

CompSci 370

Artificial Intelligence Introduction

Ron Parr

Duke University

Department of Computer Science

Course Staff (See Class Web Page)

- Ron Parr – instructor
- Graduate TAs
 - Bowen He
 - Casey Hanks
- Undergraduate TAs
 - TBD

About me

- Learned to program on my 8-bit Atari computer
- Sent a print out(!) of my Othello (reversi) playing code (in BASIC!) w/my college application
- Majored in Philosophy
- Switched to CS for graduate school
- Started at Duke in 2000

Am I Prepared?

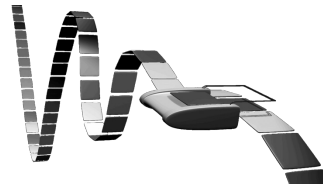
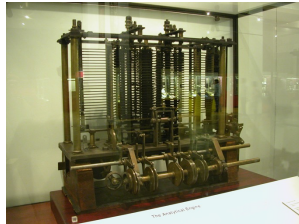
- Good programming skills:
 - We assume that you can write, debug your own programs
(If you need help programming, this class is **too hard** for you!)
 - We will use python for programming assignments
 - **We expect you to figure out** how to use Python and the command line
 - Note: Co-Lab has a class in linux and the bash shell on 1/29 that you might find helpful
- Other expectations
 - Ability to do short proofs
 - Basic probability concepts (though we will review this)
 - Basic algorithmic concepts
 - Complexity - $O()$
 - Analysis of algorithms
 - Math: Basic calculus, basic linear algebra
 - CompSci 230 is no technically required, but helpful
 - CompSci 330 also helpful

What is AI?

- For centuries, perhaps longer, people have wondered how to reproduce the smarts that people have...
- Even though we really have *no idea* how to define such things
- Defining intelligence has, itself, been a **career-long endeavor** for many scholars

Machine Intelligence Over the Centuries

- As long as people have had machines, they've wondered if they could exhibit human-like intelligence
- von Kempelen's (fraudulent) Turk (1700s), Babbage's analytical and difference engines (1800s), Turing's Turing machine (1900s)



Images from Wikipedia

Turing Test

- Computer must be indistinguishable from a human based upon written exchanges (presaged ChatGPT)
 - Does this imply intelligence?
 - How could the computer cheat?
 - Does intelligence imply a certain type of computation?
 - Could an intelligent machine still fail the test?
- Does our notion of intelligence transcend our concept of humanity?

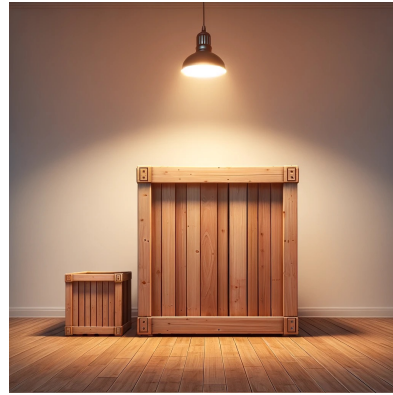
ChatGPT hasn't "solved" AI

My house has a high ceiling. I'd like to change a lightbulb on the ceiling, but can't reach it and I don't have a ladder. I have a sturdy crate that's two feet tall, but I still can't reach the bulb when standing on the crate. I have another sturdy crate that's four feet tall. I could reach the bulb from that crate, but I can't climb onto it. What should I do?



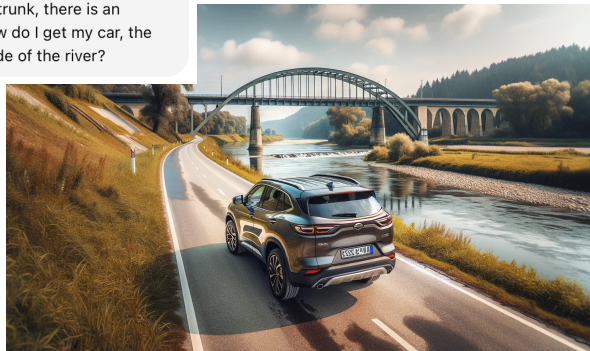
You could stack the two crates to create a more manageable height. Place the four-foot crate on the floor and then stack the two-foot crate on top of it. This way, you can climb onto the two-foot crate first and then step onto the four-foot crate. Ensure that both crates are stable and secure before climbing to avoid any risk of falling. Alternatively, if you have any other sturdy objects or pieces of furniture that can help you climb up safely, you can use those as well. If you're still uncertain or if the setup feels unsafe, consider borrowing or buying a ladder.

ChatGPT versus the obvious solution



Foundation models regurgitate (modified) plans they've seen before

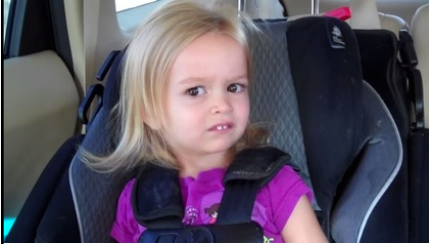
I'm in my car on the left side of a river that is spanned by a road bridge. My car has three passengers. In my trunk, there is an inflatable raft that can carry two people. How do I get my car, the raft, and all of my passengers to the right side of the river?



To solve this problem, we need to transport the car, the passengers, and the raft to the other side of the river. Here's a step-by-step solution:

1. **Cross the bridge with your car:** Drive your car across the bridge with all the passengers inside. Now, your car and the passengers are on the right side of the river.
2. **Inflate the raft:** Take the inflatable raft out of the trunk and inflate it on the right side of the river.
3. **Transport the raft to the left side:**
 - Since the raft can only carry two people, you will need to make multiple trips.
 - You can row the raft back to the left side yourself, or send two passengers with the raft to the left side.
4. **Move passengers using the raft:**
 - If you are alone on the left side, take the first passenger across to the right side using the raft. Leave them on the right side and return with the raft to the left side. Repeat this process until all passengers are on the right side.
 - If two passengers went to the left side initially, they can both row back together. Then, they can take turns rowing other passengers across, ensuring at least one person is always on each side to facilitate the crossing process.
5. **Final steps:**
 - After all passengers and the raft are on the right side, deflate the raft and pack it back into the trunk of the car.

Now, your car, the raft, and all of your passengers are on the right side of the river.



What Intelligence Isn't

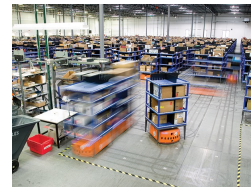
- It's not about fooling people
- For many years, many have made ongoing efforts to fool people: <http://chatbots.org/>
- ChatGPT has shown us that convincing text without genuine understanding is entirely possible

AI after Turing

- Modern AI is ~60 years old
- “AI” term proposed at 1957, Dartmouth Conference
- Has been a subject of intense study since then
 - 1960’s: Logic, search, theorem proving, perceptron
 - 70’s: Robotic & perception
 - 80’s: Expert systems, 1st industrial interest, neural nets
 - 90’s: agents, uncertainty, “AI Winter”
 - 00’s: growth of ML, NLP, usable AI systems
 - 10’s: Deep learning, industrial/commodity AI, robotics
 - 20’s: “Foudation Models” (chatGTP, Stable Diffusion, etc.)
(the long path from Turing to Turing)

Less Glamorous AI at work: AI at Amazon Warehouses

- Amazon uses **robots to move products** within its warehouses (deploys 200,000 robots)
- Amazon uses AI to **predict demand**
- May use AI to deliver products
- Consequences:
 - Pay fewer workers
 - Warehouses are packed more densely
 - Less space wasted on unpopular products
 - Combine to **increase value per sq. unit** of space



Kiva systems photo
From IEEE Spectrum 7/2/08

What about General Intelligence?

- AI didn't get traction until it focused on more specific problems
- Hard to provide "general intelligence" if you don't know what it is

- Generative AI creates an illusion of general intelligence without directly trying to solve the general intelligence problem
- We aren't teaching "general intelligence"
- We are teaching the ingredients of systems such as Amazon's warehouses and (to some extent) ChatGPT

What is covered this semester?

- See syllabus on class web page

- Note changes for this year:
 - More popular topics emphasized more
 - Less popular topics pushed to end of the semester

Major Topics *Not* (directly)Covered

- Natural Language
- Vision
- Will come up as examples of broader concepts

Class Mechanics

- Textbook: ***Artificial Intelligence, A Modern Approach***, Russell & Norvig (*fourth* edition – third is probably OK)
 - Semi-required
 - Electronic versions available
 - Please don't steal my advisor's textbook!
- Homework: 30%
 - mix of short proofs, algorithm design/analysis (15%)
 - Programming projects in teams of up to 2 (15%)
 - High level discussion OK, write-up, coding must be your team's own
- Midterm: 30%
 - Conceptual questions
 - Hopefully in person, no collaboration
 - Scheduled during class time
- Final: 40%
 - Conceptual questions
 - Hopefully in person, no collaboration
 - Scheduled according to registrar's final exam schedule

Lateness Policy

- 10% per day for w/o extensions
- Valid reasons for extensions:
 - Excused travel for official sporting events
 - Illness
 - Job interviews
 - Travel to present academic research
- Not excused:
 - Having a lot of work in other classes
 - Just “feeling stressed”
 - Extracurricular activities (clubs, startups, etc.)

Grading

- I tend to give challenging conceptual questions, and not everybody will get them
- More important for you to be **challenged** than to have a score that you can put on your refrigerator
- Don't obsess over raw scores
- At end of semester, I will decide how many points correspond to 1/3 of a letter grade
 - Will always be ≥ 3.33
 - Chosen to ensure median grade of high B+ or low A-

Attendance

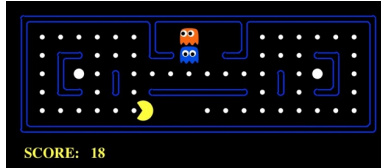
- You are not required to attend class – though it's a *good idea*
- Everything in class is covered in readings and slides
- Classes will **NOT** be recorded
- Why? Ask me offline.

Discussion Sections

- Staffed by (U)Tas
- Attendance optional
- No new material covered
- Goals:
 - Work through common issues, e.g., “Help me fix my python installation!”
 - Work through problems/examples that wouldn't fit in lecture
- We will usually post questions for discussions before discussion, and post solutions after discussion

Programming Assignments

- Based on the Berkeley Pacman framework



- Why?
 - It's *really* well-done
 - Seeing ***your own code*** run AI algorithms is **fun, motivating,** and develops your **intuitions**
 - Even debugging is instructive

Pac-Man is a registered trademark of Namco-Bandai Games, used here for educational purposes

Pacman Limitations

- Some python version restrictions/complications
- Not all algorithms make sense in this framework (life isn't a Pacman game)
- Has been around for a while
 - Pacman was new when RP was a kid
 - Temptation to cheat

Academic Honesty

- Brainstorming with friends is encouraged, but answer write up and coding must be your own work
- Don't confuse brainstorming with letting your smart friends tell you the answers
- Don't Google/chatGPT for answers!!!
- Don't troll for answers from previous semesters
- You may search for definitions

- What you turn in must be your own work!!!
(in case of programming with a partner, it must be the work of just you and your partner)

Examples of Cheating

- Simply reading solutions to similar problems found by searching
- Submitting code written by others
- Refactoring or cosmetically modifying code written by others (this is much easier to catch than you think!)
- "Borrowing" a friend's laptop and finding answers

- Note: **Uploading** to a code sharing site **is also cheating**

Consequences of Cheating

- One year, 8 people were caught submitting code from the internet as their own code
 - All cases were reported
 - Consequences included:
 - Zeroes on assignments
 - Suspension
 - Failure to graduate
 - Retraction of job offers
- Students have been caught and reported recently!

Consequences of Cheating This Year

- All cases will be reported
- A grade of zero will be given for any assignment on which cheating is detected
- At least 1/3 letter grade will be deducted from the final grade for each instance of cheating in addition to any other penalties
- Other penalties may apply, at the discretion of the instructor and/or dean

Just don't do it!

Should I worry about getting falsely accused?

- No!
- I have never had a false positive accusation
- How do I know this?
 - I don't make frivolous accusations
 - Cheating is surprisingly obvious when it happens
 - When presented with evidence, **students have always owned up**

Changes from Previous Years

- More emphasis on hot topics
- HW counts less

Not Changing

(despite a few complaints)

- Still a class about **algorithms for AI**
(not a software development class)
- Still **not a theory class**
(though some students incorrectly think it is because there's math)
- Not a machine learning class
(though more ML coverage)

How To Approach This Class

- This class can be **hard and a lot of work**, but I have taught versions of it for many years and most who are *prepared and stick with it*:
 - Earn a **reasonable grade** in the end
 - Have **fun** with projects
 - **Learn** a lot
- We are here to help!
- We do not want to be your adversaries in this process
- **Let us be your partners in learning** by allowing us to help you