CS 520: Introduction to Artificial Intelligence

Prof. Louis Steinberg

CS 520

- **Prof. Louis Steinberg**
  - **401 Hill, 445-3581, lou@cs**
  - **Office hours: Thursday 1-3pm and by appointment**
- **TA: Xiaolei Huang (xiaolei@paul)**

CS 520

- **Objective: Broad intro to the field of AI**
- **Covers**
  - **Goals, assumptions, mindset**
  - **Core topics of AI, including search, knowledge representation, reasoning, planning and learning;**
  - **AI Programming techniques in LISP and PROLOG**
  - **Selected applications of AI**
- **Prerequisites**
  - **1st order logic**
  - **Intros to graphs, algorithms, complexity**
  - **Introductions to Lisp (or Scheme) and Prolog**

CS 520

- **Expected work:**
  - **Reading, problem sets, programs**
  - **midterm, final**
- **Course home page:**
  - **http://www.cs.rutgers.edu/~lou/520/**
  - **Announcements, information, assignments, lecture notes**

Approximate Schedule (by week)

- Introduction
- Problem Solving and Search
- Programming in Lisp
- Knowledge Representation
- Programming in Prolog
- Planning

- Midterm Exam

Approximate Schedule (by week)

- Probabilistic Reasoning and Decision Making
- Machine Learning
- Genetic Algorithms and Evolutionary Computation
- Computer Vision
- Natural Language Processing

What is AI

- A set of goals
  - build an artificial intelligence
  - useful subgoals
- A class of problems
  - characteristics common to these goals
- A set of methods
  - commonly useful in solving problems like these
- A set of people

A Set of Goals

- A machine that can do anything a human can do (and requires a brain / mind)
- Subgoals of this that are useful in their own right
  - Understand natural language
  - Plan
  - Learn from experience
  - …

Build an artifact that is intelligent

Definitions of Artificial Intelligence

	Human-centered	Rationality
Thought		
Action		

Acting Humanly: The Turing Test approach

- The Turing test: A human interrogates a subject through a teletype.
  - If the human cannot tell whether the subject is a human or computer, the computer passes the test.
- To pass the Turing test a computer needs:
  - natural language processing
  - knowledge representation
  - automated reasoning
  - machine learning

Thinking Humanly: The Cognitive Modeling approach

- **Requires a model for human cognition. precise enough models allow simulation by computers.**
  - **Cognitive Science brings together Computer Science, Linguistics, Philosophy, and Psychology.**
  - **Rutgers has a center for Cognitive Science (RuCCS).**
- **Cognitive Science will not be covered in this course.**

Thinking Rationally: The laws of thought approach

- **Formal logic provides a precise notation for statements about things and their relationships.**
  - **Given sufficient memory and time, early computer programs were able to solve problems formulated in logical form, using automated reasoning and theorem proving techniques**

Thinking Rationally: The laws of thought approach

- **Obstacle**
  - **informal knowledge cannot be easily stated in logical notation.**
  - **even small logic programs may have unacceptable time and memory requirements.**

    the formal logic approach will be studied in this course

Acting Rationally: The rational agent approach

- **means to act so as to achieve one's goals given one's beliefs.**
  - **an agent is something that perceives and acts.**
  - **may require thinking rationally to decide which action will achieve one's goal.**
  - **may require natural language, vision, andlearning skills to be able to communicate with the world and generate better strategies over time.**

Advantages of the Rational Agent Approach:

- more general than the laws of thought approach.
- more amenable to study than the cognitive modeling approach.
- we will focus on the rational agent approach in this course.

Commonalities Among the Different Perspectives

- Shared belief that humans are a good source of clues about how to build an intelligent machine.
- Shared belief that theories of intelligence should be tested by implementing them in computer programs and testing them on real problems.

Universal v. Expert Abilities

- **Abilities all normal adult humans have:**
  - **Seeing, hearing, walking, talking, learning, common sense.**
- **Abilities only some human experts have:**
  - **Proving theorems, playing chess, playing a musical instrument, managing companies, negotiating agreements.**
- **The expert abilities have turned out to be easier for AI machines.**

What is AI

- **A set of goals**
  - **build an artificial intelligence**
  - **useful subgoals**
- **A class of problems**
  - **characteristics common to these goals**
- **A set of methods**
  - **commonly useful in solving problems like these**
- **A set of people**

A Class of Problems

- **Problems that require search:**
  - **No deterministic algorithm is known.**
  - **Must use “trial and error".**
  - **NP-Hard problems all have this property.**
  - **Example: Schedule courses.**
  - **Non-example: Sort a class roster.**

A Class of Problems

- **Problems that are poorly specified:**
  - **We don't know a concise, exact problem specification.**
  - **We don't know what knowledge is needed to solve the problem.**
  - **We don't have the knowledge needed to solve the problem.**
  - **Our knowledge is imprecise or inaccurate.**
  - **Example: Explain integration to a human.**
  - **Non-example: Factor an integer.**

A Set of Methods

- **Use of general inference methods**
  - **such as heuristic search, constraint propagation or resolution theorem proving.**
- **Representation of knowledge in declarative form,**
  - **such as search spaces, constraint networks or systems of logical axioms.**

A Set of Methods

- **Heuristic Search**
- **Expert Protocols**
- **Iterative Programming**
- **More task-specific methods, e.g. for learning or planning**
  - **Tend to cross tasks**
  - **“AI-complete” problems**

A Set of People

- **John McCarthy, Marvin Minsky, …**
- **Saul Amarel, Cordell Green, Terry Winograd, …**
- **Me**
- **You?**

Intelligent Agents

- **An agent**
  - **Exists in an *environment***
  - **Has *sensors* that detect *percepts* in the environment**
  - **Has *effectors* that can carry out *actions* which may change the environment**
  - **Has *goals* to achieve**
- **See Fig 2.1 in Russell & Norvig**

Spectrum of Agent Complexity

- **Reflex agents**
  - **Percepts -> actions**
  - **See fig 2.7**
- **With internal state**
  - **Have memory of past percepts, actions**
  - **See fig 2.9**
- **With Goals**
  - **Explicit representation desired state(s) of environment**
  - **See fig 2.11**
- **Utility-based**
  - ***How much* is each goal desired**
  - **See fig 2.12**

Environments

- **accessible vs. inaccessible**
- **deterministic vs. nondeterministic (stochastic)**
- **episodic vs. nonepisodic**
- **static vs. dynamic**
- **discrete vs. continuous**

Environment	Accessible	Deterministic	Episodic	Static	Discrete
Chess with a clock	Yes	Yes	No	Semi	Yes
Chess without a clock	Yes	Yes	No	Yes	Yes
Poker	No	No	No	Yes	Yes
Backgammon	Yes	No	No	Yes	Yes
Taxi driving	No	No	No	No	No
Medical diagnosis system	No	No	No	No	No
Image-analysis system	Yes	Yes	Yes	Semi	No
Part-picking robot	No	No	Yes	No	No
Refinery controller	No	No	No	No	No
Interactive English tutor	No	No	No	No	Yes

Lectures on Search

- **Formulation of search problems.**
- **Uninformed (blind) search algorithms.**
- **Informed (heuristic) search algorithms.**
- **Constraint Satisfaction Problems.**
- **Game Playing Problems.**

State Space Search: Formal Definition

- **Given:**
  - **A set of states.**
  - **An initial state.**
  - **A set of operators mapping states to states.**
    - **Preconditions of each operator.**
    - **Effects of operator.**
    - **(Cost of each operator).**
  - **A set of goal states.**
- **Find: A (minimal cost) sequence of operators that transforms the initial state into a goal state**