PEGASUS ONE

ARTIFICIAL INTELLIGENCE TOOLS AND TECHNIQUES EXPLAINED

1440 N HARBOR BLVD #900, FULLERTON, CA, 92835 | 10880 WILSHIRE BLVD., # 1101, LOS ANGELES, CA, 90024

CONTENTS

01 AI Technology Landscape

02 AI Techniques explained

04 Artificial Neural Networks

05 Natural Language Processing

THE AI TECHNOLOGY LANDSCAPE

The internet enabled global disruption; new competition emerged that replaced or challenged the market leaders in many sectors. It opened doorways and created opportunities for change on a scale not seen before. Established enterprises who had assumed that their size, infrastructure, legacy market position and supply chains posed too significant a barrier to entry for hi-tech vstartups and digital natives to threaten their dominance had to think again.

Successive waves of internet technologies, the cloud, then data analytics have washed over a growing number of industries from media and content to consumer electrical goods to even the toy and game megaretailers. Each technological innovation enabled agile digital native startups to disrupt marketplaces, removing barriers to cost, infrastructure and business risk. The latest reviews of the technology and business sector analysts all agree that AI is the next big wave of innovation and disruption. Key here is not just the technology itself, but how it is accessed.

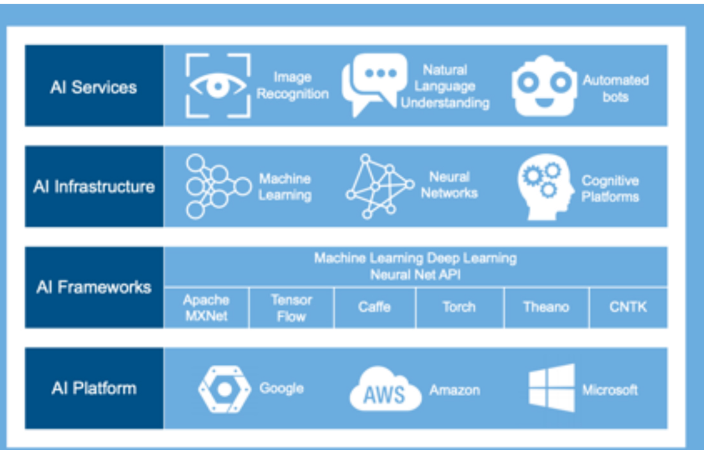

GOOGLE, AMAZON, MICROSOFT AND OTHERS, OFFER CLOUD-BASED AI TOOLS AND SERVICES, ALL OF WHICH ARE EASILY INTEGRATED INTO NEW DIGITAL PRODUCTS AND SERVICES

High capital cost is no longer a barrier to entry to state of the art AI. The digital giants from the previous generation of disruption, such as Google, Microsoft, Facebook all offer AI toolkits as free, open source software for immediate download and use. All any new disruptor needs.

is an internet connection, an original business idea and the skills to build AI.

Similarly, the incremental cost charged by the big cloud platform providers for accessing their AI frameworks is kept low.

The objective is to convince their target market, of developers and solution providers, to use their ecosystems to host high-value customer-facing digital services. Moreover, these platforms are continually extending their AI capability, productivity tools and speed of AI tool deployment to become more attractive and differentiated in an increasingly crowded marketplace.

SO, WHAT DOES THIS MEAN FOR DIGITAL TRANSFORMERS?

The scale and speed of these changes mean that any competitor can get fast, direct access to best-in-class AI technology. Digital giants, innovative start-ups, your current competitors pursuing their own transformation plans, and wholly new players yet unseen, have immediate access to mature, enterprise-ready and internet scalable AI capabilities. This combined with agile working practices, rapid prototyping and a culture that integrates AI into the daily way of working means that previously untouched markets are susceptible to incredibly fast change.

AI Techniques explained:

Some fundamental AI techniques used: Heuristics, Support Vector Machines, Neural Networks, Markov Decision Process, and Natural Language Processing.

Heuristics: Suppose we have coins with the following denominations: 5 cents, 4 cents, 3 cents, and 1 cent and we need to determine the minimum number of coins to get 7 cents. In order to solve this problem we can make use of a technique called "heuristics".

For some problems, tailored heuristics can be designed that exploit the structure present in the problem. An example of such a tailored heuristic would be a greedy heuristic for the above mentioned coin changing problem. Now a greedy heuristic would be to always choose the largest denomination possible and repeat this until we get to the desired value of 7. In our example, that means that we would start with first selecting one 5 cent coin. For the remaining 2 cents, the largest denomination we can choose is 1 cent, leaving us with the situation where we still have to cover 1 cent for which we again use 1 cent.

So our greedy heuristic gives us a solution of 3 coins (5, 1, 1) to get to the value of 7 cents. It can be easily seen that another, better, solution of only 2 coins exist using the 3 and 4 cent coins. While the greedy heuristic for the coin changing problem does not provide the best solution for this particular case, in most cases it will result in a solution that is acceptable.

Besides such tailored heuristics for specific problems, also certain generic heuristics exist. Just like neural networks, some of these generic heuristics are based on processes in nature. Two examples of such generic heuristics are Ant Colony Optimization and genetic algorithms. The first is based on how simple ants are able to work together to solve complex problems and the latter is based on the principle of survival of the fittest.

A typical problem where heuristics are applied to find acceptable solutions quickly is vehicle routing, where the objective is to find routes for one or more vehicles that have to visit a number of locations.

Support Vector Machines

The question whether an email is spam or not spam is an example of a classification problem. In these types of problems, the objective is to determine whether a given data point belongs to a certain class or not. After first training a classifier model on data points for which the class is known (e.g. a set of e-mails that are labelled as spam or not spam), you can then use the model to determine the class of new, unseen data-points. A powerful technique for these types of problems is Support Vector Machines (SVM).

The main idea behind SVM is that you try to find the boundary line that separates the two classes, but in such a way that the boundary line creates a maximum separation between the classes. To demonstrate this, we will use the following simple data for our classification problem (Figure 1 below).

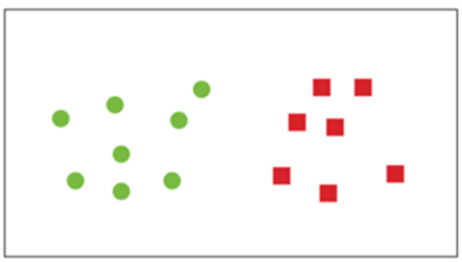

Figure 1

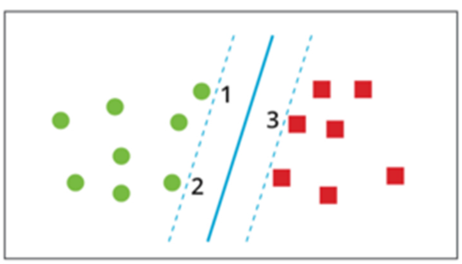

Figure 3

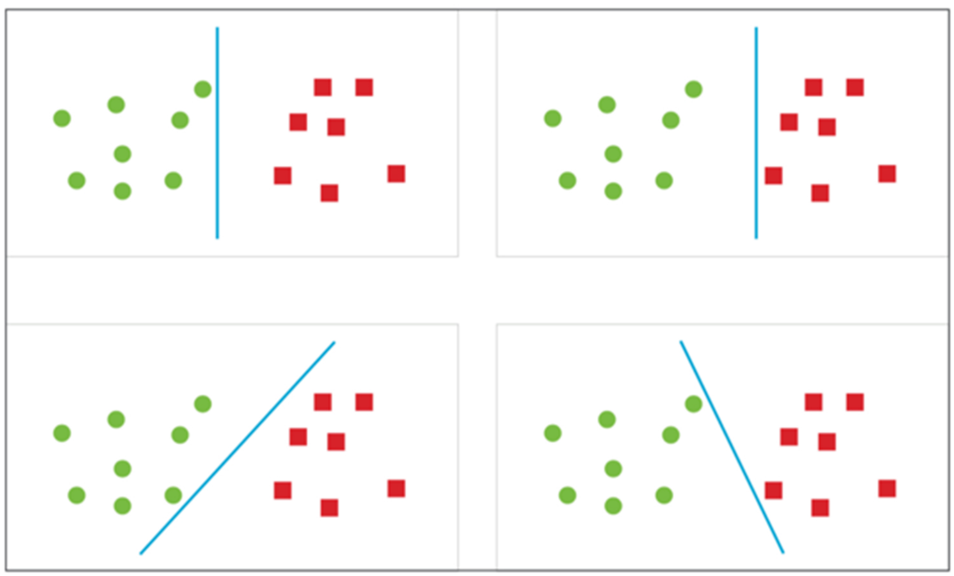

Figure 2

In this example, the green circles and the red squares could represent two different segments in a total set of customers (e.g. high potential and low potential), based on all kinds of properties for each of the customers. Any line that keeps the green circles on the left and the red squares on the right is considered a valid boundary line for the classification problem. There is an infinite number of such lines that can be drawn and 4 different examples are presented on top (Figure 2).

As stated before, with SVM you try to find the boundary line that maximizes the separation between the two classes. In the provided example, this can be drawn as Figure 3.

The two dotted lines are the two parallel separation lines with the largest space between them. The actual classification boundary that is used will be the solid line exactly in the middle of the two dotted lines. The name Support Vector Machine comes from the data points directly on either of these lines are the supporting vectors. In our example, we had 3 supporting vectors. If any of the other data points (i.e. not a supporting vector) is moved slightly, the dotted boundary lines are not affected. However, if the position of any of the supporting vectors is slightly changed (e.g. data point 1 is moved slightly to the left), the position of the dotted boundary lines will change and therefore the position of the solid classification line also changes.

In real life, data is not as straightforward as in this simplified example. We normally work with much more than two dimensions. Besides having straight separation lines, the underlying mathematics for an SVM also allows for certain type of calculations or kernels that result in boundary lines that are non-linear. SVM classification models can also be found in image recognition, like face recognition or converting handwriting to text.

Artificial Neural Networks

Artificial Neural Networks (ANN) can be described as processing devices that are loosely modelled after the neural structure of a brain. The biggest difference between the two is that the ANN might have hundreds or thousands of neurons, whereas the neural structure of an animal or human brain has billions.

The basic principle of a neural structure is that each neuron is connected with a certain strength to other neurons. Based on the inputs taken from the output of other neurons (also considering the connection strength) output is generated which can be used as input again by other neurons, see Figure 4 (left). This basic idea has been translated into an artificial neural network by using weights to indicate the strength of the connection between neurons. Furthermore, each neuron will take the output from the connected neurons as input and use a mathematical function to determine its output. This output is then used by other neurons again.

Figure 4: Graphical representation of a biological neuron (left) and an artificial neuron (right)

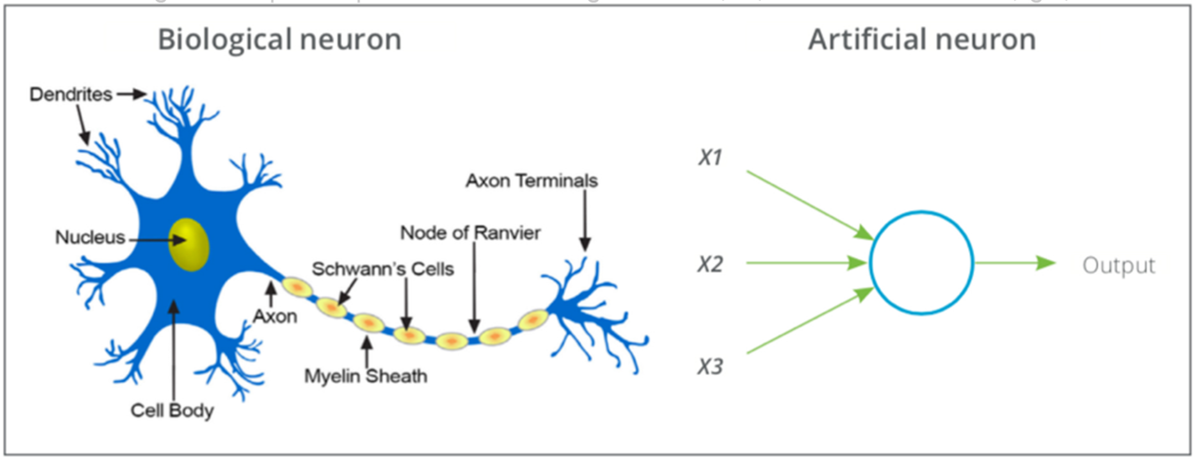

The neurons of the ANN can be structured into several layers. Figure 5 shows an illustrative scheme of such layering. This network consists of an input layer, where all the inputs are received, processed and converted to outputs to the next layers. The hidden layers consist of one or more layers of neurons each passing through inputs and outputs. Finally, the output layer receives inputs of the last hidden layer and converts this to the output for the user.

Figure 5: Schematic of a connected ANN

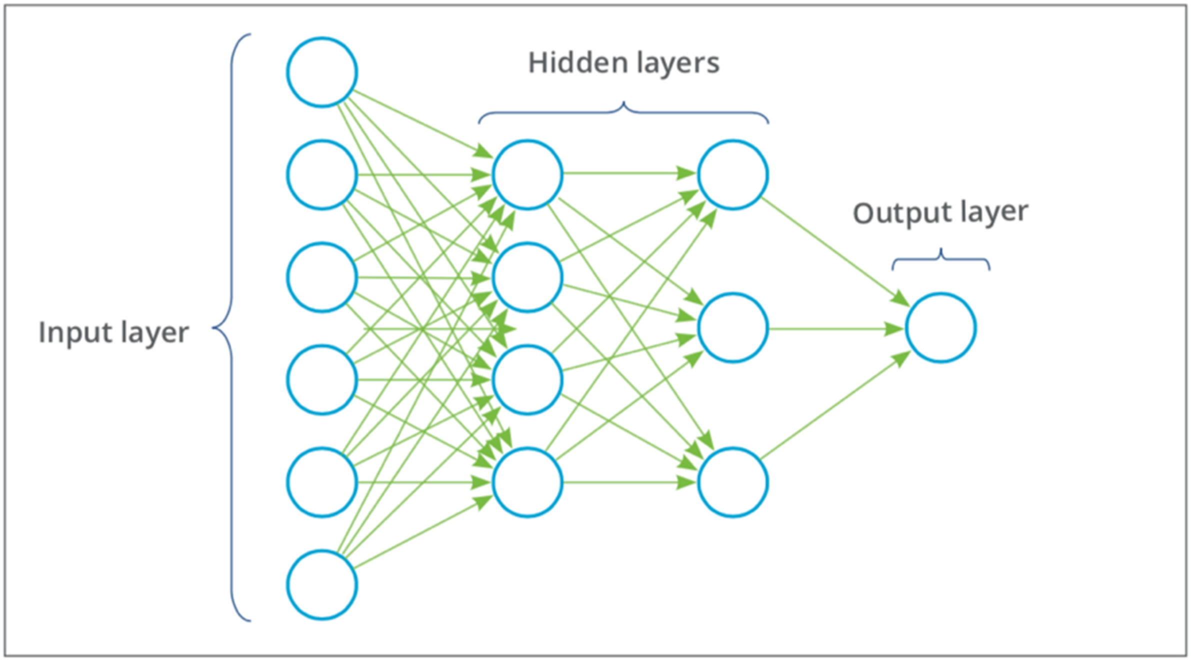

Natural Language Processing

Natural Language Processing, or NLP in short, is a term for everything from speech recognition to language generation, each requiring different techniques.

One of the first steps of NLP is lexical analysis, here a technique is used called Part-of-Speech (POS) tagging. With this technique, every word is tagged to correspond to a category of words which have similar grammatical properties, based on its relationship with adjacent and related words. Not only words are tagged, but also paragraphs and sentences. Part-of-speech tagging is mainly done with statistical models, which give probabilistic results instead of hard if-then rules, and is, therefore, more capable of processing unknown text. Also, they can cope with the possibility of multiple possible answers, instead of only one.

Concluding, the techniques used within the domain of Artificial Intelligence are, when you dive into them, just advanced forms of statistical and mathematical models. All these models cleverly put together give us tools to compute tasks, previously thought to be reserved for humans. In subsequent blogs we will dive deeper in business applications, some associated technology trends, and the top 5 risks and concerns.

Top Applications of AI: Image recognition, Speech recognition & automatic translation.

Pegasus One is a privately held, professional IT services company and a Microsoft and Amazon Partner with its U.S. headquarters in the Orange County area, and satellite offices in Mexico and India. We provide custom software development and consulting services in various technologies including Microsoft Solutions, IoT, Mobile and BI. We are a CMMI Level 5 company with high focus on quality and ROI.

email: info@pegasusone.com
phone: +1(714) 485-8104
corporate office: 1440 N Harbor Blvd #900, Fullerton, CA, 92835
sales office: 10880 Wilshire Blvd., # 1101, Los Angeles, CA, 90024