What are Neural Networks?.....

and why am I bothering to study them, anyway?

Consider traditional, serial computers...

- "Von Neumann" architectures
- Single, complex CPU
- Processes one instruction at a time
- Operates on one chunk of data at a time

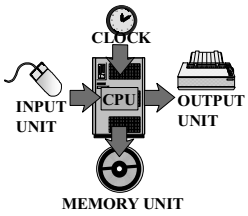

Methods to obtain greater power in serial computers:

- Increase the size of memory chunks operated on (vector processing)
- Reduce instructions necessary for each computational task (RISC architectures)
- Pump up the clock speed (supercomputers)
- Increase the complexity of what's done per clock step (math coprocessors)

Problems with serial computation

- limits of practical speed-up have nearly been reached
- despite this fact, serial computers remain incapable of many valuable tasks

Parallel Computers

- Many Von Neumann style computers, operating in parallel, and talking to one another.
- Although parallelism yields significant speed-up on some tasks:
  - communication overhead becomes a problem
  - it as proved difficult to parallelize many serial algorithms

Consider object recognition

- in arbitrary, noisy settings....

Pattern Recognition and... Art Appreciation

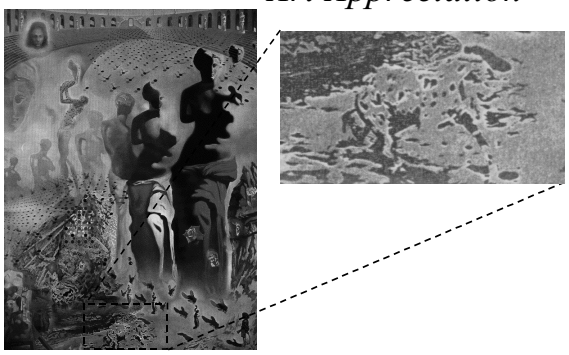

Tasks that don't easily yield to traditional computers:

- object detection in complex environments
- control of complex (time-varying) systems
- hand-written character and symbol recognition
- language understanding
- distinguishing signals from noise
- learning from experience
- recovery from damage

Interestingly....

- these tasks are of the sort that living organisms solve easily (in some cases, instantaneously)
- Q: Why?
  - A: Magic?
  - A: Differences in the way computation is performed (hardware and software)

Hardware and software in organic brains

- Brains are composed of many, simple processors operating in parallel on lots of small packets of low-level data at relatively slow speeds (1 Khz)

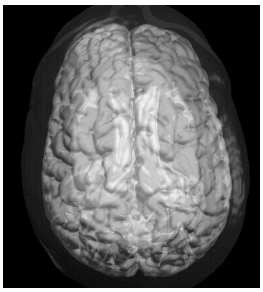

A Word of Caution:

- Artificial Neural Systems are very rough (perhaps incorrect) models of Living Neural Systems!
- However, despite possible errors in modeling real neurons, ANSs may contain important features of all massively parallel distributed computing systems.

The Scale of The Human Brain:

- The human brain contains on the order of a trillion neurons
- On average, a neuron has on the order of 10,000 synapses (connections to other neurons)...
- Giving on the order of 10 quadrillion connections!

Parts of the Brain:

- Cerebral Hemisphere: Higher Brain Functions
- Upper Brain Stem: Regulation
- Middle Brain Stem: Visual and Auditory Reflexes
- Lower Brain Stem: Critical Reflexes

The Cerebral Hemisphere

- Is arranged into distinct, 2-D layers
- This sheet-like structure is divided into distinct "lobes" that relate to specific functions:
  - Frontal Lobe: Planning
  - Parietal Lobe: Abstract Reasoning
- Temporal Lobe: Auditory Processing
- Ocipital Lobe: Visual Processing

Neuron Structure:

- Neurons consist of
  - a body (the soma),
  - a set of input lines (the dendrites), and
  - a single output line (the axon)

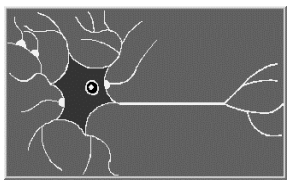

Connections

- Dendrites spread out into a tree that receives a great many inputs from axons of other neurons, through terminals (synapses)
- Axons are from 100 micrometers to 1 meter long!

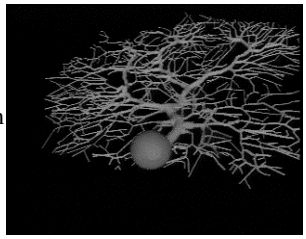

Synapses

- Synapses send electrochemical signals that alter the "activation" potential of the receiving neuron
- Two types of synaptic effects:
  - Excitatory connections, which raise potential and promote firing
  - Inhibitory connections, which lower potential and discourage firing

Mechanism of Neuron Firing:

- In its rest state, a synaptic terminal is permeable to potassium ions, that flow out and create a -70mV resting potential
- Activation builds up at synaptic terminals

Firing…

- When activation (within a certain time) passes a internally-defined threshold, the membrane becomes permeable to potassium ions, that flow in an create a +35mV pulse, which is sent down the axon

Neuron mechanisms continued...

- After the pulse, K+ ions flow out, causing a drop of potential beneath the resting potential
- After "firing", the neurons internal "pumps" return ion balance
- This is where the neuron consumes energy
- The neuron takes about a millisecond to be ready to fire again (Speed of 1 kHz!)

Avoiding Anthropomorphism:

- What are lower animal brains like:
- Reflex Nets: in animals like the hydra
  - sensors to neurons to muscles (no interconnect)
- Primitive Associative Nets: in flatworms
  - interconnects, no CNS

Simple Brains, continued

- Primitive CNSs: in earthworms
  - bundles on interconnected neurons centralized
- Primitive differentiated CNSs and peripheral nervous systems: in insects
  - CNS and PNS consist of different types of neurons

Neural Networks (aka Connectionism)

- Is an attempt to get machines to do the tasks brains do well (and typical computers do poorly) through abstraction of low-level brain function
- Is an attempt to understand brain function through simplified computer simulation

Neural Network Terminology:

- Artificial "Neurons" are also called "Nodes" or "Units"
- Nodes are connected in "Layers"
- Outputs of units in one layer are inputs to another layer
- Units compute a simple function (the "Unit Function" over there inputs to determine there output

Training

- Rather than being programmed, most neural nets are trained
- Inputs to each node are assigned "weights"
- In training, examples of correct input-output behavior are provided to the network

More Training

- Weights are adjusted by a learning algorithm to bring network behavior in line with the behavior indicated by examples
- It is hoped that the network will "generalize" from the examples to broadly accurate behavior.

Implementing Neural Networks

- Most NNs are implemented in digital simulators on serial computers
- However, NNs are ultimately best implemented on special purpose hardware
  - Digital NN chips (currently available, with some size and speed limitations)
  - Analog NNs (yielding analog speed)

More Terminology

- Supervised Learning
  - Learning with a teacher that provides solved examples
- Reinforcement Learning
  - Learning with a teacher that provides only rewards and punishment
- Unsupervised Learning
  - Learning without a teacher (internally defined goals)

Next Time:

- Given neurons as binary, temporal computational devices, how do we perform useful computations with a group of them?
  - McCullogh-Pitts Neuron Models
- How do Neural Nets Learn?
  - Hebbian weight adjustment
- Our first “real” neural net:
  - percetrons