

15-780: Graduate AI

Introduction

Instructors:

Tuomas Sandholm and Manuela Veloso

TAs:

John Dickerson and Prateek Tandon

<http://www.cs.cmu.edu/~15780>

Carnegie Mellon University

Evaluation

- 4 Homeworks 40%
- Midterm Test 20%
- Final Project 30%
- Class participation 10%

- Homeworks are done individually
- 8 total “mercy” days (no penalty) for late homeworks - cannot use more than 3 mercy days in a single homework. No credit for late homeworks with no mercy days.
- Final projects can be done in groups of up to 2 people

Resources

- Lectures
 - Presentation and discussion in class
 - Lecture slides
 - Textbook: Russell and Norvig, 3rd edition
- Instructors – office hours by appointment
- TAs – office hours will be announced

More Information


- TAs will hold recitations as needed
 - Doodle poll for best time
- Audits
 - Register by fill out audit form
 - Must do final project
- Waitlist
 - Let us know if you are on it
- Mailing list
 - 15780students @ cs.cmu.edu
 - We will send a test message

What is *Artificial Intelligence*?

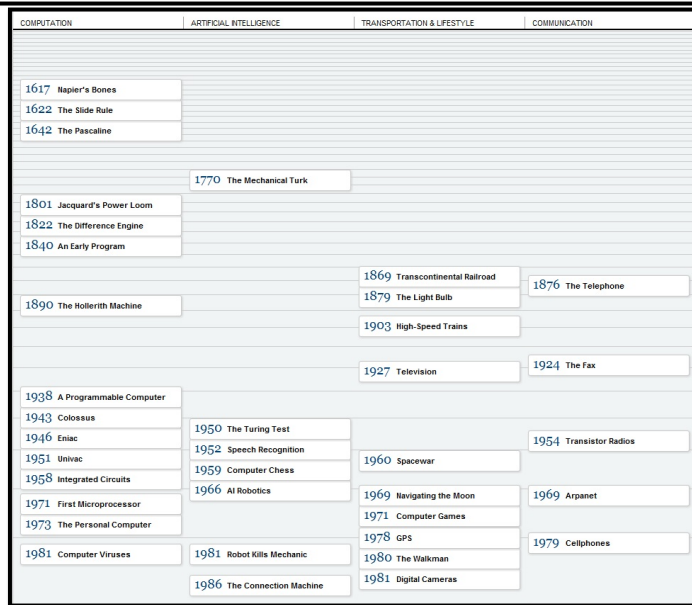
What is “intelligence” ?

Can we emulate
intelligent behavior in
machines ?

How far can we take it ?



Time Line (NYTimes 2011)



1950 The Turing Test

Alan Turing proposes the Turing test to decide if a computer is exhibiting intelligent behavior.



- Can a general machine think?
- Can machines “pass” the *imitation game*?
- Turing Test: Evaluator communicates via text channel with computer and human, must reliably identify the computer

9

Views of AI



Think like humans Cognitive Science	Think rationally Formalize inference into laws of thought
Act like humans Turing test	Act rationally Act according to laws

The Dartmouth Conference

“We propose that a two-month, ten-man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, NH. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.”

The Proponents

- John McCarthy, assistant professor of mathematics at Dartmouth (2011)
- Marvin Minsky, Harvard junior fellow in mathematics and neurology (MIT)
- Nathaniel Rochester, manager of information research at IBM, NY (2001)
- Claude Shannon, information theory, mathematician at Bell Labs (2001)



The Invited

- Trenchard More, IBM
- Arthur Samuel, IBM
- Oliver Selfridge, Lincoln Labs, MIT
- Ray Solomoff, MIT

And “two vaguely known persons from RAND and Carnegie Tech... a significant afterthought.” (Pamela McCorduck, “Machines Who Think”, page 94)

Herbert A. Simon and Allen Newell



1959 Computer Chess

Arthur Samuel's checkers program wins games against the best human players. 48 years later, the game of checkers is solved by computers.

Computer Chess at CMU

- A long history
- HiTech – Hans Berliner, Murray Campbell
- ChipTest – Feng-Hsiung Hsu
- Deep Thought – Hsu, Campbell, Anantharaman
- Deep Blue – CMU team hired by IBM Research

1997 Deep Blue

I.B.M.'s Deep Blue computer defeats Garry Kasparov, the world chess champion. Deep Blue can evaluate 200 million chess positions per second.



Jack Manning/The New York Times


Problem Solving

- Allen Newell and Herb Simon – 1950s
 - Given:
 - an initial **state**
 - a set of **actions**
 - a **goal** statement
 - Search for a **plan**, a sequence of actions that transform the initial state into a state where the goal is satisfied

Intelligent Systems


Three key steps (Craik, 1943):

1. the stimulus must be translated into an internal representation
2. the representation is manipulated by cognitive processes to derive new internal representations
3. internal representations are translated into action



1966 AI Robotics


The Stanford Research Institute creates Shakey, the first mobile robot that can reason about its surroundings. Five years later, funding is canceled when the shortcomings of the machine become apparent. Also in 1966, Joseph Weizenbaum creates Eliza, a conversational program intended to mimic a human therapist typing at a computer terminal.



Robert Reinhold/The New York Times

Intelligent Agents

- **Sensing:** vision, hearing, touch, smell
- **Cognition:** think, reason, plan, learn
- **Action:** motion, speech, manipulation




Perception – Sensors to State


- Sensors – “signal” (data) collectors from the physical world:
 - Vision, sound, touch, sonar, laser, infrared, GPS, temperature,....
- *Signal-to-symbol* challenge:
 - Recognize the *state* of the environment
 - ...wall at 2m... door on the left... green light... person in front... personX entering the room... ball at 1m and 30° East...

Learning

- Automatically generate strategies to classify or predict from training examples



Training data: Example images of object



Classification: Is the object present in the input image, yes/no?

Learning

- Automatically generate strategies to classify or predict from training examples

mpg	cylinders	displacement	horsepower	weight	acceleration	modelyear	maker
good	4	low	low	low	high	75to78	asia
bad	6	medium	medium	medium	medium	70to74	america
bad	4	medium	medium	medium	low	75to78	europa
bad	8	high	high	high	low	70to74	america
bad	6	medium	medium	medium	medium	70to74	america
bad	4	low	medium	low	medium	70to74	asia
bad	4	low	medium	low	low	70to74	asia
bad	8	high	high	high	low	75to78	america
:	:	:	:	:	:	:	:
:	:	:	:	:	:	:	:
:	:	:	:	:	:	:	:
bad	8	high	high	high	low	70to74	america
good	8	high	medium	high	high	79to83	america
bad	8	high	high	high	low	75to78	america
good	4	low	low	low	low	79to83	america
bad	6	medium	medium	medium	high	75to78	america
good	4	medium	low	low	low	79to83	america
good	4	low	low	medium	high	79to83	america
bad	8	high	high	high	low	70to74	america
good	4	low	medium	low	medium	75to78	europa

europa
75to78
medium
medium
medium
5 | medium

➔ Mpg good/bad

Predict mpg on new data

Training data: good/bad mpg for example cars

Web Search

[A.I. Artificial Intelligence \(2001\) - IMDb](http://www.imdb.com/title/tt0212720/)

www.imdb.com/title/tt0212720/

★★★★★ Rating: 7/10 - 153782 votes

A highly advanced robotic boy longs to become "real" so that he can regain the love of his human mother.

Directed by Steven Spielberg. Starring Haley Joel Osment, Jude Law.

Full cast and crew - Frances O'Connor - Synopsis - Trivia

[A.I. Artificial Intelligence - Wikipedia, the free encyclopedia](http://en.wikipedia.org/wiki/A.I._Artificial_Intelligence)

en.wikipedia.org/wiki/A.I._Artificial_Intelligence

A.I. Artificial Intelligence, also known as **A.I.**, is a 2001 science fiction drama film written, directed, and produced by Steven Spielberg, and based on Brian ...

Plot - Cast - Production - Soundtrack

[Artificial intelligence - Wikipedia, the free encyclopedia](http://en.wikipedia.org/wiki/Artificial_intelligence)

en.wikipedia.org/wiki/Artificial_intelligence

Artificial intelligence (AI) is the intelligence of machines and robots and the branch of computer science that aims to create it. **AI** textbooks define the field as "the ...

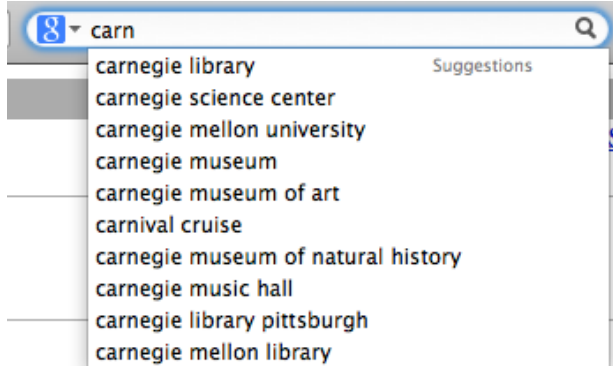
[AI Artificial Intelligence - Applications of artificial - List of artificial intelligence](#)

[AI - definition of AI by the Free Online Dictionary, Thesaurus and ...](http://www.thefreedictionary.com/AI)

www.thefreedictionary.com/AI

(Life Sciences & Allied Applications / Veterinary Science) artificial insemination. 2. (Electronics & Computer Science / Computer Science) artificial intelligence ...

Statistical Learning



Recommender Systems

Frequently Bought Together



Price For All Three: **\$40.22**

[Add all three to Cart](#) [Add all three to Wish List](#)


[Show availability and shipping details](#)

- This item:** Alan Turing: The Enigma The Centenary Edition by Andrew Hodges Paperback **\$15.54**
- Codebreakers: The Inside Story of Bletchley Park by F. H. Hinsley Paperback **\$13.15**
- The Code Book: The Science of Secrecy from Ancient Egypt to Quantum Cryptography by Simon Singh Paperback **\$11.53**

Customers Who Bought This Item Also Bought

<p>The Annotated Turing: A Guided Tour Through Alan Turing's Greatest Works → Charles Petzold ★★★★☆ (24) Paperback \$18.09</p>	<p>Codebreakers: The Inside Story of Bletchley Park F. H. Hinsley ★★★★☆ (18) Paperback \$13.15</p>	<p>Turing's Cathedral: The Origins of the Digital Age George Dyson ★★★★☆ (52) Hardcover \$19.17</p>	<p>Alan M. Turing: Centenary Edition → Sara Turing ★★★★☆ (1) Hardcover \$23.48</p>	<p>The Man Who Knew Too Much: Alan Turing and the Secret of Bletchley Park → David Leavitt ★★★★☆ (21) Paperback \$15.21</p>

Question Answering – Watson!



Multiagent Systems and Learning

- How can an agent learn from experience in a world that contains other agents too ?
 - Other agents' learning makes the world nonstationary for the former agent
- Games
 - Learn to play Nash equilibrium
 - Learn to play optimally against static opponents

Games

- Multiple agents maybe competing or cooperating to achieve a task
- Capabilities for finding strategies, equilibrium between agents, auctioning, bargaining, negotiating.


- Business
- E-commerce
- Multi-robot systems
- Investment management
-



Kidney Exchange


- In US, $\geq 50,000$ /yr are diagnosed with lethal kidney disease
- Cure = transplantation, but cadaver kidney have a long waiting list (2-5 yrs)
- Potential donors may be incompatible with patient

Kidney Exchange



Optimizing kidney exchange

- Cycle cover = optimize matched vertices under cycles of length $\leq L$
- Problem is NP-hard



RoboCup

- Robot soccer competition and symposium
- Challenge: *By 2050, a team of fully autonomous humanoid robot soccer players shall win the soccer game, complying with the official rule of the FIFA, against the winner of the most recent World Cup*
- Pioneering multi-robot domain, since the early 90s, adversarial, with fully autonomous robots
- www.robocup.org

35

CoBot – Collaborative Robot

- Tasks
 - Go to a location
 - Deliver a message
 - Escort a visitor
 - Transport object between locations
 - Semi-autonomous telepresence
 - Visitor companion
 - Deliver mail
 - Tour guide



Autonomy but Limitations

- Perceptual
 - No universal object recognizer, object finder
 - No universal speech understanding
- Cognitive
 - No complete certainty of position
 - No certainty about all decisions, what to learn
 - No static optimal knowledge representation
- Actuation
 - No ability to manipulate all objects
 - No constant mobility in all surfaces

Symbiotic Autonomy

**Autonomous Robots Tasked
“To DO MORE” than
what they actually “CAN DO”:**

Intelligent Proactive “Ask for Help”

AI of the 21st century

- The integration of **MANY pieces**
- Use of **unlimited resources**
 - Own built-in and compiled knowledge
 - Search, infer, process information, remember, adapt
 - All external digital partners (sensors, robots, cell phones, toasters, vacuum cleaners, cars, roads)
 - Digital talk, digital translate
 - The Web
 - Mine human-understandable data
 - People
 - Ask, learn from, be taught, observe, copy

Main Technical Threads

- Problem / knowledge representation
- Search, optimization, constraint thinking
- Reasoning under uncertainty
- Modeling and planning
- Multiple agents, complex decision making
- Hidden information
- Learning

Part I – No Uncertainty

- Symbolic knowledge representation
- Pure search
- Deterministic planning

Part II – Uncertainty

- Probabilistic knowledge representation
- Planning under uncertainty
- Learning
- Other agents

Schedule

- (2) Jan 14 – Jan 16 – **Introduction and agents**
 - Knowledge representation – inference
- (8) Jan 21 – Feb 13 – **Search**
 - Uninformed – SAT, CSP
 - Informed – Heuristics, MIP
- (2) Feb 18 – Feb 20 – **Planning**
 - Classical planning and path planning
- (3) Feb 25 – Mar 4 – **Probabilistic reasoning**
 - Representation
 - MDPs/POMDPs
- (1) Wednesday, March 6 – Midterm Exam

Schedule

- (2) Mar 18 – Mar 20 – **Learning**
 - Supervised and RL
- (2) Mar 25 – Mar 27 – **Introduction to games**
 - Representation, solving, complexity
- (4) Apr 1 – Apr 10 – **Multi-robot systems**
 - Adversarial, playbook planning and learning
 - Learning from demonstration, reuse, correction
- (5) Apr 15 – Apr 29 – **Game solving**
 - Normal-form games, sequential complete- and incomplete-information games
- (1) May 1 – **Discussion, open problems**