



Introduction to Artificial Intelligence

Unit # 1



Today's Agenda

- Course Contents
- Grading Policy
- Projects
- Overview of the Main Topics
- Introduction to Artificial Intelligence

Course Outline

- Overview of Artificial Intelligence
- State Space Representation
- Search Techniques
- AI in Games
- Machine Learning
- Propositional and Predicate Logic
- Probabilistic Reasoning
- Evolutionary Algorithms
- Computer Vision
- Robotics

Useful Information

- **Course Website**
- **Text Book**
 - Tim Jones, *Artificial Intelligence: A Systems Approach*, 2007.
- **Reference Books**
 - Ben Coppin, *Artificial Intelligence Illuminated*, 2004.
 - Kevin Korb and Ann Nicholson, *Bayesian Artificial Intelligence*, 2003
 - Steven Rabin, *AI Game Programming Wisdom 3*, 2005.
 - Steve Rabin, *AI Game Programming Wisdom 4*, 2008

Software Tools

- SWI-Prolog (<http://www.swi-prolog.org/>)
- GeNIe (<http://genie.sis.pitt.edu/>)
- Weka (<http://www.cs.waikato.ac.nz/ml/weka/>)
- KNIME (<http://www.knime.org/>)
- For programming assignments, you can use any standard programming language (either Java, C#, C++, etc.).

Acknowledgement

- Most of the slides of this lecture have been taken from the lecture slides of CS307 – “Introduction to Artificial Intelligence” by Dr. Sajjad Haider.

Marks Distribution (Tentative)

- Term Exam – 30 % (15 + 15)
- Final – 40 %
- Assignment - 7.5 %
- Quiz - 7.5 %
- Project – 15%

Basic Questions

(Taken from <http://www-formal.stanford.edu>)

- **What is artificial intelligence?**
- It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence.
- **Yes, but what is intelligence?**
- Intelligence is the computational part of the ability to achieve goals in the world OR “Ability to make the right decisions given set of inputs and a variety of possible actions.”
Varying kinds and degrees of intelligence occur in people, many animals and some machines.
- **Isn't there a solid definition of intelligence that doesn't depend on relating it to human intelligence?**
- Not yet. The problem is that we cannot yet characterize in general what kinds of computational procedures we want to call intelligent. We understand some of the mechanisms of intelligence and not others.

Basic Questions (Cont'd)

- **Isn't AI about simulating human intelligence?**
- Sometimes but not always or even usually. On the one hand, we can learn something about how to make machines solve problems by observing other people or just by observing our own methods. On the other hand, most work in AI involves studying the problems the world presents to intelligence rather than studying people or animals. AI researchers are free to use methods that are not observed in people or that involve much more computing than people can do.
- **Does AI aim at human-level intelligence?**
- Yes. The ultimate effort is to make computer programs that can solve problems and achieve goals in the world as well as humans. However, many people involved in particular research areas are much less ambitious.

Basic Questions (Cont'd)

- **What about IQ? Do computer programs have IQs?**
- No. IQ is based on the rates at which intelligence develops in children. It is the ratio of the age at which a child normally makes a certain score to the child's age. The scale is extended to adults in a suitable way. IQ correlates well with various measures of success or failure in life, but making computers that can score high on IQ tests would be weakly correlated with their usefulness. For example, the ability of a child to repeat back a long sequence of digits correlates well with other intellectual abilities, perhaps because it measures how much information the child can compute with at once. However, "digit span" is trivial for even extremely limited computers.

Intelligence in machines vs humans

What has happened over the last 40-50 years - to the disappointment of all those who made breathless predictions about where AI would go - is that things such as playing chess have turned out to be incredibly easy for computers, whereas learning to walk and learning to get around in the world without falling over has proved to be unbelievably difficult.

Strong vs. Weak AI

- Strong AI refers to the field of research that is interested in making computers think at a level equal to humans
- Weak AI represents the wider domain of AI technologies. Weak AI features can be added to systems to give them intelligent quantities.
- **The course focuses on weak AI.**

Strong vs. Weak AI

- The followers of **strong AI** believe that by giving a computer program sufficient processing power, and by providing it with enough intelligence, one can create a computer that can literally think and is conscious in the same way that a human is conscious.
- Many philosophers and Artificial Intelligence researchers consider this view to be false, and even ludicrous.
- **Weak AI**, in contrast, is simply the view that intelligent behavior can be modeled and used by computers to solve complex problems.

Philosophical vs. Practical AI


- Philosophical:
 - Can machines think, in principle?
 - Will machine thought be different from human thought?
- Practical
 - Collection of techniques to automatically solve particular problems that require “intelligence” (whatever that is)

Philosophical vs. Practical AI

- Philosophical AI Questions
 - Can we make something that is as intelligent as a human?
 - Can we make something that is as intelligent as a bee?
 - Can we get something that is really evolutionary and self-improving and autonomous and flexible....?
- Practical AI Questions
 - Can we save this plant \$20 million a year by improved pattern recognition?
 - Can we save this bank \$50 million a year by auto fraud detection?
 - Can we start a new industry of handwriting recognition / automated negotiation / helpdesks/?

Turing Test


- One or more human judges interview computers and human foils using terminals (so that the judges won't be prejudiced against the computers for lacking a human appearance).
- The nature of the dialog between the human judges and the candidates (i.e., the computers and the human foils) is similar to an online chat using instant messaging.
- The computers as well as the human foils try to convince the human judges of their humanness. If the human judges are unable to reliably unmask the computers (as imposter humans) then the computer is considered to have demonstrated human-level intelligence



Loebner Prize

Home Page of The Loebner Prize in Artificial Intelligence

"The First Turing Test"



- Read the following story on 2008 Loebner Prize:
http://technology.timesonline.co.uk/tol/news/tech_and_web/article4934858.ece

Critiques on Turing test

- It is widely believed today that even if a computer could pass the Turing test, it would still not truly be conscious or intelligent in the way that humans are.
- There is a need to start small and grow intelligence.
- It is more important to study the underlying principle of intelligence than to duplicate an exemplar.
- Aeronautical engineering tests do not define the goal of their field as making "machines that fly so exactly like pigeons that they can fool even other pigeons".
- The quest for 'artificial flight' succeeded when the Wright brothers and others stopped imitating birds and started using wind tunnels and learning about aerodynamics.

More Challenges

14 February 2011 Last updated at 09:31 GMT



IBM supercomputer challenges humans on TV quiz

By Maggie Shields

Technology reporter, BBC News, Silicon Valley

A supercomputer, designed by IBM, is to face two human contestants on the US quiz show Jeopardy.



IBM machines have previously taken on chess players

Artificial Intelligence Lab, IBA

Spring 2011

19

History of AI

- **1950s – The Birth of AI**
 - 1950: Alan Turing publishes "Computing Machinery and Intelligence."
 - 1956: John McCarthy coins the term, "Artificial Intelligence" at a Dartmouth computer conference.
 - 1958: John McCarthy invents the Lisp language, an AI programming language, at Massachusetts Institute of Technology (MIT).
- **1960s – The Rise of AI**
 - 1965: Joseph Weizenbaum builds ELIZA, an interactive program that carries on a dialogue in English on any topic (MIT).
 - Arthur's Clark book titled "2001: A Space Odyssey" – HAL, an intelligent computer onboard a Jupiter-bound spacecraft, murdered most of the crew out of paranoia over its own survival.
 - Knowledge representation was a strong theme during the 1960s, as strong AI continued to be a primary theme in AI research.

Artificial Intelligence Lab, IBA

Spring 2011

20

History of AI (Cont'd)

- 1970s – The Fall of AI
 - The 1970s represented the fall of AI after an inability to meet irrational expectations.
- 1990s to Today – AI Rises Again
 - Advancement in probabilistic reasoning, machine learning, intelligence agents, and computational intelligence systems.
 - 1997: IBM computer Deep Blue beats world champion Garry Kasparov in chess match.

Prophecies

- I believe that in about fifty years' time it will be possible to programme computers...to make them play the imitation game so well that an average interrogator will not have more than 70 percent chance of making the right identification after five minutes of questioning.
- — *Alan Turing, Computing machinery and intelligence, 1950.*



Prophecies (Cont'd)

- In from three to eight years, we'll have a machine with the general intelligence of an average human being.
- — *Marvin Minsky to Life magazine, 1970*

Prophecies (Cont'd)

- We can expect computers to pass the Turing test, indicating intelligence indistinguishable from biological humans, by the end of the 2020s.
- — *Ray Kurzweil, The Singularity is Near, 2005*

AI in Movies



Artificial Intelligence Lab, IBA

Spring 2011

25

DARPA Urban Challenge

- The DARPA Grand Challenge is a prize competition for driverless cars.
- The 2007 Challenge was held in California.
 - 96 KM course
 - 6 hours time duration
 - Rules included obeying all traffic regulations while negotiating with other traffic and obstacles and merging into traffic.
 - The \$2 million winner was Tartan Racing, a collaborative effort by CMU and GM.



Artificial Intelligence Lab, IBA

Spring 2011

26



The Dream

We proposed that the ultimate goal of the RoboCup Initiative to be stated as follows:

“By mid-21st century, a team of fully autonomous humanoid robot soccer players shall win the soccer game, comply with the official rule of the FIFA, against the winner of the most recent World Cup.”


Few research areas are:

- Agent architecture in general
- Computer Vision
- Combining reactive approaches and modeling/planning approaches
- Real-time recognition, planning, and reasoning
- Reasoning and action in a dynamic environment
- Sensor fusion
- Multi-agent systems in general
- Behavior learning for complex tasks
- Strategy acquisition
- Cognitive modeling in general



Ms Pac-Man Competition

- The aim of this competition is to provide the best software controller for the game of Ms Pac-Man. This is a great challenge for computational intelligence, machine learning, and AI in general.
- The mode of interaction is as follows: about 15 times per second your program will be sent a pixel map of the Ms. Pac-Man window, and it then responds with an integer indicating the direction of the joystick.
- <http://www.grappa.univ-lille3.fr/~loth/pacman.mpeg>



AI is everywhere

- Fuel injection systems in our cars use learning algorithms. Jet turbines are designed using genetic algorithms, which are both examples of AI, says Dr Rodney Brooks, the director of MIT's artificial intelligence laboratory.
- Every cell phone call and e-mail is routed using artificial intelligence, says Ray Kurzweil, an AI entrepreneur and the author of two books on the subject, *The Age of Intelligent Machines* and *The Age of Spiritual Machines*.
- Idea of ubiquitous computing is getting popular. More formally Ubiquitous computing is defined as "machines that fit the human environment instead of forcing humans to enter theirs."
- There are already autonomous lawnmowers that will wander around gardens all afternoon. The next step might be autonomous vacuum cleaners inside the house (though clutter and stairs present immediate problems for wheeled robots).
- These are all sorts of other uses for artificial animals in areas where people find jobs dangerous or tedious - land-mine clearance, toxic waste clearance, farming, mining, demolition, finding objects and robotic exploration, for example. Any jobs done currently or traditionally by animals would be a focus. We are familiar already from the Mars Pathfinder and other examples that we can send autonomous robots not only to inhospitable places, but also send them there on cheap one-way 'suicide' missions. (Of course, no machine ever 'dies', since we can restore its mind in a new body on earth after the mission.)
- In the future, drunks will be able to use cars, which will take them home like loyal horses. And not just drunks, but children, the old and infirm, the blind, all will be empowered.

Summary

- Intelligence is difficult to define, and as a result Artificial Intelligence is also hard to define.
- One definition of Artificial Intelligence is:
 - *Artificial intelligence is the study of systems that act in a way that to any observer would appear to be intelligent.*
- We are looking into the practical aspects of AI that focuses on getting inspirations from natural processes to solve real-life complex problems.