

15-780 – Introduction and History of AI

J. Zico Kolter

January 13, 2014

What is “AI”?



Some classic definitions

Buildings computers that ...

Think like humans	Think rationally
Act like humans	Act rationally

Some classic definitions

Buildings computers that ...

Think like humans	Think rationally
Act like humans	Act rationally

Requires (presumably) that we have some knowledge of how humans think; currently more the purview of cognitive science

Some classic definitions

Buildings computers that ...

Think like humans	Think rationally
Act like humans	Act rationally

Requires (presumably) that we have some knowledge of how humans think; currently more the purview of cognitive science

“May not machines carry out something which ought to be described as thinking but which is very different from what a man does?” – Alan Turing

Some classic definitions

Buildings computers that ...

Think like humans	Think rationally
Act like humans	Act rationally

The fields of logic and automated reasoning, a common theme in much original AI research

Some classic definitions

Buildings computers that ...

Think like humans	Think rationally
Act like humans	Act rationally

The fields of logic and automated reasoning, a common theme in much original AI research

Searching through logical deductions proved a very computationally intensive task, and it is unclear whether logic is a good basis for the “more vague” notions of knowledge

Some classic definitions

Buildings computers that ...

Think like humans	Think rationally
Act like humans	Act rationally

The *Turing Test*: talk to a computer or human for an hour, can you tell which one?

Some classic definitions

Buildings computers that ...

Think like humans	Think rationally
Act like humans	Act rationally

The *Turing Test*: talk to a computer or human for an hour, can you tell which one?

Interviewer: "What is $32,839 \times 128,394$?"

Computer: "I don't know!"

Some classic definitions

Buildings computers that ...

Think like humans	Think rationally
Act like humans	Act rationally

The basis for the “intelligent agent” framework in
Russell and Norvig

Some classic definitions

Buildings computers that ...

Think like humans	Think rationally
Act like humans	Act rationally


The basis for the “intelligent agent” framework in
Russell and Norvig

Unclear if this accurately captures the scope of current
AI research, or if AI researchers see themselves as
working toward this goal

The pragmatist's view

AI is that which appears in academic conferences on AI

The pragmatist's view



1980s

The pragmatist's view



1990s

The pragmatist's view




2010s

A broader definition

We won't worry too much about definitions of AI, but the following will suffice:

Artificial Intelligence is the development and study of computer algorithms that address a problem typically associated with some form of intelligence

Some notable milestones in AI



1950



1950 paper, “Computing Machinery and Intelligence,” introduces Turing test

Interviewer talks to a person and machine via a terminal interface, must determine which is which

1952-1959



Arthur Samuel develops checkers program while at IBM

Program includes a version of alpha-beta pruning in game search, position scoring, and learning from self-play

1955-1956



Alan Newell, Herbert Simon, J.C. Shaw develop Logic Theorist, which sought to prove mathematical theorems from *Principia Mathematica*

A pioneering effort in intelligence as search, found a shorter proof of one well-known theorem

1956



Dartmouth Conference on Artificial Intelligence, organized by John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon

1958



Frank Rosenblatt's Perceptron algorithm learns to recognize letters of the alphabet from images

1966



Joseph Weizenbaum develops ELIZA, a
“virtual therapist” program

ELIZA transcript from Communications of the ACM

Men are all alike.

IN WHAT WAY

They're always bugging us about something or other.

CAN YOU THINK OF A SPECIFIC EXAMPLE

Well, my boyfriend made me come here.

YOU BOYFRIEND MADE YOU COME HERE

He says I'm depressed much of the time.

I AM SORRY TO HEAR YOU ARE DEPRESSED

Its true. I am unhappy.

DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE
UNHAPPY

I need some help, that much seems certain.

WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP

Perhaps I could learn to get along with my mother.

TELL ME MORE ABOUT YOUR FAMILY

My conversation with ELIZA

Hello.

HOW DO YOU DO. PLEASE STATE YOUR PROBLEM.

I am worried that I am just speaking to a computer.

DO COMPUTERS WORRY YOU

No, but I'm trying to discern if you are a real person or a computer.

ARE YOU SAYING NO JUST TO BE NEGATIVE


You did not address my question.

YOUR QUESTION

Are you a computer?

DO COMPUTERS WORRY YOU


1966-1972



Shakey the robot developed at Stanford Research Institute


Goal is to develop a general purpose robot capable of reasoning and interacting with the world

1980-1988




Marc Raibert of CMU and MIT develops hopping robots

1982



Backpropagation for training multi-layer neural networks popularized by David Rumelhart, John Hopfield (amongst many others)

1988



Judea Pearl publishes *Probabilistic Reasoning in Intelligent Systems*, bring probability and Bayesian networks to forefront of AI

1997



Deep blue defeats Gary Kasparov in a six-game chess match (two wins, one loss, three draws)

2005-2007




Stanford and CMU respectively win 2005 and 2007 DARPA Challenges, requiring a car to drive autonomously through desert and simulated urban environments

2011



IBM's Watson defeats human opponents
on Jeopardy


AI is all around us



Face detection



Personal assistants



Software verification



Logistics planning

This course

- 15-780 is intended to be a broad but intensive introduction to a wide variety of topics in AI
- We will cover topics include: uninformed and informed search, numerical optimization, constrain satisfaction, integer programming, machine learning, classical planning, probabilistic planning and reinforcement learning, probabilistic reasoning, computer vision, robot motion planning, scheduling, natural language processing, multiagent systems... (whew)

<http://www.cs.cmu.edu/~zkolter/course/15-780-s14>


Instructors



Instructor
Zico Kolter



Instructor
Zack Rubinstein




TA
Vittorio Perera

Recommended background

- No formal pre-requisites
- But, substantial programming background is required (assignments will be in Python)
- Additional background in data structures and algorithms, linear algebra, probability will all be helpful, but not required

Textbook



- Textbook for the course is *Artificial Intelligence, A Modern Approach, 3rd Edition* (2nd edition is also ok, though you may need to map question numbers)
- Also additional material not covered in the textbook, this will be covered in slides and possibly supplemental notes

Course grading

Problem sets	45%
Mid-term	20%
Course project	25%
Class participation	10%

Course grading

Problem sets	45%
Mid-term	20%
Course project	25%
Class participation	10%

- Four problem sets, each with written and 1-2 programming assignments
- Written portion to be turned in at the *beginning* of class, programming portion turned in using Autolab
- You may discuss problems in groups, but must be written up and programmed *independently*

Course grading

Problem sets	45%
Mid-term	20%
Course project	25%
Class participation	10%

- Class participation measured through in-class polls and Piazza discussion forums
- 50% of the credit for just answering in-class polls, remaining 50% for answering them correctly and for participating in Piazza forums

Piazza

<http://www.piazza.com>

- As mentioned above, we will use Piazza for discussion and in-class polls
- All students *must* register for the course page on Piazza, and bring either a laptop or smartphone with the Piazza app to class
- Any enrolled student who has *not* signed up for Piazza by next week will be dropped from the course to make room for those on the waitlist

Late days

- You will have 5 late days to use throughout the semester, no more than 3 late days can be used on a single assignment
- Late days extend the deadline for homework by 24 hours (remember, homework due at beginning of class)
- Turn in late homework to Vittorio Perera's office, GHC 7002, sliding under the door if no one there and *write the time of submission*

Cheating

- Don't do it (please)
- CMU's academic integrity policy is here:
<http://www.cmu.edu/academic-integrity/>
- Remember, written problem sets and programming must be written up *independently*
- If you consult any outside material, you need to cite all of your references

Some parting thoughts

“Computers in the future may have only 1,000 vacuum tubes and weigh only 1.5 tons.” – Popular Mechanics, 1949

Some parting thoughts

“Computers in the future may have only 1,000 vacuum tubes and weigh only 1.5 tons.” – Popular Mechanics, 1949

“Machines will be capable, within twenty years, of doing any work a man can do.” – Herbert Simon, 1965