

15-381

Artificial Intelligence


Lecture 1: Introduction

Fall 2010

ARTIFICIAL INTELLIGENCE

Administration

- **Instructors:** Illah Nourbakhsh and Yaser Sheikh
- **TAs:** Sam Ganzfried, Benjamin Shih, Wooyoung Lee
- **Course Website:** <http://www.cs.cmu.edu/~bshih/15381>
- **Course Email-list:** <https://lists.andrew.cmu.edu/mailman/listinfo/15381-students>
- **Class Timing:** Tuesdays and Thursdays, 12:00pm-1:20pm
- **Book:** Artificial Intelligence: A Modern Approach (Russell and Norvig)
- **Optional Book:** Pattern Recognition and Machine Learning (Bishop)
- **Grading:** Assignments (30%), Midterm (20%), Final (30%), Paper reading (5%), Project (15%)



Can Machines Think?

Alan Turing, "Computing Machinery and Intelligence," *Mind*, 1950.

Mind-Body Dualism



I think,
therefore
I am

The Astonishing Hypothesis

“You, your joys and your sorrows, your memories and your ambitions, your sense of personal identity and free will, are in fact no more than the behaviour of a vast assembly of nerve cells and their associated molecules”

--- Francis Crick, 1994

“You insist that there is something a machine cannot do. If you will tell me precisely what it is that a machine cannot do, then I can always make a machine which will do just that.”

--- John von Neumann, 1948

What is intelligence?

“...making a machine behave in ways that would be called intelligent if a human were so behaving.”

--- John McCarthy, 1955

“One might ... define thinking as consisting of those mental processes that we don't understand.”

--- Alan Turing

DARTMOUTH CONFERENCE (1956)

Founding of AI

“The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.”


--- John McCarthy

DARTMOUTH CONFERENCE (1956)

Founding of AI




John McCarthy



Marvin Minsky




Claude Shannon



Oliver Selfridge

and “two vaguely known persons from RAND and Carnegie Tech... a significant afterthought.”
--- Pamela McCorduck, *Machines Who Think*



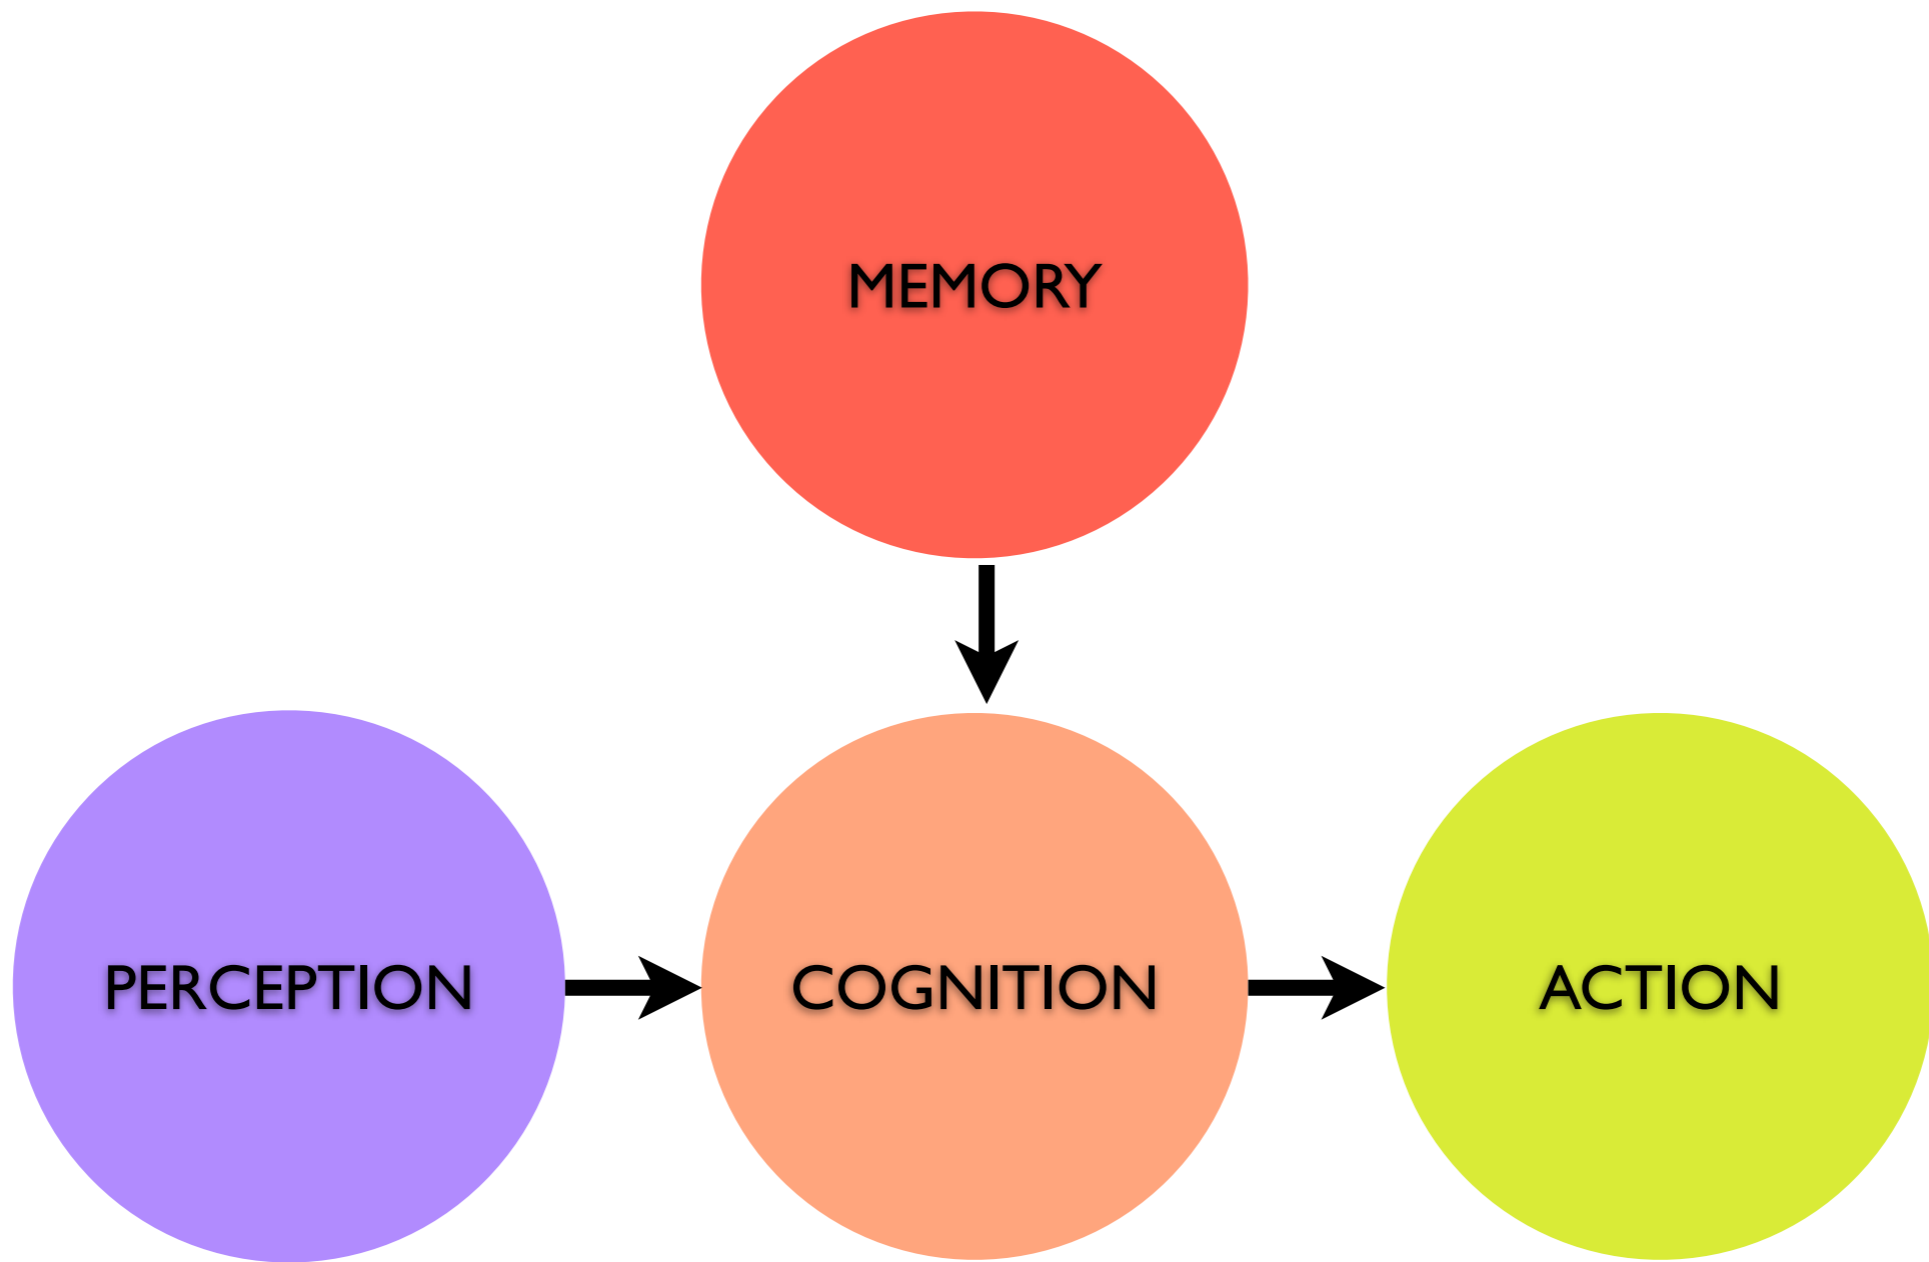
Alan Newell

Herbert Simon


**Act like
a Human**

**Act
Rationality**

- **Rational:** Optimal decision making, given current information, towards achieving a goal.




Rational Agents



$$f_{\text{rational}} = \arg \max_f V(f, E)$$

The equation defines the rational agent function f_{rational} as the function f that maximizes the score $V(f, E)$. The score V is labeled as "score", and the environment E is labeled as "environment".

THE PIANO MOVER'S PROBLEM



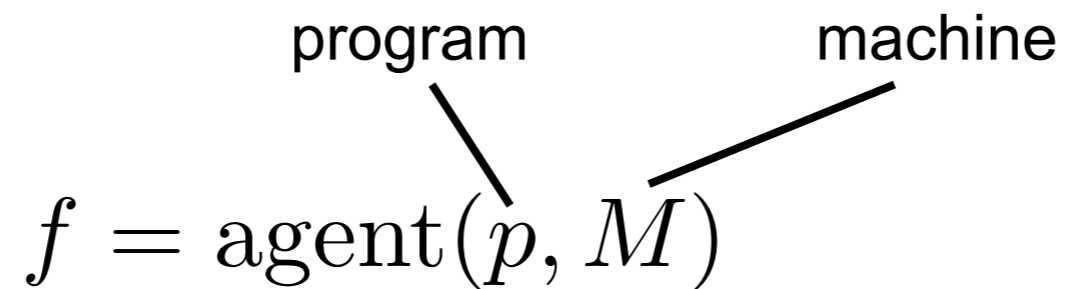
$$f_{\text{rational}} = \arg \max_f V(f, E)$$

Calculative Rationality

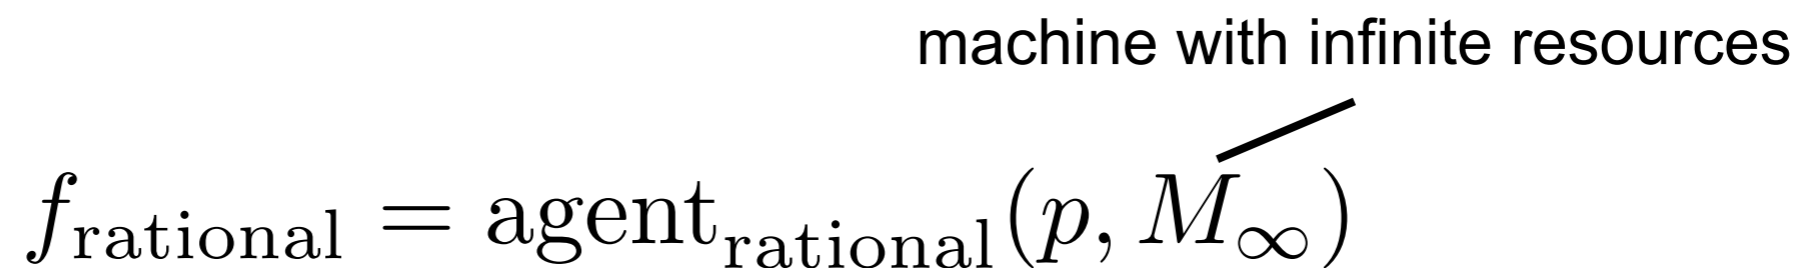
- Rational, given enough resources

$$f = \text{agent}(p, M)$$


program machine



machine with infinite resources


$$f_{\text{rational}} = \text{agent}_{\text{rational}}(p, M_{\infty})$$


THE PIANO MOVER'S PROBLEM



$$f_{\text{bounded}} = \arg \max_f V(\text{agent}_{\text{rational}}(p, M), E)$$

LOGIC: RULES FOR THOUGHT



All men are mortal
Socrates is a man
(therefore) **Socrates is mortal**



On a dark night, a policeman walks down a street

Story Credit: E. T. Jaynes, *Probability as Extended Logic*



suddenly he hears a burglar alarm

Story Credit: E. T. Jaynes, *Probability as Extended Logic*



he looks across the street,
and sees a jewelry store with a broken window

Story Credit: E. T. Jaynes, *Probability as Extended Logic*



a man wearing a mask comes crawling out, carrying a bag
which turns out to be full of expensive jewelry

Story Credit: E. T. Jaynes, *Probability as Extended Logic*

Policeman's Conclusion: The man is a burglar

RATIONAL OR IRRATIONAL?

Counter-example

The man was the owner of the jewelry store
Coming home from a masquerade party,
Just as he walked by his store a passing truck threw a
stone through the window
He was only protecting his own property.

If *A* is true, then *B* is true

A is true

B is true

If *A* is true, then *B* becomes plausible

B is true

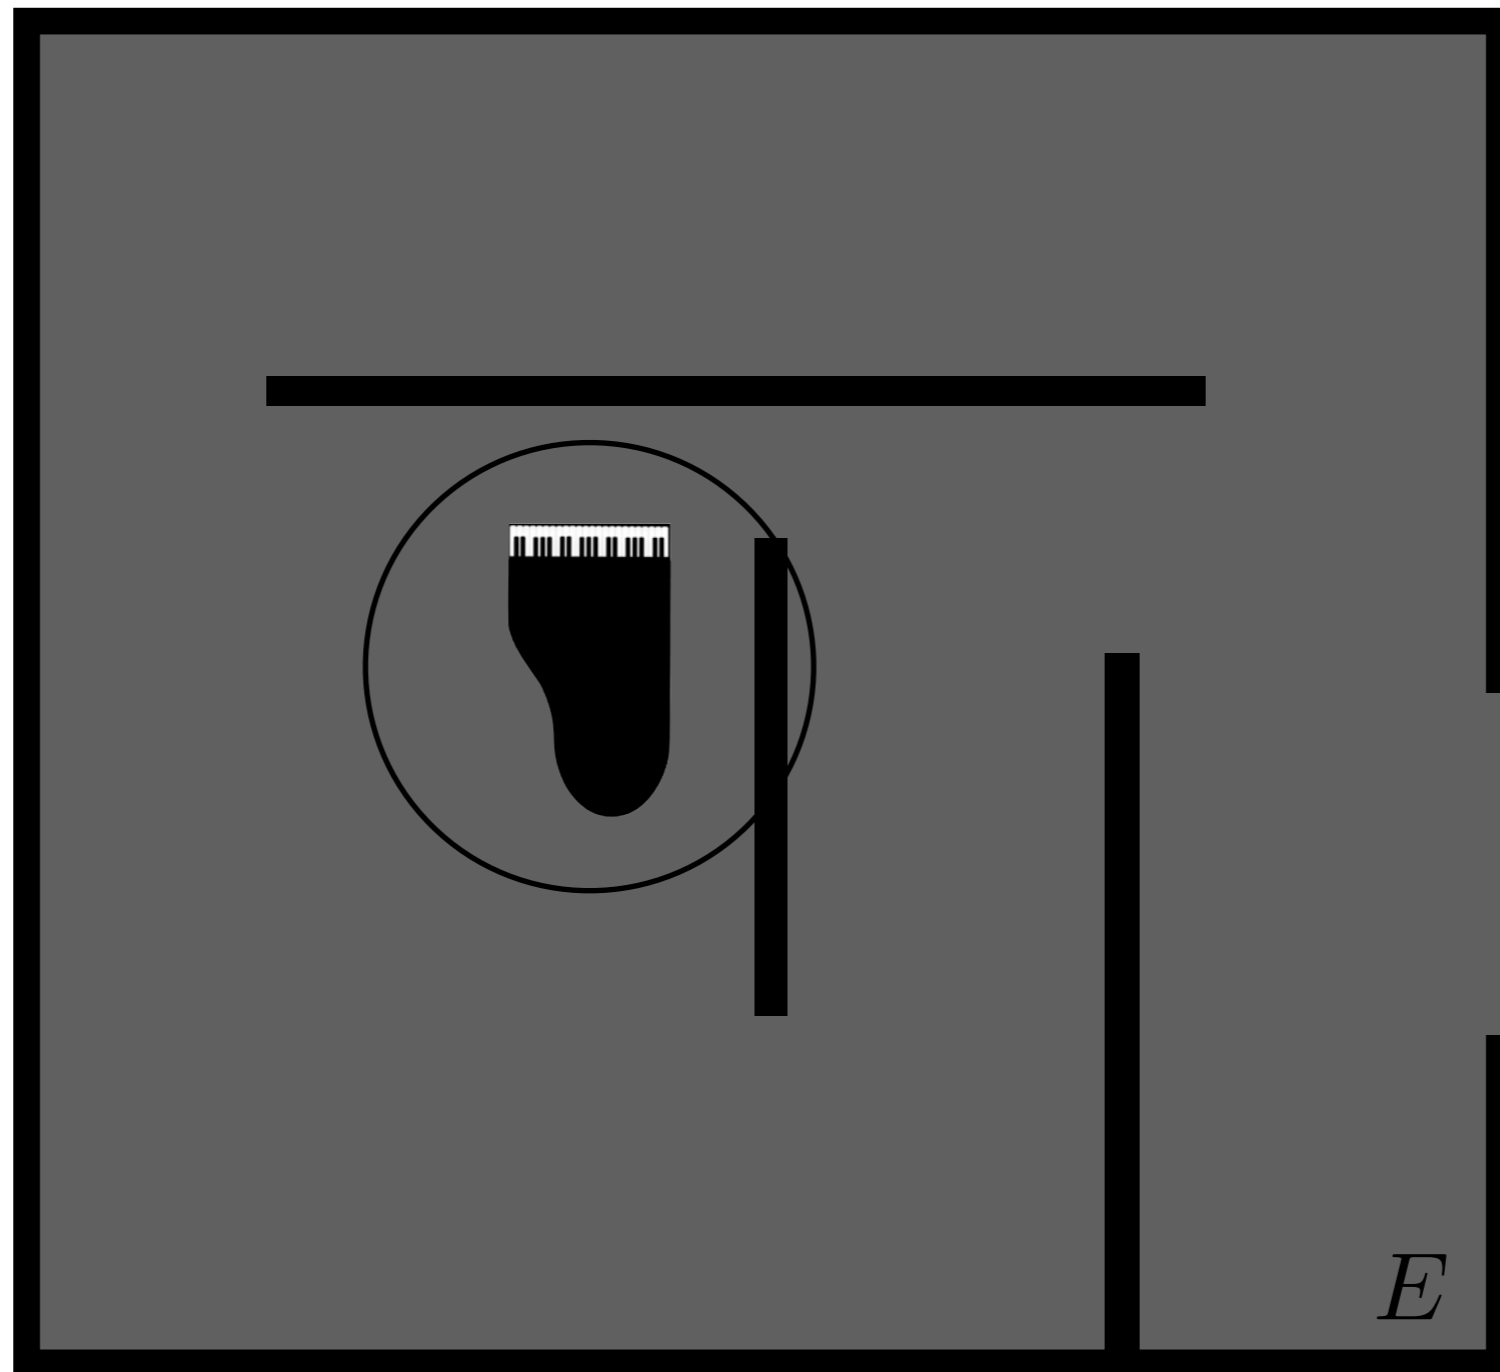
A becomes more plausible

Probabilistic Reasoning



- Sensory data is *uncertain*
- Sensory data is *incomplete*
- Bayesian Inference: Importance of *Priors*


THE PIANO MOVER'S PROBLEM IN A DARK ROOM



$O : s(E)$
 $A : \text{Translations}$
 $V : \text{Shortest path}$

$$f_{\text{bounded}} = \arg \max_f V(\text{agent}_{\text{rational}}(p, M), E)$$

VISUAL ILLUSION



COGNITIVE ILLUSIONS?

Predictable Irrationality

“as long as the task involved only mechanical skill, bonuses worked as we usually expect: **the higher the pay, the better the performance.** But when the task required even rudimentary cognitive skill ... **a higher bonus on the line led to poorer performance.**”

-- Dan Ariely *et al.*, 2009

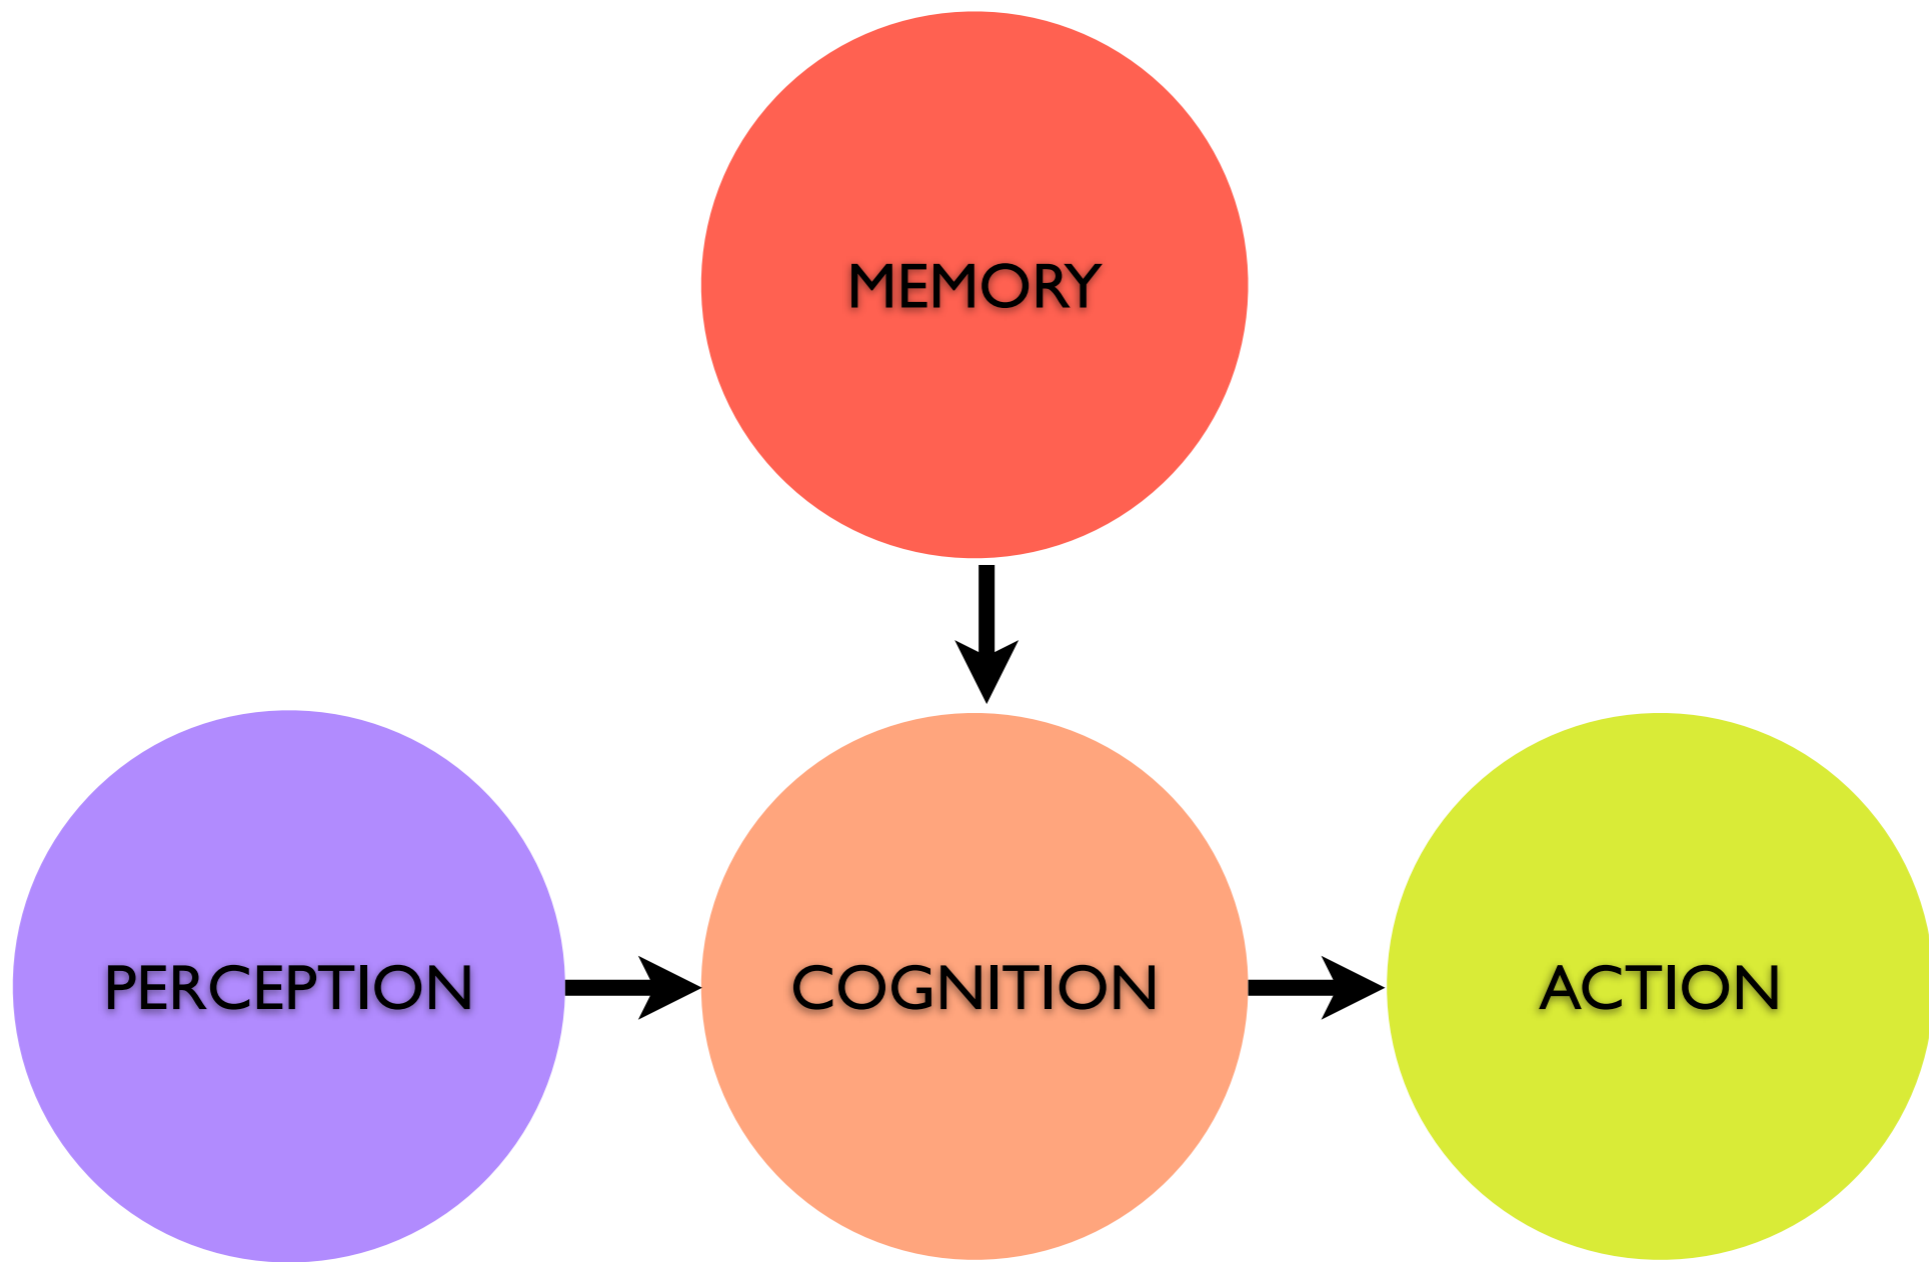
Demands of Classical Rationality

- Attach a pay-off to each possible outcome
- Attach a probability (or certainty) for each pay-off
- No room for unanticipated consequences
- Pay-offs must be *ordered*: better, same or worse than another

Bounded Rationality



- Decision when factoring in limitations of system (e.g. knowledge, mental capacity)
- “Satisficing” (satisfying and sufficient)




MORAVEC'S PARADOX

The main lesson of thirty-five years of AI research is that the hard problems are easy and the easy problems are hard.


--- Steven Pinker

Buba and Kiki




If we are predictable and mechanical,
there may be laws that govern our behavior


Biological Information Storage



Credit: Carl Sagan, *The Dragons of Eden*, 1977



Gene



Brain



Extrasomatic Knowledge



RAM



Hard drive



External Storage

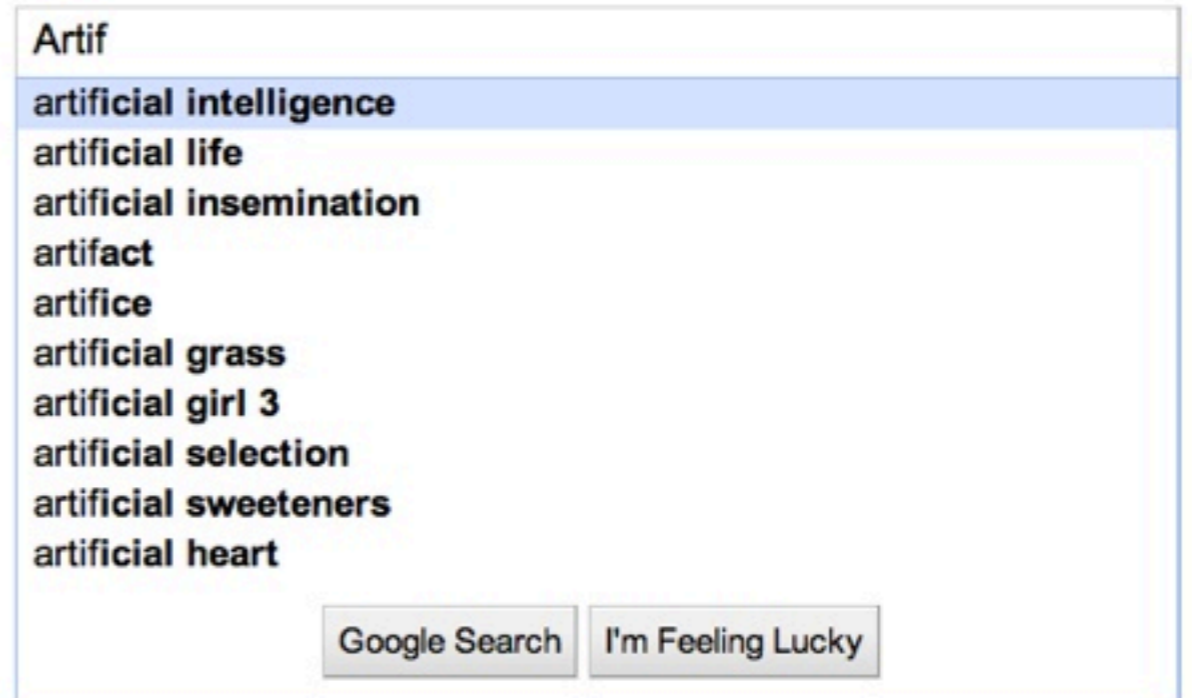
Evolution of Intelligence

- Human genes: 30 billion base pairs x 2 bits = 750 megabytes
- Human brain: 100 billion neurons x 1 bit = 12.5 gigabytes (a certain underestimate)
- Internet? ~500 exabytes (IDC, 2009)

“AI in the Postmodern Age”

--- Alexei Efros


- All interesting questions have been asked
- **Intelligence:**
 1. Remember all questions, the best answers
 2. Match current question
 3. Look up the best answer

The Google logo is displayed in its characteristic multi-colored font, with the letters 'G', 'o', 'o', 'g', 'l', and 'e' in blue, red, yellow, blue, green, and red respectively.

REPRESENTATION

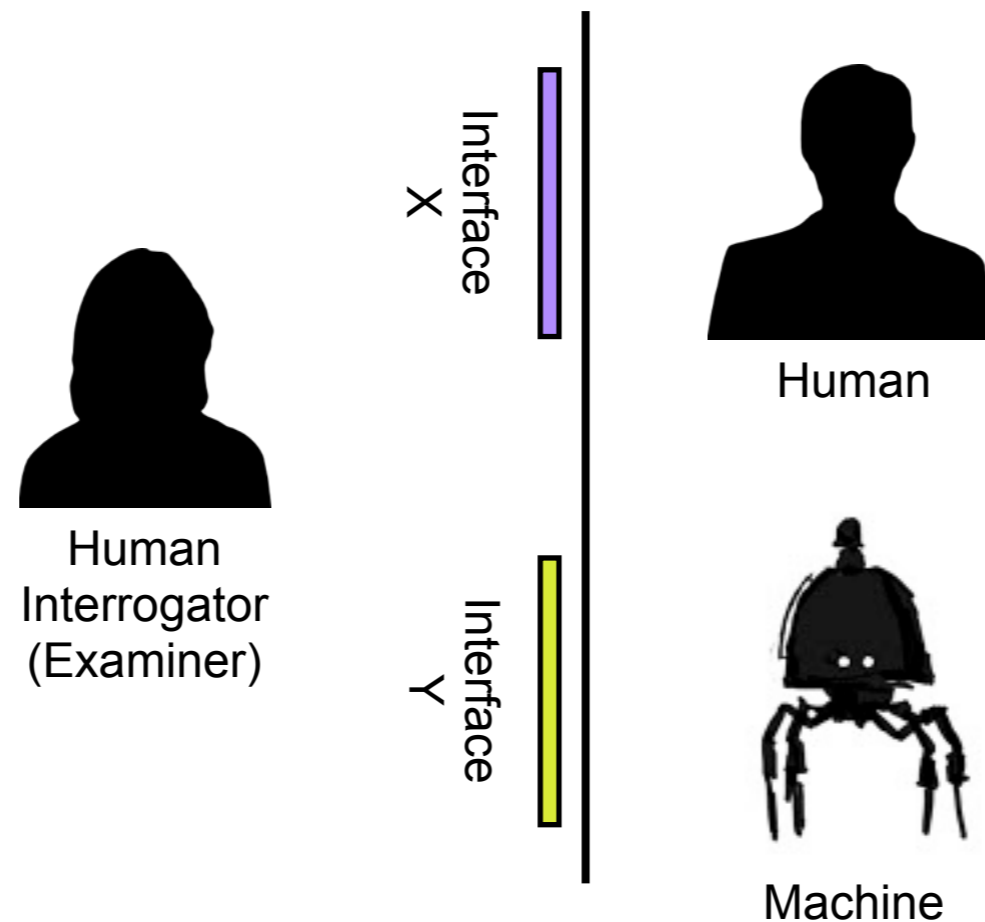
- All knowledge needs representation
 - Genetic Information
 - Behavioral Information
 - Audio Information
 - Visual Information
 - Textual Information

A major part AI is representing the problem space to allow efficient optimization for the best solution



Can Machines Think?

TURING TEST



Answer: X is a human and Y is a machine

Turing test:

70% success after 5 minutes of conversation

CAPTCHA

Select an ID and password

Yahoo! ID and Email @
Password Password Strength
Re-type Password

In case you forget your ID or password...

Alternate Email (optional)
Secret Question 1
Your Answer
Secret Question 2
Your Answer

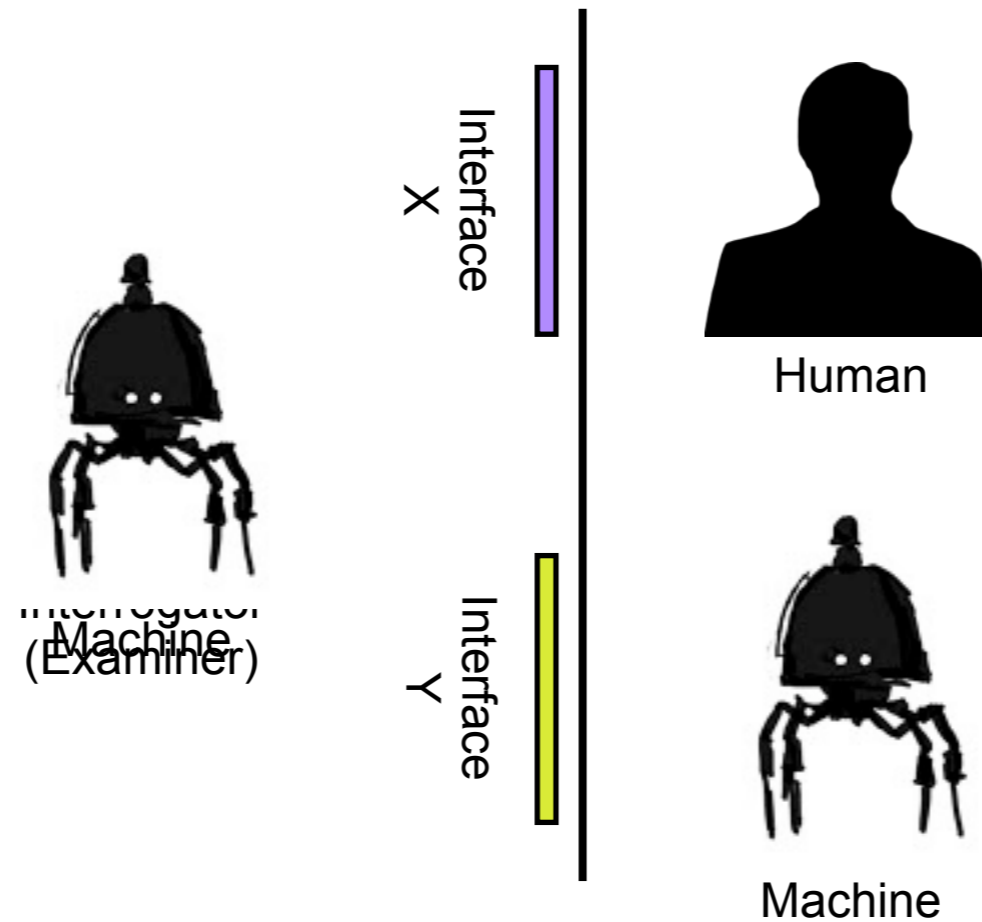
Type the code shown



By clicking the "Create My Account" button below, I certify that I have read and agree to the [Yahoo! Terms of Service](#), [Yahoo! Privacy Policy](#) and [Mail Terms of Service](#), and to receive account related communications from Yahoo! electronically.

Completely Automated Public Turing test to tell Computers and Humans Apart

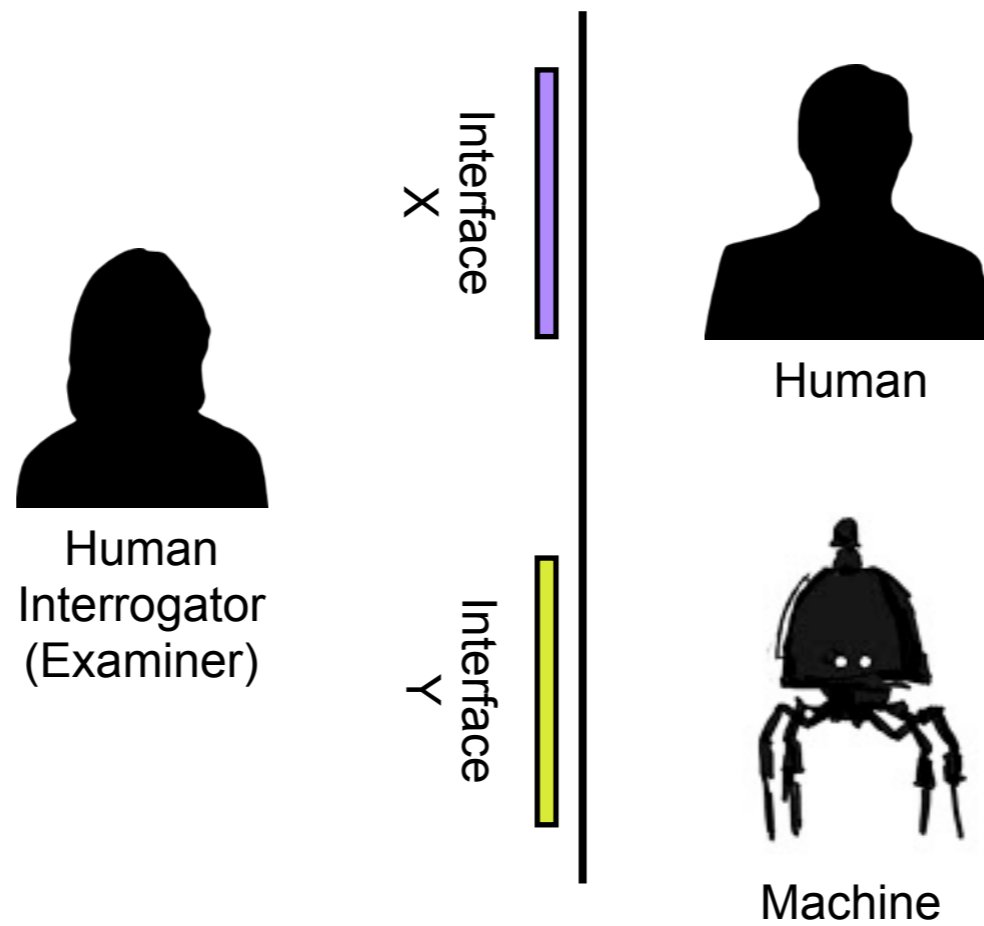
CAPTCHA



Completely Automated Public Turing test to tell Computers and Humans Apart


TURING TEST

COURSE PROJECT



TURING TEST

COURSE PROJECT



Credit: <http://xkcd.com>

Aug. 24: Introduction to Artificial Intelligence
Aug. 26: AI and the Brain + Paper selection
Aug. 31: Search I (Uninformed Search)
Sep. 02: Search II (Uninformed Search)
Sep. 07: Search III (Informed Search)
Sep. 09: STRIPS

Assignment 1 due (7.5%)

Sep. 14: Probabilistic Reasoning
Sep. 16: Bayesian Inference I
Sep. 21: Bayesian Inference II
Sep. 23: Markov Decision Processes

Assignment 2 due (7.5%)

Sep. 28: Robots
Sep. 30: Motion Planning

Oct. 05: Review
Oct. 07: **Midterm (20%)**

Oct. 12: Vision I: Photometry
Oct. 14: Vision II: Geometry
Oct. 19: Vision III: Recognition
Oct. 21: Lightning Session

Assignment 3 due (7.5%)

Oct. 26: Natural Language Processing
Oct. 28: Optimization I: Linear and Nonlinear Optimization
Nov. 02: Optimization II: Genetic Algorithms
Nov. 04: Optimization III: Stochastic Optimization, Hill Climbing

Assignment 4 due (7.5%)
Turing Test Dry Run (5%)

Nov. 09: Decision Trees, Neural Networks
Nov. 11: Classification/Clustering I
Nov. 16: Classification/Clustering II
Nov. 18: Reinforcement Learning

Turing Test Demo (10%)

Nov. 23: Human Robot Interaction I
Nov. 25: Human Robot Interaction II

Nov. 30: Review
Dec. 02: **Final (30%)**

**Defining intelligence is
the grand intellectual
challenge of the century**

Rene Descartes
Philosopher...

Aristotle
Philosopher...

Claude Shannon
Information Theorist

Warren McCulloch
Neurophysician

**John
von Neumann**
Mathematician...

Thomas Bayes
Reverend

Francis Crick
Molecular Biologist

Alan Turing
Computer Scientist

Herbert Simon
Economist...

Walter Pitts
Logician

Marvin Minsky
Cognitive Scientist

Alan Newell
Cognitive Psychologist

**“AI still has openings
for a full-time Einstein”
--- Russell and Norvig**

Reading List

- **Can Machines Think?** Alan Turing
- **A Behavioral Model of Rational Choice**, Herbert Simon
- **Knowledge-Level**, Alan Newel
- **Dragons of Eden**, Carl Sagan
- **Society of Mind**, Marvin Minsky
- **The Emotion Machine**, Marvin Minsky
- **Probability Theory as Extended Logic**, Edwin Jaynes
- **Machines Who Think**, Pamela McCorduck [History of AI research]
- **Treatise on Man**, Rene Descartes
- **Elephants Don't Play Chess**, Rodney Brooks
- **Dawn of the Age of Stochasticity**, David Mumford