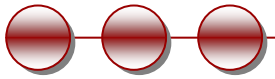

CMSC 671
Fall 2010

Tue 8/31/10
Introduction to Artificial Intelligence
Course Overview

Professor Rosa Laura Zavala,
laura.zavala@umbc.edu,
ITE 373, 410-455-8775

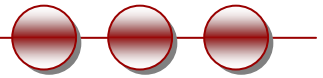

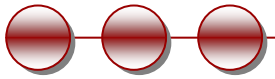

What is AI??

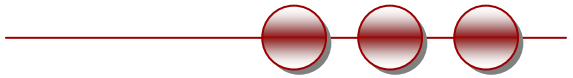

Big Questions

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

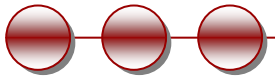

Eliza: A computer therapist

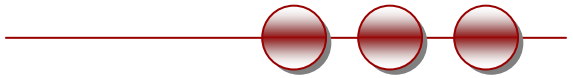

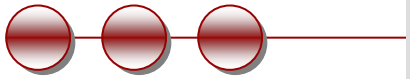

- This is a transcript of a conversation with an Eliza-like program found in Emacs
- Try typing

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?**

Eliza

- ELIZA: A program that simulated a psychotherapist interacting with a patient and successfully passed the Turing Test.
- Coded at MIT during 1964-1966 by Joel Weizenbaum.
- 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.

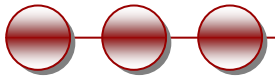

- Kenneth Colby modeled a paranoid using the same techniques circa 1968.
- PARRY has basic emotions. If it gets angry, its replies become more hostile.
- In a Turing test, expert psychiatrists were unable to distinguish PARRY's ramblings from those of real paranoids.

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?
...

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…)

Please tell me *this* isn't AI!

Parry

- 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?

History

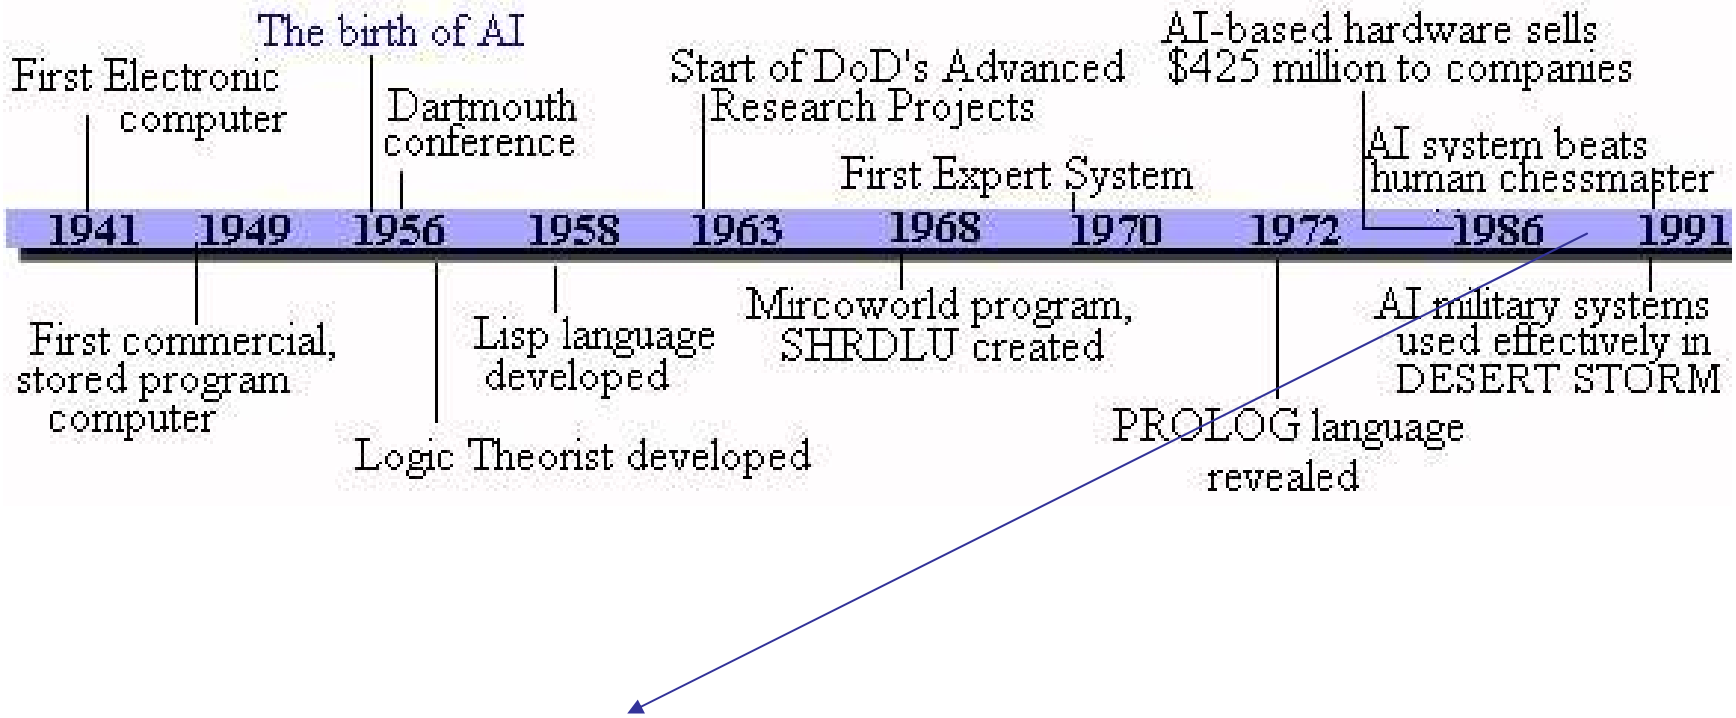

Rodney A. Brooks , Intelligence without Reason (Behavior-based AI)

Does AI involve reasoning?

- Back to our initial questions *what is AI? what is Intelligence?*

- High-level reasoning
- Behavior-based approach to robotics
  - cognition is only in the eye of an observer

History

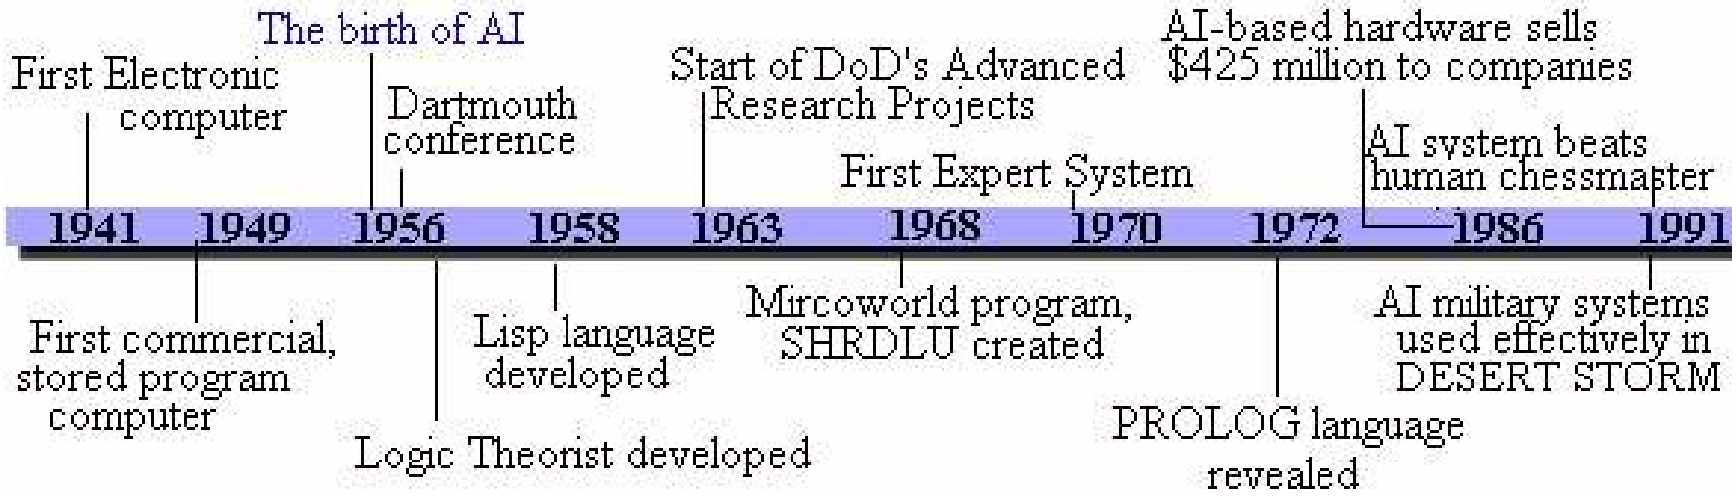

- 1997: Deep Blue beats Garry Kasparov (world champion)
- 1998: Founding of Google
- 2000: Interactive robot pets
- 2004: First DARPA Grand Challenge robot race
- 2004: Commercial recommender systems (TIVO, amazon.com)
- 2007: Checkers is solved!

Let's do some AI

Tic Tac Toe

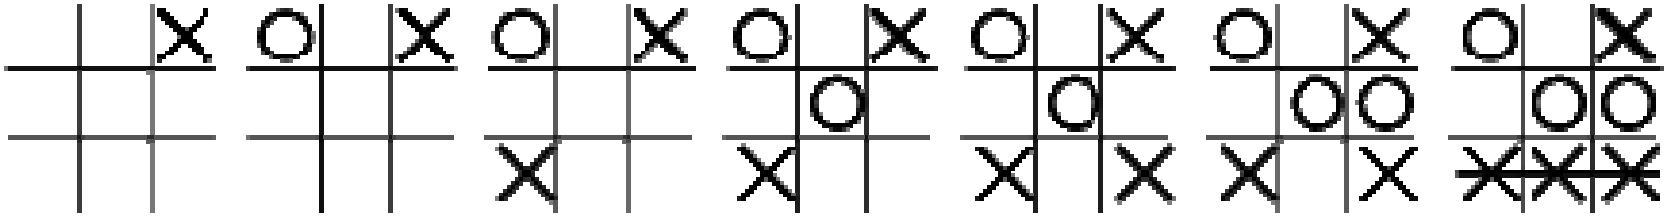

How would you build of a program that plays (optimally) Tic Tac Toe?

Let's do some AI

It is straightforward to write a computer program to play tic-tac-toe perfectly

State space complexity = 765

Game tree complexity = 26830

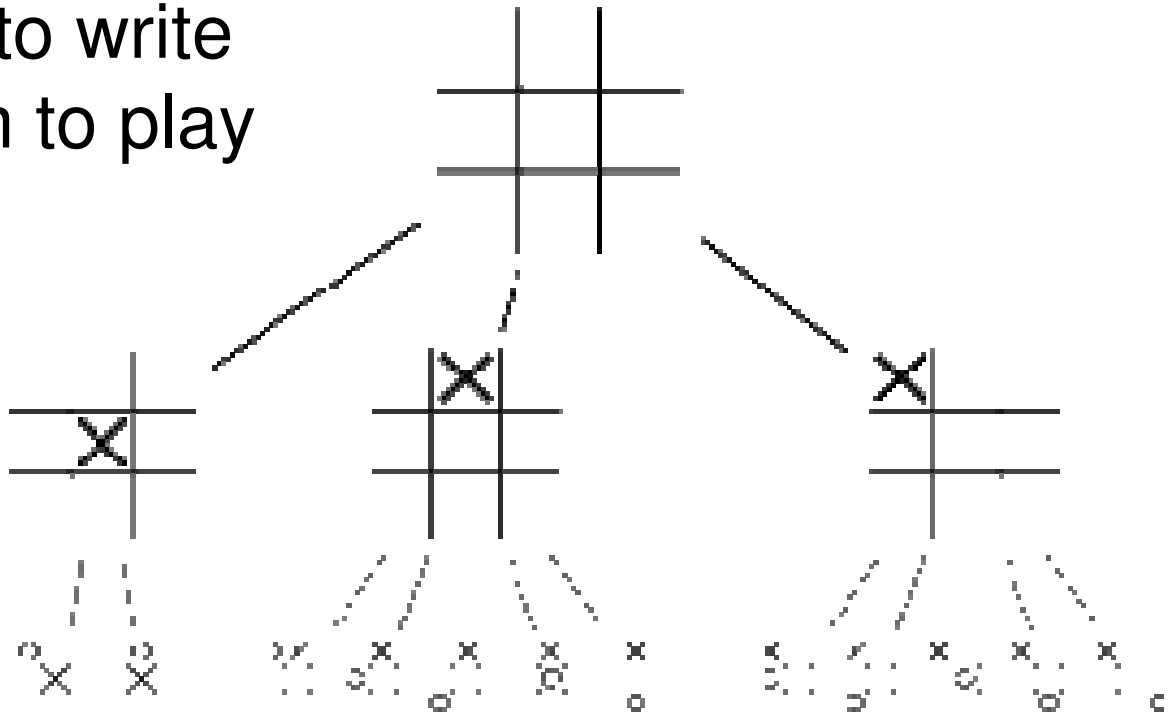

Best play from both parties leads to a draw !!!

Checkers is solved!

- Roughly 500 billion billion possible positions (5×10^{20})
- Game complexity of approx. 10^{20}
- The checkers program CHINOOK cannot lose (it can draw)

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.

Why AI?

- **Cognitive Science:** As a way to understand how the human mind works
  - e.g., visual perception, memory, learning, language, etc.

Why AI?

- **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

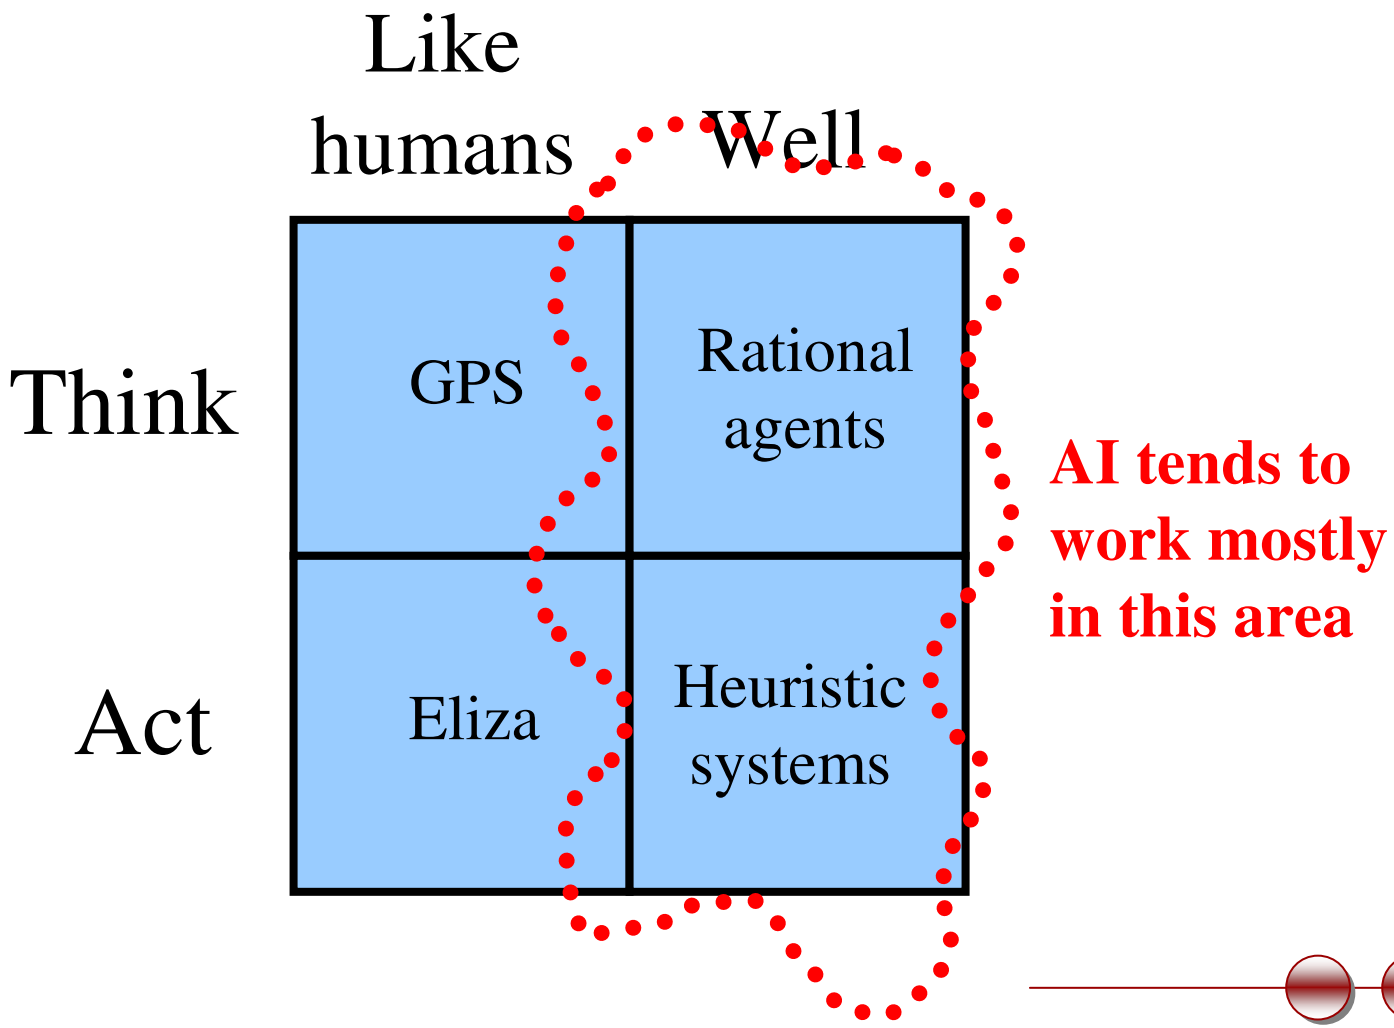

Think Well

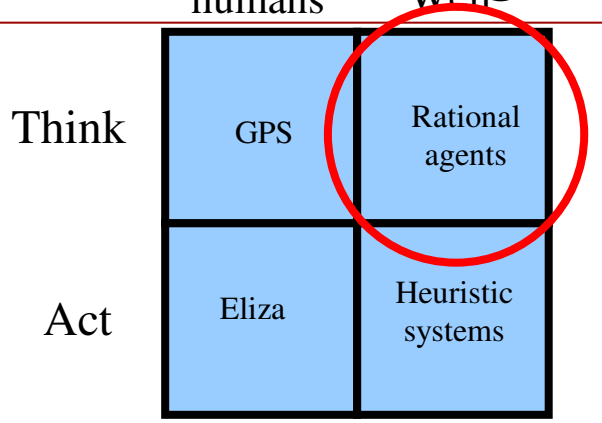

- Develop formal models of knowledge representation, reasoning, learning, memory, and problem solving, that can be rendered in algorithms.
- There is often an emphasis on systems that are provably correct, and guarantee finding an optimal solution.

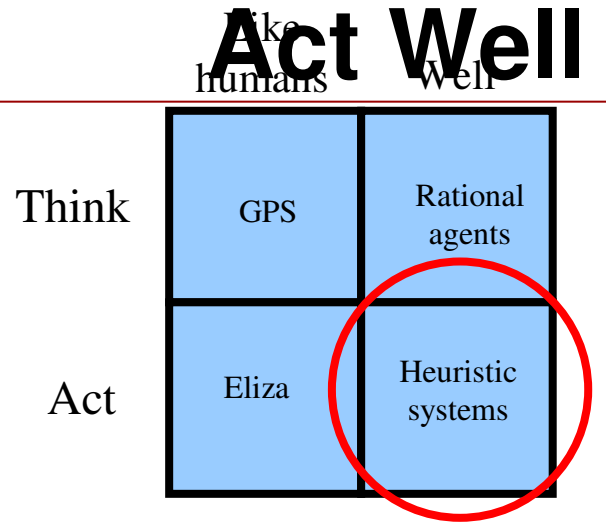

- For a given set of inputs, generate an appropriate output that is not necessarily correct but gets the job done.
- A **heuristic** (**heuristic rule**, **heuristic method**) is a rule of thumb, strategy, trick, simplification, or any other kind of device which drastically limits search for solutions in large problem spaces.

So, What is AI

- Design of agents that act rationally: act so as to achieve the best outcome or, when there is uncertainty, the best expected outcome.

Main Goals of AI

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

Foundations of AI

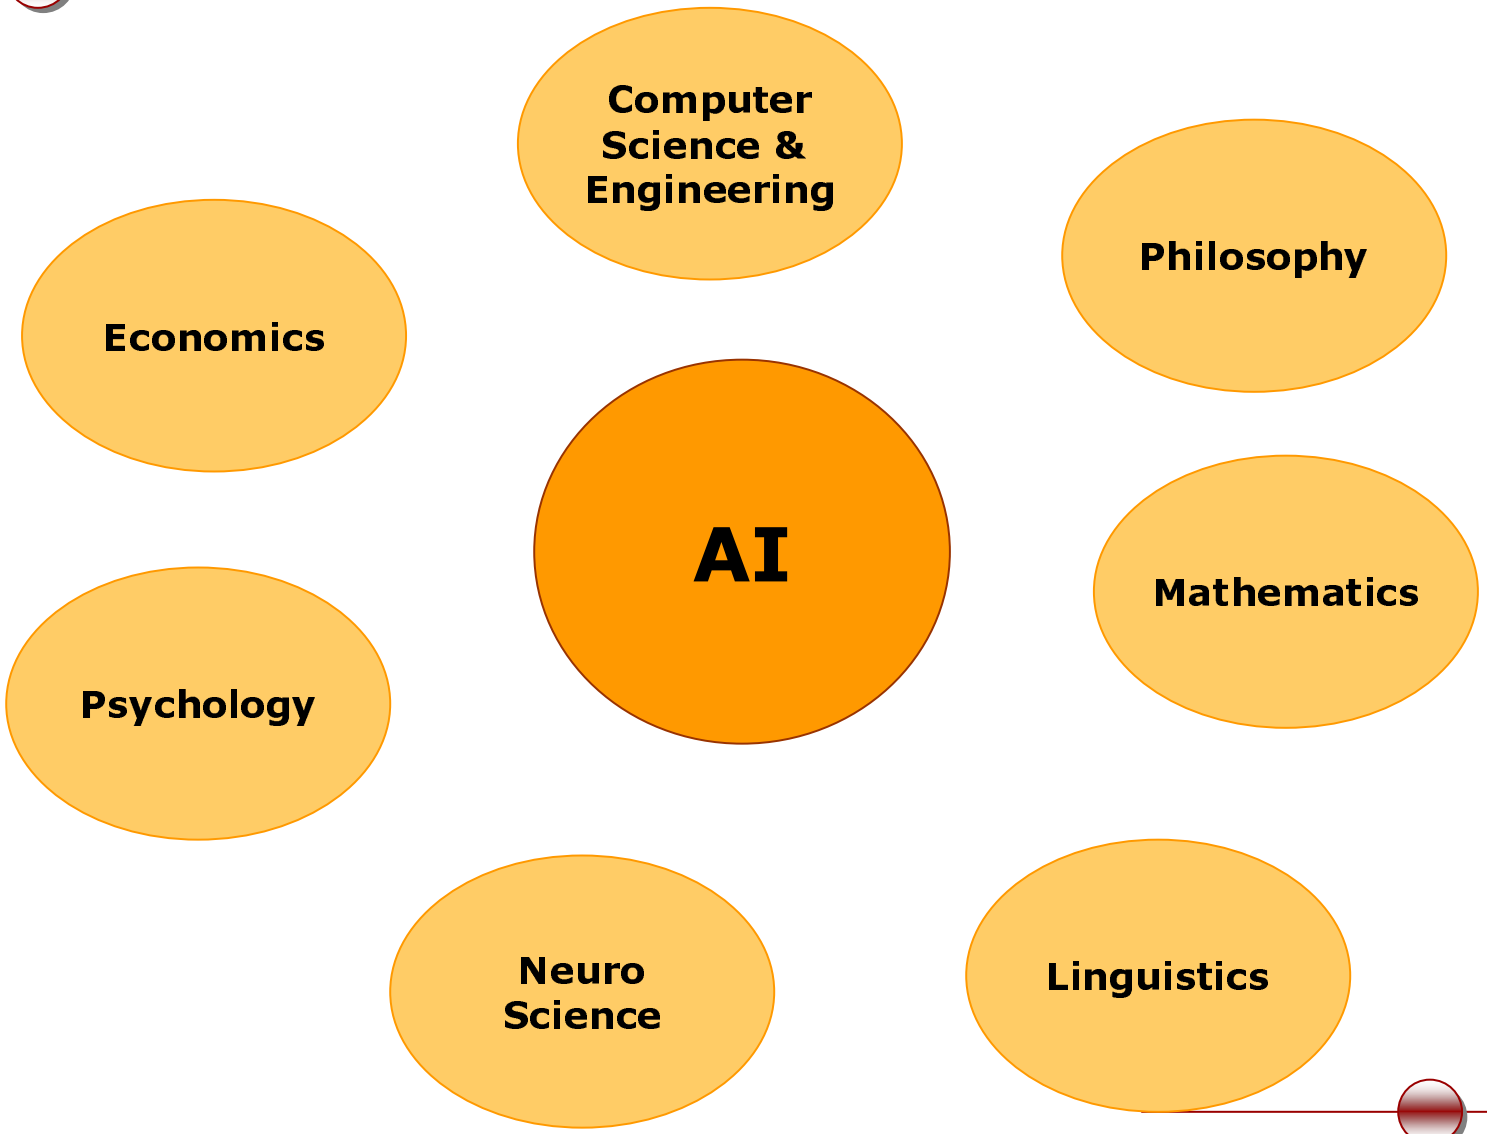

What's Easy and What's Hard?

- 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

What Can AI Systems Do?

- **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), perfect play in checkers, professional-level Go players

What Can't AI Systems Do Yet?

Exhibit true autonomy and intelligence!

- Understand natural language robustly (e.g., read and understand articles in a newspaper)
- Surf the web
- Interpret an arbitrary visual scene
- Learn a natural language
- 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

AI: A universal field

- Learning and perception
- Playing chess
- Proving mathematical theorems
- Writing poetry
- Driving a car on a crowded street
- Diagnosing diseases

AI: A universal field

- Scheduling train crews
- Automated student essay evaluation
- Packet scheduling in network routers
- Broadcast news understanding
- Vehicle diagnosis
- Robot photography

Who Does AI?

- Academic researchers
  - 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, Honeywell, Teknowledge, SAIC, MITRE, Fujitsu, Global InfoTek, BodyMedia, ...

Evolutionary Optimization

- MERL: evolving ‘bots

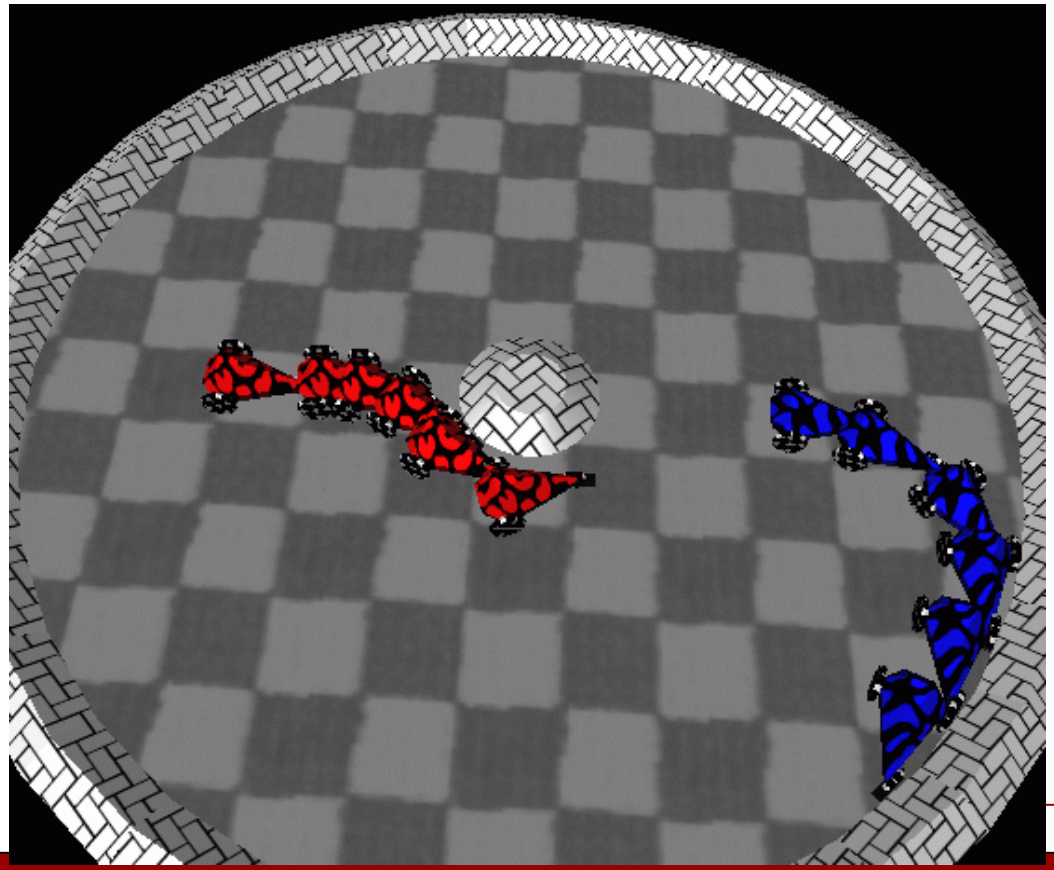

AI and Art: NEvAr

- See http://eden.dei.uc.pt/~machado/NEvAr

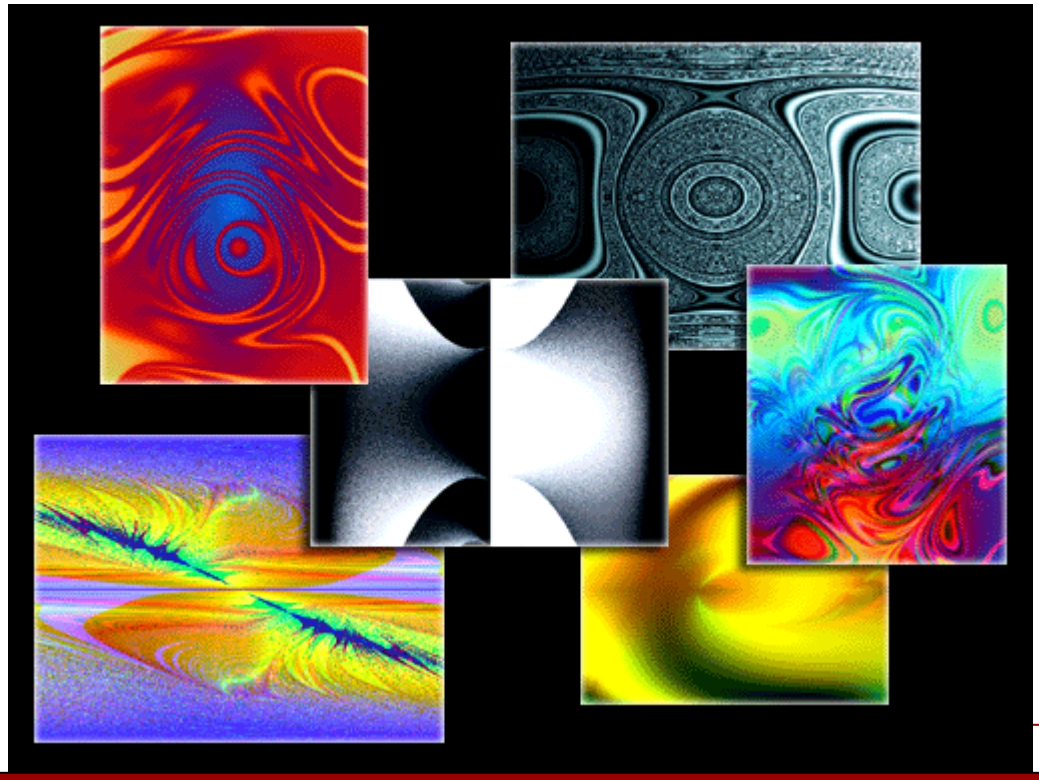

Protein Folding

- MERL: constraint-based approach

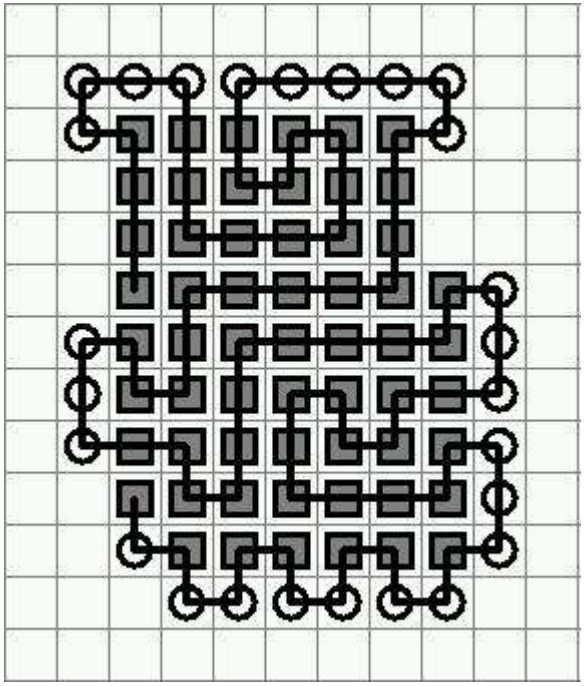

Interaction: Sketching

- MIT sketch tablet

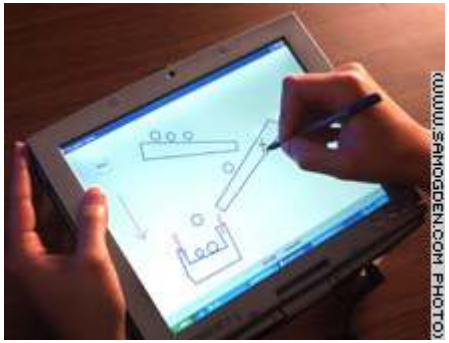

What Do AI People (and the Applications they Build) Do?

- **Represent** knowledge
- **Reason** about knowledge
- **Behave** intelligently in complex environments
- Develop interesting and useful **applications**
- **Interact** with people, agents, and the environment

Homework

- Pretest (due next week, Tuesday 9/7)

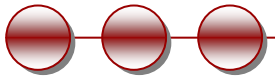

Course Overview

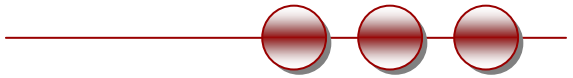

Expectations and preparation

- No prior AI experience is required
- Some knowledge of:
  - Propositional and basic first-order logic (there-exists, for-all)
  - Algorithmic analysis (big-O notation, NP-completeness)
  - Basic probability theory
- Pretest
  - Purpose: To help me assess students' knowledge

Course materials

- Course website: http://www.cs.umbc.edu/courses/graduate/671/fall10b/
  - Course description and policies (main page)
  - Course syllabus, schedule (subject to change!), and slides
  - Pointers to homeworks and papers (send me URLs for interesting / relevant websites, and I'll add them to the page!)

Course materials

- Course mailing list: cps-cmsc671@lists.umbc.edu
  - Visit http://lists.umbc.edu
  - Search for cps-cmsc671
  - Click “Subscribe” link
  - Send general questions to the list
  - Requests for extensions, inquiries about status, requests for appointments should go directly to Prof. Zavala

Preliminary grading distribution

- Homework Assignments 25%
- Midterm Exam 20%
- Final Exam 20%
- Final Project 25%
- Final Presentation 10%

Communication

- Office hours: Mondays and Wednesdays 1:30-2:30 (ITE 373)
- Appointments may also be made by request (at least 24 hours notice is best)
- Drop in whenever my door is open
- Will try to respond to e-mail within 24 hours on weekdays
- Direct general questions (i.e., those that other students may also be wondering about) to the class mailing list

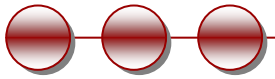

Thanks for coming -- see you next Thursday!