Towards operational excellence through **orchestrating machines** and **humans with AI**

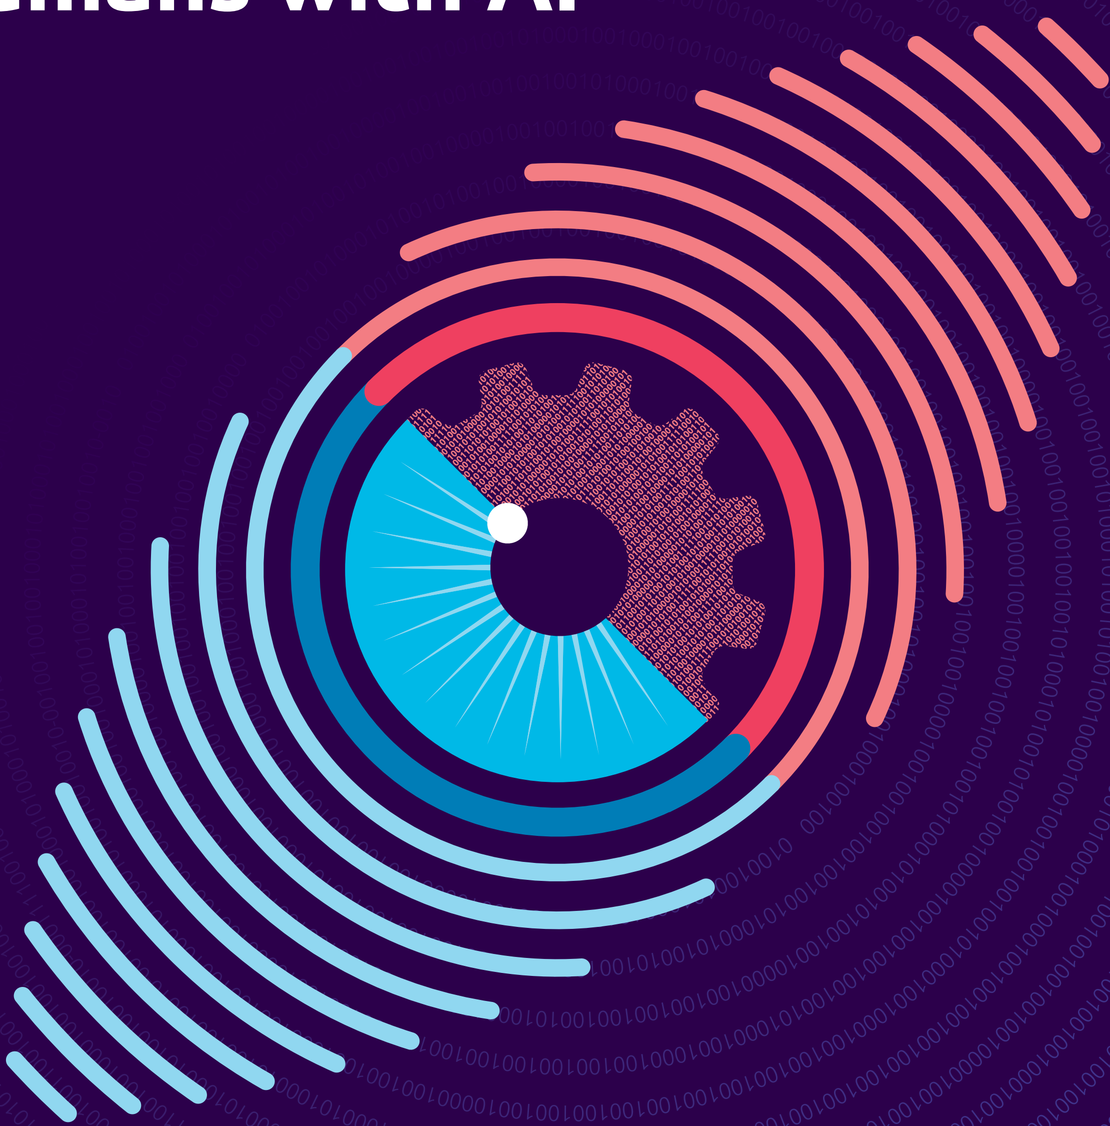

A point of view by Taoufik Amri,
Chief Data Scientist,
Capgemini's Business Services

Table of contents

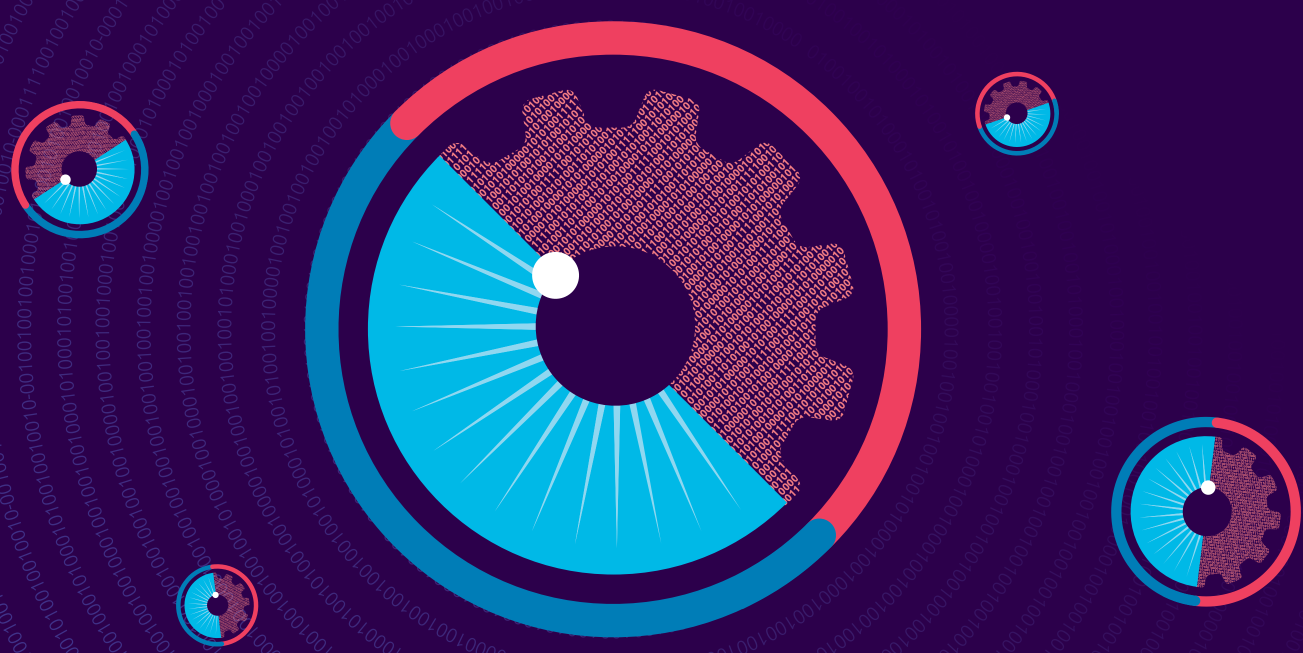

Introduction

This paper provides an overview of what intelligent automation means for business process management in the data and digital era.

Intelligent automation is based on artificial intelligence (AI), while keeping humans in the loop. Indeed, the re-engineering of processes and their optimization don't necessarily imply the replacement of human operators by machines.

Instead, we argue that intelligent automation will increasingly involve the smart orchestration of tasks between machines and humans in order to reach the operational excellence expected by our clients. In turn, this helps them to implement – what we call – the Frictionless Enterprise.

Such an orchestration will be itself governed by a less-known use of AI acting like an invisible hand. We illustrate this idea with a real-world case from a French insurance company.

> ***"Intelligent automation will increasingly involve the smart orchestration of tasks between machines and humans in order to reach the operational excellence expected by our clients."***
>
> **Taoufik Amri**
>
> *Chief Data Scientist,*
> *Capgemini's Business Services*

What is really possible with AI?

AI is mainly based on machine learning algorithms that learn from data. With this in mind, let's start with a crash course on the underlying approach is known as data science.

Crash course #1 – data science

It is a truth universally acknowledged that company performance depends on several factors, and that many of those factors are variable. If life is prone to inconsistency, so is business. Much of this is because of the unpredictability of human behavior, which is why it is interesting to explore alternative approaches to grasp them.

Data science provides a data-based approach for managing business performance and accommodating these twists and turns. To build a model, data scientists don't need to create ingenious bespoke algorithms, as is commonly believed. Instead, they need to follow a pragmatic process.

The first step is to explore data in order to analyze their fluctuations and correlations and to engineer features that exhibit them. To do this, data scientists need to select the factors that will play an important role.

Building the model is like assembling gears to create a mechanism that can be applied to data. The only systematic and consistent approach is the scientific method – in other words, an inductive and iterative process. We make assumptions from the data about the mechanisms – explaining the fluctuations and correlations we observed – and then we identify the models that could reproduce these observations.

We then check the assumption by testing on new, real data, and if the hypothesis is wrong, we have to follow this process again and again until we arrive at a good model. This process reveals a kind of "chicken and the egg" dilemma (see Figure 1) between data and model – data is needed to determine the model, and the model is necessary to leverage the data and reveal its value.

There are actually two extreme cases. The first one corresponds to the situation in which data quality is so poor that we cannot determine the right model to explain it. The other extreme case corresponds to the ideal situation, in which we have perfect data, but the fluctuations and correlations are so subtle and parsimonious that one cannot not determine which combination of models can explain them. The success of the approach depends on the talent of data scientists and their ability to be genuinely inductive, relevant to, and creative with the data.

So, who are these people, and what attributes do they have? Given their common job descriptions, we might think they are an amalgam of technologists and scientists; but in fact, it's more likely that they will primarily be scientists with business knowledge. The science comes first – these are people who conceptualize business needs with quantitative models to see what data can be leveraged.

Many such data scientists have typically studied physics, mathematics, or engineering sciences; overlaid on this is a sound knowledge of business, enabling them to leverage data and interpret the model that underlies it. Technology's supporting role is to facilitate the outcome.

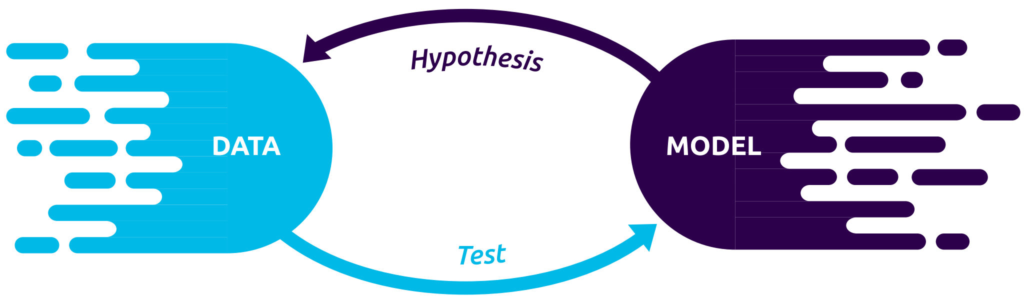

Figure 1. The data/model dilemma

"Data is needed to determine the model, and the model is necessary to leverage the data and reveal its value."

Taoufik Amri

Chief Data Scientist, Capgemini's Business Services

Crash course #2 – machine learning

An important aspect of this technology is machine learning. This is traditionally defined as a form of AI that enables computer systems to learn without being explicitly programmed.

There are three main types of learning process:

- **Unsupervised learning –** in which the machine learning algorithm learns from data that is not labeled by humans. For instance, clustering algorithms may summarize data into a small number of clusters in which data is grouped according to a common measure of its similarity
- **Supervised learning –** in which the machine learning algorithm learns from data that are labeled by humans. For example, algorithms may predict a state (no or yes, i.e., 0 or 1) or a quantity associated with a combination of data variables
- **Reinforcement learning –** in which the algorithm learns from data in order to maximize a reward. An example might be an algorithm that learns to play chess by winning or losing.

Machine learning is also called statistical learning. Contrary to natural intelligence, it needs a huge amount of data from which to learn. Although a child learns to identify cats and dogs with only a few examples, “deep learning” algorithms need many, many more.

In a business context, data scientists gather data that are representative of business operations. This data sample should be large enough to be statistically significant, enabling it to be split into three data sets (see Figure 2).

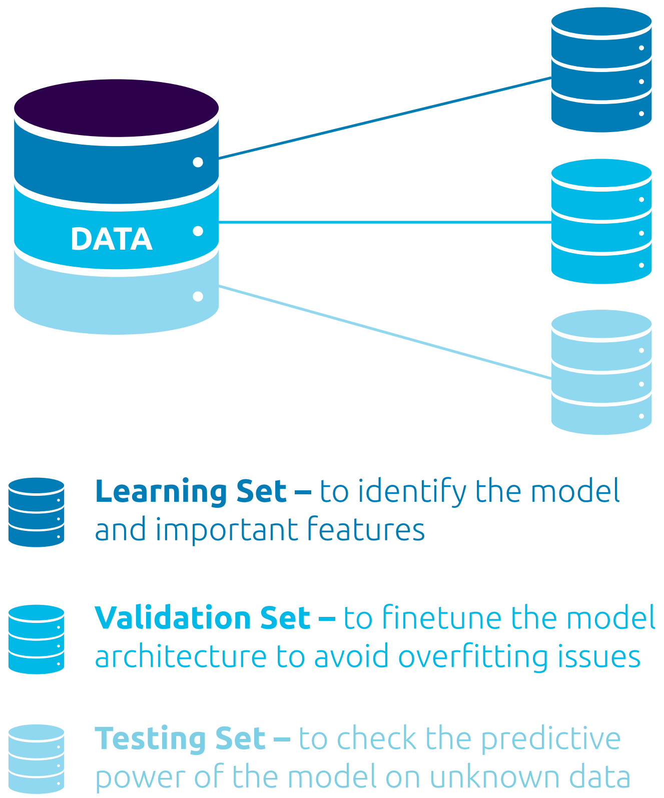

Figure 2. The three data sets needed for machine learning

> ***“Contrary to natural intelligence, [machine learning] needs a huge amount of data from which to learn.”***

Taoufik Amri

Chief Data Scientist,
Capgemini’s Business Services

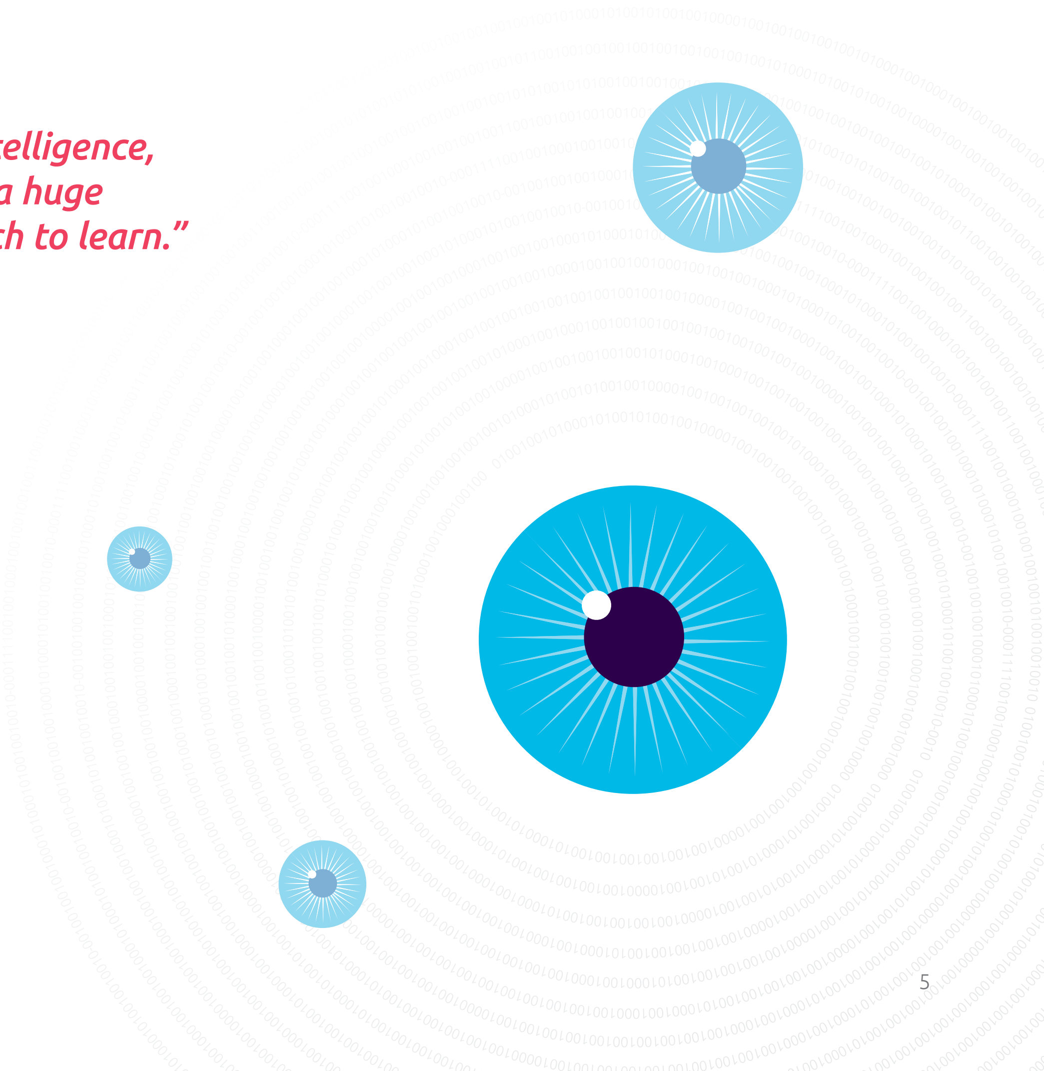

The first, the Learning Set, is the set to which we apply the scientific method that we described above. With this data set, we identify the model and the features that are important.

The second data set, the Validation Set, can then be used to fine-tune the model so as to avoid overfitting issues; and the third and final data set, which we might call the Testing Set, is used to check the predictive power of the model on data it hasn't yet encountered.

Overfitting? What's that? Before we answer this question, let's look at underfitting. This is when a model is conceived that bears little relation to the data to which it's applied. It's not able to explain anything. It's rather like the "idiot goldfish," which can neither learn nor remember. It can't explain the value of data, or account for its fluctuations.

Overfitting is the opposite – the model assiduously aims to accommodate all the data points. This makes it dangerous, because the predictions of the model for a new data input that has not been used during the learning process can be dramatically wrong. This situation corresponds to a model that has learned by heart all the data during the learning process. Faced with new data, it is simply unable to make right predictions, like a student who didn't understand what he or she learned.

The role of the Validation Set of data is to help find the model in the middle way between these two extreme cases, so that it effectively learns its input data almost, but not quite, by heart.

As machine learning algorithms are at the basis of the current state of AI, we understand that machines performing AI will systematically make errors. This is why it is still easy to distinguish between machines and humans. Let's consider the example of Captcha test. Each picture can be identified with an accuracy that is always less than 100% – say, 90%, or 0.9. If you have 3 x 3 = 9 pictures to classify, this means that the accuracy to do it well is something like (90%) ^ 9 = 38%! This is why Captcha tests are good at distinguishing between humans and robots – errors occur at expected levels.

The Turing test is similar (see Figure 3). Even if recent buzz news deals with machines succeeding at the Turing test, this is actually more about artificial stupidity than artificial intelligence. One can easily identify a machine by asking something like: "What is the logarithm of the age of universe multiplied by pi?" A machine would attempt an answer; a person would probably shrug and laugh.

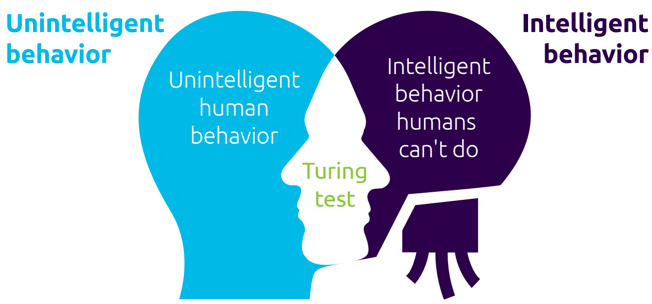

Figure 3. The Turing test

"As machine learning algorithms are at the basis of the current state of AI, we understand that machines performing AI will systematically make errors. This is why it is still easy to distinguish between machines and humans."

Taoufik Amri

Chief Data Scientist, Capgemini's Business Services

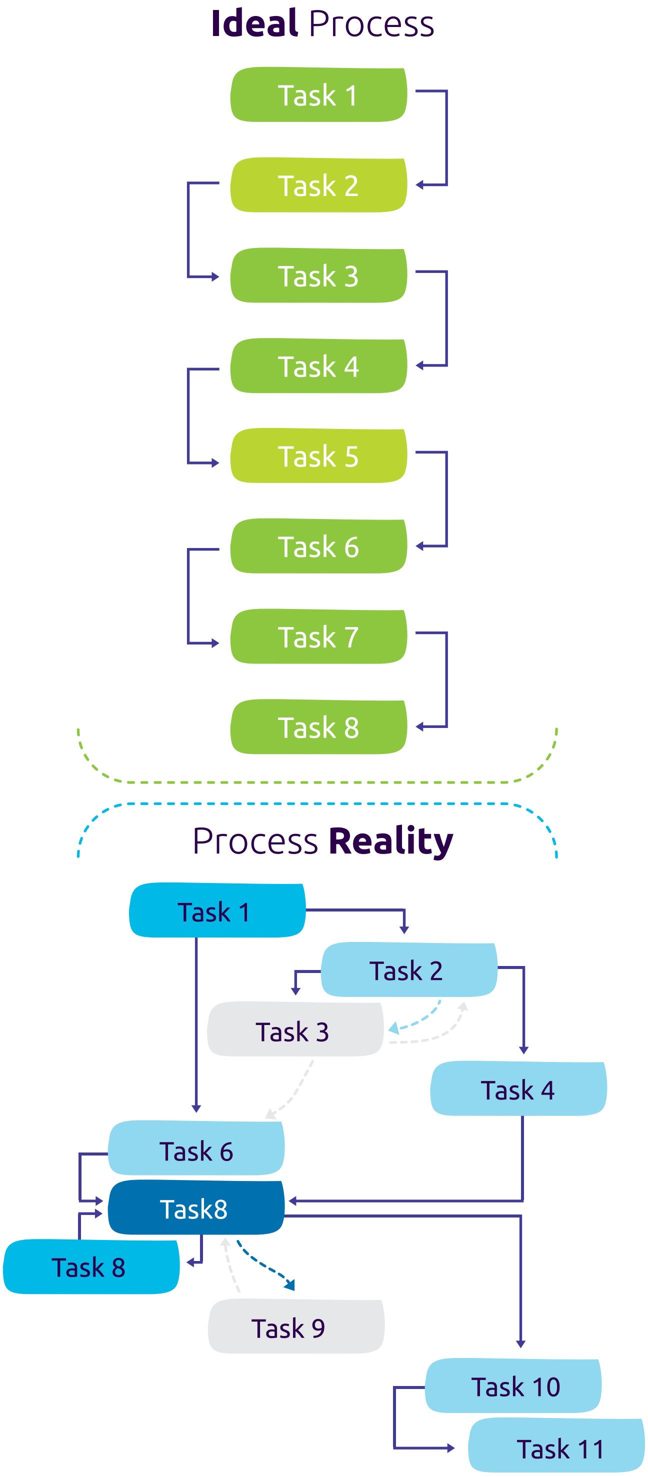

Figure 4. The ideal process vs. the process reality

Analyzing business processes with data

With digital transformation, business process management should be based on data. While this may already be the case with traditional IT systems that embody business processes, here it is not about machine learning applied to data, but about analytics applied to data logs in order to reveal the real orchestration of business operations. The purpose of machine learning is to automate a task that is only a node in a business process, which is itself represented by a graph connecting several nodes.

First, though, we need to be sure of what those business processes are – and that is not always a straightforward proposition.

Why not? Because for each process, people have a subjective and partial view, depending on their individual roles and perspectives. What's more, a process can be subject to constant change, making it difficult for its documentation to keep in step with its development. In fact, data scientists discover huge gaps between the business processes described by people working in and those revealed by process mining.

Process mining is a fast growing field in which data scientists apply analytics on data extracted from IT systems, namely event logs. These event logs are generated by all the components of an IT system that embodies a business process.

A plug and play process mining solution can be applied to event logs revealing the full and accurate picture of the business process – not the process described in the manual, nor the process as perceived either by management or by individual front-line staff, but the actual process, with all its secret add-ons, workarounds, shortcuts, dead-ends, and compromises.

The main outcome of process mining is a graphical representation of the business process – a graph with interconnected nodes, each one corresponding to a task, as shown in this diagram (see Figure 4).

The reality is of course more complex than any ideal, but it provides a sound and accurate base from which data scientists can redesign and optimize processes.

> ***"A plug and play process mining solution can be applied to event logs revealing the full and accurate picture of the business process."***

Taoufik Amri

Chief Data Scientist, Capgemini's Business Services

Improving business operations with AI

Analyzing business processes is just the start of the transformation. To improve the efficiency and effectiveness of their operations, organizations need to explore ways to introduce intelligent automation.

But they first need to get their houses in order, and when faced with the complicated reality of their own processes, the best way to achieve this is to apply a series of measures in a defined sequence, starting with the elimination of wasteful tasks before redesigning and automating those that remain.

At Capgemini, we use the Digital Global Enterprise Model (D-GEM) – our proprietary business transformation platform – which encompasses the tools and techniques for reshaping and streamlining processes, helping our clients remain competitive in a rapidly changing, business context (see Figure 5). This, in turn, enables an organization to seamlessly connects processes and people, intelligently, as and when needed to create the Frictionless Enterprise.

However, it's not only about replacing people with machines (real or virtual) – it's also about identifying tasks that can be performed better and/or faster with AI, and then applying a proven methodology for automating their work within a given process.

Process optimization also plays a part in deciding between humans and machines. It must be assisted with advanced analytics and modeling in order to assess as precisely as possible the efficiency and the value added by AI.

The Frictionless Enterprise

The Frictionless Enterprise seamlessly connects processes and people, intelligently, as and when needed. It dynamically adapts to your organization's circumstances to address each and every point of friction in your business operations.

At Capgemini, we have applied the Frictionless Enterprise to enhance cohesion across our entire suite of products and services. This enables us to respond rapidly to your changing requirements and deliver your specific business outcomes in a value-focused way.

We implement ways to detect, prevent, and overcome frictions – leveraging our latest thinking, organizational design, and intelligent solutions to achieve our goal of effortless operations.

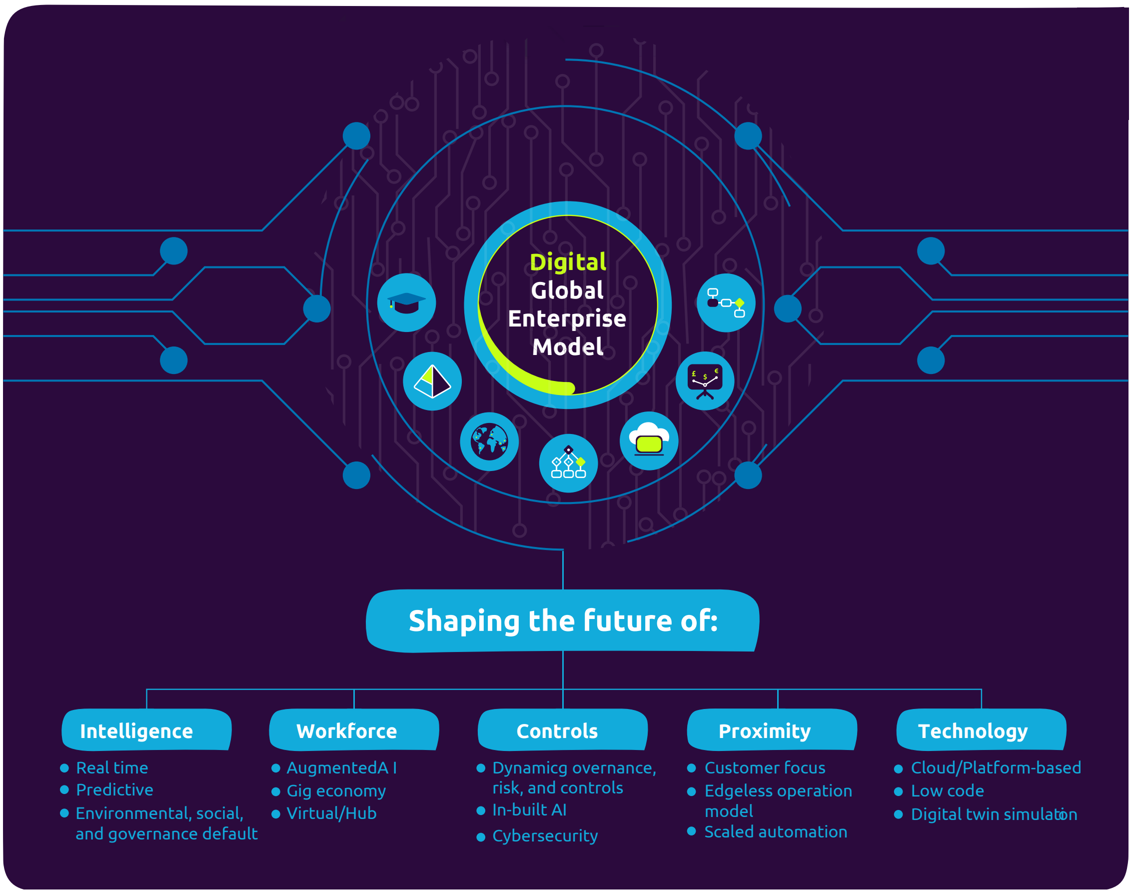

Figure 5. The Frictionless Enterprise powered by Capgemini's D-GEM platform

People or machines?

Our proven D-GEM platform helps to decide between machines and humans while optimizing a business process. It comprises three steps:

- Identifying tasks that can be performed better and/or faster with AI
- Measuring the value that AI can add in the whole process
- Designing human-in-the-loop solutions when the expected efficiency is not reached by machines alone.

So, how should organizations decide which process should be handled by which means?

In order to rationalize the decision, it is important to use the same metrics to compare the performance of people and machines. Quantitative measures that are also often probabilistic are used to characterize machine learning algorithms – for instance, metrics such as confusion matrices – can also characterize human performances, making it more straightforward to identify tasks that can be performed better with AI.

Then it is also important to measure the value that AI might add in the whole process, and not only for one task. We just looked at process mining as a means of revealing the truth about process flows. When data scientists and business experts use this real-world information to redesign processes before automating them, their first step in our three-stage methodology is to make decisions between machines and humans for each element of a task by characterizing human operators in the same way as by machine learning, using the same metrics for both.

The second step is to check than the whole process is efficient. For this, we need to carry out advanced quantitative modelling of the business process in order to simulate numerically all the possible scenarios, and to assess as accurately as possible the potential value that might be added by AI. Sometimes, a machine that is less efficient but faster than a human can bring value; this quick-and-dirty effect can only be revealed by modeling the whole business process. It depends on the position of the node within the graph representing the business process.

Finally, in some cases, these two steps cannot reach the client's expectation in terms of efficiency. The third and final stage of our methodology is to keep humans in the loop for performing such a task. For this, an AI-based operator should be added in order to orchestrate operations between machines and humans. Such an AI orchestrator is also characterized by a confusion matrix that takes into account all the possible outcomes.

In order to illustrate the confusion matrix that can be used for characterizing machines and human performances, we focus here on the case of the AI orchestrator.

This matrix provides a probabilistic measure of all the possible outcomes associated with the AI orchestrator that is based on machine learning algorithms. Such algorithms have learned to predict operations that can be totally managed by machines and those that can be handled only by human operators.

In this case, the confusion matrix is a 2 x 2 matrix. Each element of this matrix has its meaning (see Figure 6). The elements in green correspond to the true classification rates, the efficiency with which the orchestrator is doing the job. The other elements in red are the error rates with which the AI orchestrator misclassifies machines (human) operations as human (machines) operations.

The quantitative modelling of this scheme is based on this kind of matrix and enables the expected cost per operation and the efficiency of the process to be calculated. We can demonstrate that this simple human-in-the-loop design is better than machine only, at a cost that is lower than person-only.[1]

	Predicted **Machine** *Operations*	*Predicted* **Human** *Operations*
Real **Machine** *Operations*	**True Machine**	**False Human**
Real **Human** *Operations*	**False Machine**	**True Human**

Figure 6. The confusion matrix for an AI orchestrator between machines and humans

"Our proven D-GEM platform helps to decide between machines and humans while optimizing a business process."

Taoufik Amri

Chief Data Scientist, Capgemini's Business Services

AI orchestration for a French insurance company

Let's look at how the logic of the argument outlined above plays out in practice.

Challenge

At Capgemini, we've been working with a large insurance company based in France. Each year, around 50,000 dentists from across the country produce in the region of 250,000 dental care quotes, which the company needs to assess according to their health cover, and then provide a response to the customer. To date, there have been no pre-defined templates for dentists to make their quotes. To handle the workload, the insurer has employed 25 full-time employees in five regional centers across France.

The previous manual process (see Figure 7) worked as follows: the insurer's employees would read the quotes, input them into the legacy system, process the validation of the payback, and then communicate the outcome to the customer making the submission. On average, each quote took nine minutes to be processed.

Solution

Our intermediate proof of concept (PoC) has introduced an RPA solution in which data from the majority of quotes is read and extracted via optical character recognition (OCR), just before a data checking ensured by human operators. The rest of the process – input, validation, and feedback – is now handled via RPA. As a result, individual task time now stands at three minutes – one-third of the original time – and only needs the intervention of 12 full-time employees (just under a half) (see Figure 7).

We're now helping our client to go a stage further. An artificially intelligent supervisor, which we call our "AI orchestrator," is being introduced to manage the hand-off between the machine – the RPA system – and human intervention. This orchestrator uses a machine learning algorithm to classify operations that can be managed by the machine (around 85% of the total) and those that will need to be processed by people (the remaining 15%). The learning process of such an orchestrator can be unsupervised, supervised, or supported by reinforcement, depending on the nature of the data and the task to be automated. Here, it is supervised by the results of quality controls.

The AI orchestrator won't always make correct decisions about individual operations, and may allocate an operation to a person that could have been handled by the machine. In such instances, it's likely that the process efficiency will rise – but of course, the cost and time associated with that task will also rise. It's a trade-off – and over time, everything can be precisely calculated with our framework.[1]

Outcome

In the case of our insurance client, the framework we have established to handle machine-driven and human tasks has enabled us to anticipate that our AI orchestrator model will increase time savings to 80%. In other words, a task that previously took nine minutes on average could now be completed in just one minute (see Figure 7). Meanwhile, staff engagement can be reduced still further, to just six full-time employees. What's more, overall quality is expected to rise by more than 87%.

Implementation of these PoC exercises has been rapid. Once the model is established, the main task is to coach the AI orchestrator using data sets as described earlier in this paper – and the more data, the better.

A detailed technical paper on this application discusses this topic in more detail.

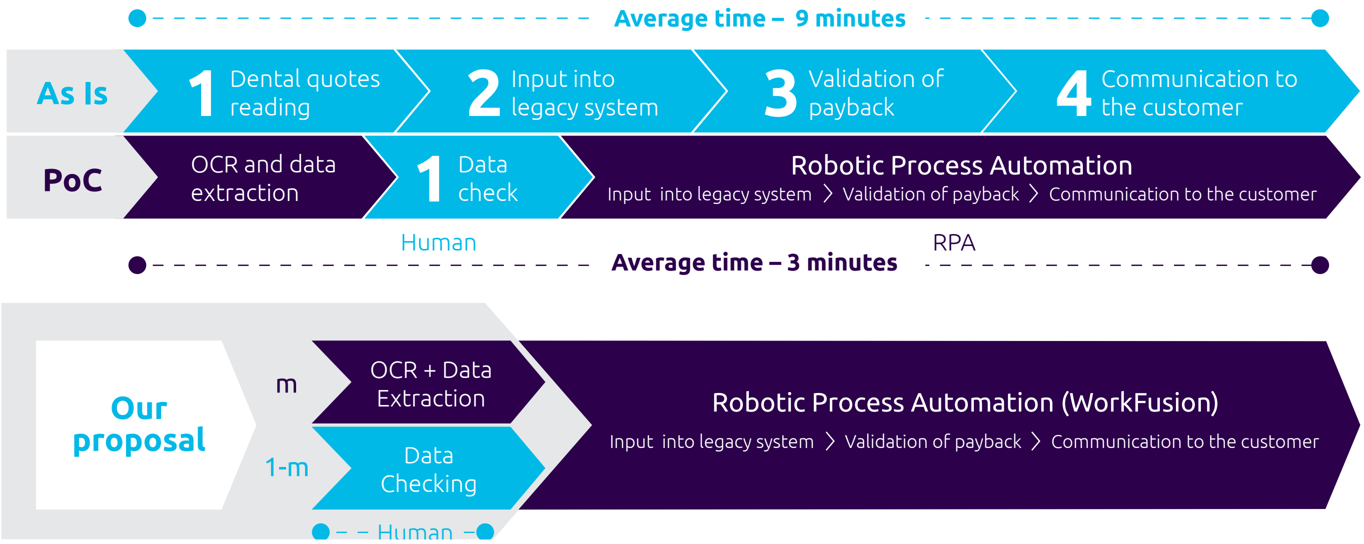

Figure 7. AI orchestration for a French insurance company

A realistic and optimistic vision

The introduction of automation and AI is often regarded these days as synonymous with the large-scale replacement of people with machines.

As this paper argues, this need not be the case. It is entirely possible to orchestrate activities between people and machines; in fact, it's not only possible, but preferable. By developing models and frameworks that re-engineer processes for the digital age, we can deliver business outcomes that are superior than could be achieved by either machines or humans on their own.

It's a different vision, which, as is so often the case at Capgemini, is rooted in practicality and circumstances in its bid for the operational excellence expected by our clients.

That's not all. It's also a more realistic and optimistic vision of the future of work in the age of AI that, happily, helps organizations meet their obligations not just to their customers and to their balance sheets, but also to wider society, which completely reinforces Capgemini's brand promise of "People matter, results count."

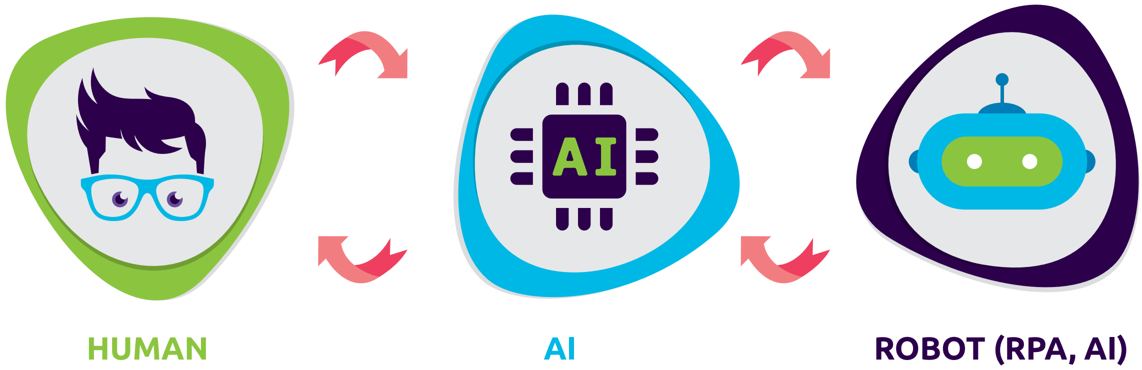

Figure 8. Orchestrating machines and humans with AI

> ***"The introduction of automation and AI is often regarded these days as synonymous with the large-scale replacement of people with machines. As this paper argues, this need not be the case. It is entirely possible to orchestrate activities between people and machines; in fact, it's not only possible, but preferable."***

Taoufik Amri

Chief Data Scientist,
Capgemini's Business Services

About the author

Taoufik Amri is a former quantum physicist. He graduated from the Ecole Normale Supérieure in France with a Ph.D. in quantum and statistical physics. As the principal data scientist for Capgemini's Business Services, Taoufik advises Capgemini's clients on implementing ready-to-use deep tech and intelligent automation solutions to dramatically accelerate the optimization of our clients' business operations. Taoufik also identifies tasks that can be performed better and/or faster with AI, measures the value added by AI with advanced quantitative business process models, and designs human-in-the-loop solutions.

Connect with us:

@CapgeminiBusSvc

www.linkedin.com/showcase/capgemini-business-services

businessservices.global@capgemini.com

About Capgemini

Capgemini is a global leader in consulting, digital transformation, technology and engineering services. The Group is at the forefront of innovation to address the entire breadth of clients' opportunities in the evolving world of cloud, digital and platforms. Building on its strong 50-year+ heritage and deep industry-specific expertise, Capgemini enables organizations to realize their business ambitions through an array of services from strategy to operations. Capgemini is driven by the conviction that the business value of technology comes from and through people. Today, it is a multicultural company of 270,000 team members in almost 50 countries. With Altran, the Group reported 2019 combined revenues of €17 billion.

Visit us at
www.capgemini.com

People matter, results count.