

Intro to Artificial Intelligence CSC384h, Winter 2009 Hojjat Ghaderi

- Contact
 - hojjat@cdf.toronto.edu, Office: SF3208
 - Office Hours: Mon and Wed 12–1:30pm in SF3207
- Lectures:
 - Monday, Wednesday, & Friday 11–12pm SS2110.
- Tutorials
 - Usually no tutorials, instead examples and question answering will be part of the lectures.
 - May have a Prolog tutorial if needed.

CSC384: Prerequisites ...


- Prerequisites will not be checked for this course, except for the CGPA (cumulative grade point average) condition.
 - You don't need to request a waiver.
 - You should have a stats course either the standard STA 247/255/257 or at least something like STA 250.
 - You need to have some familiarity with Prolog, CSC324 is the standard prerequisite. We will provide 3–4 tutorials on Prolog.
 - In all cases if you do not have the standard prerequisites **you will be responsible** for covering any necessary background on your own. We can't provide any assistance with this.

CSC384h: Textbook

- Recommended Text:

Artificial Intelligence: A Modern Approach by Stuart Russell and Peter Norvig. 2nd Edition, 2003.

 - 2 copies of are on 24hr reserve in the Engineering and Computer Science Library.
 - Recommended but not required.
 - Lecture notes cover much of the course material and will be available online before class. Print a copy and bring to the class!
- Additional Reference:
 - **Computational Intelligence: A Logical Approach** by David Poole, Alan Mackworth and Randy Goebel.



CSC384: Website

- Course web site

<http://www.cs.toronto.edu/~hojjat/csc384h/>

 - It will be the primary source of more detailed information, announcements, etc.
 - Check the site often (at least every one or two days).
 - Updates about assignments, clarifications etc. will also be posted on the web site.
- Course bulletin board (will not be moderated)

<https://csc.cdf.toronto.edu/bb/YaBB.pl?board=CSC384H1S>

CSC384: E-mail/board policies...

- The course bulletin board will not be moderated.
- It can be used to communicate with your fellow students.
- Do not send questions there that you want answered by the instructor. Send e-mail directly.
- For each assignment, a TA will be assigned to answer questions. **Please send your questions about each assignment to the TA.**
- Answers that would be important to everyone will be posted to the web site.
- Send only Plain Text (no HTML/MIME) using [your CDF accounts](#).
- Start the subject of all your emails with "[CSC384]". (Please see <http://www.cs.toronto.edu/~hojjat/csc384h/contactpolicy.htm>)
- A silent policy will take effect 24 hours before each assignment is due, i.e. no question related to the assignment will be answered during this period.

CSC384: Important Dates (tentative)

- Mon Jan 12, Assignment 1 out
- Sun Jan 18, Last Day to Add the course
- Mon Feb 2, A1 due, A2 out.
- Mon Feb 9, **Term Test 1**
- Feb16-20 Reading week
- Mon Feb 23 A2 due, Assignment 3 out.
- Mon Mar 2 **Term Test 2**
- Sun Mar 08 Last Day to Drop the course
- Wed Mar 11 Assignment 3 due, Assignment 4 out.
- Mon Mar 23 **Term Test 3**
- Mon Apr 6 Assignment 4 due.
- Wed Apr 8 **Term Test 4** (*last day of course*)
- Thr Apr 9 End of winter classes

CSC384: How you will be graded

- Course work
 - 4 Assignments (mostly programming, some short answer)
 - 4 term tests (50 mins each).
 - No final exam
 - Assignments are worth a total of 40%: [10% each].
 - Term tests are worth a total of 60%: [15% each].
- Late Policy/Missing Test
 - No late assignment is accepted.
 - Missed Test/Assignment with a medical excuse will be given a mark based on the student's and the class's performance on all tests/assignments.
- Plagiarism (handing of work not substantially the student's own) will not be tolerated!

More@ <http://www.cs.toronto.edu/~fpitt/documents/plagiarism.html>

What is Artificial Intelligence?

- Chapter 1 of textbook.
- What is **AI**?
- What is **intelligence**?
- What features/abilities do humans (animals? animate objects?) have that you think are indicative or characteristic of intelligence?
- Who is AI community? How do they define AI?

Webster says: a. the capacity to acquire and apply knowledge. b. the faculty of thought and reason. ...



Subareas of AI

- Perception: vision, speech understanding, etc.
- Machine Learning, Neural network
- Robotics
- Natural language understanding
- Reasoning and decision making (**our focus**)
 - Decision making (*search, planning, decision theory*)
 - Knowledge representation
 - Reasoning (*logical, probabilistic*)

CSC384: Intro to Artificial Intelligence

- Knowledge Representation (**Chapter 7–10**)
 - First order logic for more general knowledge.
 - Knowledge represented in declarative manner.
- Planning (**Chapter 11–12**)
 - Predicate representation of states.
 - Planning graph

What We Cover In CSC384

- Search
 - Heuristic Search. (**Chapter 3,4**)
 - Search spaces
 - Heuristic guidance
 - Backtracking Search (**Chapter 5**)
 - “Vector of features” representation
 - Case analysis search.
 - Game tree search (**Chapter 6**)
 - Working against an opponent.

CSC384: Intro to Artificial Intelligence

- Uncertainty (**Chapter 13–16**)
 - Probabilistic reasoning, Bayes networks
 - Utilities and influence diagrams.
 - Temporal probabilistic reasoning.

Further Courses in AI

- CSC321H “Introduction to Neural Networks and Machine Learning”
- CSC401H1 “Natural Language Computing”
- CSC411H “Machine Learning and Data Mining”
- CSC412H1 “Uncertainty and Learning in Artificial Intelligence”
- CSC420H1 “Introduction to Image Understanding”
- CSC485H1 “Computational Linguistics”
- CSC486H1 “Knowledge Representation and Reasoning”
- CSC487H1 “Computational Vision”

Human intelligence

- Is imitating humans the goal?
- Pros?
- Cons?

AI Definitions by Russell + Norvig

	Like humans	Not necessarily like humans
Think	Systems that think like humans	Systems that think rationally
Act	Systems that act like humans	Systems that act rationally

Human intelligence

- The Turing Test:
 - A human interrogator. Communicates with a hidden subject that is either a computer system or a human. If the human interrogator cannot reliably decide whether or not the subject is a computer, the computer is said to have passed the Turing test.
- Turing provided some very persuasive arguments that a system passing the Turing test is intelligent.
- However, the test does not provide much traction on the question of how to actually build an intelligent system.

Human intelligence

- In general there are various reasons why trying to mimic humans might not be the best approach to AI:
 - Computers and Humans have a very different architecture with quite different abilities.
 - Numerical computations
 - Visual and sensory processing
 - Massively and slow parallel vs. fast serial

	Computer	Human Brain
Computational Units	1 CPU, 10^8 gates	10^{11} neurons
Storage Units	10^{11} bits RAM 10^{12} bits disk	10^{11} neurons 10^{14} synapses
Cycle time	10^{-9} sec	10^{-3} sec
Bandwidth	10^{10} bits/sec	10^{14} bits/sec
Memory updates/sec	10^9	10^{14}

Rationality

- The alternative approach relies on the notion of **rationality**.
- Typically this is a precise mathematical notion of what it means to *do the right thing* in any particular circumstance. Provides
 - A precise mechanism for analyzing and understanding the properties of this ideal behavior we are trying to achieve.
 - A precise benchmark against which we can measure the behavior the systems we build.

Human intelligence

- But more importantly, we know very little about how the human brain performs its higher level processes. Hence, this point of view provides very little information from which a scientific understanding of these processes can be built.
- However, Neuroscience has been very influential in some areas of AI. For example, in robotic sensing, vision processing, etc.


Rationality

- Mathematical characterizations of rationality have come from diverse areas like logic (laws of thought) and economics (utility theory how best to act under uncertainty, game theory how self-interested agents interact).
- There is no universal agreement about which notion of rationality is best, but since these notions are precise we can study them and give exact characterizations of their properties, good and bad.
- We'll focus on acting rationally
 - this has implications for thinking/reasoning

Computational Intelligence

- *AI tries to understand and model intelligence as a computational process.*
- Thus we try to construct systems whose **computation** achieves or approximates the desired notion of rationality.
- Hence AI is part of Computer Science.
 - Other areas interested in the study of intelligence lie in other areas or study, e.g., cognitive science which focuses on human intelligence. Such areas are very related, but their central focus tends to be different.

Agent Schematic (I)



- This diagram oversimplifies the internal structure of the agent.

Agency

- It is also useful to think of intelligent systems as being **agents**, either:
 - with their own goals
 - or that act on behalf of someone (a “user”)
- An *agent* is an entity that exists in an *environment* and that *acts* on that environment based on its *perceptions* of the environment
- An *intelligent agent* acts to further its own interests (or those of a user).

Agent Schematic (II)



- Require more flexible interaction with the environment, the ability to modify one’s goals, knowledge that be applied flexibly to different situations.

Degrees of Intelligence

- Building an intelligent system as capable as humans remains an elusive goal.
- However, systems have been built which exhibit various specialized degrees of intelligence.
- Formalisms and algorithmic ideas have been identified as being useful in the construction of these “intelligent” systems.
- Together these formalisms and algorithms form the foundation of our attempt to understand intelligence as a computational process.
- *In this course we will study some of these formalisms and see how they can be used to achieve various degrees of intelligence.*

Next Time

- Start talking about search. See web page for readings and slides.
- If you are not comfortable with Prolog take a look at the web page “Prolog Resources Link”. There is a simple set of Prolog examples, a link to an on-line tutorial, and other material than you can use.
 - *You are expected to know enough Prolog (or be able to learn enough Prolog) to do the assignments.*

AI Successes

- In 1997 Deep Blue defeated world chess champion Garry Kasparov in six games.
 - But Deep Blue can still make novice mistakes!
- World champion checkers player Chinook.
- World champion Backgammon player learned how to play.
- In 1999, a NASA AI agent ran a satellite beyond Mars for over a day, without ground control.
-