Introduction to Artificial Intelligence
CSCI 3202
Fall 2007

Greg Grudic

Questions?

This Class

- Intro continues…
- Amateur AI philosophy…

What is AI?

Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

Warning, I advocate for “acting rationally” based on Machine Learning

- but I am willing to hear other arguments and change my mind

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

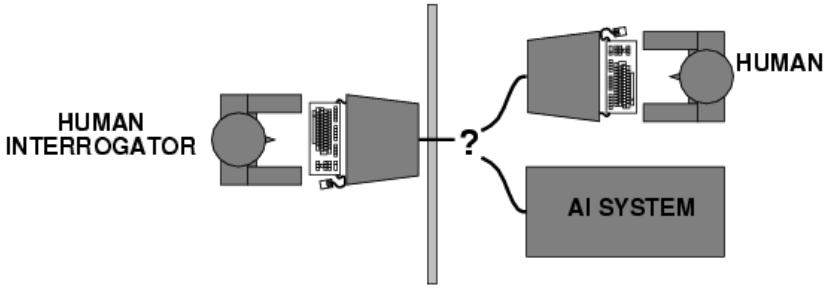

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning

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?
  3. Should thinking need to be associated with actions?

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
  - Problem: How do we know the agent is doing this?
- 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:

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

- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveats:
  - computational limitations make perfect rationality unachievable
    - design best program for given machine resource
  - Can we ever know if an agent is acting rationally?

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 (fast enough?)
- 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
- 1952—69 — Great enthusiasm for AI!
- 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 (Machine Learning) return to popularity
- 1990-- — Machine Learning, Statistics and Mathematics join forces
- 1987-- — AI becomes a science
- 1995-- — The emergence of intelligent agents

Some state of the art AI

- 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

My Personal View of AI

- I want to build a robot that will
  - Clean my house
  - Cook when I don't want to
  - Wash my clothes
  - Cut my grass
  - Fix my car (or take it to be fixed)
  - i.e. do the things that I don't feel like doing…
- Therefore: AI is (to me) the science of building machines (agents) that act rationally with respect to a goal.

Agent: *sensing*, *computation*, and *action*

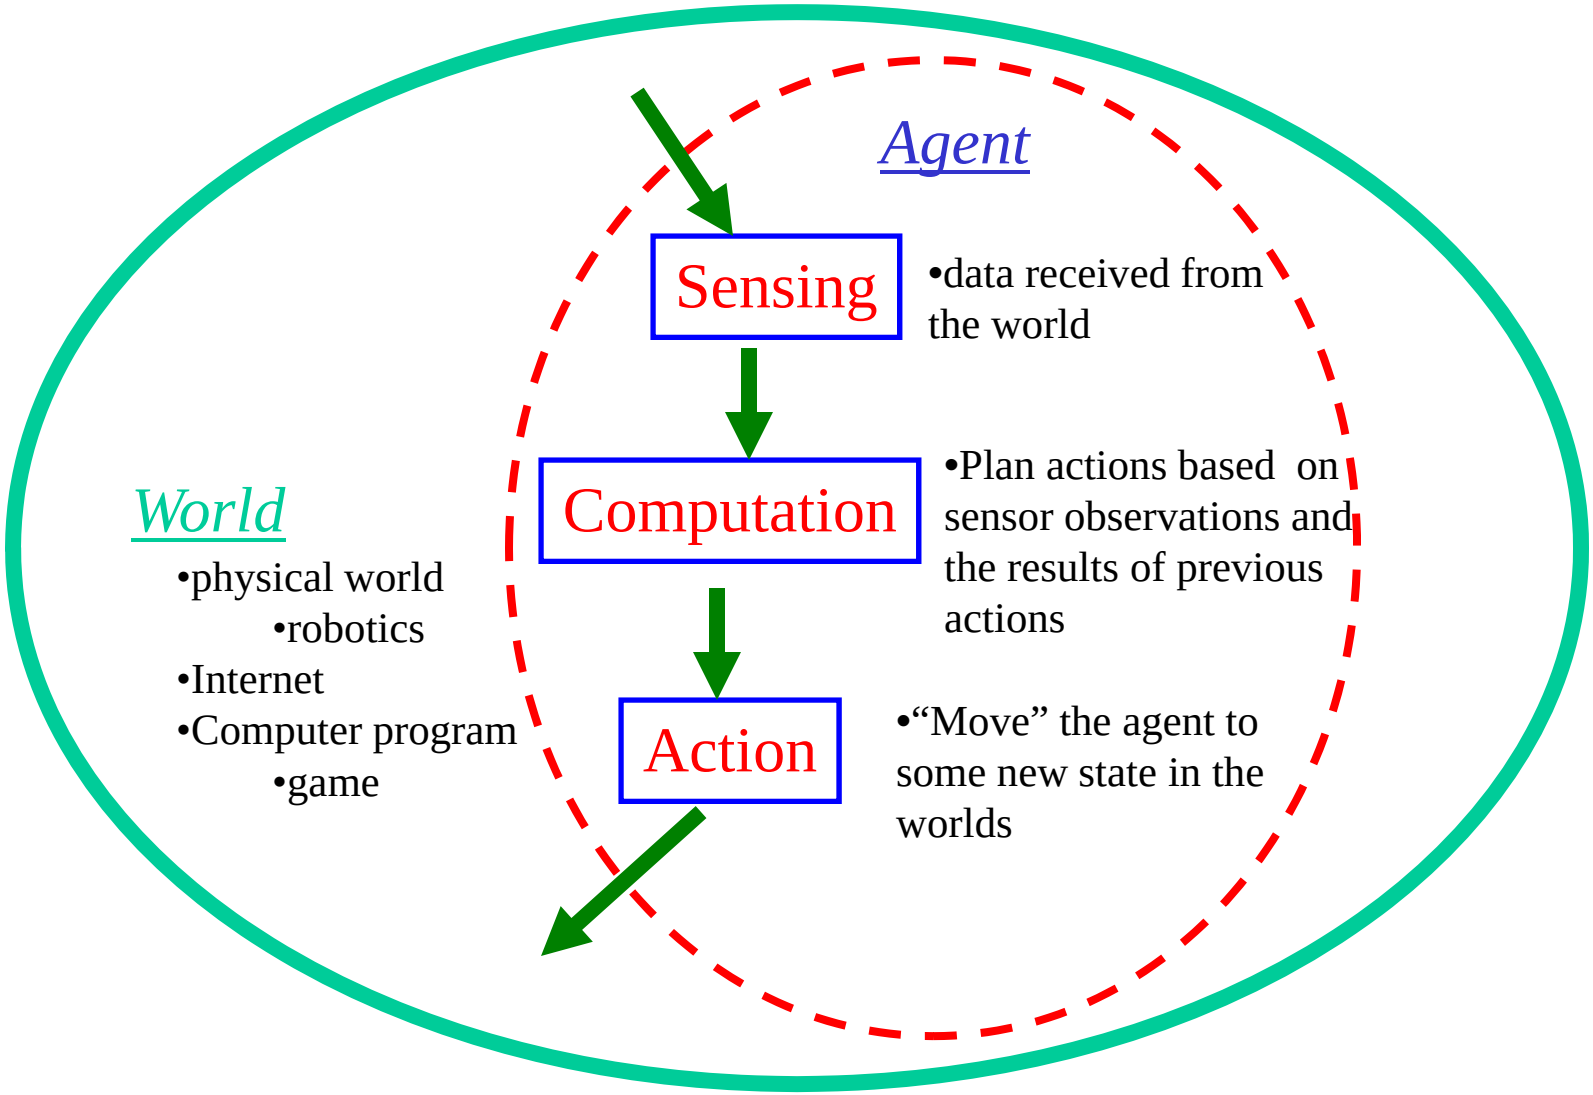

What is a Rational Agent?

- An agent is an entity that senses, computes and acts in some world
- A rational agent is one that does the right thing
  - The right thing: that which is expected to maximize goal achievement (*accomplishing tasks that Greg doesn't feel like doing*), given the available information
- This is **not** a new idea:
  - Aristotle (Nicomachean Ethics): *Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good*

Elements of AI

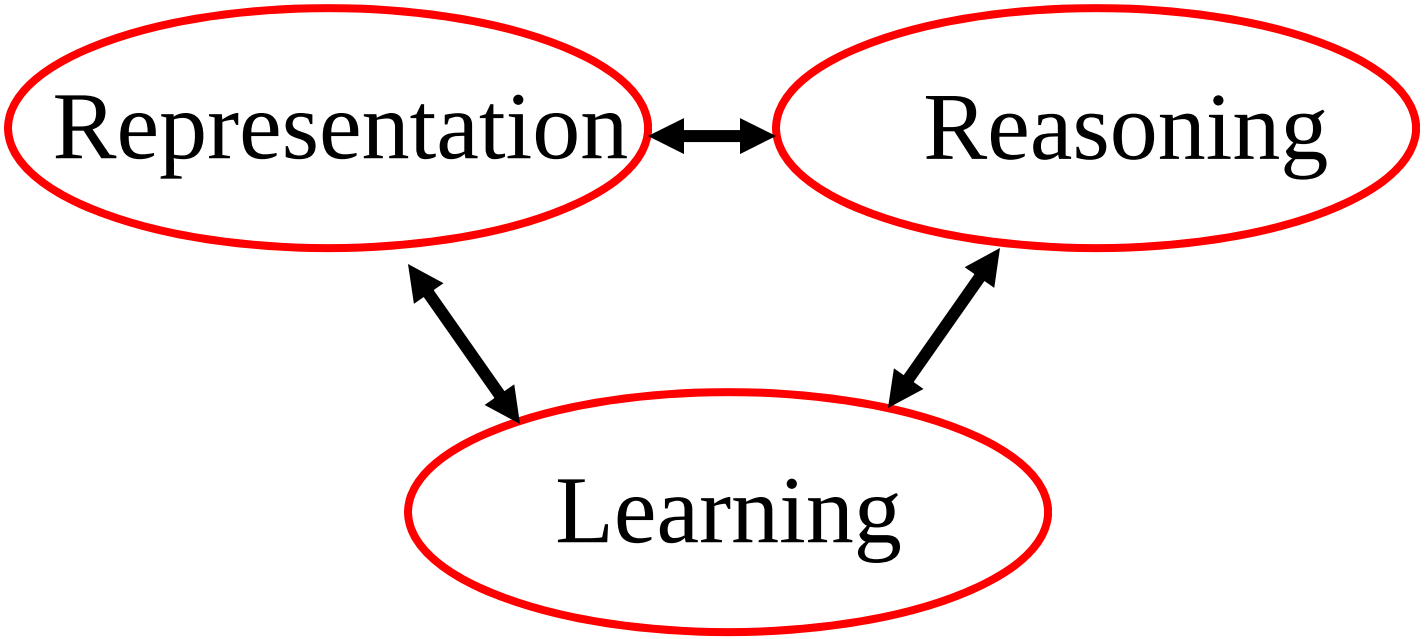

(My) Elements of AI

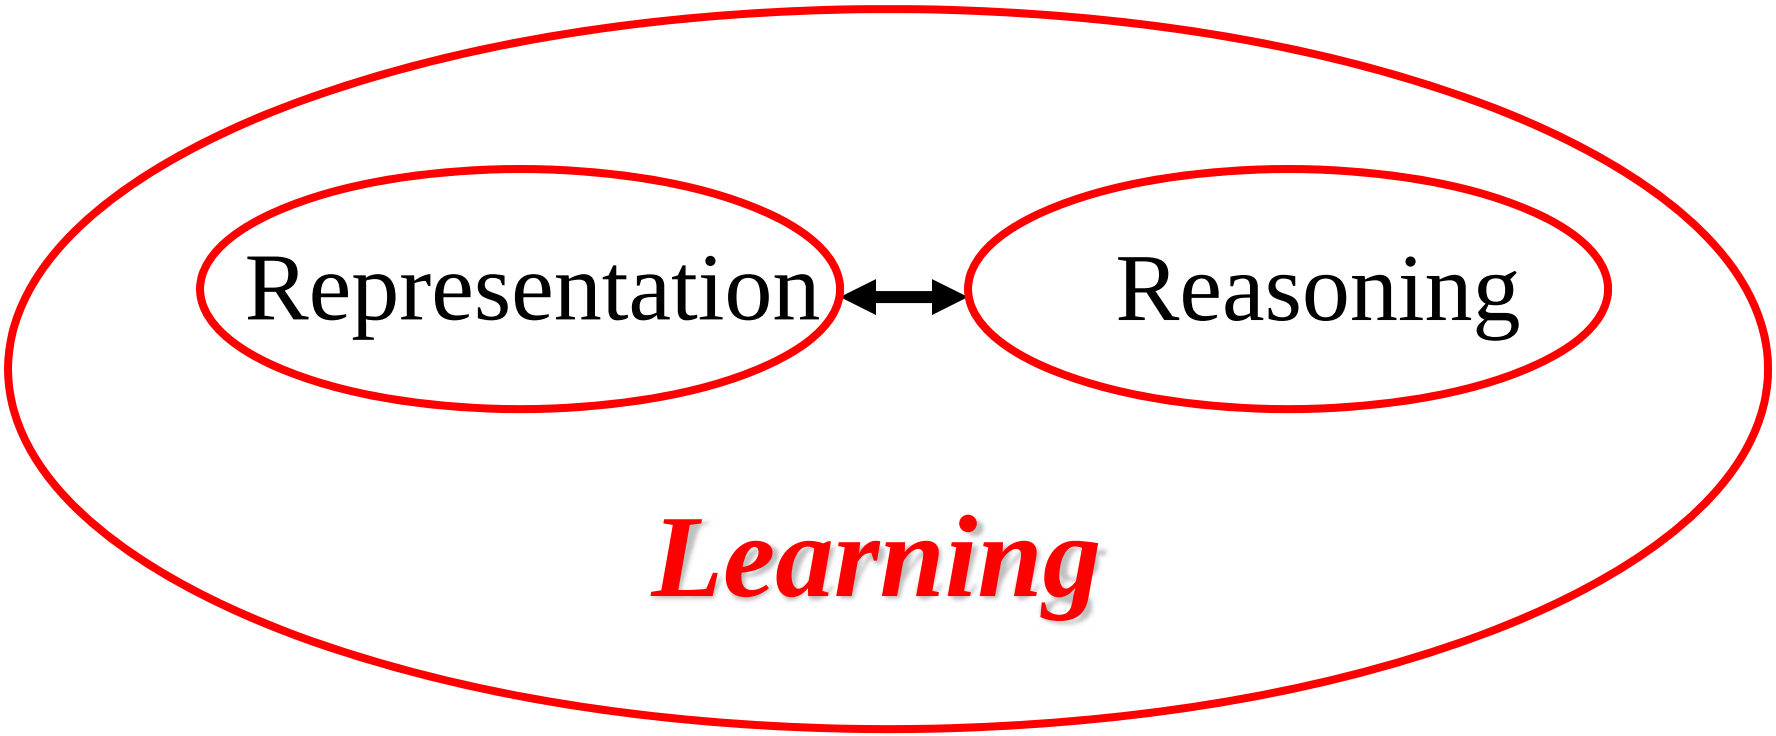

Why Must Representation and Reasoning be Encompassed by Learning?

- Fundamental lesson of AI (learned in the 1980's):
  - It is not possible to hand code knowledge about anything but the most trivial problem domains!
    - Uncertainty is a key problem!
  - Expert Systems: largely failed because an expert (e.g. doctor) doesn't know how to formalize (code) what makes her an expert!
  - For Example: I'm an expert on chairs but I can't (and no one can!) write a program that identifies chairs in an image
    - **However, ML techniques can!**
- How can I reason rationally about a world I cannot encode knowledge about?
- I do not believe that an agent can gain knowledge about a world without sampling it and learning from those samples….

AI Agent: A Different Perspective

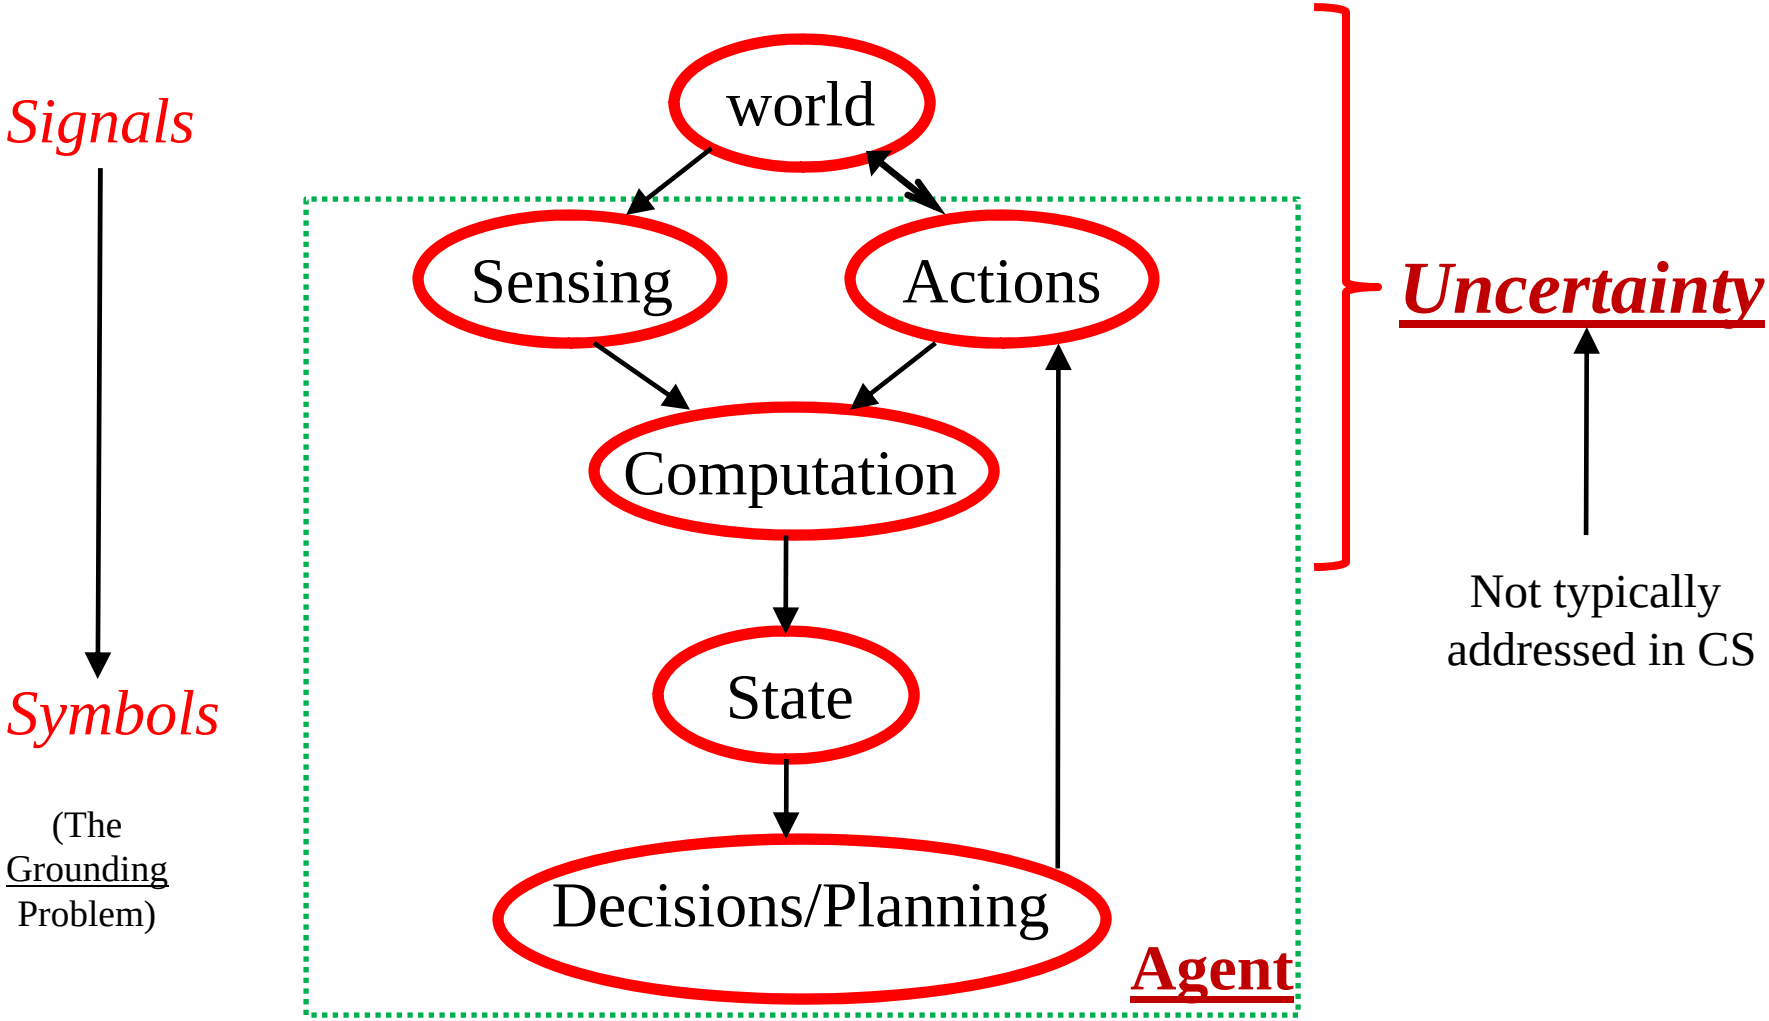

Why is Machine Learning Important?

- Machine Learning is a Principled Methodology for dealing with uncertainty (noise) in
  - world
  - sensors
  - computation
  - action

Where can Machine Learning Algorithms be Found?

- Marketing
  - Who should a company target for advertising?
- Profiling
  - Is passenger 57 likely to hijack the plane?
- User interfaces
  - Making it easier to interact with a PC by anticipating what I am doing.
- Document characterization
  - Searching the web for things of interest.
- Bioinformatics
  - Human genome project
    - Which gene is responsible for the cancer that runs in my family?
- Data mining
  - "Data doubles every year", Dunham 2002
  - ML algorithms are used to make sense of this data
- Economics, medical diagnosis, robotics, computer vision, manufacturing, inventory control, elevator operation….

What is Machine Learning?

- "The goal of machine learning is to build computer systems that can adapt and learn from their experience."
  - Tom Dietterich

- What does this mean?
- When are ML algorithms NOT needed?

A Generic System (Agent?)

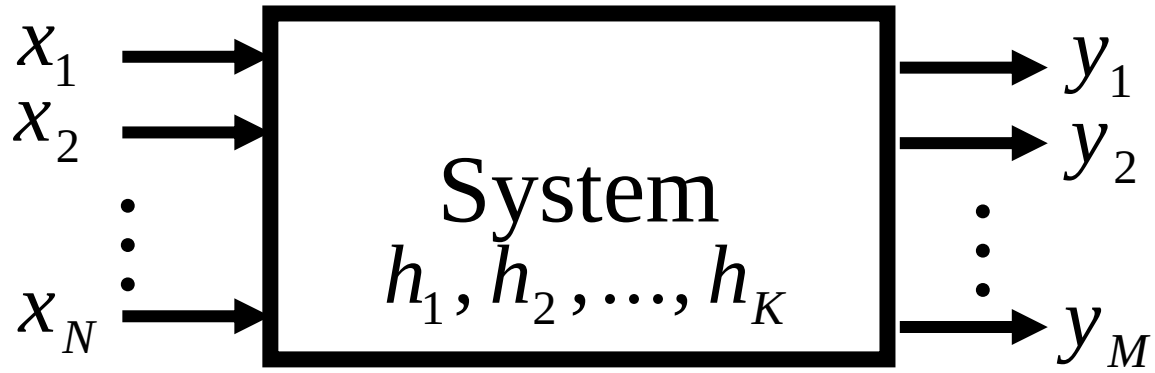

Input Variables: $\mathbf{x} = (x_1, x_2, \ldots, x_N)$

Hidden Variables: $\mathbf{h} = (h_1, h_2, \ldots, h_K)$

Output Variables: $\mathbf{y} = (y_1, y_2, \ldots, y_K)$

Another Definition of Machine Learning

- Machine Learning algorithms discover the relationships between the variables of a system (input, output and hidden) from direct samples of the system

- These algorithms originate form many fields:
  - Statistics, mathematics, theoretical computer science, physics, neuroscience, etc

When are ML algorithms NOT needed?

- When the relationships between all relevant system variables (input, output, and hidden) is adequately understood!

- This is NOT the case for many *complex* real systems!

Main Subfields of Machine Learning

- Supervised learning
  - Classification
  - Regression
- Semi-Supervised (Transduction) learning
- Active learning
- Reinforcement Learning
- Unsupervised Learning

Learning Classification Models

- Collect Training data
- Build Model: happy = **f**(feature space)
- Make a prediction

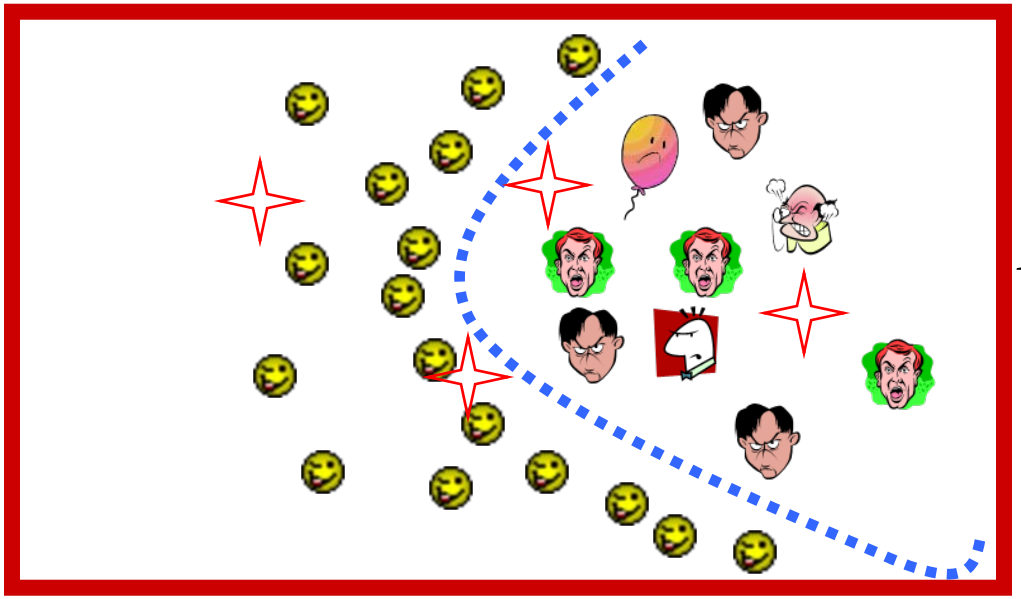

High Dimensional Feature Space

Learning Regression Models

- Collect Training data
- Build Model: stock value = **f**(feature space)
- Make a prediction

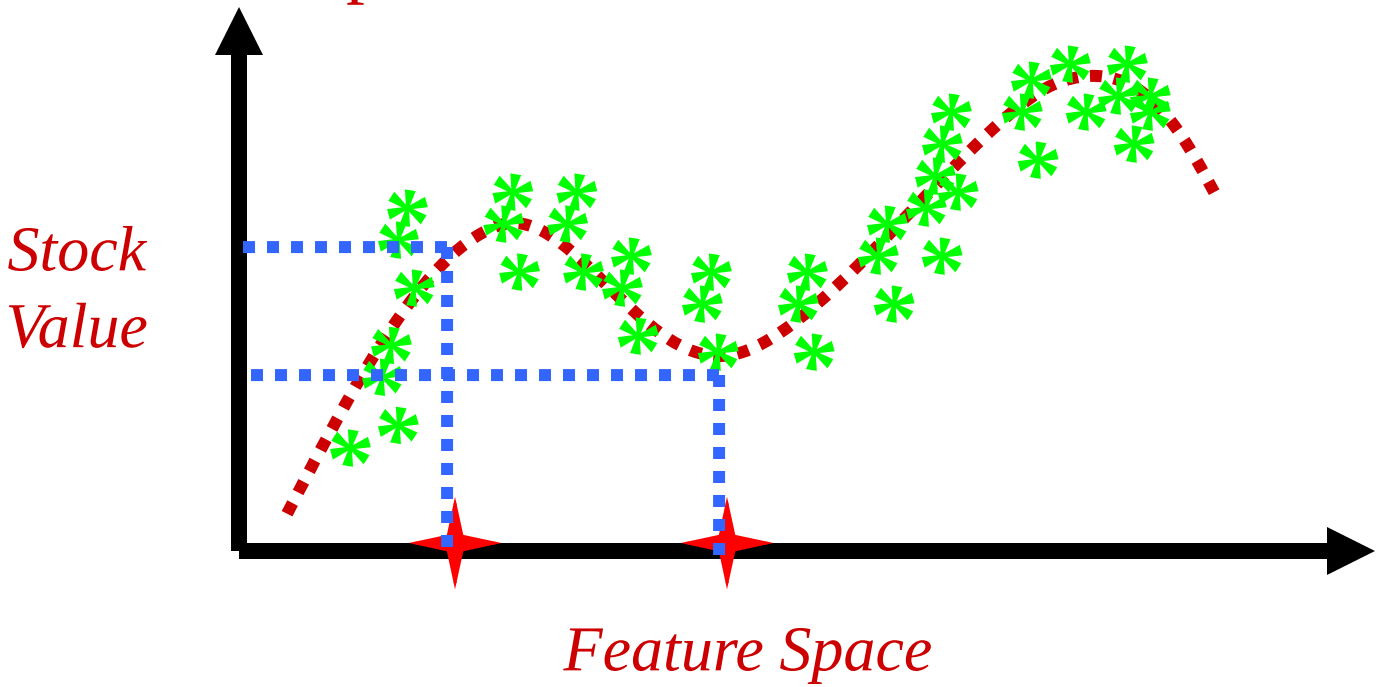

Search

- AI can be thought of as
  1. Specification of a GOAL
     - Optimization criteria…
  2. Method for searching action and sensor space to achieve the goal
- Two types of searches
  - Symbolic (logic, reasoning, etc)
  - Numeric – establish a continuous search space (topology)
- Search in the real world is hard….
  - Efficient solutions require constraints in search space
- Machine Learning is one framework for efficiently constraining search…

Planning

- Start with an assumed structure in the problem space
  - e.g. robot in a Cartesian World (3-D map) wants to get to a GPS goal position from some start GPS position
- Structure is used to plan a sequence of actions from some initial state to a goal state.

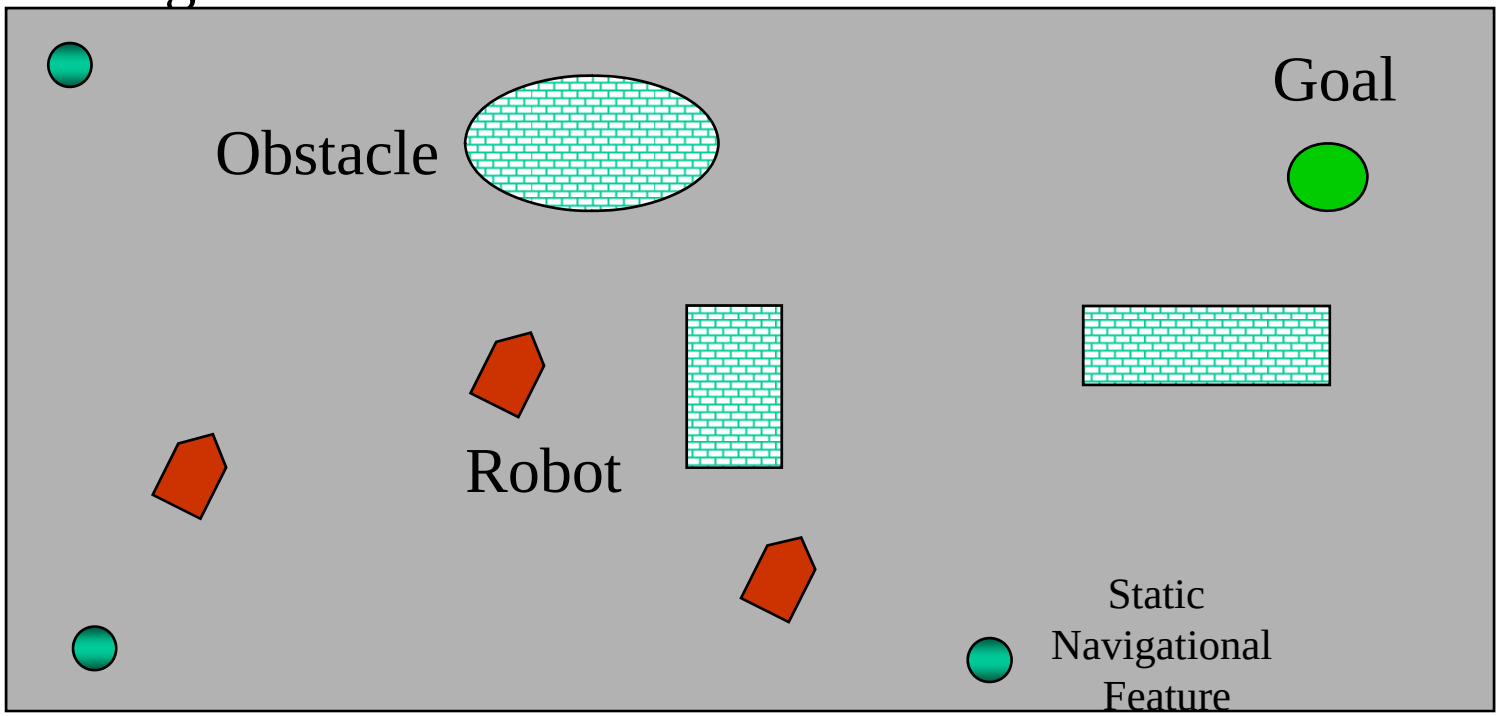

Optimal Decision Theory

- Acting under uncertainty
  - Measuring uncertainty in complex environments is the domain of Machine Learning
- Given all the available information, what is the optimal decision (or action) that the agent should take?

Computer Vision

- The camera is our best sensor for physical human environments…
  - Humans are extremely good at interpreting the world visually
  - AI systems that work in the human physical world need to utilize visual data
- Computer vision uses realistic constraints and knowledge of camera geometry to infer knowledge about the world from 2D images

Robotics

- Robotics is AI in the physical world
- It is the hardest subfield of AI because robots must sense and act in the (uncertain) physical world
  - AI inside a computer (internet) is much more constrained
- The computer revolution has changed the world….
- **However**, the robotics revolution, when it happens, will make the computer revolution pale in comparison