FOUNDATIONS OF ARTIFICIAL INTELLIGENCE

Textbook

- This is a great book
  - S. Russell and P. Norvig *Artificial Intelligence: A Modern Approach* Prentice Hall, 2003, Second Edition.
  - Most widely used in U.S. universities
  - It's so good….

- Artificial Intelligence: A New Synthesis

  Nils J. Nilsson. Morgan Kaufmann.

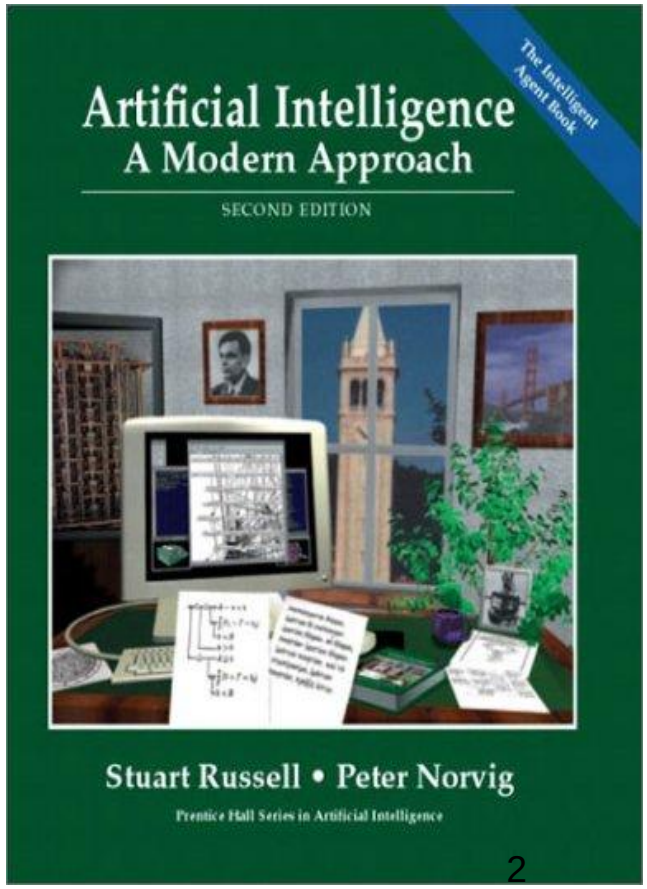

Outline

- Course overview
- What is AI?
- A brief history
- The state of the art

Course overview

- Introduction and Agents (chapters 1,2)
- Logic (chapters 7,8,9)
- Expert Systems (introduction)
- Uncertainty (chapters 13,14)
- Learning (chapters 18,20)
- Search (chapters 3,4,5,6)
- Games
- Planning (chapters 11,12)

Programming Language: Prolog, Scheme

Big questions

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

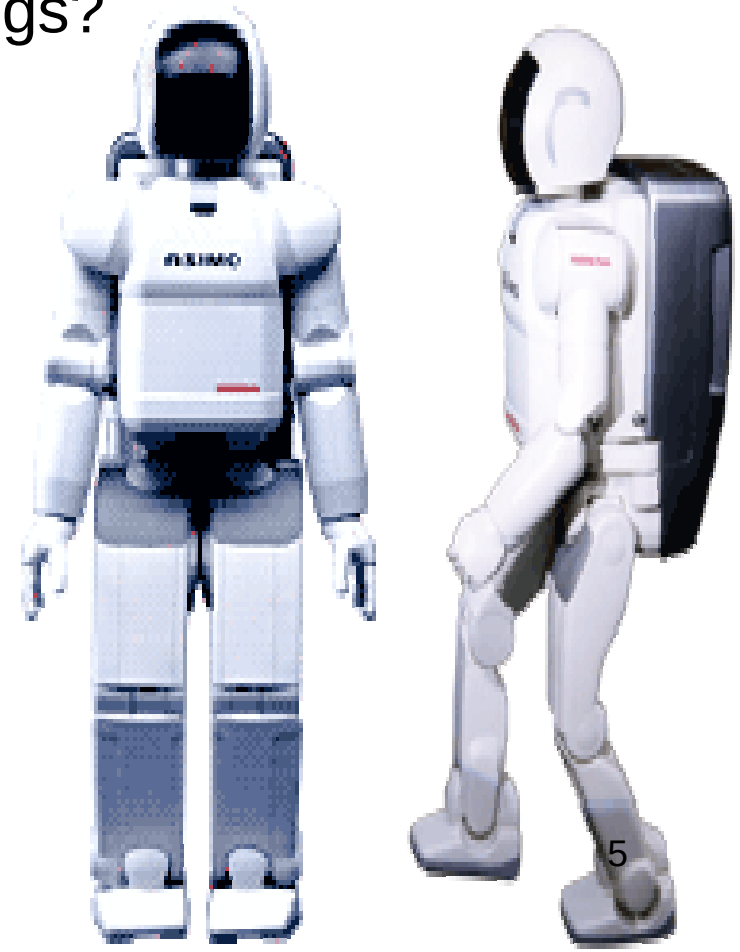

Why AI?

- **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.

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?

What is Artificial Intelligence?

- *"The exciting new effort to make computers think… machines with minds, in the full and literal sense."* *(Haugeland 1985)*

- *"The study of mental faculties through the use of computational models."* *(Charniak & McDermott 1985)*

- *"AI is the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the characteristics we associate with intelligence in human behaviour – understanding language, learning, reasoning, solving problems and so on."* *(Barr & Feigenbaum, Jackson chapter 2)*

Turing Test

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" **à** "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game

- Separate 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.

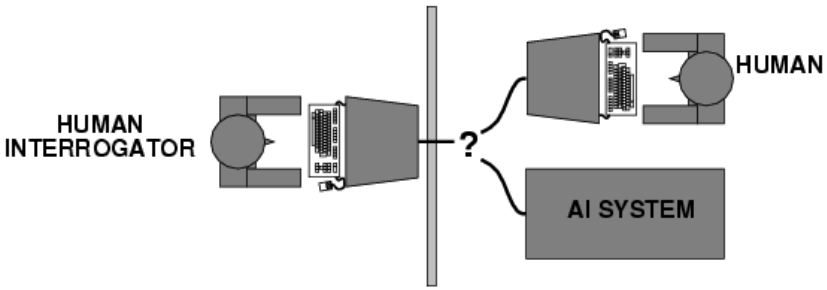

Turing Test

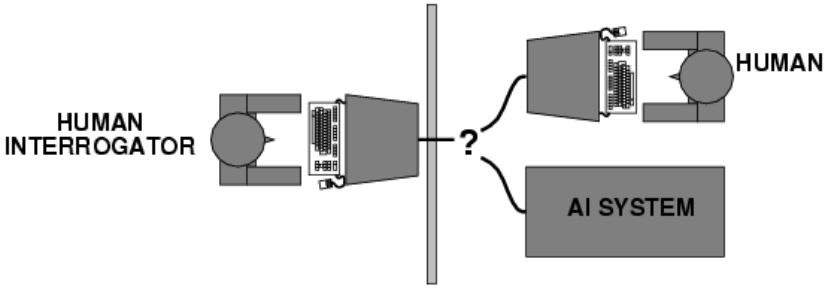

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning

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
- 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…)

Applications

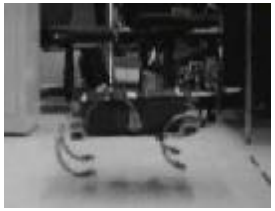

RHex Hexapod
Dr. Bassel ALKHATIB

Honda's Asimo

ATR's DB Android

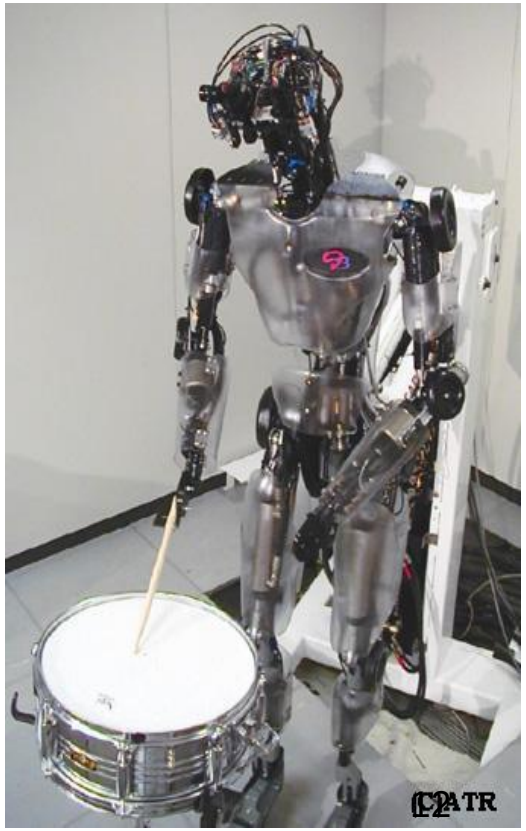

Chess: Deep Junior (IBM) tied Kasparov in 2003 match

2003 Kasparov v. Deep Junior ends in a 3-3 Draw

AI Techniques

- Rule-based
- Fuzzy Logic
- Neural Networks
- Genetic Algorithms
- Exhaustive search
- Expert Systems
- Logic

What is AI?

- The use of computers to solve problems that previously could only be solved by applying human intelligence…. thus something can fit this definition today, but, once we see how the program works and understand the problem, we will not think of it as AI anymore (David Parnas)

What is AI?

Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

The textbook advocates "acting rationally"

Thinking humanly: cognitive modeling

- 1960s "cognitive revolution": information-processing psychology
- Requires scientific theories of internal activities of the brain
- -- How to validate? Requires

  1) Predicting and testing behavior of human subjects (top-down)

  or 2) Direct identification from neurological data (bottom-up)
- Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI

Thinking rationally: "laws of thought"

- Aristotle: what are correct arguments/thought processes?
- Several Greek schools developed various forms of *logic*: *notation* and *rules of derivation* for thoughts; may or may not have proceeded to the idea of mechanization
- Direct line through mathematics and philosophy to modern AI
- Problems:
  1. Not all intelligent behavior is mediated by logical deliberation
  2. What is the purpose of thinking? What thoughts should I have?

Acting rationally: rational agent

- Rational behavior: doing the right thing
- The right thing: that which is expected to maximize goal achievement, given the available information
- Doesn't necessarily involve thinking – e.g., blinking reflex – but thinking should be in the service of rational action

Rational agents

- An agent is an entity that perceives and acts
- This course is about designing rational agents
- Abstractly, an agent is a function from percept histories to actions:

$$[f: P^* \rightarrow A]$$

- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: computational limitations make perfect rationality unachievable

  à design best program for given machine resources

Foundations of AI

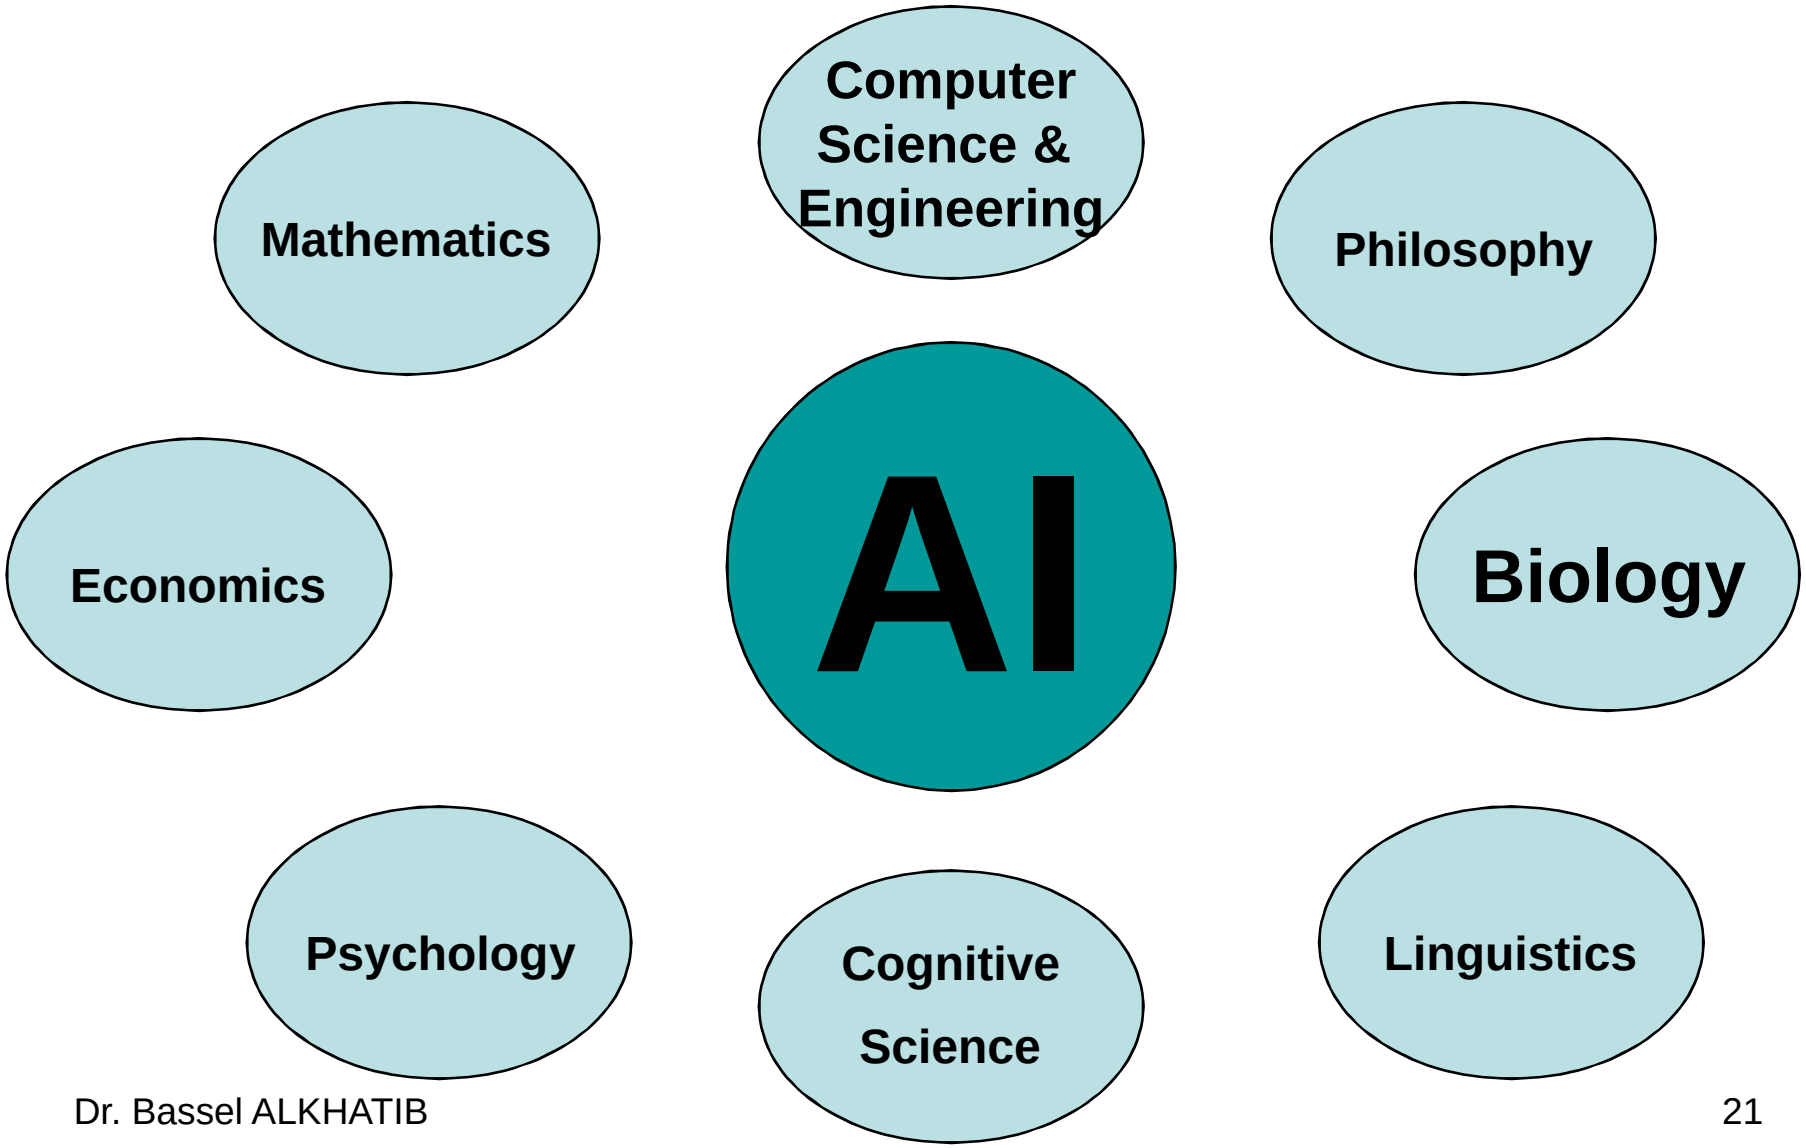

AI prehistory

• Philosophy	Logic, methods of reasoning, mind as physical system foundations of learning, language, rationality
• Mathematics	Formal representation and proof algorithms, computation, (un)decidability, (in)tractability, probability
• Economics	utility, decision theory
• Neuroscience	physical substrate for mental activity
• Psychology	phenomena of perception and motor control, experimental techniques
• Computer engineering	building fast computers
• Control theory	design systems that maximize an objective function over time
• Linguistics	knowledge representation, grammar

Abridged history of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1965 Robinson's complete algorithm for logical reasoning
- 1966—73 AI discovers computational complexity
  Neural network research almost disappears
- 1969—79 Early development of knowledge-based systems
- 1980-- AI becomes an industry
- 1986-- Neural networks return to popularity
- 1987-- AI becomes a science
- 1995-- The emergence of intelligent agents

History

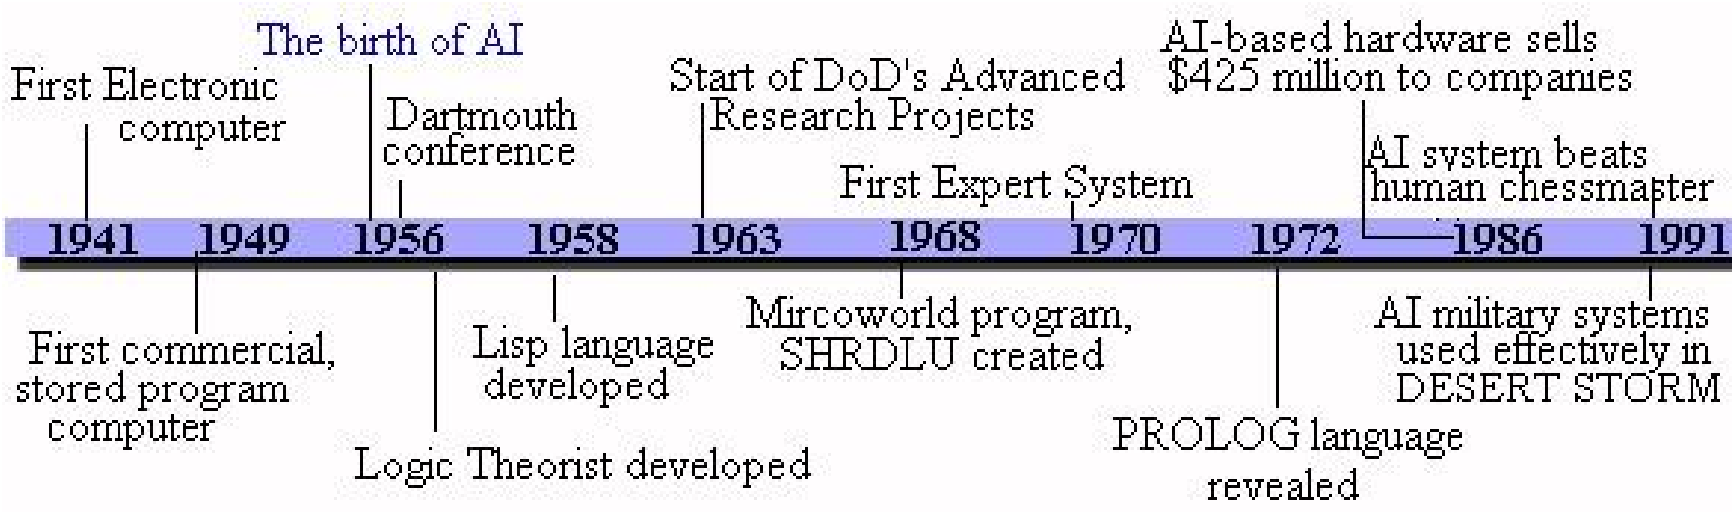
The birth of AI
First Electronic computer
Dartmouth conference
Start of DoD's Advanced Research Projects
First Expert System
AI-based hardware sells $425 million to companies
AI system beats human chessmaster
1941
1949
1956
1958
1963
1968
1970
1972
1986
1991
First commercial, stored program computer
Logic Theorist developed
Lisp language developed
Mircoworld program, SHRDLU created
PROLOG language revealed
AI military systems used effectively in DESERT STORM

State of the art

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- Proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- `Proverb` solves crossword puzzles better than most humans

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
- **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
- 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!

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, UMBC, Stanford, Berkeley, MIT, UIUC, UMd, U Alberta, UT Austin, ...
- 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, ...

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

Robotics

- SRI: Shakey / planning ..\movies\sri-Shakey.ram
- SRI: Flakey / planning & control ..\movies\sri-Flakey.ram
- UMass: Thing / learning & control
  ..\movies\umass_thing_irreg.mpeg
  ..\movies\umass_thing_quest.mpeg
  ..\movies\umass-can-roll.mpeg
- MIT: Cog / reactive behavior
  ..\movies\mit-cog-saw-30.mov
  ..\movies\mit-cog-drum-close-15.mov
- MIT: Kismet / affect & interaction
  ..\movies\mit-kismet.mov
  ..\movies\mit-kismet-expressions-dl.mov
- CMU: RoboCup Soccer / teamwork & coordination
  ..\movies\cmu_vs_gatech.mpeg

Applications

- AI and Data Integration
- AI and the Internet
- Art and Creativity
- Information Extraction

- A sample from IAAI-03:
  - Scheduling train crews
  - Automated student essay evaluation
  - Packet scheduling in network routers
  - Broadcast news understanding
  - Vehicle diagnosis
  - Robot photography
  - Relational pattern matching

AI & art: NEvAr

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

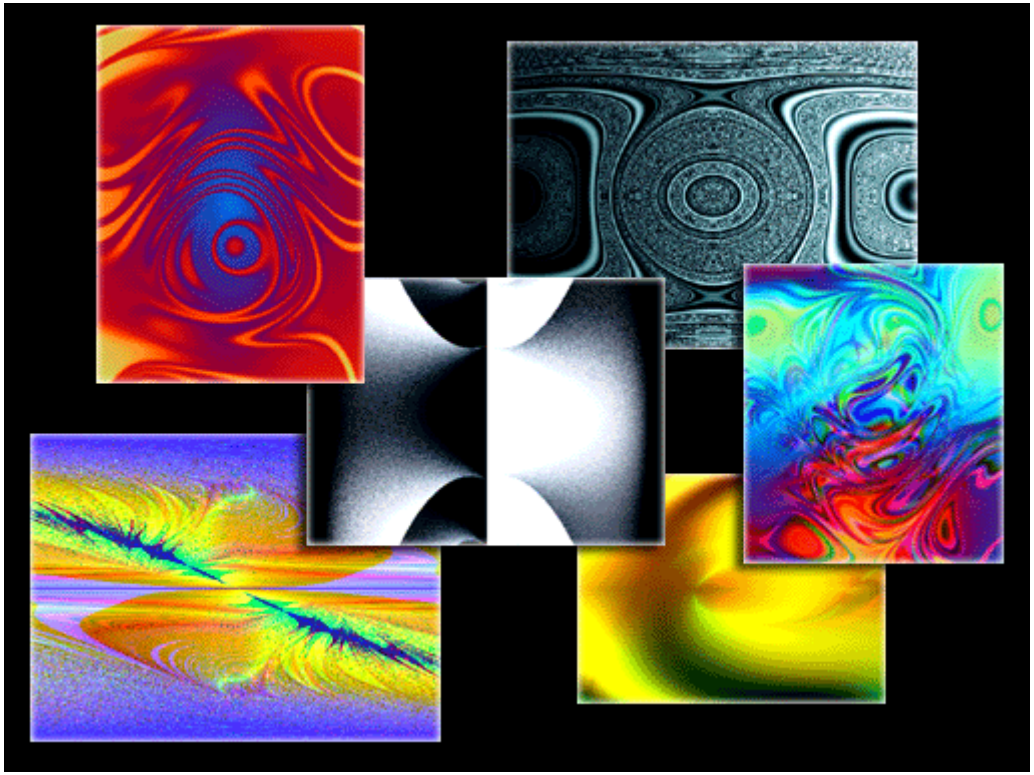

Interaction: Sketching

- MIT sketch tablet

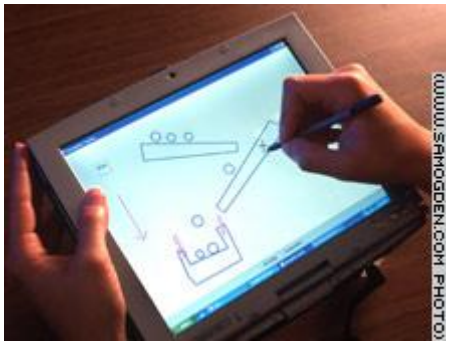

Other topics/paradigms

- Intelligent tutoring systems
- Agent architectures
- Mixed-initiative systems
- Embedded systems / mobile autonomous agents
- Machine translation
- Statistical natural language processing
- Object-oriented software engineering / software reuse