

Chapter 20 Section 5 --- Slide Set 1

Prologue to neural networks: the neurons in the human brain

Additional sources used in preparing the slides:

Various sites that explain how a neuron works

Robert Wilensky's slides: <http://www.cs.berkeley.edu/~wilensky/cs188>

Inspiration: The human brain

- **We seem to learn facts and get better at doing things without having to run a separate “learning procedure.”**
- **It is desirable to integrate learning more with doing.**

Understanding the brain (1)

“ Because we do not understand the brain very well we are constantly tempted to use the latest technology as a model for trying to understand it. In my childhood we were always assured that the brain was a telephone switchboard. (“What else could it be?”) I was amused to see that Sherrington, the great British neuroscientist, thought that the brain worked like a telegraph system. Freud often compared the brain to hydraulic and electro-magnetic systems. Leibniz compared it to a mill, and I am told that some of the ancient Greeks thought the brain functions like a catapult. At present, obviously, the metaphor is the digital computer.”

-- John R. Searle

(Prof. of Philosophy at UC, Berkeley)


Understanding the brain (2)

“ The brain is a tissue. It is a complicated, intricately woven tissue, like nothing else we know of in the universe, but it is composed of cells, as any tissue is. They are, to be sure, highly specialized cells, but they function according to the laws that govern any other cells. Their electrical and chemical signals can be detected, recorded and interpreted and their chemicals can be identified, the connections that constitute the brain’s woven feltwork can be mapped. In short, the brain can be studied, just as the kidney can.

**-- David H. Hubel
(1981 Nobel Prize Winner)**

The brain


- The brain doesn't seem to have a CPU.
- Instead, it's got lots of simple, parallel, asynchronous units, called *neurons*.
- There are about 10^{11} neurons of about 20 types



Neurons

- Every neuron is a single cell that has a number of relatively short fibers, called *dendrites*, and one long fiber, called an *axon*.
 - The end of the axon branches out into more short fibers
 - Each fiber “connects” to the dendrites and cell bodies of other neurons
 - The “connection” is actually a short gap, called a *synapse*
 - Axons are transmitters, dendrites are receivers
- There are about 10^{14} connections


Neuron



How do neurons work

- **The fibers of surrounding neurons emit chemicals (neurotransmitters) that move across the synapse and change the electrical potential of the cell body**
 - **Sometimes the action across the synapse increases the potential, and sometimes it decreases it.**
 - **If the potential reaches a certain threshold, an electrical pulse, or action potential, will travel down the axon, eventually reaching all the branches, causing them to release their neurotransmitters. And so on ...**


How do neurons work (cont'd)



How do neurons change

- **There are changes to neurons that are presumed to reflect or enable learning:**
 - The synaptic connections exhibit *plasticity*. In other words, the degree to which a neuron will react to a stimulus across a particular synapse is subject to long-term change over time (*long-term potentiation*).
 - Neurons also will create new connections to other neurons.
 - Other changes in structure also seem to occur, some less well understood than others.

Neurons as devices



- Neurons are slow devices.
 - Tens of milliseconds to do something. (1ms – 10ms cycles time)
 - Feldman translates this into the “100 step program constraint”: Most of the AI tasks we want to do take people less than a second. So any brain “program” can’t be longer than 100 neural “instructions.”
- No particular unit seems to be important. Destroying any one brain cell has little effect on overall processing.

How do neurons do it?

- **Basically, all the billions of neurons in the brain are active at once. So, this is truly *massive parallelism*.**
- **But, probably not the kind of parallelism that we are used to in conventional Computer Science.**
 - **Sending messages (i.e., patterns that encode information) is probably too slow to work.**
 - **So information is probably encoded some other way, e.g., by the connections themselves.**

AI / Cognitive Science Implication

- **Explain cognition by richly connected networks transmitting simple signals.**
- **Sometimes called**
 - **Connectionist computing**
(by Jerry Feldman)
 - **Parallel Distributed Processing (PDP)**
(by Rumelhart, McClelland, and Hinton)
 - **Neural networks (NN)**
 - **Artificial neural networks (ANN)**
(emphasizing that the relation to biology is generally rather tenuous)

From a neuron to a perceptron

- **All connectionist models use a similar model of a neuron**
- **There is a collection of units each of which has**
 - a number of weighted *inputs* from other units
 - inputs represent the degree to which the other unit is firing
 - weights represent how much the units wants to listen to other units
 - a *threshold* that the sum of the weighted inputs are compared against
 - the threshold has to be crossed for the unit to do something (“fire”)
 - a single *output* to another bunch of units
 - what the unit decided to do, given all the inputs and its threshold

Notes

- The perceptrons are continuously active
 - Actually, real neurons fire all the time; what changes is the *rate of firing*, from a few to a few hundred impulses a second
- The weights of the perceptrons are not fixed
 - Indeed, learning in a NN system is basically a matter of changing the weights

Comments on neural networks

- **Parallelism in AI is not new.**
 - **spreading activation, etc.**
- **Neural models for AI is not new.**
 - **Indeed, is as old as AI, some subdisciplines such as computer vision, have continuously thought this way.**
- **Much neural network works makes biologically implausible assumptions about how neurons work**
 - **backpropagation is biologically implausible**
 - **“neurally inspired computing” rather than “brain science.”**

Comments on neural networks (cont'd)

- **None of the neural network models distinguish humans from dogs from dolphins from flatworms. Whatever distinguishes “higher” cognitive capacities (language, reasoning) may not be apparent at this level of analysis.**
- **Relation between neural networks and “symbolic AI”?**
 - **Some claim NN models don't have symbols and representations.**
 - **Others think of NNs as simply being an “implementation-level” theory. NNs started out as a branch of statistical pattern classification, and is headed back that way.**

Nevertheless

- NNs give us important insights into how to think about cognition
- NNs have been used in solving lots of problems
 - learning how to pronounce words from spelling (NETtalk, Sejnowski and Rosenberg, 1987)
 - Controlling kilns (Ciftci, 2001)