CS-568 Deep Learning

Nazar Khan

PUCIT

Introduction to Deep Learning

Intelligence

Moravec's paradox[a]: The most difficult things to teach a computer are the ones that a two-year old has already learned – talking, listening, seeing, smelling, walking, grasping, memory and recall, thinking.

[a] https://en.wikipedia.org/wiki/Moravec%27s_paradox

- Adult humans have uniquely human attributes – metaphor, poetry, satire.
- Birds can fly through very small holes at full-speed *while fighting each other*.
- Blind bats can locate, identify and catch their *flying* prey by sending, receiving and analysing sound waves.

Intelligence

> 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 biological brains and how they operate exactly is embarrassingly limited.
- We are clueless regarding the most fundamental questions.
  - What is intelligence? Are you intelligent if you can't make a mistake?
  - Where in our brains does intelligence lie?
  - What is our brain?
  - Are our brains *just* computational devices or do they *do something more*?
  - What is consciousness?

The Brain

The average human brain has about **86 billion neurons** (or nerve cells) and **many more neuroglia** (or glial cells) which serve to support and protect the neurons [and perhaps even assist in their functionality]. Each neuron may be connected to up to **10,000 other neurons**, passing signals to each other via as many as **1,000 trillion synaptic connections**, equivalent by some estimates to a **computer with a 1 trillion bit per second processor**. Estimates of the human brain's memory capacity vary wildly from **1 to 1,000 terabytes** (for comparison, the 19 million volumes in the US Library of Congress represents about 10 terabytes of data). Source: `https://human-memory.net/brain-neurons-synapses/`

Notice the vague terminology – about, may be, up to, as many as, by some estimates, from 1 to 1000 terabytes.

So what is this course about?

- Modelling what we do not understand seems foolish.
- However, there do exist numerous *practical* techniques that give machines the *illusion of being intelligent*.
- This is the domain of artificial intelligence, statistical pattern recognition, machine learning and deep learning.
- 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 machine learning, computer vision, natural language processing 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 encouraged[1] to brush up on their knowledge of
  - probability and statistics
  - calculus (differentiation, partial derivatives, chain rule)
  - linear algebra (vectors, matrices, dot-product, orthogonality, eigenvectors, SVD)
- This is also a *code-heavy* course.

[1] ordered

Prerequisites

The only way to benefit from this course is to be prepared to *spend lots of hours reading the slides, text books, tutorials and attempting exercises* preferably alone or with a class-fellow.

Learn to be mostly alone for the next five months.
Learn to be (reasonably) selfish.
Your social life will be adversely affected. It should be! It's time to grow up.
You will need to work harder than ever before.
And even that might not be enough!

Administrative Stuff

Passing this course with at least a B+ grade is necessary for students planning to undertake research[a] in the CVML group.

[a] No space for PhD students at the moment

Course web-page:
`http://faculty.pucit.edu.pk/nazarkhan/teaching/Spring2020/CS568/CS568.html`

Texts and resources:

- *Deep Learning* by Ian Goodfellow and Yoshua Bengio and Aaron Courville `http://www.deeplearningbook.org/`
- *Pattern Recognition and Machine Learning* by Christopher M. Bishop (2006)
- `https://pytorch.org/tutorials/`

Administrative Stuff

Lectures:
Monday and Wednesday, 8:15 am – 9:40 am in AlKhwarizmi lecture theater.
Grading scheme:

Assignments	Quizzes	Mid-term	Final
20%	5%	35%	40%

- There will be no make-up for any missed quiz.
- Worst score on quizzes will be dropped.
- Worst score on assignments will be dropped.

Programming Environment

We will be using

- Python as our programming environment.
- PyTorch as our deep learning framework.
- Google Colaboratory as our GPU enabled machine on the cloud.
- Jupyter notebooks as our interactive tutorials.

Deep Learning vs. Machine Learning vs. AI

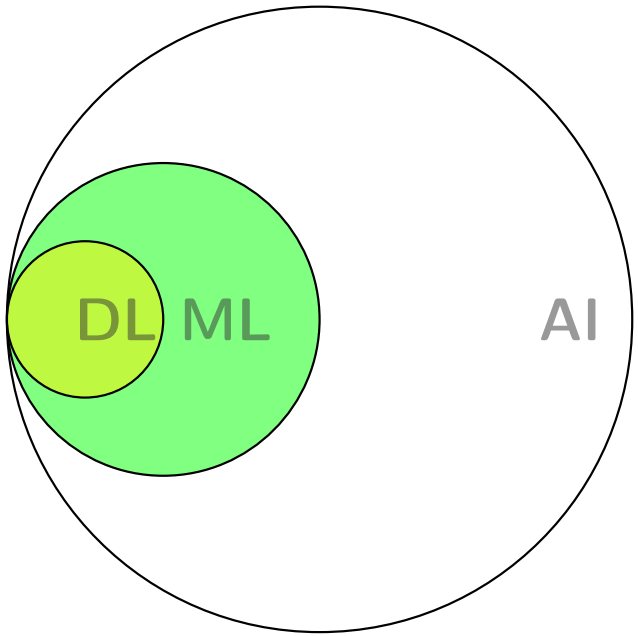

ML and AI problems are increasingly being solved using DL.
Laymen have started considering DL to be the same as AI.

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.

Gartner Hype Cycle

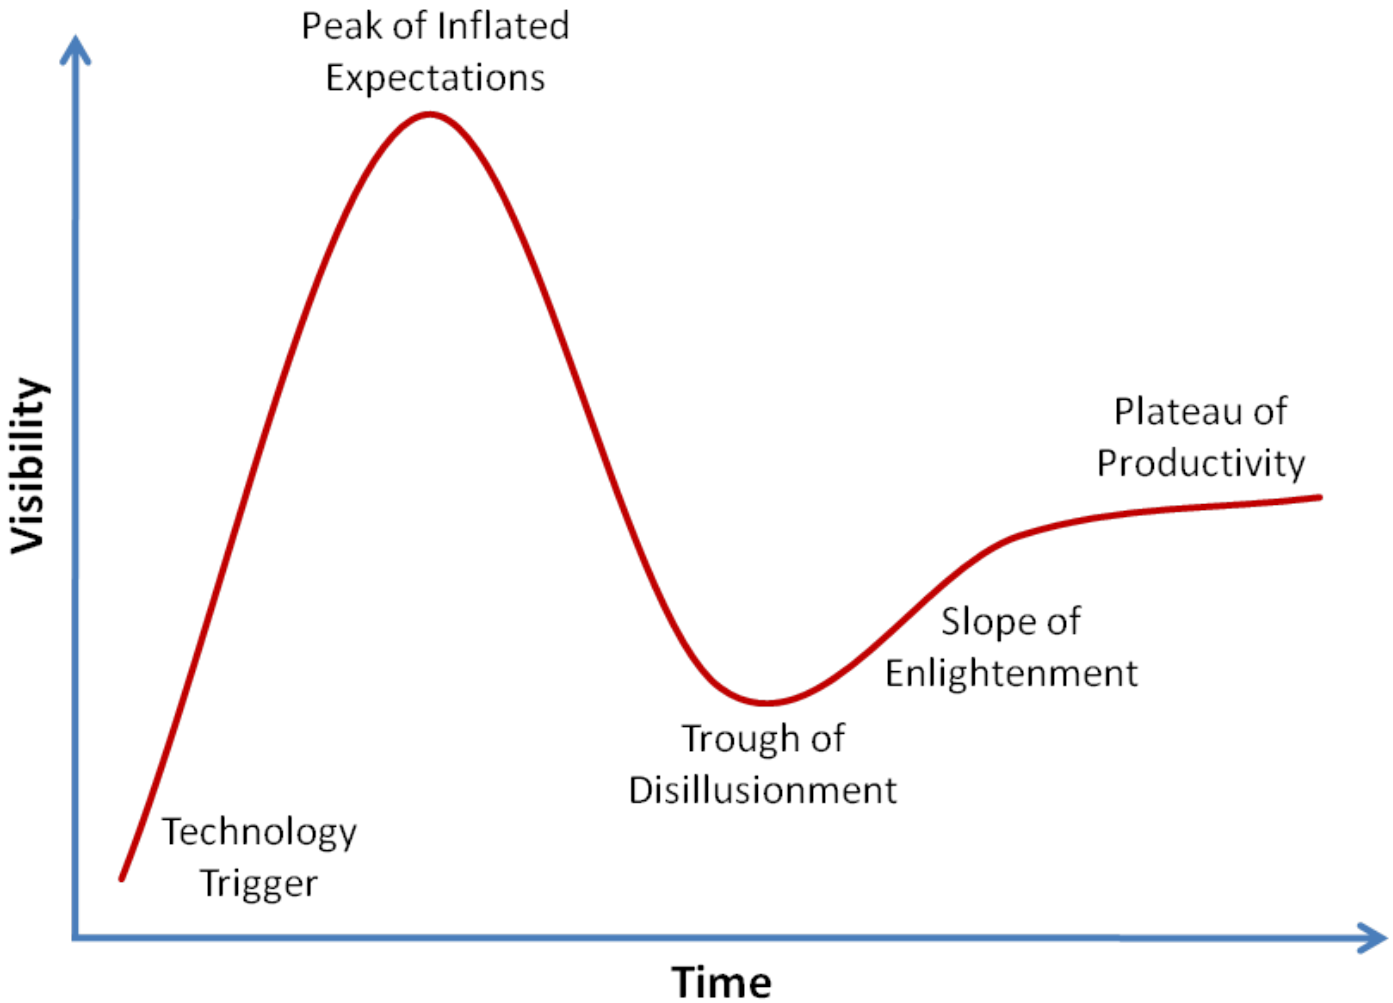

Gartner Hype Cycle for Artificial Intelligence, 2019

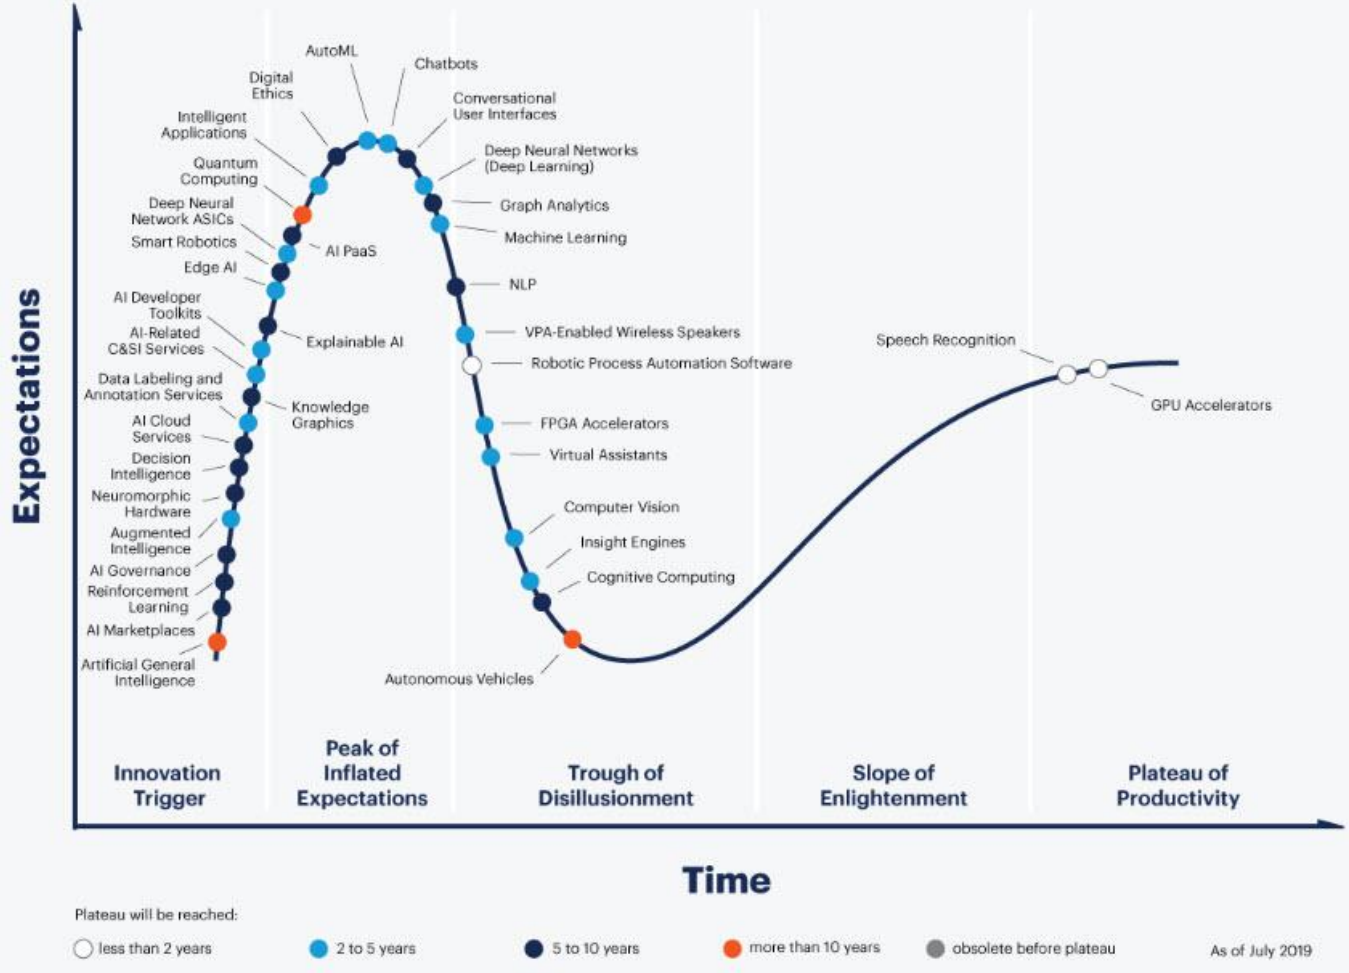

gartner.com/SmarterWithGartner

Source: Gartner

© 2019 Gartner, Inc. and/or its affiliates. All rights reserved.

Gartner

Introduction

Machine Learning and Pattern Recognition are different names for essentialy 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

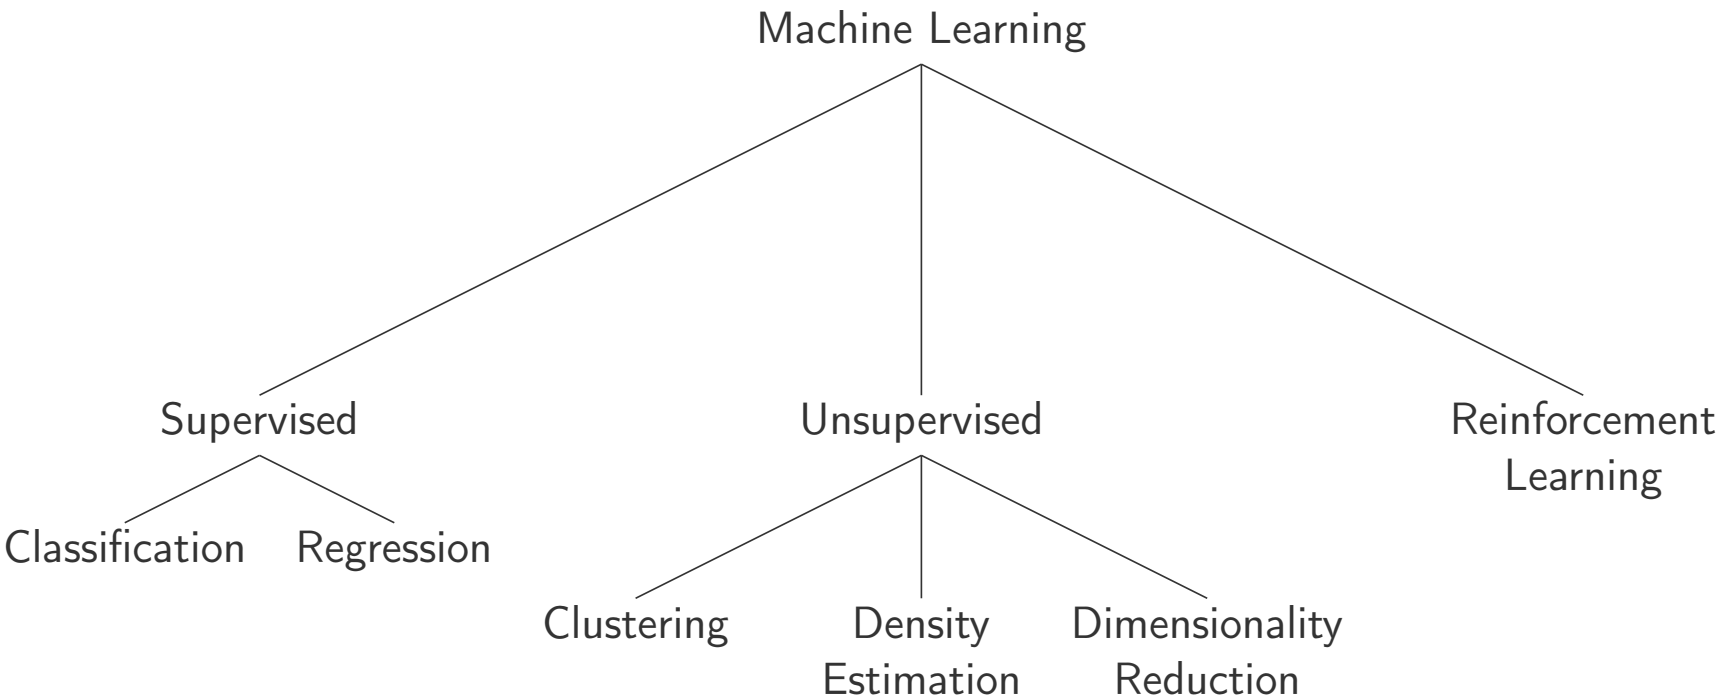

Supervised Learning

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

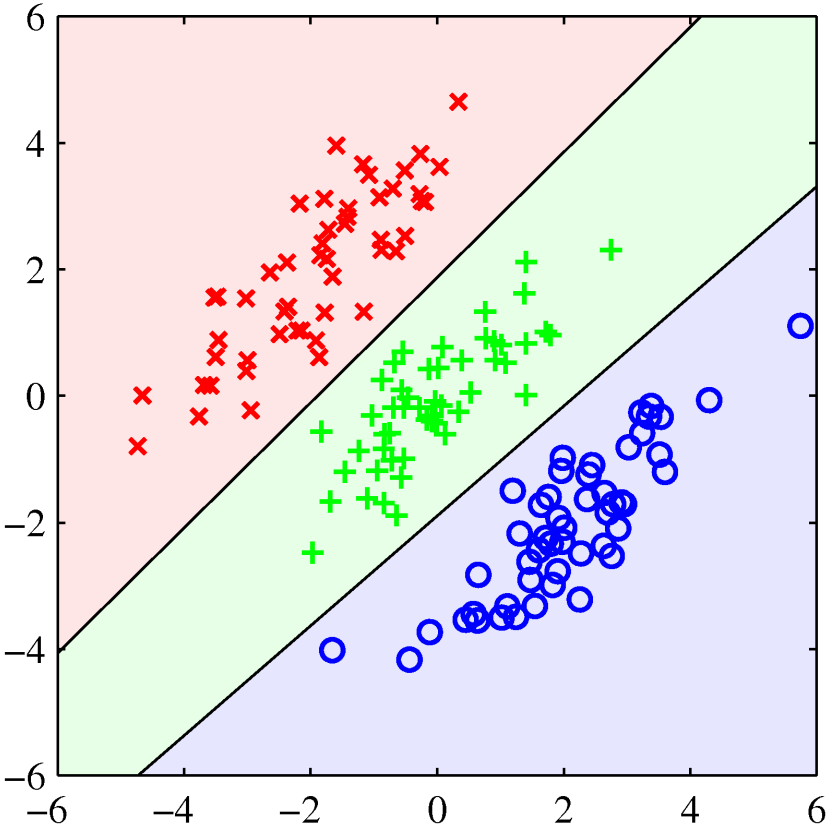

Classification

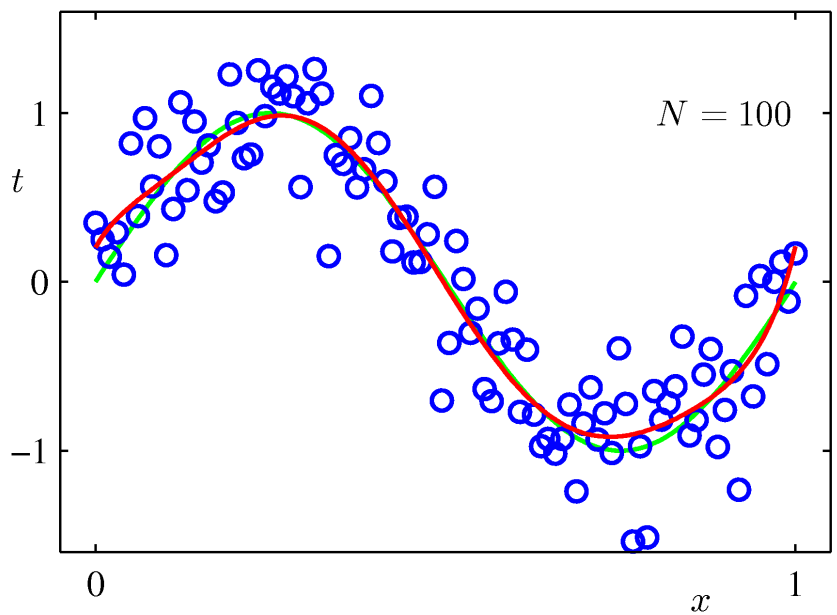

Regression

Unsupervised Learning

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

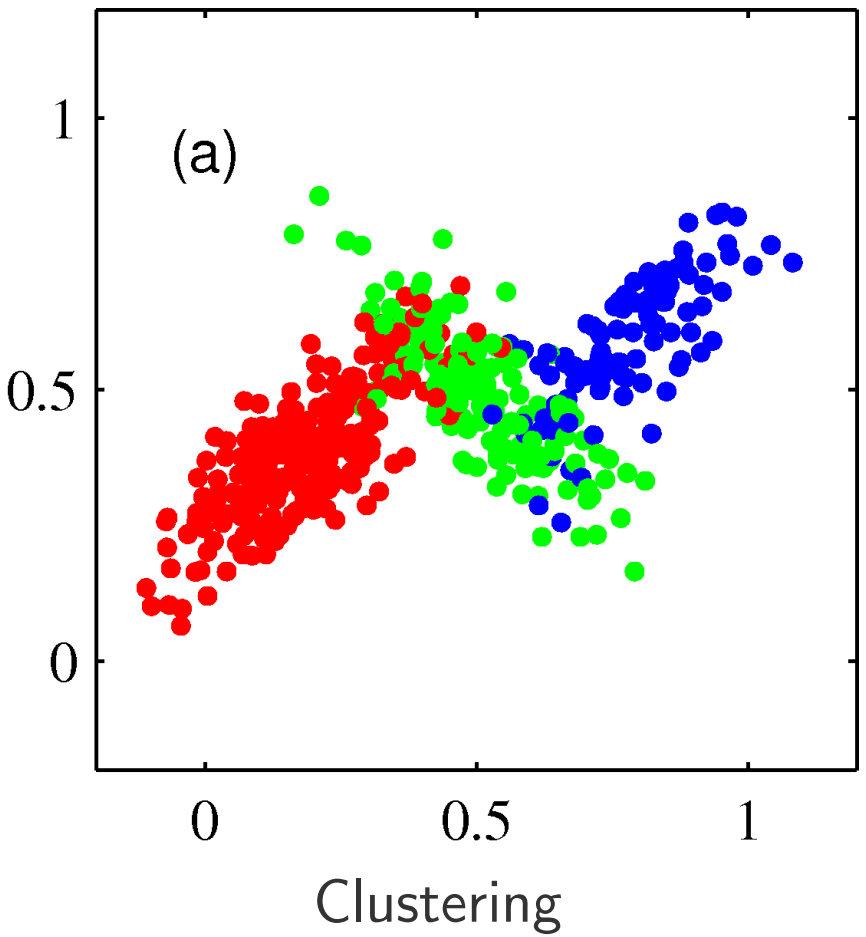

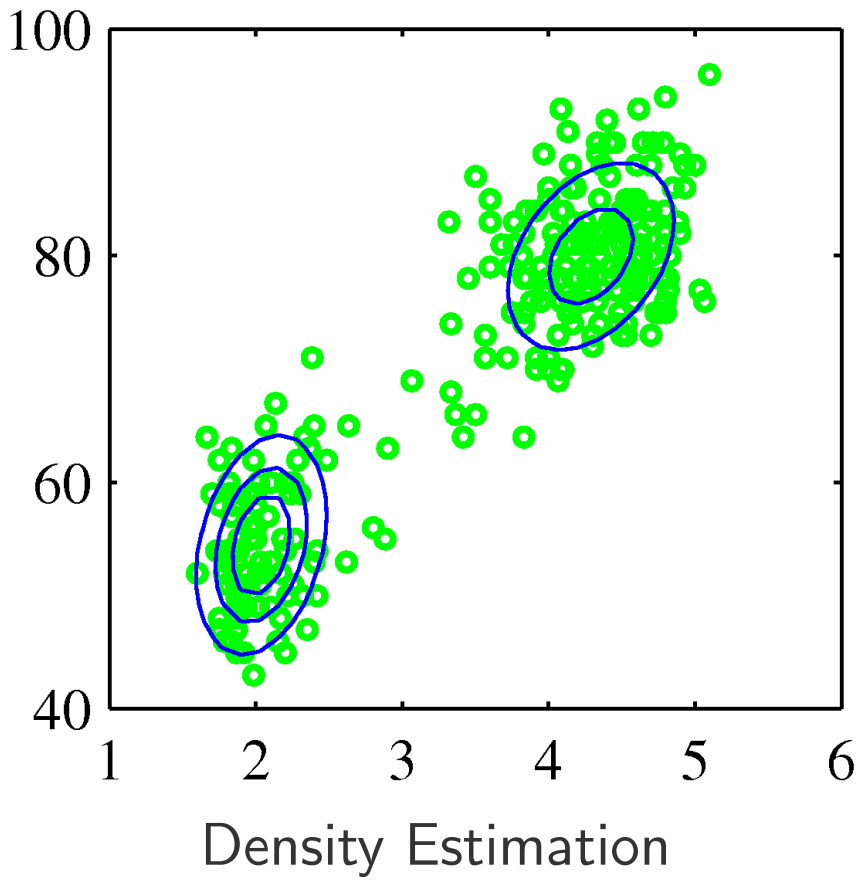

Reinforcement Learning

- Find actions that maximise a reward within an environment.

Figure: Based on the current state of the game (environment), each action of the player changes the state and yields a reward – points or death. The player learns to reinforce taking actions that lead to positive reward and not taking actions that lead to negative reward. Source: https://www.freecodecamp.org/news/a-brief-introduction-to-reinforcement-learning-7799af5840db/

Applications of Deep Learning

- Computer Vision
  - Face detection and recognition
  - Expression recognition and synthesis
  - Handwriting recognition
  - Image and video understanding
- Speech Recognition
- Natural Language Processing
  - Word similarity
  - Translation
  - Text generation