Rise of the machines
"They grow up so fast!"

Expressions like artificial intelligence (AI), machine learning (ML) or deep learning are often used in newspaper headlines today. This is almost always to present some new thing that robots can make and how they are becoming more intelligent every year, creating new opportunities and threats at the same time. After reading all these reports, one might be tempted to ask: what is a robot, actually? What does it mean for a robot to learn, can it think? Our goals in this paper are to make this topic easier to understand, show why you should care and provide insights on how to build a portfolio that benefits from current technological changes.

To put things in relationship with each other clearly and from the start: artificial intelligence is the overarching field, encompassing subfields like machine learning, machine vision, language processing and robotics. Deep learning is one of the methods used in machine learning. Very often we see these terms being used interchangeably. What we want to focus on in this article is machine learning.

Let us consider a short bit of history with the word robot, which has a strong legacy: it is accepted to have been coined by Czech author Karel Čapek almost a century ago in a play that presented artificially created servants overthrowing their masters. In Czech, the word *robota* means labour, in the sense of forced labour[1]. The evolution of robotic "brains", allowing them to learn from past experience, was pioneered by Dr. Arthur Samuel in the 1950's and 60s in the US. He is considered the father of machine learning[2]. He tested his ideas on a computer game of checkers, creating a system that responded to player choices with moves that its memory had previously registered as leading to victory[3]. This first approach, combined with the limited power of computers at the time, did not allow Samuel's software to beat experienced human players. A more recent example of applied machine learning is provided in the video below: it shows how a computer learns to play the videogame Super Mario, based on how far it made it in the game with its previous moves[4].

[1] RoboticsTrends, 2016 : http://www.roboticstrends.com/article/origin_of_the_word_robot
[2] Stanford, n.d : http://infolab.stanford.edu/pub/voy/museum/samuel.html
[3] Samuel, 1969: Some Studies in Machine Learning Using the Game of Checkers
[4] SethBling, 2015: https://www.youtube.com/watch?v=qv6UVOQ0F44. *For those of you interested in seeing what the code that allows the computer to adapt actually looks like, the creator has made it available under https://pastebin.com/ZZmSNaHX.*

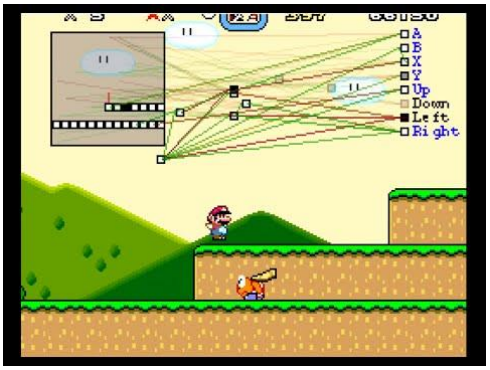

This illustration helps us understand machine learning better: it is the process through which a computer can learn from past experience to adapt its behaviour and reach its predetermined goal. In Super Mario, the goal was to go as far as possible. This means that instead of having to manually program a new set of directions for each new situation a computer faces, the creator of such a computer rather provides it with a framework of rules that should apply to all future situations. These rules are then put to the test in different situations through trial and error. The data the computer gathers during these tests are analysed to derive which adjustments within the framework bring it closer to the goal and which ones don't. The rules are then improved by the computer by incorporating the positive factors and discarding the negative ones. This reminds us of biological evolutionary processes with random mutations, some of which survive and some of which do not. The goal is to be able to create a model that can respond adequately to new tests it hasn't faced yet thanks to what it learned in previous tests.

There is a lot of hype around machine learning at the moment, with people already speaking of artificial consciousness and even granting citizenship to humanoid devices[5]. To be clear, we are still very far away of a scenario like in the sci-fi movie Ghost in the Shell, where sentient AIs escape their owners and ask for political asylum to live freely. Technically speaking, even the most complex machine learning algorithms still remain mathematical optimisation processes at their core. But despite this, it is true that machine learning applications do promise to bring significant changes to our world, and there are plenty of examples.

Many of us have heard about the virtual assistants offered by large US tech firms. Siri, Alexa, Bixby or Cortana are all names given to software that recognises our natural language, processes our demands and responds with natural language[6]. They do not only find information, they can also be used to make our lives more comfortable at home when they control smart appliances[7]. Examples at other large tech companies would be Netflix's watchlist, which recommends new shows based

[5] Forbes, 2017 : https://www.forbes.com/sites/zarastone/2017/11/07/everything-you-need-to-know-about-sophia-the-worlds-first-robot-citizen/#2dfecd846fa1
[6] UCSB, 2017: http://www.news.ucsb.edu/2017/018092/iphone-siri-and-artificial-intelligence
[7] Wirecutter, 2017: https://thewirecutter.com/reviews/what-is-alexa-what-is-the-amazon-echo-and-should-you-get-one/#what-can-alexa-skills-do

on what we have watched before or Facebook, which displays ads tailored to its users based on their profile and actions.

In healthcare, we see companies like RetinAI Medical develop algorithms that support doctors in diagnosing more precisely what their patients suffer from[8]. Large pharma companies also use algorithms to speed up drug discovery by running tests on different compound combinations and adapting these according to the ones that yield the best results[9]. Automated processes that allow for time gains also apply to reporters. News Agency Reuters developed their own set of algorithms to scan activity on Twitter and determine which tweets are newsworthy, and beyond this, whether they are true or not. It pre-filters 700 million daily tweets to give a narrowed down list to Reuters journalists who then decide what to publish. This allows them to be among the first to report on events when they happen[10]. A last example would be in payments. Companies that offer payment services make use of machine learning to automatically detect and take adequate actions when they see abnormal purchasing behaviour with a credit card. There are just too many transactions made every second for a human to be able to control all of them, but computers can and they make the human's task more manageable with their speed[11].

People who own these machines obtain a strong competitive edge on those who don't. Machine learning helps automate many tasks which were previously performed by humans, leading to speed gains and cost reductions. It also offers automatic fine-tuning, bringing increased comfort to users of many services and more effectiveness to predictive models. These are some of the advantages brought by machine learning today. We don't know yet what it will be capable of in the future and which novel applications will disrupt our ways of living. This could actually prove to be negative for society in the longer run, explaining why we hear voices warning against the technology's potential for misuse. We will highlight some of the issues with machine learning in the following lines.

The obvious culprit, automation will lead to job losses for many people[12]. The question is whether these people will be able to adapt to new jobs that could emerge with technological innovation, assuming there are enough new jobs to cover the ones made obsolete by automation. Left unchecked, this would lead to rising inequality in our societies, which for investors means a lot of uncertainty in terms of regulatory environment. Will these companies be taxed more, will there be taxes

[8] RetinAI, 2018: https://www.retinai.com/
[9] Reuters, 2017: https://www.reuters.com/article/us-pharmaceuticals-ai-gsk/big-pharma-turns-to-ai-to-speed-drug-discovery-gsk-signs-deal-idUSKBN19N003
[10] Reuters, 2017: https://agency.reuters.com/en/insights/articles/articles-archive/reuters-news-tracer-filtering-through-the-noise-of-social-media.html
[11] Gizmodo, 2015 : https://gizmodo.com/how-banks-use-machine-learning-to-know-a-crooks-using-y-1744771152
[12] McKinsey, 2016: https://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/where-machines-could-replace-humans-and-where-they-cant-yet

on AI? The broader questions are, however, what long-run impact would this evolution have on economic systems worldwide, what megatrends could it initiate?

Another issue is that the model needs data to learn and will adapt to the data it is fed, so if there is no data or if the data is of poor quality and/or biased, there will be issues with the results. This is what is happening to IBM's Watson in some domains in healthcare. It cannot help with predictions in some aspects of cancer research because there is not enough data and not enough specialists to feed it relevant information[13]. An example of the issue with "garbage in, garbage out" in models is the case of COMPAS, a software used in Wisconsin to predict the risk of recidivism for people accused of different crimes. It has been demonstrated that it tends to rate black defendants as higher risk than what they represent and the opposite for white defendants. The reason for this is that it has been fed with data that is in itself biased, so it reflects what it learnt[14]. This illustrates that computers might be smart, but are not wise yet: they only optimise what they have been tasked to do.

Some think that computers will rise above these hurdles someday, becoming sentient and outsmarting us in the process. There is talk about the day where an AI will be more intelligent than humans and will create other AIs that are even more intelligent than itself, leading to an intelligence explosion that could leave humans far behind[15]. We do not need to start the discussion on killer robots. All this shows that machine learning touches something sensitive at our core and raises more questions than answers at this stage, as is probably the case with any radical technology. So how does one build a portfolio to benefit from the tech when there are so many uncertainties?

As was the case with our paper on blockchain technology, we believe the next SAP, Google or Facebook will use a combination of today's and tomorrow's hot technologies in a novel way that will revolutionize the way we do business and live our daily lives. Many entrepreneurs have understood this and put a mention of machine learning in their pitch decks to catch investors' eyes, and as was the case for blockchain projects, it is critical to be able to separate high-potential projects from noise. A team of engineers will have better chances at understanding the technology and if it is really good or not, but they will also profit from the views of people with a business mindset to see the innovation's commercial uses and how to approach partners and clients. This highlights again the need for thorough due diligence by a team with experience in various industries, allowing investors to add carefully selected companies to a diversified portfolio that benefits from technological changes.

[13] MIT Technology Review, 2017: https://www.technologyreview.com/s/607965/a-reality-check-for-ibms-ai-ambitions/
[14] NY Times, 2017: https://www.nytimes.com/2017/10/26/opinion/algorithm-compas-sentencing-bias.html
[15] Medium, 2017: https://medium.com/machine-learning-for-humans/why-machine-learning-matters-6164faf1df12