Can Machines Learn?

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

A Desktop

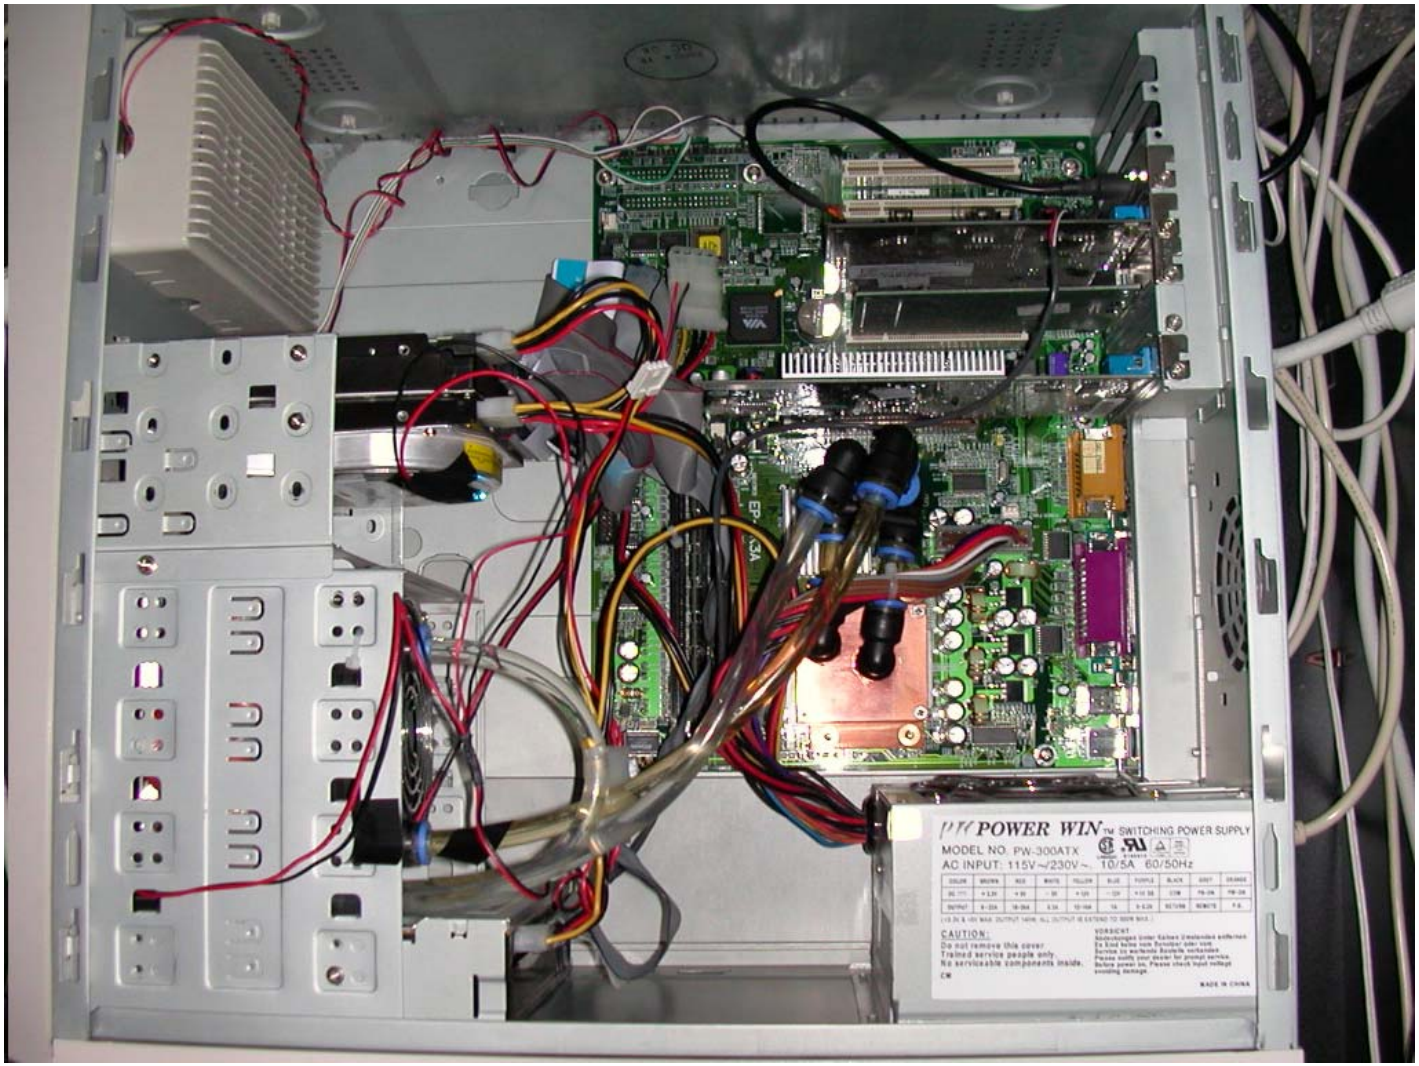
POWER WIN
SWITCHING POWER SUPPLY
MODEL NO. PW-300ATX
AC INPUT: 115V~/230V~
10/5A 60/50Hz
CAUTION:
Do not remove this cover.
Trained service people only.
No serviceable components inside.

A Computer Program

```
import advnetsys.adbc.*;

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

public class Selections extends Thread {

    public Selections() {
            Object object = null;

     try {
            Mpe3000 mpe3000 = new Mpe3000("172.24.64.78:30803:mgr/sooper.sysadmin");
              Database iosdb = new Database("music.adbctest.sysadmin",5,mpe3000);
        iosdb.connect("MGR");
     // This method "dumps" the Selections dataset
                    TurboBuffer inv_location = null;
                    inv_location = iosdb.createTurboBuffer("Selections");
                    inv_location.setFetchSize( 1000 );
                    inv_location.setColumns("@");

                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

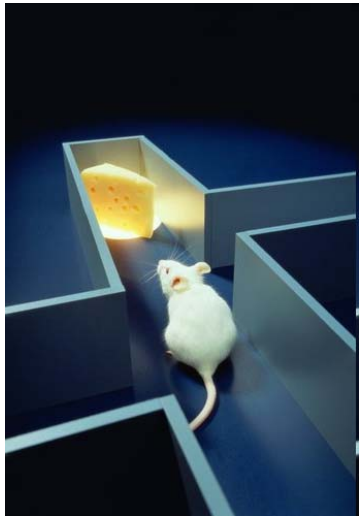

Unsupervised learning

Supervised Learning

- Example: digit recognition (postal code)

- Simplest approach: memorization

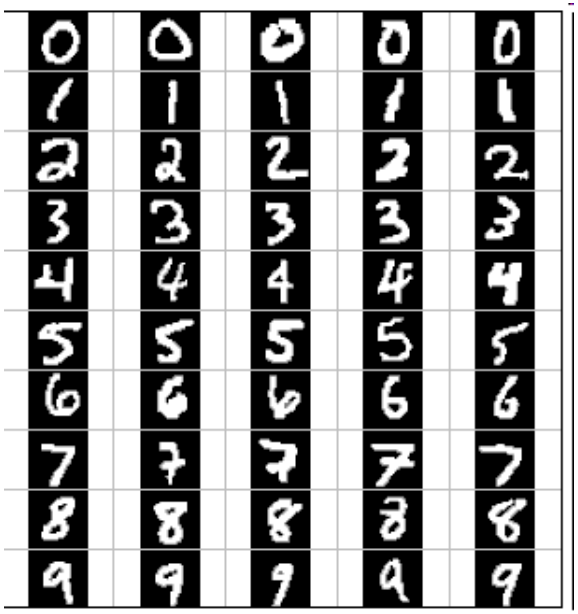

Supervised Learning

- Nearest neighbour:

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

Prediction

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

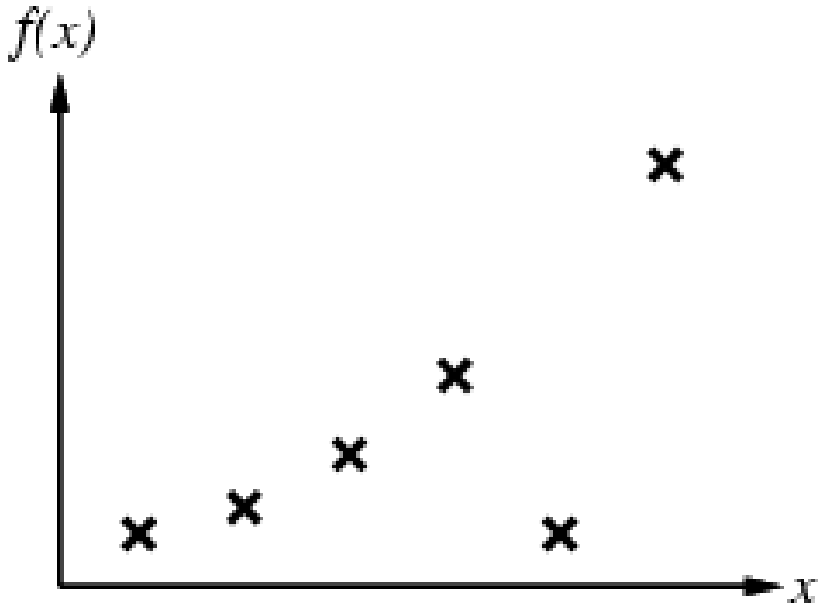

Prediction

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

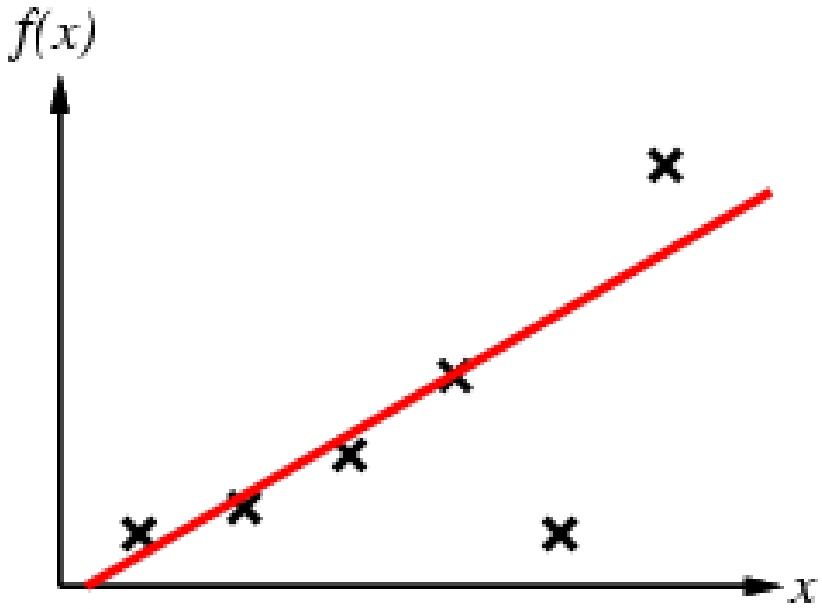

Prediction

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

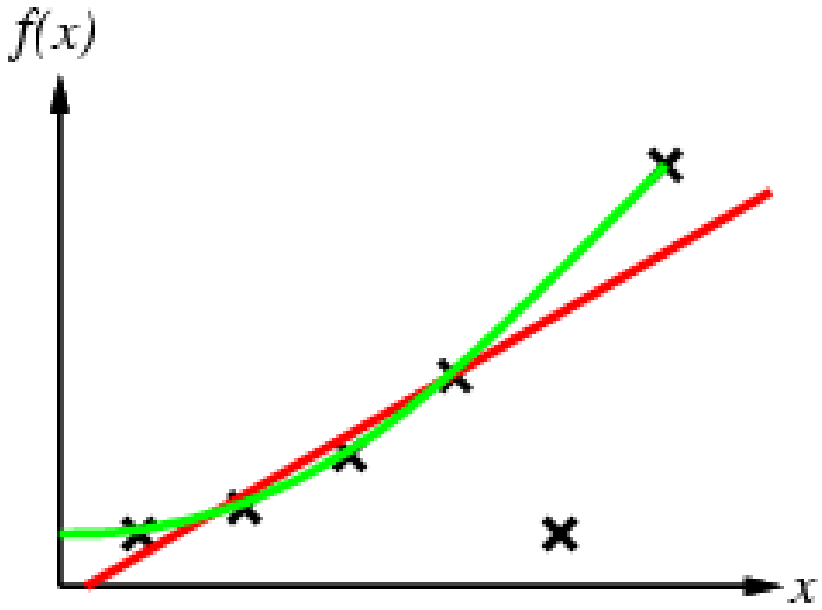

Prediction

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

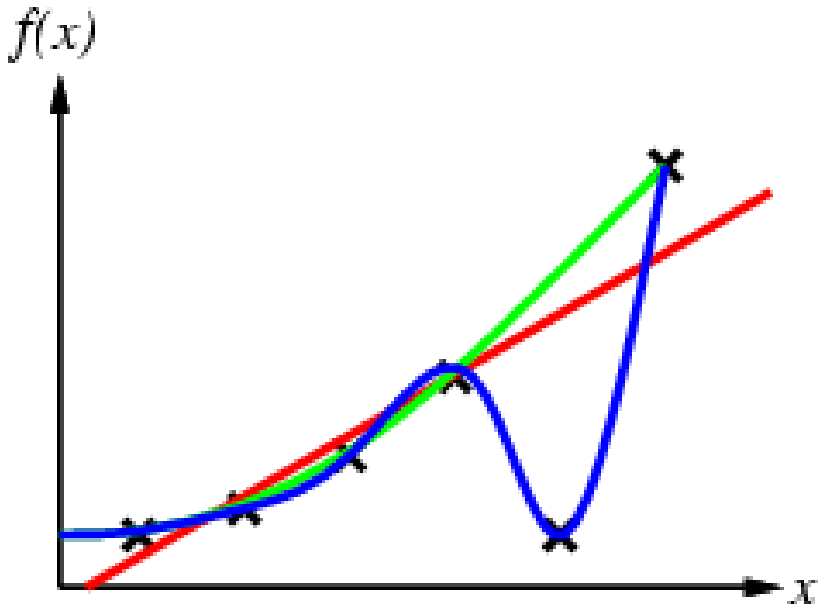

Prediction

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

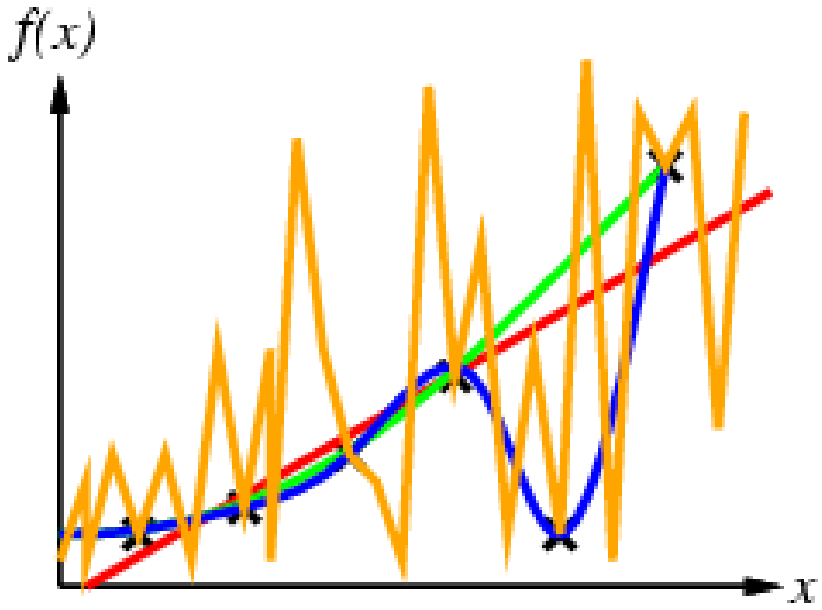

Generalization

- Key: a good hypothesis will generalize well (i.e. predict unseen examples correctly)

- Ockham's razor: prefer the simplest hypothesis consistent with data

Reinforcement Learning

- Differs from supervised learning

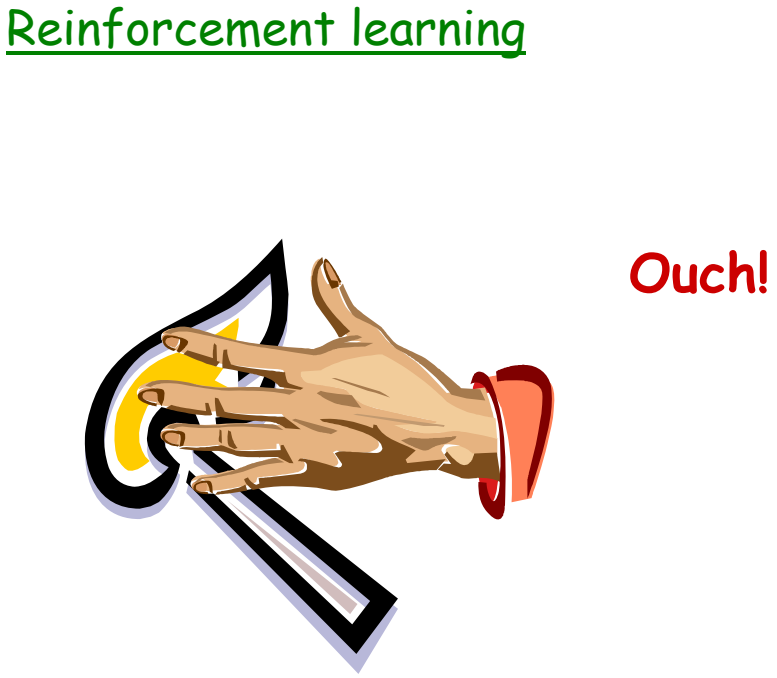

Animal Psychology

- Negative reinforcements:
  - Pain and hunger
- Positive reinforcements:
  - Pleasure and food
- Reinforcements used to train animals

- Let’s do the same with computers!

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

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

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

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

Thank You

Questions?