Rule-Based AI

Vinnie Monaco

Naval Postgraduate School

Outline

- Background
- Early attempts
- Expert systems
- Conclusions

Rules

- If X then Y
  - If it is a federal holiday, then there are no classes.
- All X are Y
  - Baseballs are round
- To do X, one must first do Y
  - Requesting leave requires submitting a form

Inference

Rules and facts

- If something is round, then it can roll away.
- A baseball is round.

Conclusion

- A baseball can roll away.

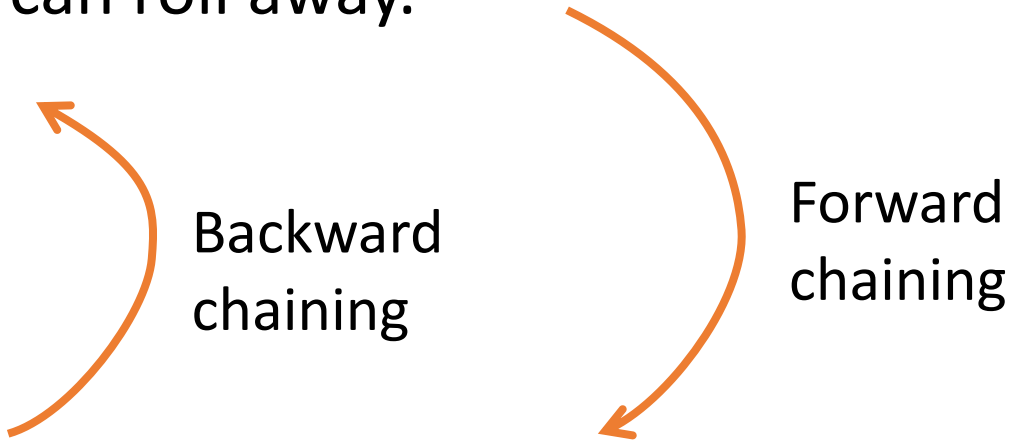

Components of a rule-based AI

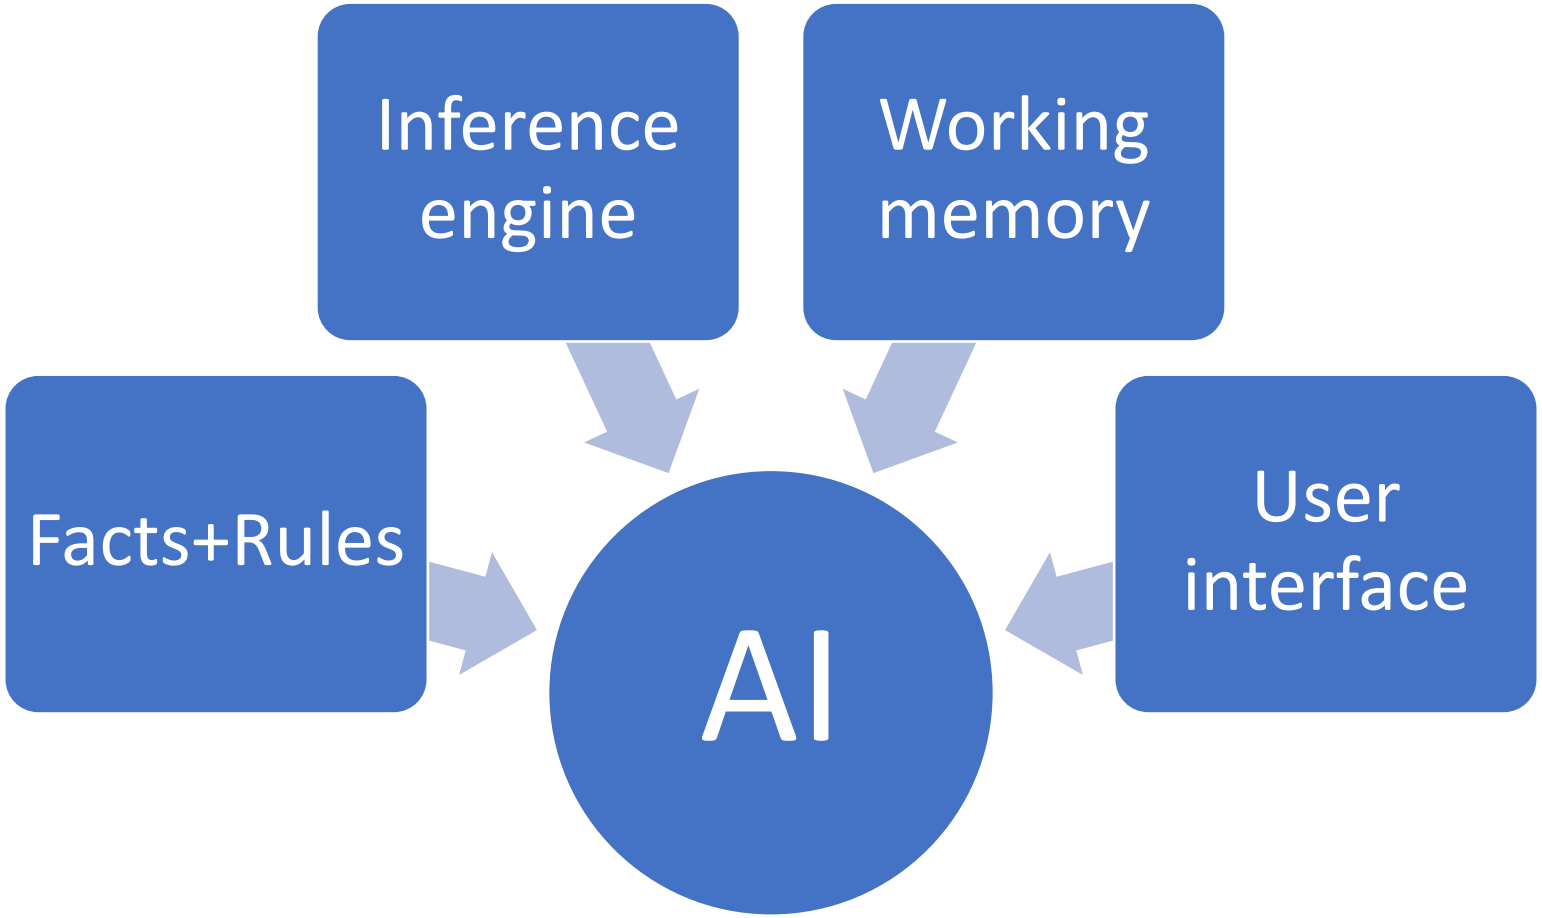

Tools of the trade

Prolog

```
roll_away(X) :- shape(X, round).
shape(baseball, round).
```

Knowledgebase

```
?- roll_away(baseball)
```

Goal/query

Early attempts

- 1959: Checkers
  - Arthur Samuel
- 1959: GPS
  - General Problem Solver
- 1961: SAINT
  - Symbolic Automatic Integrator
- 1962: ANALOGY
  - A is to B as C is to ?
- 1964: STUDENT
  - If the number of customers Tom gets is twice the square of 20% of the number of advertisements he runs, and the number of advertisements is 45, then what is the number of customers Tom gets?

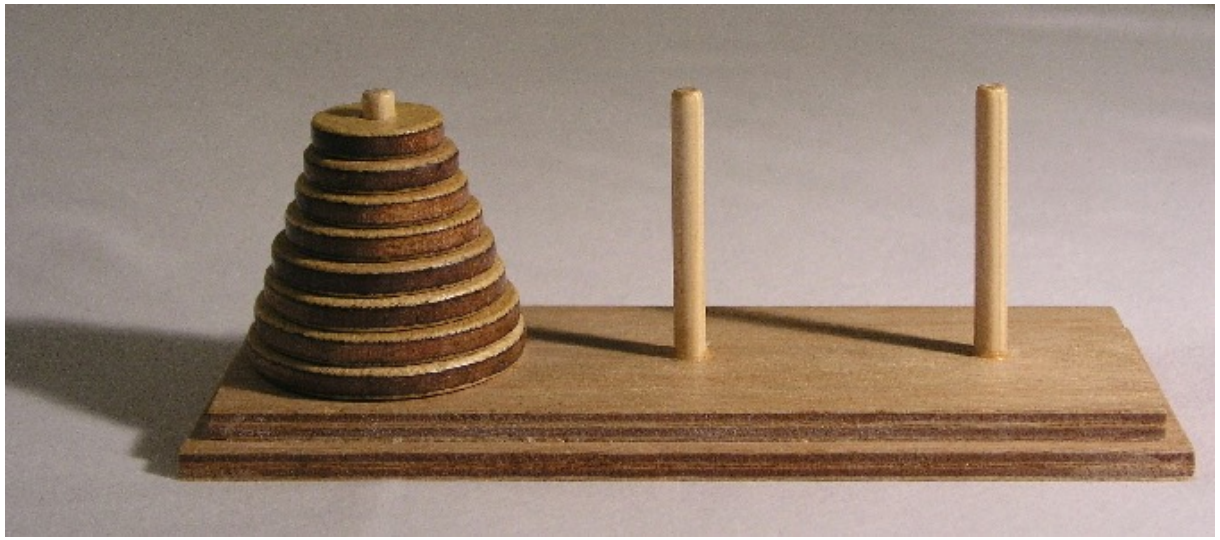

The rise of Good old-fashioned AI (GOFAI)

“at the end of the century [i.e., by 2000], the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted”

- Alan Turing, 1950

“within a generation ... the problem of creating 'artificial intelligence' will substantially be solved.”

- Marvin Minsky, 1967

“In from three to eight years we will have a machine with the general intelligence of an average human being.”

- Marvin Minsky, 1970

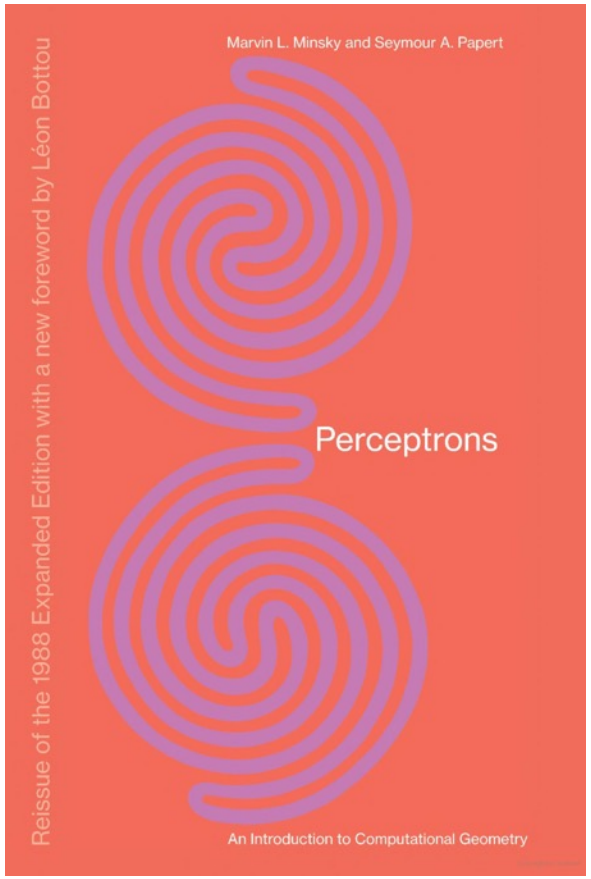

Minsky-Papert 1969

The 1st AI winter: 1974-1980

1973 “Controversy” debate following the Lighthill Report

Expert systems

Dendral

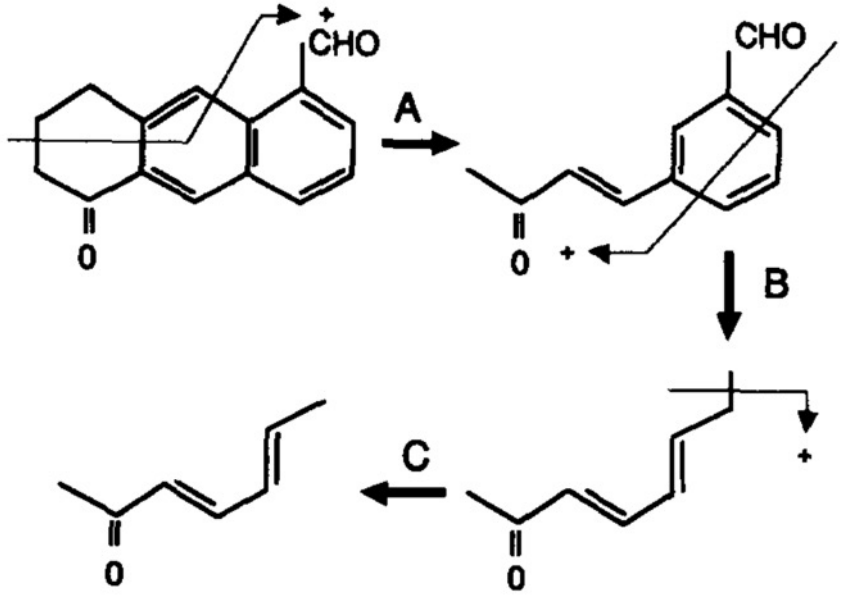

MYCIN

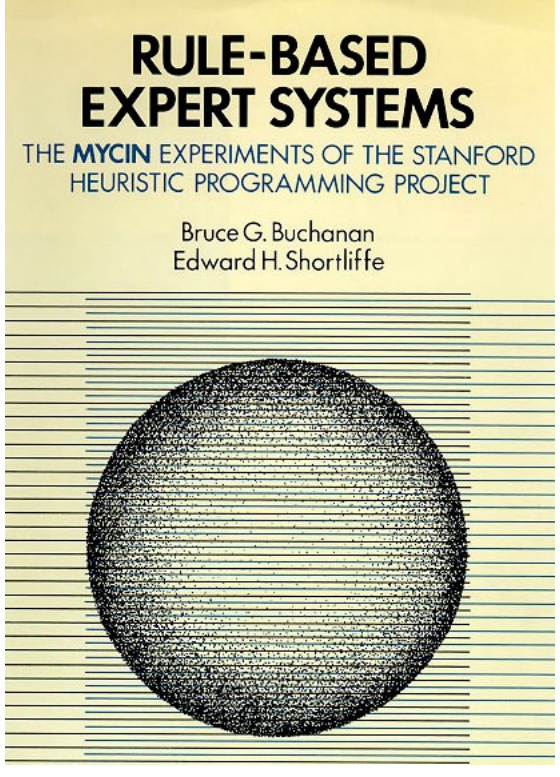

XCON

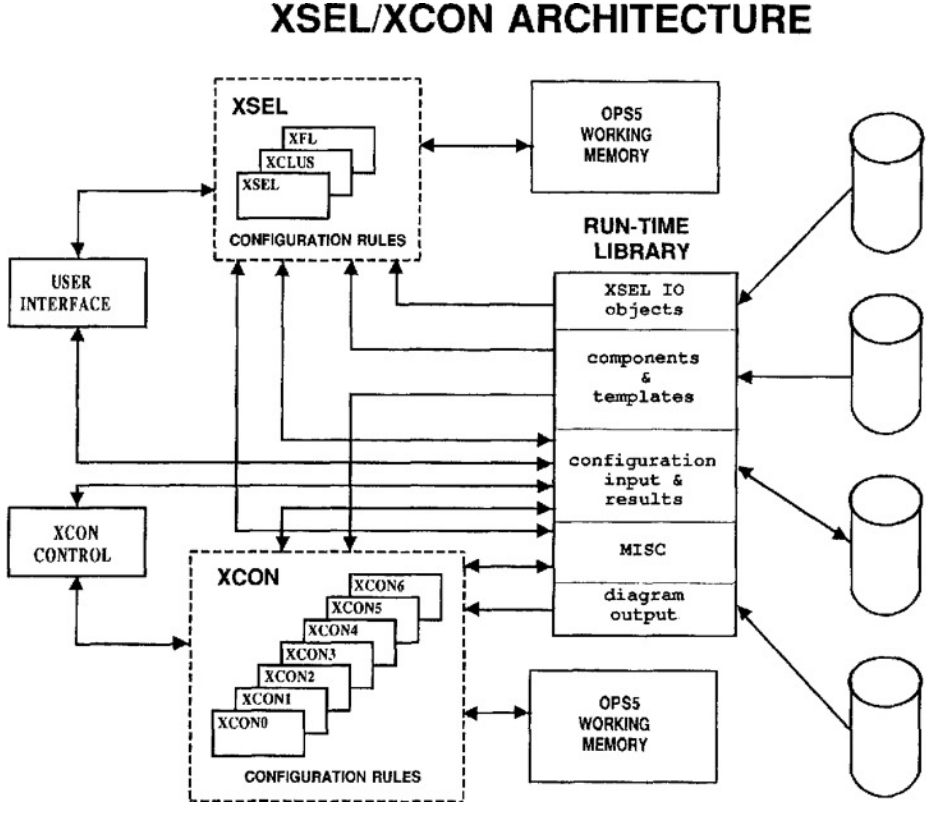

Japan's "Fifth Generation Computer Systems"

- The next step in computing
  - Vacuum tubes → transistors → integrated circuits → microprocessors → parallelism
- Began 1982, lasted 10 years
- $400 million effort
- Failure or ahead of it's time?

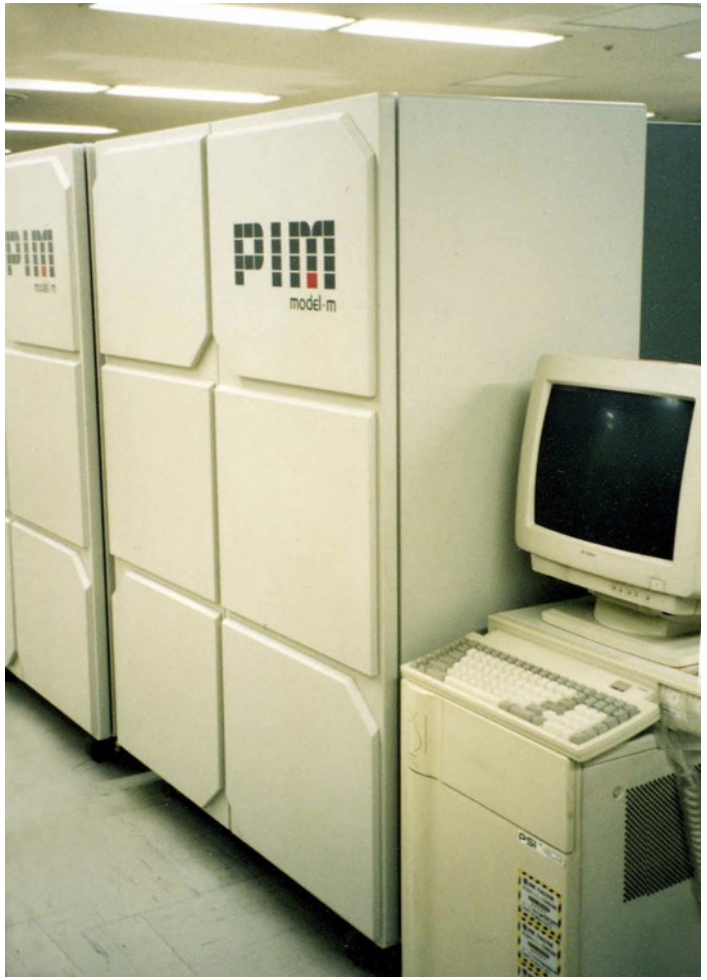

Parallel Interface Machine (PIM)

The 2nd AI Winter: 1987-1993

The New York Times

Setbacks for Artificial Intelligence

Companies Are Hurt By Poor Decisions

By ANDREW POLLACK

Special to The New York Times

SAN FRANCISCO, March 3 — A major retrenchment is occurring in the artificial intelligence industry, dashing the hopes of many companies that thought they would prosper by providing the technology to make computers "think."

Some of the setback stems from the failure of artificial intelligence to quickly live up to its promise of making machines that can understand English, recognize objects or reason like a human expert — to be used for such purposes as diagnosing machinery breakdowns or deciding whether to authorize a loan. Despite this, the technology is making slow but steady progress, and now is being subtly incorporated into more conventional computer programs.

Weak Business Moves

Far more of the artificial intelligence industry's problems result from poor business decisions by companies that were heavily weighted with technologists rather than business minds. Their main mistake was trying to use special computers for artificial intelligence; the machines were too expensive and did not mesh well with those used by potential customers.

"People believed their own hype," said S. Jerrold Kaplan, co-founder of one leading artificial intelligence company, Teknowledge Inc., and now

Trouble for Many of the Artificial Intelligence Companies

As the industry realigns, the companies that relied on special purpose machines are languishing.

Company / Headquarters	Description
EXPERT SYSTEM DEVELOPMENT TOOLS	
Teknowledge Palo Alto, Calif.	Four quarters of losses. 60 workers of 220 laid off. 1987 sales: $20 million.
Intellicorp Mountain View, Calif.	Six quarters of losses. 30 workers of 200 laid off. 1987 sales: $20 million.
Carnegie Group* Pittsburgh	Losses. 20-40 workers of 200 laid off. 1987 sales: $12 million.
Inference* Los Angeles	Losses. 20 workers of 130 laid off. 1987 sales: $12 million.
MACHINE MANUFACTURERS	
Symbolics Cambridge, Mass.	Continuing losses. Third round of layoffs last fall. Ousted chairman and founder. 1987 sales: $104 million.
Lisp Machine Andover, Mass.	Filed for bankruptcy last year. 1986 sales: $12 million.
Xerox Stamford, Conn.	Sluggish sales; recent reorganization of its artificial intelligence business.
Texas Instruments Dallas	Big push in artificial intelligence. Announced a chip containing Lisp to go into Macintosh.
EXPERT SYSTEM APPLICATIONS	
Syntelligence* Sunnyvale, Calif.	No layoffs. 1987 sales: $9 million.
Applied Expert Systems* Cambridge, Mass.	Layoffs. 1987 sales: $4 million.
Palladian* Cambridge, Mass.	Layoffs. Ousted chairman and founder. 1987 sales: $6 million.

New York Times, March 4, 1988

- Commercialization slowed
- LISP machines underperformed
- Fifth Generation project ended

Expert system limitations

- They cannot learn (easily)
- Restricted by the size of the knowledgebase
- Expensive to maintain
- **Lack understanding of what human expertise really is**
- Not possible to explicitly define some rules

A difficult task for rule-based AI

Recognize sunglasses if:

- pixel at 100,150 is pink
- pixel below is pink
- pixel to right is black
- …

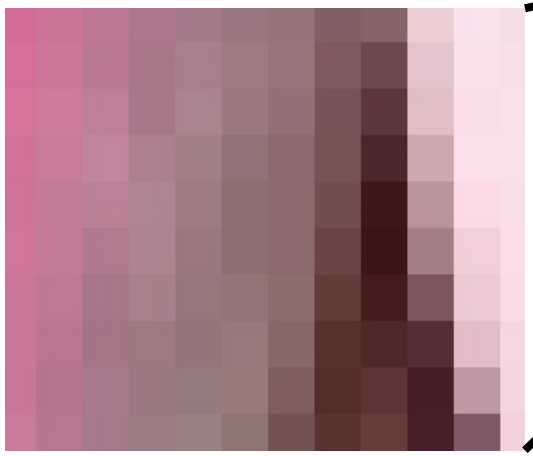

The context problem

Winograd Schema Challenge

The city councilmen refused the demonstrators a permit because they **feared** violence.

The city councilmen refused the demonstrators a permit because they **advocated** violence.

Ongoing effort: Cyc

- Started by Doug Lenat in 1984
- An attempt to solve the “common sense problem”
- 25 million common sense rules
  - 1000 persons-years to build
- Just recently started to commercialize

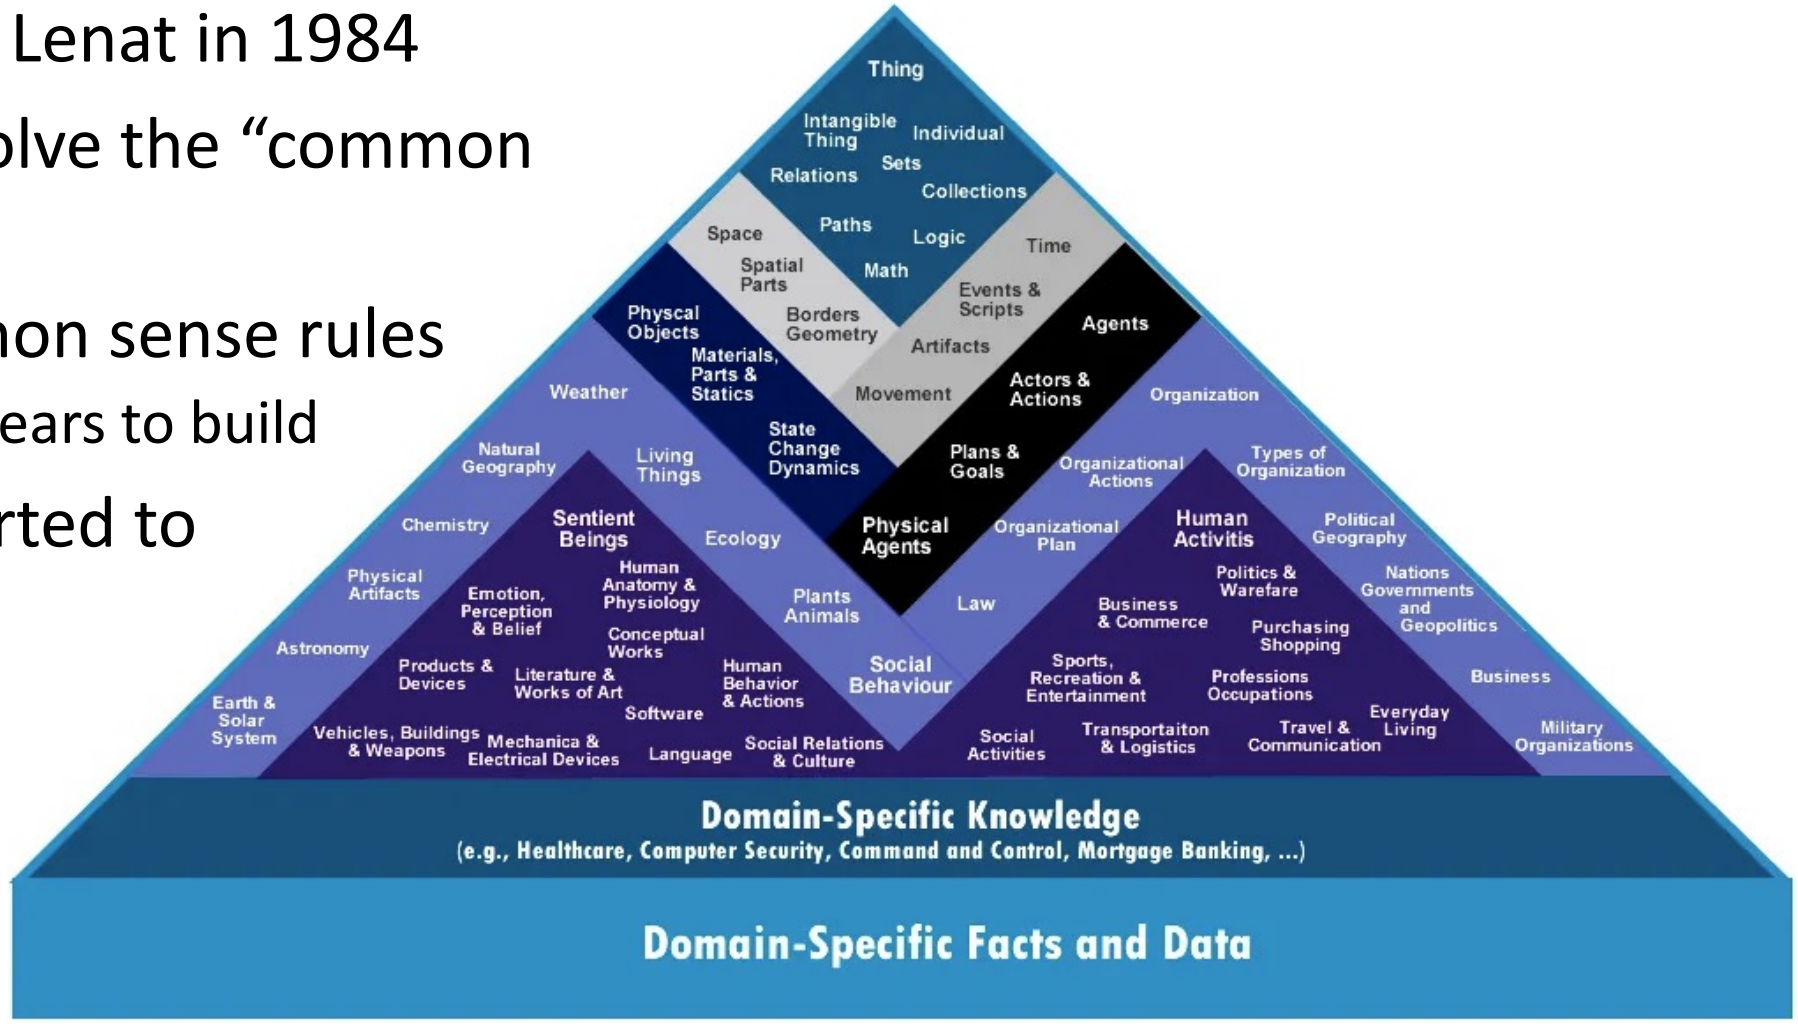

Rule-based AI vs neural networks

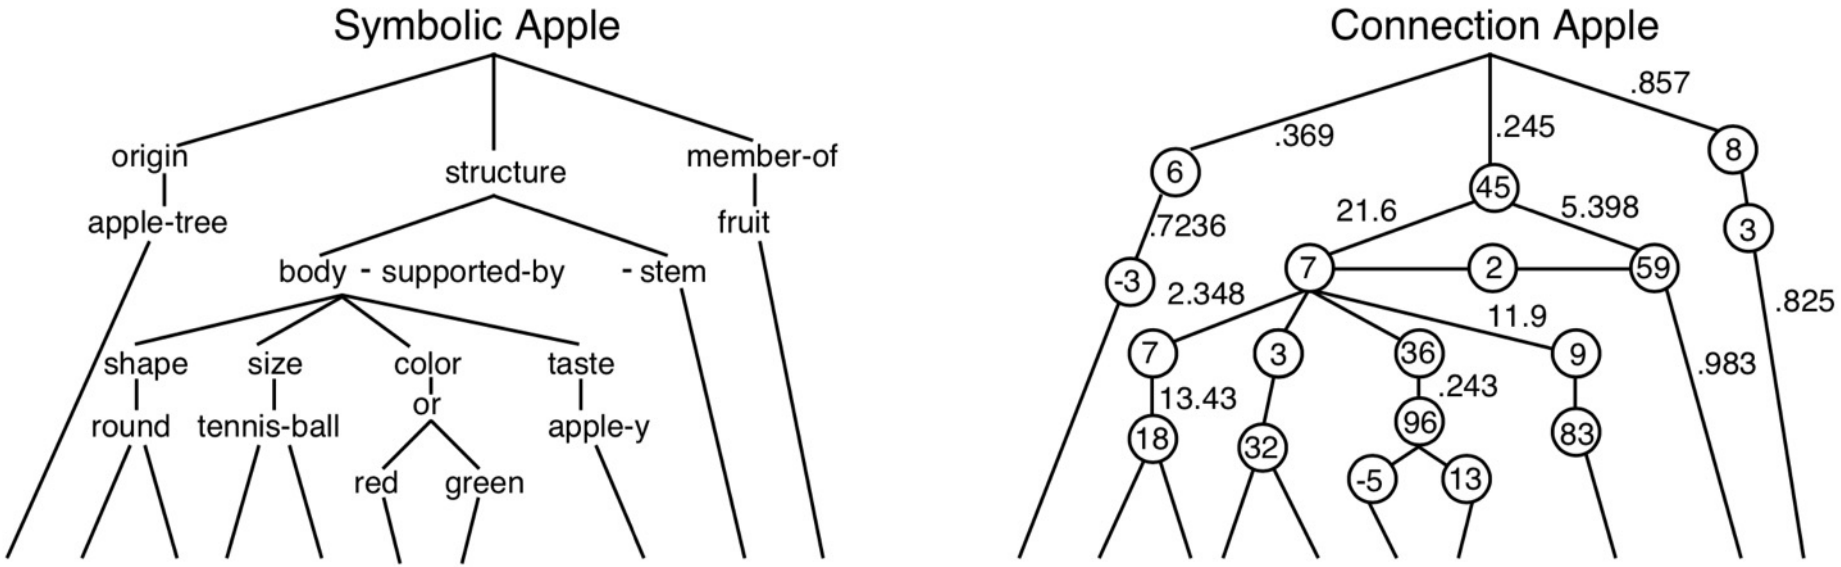

Advantages of rule-based AI

- Computationally inexpensive
- Explainable
- Well-suited for symbol manipulation and reasoning

Current trends

DNC: Differentiable Neural Computer, 2016

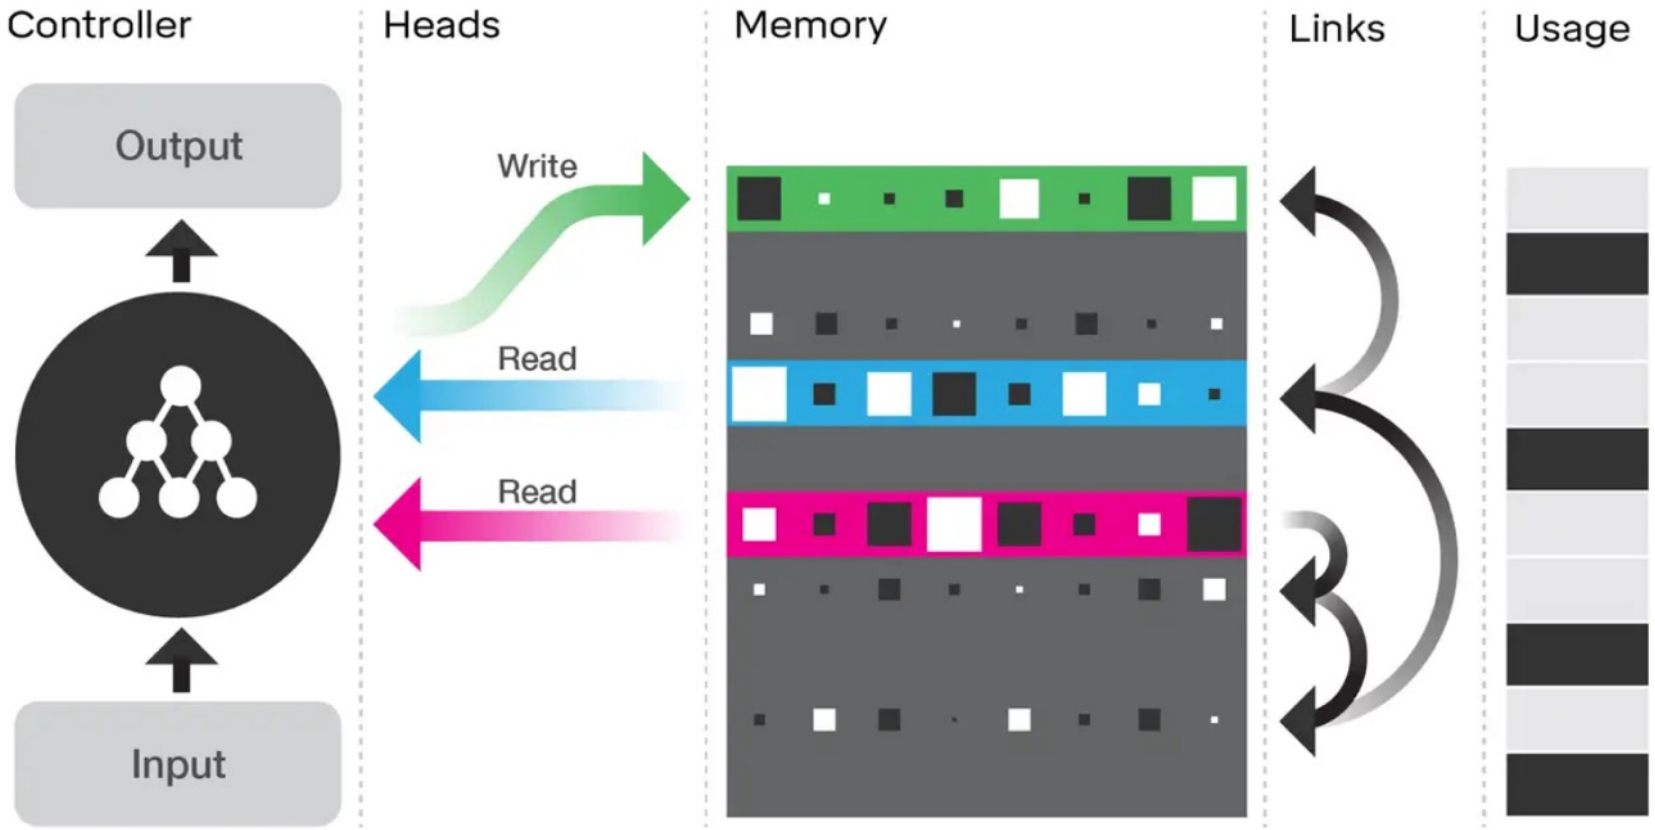

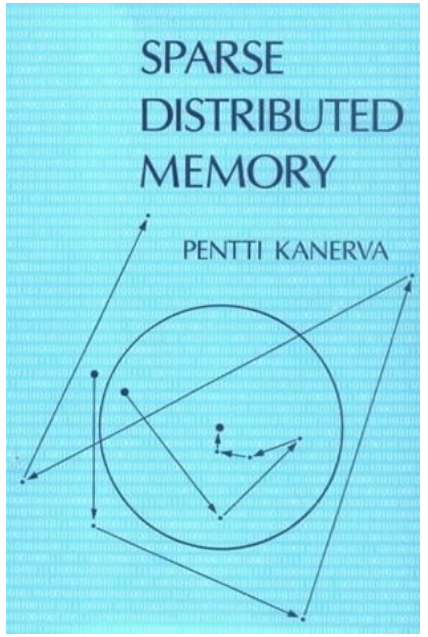

1988

Suggested reading

Why AI is Harder Than We Think

Melanie Mitchell
Santa Fe Institute
Santa Fe, NM, USA
mm@santafe.edu

Abstract

Since its beginning in the 1950s, the field of artificial intelligence has cycled several times between periods of optimistic predictions and massive investment ("AI spring") and periods of disappointment, loss of confidence, and reduced funding ("AI winter"). Even with today's seemingly fast pace of AI breakthroughs, the development of long-promised technologies such as self-driving cars, housekeeping robots, and conversational companions has turned out to be much harder than many people expected. One reason for these repeating cycles is our limited understanding of the nature and complexity of intelligence itself. In this paper I describe four fallacies in common assumptions made by AI researchers, which can lead to overconfident predictions about the field. I conclude by discussing the open questions spurred by these fallacies, including the age-old challenge of imbuing machines with humanlike common sense.

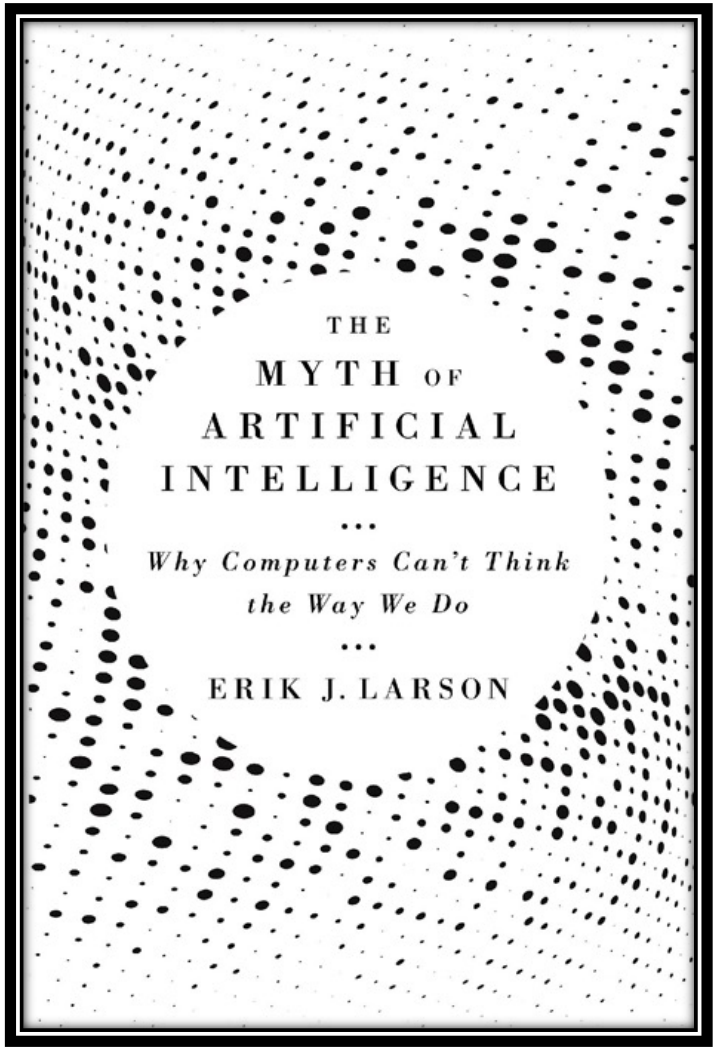

Questions?

vinnie.monaco@nps.edu
https://vmonaco.com