Introduction to Artificial Intelligence CS430

Instructor: Tom Dietterich
221C Dearborn Hall
Teaching Assistant: Hongli Deng
108 Hovland Hall (office hours)

Acknowledgement: Thanks to Devika Subramanian of Rice University for some course materials

Textbook

- Russell & Norvig: Artificial Intelligence A Modern Approach (2nd Edition)

"The publication of this textbook was a major step forward, not only for the teaching of AI, but for the unified view of the field that this book introduces. Even for experts in the field, there are important insights in almost every chapter." (Amazon.com review)

Course Plan

- Fundamentals: representation, reasoning, and learning
  - unified representation: Bayesian networks
- Application Areas of Intelligent Systems
  - Natural Language Processing
  - Vision and Speech
  - Robotics

Assignments

- Daily Reading Assignments
- Weekly Homework Assignments
- Bi-weekly Programming Assignments
- Two Midterm Exams
- Final Project: Learning Spam Filter
  - Alternative projects may be proposed.
  - Teams of 1-3 people (prefer 2 people)

Course Objectives (1)

- Master Bayesian networks for knowledge representation
- Understand two Bayesian network reasoning methods
  - Belief propagation
  - Particle filters
- Understand one Bayesian network learning method

Course Objectives (2)

- Be able to apply Bayesian networks for language modeling (specifically, for email spam detection)
- Apply algorithms for learning the networks
- Understand how Bayesian networks are applied to vision, speech, robotics, etc.

What is Artificial Intelligence?

- Computer Science
  - Methods for applying computers to problems
  - Study of the fundamental limits of computation
- Artificial Intelligence
  - Methods for applying computers to problems that require "intelligence"
  - Study of the fundamental limits of "intelligent" behavior by computers

What is Intelligence?

	"Like People"	"Rationally"
Think	Cognitive Science	Laws of Thought
Act	Turing Test	Rational Agents

Act Like Humans: The Turing Test

- Can Computer fool a human interrogator?

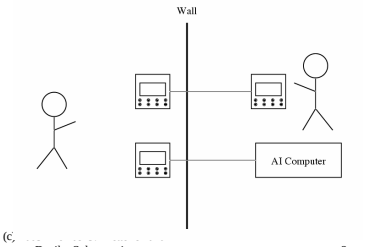

Abilities Required for Turing Test

- Natural Language Processing (understanding, generation)
- Automated Reasoning
- Learning
- Knowledge Representation and Storage
- Vision (for "total turing test")
- Robotics (for "total turing test")

Problem: Tends to focus on human-like errors, linguistic tricks, etc. Does not product useful computer programs

Think Like Humans: Cognitive Science

- Goal: Develop precise theories of human thinking
- Cognitive Architecture (e.g., SOAR, ACT-R)
  - Software Architecture for modeling human performance
  - Describe task, required knowledge, major subgoals
  - Architecture follows human-like reasoning
  - Makes testable predictions: Time delays during problem solving, kinds of mistakes, eye movements, verbal protocols, learning rates, strategy shifts over time, etc.
- Problems:
  - Identifiability: It may be impossible to identify the detailed structure of human problem solving using only externally-available data. "Optimal" performance is an excellent predictor of human performance in most routine tasks.

Thinking Rationally: The Logical Approach

- Ensure that all actions performed by computer are justifiable ("rational")

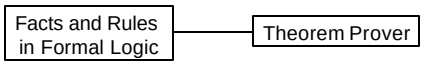

- Rational = Conclusions are provable from inputs and prior knowledge
- Problems:
  - Representation of informal knowledge is difficulty
  - Hard to define "provable" plausible reasoning
  - Combinatorial explosion: Not enough time or space to prove desired conclusions.

Acting Rationally: Rational Agents

- Claim: "Rational" means more than just logically justified. It also means "doing the right thing"

Rational agents do the best they can given their resources

Rational Agents

very few resources		lots of resources
no thought "reflexes"	limited, approximate reasoning	Careful, deliberate reasoning

- Adjust amount of reasoning according to available resouces and importance of the result
- This is one thing that makes AI hard

Areas of Study in AI

- Reasoning, optimization, resource allocation
  - planning, scheduling, real-time problem solving, intelligent assistants, internet agents
- Natural Language Processing
  - information retrieval, summarization, understanding, generation, translation
- Vision
  - image analysis, recognition, scene understanding
- Robotics
  - grasping/manipulation, locomotion, motion planning, mapping

Where are we now?

- SKICAT : a system for automatically classifying the terabytes of data from space telescopes and identifying interesting objects in the sky. 94% classification accuracy, exceeds human abilities.
- Deep Blue: the first computer program to defeat champion Garry Kasparov.
- Pegasus: a speech understanding program that is a travel agent (1-877-LCS-TALK).
- Jupiter: a weather information system (1-888-573-TALK)
- HipNav: a robot hip-replacement surgeon.

Where are we now?

- Navlab: a Ford escort that steered itself from Washington DC to San Diego 98% of the way on its own!
- google news: autonomous AI system that assembles "live" newspaper
- DS1: a NASA spacecraft that did an autonomous flyby an asteroid.
- Credit card fraud detection and loan approval
- Search engines: www.citeseer.com, automatic classification and indexing of research papers.
- Proverb: solves NYT puzzles as well as the best humans.

Surprises in AI research

- Tasks difficult for humans have turned out to be "easy"
  - Chess
  - Checkers, Othello, Backgammon
  - Logistics planning
  - Airline scheduling
  - Fraud detection
  - Sorting mail
  - Proving theorems
  - Crossword puzzles

Surprises in AI research

- Tasks easy for humans have turned out to be hard.
  - Speech recognition
  - Face recognition
  - Composing music/art
  - Autonomous navigation
  - Motor activities (walking)
  - Language understanding
  - Common sense reasoning (example: how many legs does a fish have?)