ECS 271 Machine Learning

Rao Vemuri

Spring 2003

ECS 271: Outline

- Introduction
- Concept Learning
- Decision Tree Learning
- Artificial Neural Nets
- Computational Learning Theory
- Bayesian Learning
- Genetic Algorithms + other topics

Grading

- 45% for a term paper
- 25% Mid term (in class, Closed book)
- 15% {Home Work Assignments}
- 15% Final (Take Home, 24 hr turn-around)
  - Each student reads and critiques two term papers from other students. These critiques have no influence on how I grade the term papers.

Outside Reading

- The text book is really at an upper division UG level. You are expected to read outside papers to boost the theoretical background.

Term Paper

- Choose the topic. Any problem with which you are **very familiar** would be a good topic
- Choose a ML Technique discussed in the book
- Search the WWW for a good data set
- Implement the method, get the results
- Write a report (see class web site for more details)
- This is a 4-credit course. I expect you to devote a minimum of 20 hours on this project alone and another 20 hours for the rest of the course.

What is Machine Learning?

- Learning is
  - making *useful* changes in our minds (Minsky)
  - constructing or modifying *representations* of what is being *experienced* (Michalski)
  - knowledge acquisition in the absence of explicit programming (Valiant)
  - An adaptive process that enables a system to do the same type of task more effectively the next time around (Simon)
    - "type" means "drawn from the same population"

Why ML? Why Now?

- Recent progress in algorithms and theory
- Growing flood of on-line data
- Growing computational power
- Easier to build than similar programmed systems
  - Search engines
  - Computer games (TD-gammon)
  - User interfaces/Personal assistants (paper clip asst.)
  - Robot exploring Mars/ Vacuum Cleaner
- Better understanding of human/animal learning

Opportunities for Tomorrow

- Learn across full mixed-media data
- Learn across multiple databases
  - Internal databases
  - Web and Newsfeeds
- Learn by active experimentation
- Learn to make decisions rather than predictions

Types of Learning

- Rote Learning (memorization)
- Learning from examples (*generalize* from specific instances)
- Learning by being told (*compile* abstract advise)
- Learning from scalar feedback (*reinforcement* learning)

What is the Learning Problem?

- Given
  - a task T
  - a performance measure P
  - some experience E
- Learn to do the task
- Example: Learn to play T = checkers until P = 90% of the campus-wide tournaments are won, by using E = the opportunity to play against player A.

How Do You Learn Checkers?

- What exactly is the meaning of "learning to play checkers"?
  - Given a board, what is the next move?
  - Given a board, what is the best move from among {legal moves}?
- The {legal moves}defines a large search space. From this space we choose a move. That is, we are learning a target function V, given the board.
  - Target function, V: Board → Real Number
  - ChooseMove: Board → Move

Target Functions

- Define a function that assigns a score to the board
- The better the board, the better the score assigned.
- Then our search is really to maximize the score.

V for Checkers

- If *b* is the final board state
  - *V(b)* = 100, for a winning state
  - *V(b)* = -100, for a losing state
  - *V(b)* = 0, for a draw
- If *b* is NOT the final board state
  - *V(b) = V(b`),* where *b`* is the best final board state that can be reached from *b* while playing optimally (This is non much of a help)

Operational Definition

- What we want is a function that we can compute.
- What exactly is the value V of a board?
  - We do not know. The best we can do is approximate with something we can compute
  - The most friendly computable functions are polynomials

$$V'' = w_0 + w_1 x_1 + \ldots w_n x_n$$

Representation of Learned Function, V”

- V” is an approximation to the unknown V
- Xi’s are features of the board
- X1= # of black pieces
- X2 = # of red pieces
- X3 = # of black kings
- X4 = #of red kings
- X5 = # of blacks threatened by reds
- X6 = # of reds threatened by blacks

A Concept Learning Problem

- Consider a task
- The task is to learn a concept described by Boolean attributes
- Both positive and negative examples are provided by a teacher
  - The concept to be learned is “dog”

$$\text{is\_mammal} \wedge \text{has\_claws} \wedge \neg\text{can\_fly} \wedge \text{has\_tail} \wedge \text{can\_bark}$$