

## Psychology 354

### Elements Of Connectionist Cognitive Science

#### Connectionist Philosophy Association Decision Representation and Limitations

### Cartesian Reaction

- The Cartesian philosophy that inspired the classical approach championed logicism and nativism
- "It is an established opinion among some men that there are in the understanding certain innate principles; some primary notions, characters, as it were, stamped upon the mind of man, which the soul receives in its very first being, and brings into the world with it" (Locke, 1706, p. 17)
- Locke was one of the pioneers of empiricism, a reaction against Cartesian philosophy



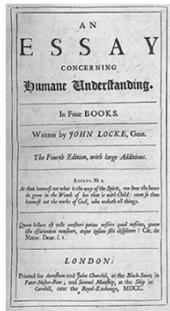
Descartes



John Locke

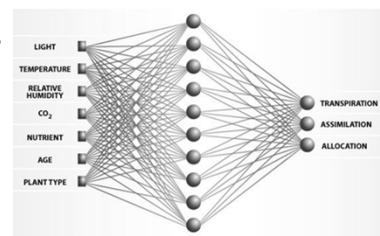
### Nurture Over Nature

- Locke's goal was to replace Cartesian nativism with empiricism, the view that the source of ideas was experience
- Locke aimed to show "how men, barely by the use of their natural faculties, may attain to all of the knowledge they have without the help of any innate impressions" (1706, p. 17).
- Locke argued for experience over innateness, for nurture over nature
- This same view motivates connectionist cognitive science



### Connectionist Elements

- Artificial neural networks are "neuronally inspired" networks of simple processors that operate in parallel
- There are several core properties of such networks that are used to contrast connectionist models, inspired by the brain, with classical models, inspired by the digital computer
- Let us consider some of these elements



## First Building Block: Association

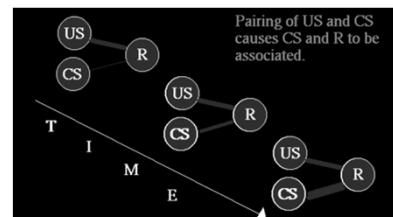
- One of the key building blocks for a connectionist system is a method for storing associations between input and output pattern
- Based on association psychology's associations between ideas
- Laws of association
  - Contiguity or Habit
  - Similarity
  - Contrast



William James

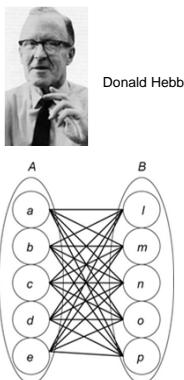
## Classical Conditioning

- Classical conditioning is an example of the associative law of contiguity at work
- The conditioned stimulus is not associated with the unconditioned response
- After repeated pairings with an unconditioned stimulus (contiguity), the conditioned stimulus becomes associated with the desired response



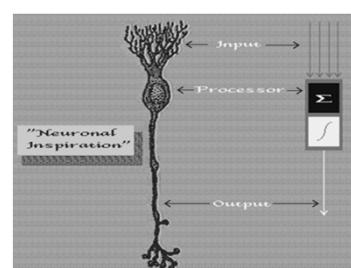
## Connectionist Association

- Connectionists use the law of contiguity, called the Hebb rule, to create simple content addressable memory systems
- "When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes place in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased" (Hebb, 1949)
- Such memory systems make interesting, content-based, human-like errors



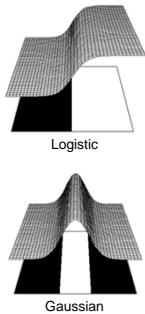
## Element 2: Decision

- Neurons do not merely associate with one another
- They make decisions – all or none responses – about their incoming signals
- A second connectionist element is a nonlinear activation function
- This function serves to make a decision about incoming signals



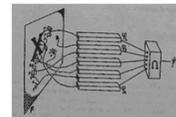
## Example Functions

- Modern artificial neural networks use nonlinear activation functions
- The most popular is the logistic equation, which defines the sigmoid approximation to an all-or-none function that defines an integration device
- An important one developed in my own lab uses a Gaussian equation to create a value unit
- Both functions "squash" input signals that range from negative to positive infinity into activation values that range between 0 and 1



## Perceptron As Example

- A perceptron, invented by Rosenblatt in the 1950s, is a distributed memory system whose output unit or units use a nonlinear activation function
- A perceptron can be trained to make all-or-none classifications of input patterns



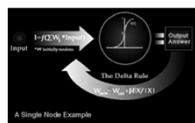
Perceptron



Frank  
Rosenblatt

## Learning From Mistakes

- Perceptrons, and the modern networks that followed, learn from their mistakes
- The difference between desired and actual output defines error
- Error is used to modify connection weights in such a way that error will be smaller when a trained pattern is presented again



## Musical Perceptron

- Task: learn a jazz chord progression (II-V-I)
- Perceptron's 12 input and 12 output units represent piano keys
- How?: Present current chord – network activates notes of next chord to play – train until correct sequence has been learned

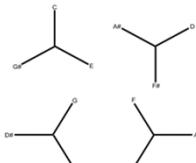
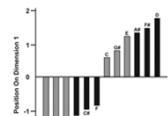
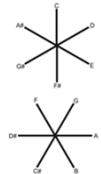
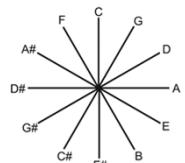


Key	II	V	I
C	Dm7	G7	Cmaj7
A#	Cm7	E7	Amaj7
G#	Dm7	D7	Gmaj7
F#	Gm7	C7	F#maj7
E	F#m7	B7	Emaj7
D	E7	A7	Dmaj7
C	Dm7	G7	Cmaj7

Table 4-2. A progression of II-V-I progressions, descending from the key of C major. The chords in each row are played in sequence, and after playing one row, the next row is played.

## Strange Circles

- How does a perceptron represent the jazz progression?
- An analysis of its connection weights reveals that it uses strange circles – circles of major seconds and major thirds, not the circle of fifths
- This shows how neural networks can provide representational insights



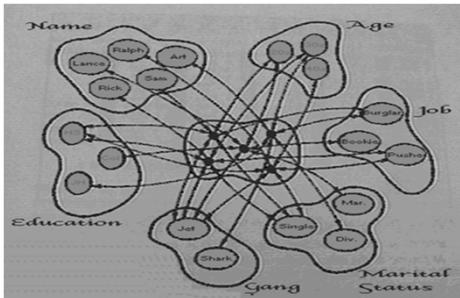
## Distributed Advantages: Jets vs Sharks

- Consider the following information, which a Classical model would encode as a serial set of records
- Connectionist researchers argue for a different kind of representation, a distributed representation, which offers certain advantages



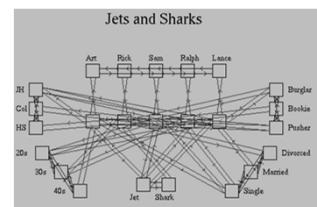
NAME	GANG	AGE	SCHOOLING	MARITAL	JOB
Gene	Jets	20s	College	Single	Pusher
Ralph	Jets	30s	JH	Single	Pusher
Phil	Sharks	30s	College	Married	Pusher
Ike	Sharks	30s	JH	Single	Bookie

## Jets/Sharks Network



## Jets/Sharks Processing

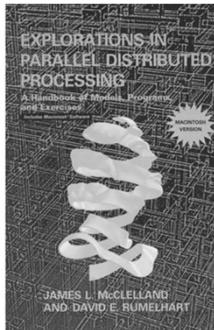
- Nodes organized in pools
- Inhibitory weights within pool
- Excitatory weights between pools
- Individual units point to property units
- Signals pass through network, and IAC network finds a stable pattern of processor activity



## Connectionist Advantages

Let's use a computer simulation to demonstrate the following advantages of distributed representations:

- Retrieve information from partial description
- Correct information from incorrect input
- Retrieve information not explicitly encoded
- Retrieve information after network damage



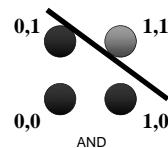
## Perceptron Limitations

- In their book *Perceptrons*, Minsky and Papert proved that there were limits to what perceptrons could learn because they can only make a single cut to a pattern space
- Traditional perceptrons can only make distinctions that are linearly separable, like AND, and cannot deal with linearly nonseparable problems, like XOR

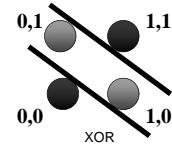


Marvin Minsky

Seymour Papert



AND



XOR

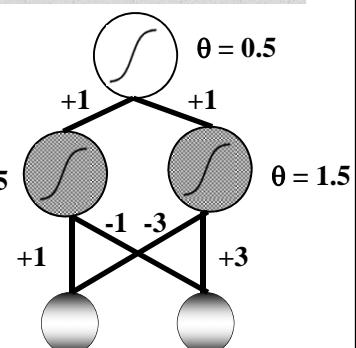
## Element 3: Intermediate Layers

- The solution to this problem is to add layers of intermediate processors, called hidden units
- The result is the multilayer perceptron
- Hidden units also use nonlinear activation functions
- This permits a sequence of complex decisions to be made by an artificial neural network
- This required new error-driven learning rules to be discovered, which only happened around the mid 1980s



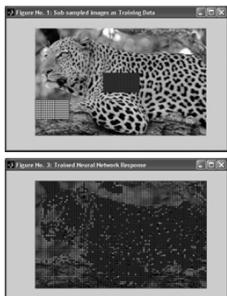
## Layers and Linear Nonseparability

- XOR is not a linearly separable problem
- This is because more than 1 cut is required in its pattern space
- A multilayered perceptron, though,  $\theta = 0.5$  is quite capable of learning how to solve this problem
- The hidden units transform the space into one that the output unit can easily carve



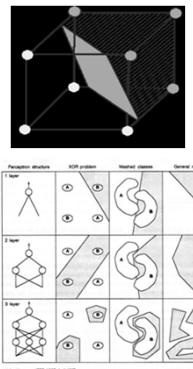
## What Can Networks Do?

- What kinds of tasks can modern networks perform?
- Are these networks powerful enough to be of interest to cognitive scientists?
- In the example on the right, the red dots in the lower image indicate the parts of the picture that a network is classifying as being "leopard" after training on the sampled pixels indicated in the upper image



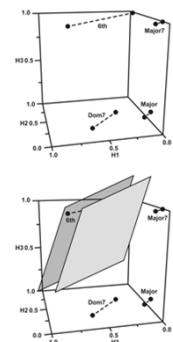
## Pattern Recognition

- Networks are frequently used to classify patterns
- They carve a pattern space into decision regions
- Patterns are classified according to these decision regions



## Example: Chord Classification

- An example of this network ability is the multilayer network described in the text that classifies types of musical chords
- The hidden units simplify the pattern space permitting it to be easily carved



## How Powerful?

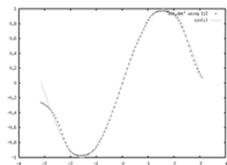
Lippmann (1987) has shown that a network with only two layers of hidden units can be an *arbitrary pattern classifier*. “No more than three layers [of connections] are required in perceptron-like feedforward nets” (p. 16).



Richard P. Lippmann

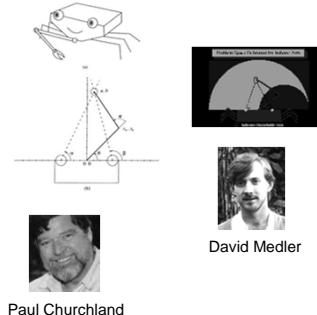
## Function Approximation

- A second task for a network is to estimate the value of a function on the basis of a set of predictor variables
- This is like pattern classification, but is continuous (and not discrete)
- In the example on the right, a network is given an x-value as input, and it outputs an estimate of the value of  $\sin(x)$  (red dots)



## An Example

- Churchland's 2D crab must convert two eye angles into two joint angles in order to grasp a sighted target
- David Medler trained a committee of networks to accurately perform this function mapping



Paul Churchland

David Medler

## How Powerful?

It has been proven that networks can be *universal function approximators*. “If we have the right connections from the input units to a large enough set of hidden units, we can always find a representation that will form any mapping from input to output” (Rumelhart, Hinton, & Williams, 1986).



David Rumelhart



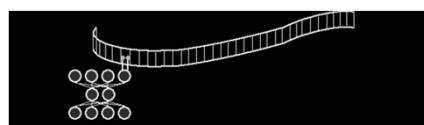
Geoff Hinton



Ronald J. Williams

## Networks Vs UTMs

- We have already seen that cognitive scientists are interested in very powerful computational systems (i.e., UTMs)
- Are PDP networks as powerful as a UTM?
- If not, then we would not be interested in them!



## An Old Proof

- McCulloch and Pitts proved that one could build a UTM tapehead from a network.
- “To psychology, however defined, specification of the net would contribute all that could be achieved in that field” (1943).



Warren McCulloch



Walter Pitts

## Connectionist Advantages

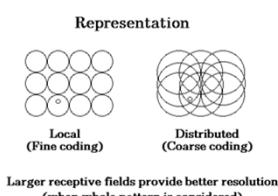
- “The neural network revolution has happened. We are living in the aftermath” (Hanson & Olson, 1991, p. 332)
- Why have artificial neural networks become so popular?
- What kinds of contributions can these networks make to cognitive science?
- What kinds of problems do they face?



Stephen José Hanson

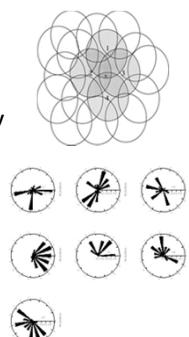
## Example: Coarse Coding

- An important distributed representation is called coarse coding
- Coarse coding requires that a property be encoded by a set of detectors
- Usually the detectors will have overlapping sensitivities
- Many examples of this type of coding are found in the human visual system
  - Hyperacuity
  - Colour perception
- Artificial neural networks often develop coarse codes – distributed representations – to mediate stimulus-response regularities



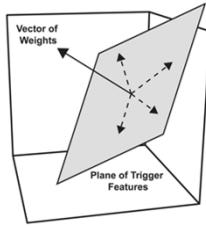
## Advantages of Distributed Representations

- What are the advantages of coarse coding or distributed representations?
- They increase the accuracy of inaccurate components
- They permit networks to generalize to new instances
- They permit networks to degrade gracefully
- They permit networks to be damage resistant
- They result in proposals for new kinds of representations for solving cognitive information processing problems
- Even critics of connectionism admit that “the study of connectionist machines has led to a number of striking and unanticipated findings; it’s surprising how much computing can be done with a uniform network of simple interconnected elements” (Fodor & Pylyshyn, 1988, p. 6).



## Advantages Require Analysis

- To find such advantageous properties one must spend a great deal of time interpreting the internal structure of networks
- What kinds of representations do networks discover?
- What kinds of algorithms mediate input/output relationships?
- “The major lesson of neural network research, I believe, has been to thus expand our vision of the ways a physical system like the brain might encode and exploit information and knowledge” (Clark, 1997. P. 58).
- There is no more room for “Gee Whiz Connectionism”
- Network interpretation is a key theme of the chapter!



## Some Classical Problems

- Poor for ill-posed problems
- Not damage resistant
- Does not degrade gracefully
- Serial -- therefore slow
- Not biologically plausible!
- Connectionists argue that these problems are solved when you abandon the notion of the digital computer, and instead turn to the kind of information processing done by the brain



## Classical vs Connectionism

	Classical Cognitive Science	Connectionist Cognitive Science
Core Ideas	<ul style="list-style-type: none"> <li>• Mind as a physical symbol system</li> <li>• Mind as a digital computer</li> <li>• Mind as a rule-based system</li> <li>• Mind as creator and manipulator of models of the world</li> <li>• Mind as sense-think-act processing</li> </ul>	<ul style="list-style-type: none"> <li>• Mind as information processor, but not as a digital computer</li> <li>• Mind as a parallel computer</li> <li>• Mind as pattern recognizer</li> <li>• Mind as a statistical engine</li> <li>• Mind as biologically plausible mechanism</li> </ul>
Preferred Formalism	Symbolic logic	Nonlinear optimization
Tacit Assumption	Nativism, native system	Empiricism
Type of Processing	Symbol manipulation	Pattern recognition
Prototypical Architecture	Production system (Newell, 1973)	Multi-layer perceptron (Rumelhart, Hinton, & Williams, 1986)
Prototypical Domain	<ul style="list-style-type: none"> <li>• Language</li> <li>• Problem solving</li> </ul>	<ul style="list-style-type: none"> <li>• Perception</li> <li>• Information learning</li> <li>• Perceptual categorization</li> </ul>
Philosophical Roots	<ul style="list-style-type: none"> <li>• Hobbes</li> <li>• Descartes</li> <li>• Leibniz</li> <li>• Clark</li> </ul>	<ul style="list-style-type: none"> <li>• Aristotle</li> <li>• Locke</li> <li>• Hume</li> <li>• James</li> </ul>
Some Key Modern Theorists	<ul style="list-style-type: none"> <li>• Chomsky</li> <li>• Dennett</li> <li>• Fodor</li> <li>• Pylyshyn</li> </ul>	<ul style="list-style-type: none"> <li>• J.A. Anderson</li> <li>• Hinton</li> <li>• Seiden</li> <li>• McClelland</li> </ul>
Some Pioneering Works	<ul style="list-style-type: none"> <li>• Plans And The Structure Of Behavior (Miller, Galanter, &amp; Pribram, 1960)</li> <li>• Aspects Of The Theory Of Symbol (Anderson, 1983)</li> <li>• Human Problem Solving (Newell &amp; Simon, 1972)</li> </ul>	<ul style="list-style-type: none"> <li>• Principles Of Neurodynamics (Rosenblatt, 1959)</li> <li>• Parallel Models Of Associative Memory (Hawkins &amp; Anderson, 1989)</li> <li>• Parallel Distributed Processing (McClelland &amp; Rumelhart, 1981)</li> <li>• Parallel Distributed Processing (McClelland &amp; Rumelhart, 1986b)</li> </ul>