CompSci 370

Artificial Intelligence Introduction

Ron Parr

Duke University

Department of Computer Science

Course Staff (Office Hours - See Class Web Page)

- Ron Parr – instructor
- Graduate TAs
  - Jack Goffinet
  - Dillon Sandhu
- Undergraduate TAs
  - Qiaoyi Fang
  - Ankit Jajoo
  - Ryan Krakower
  - Dev Seth

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
- Once taught this class to just 6 students

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 figure out how to use Python and the command line

- Other expectations
  - Ability to do short proofs
  - Basic probability concepts (though we will review all of this)
  - Basic algorithmic concepts
    - Complexity - O()
    - Analysis of algorithms
  - Math: Basic calculus, basic linear algebra
  - CompSci 230 is essential
  - CompSci 330 also helpful

How is 270 different from 370?

- It’s the same class
- Course numbering change to reflect difficulty of class, amount of CS experience expected

- Hope to avoid high attrition w/clear expectations

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)

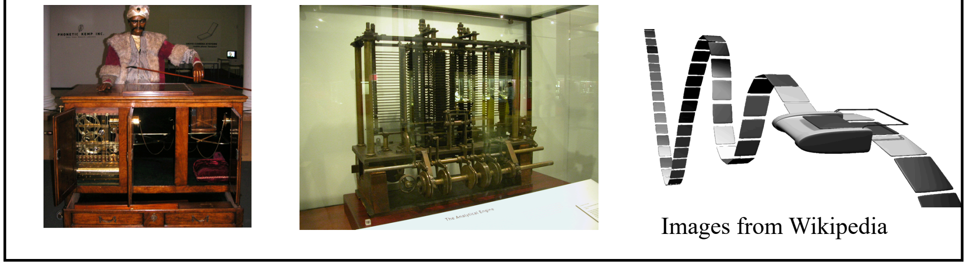

Images from Wikipedia

Turing Test

- Computer must be indistinguishable from a human based upon written exchanges
  - 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?

What Intelligence Isn’t

- It’s not about fooling people
- Fooling people is (in some cases) easy, e.g., eliza: http://www.manifestation.com/neurotoys/eliza.php3
- (built in to emacs meta-x doctor)
- More recent efforts: http://chatbots.org/

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: Up to you!

AI in Your Life

- Game playing - chess, Go, jeopardy, Starcraft
- Voice recognition – Siri, Alexa, Google Assistant
- Recommender systems – Netflix, amazon
- Handwriting recognition
- Scene, object, face recognition: Face ID, MS seeing AI, image search (objects and faces)
- Automated logistics – UPS, US military
- Space exploration
- Automated science & medicine
- Robotics & Autonomous Vehicles

Example: AI at Amazon Warehouses

- Amazon uses robots to move products within its warehouses (deploys 200,000 robots)
- Amazon uses AI to predict demand
- 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

But Where's the 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
- Are we mimicking intelligence or getting closer to it by focusing on specific problems?

The sad (reassuring?) truth about modern AI

- Good news: Fears about the robot apocalypse are (for now) overblown
- Bad news:
  - Not because we're clever about preventing it
  - Because we aren't tackling:
    - Awareness
    - Deep understanding
    - High level reasoning
    - Robustness

Original terminator movie poster
Image downloaded from wikipedia

Eye on the prize

- AI’s narrower focus has earned the field credibility and practical successes, yet
- Some senior researchers complain that we have taken our eye off the prize:
  - Too much focus on specific problems
  - Lack of interest in general intelligence

- Are we ready to tackle general intelligence?
- Not this semester ☺

What is covered this semester?

- Search
  - Uninformed search, informed search, CSPs, classical planning
- Game Playing
  - minimax, alpha-beta search
- Logic and Knowledge Representation
  - Propositional logic, first order logic, theorem proving
- Reasoning under uncertainty
  - probability, Bayes nets, HMMs & tracking
- Probabilistic planning and reinforcement learning
- Introduction to machine learning
- Introduction to game theory

Major Topics *Not* Covered

- Natural Language
- Vision, except as application of computer vision

Class Mechanics

- Textbook: *Artificial Intelligence, A Modern Approach*, Russell & Norvig (*third* or *fourth* edition)
  - Semi-required
  - On Amazon, electronic version (3rd edition) also available
  - Please don’t steal my advisor’s textbook!

- Homework: 40%
  - High level discussion OK, write-up, coding must be your own (see matrix on class web page)
- Midterm: 30%
  - On Sakai (75 continuous minutes in 24 window), no collaboration
  - Starts during class time
- Final: 30%
  - On Sakai (180 continuous minutes in 24 window), no collaboration
  - Starts during scheduled final exam time

- Homework will be a mix of short proofs, algorithm design/analysis, and programming projects

Grading

- I tend to give challenging problems, 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
  - Typically chosen to ensure median grade of B+ or A-

Attendance

- You are not required to attend class – though it’s a ***good idea***
  - Videos and slides will be available
  - Everything in class is covered in readings
- You will not hurt my feelings if you skip class
- When you attend class:
  - Please be present in mind ***and*** body
  - Please don’t:
    - Text/email
    - Surf the web
    - Play games
- About multitasking:
  - You might think you can do this w/o cost, but you can’t
  - You aren’t fooling anybody

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
- Will aim to post materials and record sessions

Programming Assignments

- Based on the Berkeley Pacman framework

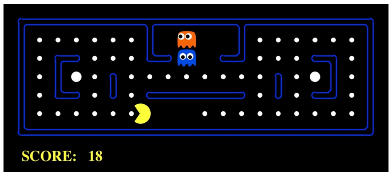

- Why?
  - It’s *really* well-done
  - Seeing ***your own code*** run AI algorithms is fun, motivating, and develops your intuitions
  - Even debugging is instructive

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

Pacman Limitations

- Recently ported to Python 3.x (may have issues with versions higher than 3.7)
- 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 for answers!!!
- Don’t troll for answers from previous semesters
- You may Google for definitions
- What you turn in must be your own work!!!

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
- We will catch efforts to hide copying of code!

- 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

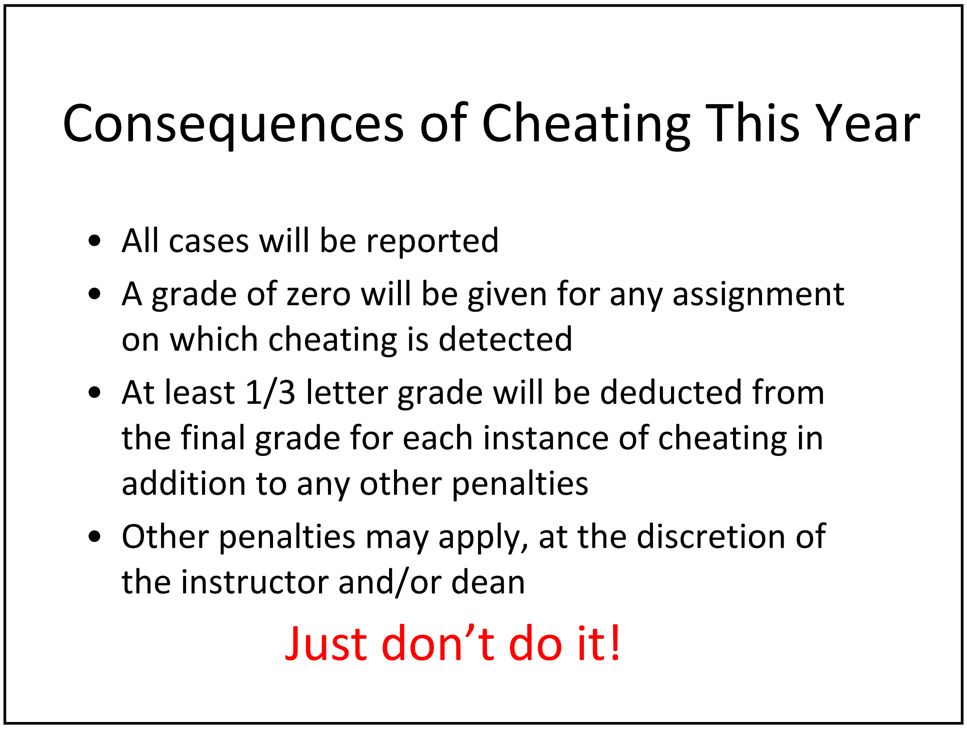

Consequences of Cheating This Year

- 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!

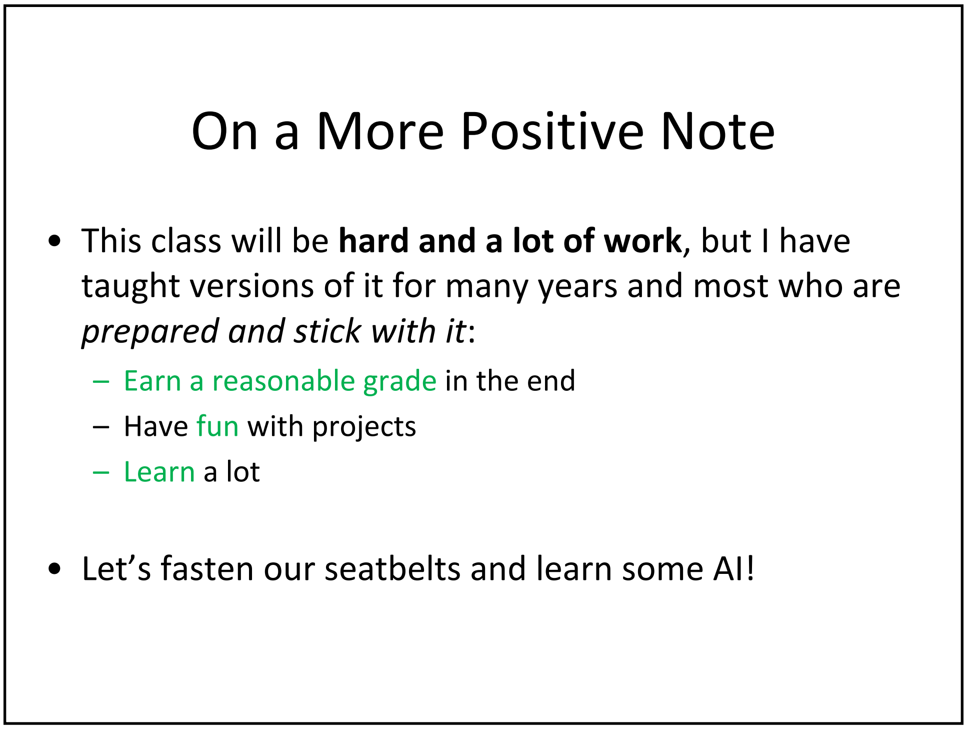

On a More Positive Note

- This class will 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
- Let’s fasten our seatbelts and learn some AI!