

# CMSC 471/671

## Fall 2006

Tim Finin, [finin@cs.umbc.edu](mailto:finin@cs.umbc.edu)

### Today's class

- Course overview
- Introduction
  - Brief history of AI
  - What *is* AI? (and why is it so cool?)
  - What's the state of AI now?

## Course Overview

### Course materials

- Blackboard under 471/671 and <http://www.cs.umbc.edu/471/fall06/>
  - Course description and policies (main page)
  - Course syllabus, schedule (subject to change!), and slides
  - Homework assignments
  - Sample code
  - Links to AI papers and resources
- Blackboard discussion list
  - We'll use this in lieu of a mailing list
  - Read it several times a week
  - Look for answers to questions before you ask

## Homework and grading policies

- Eight to ten homework assignments (mix of written and programming)
- One-time extensions of up to a week will generally be granted *if requested in advance*
- Last-minute requests for extensions will be denied
- Late policy: being refined, see web next week
- **NOTE ON READING:** *Please do the reading before each class!*

## Programming

- Learning Prolog will be part of the course
  - Why?
  - Sicstus Prolog is installed on gl.umbc.edu
  - We'll use SWI-Prolog, tho
  - It should be installed on gl.umbc.edu next week
  - It's free (GPL) and runs on Windows, OSX and Linux
- No Lisp
  - Why not?
- Some assignments may require using other systems
  - E.g., C5 decision tree learning system, Jess production rule system

## Exams

- Midterm exam
  - In class in mid October (e.g., 10/17)
  - About 15% of grade
- Final exam
  - At regularly scheduled time
  - About 25% of grade
  - Comprehensive, but with an emphasis on the last half of the material (e.g., 30/70 split)

## 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 to do so
  - Do not interfere with other students' academic activities
- Consequences include (but are not limited to) a reduced or failing grade on the assignment, or in the class

## Instructor availability

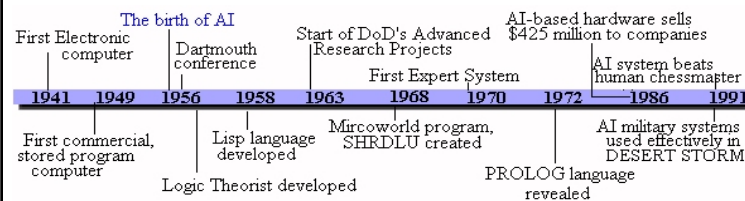
- Professor Finin
  - Official office hours: Tues. 1:00-2:00, Thr 1:00-2:00 (ITE 329)
  - Appointments may also be made by request (24 hours notice is best)
  - Drop in whenever my door is open
  - Direct general questions (i.e., those that other students may also be wondering about and that Google can't answer) to Blackboard discussion list
  - We will try to respond to postings on the discussion list or private email messages within 24 hours
- TAs
  - Justin Martineau (100%)
  - Brandon Wilson (50%)
  - See blackboard site

## 471 vs. 671?

- Grad students will be asked to do more work
  - Extra parts on assignments
  - Possible grad project
- Undergrads may do some of these for extra credit
- The two populations will be graded separately

## What is AI??

## History



## Eliza

- ELIZA: A program that simulated a psychotherapist interacting with a patient and successfully passed the Turing Test.
- Coded at MIT during 1964-66 by Joel Weizenbaum in FORTRAN
- First script was DOCTOR.
  - The script was a simple collection of syntactic patterns not unlike regular expressions
  - Each pattern had an associated reply which might include bits of the input (after simple transformations (my → your))
- Weizenbaum was shocked at reactions:
  - Psychiatrists thought it had potential.
  - People unequivocally anthropomorphized.
  - Many thought it solved the NL problem.

## Eliza

- This is a transcript of a conversation with an Eliza-like program found in Emacs

- Try typing

*Meta-x  
doctor*

in your favorite Emacs.

- I am the psychotherapist. Please, describe your problems. Each time you are finished talking, type RET twice.
- i'm having a hard time with my students
- Why do you say that?
- well, they keep asking me for money
- Maybe your plans have something to do with this.
- what plans?
- Is it because of your life that you say what plans?
- No, it's because I didn't know what you were talking about
- Is it because you didn't know what I was talking about that you came to me?
- no, it's because i wanted to see if you could pass the turing test
- Does it bother you that you wanted to see if I could really pass the turing test?
- .....
- How do you reconcile problems at school?
- This is a waste of time. Good bye
- My secretary will send you a bill.

## Colby's PARRY

- Kenneth Colby modeled a paranoid using the same techniques circa 1968.
- PARRY has basic emotions. If it gets angry, its replies become more hostile.
- In a Turing test, expert psychiatrists were unable to distinguish PARRY's ramblings from those of real paranoids.

H: Good afternoon. My name is Doctor Jones.  
P: Pleased to meet you.  
H: How are you?  
P: All right I guess.  
H: Could you tell me your name?  
P: Pat Smith.  
H: I though you might have some problems.  
P: I've been upset lately.  
H: By what?  
P: People get on my nerves sometimes.  
H: Tell me more.  
P: Do you know anything about bookies?  
...

## Parry meets Eliza

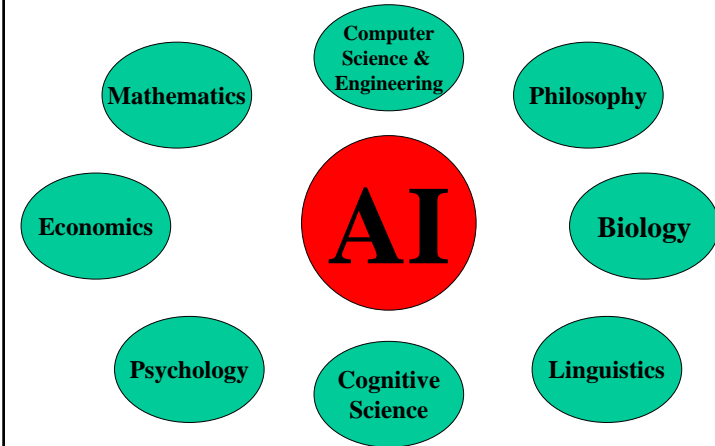
*Please tell me \*this\* isn't All!*

P: I don't understand your motives.  
E: Why don't you understand my motives?  
P: Shouldn't I be?  
E: Why do you say that just now?  
P: That's just the way it is.  
E: Suppose that's the way it were not.  
P: So?  
E: I am not sure I understand you fully.  
P: You harp on the same questions too much.  
E: What does that similarity suggest to you?  
P: I don't get you.  
E: Do you wish that you did get I?

## Recent AI History

- AI has had it's ups and downs
- It's ascending now, IMHO
- Hot topics:
  - Text mining, natural language technology
  - Anything on the web, e.g., Semantic Web, Google's iindex as our giant shared knowledge base
  - Machine learning
  - Autonomous vehicles
  - AI for games?

## Foundations of AI



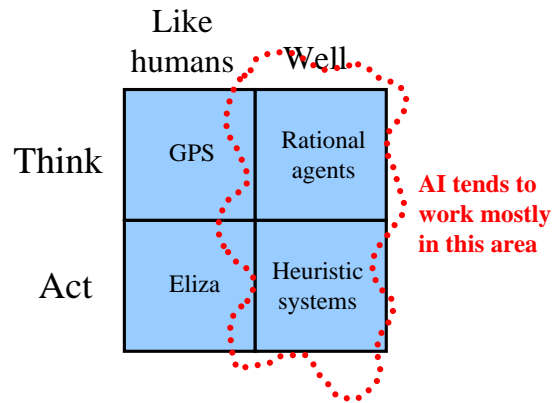
## Big questions

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

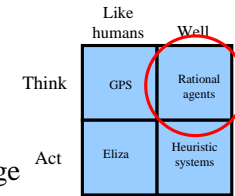
## Why AI?

- **Engineering:** To get machines to do a wider variety of useful things
  - e.g., understand spoken natural language, recognize individual people in visual scenes, find the best travel plan for your vacation, etc.
- **Cognitive Science:** As a way to understand how natural minds and mental phenomena work
  - e.g., visual perception, memory, learning, language, etc.
- **Philosophy:** As a way to explore some basic and interesting (and important) philosophical questions
  - e.g., the mind body problem, what is consciousness, etc.

## Possible approaches



## Think well

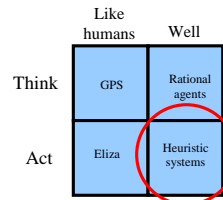


- Develop formal models of knowledge representation, reasoning, learning, memory, problem solving, that can be rendered in algorithms.
- There is often an emphasis on a systems that are provably correct, and guarantee finding an optimal solution.

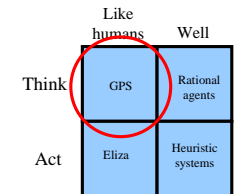
## Act well

- For a given 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, simplification, 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



## Think like humans



- Cognitive science approach
- Focus not just on behavior and I/O but also look at reasoning process.
- Computational model should reflect “how” results were obtained.
- Provide a new language for expressing cognitive theories and new mechanisms for evaluating them
- GPS (**General Problem Solver**): Goal not just to produce humanlike behavior (like ELIZA), but to produce a sequence of steps of the reasoning process that was similar to the steps followed by a person in solving the same task.

## Act like humans

- Behaviorist approach.
- Not interested in how you get results, just the similarity to what human results are.
- Exemplified by the Turing Test (Alan Turing, 1950).
- Has applications in interactive entertainment (e.g., computer games)

	Like humans	Well
Think	GPS	Rational agents
Act	Eliza	Heuristic systems

## What's easy and what's hard?

- It's been easier to mechanize many of the high-level tasks we usually associate with "intelligence" in people
  - e.g., symbolic integration, proving theorems, playing chess, medical diagnosis
- It's been very 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 between the two categories?

## Turing Test

- Three rooms contain a person, a computer, and an interrogator.
- The interrogator can communicate with the other two by 'teleprinter'.
- The interrogator tries to determine which is the person and which is the machine.
- The machine tries to fool the interrogator into believing that it is the person.
- If the machine succeeds, then we conclude that the machine can think.

## The Loebner contest

- A modern version of the Turing Test, held annually, with a \$100,000 cash prize.
- Hugh Loebner was once director of UMBC's Academic Computing Services (née UCS)
- <http://www.loebner.net/Prizef/loebner-prize.html>
- Restricted topic (removed in 1995) and limited time.
- Participants include a set of humans and a set of computers and a set of judges.
- Scoring
  - Rank from least human to most human.
  - Highest median rank wins \$2000.
  - If better than a human, win \$100,000. (Nobody yet...)

## What can AI systems do?

Here are some example applications

- **Computer vision:** face recognition from a large set
- **Robotics:** autonomous (mostly) automobile
- **Natural language processing:** simple machine translation
- **Expert systems:** medical diagnosis in a narrow domain
- **Spoken language systems:** ~1000 word continuous speech
- **Planning and scheduling:** Hubble Telescope experiments
- **Learning:** text categorization into ~1000 topics
- **User modeling:** Bayesian reasoning in Windows help (the infamous paper clip...)
- **Games:** Grand Master level in chess (world champion), checkers, etc.

## What can't AI systems do yet?

- Understand natural language robustly (e.g., read and understand articles in a newspaper)
- Surf the web and find interesting knowledge
- Interpret an arbitrary visual scene
- Learn a natural language
- Play Go well
- Construct plans in dynamic real-time domains
- Refocus attention in complex environments
- Perform life-long learning

**Exhibit true autonomy and intelligence!**



Carnegie Mellon



### Who does AI?



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



Honeywell



Microsoft

MITRE

FUJITSU

## AI at UMBC in CSEE

- Maple Lab (desJardins)
  - Multiagent systems, planning, machine learning
- Coral Lab (Oates)
  - Machine learning, robotics, cognitive science
- Ebiquity Lab (Finin, Peng, Joshi, Yesha)
  - Semantic web, multiagent systems, pervasive computing, text mining
- Institute for Language and Information Technology (Nierenberg, McShane, Beale)
  - NLP, information extraction, machine translation, intelligent tutors
- DIADIC Lab (Kargupta)
  - Datamining, bioinformatics