

CSCI 400: Special Topics (Artificial Intelligence)

David Keil, Framingham State University, Fall 2013

SYLLABUS DRAFT

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.

Course description

An introduction to cognitive science and efforts to implement intelligence in computer hardware and software, with application areas. Topics include search, knowledge, reasoning, uncertainty, adaptation, learning, embodied interaction, future prospects, and philosophical considerations. Concepts will be demonstrated with software.

Prerequisites: Junior status as a Computer Science Major. Course will use many concepts from CSCI 271 Data Structures and MATH 292 Discrete Math I.

Required reading

- S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 2nd ed., Prentice Hall, 2003.
- Handout material, including slides: see course site.

Reading and studying text material related to the course are musts.

Meeting times

MWThF 12:30-1:20

Hemenway Hall 132 (annex)

To contact instructor:

Office hours (Hemenway Hall 318A):

M 2:30-3:30 p.m., W 9:30-10:30 a.m.,

Th 11:30-12:30 p.m., others by appointment

Office: Hemenway Hall 318A

Telephone: (508) 626-4724

Email: dkeil@framingham.edu

URL: www.framingham.edu/~dkeil

Course overview

Other possible names for the field of artificial intelligence are “Solving hard computational problems” and “Rational adaptive computational behavior.” By hard problems, we mean ones subject to combinatorial explosion in the number of possible solutions to consider, or *interactive* problems that require learning and adaptation. In this course, we emphasize interaction and the notion of *rational agents* – computational entities that take good actions in response to their percepts (inputs).

Our interest includes *models* of intelligent *agents* and of *systems* capable of intelligent behavior. Intelligence requires *adaptation* to an *environment* via *behavior* that is designed to change the environment. Intelligent behavior is most typically *interactive* and is often, *social* behavior of systems that comprise multiple actors.

Intelligent behavior is often associated with *algorithmic* problems that involve planning several interaction steps ahead, so that the number of possible situations to consider explodes the farther ahead we look. The task for the intelligent agent is to find a approximately optimal action among a huge number of possible actions

To counter the intractability of many optimization problems, an approach of *bounded rationality* is suggested, so that *satisficing* rather than *optimality* is the goal. Intelligent behavior, therefore, is an approximation of perfect behavior; it is “doing the right thing” in the sense of getting close enough to exactly the right thing. It is the time constraints of real-time interaction that forces the resort to satisficing.

AI seeks to model cognition, including knowledge acquisition, storage, and application. Tools include *logics* (propositional, predicate, modal and temporal), and *probability theory*.

We will refer to the related fields of neuroscience, cognitive science, psychology, philosophy, control theory, and decision theory.

The trend of AI research has evolved from *reasoning* in deterministic *toy environments* (e.g., board games), to *inference* using predicate logic, to

planning, uncertainty, and belief, to interacting with an environment, to coordination of multiple agents in dynamic environments.

Grades and classroom format

Please see the paper, "What we do in my classroom," attached, which is part of this course's syllabus. See especially guidelines there for assignments, grading, and collaboration.

What we will investigate

Students are asked to consider the following questions and to develop their own points of view.

1. Is the general problem for any cognitive system to give *timely responses to percepts from an environment*?
 2. Can intelligence be summarized as *reasoning; ontogenetic learning (single agent); phylogenetic learning (evolutionary); and sociogenetic learning (distributed)*?
 3. May AI be summarized as *rational adaptive computational behavior*?
 4. *Timely responses to percepts from an environment* is the general problem for any cognitive system.
 5. Can many problems solved using intelligence may be reduced to the *search* of an exponential-sized *state space*?
 6. Is it AI's task to give *approximate solutions to NP-hard problems*?
 7. What are ways in which knowledge is represented?
 8. How can knowledge be added to a knowledge base?
 9. Is *probabilistic reasoning* part of intelligence?
 10. What is the form of knowledge in partially observable and non-deterministic environments?
 11. What type of learning occurs in static environments?
 12. Is *adaptation* a higher form of intelligent behavior than generalization?
 13. Are *sociogenetic* (distributed) and *phylogenetic* (evolutionary) learning higher forms than *ontogenetic* (single-agent)?
 14. Are situated and embodied forms of intelligence more robust than reasoning-based systems?
 15. Is intelligence *emergent* and *decentralized*?
16. Do persistent, dynamic, physical environments pose problems solvable only with multi-agent systems and distributed AI?
 17. What is the role of indirect interaction in intelligence?
 18. Could clarity in research in cognitive science be served by relaxing the association of intelligence with humanness and by reframing AI and cognition as *rational adaptive behavior*?
 19. Does the notion of *bounded optimality* (a property of the best *program* achievable to solve a problem that entails adaptation) offer a sound theoretical foundation for AI research?

Learning objectives

Successful students will show that they can do the following:

- 1a. Describe the computational-representational understanding of mind
- 1b. Distinguish classes of problem environment
- 1c. Describe neural and transition-system models of computation
- 1d. Describe a reflex agent within the rational-agent model of AI
- 2a. Define goal-based state-space search
- 2b. Construct a game tree and perform a goal-driven analysis of it
- 2c. Explain how heuristics are used to provide adequate solutions to hard search problems
- 2d. Explain what hard computational problems means formally and in practice
- 3a. Use inference in propositional and predicate logic
- 3b. Describe methods of representing and using knowledge
- 3c. Distinguish knowledge-based from goal-driven agents
- 4a. Describe ways to operate under conditions of uncertain knowledge
- 4b. Use some concepts of probability theory
- 4c. Demonstrate Markov models or Bayesian reasoning
- 5a. Describe and use methods of symbol-based *supervised learning*
- 5b. Describe the *connectionist* approach to AI
- 5c. Describe evolutionary computation
- 5d. Explain concepts of natural-language processing
- 6a. Identify problems that require adaptation
- 6b. Describe bounded rationality

- 6c. Describe methods of reinforcement learning
- 6d. Explain features of robotic systems
- 7a. Describe multi-agent solution approaches
- 7b. Explain the relationship of intelligence to indirect interaction, self organization, and emergent behavior
- 8a. Discuss philosophical issues raised by the notion of artificial intelligence
- 8b. Discuss the weak and strong AI theses, and future prospects for AI
- 8c. Explain the notion of bounded optimality
- 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

Research paper

Each student will propose in writing a topic for a short research paper; submit a preliminary version that responds to comments about the proposal; present the paper in class; and submit a final version at the end of the course. Subject to proper formatting, textual presentation, and documentation, papers will be made available on the Web.

Grading weights

Objectives attained	35 %
Participation	10
Assignments	10
Research paper	10
Problem-solving quizzes	10
Multiple-choice quizzes	10
Final exam and summary quiz	<u>15</u>
	100

Research interests of instructor

The instructor's research includes investigation of the power of multi-stream, multi-agent, and indirect interaction. See instructor's web site for publications related to AI and indirect interaction.

Accommodations

"Students with disabilities who request accommodations are to provide Documentation Confirmation from the Office of Academic Support within the first two weeks of class. Academic Support is located in the Center for Academic Support and Advising (CASA). Please call (508) 626-4906 if you have questions or if you need to schedule an appointment." (See <http://www.framingham.edu/CASA/Accommodations/accomm.htm>.)

Course Plan

Dates	Topic	Russell-Norvig chapters
9/7 - 9/15	<i>Introduction</i>	Appendices A-B
9/16 – 9/21	1. Cognition and computation	1-2; handouts ¹
9/22 – 9/26	2. State-space search	3-6
10/5	<i>Problem-solving quiz on topics 1-2</i>	
10/3 – 10/7	3. Knowledge representation and rule-based inference	7-11; 12.1-12.2
10/10 – 10/17	4. Uncertainty and probabilistic reasoning	12.3-12.6; 13-17
10/19	<i>Research reports (abstracts)</i>	
10/20	<i>Problem-solving quiz on topics 3-4</i>	
10/24 – 10/28	5. Supervised learning	18-20; 23-24
10/31	<i>Make-up quizzes on topics 1-4</i>	
11/2 – 11/10	6. Reinforcement learning and adaptation	17.5; 21-22
11/11	<i>Research reports (preliminary drafts)</i>	
11/14	<i>Problem-solving quiz on topics 5-6</i>	
11/16 – 11/21	7. Distributed AI and multi-stream adaptive interaction	12.7; 17.6; 25; handouts ^{2,3}
11/23 – 11/30	8. Philosophical considerations and future prospects	26-27; handouts ^{4,5,6}
12/1 – 12/14	Summary	
12/2	<i>Research reports (final drafts)</i>	
12/5	<i>Problem-solving quiz on topics 7-8</i>	
Fri., 12/16, 8:00 – 11:00am	<i>Final exam (Topics 1-8)</i>	

¹ Thagard, pp. 3-15; Simon, pp. 675-691; Weizenbaum pp. 128-131.

² D. Keil, Indirect interaction in evolving adaptive multi-agent systems (2006); Decentralization and stigmergy (2008);

³ D. Keil and D. Goldin, Models of Self-Organizing Systems: Learning from the Ants and the Bees (2007)

⁴ D. Keil, Notes for debate on AI (2005)

⁵ J. Markoff, A fight to win the future: Computers vs. humans

⁶ R. Kurzweil, The coming merging of mind and machine (2002).