

Machine Learning Overview

Sargur N. Srihari

srihari@cedar.buffalo.edu

What is Machine Learning?

- Principled way of building high performance information processing systems
 - search engines, adaptive user interfaces, personalized assistants, web bots, and scientific applications
 - from medical records: treatments most effective
- How to construct computer programs that automatically improve with experience
 - houses learn to optimize energy costs
 - computer games
- Programming tasks that humans perform well but difficult to specify algorithmically
 - Language Related Technologies
 - IR, NLP, DAR, ASR

ML vs PR

- ML has origins in Computer Science
- PR has origins in Engineering
- They are different facets of the same field

ML and AI

- If an expert system--brilliantly designed, engineered and implemented--cannot learn not to repeat its mistakes, it is not as intelligent as a worm or a sea anemone or a kitten.
-Oliver G. Selfridge, from *The Gardens of Learning*.
- "Find a bug in a program, and fix it, and the program will work today. Show the program how to find and fix a bug, and the program will work forever."
- Oliver G. Selfridge, in [AI's Greatest Trends and Controversies](#)

ML opens new uses of computers


new levels of competence and customization

better understanding of human learning abilities/disabilities

Broad Definition of Learning

- Include most tasks conventionally called “learning”
- Encompass computer programs that improve from experience in straightforward ways:
 - DBMS that allows users to update entries which improves answering database queries
- Goal is to:
 - Define precisely a class of problems that forms interesting forms of learning, explore algorithms to solve such problems, understand fundamental structure of learning problems and processes

Example Problem: Handwritten Digit Recognition



- Handcrafted rules will result in large no of rules and exceptions
- Better to have a machine that learns from a large training set

Wide variability of same numeral

Achievements in Machine Learning

- Pattern Recognition
 - Recognizing spoken words, handwritten words
- Data Mining
 - Predict recovery rates for pneumonia patients
 - Detect fraudulent use of credit cards
- Artificial Intelligence
 - Drive autonomous vehicles on highways
 - Backgammon: human world champion level
- Information retrieval
 - Search Engines
 - Text Classification



Some Successful Applications of Machine Learning

- Learning to recognize spoken words
 - Speaker-specific strategies for recognizing primitive sounds (phonemes) and words from speech signal
 - Neural networks and methods for learning HMMs for customizing to individual speakers, vocabularies and microphone characteristics

Some Successful Applications of Machine Learning

- Learning to drive an autonomous vehicle
 - Train computer-controlled vehicles to steer correctly
 - Drive at 70 mph for 90 miles on public highways
 - Associate steering commands with image sequences

ALVINN [Pomerleau] drives 70 mph on highways



Scientific Application of Machine Learning

- Learning to classify new astronomical structures
 - Very large databases to learn general regularities implicit in the data
 - Classify celestial objects from image data
 - Decision tree algorithms are now used by NASA to classify all objects in sky survey which consists of 3 terabytes of image data

Well-defined Learning Problem

- Identify three features
 - class of tasks
 - measure of performance to be improved
 - source of experience

Well-Posed Learning Problems

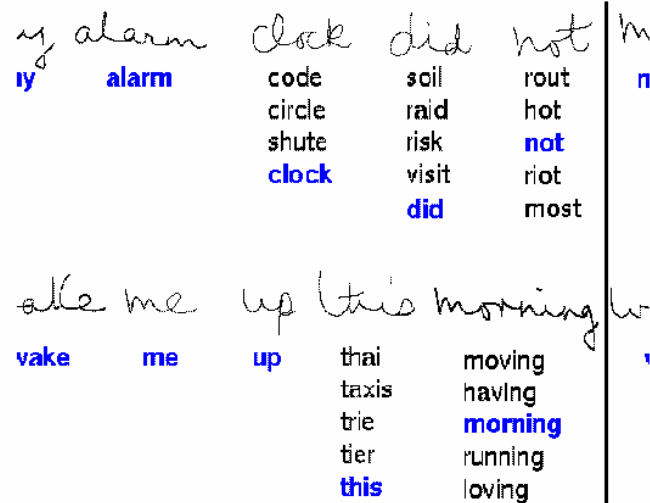
- A computer program is said to learn from *experience* E
- with respect to *class of tasks* T and *performance measure* P ,
- if its performance at tasks T , as measured by P , improves with experience E .

A Robot Driving Learning Problem

- **Task T**
 - driving on public, 4-lane highway using vision sensors
- **Performance measure P**
 - average distance traveled before an error (as judged by human overseer)
- **Training experience E**
 - a sequence of images and steering commands recorded while observing a human driver

Handwriting Recognition

- **Task T**
 - recognizing and classifying handwritten words within images
- **Performance measure P**
 - percent of words correctly classified
- **Training experience E**
 - a database of handwritten words with given classifications



Text Classification Problem

- **Task T**
 - assign a document to its content category
- **Performance measure P**
 - Precision and Recall
- **Training experience E**
 - Example pre-classified documents

The ML Approach

1. Data Collection

Large sample of data of how humans perform the task

2. Model Selection

Settle on a statistical model of the process

3. Parameter Estimation

Calculate parameter values by inspecting the data

Using learned model perform:

4. Search

Find optimal solution to given problem

Relationships to Other Fields

- Key algorithms and theory that form the core of machine learning draw heavily from concepts and results from:
 - Statistics
 - Artificial Intelligence
 - Philosophy
 - Information theory
 - Biology
 - Computational complexity
 - Control theory

Disciplines: Influence on Machine Learning

- Artificial intelligence
 - learning symbolic representations of concepts
 - machine learning as a search problem
 - learning as an approach to improving problem solving
- Pattern recognition
 - Bayes' theorem for calculating hypothesis probabilities
 - Naïve Bayes classifier

Disciplines: Influence on Machine Learning

- Computational complexity theory
 - Theoretical bounds on complexity of different learning tasks measured in terms of no of training samples, no of mistakes

Disciplines: Influence on Machine Learning

- Control theory
 - procedures to control processes to optimize predefined objectives and predict next state of process controlled
- Information theory
 - measures of entropy, information content
 - minimum description length approaches to learning
 - optimal codes and relationship to optimal training sequences for encoding hypothesis

Disciplines: Influence on Machine Learning

- Philosophy
 - Occam's razor--simplest hypothesis is best (decision trees)
 - justifying generalizing beyond observed data

Disciplines: Influence on Machine Learning

- Psychology and neurobiology
 - power law of practice-- human response improves with performance
- Statistics
 - characterization of errors (eg, bias and variance) that occur when estimating the accuracy of a hypothesis based on a limited sample of data
 - confidence intervals, statistical tests

Perspectives & Issues

- Machine Learning involves searching a very large space of possible hypotheses that fits observed data and any prior knowledge held by the observer

Issues in Machine Learning

- What algorithms exist?
- How much training data is sufficient?
- When and how can prior knowledge held by the learner guide the process of generalizing from examples?

Issues in Machine Learning

- What is the best strategy for choosing a useful next training experience?
- What is the best way to reduce the learning task to one or more function approximation systems?
- How can the learner automatically alter its representation to represent and learn the target function?

Summary

- Machine learning algorithms have great practical value in a variety of application domains
- Machine learning draws on ideas from a diverse set of disciplines
- A well-defined learning problem requires a well-specified task, performance metric, and source of experience
- Designing a machine learning approach involves a number of design choices
- Learning involves search

Student Background

- For undergraduate/graduate students in
 - Computer science and engineering
 - Statistics
 - Social sciences
- Two guidelines
 - Accessible to undergraduates
 - Material for PhD students before doing doctoral research in machine learning

The Material

- Balance of theory and practice
 - How does learning performance vary with number of samples
 - Which learning algorithms are appropriate for various learning tasks
 - Emphasis on Bayesian methods
 - Becoming more popular because of increased computational power and higher accuracy

Course Topics

1. Introduction
2. Probability Distributions
3. Linear Regression
4. Artificial Neural Networks
5. Kernel Methods: SVM
6. Graphical Models
7. Mixture Models and EM
8. Approximate Inference: Sampling
9. Sequential Data

Books

- Text Book:
 - Chris Bishop
 - Pattern Recognition and Machine Learning, 2006
 - Springer
- Reference Book:
 - Tom Mitchell
 - Machine Learning, 1997
 - WCB McGraw Hill