

# CS 4100/5100: Foundations of AI

## Introduction <sup>1</sup>

Instructor: Rob Platt  
r.platt@neu.edu

College of Computer and information Science  
Northeastern University

September 5, 2013

---


<sup>1</sup>In general, my slides sometimes draw images and/or text from Gillian Smith's slides (<http://sokath.com/main/>).

# About me


I work in robotics. In particular: robot manipulation, planning, perception, control under uncertainty.




Porterbot, MIT



Atlas, Boston Dynamics



Robonaut 2, NASA



UMass Torso, UMass  
Amherst

# Topics to be covered in this course

- ▶ Search
  - ▶ uninformed search
  - ▶ heuristic search
  - ▶ constraint satisfaction
  - ▶ adversarial search
- ▶ Propositional logic
- ▶ Reasoning under uncertainty
  - ▶ Bayes networks
  - ▶ inference, approximate inference
- ▶ Decision theory
- ▶ Machine learning


# Course evaluation

- ▶ (10%) Quizzes held at the beginning of each class
  - ▶ Cover basic information from the reading that was due that day.
  - ▶ Approx 10 quizzes total. I will drop the lowest two quiz grades.
- ▶ (50%) Homework assignments
  - ▶ Approx 5 homework assignments.
  - ▶ I will drop the lowest homework assignment.
- ▶ (20%) Midterm exam
- ▶ (20%) Final exam

# Contact Info

- ▶ Professor: Rob Platt, [r.platt@neu.edu](mailto:r.platt@neu.edu)
- ▶ Office hrs: TBD or by appointment
  
- ▶ TA: Maryam Aziz
- ▶ Office hrs: Monday 8am - 10am, [azizm@ccs.neu.edu](mailto:azizm@ccs.neu.edu), 472 West Village H
  
- ▶ TA: Yupeng Gu
- ▶ Office hrs: TBD, [gyp7364@gmail.com](mailto:gyp7364@gmail.com), 208 West Village H

# Textbook



## Artificial Intelligence: A Modern Approach Third Edition

Stuart Russell, Peter Norvig


# The difference between talking and cheating

You **ARE** allowed to: talk to other students about concepts in the reading.

You **ARE NOT** allowed to: discuss or share the solutions to homework assignments with anyone besides me or the TA.

If you violate these rules, I will probably find out about it (using Moss). If I do, then you will be penalized and I may refer you to the college for disciplinary action.

# Questions?






# What is AI: the Turing Test

- ▶ The question, "can machines think?" is ill defined
- ▶ Define intelligence in terms of how the machine *\*acts\**.
- ▶ But, how do we measure how nearly human a machine acts?
  - ▶ One answer: The Turing Test (*i.e.* the "imitation game").

# The Turing Test



- ▶ Player C is human.
- ▶ Either player A is human and B is a machine or *vice versa*.
- ▶ Player C cannot see players A or B and must communicate only by means of written communication.
- ▶ The machine is programmed to act like a human.
- ▶ Can player C determine which of A or B is the machine?

# What is AI: other visions

- ▶ Cognitive modelling: create models of how we think the human mind functions.
  - ▶ for example: we might create a cognitive model of how we \*think\* human motor dexterity works (based on biological evidence). We might try to create robotic dexterity this way.
- ▶ Logic: create the "logical machine" that can make deductions from observed evidence.
  - ▶ 1929: Presburger showed that the theory of natural numbers was decidable
  - ▶ 1931: Godel showed that any sufficiently strong axiomatic system is not complete
- ▶ Rational agent: create an agent that acts so as to achieve some goal or optimize some objective function.
  - ▶ for example: we might create a robotic mouse. The objective of the mouse is to locate and consume cheese. An intelligent mouse system would calculate how best to achieve this goal.

# A short (and scattered) history of AI

Herbert Simon (1957):

"It is not my aim to surprise or shock you – but the simplest way I can summarize is to say that there are now in the world machines that think, that learn, and that create. Moreover, their ability to do things is going to increase rapidly until – in a visible future – the range of problems they can handle will be coextensive with the range to which the human mind has been applied."


## A short (and scattered) history of AI

- ▶ 1952, Arthur Samuel: developed a computer program that could play checkers (using an early version of alpha-beta pruning).
- ▶ 1960, Donald Michie: MENACE (Matchbox Educable Noughts And Crosses Engine) – a program that learned to play an optimal game of tic-tac-toe (an early version of RL).




# A short (and scattered) history of AI

- ▶ 1962, Frank Rosenblatt: Perceptrons. Proof that simple neural network learning algorithms can learn any pattern that can be represented by the network.
  - ▶ But, it turns out that many interesting functions cannot be represented this way: xor, for example.
- ▶ late 1960s, Nilsson *et. al.*: Shakey project at SRI (occasioned the development of STRIPS, an early planner).



## A short (and scattered) history of AI



- ▶ 1972, Terry Winograd: SHRDLU, natural language understanding in the context of a "blocks world".

# Problems scaling up

Although AI had many early successes, the problem of scaling up to "real world" problems was more difficult than many predicted.

- ▶ A key problem is that the theory of computational complexity was not well understood.



# Expert systems

Rule-based systems. Systems that reason logically from a large "expert" rule set.

- ▶ 1971, Feigenbaum, DENDRAL: a program that could infer molecular structure from information provided by a mass spectrometer.
- ▶ Shortliffe, MYCIN: a program that could diagnose blood infections.
  - ▶ After asking a long series of questions, MYCIN would produce a diagnosis and a corresponding confidence number.
  - ▶ MYCIN outperformed clinicians in Stanford medical school!

# AI today

Today, AI research has many applications including:

- ▶ Robotics
- ▶ Games
- ▶ Computer Vision
- ▶ Information retrieval
- ▶ machine translation
- ▶ natural language understanding

# Stanley




- ▶ won the 2005 DARPA desert grand challenge
- ▶ drove 150 miles on roads in the Mojave desert autonomously

# Watson



- ▶ 2011: beat the human world champions of the day

# What is an *agent*?



- ▶ The idea that an independent entity interacts w/ the external world.
  - ▶ Sensors: input
  - ▶ Actuators: output
  - ▶ Objective: achieve some goal or optimize some performance measure.

# Types of environments

- ▶ Fully observable vs partially observable (Backgammon vs Poker)
  - ▶ Do the agent's sensors tell it everything it needs to know (fully observable) or does the agent need to remember things in order to optimize its performance (partially observable)?
- ▶ Single agent vs multi agent
- ▶ Deterministic vs stochastic (Chess vs Backgammon)
  - ▶ Do the agent's actions always have the same effects (deterministic) or is the outcome stochastic?
- ▶ Episodic vs sequential
- ▶ Static vs dynamic
- ▶ Discrete vs continuous (Chess vs Robot)
  - ▶ Is the environment most easily described by continuous or discrete state variables?
- ▶ Known vs unknown