

David Keil, Framingham State University

CSCI 300 Artificial Intelligence

Introduction

1. What this course offers
2. Course plan
3. How the course delivers

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

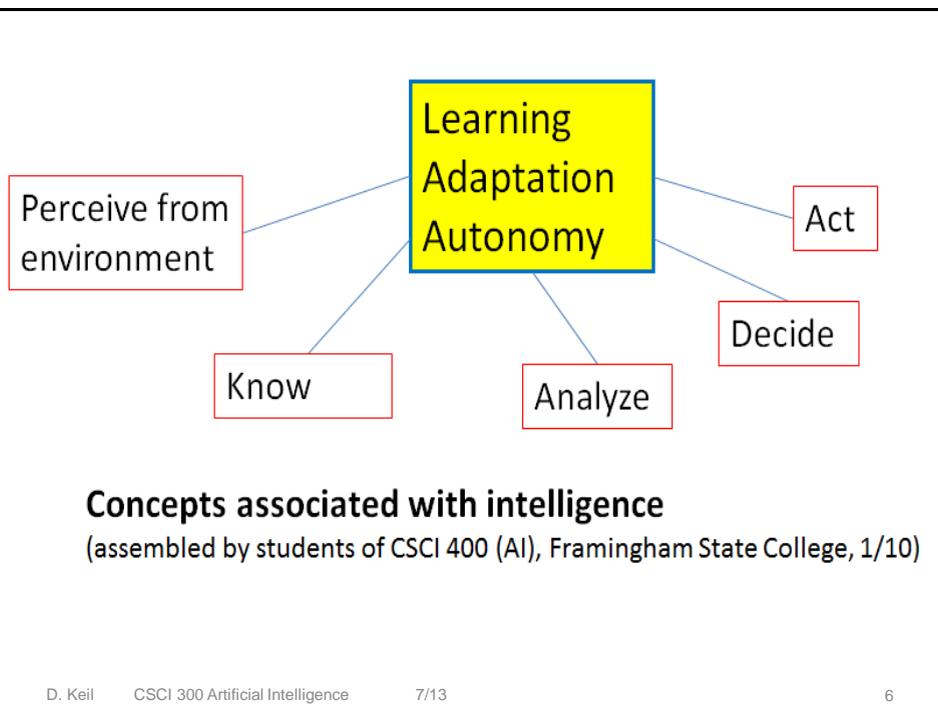
- May AI be summarized as *rational adaptive computational behavior*?
- What is intelligence about?
- *Key question: what do you need to know about AI, as a computer professional and as a citizen?*

1. What this course offers

- In which CS concentration are you?
- Other relevant background?
- What do you expect from this course?

Is intelligence required to...

- Play chess?
- Play Jeopardy?
- Jump from branch to branch?
- Play tennis?
- Play volleyball?



Domain of concern

- Is AI the set of problems that have ever been *considered* hard?
- *Is AI* - ?
 - Inference
 - Interaction
 - Pattern matching
 - Planning
 - Social
 - Imitating humans
 - Search for solutions
 - Reasoning
 - Using rules of thumb
 - Learning, adaptation
 - Solving intractable problems

What this course is about

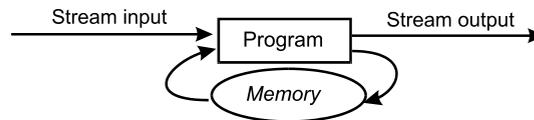
- We are studying here systems that *do the right thing* to solve *difficult computational problems*
- Our textbook suggests the concept of rational *agents* that take actions on *percepts* from their environments
- Environment types, e.g., *functional* or *interactive*, determine agent types
- We progress from problems with straightforward solutions to problems in more difficult environments

Three computing paradigms

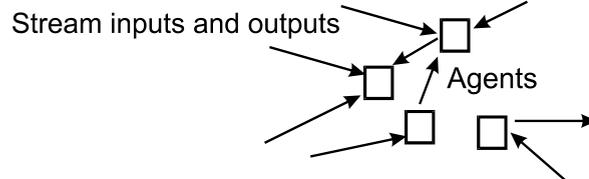
- Algorithmic



- Sequential interactive



- Multi-stream interactive



What are “difficult problems”?

Hard computational problems could mean

- *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

These are problems not efficiently solvable by traditional approaches to algorithm design, such as brute force, divide and conquer, greedy, or dynamic-programming

Is intelligence our topic?

- *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
- Topics 1 and 8 will discuss *mind*
- *Q*: Should this course be called, “Artificial Cognitive Systems”?

The discipline of 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
- The discipline called “AI” started in about the 1950s but has earlier origins

Four approaches to AI

AI systems are defined alternatively as ones that:

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

The 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

Russell-Norvig's definition of AI

- AI is “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...”
- *Questions:*
 - Doesn't *any program* act on inputs?
 - Can an agent's *actions* affect its later *percepts*?

AI as rational adaptive computational behavior

AI systems are:

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

Can problems be classified as AI or not-AI?

AI

Chess playing

Integrating algebraic expressions

Voice recognition

Finding concepts from examples

Rule-based inference

Not AI

Tic tac toe

Addition

Searching arrays

Sorting

Voice generation

Toy vs. real-world problems

- *Toy* (no one would buy a solution):
 - Vacuum world (move if clean, else suck dirt)
 - 8-puzzle (repeatedly move square piece to empty square in 3 x 3 grid)
 - 8-queens (chess; arrange 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

- Require use of *logical inference*
- Processing relates to *meaning*
- Solutions use *heuristics* (rules of thumb)
- *Intractable*, requiring suboptimal solutions
- Involve *interaction* and require *adaptation*
- Similarity to problems addressed by life forms

Early history of AI

- Neuron research by McCulloch and Pitts, 1943
- *Explosion of hope, 1950s and 1960s:*
 - Dartmouth workshop, 1956 (J. McCarthy, M. Minsky, C. Shannon, Rochester, A. Newell, H. Simon, A. Samuel)
 - LISP language
 - General Problem Solver; theorem provers
 - Blocks worlds
 - Perceptrons (neural nets)

Expert systems

- Knowledge-based *expert systems* separated *knowledge* (rules) from *reasoning* (inference)
- *Example:* Dendral (1969) inferred molecular structure from spectral data,
- Japanese *Fifth Generation* project and U.S. consortium (1980s) drew billions in investment for knowledge-based systems

Physical symbol system hypothesis

- “The necessary and sufficient condition for a physical system to exhibit general intelligent action is that it be a physical symbol system”
- A. Newell and H. Simon, 1976

Approaches to AI

- Physical symbol system hypothesis inspired:
 - Empirical and constructivist approach
 - Research on forms for representation and on search algorithms
 - A notion of disembodied cognition
- *Probabilistic and stochastic* approaches (*Markov models, Bayesian networks*) followed
- Rise of *intelligent-agent* model, mid-1990s, focused on interaction with environment

Constructivist approach

- An agent may construct its own knowledge from percepts
- Emphasis shifted from representation and search toward learning and adaptation
- Simple local adaptations enable a complex system to shape response to environment
- *Agent and evolutionary approach*: From independent action of agents, intelligent behavior emerges by natural selection

2. Course plan

1. 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

Multi-topic objectives

- 0a. Show knowledge of facts and concepts
- 0b. Summarize the semester's learning
- 0c. Carry out documented research on AI
- 0d. Participate in class activities throughout the semester
- 0e. Solve problems as part of a team
- 0f. Present results in the classroom

Background concepts

- 0.1a Recognize basic concepts of precalculus*
- 0.1b Write the truth table for a propositional-logic formula or logic circuit*
- 0.2 Design a looping algorithm*
- 0.3a Find a path in a graph*
- 0.3b Explain the relation between the logarithm function and the heights of trees*
- 0.4 Explain basic notions of combinatorics*

Objectives for topics 1-4

1. Compare human cognition with computational or agent models of perception-action
2. Explain how heuristics offer ways to pursue goals in exponentially large search spaces
3. Describe the representation and use of knowledge in inference-based problem solving
4. Apply probability theory to describe agents operating in uncertain environments

Objectives for topics 5-8

5. Describe different ways to supervise agents to learn and improve their behavior
6. Explain adaptive learning from the environment
7. Explain the relation between distributed artificial intelligence and emergent behavior
8. Defend a theory of mind, relating it to ethical issues raised by artificial cognitive systems

Summary objective

9. Distinguish stages in the development of artificial-intelligence research and applications

3. How the course will deliver

- What learning environments have been best for you?

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

A possible way to explore AI

- *Voice input* for Google on mobile phones
- *Siri interface* for iPhone
- *Programming a humanoid robot*, such as Nao, to interact with humans; e.g., children
- *Advantages of such an approach*: connects *theory to reality* and *human interaction*
- *Disadvantages*:
 - Would have to rewrite textbooks and slides
 - Would risk losing gains of past AI research

Course organization

- We have 8 *topics*
- Each topic has an *objective*
- Each topic has 3-5 *subtopics*
- Each subtopic has 1-2 *desired outcomes*
- Some outcomes are considered *core*
- A set of *problems* for each subtopic outcome is available
- Exercises and quizzes focus on these

Exercises, group work, and presentations

- *Exercises* consist of individual and group problem solving
- Each group *presents* one solution per topic
- I evaluate each group's solution to each subtopic problem
- Groups review and submit individuals' exercises, problem solutions for checkoff
- Group solutions and presentations count 15% of final grade; individual solutions 5%

Research paper

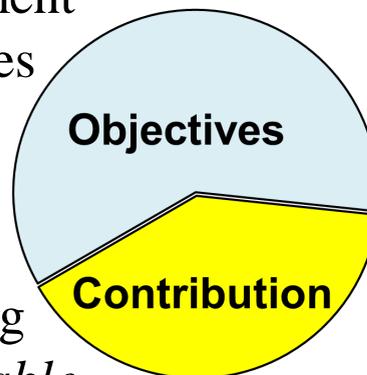
- You will do research that helps *you* define a kind of intelligence or cognition, or describe an application of the ideas in this course
- Everyone will report back a number of times in the semester about their research

Showing attainment of objectives

- After doing exercises on a topic, a student may show attainment of an outcome/objective by answering a quiz problem on it, in writing, in class
- 3-4 opportunities will be available per subtopic
- 60% of final assessment and grade is based on showing attainment of objectives and outcomes

Assessment and grading

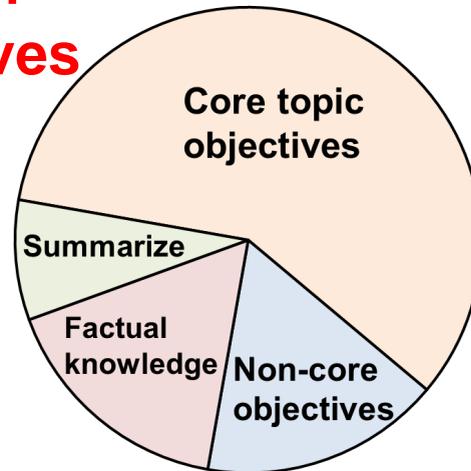
- *To measure:*
 - Individual achievement of learning objectives
 - Contribution to the learning of the class
- *Breakdown:* 60/40
- *Assumptions:* Learning is *shared* and *measurable*



Assessment of learning objectives

Assumptions:

- *Application of concepts is measurable via core and other topic objectives*
- *Facts about concepts matter*
- *We learn by summarizing and reflecting*



Assessment of contribution and participation

We assume that learning happens by:

- *Doing and sharing inquiry*
- *Being present*
- *Solving problems together*
- *Activity throughout the semester*



Final exam and summary quiz

- On final exam day, each student will be asked to present elements of her/his research or semester project
- During the last week of classes, we'll have a *summary quiz* of multiple-choice questions and multi-topic problems

Semester grade

Application of concepts	
core topic objectives	35
other topic objectives	10
Knowledge of facts	10
Summary and reflection	5
Written contribution	15
Presenting results in person	10
Group activity	10
Attendance	<u>5</u>
	100

References

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

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