

CS494/594: Artificial Intelligence

Fall 2013

Tuesday/Thursday, 11:10 AM – 12:25 PM

Instructor: Dr. Lynne E. Parker

TA: Hao Zhang



"Artificial Intelligence is the study of how to make real computers act like the ones in the movies."

--Anonymous

Outline

- Overview syllabus and class policies
- Course Overview
- Introduction to AI (Chapter 1)
 - What is AI?
 - A brief history
 - The state of the art
- Intelligent Agents (Chapter 2)
 - Agents and environments
 - Rationality
 - PEAS (Performance measure, Environment, Actuators, Sensors)
 - Environment types
 - Agent types (next time)

Overview of Syllabus and Class Policies

(See handout)

Course Overview

- Introduction to AI and Intelligent Agents
- Problem-solving
 - Heuristic Search
 - Constraint Satisfaction
 - Searching with non-deterministic actions and partial observability
 - Online search
 - Adversarial search (games)
 - Constraint propagation
- Knowledge, reasoning, and planning
 - First order logic
 - Classical planning
 - Planning/acting in real world
 - Knowledge representation
- Dealing with uncertainty – probabilistic reasoning
- Natural language processing
- (if time allows) Perception and robotics
- Philosophical and ethical issues
- Practical applications

What is AI?

Systems that think like humans

“The automation of activities that we associate with human thinking – activities such as decision-making, problem solving, learning, ...”

(Bellman, 1978)

Systems that think rationally

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

(Charniak and McDermott, 1985)

Systems that act like humans

“The art of creating machines that perform functions that require intelligence when performed by people”,
(Kurzweil, 1990)

Systems that act rationally


“AI...is concerned with intelligent behavior in artifacts.”

(Nilsson, 1998)

Acting humanly: The 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 6 major components of AI: knowledge representation, automated reasoning, natural language understanding, machine learning, computer vision, robotics

Problem: Turing test is not **reproducible**, **constructive**, or amenable to **mathematical analysis**

Thinking humanly: Cognitive modeling approach

- 1960s “cognitive revolution”: information-processing psychology replaced prevailing orthodoxy of behaviorism
- Requires scientific theories of internal activities of the brain
 - What level of abstraction? “Knowledge” or “circuits?”
 - How to validate? Requires:
 - 1) Predicting and testing behavior of human subjects (top-down)
 - 2) or, Direct identification from neurological data (bottom-up)
- Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI
- Both share with AI the following characteristic:
 - *The available theories do not explain (or engender) anything resembling human-level general intelligence*
- Hence, all three fields share one principal direction!

Thinking rationally: “Laws of Thought” approach

- Normative (or prescriptive) rather than descriptive
- 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 approach

- **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
- Aristotle (Nicomachean Ethics):
Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good

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:


$$f : P^* \rightarrow A$$

- 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

Foundations of AI


- Philosophy (428 BC – Present)
 - Can formal rules be used to draw valid conclusions?
 - How does the mental mind arise from a physical brain?
 - Where does knowledge come from?
 - How does knowledge lead to action?

- Aristotle, Leonardo da Vinci, Pascal, Descartes, etc.




Foundations of AI (con't.)

- Mathematics (~800 – present)
 - What are the formal rules to draw valid conclusions?
 - What can be computed?
 - How do we reason with uncertain information?
- Logic, computation (algorithms), probability




Foundations of AI (con't.)

- Economics (1776-present)
 - How should we make decisions so as to maximize profit?
 - How should we do this when others may not go along?
 - How should we do this when the payoff may be far in the future?
- Utility, decision theory, game theory, operations research, Markov decision processes




Foundations of AI (con't.)

- Neuroscience (1861-present)
 - How do brains process information?




- Moore's law predicts that CPU's gate count will equal brain's neuron count around 2020.
- *But, even though a computer is a million times faster in raw switching speed, the brain is actually 100,000 times faster at what it does.*




Foundations of AI (con't.)

- Psychology (1879 – present)
 - How do humans and animals think and act?
- Behaviorism, cognitive psychology, cognitive science



Foundations of AI (con't.)

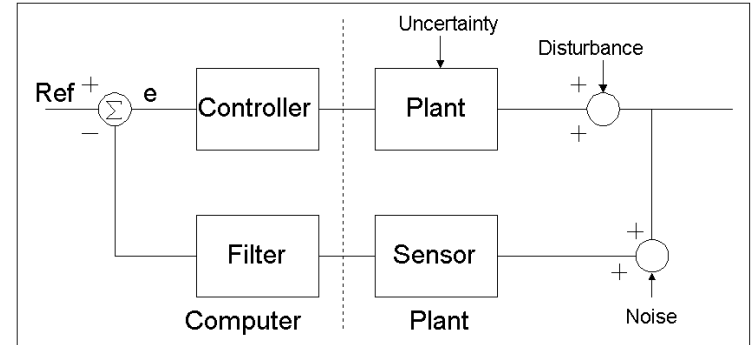
- Computer engineering (1940 – present)
 - How can we build an efficient computer?



- AI requires: (1) intelligence, (2) an artifact (i.e., a computer upon which the intelligence is generated)

Foundations of AI (con't.)


- Control theory and Cybernetics (1948 – present)
 - How can artifacts operate under their own control?



- Control theory: Maximizing an objective function over time
 - Uses calculus and matrix algebra, which lend themselves to systems that are describable by fixed sets of continuous variables;
 - Exact analysis typically feasible only for linear systems
- AI: Designing systems that behave optimally
 - Founded as a way to “escape” from limitations of the mathematics of control theory
 - Use of logical inference and computation allows AI to consider problems such as language, vision, and planning, which are outside the field of control theory

Foundations of AI (con't.)

- Linguistics (1957-present)
 - How does language relate to thought?
- Computational linguistics, natural language processing, knowledge representation



Summary of 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
 - formal theory of rational decisions
- Neuroscience
- Psychology
 - adaptation
 - phenomena of perception and motor control
 - experimental techniques (psychophysics, etc.)
- Control theory
 - homeostatic systems, stability
 - simple optimal agent designs
- Linguistics
 - knowledge representation
 - grammar

Potted history of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1952-69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1966-74 AI discovers computational complexity
Neural network research almost disappears
- 1969-79 Early development of knowledge-based systems
- 1980-88 Expert systems industry booms
- 1988-93 Expert systems industry busts: "AI Winter"
- 1985-95 Neural networks return to popularity
- 1988+ AI + scientific method. Resurgence of probability; general increase in technical depth. "Nouvelle AI": Artificial life, GAs, soft computing
- 1995+ Agents, agents everywhere ...
- 2001+ Availability of very large data sets

What can AI do today?

- Robotic Vehicles:

- (2005) Stanley (the driverless vehicle, from Stanford) went 132 miles through the Mojave desert to win DARPA Grand Challenge



- (2006) Boss (the driverless vehicle from CMU) drove autonomously in urban environment to win DARPA's Urban Challenge



What can AI do today? (con't.)

- Robotic Vehicles (con't):

- Google's Driverless Car – as of 2012, vehicle had logged over 300,000 miles of accident-free driving on highways and roads



What can AI do today? (con't.)

- Speech recognition:

- United Airlines can book flights using only autonomous speech recognition and dialog management



- Siri – Apple's intelligent personal assistant that uses natural language user interfaces to answer questions, make recommendations, and perform actions



What can AI do today? (con't.)

- Autonomous planning and scheduling:
 - NASA's Remote Agent program became 1st onboard autonomous planning program to control the scheduling of operations for a spacecraft



What can AI do today? (con't.)

- Game playing:

- 1997: IBM's Deep Blue became 1st computer to defeat world champion in a chess match



- 2011: IBM's Watson won 1st prize (\$1M) on Jeopardy! in competing against top (human) former winners



What can AI do today? (con't.)

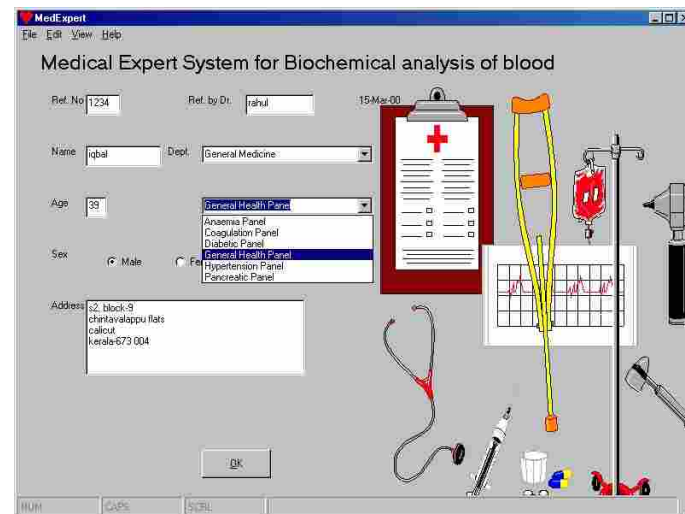
- Spam-fighting:
 - (Today) Learning algorithms classify >1 Billion messages each day as spam



What can AI do today? (con't.)

- Diagnosis:

- Medical diagnosis programs based on probabilistic analysis have been able to perform at the level of an expert physician in several areas of medicine



What can AI do today? (con't.)

- Logistics Planning:

- U.S. military deployed a Dynamic Analysis and Replanning Tool (DART) in 1991, for automated logistics planning and scheduling, generating plans in hours that previously would have taken weeks




What can AI do today? (con't.)

- Robotics:

- Many surgeons now use robotic devices in surgery (e.g., da Vinci robot)



- Roomba cleans floors (over 2M sold)



- PackBot handles hazardous materials



What can AI do today? (con't.)

- Language understanding and problem solving:
 - PROVERB (1999) is a computer program that can solve crossword puzzles better than most humans

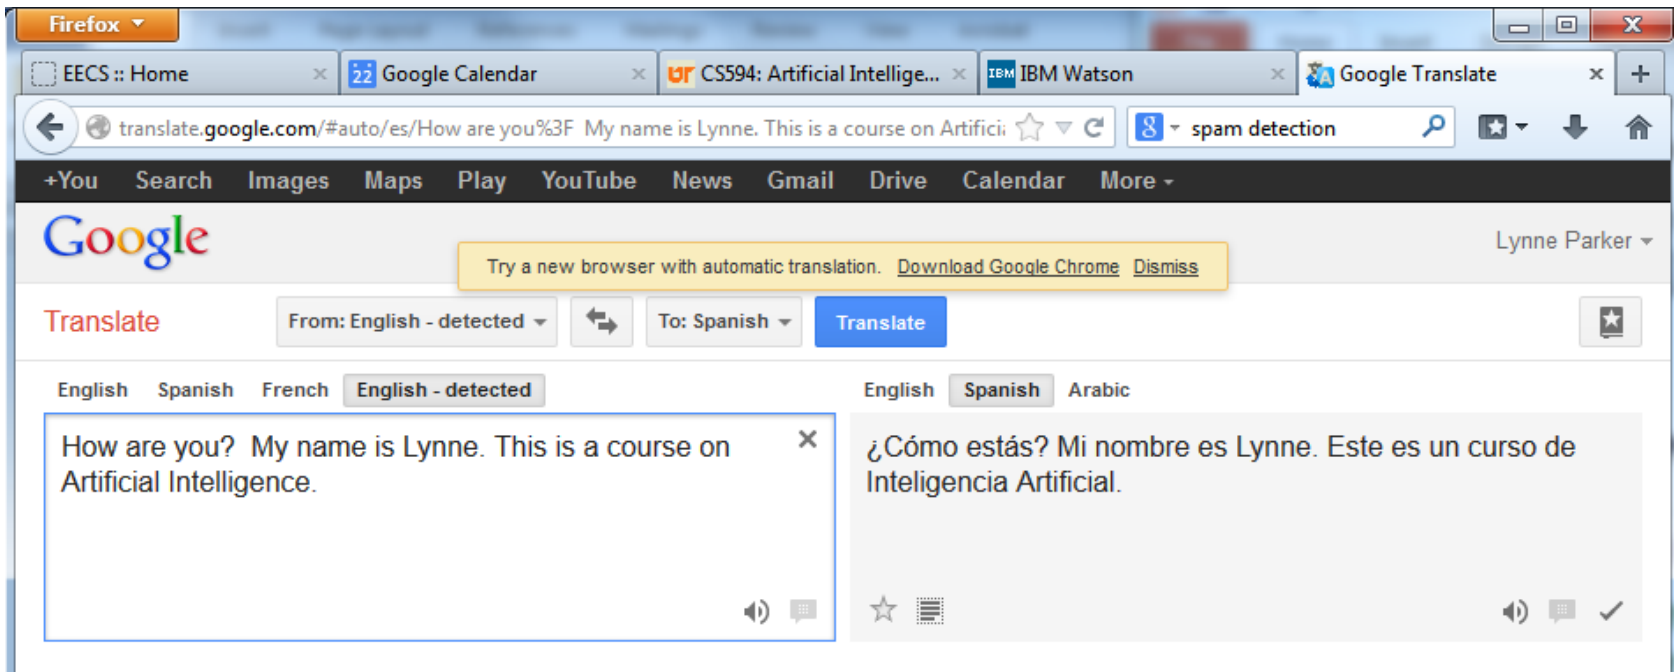


© 2007. The New York Times

What can AI do today? (con't.)

- Machine translation

- Computer program automatically translates between various, using machine learning techniques (e.g., SDL Language Weaver's software, Google Translate, etc.)



Your Assignment before next class...

- Read Chapters 1-2 of Russell and Norvig
- See you next time...