Neural Networks

CPS 270
Ron Parr

Why Neural Networks?

- Maybe computers should be more brain-like:

	Computers	Brains
Computational Units	10^8 gates/CPU	10^{11} neurons
Storage Units	10^9 bits RAM 10^{12} bits HD	10^{11} neurons 10^{14} synapses
Cycle Time	10^{-9} S	10^{-3} S
Bandwidth	10^{10} bits/s*	10^{14} bits/s
Compute Power	10^9 Ops/s	10^{14} Ops/s

Comments on Blue Gene

- Blue Gene: World's Fastest Supercomputer
- Currently 70 Teraflops -> 360 Teraflops
- 10^{13} -> 10^{14} Ops/s (brain level?)
- 16 TB memory (10^{13} bits)
- 14 Megawatts power ($1M/year in electricity)

- http://sc-2002.org/paperpdfs/pap.pap207.pdf
- http://www.research.ibm.com/bluegene/BG_External_Presentation_January_2002.pdf

Neural Network Motivation

- Individual neurons are slow, boring
- Brains succeed by using massive parallelism
- Copy what works
- Raises many issues:
  - Is the computational metaphor suited to the computational hardware?
  - How do we know if we are copying the important part?
  - Are we aiming too low?

Artificial Neural Networks

- Develop *abstraction* of function of actual neurons
- Simulate large, massively parallel artificial neural networks on conventional computers
- Some have tried to build the hardware too
- Try to approximate human learning, robustness to noise, robustness to damage, etc.

Use of neural networks

- Trained to pronounce English
  - Training set: Sliding window over text, sounds
  - 95% accuracy on training set
  - 78% accuracy on test set
- Trained to recognize handwritten digits
  - >99% accuracy
- Trained to drive
  (Pomerleau's no-hands across America)

Neural Network Lore

- Neural nets have been adopted with an almost religious fervor within the AI community several times
- They are often ascribed near magical powers by people, usually those who know the least about computation or brains
- For most AI people, the magic is gone, but neural nets remain extremely interesting and useful mathematical objects

Artificial Neurons

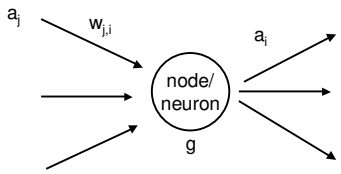

$$a_i = g(\sum_j w_{j,i} a_j)$$

g can be any function, but usually a smoothed step function

Threshold Functions

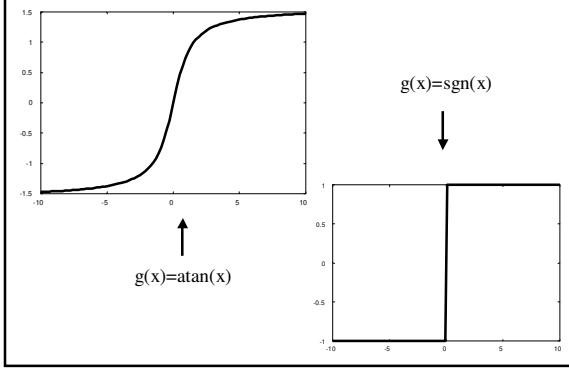

Network Architectures

- Cyclic vs. Acyclic
  - Cyclic is tricky, but more biologically plausible
    - Hard to analyze in general
    - May not be stable
    - Need to assume latches to avoid race conditions
  - Hopfield nets: special type of cyclic net useful for associative memory
- Single layer (perceptron)
- Multiple layer

Feedforward Networks

- We consider acyclic networks
- One or more computational layers
- Entire network can be viewed as computing a complex non-linear function
- Typical uses in learning:
  - Classification (usually involving complex patterns)
  - Function approximation (general case of pattern classification, where target is a general function)

Perceptron

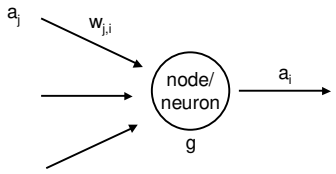

For now, we assume that g is a simple step function (sgn)
Assume only 1 neuron and one output.

Perceptron Learning

- We are given a set of inputs $A_1 \ldots A_n$
- $T_1 \ldots T_n$ is a set of target outputs (boolean)
- W is our set of weights
- $Net(A_i,W)$=output of perceptron given input A_i and weights W.
- $Error(A_i,W) = T_i - Net(Ai,W)$
- Goal: Pick W to optimize:

$$\min_W \sum_i error(A_i,W)$$

Update Rule

Repeat until convergence:

$$\forall i \forall j : W_j \leftarrow W_j + \alpha a_j Error(A_i,W)$$

"Learning Rate"

- i iterates over samples
- j iterates over weights

http://neuron.eng.wayne.edu/java/Perceptron/New38.html

Perceptron Learning Properties

- Good news:
  - If there exists a set of weights that will correctly classify every example, the perceptron learning rule will find it
- Bad news:
  - Perceptrons can represent only a small class of functions, "linearly separable," functions

Linearly Separable Functions

What is a perceptron really doing?

It checks if a linear combination of the inputs is greater than a threshold.

$$w_1 a_1 + w_2 a_2 \ldots w_n a_n > 0 ?$$

Perceptron asks: What side of a hyperplane does A lie on?

Q: How can we change from >0 to >C for arbitrary C?

Visualizing Linearly Separable Functions

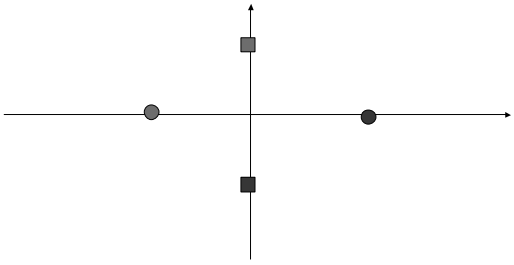

Is red linearly separable from green?
Are the circles linearly separable from the squares?

Observations

- Linear separability is fairly weak
- We have other tricks:
  - Functions that are not linearly separable in one space, may be linearly separable in another space
  - If we twiddle our inputs to our neural network, then we change the space in which we are constructing linear separators
  - Every function has a linear separator (in some space)
- Perhaps other network architectures will help

Multilayer Networks

- Once people realized how simple perceptrons were, they lost interest in neural networks for a while
- Multilayer networks turn out to be much more expressive (with a smoothed step function)
  - With 2 layers, can represent any continuous function
  - With 3 layers, can represent many discontinuous functions
- Tricky part: How to adjust the weights

Smoothing Things Out

- Consider single-layer case first
- Idea: Do gradient descent on a smooth error function
- Error function is sum of squared errors

$$E = 0.5 \sum_i error(A_i, W)^2$$

- i iterates over samples
- j iterates over weights

$$\frac{\partial E}{\partial w_j} = \sum_i error(A_i, W) \frac{\partial g(\sum_k w_k a_{ik})}{\partial w_j}$$

input k for sample i

$$= \sum_i error(A_i, W) g'(\sum_k w_k a_{ik}) a_{ij}$$

Back-prop summary

- Decomposes into a set of local computations
- Strong arguments that this can be parallelized
- Biologically plausible
- Gradient descent will find a local minimum of the error function
  - Is this good enough?
  - Can learn non-linear functions, but still very sensitive to input format

Neural Network Conclusions

- Many applications for pattern recognition tasks
- Very widely deployed AI technology
- Connection to biology still uncertain
- Results are hard to interpret (compare with decision trees)