Introductory Applied Machine Learning

Nigel Goddard
School of Informatics

Semester 1

The **primary aim** of the course is to provide the student with a set of practical tools that can be applied to solve real-world problems in machine learning.

Machine learning is the study of computer algorithms that improve automatically through experience [Mitchell, 1997].

In many of today's problems it is

very hard to write a correct program

but very easy to collect examples

Idea behind machine learning:
from the examples, generate the program

Spam Classification

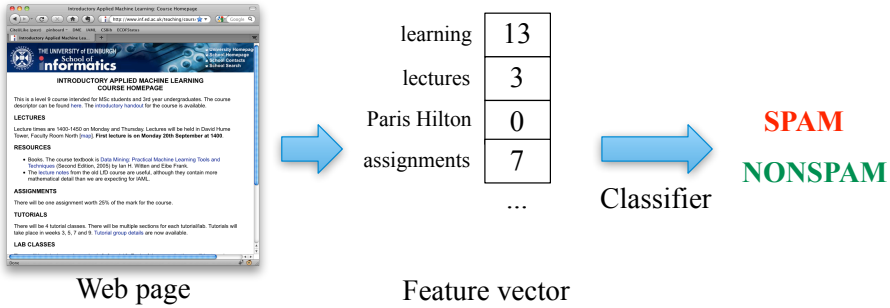

Image Processing

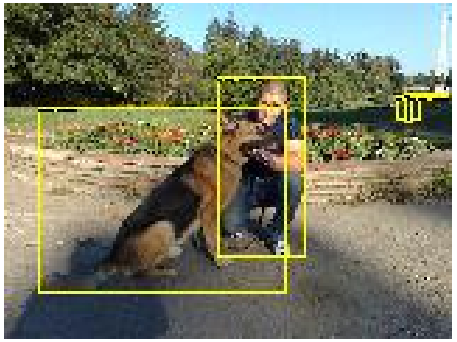

- Classification: Is there are dog in this image?
- Localization: If there is a dog in this image, draw its bounding box
- See: http://host.robots.ox.ac.uk/pascal/VOC/

Primate splice-junction gene sequences (DNA)

```
CCAGCTGCATCACAGGAGGCCAGCGAGCAGGTCTGTTCCAAGGGCCTTCGAGCCAGTCTG EI
GAGGTGAAGGACGTCCTTCCCCAGGAGCCGGTGAGAAGCGCAGTCGGGGGCACGGGGATG EI
TAAATTCTTCTGTTTGTTAACACCTTTCAGACTTATGTGTATGAAGGAGTAGAAGCCAAA IE
AAACTAAAGAATTATTCTTTTACATTTCAGTTTTTCTTGATCATGAAAACGCCAACAAAA IE
AAAGCAGATCAGCTGTATAAACAGAAAATTATTCGTGGTTTCTGTCACTTGTGTGATGGT N
TTGCCCTCAGCATCACCATGAACGGAGAGGCCATCGCCTGCGCTGAGGGCTGCCAGGCCA N
```

- Task is to predict if there is an IE (intron/exon), EI or N (neither) junction in the centre of the string
- Data from ML repository: http://archive.ics.uci.edu/ml/

Financial Modeling

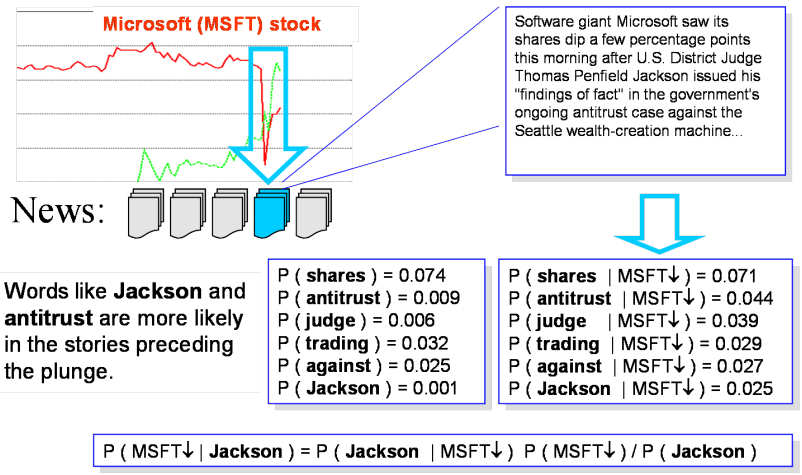

[Victor Lavrenko]

Collaborative Filtering

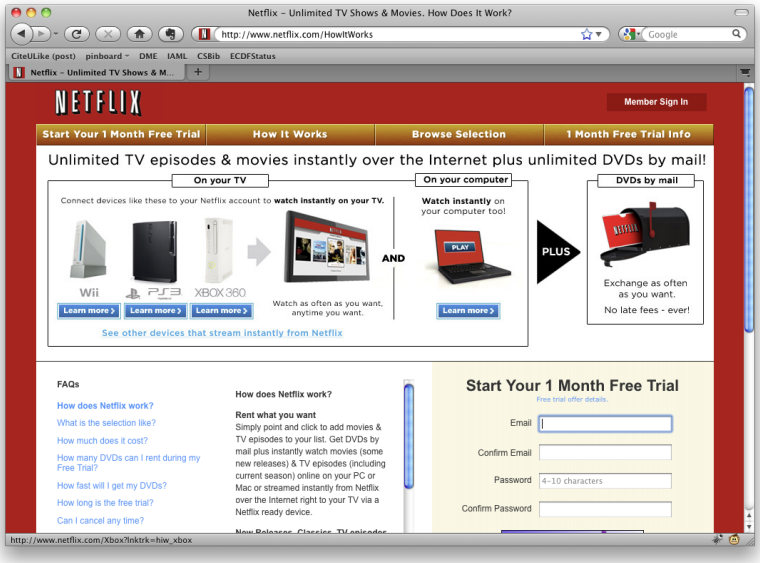

More applications

- Science (Astronomy, neuroscience, medical imaging, bio-informatics)
- Environment (energy, climate, weather, resources)
- Retail (Intelligent stock control, demographic store placement)
- Manufacturing (Intelligent control, automated monitoring, detection methods)
- Security (Intelligent smoke alarms, fraud detection)
- Marketing (targetting promotions, ...)
- Management (Scheduling, timetabling)
- Finance (credit scoring, risk analysis...)
- Web data (information retrieval, information extraction, ...)

Overview

- What is ML? Who uses it?
- Course structure / Assessment
- Relationships between ML courses
- Overview of Machine Learning
- Overview of the Course
- Maths Level
- Reading: W & F chapter 1

Acknowledgements: Thanks to Amos Storkey, David Barber, Chris Williams, Charles Sutton and Victor Lavrenko for permission to use course material from previous years. Additionally, inspiration has been obtained from Geoff Hinton's slides for CSC 2515 in Toronto

Administration

- Course text: Data Mining: Practical Machine Learning Tools and Techniques (Second/Third Edition, 2005/2011) by Ian H. Witten and Eibe Frank
- All material in course accessible to 3rd- & 4th-year undergraduates. Postgraduates also welcome.
- Lectures: 50% online, with quiz and review
- Assessment:
  - Assignments (2) (25% of mark)
  - Exam (75% of mark)
- 4 Tutorials and 4 Labs
- Course rep
- Plagiarism

`http://web.inf.ed.ac.uk/infweb/admin/policies/guidelines-plagiarism`

Machine Learning Courses

IAML Basic introductory course on supervised and unsupervised learning

MLPR More advanced course on machine learning, including coverage of Bayesian methods (Semester 2)

RL Reinforcement Learning.

MLP Real-world ML. This year: Deep Learning.

PMR Probabilistic modelling and reasoning. Focus on learning and inference for probabilistic models, e.g. probabilistic expert systems, latent variable models, Hidden Markov models

- Basically, IAML: Users of ML; MLPR: Developers of new ML techniques.

Overview of Machine Learning

- Supervised learning
  - Predict an output y when given an input $\mathbf{x}$
  - For categorical y : *classification*.
  - For real-valued y : *regression*.
- Unsupervised learning
  - Create an internal representation of the input, e.g. clustering, dimensionality
  - This is important in machine learning as getting labels is often difficult and expensive
- Other areas of ML
  - Learning to predict structured objects (e.g., graphs, trees)
  - Reinforcement learning (learning from "rewards")
  - Semi-supervised learning (combines supervised + unsupervised)
  - We will not cover these at all in the course

Supervised Learning (Classification)

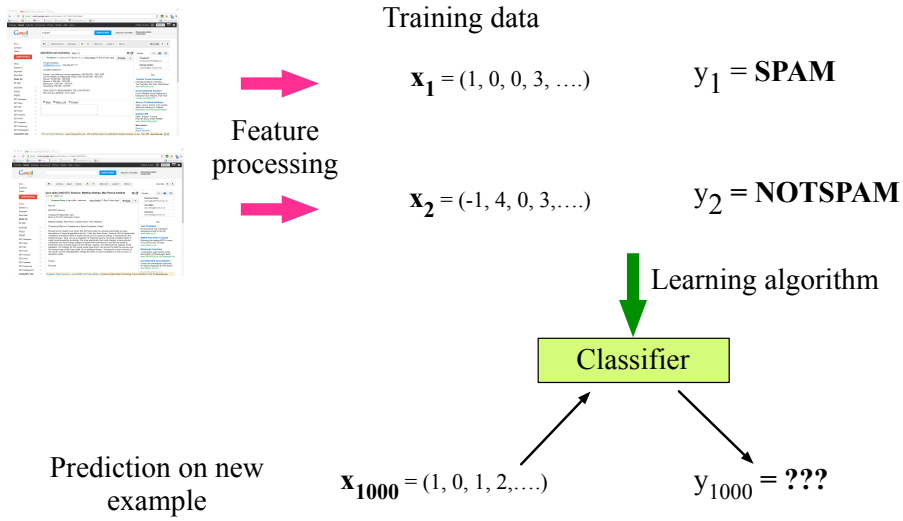

Supervised Learning (Regression)

In this course we will talk about linear regression

$$f(\mathbf{x}) = w_0 + w_1 x_1 + \ldots + w_D x_D$$

- $\mathbf{x} = (x_1, \ldots, x_D)^T$
- Here the assumption is that $f(\mathbf{x})$ is a linear function in $\mathbf{x}$
- The specific setting of the parameters $w_0, w_1, \ldots, w_D$ is done by minimizing a score function
- Usual score function is $\sum_{i=1}^{n}(y^i - f(\mathbf{x}^i))^2$ where the sum runs over all training cases
- Linear regression is discussed in W & F §4.6, and we will cover it later in the course

Unsupervised Learning

In this class we will focus on one kind of unsupervised learning, clustering.

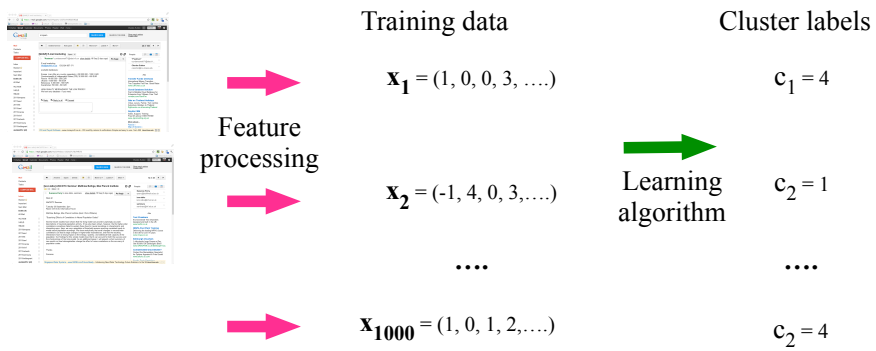

General structure of supervised learning algorithms

Hand, Mannila, Smyth (2001)

- Define the **task**
- Decide on the **model structure** (choice of inductive bias)
- Decide on the **score function** (judge quality of fitted model)
- Decide on **optimization/search method** to optimize the score function

Inductive bias

- Supervised learning is inductive, i.e. we make generalizations about the form of $f(\mathbf{x})$ based on instances $\mathcal{D}$
- Let $f(\mathbf{x}; L, \mathcal{D})$ be the function learned by algorithm L with data $\mathcal{D}$
- Learning is impossible without making assumptions about f !!

The futility of bias-free learning

1

0

???

The futility of bias-free learning

- *A learner that makes no a priori assumptions regarding the target concept has no rational basis for classifying any unseen examples* (Mitchell, 1997, p 42)
- The *inductive bias* of a learner is the set of prior assumptions that it makes (we will not define this formally)
- We will consider a number of different supervised learning methods in the IAML; these correspond to different inductive biases

Machine Learning and Statistics

- A lot of work in machine learning can be seen as a rediscovery of things that were known in statistics; but there are also flows in the other direction
- The emphasis is rather different. One difference is a focus on *prediction* in machine learning vs *interpretation* of the model in statistics
- Until recently, machine learning usually referred to tasks associated with artificial intelligence (AI) such as recognition, diagnosis, planning, robot control, prediction, etc. These provide rich and interesting tasks
- Today interesting machine learning tasks abound.
- Goals can be autonomous machine performance, or enabling humans to learn from data (data mining).

Provisional Course Outline

- Introduction (Lecture)
- Basic probability (Lecture)
- Thinking about data (Online/Quiz/Review)
- Naïve Bayes classification (Online/Quiz/Review)
- Decision trees (Online/Quiz/Review)
- Linear regression (Lecture)
- Generalization and Overfitting (Lecture)
- Linear classification: logistic regression, perceptrons (Lecture)
- Kernel classifiers: support vector machines (Lecture)
- Dimensionality reduction (PCA etc) (Online/Quiz/Review)
- Performance evaluation (Online/Quiz/Review)
- Clustering (k-means, hierarchical) (Online/Quiz/Review)

Maths Level

- Machine learning generally involves a significant number of mathematical ideas and a significant amount of mathematical manipulation
- IAML aims to keep the maths level to a minimum, explaining things more in terms of higher-level concepts, and developing understanding in a procedural way (e.g. how to program an algorithm)
- For those wanting to pursue research in any of the areas covered you will need courses like PMR, MLPR

Why Maths?

- IAML is focused on intuition and algorithms, not theory
- But sometimes you need mathematical notation to express the algorithms precisely and concisely
- e.g., We represent training instances via vectors ($\mathbf{x} \in \mathbb{R}^k$), and linear functions of them as matrices
- Your first-year courses covered this stuff
  - But unlike many Informatics courses, we actually use it!

Functions, logarithms and exponentials

- Defining functions.
- Variable change in functions.
- Evaluation of functions.
- Combination rules for exponentials and logarithms.
- Some properties of exponential and logarithm.

Vectors

- Scalar (dot, inner) product, transpose.
- Basis vectors, unit vectors, vector length.
- Orthogonality, gradient vector, planes and hyper-planes.

Matrices

- Matrix addition, multiplication
- Matrix inverse, determinant.
- Linear transformation of vectors
- Eigenvalues, eigenvectors, symmetric matrices.

Calculus

- General rules for differentiation of standard functions, product rule, function of function rule.
- Partial differentiation
- Definition of integration
- Integration of standard functions.

Probability and Statistics

We will go over these next time, but useful if you have seen these before.

- Probability, events
- Mean, variance, covariance
- Conditional probability
- Combination rules for probabilities
- Independence, conditional independence