


## Artificial Intelligence

### Class 1: Course Overview



Dr Cynthia Matuszek (Dr M)  
cmat@umbc.edu

Slides adapted with thanks from: Dr. Marie deJardin, Dr. Tim Finin, Drs. Paula & David Matuszek

## Course Staff

- **Professor: Dr. M**
  - cmat@umbc.edu, ITE 331
  - Office hours: M 10-11, W 9:30-10:30, or by appointment
    - Concepts from lectures, general concerns, projects ...
- **TA: Erfan Noury Qaralajar,**
  - erfanel@umbc.edu, ITE TBD
  - Office hours: TBD
    - Homework assistance, coding assistance, general concerns...

## My Research

- **Robotics**
  - How can we go from industrial robots to useful robots in human environments? (Schools, cars, homes...)
- **Natural Language Processing**
  - How can computers learn to understand and speak human languages (English)?
- **Artificial intelligence**
  - How to get computers to behave in ways that we would consider to be "intelligent"
- **Human-Robot Interaction (HRI)**

## Today: Intro & Overview

- **Review of syllabus and schedule**
  - Academic honesty
  - Expectations
- **Brief history of AI**
  - What is AI? (and why is it so cool?)
  - What's the state of AI now?
- **Topics we'll cover**
- **What is 'intelligence'?**

From handout:  
<http://tiny.cc/ai-schedule>  
<http://tiny.cc/ai-class>  
<http://tiny.cc/ai-piazza>

## Classroom Policies

- Be courteous to classmates and instructors.
- No devices in use except when specified.
  - You don't learn as much.
  - People around you don't learn as much.
  - <http://tiny.cc/devices-in-class>
- No food or drink in this classroom.
  - Water is fine.

## Grading

*A little bit of stern stuff*  
 These policies are firm.  
 Don't ask me to make exceptions after the fact;  
 don't fuss about something being "just a little" late;  
 don't blame your internet.


|                     |     |
|---------------------|-----|
| Class participation | 5%  |
| Midterm             | 15% |
| Homework            | 30% |
| Quizzes and surveys | 2%  |
| Project             | 28% |
| Final exam          | 20% |

- **Grade questions:**
  - 24-hour "cooling" period
- **Grade changes/regrades:**
  - Requests to **professor and TA**
  - TA cannot change grades!

**Pop quiz: Can Dr M add?**

## Participation

- Attend class.
  - The program (and I) expect you to be here for the whole semester.
- Speak up.
  - Answer questions
  - Ask questions
  - Tell us your thoughts
  - There are lots of opportunities to talk here!
- Be active on Piazza.
  - Ask and answer questions.
  - Post links to interesting material.



7

## 5-6 Homework Assignments

- Written, problem set, and programming
  - Due at 11:59pm the day before class
  - Late: 25% off /day
- Assignments will be turned in **electronically**
  - Blackboard / online forms / email
  - Assignment will specify
  - 10% penalty for not following turn-in instructions
  - Example: Wrong file type
- Questions? **Piazza, then TA**

8

## Time Management

- Some things can be rescheduled
  - E.g., overlapping exams
  - If enough people have them
- Individual extensions *may* be given:
  1. With reasonable cause
  2. When made in advance
- Please talk to me!



9

## Some Advice

- Grad School is a hard transition
  - Everyone is smart now!
  - New expectations about writing, time management, and behavior
    - I know you don't believe me right now; that's cool
- Moving to a new place/field is hard
- There's a lot we can do to help – *if we know*.

10

## Reading

- Pre-readings: Do these **before that class**
  - First ~10 minutes will be Q&A about them
- Readings: Do these after class
  - More detail on concepts

Before class      After class

| CLASS | DATE | TOPIC                       | PRE-READINGS                               | READINGS     | HOMEWORK   | HANDOUTS / NOTES |
|-------|------|-----------------------------|--|--------------|--|------------------|
| 1     | 9/1  | Introduction and overview   | Class web page<br>Integrity policy         | Ch. 1        | Intro, HW2, HW3<br>Practice problem set<br>HW1, HW2  | Intro A, 13Ea    |
| 2     | 9/8  | Agents                      | 2.1, 2.2 (non-2.2.2);<br>Akin 2.3.1, 2.3.2 | Ch. 2        | Survey the<br>Practice problem set<br>HW 2, HW3, HW4 |                  |
| 3     | 9/8  | Problems solution to search | 3.1 (HW6, 3.1.1)                           | Ch. 3, 3.3.3 |  |                  |

11

## Academic Integrity

- Instructor's responsibilities:
  - Be respectful
  - Be fair
  - Be available
  - Tell the students what they need to know and how they will be graded
- Students' responsibilities:
  - Be respectful
  - Do not cheat, plagiarize, or lie, or help anyone else do so
  - Do not interfere with other students' academic activities

12

## Academic Integrity Policy

- “By enrolling in this course, each student assumes the responsibilities of an active participant in UMBC’s scholarly community, in which everyone’s academic work and behavior are held to the highest standards of honesty. **Cheating, fabrication, plagiarism, and helping others to commit these acts are all forms of academic dishonesty, and they are wrong.** Academic misconduct could result in disciplinary action that may include, but is not limited to, suspension or dismissal.”

[Statement adopted by UMBC’s Undergraduate Council and Provost’s Office]

13

## Integrity: Plagiarism

- Representing someone else’s work as your own is plagiarism.
  - **What if the reference is in the bibliography?**
    - If you didn’t explicitly quote the text you used *and* cite the source *where* you used the text, it is plagiarism.
  - **What if I only use some of the words?**
    - Scattering some of your own words and rephrasing isn’t enough. If the ideas are not restated entirely in your own words, it is plagiarism.

14

## Integrity: Plagiarism

- More Examples
  - **The introduction and background material are borrowed; all of the research is original.**
    - If somebody else’s words appear in any document that you claim is written by you, it is plagiarism.
  - **It was a draft or not an official assignment**
    - If you represented somebody else’s words as your own, even in an informal context, it is plagiarism.
  - **“But the professor told me to use that source!”**
    - Unless you are explicitly told to copy a quote from a source, you must write your answers *in your own words*.

15

## Integrity: Abetting

- This includes putting someone’s name on something when they didn’t work on it.
  - “This is just everyone on our team” is wrong.
- Know what your project partners are doing.
  - Their cheating can hurt you.
- Helping another student to cheat, falsify, or plagiarize will result in you receiving **the same penalty**.

16

## Integrity: What To Do

- You can **always** bring it to me.
- Cheating from you / in your group / etc.:
  - You **may** talk to them about it first
    - Unless it’s too late (it’s been turned in, the test is over)
      - Then you are abetting unless you report
    - Some people may get sneakier instead of improving
  - You **do not have to** talk to anyone but me.

17

## Integrity: Penalties

- Penalties depend on the offense and whether it recurs
- The **minimum** penalties are:
  - Receiving a zero on an assignment
  - Being required to redo the assignment, without credit, in order to pass the class
- Additional penalties may include:
  - Receiving a full grade reduction in the class
  - Failing the class without possibility of dropping it
  - Suspension or expulsion from the university

18

## About Groupwork

- Study groups are **encouraged!**
  - Talking about the homework is completely acceptable
  - Just don't share code
- Programming must be done **individually**
  - Programs must be written entirely by you
  - Copying another person's code is never acceptable
  - You can help debug
- Some homework is for 2-3 students working together
  - The assignment will say so; otherwise, it's individual.

19

## Availability & Communication

- **Post all questions to Piazza** (unless it violates integrity)
  - We will try to respond to Piazza posts immediately
  - Email takes 24-48 hours
- **Always send email to professor and TA**
  - Piazza, then TA, then prof+TA
- Office hours
- Drop by when my door is open
  - If I'm busy (often), we'll make an appointment
  - I will remain after class when I can

20

## Schedule

- You will check this pretty much every class

### CMSC 671: Principles of Artificial Intelligence

Syllabus • Schedule • Academic Integrity • Piazza Page

Please note:

• All readings refer to chapter sections in the textbook.

• Pre-readings should be done before class.

• All the reading material is covered on exams, so please make sure to bring questions about

• Due dates are the night before the listed class. For example, homework 1 is due on 9/19 at


The schedule:

This is a tentative schedule. Topics, reading assignments, homeworks, and exam dates are subject to change.

| CLASS | DATE | TOPIC                     | PRE-READINGS                                | READINGS | HOMEWORK   | HANDOUTS / NOTES |
|-------|------|---------------------------|---|----------|--|------------------|
| 1     | 9/1  | Introduction and overview | Class with your instructor and/or           | Ch. 1    | None assigned yet<br>Piazza questions out<br>10/1-10/4       | Intro & URLs     |
| 2     | 9/8  | Agents                    | 3.1, 2.2 Intro, 2.2.1;<br>4.1.1-4.1.3-2.3.2 | Ch. 2    | Survey due<br>Piazza questions due<br>By 11:59 PM--see above |                  |

21

## What is AI?



22

## Artificial Intelligence

- Key types
  - Strong AI: mental/thought capabilities equal to (or better than) human
  - Weak (bounded) AI: intelligent actions or reasoning in some limited situations
- “Human-level” intelligence
  - In what situation?
  - Internally?
- Self-awareness / Consciousness

23

## “Intelligence” is Problematic

- These are problematic.
  - How do we measure it?
  - What's an ‘intelligent action’?
    - In practice, ‘previously human only’
- Is there something ineffable missing?
  - What?
- How do we test?




24

## AI: A Vision

- Could an intelligent agent living on your home computer...
  - **Manage your email**
  - Coordinate your **work and social activities**
  - Help **plan your vacations**
  - **Watch your house** while you take those vacations?







25

## Main Goals of AI

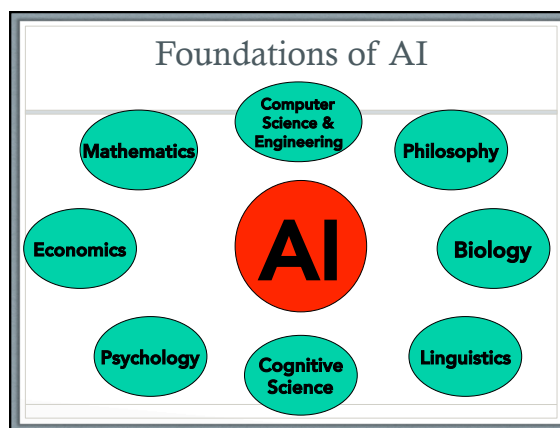
- **Represent** and **store** knowledge
- **Retrieve** and **reason** about knowledge
- **Behave** intelligently in complex environments
- **Learn** from environment and interactions
- Develop interesting and useful **applications**
- **Interact** with people, agents, and environment

26

## Why AI?

- **Engineering**
  - To get machines to do a wider variety of useful things
    - Understand spoken natural language
    - Recognize individual people in visual scenes
    - Find the best travel plan for your vacation
- **Cognitive Science**
  - Help understand how natural minds work
    - Visual perception, memory, learning, language, etc.
- **Philosophy**
  - As a way to explore interesting (and important) philosophical questions

27



## Big Questions

- Can machines think?
- If so, how?
- If not, why not?
- What does this say about human beings?
- What does this say about the mind?

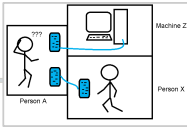
36

## What's Easy and What's Hard?

- It's easi(er) to mechanize high-level tasks
  - Symbolic integration
  - Proving theorems
  - Playing chess
  - Medical diagnosis
- It's hard to mechanize tasks that lots of animals can do
  - Walking around without running into things
  - Catching prey and avoiding predators
  - Interpreting complex sensory information (e.g., visual, aural, ...)
  - Modeling the internal states of other animals from their behavior
  - Working as a team (e.g., with pack animals)
- Is there a fundamental difference?

37

## Turing Test



- Three rooms:
  - 1 person, 1 computer, and 1 interrogator
  - The interrogator can communicate with the other two
  - The interrogator tries to decide which is the person
  - Both try to convince the interrogator they are the person
- If the machine succeeds, the machine can think  
 ...Right? (no)

Image: <http://pinofreethinkers.org/2012/06/22/turing-tremendous-talent-and-trenchant-test/turing-test/>

## The Loebner Contest

- A modern version of the Turing Test, held annually
  - \$100,000 cash prize.
  - Hugh Loebner was once director of UMBC's Academic Computing Services (née UCS)
- Restricted topic (removed in 1995) and limited time.
- Participants: set of humans, set of computers, set of judges.
- Scoring
  - Rank from least human to most human.
  - Highest median rank wins \$2000

39

## What Can AI Systems Do Now?



|   |  |
|---|--|
| <ul style="list-style-type: none"> <li>• <b>Computer vision:</b> face recognition from a large set</li> <li>• <b>Natural language processing:</b> machine translation</li> <li>• <b>Expert systems:</b> medical diagnosis in a narrow domain</li> <li>• <b>Spoken language systems:</b> ~1000 word continuous speech</li> <li>• <b>Planning and scheduling:</b> Hubble Telescope experiments</li> </ul> | <ul style="list-style-type: none"> <li>• <b>Robotics:</b> autonomous (mostly) automobile</li> <li>• <b>User modeling:</b> Bayesian reasoning in Windows help (the infamous paper clip...)</li> <li>• <b>Games:</b> Grand Master level in chess (world champion), perfect play in checkers, Go</li> <li>• <b>Search:</b> You've used Google.</li> <li>• <b>Learning:</b> So much learning.</li> </ul> |
|---|--|


## What Can't AI Systems Do Yet?

- Understand natural language robustly
- Learn a natural language
- Surf the web
- Interpret an arbitrary visual scene
- **Play Go as well as the best human players**
- Construct plans in dynamic real-time domains
- Refocus attention in complex environments
- Perform life-long learning



Exhibit true autonomy and intelligence?

41












### Who Does AI?

- Academic researchers (perhaps the most Ph.D.-generating area of computer science in recent years)
  - Some top schools: CMU, Stanford, Berkeley, MIT, UW, UMd, U Alberta, UT Austin, ... (and, actually, UMBC!)
- Government and private research labs
  - NASA, NRL, NIST, IBM, AT&T, SRI, ISI, MERL, ...
- Lots of companies!
  - Google/Alphabet, Microsoft, Amazon, Honeywell, Teknowledge, SAIC, MITRE, Fujitsu, Global InfoTek, BodyMedia, ...









42

# Applications

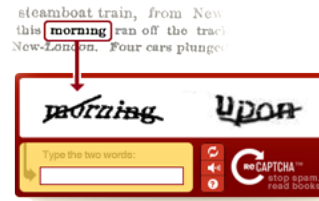
43

## Game Playing



44

## Text/Sketch Recognition



45

## User Modeling & NLP



46

## Robotics



47


## Knowledge Representation

Watson



48

## Evolutionary Art

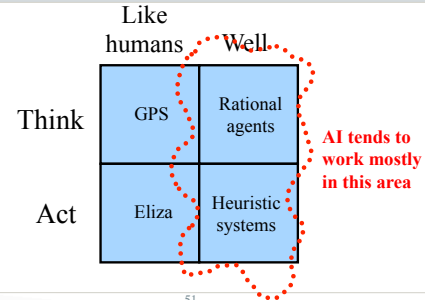


## Computer Vision



50

## Possible Approaches



51


## Thinking Well



- Develop formal models of knowledge representation, reasoning, learning, memory, and problem solving, that can be rendered in algorithms.
- There is often an emphasis on systems that are provably correct, and guarantee finding an optimal solution.

52

## Acting Well



- For a set of inputs, generate an appropriate output that is not necessarily correct but gets the job done.
  - A **heuristic (heuristic rule, heuristic method)** is a rule of thumb, strategy, trick, or any other kind of device which drastically limits search for solutions in large problem spaces.
  - Heuristics do not guarantee optimal solutions; in fact, they do not guarantee any solution at all: **all that can be said for a useful heuristic is that it offers solutions which are good enough most of the time.**
- Feigenbaum and Feldman, 1963, p. 6

53

## Thinking Like Humans



- Cognitive science approach
- Focus not just on behavior and I/O
  - Also look at reasoning process.
- Computational model reflects “how” results were obtained
- Provide a new language for expressing cognitive theories and new mechanisms for evaluating them
- GPS (General Problem Solver):
  - Not just to produce humanlike behavior, but to produce a sequence of steps of the reasoning process similar to the steps followed by a person

54

## Acting Like Humans




- Behaviorist approach.
- Not about how you get results, just the similarity to what human results are.
- Exemplified by the Turing Test

55



## What about Statistical Methods?

|       | Like humans | Well              |
|-------|-------------|-------------------|
| Think | GPS         | Rational agents   |
| Act   | Eliza       | Heuristic systems |



56

## For Next Time

- Due at 11:59pm before next class:
  - Fill out the survey
    - Posted before weekend
    - Where? ☺
  - Read academic integrity statement
  - Read syllabus
  - Sign up for Piazza and join this class
- Look at the reading lists
- Do pre-reading for next time!

57