

Artificial Intelligence: An Introduction

Mark Maloof

Department of Computer Science
Georgetown University
Washington, DC 20057-1232

<http://www.cs.georgetown.edu/~malooof>

August 30, 2017

What is Artificial Intelligence?

Summon the demon!

What is Artificial Intelligence?

- ▶ augmented intelligence
- ▶ business intelligence
- ▶ cognitive analytics
- ▶ cognitive computing
- ▶ cognitive technologies
- ▶ expert systems
- ▶ knowledge-based systems
- ▶ intelligent agents
- ▶ intelligent systems

McCarthy et al., 1955

- ▶ “The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.”

Haugeland, 1985

- ▶ “The exciting new effort to make computers think...machines with minds, in the full and literal sense.”

Charniak and McDermott, 1985

- ▶ “...the study of mental faculties through the use of computational models.”

Rich and Knight, 1992, 2009

- ▶ “The study of how to make computers do things at which, at the moment, people are better.”

- ▶ “Artificial intelligence, broadly (and somewhat circularly) defined, is concerned with intelligent behavior in artifacts. Intelligent behavior, in turn, involves perception, reasoning, learning, communicating, and acting in complex environments.”

Russell and Norvig's Four Approaches


1. Think like a human
2. Act like a human
3. Think rationally
4. Act rationally

Think Like A Human

- ▶ “...machines with minds, in the full and literal sense”
- ▶ Put simply, program computers to do what the brain does
- ▶ How do humans think?
- ▶ What is thinking, intelligence, consciousness?
- ▶ If we knew, can computers do it, think like humans?
- ▶ Does the substrate matter, silicon versus meat?
- ▶ Computers and brains have completely different architectures
- ▶ Is the brain carrying out computation?
- ▶ If not, then what is it?
- ▶ Can we know ourselves well enough to produce intelligent computers?

Act Like A Human

Turing Test



Obligatory xkcd Comic



The Brilliance of the Turing Test

- ▶ Sidesteps difficult questions:
 - ▶ What is intelligence?
 - ▶ What is thinking?
 - ▶ What is consciousness?
- ▶ If humans can't tell the difference between human intelligence and artificial intelligence, then that's it
- ▶ Proposed in 1950, Turing's Imitation Game is still relevant

Think Rationally

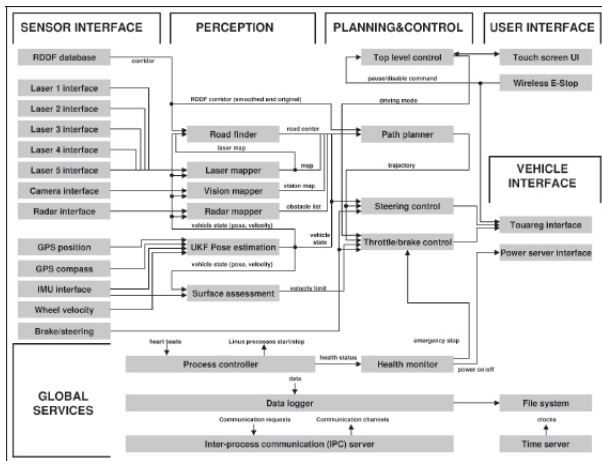
- ▶ Think rationally? Think logic!
- ▶ Put simply, write computer programs that carry out logical reasoning
 - ▶ Logic: propositional, first-order, modal, temporal, ...
 - ▶ Reasoning: deduction, induction, abduction, ...
- ▶ Possible problem: Humans don't really think logically
- ▶ Do we care? Strong versus weak AI
- ▶ One problem: often difficult to establish the truth or falsity of premises
- ▶ Another: conclusions aren't strictly true or false

Act Rationally


- ▶ Act rationally? Think probability and decision theory!
- ▶ “A rational agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome” (Russell and Norvig, 2010, p. 4)
- ▶ `<jab>` “when there is uncertainty” `</jab>`
- ▶ When *isn't* there uncertainty?
- ▶ Predominant approach to AI (for now)

Video: The Great Robot Race






Stanley




Source: Thrun (2010, Figure 7)

Stanley



Source: Thrun (2010, Figure 9a)



What is This Course About?

- ▶ Summoning the demon
- ▶ In some sense, it's an advanced course in data structures and algorithms
- ▶ You know how to store and manipulate traditional data
 - ▶ How do you represent causal relationships between events?
 - ▶ How do you reason about uncertain events?
 - ▶ How do you reason deductively?
 - ▶ How do you learn deductive rules?
- ▶ You know a lot about algorithms with polynomial running times
 - ▶ How do you deal with problems that require algorithms with exponential running times?

Learning Goals

- ▶ By the end of the semester, you'll be able to:
 - ▶ explain the main philosophical stances of artificial intelligence
 - ▶ describe, compare, and contrast the main formalisms used in artificial intelligence
 - ▶ describe, compare, and contrast algorithms used in artificial intelligence for reasoning and for learning
 - ▶ describe and critique thought experiments that support and refute strong artificial intelligence
 - ▶ devise and implement solutions to tasks requiring symbolic computation using the Lisp programming language
 - ▶ devise and implement solutions for computational tasks using Lisp as a functional programming language

Then There are the Big Questions

- ▶ Is the brain a Turing machine?
- ▶ Can machines think?
- ▶ Can Turing machines be minds?
- ▶ Can digital computers be conscious?
- ▶ Can we understand ourselves well enough to build intelligent machines?
- ▶ Should intelligent robots have rights?
- ▶ Are we summoning the demon?

Course Plan and Logistics

- ▶ Class Web page:
 - ▶ <http://www.cs.georgetown.edu/~malooof>
 - ▶ We'll use Bb for discussion
 - ▶ We'll use Autolab for project submission
- ▶ Five projects: intro to Lisp, search, deduction, uncertain reasoning, learning
- ▶ Midterm and final exams
- ▶ Office hours: TR 12:00–1:30 PM (or by appointment)
- ▶ Subscribe and use Bb for discussion about lectures and projects
 - ▶ Post using English, math, algorithms, or Lisp
 - ▶ No project solutions or partial solutions unless I post it
 - ▶ Posts can be anonymous, and they're not moderated
- ▶ Use e-mail for personal matters

Academic Integrity

- ▶ First and foremost, Georgetown students do honest work
- ▶ A small percentage gets into trouble when they get too busy—it's always P4
- ▶ Never submit someone else's work as your own without attribution
- ▶ You can get help from your fellow students:
 - ▶ Can I get your notes from yesterday's class?
 - ▶ What's that command that starts lisp?
 - ▶ How do you pass a comparator into the member function?
 - ▶ What's the name of that stupid machine in the cloud?
 - ▶ What does this error mean?
- ▶ But you can't get help with graded assignments:
 - ▶ How are you representing nodes for P2?
 - ▶ How did you implement that insert function?
 - ▶ Can I see your prove function?
- ▶ Use class, office hours, and the discussion board for these types of questions
- ▶ Details on the Web site

Why Lisp?

- ▶ Great for symbolic computation
- ▶ Great for functional programming
- ▶ Learning Lisp will make you a better programmer in other languages

Why Lisp?

- ▶ Great for symbolic computation
- ▶ Great for functional programming
- ▶ Learning Lisp will make you a better programmer in other languages
- ▶ Phillip Greenspun's Tenth Rule of Programming: "Any sufficiently complicated C or Fortran program contains an ad hoc, informally-specified bug-ridden slow implementation of half of Common Lisp."
- ▶ I examined Clojure but...

Symbolic Computation

- ▶ Differentiation:

- ▶ General form: $a^d \Rightarrow da^{d-1}$

- ▶ Application: $n^2 \Rightarrow 2n$

- ▶ DeMorgan's Law:

- ▶ General form: $\neg(\phi \wedge \psi) \Rightarrow \neg\phi \vee \neg\psi$

- ▶ Application: $\neg(\neg p \wedge q \rightarrow r) \Rightarrow \neg\neg p \vee \neg(q \rightarrow r)$

Foundations of AI

- ▶ Computer Science: machines, languages, formal theories of computation
- ▶ Mathematics: limits of computation, probability, statistics, Bayes' rule, experimental design
- ▶ Philosophy: logic, epistemology, philosophy of mind, theories of consciousness
- ▶ Psychology: cognitive models, human problem solving and learning, etc.
- ▶ Economics: decision theory, game theory
- ▶ Neuroscience: perception, neural models
- ▶ Linguistics: human language, translation, grammars
- ▶ Biology: adaptive systems, evolving systems

A Brief History of AI

- ▶ Aristotle (384–322 BC): Viewed syllogisms as the cognitive basis for rational thought, and not just as a model to understand idealized rational thought
 - ▶ Syllogism: If A is a B, and B is a C, then A must be a C
- ▶ Descartes (1596–1650): Had a very mechanistic view of the brain
- ▶ Laplace (1749–1827): commented that if he knew the positions of all things in the universe, he could predict the future
- ▶ Lady Lovelace (1842): “The Analytical Engine has no pretensions to *originate* anything. It can do *whatever we know how to order it to perform*”

A More Modern History of AI

- ▶ Gestation (1943–1955): McCulloch-Pitts neurons, Hebbian learning, Can machines think? (Turing, 1950)
- ▶ Birth (1956): Dartmouth conference
- ▶ Great expectations (1952–1969)
- ▶ Humble pie (1966–1973)
- ▶ Knowledge-based systems and weak methods (1969–1979)
- ▶ AI, the industry (1980–present)
- ▶ Re-emergence of neural networks (1986–present)
- ▶ Emergence of probability (1988–present)
- ▶ AI as an empirical science (1987–present)
- ▶ Giving up on the grand prize (1990–present)
- ▶ Rationality rules the roost (1995–present)
- ▶ AI in the age of the Internet and the Web (2001–present)
 - ▶ massive processing power, access to massive amounts of data

Computer Science versus Artificial Intelligence

- ▶ No easy answer
- ▶ Algorithmic versus Heuristic?
 - ▶ “a collection of algorithms that are computationally tractable, adequate approximations of intractably specified problems” (Partridge and Hussain, 1991, p. 5).
- ▶ Exact versus Approximate?
- ▶ Optimal versus Approximate?
 - ▶ People in AI also care about exact and optimal solutions
 - ▶ People in CS also care about approximate solutions


The CS Approach to Tic-Tac-Toe

- ▶ Construct a table with two columns
- ▶ The entries in the left column are board configurations
- ▶ The entries in the right are the best move from the configuration in the left column
- ▶ When the human player moves, the computer looks up the board configuration in the first column and returns the configuration in the second
- ▶ How well will the computer play tic-tac-toe?

current	...	<table border="1"><tr><td>X</td><td>X</td><td></td></tr><tr><td></td><td>O</td><td></td></tr><tr><td>O</td><td></td><td></td></tr></table>	X	X			O		O			...	<table border="1"><tr><td>X</td><td>X</td><td></td></tr><tr><td>X</td><td>O</td><td></td></tr><tr><td>O</td><td></td><td></td></tr></table>	X	X		X	O		O			...
X	X																						
	O																						
O																							
X	X																						
X	O																						
O																							
next	...	<table border="1"><tr><td>X</td><td>X</td><td>X</td></tr><tr><td></td><td>O</td><td></td></tr><tr><td>O</td><td></td><td></td></tr></table>	X	X	X		O		O			...	<table border="1"><tr><td>X</td><td>X</td><td>X</td></tr><tr><td>X</td><td>O</td><td></td></tr><tr><td>O</td><td></td><td></td></tr></table>	X	X	X	X	O		O			...
X	X	X																					
	O																						
O																							
X	X	X																					
X	O																						
O																							

The AI Approach to Tic-Tac-Toe

- ▶ Write a heuristic function that measures the “goodness” of a board configuration
- ▶ For a given board configuration, generate the next possible moves
- ▶ Use the heuristic function to evaluate each next possible move
- ▶ Select the one that is best
- ▶ How well will the computer play tic-tac-toe?



Why Bother with the AI Approach?

- ▶ For complex problems, we can't build a table of perfect moves
- ▶ We can't enumerate the possibilities
- ▶ Shannon (1950) estimated that enumerating all legal configurations for chess would require 10^{90} years of computation
- ▶ The table would require 10^{43} entries
- ▶ Nowadays, computers using the AI approach, such as IBM's Deep Blue, easily beat grandmasters

Characteristics of AI Problems

- ▶ Fully observable, partially observable
- ▶ Single agent, multi-agent, adversarial agent
- ▶ Deterministic, stochastic
- ▶ Static, dynamic
- ▶ Discrete, continuous

A Parting Shot: Tesler's Theorem

- ▶ “Intelligence is whatever machines haven't done yet.”
- ▶ Commonly quoted as “AI is whatever hasn't been done yet.”
- ▶ Also known as the AI effect

Artificial Intelligence: An Introduction

Mark Maloof

Department of Computer Science
Georgetown University
Washington, DC 20057-1232

<http://www.cs.georgetown.edu/~malooof>

August 30, 2017

References I

- E. Charniak and D. McDermott. *Introduction to Artificial Intelligence*. Addison-Wesley, Reading, MA, 1985.
- J. Haugeland. *Artificial intelligence: The very idea*. MIT Press, Cambridge, MA, 1985.
- J. McCarthy, M. I. Minsky, N. Rochester, and C. E. Shannon. A proposal for the Dartmouth summer research project on artificial intelligence, 1955. URL <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>. [Online; accessed 7 August 2014].
- N. J. Nilsson. *Artificial Intelligence: A New Synthesis*. Morgan Kaufmann, San Francisco, CA, 1998.
- D. Partridge and K. M. Hussain. *Artificial intelligence and business management*. Ablex Publishing, Norwood, NJ, 1991.
- E. Rich and K. Knight. *Artificial intelligence*. McGraw-Hill, New York, NY, 2nd edition, 2009.
- E. Rich, K. Knight, and S. B. Nair. *Artificial intelligence*. Tata McGraw-Hill, New Delhi, 3rd edition, 2009.
- S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Upper Saddle River, NJ, 3rd edition, 2010.
- C. E. Shannon. Programming a computer for playing chess. *Philosophical magazine*, 41(314), 1950. URL <http://vision.unipv.it/IA1/ProgrammingaComputerforPlayingChess.pdf>.
- S. Thrun. Toward robotic cars. *Communications of the ACM*, 53(4):99–106, 2010. URL <http://cacm.acm.org/magazines/2010/4/81485-toward-robotic-cars/>.
- A. M. Turing. Computing machinery and intelligence. *Mind*, LIX(236):433–460, 1950. URL <http://mind.oxfordjournals.org/content/LIX/236/433>.