

Representation and Computation

Syntactic Models of Reasoning

- Treat reasoning as a process of operating on symbols purely in terms of the characteristics of the symbols themselves (syntax), not their content (semantics)
- Formalization project in logic and mathematics
- Provides a means for making reasons into causes
 - The rules of reasoning respect the norms of inference but do not concern themselves with content
 - A “syntactic engine” that emulates a “semantic engine”



Language of Thought: Mentalese


- Thinking requires a representational system capable of representing anything that might be thought
- This representation must support all the inferences a person might make
- Only syntactically structured languages have the power to support a broad range of inferences
 - Arriving at beliefs on the basis of perception
 - Acquiring concepts through hypothesis testing
 - Decision making through evaluating possible actions

Computation and Computable Functions

- Functions that are evaluable or decidable through the execution of rote procedures or algorithms
- Turing: decidable functions can be evaluated through a finite number of steps using a Turing Machine
- A Turing Machine (and a digital computer) are automated syntactic devices

- Plausible hypothesis: human minds are also automated computational systems

Production Systems



A mind is equipped with an internal memory and a set of rules which can add and delete items from memory (and respond to external inputs or create outputs)

Learning as Hypothesis Testing

- To learn one must explicitly represent a hypothesis and test it with evidence
 - Hypothesis: past-tense formed by adding *-ed* to a verb
 - *Feared* is the past tense of *fear*. Evidence confirms
 - *Goed* is the past tense of *go*. Evidence falsifies
- Chomsky: The range of possible hypotheses is too large for a child to test a sufficient range so as to narrow in on the correct one
 - Hence, must be born with an innate linguistic capacity that sets the range of possible hypotheses (grammars)
- Fodor: where could any hypotheses come from to be begin with?
 - The learner must already be able to state all the hypotheses to be tested



Meaning and the Language of Thought

- Insofar as the mind operates as a syntactic engine, it does not depend on the meanings of representations
 - Fodor advocates methodological solipsism
- But our mental representations do have content (are intentional)
 - Even if that doesn't matter for how the mind works
- How to connect mental representations to their content?
 - Content is what would produce the representation under ideal conditions
 - Content is what the representations were selected for in evolution or learning
 - Fodor: asymmetric dependence: cows cause "horse" only if horses do



HORSE




The Only Game in Town




- When Fodor published the *Language of Thought* (1975) he viewed himself as providing an exposition of the commitments of contemporary cognitive psychology
 - The only program that seemed to have any hope of explaining the richness of human behavior had to assume a language-like mental representation system and syntax-sensitive rules
 - Otherwise the mind would be as mysterious as Descartes had assumed
 - The only competitor, associationism (aka behaviorism) was dead—proven to be inadequate
- But do the dead stay dead?

A non-Symbolic Alternative: Artificial Neural Networks (Connectionism)



Biological Neurons



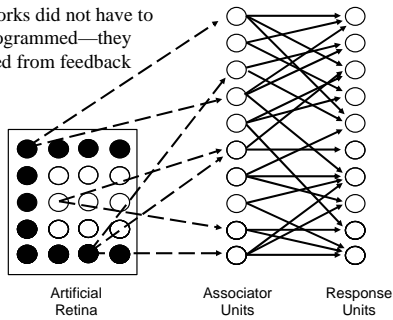
Artificial Neurons

McCulloch and Pitts (1943) saw how to build sentential logic networks out of artificial neurons: negation, and-gates, or-gates

Pitts and McCulloch (1947) saw the potential to model perception, etc. with less structured networks

Rosenblatt's (1962) Perceptrons

Networks did not have to be programmed—they learned from feedback






Minsky & Papert and the Demise of Perceptrons




Exclusive Or

A	B	A xor B
1	1	0
1	0	1
0	1	1
0	0	0

Failure of linear separability








Connectionism Returns



New learning rules: Delta Rule and Backpropagation



Solving XOR with Backpropogation





Nets learn to talk: NETtalk

Corpus presented to network

Started with random weights

Error backpropogated through network to adjust weights




Figure 1: Schematic drawing of the NETtalk network architecture. A window of letters in an English text is fed to an array of 203 input units. Information from these units is transformed by an intermediate layer of 80 "hidden" units to produce patterns of activity in 26 output units. The connections in the network are specified by a total of 18629 weight parameters (including a variable threshold for each unit).


Fodor and Pylyshyn's Response



- Thought, like language, is productive and systematic
 - It is always possible to think a new thought never before thought
 - Any given thought stands in relation to other possible thoughts such that if one can be thought, so can the other
- This can be explained (and only be explained) if representations have a combinatorial syntax and semantics
 - Productivity explained through recursion
 - Systematicity explained by compositional grammar
- Connectionism at best a story of how Language of Thought might be implemented in the brain

Connectionist Responses

1. Implementation might matter: Minds degrade gracefully, work with incomplete inputs, etc., and these are virtues of connectionist networks
2. Connectionist networks can implement a sufficient degree of systematicity and productivity without compositional syntax (Jordan's RAAM networks)
3. Connectionist networks can deal with natural language without implementing a compositional syntax: Elman's simple recurrent networks



Computation and Expertise

- Learning to follow rules is at best a step along the way to acquiring expertise in a domain
- Hubert Dreyfus: skilled performance involves “seeing”, not computing
 - The performance of a chess grandmaster drops very little if forced to make moves in seconds
- Perception remains a very hard problem for AI (and seemingly for human brains)
 - Can it be explained without assuming computation over representations?
