CS W4701
Artificial Intelligence

Fall 2013

Chapter 1:

Introduction

Jonathan Voris

(based on slides by Sal Stolfo)

4701

- Course home page:

http://www.cs.columbia.edu/~jvoris/AI

- Textbook: S. Russell and P. Norvig *Artificial Intelligence: A Modern Approach* Prentice Hall, 2003, Third Edition

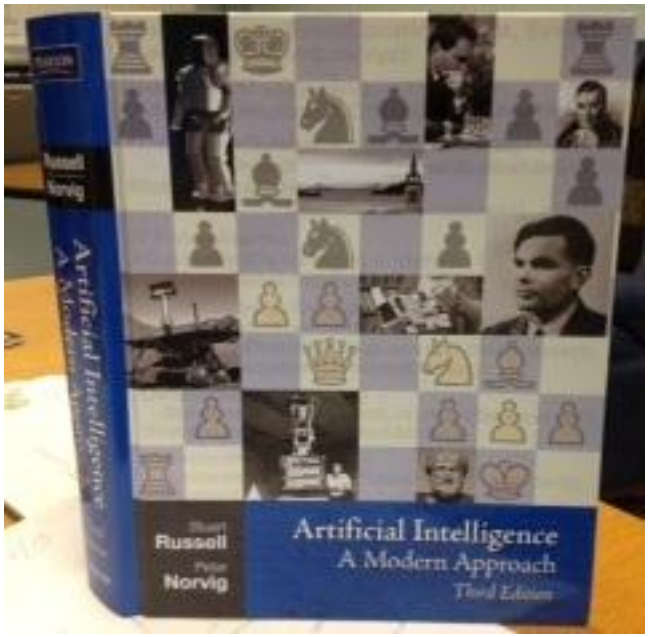

Who are we?

- Instructor: Jonathan Voris
  - Office: CEPSR 612
  - Email: jvoris_at_cs_dot_columbia_dot_edu

Who Are We?

- Teaching Assistants:
  - Tingting Ai
    - Email: ta2355 at columbia dot edu
    - Office Hours Mon 4-6, Wed 10-Noon
  - John Sizemore
    - Email: jcs2213 at columbia dot edu
    - Office Hours: Tue TBD
  - Keqiu Hu
    - Email: kh2567 at columbia dot edu
    - Office Hours: Thu 4-6

Who Are We?

- Teaching Assistants:
  - Ya Cai
    - Email: yc2788 at columbia dot edu
    - Office Hours: Tue 4:00-6:00
  - Fang-Hsiang Su
    - Email: mikefhsu at cs dot columbia dot edu
    - Office: CEPSR 608
    - Office Hours: Thu 5:00-7:00

More about Jon

- Graduated from the Ph.D. Program at Polytechnic Institute of New York University
- Currently a Postdoctoral Research Scientist in the Columbia Intrusion Detection Systems Lab
- Research in computer security, privacy, and usability
- Web site: https://www.cs.columbia.edu/~jvoris/

Outline

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

Tentative Grading Scheme

- 4 Projects/1 Midterm/1 Final

Assignment/Test	Percentage
Project #1	15%
Project #2	15%
Project #3	15%
Project #4	15%
Midterm	15%
Final	20%
Class Participation	5%

Grading Policy

- Submit:
  - Source code
  - Documentation
  - Test run input and output
- General criteria:
  - Correctness: 75%
  - Design/Structure: 15%
  - Documentation: 10%

Late Policy

- 10 points deducted immediately
- 10 points per day thereafter
  - 10/24 points deducted per hour
- No submissions accepted after the next project due date

Academic Honesty

- Read and understand the department's policy here:
  - http://www.cs.columbia.edu/education/honesty

Course Overview

- Introduction and Agents (chapters 1,2)
- Search (chapters 3,4,5,6)
- Mathematical Logic (chapters 7,8,9)
- Uncertainty (chapters 13)
- Learning (chapters 18,20)

- 4 Projects/1 Midterm/1 Final

What is Artificial Intelligence?

- Loaded question…let’s back up a little here

What is Artificial Intelligence?

- Loaded question…let’s back up a little here
- What is a computer?

What is Artificial Intelligence?

- Loaded question…let's back up a little here
- What is a computer?
  - Hrmm…
  - Let's start with these easier questions….

What is a Hammer?

What is a Hammer?

A hammer is an AMPLIFIER for….

What is a Phone?

What is a Phone?

A phone is an AMPLIFIER for….

What is a Car?

What is a Car?

A car is an AMPLIFIER for….

What is a Computer?

What is a Computer?

A computer is an AMPLIFIER for….

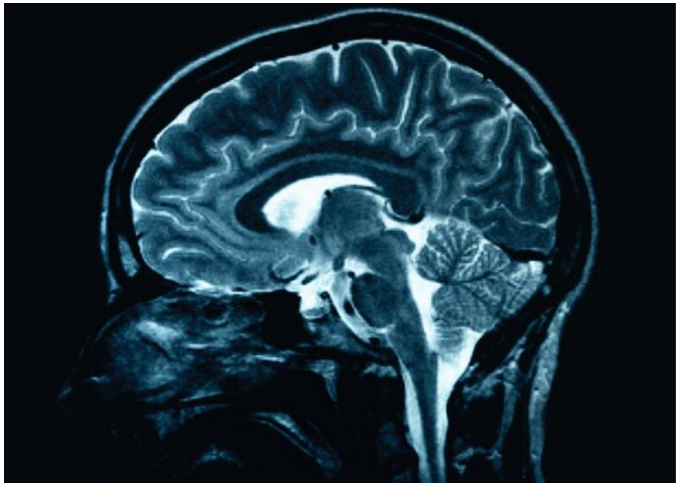

The Brain!

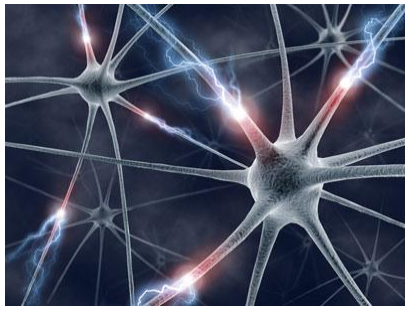

- 50-100B of these:
- 10,000’s connections each!
- ~10B critical pyramidal cells involved with cognition
- 1000 trillion (1 quadrillion) connections!

- Why is it so wrinkled?

- Frontal lobes oversized by mammalian standards
- Vision processing oversized
- Three times larger than next average mammal

The Brain!

- What does it do?
- Remembers stuff:

- Semantic memory: General knowledge, trivia and facts are stored in the temporal lobe and the cortex.
- Episodic memory: New data and recent events are stored in the prefrontal cortex and the temporal lobe.
- Working memory: Information and knowledge required for daily life – such as telephone numbers and learned skills like driving -- are stored in the prefrontal cortex.
- Procedural memory: Secondhand skills, things we take for granted, such as walking and cycling, are stored in the cerebellum.

The Mind!

- Is it physical?
- Where is it?
- Is it real or imagined? ☺
- Mind-Body Problem
  - Descartes, Plato, Aristotle, Asian philosophy
  - Dualism: separate from each other
  - Monism: rationalists, two aspects of an underlying reality
- Let’s just deal with our own reality for now….

What is AI?

Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

The textbook advocates "acting rationally"

What is AI?

- Views of AI fall into four categories:

Thinking Humanly	Thinking Rationally
Acting Humanly	Acting Rationally

Acting humanly: Turing Test

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

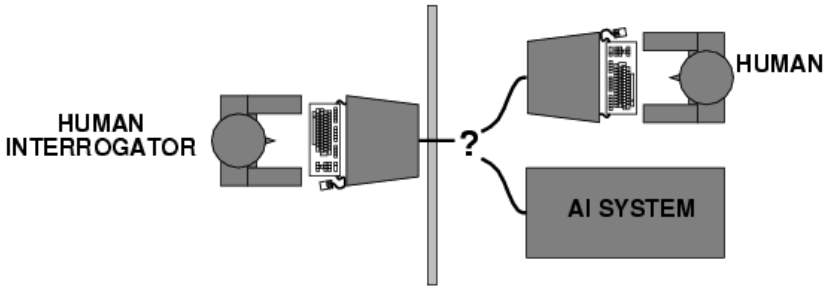

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning

Turing Test: Relevance?

- Strengths of this approach?
- Weaknesses of this approach?

Turing Test: Relevance?

- Strengths of this approach?
- Weaknesses of this approach?
- Practical applications?

Turing Test: Relevance?

- Strengths of this approach?
- Weaknesses of this approach?
- Practical applications?

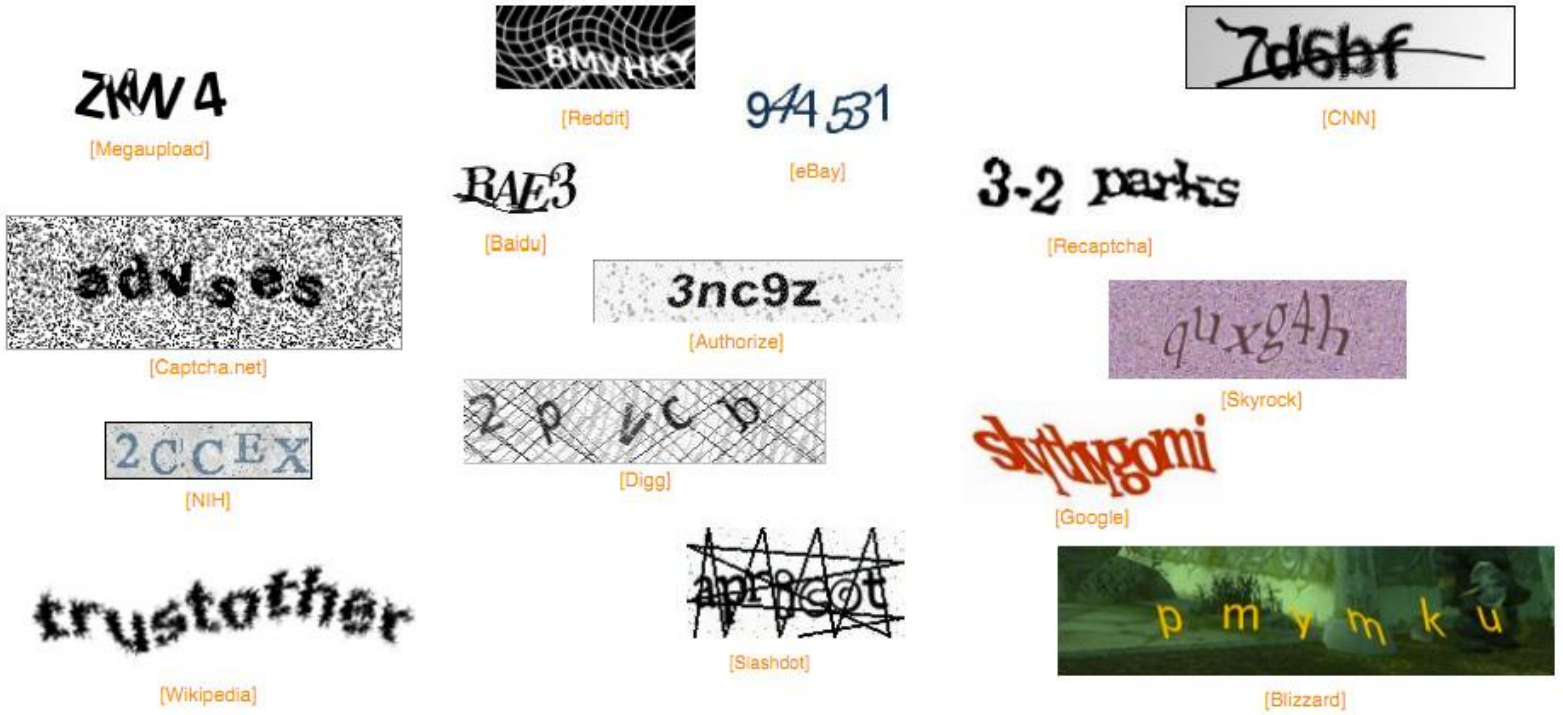

What is AI?

- Views of AI fall into four categories:

Thinking Humanly	Thinking Rationally
Acting Humanly	Acting Rationally

Thinking humanly: cognitive modeling

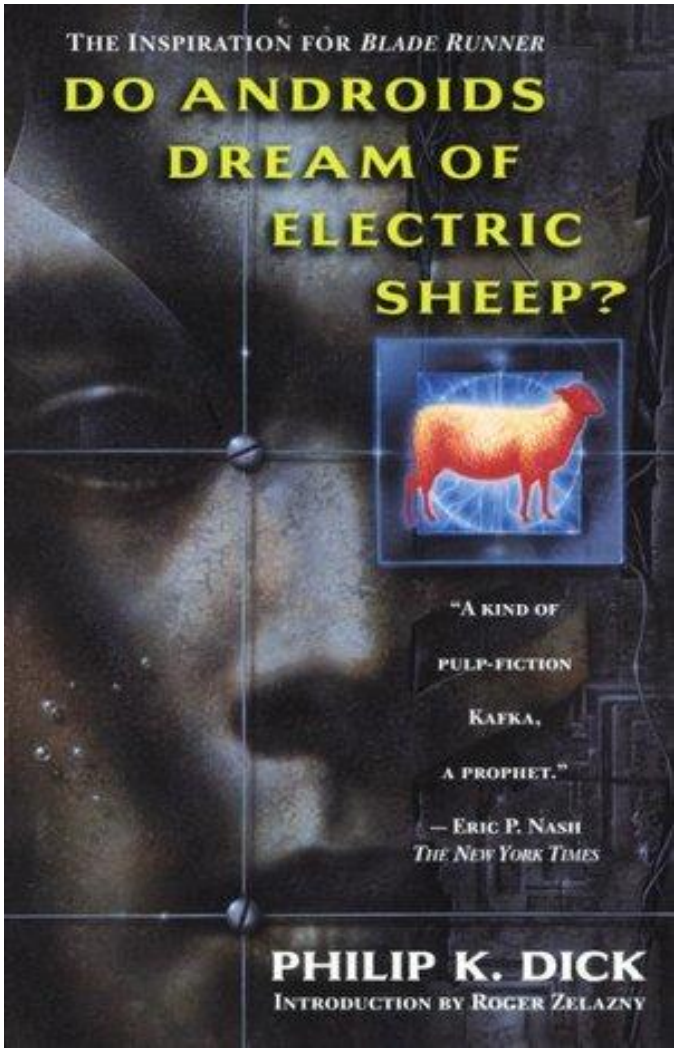

Thinking humanly: cognitive modeling

Thinking humanly: cognitive modeling

- Well, how do humans think anyway?
  - Introspection
  - Psychological / physiological experiments
- Develop *theory* of human thought
- Implement *theory* as a *program*
  - Its as simple as that!
- Check similarities to human behavior

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 humanly: cognitive modeling

- Relationship between “good” algorithms and “human” algorithms
- Separation has formed two distinct fields:
  - Cognitive Science
  - Cognitive Neuroscience

What is AI?

- Views of AI fall into four categories:

Thinking Humanly	**Thinking Rationally**
Acting Humanly	Acting Rationally

Thinking rationally: "laws of thought"

- Aristotle: what are correct arguments/thought processes?
  - Sought *irrefutable* reasoning
  - Correct premises -> correct conclusions
- 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

Thinking rationally: "laws of thought"

- Logicism: Direct line through mathematics and philosophy to modern AI
- 19th Century: Precise notations for logical relations
  - Arithmetic for things
- 1965: Problem solving programs
- Problems:
  1. Expressing things in formal notation
  2. Uncertainty
  3. Theory vs practice
     1. Efficiency
     2. Ordering

What is AI?

- Views of AI fall into four categories:

Thinking Humanly	Thinking Rationally
Acting Humanly	**Acting Rationally**

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:
- Logicism

$$[f: \mathcal{P}^{\star} \rightarrow \mathcal{A}]$$

- 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

  → design best program for given machine resources

The Secret Origin of AI

- AI has a long academic pedigree
- Was born out of ideas from many fields

The Secret Origin of AI

- Philosophy: pre-computer AI
  - Logic
  - Methods of reasoning
  - Mind as physical system
  - Foundations of learning
  - Language
  - Rationality
- Ramon Lull 1315: Can a device reason?
- Already talked about Aristotle
- Da Vinci: plans for an unbuilt mechanical calculator
- Early 1600s: Schickard, Pascal, Leibinz actually construct them
- Mind-body problem

The Secret Origin of AI

- Logical positivism
  - All knowledge formed by *logical theories* which connect *observation*
  - Sound familiar?
  - Read works of:
    - Hume
    - Wittgenstein
    - Russell

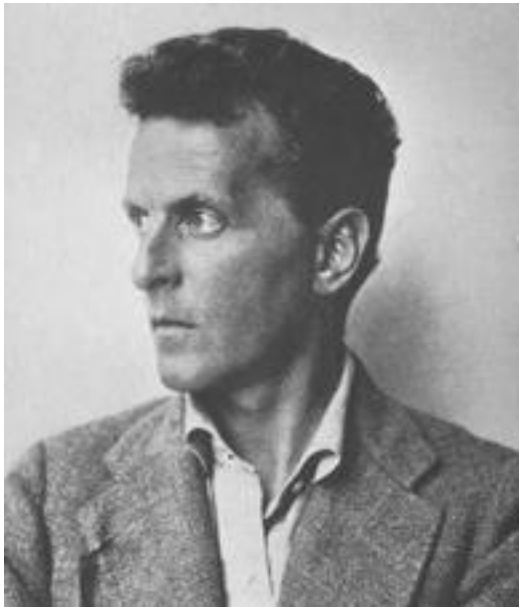

The Secret Origin of AI

- Mathematics: transition of AI into a formal science
  - Formal representation and proof algorithms
  - Computability, (un)decidability, (in)tractability
  - Probability / uncertainty
- Math and logic: 1800s: George Boole
- Algorithms: Euclid & GCD
- 1931: Goedel's incompleteness theory
  - Some things are true but unprovable
- 1900s: Turing is the man
  - Some functions cannot be programmed
- 1700s: Bayes – Calculating probabilities from observations

The Secret Origin of AI

- Economics: practical aspects of AI
- 1776: Smith – Wealth of Nations – turned economics into a science
  - "Invisible Hand"
- **Utility**: Association of value to event outcomes
- Decision theory: probability + utility
  - Framework for analyzing decisions made with imperfect knowledge
- 1944: von Neumann and Morgenstern: Game Theory
  - Ways to model interactions between rational agents

The Secret Origin of AI

- A few other fields:
- Neuroscience
  - How does the brain work?
  - Physical substrate for mental activity
- Psychology
  - How do humans think, and why?
  - Phenomena of perception and motor control
  - Experimental techniques
- Computer engineering
  - Improving hardware
- Control theory
  - Design systems that maximize an objective function over time
- Linguistics
  - How does language affect thought (or vice versa)?
  - Knowledge representation
  - Grammar

A Brief History of AI

- 1943-1955: Getting warmer…
- 1943: McCulloch & Pitts neuron model blends
  - Neurology
  - Propositional logic
  - Theory of computation
- 1950: First “neural network”
  - Computer simulating 40 neurons
- Turing in the house
  - "Computing Machinery and Intelligence"
  - “Child program” idea

A Brief History of AI

- 1956: The birth of AI
- Dartmouth College
- Two month workshop
- 10 attendees
- Newell and Simon's "Logic Theorist" Program
  - Creates a shorter proof for one of Russell's theorems
  - Original paper rejected!
- Firmly planted AI in CS
- Introduced all the field's main players

A Very Brief History of AI

- 1952-1969: This new AI thing is great!
- Low expectations
- Early “successes” showing computers >> calculators
- Physical Symbol System hypothesis: intelligent systems must operate on structures of symbols
- IBM: Early AI programs
  - Samuel’s checkers program outwits its creator
  - Gelernter's Geometry Engine
- 1958: Lisp is born at MIT! (much more about this soon)
- Study of domain limited *microworlds*
  - Calculus integration
  - Algebra problems
  - “Blocks world”

A Very Brief History of AI

- 1966-1973: Maybe not all its cracked up to be…
- Herbert Simon: computer will beat chess champion by late 60s
  - It actually took until the late 90s – whoops…
- Research runs into problems of *scalability* and *complexity*
- Solutions to simple problems don't apply well to more challenging ones
- Syntax != knowledge
- Example: translation
  - "The spirit is willing but the flesh is weak"
  - "The vodka is good but the meat is rotten"
- Difference between theory and practice
- No progress in genetic algorithms
  - Data represented poorly
- Governments start to cut funding
- Neural network research almost disappears

A Very Brief History of AI

- 1969-1979: Knowledge based systems
- Instead of generic problem solvers
  - Attack narrower problems
  - Leverage domain specific information
  - Knowledge from lots of special rules
- 1969: DENDRAL program for inferring molecular structure
- Expert systems: attempts to copy human decision making
  - MYCIN: medical diagnosis
- Return of translation efforts
  - Natural Language Processing
  - Less emphasis on language and more on knowledge representation

A Very Brief History of AI

- 1980 – Now: AI sells out
- Companies realize this AI thing can save them money
- R1: Expert system for computer orders
  - A computer in charge of ordering computers!
- DuPont develops and uses hundreds of expert systems
- Nations launch competing AI efforts
  - 1981: Japan's "Fifth Generation" project
  - US: Microelectronics and Computer Technology Corporation (MCC)
  - UK restores funding
  - None of these efforts were fruitful by the way…
- 1988: AI is a billion dollar industry
  - Yet companies still have trouble delivering on claims

A Very Brief History of AI

- 1986 – Now: AI grows up
- The return of neural networks – split into
  - Modeling organic neurons
  - Creating and analyzing effective networks
- Embraces scientific methods
  - Formal hypotheses
  - More rigorous experiments
  - Statistical analysis of results
  - Replication!
- Developments in use today:
  - Data mining
  - Bayesian networks

“My CPU is a neural net processor; a learning computer. But Skynet pre-sets the switch to read-only when we're sent out alone.”

A Very Brief History of AI

- 1995 – Now: Intelligent agents
- New debate over whether AI should
  - Continue to focus on concrete domain specific tasks, or
  - “Human level AI”
- The Internet comes in to play
  - Search algorithms
  - Recommendation systems
- The rise of “bots”
  - Where is Turing when you need him!?

A Very Brief History of AI

- 2001 – Now: Rise of big data
- Advent of the internet provides plenty of data to practice on
- Maybe the problem isn’t bad algorithms, but insufficient data?
  - Certain algorithms perform much better with huge data
- AI is here to stay

Where is AI now?

- In games:
  - Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
  - `Proverb` solves crossword puzzles better than most humans
  - Jeopardy!! Watson wipes the floor with the humans in 2011
- In cars:
  - No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- In math:
  - Proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- In war:
  - 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
  - Savings repaid DARPA's entire investment in AI research
- In space:
  - NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- AI Knowledge based system components are embedded in many real-world applications

Where is AI now?

Everywhere!