Business & Decision

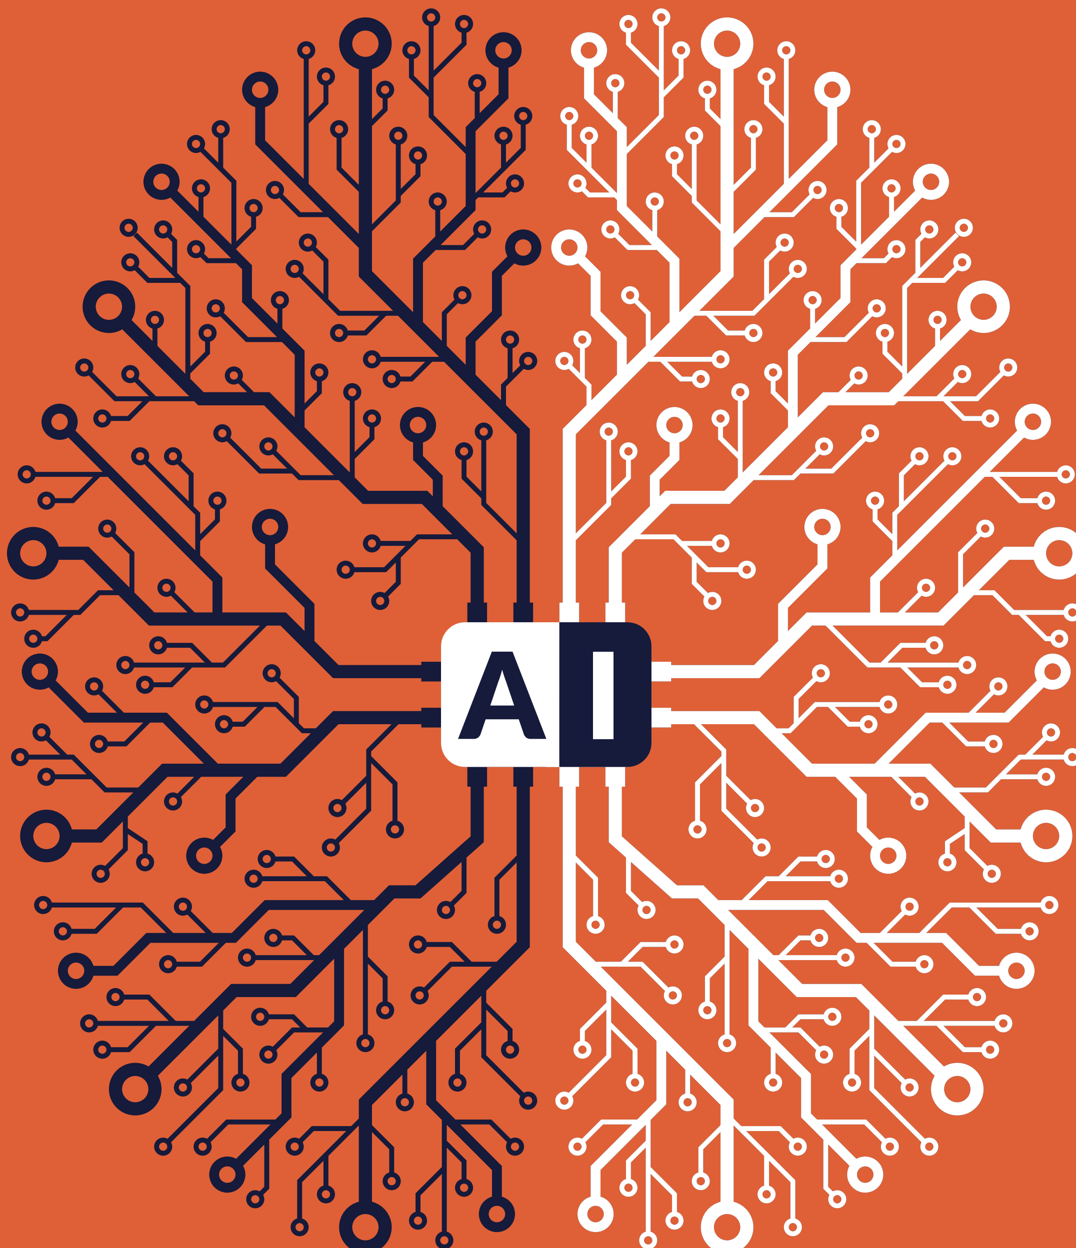

ARTIFICIAL INTELLIGENCE

Stay in control of your future!

Foreword

If there is one topic that really ignites passion and fuels all ideas and discussions in the world of new technologies, it's Artificial Intelligence. The acronym AI has become firmly rooted in our language. Beyond the cutting edge of technology, the stakes are sky-high and mainly on a societal level.

Ardent supporters maintain that it is the biggest technological advancement since the invention of the steam engine, while detractors claim it marks the end of individual and collective freedoms.
These extreme points of view, ignoring the nuances, insidiously create a sense of mystery around its origins, its use and its benefits.

As an expert in the field, Business & Decision is well-placed to shed some light on Artificial Intelligence. By sharing real-life experiences, Business & Decision aims to combine its values as a company with the practical use of this new technology.
That sums up the goal of this white paper.
You can't have AI without data, the fuel of AI. It's through teaching and transparency that fears and doubts around the use of data and algorithms can be allayed.

AI has its risks, there's no doubt about it. Our role is to highlight them, mitigate them and provide warnings if necessary. But we cannot lose sight of the fact that AI is a tool that must stay in service of people. We firmly believe that the risks can be assessed and the pitfalls avoided, for organisations in particular, when the principles of responsibility, ethics and inclusions are at the heart of AI projects.

The benefits of Artificial Intelligence are countless and still under-acknowledged. In the fields of health, ecology and carbon footprints, urban policy, industry and services, this unprecedented vehicle for innovation opens up prospects of improvement for everyone. That is, provided that it is used with total transparency and in a way that inspires trust.

On the long journey learning about all that Artificial Intelligence has to offer, we share our experience and knowledge with our clients, partners and experts. We will learn from our doubts and, sometimes, from mistakes. But we will never lose sight of the principles and values that guide what we do.

Valérie Lafdal,
CEO, Business & Decision France
Deputy CEO, Business & Decision Group

1 *A Definition of Artificial Intelligence*

2 *Intensive Learning*

3 *Risks and Opportunities in AI*

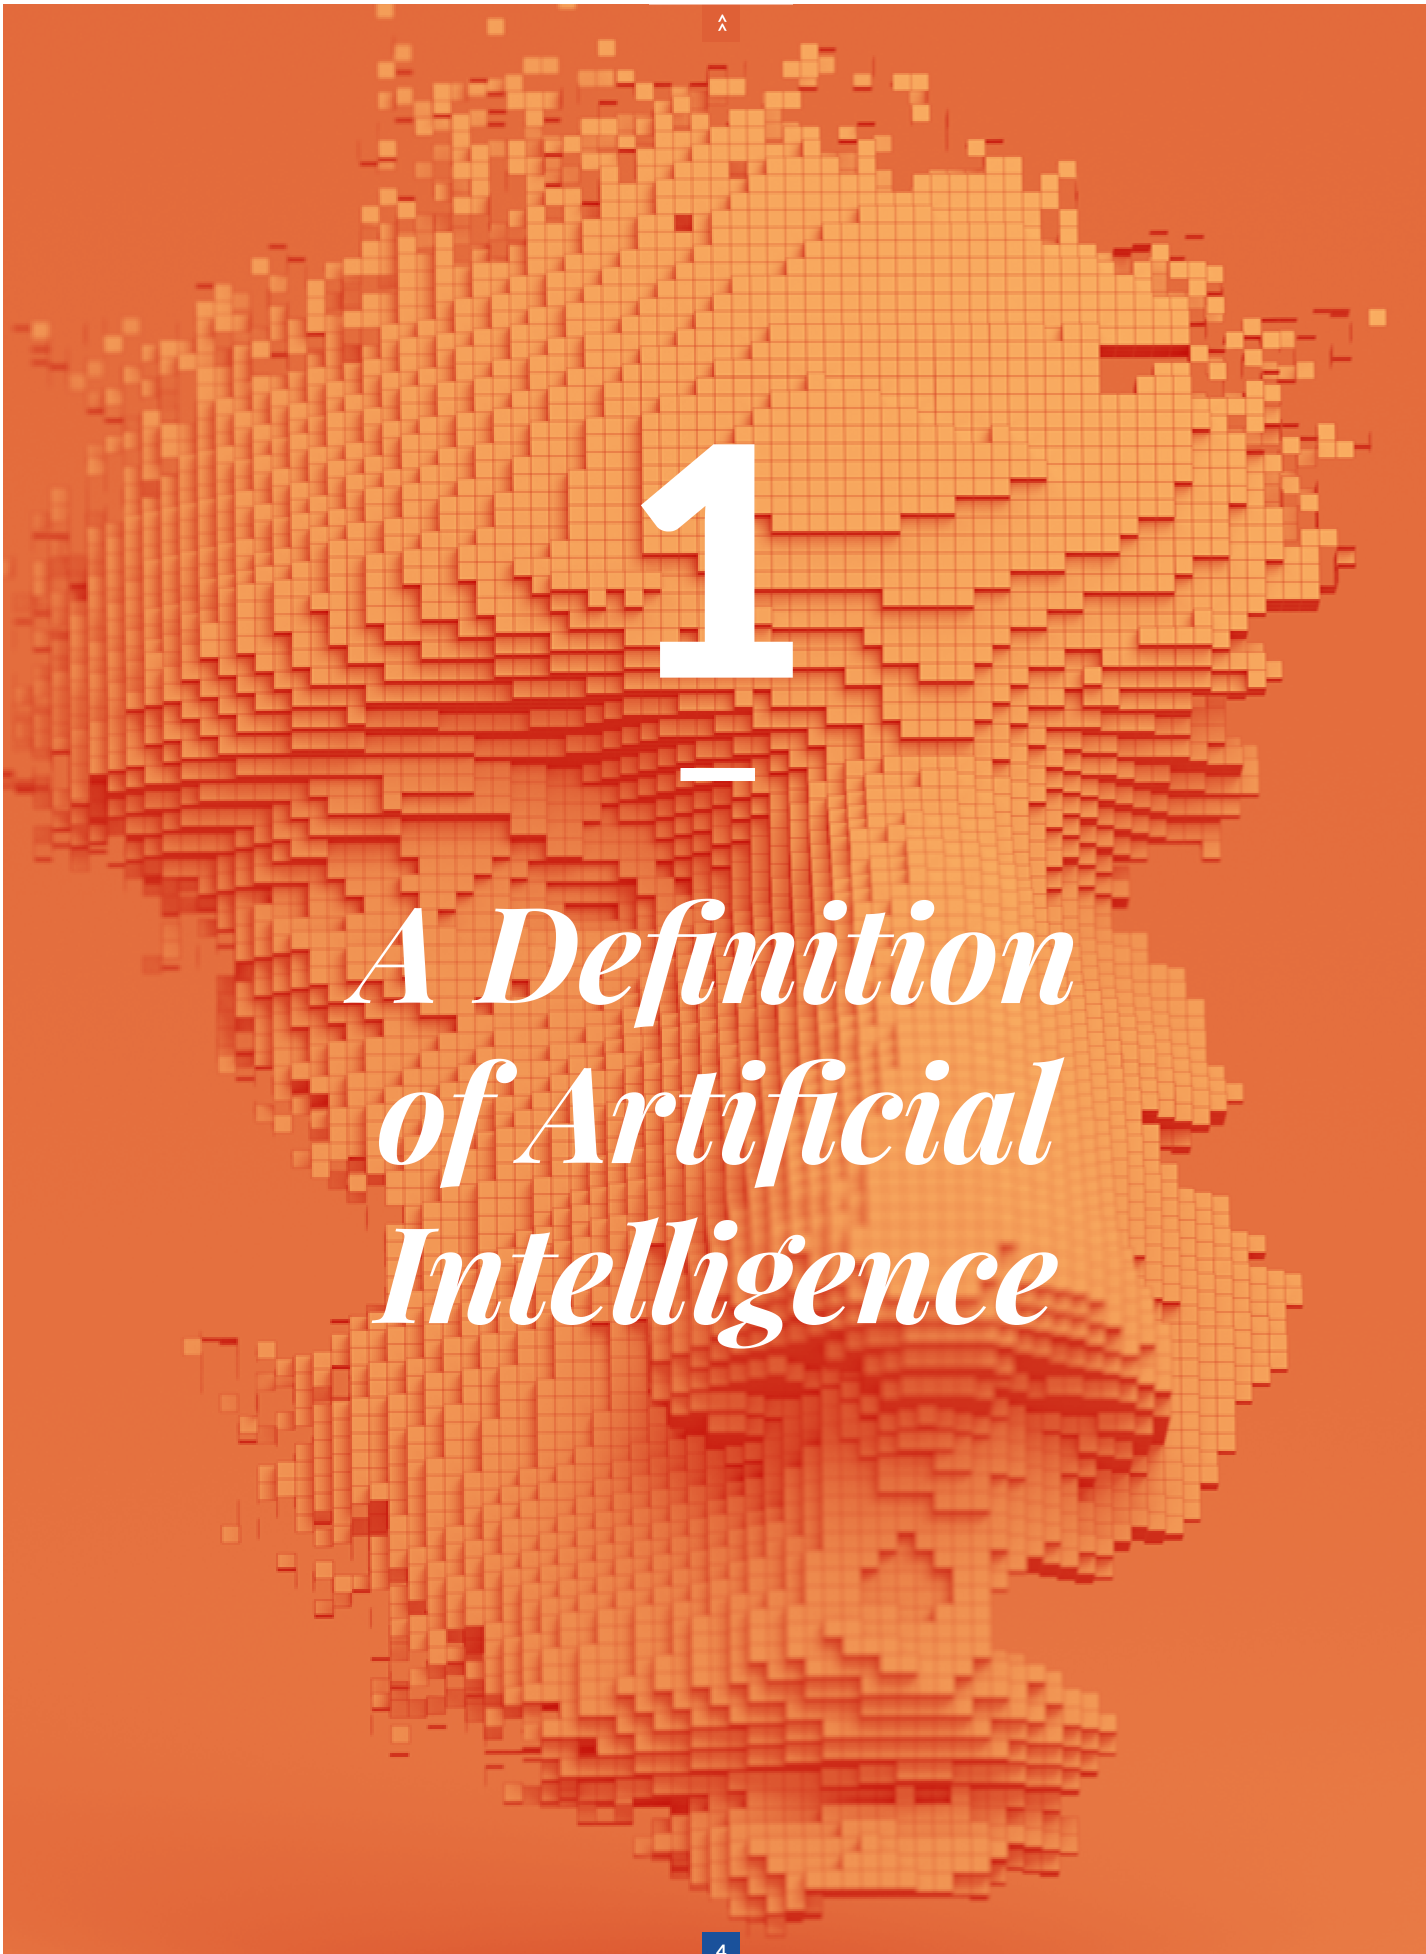

1

A Definition of Artificial Intelligence

A Definition of Artificial Intelligence

In 2020, Artificial Intelligence (AI*) Specialist was the number one emerging job in LinkedIn's Emerging Jobs Report for the USA. Meanwhile, in France, the role of Artificial Intelligence Engineer was ranked in second place among the most in-demand jobs[1]. Recruitment of AI specialists has increased by 74% over the last four years. A trend which is set to continue: global spending is expected to double over the next four years to exceed 110 billion dollars in 2024[2]. France has been the number one country in Europe for investment in AI since 2019[3]. But why is Artificial Intelligence, a term first coined in the '50s, attracting so much interest today? Let's look back at some of the underlying principles.

RECRUITMENT OF AI SPECIALISTS HAS INCREASED BY 74% OVER THE LAST 4 YEARS.

The Birth of Artificial Intelligence

When we refer to Artificial Intelligence, what do we actually mean? According to the definition given by the prestigious French dictionary Larousse Encyclopaedia[4], it is a *"set of theories and techniques used to create machines capable of simulating human intelligence"*. We need to take a look back at the history of AI, sorting fact from fiction, to better understand the risks and opportunities around it in the 21st century.

Background in the 1940s and 1950s

It was back in the mid-1940s that neurologists Warren McCulloch[5] and Walter Pitts[6] published their first essays on what they considered, at the time, to be mathematical modelling of the brain.

Their work inspired a number of researchers such as Franck Rosenblatt[7], John McCarthy[8] and Marvin Minsky[9], who aspired to model and electronically reproduce brain functioning. In light of the early machine learning results obtained with artificial neurons (later named 'perceptrons' by Rosenblatt), researchers became convinced towards the end of the 1950s that putting these electronic neurons into a 'network' would make it possible to imitate the way the brain functions in order to solve infinitely complex problems.

Why the term 'Artificial Intelligence'?

At the time, in order to fund their research, McCarthy and Minsky came up with the term 'Artificial Intelligence', which was introduced at a conference in Dartmouth (Dartmouth Summer Research Project on Artificial Intelligence) during the summer of 1956. AI was then not only officially created as an academic discipline, but also sparked great interest amongst American investors.

A false start?

Following the conference in Dartmouth, McCarthy and Minsky received very substantial funding. But this surge of interest would only last until the end of the 1960s, when research in the field started to sink gradually into what is known as 'the first AI winter'. Why? Because the results up to that point were extremely disappointing in comparison with the promises initially made in Dartmouth.

There are several explanations for this: on top of the unavailability of sufficient computing power, AI demonstrated deficiencies directly linked to the intrinsic limitations of a machine compared to human intelligence, as Luc Julia explains (see box on page 7). The reasons behind this over-ambitious approach are now somewhat clearer and it is important, to avoid falling into another AI winter, to have a good understanding of them.

THE 1970S MARKED THE 1st AI WINTER.

First of all, the initial mathematical modelling was no doubt much too simplistic. The further research on the brain and neuroscience progresses, the more unexpected complexities and unanswered questions emerge. And even if we were able to fully understand the brain's structure, which is far from the case, we would still be missing the most important thing: understanding of its function, that is to say, in very simple terms, the 'software' that makes it work.

To illustrate this point, we can take the example of the hypothesis, increasingly accepted by some biology researchers, that millions of years of evolution is written into each and every one of our cells, including, notably, our neurons. I'm sure you'd agree that if this is the case, it seems that it would be impossible to model it with a finite quantity of memory. That is why many AI researchers have recently moved away from McCarthy and Minsky's initial concept, although the name 'Artificial Intelligence' continues to be used.

After a period dominated by expert systems* in the '80s, AI went through another slump - known as the 'second AI winter' - between 1987 and 1993, related to how it was perceived in the eyes of investors and government agencies. It wasn't until the 2000s, and more specifically 2012, that Deep Learning* (see glossary) brought new momentum to the discipline.

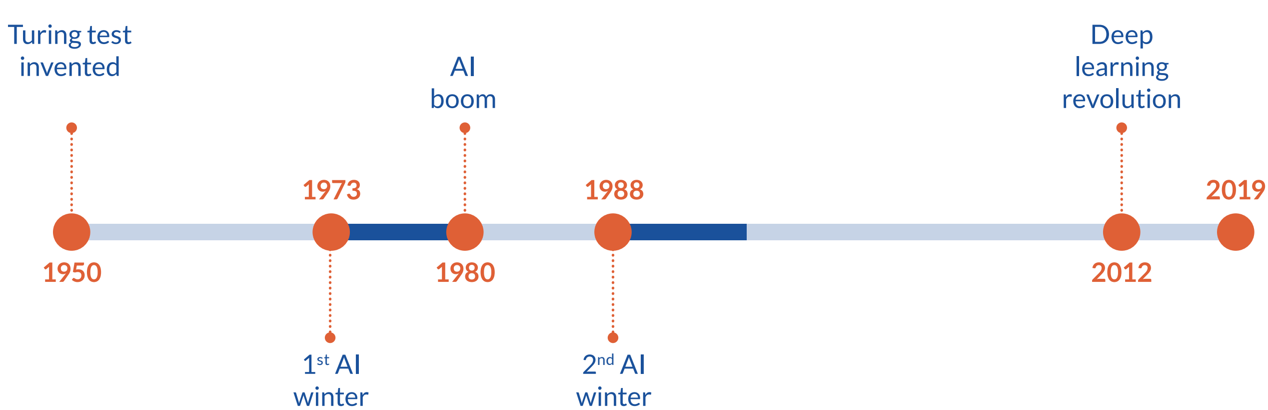

The Expert's View

Luc Julia,
CTO of Samsung Electronics,
co-creator of Siri and author of *There is no such thing as Artificial Intelligence*
First Éditions - 2019

Take the example of the self-driving or autonomous car: level 5 autonomy [editor's note: driving completely independently without the help of a driver in all circumstances] will never exist. Why? Because there will always be unforeseen situations that humans can adapt to while machines, on the other hand, can't. A self-driving car on the busy Arc de Triomphe roundabout in Paris at 6pm won't go anywhere because the rules of the road are not followed there. Machines never invent; all they do is apply what they have been taught or what they have observed in the models* they have been given.

What are the differences between humans and AI?

Contrary to common misconceptions, Artificial Intelligence is not intended for (or even capable of) replacing humans. To date, there are three major differences between humans and AI:

Never-before-seen • Opinion • Ethics

Here are the three major limitations of AI for **Michael Deheneffe**, Strategy & Innovation Director, Business & Decision:

"AI is not capable of reacting to something it has never seen before: it can't react if it hasn't learned how. This means that a single grain of sand that slips between the cogs has the potential to disrupt its whole ecosystem. The second limitation is around the idea of opinion, or the fact that AI doesn't have a point of view. Humans understand and create; AI learns and executes. In reality there is no true intelligence in AI. Finally, Artificial Intelligence has no concept of ethics or justice. It is incapable of differentiating between good and bad, and, accordingly, of deciding for itself on the best course of action."

Definition of AI

In view of the major differences between artificial and human intelligence, it seems that, in the short to medium term, AI can only exist in support of humans. A position that the modern definition of AI tends towards.

Stuart Russell & Peter Norvig's definition

Based on the definition published in 2010 in the textbook *Artificial Intelligence: A Modern Approach* by Stuart Russell & Peter Norvig, Artificial Intelligence is a discipline devoted to building intelligent agents*, which are entities (physical and/or software), with some degree of autonomy, capable of perceiving their environment (through sensors) and acting upon it (via actuators). They can also analyse and make decisions with a view to fulfilling objectives.

Stuart Russell & Peter Norvig's book specifies that AI as defined here works within a specific environment, and this is also sometimes defined as 'artificial narrow intelligence', or 'weak AI', as opposed to the concept of 'artificial general intelligence' or 'strong AI'. The idea of strong AI does not really have any significance in the real world, at least as research stands currently, since AI, as we showed previously, does not work outside of its environment. Changing the environment of an Artificial Intelligence system necessitates redoing at least a very large part of its learning within this new environment. We will therefore work on the assumption from here on out that **all AI is weak AI**.

Artificial Intelligence according to Business & Decision

In light of the above, Business & Decision proposes the following definition of Artificial Intelligence:
*"Discipline devoted to the building of agents **that interact** with the environment and can **learn** from initial data provided and/or data collected during interactions to carry out a function **in a specific environment**, with some degree of autonomy."*

In summary, AI captures data, analyses it and interacts with the environment (according to Peter Norvig's definition). But we add to that the ideas **of interaction** and of capacity **for learning** in a **specific environment**. If we think of it this way, Artificial Intelligence is based on three components: **interaction, analysis and learning** to support decision-making.

"AI is humans giving machines four skills.
Perception first of all, to allow it to recognise shapes, images, text, etc.
Then interaction, which can be done in a number of ways, from simply displaying messages on a screen to the most advanced humanoid robots.
These first two abilities enable the AI system to 'live' in the real world.
The two other abilities are related to the field of cognition. Firstly, analysis, which can be at three levels: descriptive, predictive and prescriptive. And finally, learning.
This last ability is fundamental as it is the basis for all of the others.
In reality, AI systems are only able to do one thing: learn."

Mick Levy,
Business Innovation Director, Business & Decision

The Expert's View

Luc Julia,
CTO of Samsung Electronics, co-creator of Siri
and author of *There is no such thing as Artificial Intelligence* -
First Éditions - 2019

"The term AI was used for the first time in 1956. But unfortunately, the scientists set out to tackle the most difficult problem, a problem that we still haven't solved: natural language understanding* (NLU or NLP).. This resulted in the research being abandoned in the '60s and we entered the first AI winter. So we made expert systems for 30 years, until 1997, the year an expert system was able to beat Gary Kasparov at chess for the first time. Then the internet enabled large volumes of data to be generated, allowing research on neural networks to be resumed. In 2016, the world Go champion was also defeated, by Google's DeepMind. This doesn't mean however that we can liken AI to intelligence; human intelligence is a continuum. The brain is infinite. So no, Artificial Intelligence is not going to replace humans. Descartes said: "Language is unique to humans". So it doesn't belong to machines! In this regard, the Artificial Intelligence that people are scared of, that is supposed to be able to seize control from humans, doesn't exist! There is, however, a form of AI that exists and that we have been working on for 30 years. While artificial general intelligence or 'strong' AI doesn't exist, Artificial Intelligence, in the sense of weak or narrow AI, does. Even using the mathematical and statistical* techniques of today, we will never manage to get even close to human intelligence. At best, we can talk about augmented intelligence - or intensive learning - through the tools that we use to improve, physically or intellectually, like with a hammer: it is up to the human being to decide to put it to good or ill use... and the same for Artificial Intelligence."

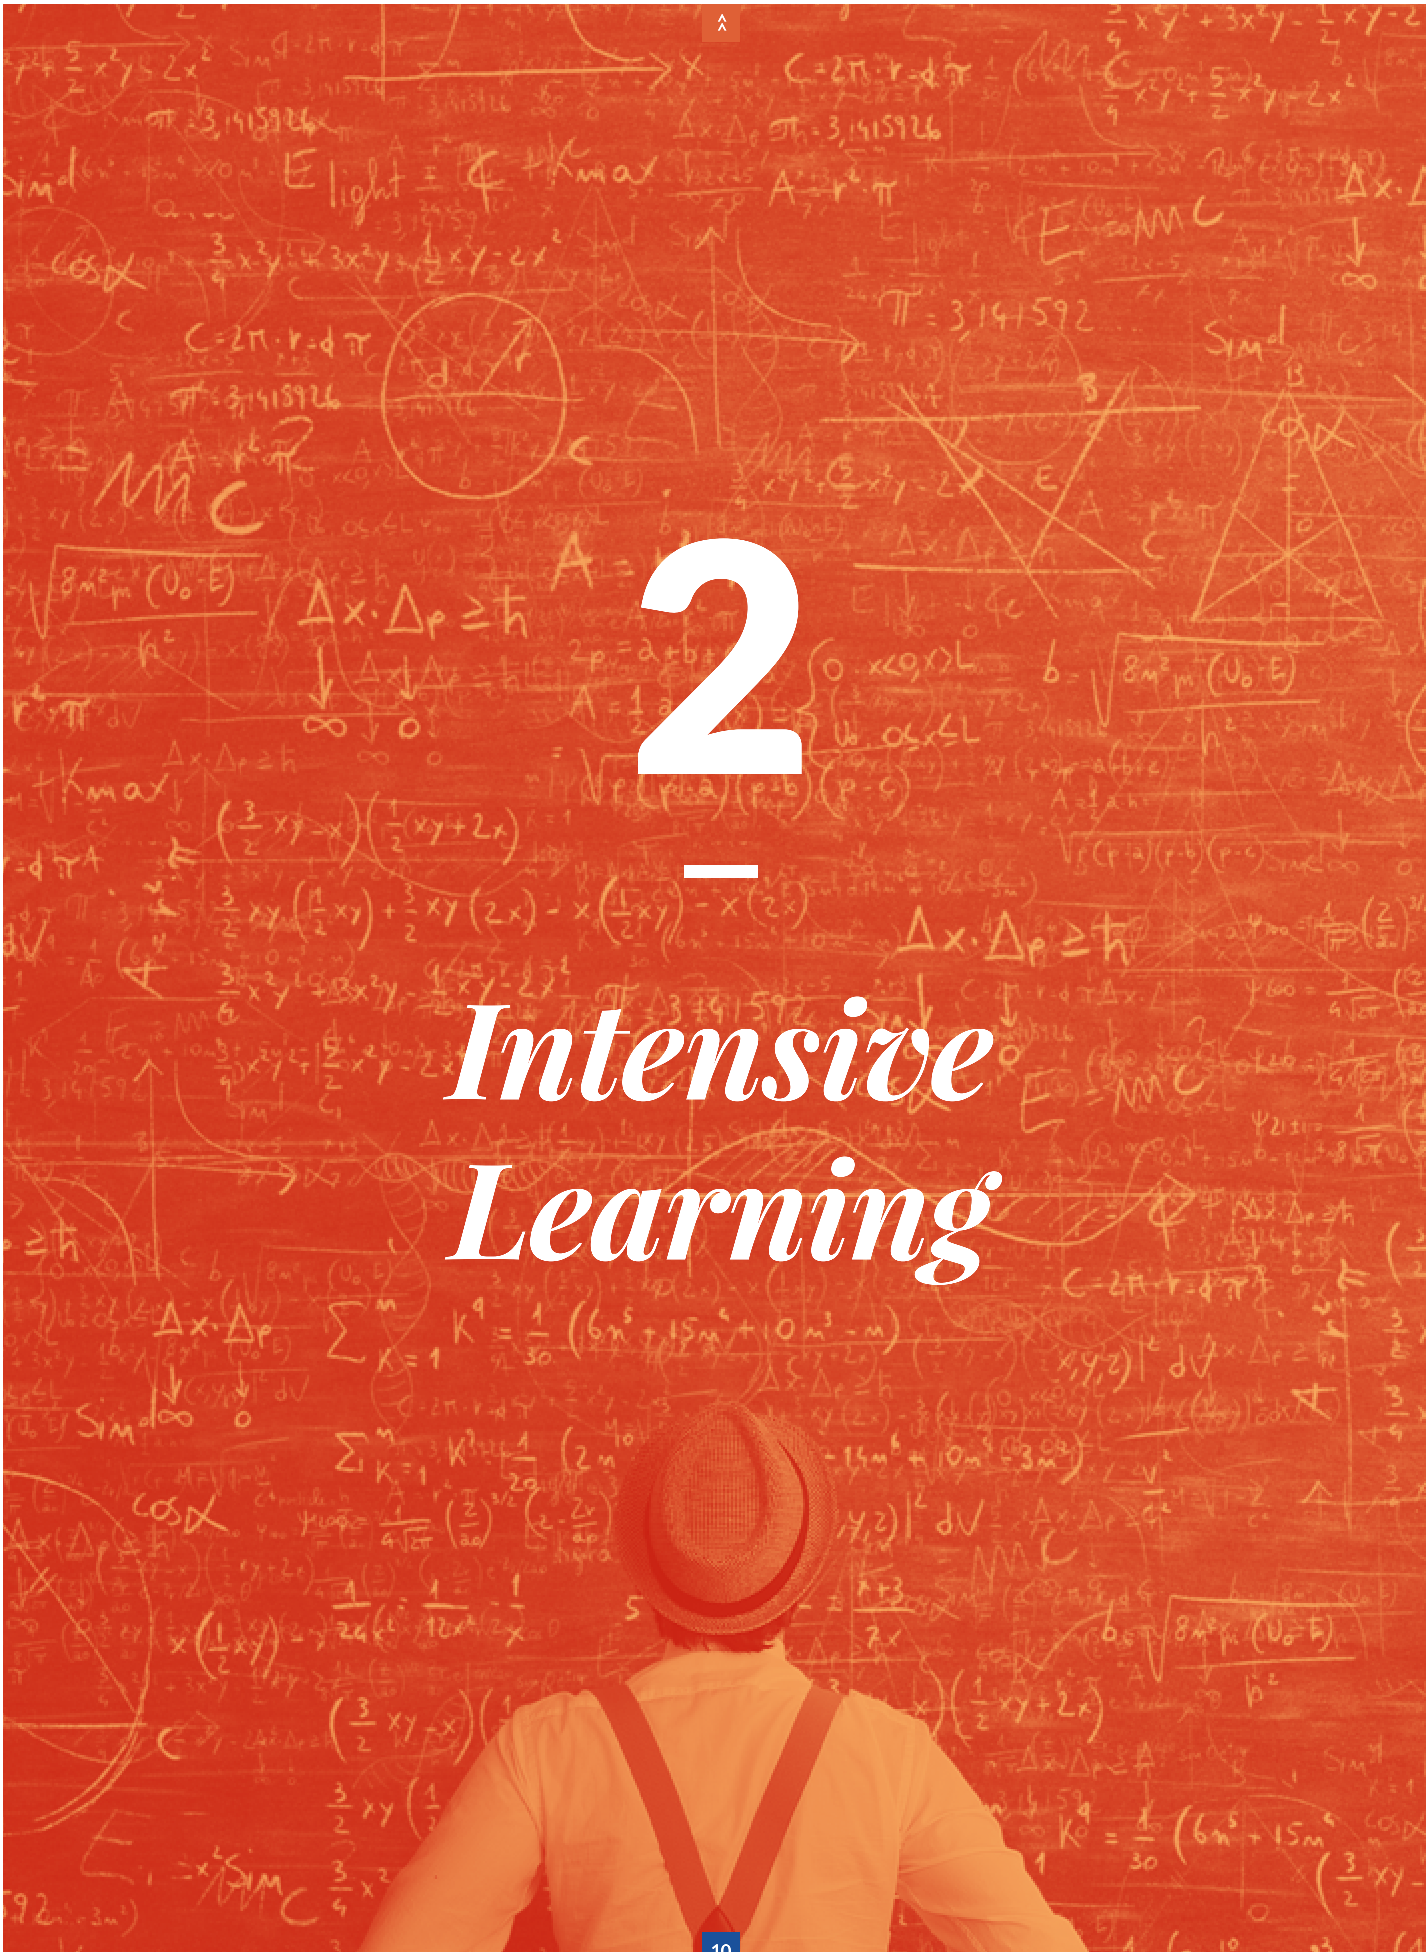

2

Intensive Learning

Intensive Learning

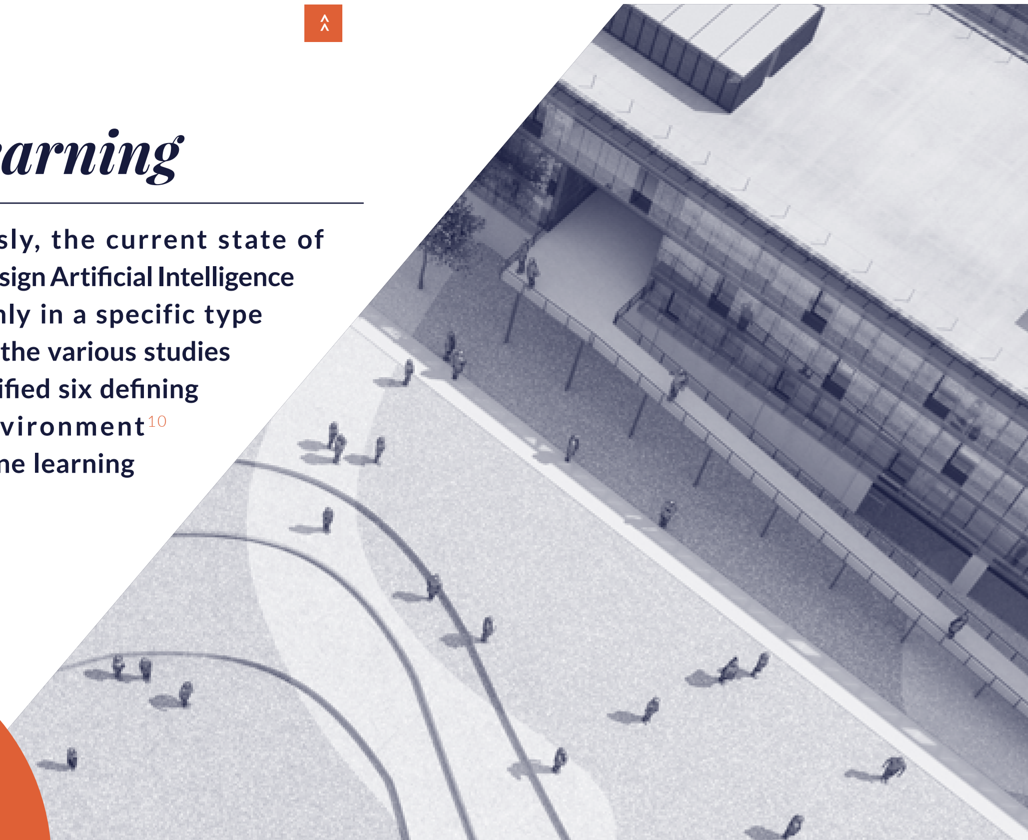

As highlighted previously, the current state of knowledge enables us to design Artificial Intelligence systems that function only in a specific type of environment. Based on the various studies conducted, we have identified six defining properties of an AI environment[10] and three types of machine learning techniques.

THE ENVIRONMENT OF AN ARTIFICIAL INTELLIGENCE SYSTEM IS VITAL AS IT IS AN INTEGRAL PART OF THE DESIGN AND LEARNING PROCESS.

AI Environments

The environment of an Artificial Intelligence system is vital as it is an integral part of the design and learning process. AI works according to the environment in which it was created. That's why it is important to consider the environment when designing Artificial Intelligence systems to make them effective with regard to the desired objectives. Separated from their environment, they lose all their meaning and, most importantly, their effectiveness.

In order to 'correctly' build an AI solution that will meet an organisation's challenges, it is necessary to take into account six main properties that can affect its environment (and therefore the AI system itself).

1 | Accessible/fully observable versus inaccessible/partially observable environments

An environment is said to be accessible (or fully observable) if **all data** on the environment is accessible to the AI system, for example with chess or the game Go. In real-life situations, this is usually not the case, with part of the environment often remaining hidden from the AI system.

2 | Deterministic versus stochastic environments

If the current state and the action taken by the AI system allow it to predict with certainty what the state of the environment will be following the action - for example, turning off a motor stops a machine - we say that the environment is deterministic.

Conversely, if the action performed by the AI system could lead to a number of possible results - for example, the braking of a car where the effectiveness depends on the road conditions - this is called a stochastic environment.

Environments that involve a very large number of different factures are stochastic by their very nature, since the multiplication of factors leading to millions of variables makes complete control of the situation by the AI system difficult or even impossible. The algorithm must therefore work as though what it cannot control is subject to chance, which heavily impacts the type of learning technique.

3 | Episodic versus sequential environments

Behind this concept lies an important characteristic with regards to the way an AI system can process data and actions. In an episodic environment, every business action by the AI system is divided into ‘mini-episodes’ where only the current state of the environment can be taken into account.

In other words, each action by the AI system is a new episode resulting from a new set of perceptions totally independent of the previous and following ones. A typical example is that of a machine sorting good items from bad ones without any relation to the previous items. In this situation, time is just one variable amongst others: the concept of anticipation therefore doesn’t really have any importance in an episodic environment. As soon as there can be consequences beyond the state immediately following the next action, we say that the AI system is operating in a ‘sequential’ or ‘temporal’ environment. A large number of everyday applications take place within a sequential environment where the variable of time plays a central role.

4 | Static versus dynamic environments

If the environment in which an AI system operates does not change over time, but only as a result of the actions taken by the agents present, we call it a static environment. On the other hand, when time plays a role in the environment changing, for example with uncontrolled moving objects, we call it a dynamic environment.

5 | Discrete versus continuous environments

The concept here is close to that used in mathematics. If the environment has a countable series of possible states, we call it a discrete environment.

Otherwise, we have a continuous environment. A plane seat, for example, is typically a continuous environment as there is an infinite, uncountable number of possible seat positions.

6 | Single agent versus multi-agent environments

This simply refers to the number of agents (in the general sense of things which can act) that can influence the environment taken into account by the AI system. The single agent case is not often relevant since only the AI system can act upon the environment (a situation that, when it comes to it, rarely occurs).

There are different types of multi-agent environments, depending on whether the agents are competing with each other or collaborating.

The Expert's View

Didier Gaultier,
Data Science and AI Director, Business & Decision

The properties of the environments in which AI systems operate are numerous, sometimes interlinked and often complex. The unobserved can act upon the environment of an AI system at any time and thus modify the direct effect of its actions, so it tends to make the environment stochastic.
Likewise, when an environment is sequential, every action by the AI system can have possible repercussions on subsequent ones. The AI system therefore needs to be given the ability to predict in order to enable it to anticipate situations that could occur much later on. A simple example is a vehicle increasing its speed. This will inevitably change the braking time and distance later if, for example, it comes across an obstacle.
The complexity of the development of an AI system is directly related to the complexity of the environment; if you have an environment which is partially observable, stochastic, sequential *and* dynamic - which is the case in many real-life situations - an AI system's task can prove to be deeply complex. With current technology, an AI system cannot be dissociated from its environment. It is therefore vital to know how to analyse the type of environment you are faced with before deciding which type of learning technique and algorithm* should be used.

Machine Learning Techniques

Like with Artificial Intelligence environments, there is more than one single type of machine learning. There are three main categories, each with their own characteristics and applications.

Supervised learning

HOW IT WORKS

This type of machine learning, which most likely owes its origins to Carl Friedrich Gauss in 1795, is now by far the most commonly-used type of learning technique in Artificial Intelligence.

There are two different types of supervised learning depending on the type of variables* to be predicted. If it is one or more numerical variables*, it is known as 'regression'. If, on the other hand, it is a categorical variable* (which takes values from a predefined list), it is known as classification or a classifier.

The concept of labelled (or annotated) examples is fundamental for this type of algorithm, as it involves training the model using examples for which the expected output is already known. We say that they are labelled because we give the algorithm the output in the form of a label added to a piece of data. In reality, this label is contained within the target variable to be predicted, but **only for the known items** that serve as a starting point (which as a whole make up what we call the training dataset). Then, the algorithm predicts the value of the unknown items itself.

In classification, the model has to predict one or several outputs from predetermined lists. This is known as supervised classification (each possible value of an output corresponds to a predefined category). It has a huge number of possible applications ranging from the prediction of a simple numerical value to extremely complex robotics operations.

"SUPERVISED LEARNING IS A MACHINE LEARNING TASK CONSISTING OF LEARNING A PREDICTIVE FUNCTION BASED ON ANNOTATED EXAMPLES."*

SOURCE: WIKIPEDIA

EXAMPLE USE CASE

Port optimisation is one of the most interesting, even remarkable, examples of supervised learning. We have implemented it for one of the largest cargo ports in Europe.

The loading and unloading of container ships requires flawless organisation and logistics. If the port operator doesn't finish these tasks in time (for example, if the container ship has had to wait in unexpected queues), the port will be subject to multiple fines, not to mention the fact that delays tend to have a knock-on effect on the loading (or unloading) of subsequent ships. The total amount of these fines can end up being astronomical if port operations are constantly delayed.

However, it is always difficult to predict the estimated time of arrival (ETA) of a vessel at the port. Multiple factors such as the weather, route congestion, badly-managed queues or simply a lack of forecasting can completely disrupt a port's logistics. That's where Artificial Intelligence comes in.

The project implemented by Business & Decision consisted of using a supervised algorithm to, as a first stage, make a precise prediction of the ETA. There were of course a large number of types of data used such as the weather, route congestion, the different queues, etc.

There were also a huge number of benefits gained. These included: a dramatic decrease in delays and the multiple fines that result from them as well as a mutually beneficial situation for all parties involved in the running of the port (the port, the operator, shipping companies, customs, etc.).
From a financial perspective, the savings made through the AI system in a case like this can be well over a million euros per year. This means that a project like this can become profitable in the space of a few quarters.

WHEN TO USE

To sum up, supervised learning should be used:

- When we need to make a prediction (variable to be predicted) where the features are fully known, whether it is a value from a list or a numerical value (in that case, the variable to predict must be structured).

- When we have a sufficient known historical dataset of the variable to be predicted. This is a labelled historical dataset (where we know the value to be predicted for a training dataset).

APPLICATIONS

There are many possible applications of supervised learning. To name a random selection:

Approving bank loans	Calculating risks	Estimating the effectiveness of marketing campaigns
Predicting equipment failures	Calculating insurance premiums	Diagnostic support
Detecting anomalies	Predicting attrition (loss of customers)	Determining market prices, sales prices for a trading platform, etc.
Detecting fraud		

The special case of time-series data

Supervised learning takes a particular form when used to study the evolution of data over time. In this case, the data is a ‘time series.’ The algorithms used for supervised learning involving time-series data are different to ‘classic’ supervised algorithms. There are also greater constraints on the training data. The data must satisfy the ‘interval data’ rule, which states that the measurements must be taken at regular intervals (every minute or every hour, for example). Not following this basic rule very often leads to data becoming difficult or impossible for algorithms specialised in time series to handle. As seen previously, the time axis is of particular importance in AI. If it is altered or inaccurate, this can very quickly prevent an AI system from functioning correctly[11].

Unsupervised learning

HOW IT WORKS

Unlike in supervised learning, in this type of technique the training data is not labelled (or annotated) in advance. The algorithm must therefore work out for itself what work needs to be done. This learning style can be traced back to Karl Pearson in 1901. It is primarily based on detecting similarities and differences between individuals. There are a number of different variations.

EXAMPLE USE CASE

The example use case of **retail outlet typologies** given here was carried out by Business & Decision for several retail chains. The first step is creating a database of all of a chain's outlets. In this data hub will be included the physical features of an outlet and those of its customer catchment area: proximity to public transport facilities, the socio-economic characteristics of the population, historical sales data, etc. By using, for example, anonymised data to comply with the GDPR, it is even possible to integrate the characteristics of people passing through the area, visitor numbers, and significant events that could influence the performance of the outlet (the list given here is not exhaustive).

Once this data has been carefully added to the data hub, an unsupervised algorithm will segment the outlet database. How finely the data is divided up can be adjusted according to the end goal. In this case, a B2C chain succeeded in doubling one of its major outlets in size thanks to this process. How? By applying first unsupervised, then supervised analysis of sales potential. The information used by the supervised algorithm to make its predictions included the segments given first by the unsupervised algorithm and the data in the data hub. It turned out to be possible, through optimising the capacity of the outlet (physical features) and its attractiveness, to double the number of daily sales.

WHEN TO USE

To sum up, unsupervised learning should be used:

- When we need decision support or prediction where the features are not fully known and we are looking to create new, relevant, structured indicators.
- When the features involved are more important than a set of historical data on the variable to be predicted.
- When we have large amounts of information, for example in multi-factorial environments.
- When the problem would be simplified or resolved by finding new features that are simple and shared across different items and grouping items according to their features.

APPLICATIONS

The possible applications of unsupervised analysis are endless. It is possible, for example, to gain a better understanding of why a certain client has moved upmarket more easily than another. The causes of proliferation of a virus can be analysed using this type of algorithm (because they often involve multiple factors), thereby paving the way for a supervised analysis to make epidemiological predictions.

One of the main applications in B2C is smart customer segmentation for the launch of marketing campaigns. The power and precision of unsupervised algorithms sometimes goes beyond expectations. For example, in 2015, a major US retailer (Target near Minneapolis) was able to predict that a 16-year-old customer was pregnant (despite not having this information directly). It sent her offers based on this (maternity clothes, etc.), prompting anger from her parents towards the store... until they eventually discovered that their daughter really was pregnant.

Beyond these examples, unsupervised learning must often be seen as a preliminary analysis that will allow subsequent supervised analysis to work better. Unlike with supervised analysis, it doesn't require you to have labelled historical data.

In the context of detecting insurance fraud, Business & Decision was able to help a client that didn't have historical data on fraud and enable supervised prediction to be carried out. Unsupervised learning with different anomaly scores was able to separate into certain segments the files that needed to be studied manually to potentially find cases of fraud. Once analysed, these clusters revealed over €100,000 in proven fraud. But most importantly, they were then able to serve as a training database for supervised analysis.

THE POWER AND PRECISION OF UNSUPERVISED ALGORITHMS SOMETIMES GOES BEYOND EXPECTATIONS.

We must also point out that in most cases, the results of unsupervised analysis must go through a very in-depth process of interpretation by a data scientist. In the example above, it was through the work of humans that the real cases of fraud within the clusters were identified and a database for supervised learning was eventually obtained.

Some randomly selected examples of applications:

Automated data exploration	Targeting marketing campaigns	Carrying out studies or surveys
Simplifying databases (reducing the number of variables)	Identifying types of anomalies, failures or fraud	Splitting products or outlets into categories
Segmenting a customer base	Setting up management dashboards	Identifying possible types of fraud
		Profiling customers who leave (or will leave) your brand

Reinforcement learning

HOW IT WORKS

Reinforcement learning was first introduced by Richard Ernest Bellman in 1957. It's a fascinating type of machine learning in more ways than one, since it's the one that enabled the Alpha-Go programme to beat Fan Hui, the French Go champion, in October 2015, then the world champion, Ke Jie, in May 2017.

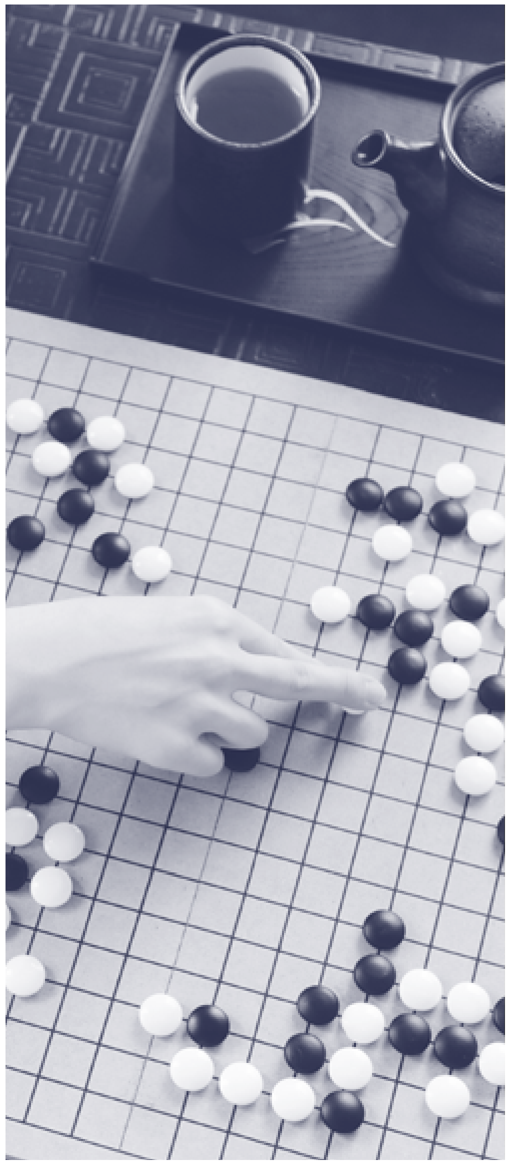

This approach is fascinating for many marketers in that it serves as the basis for a number of Next Best Offer or Next Best Action recommender systems. It also (sometimes incorrectly) referred to as 'self learning' or 'active learning'. In reinforcement learning, it is essential that the agent is 'coached' for it be able to make progress. In the case of a chess match, the coaching comes from the information 'match won' or 'match lost.' The AI system must lose millions of times (generally more) before it wins even one match.

In this type of learning technique, unlike the two previous types (supervised and unsupervised) the system learns as it goes along (instead of learning before it is used). There are many, constantly developing applications. In practice, they are limited by the fact that errors are part and parcel of this technique's DNA. As soon as a new situation arises, the algorithm will have to explore a whole universe of possible errors and advance blindly before it finds the most appropriate response.

If you were to put this type of algorithm in control of a self-driving car, it would have to drive the car into a ditch (and survive it) dozens of times in order to learn that a ditch is an obstacle to be avoided. On the other hand, in a relatively controlled environment, like a board game or recommending products to customers, it can give outcomes that haven't been seen before and that weren't necessarily expected (as was the case with Alpha-Go, which has become unbeatable).

In the case of a marketing recommender system, the coach is no other than... the customers themselves. The information used to coach the system is, in fact, the act of buying. If the client buys something, this sends positive feedback to the algorithm. If they do not, that information is also taken into account. Some recommender systems can be pre-trained before they are actually put to use. Recommending the wrong product to a customer from time to time is, at the end of the day, an 'acceptable' price to pay to obtain an algorithm capable of continually adapting to the market.

EXAMPLE USE CASE

Let's take the example of an NBO-type (Next Best Offer) online offer recommender system on the web shop of a major retailer, created by Business & Decision. The recommender system presents visitors with an initial offer (which they can click on). Then, they are asked to perform a search. As soon as the customer enters the first features they want to search for, the recommender system offers them several personalised offers which better match their preferences and requirements.

This type of NBO system was run alongside and in competition with a rule-based system, a type of system more traditionally used in this kind of application. **After a few months, in comparison with the rule-based system, the retailer doubled conversion rates and tripled income generated online**: not only was the website making twice as many sales, but this was also with a considerably higher average basket value. The rule-based system was therefore 'disconnected' to let the recommender system do its job.

IN A FEW MONTHS, CONVERSION RATES WERE DOUBLED AND THE TURNOVER GENERATED ONLINE WAS TRIPLED.

One of the advantages of this solution is that it is capable of adapting to changes in the market, since it never stops improving. And it makes what every business dreams of a reality, as it still improves every time a customer doesn't buy something!

WHEN TO USE

For reinforcement learning techniques to be used, all of the following conditions must be met:

- If an action fails, the consequences must not be too serious (with the current state of technology, the use of this type of learning technique in the control of fast-moving machines or vehicles must be avoided).
- This type of machine learning requires a coaching system to be present at all times to indicate, after a certain number of actions, if it has made a mistake or not (a customer buys something, a match is won, a human coach, etc.).
- There must be a relatively long period of time available for training, during which the system will go through an excessively high number of failures.
- The environment can be deterministic or stochastic; reinforcement learning is good at managing and making use of uncertainty over the result.

APPLICATIONS

We are far from having explored all possible applications, especially as technology is regularly advancing in this field. We should mention that this type of machine learning, like the recommender system given as an example, can sometimes be used alongside another type during the initial training phase. It should also be noted that the training phase can be 'frozen' at a certain point if it is believed that the environment is not going to evolve much more going forward.

Common applications include:

Controlling 'slow' types of autonomous moving machines that don't present any major risk of damage in the event of a collision during training. A typical example is a robotic vacuum or another cleaner, which, thanks to this approach, will be able to learn the layout of its surroundings. There is no danger since, with its slow speed, its sensors will warn it about the type of obstacles before there is any risk of a collision. Other examples of this type of device include robotic lawn mowers and floor polishers.

A lift control or automated queue management system can be managed according to this principle (provided that the consequences of failure are limited to longer waiting times, and don't include physical damage or breakdowns).

WE ARE FAR FROM HAVING EXPLORED ALL POSSIBLE APPLICATIONS.

There are many applications in robotics that allow a robot to learn to walk, run, jump, etc.

All the programmes capable of playing various board games also fit into this category.

Prioritising tasks is another area where this type of learning technique can be used.

The main application of this type of algorithm remains of course in recommender systems in marketing. There are Next Best Offer systems, the most common, which generally suggest one or more products for purchase, and Next Best Action systems. The latter go one step further, since they can suggest something other than a purchase, for example speaking to an agent or salesperson in order to resolve or anticipate a problem.

Finally, we should mention applications, still in the research phase, in the control of semi-autonomous vehicles. The reinforcement learning is 'frozen' after a certain period of time, meaning that there is no chance of the system exploring new solutions that could be dangerous while in use. This technique has notably been tested by Tesla, but remains largely confined to research, especially in combination with neural networks.

The special case of deep learning

HOW IT WORKS

The first research on deep learning, based on what are known as 'convolutional*' neural networks, dates back to the 1970s. However, the concept didn't start to really be taken seriously by the community until 2012, following the work done by Yann Le Cun starting in the 2000s. The way they work is based on similarities with the visual cortex.

Applications in the analysis of non-structured data are near-infinite. They range from image and video recognition to interpretation of semantic and textual data and vocal recognition. Deep learning is the technology that brought about the revival in the field of Artificial Intelligence. It should be noted that deep learning can fit into the supervised, unsupervised or reinforcement learning categories.

EXAMPLE USE CASE

The following example use case, carried out by Business & Decision, is an unusual one in that it took place in the world of fashion. It's the case of a somewhat unusual recommender system as it can identify, based on a photo of a customer, items that could fit with their dress style.

It is different in this way to traditional recommender systems (called 'look alike' systems) which only take into account past purchases by other customers and internet-users. These can certainly give valuable insights, but, as they can only derive these from purchases that have already been made, they are not able to cover all possibilities, especially in fields such as fashion where the combinations are virtually infinite. Deep learning, on the other hand, is able to make unconventional but nonetheless relevant suggestions based on style, colour, size, texture, etc.

WHEN TO USE

One of the characteristics of deep learning is that it requires, in the supervised learning style, a very large training database in order to be able to give relevant results. It is essential for this training data to be labelled. In short, if we want it to recognise images of objects or animals, each image we give it during the training phase will need to be assigned a label indicating what type of object or animal it is.

This is currently the most widespread and usable application of deep learning. There are other types of deep learning, such as unsupervised or even self-supervised learning. But they remain in the phase of basic research and so will not be covered here. Once the initial training phase is finished, it is possible to feed the deep learning system unlabelled images. It will then be able to identify these images, provided of course that enough images of this type were provided to it during the training phase.

It is important to remember that deep learning doesn't improvise and will not know how to react to new situations. It's for this reason that we refer to weak Artificial Intelligence (as opposed to what is known as strong AI which would theoretically be able to adapt to different contexts, but which doesn't exist at the current stage of technological development).

IT IS IMPORTANT TO REMEMBER THAT DEEP LEARNING DOESN'T IMPROVISE. IT DOES NOT REACT EFFECTIVELY IN NEW AND UNKNOWN SITUATIONS.

It should also be noted that the training process for deep learning requires huge levels of computing power. This power is needed to carry out operations as simple as addition and multiplication but at enormous quantities. Something to bear in mind is that a normal 4K image in itself has over eight million pixels*. Yet deep learning will sometimes need to process well over 30 per second, more than this in colour - tripling the number of pixels - while each pixel will in turn generate a mind-boggling number of operations.

Consequently, even the most powerful CPUs* (Central Processing Units) are proving increasingly insufficient to carry out these massive quantities of calculations. For this reason, they are often replaced by GPUs* (Graphics Processing Units), which are capable of doing huge amounts of calculations simultaneously. Google has even designed TPUs* (Tensor Processing Units), which are GPUs highly specialised in this task. This means the training phase in deep learning must more often than not be done on specialised big data infrastructures. It should be noted that the services offered by different cloud providers often include special services with computing power and GPUs suitable for training. Fortunately, once this preliminary phase has been carried out, the deployment and use of deep learning requires considerably less computing power, and can be carried out using more conventional and, most importantly, less energy-intensive architecture.

APPLICATIONS

There is such a vast range of applications of deep learning that an exhaustive list cannot be given here.

Here is a random selection, however:

- Natural Language Processing (NLP) or Text Mining.
- Correcting spelling and translating texts.
- Image recognition in industry to automate processes.
- Controlling autonomous or semi-autonomous vehicles (cars, boats, planes, helicopters, drones).
- Support with musical composition.
- Support with graphic design.
- Chatbots*.
- Voicebots* and personal assistants like Alexa, Siri, Cortana, Google Assistant, etc.

- Video surveillance for public places (train stations, airports) or industrial property.
- Monitoring the growth of bacterial cultures or viruses for the creation of vaccines.
- Medical imaging and diagnostic support.
- Board games.
- Robotics (the creation of autonomous robots).
- Optimising manufacturing processes.
- Analysing the content of meetings and creating summary documents.

Comparative Table of Machine Learning Techniques

TYPE OF MACHINE LEARNING TECHNIQUE

	Supervised	Unsupervised	Reinforcement	Deep - Supervised	Deep - Unsupervised	Deep - Reinforcement
Type of input data	Structured	Structured	Structured	Unstructured	Unstructured	Structured and unstructured
Continual learning	Unsupervised	Unsupervised	Yes	Unsupervised	Unsupervised	Yes
Transfer learning*	Unsupervised	Unsupervised	Unsupervised	Yes	Yes	Yes
Complexity	Low to medium	Medium	High	Very high	Extremely high	Extremely high
Time needed for learning	Seconds to hours	Minutes to days	Continual	Minutes to days	Minutes to days	Continual
Transparency and explainability	Very good	Good	Low	Very low	None	None
Risk of bias	High	Low to medium	Low to medium	High	Medium	Medium
Type of hardware	CPU	CPU	CPU	CPU, TPU or GPU	CPU, TPU or GPU	CPU, TPU or GPU

The AI of Tomorrow: What is the Current State of Research?

French AI start-ups are the European champions when it comes to investments. In 2019, they raised double the amount of funds they raised in 2018[12], that is 1,268 million dollars (or over 900 million euros). It is worth noting that three quarters of the investments were from France and only 16% from the rest of Europe.

> IN 2019, FRENCH AI START-UPS RAISED DOUBLE THE AMOUNT OF FUNDS THEY RAISED IN 2018[12], CORRESPONDING TO 1,268 MILLION DOLLARS.

The global AI market is booming; the amounts raised by specialised start-ups tripled between 2010 and 2016. In just six years, investments worldwide went from 600 million to 1.8 billion dollars[13] (around 1.3 billion euros). Unsurprisingly, it's US companies that are in the lead - especially the GAFAM companies (Google, Apple, Facebook, Amazon and Microsoft) - with a total of 62% of the investments in the sector, compared to 3% for France.

However, despite these disparities, all sectors, all countries and all companies should benefit from the advantages of Artificial Intelligence. But it's important to bear in mind that the risk of increasing these disparities is far greater for those that don't invest in AI soon enough... or don't go about it the right way.

The Expert's View

Françoise Soulié-Fogelman,
Scientific Advisor, Hub FranceIA

Small data is one of the most important topics today. For the best algorithms, a very large number of data must be collected, connected then labelled. The challenge is therefore finding algorithms that are much more economical with data. This is currently AI's main weakness. In real terms, how do I go about making sure that my AI system has the lowest error rate possible? Another important topic has to do with the need to audit my system and, in particular, have access to detailed monitoring of the data. But as it stands, we don't retain any documented record of this data. Systems must be highly auditable. It's a fundamental point that currently poses a problem. The future of research in AI no doubt lies partly in self-supervised learning. We will have to succeed in being able to use small data. We can't continue to settle for big data guzzling algorithms when we know that, for example, only Facebook and some Chinese companies can really do facial recognition because they have several billions of images available to them! Finally, there is an ethical rule that says that everyone should be able to take part. Yet today data is a barrier to enter AI for start-ups and small companies. One of the keys is small data. That's an area researchers should focus on.

3

Risks and Opportunities in AI

Risks and Opportunities in AI

AI applied in business contexts is an essential growth accelerator but is above all a necessary extension of human skills. Contrary to the definition given at the Dartmouth conference in 1956, AI is not meant to replace humans, but to supplement and support them!

In 2020, the value of the AI market was estimated at 17.3 billion dollars[14] (almost 12.5 billion euros). The estimate for 2025 is almost 90 billion dollars (around 65 billion euros). In five years, the value of the market will have more than quadrupled. So be careful not to miss the AI boat!

IN 5 YEARS, THE VALUE OF THE AI MARKET WILL HAVE MORE THAN QUADRUPLED[14].

"We have to find the right combination of human and Artificial Intelligence. AI will allow us to move towards augmented intelligence. The idea is for the final decision to always stay on the human side whenever they have accountability. We must also accept the idea that humans are not capable of analysing billions of pieces of information, and that the use of AI will add a great deal of value to human intelligence. So we should think of AI as an excellent assistant to humans."

Mick Levy,
Business Innovation Director, Business & Decision

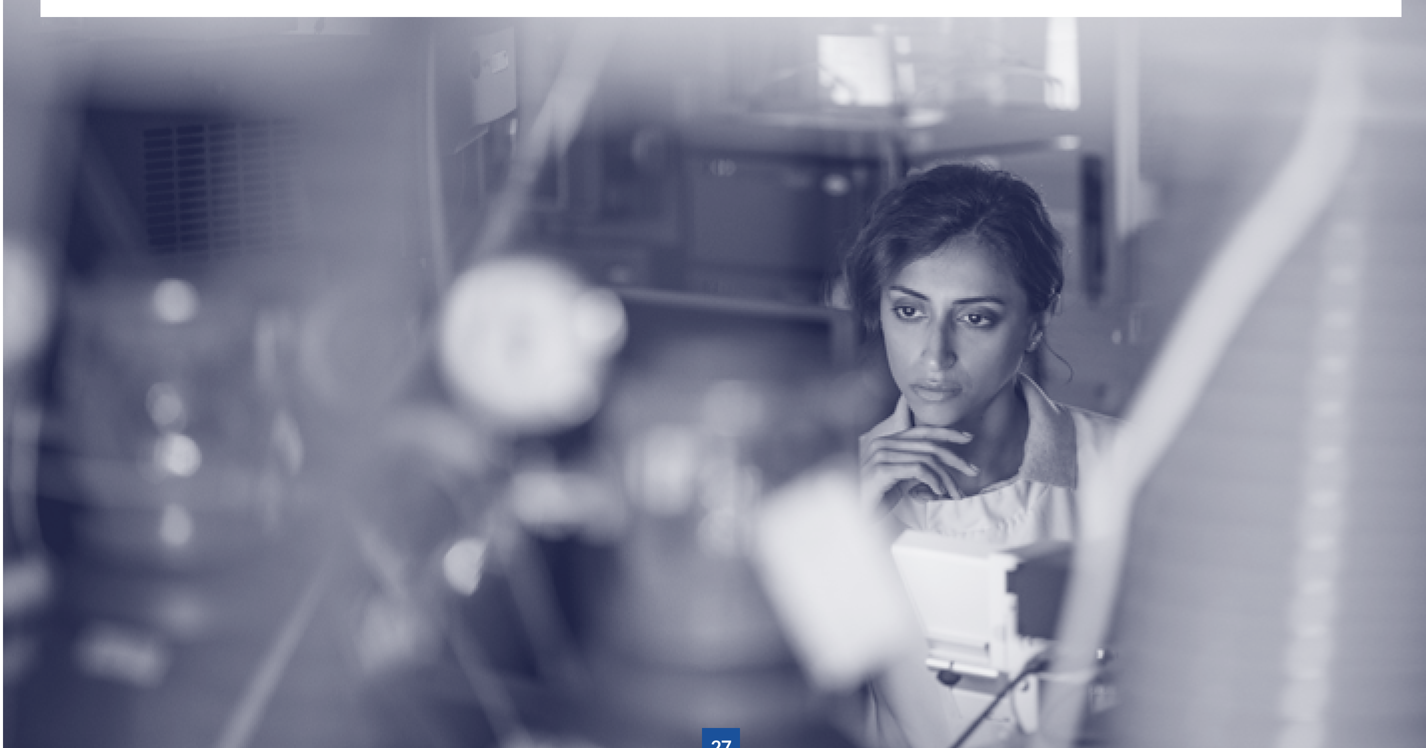

Why Businesses Should Go For It

The four main families of AI use cases

There are four ways to extract maximum value from the data available and transform it into substantial benefits.

- **The first** is based on improving operational efficiency and doing things better. This category includes all forms of logistical and operational optimisation, cost optimisation, production efficiency, cost savings, accelerating processing, etc.

- **The second** involves strengthening relationships with customers, personalising messages and offers to the point of making ties with customers more intimate and more relevant. To put it simply, doing more.

- **The third** lies in anticipating and managing risks. This category also includes the areas of regulatory compliance and fighting against fraud.

- **The fourth**, and last, has to do with the creation of new products and services that wouldn't be possible without AI. This combination of data and AI constitutes a significant foundation for innovation for businesses and local authorities.

At a time when data is the new black gold of the 21st century, harnessing its value through AI is becoming a priority for all business functions in all industries! Companies will now have to get AI-ready and not fall behind when it comes to the implementation of Artificial Intelligence.

Although AI is starting to be integrated occasionally into software packages for business use, whenever it is involved with what we can call the 'core business' of a company, the general trend is for companies to build their own bespoke systems. One reason for this is so that they can match it as closely as possible with the company's environment and data ecosystem. Another reason is to maintain control over their own data, this type of application requiring a certain centralisation of data and the use of business data lakes and data hubs*.
We are at the dawn of the business AI era!

WHILE DATA CREATED BY BUSINESSES ONLY ACCOUNTED FOR 30% OF THE TOTAL IN 2016, THIS PROPORTION IS SET TO DOUBLE BETWEEN NOW AND 2025[15].

Benefits of AI

The applications of Artificial Intelligence can bring a good return on investment. AI is first and foremost a fantastic decision-making tool which can expand on the BI (Business Intelligence) of today. For example, following the recommendations of an AI system enabled a retail chain to double the turnover of its largest outlet.

AI is also a fantastic tool for managing and monitoring activity. At a time when debit card fraud is exploding, AI can help to detect fraud and prevent it in the vast majority of instances. On a production line, as long as the data is available, AI can detect in advance if a part is going to turn out to be faulty. Spotting a few fraudulent transactions among billions or picking up on anomalies in hundreds of millions of products is an area where AI excels and proves itself to be an essential ally for humans.

Another example: prediction is an area where AI proves very useful. Today, an AI system can predict, for example, the exact time a container ship will arrive at a port, making it possible to optimise all the port's logistics and save tens of millions of euros in the process.

Finally, as technology becomes more mature, AI-based prescriptive tools are starting to appear. An AI system will be able to, for example, suggest that a part from the cylinder block of a crane is replaced before it stops working, avoiding the site being completely shut down for several days.

AI SHOULD BE ABLE TO INCREASE THE PROFITABILITY OF BUSINESSES BY MORE THAN A THIRD IN 2035[16].

Therefore, with almost seven in ten employed workers in France[17] saying they lack motivation for their work (a figure which jumps to 85% worldwide), AI seems an effective means of support, especially when it comes to taking over the most monotonous, perfunctory and repetitive tasks. That's where RPA or Robotic Process Automation comes into play, made possible by Artificial Intelligence. For companies, it's an important tool to avoid burn-out and bore-out... as well as the resulting loss of income, estimated at nearly 7 billion dollars.

The Expert's View

Ada Sekirin,
CEO Europe, Business & Decision

Companies that give deeper consideration to ethics will be more competitive. The general public is particularly mindful of this issue. Accordingly, ethics and protection of the data used for commercial purposes could become a competitive factor in the future. Today, businesses are prepared to use AI from a marketing point of view, but without replacing human intelligence with AI.

Developing a Trustworthy AI System

What are the risks of AI?

Of course, these benefits are all very real and make AI very attractive to companies. However, it is not without risk. *According to Françoise Soulié-Fogelman*, Scientific Advisor at Hub FranceIA, AI presents three main risks:

- **Algorithmic bias**

While there is relatively good awareness of racist and sexist types of bias[18], there remain a large number of unknown types, as human reasoning is naturally subject to biases. The result is that discrimination of all kinds is hard to shake. But is it up to AI to correct a vision of the world that comes from real data? It is really more a matter of detecting the discrimination created by this AI and measuring it: data scientists most likely can't correct the data, but they can measure the results that come from it. That's why they will, in the end, have to be responsible for the models they produce. This means that the current role of data scientist will start to be joined by specialist roles dedicated to very specific tasks: analysing bias, discriminations, ethics, etc.

The Expert's View

Didier Gaultier,
Data Science and AI Director, Business & Decision

When GDPR is an obstacle to evaluating and testing AI systems

Complex AI systems, as we've seen, have a transparency problem with their decisions; the more complex they are and the more they use big data, the less transparent and traceable they are. One of the solutions to this involves using statistical tests. For example, to test for gender-based discrimination in an AI process, we can input randomly-selected representative samples of men and women into the model* and carry out statistical analysis of the AI system's behaviour, comparing the output against these samples.

So far, so good. But how can we test for discrimination based on politics, religion, ethnicity or trade union membership? It's not possible because the GDPR prohibits collecting and storing files containing this type of information, known as sensitive data. In Europe today, considering the regulations in force and the complexities of AI, we are at an impasse when it comes to resolving this type of problem. It is reasonable to think that adapting the GDPR to take AI into account will be vital in the future.

- **Transparency**

Over 70% of French people would like more information on the positive impact of AI in health and education[19]. But in the event of negative impact, it is important to put controls in place and inform people about the risks.

The Expert's View

Michael Deheneffe,
Strategy & Innovation Director, Business & Decision

The current crisis is very revealing when it comes to expectations and priorities. It's shining a light on behaviours. This has greatly increased the demand for transparency, as much from organisations as from individuals. 77% of French people thought that communication from brands during the public health crisis about how their products could help them through the situation was very important. 50% even made it a factor in their future purchasing decisions[20]. In the future, some fraudulent practices, for example, will be unacceptable. AI will therefore be challenged with regard to its usefulness and its objectives. Accordingly, AI will have to be used to optimise what has value and minimise what does not. That is what companies should move towards: AI where no one loses!

- **Liability**

The team working on AI in a company should be multidisciplinary to involve as many different points of view as possible and thus manage the risk of any issues where the company would be liable. That's why AI-related projects are often jointly managed by different departments, with the main example being a collaboration between IT and other departments in 23% of cases. IT manages these projects on its own in only 19% of cases[21], and more rarely (11%), this role falls to data science departments. There has been some major progression here since in 2019, 33% of AI and machine learning* projects were managed by the IT department alone.

IN 2019,
33% OF AI AND MACHINE LEARNING* PROJECTS WERE MANAGED BY THE IT DEPARTMENT ALONE.

The Expert's View

Françoise Soulié-Fogelman,
Scientific Advisor, Hub FranceIA

The first question companies should ask themselves is about risks: does the product they are going to sell present risks for users? To answer it, a multidisciplinary approach is needed within the company, with people at different levels of the organisation: managers, risk managers, data scientists, IT specialists, etc. We need to break down the barriers between different departments! This plurality of opinions and knowledge makes the topic quite complex. We need to give training on trustworthy AI and establish what the responsibilities are faced with such a multi-faceted topic.
Once the team has been defined, the project goes through a design phase, then the production phase. At this point, it is vital to monitor the product throughout its life cycle with one constant consideration: ethics! In this respect, AI entails implementing a monitoring system and a tool capable of automating this continuous ethical evaluation for all products. This process is an integral part of the understanding between companies and consumers.
But today, it is evident that while the majority of French start-ups and large groups have started to consider the subject of ethical and responsible AI, many SMEs and mid-sized businesses are late to the party! Yet, they account for 80% of jobs. Before advancing with their Artificial Intelligence projects, they need to increase their digital maturity!

Assessing Your Maturity in Terms of AI

One of the factors to take into account in the implementation of AI projects is the company's maturity. A company or organisation's maturity in terms of AI is the combined result of a multitude of different factors, with some of the main ones being culture, structure, human and technological resources, data governance, infrastructure, projects, etc.

Business & Decision's experience in the area generally confirms that maturity in terms of data in general and in terms of AI are closely linked. Considering that data is the fuel of AI (see our last white paper, *Data Ethics, AI Ethics – The 2 faces of a responsible future*) and that bad quality of data reduces the chances of success to zero in the majority of AI projects, this result is hardly surprising.

Data quality is certainly a necessary condition for successfully carrying out AI projects, but it isn't enough on its own. There are certain conditions that must be met as a minimum:

A very good understanding of industry challenges.

Skilled data scientists and data engineers.

Access to quality data (internal and external) that is related to these challenges.

Sufficiently high-performing infrastructure (cloud or on-premises) to meet the required levels of storage, memory and computing power, especially during the design phase of the AI project (the training phase consumes large amounts of power and data).

Data governance in place and a company culture favourable to AI.

Business & Decision has provided an evaluation table so that you can self-evaluate your company. You just need to try to identify which level of data and AI maturity your company is at in the table. As we previously highlighted in our last white paper, *Data Éthique / IA Éthique : les deux visages d'un futur responsable*, there is no wrong answer. The important thing when making this journey is knowing where you're starting from, so that you can allocate the right resources and put in the right amount of effort.

DATA MATURITY MODEL

	People	Processes	Technology	Data & AI	
Integrated	• Culture of data quality • Data governance • Key Data Stewards	• Optimised data exploitation process • Data quality is a key factor in all core business and IT department processes	• Collaborative platform for management of semantic layers • Operational MDM tools	• A large number of viable AI projects. in the production phase. The majority of AI projects are successfully completed to deadlines	• Data is treated as a key asset • Data creates large amounts of value through AI • Controlled expansion of data assets
Proactive	• Proactive behaviour with regards to data quality • Key positioning of CDOs • Recognised Data Stewards	• Data quality incorporated in core business processes • Centralised metadata • Data quality strategy	• Tool dedicated to data governance Centralised metadata management platform	• Several AI projects in the production phase, some significantly delayed	• Information generally coherent and shared • Shared view (IT and other departments) of data assets
Controlled	• Data Stewards in place • Introduction of a CDO • Responsibilities of Business Process Owners made clear	• Documented data dictionary and business rules • Shared management of metadata	• Data quality platform • Analytical MDM & metadata sharing • Business glossaries	• Some AI projects in the production phase, most significantly delayed, some projects abandoned before the end	• Measurement of quality but with disparities between different domains • Value of data quality shared • Information managed by non-IT departments
Managed	• Data quality perceived as necessary • Appointment of Data Stewards	• Existence of a partial business glossary • Data quality audit • Methodologies occasionally put in place	• Data quality tool • Data cleansing • Partial metadata repository	• Some AI projects but very few in the production phase and the vast majority are halted partway through	• Ability to evaluate the level and impact of data quality • Data quality is recognised as an important business consideration
Unpredictable	• Lack of awareness about data quality • Responsibilities not clearly defined	• No common language • Project by project • No metadata management	• No tools.	• No or very few AI projects, and all AI projects are halted partway through	• Level of quality not known

Integrating and Working With AI in the Business

By transforming their organisation to better prepare their employees for human-machine collaboration, businesses could increase their turnover by 38%[22]. This makes it a vital topic for organisations.

In the industry of the future, humans and AI will need to combine their strengths. On the AI side, automation and the analysis of large volumes of data, and on the human side, creativity, reactivity, empathy, emotion and ethics. Businesses could introduce a 'grid' to determine when to use AI and when to use humans.

Leadership	Creativity	Empathy	Ethics	Learning	Inference	Maintenance	Interaction	Communication	Avatars	Transactions	Iterations	Applications	Calculations
Activities reserved for humans				Activities where humans and machines collaborate						Activities allotted to machines			

Source: *Human + Machine: Reimagining Work in the Age of AI*, by Paul R. Daugherty and H. James Wilson

Setting up ethical and responsible AI

What is ethical AI?

Maxence Dhellemmes, R&D Director, Business & Decision

"Ethical AI must be kind and fair for the user while also respecting humans and the planet. Ethical AI is naturally an ideal. It's a target we have to set ourselves, putting in place safeguards along the way to ensure we're heading in the right direction. Ethical AI is an innovation; it adds value for the end user. To sum up, ethical AI lies in benevolent AI that is in tune with the needs of users. So for companies, this kind of positioning means developing the business while staying in line with 'human' expectations."

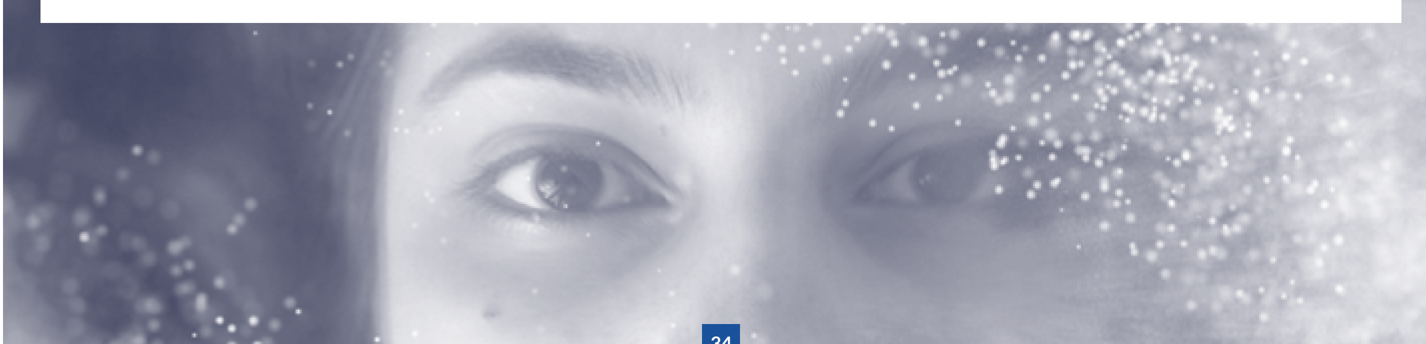

Ethical AI for a business is built around six main actions:

Appointing a **Chief Ethics Officer** alongside the Data Protection Officer (DPO). A partnership that will guarantee the integrity and quality of data.

Alongside monitoring, there must be a code of conduct or an internal **charter** that all data specialists are obliged to be familiar with and to follow.

Mirroring the concept of Privacy by Design, companies must move towards an **Ethics by design** model to incorporate ethics into the design of AI systems.

It is essential that **humans stay number one**. This means that there must be a process in place to allow humans to take back control from AI systems at any point in time.

Carrying out **regular audits of algorithms**.

Where AI is concerned, the business must be 'human first' as well as 'human centric.'

The Expert's View

Ada Sekirin,
CEO Europe, Business & Decision

The main challenge lies in making the creators of algorithmic models aware of their responsibilities. While today AI is used primarily in areas like marketing, e-commerce, customer relations and the media, tomorrow, the next wave will apply to production processes, Industry 4.0, etc. - particularly difficult areas. That is why it is necessary to adopt a fundamentally scientific approach to give a framework. The current lack of a generally accepted global or generic set of ethics makes programming algorithms much more complex as it requires taking a certain position. A 'sticking point' that could stand in the way of the development of AI in business, the use of which remains very limited today. The debate around AI must therefore be based on concrete use cases in order to 'unstick' the situation and move forward.

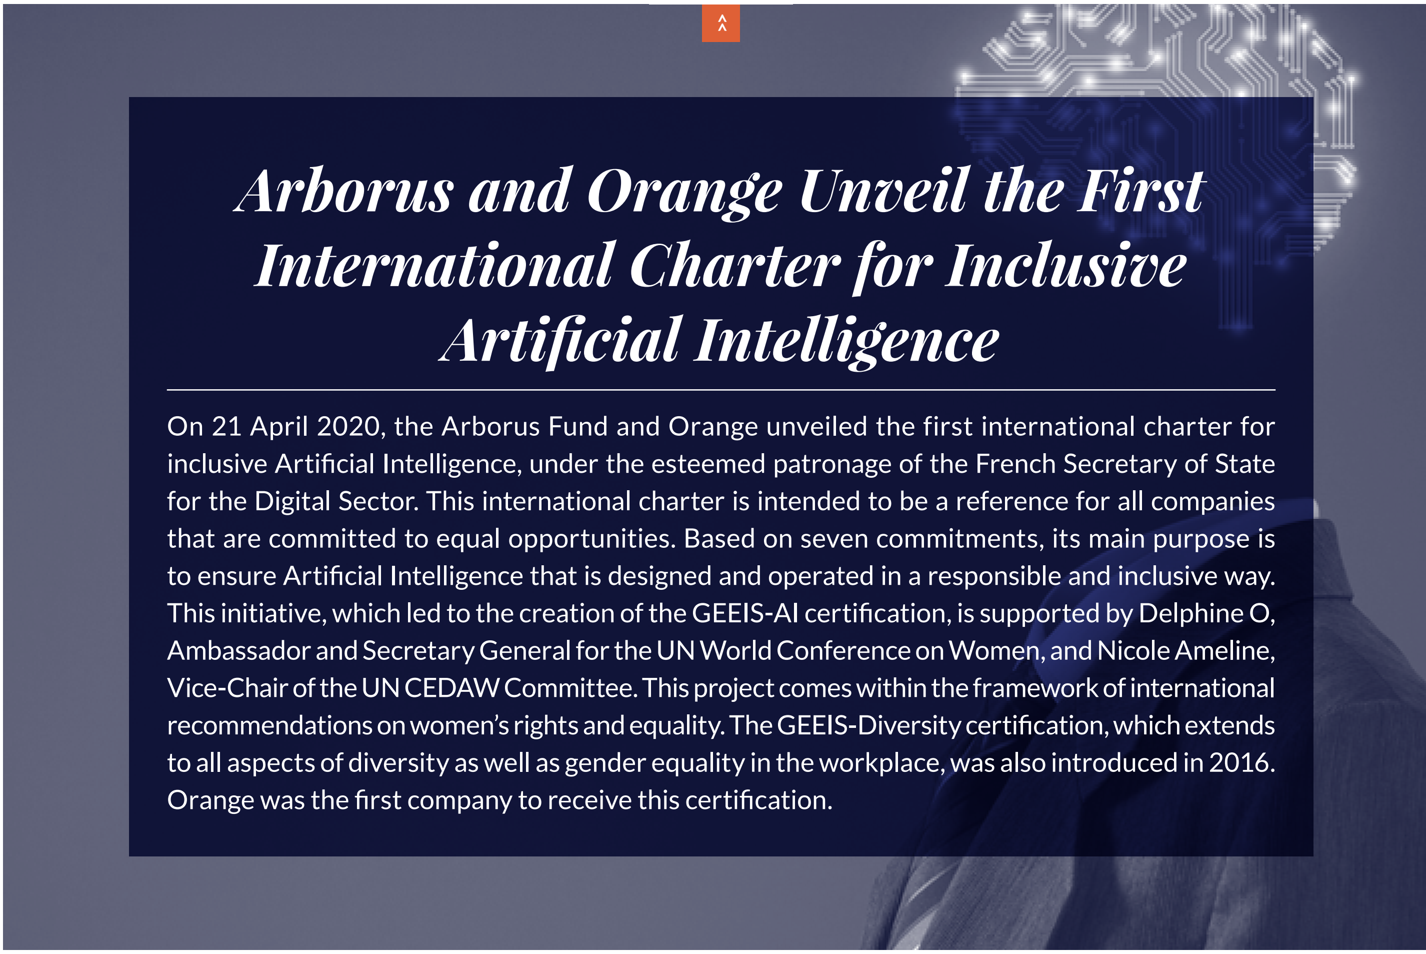

Arborus and Orange Unveil the First International Charter for Inclusive Artificial Intelligence

On 21 April 2020, the Arborus Fund and Orange unveiled the first international charter for inclusive Artificial Intelligence, under the esteemed patronage of the French Secretary of State for the Digital Sector. This international charter is intended to be a reference for all companies that are committed to equal opportunities. Based on seven commitments, its main purpose is to ensure Artificial Intelligence that is designed and operated in a responsible and inclusive way. This initiative, which led to the creation of the GEEIS-AI certification, is supported by Delphine O, Ambassador and Secretary General for the UN World Conference on Women, and Nicole Ameline, Vice-Chair of the UN CEDAW Committee. This project comes within the framework of international recommendations on women's rights and equality. The GEEIS-Diversity certification, which extends to all aspects of diversity as well as gender equality in the workplace, was also introduced in 2016. Orange was the first company to receive this certification.

Can we (and should we) regulate AI?

On 19 February 2020, the European Commission unveiled its underlying principles and actions for a digital transformation that works for all, reflecting the best of Europe: open, fair, diverse, democratic and confident. The objectives are threefold: putting people first, opening up new opportunities for businesses and boosting the development of trustworthy technology. In this way, Europe hopes to foster an open and democratic society through a vibrant and sustainable economy.

A DIGITAL TRANSFORMATION THAT WORKS FOR ALL.

To shape Europe's digital future, the European data strategy is based on the human-centric development of Artificial Intelligence (AI). For Thierry Breton, Commissioner for Internal Market: "Our society is generating a huge wave of industrial and public data, which will transform the way we produce, consume and live. I want European businesses and our many SMEs to access this data and create value for Europeans – including by developing Artificial Intelligence applications. Europe has everything it takes to lead the 'big data' race, and preserve its technological sovereignty, industrial leadership and economic competitiveness to the benefit of European consumers."

For more information about the European strategy for ethical and responsible AI:
https://ec.europa.eu/commission/presscorner/detail/en/ip_20_273

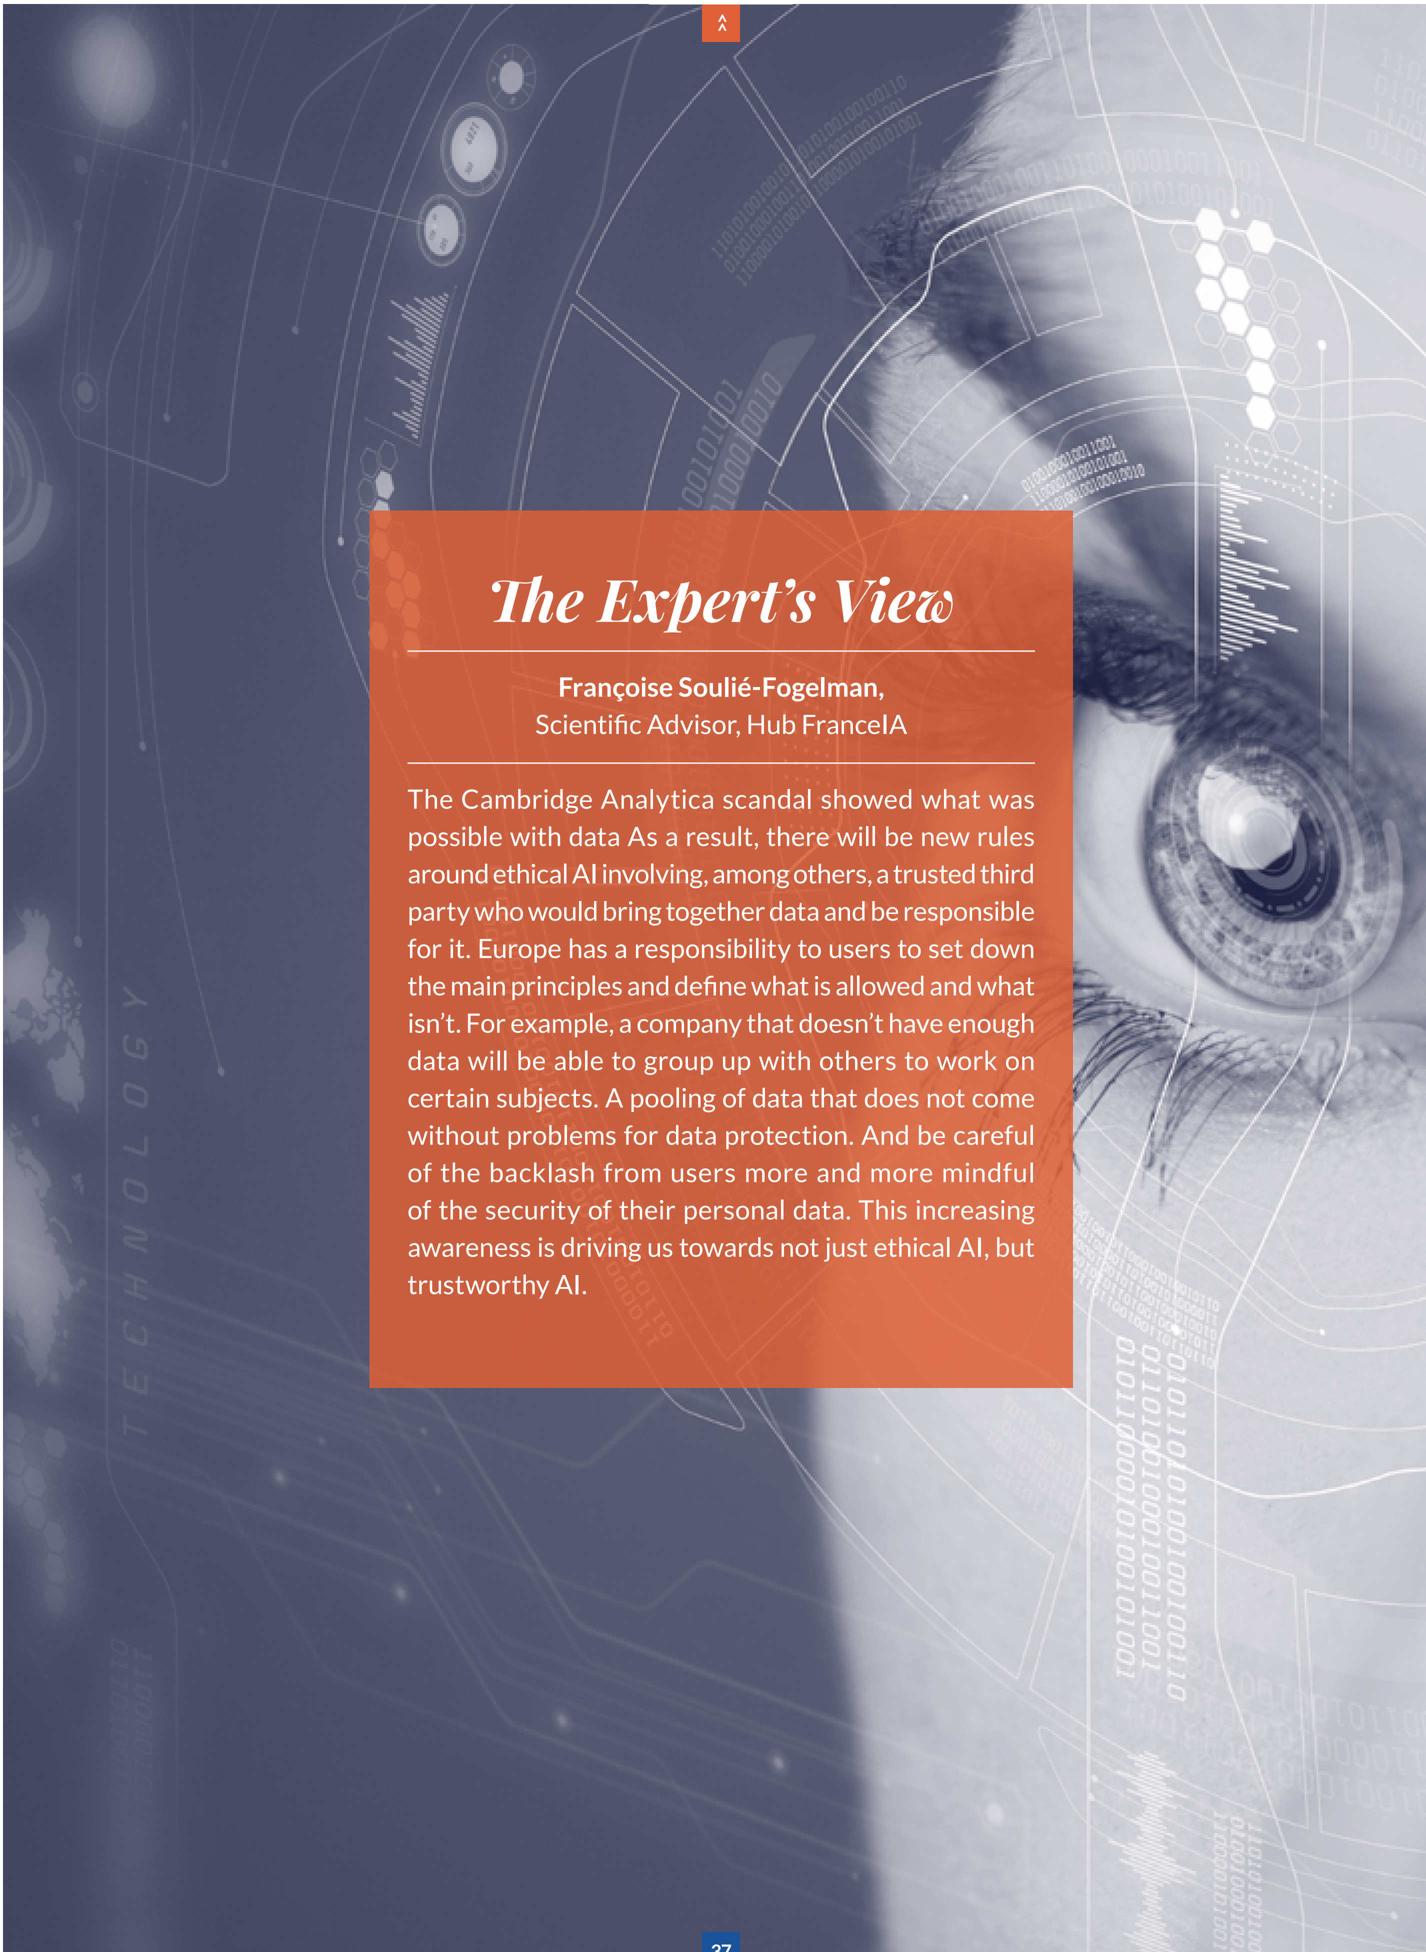

The Expert's View

Françoise Soulié-Fogelman,
Scientific Advisor, Hub FranceIA

The Cambridge Analytica scandal showed what was possible with data As a result, there will be new rules around ethical AI involving, among others, a trusted third party who would bring together data and be responsible for it. Europe has a responsibility to users to set down the main principles and define what is allowed and what isn't. For example, a company that doesn't have enough data will be able to group up with others to work on certain subjects. A pooling of data that does not come without problems for data protection. And be careful of the backlash from users more and more mindful of the security of their personal data. This increasing awareness is driving us towards not just ethical AI, but trustworthy AI.

Conclusion

Between now and 2024, the global Artificial Intelligence market is set to experience average annual growth of more than 20%[23]. Examples of the most popular use cases among companies and organisations today include: decision support for customer service, recommender systems, optimising logistics, diagnostic support in healthcare, smart cities, cybersecurity, IT automation and the management of sustainable energy.

Why this exponential growth? Ritu Jyoti, Program Vice President for Artificial Intelligence Research at IDC, emphasises that it is because AI enables companies to *"stay competitive in the digital economy. They will adopt AI - not just because they can, but because they have to."* There are two objectives: improving the client experience and optimising staff efficiency to move towards an 'augmented employee.'

France has understood the importance of Artificial Intelligence: the government made AI a long-term priority in the recovery plan it released in September 2020. It is clear that Artificial Intelligence is more than just a marketing 'gadget'; this is technology that is increasingly high-performing, polished and vital to company strategy. **That said, even with 34% of medium-sized and large French companies using AI[24], how many of them are really mature?** On the one hand will be the companies that implemented AI projects in an ethical and responsible way, built on solid data governance. On the other hand will be those that didn't invest enough and developed AI systems in a rush with badly-defined data governance, over time exposing themselves to a potentially severe backlash from customers, employees and shareholders.

Adopting AI is above all a value-generating and benevolent strategy for companies in the medium to long term. It involves changing working methods and cultures. This impetus requires close involvement of general management and also needs to be supported. Whichever way you look at it, it is certainly not simply a matter of technology.

The Last Word

Luc Julia, CTO of Samsung Electronics and co-creator of Siri

*"It is difficult to create a regulation on AI since it covers such a vast area.
That's why we need to educate to improve understanding of it.
But the millions of calculations carried out by Artificial Intelligence make this a difficult task.
Europe was a pioneer in terms of regulations around respecting personal data.
The California Consumer Privacy Act (CCPA) follows the lead of its GDPR.
Individuals must gain an understanding of the impact and, once informed, make the right decision for them. History has shown us that industrial revolutions always bring employment problems at one point. In 1790, the hand weavers of Lyon's silk industry were replaced by power looms.
In 1820, Lyon had the most dynamic job market in France because the machines had created a thriving economy which generated new jobs.
The best support is education. AI is going to change our jobs but it's not going to replace us!"*

Data Science and AI Glossary

Agent: in Artificial Intelligence, an intelligent agent (IA) is an entity (physical and/or software) with some degree of autonomy, capable of perceiving its environment (through sensors) and also acting upon it (via actuators) in order to achieve goals. (Wikipedia)

Deep learning: a set of machine learning methods that attempt to model high-level abstractions in data by using architectures composed of non-linear transformations [editor's note: such as artificial neural networks]. These techniques have enabled significant and fast progress in the fields of audio and visual signal analysis, especially computer vision, vocal recognition and automated language processing. In the 2000s, this progress generated considerable private, university and public investment, especially from the GAFAM companies (Google, Apple, Facebook, Amazon, Microsoft). (Wikipedia)

Supervised learning: machine learning task consisting of learning a predictive function based on annotated examples. (Wikipedia)

Chatbot: software agent that converses with a user. The user is invited to formulate their request in natural language. It is refined by an exchange based primarily on questions and answers pre-programmed in a knowledge base. If the conversation is oral and not written, this is known as a voicebot.

CPU (Central Processing Unit): component in many electronic devices including computers that executes instructions comprising a computer program. Along with memory, it is one of the functions that has existed since the earliest computers. A processor on a single integrated circuit is a microprocessor. (Taken from Wikipedia)

Data engineering: discipline focused on building and maintaining the organisation's data pipeline systems (O'Reilly). This discipline encompasses types of technology characteristic of big data, DataOps and MLOps.

Data lake: data storage method used for big data. The data is kept in its natural format, with little done to it. Data lakes favour quick and large-scale storage of heterogeneous data, adopting a cluster architecture. They are not optimised for SQL queries or traditional relational DBMSs, unlike Data Hubs, in which the data is transformed, improved in quality, checked, sorted and optimised for exploitation by data scientists and non-specialists.

Data science: an interdisciplinary field focused on extracting knowledge from data sets (Wikipedia). It is made up of a number of disciplines, namely: data engineering, industry knowledge, statistics, machine learning and Artificial Intelligence (Business & Decision 2020).

DataOps: process-oriented methodology, used by data engineering teams, to improve the quality and reduce the cycle time of data analytics. While DataOps began as a set of best practices, it has now matured to become a new approach to data engineering.

DevOps: set of practices that combines software development (Dev) and IT operations (Ops). It is largely characterised by the promotion of automation and monitoring of all stages of software creation, from development, integration, testing and delivery to deployment, operation and infrastructure maintenance. The DevOps principles advocate short systems development lifecycles, an increase in the frequency of deployments and continuous delivery to better achieve the company's financial objectives. (Wikipedia)

GPU (Graphics Processing Unit): processing unit, found in the form of an integrated circuit (or chip) on a video card or motherboard or directly linked to the same integrated circuit as the microprocessor, that carries out the image and video processing functions. GPUs generally have a highly parallel structure making them effective for a wide range of graphics tasks such as 3D rendering, video memory management and video signal processing, but also for matrix operations and, importantly, training artificial neural networks. (Taken from Wikipedia)

AI or Artificial Intelligence: discipline devoted to the building of autonomous agents that can analyse and make decisions (Peter Norvig 2010). Also known as augmented intelligence, it is a branch of data science devoted to the building of agents that interact with the environment and can learn from initial data provided and/or data collected during interactions to carry out a function in a specific environment, with some degree of autonomy (Business & Decision 2020).

AIOps – also known as MLOps or ProcessAI: methodology used by data science teams to effectively manage, deploy and maintain AI projects within businesses. This methodology, consisting primarily of a set of good practices, was born out of the intersection between the widespread CRISP method, long used in data science, and DevOps practices.

Machine learning: an area of study within data science that uses mathematical and statistical approaches to give computers the ability to learn from data using algorithms. (Wikipedia)

Model: in data science or statistics, this is an approximative mathematical description of the mechanism which generated the observations and which responds to certain hypotheses. It is generally expressed as a function (or a family of functions) of input variables X1... Xn. Each member of the family is a possible approximation of reality. Thus, the inference lies in determining if one of the members is sufficiently compatible with all the data.

NLP (or NLU, ***Natural Language Processing*** or ***Understanding***, also known as ***Text Mining***): multi-disciplinary field involving linguistics, computer science and Artificial Intelligence, which aims to create natural language processing tools for various applications, for example chatbots. (Wikipedia)

Pixel: smallest basic element of a raster image or video. The name comes from the term 'picture element.' The number of horizontal and vertical pixels in an image gives us what we call the resolution. For example, a 4K TV image in Ultra High Definition (UHD) has 3,840 pixels horizontally and 2,160 vertically, that is four times the Full HD definition of 1,920 x 1,080.

CNN (Convolutional Neural Networks): in deep learning, a convolutional neural network is a type of artificial neural network in which the connectivity pattern between neurons is inspired by the animal visual cortex.

General Data Protection Regulation (GDPR): European Union regulation which serves as the common, definitive legislation on the protection of personal data. It strengthens and unifies data protection for individuals within the European Union. (Wikipedia)

RPA (***Robotic Process Automation***): technology whereby agents are created that learn from the behaviour of a user, especially on a graphic interface. In the classic process automation approach, an IT developer writes a computer programme that carries out a certain number of tasks and interacts with the application programming interface (API). In robotic process automation, the system learns the list of tasks to be automated by observing the behaviour of human users. Robotic process automation is used in businesses to automate repetitive tasks.

Statistics: the discipline studying phenomena through collecting, processing and analysing data, interpreting the results and presenting them in a way that makes the data comprehensible to all. (Wikipedia)

Expert system: a computer system emulating the decision-making ability of a human expert in a particular domain. It is one of the ways of working towards Artificial Intelligence. More specifically, an expert system is a piece of software capable of responding to questions by reasoning using known facts and rules. A major use of expert systems is as decision support tools (Wikipedia). TPU (***Tensor Processing Unit***): application-specific integrated circuit specialised in matrix operations, belonging to the family of GPUs, developed by Google specifically to accelerate Artificial Intelligence systems using neural networks.

Transfer Learning: technique which enables an AI system to be pre-trained in one environment and then used in a related but slightly different one.

The training does still need to be finished in the exact environment in which the AI system will need to operate. But this phase is much shorter and uses much less energy than the initial training phase.

Boolean variable: variable that can only have two possible values, true or false, one or zero. It is named after the creator of an algebraic system of logic: George Boole, who was behind the basic principles that make all modern computers and digital electronics work. A Boolean variable is said to be a structured variable.

Categorical variable: variable that can take one of a predefined set of values, for example a car brand. The term 'qualitative variable' is also used. A categorical variable is said to be a structured variable.

Unstructured variable: variable that is not numerical, categorical, Boolean or a date. It can be a text variable (string), containing an audio or video signal, or a machine log.

Numerical variable: variable which takes the form of a number, either whole (this is known as a discrete numerical variable) or not (this is generally known as continuous numerical variable). The term 'quantitative variable' is also used. A numerical variable is said to be a structured variable.

Structured variable: numerical, categorical, Boolean or date/time variable.

References

1 Annual emerging jobs report for France, LinkedIn, December 2019 – https://business.linkedin.com/fr-fr/talent-solutions/emerging-jobs-report/emerging-jobs-report-france

2 According to IDC, August 2020 – https://www.idc.com/getdoc.jsp?containerId=prUS46794720

3 According to a France Digitale study - Roland Berger, Maddyness, November 2019 – https://www.maddyness.com/2019/11/04/france-leader-investissements-ia/

4 Larousse encyclopaedia – https://www.larousse.fr/encyclopedie/divers/intelligence_artificielle/187257

5 Warren McCulloch: https://fr.wikipedia.org/wiki/Warren_McCulloch

6 Walter Pitts: https://fr.wikipedia.org/wiki/Walter_Pitts

7 Franck Rosenblatt: https://fr.wikipedia.org/wiki/Frank_Rosenblatt

8 John McCarthy: https://fr.wikipedia.org/wiki/John_McCarthy

9 Marvin Minsky: https://fr.wikipedia.org/wiki/Marvin_Minsky

10 See *Artificial Intelligence: A Modern Approach* (3rd Edition) by Stuart Russel & Peter Norvig, Pearson

11 For more on this subject, see our first white paper (in French): *Data Éthique, IA Éthique, les deux visages d'un futur responsable* – https://fr.blog.businessdecision.com/livre-blanc-data-ethique-ia-ethique-futur-responsable/

12 According to a France Digitale study - Roland Berger, Maddyness, November 2019 – https://www.maddyness.com/2019/11/04/france-leader-investissements-ia/

13 According to the *Future of Artificial Intelligence* report produced by Hub Institute in partnership with IBM, 2018 – https://hubinstitute.com/Report/hubreport-future-ai

14 According to a prediction of artificial intelligence market revenue worldwide between 2016 and 2025, in millions of dollars - https://fr.statista.com/statistiques/638963/revenus-marche-intelligence-artificielle-monde/

15 According to an IDC study, 2018 – https://www.decideo.fr/Big-Data-les-donnees-sont-partout-mais-insuffisamment-exploitees-_a10425.html

16 As per an Accenture study, 2018 – https://www.ladn.eu/nouveaux-usages/etude-marketing/intelligence-artificielle-tout-le-monde-va-en-profiter-surtout-les-americains/

17 https://www.journaldunet.com/solutions/reseau-social-d-entreprise/1486875-en-2020-quelle-sera-la-place-de-l-intelligence-artificielle-dans-le-secteur-du-e-commerce/

18 For more on this subject, see our first white paper (in French): *Data Éthique, IA Éthique, les deux visages d'un futur responsable* – https://fr.blog.businessdecision.com/livre-blanc-data-ethique-ia-ethique-futur-responsable/

19 Observatoire de la Notoriété et de l'Image de l'Intelligence Artificielle en France (Observatory on the Reputation and Image of Artificial Intelligence in France), Impact AI - Ifop, 2019 - https://www.ifop.com/wp-content/uploads/2019/12/CP-Impact-AI-Barom%C3%A8tre-2019-VDEF.pdf

20 According to the Edelman Trust Barometer, April 2020 - https://www.elanedelman.com/research/trust-and-the-brand

21 ZDnet, April 2020 - https://www.zdnet.fr/actualites/les-projets-autour-de-l-intelligence-artificielle-sont-plus-matures-en-2020-39901613.htm

22 According to an Accenture study, 2018 – https://www.accenture.com/fr-fr/company-news-release-artificial-intelligence-workforce

23 According to IDC, August 2020 – https://www.idc.com/getdoc.jsp?containerId=prUS46794720

24 According to a study by JDN, Club Décision DSI and IT Research, September 2020 – https://www.itresearch.fr/un-tiers-des-moyennes-et-grandes-entreprises-francaises-utilisent-lia/

Acknowledgements

This Business & Decision white paper was edited by Didier Gaultier, Data Science & AI Director, with:

Valérie Lafdal
CEO for France and Deputy CEO for the Group

Ada Sekirin
CEO Europe

Michael Deheneffe
Strategy & Innovation Director

Mick Levy
Business Innovation Director

Our special thanks for the invaluable insights and contributions of:

Faycal Boujemaa
Technology Strategist - Orange Labs Research

Philippe Ensarguet
CTO - Orange Business Services

Jérôme Goulard
Head of CSR, Ethics & Diversity – Orange Business Services

Pascal Hauret
Performance Manager - Orange Business Services

Luc Julia
CTO & Senior Vice President of Innovation – Samsung

Françoise Soulié-Fogelman
Scientific Advisor - Hub FranceIA

With help from:

Romain Bernard
Data Science Manager for Northern France & Ile-de-France

Jérôme Dewever
MDM Consulting & Expertise Manager

Maxence Dhellemmes
R&D Director

Christophe Hatterer
Data Governance Manager

Cédric Missoffe
Director of the Consulting & Expertise agency

Mondher Sendi
AI Expert - Computer Scientist

Stéphane Walter
Big Data Consulting & Expertise Manager

Design:
Bespoke agency and Violaine Cherrier

Business & Decision - RCS B 384 518 114 - Photo credit: iStock

Business & Decision
Cœur Défense A, 110 Esplanade Général de Gaulle, 92931 Paris La Défense Cedex
www.businessdecision.com
blog.businessdecision.com

Business & Decision