What are Neural Networks?

Modeled off the networks in our own brains, Neural Networks, or Deep Learning as it is sometimes known, is a branch of Machine Learning capable of efficiently learning from large amounts of data. This data can come in a variety of forms including images, number collections, words & even time series/sequential data such as video, sentences or any other type of data sequences.

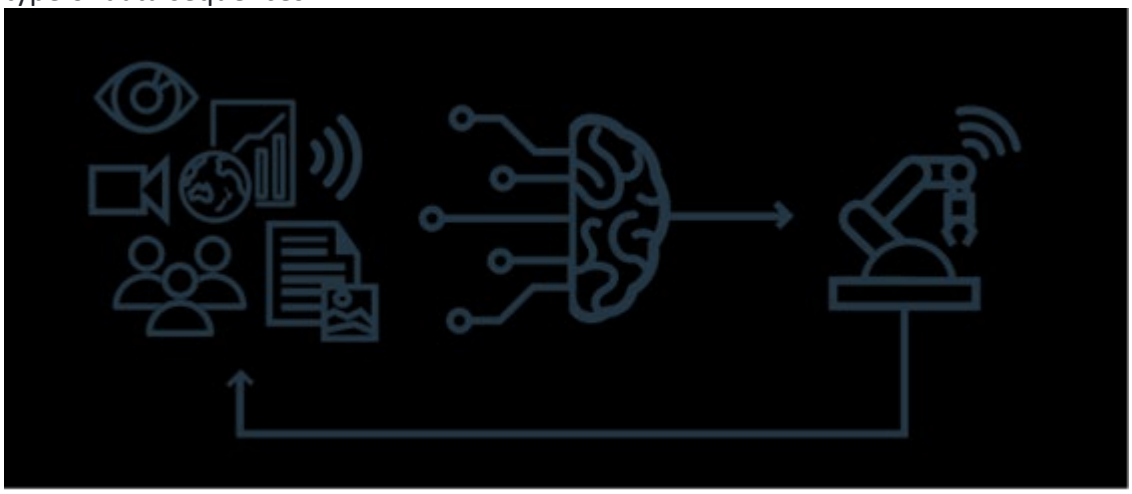

IMAGE: machinelearning.png

Despite being developed in the 1940's Neural Networks have only started to flourish in the last decade thanks to the availability of big data and processing power. What is more is you have probably already come across and used Neural Networks without even knowing. They are the reason spam emails go straight to the junk folder and the same reason you get Netflix and YouTube recommendations that align well with what you want to watch or hear.

How is Machine Learning used?

We can also see this branch of Machine Learning used in fraud detection, e-commerce and other cost vs need based business models, for example in Uber's surge pricing.

The Uber app uses Neural Networks to adjust its algorithm in real time in response to the huge amounts of data constantly flowing through the app. Everything from time of day, location, city, average distance traveled, event information, traffic information, traffic patterns, car ownership and income for geographic area is all considered, calculated and acted on. With more users than drivers, the algorithm adjusts pricing to try to entice offline drivers to want to start taking passengers. At the same time, this encourages passengers to hold off on their journey if they can. Eventually the supply and demand balance out and the app will eventually return to more regular pricing. All decisions made thanks to the learnings, calculations and then decision-making abilities of a computer.

How do Neural Networks work?

Neural Networks enable machines to learn from data, so the only human involvement required is to provide the data, both for training and for use in the real world. The algorithms that train the networks are 'optimisation algorithms' that are designed to efficiently minimise cost functions by adjusting (up to) millions of model parameters, the outcome is that key features in the data that correlate with the desired outcome are automatically detected. Being able to process huge amounts of data and use it to create effective decisions, as with Uber, is what makes Machine Learning so effective.

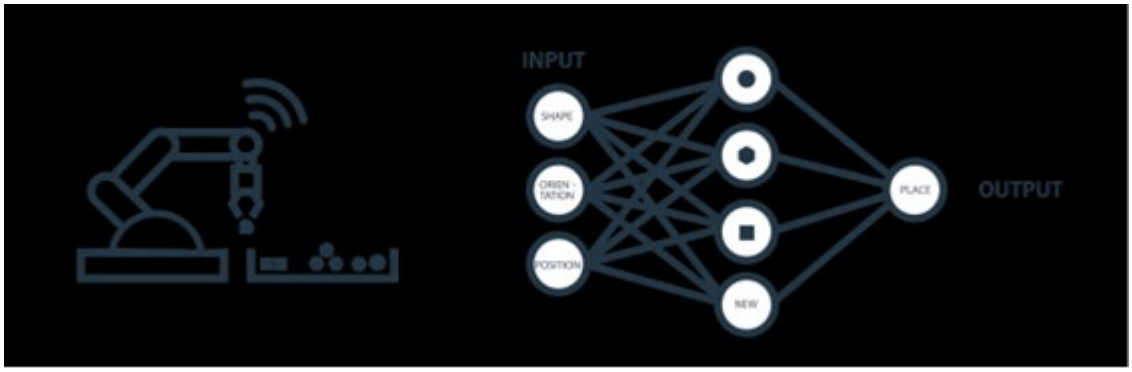

IMAGE: neural-networks-png-03

Machine Learning removes the need for specialist programs from having to be written that attempt to capture relevant features in data to solve some problem. Instead, Machine Learning methods will learn from the data which features are relevant to meeting the desired outcome. Instead of focusing on a few specific features (as humans do), Machine Learning methods will use all of the available information in the data to make predictions. This usually results in a more robust system because when a critical feature is missing, the human designed – specialist programs tend to fail as they tend to focus on only a few obvious features, whereas Machine Learning tends to capture even the most subtle features from which the model can still make accurate predictions.

Neural Networks work better than these specialist programs as they take a series of inputs, process these and present you with an output, which over time and with more data becomes increasingly accurate. Each network is simply a mathematical function. Put simply there are three parts to a Neural Network, an input, an output & inner layers in the middle. The network can be made up of millions of smaller functions called neurons, the inputs and outputs of these neurons connect together to make the inner layers of the network. Its then these inner layers of the network that process the input data to determine the output, with each inner connection having a weight associated with it that determines how much the input at this connection will influence the output. The value of the weights is what is learned during the training process.

Training & Network Types

Setting up a Neural Network

Of the three main parts that make up a Neural network, the inner layer, which processes the input data, determines the output and output type. Outputs can be split into three main forms.

- Categorical – to discriminate between different classes only one class can be true
- Multiple Category – Example, network for detecting dogs and hats. Both classes can be true if a dog is wearing a hat.
- Regression – Predicting the location of a feature in an image, e.g., finding a car in a photo

There are three ways Neural Networks can be trained; supervised learning, unsupervised learning and reinforcement learning. During the training phase, adjustments to the parameters of the network minimize the prediction error. Once your Neural Network is set up and trained, it can then be used to make predictions on new, previously unseen data.

Supervised Learning

Supervised learning is a type of training that requires a human (supervisor) to provide examples of input data and desired output (the label), allowing the machine to learn how to

process the input to get a desired output. An example of this type of training would be giving the network a selection of images and telling it that the input is a selection of different cuts of meat, with hindquarter and ribs being the labels. The network then learns the features that distinguish between these two types of cuts. For instance, a possible feature of a rib cut could be the sequence of parallel lines (the ribs), for the hindquarter it could be the presence of a H-bone. In many applications obvious or key features may be obscured in the image. Therefore, a well-trained Neural Network will use many other features to draw its conclusion. Once sufficiently trained, the network will then be able to classify any new image inputs into cuts of ribs and cuts of hindquarter.

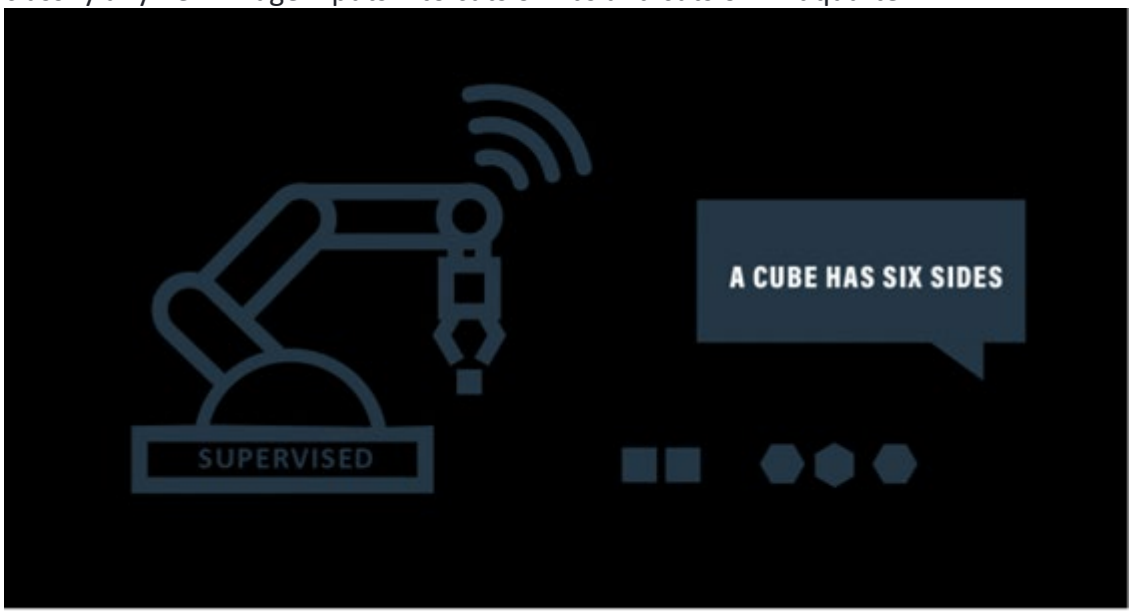

IMAGE: robot-learning-03

Unsupervised Learning

Another type of Neural Network training is unsupervised learning, where the network will be trained with unlabelled data. This form of learning is normally used for classification tasks. For example, if you were to tell the network there are two distinct classes to find, the model will try to discriminate the data into two groups I.e classifying visitors to an office into two categories. There is no way of knowing what features the network will use to discriminate into groups. There are a number of features it can take into account including; gender, skin colour, weight, height etc. But, because it does not know what data belongs with what label, it may separate data into groups by using any combination of features and not necessarily the ones we are after - a very real problem in unsupervised learning.

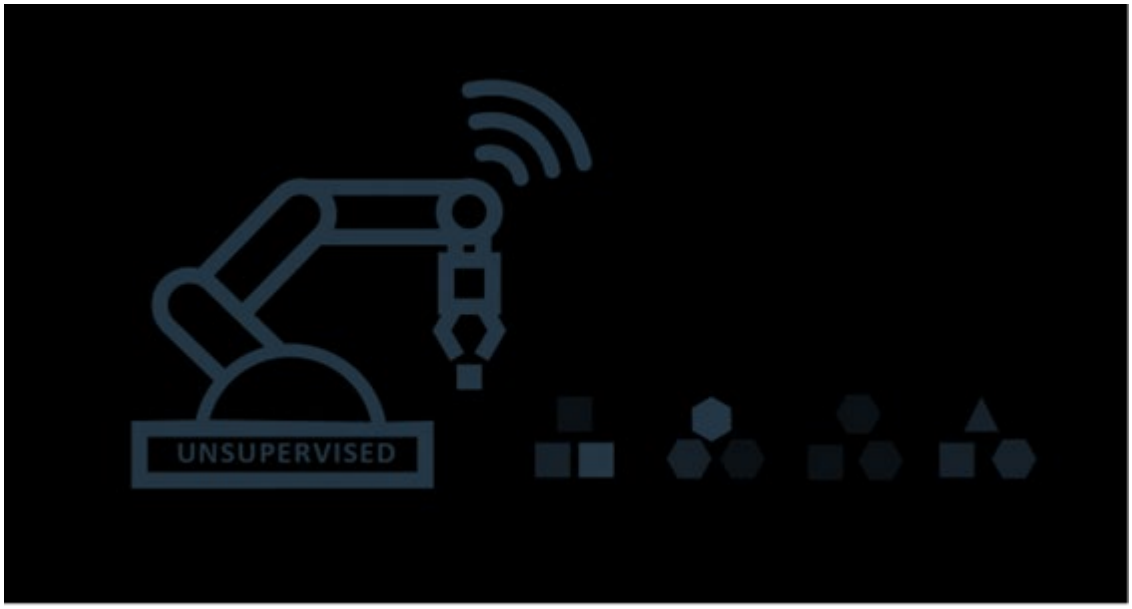

IMAGE: robot-learning-01

Reinforcement Learning

The third type of learning is reinforcement learning and works on the principle of feedback. As an example, if a machine was to learn how to race a car in a game, it would be provided feedback automatically based on how far the car drove without crashing or how quickly it managed to complete the race. The machine would then make thousands of copies of the network, each with small mutations to the parameters and will then test each one of those networks. Those that survive the furthest have mutations made from them, while the unsuccessful tests drop off and so on. The cycle continues until we have a machine that is excellent at racing a car in a particular game. No humans are involved and the machine can sit there and learn on its own. The machine that beat humans at Go – the Chinese strategy game once considered too advanced for computers to beat, was won by a machine that had played itself millions of times, testing and continuing with strategies that worked and removing strategies that didn't. It managed this all within 4 hours, much faster than the upwards of 10000 hours its human counterparts would have had to play to achieve a similar feat.

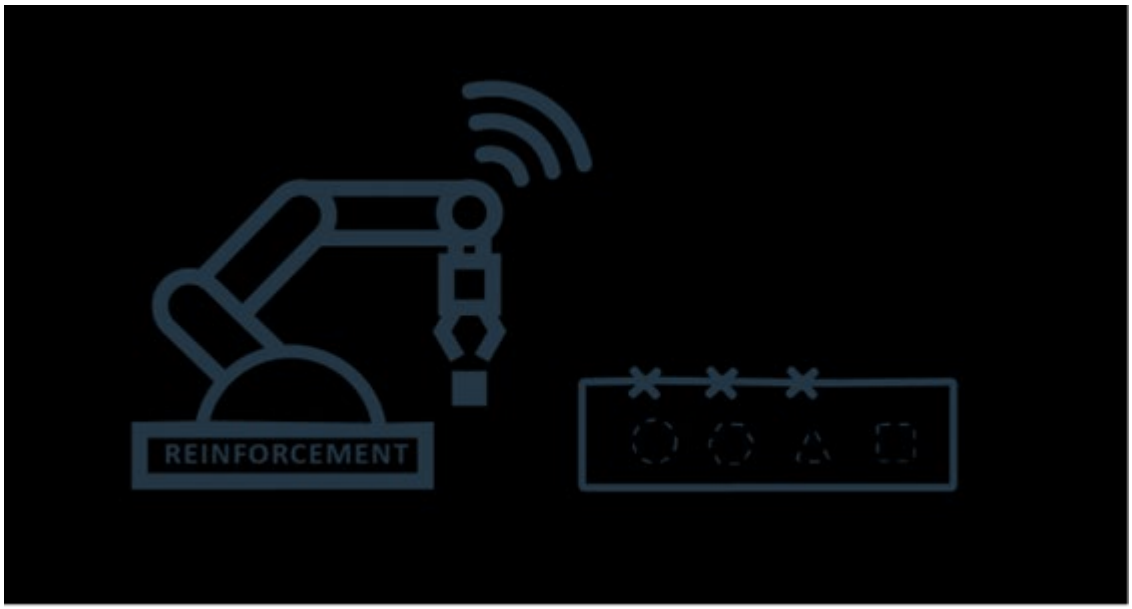

IMAGE: robot-learning-02

Neural Networks in Automation

“Applications already making an impact”

Machine Learning is very useful for vision technology – where camera guided robot automation is required. Machine Learning, when combined with automation is useful in areas such as Picking where a robot may come across new objects regularly or where programming a robot for all possible scenarios is impractical. Materials Handling and Industrial Automation will both see huge benefits from incorporating Machine Learning into systems, helping streamline and improve existing practices within those industries. Data is the last major area that will see huge impact from Machine Learning, as it will streamline automated processes in understanding and processing huge amounts of data, as well as finding patterns and trends.

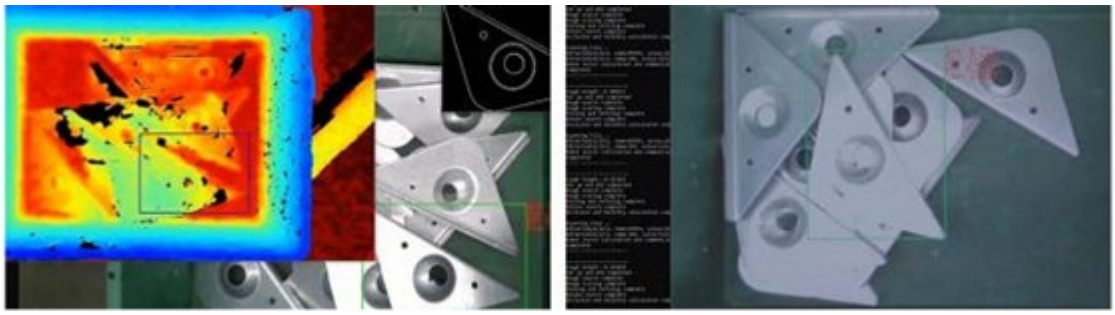

IMAGE: Bin-Picking2

This Bin Picking Robot from Scott, uses RGB data & Machine Learning to detect and pick brackets

Being able to make decisions in real time is one of the big reasons why Machine Learning will have such a huge influence on how automation is approached in a variety of industries. In healthcare, symptoms and patient data can be input for analysis, which a doctor can then review as findings from the network output, thus speeding up complex and delicate diagnostics. Neural Networks have also proved to perform better in detecting a range of illnesses in scans when trained correctly. Care industries will also see advances in assistive technology such as companion robots. Machine Learning means these robots will be able to learn and respond better to their environments, providing a higher level of assistance and care for those they help.

The ability to adapt to make appropriate & accurate decisions based on diverse data and inputs is what makes Machine Learning such a valuable asset in Materials Handling. Complex production lines and other variable processes can be automated more efficiently with less error margins and down time. For example, a robot working in a cold storage room picking & placing boxes may come across boxes with varying degrees of frost or ice. Machine Learning models trained on a large data set will be able to cope with the diversities presented, enabling the robot to recognize the unfamiliar shape as product and then correctly pick the box.

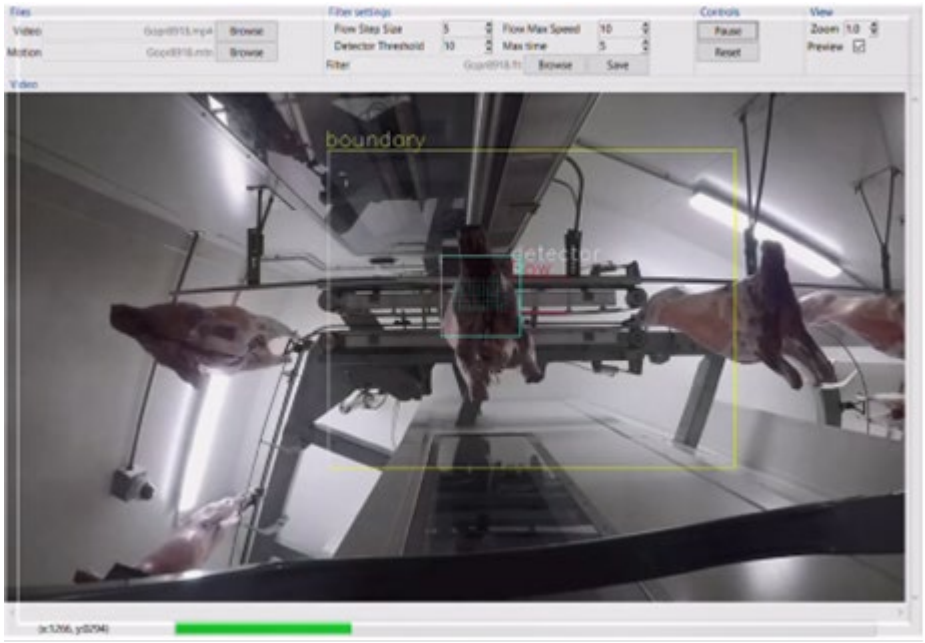

IMAGE: Vision-system-is-used-for-data-collection-in-forequarter-cut-paddywack-detection

Vision system is used for data collection in forequarter cut Paddywack detection

Scott has used Machine Learning in large-scale meat processing. In this area, Neural Networks outperformed classical vision techniques because the models were capable of adapting to the immense diversity of data in the meat industry, thus improving outcomes for producers. One such use was in predicting positions of a specific cut on a piece of meat. This cut was more difficult than most as its position is often partially obscured from the field of view of the vision system because of the diversity in the size and shape of the meat. These challenges meant that creating this cut accurately, especially in an automated environment, was difficult. By using a Neural Network trained on a diverse data set, the system has attained a cutting accuracy never before seen in the industry.

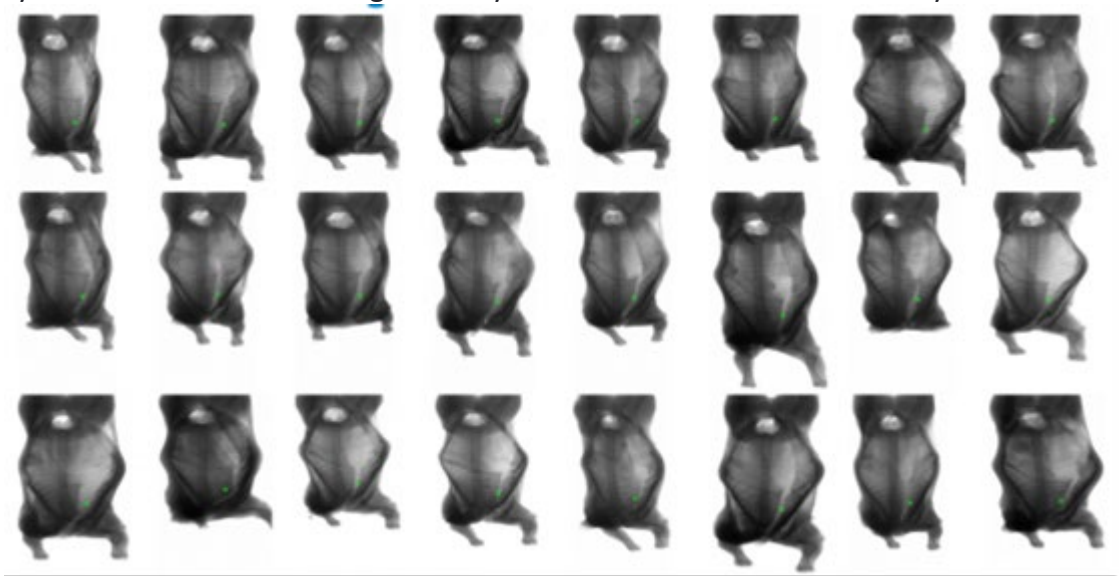

IMAGE: Rib-Prediction-with-Convolutional-Neural-Network

Convolutional Neural Network predicts rib placement using X-Ray data for accurate cutting (link to X-Ray Grading)

Similar Machine Learning was used, in combination with vision technology, for a mining application in which Scott was involved. This application saw Neural Networks trained to distinguish between contaminants and gold ore. The machine learnt the difference between ore pixels and other pixels, helping it determine through analysis of the data, which pixels were statistically more likely to be ore and those more likely to be contaminants. For this application, Machine Learning was ideal due to both the massive diversity and the amount of data available for training. Once trained it could correctly detect in real time, deciding what was and was not a contaminant, something that would be impractical using traditional Computer Vision techniques. The resulting decisions and data of this application could additionally work as future reinforcement training for further detections.

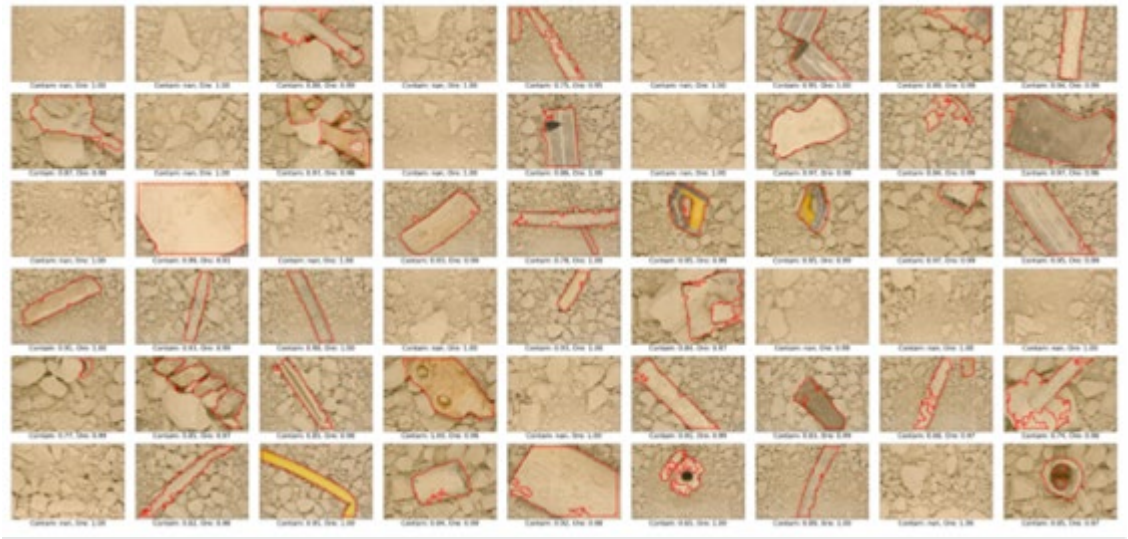

IMAGE: Contaminant-Detection-in-Gold-Ore
Contaminant detection in gold ore, contaminant's shown with a red outline

While we have yet to fully realise or unleash the full potential of Machine Learning in industrial robotics and automation, advances in detection, analysis and problem solving have already simplified processes in a multitude of industries, including materials handling, meat processing and mining. Applications that utilise Machine Learning currently are making an impact and the development and learnings gained from these existing systems will continue to shape how we use Machine/Deep Learning and Neural Networks in robotics & automation going forward - making the future both promising and exciting.

For additional information on how neural networks work watch Mark Rober (ex NASA engineer) explains and builds a Neural Network to predict baseball signals and Machine Learning and Artificial Intelligence Crash Course.

For more resources on Neural Networks in Automation see:
Machine Learning in Robotics HYPERLINK "https://emerj.com/ai-sector-overviews/machine-learning-in-robotics/"– HYPERLINK "https://emerj.com/ai-sector-overviews/machine-learning-in-robotics/" 5 Modern Applications by EMERJ, an AI Research & Advisory Company
Applying Artificial Intelligence and Machine Learning in Robotics from the Robotics Industries Association
Entering a New Robotics Age with Machine Learning by Electronic Design