The Speaker
Dr Peter Grossi

Graduated Electronics, Maths and Physics 1969.
With Ferranti until 1989 developing computer systems, mostly for defence.
Formed own company and operated as technical and management consultant until 2014.
Obtained PhD for research into mathematical modelling of organisational performance in 2007.
Member of the IQA and it’s successor the CQI for many years, but retired from it when I retired.
Continuing member of the IET to keep in touch with engineering developments.

My interest in AI started with a seminal book “Neural Computing – an introduction” by R. Beale and T. Jackson soon after it was published in 1990.

I produced small AI models in my office to develop my practical understanding, and followed developments through on-line technical discussions, training presentations and technical bulletins.

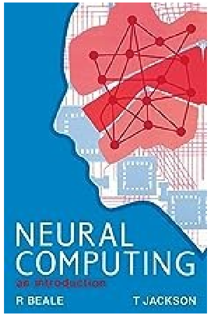

Agenda

What does it do?
Will it replace conventional programming?
Why should I want one?
Where do I find one?
What does it cost?
Where did the idea come from?
How does it work?
How is it programmed?
What are it's strengths and limitations?
How do you know it works?
Review of some current applications

What does it do?

It is critical to understand the difference between two things:

1. A machine that looks intelligent by reflecting the intelligence of those that designed, built and programmed it.
2. A machine that learns for itself how to solve problems and then does so on demand without detailed instruction.

The former is conventional programming, and the latter is Artificial Intelligence.

This presentation is concerned with the latter.

What does it do?

Almost every detail of our lives is governed by patterns.

Faces and voices
Speech and writing
Shopping habits
Travelling habits
Energy consumption
Infrastructure planning
Weather (short-term)
Environment (long-term)
Medical diagnosis and prescription
Drug development
Investments
Maps and navigation
Farming and food production
Business management
Conservation
Etc
Etc

What does it do?

An AI system therefore has at it's heart a pattern recognition module.

For this to work as our minds do it cannot be programmed in a conventional way: it has to learn for itself how to recognise patterns in data.

This may seem fanciful, and it took the best mathematicians over 50 years to work out a way of doing it. But the results now speak for themselves.

This is usually realised by a module that works in a similar way to a biological brain.

An AI system would also contain conventional programming and engineering to process the inputs and outputs.

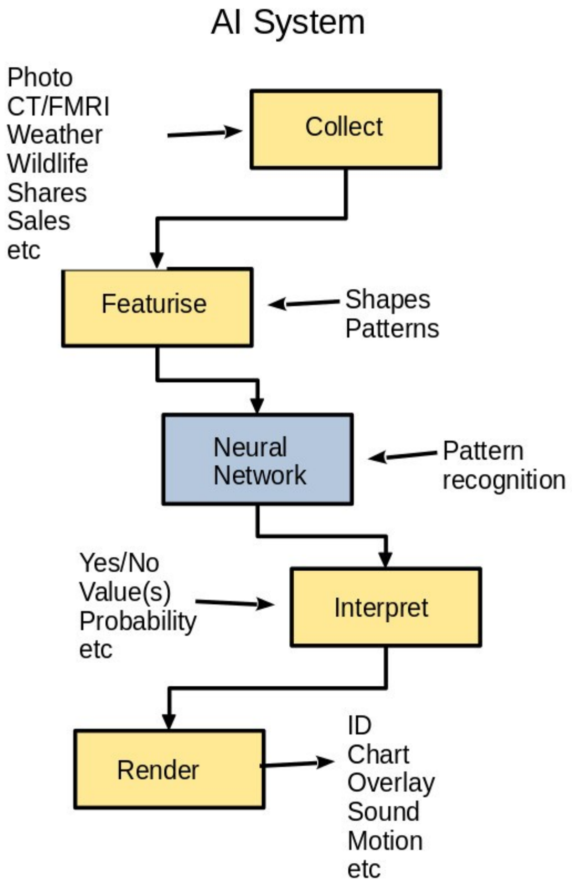

Will it replace conventional programming?

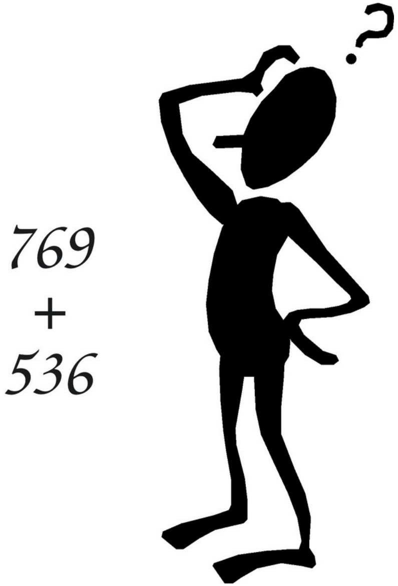

The short answer is No.

Like a biological brain, it is not efficient for deterministic calculations where there is a clearly defined series of calculations required to produce a value.

AI is suited where a trained person would normally be required to cast a judgement on something. Such as recognising someone’s face, or a tumour from a scan.

Most day-to-day requirements for computing applications are deterministic: we can pre-define all the steps required to get the result we want.

We are looking to extend this capability, not to replace it.

Will it replace conventional programming?

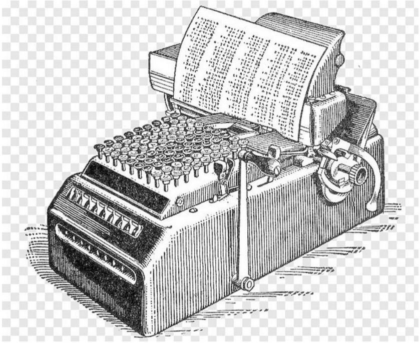

But AI is an extremely important development.

The introduction of AI is as big a step in computing as the introduction of digital computing was from the 1960s.

It extends the capability of computers into areas previously requiring expert human assessment.

Why should I want one?

The short answer is that you don't.

Like a database, AI is embedded in an application that you see only as a whole. You may not even know that you are using AI, even when it is there.

The long answer is that you most likely have already been the beneficiary (or victim) of them.

Where can I find one?

Help desk on-line chatbots.
Face identification on cameras.
Handwriting recognition (an early application).
Speech-to-text.
Market analysis and advert targeting.
Medical diagnosis (NHS £123M funding so far).
Photosimulation – artificial characters replacing actors.

Some of these are good.
They can extend trained professionals’ capacity, save time and reduce costs.

Some are potentially dangerous.
From analysing on-line activity they can recognise people’s weaknesses, prejudices and fears.
They can create content to engage an emotional and possibly dangerous response.

What does it cost?

As a user it costs nothing, as it is provided as part of a system.

As a company resource it is part of the system and software development cost, although there is a significant difference in the development program from conventional methods.

But...

What does it cost?

There is a more serious, even pernicious cost.

Whenever you provide personal details, even to register for a "free" service, your details are collected, aggregated with others' and sold on.

It's not supposed to happen, but it does.

Your on-line activity is also recorded – what you buy, what you search for etc.

This is used for targeted marketing, not all of which is entirely honest.

Companies also use "scraping", where your social media streams are analysed, even your pictures copied.

Personal profiles are constructed without your consent or knowledge.

So the cost of AI is not just the abuse of your identity, which is otherwise already well established, but the way in which it can be used to find subtle ways of influencing you or those whose relationships you value.

Where did the idea come from?

It has been known for a long time that the biological brain has a network of tiny nodules (neurons) interconnected with electrically sensitive fibres (dendrites).

More recently it is found that there are about 86 billion neurons in the average human brain, each connected by about 7000 dendrites.

From about the 1940s, mathematicians pondered whether this could be emulated in a mathematical matrix.

By about 1988 researchers had access to sufficient computing power to try some ideas, based upon the neuron/dendrite structure.

By 1990 small networks had been produced which showed potential for self learning, but these were of limited interest to anybody other than curious mathematicians.

By about 2000 some useful digital applications were produced. The earliest was handwriting recognition (writing in boxes, rather than continuous script at this point).

These applications learned (like the biological brain) by example instead of deterministic programming. This opened up a whole realm of new capabilities, which have developed exponentially since.

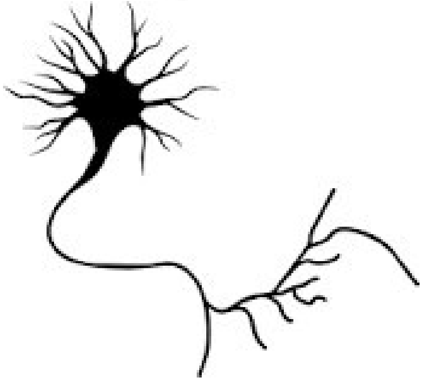

How does it work?

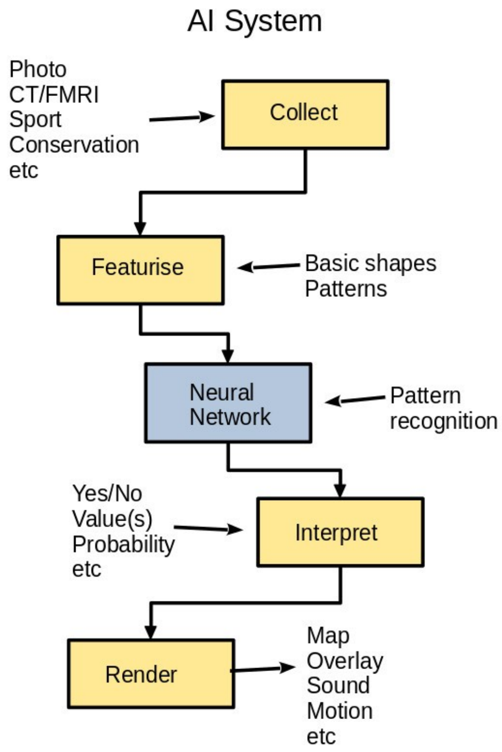

The heart of an AI system is an Artificial Neural Network.

Commonly known as a Neural Network or Neural Net.

It usually consists of elements that emulate biological neurons, and known as Perceptrons.

They are arranged as a cascade, where each “layer” does a simple calculation and feeds the elements in the succeeding layer.

The number of layers and the width of each layer depends on the complexity of the application, but may number in the thousands.

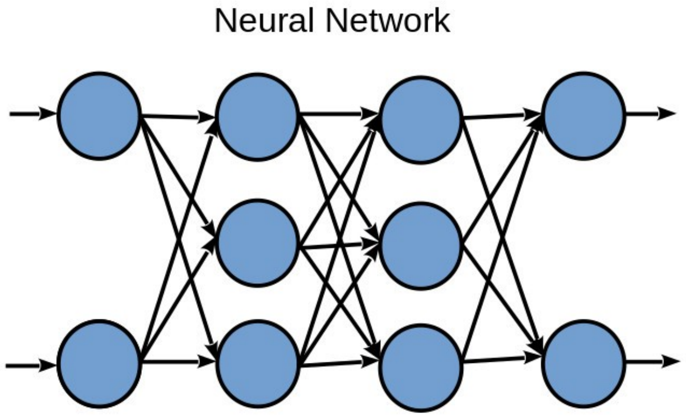

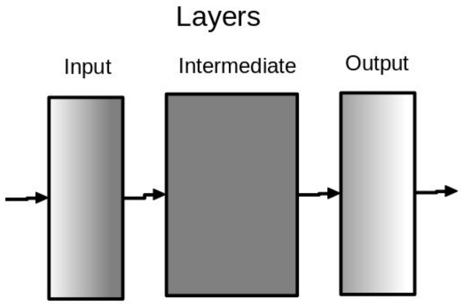

How does it work?

The drawings highlight a single intermediate layer Perceptron and it’s parts.

The first part combines all the inputs from the previous layer, and does something simple to arrive at a single number.

The second part may be one of several types, and arrives at another number.

The third part distributes the new number to all the Perceptrons in the succeeding layer. Each one is first multiplied by a unique weighting factor.

The input and output layers comprise of simplified elements, having only weighting or combining functions respectively.

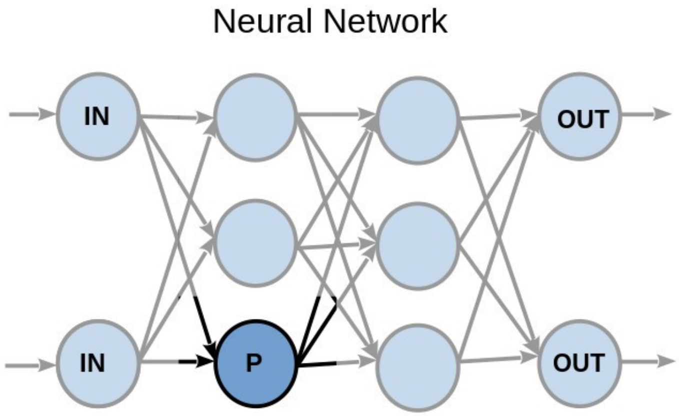

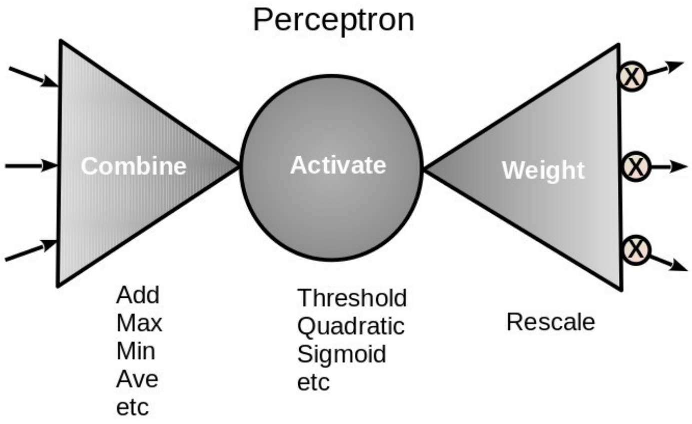

How does it work?

But all the interconnections have individual weights, which can be separately changed, so they represent the memory of the network.
And as these numbers affect the output response to any input, they control what it does.
So these numbers also act like a computer program that can be changed to do different things.
This is totally different from conventional programming, where program instructions and data are separate.

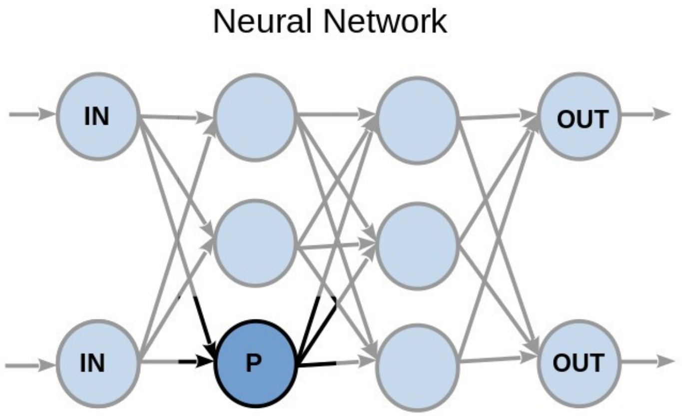

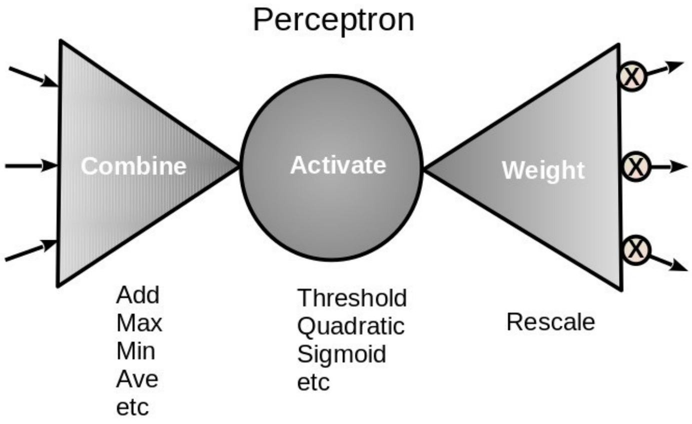

How is it programmed?

A neural network is not programmed as such. It is trained, using a sort of Hebbian learning system.

In this way, an input pattern is provided as a test, and the output from the network is compared with that required.

If the output is not as required, feedback is generated that selectively modifies the internal weights of the network.

Each weight is modified according to it’s value and the significance it has to the output that is in error.

This gets a bit complicated, but the diagram gives a broad idea of how it works.

In response to an error, only very small adjustments are made at any one time.

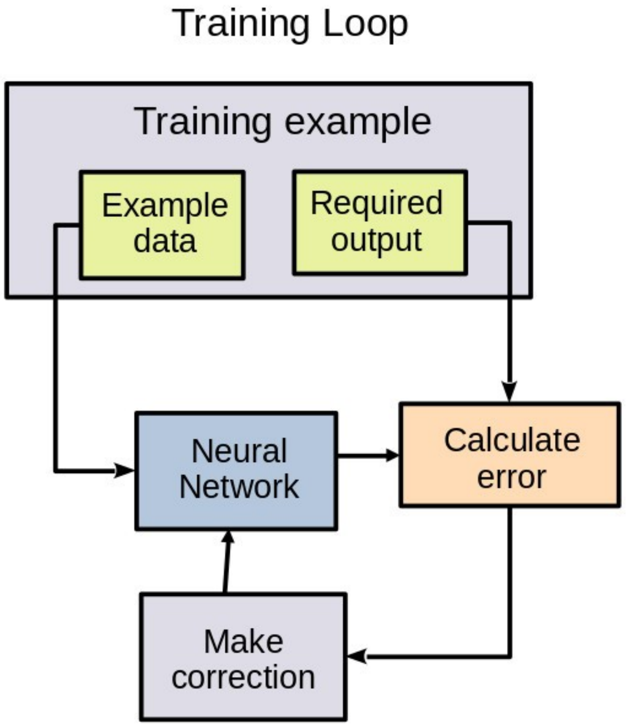

How is it programmed?

Mathematically it is like trying to find the lowest point in a mountain range

You cannot simply run down the slope you are on, because the lowest point locally might not be the lowest point overall.

So you have to start from many points, taking each a step at a time, eliminating the ones that bottom out first.

For a reliable result you need to start from many different places, and spread over the entire range.

How is it programmed?

A useful network will need to be trained on many different examples. For example to recognise handwriting it will need many different written examples of every letter or number.

Each time the network gets an output wrong, every number in the network gets changed to a greater or lesser extent.

So after trying all the examples once, they all have to be gone through again. And again, perhaps thousands of times until they are all reliably recognised.

Humans learn in a similar way. Do you remember learning your Times Table? How many times did you recite it before it stuck?

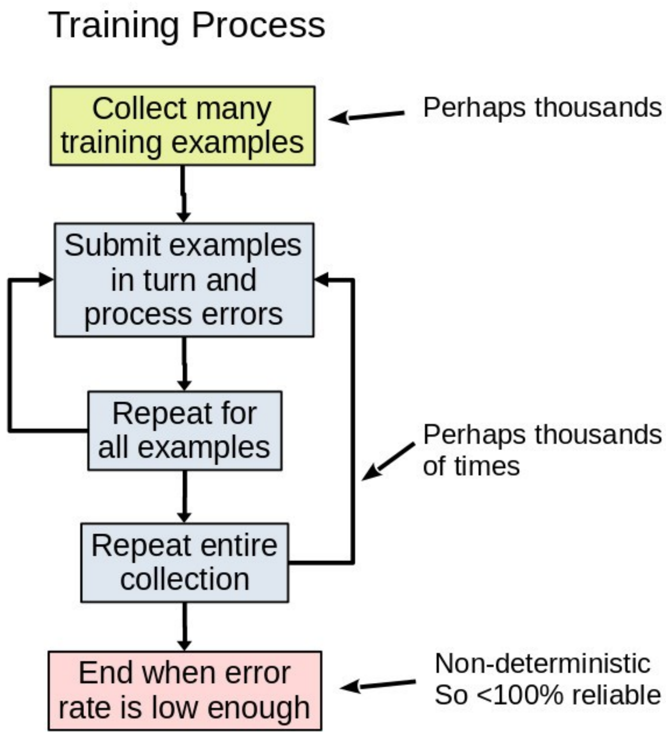

How is it programmed?

A typical learning curve for an AI system shows that the progress is irregular.
The blue profile shows the random errors on first pass of 20 samples, and the overall average at the right edge.
The red, yellow and green profiles show the errors after stages with many passes.
Note that the highs and lows move around, but the overall average and variation reduces.
In a typical system the number of samples and passes may each be in the thousands.
If the system has not been designed suitably the later passes will remain lumpy and possibly switch between extremes instead of becoming smooth.

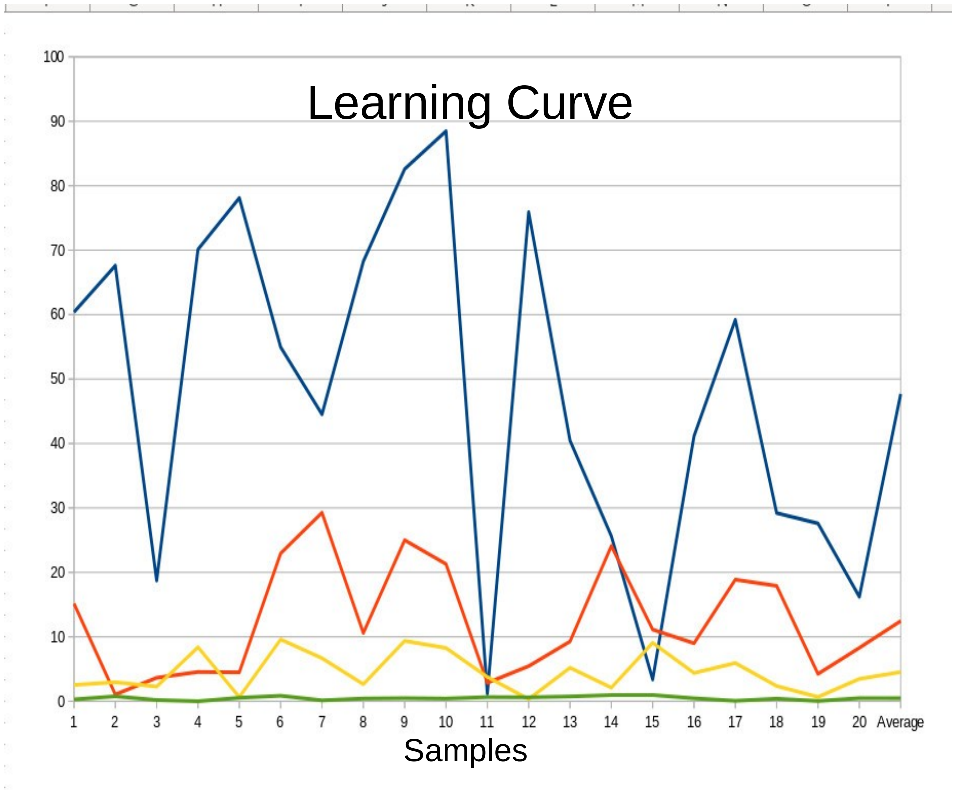

What are it’s strengths?

- When the network is trained properly it should recognise a pattern not seen in training (e.g. a character written in an unfamiliar hand), or at least as likely as by an expert human.
- It can be used to find patterns in data that nobody has noticed.
- When the network is issued to run on a live system it is very compact and efficient; complex patterns are recognised with comparatively small processor effort.
- It can be implemented on existing graphic processing chips that accelerate bulky but simple arithmetic.

What are AI’s limitations?

An AI system is subject to the same problems as any conventional computer system, such as whether the requirement is properly understood and the human or mechanical interface properly conceived.

AI is not deterministic. It is used to provide an assessment rather than a calculation. It can be more reliable than a trained expert, but still not 100% except in the simplest requirements, where a conventional approach may be better suited anyway.

Two important limitations lie in the neural network structure and the training. Both are very different from conventional systems and present their own problems.

What are AI's limitations?

Deciding the size and shape of the network is a black art: there are guidelines but unfamiliar projects may require some experimentation.

The arithmetic functions used for the combinations and activations in the Perceptrons are also critical, and may not be easy to determine without some experimentation.

The training is critical. The operational capability and reliability depends on the relevance, scope and quantity of training data.

Sometimes training data can be gleaned from archives (old medical records, weather maps, security videos, shopping habits etc) but sometimes it has to be collected specially for the application, which can take time and resources.

Training data may be subject to unconscious selection or bias, or selective availability. Any bias will affect the operation and must be honestly recognised as a fault.

So AI can in some cases have unfortunate human tendencies, including bias and laziness.

How do you know it works?

The internal workings of a neural network are entirely inscrutable.
You cannot trace errors back into the code to diagnose and correct them.
The performance can only be assessed statistically – does it perform to the standard set (e.g. compared with trained experts)?

To quality-assure an AI system, a technique similar to assessing people needs to be used.
Does the training data cover the entire scope of the operational requirement?
Is there enough of it?
Does the “learning curve” show progressive improvement to the required standard?
And, after initial use: has it performed as expected in service?

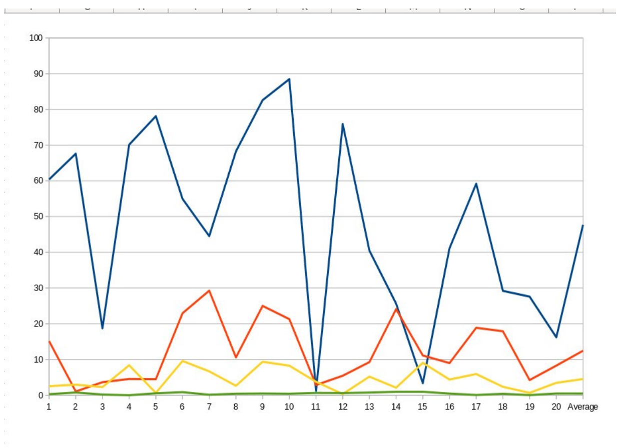

Some Current Applications

Financial
- Stock investment – predicting events and automating transactions
- AlJazeera “Money Bots” Sept 2023
- Fraud and black market recognition

Medical imaging
- Tumour hunting
- FMRI (brain activity) analysis

Medical diagnosis and prescription
- Complex symptoms and interactions

Robotics
- Prosthetics, haptics

Other imaging
- Wildlife conservation – population surveys from aerial photos
- Autonomous drones – auto navigation
- Facial and voice recognition – security and criminal detection

Investigative Genetic Genealogy (IGG)
- Whale tracking
- Criminal investigations

Marketing and influencing
- Demographic/regional analysis and targeting
- Personal profiling by social media, websites and news media scraping