

Artificial Intelligence

An introduction

Artificial Intelligence (AI) - an intro

- AI is everywhere, all the time: ubiquitous
 - Web search
 - Text auto-completion
 - TTS
 - Generative systems
 - Games
- Goals
 - Some applications of AI
 - What is AI
 - Know about some schools of thought and some historically relevant people and events
 - What is the Turing test
 - The General Problem Solver (GPS)
 - What are rational agents
 - About building intelligent agents



Artificial Intelligence (AI) - an intro

- Branch of Computer Science (CS)
- Plural definitions of AI, not singular
 - As one consequence, it is hard to regulate
- What is "intelligence"?
 - To answer the question "Can machines think?", Alan Turing proposed the "Turing test"
 - First described in the 1950 paper "Computing Machinery and Intelligence"
 - Full reprint as chapter 3 of the book "Parsing the Turing Test":
 - <https://link.springer.com/book/10.1007/978-1-4020-6710-5>
- General understanding:
 - Intelligence = Perceive + Analyze + React
- AI?
 - John McCarthy, 1955: "the goal of AI is to develop machines that behave as though they were intelligent"
 - So, "automation of intelligent behavior"?

Artificial Intelligence (AI) - Event

- where, when? @Dartmouth college, workshop, June to August, 1956
- who [organizers]? John McCarthy, Marvin Minsky, Nathaniel Rochester, Claude Shannon
- goal: explore if/how "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"
 - Automatic computers: computers for automation
 - NLP for computer interaction
 - Neuron nets
 - Theory of calculations limits
 - Machine self-improvement
 - Abstractions by machines
 - Randomness and creativity



What is AI? (again)

- What is Intelligence?
- What is AI?
 - The area of Computer Science (CS) that studies how machines can perform tasks, in a way we would normally classify as "intelligent"
 - Subjective
 - Too broad
 - The area of CS that studies how machines can perform tasks, mimicking human intelligence
 - Much of what we have today would not be AI
 - Bridge to neuroscience
 - Midfield?
 - ?



The Turing Test

- A test to see if a computer "can think" to mimic human behavior
- "intelligence" = the ability to achieve human-level intelligence during a conversation
- If a machine can trick a human interrogator into thinking it is human, then it is "intelligent"
- Test design
 - 2 separate sides
 - Human "interrogator", interacts with two "respondents", via text interface
 - physical simulation is unnecessary for intelligence
 - One "respondent" is human, the other is a machine
 - The "interrogator" does NOT know which is which
- Implies: NLP, KR, reasoning, learning
- Total Turing Test? (adds vision + movement)

Approaches, people

- Symbolists - intelligence is "symbol manipulation" - GOFAI
 - names: Marvin Minsky, John McCarthy, Allen Newell, Herbert A. Simon
 - Logic, Philosophy
 - deduction: $A \text{ is human} + \text{all humans are mortal} = ? \text{ (A is mortal)}$
 - induction: $A \text{ is human} + ? = A \text{ is mortal (all humans are mortal)}$
- Connectionists - inspired by the structure and function of biological brains
 - names: Geoffrey Hinton, Yann LeCun, Yoshua Bengio
 - Neuroscience
 - Artificial Neural Networks (ANNs)
- "Evolutionaries" - inspired by principles of biological evolution
 - names: John Holland, Ingo Rechenberg, Hans-Paul Schwefel
 - Biology
 - Genetic algorithms
- Bayesians - use Bayesian statistics and probability
 - names: Judea Pearl (Bayesian networks), Andrew Gelman
 - Statistics
- "Analogizers" - focus on analogy-based methods
 - names: Douglas Hofstadter
 - Psychology
 - Techniques to find similarities between examples



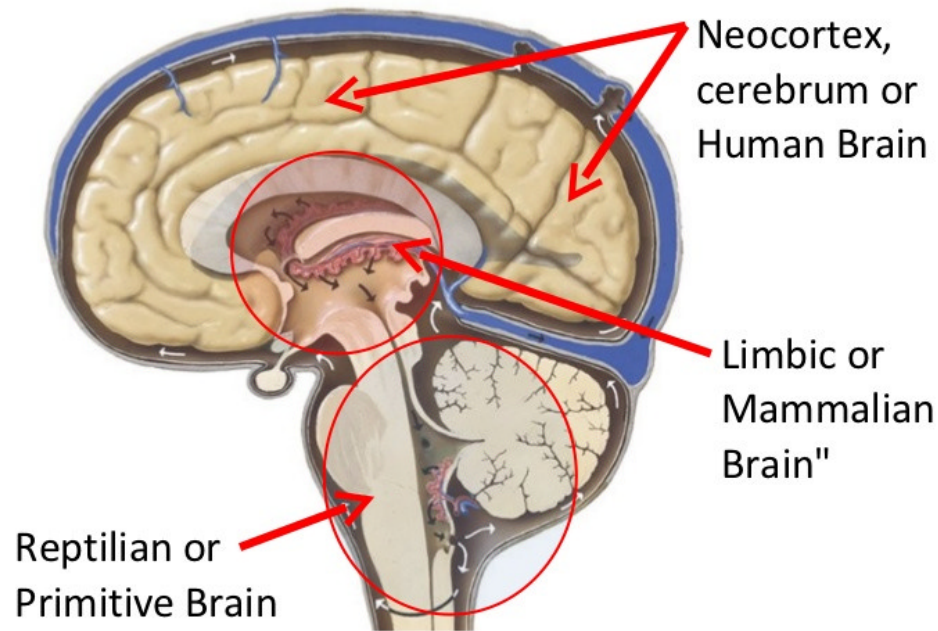
Brain, Body, Mind

- Descartes [1596, 1650]
 - Body vs. Mind
- Phineas Gage [1823, 1860]
 - Body is mind?
- Brain
 - Brainstem (heart control, temperature control) => midbrain (motor functions, basic needs) => limbic system (behavior, emotions) => neocortex (higher-level thoughts)
- From data to "intelligence"
 - Data => *processing* => information => *cognition* => knowledge => *pattern extraction* => understanding => *inference* => intelligence
- Automation is a real need (AI can help)
 - High volume, multiple-sources, unorganized, ever-changing, time-sensitive



Brain, Body, Mind

Your Three Brains



System	found in, for example	behaviors	emotions
Reptilian	reptiles, mammals	reflex, basic survival	none, instinctual drives and reactions
Limbic	mammals	associative	automatic, "instant", unreasoned
Neocortex	humans	deterministic understandings, can think probabilities, can assume neutral positions	controlled, with a "latency", complex emotions

Branches

- Narrow, problem specific, intelligence
 - Logic-based AI
 - statements in logical form, expressing facts and rules
 - Search
 - Path-finding
 - Problem-solving over well-specified state-spaces
 - Heuristics
 - Knowledge Representation
 - Genetic Programming
- General intelligence
- Machine Vision
- Machine Learning (ML)
- Natural Language Processing (NLP)
- Natural Language Generation
- Artificial Creativity



How do <agents> "think"?

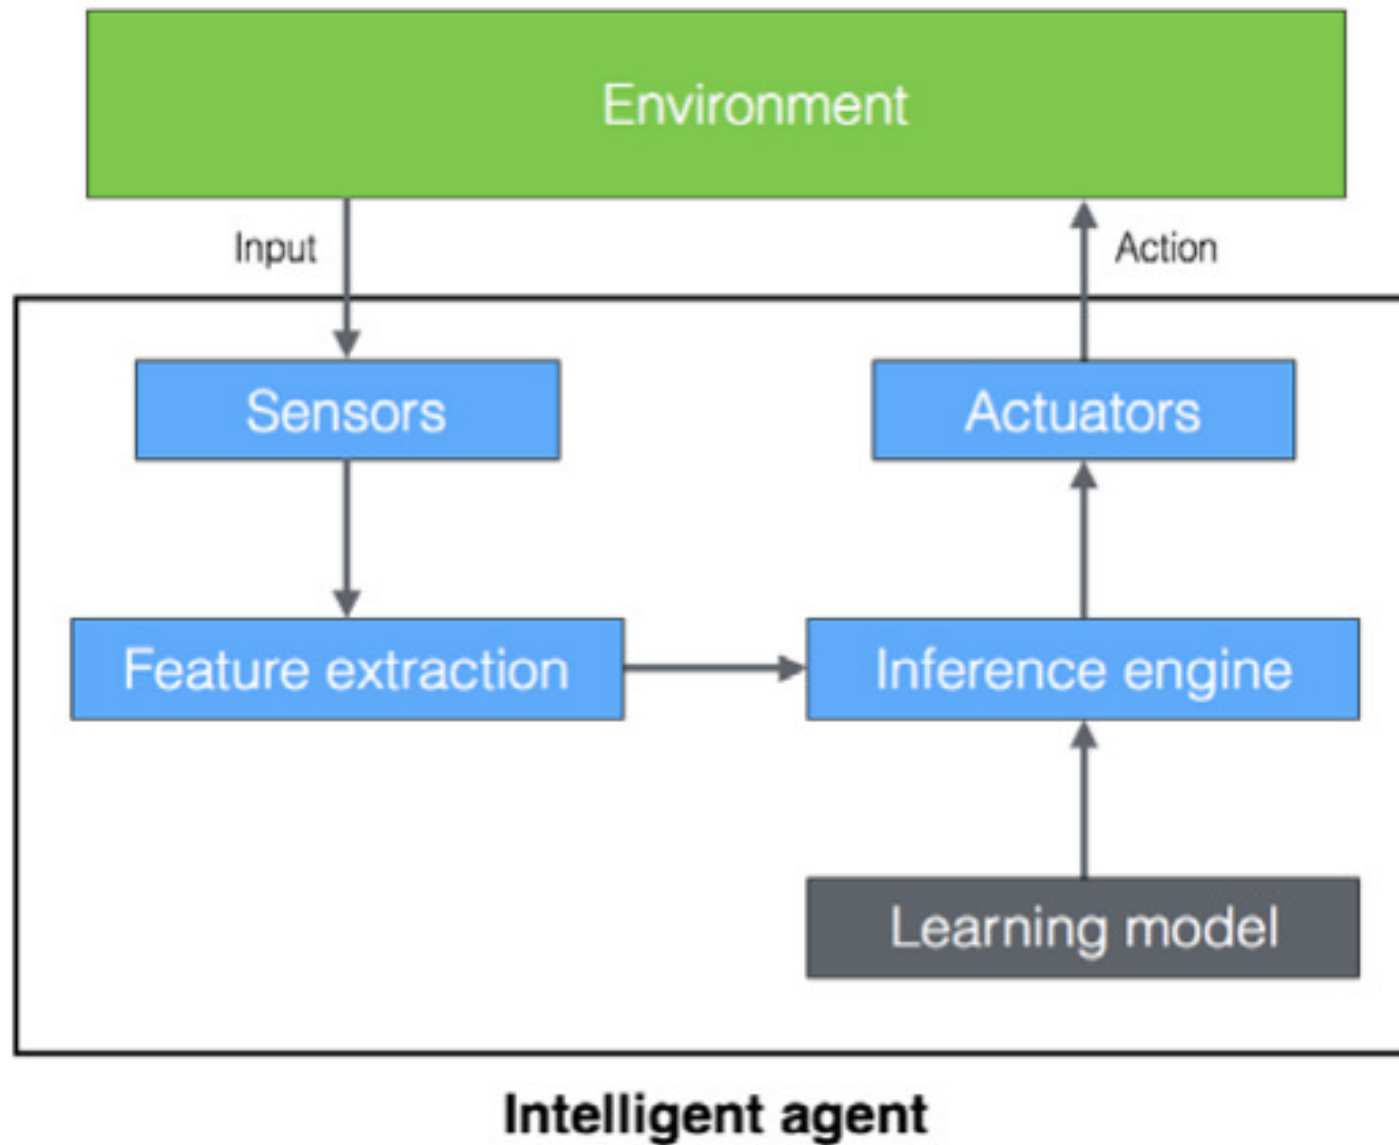
- Humans
 - Dual-Process theory
 - automatic, fast vs. slow, deliberate
 - Cognitive load theory
 - memory hierarchy
 - Multiple intelligences theory
 - plural types of "intelligence"
 - Computational theory of the mind
 - Mind as a computer à la Symbolists
 - Parallel Distributed Processing
 - à la Connectionists (aka Connectionism)
- Artificial agents
 - Rational behavior?
 - "Follow rules and their logical implications to achieve a desirable outcome" (desirable?)



The "GPS"

- General Problem Solver (~1959)
- A symbolists' system
 - high-level, explicit representations of problems and logic-based methods for solving them
- Theory of human problem solving stated in the form of a simulation computer program
- Goal: "universal problem-solving machine", using the same base algorithm for every problem (!)
- Created with IPL (Information Processing Language)
- Every problem
 - set of well-formed formulas
 - represented on a directed graph
- Could solve
 - Proofs of mathematical theorems in geometry and logic
 - Not real-world problems

Building "intelligent" agents



References

- About Descartes
 - <https://iep.utm.edu/descmind/>
- Alan Turing on "intelligent machines" @
https://github.com/amsm/20190623_on_turing_birthday/raw/master/sci_paper_alan_turing_1948_intelligent_machinery.pdf
- The Quest for AI @
<https://books.google.com/books?id=nUJdAAAAQBAJ>
- https://en.wikipedia.org/wiki/General_Problem_Solver

