

CS-XXX: Introduction to Machine Learning

Introduction

Arbish Akram

Department of Computer Science
Government College University

Slides by Dr. Nazar Khan

August 16, 2026

Preliminaries

- ▶ The ability of biological brains to sense, perceive, analyse and recognise patterns can only be described as stunning.
- ▶ They also have the ability to learn from new examples with or without being taught.
- ▶ Mankind's understanding of how biological brains operate exactly is embarrassingly limited.
- ▶ However, there do exist numerous *practical* techniques that give machines the *appearance* of being intelligent.
- ▶ This is the domain of statistical pattern recognition and machine learning.

Preliminaries

- ▶ Instead of attempting to mimic the complex workings of a biological brain, this course
 - ▶ aims at explaining mathematically well-founded techniques for analysing patterns and learning from them, and is therefore
 - ▶ a *mathematically involved* introduction into the field of pattern recognition and machine learning.
- ▶ It will prepare you for further study/research in
 - ▶ Pattern Recognition
 - ▶ Machine Learning
 - ▶ Computer Vision
 - ▶ Big Data Analytics


and others areas attempting to solve Artificial Intelligence (AI) type problems.

Prerequisites

- ▶ The course is designed to be self-contained.
- ▶ *Required mathematical details will be covered in the lectures.*
- ▶ However, this is a *math-heavy course*. Students are encouraged¹ to brush up on their knowledge of
 - ▶ probability and statistics
 - ▶ calculus (differentiation, partial derivatives)
 - ▶ linear algebra (vectors, matrices, dot-product, orthogonality, eigenvectors, SVD)
- ▶ The only way to benefit from this course is to be prepared to *spend lots of hours reading the text book and attempting its exercises* preferably alone or with a class-fellow.

¹ordered

Deep Learning vs. Machine Learning vs. AI



ML and AI problems are increasingly being solved using DL.

Laymen have started considering DL to be the same as AI.

AI: software that solves problems by itself.

ML: algorithms and models that *learn* from processed data.


DL: *neural networks* that *learn better* from *less processed* data.

Introduction

Machine Learning and Pattern Recognition are different names for essentially the same thing.


- ▶ Pattern Recognition arose out of Engineering.
- ▶ Machine Learning arose out of Computer Science.
- ▶ Both are concerned with automatic discovery of regularities in data.
- ▶ Regularity implies order. Learning implies exploiting order in order to make predictions.

Machine Learning




Supervised Learning

- ▶ **Classification:** Assign x to *discrete* categories.
 - ▶ Examples: Digit recognition, face recognition, *etc.*
- ▶ **Regression:** Find *continuous* values for x .
 - ▶ Examples: Price prediction, profit prediction.




Classification




Regression

Unsupervised Learning

- ▶ **Clustering:** Discover groups of similar examples.
- ▶ **Density Estimation:** Determine probability distribution of data.
- ▶ **Dimensionality Reduction:** Map data to a lower dimensional space.



Clustering




Density Estimation

Reinforcement Learning

- ▶ Find actions that maximise a reward. Example: chess playing program competing against a copy of itself.
- ▶ Active area of ML research.
- ▶ We will not be covering reinforcement learning in this course.

Classical Algorithms vs. Machine Learning

Problem: Given an image x of a digit, classify it between $0, 1, \dots, 9$.



Non-trivial due to high variability in hand-writing.

Classical Algorithms vs. Machine Learning

Classical Approach: Make hand-crafted rules or heuristics for distinguishing digits based on shapes of strokes.

Problems:

- ▶ Need lots of rules.
- ▶ Exceptions to rules and so on.
- ▶ Almost always gives poor results.

Classical Algorithms vs. Machine Learning

ML Approach:

- ▶ Collect a large *training set* $\mathbf{x}_1, \dots, \mathbf{x}_N$ of hand-written digits with known labels $t_1, \dots, t_N$.
- ▶ Learn/tune the parameters of an *adaptive* model.
 - ▶ The model can adapt so as to reproduce correct labels for all the training set images.

Classical Algorithms vs. Machine Learning

- ▶ Every sample x is mapped to $f(x)$.
- ▶ ML determines the mapping f during the *training phase*. Also called the *learning phase*.
- ▶ Trained model f is then used to label a new *test image* x_{test} as $f(x_{\text{test}})$.

Terminology

- ▶ *Generalization*: ability to correctly label *new* examples.
 - ▶ Very important because training data can only cover a tiny fraction of all possible examples in practical applications.
- ▶ *Pre-processing*: Transform data into a new space where solving the problem becomes
 - ▶ easier, and
 - ▶ faster.

Also called *feature extraction*. The extracted features should

- ▶ be quickly computable, and
- ▶ preserve useful discriminatory information.