


# Artificial Intelligence

## Class 1: Course Overview



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

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

1

## Course Staff

- Professor: Cynthia Matuszek (Dr. M)
  - cmat@umbc.edu, ITE 331
  - Office hours: **M 11 - 12, F 1:30 - 2:30**, or by appointment
    - Concepts from lectures, general concerns, projects ...
- TA: Pat Jenkins
  - pjenk1@umbc.edu, ITE 334
  - Office hours: **Tu 2:30 - 3:30, W 12 - 1**
    - Homework assistance, coding assistance, general concerns...
    - Most homework help

2

## 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)

3

## Today: Intro & Overview

- Review of syllabus and schedule
  - Academic honesty
  - Expectations and conduct
  - Policies, grading, etc.
- Brief history of AI
  - What is AI? (and why is it so cool?)
  - What's the state of AI now?
- What is 'intelligence'?

*but important*

4


## Resources

- The syllabus is longish, but important:
  - <http://tiny.cc/ai-class>
- Lecture topics and reading on the schedule:
  - You'll want to check this every single class.
  - <http://tiny.cc/ai-schedule>
- We will use Piazza extensively:
  - So get your account sorted right away!
  - <http://tiny.cc/ai-piazza>

5

## 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 in this classroom.



6

## Harassment and Inclusivity

**All students are entitled to a safe, respectful, and inclusive learning environment both inside and outside the classroom.**

- No discrimination, exclusion, or harassment
- Respectful, inclusive discussion
  - Use one another's preferred names, pronouns, etc.
- If there are problems
  - Talk to me, the TA, or someone else
  - There are resources on the syllabus

7

## Grading

- Grades in Blackboard
  - Know your grades
  - but also*
  - Keep track of what's left
- Grade questions:
  - 24-hour "cooling" period before *any* discussion
- Grade changes/regrades:
  - Requests to **professor and TA**
  - TA cannot change grades!

Class participation*	15%
Midterm	15%
Homework	30%
Project	20%
Final exam	20%

8

8

## Participation

- Attend class.
  - The program (and I) expect you to be here for the **entire** semester.
- Speak up.
  - Ask & answer questions
  - Tell us your thoughts
- Be active on Piazza.
  - Ask and answer questions.
  - Post links to interesting material.
- Do any take-home quizzes and/or surveys.



9

9

## 5-6 Homework Assignments

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

10

10

## 5-6 Homework Assignments

- Written text, problem sets, and programming
  - **D** **These policies are firm.**
  - **L** **Don't ask for exceptions after the fact.**
  - **L** **Don't tell us it's "just a little bit" late.**
  - **Ass** **You have multiple turnins. Use them!**
- Assignment will specify Blackboard, forms, or email
- Sometimes 10% penalty for not following instructions
  - Example: Wrong file type
- Questions? **Piazza, then TA**

11

11

## 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
- Talk to me!



12

12

## A Word About This Class

- The biggest complaint people have had:

*I needed an easy class to go with OS and algorithms and this class was supposed to be easy*

- This **is not** an easy class
  - OS, AI, and algorithms? Not wise.
- We cover a lot of ground – you must keep up
  - Time-consuming** and **theory-heavy**
  - Difficulty** depends on your background

13

13

## And Some Advice

- Grad School is a hard transition
  - Everyone around you 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

14

14

## Reading

- Pre-readings: Do these **before that class**
  - First ~10 minutes will be Q&A about them

Week	Topic	Readings	Lecture	Discussion	Assignment
1	Introduction to AI	...	...	...	...
2	Search	...	...	...	...
3	Machine Learning	...	...	...	...
4	Natural Language Processing	...	...	...	...

- Readings: Do these after class
  - More detail on concepts

15

15

## Academic Integrity

ABOUT the whole thing where CS is fundamentally about resources it reinventing wheels, but why they need to do some stuff to learn because it's not always out there

- 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

16

16

## 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]

17

17

## Integrity: Plagiarism

- Representing someone else's work as your own is plagiarism.
- What 'counts' is cultural. It is probably different now from what you are used to.
  - This means you must be very careful!
- We use research and publication standards.

18

18

## Integrity: Plagiarism

- **We are about to talk about plagiarism a whole lot, because I want you to succeed. Don't tune out.**
- We use research and publication standards.

19

19

## Integrity: Plagiarism

- Examples
  - **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.

20

20

## 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*.

21

21

## Integrity: Plagiarism

- Any time you find yourself copy-pasting—even a **sentence**—you are plagiarizing.
- Copying code from **any** source without citations is plagiarizing.
- The 1<sup>st</sup> time, you **must redo the assignment for 0 credit**.
- You may also:
  - Get a full grade reduction in the class
  - Fail the class without possibility of dropping it
  - Be suspended or expelled from university

22

22

## 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**.

23

23

## Integrity: What To Do

- You can **always** bring it to me.
- Cheating by others / 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
  - This is not the preferred option
- You **do not have to** talk to anyone but me
- If you think *you* made a mistake, bring it to me.

24


24

## 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 discuss conceptually; never look at code
- Some homework is for 2-3 students working together
  - The assignment will say so; otherwise, it's individual.

25

## Availability & Communication

- **Post all questions to Piazza** (without code)
  - 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

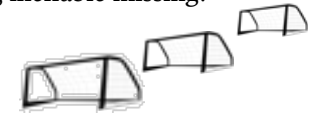
26

## All right! Fun stuff! What is AI?



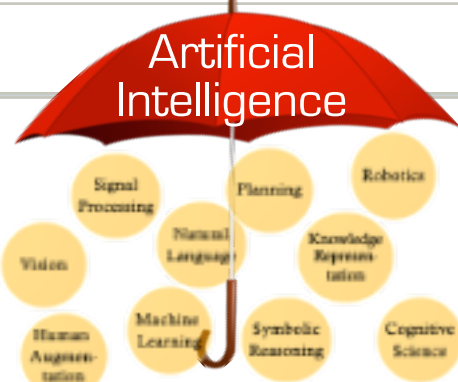
27

## “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? 

28

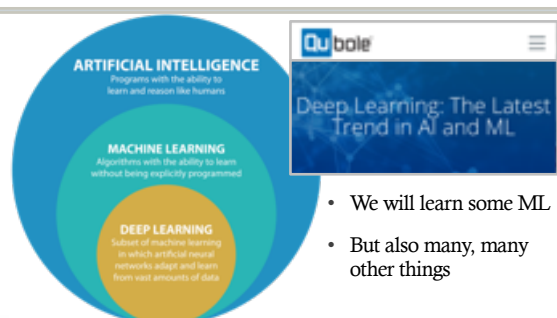
## Artificial Intelligence



**And many more**

29

## But Deep Learniiiiinnng



- We will learn some ML
- But also many, many other things

30

## 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
- Also: scruffy/neat, symbolic/statistical, ...
- “Human-level” intelligence
  - In what situation?
  - Internally?
- Self-awareness / Consciousness

31

## 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?**
  - **Drive you** to that vacation?
  - Do your **laundry** while you are away?



32

## Main Goals of AI

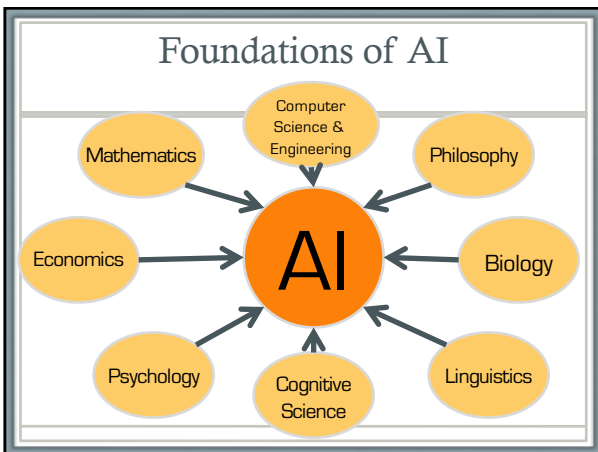
- **Represent** and **store** knowledge
- **Retrieve** and **reason** about knowledge
- **Behave** intelligently in complex environments
- **Learn** from environment and interactions
- Develop interesting and useful **applications**
- **Interact** with people, agents, and environment

33

## 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

34



35

## 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?

43


## 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?

44

44

## 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: filipinfreethinkers.org/2012/06/23/turings-tremendous-talent-and-trenchant-test/turing-test*


45

## What Can AI Systems Do Now?

- **Computer vision:** face recognition from a large set
- **Natural language processing:** machine translation
- **Expert systems:** medical diagnosis in a narrow domain
- **Spoken language systems:** ~1000 word continuous speech
- **Planning and scheduling:** Hubble Telescope experiments
- **Robotics:** autonomous (mostly) automobile
- **User modeling:** Bayesian reasoning in Windows help (the infamous paper clip...)
- **Games:** Grand Master level in chess (world champion), perfect play in checkers, Go
- **Search:** You've used Google.
- **Learning:** So much learning.



47


## What Can't AI Systems Do Yet?



- Understand natural language robustly
- Learn a natural language **Exhibit true autonomy and intelligence?**
- 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

48

48







  


  
**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, ...

49

49

# Applications

50

50

### Game Playing



51

51

### Text/Sketch Recognition

steamboat train, from New  
this **morning** ran off the track  
New-London. Four cars plunged



52

52


### User Modeling & NLP



53

53

### Robotics



54

54

### Knowledge Representation

Watson




55

56

55

### Evolutionary Art

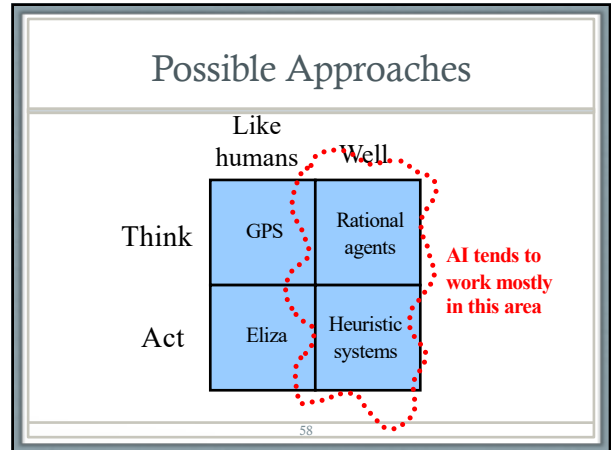


56





57



58

## Thinking Well

	Like humans	Well
Think	GPS	Rational agents
Act	Eliza	Heuristic systems

- 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.

59

## Acting Well

	Like humans	Well
Think	GPS	Rational agents
Act	Eliza	Heuristic systems

- 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

60

## Thinking Like Humans

	Like humans	Well
Think	GPS	Rational agents
Act	Eliza	Heuristic systems

- 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

61

## Acting Like Humans


	Like humans	Well
Think	GPS	Rational agents
Act	Eliza	Heuristic systems

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

62

## What about Statistical Methods?

	Like humans	Well
Think	GPS	Rational agents
Act	Eliza	Heuristic systems



63

63

## For Next Time

- Before next class:
  - Read academic integrity statement
  - Read syllabus
  - Sign up for Piazza and join this class
  - **Do pre-reading for next time!**
- Look at the reading lists

64

64