
Machine Learning for Quantum Computing Specialists

Daniel Goldsmith and M M Hassan Mahmud

Digital Catapult

February 2024

Abstract

Quantum machine learning (QML) is a promising early use case for quantum computing. There has been progress in the last five years from theoretical studies and numerical simulations to proof of concepts. Use cases demonstrated on contemporary quantum devices include classifying medical images [1] and items from the Iris dataset [2], classifying [3] and generating [4] handwritten images, toxicity screening [5], and learning a probability distribution [6]. Potential benefits of QML include faster training [2] and identification of feature maps not found classically [7]. Although, these examples lack the scale for commercial exploitation, and it may be several years before QML algorithms replace the classical solutions, QML is an exciting area.

This article is written for those who already have a sound knowledge of quantum computing and now wish to gain a basic overview of the terminology and some applications of classical machine learning ready to study quantum machine learning. The reader will already understand the relevant relevant linear algebra, including Hilbert spaces, a vector space with an inner product.

1 Introduction

This review is written for those who already some quantum computing knowledge and need to gain a basic overview of machine learning before learning about quantum machine learning (QML). There has been progress in the last five years from

theoretical studies and numerical simulations to proof of concepts. Use cases demonstrated on contemporary quantum devices include classifying medical images [1] and items from the Iris dataset [2], classifying [3] and generating [4] handwritten images, toxicity screening [5], and learning a probability distribution [6]. Potential benefits of QML include faster training [2] and identification of feature maps not found classically [7].

You may be encouraged to read that “a quantum feature map may enable a quantum kernel classifier on NISQ machines”, but left wondering what a “feature map”, a “kernel” and a “classifier” are. You may find it promising that “a possible early use case for quantum machine learning is to use quantum computers to generate the latent spaces for GANs” and be puzzled by the terms “GAN model”, and “latent space”. This review explores these and other important machine learning concepts at a high level. Key concepts are highlighted in **bold** when introduced. This is not an exhaustive text book on Machine Learning and you are signposted to additional sources of information to supplement this review. In the interests of brevity we pass over some popular machine learning algorithms like [Random forests](#) ¹.

This review was written by a quantum computing technologist who six months ago knew that Machine Learning could separate pictures of cats from pictures of dogs, but not much more. Since then I have been guided by my expert mentor, Dr M M Hassan Mahmud who has given me great advice and tuition, and has carefully reviewed this document. You are in safe hands!

1.1 Quantum Computing has given you a good start to machine learning

You may already understand most of the mathematics needed for machine learning from quantum computing. You probably understand the linear algebra on which much of machine learning is built, because the same linear algebra is used in quantum computing, albeit with complex numbers. Most likely you already understand a Hilbert space as a vector space equipped with an inner product because a quantum state can be represented as a vector in Hilbert space. The inner product is a key concept in machine learning, and you have probably already has seen that the inner product of two quantum states ψ and ϕ , $\langle\psi|\phi\rangle$, gives a measure of the overlap, or similarity between the two states. You will understand that a quantum state can be

¹In this article we adopt a hybrid approach of citing references because of the intended mode of dissemination as an online article. For well known concepts from computing, mathematics or quantum physics, we just use an in text hyperlink to an online resource. For books or research papers, we include proper citations.

represented as a vector, and an operation on a quantum state can be represented as a matrix that produces a new vector representing a different quantum state. Matrices and vectors are ubiquitous in machine learning. The symbols for matrices and vectors are shown in **bold** in this review.

You will understand that a matrix \mathbf{A} is Hermitian if $\mathbf{A} = \mathbf{A}^\dagger$, where \mathbf{A}^\dagger represents taking the transpose of the matrix, and then the complex conjugate of each matrix element. In machine learning the equivalent to the Hermitian of a matrix is a symmetric matrix, as complex numbers are not normally relevant. A matrix is symmetric if $\mathbf{A} = \mathbf{A}^T$ where \mathbf{A}^T represents the transpose of a matrix. We will see that covariance matrices are symmetric.

You will understand that a square matrix $\mathbf{A} \in \mathcal{R}^{n \times n}$ has eigenvalues $\lambda_i \in \mathcal{R}$ and corresponding eigenvectors $\mathbf{x}_i \in \mathcal{R}^n$ such that:

$$\mathbf{A}\mathbf{x}_i = \lambda_i\mathbf{x}_i \tag{1}$$

You will also have some knowledge of probability. If you would like to brush up on the underlying maths, then we recommend the on-line book [Mathematics for Machine Learning](#) [8]. In fact, add a few sections on complex numbers, and this book could be called “*Mathematics for Machine Learning and Quantum Computing*”.

You may also be familiar with [Python](#), and may have used [Jupyter](#) notebooks. Many Machine learning libraries such as [Pytorch](#) and [Tensorflow](#) use Python. If you have not worked with Python a good introductory book is “*Python Crash Course*” [9]. If you install [Anaconda](#) then Python and Jupyter will be downloaded as part of the installation. Jupyter notebooks are straightforward to pick up. Python code can be written directly into code cells, and text, including mathematical formulae, can be written in [Markdown](#) cells. The code can be changed directly on your computer to gain insight into the example.

This review provides links to Jupyter notebooks hosted on [the GitHub repository for our quantum technology access programme \(QTAP\)](#) [10] showcasing the comprehensive and easy to learn [Scikit Learn](#) machine learning package. Machine Learning is a practical subject and it is really useful to work on real examples. [Kaggle](#) has over 50,000 public datasets available for free that you can practice on.

1.2 Quantum machine learning - some caveats

Most likely you are reading this review because you are interested in Quantum Machine Learning (QML). We will point out use cases for QML. Be aware though, that is not a complete review of the current status of QML. The examples illustrate the text and point to additional references. Hopefully this will whet your appetite for more.

When reading the literature on QML you need to be aware of the following caveats:

1. We are in the age of [Noisy Intermediate Scale Quantum devices](#) (NISQ) with high error rates and few qubits.

2. In theory these errors can be corrected using **logical** qubits where a quantum state is held across many **physical qubits** and a corrupted physical qubit can be detected and corrected by referring to the other uncorrupted physical qubits, but there are significant challenges to using quantum error correction for useful computations.

3. At some stage in the future the expectation is that we will have the **fault tolerant** quantum computers with large numbers of error corrected logical qubits predicted by Preskill [11]. Increasing the number of physical qubits and getting error correction working properly will take at least a decade, and probably much longer.

4. There is no commercially available [quantum memory](#) to store quantum states, meaning that it is difficult to carry out quantum calculations that take longer than the short qubit decoherence times. Without quantum memory, converting machine learning data into quantum states will be very difficult on today's noisy devices.

5. Many quantum computing algorithms, such as Grover's search algorithm [12], the Harrow Hassidim Lloyd (HHL) [13] algorithm for matrix inversion and Quantum Phase Estimation [14], that might be useful for QML, are not feasible until fault tolerant quantum computers become widely available.

2 What is Artificial Intelligence?

Artificial Intelligence (AI) dates back to the earliest days of computing. AI can be defined as “*the set of technologies that solve problems that require human like intelligence*”. What is regarded as “*human like intelligence*” ratchets up over time, and as we will see, much deployed AI is now just regarded as software. So the unofficial definition of AI is “*anything that humans can do but computers cannot do ... yet!*”.

The book [Artificial Intelligence, A Modern Approach](#) [15] gives a good overview of AI. Our colleague at Digital Catapult, Dr Robert Elliot Smith, has written an important book “[Rage Inside The Machine](#)” [16]. As well as highlighting the involvement of AI algorithms in driving intolerance, Rob Smith puts AI and ML into their historical context, which I found valuable. Another very useful book is “[Machine Learning with Quantum Computers](#)” by Maria Schuld and Francesco Petruccione [17]. This book covers both machine learning and quantum computing, and looks at quantum machine learning algorithms likely to be useful in the future. Maria Schuld has posted insightful videos on YouTube, for example, [Is quantum advantage the right goal for QML?](#)

I also found the book “[Machine Learning for Absolute Beginners](#)” [18] by Oliver Theobald useful as it gave a good basic overview of ML.

AI has its roots in computer science and draws from diverse fields including logic, probability theory, statistics, neuroscience, information theory, cognitive science, linguistics and even biology.

Machine Learning is a subset of AI, and machine learning algorithms have the ability to learn from patterns in data without being explicitly programmed. There is a useful MIT OpenCourseWare Introduction to Machine Learning available on [YouTube](#).

A further subset of machine learning is **Deep Learning**, in which **artificial neural networks** adapt to learn from vast amounts of data. Deep learning methods have had tremendous success in the real world in recent years, and a lot of the current excitement around AI/ML can be attributed to their success. There is another great MIT OpenCourseWare Introduction to Deep Learning also available on [YouTube](#).

The relationship between the different components of Artificial Intelligence is shown in [Figure 1](#):

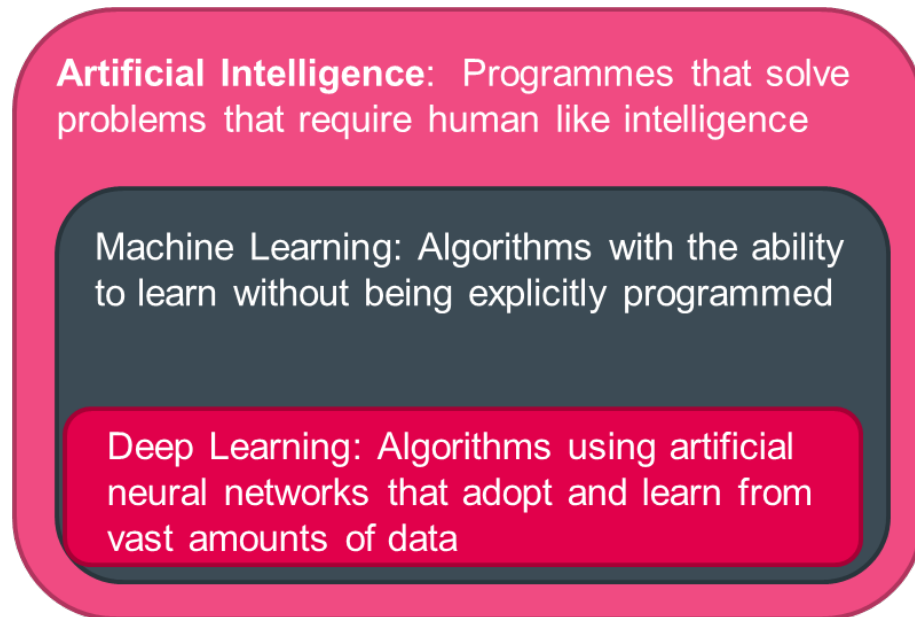


Figure 1: How Artificial Intelligence, Machine Learning and Deep Learning fit together

2.1 My first experience of AI - the VAT “expert system”

When I was working as an auditor in the early nineties we used an AI VAT “**expert system**” to give our clients VAT advice. We interviewed the client, and asked pre-defined questions like “do you use cars for business purposes?”, “what is your turnover?”, input the answers into our COMPAQ microcomputer shown in Figure 2 and printed out the results on a dot matrix printer to give to the client.

With hindsight our “expert system” suffered from all the issues of early expert systems identified by Rob Smith in his book [16]. It needed a lot of time from very expensive VAT and computing experts to design and build. It was fragile, and quite often the carefully studied output simply advised the client to consult with human experts for advice. To our employer, this may have been a feature rather than a bug, particularly if the client paid for extra VAT consultancy. However, we felt that we were on the leading edge of advanced digital technology, as we refined our IT skills playing golf on the microcomputer after working hours.

This “expert system” needed to be explicitly programmed for each and every situation it would encounter. In the next example we will consider a simple **machine learning model** that **learns** from a pattern in the data to predict values for unseen data.



Figure 2: The COMPAQ microcomputer - Thanks to the Norsk Tekisk Museum CC BY-SA 3.0

A **machine learning model** consists of an algorithm which has been optimised on training data to make useful predictions of unseen data.

2.2 An introductory example - linear regression

I have placed a [Jupyter Notebook](#) on GitHub demonstrating an example of linear regression using the Scikit Learn package. This is an example of **supervised learning**, where the algorithm predicts values based on examples provided. Although relatively straightforward, this example illustrates many of the points relevant for more complex machine learning models. This algorithm pre-dates computers, since for small data sets the calculations can be done by hand.

2.2.1 Aim of the linear regression example

The aim of this basic machine learning example is to predict the company turnover from the number of employees.

2.2.2 Loading the data using pandas

In my Notebook [Pandas](#), a Python data analysis tool, is used to load and process the dataset: a list of UK quantum technology companies with the company turnover, numbers of employees and latitudes and longitudes. The dataset is loosely based on real UK data. A plot of employees against turnover from the dataset is shown in [Figure 3](#):




Figure 3: Plot of employee data against company turnover: the raw data

2.2.3 Machine Learning terminology used to describe the data

The company turnover, numbers of employees, latitude and longitude are called **features**. Each line of the file is called a **sample**. Each sample can be regarded as a 4 dimensional vector, with one dimension for each feature, allowing concepts from linear algebra to be used. The set of all feature vectors is known as the **design matrix**. In this example the latitudes and longitudes are not used. The company turnover is a **label** in this example and the number of employees is an **input**. In any data analysis “Garbage In, Garbage Out” applies, so the data is first listed, plotted on a graph to check for outliers and any records with unavailable data are removed before proceeding. Whilst data analysis, and [data science](#) more generally, are very important to machine learning in practice, it is too large a subject to be considered further in this review.

2.2.4 Training and test data

Next, the data is split into two components, **training data** which makes up 70% of the data, and **test data**, which makes up the remaining 30%. The **model** is **trained** with the training data, and then makes predictions which are validated against the **test data**, showing how good the model is at **generalising** from the training data to the test data, a key feature of all models.

More complicated models need [validation data](#), as well as [training and test data](#).

2.2.5 The hypothesis space

Every machine learning problem has a family of candidate solutions chosen by the ML engineer. Each family of solutions, for example, polynomial functions, trigonometric functions and linear functions, is called a **hypothesis space**. The set of hypothesis spaces are limited only by our imagination. The aim of a particular ML algorithm is to use the training data to find the best solution in a given hypothesis space.

There is no way of being sure which is the best hypothesis space and algorithm to use. In fact, the “**no free lunch**” theorem implies that there is no single best machine learning algorithm. All algorithms work well in some cases and badly in other.

In this example, if we write x as the number of employees and \hat{y} as the predicted turnover, then the hypothesis spaces are different families of **predictor functions** that returns a prediction $\hat{y} = f(x)$ from an input x .

2.2.6 The inductive bias

To get a tractable ML algorithm it is necessary to make specific assumptions about the form of the mapping from inputs to labels as shown in Figure 4. These assumptions are called the **inductive bias**. The hypothesis space encodes these assumptions mathematically. Just as we all have **sub-conscious biases** which means that we may make incorrect decisions without realising it, we may also incorrectly tie our models to a mistaken assumption.

In this example it is assumed that the predictor function is linear, with $f(x) = w_1x + w_0$. The model needs to estimate w_1 , the slope of the line, and w_0 the offset.

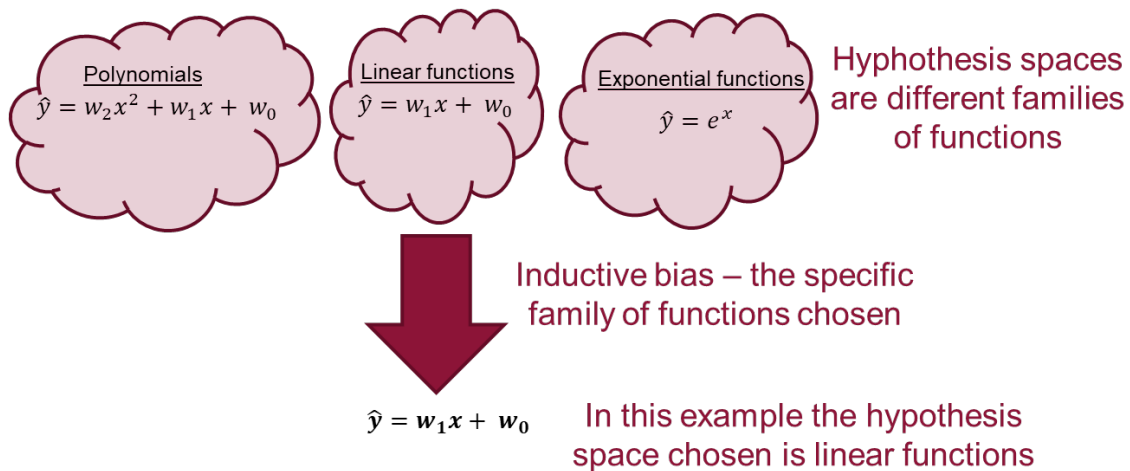


Figure 4: The inductive bias identifies the hypothesis space to be used from all possible hypothesis spaces

It is normal in ML to regard the parameters of a model as a vector θ . Here $\theta = [w_1, w_0]^T$. I have followed the common convention in ML of writing a column vector as the transpose of the row entries, which saves space. The predictor function returns the prediction from an input x and parameters θ and is written as $f(x, \theta)$.

2.2.7 Choosing and optimising a loss function

Once we have chosen our hypothesis space we need to find the best parameters, or hypothesis, to use. To find the best parameters we need a way to decide whether one hypothesis is better than another. This is done through a **loss function** which measures the quality of prediction made by our ML model. The loss function $\mathcal{L}(\theta)$ determines the usefulness of the predictions, for example, in estimating turnover. So,

choosing a loss function is a crucial decision when setting up a ML problem. In this example, the loss function for each sample is chosen as the square of the distance between the actual turnover y and the predicted turnover \hat{y} :

$$\mathcal{L}(y, \boldsymbol{\theta}) = \|y - \hat{y}\|^2 = \|y - f(x, \boldsymbol{\theta})\|^2$$

The loss function is shown in Figure 5

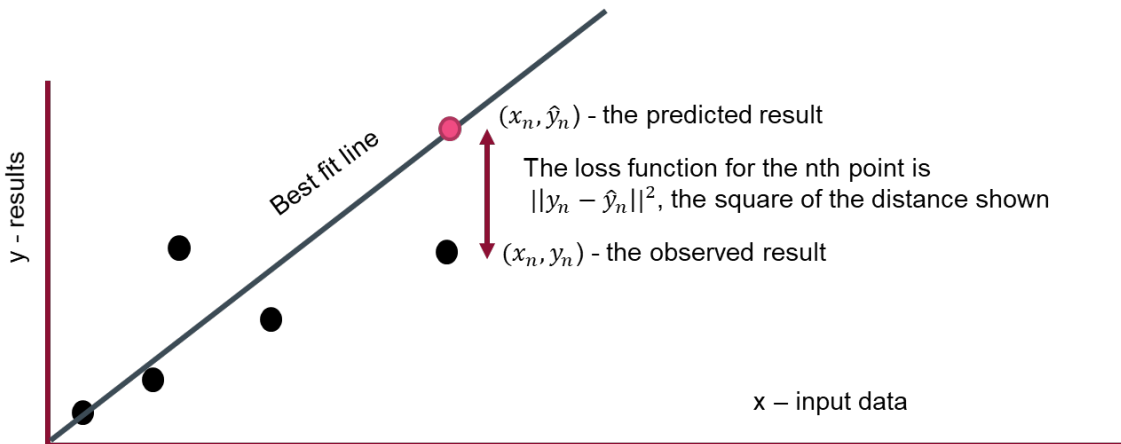


Figure 5: The loss function shown diagrammatically

The parameters are chosen to minimise the sum of the loss function over all the samples in the training set,

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_i \mathcal{L}(y_i, \boldsymbol{\theta})$$

which is a “least squares” **optimisation** problem. There is an exact solution under certain conditions.

If we consider the n th sample we can write x_n as a vector \mathbf{x}_n . Here $\mathbf{x}_n = [x_n, 1]^T$. This means that we can neatly write $f(x_n) = \mathbf{x}_n^T \boldsymbol{\theta}$, the vector dot product, or inner product, of the input vector and the parameter vector. If there are N samples we can write the **design matrix** as $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n, \dots, \mathbf{x}_N]^T \in \mathcal{R}^{N \times 2}$, and the observed results vector $\mathbf{y} = [y_1, \dots, y_i, \dots, y_N]^T \in \mathcal{R}^N$. All these different matrices and vectors are illustrated in Figure 6

Using this notation it is shown in Chapter Nine of [Mathematics for Machine Learning](#) [8] that by differentiating $\boldsymbol{\theta}$ a minimum of the loss function is found when

$$\boldsymbol{\theta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$$

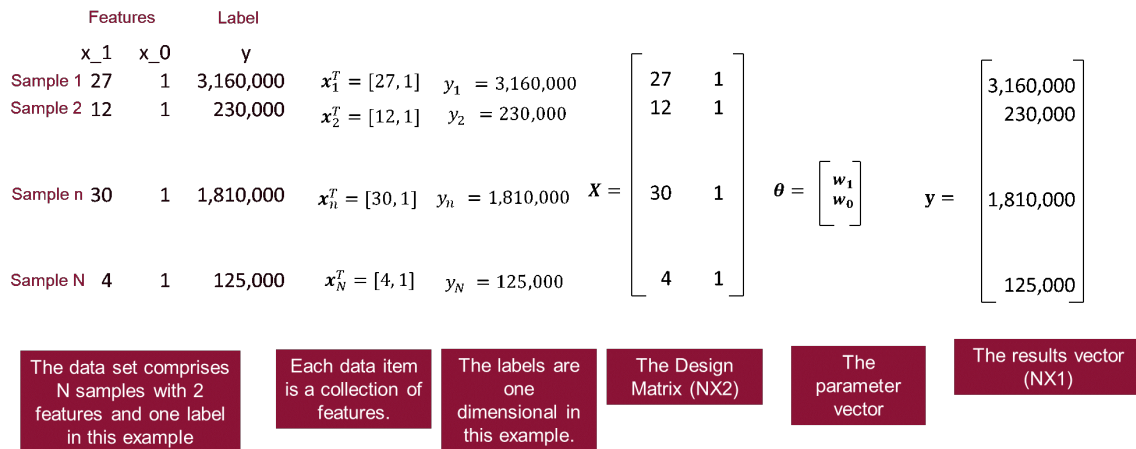


Figure 6: This diagram show the dataset with features and labels, the results vector, design matrix, and the parameter vectors

In other words this is a matrix inversion problem where an equation of form $\mathbf{A}\boldsymbol{\theta} = \mathbf{b}$ with $\mathbf{A} = \mathbf{X}^T\mathbf{X}$ and $\mathbf{b} = \mathbf{X}^T\mathbf{y}$ needs to be solved. Figure 7 below shows the dataset, and the best fit line for the training data against the test data calculated by SciKit Learn:

2.2.8 Useful linear regression documentation

Details of the LinearRegression class used are found on the SciKit Learn (abbreviated to sklearn) [documentation](#). There are notes on linear models in the SkiKit Learn [user notes](#).

2.2.9 QML application - linear regression

The matrix inversion problem above could, in principle, be solved by the quantum HHL algorithm. However, this is likely to require fault tolerant quantum computers, so the [caveats](#) above are relevant.

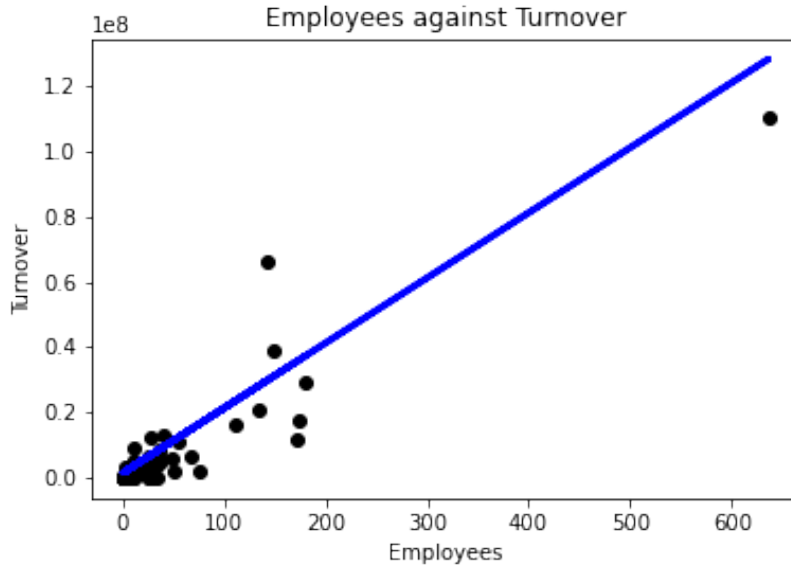


Figure 7: Plot of employee data against company turnover: a best fit straight line from the training data against the test data

2.3 Overfitting

In the example above we predicted unseen data points using a best fit straight line $f(x) = w_1x + w_0$. We might think that we could increase the accuracy of our prediction by changing our **inductive bias** to use a different function from a different **hypothesis space**, a polynomial:

$$f(x) = w_6x^6 + w_5x^5 + w_4x^4 + w_3x^3 + w_2x^2 + w_1x + w_0$$

This predictor function has more parameters, and so might fit the dataset better. In the second [example](#) I have done exactly that, using some manufactured data. However, there is a snag. Now, the best-fit curve shown in [Figure 8](#) has too many parameters: it fits the training data well, but does not generalise to the unseen test data as well as the simple straight line. This is called **over-fitting**. For some datasets the fit to the polynomial might give better generalisation than the straight line, but here the performance is worse.

One solution to overfitting is **regularisation**. Early stopping, feature selection and ensembling are other solutions widely used. With regularisation the predictor function could be simplified, using a straight line, not a polynomial. Alternatively




Figure 8: Plot of employee data against company turnover : the polynomial has over-fitted the data

a penalty term $\lambda \|\theta\|^2$ can be applied to the loss function. This works because the magnitude of the parameter values can become larger with overfitting.

2.4 Clustering

The examples so far have been examples of **supervised learning**: we have fed our model training examples, with input values and labels (employee numbers), and then our model has been able to predict the label for unseen data.

If an AI system is to learn about its environment, it must derive patterns in the data without any examples. This is called **unsupervised learning**, and is often very useful for exploratory data analysis. An example of this is **clustering**.

In the third [example](#) we use **clustering** to investigate the latitudes and longitudes of the UK quantum technology companies that we met first in the linear regression example.

2.4.1 Aim of the clustering example

The aim of the clustering example is to group the companies into clusters so that companies in the same cluster are close together.

2.4.2 K-Means clustering

A tool called [K-Means clustering](#) is used to group the data into 10 clusters, as shown in Figure 9. In this algorithm $K = 10$ initial centroids are chosen randomly. Data points are assigned to clusters so that the **distance** from data points to the initial centroids is minimised. Then $K = 10$ new centroids are chosen at the centre of the new clusters. The data points are reassigned, and the process continues iteratively, until there are no further re-assignments of data points to new clusters.

A key point is that the “**distance**” is the **Euclidean distance** in two dimension, where the distance between vectors $\mathbf{a} = (a_x, a_y)$ and $\mathbf{b} = (b_x, b_y)$ is $\sqrt{(a_x - b_x)^2 + (a_y - b_y)^2}$.

The “**centroids**” of the clusters found are shown as crosses.




Figure 9: Plot of UK quantum technology company locations, and their K-means centroids, with $K = 10$

If we look at the plot of these centroids on a map of the UK shown in Figure 10 we see that the quantum technology companies are clustered around Quantum Hubs at Birmingham, Glasgow, Oxford and important Universities at Cambridge,

London, Bristol, Belfast, Liverpool and Manchester. There are also companies in the South West and North East. The machine learning algorithm has found these clusters without any external help.

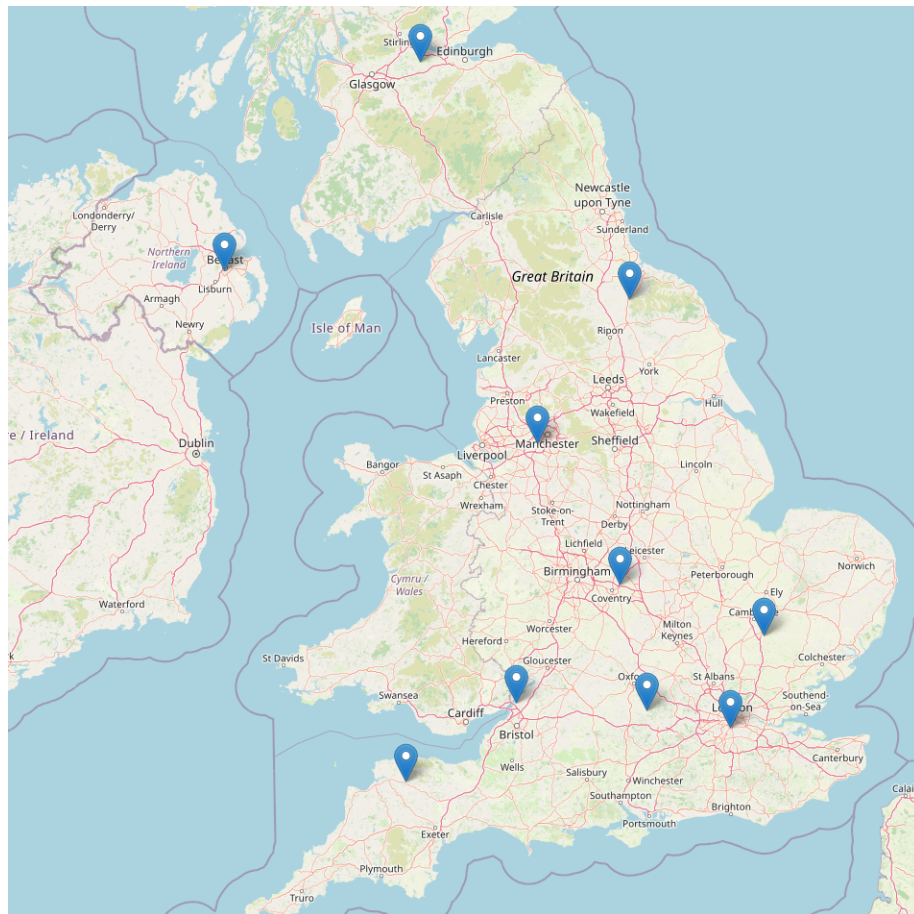


Figure 10: Plot of the K=10 Means centroids for UK quantum technology companies on a map of the UK

There is more information on [SciKit learn](#). If you would like to practice your clustering there are some artificial and real world datasets provided by [Thomas Barton](#) available on [GitHub](#).

2.5 Increasing dimensionality - the kernel trick

The next example is a **binary classification** supervised learning problem. I have placed a [Jupyter Notebook example](#) on GitHub.

2.5.1 Aim of the support vector classifier example

The aim of a support vector classifier is to predict labels for unseen data. The input is a two dimensional co-ordinate vector, the labels are binary $y_i \in \{-1, 1\}$, and as explained below, the data points can be separated by a linear threshold.

2.5.2 The support vector classifier

In Figure 11 different colours have been given to training data points with different labels.




Figure 11: Plot of the training data for the support vector classifier example

The algorithm used, the **support vector classifier**, is based on an inner product, which gives a measure of similarity between two vectors. In this example the inner product is defined as the dot product and the inner product of two vectors \mathbf{x} and \mathbf{x}' is:

$$\langle \mathbf{x}, \mathbf{x}' \rangle = \mathbf{x}^T \mathbf{x}'$$

An inner product is a map from two vectors in the data domain \mathcal{X} to a real number which meets certain [axioms](#).

$$\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{R}$$

Just like the inner product in quantum mechanics, the inner product in Machine Learning gives a measure of similarity between two vectors. Using the inner product allows access to geometrical concepts such as the angle between two vectors, and their length.

In this example a **support vector classifier** uses the dot product to find a **threshold**, a line to classify the new input data. Any data above the threshold line is classified as red, and anything below the line is classified blue. Figure 12 shows that the support vector classifier gives accurate predictions for the test data. The graph also shows the threshold and the **margin**, the distance between the threshold and the closest observation to the threshold. The margin is shown as two dashed lines.




Figure 12: Plot of predictions showing the threshold, and margins

In support vector classification, the threshold is found by maximising the margin. The inner product plays a fundamental role in solving this optimisation problem.

2.6 The support vector machine

In all the examples so far there is an explicit assumption that similar points are “close” in two dimensional space, either because they are separated by a small Euclidean distance, or have a normalised inner product for the data position vectors close to one.

In the next example similar points are not necessarily close in two dimensions. I have placed a [Jupyter Notebook](#) on GitHub illustrating the **Kernel trick**, which is a general method to solve classification problems where the data points can’t be separated by a linear classifier, as shown in Figure 13.




Figure 13: Plot of training data for the support vector machine example

2.6.1 Aim of the support vector machine example

The aim of the support vector machine example is to separate data points that can’t be separated by a linear classifier.

2.6.2 The kernel trick

As explained in “*Machine Learning with Quantum Computers*” by Maria Schuld and Francesco Petruccione [17] the **support vector machine** algorithm requires a **kernel**, an inner product which gives a measure of similarity between two vectors, but in this case the inner product is between two vectors in a higher dimension feature space \mathcal{F} defined by a feature map $\phi : \mathcal{X} \rightarrow \mathcal{F}$. The kernel κ for two vectors \mathbf{x} and \mathbf{x}' is defined as:

$$\kappa(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle$$

In this example a feature map ϕ is defined as

$$\phi((x_1, x_2)^T) = (x_1, x_2, 0.5 * (x_1^2 + x_2^2))^T$$

The feature map is used to transform the dataset from a two dimensional plane into three dimensional space, as shown in Figure 14. The points are now separable. The threshold is a two dimensional plane perpendicular to the z-axis.

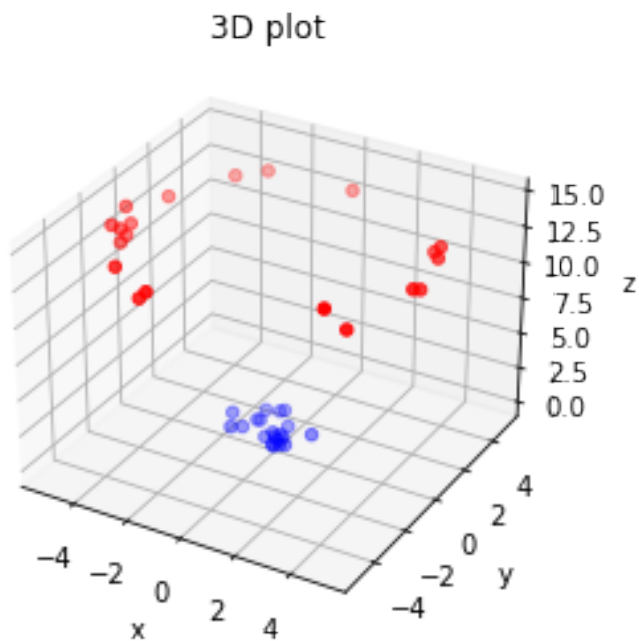


Figure 14: Plot of data in three dimensions

It is not necessary to actually do a mapping and hold the information in the extra dimensions introduced. All that is needed is to compute the kernel function between two points x and x' and return the result, which is a real number. This is called the **kernel trick**. This can be seen in the Notebook where a new kernel function has been written to be called by the SciKit Learn support vector classifier, which does not “know” about the new dimension introduced. The accurate predictions by the support vector machine classifier are shown in Figure 15.




Figure 15: Plot of predictions

If we added more dimensions to our example the threshold would become a **hyperplane** in several dimensions. There is a great [YouTube video](#) on Support Vector machines.

2.6.3 The quantum kernel classifier

There is a link between the kernel function and the quantum inner product. The kernel function is defined as:

$$\kappa(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle$$

This is very similar to the quantum inner product:

$$\langle \phi(\mathbf{x}) | \phi(\mathbf{x}') \rangle$$

2.6.4 The quantum kernel classifier in the NISQ era

The quantum inner product can be estimated by a simple circuit, and as mentioned in the introduction, “a quantum feature map may enable a quantum kernel classifier on NISQ machines”. However, there are still considerable challenges in translating the classical data to a quantum form. Maria Schuld has written an insightful [paper](#) about quantum kernel methods [19].

Useful progress has been made on NISQ machines. A 2022 [paper](#) by Boris Albrecht and others [5] describes a quantum kernel running on a 32 qubit neutral atom quantum device. In this toxicity screening experiment, molecules were labelled as either toxic or not toxic, and the QML algorithm classified unseen molecules, with input data being a graph representing the arrangement of atoms in the molecule. They achieved results comparable to the best classical algorithms.

The quantum kernel used is explained in a 2021 [paper](#) by Louis-Paul Henry and others [20]. Each atom in the molecules screened was represented by a qubit, and the graph of the molecular structure determined which qubits were allowed to interact. The quantum state of the system was updated by applying a series of Hamiltonians. The precise form of the Hamiltonian pulse sequence was found through training. An observable was measured at the end of the updates, yielding a bit string. From this bit string a probability distribution $\mathcal{P}(\mathcal{G})$ was built up for each graph by repeating the measurements several times.

Once probability distributions had been obtained for each graph, the kernel of two graphs \mathcal{G} and \mathcal{G}' was calculated classically as:

$$\kappa(\mathcal{G}, \mathcal{G}') = \exp[JS(\mathcal{P}, \mathcal{P}')] \in [0.5, 1]$$

JS is the Jensen-Shannon divergence, a measure of the distance between the two probability distributions \mathcal{P} and \mathcal{P}' . This kernel was then used in a support vector machine, as described above.

In a 2018 [paper](#) Vojtech Havlicek and others [21] describe using 5 qubit superconducting quantum devices to experimentally realise two other quantum kernels.

2.7 Reducing problem dimensionality - Principal Component Analysis

In the example above we **increased the dimensionality** of the clustering problem to make points linearly separable and then used a support vector machine to find the threshold hyperplane. In the next example we will use **Principal Component Analysis (PCA)**, an unsupervised algorithm, that **decrease the dimensionality** of the data to find the key features as shown in Figure 16. The mathematics used is interesting, but not essential, and [an Appendix](#) explains how the co-variance matrix is diagonalised, and the principal components are then identified as the eigenvectors with the largest eigenvalues.




Figure 16: In Principal Component Analysis the original data is compressed, retaining the key information

If we plot a graph of the data against the principle components then we can see patterns that we would otherwise miss. An [example from SciKit Learn](#) used the famous IRIS dataset where four features of the flowers are considered: sepal length, sepal width, petal length and petal width. When the dimensionality of the problem

is reduced to two dimensions by PCA, it can be seen in Figure 17 that there are three clusters of data, representing three different species of Iris - Setosa, Versicolour and Virginica.




Figure 17: Results of the PCA analysis of the IRIS dataset from SciKit Learn

There is more information about PCA on [SciKit Learn](#).

2.8 Deep learning

Deep learning is inspired by biological models, such as the brain, in which signals are processed by individual neurons connected to thousands of other neurons, which “fire” dependent on the signals received. In deep learning these neurons are modelled by networks. A very simple network is shown in Figure 18:

In this example there are four inputs and each of these is adjusted by a weight. The weight adjusted inputs are summed, and a bias b is applied. If the input is a vector \mathbf{x} and the weights are given by \mathbf{W} so that the sum z is given by

$$z = \mathbf{W}^T \mathbf{x} + b$$

The output will then depend on some activation function, $g(z)$. Rectified linear units use the activation function $g(z) = \max(0, z)$ so there is a linear output if the sum




Figure 18: A simple network with four inputs

is greater than zero. Other functions, such as the sigmoid $S(z)$, which produces an output between zero and one are also used in some problems and have an output

$$S(z) = \frac{1}{1 + e^{-z}}$$

The power of deep learning comes from using large numbers of nodes connected together. In Figure 19 two **hidden layers** are added and each arrow is associated with a weight. The idea behind this is that the hidden layers might enable aspects of the image, like lines, to be trigger activations in the first layer, and these are combined so that geometrical shapes, like triangles, trigger activations in the second layer, and so on with more abstract levels of detail, until objects like hand written numbers can be classified.

In a supervised learning classification problem a loss function is set to zero if the classification is successful, and to one if it fails. This loss function is a complex function of the various weights and activation functions in the model. **Backpropagation** can be used to adjust the model weights according to the gradient of the loss function with respect to the given parameter. A feedback loop can be used to modify parameters based on the results on previous calculations, to minimise the loss function.

Larger versions of these types of networks have been very successful at a wide range of problems including complex image and video analysis and generation, solving complex decision making problems at or beyond human level capability, natural language processing and generation.

If you would like to work through an example there is a great [PyTorch Tutorial](#) that uses the FashionMNIST dataset to train a neural network. The model predicts




Figure 19: A network with two hidden layers

if an input image belongs to one of 10 classes such as T-shirt, Trouser and Pullover. The tutorial starts with an introduction to Tensors, which are used to hold a representation of the network. Then the tutorial walks through loading and transforming the data, building a neural network, and **automatic differentiation** to optimise the model parameters. The dataset comprises 28×28 grayscale image and an associated label from one of the 10 classes. The model has an input has 784 features, one for each grayscale value, and hidden layers with 512 nodes.

2.9 Generative Adversarial Networks

You may have read that **deepfakes** can be generated by Generative Adversarial Networks, or GANS. A deepfake is a digital object such as an image, or video clip, produced by Machine Learning techniques. The object, like the face in Figure 20 appears realistic, but is a deepfake generated by an algorithm:

A GAN is made up of two deep learning models that compete against each other. A number of real examples, such as images, are available. As shown in In Figure 21,



Figure 20: A deepfake image produced by bod lnga klang, Public domain, via Wikimedia Commons

the Generator Model generates new examples and the Discriminator model classifies the examples as real or fake. Both models are trained. The Generator is trained to produce fakes that fool the Discriminator. The Discriminator is trained to identify real examples from fakes. Ultimately the Generator can produce examples that seem very realistic, as seen in Figure 20.

In Figure 21 above the latent space is a random seed that ensures that each Generator example is different.




Figure 21: A GAN. The Generator is optimised to produce fakes that fool the Discriminator. The Discriminator is optimised to identify real examples from fakes

2.9.1 A quantum latent space

As stated in the introduction “a possible early use case for quantum machine learning is to use quantum computers to generate the latent spaces for GAN models”. A Boson Sampling quantum computer can be used to generate the latent space, with promising results.

An alternative approach, set out in a 2019 [paper \[22\]](#) adds a Quantum Boltzman machine to the Discriminator to improve performance.

3 To conclude

This review has described how the “human like intelligence” of **Artificial Intelligence** (AI) has evolved from **Expert Systems** to **Machine Learning** (ML) algorithms that can learn from data. In a simple **linear regression example** we looked at a **hypothesis space** of linear predictor functions selected from a large range of possible prediction functions by the **inductive bias**.

We defined a **loss function** to measure the quality of the solution, and observed that many ML problems are optimisation problems in which parameters are varied to minimise the loss function. We showed how having too many parameters can lead to **overfitting**, where the predictions tracked noise in the training data too closely, and did not generalise properly.

We then moved from the linear regression **supervised learning** example, where labels are provided for each example, to an **unsupervised learning clustering**

example, where the algorithm learnt to group the points into clusters without pre-defined labels for the data.

Next, we looked at a **support vector classifier** which found a linear threshold between two groups of points. This example was generalised to a **support vector machine** where the data set, when plotted in two dimensions, could not be separated by a linear threshold. Instead the Kernel trick was used to introduce a higher dimension in which the data points could be separated by a hyper-plane.

Instead of increasing the number of dimensions, in **Principal Component Analysis** (PCA) the number of dimension was reduced. Linear algebra techniques were used to diagonalise the covariance matrix, and we recognised that the eigenvectors with the largest eigenvalues (the diagonal entries) contained most of the information about the data, allowing a simplified analysis in two or three dimensions.

Finally, we look at **deep-learning**, which is inspired by biology-like networks, and generalised this to a **Generative Adversarial Networks**, where a **Generator** is locked in a games-theory inspired competition with a **Discriminator**. Even a human can struggle to identify the Generator's fakes from the real images used in training.

We have also shown where quantum computing might be relevant, and raised some important **caveats** about the readiness of the technology.

We hope that you have found this review valuable. If you have any questions or feedback please contact the quantum team at Digital Catapult quantum@digicatapult.org.uk

Acknowledgements

Many thanks to Robert Elliott Smith for his helpful comments.

Appendix - The mathematics of Principal Component Analysis

This Appendix explains how the covariance matrix is diagonalised, and the principal components are identified as the eigenvectors with the largest eigenvalues.

Following [Mathematics for Machine Learning](#) [8] we consider a set of N data vectors x_n , each of dimension D , with mean zero. The dataset can be written $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_n, \mathbf{x}_N\}$ where $\mathbf{x}_n \in \mathcal{R}^D$. Because the mean is zero the data covariance matrix is a $D \times D$ matrix defined as:

$$\mathbf{S} = \frac{1}{N} \sum_{n=1}^N \mathbf{x}_n \mathbf{x}_n^T$$

The data covariance matrix is a symmetric matrix. The diagonal elements represent the variance, a measure of how much variability there is in each dimension of the data. The other matrix elements represent the covariances, a measure of how the variability in one data dimension is linked to variability in the other. This matrix is very similar to the density matrix ρ in quantum mechanics where the sum of the diagonals is normalised to one, and the off-diagonal terms hold details of correlations:

$$\rho = |\psi\rangle\langle\psi|$$

There is a well known theorem from linear algebra that a symmetric matrix like \mathbf{S} can be diagonalised and we can always write:

$$\mathbf{S} = \mathbf{P}\mathbf{D}\mathbf{P}^{-1}$$

\mathbf{D} is a diagonal matrix and it is often the case that some of the diagonal entries are much larger than others. Conventionally \mathbf{D} is written with the largest eigenvector λ_1 at the top, and the rest in descending order:

$$\mathbf{D} = \begin{bmatrix} \lambda_1 & 0 & \cdot & \cdot & \cdot \\ 0 & \lambda_2 & 0 & \cdot & \cdot \\ 0 & 0 & \lambda_3 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ 0 & \cdot & \cdot & \cdot & \lambda_D \end{bmatrix}$$

If we multiply both sides of the equation $\mathbf{S} = \mathbf{P}\mathbf{D}\mathbf{P}^{-1}$ above by \mathbf{P} from the right then we get.

$$\mathbf{S}\mathbf{P} = \mathbf{D}\mathbf{P}$$

and considering the column vectors of \mathbf{P} : $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_M$ then we see these satisfy an eigenvector equation:

$$\mathbf{S}\mathbf{p}_i = \lambda_i \mathbf{p}_i$$

The column vectors \mathbf{p}_i are all eigenvectors of \mathbf{S} . Those associated with the largest diagonal entries are the **principal components**. Moreover these eigenvectors are row vectors of \mathbf{P}^{-1} since it is known that $\mathbf{P}^T = \mathbf{P}^{-1}$.

If we take the first M rows of \mathbf{P}^{-1} this $M \times N$ matrix will project the data onto a lower dimension $M \times M$ subspace whilst retaining much of the information. The matrix multiplication is shown in Figure 22




Figure 22: A vector \mathbf{z} in reduced M dimensional space can be produced by multiplying a D dimensional data vector \mathbf{x} by the transpose of the first M eigenvectors of the covariance matrix.

References

- [1] Natansh Mathur, Jonas Landman, Yun Yvonna Li, Martin Strahm, Skander Kazdaghli, Anupam Prakash, and Iordanis Kerenidis. Medical image classification via quantum neural networks. <http://arxiv.org/abs/2109.01831>, 9 2021.
- [2] Amira Abbas, David Sutter, Christa Zoufal, Aurélien Lucchi, Alessio Figalli, and Stefan Woerner. The power of quantum neural networks. *Nature Computational Science*, 1:403 – 409, 2020.
- [3] Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors. <https://arxiv.org/abs/1802.06002>, 2 2018.
- [4] He Liang Huang, Yuxuan Du, Ming Gong, Youwei Zhao, Yulin Wu, Chaoyue Wang, Shaowei Li, Futian Liang, Jin Lin, Yu Xu, Rui Yang, Tongliang Liu, Min Hsiu Hsieh,

- Hui Deng, Hao Rong, Cheng Zhi Peng, Chao Yang Lu, Yu Ao Chen, Dacheng Tao, Xiaobo Zhu, and Jian Wei Pan. Experimental quantum generative adversarial networks for image generation. *Physical Review Applied*, 16:024051, 8 2021.
- [5] Boris Albrecht, Constantin Dalyac, Lucas Leclerc, Luis Ortiz-Gutiérrez, Slimane Thabet, Mauro D’Arcangelo, Vincent E. Elfving, Lucas Lassablière, Henrique Silvério, Bruno Ximenez, Louis-Paul Henry, Adrien Signoles, and Loïc Henriët. Quantum feature maps for graph machine learning on a neutral atom quantum processor. <http://arxiv.org/abs/2211.16337>, 11 2022.
- [6] Marcello Benedetti, Brian Coyle, Mattia Fiorentini, Michael Lubasch, and Matthias Rosenkranz. Variational inference with a quantum computer. *Physical Review Applied*, 16, 3 2021.
- [7] Michele Grossi, Noelle Ibrahim, Voica Radescu, Robert Loredó, Kirsten Voigt, Constantin von Altrock, and Andreas Rudnik. Mixed quantum–classical method for fraud detection with quantum feature selection. *IEEE Transactions on Quantum Engineering*, 3:1–12, 2022.
- [8] Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon On. Mathematics for machine learning. <https://mml-book.com>. Accessed: 2023-01-21.
- [9] Eric Matthes. *Python Crash Course - A Hands-on, Project-Based Introduction to Programming*. No Starch Press, San Francisco, second edition, 2019.
- [10] Daniel Goldsmith. Quantum technology access programme github repository. <https://github.com/digicatapult/QTAP>, 20223.
- [11] John Preskill. Fault-tolerant quantum computation. <https://arxiv.org/abs/quant-ph/9712048>, 12 1997.
- [12] Lov K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters*, 79:325–328, 6 1997.
- [13] Aram W. Harrow, Avinandan Hassidim, and Seth Lloyd. Quantum algorithm for solving linear systems of equations. *Physical Review Letters*, 103, 11 2008.
- [14] A. Yu. Kitaev. Quantum measurements and the abelian stabilizer problem. <http://arxiv.org/abs/quant-ph/9511026>, 11 1995.
- [15] Stuart Russell and Peter Norvig. Artificial intelligence a modern approach, third edition. <https://zoo.cs.yale.edu/classes/cs470/materials/aima2010.pdf>. Accessed: 2023-01-21.

- [16] Robert Elliott Smith. *Rage Inside the Machine: The Prejudice of Algorithms, and How to Stop the Internet Making Bigots of Us All*. Bloomsbury Business, London, 2019.
- [17] Maria Schuld and Francesco Petruccione. Machine learning with quantum computers. <https://link.springer.com/book/10.1007/978-3-030-83098-4>. Accessed: 2023-01-21.
- [18] Oliver Theobald. *Machine Learning for Absolute Beginners*. Scatter Plot Press, third edition, 2022.
- [19] Maria Schuld. Supervised quantum machine learning models are kernel methods. <https://arxiv.org/abs/2101.11020>, 2021.
- [20] Louis-Paul Henry, Slimane Thabet, Constantin Dalyac, and Loïc Henriët. Quantum evolution kernel : Machine learning on graphs with programmable arrays of qubits. *Physical Review A*, 104, 7 2021.
- [21] Vojtech Havlicek, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Supervised learning with quantum enhanced feature spaces. *Nature*, 567:209–212, 4 2018.
- [22] Eric R. Anschuetz and Cristian Zanoci. Near-term quantum-classical associative adversarial networks. *Physical Review A*, 100, 5 2019.