Introduction to machine learning

Herman Kamper

Let's write a spam filter

Simple first attempt in Python:

```
def check_spam(email):
    spam = False
    for word in email:



    return spam
```

More advanced attempt:

```
def check_spam(email):
    spam_score = 0
    for word in email:
        if word in ["inheritance", "win", "uncle", ...]:




    return spam
```

Machine learning

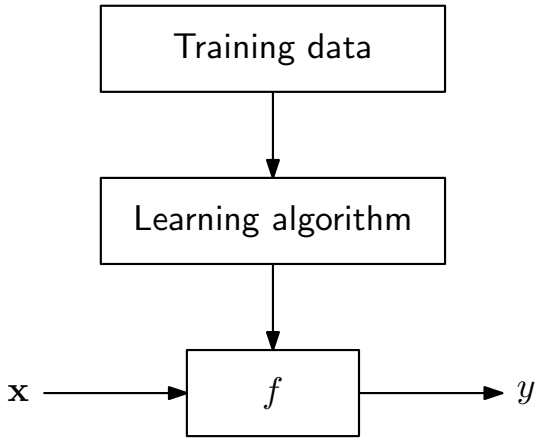

The model f is a function that maps input $\mathbf{x}$ to output y.

Inside the model f, we have parameters that we need to learn. I often denote the parameters as $\mathbf{w}$ or $\boldsymbol{\theta}$.

We fit the parameters through a learning algorithm that adjusts the parameters based on training data.

At test-time when we use the model, we hope it generalises: that it makes correct predictions for inputs $\mathbf{x}$ that are not exactly the same as those seen during training.

Two types of learning

Supervised learning[1]

- Involves learning a model that makes an output prediction given some input.
- Typically trained on $(\mathbf{x}, y)$ input-output pairs.
- E.g. a learned spam filter such as the one above, or a model that takes a chest X-ray and predicts whether a patient has tuberculosis or not.

Unsupervised learning

- Involves learning from unlabelled data, i.e. you only have inputs $\mathbf{x}$ without any corresponding output.
- E.g. determine whether there are groups (or clusters) of households with similar electricity usage behaviour; in this case the model only has access to electricity usage measurements without any output label.

[1] There are also other types of learning that you can look up on Wikipedia if you are interested: self-supervised learning (which is closely related to unsupervised learning), semi-supervised learning, and reinforcement learning.

Supervised learning: Classification

X-ray images of lungs:[2]

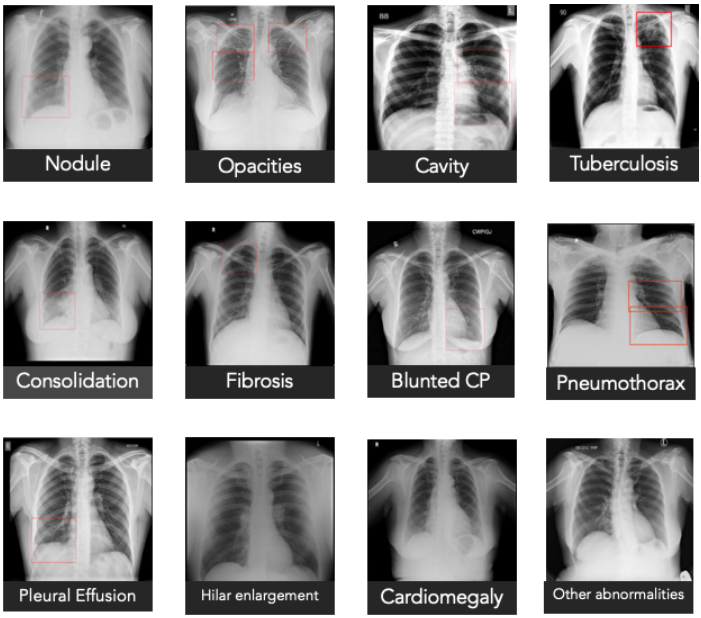

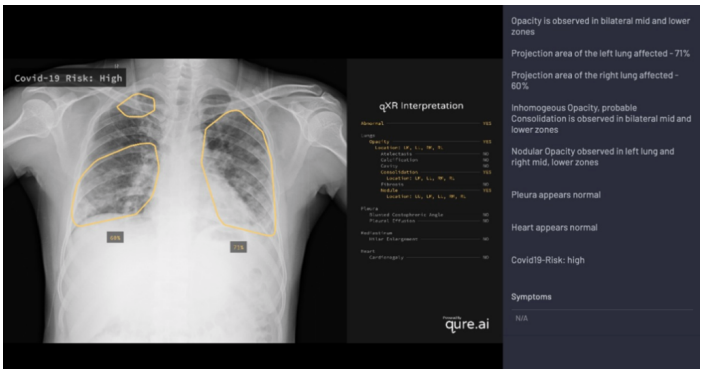

[2]Figures from MIT Technology Review.

Unsupervised learning: Clustering

Unsupervised learning: Clustering

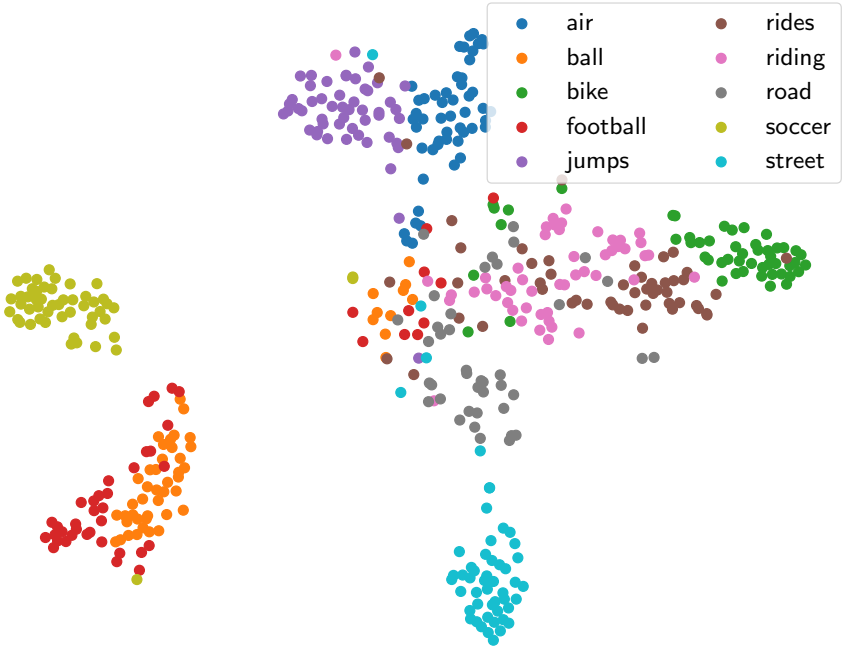

Unsupervised learning: Dimensionality reduction

Application of compression:[3]

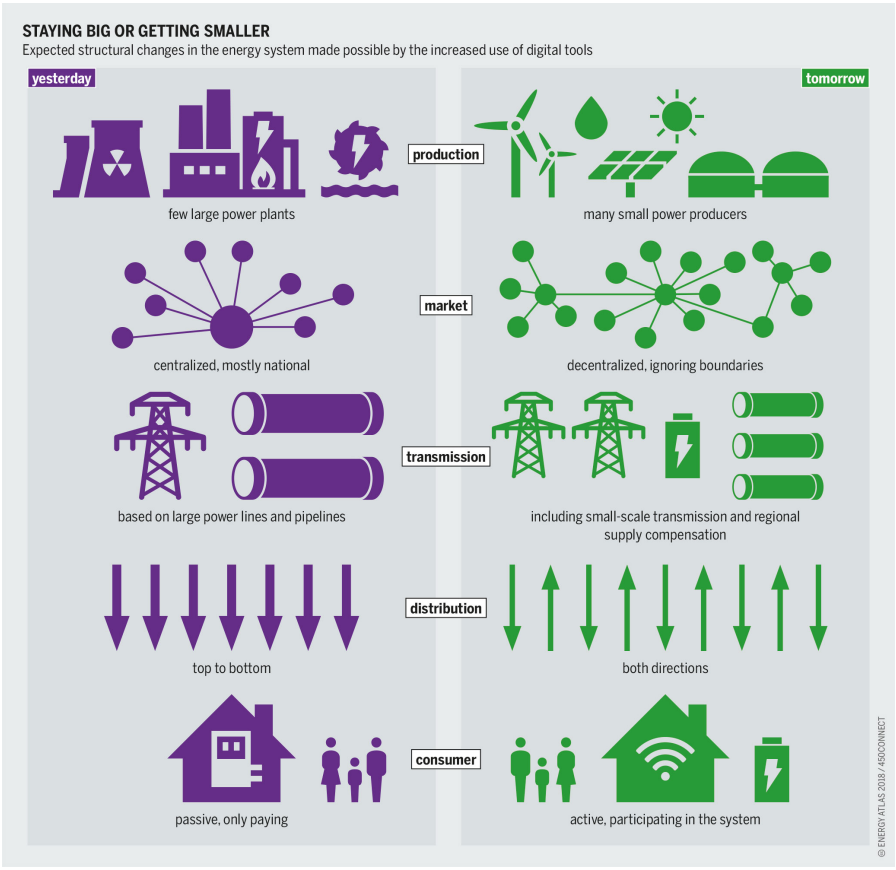

[3]Figure from Wikipedia.

Machine learning, statistics, data science/analytics

Machine learning

- Arthur Samuel, 1959: "Field of study that gives computers the ability to learn without being explicitly programmed."
- Tom Michell: "The field of Machine Learning seeks to answer the question 'How can we build computer systems that automatically improve with experience, and what are the fundamental laws that govern all learning processes?' "
- Yoshua Bengio: "Machine learning research is part of research on artificial intelligence, seeking to provide knowledge to computers through data, observations and interacting with the world. That acquired knowledge allows computers to correctly generalize to new settings."
- Kevin P. Murphy: "The goal of machine learning is to develop methods that can automatically detect patterns in data, and then to use the uncovered patterns to predict future data or other outcomes of interest."

Statistics

- "Analysis, interpretation and presentation of data." (Wikipedia)
- More emphasis on explaining than prediction: "What does the data say?"

Data science/analytics

- Combination of machine learning and statistics.
- But it also goes beyond that: It deals with the engineering infrastructure for acquiring, storing and (pre-)processing data.

DatA414 course information

Lectures, practicals and tutorials

Lectures:

- Monday 12:00 (A406)
- Tuesday 12:00 (A306)
- Wednesday 09:00 (A306)

Practicals and tutorials:

- Thursday 11:00 to 13:00
- Alternate between practicals and tutorials (approximately).
- Practicals: M2004
- Tutorials: A306
- Practical 1: Already out for this week (it is long).
- Practicals will be in Python (with NumPy and scikit-learn).

Office hour:

- Tuesday 14:00 to 15:00 (K509, MediaLab)
- If this moves I'll change the top of SUNLearn.
- I am very bad with emails :(

Lecture structure and what is examinable

- Lectures will be mostly on notes like this one.
- I will scribble on the notes, so bring them along (printed or in electronic form), or get a notebook.
- Will release notes at the start of each week on SUNLearn.
- There is a link to all of the previous years' videos on SUNLearn.
- The end of each note will link to the videos that were covered and the sections from ISLR (e.g. see the last page of this note).
- Everything from the sections in ISLR, the videos and the released notes will be examinable, except if I indicate otherwise (but I will be clear!).

Text book

Official DatA414 course text book (referred to as ISLR in the notes):

- G. James, D. Witten, T. Hastie and R. Tibshirani, *An Introduction to Statistical Learning*, 2nd ed., 2021.

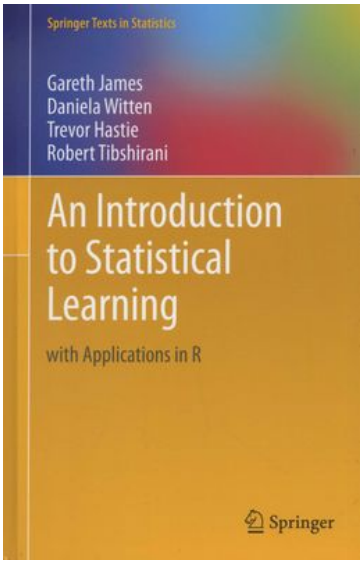

A number of other text books can also be helpful as references (but can be intense). I will not examine content from these:

- T. Hastie, R. Tibshirani and J. Friedman, *The Elements of Statistical Learning*, 2nd ed., 2017.
- C. Bishop, *Patten Recognition and Machine Learning*, 2006.
- K. P. Murphy, *Machine Learning: A Probabilistic Perspective*, 2012.

Assessments

- Final mark (flexible assessment):

$$\mathrm{FM} = 0.2 \cdot \mathrm{SM} + 0.3 \cdot \mathrm{A1} + 0.5 \cdot \mathrm{A2}$$

- Semester mark (SM) will be made up of assignments and small tests.
- Will announce small tests in advance (helps with learning).
- Read through the module framework on SUNLearn.

Videos covered in this note

- Intro to ML 1: What is machine learning? (10 min)
- Intro to ML 2: Supervised and unsupervised learning (6 min)

Reading

- ISLR 2.1.4