Artificial intelligence: scientific and engineering perspectives

Etienne Barnard

Overview

- AI: the vision
  - Scientific / philosophical
  - Engineering
- Approaches
  - Symbolic (represent & search)
  - Statistical (classifiers, neural networks &c)
- Applications
- Achievements & prospects

Scientific / philosophical goals of AI

- Understand cognitive capabilities of biological organisms
  - Brains, neurons, neurotransmitters, …
- Resolve mind/body dualism
  - Consciousness, determinism, mental vs physical worlds

Engineering goals of AI

- Develop tools that can mimic capabilities of “intelligent” entities (e.g. humans)
  - Sensory responsiveness
    - Speech, vision, …
  - Manipulation (control) of environment
    - Planning, actions, feedback
  - Complex data processing
    - Retrieval, filtering, creativity

Approaches to AI

- Symbolic AI:
  - Intelligence = manipulation of symbols
  - AI: symbolic representation of knowledge, search through representations
  - (Two categories: rationality vs. imitation)
- Statistical AI:
  - Intelligence = extraction of statistical regularities in data
  - AI: simulation of biological “information filters”

The symbolic approach (“classical AI”)

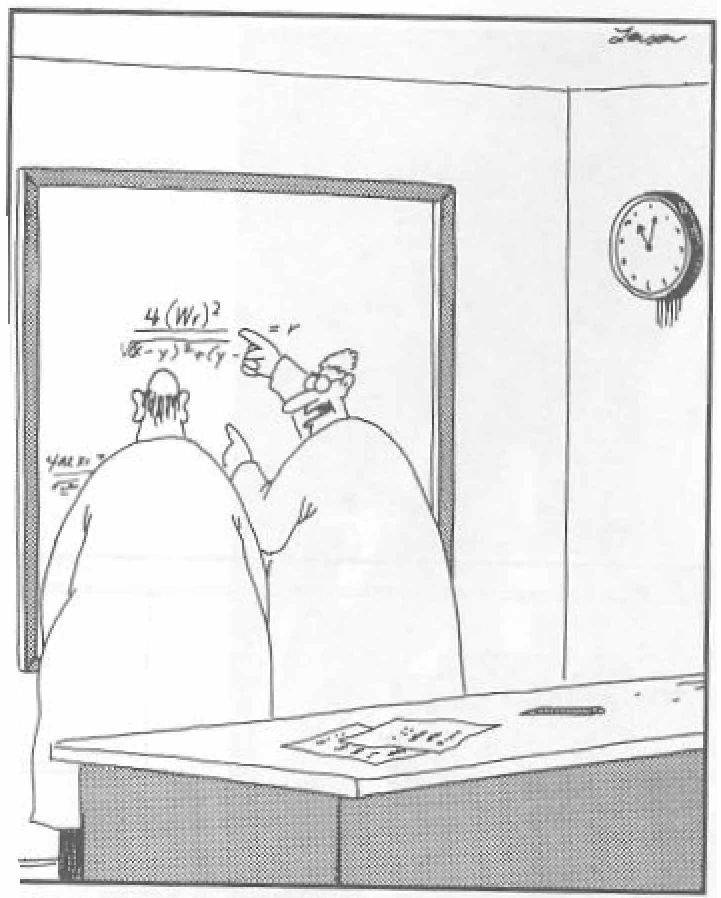

“Yes, yes, I *know* that, Sidney…*every*body knows *that!*…But look: Four wrongs *squared,* minus two wrongs to the fourth power, divided by this formula, *do* make a right.”

Knowledge representation

- Goal: develop data structures that make it easy to represent diverse, “real-world” knowledge
- Typical approaches
  - Logic (propositional, predicate, HO)
  - Semantic networks (nodes, named links, hierarchies)
- Major challenge is automation of knowledge capture (“knowledge engineering”)

Search

- Goal: find efficient ways to operate on data structures to extract conclusions
- Approaches:
  - "Uninformed": systematic enumeration of all possibilities *(typically exponentially expensive)*
  - "Informed": use problem-specific information to accelerate process
    - Exact solutions: A*- search
    - Approximate solutions: greedy search
- For most realistic problems, approximation required
    - Main challenge: design of "heuristic function"

The statistical approach

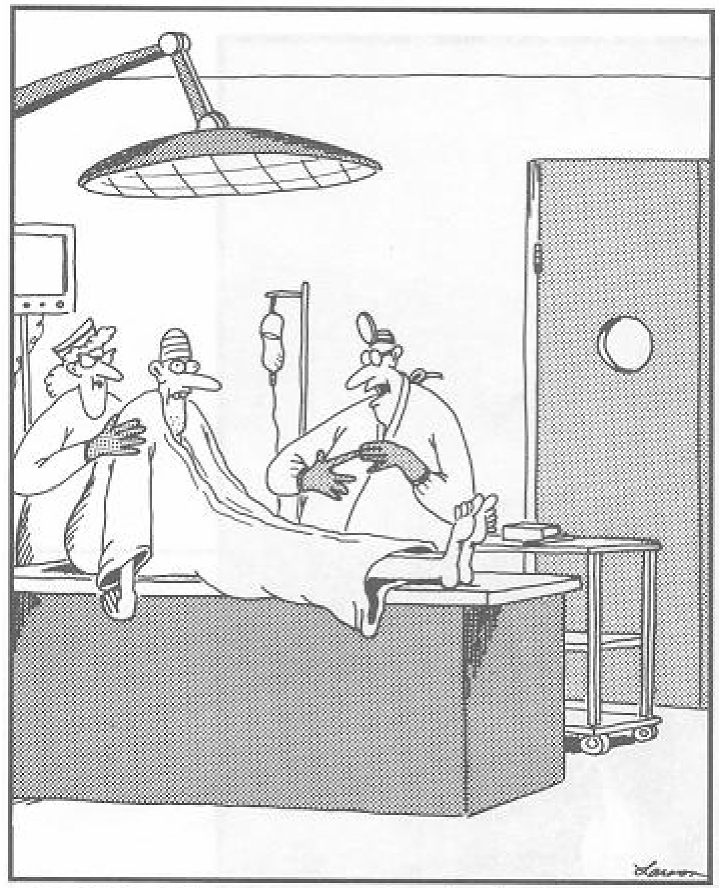

"OK, Mr. Dittmars, remember, that brain is only a temporary, so don't think too hard with it."

Classification / regression

- Goal: “learn” approximation of classification / function from examples
- Approach:
  - Select functional form (parameterization), training & evaluation data, measure of fit
  - Optimize parameters (conventional / novel approaches) – fit of training data
  - Evaluate on distinct data

Classification / regression (2)

- Categories of solutions (functional forms):
  - Linear
  - “Instance-based” (nearest-neighbor, kernel, support vector machines)
  - Neural networks (hierarchy of simple non-linear elements)
  - Decision trees
  - …

Density estimation

- Goal: estimate probability density of categories of events / tokens
  - Provably optimal basis for decision making
  - Provably exponentially “hard” under “realistic” circumstances
- Core algorithms similar to estimation of classifiers / regressors
- Key innovation is representation of independence
  - Bayesian networks and generalizations

Feature extraction

- Goal: represent relevant information numerically
- Criteria:
  - Compactness of single categories
  - Separability of different categories
- Examples:
  - Vision: intensities in appropriately chosen geometrical regions
  - Cell biology: identities of amino acids in a window
  - Web search: number, status of linked pages
- For many practical systems, discovery of suitable features is key to success

Solving problems with AI

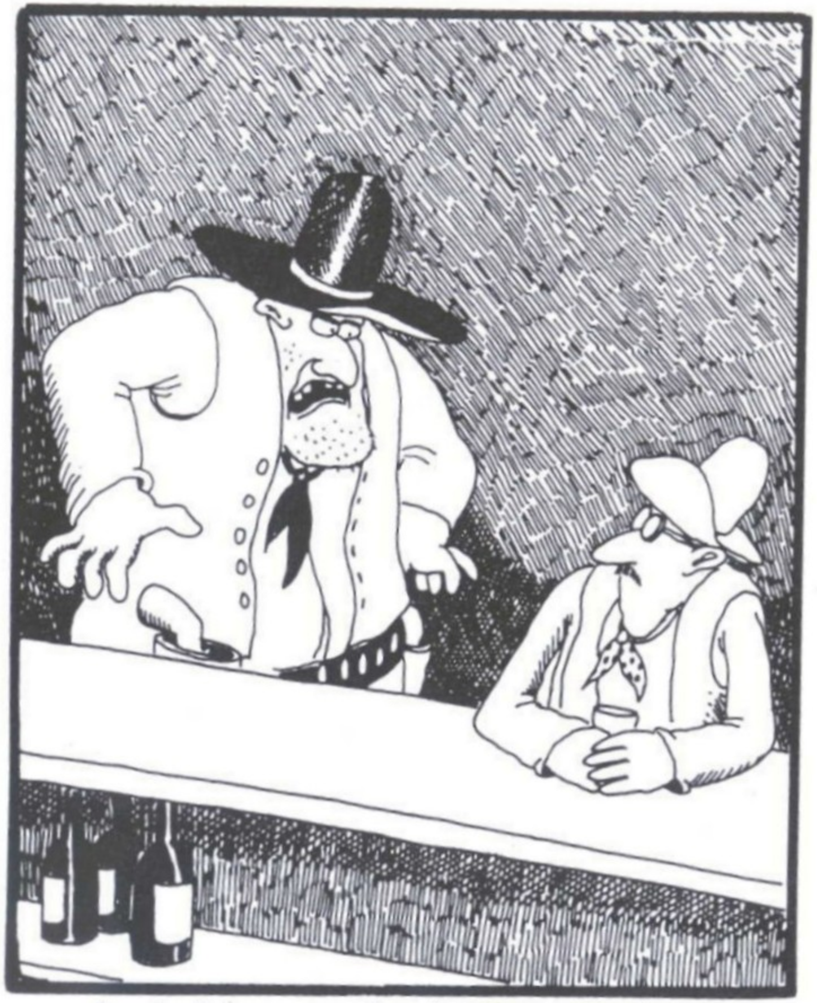

"I asked you a question, buddy . . . What's the square root of 5,248?"

Application: scheduling

- Goal: assign set of **resources** to sequence of **tasks**, subject to **constraints** in order to optimize **goal function**
- Examples: campaign planning, design of timetables
- Typical solution:
  - Represent as nodes and operators
  - Search through assignments with domain-independent (most-constrained, least constraining, …) + domain-dependent heuristics

Scheduling: state of the art

- Large problems (1000s of resources and tasks, 100s of constraints) regularly solved
  - Hubble scheduling of observations – 3 weeks $\rightarrow$ 10 minutes
  - Standard practice for logistics & travel companies
  - Military planning has “paid the bill” for AI

Application: speech recognition

- Goal: understand human speech
- Examples: control, dictation, telephone-based information retrieval
- Typical solution:
  - Density estimate of acoustics of speech sounds (phonemes)
  - Dictionary to represent pronunciations of words
  - Grammar or statistical model of expected utterances

Speech recognition: State of the art

- Desktop dictation
  - > 99% accuracy on unconstrained text
  - ±1 hour of speaker training
  - Widely used in specialized areas (e.g. legal / medical) or where typing is a problem
- Telephone-based information access
  - Speaker-independent, task specific
  - Answers $>10^6$ telephone calls per day
  - Major potential in developing world?

Looking back, looking ahead

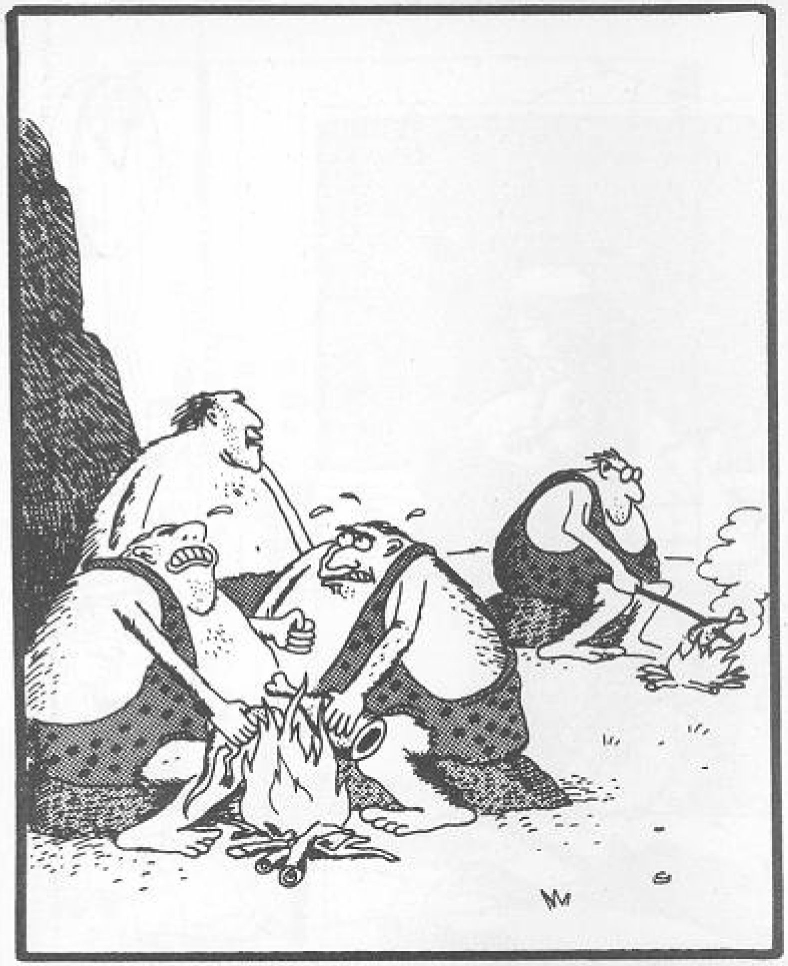

"Hey! Look what Zog do!"

A history of overselling

- “*Within ten years a digital computer will be the world's chess champion*”

A history of overselling

- “*Within ten years a digital computer will be the world's chess champion*”
  - – Newell, 1957

A history of overselling

- “*Within ten years a digital computer will be the world's chess champion*”
  – Newell, 1957
- “*There are now in the world machines that think, that learn and create*”

A history of overselling

- “*Within ten years a digital computer will be the world's chess champion*”
  - – Newell, 1957
- “*There are now in the world machines that think, that learn and create”*
  - – Simon & Newell, 1958

Achievements

- Grandmaster-level chess (and valued competitor in computer games)
- Standard part of logistics planning
- Telephone agent; security inspector; credit verifier
- Google!

The pessimistic view

- Computers will never match human intelligence because:
  - They are not “grounded” (*Searle*, *Dreyfus*)
  - They lack the appropriate quantum-mechanical machinery (*Penrose*)
  - Dualism was right after all (*Chalmers*)

The optimistic view (1)

- Tasks that have been considered *intelligent*:
  - Solving algebraic problems
  - Playing chess
  - Understanding speech
- AI keeps moving the frontier!
- Incremental approach will take us past *X*, for any *X*

The optimistic view (2)

- Intuitively definable tasks that we cannot accomplish artificially:
  - Turing test
  - Survival of mosquito
- Fundamental gap in understanding
  - *[?? Induction / learning how to learn ??]*

What next?

- Continued uptake of “AI” methods & tools in computing
  - Search, relational methods, neural networks $\rightarrow$ Bayesian networks, agent-based processing
- Growth of AI as more tasks must be automated
  - Internet bots, ubiquitous networks, security networks, 24 x 7 customer service, …
- A return to fundamental research

References:

- “*The age of intelligent machines*”, R Kurzweil, MIT Press, 1990 (**Popular**)
- “*Artificial Intelligence: A Modern Approach*”, 2nd ed., S Russell & P Norvig, Prentice Hall, 2003 **(Overview)**
- *IEEE Trans. on Neural Networks*, *Machine Learning*, *IEEE Trans. on PAMI*, *Artificial intelligence* **(Technical)**