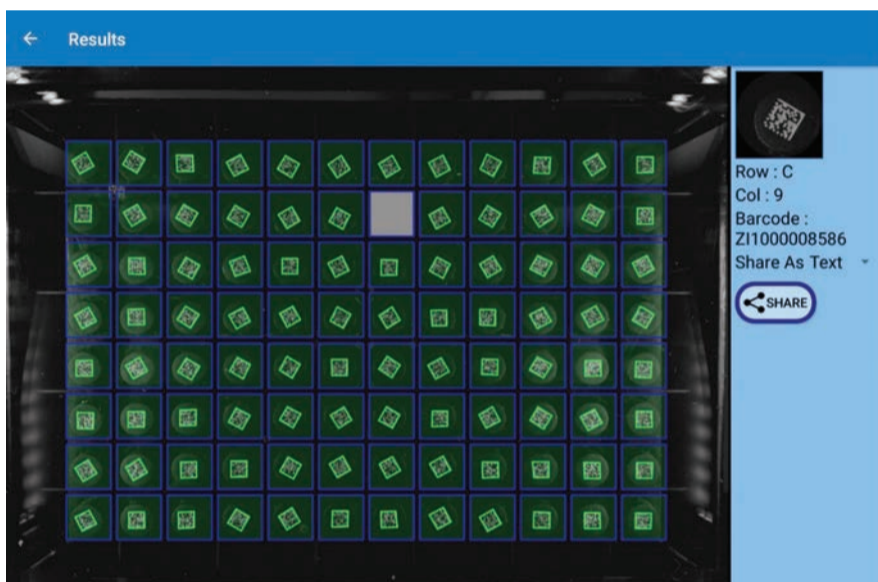


What is AI and how can it be applied to scientific instruments?

Dr Alex Beasley, Ziath Ltd

Senior Electrical Design Engineer Dr Alex Beasley looks at current applications of Artificial Intelligence as applied to real problems in scientific instruments. He demonstrates how a Neural Network approach can be deployed to analyse real-world images and determine key properties from the data with far more accuracy and far faster than traditional detection techniques.



In recent years, the scientific community has increasingly looked at Artificial Intelligence (AI) as a future tool that can deliver great benefits when applied to operations and measurements that instruments undertake. But what exactly is AI?

Firstly, let's define a few terms you may have heard of: Machine Learning is the process to create Artificial Intelligence (AI). Machine learning can be applied through several different mechanisms which include 'fuzzy logic', 'discriminant analysis' and 'neural networks'. Due to their ability to process computationally intensive problems, neural networks are the basis of most commercially-viable AI that could be applied to the type of problems that scientific instruments try to solve.

The limit to implementing a neural network is simply the number of 'nodes', or possible connections, in the processor being used. The number of nodes in the human brain is orders of magnitudes greater than even the most powerful processors can muster. The much-vaunted 'singularity' where a human consciousness could be uploaded in silico is still only the stuff of science fiction. However, in the real world, neural networks have an important role to play.

To successfully apply an AI solution to a given problem, we must first characterise the problem. This is usually a mathematical solution to a specific problem of analysing data, such as a smoothing algorithm, or data trend identification. In designing a neural network, the network is constructed of a number of 'layers', where each layer represents a separate mathematical operation. The combination of these mathematical functions on a data stream allows for data features to be extracted. That might be the useful answer that is sought, or it might be an intermediate stage that undergoes further processing. For example, a network may establish two or three different ways of averaging a given set of data. The outputs could then be fed to a comparator which selects the best fit answer from the three presented to it, based on percentage probabilities of it being correct.

In creating the neural network, it will initially act as a 'black box' - that is to say, that within the network, typically, there are no pre-determined nodal interconnections; no nodes or possible paths through the network have been assigned weights or biases. This is the starting point for the AI.

Determining the dataset to be used is critical to the characterisation of the problem that will be solved using AI. In our case it is a .PNG image file, but any data set could be the starting point. AI is particularly useful for image analysis and has found application already in histocytology for the determination of cervical cancer cell types in smear tests. The AI algorithm has been shown to be nearly 100% accurate in identifying potential tumour cells and it works just as well at 5pm on a Friday as it does at 9am on Monday and does not suffer from fatigue, unlike the human technicians who currently perform these tests.

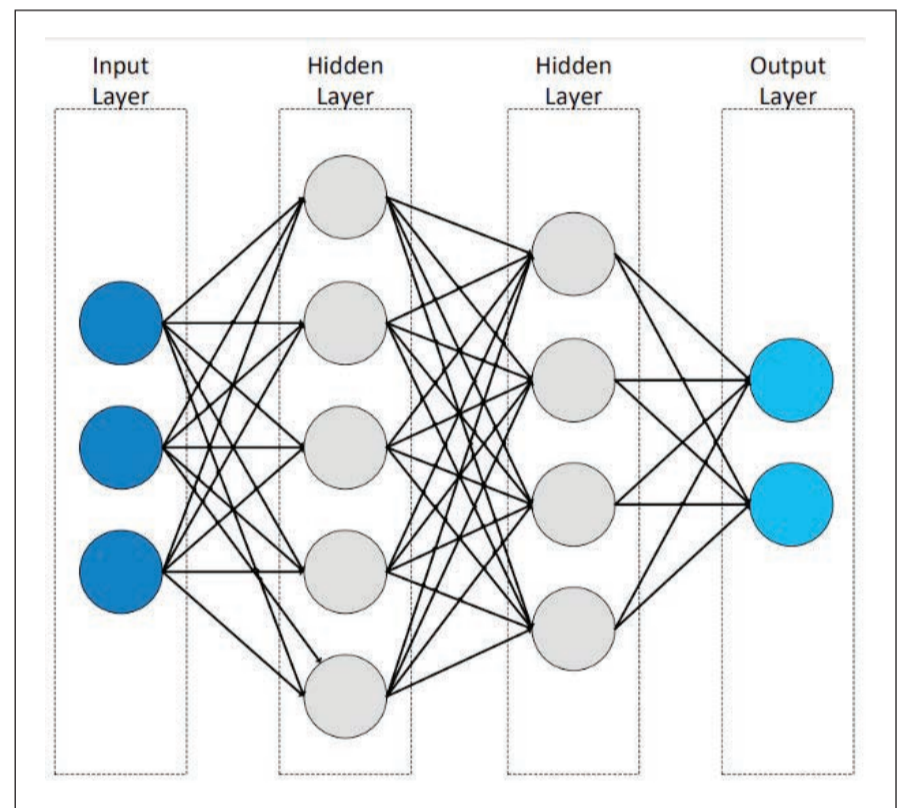


Figure 1. Simplified network diagram showing input, output and two hidden layers. The nodes (circles) represent mathematical functions and the arrows represent data flow through the network.

The Input Layer of the neural network is defined by the size of the input datastream. In our case it is an 8-bit 50 x 50-pixel image where each pixel has a possible value between 0 - 225. Each layer that this data is fed through, in sequence, is a separate mathematical argument. The combination of the mathematical functions yields an overall greater function, such as image classification.

The role of the Classification network is to decide, or classify, the content of the image according to a pre-determined set of parameters. For example, is the output zero, one or greater than one? It could also be a simple yes/no answer, or in our case, is a barcode present or not? In our case, the classification network produces two answers - the probability of the image containing a barcode and the probability the image does not contain a barcode. The output is fed to a comparator that examines the result and yields the answer to the question: does our image contain a barcode? If the network is wrong, then it must be 're-trained'. In a sophisticated network, techniques such as the 'Adam function' may be used to explore an area of 'solution space' - in graphical terms, the area which is bounded by the locus of all possible valid solutions to the mathematical arguments.

In practice, the network is 'trained' by presenting it with large numbers of known examples, usually many thousands. The nature of the control of a neural network allows the network to adjust the weights and biases of each node to 'grow' the network in a direction which more closely resembles the allowed solutions. In our case we present the network with thousands of different barcodes, some good, some bad and some downright ugly! Our network adjusts itself until, eventually, it can discriminate between a good and a badly printed barcode. With further training, it can determine if no bar code is present at all. That sounds obvious and easy, but in practice many factors in the lab such as overhead illumination, proximity to windows, time of day and reflections from adjacent populated wells can all cause the imager to perceive a barcode where in fact, none is present.

Once trained, the network must be verified. A Verification Set of data, approximately 20% of the size of the Training Set, containing new, relevant images is used to score the performance of the network.

In practice, networks can have a tendency to be 'over-trained', such that they are marvellous at processing the training set but fail miserably when presented with new data. To mitigate this 'over-training' problem a number of techniques can be employed: such as augmenting the training data and forcing the partially trained network to lose current values of weights and biases (known as 'dropout'). Data augmentation allows us to perform transformations, rotations and other image pre-processing operations on our Training Set. By augmenting the data it is ensured that the variety of images used by the network for training is increased, hence lowering the tendency to create an 'over-trained' network. Dropout is a second process to help reduce 'over-training'; where a percentage of the trained weights/biases are randomly removed between training epochs. The neural network then must re-learn the weights when it is presented with another iteration of the Training Data. By preventing 'over-training', when the network is presented with brand new data, say from the Verification Set, it is able to score just as highly when compared to the Training Set. Once a network is able to score highly against brand new data it is considered ready for deployment.

There are a number of frameworks that can be used for the development and training of neural networks; such as Tensorflow, PyTorch, Torch and Keras for example. The frameworks run using common programming languages such as Python. The frameworks provide an end-to-end open-source platform for machine learning with a suite of tools and libraries that lets developers easily build and deploy Machine Learning powered applications, such as scientific image analysis.




Figure 2. Neural networks combine a series of mathematical logic operations in a particular sequence to achieve the desired result.

Ziath has chosen to implement such a powerful neural network in the latest control software, DP5, for their CMOS camera-based 2D-barcode tube readers. This new AI functionality underpins the Ziath software's empty-well, or missing tube, determination using exactly the method described above. The output is a determination that there either is, or is not, a tube present in each location.

The R&D team at Ziath are now working to move the neural network onto an FPGA. This is a Field-Programmable Gate Array. Essentially this device combines billions of transistors into a platform from which custom architectures can be created to solve the user specific problem. FPGAs have a number of advantages over other processing technologies; they provide a flexible platform with the ability to update its configuration in the field, they yield a very high performance compared to other processors, and they allow for rapid development turn-around compared to custom silicon. Translating the neural network in this way will determine the architecture, size, capacity and cost of the FPGA. Even so, the modest device intended for Ziath's application can still run many times faster than the embedded PC used to address it, so an even bigger, more powerful network could be accommodated at some future point, if required. This FPGA is carrying out giga-operations every second running at a clock speed of 125MHz, which enables it to analyse a 96-position SBS format rack in a little under 8 milliseconds. It could run faster on a larger FPGA, but the cost goes up commensurately, and the system is limited by other bottlenecks such as data transfer between devices.

As noted above, Machine learning is very good for certain image analysis problems, so called machine vision. It can identify certain telling features at very high resolution which the human eye would find difficult to resolve. In addition, deconvolution of overlying data such as spectra or chromatograms lend themselves to an AI solution. In the lab potentially any large data set which requires pattern-recognition could be undertaken by AI. In other areas, methods are in development that will reduce the amount of training data need to condition a neural network, making their application easier and quicker. The promise of AI in the lab is faster, more sensitive and cheaper data-processing than is possible today.

Dr Alex Beasley CEng. MIET MIEEE is Senior Electronic Design Engineer with Ziath Ltd at its Cambridge, UK, headquarters.

