Can Machines Learn?

Pascal Poupart
Associate Professor
David R. Cheriton School of Computer Science
University of Waterloo

A Desktop

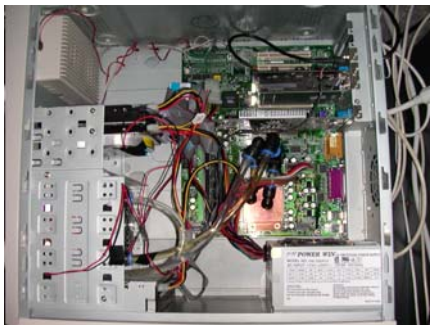

A Computer Program

```
import adbmstsys.adbc.*;

import java.lang.*;
import java.util.*;
import java.io.*;

public class Selections extends Thread {

    public Selections() {
        Object object = null;

    try {
            Mps3000 mps3000 = new Mps3000("172.24.64.78:30003:mgr/scopar.rysadmin");
            Database iosdb = new Database("mysid.adbctest.rysadmin",0,mps3000);
        iosdb.connect("RDB");
    // This method "dumps" the Selections dataset
                    TurboBuffer inv_location = null;
                    inv_location = iosdb.createTurboBuffer("Selections");
                    inv_location.setFetchSize( 1000 );
                    inv_location.setColumns("0");

                System.out.println("Selections Detail");
                System.out.println(" ");
// Passing NO parameter to the turbobuffer initiateRead() method will perform a serial read.
                    inv_location.initiateRead();
                    boolean notfound = true;
                    while ( inv_location.next() ) {
                System.out.println("Album Code =       " + inv_location.getString("AlbumCode"));
                System.out.println("Selection Name =   " + inv_location.getString("SelectionName"));
                System.out.println("Composer Name =    " + inv_location.getString("ComposerName"));
                System.out.println("Performer =        " + inv_location.getString("Performers"));
                System.out.println(" ");
                System.out.println(" ");
                    }
                } catch (ADBCException ex) { System.out.println("TurboIMAGE error " + ex.getMessage());
                                             ex.printStackTrace();}
    }
    public static void main( String[] args ) {
            Selections Selections = new Selections();
    }
}
```


Machine Learning

- Arthur Samuel (1959): Machine learning is the field of study that gives computers the ability to learn without being explicitly programmed.
- Tom Mitchell (1998): A computer program is said to **learn** from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.

Three categories

Supervised learning

Reinforcement learning

Unsupervised learning

Supervised Learning

- Example: digit recognition (postal code)

- Simplest approach: memorization

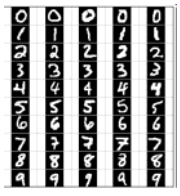

Supervised Learning

- Nearest neighbour:

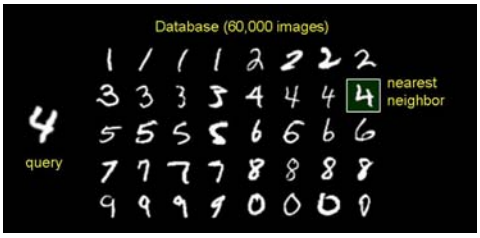

7

More Formally

- Inductive learning:
  - Given a training set of examples of the form (x,f(x))
    - x is the input, f(x) is the output
  - Return a function h that approximates f
    - h is called the hypothesis

8

Prediction

- Find function **h** that fits f at instances x

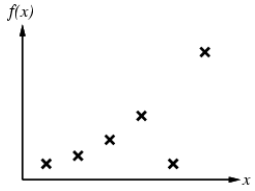

9

Prediction

- Find function **h** that fits f at instances x

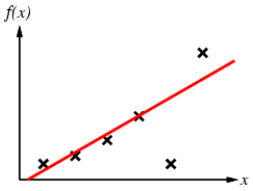

10

Prediction

- Find function **h** that fits f at instances x

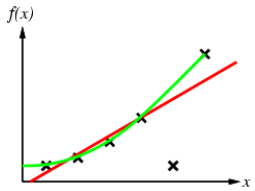

11

Prediction

- Find function **h** that fits f at instances x

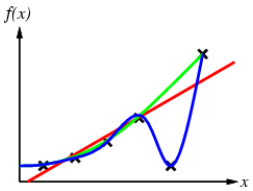

12

Prediction

- Find function **h** that fits f at instances x

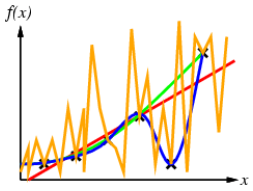

13

Generalization

- Key: a good hypothesis will generalize well (i.e. predict unseen examples correctly)
- Ockham's razor: prefer the simplest hypothesis consistent with data

14

Reinforcement Learning

- Differs from supervised learning

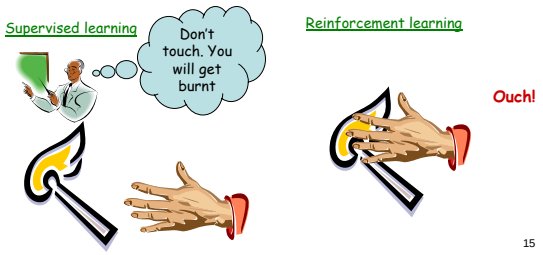

15

Animal Psychology

- Negative reinforcements:
  - Pain and hunger
- Positive reinforcements:
  - Pleasure and food
- Reinforcements used to train animals
- Let's do the same with computers!

16

Backgammon

- TD-Gammon:
  - Gerald Tesauro (1995)
  - Computer program
  - Best backgammon player!
- Play many games in simulation against itself
  - +1 for each win
  - -1 for each loss

- Optimization problem: find strategy that maximizes cumulative score

17

Helicopter Control

- Difficult to control:
  - Highly unstable

- Andrew Ng (Stanford, 2006):
  - Autonomous control by reinforcement learning
  - Step 1: learn neural net simulator based on flight data with human pilot
  - Step 2: optimize controller based on reinforcements for following a predefined trajectory

18

Applications of Machine Learning

- Speech recognition
  - dictation software
- Natural Language Processing
  - Text categorization
  - Information Retrieval
- Data Mining
  - Customer profiling
- Robotic Control
  - Mobile robots
  - Soccer playing robots

19

Vision

- Meta-programming: program computers to learn by themselves
- Lifelong machine learning: machines that continuously learn
- Transfer learning: machines that generalize their experience to new situations

- Challenges:
  - Computational complexity
  - Sample complexity

20

Thank You

Questions?

21