CSCI 400 Artificial Intelligence
David Keil, Framingham State University

Introduction

1. What this course is about
2. Course plan and details

Invitation

- Would you like to explore the limits of computers' ability to solve the kinds of problems that humans solve?
- Would you like to discover what the mechanisms of cognition, learning, knowledge, and belief are about?
- To know how a robot works?
- To examine the arguments for and against machine consciousness?
- Join us in this course and look into the future.

Inquiry

- May AI be summarized as *rational adaptive computational behavior*?
- Each of us may get some new ideas from studying AI; the "Inquiry" items listed on the syllabus are some questions that we will investigate

1. What this course is about

- We are studying here *solving difficult computational problems*
- Hence searching collections, sorting arrays, and playing tic-tac-toe are *not* problems requiring intelligence, though they may give us ideas about solving some of them
- Some AI problems are *algorithmic*, some *interactive*
- We emphasize interaction, equivalent to what Russell and Norvig call *agent* problem solving

Defining "difficult problem"

- Hard computational problems means *NP-hard*, i.e. functional problems that are expected to require time exponential in the size of the input, with combinatorial explosion; or *interactive* (service or mission) problems that are NP-hard at each step
- Hence these are problems not efficiently solvable by brute force, divide and conquer, greedy, or dynamic-programming approaches to algorithm design

Domain of concern

- Should this course be called "hard information processing problems"?
- Is AI the set of problems that have ever been *considered* hard?
- AI is:
  - Inference – Search for solutions
  - Interaction – Reasoning
  - Planning – Learning, adaptation
  - Pattern matching – Using heuristics
  - Social – Imitating humans
  - Solving NP-hard problems

Intelligence and mind

- Intelligence is much more than imitating humans; is not something quantified, as with IQ tests
- We don't define intelligence, mind, or consciousness in this course, though we discuss definitions that have been offered
- AI overlaps with neuroscience, computer science, cognitive science, psychology, sociology, economics, control theory, and operations research
- AI is part of computer science but is useful to other fields as well

Russell-Norvig's definition of AI

- *AI:* "the study of agents that receive percepts from the environment and perform actions. Each such agent implements a function that maps percept sequences to actions..."
- Instructor accepts
  - sentence 1, except that *any* program does this
  - sentence 2, except that the *semantics* are wrong, because actions really derive from the *state of the environment*, and any percepts after the first one may be the results of previous actions

Four approaches to AI

AI systems are defined alternatively as ones that:

1. *Act like humans* (e.g., Turing test)
2. *Think like humans* (machines with minds; cognitive science)
3. *Think rationally* (logicist tradition of AI, focusing on formal reasoning)
4. *Act rationally* (rational-agent approach)

Rational-agent approach is the most general, because inference is only one way to generate good action, and the purpose of an action may be different from the purpose a human might have

Rational adaptive computational behavior

AI systems are:

- *Rational*: using inference to maximize expected reward
- *Adaptive*: learning from the environment
- *Computation*: processing of symbols
- *Behavior*: taking action to achieve goals

Classifying problems as AI problems

AI	Not AI
Chess playing	Tic tac toe
Integrating algebraic expressions	Addition
Voice recognition	Searching arrays
Finding concepts from examples	Sorting
Rule-based inference	Voice generation

Research and reality

- The setting for the real intelligence that is agreed to exist today is the human *brain*
- In modeling cognition computationally, we rely on a foundation of *models of computation* and *notions of intractability*
- Intelligent behavior is associated with finding satisfactory solutions to *hard* (intractable) *computational problems*

Toy vs. real-world problems

- *Toy* (no one would buy a solution):
  - Vacuum world (move if in clean place, else suck dirt)
  - 8-puzzle (repeatedly move square piece to empty square in 3 x 3 grid)
  - 8-queens (arrange on chessboard so none threatened)
- *Real-world* (of practical use):
  - Routing (travel, packet communication)
  - Natural-language understanding
  - Expert inference
  - Robotic navigation
  - Voice recognition

Typical features of AI problems

- Use inference
- Processing relates to meaning
- Use heuristics
- Intractable requiring suboptimal solutions

Early history of AI

- Artificial neuron research by McCulloch and Pitts, 1943
- *Explosion of hope:*
  - Dartmouth workshop (J. McCarthy, M. Minsky, C. Shannon, Rochseter, Newell, H. Simon, A. Samuel), 1956
  - LISP, General Problem Solver, theorem provers
  - Physical symbol system hypothesis
  - Blocks worlds
  - Perceptrons (neural nets)

AI since 1969

- Knowledge systems, expert systems
- Dendral (1969) inferred molecular structure from spectral data, separated knowledge (rules) from reasoning (inference)
- Japanese “Fifth Generation” project and U.S. consortium (1980s) brought billions in investment
- Transition to rigorous scientific foundations with Hidden Markov Models, Bayesian networks
- Rise of intelligent-agent model, mid-1990s, focused on interaction with environment

Approaches to AI

- Physical symbol system hypothesis inspired:
  - Empirical and constructivist approach
  - Research revolved around forms for representation and search algorithms
  - Notion of disembodied cognition
- Probabilistic and stochastic approaches

Constructivist approach

- Emphasis shifted from representation and search toward learning and adaptation
- Simple local adaptations enable complex system to be shaped in response to environment
- *Agent and evolutionary approach:* From independent action of agents, intelligent behavior emerges by natural selection

2. Course plan and details

Topics:

1. Adaptation, cognition, and computation
2. State-space search
3. Knowledge representation and inference
4. Uncertainty and probabilistic reasoning
5. Supervised learning and natural language
6. Reinforcement learning and adaptation
7. Distributed AI and multi-stream interaction
8. Philosophical considerations and future prospects

Organization of course

- Textbooks in AI show that the field has many branches that appear separate
- Too many branches exist for us to explore in a class, as a textbook does; too much detail
- *Solution:* to group the branches by approach and problem type; by features such as interaction, uncertainty, inference, embodiment, and planning

Objectives

Objectives are listed for each topic, specifying capabilities that a successful student will show during the semester

Following are *multi-topic objectives*:

9a. Distinguish stages in the development of AI research and application

9b. Carry out documented research on a topic in AI

9c. Work in a group, producing written results

9d. Make a presentation about part of the course material

Classroom format

- Emphasis is *discussion and interaction*
- *Slides and presentation* summarize content of the course
- We will ask each other questions
- Classroom is a focused professional environment
- *Participation* matters
- *Resources:* classroom, textbook, handouts, student research

Grades and assessment

- Objectives assessed by problem-solving quizzes determine 40% of grade
- Maximum quiz-question score attained for an objective determines the level recorded for final grade
- *Research paper* is in stages
- *Assignments* apply knowledge
- *Participation* is measured by questioning and responsiveness in class; not by attendance

Semester grade

Objectives attained	40	%
Assignments	10	
Multiple-choice quizzes	10	
Problem-solving quizzes	10	
Research paper	10	
Final exam	10	
Participation	10	
	100	%

Concepts

abstraction
action
adaptation
agent
algorithm
control theory
data structure
emergence
exhaustive search
function
generalization
good-old-fashioned AI
heuristic
inference
intelligence
interaction
knowledge
learning
natural language
NP-hard
partially observable
percept
physical symbol system hypothesis
probability
reasoning
satisfice
search
state-space search
symbol
theorem
toy problem
Turing test
voice recognition

References

George Luger. *Artificial Intelligence*. Addison Wesley, 2005.

Stuart Russell and Peter Norvig. *Artificial Intelligence : A Modern Approach, 2nd ed.* Prentice Hall, 2003.