Artificial Neural Networks

Q550: Models in Cognitive Science

"Intelligence is 10 million rules." --Doug Lenat

The human brain has about 100 billion neurons. With an estimated average of one thousand connections between each neuron and its neighbors, we have about 100 trillion connections, each capable of a simultaneous calculation... (but) only 200 calculations per second... With 100 trillion connections, each computing at 200 calculations per second, we get 20 million billion calculations per second. This is a conservatively high estimate... by the year 2020, (a massively parallel neural net computer) will have doubled about 23 times (from 1997's $2,000 modestly parallel computer that could perform around 2 billion connection calculations per second) ... resulting in a speed of about 20 million billion neural connection calculations per second, which is equal to the human brain.

Ray Kurzweil, "The Age of Spiritual Machines", 1999

Biologically Inspired Models

- We want to add biological constraints to our model as well as behavioral (output data) constraints
- Our molar level tends to be the neuron
- If neurons are the hardware that the system we're trying to model is operating with, let's consider what we know about them

Some random facts from Jeopardy…

Some Biological Facts

- 99% of our DNA is the same as that of a chimp
- The nervous system consumes 25% of the body's energy
- We typically think of the neuron as the basic unit of thought:
  - Dendrite (input)
  - Cell body/soma (integrator)
  - Axon (communication line)
  - Terminal (output)
  - Synapses (connections)
  - Neurotransmitters (messages)

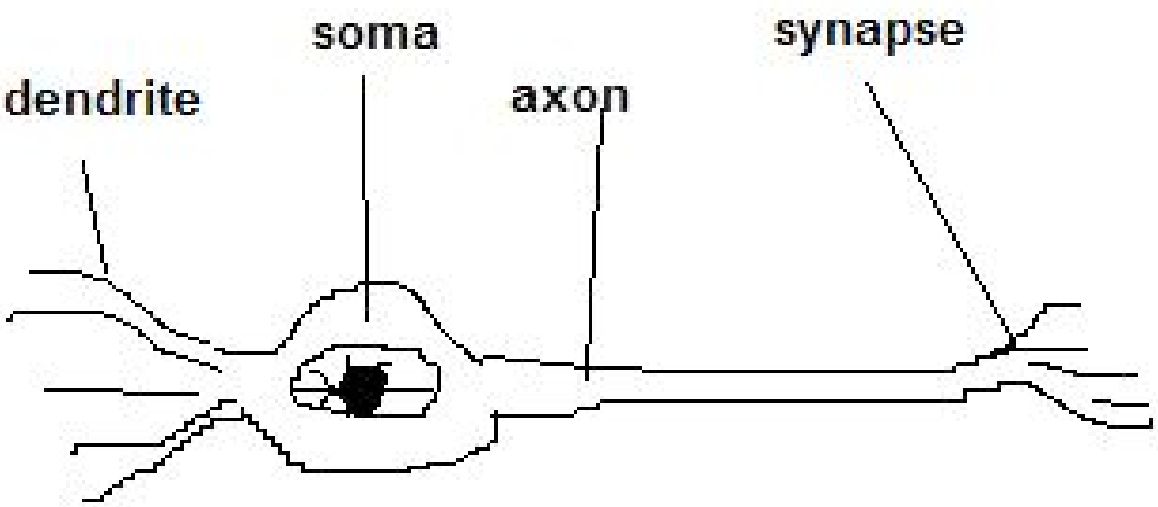

A Single Neuron

$$net = \sum_{i=0}^{n} Wi\ Xi$$

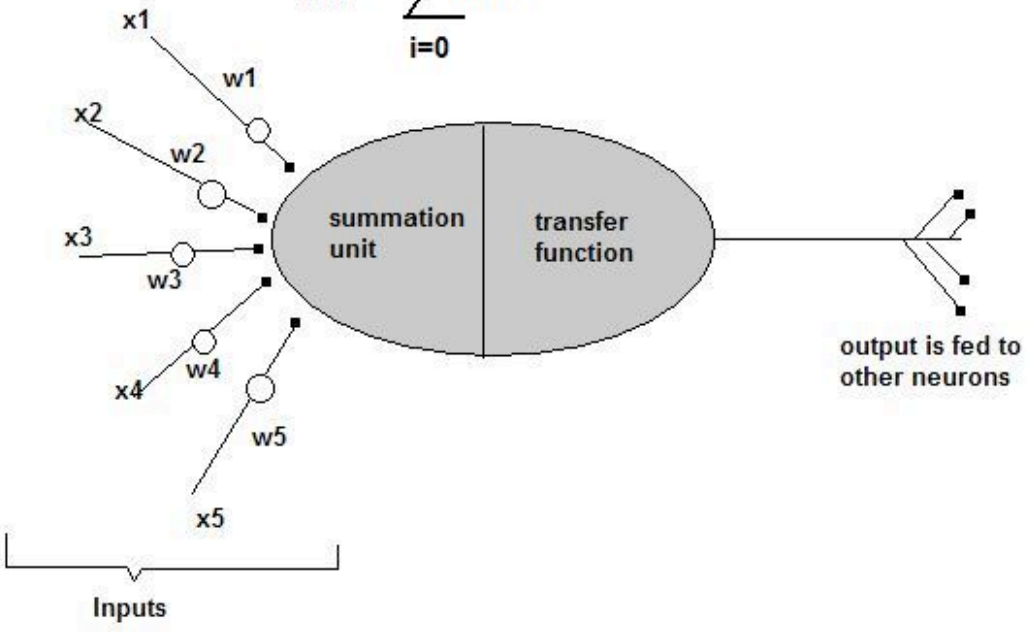

Some Biological Facts

- There is competition for connections among neurons, and unconnected cells die (50% before birth)
- 10^{11} neurons are left in the brain after birth (with no new generation)
- 10^{15} neural connections. Each neuron connects to a small fraction of others (few hundred or thousand)
- Unused or unconnected neurons die
- Excitatory connections increase a cells firing potential; inhibitory connections reduce it
- Electrical response of neuron is an all-or-none potential or spike

Some Biological Facts

- A single cell can have several thousand synapses on it
- The inputs from different synapses approximately add at the cell body
- Function relating integrated activation to firing rate is a nonlinear sigmoidal function

OK, now lets incorporate some of these constraints into our models…(starting with simple tasks)

Types of Neural Nets:

1. **Feedforward**
   - Perceptrons, linear associators
2. **Recurrent**
   - Hidden layer's state depends on it's state at a previous time
   - SRNs, Hopfield nets
3. **Stochastic**
   - Boltzman machines (noisy networks)

Types of Training:

1. **Supervised**
   - Give data and correct response
2. **Unsupervised**
   - Give data only
3. **Reinforcement**
   - Data not provided, but determined by the models interaction w/ environment. Goal is to discover a *policy* for selecting actions that minimizes long-term costs

Autoassociative vs. **heteroassociative**

Learning rules: delta rule, backprop, gradient descent, evolutionary, expectation maximization

Datasets online:

ANNs are very good a pattern recognition/classification

On our website, you'll find some sets of training and testing exemplars to try the algorithms out on

- Feature lists for capital letters (from Rumelhart & McClelland)
- Handwritten letters, numbers, and math symbols (NIST and some of my own)
- Fingerprint vectors (NIST)
- Elman's "cat chase mouse" artificial language
- VSM vectors on Reuters's documents

Rumelhart & McClelland (1981) PR

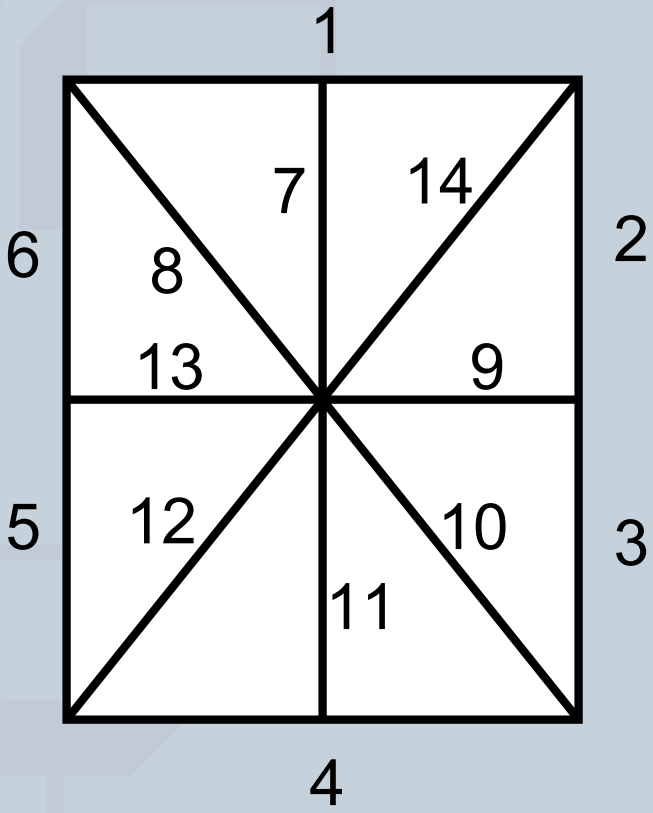

Uppercase letters are represented by binary vectors of these 14 features

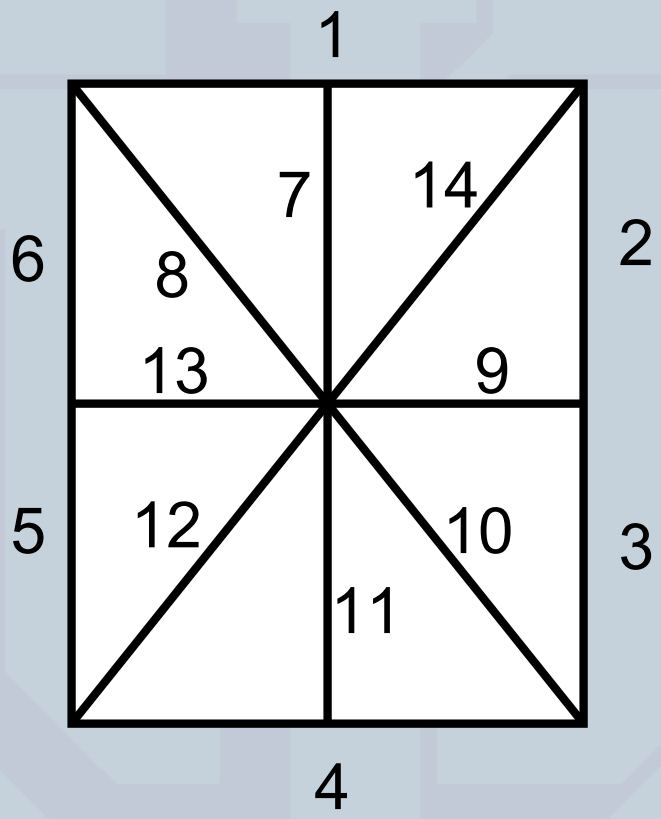

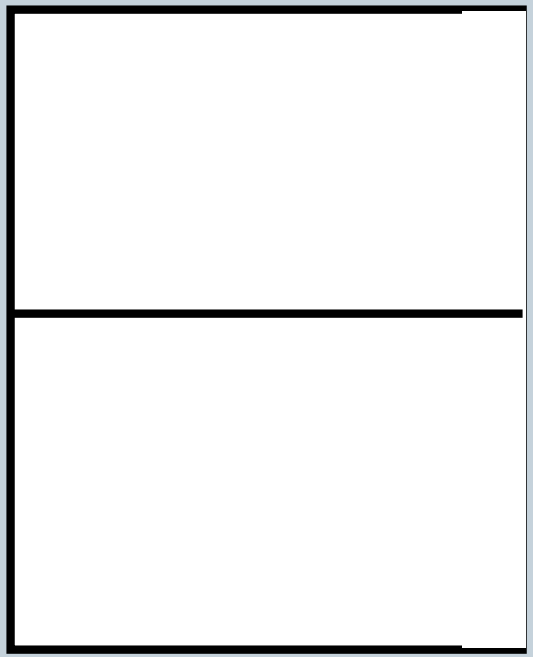

E = [1 0 0 1 1 1 0 0 1 0 0 0 1 0]

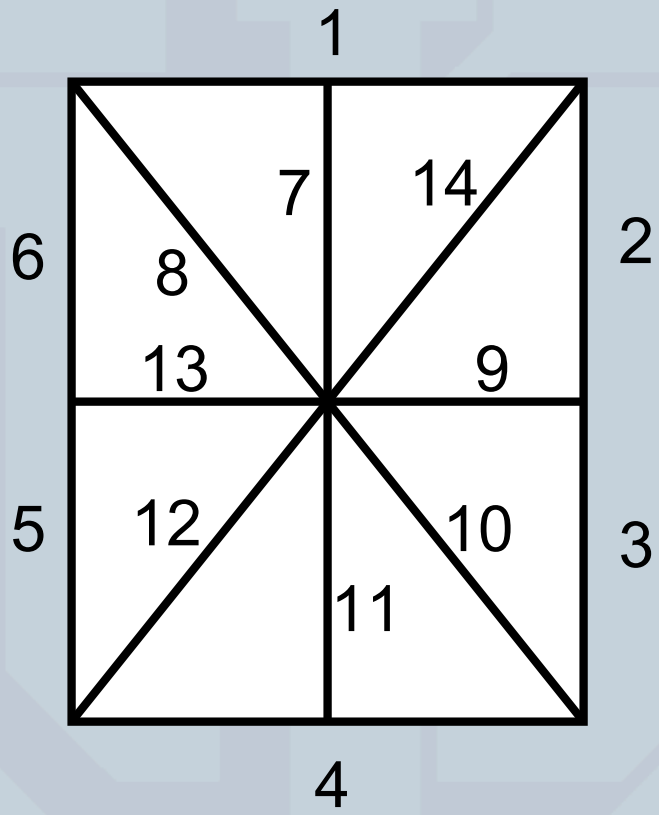

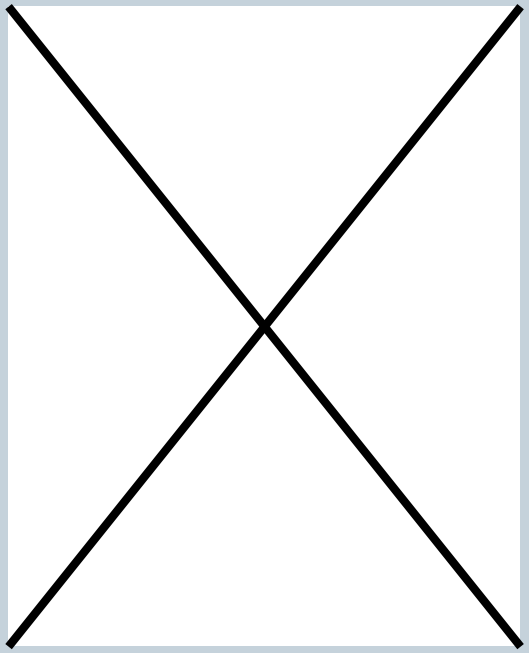

X = [0 0 0 0 0 0 0 1 0 1 0 1 0 1]

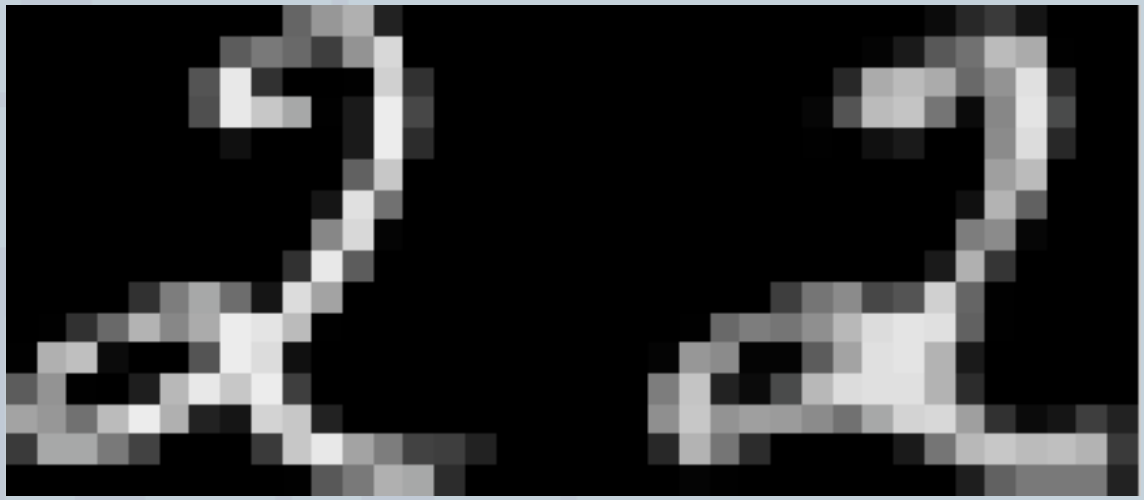

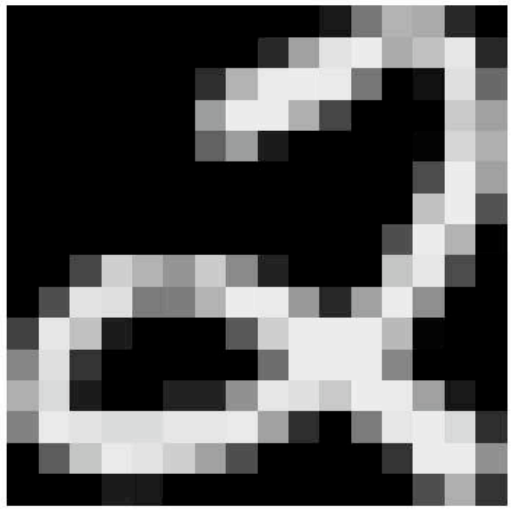

→ 0 0 1 1 0 1 0 1 0 1 0 1 1

.95 .85 .79 .23 .04 .01 ..

Single-Layer Perceptron

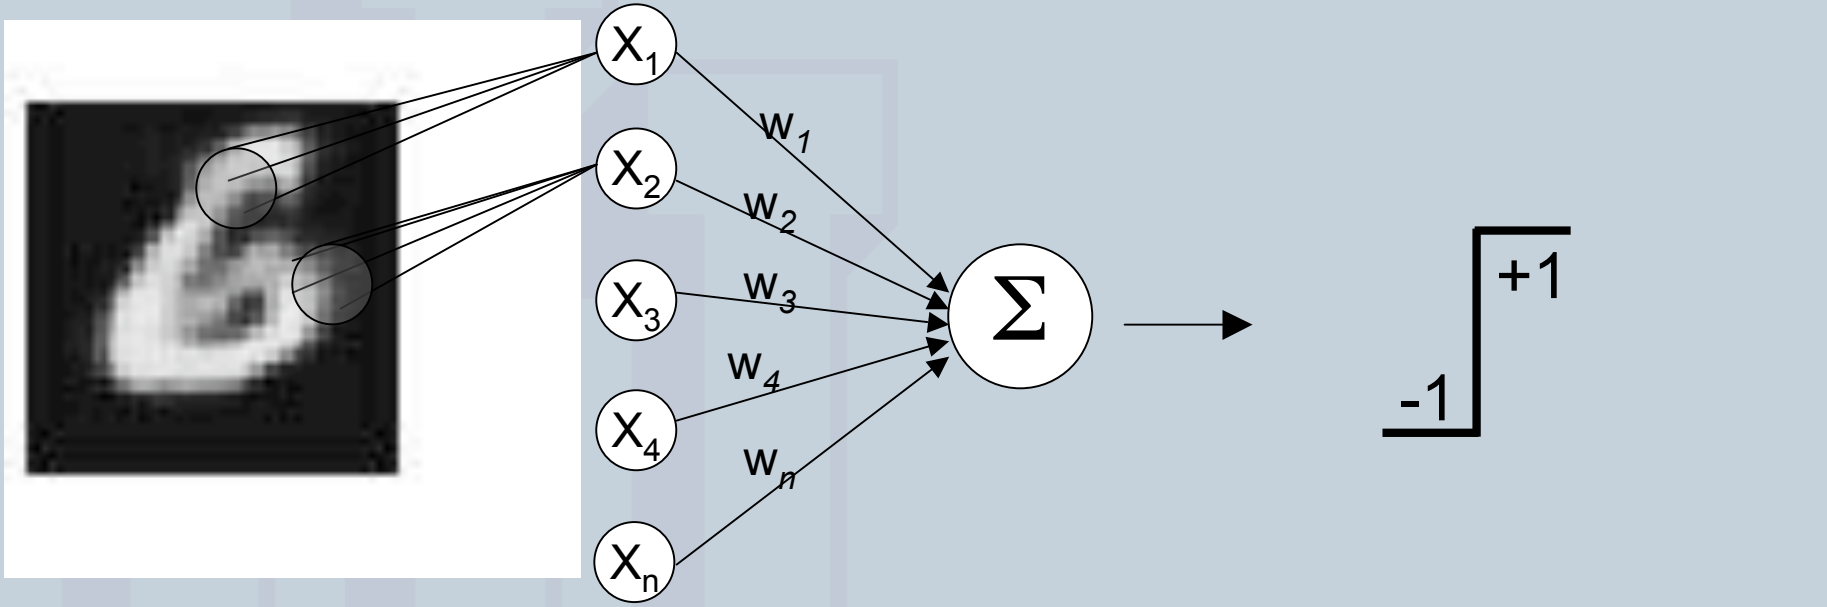

Artificial Retina

Input Nodes

Input Summation

THD
(McCulloch-Pitts neuron)

If $y = t$, do nothing
If $y \neq t$, then delta update:

$$\Delta w_{ij} = \alpha(t_i - y_i)x_j$$

Multiclass Single-Layer Perceptron

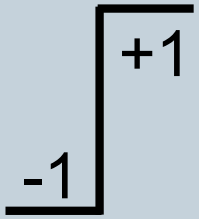

The death of perceptrons

Minsky & Papert (1969) *Perceptrons*

- They can only learn categories that are linearly separable
- They cannot do XOR

This led to a decline in funding, and it took more than a decade for neural nets to make a comeback with Grossberg's work

Multilayer Perceptrons

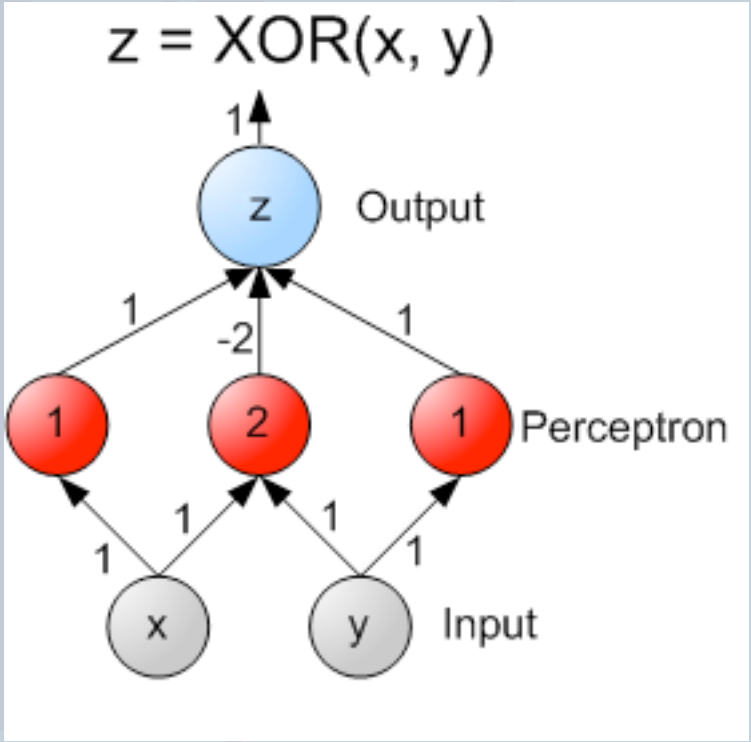

Multilayer Perceptrons

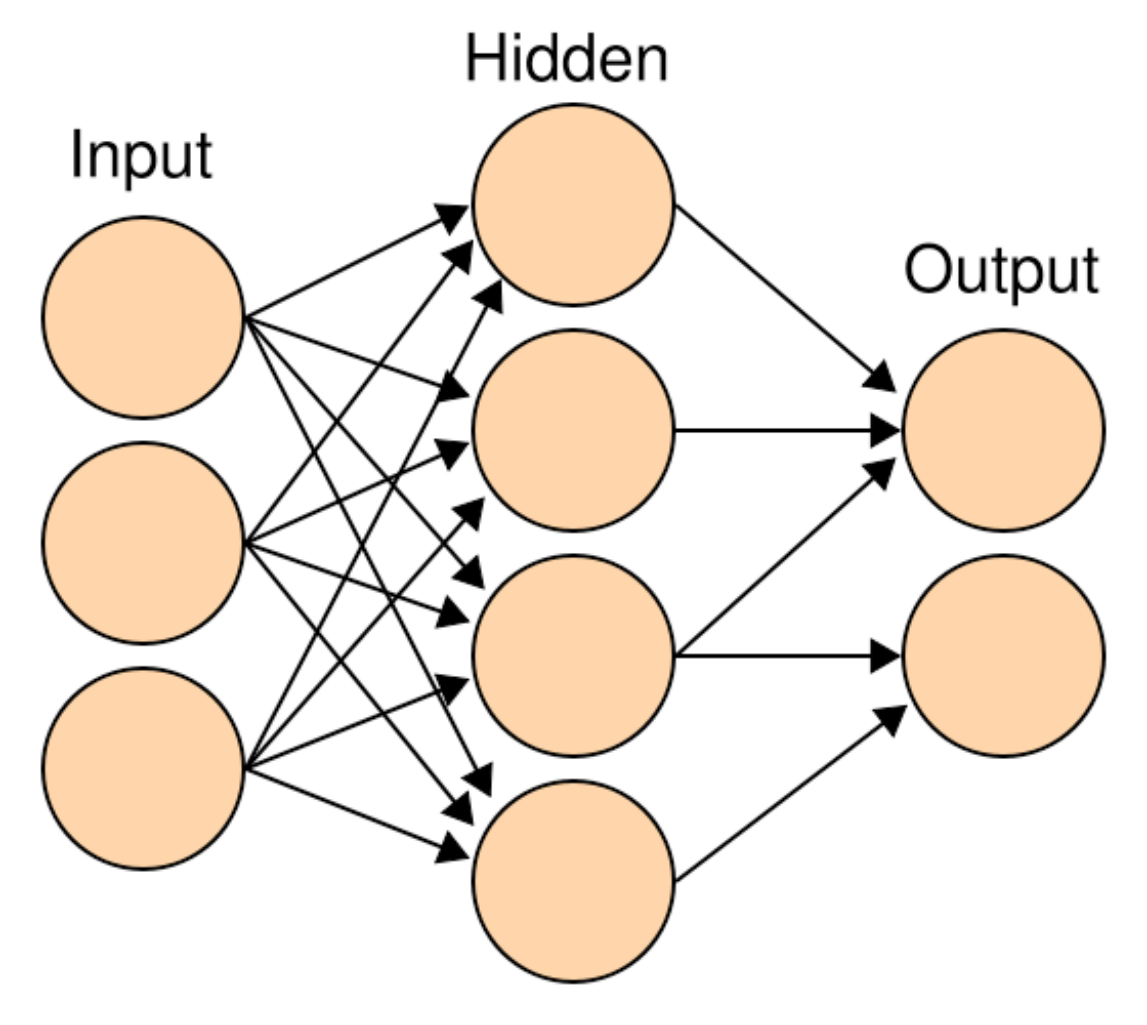

But let's use a smarter update rule based on gradient descent and "assigning blame where blame is due"

a.k.a. Error Backpropagation

PDP++

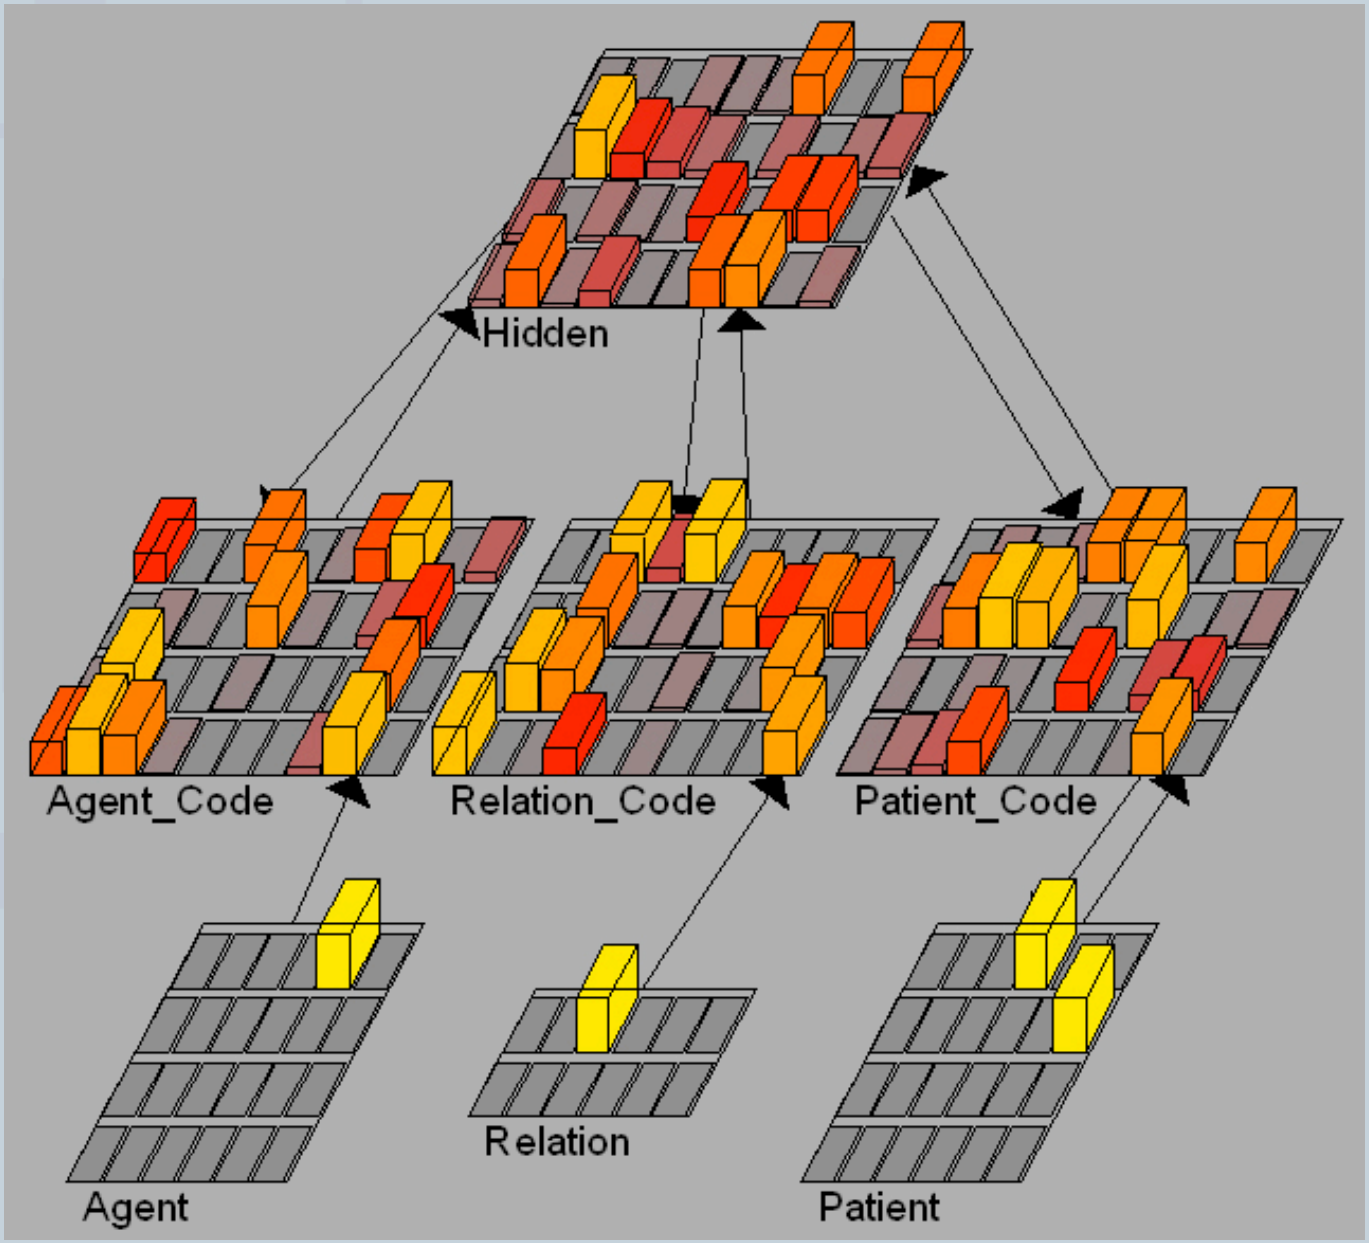
Hidden
Agent_Code
Relation_Code
Patient_Code
Agent
Relation
Patient

O’Reilly: Fitting behavioral data without biological constraints is of questionable value in understanding how the brain actually subserves behavior

Ratcliff: but few neurally plausible models fit data as well as global memory models