Computer Science 6100/4100: Machine Learning

RPI, Fall 2008
Instructor: Sanmay Das

What is Machine Learning?

- "Enabling computers to learn from data"
- Supervised learning: generalizing from seen data to unseen
- Unsupervised learning: Finding patterns in input data
- Reinforcement learning: learning how to act

Where Does This Fit in AI?

- What is the goal of AI?
  - Two dichotomies: thinking/acting and humans/rationality (Russell & Norvig [RN])
- Turing test: acting like humans (restricted domain)
- Alternative goal: acting rationally?

Rational Behavior

- [RN]: Maximize goal attainment
- Easier when we specify utility functions
- Perfect rationality is problematic:
  - Computational limitations
  - Definition of optimality

Intelligent Agents [RN]
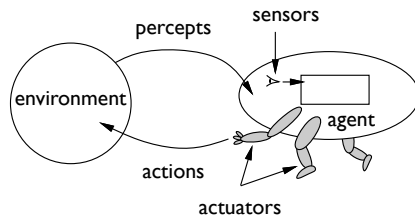
- Agents include humans and artificial agents
- Agent function maps percept histories to actions $f : P* \rightarrow A$
- Agent program runs on the physical architecture to produce f

Human Beings and Humanoid Robots

- Environment: ?
- Sensors/Percepts: ?
- Actuators/Actions: ?
- Performance measure: ?

A Trading Agent

- Environment: ?
- Sensors: ?
- Percepts: ?
- Actuators: ?
- Actions: ?
- Performance measure: ?

Supervised Learning

- Induction
- Simplest form: true function
  $$f : X \rightarrow Y$$
- You are given pairs generated from f
  $$(x_1, y_1), (x_2, y_2), \ldots (x_n, y_n)$$
- Learn h close to f

Learning a Hypothesis

- How do we learn *h*? Algorithms from machine learning and statistics
- What can we prove about *h*? Statistical/computational learning theory
- Different learning algorithms have different *inductive biases*
- Example...good loans

Hypothesis Spaces

- Examples: decision trees, linear classifiers
- Finding a hypothesis can be thought of as search in the hypothesis space
- Hypothesis space is part of the inductive bias

Unsupervised Learning

- No explicit outputs
- Build a model of the inputs in some way
  - Probabilistic model of feature distribution
  - Clustering

Reinforcement Learning

- Agent interacts with the world
- Receives feedback in the form of rewards (or costs)
- Must choose which actions to take
- Major issues:
  - Delayed reward/credit assignment
  - Exploration/Exploitation

Elements of RL

- Sutton & Barto [SB]:
  - Policy
  - Reward function
  - Value function
  - Model?

Markov Decision Processes (Problems)

- State space: S
- Initial state: S_0
- Action space: A
- Transition model: $T : S \times A \to S \times [0, 1]$
- Reward function: $R : S \to \mathbb{R}$

Utility Theory

- Going from the real-world to sensible reward functions
- Parallelism: in economics, utility theory is useful in abstracting over preferences
- Thought question: which of these two options would you prefer? $1000 with 50% prob., or $400 for sure?
- Basis for the insurance industry!

Learning in Economics

- Agents that participate in markets are assumed to be rational
- This means they solve interesting learning and decision-making problems
- Change the focus to understanding how the interaction of rational players leads to system-wide dynamics...
  - Two restaurants
  - Kyle's model

Syllabus and Course Policies

- Come to class
- Participate
- 3 projects (can be done in pairs)
- 2 in-class exams
- Check the website regularly!

A Note on Math

- Calculus
- Ability to play with matrices
- Probability!
  - Uniform, Gaussian distributions, Bayes rule
    - Let's do a quick problem...

Problem 1

- MBC Instruments has designed a new test for Horrible Disease. This test is correct with 99% accuracy. 1 in 10000 people in the general population has the disease. SBC took the test and it came out positive. Is it more likely that SBC has the disease or doesn't? Calculate the probability and turn in your work

Statistics

- Difference between standard deviation and standard error?