

Introduction to Artificial Intelligence



- AI's Successes
- What is AI?
- Rationality
- A bit of History



Exciting Time for AI


<http://aaai.org/AITopics/html/current.html>

- Expert Systems
 - Organic chemistry, Medicine, Geology, Configuring Computers, ...
 - User Help Systems [Microsoft]
- Web search engines
 - Google, Yahoo!, ...
- Natural Language Understanding
 - Eliza, Ask Jeeves, spell checkers, grammar checkers, ...
 - Real-time English-German speech translation!

AI Playing Games


■ CHESS:

- *Deep Blue* BEAT Kasparov (May '97)
- Matches expert level performance
 - "intelligent & creative" play
- . . . thinks differently from human expert by examining 200million board-positions/move




■ Checkers:

- *Chinook* is 1994 World champion CHECKERS Player!
- 2007: proven to be draw!



■ Othello

- Logistello beat Human Champion, 6-0 !



■ Backgammon

- TD-Gammon is competitive with World Champion



■ Poker

- Comparable with best humans! (2007)



Man vs Machine Poker

- Comparable with top human players (2007)
- Attracted international media attention.

“We won, not by a significant amount, and the bots are closing in.”
– Phil Laak



“I really am happy it's over. I'm surprised we won ... It's already so good it will be tough to beat in future.”
– Ali Eslami



The Sydney Morning Herald

smh.com.au

San Francisco Chronicle

The New York Times



CBS NEWS




The Washington Post

Playing Go




- First program to ever beat a human Go expert!
- Scientific American, June 2007


“UCT may prove useful for targeting advertisements on the Web, finding the best settings for an industrial plant or optimizing channel allocation in cellular systems.” – Scientific American



- UCT used to win 2007 General Game Playing (former UofA student)



Language "Understanding" + Problem Solving




- PROVERB: solves cross-word problems using
 - many information sources
 - (geography, puns, . . .)
 - word fillers
 - past puzzles
 - . . .
- In NewYorkTimes puzzles
 - 95.3% words correct
 - 98.1% letters correct
 - ⇒ high-quality cruciverbalist!
- <http://www.oneacross.com>

Robotics

- Ping pong player
- Honda robot
- Automated car #1: 1996
 - "ALVINN": vision trained to steer car
 - "No hands across America"
 - over 2,805 miles, in control for 98% !
 - Different terrain, road conditions, weather,
- DARPA Challenge Oct 2005:
 - 132 miles
 - Completely autonomous
 - Stanley



Theorem Proving (Nov'96)



- First "creative" proof by computer
- Solved 60 year open problem.
 - Robbins' problem in finite algebra
 - <http://www.mcs.anl.gov/home/mccune/ar/robbins>
- Qualitative difference from previous brute-force results
- . . . unlike [Appel/Haken] computer proof of "4 color" theorem



Decision theory and Statistical User-models

- Lymph node pathology
Explain response. . .
- Microsoft Office '97 + Answer Wizard
Diagnosis/reasoning using Bayesian
models
<http://www.research.microsoft.com/research/dtg/>
- Restricted natural language parsing




Logistics Planning

- Gulf War, 1991
 - 50,000 people, cargo, vehicles
 - various start and designation points
- "DART" (AI System)
 - produced plan in hours;
 - by hand: many weeks
- DARPA:
"this single application paid back 30 years of funding" !



Autonomous Intelligent Systems (NASA)



- Next generation spacecrafts
"DS1" (launched mid98)
- Control space-craft and carry out objectives, while maintaining internal systems, ... over decades.
- Fast real-time, on-line performance
- Automatic planning and execution model
- General deduction from first principles
synthesize timely responses to anomalous & unexpected situations

<http://ic-www.arc.nasa.gov/ic/projects/mba/>


Machine Learning Successes: Computer learns...




AMERICAN
EXPRESS



- to find ideal customers
 - Credit Card approval (AMEX)
 - Humans $\approx 50\%$; ML is $>70\%$!
- to find best person for job
 - Telephone Technician Dispatch [Danyluk/Provost/Carr 02]
 - BellAtlantic used ML to learn rules to decide which technician to dispatch
 - Saved \$10+ million/year
- to predict purchasing patterns
 - Victoria Secret (stocking)
- to help win games
 - NBA (scouting)
- to catalogue celestial objects [Fayyad et al. 93]
 - Discovered 22 new quasars
 - $>92\%$ accurate, over tetrabytes




ML#2: Sequential Analysis




- **BioInformatics 1: identifying genes**

- Glimmer [Delcher et al, 95]
- identifies 97+% of genes, automatically!



- **BioInformatics 2: Predicting protein function, ...**



- **Recognizing Handwriting**

Handwritten text examples:
 Saw bushes 1-1-0
 Bought a knife 0-0-0
 for Skating, yim 0-1-0
 for from black 0-0-0
 at the play 0-9-0

US postal service zip code reader

Handwritten zip code examples:
 1-2-02-0
 3-2-02-0
 5-0-0-0
 3-1-0
 1-1-1-1
 1-1-2-0

- **Recognizing Spoken Words**



- "How to wreck a nice beach"

Ticketmaster, Speechworks, Bell, Verbmobile [translation]



ML#3: Control

- **TD-Gammon** (Tesauro 1993; 1995)
 - World-champion level play by **learning** ...
 - by playing millions of games against itself!
- **Drive autonomous vehicles** (Thrun 2005)
 - DARPA Grand Challenge
- **Printing Press Control** (Evans/Fisher 1992)
 - Control rotogravure printer, prevent groves, ... specific to each plant
 - More complete than human experts
 - Used for 10+ years, reduced problems from 538/year to 26/year!
- **Oil refinery**
 - Separate oil from gas
 - ... in 10 minutes (human experts require 1+ days)
- **Manufacture nuclear fuel pellets** (Leech, 86)
 - Saves Westinghouse >\$10M / year
- **Adaptive** agents / user-interfaces



Industry Interest in ML

"If you invent a breakthrough in artificial intelligence, so machines can learn, that is worth 10 Microsofts."
– Bill Gates, New York Times, 1/03/2004

Microsoft[®]

Google[™]

YAHOO![®]

IBM

intel[®]

hp[®]
i n v e n t

Sun[®]
microsystems

WAL★MART[®]

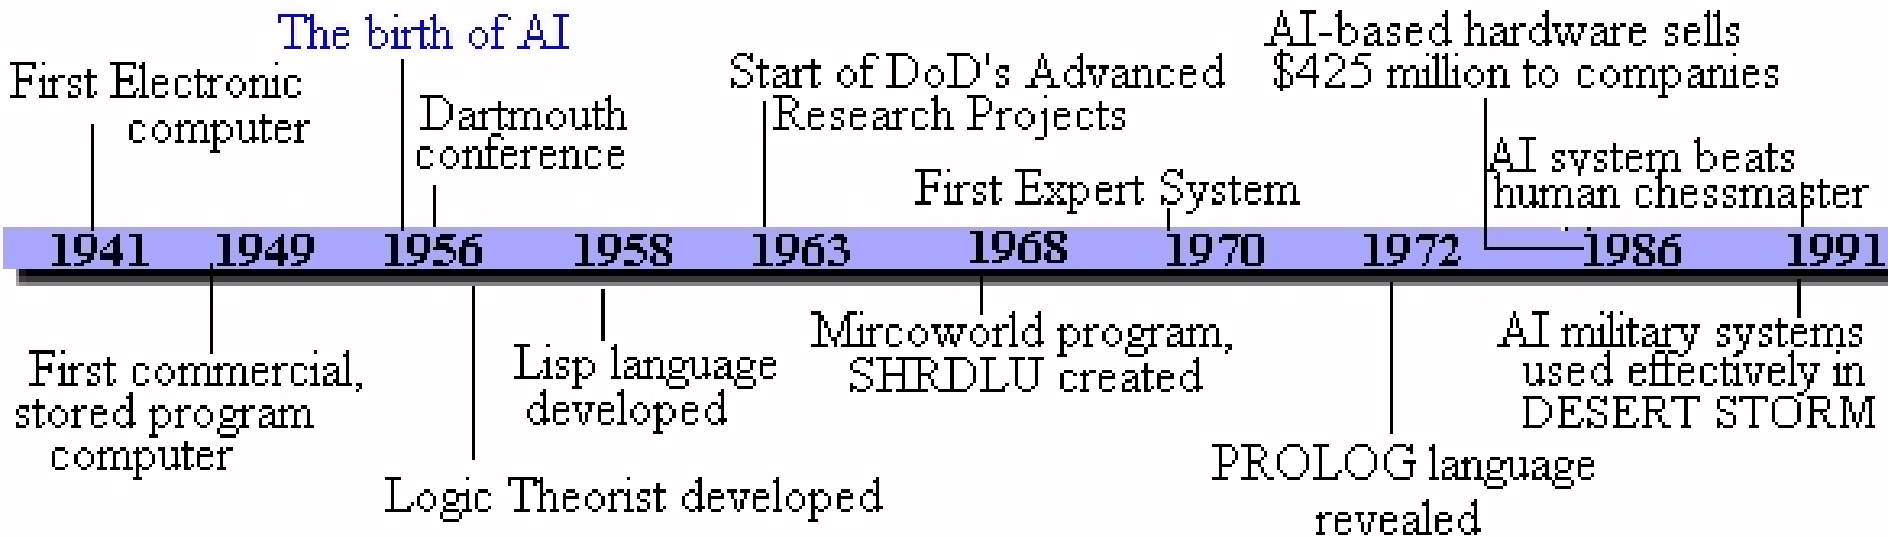
SIEMENS

AMERICAN
EXPRESS

amazon.com[®]

 **BOEING**

AI History



- Foundations...
- Technology push
- Artificial domains
- Real applications



Trends Since 90's

- Relying less on logic;
more on *probability theory and statistics*
- More emphasis on *objective performance evaluation*
- Intelligent (rational) Agents
- Accomplishments in
 - Game playing: Deep Blue, Chinook, Logistello, ...
 - Space Probe
 - Biological sequence analysis
 - OCR
 - Consumer electronics



Questions ... for audience!

- What is Artificial Intelligence?
 - ? artificial HUMAN intelligence ?
 - ? intelligent ARTIFACTS ?
 - ? applied EPISTEMOLOGY ?
 - Science? Math? Psychology?
Philosophy? . . .
- What does AI involve?
 - ... not involve?
 - ? Reasoning/ProblemSolving/"Thinking" ?
 - ? Game playing ?
 - ? (Intelligent) Agents ?
 - ? Image Understanding ?
 - ? Robotics ?
 - ? Natural Language / Communication ?
 - ? Learning ?
- Can it succeed?
- Why should we care if it does?
 - ? Previous successes ?
 - ? Anticipated successes ?
- What are open problems?



What is Artificial Intelligence?

AI studies systems (eg, computer programs)
that ...

1. Act like humans	4. Act rationally
2. Think like humans	3. Think rationally





1. Act Like Humans

1. Act like humans	4. Act rationally
2. Think like humans	3. Think rationally

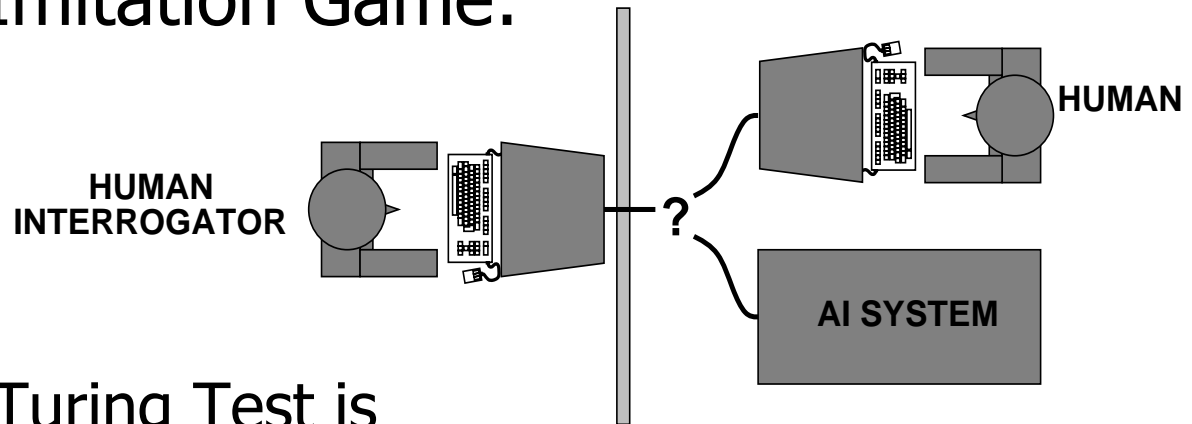
- AI = ...
creating machines that perform
“intelligent functions”
... only by intelligent humans

- Methodology:
Find an intellectual task
and
make a computer do it

- Prove a theorem
- Play chess
- Plan a surgical operation
- Diagnose a disease
- Navigate in a building

Turing Test

- *NOT:* "Can machines think"
- *Instead:* "Can machines behave intelligently?"
- Operational test of intelligent (human like) behavior: Imitation Game:



- Problems: Turing Test is
 - not reproducible,
 - not constructive,
 - not amenable to mathematical analysis



2. Think Like Humans

1. Act like humans	4. Act rationally
2. Think like humans	3. Think rationally

- 1960s “cognitive revolution”:
information-processing psychology replaced
behaviorism
- Requires scientific theory of internal activities of the brain
 - What level of abstraction?
“Knowledge” or “circuits”?
 - How to validate? Requires
 1. Predicting and testing behavior of human subjects (top-down) vs
 2. Direct identification from neurological data (bottom-up)
- Both approaches
(≈ Cognitive Science and Cognitive Neuroscience)
are now distinct from “core AI”



Examples

- Garden-Path Sentence:
 - The horse raced past the barn fell.
- Center-embedding:
 - The cat that the dog that the mouse that the elephant admired bit chased died.
 - The elephant admired the mouse that bit the dog that chased the cat that died.
- Multiplication: $3242342 * 2342342$

Do we want to duplicate human imperfections?



Artificial Human Intelligence??

- How to build good FLYING machine
 - [?] Synthesize birds?
 - [?] Understand aerodynamic principles
- Best FLYING Machine = Airplane
≠ Improved Bird
- How to build good THINKING machine
 - [?] Synthesize human brains?
 - [?] Understand reasoning principles
- Best THINKING Machine ?=? Improved Human Brain
[. . . DeepBlue . . .]

3. Think Rationally:

Laws of Thought

1. Act like humans	4. Act rationally
2. Think like humans	3. Think rationally

- Normative (or prescriptive), not descriptive
- Aristotle:
 - What are correct arguments/thought processes?
 - Several Greek schools developed forms of *logic*: notation and rules of derivation for thoughts.
- Problems:
 - Not all intelligent behavior is mediated by logical deliberation
 - What is the purpose of thinking?
What thoughts *should* I have?

4. Act Rationally

1. Act like humans	4. Act rationally
2. Think like humans	3. Think rationally

- Rational behavior: *doing the right thing*
- “The right thing” \equiv
action that is expected to
maximize goal achievement,
given the available information
- Doesn't necessarily involve thinking
(eg, blinking reflex)...
but
thinking should be in service of
rational action



Rational Agents

- AI \equiv Study of design of **rational agents**
- **agent** = thing that acts in environment
- **Rational agent** = agent that acts rationally:
 - actions are appropriate for goals and circumstances
 - makes appropriate choices given
 - perceptual limitations and finite computation
 - (flexible to changing environments and goals; learns from experience)
- Details later...

Goals of Artificial Intelligence

- Scientific goal:
 - understand principles that make rational (intelligent) behavior possible, in natural or artificial systems.
- Engineering goal:
 - specify methods for design of useful, intelligent artifacts.
- ~~■ Psychological goal:
 - understanding/modeling people
 - cognitive science (not this course)~~





Goals of This Course

- Introduce key methods & techniques from AI
 - searching
 - reasoning and decision making
 - logical
 - decision theoretic (probabilistic)
 - ? learning (covered in detail in CMPUT466)
 - ? language processing/understanding,
 - ? vision/robotics
 - . . .
- Understand applicability and limitations of these methods



Approach of Course

- Our approach:
 - Characterize Environments
 - Identify rational agent that is most effective for each environment
- Study increasingly complicated agent architectures requiring
 - increasingly sophisticated representations,
 - increasingly powerful reasoning strategies

To help set expectations...

State of the Art:

- Play a decent game of table tennis
- Drive safely along a curving mountain road
- 😊 ■ Drive safely along White Avenue
- Buy a week's worth of groceries on the web
- 😊 ■ Buy a week's worth of groceries at Safeway
- Play a decent game of bridge
- 😊 ■ Discover and prove a new mathematical theorem
- 😊 ■ Design and execute a research program in molecular biology
- 😊 ■ Write an intentionally funny story
- Give competent legal advice in a specialized area of law
- 😊 ■ Translate spoken English into spoken German in real time
- 😊 ■ Converse successfully with another person for an hour
- 😊 ■ Perform a complex surgical operation
- 😊 ■ Unload any dishwasher and put everything away



Summary

- AI Successes
- What is AI?
- Rationality
- A bit of History
- Expectations...