Deloitte.

Artificial Intelligence

March 2018

What is AI? What is Machine Learning? What is Cognitive Analytics? How do these terms relate, or differ, from one another? In general terms, AI refers to a broad field of science encompassing not only computer science but also psychology, philosophy, linguistics and other areas.

Contents

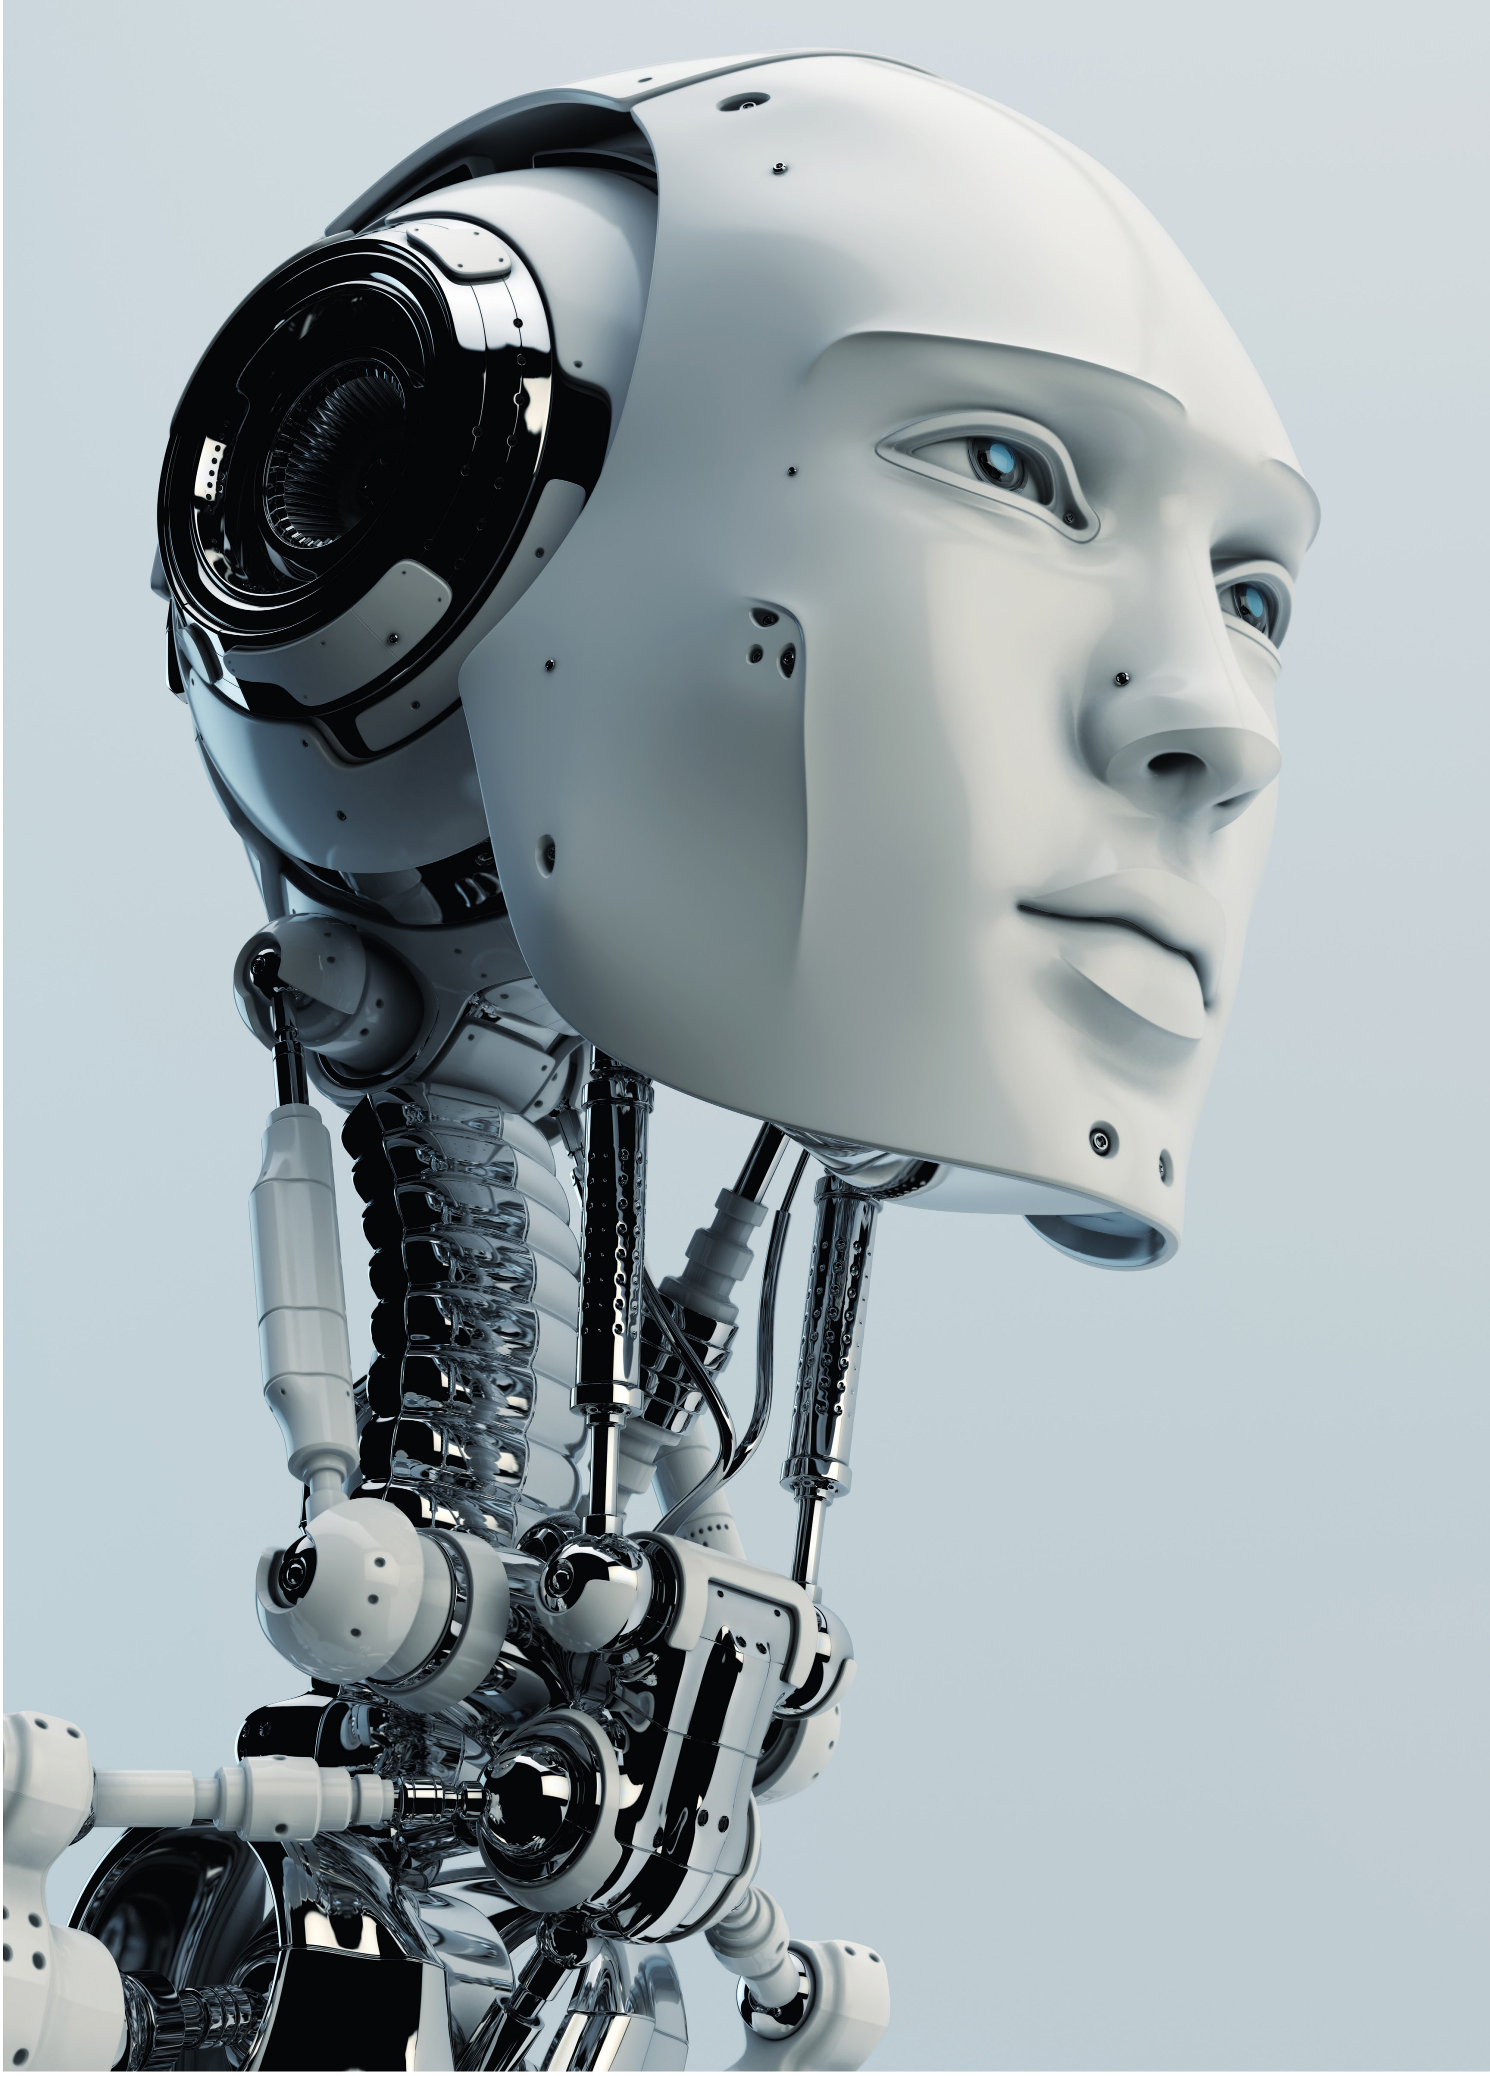

Artificial Intelligence defined

The topic of Artificial Intelligence is at the top of its Hype curve[1]. And there are many good reasons for that; it is exciting, promising and a bit scary at the same time. Various publications are claiming that AI knows what we want to buy, it can create Netflix series, it could cure cancer and it may eventually take our jobs or even destroy mankind.[2]

The problem and at the same time opportunity with AI is that it's not very well defined. If we would show the navigation system of our car to someone living in 1980, he or she would probably consider it as a form of Artificial Intelligence, whereas we nowadays would probably not. We are seeing the same with speech and image recognition, natural language recognition, game engines and other technologies that are becoming more and more common and embedded in every-day technology.

On the other hand, various technology solution providers are taking the opportunity to rebrand their existing solutions to AI, to take advantage of the huge hype and that the market is experiencing and the resulting press coverage. If we have a built a machine learning model that predicts customer demand, a solution that has been existing for years, we would have called it "data mining" in the past and we now see it rebranded as "artificial intelligence". This is adding to the confusion and may very well lead to inflated expectations.

Nevertheless, recent developments in AI are impressive and exciting. But also overestimated and misunderstood. In order to split hype from reality and help forming a view on this market, we will publish a series of articles explaining the world of AI, zooming in to the techniques that are associated with AI, the most appealing business applications and potential issues we can expect.
In this first article we will start with the beginning, by explaining AI and associated terms in five definitions. What is AI? What is Machine Learning? What is Cognitive Analytics? How do these terms relate, or differ, from one another?.

Artificial Intelligence (AI)

In general terms, AI refers to a broad field of science encompassing not only computer science but also psychology, philosophy, linguistics and other areas. AI is concerned with getting computers to do tasks that would normally require human intelligence. Having said that, there are many point of views on AI and many definitions exist. Below some AI definitions which highlight key characteristics of AI.

"AI refers to a broad field of science encompassing not only computer science but also psychology, philosophy, linguistics and other areas"

[1] https://www.gartner.com/document/3380751/
[2] http://www.bbc.com/news/technology-37713629
https://futureoflife.org/ai-open-letter/

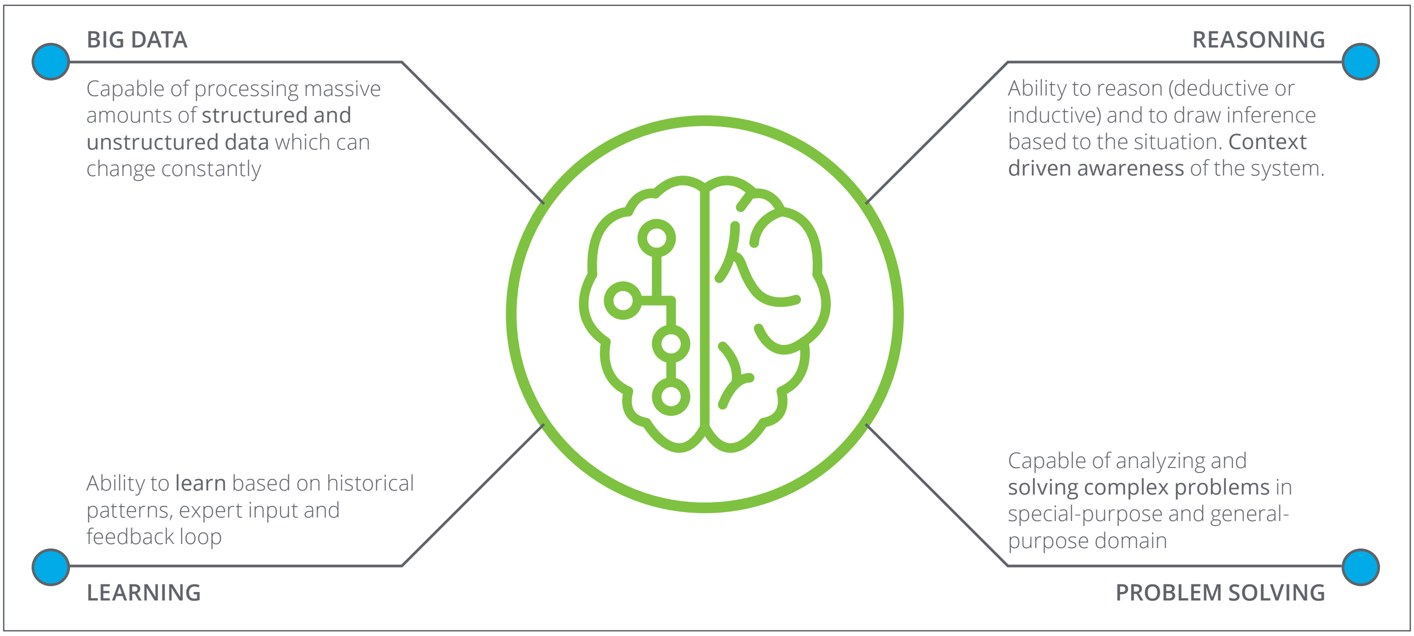

Figure 1: Key characteristics of an AI system

Some general definitions:

- "Artificial intelligence is a computerized system that exhibits behavior that is commonly thought of as requiring intelligence."[3]
- "Artificial Intelligence is the science of making machines do things that would require intelligence if done by man." [4]

The founding father of AI Alan Turing defines this discipline as:

- "AI is the science and engineering of making intelligent machines, especially intelligent computer programs". [5]

In all these definitions, the concept of intelligence refers to the ability to plan, reason and learn, sensing and building some kind of perception of knowledge and communicate in natural language.

Narrow AI vs General AI

A chess computer could beat a human in playing chess, but it couldn't solve a complex math problem. Virtually all current AI is "narrow", meaning it can only do what it is designed for. This means for every problem, a specific algorithm needs to be designed to solve it. Narrow AI are mostly much better at the task they were made for than humans, like face recognition, chess computers, calculus, translation. The holy grail of AI is a General AI, a single system that can learn about any problem and then solve it. This is exactly what humans do: we can specialize in a specific topic, from abstract maths to psychology and from sports to art, we can become experts at all of them.

An AI system combines and utilizes mainly machine learning and other types of data analytics methods to achieve artificial intelligence capabilities.

---

[3] Preparing for the Future of Artificial Intelligence, NSTC, 2016
[4] 6. Raphael, B. 1976. The thinking computer. San Francisco, CA: W.H. Freeman
[5] http://www-formal.stanford.edu/jmc/whatisai/node1.html

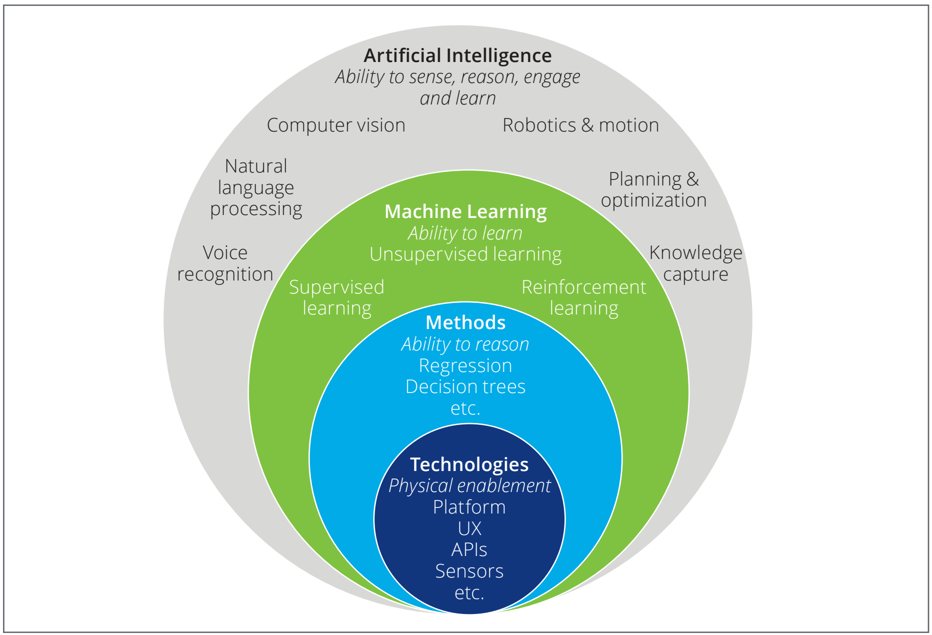

Figure 2: relation between AI, Machine Learning and underlying methods and infrastructure

Machine Learning

Machine learning is the process whereby a computer distills meaning by exposure to training data[6]. If for example you want an algorithm to identify spam in e-mails, you will have to train the algorithm by exposing it to many examples of e-mails that are manually labeled as being spam or not-spam. The algorithm "learns" to identify patterns, like occurrence of certain words or combination of words, that determines the chance of an e-mail being spam. Machine learning can be applied to many different problems and data sets. You can train an algorithm to identify pictures of cats in photo-collections, potential fraud cases in insurance claims, transform handwriting into structured text, transform speech into text etc. All these examples would require labeled training sets.

Depending on the technique used, an algorithm can improve itself by adding a feedback loop that tells it in which cases it made mistakes.

The difference with AI however is that a machine learning algorithm will never "understand" what it was trained to do. It may be able to identify spam, but it will not know what spam is or understand why we want it to be identified. And if there is a new sort of spam emerging, it will probably not be able to identify it unless someone (human) re-trains the algorithm.

Machine learning is at the basis of most AI systems. But while a machine learning system may look "smart", in our definition of AI it is in fact not.

Cognitive Analytics

Cognitive Analytics is a subset of A.I. that deals with cognitive behavior we associate with 'thinking' as opposed to perception and motor control. Thinking allows an entity to obtain information from observations, learn and communicate.

A cognitive system is capable of extracting information from unstructured data by extracting concepts and relationships into a knowledge base. For example, from a text about Barack Obama, the relations from Figure 3 can be extracted using Natural Language Processing. 80% of all company data is unstructured and current Cognitive Analytics systems can search all of it to find the answer to your question.

[6] Stephen Lucci, 2016, Artificial intelligence in the 21st century : A living introduction

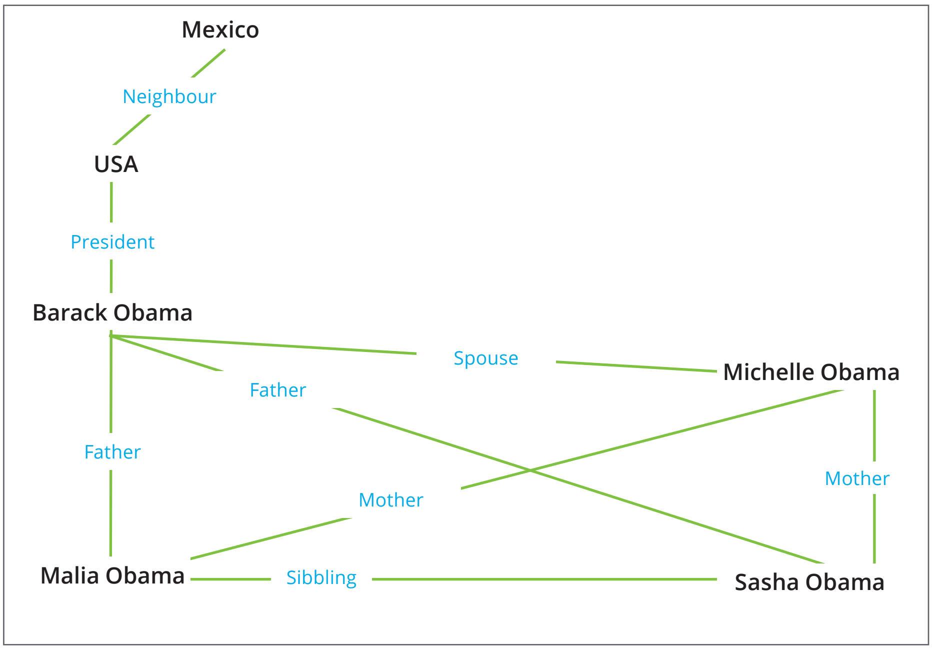

Figure 3: A knowledge base extracted from text

Learning enables the Cognitive System to improve over time in two major ways. Firstly, by interacting with humans, and obtaining feedback from the conversation partner or by observing two interacting humans. Secondly, from all the data in the knowledge base, new knowledge can be obtained using inference.

Another important aspect of Cognitive Analytics is the ability to use context. Context enables a Cognitive Analytics system to infer meaning from language. For example, a chatbot can take into account the conversation history to infer who is referred to by the word he:

User: Who is Obama's wife?
AI: Michelle Obama.
User: How old is *he*?
AI: Barack Obama is 55 years old.

Figure 4: Example conversation of a cognitive system

For this simple exercise, the system needs to be aware of names that represent people, relationships between people, gender and the common sense to infer that Obama refers to Barack Obama. All of this contextual information is required to make the right inferences to answer both questions.

Since Cognitive Systems are aware of context, can understand unstructured data and reason about information, they can communicate with humans as well. This enables the system to understand a question posed in English, no longer requiring the time-consuming process of converting the question into a format the computer can work with. For example, a call center representative cognitive system can quickly answer a customer's question about camping gear by using information from product descriptions, customer reviews, sales histories, topical blogs, and travel magazines.[7] Cognitive Systems can understand and communicate through many mediums, including speech, image, video, sign language, graphs or any combination of these.

Robotics

AI is an important enabling factor in design and operationalizing smart robots and other process automation applications. In its most simple form, a robot may be a machine that is programmed to perform a simple task, by following step-by-step instructions. It could consist of a rule-based engine that explicitely tells the system what to do when a certain condition occurs. A robot in a car factory Is programmed like that and hardly considered "intelligent". But robotics exist in a variety of much more intelligent shapes, ranging from unmanned autonomous vehicles (UAV's), drones, smart vacuum cleaners to intelligent chatbots and smart assistants etc. How advanced robots are is vivid if we look at robots developed by Boston Dynamics[8] and MIT's Cheetah II[9]. Other example is Amelia[10], an intelligent assistant with NLP capabilities. Key aspect of robotics is that it combines hardware (mechanical parts, sensors, screens etc.) with intelligent software and data to perform a task for which certain level of intelligence is required (e.g. orientation, motion, interaction etc.).

Smart Machines

The major theme in using the term "Smart Machines" is autonomy. Smart Machines are systems that –to some extend- are able to make decisions by themselves, requiring no human input. Cognitive Analytics systems can be Smart Machines, as well as robots, or any kind of AI, as long as it adheres to this rule. In the case of a robot, autonomy could consists of a capability to plan where it wants to go, what it wants to achieve and how to overcome obstacles. Rather than being human-controlled or simply following instructions, it could achieve higher-level goals like getting groceries, inspecting buildings and so forth. This is enabled by planning methods, self-preservation instincts on top of the skills that a normal robot already requires.
In the case of a Cognitive System, it will pro-actively try to learn new facts, gauge opinions and learn new common sense rules by engaging in active conversation with humans, asking questions and double-checking them with data found online. It will also actively inform decision makers about changes it has observed, for example if the opinion of customers on social media suddenly makes a swing. It could even act upon these changes, in the example engaging with the customers or sharing the positive opinions on the social media outlets of the company.

Since Smart Machines are autonomous and intelligent, they might start communicating among themselves. This leads to multi-agent systems that can make trades to improve their utility. The building-inspecting robot can ask a drone to inspect the roof for him, trading this favor for another favor, like transporting goods or simply currency.

A Cognitive System that becomes a Smart Machine can specialize in a specific area, becoming an expert in that area. Now, other Smart Machines can ask it for information in that area, and it will be able to provide more relevant answers more quickly than a general Cognitive System that is not specialized. Information brokers like this improve the overall utility of the whole network of Smart Machines.

Conclusion

The terms Machine Learning, Cognitive, Robotics and smart machines are used often in relationship to AI, or sometimes even as synonyms. AI is a complex field of interest, with many shapes and forms. Therefore we have tried to shine some light on the most used terminology. In subsequent blogs, we will dive deeper in techniques behind AI systems, business applications, some associated technology trends and the top 5 risks and concerns.

[7] Deloitte, Cognitive analytics™ The three-minute guide
[8] http://www.bostondynamics.com/
[9] https://biomimetics.mit.edu/research/dynamic-locomotion-mit-cheetah-2
[10] http://www.ipsoft.com/amelia/

Artificial Intelligence techniques explained

In order to 'demystify' Artificial Intelligence, and in some way get more people involved in AI, we are publishing a series of articles explaining the world of AI, zooming in on the techniques that are associated with AI, the most appealing business applications, and potential issues we can expect.

The first blog article explained some of the most commonly used definitions of AI. In this second article we will explain 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".

Webster[1] defines the term Heuristic as "involving or serving as an aid to learning, discovery, or problem-solving by experimental and especially trial and error methods". In practice, this means that whenever problems get too large or too complex to find the guaranteed best possible solution using exact methods, heuristics are a way to employ a practical method to find a solution that is not guaranteed to be optimal, but one that is *sufficient* for the immediate goals.

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[2] and genetic algorithms[3]. 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.

---

1 https://www.merriam-webster.com/dictionary/heuristic
2 https://en.wikipedia.org/wiki/Ant_colony_optimization_algorithms
3 https://en.wikipedia.org/wiki/Genetic_algorithm

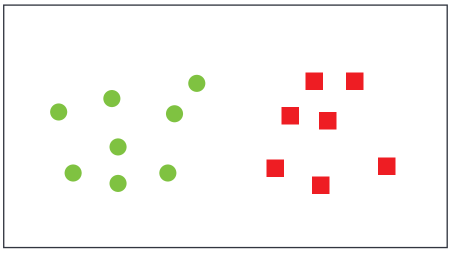

Figure 1

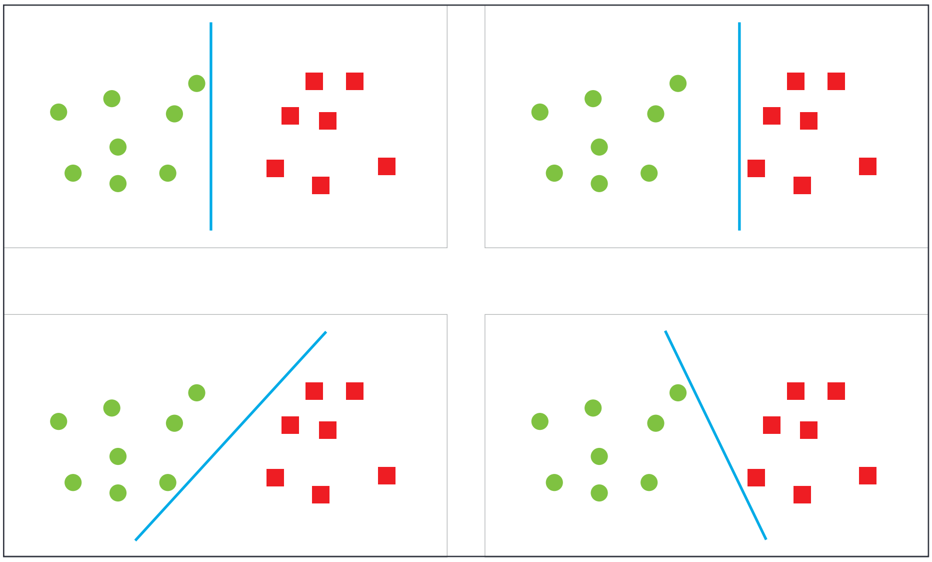

Figure 2

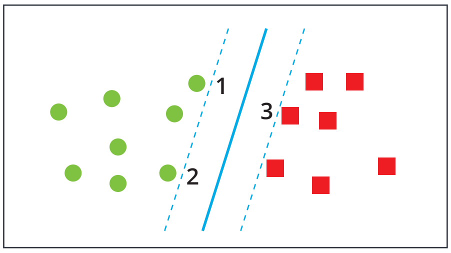

Figure 3

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 labeled 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[4] (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).

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.

[4] https://en.wikipedia.org/wiki/Support_vector_machine

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.

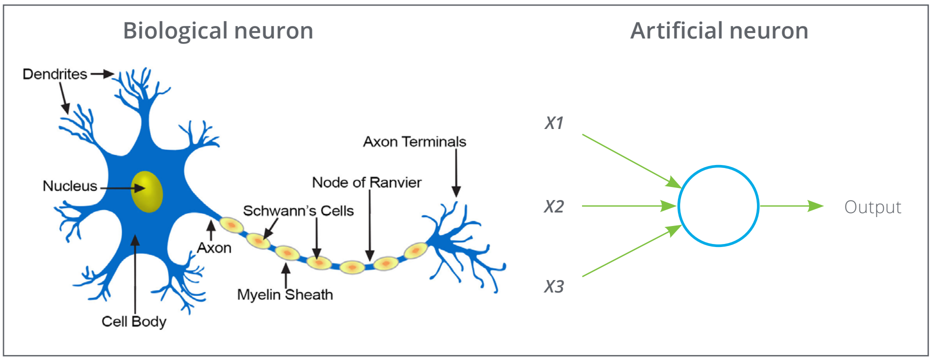

Figure 4: Graphical representation of a biological neuron (left) and an artificial neuron (right)

Artificial Neural Networks

Animals are able to process (visual or other) information from their environment and react adaptively to a changing situation. They use their nervous system to perform such behavior. Their nervous system can be modeled and simulated and it should be possible to (re)produce similar behavior in artificial systems. Artificial Neural Networks (ANN) can be described as processing devices that are loosely modeled after the neural structure of a brain. The biggest difference between the two is that the ANN might have hundreds or thousands 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) an output is generated which can be used as input again by other neurons, see Figure 1 (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.

Where in the biological brain learning takes place by strengthening or weakening the bonds between different neurons, in the ANN the learning takes place by changing the weights between the neurons. By providing the neural network with a large set of training data with known features the best weights between the artificial neurons (i.e. strength of the bond) can be calculated in order to make sure the neural network best recognizes the features.
The neurons of the ANN can be structured into several layers[5]. 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 2 shows an example network where all neurons in one layer are connected to all neurons in the next layer. Such a network is called *fully connected*. Depending on the type of problem you want to solve different connection patterns are available. For image recognition purposes, typically

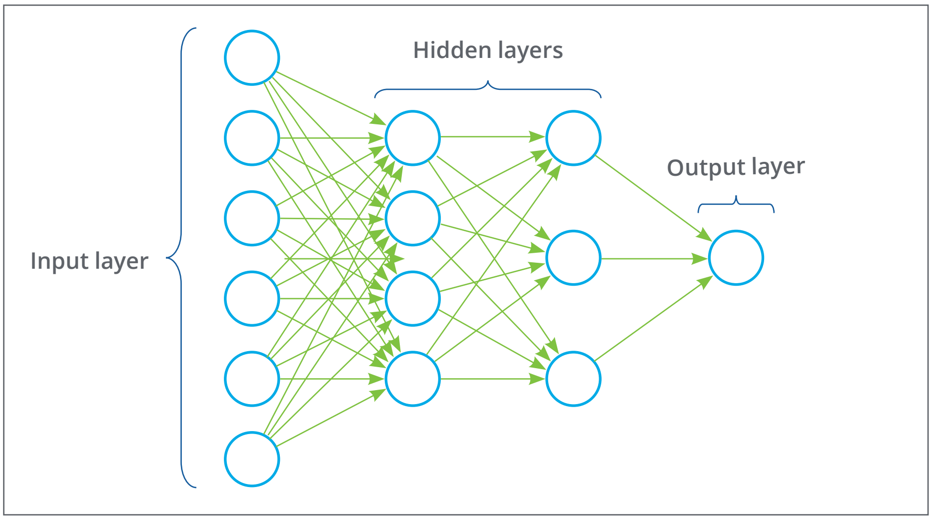

Figure 5: Schematic of a connected ANN

[5] http://neuralnetworksanddeeplearning.com/chap1.html

Convolutional networks are used, where only groups of neurons from one layer are connected to groups of neurons in the next layer. For speech recognition purposes, typically *Recurrent networks* are used, which allow for loops from neurons in a later layer back to an earlier layer.

Markov Decision Process

A Markov Decision Process (MDP) is a framework for decision-making modeling where in some situations the outcome is partly *random* and partly based on the input of the decision maker. Another application where MDP is used is planning, where the planning is optimized. The basic goal of MDP is to find a *policy* for the decision maker, tells him which particular action should be taken at which state.

An MDP model consists of the following parts[6]:

- A set of possible states: for example, this can refer to a grid world of a robot or the states of a door, door open and door closed.
- A set of possible actions: a fixed set of actions a robot for example can take, such as going north, left, south or west. Or with respect to a door, closing or opening door.
- Transition probabilities: this is the probability of going from one state to another state. For example, what is the probability that the door is closed, after the action of closing the door is performed.
- Rewards: these are used to guide the planning. With respect to the robot and the grid example, a robot may want to move north to reach its destination. Actually going north will result in a higher reward.
- Once the MDP is defined, a policy can be trained using "Value iteration" or "Policy Iteration". These methods calculate the expected rewards for each of the states. The policy then gives the best action that can be taken from each state.

As an example, we will define a grid which can be seen as an ideal, finite world for a robot[7]. This example grid is shown in Figure 6.

The robot can move (action) from each position in the grid (state) in four directions, namely north, left, right and south. The probability that the robot goes into the desired direction is 0.7 and 0.1 if it goes towards any of the other 3 directions. A reward of -1 (i.e. a penalty) is given if the robot bumps into a wall and doesn't move. Also, additional rewards and penalties are given if the robot reaches the cells that are colored green and red, respectively. Based on the probabilities and rewards a policy (function) can be made using the initial and final state.

Another example where MDP can be used is the inventory planning problem, where a stock keeper or manager each week has to decide how many units have to be ordered. The inventory planning can be modeled as an MDP, where the states can be considered positive inventory and shortages. Possible actions are for instance ordering new units or backlogging to the next week. Transition probabilities can be considered as the action that will be taken based on the demand and inventory for the current week. Rewards, or in this case, costs are typically unit order costs and inventory costs.

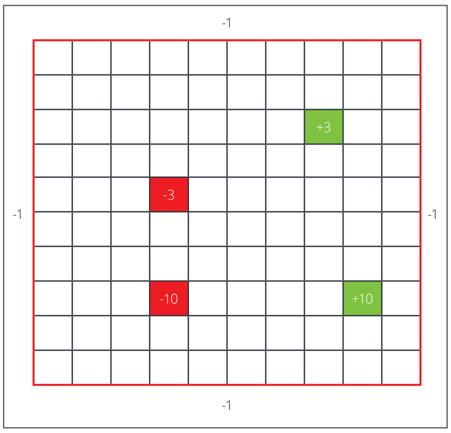

Figure 6: Example – grid world of a robot

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. A few of the important techniques will be explained below, which are Part-of-Speech tagging, Named Entity Recognition, and Parsing.

Let us examine the sentence "John hit the can." 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. A technique which is often used for tagging is a Hidden Markov

[6] https://www.cs.rice.edu/~vardi/dag01/givan1.pdf
[7] http://artint.info/html/ArtInt_224.html

Model (HMM). An HMM is similar to the Markov Decision Process, where each state is a part of speech and the outcome of the process is the words of the sentence. HMM's 'remember' sequences of words that came before and can, based on this, make better estimates of what POS a word is. For example: 'can' in 'the can', is more likely to be a noun than a verb. The end result is that the words are tagged as followed: 'John' as a noun (N), 'hit' as a verb (V), 'the' as a determiner (D) and 'can' as a noun (N) as well.

Named Entity Recognition or NER, is similar to POS tagging. Instead of tagging words with what function the word has in the sentence (POS), words are tagged with the type of entity the word represents. These entities can be for example person, company, time, location, etc. But also more specialized entities such as gene, or protein. Although an HMM can also be used for NER, the technique of choice is a Recurrent Neural Network (RNN). An RNN is a different type of neural network as discussed earlier, but it takes sequences as input, that is a number of words in a sentence, or complete sentences, and remembers the output from the previous sentence.[8] In the sentence we are looking at, it will recognize John as the entity 'person'.

A final technique to be discussed is called Parsing (Syntactic Analysis). With parsing, the text is analyzed for grammar and the arranging of the words so that the relationship between the words is clear. This is done by using the Part-of-Speech tag from the lexical analysis and then grouping these into small phrases, which in turn can also be combined with other phrases or words to make a slightly longer phrase. This is done until the goal is reached: to have used every word in the sentence. The rules how the words can be grouped are called the grammar and can take a form like this: D+N = NP, which reads: a Determiner + Noun = Noun Phrase. The final result is depicted in the figure 7.

Conclusion

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.

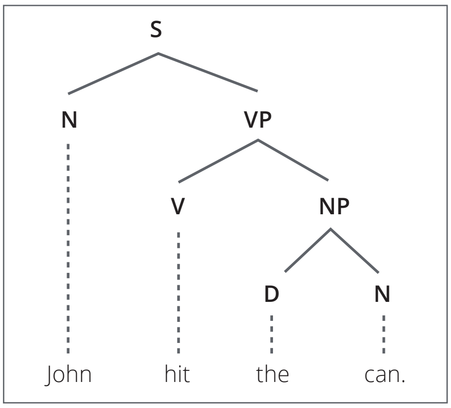

Figure 7

[8] More on RNN: http://karpathy.github.io/2015/05/21/rnn-effectiveness/

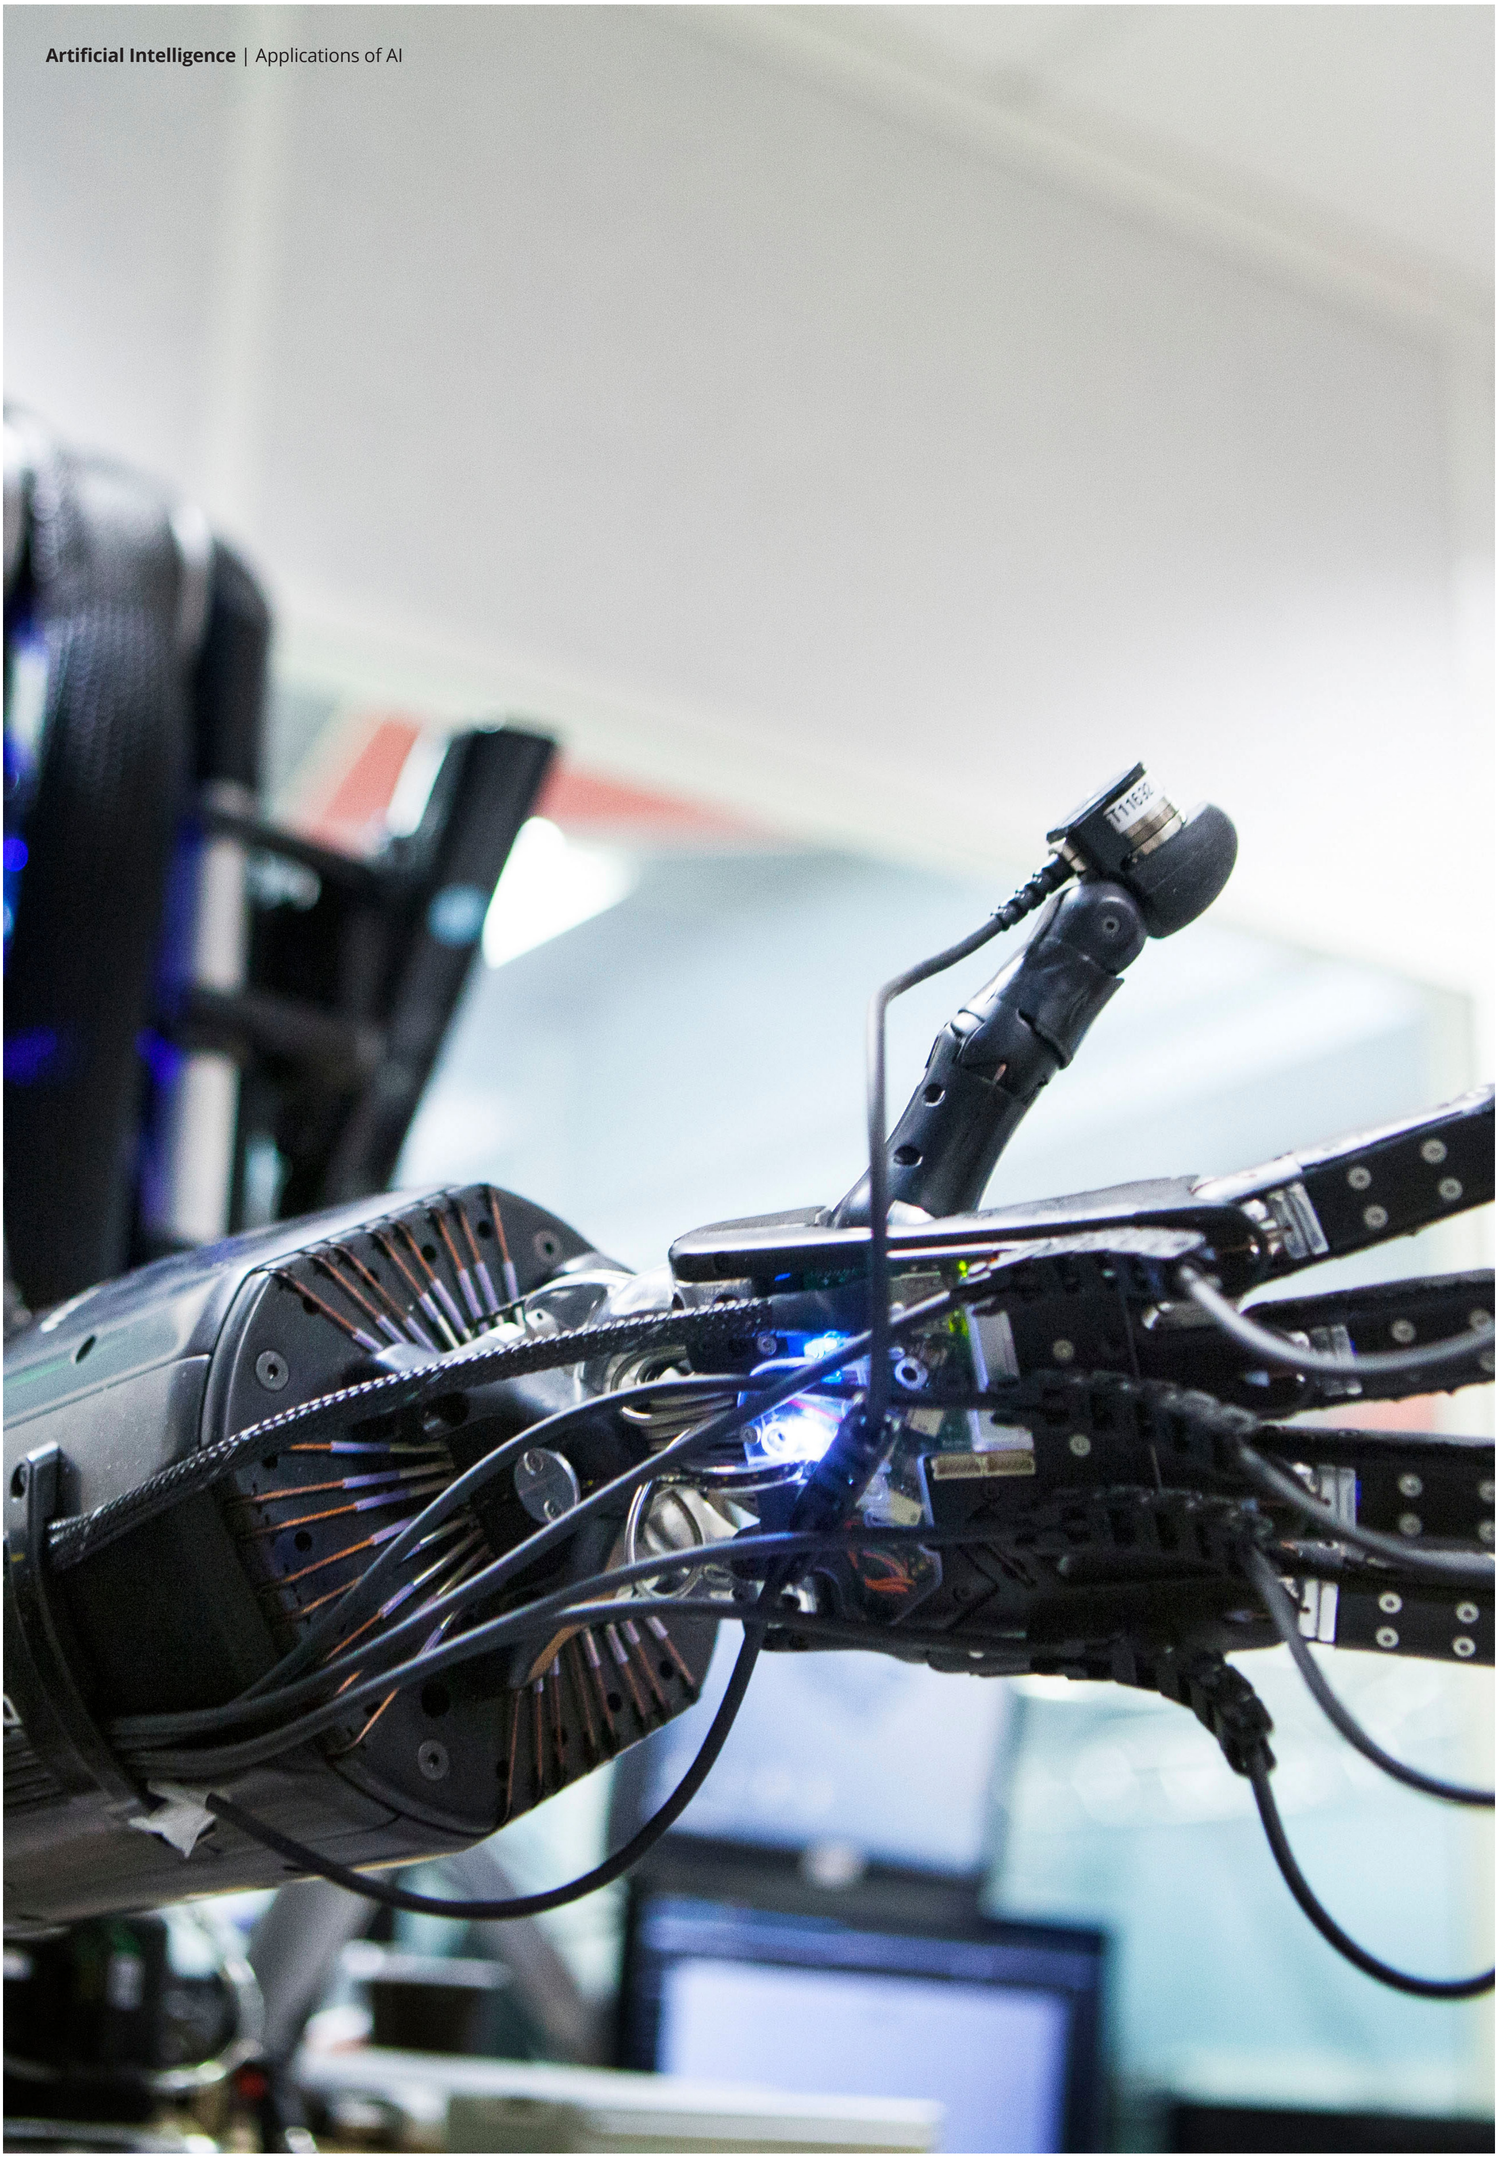

Applications of AI

In order to 'demystify' Artificial Intelligence (AI), and in some way get more people involved in it, we are publishing a series of articles explaining the world of AI, zooming in on the techniques that are associated with it, the most appealing business applications, and potential issues.

In our second blog article, we covered some fundamental AI techniques. In this article we take the next step and will discuss 5 applications of AI: Image recognition, Speech recognition, automatic translation, Q&A and Games.

Introduction

For a lot of tasks, computers are becoming or are already better than humans. The calculator on your cellphone for example, can calculate the most difficult multiplications in an instance that would take people a few minutes and a piece of paper and pen. It has already been a long time since the world champion of chess was able to beat the strongest chess computers. In the past few years however, computers started doing typically human things such as seeing and hearing, which they do as good as, if not better than people. Google can recognize exactly which things are present on an image and Siri can understand what you are saying, to name a few examples. In the coming years the intelligence of these applications will increase in a rapid pace, due to developments in the field of neural networks and increased computational power.

The increased accuracy, availability, and ease of implementation of artificial intelligence methods creates opportunities for companies to apply them in their business. As an example, insurance companies use AI to read claims from their clients, understand if the claim is difficult or easy, and it can give a recommendation on how to handle the claim. The insurance employee then only needs to do a quick check before approving the recommendation. This can save a lot of time and increase the quality of the work. This is just one example. In the remainder of this article we will discuss 5 exciting applications in which we will see great development in the coming years:

- Image recognition
- Speech recognition
- Translation
- Q&A
- Games

These developments will make applications cheaper and more accurate, opening the door for business to use them.

Image recognition

Recognizing images is an easy task for most of us. We don't have any trouble differentiating a car from a tiger or recognizing that a car is still a car when you observe it from the front instead of from the side. This task has been proven considerably more difficult for computers, but recent progress in image recognition accuracy has resulted in interesting applications. Because different vendors like Google and IBM are offering their preprogrammed algorithms open source and software libraries like Tensorflow[1] make it possible to construct your own algorithms, visual recognition is becoming more accessible for the public.

[1] https://www.tensorflow.org/

Well-known applications of image recognition are Google's shopper app[2] or facial recognition for security cams. Such applications are already using image recognition on a daily basis, however there has been a lot of development in other areas over the last few years. IBM Watson, which we know from playing Jeopardy[3], has developed its image recognition skills in the field of medicine. IBM Research has been working on deep learning techniques for computer vision[4] that could be used to recognize whether skin irregularities are melanoma. They created an ensemble of methods that can segment skin lesions and methods that can detect the area and surrounding tissue for melanoma and tested it on a large publicly available dataset, which they describe in a pre-print of the article[5]. The vision of IBM is that at a certain point medical staff can send a picture of skin irregularities to Watson, the same way that they send blood sample to the lab.

Facial recognition, which we mainly know from security cameras, has also been developed in other areas[6]. A survey of 150 retail executives by Computer Services Corporation, which was held in the UK in 2015, suggested that a quarter of all British shops use facial recognition software. The software is used for security, as one might expect, but also to track customers to observe their behavior as an effect of product displays, or the traffic flow in the store. This is a familiar concept in web shops, where A/B testing can be used to see which website display yields the best profits, however this suggests that facial recognition tools can be used to orchestrate these tests live in a store. Different software development companies are offering facial recognition for retailers[7/8]. They apply specific algorithms that use facial landmarks to recognize and distinguish between faces, which can be saved and later matched to enhance customer experience and personalize service.

Figure 1: A mannequin of Saks Fifth Avenue, an early adaptor of the facial recognition software in retail[22].

[2] https://google-shopper.en.softonic.com/android
[3] http://www.techrepublic.com/article/ibm-watson-the-inside-story-of-how-the-jeopardy-winning-supercomputer-was-born-and-what-it-wants-to-do-next/
[4] https://www.ibm.com/blogs/research/2016/11/identifying-skin-cancer-computer-vision/
[5] https://arxiv.org/abs/1610.04662[3] http://artint.info/html/ArtInt_224.html
[6] https://www.theguardian.com/cities/2016/mar/03/revealed-facial-recognition-software-infiltrating-cities-saks-toronto
[7] https://www.scnsoft.com/case-studies/facial-recognition-for-retail
[8] https://www.facefirst.com/industry/retail-face-recognition/

Speech recognition

Speech recognition is an AI application that recognizes speech and can turn spoken words into written words. It is hardly used on its own but it is largely used as an addition to Chatbots, virtual agents and mobile applications. Well known examples are Apple's Siri, Google Home and Microsoft's Alexa. Speech recognition started already in 1952 with 'Audrey'[9]. Audrey was able to recognize digits spoken in a single voice, which is quite impressive given the computers back then. Today we have applications on our phone and in our car that can respond to our voice.

Not only the amount of applications with voice recognition capacity has increased, also the accuracy of voice to words has improved dramatically over the last few years according to KPCB[10].

One of the business applications that has gained quite a lot of ground is the use of speech recognition in health care[11]. A lot of physicians are working with an electronic health record (EHR) to document patient information, however this has been said to delay the consults and to restrict the patient narrative. Using speech recognition, patient documentation can be recorded in a flexible and fast manner, which allows the physician to pay more attention to the patient. This solution is already offered by different vendors such as Nuance[12] and M*Modal[13].

The rapid improvement of speech recognition accuracy offers a lot of opportunities in the near future. Having all our soft- and hardware voice controlled might not be as far away as many people think.

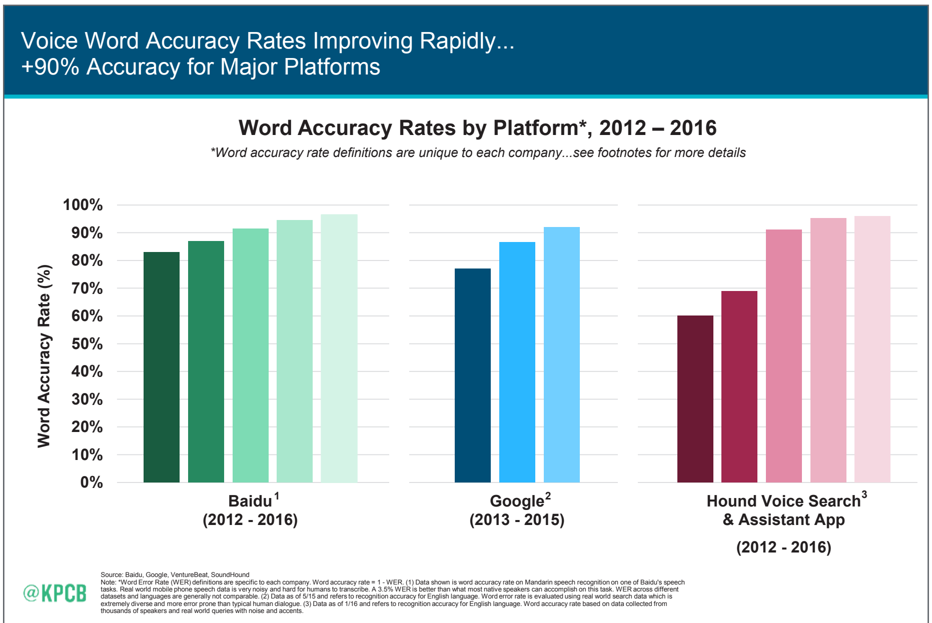

Figure 2: A table from the KPCB report of internet trends about the incline in accuracy of speech recognition10

Translation

A different topic with large business implications is automatic translation. This topic can be defined as the process of translating text from one language to another by using software. Traditionally, translation was done by substituting each word by its closest counterpart in the other language. While this works reasonably well for single words, a pair of words or sentences are generally harder to process correctly due to the fact that relations between words are important for the meaning of a sentence, but such nuances cannot be captured when each word is analyzed separately.

The usage of deep learning has had a significant impact on the quality of machine translations by completely shifting the paradigm. Rather than working in a rule-based way, powered by human decision making, translation using a neural network is completely based on mathematics. On relatively basic texts, the GNMT system translations approach the quality of human translators. An experiment even showed that when you translate English to Korean and subsequently translate English to Japanese, the model is able to translate Korean to Japanese reasonably well, without any prior training focused on the formal link between the two languages. One article even asked the question: "have computers invented their own internal language?"[8].

---

9 http://www.pcworld.com/article/243060/speech_recognition_through_the_decades_how_we_ended_up_with_siri.html
10 https://www.slideshare.net/kleinerperkins/2016-internet-trends-report/119-KPCB_INTERNET_TRENDS_2016_PAGE119Voice
11 http://www.healthcareitnews.com/news/speech-recognition-proving-its-worth
http://www.cleverbot.com/
12 https://www.nuance.com/healthcare/physician-and-clinical-speech/dragon-medical.html
13 https://mmodal.com/speech-solutions/

The impact of quality translations in a global economy are enormous. With business translations originally dominated by conversions between European languages, the need for translations to Chinese, Japanese, and Korean is increasing[9]. A simple example is one that Uber was investigating, where automatic translation takes place between you and your local Uber driver, who can only communicate in Chinese.

Question answering

Q&A agents or Chatbots are another example of applying AI to language. When talking about the ability of having conversations, distinctions are made in the domain and the way of generating an answer of the agent. A chatbot can be focused on answering questions in an open or closed domain. When it operates in an open domain, it should be able to answer general questions that can concern any topic (see for example cleverbot). This is generally harder than a closed domain, which concerns only a limited amount of topics. Closed domains, however, have very good business application such as answering questions at helpdesks. A couple of years ago, there was a breakthrough in question answering interest, when IBM Watson beat humans in a game of Jeopardy, a well-known American quiz show[10]. More recently another breakthrough was made by Google, who can now give chatbots the ability to have a short term memory, which gives the chatbot the ability to mimic real-life conversations more realistically[11].

The ability of machines to recognize intent (or the purpose) from a question and to answer it in a variety of ways, is again something that can be seen in many business applications. After an intent is distilled from the command of a user, it can be linked to a specific follow up action. This action can range anywhere from asking a return question to the retrieval of information from the internet. In the area of customer service, Chatbots are quickly becoming the norm, one example being IPsoft's Amelia[12]. Standard queries are already handled automatically, with only the difficult ones being forwarded to human decision makers. Question answering has also made an introduction in the field of Law, where lawyers can pose questions in natural language to an intelligent assistant about legal cases. The assistant can respond to the query with the relevant passage, drawn from high quality legal documentation[13].

Game/ Solver

One of the most exciting applications of AI lies in playing games. Playing a game well requires you to not only know the rules, but to calculate the next possible moves within these rules, and finally make a careful judgement on which move would give you to best chance to win. If computers can play games as well as human players, there are no reasons why they cannot learn any other difficult task that people do in their daily work (although human supervision probably remains needed).

Recently there was a big step forward in the field of games when the world Champion of Go was beaten by a computer for the first time. Go is a game that cannot be brute-force calculated, since the number of possible moves is higher than the number of stars in the universe. The top Go players of the world rely for a large part on their intuition to come to the best moves. Google's AlphaGo (a neural network based go-engine), however, learned how to play like a top human player by studying millions of human games. It then became even stronger by playing against another version of itself millions of times, which finally enabled it to beat the world champion. If computers can beat human players in one of the most complicated games that currently exists, then where do the possibilities for AI stop?

One big advantage people still have over computers, is that we can take our knowledge and training in one area, and apply it to a new task or area. For example, good go players can apply their way of thinking to solve their daily problems in their jobs. AlphaGo cannot do this: it is only good at playing go and nothing else. When you make it learn something else, like chess, it will lose its ability to play go. Recently however, a first step was taken in overcoming this problem: neural networks are now able to remember to most important knowledge from one game, and at the same time learn a new game (link). Google Deepmind wrote a new algorithm that allowed a neural network to learn 10 Atari games at the same time, and play them with human performance.

Once this field will be more developed, computers will be able to perform series of difficult tasks that at the moment only people can perform. Google themselves use it to lower the energy bills of their large datacenters. The AI controls over 120 variables in Google's datacenters, such as the windows, fans and cooling systems, optimizing for energy usage while keeping computing performance up. The optimization potentially lowers Google's energy bill for hundreds of millions over several years. Another application lies in healthcare, where an app of Deepmind saves nurses over two hours of time per day by warning about upcoming acute kidney failure. These are two applications, but since this is a newly developed field there is a huge potential for more. Think about predicting stock prices, optimizing the layout of distribution centers. Imagination and available data are the limit.

Conclusion

We have discussed a few applications of AI that will become more advanced the coming years and that will probably become more and more common in our daily lives. In the series of AI blogs we have tried to shed light on applications of AI that everyone will encounter at some point and give some examples of where these applications might lead to in the near future. In the next blog we will address 5 trends that leapfrog AI adoption.

Five technology trends that leap-frog Artificial Intelligence

In our first two articles about AI terminology[1] and techniques, we have discussed the definitions and techniques that make up AI. But developments in AI do not stand on its own. The broader field of technology development have a huge interrelationship with the adoption of AI, making it more accessible on one hand and driving the need to use AI on the other hand

In this article, we will elaborate on a few broader technology developments that are leap-frogging the adoption of AI. Trends or factors of growth that bring AI to more people and organizations, of which some of the more important trends are described in this article.

1. Cloud

One of the first trends that enables the rapid growth of AI is cloud computing. As explained in our previous article[2], AI techniques are based on complex mathematical models and require large amounts of training data (examples) to learn their intelligent capabilities. Therefore building, improving and running AI applications requires immense computing power. Cloud technology offers that in a flexible and scalable environment at relatively low-cost and without huge initial investments.[3]

In addition, the IT infrastructure of large corporates is often too big and too rigid to experiment with AI applications across and within the enterprise platforms. AI cloud services such as Amazon AWS AI, Microsoft Cortana, IBM Bluemix/Watson, Google Cloud Machine Learning and HPE HavenOnDemand allow you to quickly build and run applications. With cloud 'Analytics as a Service' (AaaS) offerings, organizations can experiment with AI and begin to build intelligent applications without harness existing IT infrastructure.

2. Big Data

Second amongst the recent accelerators for AI is **big data**, or more specifically: large, fast, and/or unstructured data. Think of all the information in images, text, sensor data or other data generated by for example mobile devices. Right now, 80% of all company data is unstructured and increases much faster in size than its structured counterpart[4]. In recent years, technology has become widely available to capture, store and process that data. Therefore, many companies have invested in building "Data Lake" platforms manage their Big Data.

The potential for applications of this unstructured data is huge and yet largely untapped. AI techniques make it possible to process and analyze unstructured data, allowing business to obtain valuable insights from the information and improve their decision making. It can potentially

---

1 Part 1: Artificial Intelligence defined - https://www2.deloitte.com/nl/nl/pages/data-analytics/articles/part-1-artificial-intelligence-defined.html
2 Part 2: Artificial Intelligence Techniques Defined: https://www2.deloitte.com/nl/nl/pages/deloitte-analytics/articles/part-2-artificial-intelligence-techniques-explained.html
3 The Cloud: technical advantages - https://www2.deloitte.com/au/en/pages/deloitte-private/articles/cloud-computing-deloitte-private.html https://www.nuance.com/healthcare/physician-and-clinical-speech/dragon-medical.html
4 Dark analytics: illuminating opportunities hidden in unstructured data: https://dupress.deloitte.com/dup-us-en/focus/tech-trends/2017/dark-data-analyzing-unstructured-data.html

find patterns and complex relationships by shifting through billions of observations. For example, in the process of assessing insurance claims, intelligent AI applications can automatically understand natural language from texts and analyze images such as photographs. By using these techniques, AI applications have the potential of detecting fraud earlier, improve the quality and consistency of the claim assessment and make the process more efficient. This is only one out of many example use cases.

While on the one hand AI is the solution to analyze large amounts of unstructured data, AI on the other hand needs big data in order to become 'intelligent'. As already mentioned in the previous section and our previous article, AI applications need to be trained with many examples. Think of this as personal development: during your life, you learn from experience (examples given by your environment), which enable you to perform certain tasks and get better, leading to more expertise. This can be illustrated by an example of Tesla's Fleet Learning[5], where all Tesla's continuously share the data with the central intelligence system, such as map data: where is driven when - and at what speed. Once a significant number of cars report on a changing condition (due to i.e. poor road conditions, works, etc.) the system is updated allowing others cars to learn to anticipate on the changed condition. Hence the relationship between AI and big data is two-fold: big data is a prerequisite for AI and AI is the solution to process unstructured data and derive insights from it.

3. APIs

Perhaps the easiest way to start building intelligent applications is by using Application Programming Interfaces (APIs). An API is a piece of out-of-the-box functionality that can be called from another program or App. If for example, your app requires face recognition, you can call an API rather than program it yourself. Many of the large technology firms offer APIs in the field of computer vision, speech recognition, and natural language processing (NLP) or other cognitive domains on their cloud platforms. Intelligent APIs are pre-trained and pre-configured models for a certain task and serve as gateways to AI applications. This can be illustrated with the Visual Recognition[6] API from IBM. When this API is called and receives an image of a car, it recognizes the car, and perhaps other objects that are on the provided image. The recognized objects are returned to the user with a certain amount of confidence for each class of objects.

Furthermore, APIs do not need to be used as standalone services. They can serve as building blocks for combined intelligent applications. For example: building a speech translator assistant requires a speech-to-text API, a translation API and a text-to-speech API to return the translation in the other language. This modular characteristic of APIs makes them very useful for a wide range of AI applications. Using cognitive APIs is the easiest and quickest way to start building and integrating an AI application.

4. Open Source

Although APIs are a great way to start building AI applications, they are very specifically aimed at achieving a certain task. When you need to perform a machine learning task that isn't available through an API, you can build one yourself. This requires knowledge, complex algorithms and frameworks.

Today, increasingly more AI algorithms and frameworks are available as Open Source, meaning they are publicly available and often at no license cost. Consequently, developers of AI applications can rely on the knowledge and previous work of a large user base. That makes this trend our fourth accelerator of the AI uptake. An example of open source AI software is TensorFlow from Google.

TensorFlow is an open source machine learning library with many different algorithms and frameworks[7]. Open source training sets are available these days which allow for loading Intelligence directly into your AI application[8].

A simple comparison between an API and the underlying open source framework can be illustrated by the following example. Google's Speech API can be used to develop applications that receive audio from its user and convert it to plain text. To do this, Google's API uses techniques such as deep learning and is trained with millions of examples. Next to this API, Google also made the source code for the (deep learning) neural network itself publicly available. However, with only the source code, no audio can be converted

5 PTesla's Fleet Learning - http://gas2.org/2017/02/14/fleet-learning-capacity-key-teslas-lead-self-driving-race/
6 Visual Recognition by IBM - https://www.ibm.com/watson/developercloud/visual-recognition.html
7 TensorFlow by Google - https://www.tensorflow.org/
8 Open Data for Deep Learning - https://deeplearning4j.org/opendata

to text. The model needs to be trained with many examples, but can be set very specifically according to the needs of the user. This can gain the competitive advantage organizations are looking for.

As a final note on open source, AI applications can also use a combination of APIs and open source technologies to rapidly build up an application from multiple pre-built modules, which speeds up the development process and uses the best of both worlds. Altogether, open source AI technologies broaden the user base and the set of use cases for AI, and create competition for large commercial AI vendors so that prices remain low. At the same time, some of the large vendors take Open Source AI code and package it into a commercial application, adding support, maintenance, training etc.

5. IoT and standardization

The last important accelerator for AI is Internet of Thing (IoT), meaning all the (mobile) devices, vehicles or sensors that are connected to the internet. All these devices together generate a massive amount of fast, semi-structured data, which can 'feed' and improve AI applications. A few examples of these are: self-driving vehicles, homes with intelligent thermostats[9], intelligent pacemakers from patients that provide doctors with real-time insights or a parking garage that recognizes your car[10]. As more and more devices are connected, standardization of data flows, -formats and -services is needed, so that these devices can interact properly. Standardization in IoT is still in development[11]. Two important building blocks for standardization in the AI field are the JSON data format and REST principles for APIs[12].

JSON (JavaScript Object Notation)[13] is a data format that is easy to write and understand for humans and easy to generate and parse for computers. REST stands for Representational State Transfer and is an architecture/approach for web services which allows easy evolution of API driven services.

With more data generated by IoT devices and easier development and integration due to standardized formats and principles, AI solutions are emerging more broadly.

Conclusion

Recent technological trends drive the broad adoption of AI. With the cloud as platform and APIs as building blocks for intelligent applications, AI is available for more people and organizations than ever before. The rise of unstructured data opens opportunities for AI technologies which can give companies a competitive advantage. Also, AI applications are better than in the past thanks to better availability of training data. Moreover, with open source AI technologies, customized AI solutions can be developed quickly using pre-built modules and using wisdom of the crowd. Finally, IoT will drive standardization further which allows for even more data becoming available and better integration of devices.

With the recent technological trends described in this article, we believe that the right time has come for organizations to start developing their first AI use cases.

---

9 Combining AI with IoT - https://www.techemergence.com/artificial-intelligence-plus-the-internet-of-things-iot-3-examples-worth-learning-from/

10 The Edge: the smartest building in the world - https://www.bloomberg.com/features/2015-the-edge-the-worlds-greenest-building/

11 2016 : A year of advancing Internet of Things (IoT) applications: https://innovation-in-manufacturing.deloitte.com/2016/12/28/2016-a-year-of-advancing-internet-of-things-iot-applications/

12 Open Connectivity Foundation (OCF): https://openconnectivity.org/

13 JSON - http://www.json.org/

AI opportunities for the future

As mentioned in our previous blogs, Artificial Intelligence (AI) is said to be one of the most disruptive technologies impacting several industries and businesses. Researchers and practitioners compare the emergence of AI with the industrial revolution of the last century[1]. AI comprises of a set of technologies that will change business as we know it. More and more tasks will be automated, but the largest impact will follow from brand-new business models and the rise of intelligent services that did not exist before.

To conclude our series of blogs on AI, this edition will highlight 5 exemplary business opportunities enabled by use of AI. We provide these examples to have a discussion around questions such as: what does incorporating AI in your business can offer? What value can be delivered by AI? And how to measure the value of AI? We know that impact of AI can be expressed in different ways, e.g. as depicted in table below:

Value	Example
Cost reduction	High level of intelligent automation in call centers for insurers
Optimal service efficiency	Efficiency in healthcare by reducing contact moments between patients and general practitioners (implementing smart chatbots to perform triage and answer most common questions)
New flows of revenue	Cross-sell and up-sell of products in e-commerce via intelligent recommendation systems and via high-level of personalization
Customer satisfaction	New services such as grocery shops providing health advise to customers

We will now go through some examples across various industries and examine the nature of AI implemented in these cases.

[1] https://medium.com/mmc-writes/the-fourth-industrial-revolution-a-primer-on-artificial-intelligence-ai-ff5e7fffcae1

Application 1: AI in disease diagnosis and illness treatment

Cost reduction by AI seems to be an important driver to counter the ever increasing cost of healthcare. However the opportunity to improve the effectiveness of healthcare by AI driven diagnostics and treatment plans is much bigger.

Opportunities range from drug design to patient diagnosis and to create personalized treatment plans (almost on DNA level).

Just as an example, Infervision uses Deep Learning techniques on patient data derived from X-Ray, CT, MRI, text descriptions of symptoms and diagnostic reports, to construct automatic diagnostic recommendations[2]. Multiple top-notch hospitals in China already work with this recommendation system, identifying (lung) cancer in an early stage.
Another example is the use of Deep Learning to classify skin cancer. Researchers already developed a system that can classify skin cancer with a precision comparable to dermatologists[3]. Transferring this technology to mobile devices can make dermatologic care available for more people than ever before.

Application 2: AI in fashion design and customer interaction

Not the first topic that comes to mind regarding AI, but the technology is expected to influence the fashion industry as well. A long term, but disruptive, opportunity is made possible by GANs: Generative Adversarial Networks. The algorithm uses two neural networks: a generator and a discriminator. The logic of GAN is that the generator generates images of which the discriminator believes them to be real (using an iterative approach). The generated images are so alike to the training set, they would fit right in. Using this technique a GAN can act as a true fashion designer by creating brand new items, in the same style as the style of the items on which it is trained. The Amazon team Lab126 has already started to test the possibilities in this domain[4]. For this specific example, in terms of value, greater customer satisfaction and even competitive advantage can be achieved.

There are also other examples. Fore instance AI driven chatbots equipped with NLP and image recognition capabilities can interact with customers in an intuitive way and advise them on products. Imagine you taking a picture of your shirt and sharing it with the chatbot who then comes up with advise on newest trendy sneakers matching your outfit.

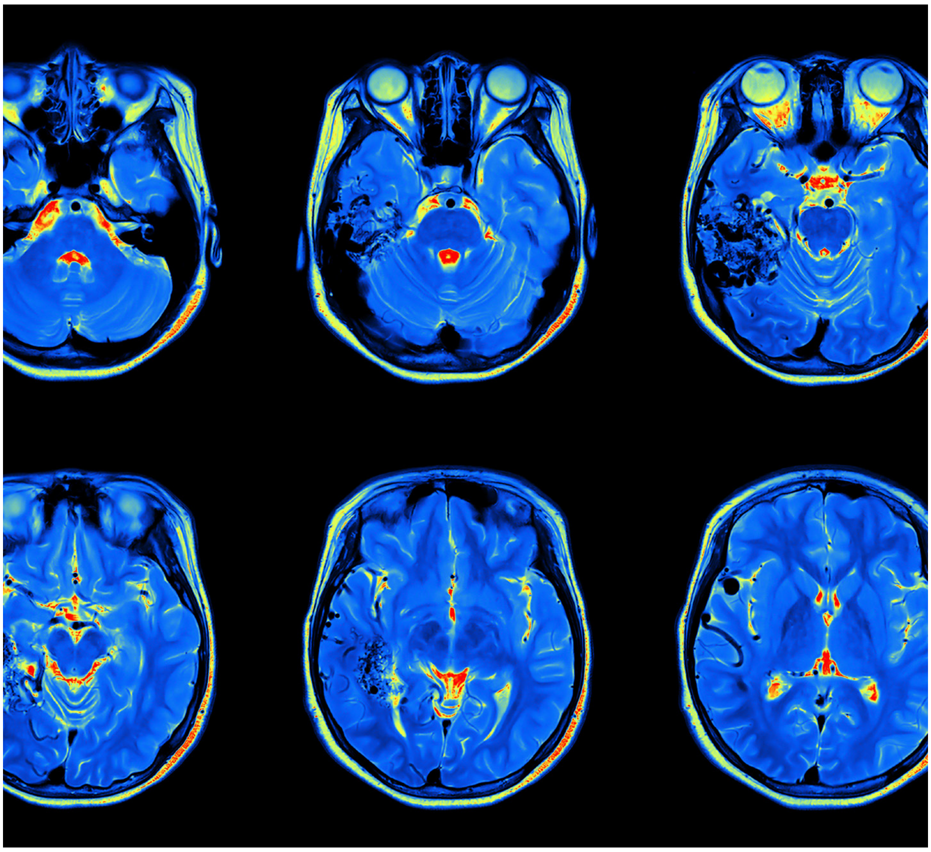

[2] https://www.forbes.com/sites/jenniferhicks/2017/05/16/see-how-artificial-intelligence-can-improve-medical-diagnosis-and-healthcare/#68a5ef462239

[3] http://www.nature.com/nature/journal/v542/n7639/full/nature21056.html?foxtrotcallback=true

[4] https://www.technologyreview.com/s/608668/amazon-has-developed-an-ai-fashion-designer/?utm_content=buffer4008a&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer

Application 3: AI in cyber crime and fraud detection

The impact of cybercrime on global business is enormous. Next to intangible damage for a company (e.g. brand damage), costs can be huge. It is estimated by Forbes that the global costs for cybercrime will be 6 trillion dollars by 2021. A large proportion of this cost is in credit card fraud and although fraud detection techniques exists for years, they are still not sufficient enough to prevent this from happening. New AI techniques, like Recurrent Neural Networks that were previously not available can be the solution to detect fraud in earlier stages[8]. Equipped by trained RNN's, fraud detection systems can scan thousands of transactions instantly and predict/classify them into buckets (e.g. ranging from high to low probability of fraud). This type of system could help save time by focusing the cases where probability is high for fraud. A word of caution in this case is also applicable. Utilizing these type of algorithms requires a solid governance system to monitor the output and limit potential bias as much as possible. It's very well known that due to pre trained character of the algorithms, bias can be serious flaw in the system.

Criminal identification

With the recent launch of the new iPhone X, Apple brings face recognition to many of us worldwide. In the next years iPhone users will be able to unlock their iPhone by looking into the front camera. But authentication of personal content is not the only application for face recognition. Governments and security services use face recognition to identify citizens and track down criminals. Recently the police in China captured twenty-five wanted criminals at a Chinese beer festival based on photos taken at the entrance of the festival[5]. Not only facial recognition can help in tracking down criminals, emotion analysis can deliver extended value regarding this opportunity. Being already in use, NTchLab created software that identifies if someone is feeling stressed or angry. Put into use in e.g. a grocery store, detecting a stressed person will alert security to pay extra attention[6].

Figure 1: Targeted Advertising Demo

Application 4: AI in personalized advertising and support

Imagine walking into a store and your favorite brand or product is at discount the same day. This can be a coincidence and your lucky day, but it's an advantage for the store owner as you're probably more willing to buy the product. All of this is possible with AI techniques like face-recognition. A personalized advertising application as described above is trained with many examples, knowledge of people and their shopping behavior. It results in an intelligent application that can identify the type of customers, their emotions and possible shopping preferences. Personalized advertising is a proven method to generate more revenue. Since these techniques are already available today, this opportunity is a high value/short term win.

Voice ordering

England largest online grocery shop Ocado is a true pioneer when it comes to AI-technology in a customer's daily shopping routine. With a cutting edge robotics system and an autonomous delivery system in store, they recently released a new Amazon feature: voice control[7]. When in possession of an Amazon Echo (Alexa) one can simply call out new items and Alexa will add them to an existing shopping list. Although this feature is already on the market, its added value will be something for the future. Currently, a shopping list has to be initiated manually and Echo only listens to full and scripted sentences. So a phrase such as: 'He Echo, we're out of milk!' is still a long way off.

---

[5] http://www.dailymail.co.uk/sciencetech/article-4851564/Facial-recognition-detects-criminals-beer-festival.html?ITO=1490&ns_mchannel=rss&ns_campaign=1490

[6] http://www.dailymail.co.uk/sciencetech/article-4494994/New-technology-reads-emotions-potential-terrorists.html

[7] https://www.engadget.com/2017/06/29/cargopod-ocado-oxbotica-driverless-delivery-van/

Application 5: AI in logistics and on-time delivery

With immense growth of e-commerce and on-line shopping, one of the biggest challenges for ecommerce and many logistics companies is to reduce the costs of delivery in the "last mile" while maintaining quality and service. Some of the most innovative companies are going through a paradigm shift using AI. Some great examples are out there where driverless cars are being utilized for autonomous delivery and real-world tests are right now on-going (7). There are also more practical areas where AI plays an essential role in the logistics. Think about on-time delivery, planning and time-estimation used by logistic companies through their track-and-trace systems. A good example of how AI is utilized is an algorithm developed by Deloitte Netherlands to predict timeslot of delivery. This model is able to make predictions for future deliveries based on delivery route patterns. These patterns are found using AI techniques. Proven cost reduction and significant improved customer satisfaction are the big wins here.

Conclusion

We have discussed some of the applications of AI and can conclude that that there is a major opportunity for the business to improve, gain efficiency and define new business models by utilizing the power of AI.

This was the 5th and last edition of our series on AI. We started the series on AI by providing an overview of terms and definitions of AI[9]. In the 2nd blog we have discussed some of the fundamental techniques which are at the heart of AI[10]. In the 3rd article we went on and gave five applications of AI: Image recognition, Speech recognition, automatic translation, Q&A and Games. Then we continued and analyzed five technology trends that leap-frog AI. This was published in our 4th blog. And then finally, in this 5th release we have provided some practical applications of AI in day-to-day life.

AI is and will remain a dynamic and broad filed encompassing several areas of expertise, ranging from computer science to mathematics, neuroscience all the way up to philosophy and even biology. We hope that the topics discussed in our 5x5 on AI edition will help the reader to get an end-to-end overview of AI in a more practical sense rather than a bullet proof complete scientific exercise.

8 https://www.forbes.com/sites/theyec/2017/07/13/the-true-cost-of-cybercrime-for-businesses/#61c4114a4947
9 https://www2.deloitte.com/nl/nl/pages/deloitte-analytics/articles/part-1-artificial-intelligence-defined.html
10 https://www2.deloitte.com/nl/nl/pages/deloitte-analytics/articles/part-2-artificial-intelligence-techniques-explained.html
11 https://www2.deloitte.com/nl/nl/pages/deloitte-analytics/articles/part-3-applications-of-artificial-intelligence.html
12 https://www2.deloitte.com/nl/nl/pages/deloitte-analytics/articles/part-4-five-technology-trends-that-leap-frog-artificial-intelligence.html?id=nl:2sm:3li:eng_da_corp:AI4

Authors

Stefan van Duin
Director
Tel: + 31 88 288 4754
Mobile: + 31 61 234 4457
Email: SVanDuin@deloitte.nl

Stefan van Duin is Director at Deloitte Consulting and leading the Data Analytics team, a team of over 100 professionals specialized in Big Data, Data Visualisation, Machine Learning and Artificial Intelligence. Stefan is convinced that more and more companies consider data analytics as a core competency, to develop better products, services and business models, or to personalize the client experience. Stefan has made it his mission to help companies on their journey.

Naser Bakhshi
Senior Manager
Tel: + 31 88 288 3874
Mobile: +31 61 201 1665
Email: NBakhshi@deloitte.nl

Naser Bakhshi is Senior Manager Advanced Analytics in the Deloitte Dutch Analytics service line. His primary focus is on Artifical Intelligence, Machine Learning and Advanced Analytics where he focuses on AI driven solution for monitoring, predictions and process optimization. Moreover, Naser is an experienced transition manager with experience in defining advanced analytics strategy (incl. business case definition), analytics roadmap, implementation of Analytics Centre of Excellence (incl. defining operating models for ACE, end-to-end processes and RACI, drafting Analytics architecture and data acquisition and quality check processes. Combining hard analytical skills with a strong feeling for organizational change and sensitivity make Naser suitable for supporting clients in becoming insight driven.

Co-authors
Chapter 2: Titus Sloet tot Everlo & Hicham el Bouazzaoui.
Chapter 3 : Carmen Wolvius en Sjors Broersen
Chapter 4 : Johan van der Veen en Thomas Heeneman
Chapter 5: Joanne Lijbers, Carmen Wolvius en Hicham El Bouazzaoui.

Sources

Chapter 1:

1. https://www.gartner.com/document/3380751/
2. http://www.bbc.com/news/technology-37713629
   https://futureoflife.org/ai-open-letter/
3. Preparing for the Future of Artificial Intelligence, NSTC, 2016
4. 6. Raphael, B. 1976. The thinking computer. San Francisco, CA: W.H. Freeman
5. http://www-formal.stanford.edu/jmc/whatisai/node1.html
6. Stephen Lucci, 2016, Artificial intelligence in the 21st century : a living introduction
7. Deloitte, Cognitive analytics™ The three-minute guide
8. http://www.bostondynamics.com/
9. https://biomimetics.mit.edu/research/dynamic-locomotion-mit-cheetah-2
10. http://www.ipsoft.com/amelia/

Chapter 2:

1. https://www.merriam-webster.com/dictionary/heuristic
2. https://en.wikipedia.org/wiki/Ant_colony_optimization_algorithms
3. https://en.wikipedia.org/wiki/Genetic_algorithm
4. https://en.wikipedia.org/wiki/Support_vector_machine
5. http://neuralnetworksanddeeplearning.com/chap1.html
6. https://www.cs.rice.edu/~vardi/dag01/givan1.pdf
7. http://artint.info/html/ArtInt_224.html
8. More on RNN: http://karpathy.github.io/2015/05/21/rnn-effectiveness/

Chapter 3:

1. www.deloittesmartreports.com
2. https://www.tensorflow.org/
3. https://google-shopper.en.softonic.com/android
4. https://www.techrepublic.com/article/ibm-watson-the-inside-story-of-how-the-jeopardy-winning-supercomputer-was-born-and-what-it-wants-to-do-next/
5. https://www.ibm.com/blogs/research/2016/11/identifying-skin-cancer-computer-vision/
6. https://arxiv.org/abs/1610.04662
7. https://www.scnsoft.com/case-studies/facial-recognition-for-retail
8. https://www.facefirst.com/industry/retail-face-recognition/
9. https://www.pcworld.com/article/243060/speech_recognition_through_the_decades_how_we_ended_up_with_siri.html
10. https://www.slideshare.net/kleinerperkins/2016-internet-trends-report/119-KPCB_INTERNET_TRENDS_2016_PAGE119Voice
11. http://www.healthcareitnews.com/news/speech-recognition-proving-its-worth
12. https://www.nuance.com/healthcare/physician-and-clinical-speech/dragon-medical.html
13. https://mmodal.com/speech-solutions/
14. https://techcrunch.com/2016/11/22/googles-ai-translation-tool-seems-to-have-invented-its-own-secret-internal-language/
15. http://www.cleverbot.com/
16. https://www.ipsoft.com/amelia
17. https://www.ibm.com/blogs/watson/2016/01/ross-and-watson-tackle-the-law/
18. https://en.wikipedia.org/wiki/AlphaGo_versus_Lee_Sedol
19. https://deepmind.com/blog/innovations-alphago/
20. https://arxiv.org/abs/1712.01815
21. http://www.telegraph.co.uk/technology/2017/03/15/googles-deepmind-ai-learns-like-human-overcome-catastrophic/
22. https://www.bloomberg.com/news/articles/2016-07-19/google-cuts-its-giant-electricity-bill-with-deepmind-powered-ai
23. http://www.wired.co.uk/article/deepmind-nhs-streams-deal

Chapter 4:

1. https://dupress.deloitte.com/dup-us-en/focus/tech-trends/2017/dark-data-analyzing-unstructured-data.html
2. http://gas2.org/2017/02/14/fleet-learning-capacity-key-teslas-lead-self-driving-race/
3. https://www.ibm.com/watson/services/visual-recognition/
4. http://www.healthcareitnews.com/news/speech-recognition-proving-its-worth
5. https://www.nuance.com/healthcare/physician-and-clinical-speech/dragon-medical.html
6. https://mmodal.com/speech-solutions/
7. https://www.tensorflow.org/
8. https://www.techemergence.com/artificial-intelligence-plus-the-internet-of-things-iot-3-examples-worth-learning-from/
9. https://www.bloomberg.com/features/2015-the-edge-the-worlds-greenest-building/
10. https://innovation-in-manufacturing.deloitte.com/2016/12/28/2016-a-year-of-advancing-internet-of-things-iot-applications/
11. https://openconnectivity.org/
12. http://www.json.org/

Chapter 5:

1. https://www.forbes.com/sites/jenniferhicks/2017/05/16/see-how-artificial-intelligence-can-improve-medical-diagnosis-and-healthcare/#68a5ef462239
2. http://www.nature.com/nature/journal/v542/n7639/full/nature21056.html?foxtrotcallback=true
3. https://www.technologyreview.com/s/608668/amazon-has-developed-an-ai-fashion-designer/?utm_content=buffer4008a&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer
4. http://www.dailymail.co.uk/sciencetech/article-4851564/Facial-recognition-detects-criminals-beer-festival.html?ITO=1490&ns_mchannel=rss&ns_campaign=1490
5. http://www.dailymail.co.uk/sciencetech/article-4494994/New-technology-reads-emotions-potential-terrorists.html
6. https://www.engadget.com/2017/06/29/cargopod-ocado-oxbotica-driverless-delivery-van/