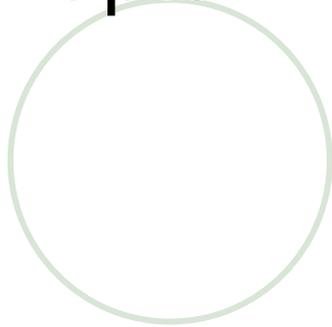




What is A.I.?

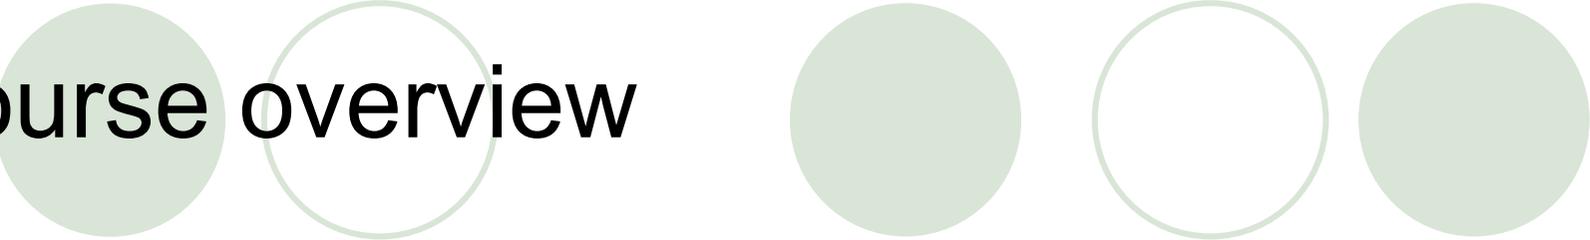
Chapter 1



Outline

- Course overview
- What is AI?
- A brief history
- The state of the art

Course overview



- Introduction and Agents (chapters 1, 2)
- Search (chapters 3, 4, 5, 6)
- Logic (chapters 7, 8, 9)
- Uncertainty (chapters 13, 14)
- Learning (chapters 18, 20)
- Natural Language Processing (chapters 22, 23)

What is AI?

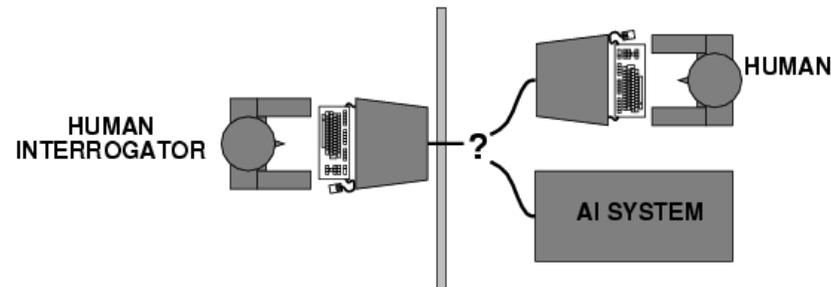
Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

The textbook advocates "acting rationally"

Acting humanly: Turing Test

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" → "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game



- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning

Thinking humanly: cognitive modeling

- 1960s "cognitive revolution": information-processing psychology
- Requires scientific theories of internal activities of the brain
 - How to validate? Requires
 - 1) Predicting and testing behavior of human subjects (top-down)
 - or 2) Direct identification from neurological data (bottom-up)
- Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI

Thinking rationally: "laws of thought"

- Aristotle: what are correct arguments/thought processes?
- Several Greek schools developed various forms of *logic: notation and rules of derivation* for thoughts; may or may not have proceeded to the idea of mechanization
- Direct line through mathematics and philosophy to modern AI
- Problems:
 1. Not all intelligent behavior is mediated by logical deliberation
 2. What is the purpose of thinking? What thoughts should I have?

Acting rationally: rational agent

- **Rational** behavior: doing the right thing
- The right thing: that which is expected to maximize goal achievement, given the available information
- Doesn't necessarily involve thinking – e.g., blinking reflex – but thinking should be in the service of rational action

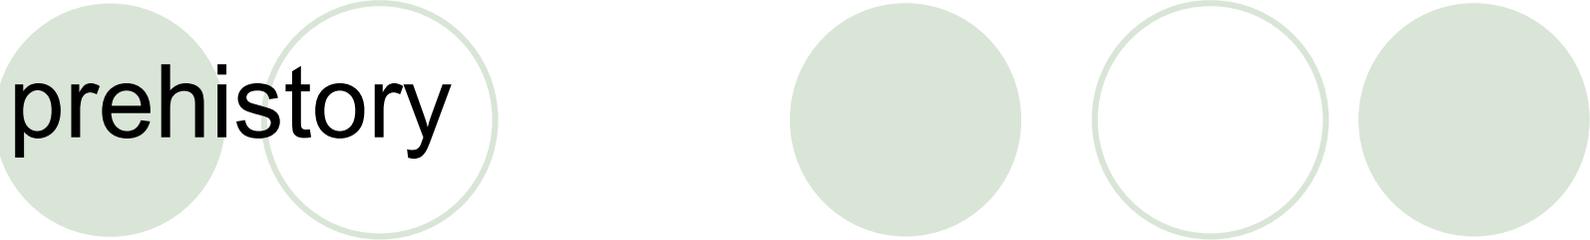
Rational agents

- An **agent** is an entity that perceives and acts
- This course is about designing rational agents
- Abstractly, an agent is a function from percept histories to actions:

$$[f: \mathcal{P}^* \rightarrow \mathcal{A}]$$

- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: computational limitations make perfect rationality unachievable
 - design best **program** for given machine resources

AI prehistory

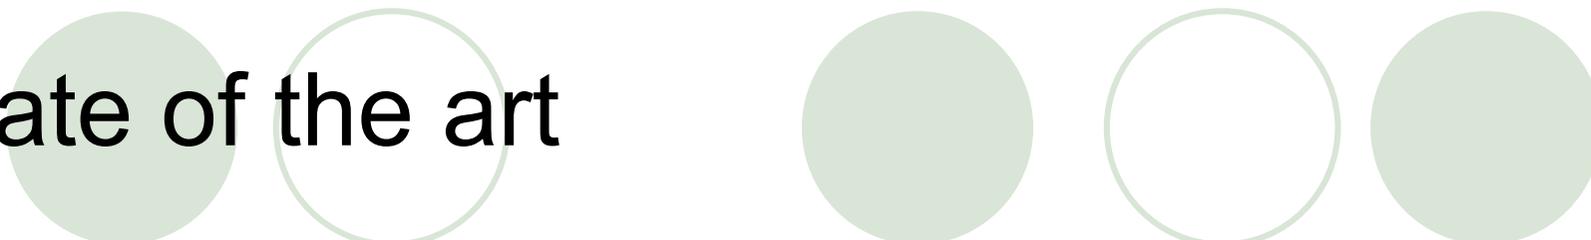


- Philosophy Logic, methods of reasoning, mind as physical system foundations of learning, language, rationality
- Mathematics Formal representation and proof algorithms, computation, (un)decidability, (in)tractability, probability
- Economics utility, decision theory
- Neuroscience physical substrate for mental activity
- Psychology phenomena of perception and motor control, experimental techniques
- Computer engineering building fast computers
- Control theory design systems that maximize an objective function over time
- Linguistics knowledge representation, grammar

Abridged history of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1952–69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1965 Robinson's complete algorithm for logical reasoning
- 1966–73 AI discovers computational complexity
Neural network research almost disappears
- 1969–79 Early development of knowledge-based systems
- 1980– AI becomes an industry
- 1986– Neural networks return to popularity
- 1987– AI becomes a science
- 1995– The emergence of intelligent agents
- 2001– The availability of very large data sets

State of the art



- Robotic vehicles can drive autonomously in most situations (95%). CMU's BOSS can drive through an urban environment, following traffic rules and avoiding pedestrians.
- Customers can call United Airlines to book flights, or use Google Voice to translate their native speech into other languages
- Deep Blue beat the world-reigning chessmaster Gary Kasparov in 1997, and computers have continued to convincingly beat humans in recent years.
- Learning algorithms help to classify spam mail, helping all email users save time, sorting out over 80-90% of mail as spam traffic.
- Both military and commercial sectors employ AI to handle logistics. Aircraft routing and convoy logistics A.I. are used to coordinate the movement of massive numbers of supplies and units according to constraints