VICTORIA UNIVERSITY OF WELLINGTON
Te Whare Wananga o te Upoko o te Ika a Maui

School of Engineering and Computer Science

COMP 307 — Lecture 04

Machine Learning 1: Fundamentals

Mengjie Zhang (Meng)

mengjie.zhang@ecs.vuw.ac.nz

Outline

- Why Machine Learning
- What is machine learning about
- Types of machine learning
- Machine learning algorithms
- Training set vs test set
- Generalisation

Why Machine Learning

- Why do we (human) learn?
- To make smarter machines (systems)
  - improve performance, without (or with little) human intervention
  - robust behaviour in noisy environments
  - "learn about the world" in order to act sensibly
- To understand intelligence (by building it)

- Our COMP, ECEN, NWEN and SWEN people used the most)
  – COMP: IVC, graphics, WI, EC, prob ML
  – NWEN: Network opt., intrus. det., WIFI cover. energy opt.
  – SWEN: SW testing, defect prediction, Web serv. composition
  – ECEN: RoboCup, robot vision, mechatro., sig. processing

What is Machine Learning

- Machine learning is concerned with the design and development of algorithms and techniques that allow computers to "learn"
- "Machine learning is the study of computer algorithms that improve automatically through experience"
- Any system which changes itself
- Any system which improves its performance over time
- "Making sense of the world"
- "Finding patterns and commonalities in experience"
- others ... ...

Two Approaches

- Using machine learning to build/train intelligent agents
  - Building an expert system by training on preclassified examples
  - Building a voice recognition system by training on large data sets
  - Building a face detection system by training a recogniser
  - *Agent does not learn while working, learning can be very slow*
- Building agents that learn from experience and improve their performance
  - Voice recognition systems that get better with experience
  - Spam filtering system that learns from ongoing user feedback
  - Household robot that learns what the owners want
  - *Agent learns while working, learning must be fast*
- We will mainly focus on the concepts and principles

Inputs and Outputs of Learning Systems

- What is being learned (and how is it represented)?
  - Classifiers / Predictors
  - Concept descriptions
  - Models of the world
  - Rules for choosing actions
  - (Hidden) patterns / features
  - ... ...
- What is it learned from? (and how is it represented)?
  - Set of instances
  - Sequence of actions / events
  - Labeled / unlabeled / reward
  - Batch or incremental

Types of Learning Systems

- One helpful categorisation:
  - Supervised learning
  - Unsupervised learning
  - (Semi-supervised learning)
  - Reinforcement learning

Supervised Learning

- Given: instances of inputs and target outputs (labels)
- Generate: a function that maps inputs to desired outputs
- Predict: the correct output for a new (unseen) input
- Examples:
  - Learn rules for mortgage approval from records of past decisions
  - Learn to recognise words from speech of handwriting
  - Learn a description or rule for postal(zip) code recognition
  - Learn patterns/trends for predicting the stock market/weather/trafic ...
  - Learn patterns/features from fingerprint images for terrorist detection at airports
  - Learn a model/experienced formula from a sets of parameter values in real world application
- Most widely explored type of machine learning
- Many different approaches (we will focus)

Unsupervised and Reinforcement Learning

- **Unsupervised Learning"**
  – Given: set of unlabelled instances
  – Infer: subsets of "similar" instances (disjoint or overlapping)
  – Examples
    - Find clusters in high-dimensional data
    - Construct species hierarchy
    - Group results from search engine into categories to assist user to refine search
    - Identify groups of genes that have similar properties
- **Semi-supervised learning**: A mixture of supervised learning and unsupervised learning
- **Reinforcement Learning:**
  – Given: Sequence of actions and events, and reward/penalty occasionally
  – Infer: results for choosing best actions
  – Examples: Robot navigation tasks, Multiple lift controller, ...

Machine Learning Tasks

- Classification/Prediction: map data into predefined groups or classes (supervised learning, as above)
- Regression: map a data item to a real valued variable.
  — It involves the learning of the function that does this mapping. (supervised learning, more later)
- Clustering: unsupervised learning. Determine the similarity among the data on predefined attributes, and categorise the most similar data into groups.
- Association Rules (Link analysis)
  - identify specific types of data associations,
  - often used in the retail sales community to identify items that are frequently purchased together
  - Data mining. More in COMP422.

Main Learning Paradigms/Techniques

- Case based learning (or instance based learning): Use specific cases or experiences and rely on flexible matching methods to retrieve similar cases.
  – Example: *Nearest neighbour*
- Induction learning: Induce a general rule from a set of examples.
  –Example: *decision trees*
- Statistical (probability based) learning:
  - Naive Bayes (next lecture)
  - Support Vector Machines (briefly)
  - Bayesian Belief Networks (COMP421)
- Analytic learning systems: Represent knowledge as rules in logic form
  – Example: *Horn clauses*.

Learning Paradigms/Techniques (Continued)

- Connectionist learning: based on human brain behaviour
  - *artificial neural networks*
  - *parallel distributed processing systems*
  - (more later)
- Genetic/evolutionary learning: based on the mechanism of natural selection and natural genetics.
  - *Genetic algorithms*: evolve *bit strings* or *chromosomes*
  - *Genetic programming*: evolve computer programs
  - *PSO, EMO, LCS, ...*

Supervised Learning Systems

- Simple systems:
  - Representation: feature vectors
  - no missing values
  - no errors
  - sufficient features and sufficient examples
- Complex
  - Representation: multiple components and relationships
  - missing values
  - noisy data
  - limited examples

A Typical Supervised Learning System

- Presented with a set of training examples, some positive and some negative
- Need to come up with a rule/pattern that distinguishes the positive examples from the negative ones

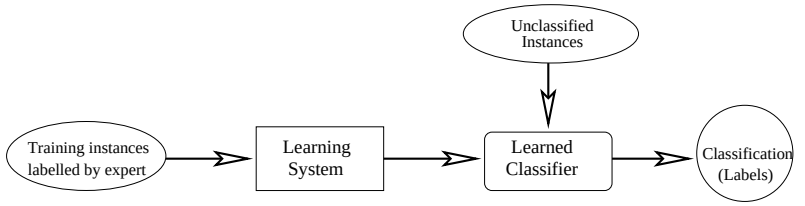

- Training Set: a collection of instances from which a classifier is induced.
- Test Set: A collection of instances which were never used for learning the classifier.
  – For measuring the performance of the learnt classifier.

Generalisation

- One of the major advantages of a learning system
- It refers to the ability to learn the useful patterns (e.g. classifiers) from the training data set and to apply the learned patterns to the test (unseen) data.
- The generalisation ability depends on how well the learning system has modelled the relationships in the training set.
- If the training set contains all the possible relationships between all the cases, then the learned program, once trained, should give good performance on the test data.
- There are two important issues here: *overtraining* or *overfitting*, and the *training set size*.

Summary

- Basic concepts of machine learning
- Categories of machine learning
- Common machine learning tasks
- Main machine learning paradigms/approaches
- Training set vs test set
- Generalisation
- Next lecture: Nearest neighbour and naive Bayes methods for classification
- Suggested reading: online materials and sections 20.4 and 20.2 (2nd edition) or sections 18.8 and 20.2 (3rd edition)