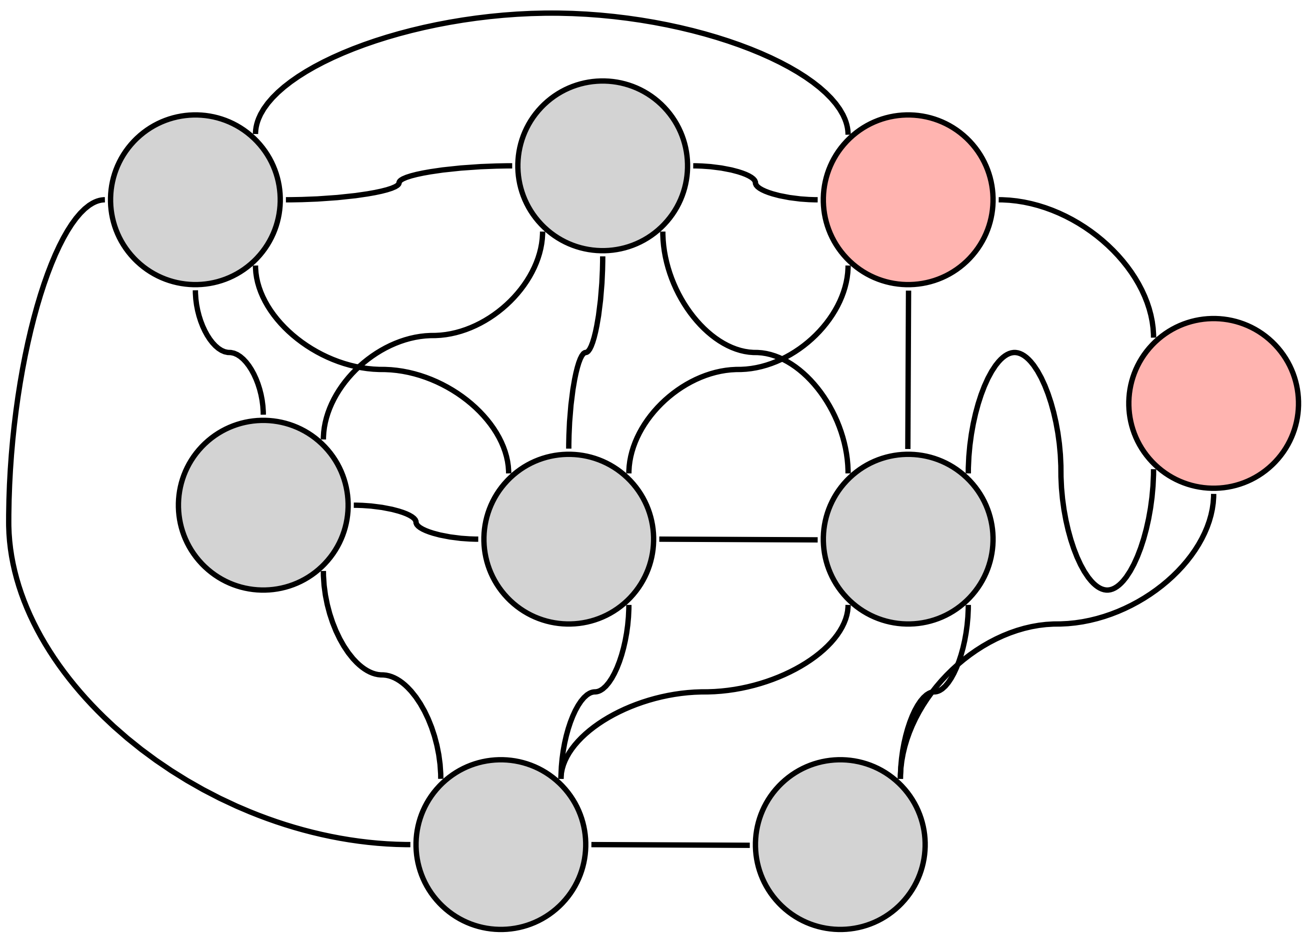

Memory in an ANN

- During processing, unit activation changes at each clock tick
- During training, weights typically change slowly
- Activations . . . Short-term memory
- Weights . . . Long-term memory

To Do List

- Read preface and get going on Chapters 1 and 2 of the book (Covers first several labs)
- We will be doing simple paper and pencil exercises in the first lab (but bring laptops)
- and getting to know the customized software (still under development!) we will be exploring together