Overview

- Learning
- Background: the human nervous system
- Overview of neural network models
- Application of neural networks

Learning

- Adapt through interaction with the world: rote memory to developing a complex strategy
- Types of learning:
  1. Supervised learning (dense feedback)
  2. Unsupervised learning (no feedback)
  3. Reinforcement learning (sparse feedback), etc.
- Advantages (two, among many):
  1. Fault tolerance
  2. No need for a complete specification to begin with
- Becoming a central focus of AI.

Neural Networks

Neural networks is one particular form of learning from data.

- simple processing elements: named units, or neurons
- connective structure and associated connection weights
- learning: adaptation of connection weights

Neural networks mimic the human (or animal) nervous system.

Many Faces of Neural Networks

- Abstract mathematical/statistical model
- Optimization algorithm
- Pattern recognition algorithm
- Tools for understanding the function of the brain
- Robust engineering application

The Central Nervous System

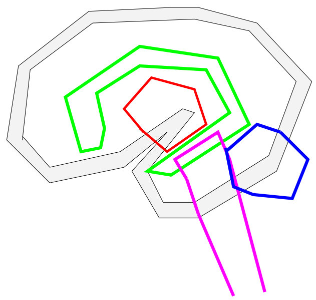

- Cortex: thin outer sheet where most of the neurons are.
- Sub-cortical nuclei: thalamus, hippocampus, basal ganglia, etc.
- Midbrain, pons, and medulla, connects to the spinal cord.
- Cerebellum (hind brain, or little brain)

Function of the Nervous System

Function of the nervous system:

- Perception
- Cognition
- Motor control
- Regulation of essential bodily functions

The Central Nervous System: Facts [a]

Facts: human neocortex

- Thickness: 1.6mm
- Area: 36cm $\times$ 36cm (about 1.4 ft^2)
- Neurons: 10 billion (10^{10})
- Connections: 60 trillion (6×10^{13}) to 100 trillion
- Connections per neuron: 10^4
- Energy usage per operation: 10^{-16}J (compare to 10^{-6}J in modern computers)

[a] *Neural networks: a comprehensive foundation* by Simon Haykin (1994), and *Foundations of Vision* by Brian Wandell (1995). May slightly differ from those in Russel & Norvig. No need to memorize these figures.

How the Brain Differs from Computers

- Densely connected.
- Massively parallel.
- Highly nonlinear.
- Asynchronous: no central clock.
- Fault tolerant.
- Highly adaptable.
- Creative.

Why are these crucial?

Neurons: Basic Functional Unit of the Brain

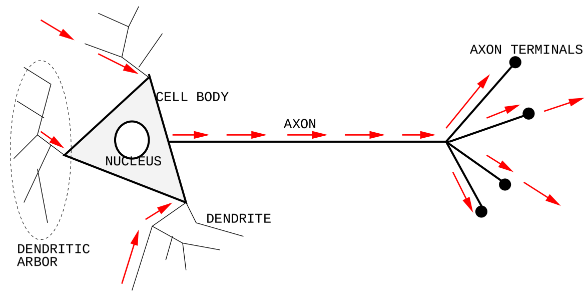

- Dendrites receive input from upstream neurons.
- Ions flow in to make the cell positively charged.
- Once a firing threshold is reached, a spike is generated and transmitted along the axon.
- Axon terminals release neurotransmitters to relay the signal to the downstream neurons.

Abstraction of the Neuron in Neural Networks

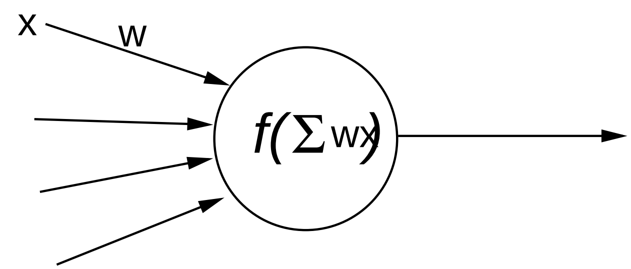

- Input
- Connection weight
- Transfer function: $f(\cdot)$

Typical transfer functions: step-function or sigmoid.

Key Points

- Types of learning
- Neural networks: basics
- The central nervous system: how it differs from conventional computers.

Next Time

- Perceptrons
- Multilayer perceptrons and backpropagation

Homework: Due 11/13/02 (Wed) in class

1. Unify (if possible) the following pairs of predicates and give the resulting substitutions. b is a constant.

   (a) $P(f(a), g(x))$ and $P(y, y)$

   (b) $P(x, f(x), z)$ and $P(g(y), f(g(b)), y)$

   (c) $P(x_1, g(x_1), x_2, h(x_1, x_2), x_3, k(x_1, x_2, x_3))$ and $P(y_1, y_2, e(y_2), y_3, f(y_2, y_3), y_4)$

2. Determine whether the following clauses have factors. If yes, give the factors.

   (a) $P(x) \lor Q(y) \lor P(f(x))$

   (b) $P(a) \lor P(b) \lor P(x)$

   (c) $P(x) \lor P(a) \lor Q(f(x)) \lor Q(f(a))$

3. Given the following axioms and the conclusion, (1) convert the English sentences into standard form (prenex normal form, then conjunctive normal form) in first-order logic, and (2) using resolution, show that the conclusion is a logical consequence of the axioms.

   1. Every coyote chases some roadrunner.
   2. Every roadrunner who says “beep-beep” is smart.
   3. No coyote catches any smart roadrunner.
   4. Any coyote who chases some roadrunner but does not catch it is frustrated.
   5. **Conclusion** If all roadrunners say “beep-beep”, then all coyotes are frustrated.

Homework: cont’d

4. Show the following using definitions and axioms in pp.421–423:

   (a) $P(A|B \land A) = 1$

   (b) $P(A|B, C) = \frac{P(A,B|C)}{P(B|C)}$

5. Consider exercise 14.3, p.433 in the textbook. How accurate should the test – $P(TestPositive|Disease)$ and $P(\neg TestPositive|\neg Disease)$ – be so that the probability that you actually have the disease if you tested positive, i.e. $P(Disease|TestPositive)$, is over 90%?

6. Using the belief network given in p.439 of the textbook (Figure 15.2), calculate the probability that only John calls after he hears the alarm go off, and when there was actually a burglary, but not an earthquake.

Submit a written solution (on paper) to the instructor before the class on 11/18/2002 (Monday). Show all of your work. **No extensions nor late submissions are allowed.**