CS 486/686
Introduction to Artificial Intelligence

Alice Gao

Lecture 1

Based on work by K. Leyton-Brown, K. Larson, and P. van Beek

Outline

Learning goals - CS 486/686 Lecture 1

By the end of the lecture, you should be able to

- Get to know a bit about Alice and one or more classmates.
- Name an application of AI. Name a topic in this course.
- Describe tips for succeeding in this course.
- Describe the four definitions of AI. Explain why we will pursue one over the other three.

Outline

Who am I?

My name is Alice Gao. I grew up in Beijing, China, and have lived in Vancouver, Toronto, Boston, Cambridge (UK), New York City, and Kitchener.

Research: artificial intelligence, game theory, education, and peer grading.

My work/education history:

- Lecturer, Computer Science, University of Waterloo.
- Postdoc, Computer Science, UBC.
- Ph.D., Computer Science, Harvard University.
- Undergraduate, Computer Science and Mathematics, UBC.

Hobbies: board games, escape room games, hiking, swimming, and traveling.

Meet your peers

- In the next 2 minutes, introduce yourself to someone you don't know.
- Talk about courses, co-op, summer activities, dorms, extracurricular activities, graduation, jobs, etc.
- I encourage you to sit in a different section of the classroom every lecture and get to know the people around you.

Outline

The State of Art of AI

What can AI do today?

- Little success on the grand goal (building a general intelligence agent)
- Lots of success in restricted domains

Checkers

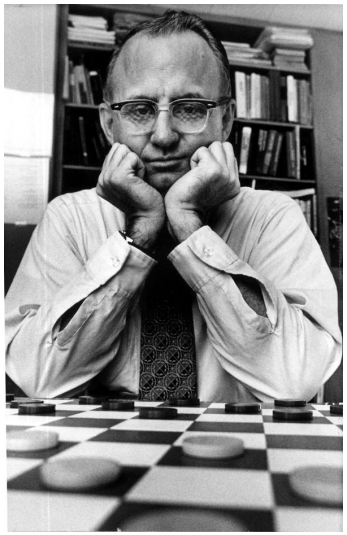

Checkers

- 500 billion billion possible positions (5×10^{20})
- Marion Tinsley, the world champion of checkers.
- Chinook, Jonathan Schaeffer, University of Alberta.
- Tinsley vs Chinook in 1992 and 1994.
- Schaeffer, Jonathan, et al. "Checkers is solved." science 317.5844 (2007): 1518-1522.

CQ: Checkers

CQ: Assuming that both players play checkers perfectly, the player, who goes first,

(A) has a strategy to guarantee a win.

(B) has a strategy to guarantee a draw.

Chess

Chess

- More than 10^{100} positions
- Deep Blue, IBM
- Beat world champion in 1997
- Strongest chess engines: Stockfish, Houdini, Komodo, ...
- Program search depth: 20; Human search depth 3-4

CQ: Chess

CQ: Deep Blue was the first computer to beat a reigning world chess champion. Which Russian did Deep Blue beat in May 1997?

(A) Vesselin Topalov
(B) Bobby Fischer
(C) Garry Kasparov
(D) Boris Spassky

Go

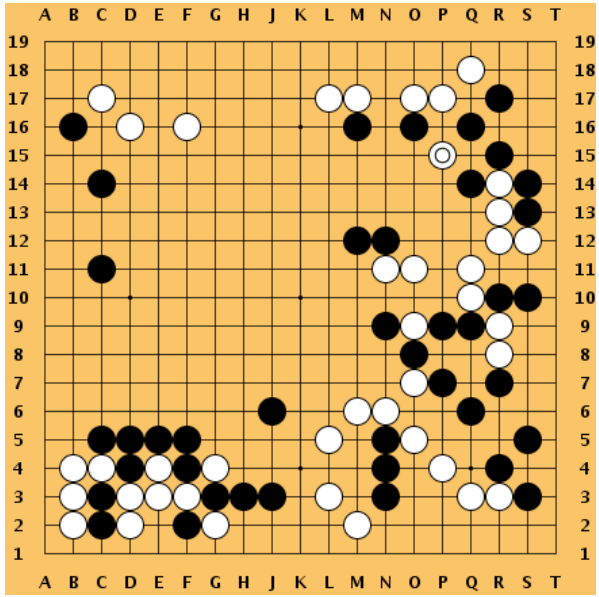

AlphaGo v.s.

Go

- More than 10^{360} positions
- AlphaGo, Google DeepMind
- AlphaGo v.s. Lee Sedol (9-dan rank) in March 2016.
- Silver, David, et al. "Mastering the game of Go with deep neural networks and tree search." nature 529.7587 (2016): 484.

CQ: Go

CQ: What was the outcome of the 5-game match between AlphaGo and Lee Sedol in March 2016?

(A) 5-0

(B) 4-1

(C) 3-2

Poker

(a) Michael Bowling, UofA

(b) Tuomas Sandholm, CMU

Poker

- Play with uncertainty. Must model opponent(s). Care about long-term payoff.
- Latest news from U of A:
  Bowling, Michael, et al. "Heads-up limit hold'em poker is solved." Science 347.6218 (2015): 145-149.
  DeepStack defeated professional poker players at heads-up no-limit Texas hold'em.
- Latest news from CMU:
  Brown, Noam, and Tuomas Sandholm. "Superhuman AI for heads-up no-limit poker: Libratus beats top professionals." Science (2017): eaao1733.

Jeopardy!

"AI for $100, Alex."
"This popular TV quiz show is the latest challenge for IBM."
"What is Jeopardy?"

Jeopardy

- Watson, IBM
- Beat Brad Rutter and Ken Jennings in 2011.
- Question delivered in text, had to generated answer in a few seconds. Stored 200 million pages locally (No internet allowed).
- Now used for healthcare.
- Full story https://tek.io/2lKMQIe

Autonomous Cars

2005 DARPA Grand Challenge

(a) Stanley

(b) Kat-5

(a) TerraMax

(b) H1ghlander

(c) Sandstorm

2005 DARPA Grand Challenge

- 212km course near California/Nevada state line.
- 5 out of 23 vehicles successfully completed the course.
- Narrow tunnels, sharp turns, and a winding mountain pass

CQ: 2005 DARPA Grand Challenge

CQ: In the 2005 DARPA Grand Challenge, out of the five vehicles that completed the 212km course, which vehicle won the challenge by taking the least amount of time?

(A) Stanley by Stanford University

(B) Kat-5 by The Grey Insurance Company

(C) TerraMax by Oshkosh Truck Corporation

(D) H1ghlander by Carnegie Mellon University

(E) Sandstorm by Carnegie Mellon University

Many other applications of AI

- FCC Spectrum Auction https://bit.ly/2oQC6dg
- Vacuum robots https://bit.ly/2wWAC5q
- Spam filtering https://bit.ly/2rNLXDW
- Automated planning and scheduling for transportation during Persian Golf Crisis in 1991 https://bit.ly/1LSEetu
- Automated phone systems https://ibm.co/2id0Wkp

Topics in CS 486/686

- Search
- Constraint Satisfaction Problems
- Planning Under Certainty
- Reasoning Under Uncertainty
- Decision Making Under Uncertainty
- Learning

Outline

Course Administration

CS 486/686 Introduction to Artificial Intelligence

2 sections:

- 1:00-2:20 Mon/Wed MC 1056
- 2:30-3:50 Mon/Wed MC 1056

Instructor:

- Alice Gao (a23gao@uwaterloo.ca, DC 3117)

TAs:

- Benjamin Armstrong, Aravind Balakrishnan, Camilo Andres Munoz Bravo, Atrisha Sarkar, Colin Michiel Vandenhof

Course Resources

Website: https://www.student.cs.uwaterloo.ca/~cs486

Sign up for Piazza here

Learn: https://learn.uwaterloo.ca/d2l/home/403295

- Clicker registration, assignment submissions, and grades

Textbooks:

- No required textbook. Lectures follow the Russell and Norvig book closely.
- Artificial Intelligence: A Modern Approach by S. Russell and P. Norvig (3rd Edition)
- Artificial Intelligence: Foundations of Computational Agents, D. Poole and A. Mackworth (available online)

Grading Scheme

CS 486

- Clickers: 5%
- Quizzes: 20%
- Assignments: 30%
- Final: 45%

CS 686

- Quizzes: 15%
- Assignments: 25%
- Final: 40%
- Project: 20%

CQ: What do you think of clicker questions?

CQ: What do you think of clicker questions?

(A) I like them and I think they are useful.

(B) I don't like them, but I think they are useful.

(C) I don't like them and I think they are useless.

(D) I am indifferent between using and not using them.

(E) None of the above.

CQ: Why does Alice want to use clickers?

CQ: Why does Alice want to use in-class clicker questions and make them count for 5% of the final grade?

(A) To see if students are awake.

(B) To develop good exam questions.

(C) To test the wisdom of the crowds effect.

(D) To force students to attend lectures.

(E) To encourage active learning in class.

Dealing with Clicker Questions 5%

Policy for clicker marks

- For each question, 2 points for responding and 1 point for choosing the correct answer.
- Only retain best 75% of the clicker marks.

Tips for dealing with clicker questions

- Don't stress about them. They are meant to be low-stake.
- Want you to think and work through problems.
- Good questions may appear on exams.
- Wisdom of crowds effect: Feel free to discuss with your neighbours.

Dealing with Quizzes 20% or 15%

- Spread out the midterm over the term.
- Approximately 1 quiz per week. (1.5% to 2% per quiz)
- 10 minutes in the middle of class.
- Predominantly multiple-choice questions. Occasionally contain short answer questions.
- More conceptual than application questions

Dealing with Assignments 30% or 25%

- 4-5 assignments. 1 assignment every 2.5-3 weeks.
- 1-2 questions per assignment
- At least one question per assignment involves programming.

Dealing with the Project 20%

Key dates

- Your project proposal is due on Monday October 15th, 2018.
- Your project report is due on Monday December 3rd, 2018.

See the project page on the website for more information.

If you are stuck or confused, come talk with me or with the TAs.

Study tips

- If you come to lectures, make a decision to focus, to engage, and to make the most of your time here.
- Question everything I say. Ask yourself why. Write down things that I don't write down.
- You learn the most from doing the assignments.
- Study regularly based on the learning goals. Don't cram.
- Struggling is necessary for learning.

Outline

What is Artificial Intelligence?

Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally

Humans v.s. Rationality

Compare to human performance	Compare to an ideal concept of intelligence
Systems that act **like humans**	Systems that think **rationally**
Systems that act **like humans**	Systems that act **rationally**

Thinking v.s. Acting

Thought processes and reasoning	Systems that **think** like humans	Systems that **think** rationally
Behaviour	Systems that **act** like humans	Systems that **act** rationally

Thinking Humanly

The Cognitive Modeling Approach

- Few examples of intelligence
- How do humans think?
  - Introspection
  - Brain imaging
- Cognitive science

Acting Humanly

The Turing Test Approach

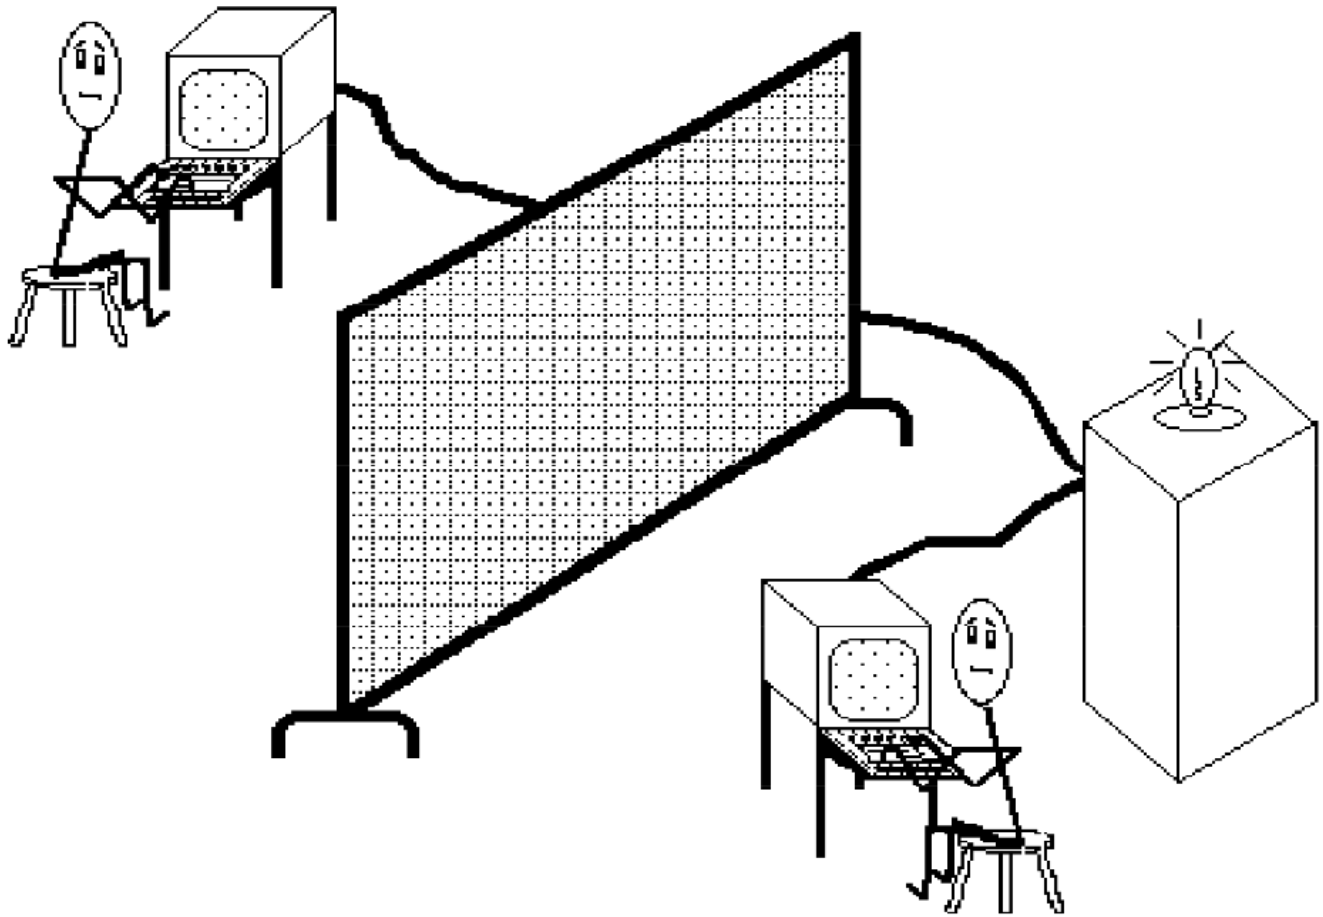

The Turing Test

- An operational definition
- The Turing Test and the Total Turing Test
- Gave rise to six core areas of AI

Rationality

- Rationality: an abstract "ideal" of intelligence, rather than "whatever humans do"
- A system is rational if it does the "right thing," given what it knows.

Thinking Rationally

The Laws of Thought Approach

- Greek philosophers invented logic.
- The logicist tradition
- Two obstacles for using this approach in practice

Acting Rationally

The Rational Agent Approach:

- Agent means todo.
- The goal of a rational agent
- What behaviour is rational?

CQ: Which definition of intelligence would you adopt?

CQ: If you were an Artificial Intelligence researcher,
which of the following definitions of intelligence would you adopt?

(A) Systems that think like humans

(B) Systems that act like humans

(C) Systems that think rationally

(D) Systems that act rationally

Which definition of intelligence did we adopt?

A system is intelligent iff it acts rationally.

Why do we care about behaviour instead of thought processes and reasoning?

Which definition of intelligence will we adopt?

A system is intelligent iff it acts rationally.

Why do we measure success against rationality instead of against humans?

Revisiting the learning goals

By the end of the lecture, you should be able to

- Get to know a bit about Alice and one or more classmates.
- Name an application of AI. Name a topic in this course.
- Describe tips for succeeding in this course.
- Describe the four definitions of AI. Explain why we will pursue one over the other three.