



**AI**

**MONITOR**

# CONTENTS

---

Nobody cares about AI . . . . .	5
Strategy . . . . .	6
Top-down-bottom-up. . . . .	7
Do something with data. . . . .	8
Crystal ball . . . . .	8
Stay close to your core . . . . .	9
It's a race . . . . .	9
Execution . . . . .	11
Data doesn't tell you anything . . . . .	11
Expectation management. . . . .	12
The feedback loop . . . . .	12
Capabilities . . . . .	13
The first mover disadvantage. . . . .	14
AI is the solution, what is your problem? . . . . .	14
Take back control . . . . .	15
Skills . . . . .	16
Ambition vs reality . . . . .	19
The world is your playground. . . . .	20
The business of play . . . . .	20
Nobody knows your business like you. . . . .	20
Data . . . . .	21
Let's build a lake . . . . .	21
Sharing is caring . . . . .	21
Show me the money. . . . .	22
Accuracy is overrated . . . . .	22
Governance . . . . .	23
Love and hate. . . . .	24
Infiltrate the processes . . . . .	24
Risky business. . . . .	24
Finding a home . . . . .	24
Trust . . . . .	25
Google doesn't know everything. . . . .	26
Legal & reputation . . . . .	26
In summary . . . . .	27
About the author the research . . . . .	28



Over the last months BRITEFLO sat down with more than twenty Belgian CIOs to discuss their practical experiences with AI. They all understand this technology has a huge potential. At the same time the road ahead is not fully charted, there are many questions to answer and decisions to be made. In the end there is no better way to learn than by sharing and discussing with peers and partners. That is what this document is about.

# NOBODY CARES ABOUT AI

If one thing has become abundantly clear through this BRITFLO+CIONET initiative, then it is this: nobody cares about AI. Or more precisely: what matters is the business value, not the technology in and by itself. Yes, AI can help make decisions by finding patterns that no human could ever detect. Yes, AI can do this on huge amounts of data and in real-time. But how do you embed this in an end-to-end business process throughout your company? Where is the business value? How do you get everyone on board? What are the legal ramifications? Which capabilities do you need to develop? In summary: how can you infuse AI everywhere in the complex, multi-dimensional context of your company? That is what matters. What we have here is the opportunity to fundamentally reshape all aspects of our business, to make use of this technological surge to transform how work gets done. The path to success is by making this a core part of your strategy. It will require the consideration of people, process, and culture and therefore the mobilization of all parts of your organization, and board-level steering that can't just be the sole responsibility of the CIO.

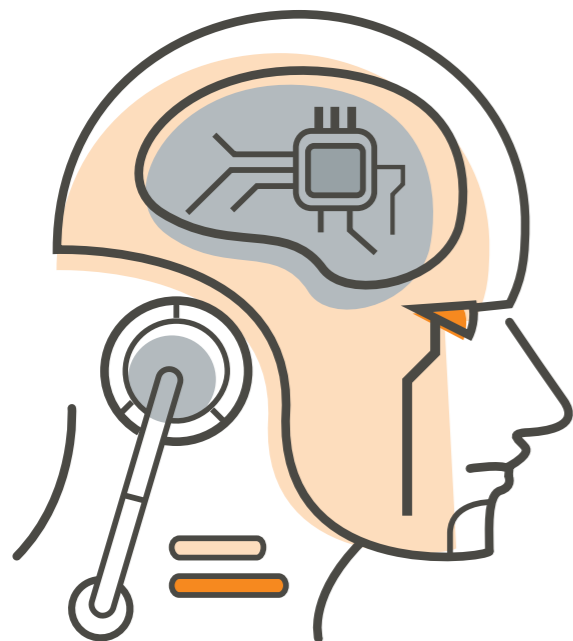
We also face a significant hyperbole around AI, an expectation that this magical technology can allow you to boldly go where no business has gone before. You may have encountered the statement by Arthur C. Clarke that

---

*“Any sufficiently advanced technology is indistinguishable from magic”*

---

Rest assured that AI is not magic. Surely 'data-science' would be a better term because what we have here is a set of tools and techniques that do seem daunting at first, but whose capabilities can be understood by everyone.



# STRATEGY

Where many companies have the vision, only a few also have a concrete strategy in which they have articulated the concrete steps necessary to capture the value. Banks are leading the way. KBC has a data science directorate reporting directly to the CEO. At BNPP the Chief Data Officer is member of the executive committee. Because they understand the impact of data and AI, banks have elevated this topic to the highest level in their organisation. Big pharma have a similar approach: at UCB 'AI in the core' is seen as a transformative force, they are not just looking for incremental improvement. These companies have given AI a specific place in their business strategy. They understand the impact it can have on their business model. They also have a plan.

The big benefit is clarity, focus and funding. Everyone in the organization understands how data and analytics of some sort, including AI, are important for the business. It is not hidden away somewhere else; it is a top-level concern. All other levels understand that they better pay attention because data matters to their company and to their career. Everyone is awake.

What these leading companies have in common is that they had a head start in digital. Creating data from data has been their business for a long time and now they are taking the next step. For many others the focus today is still on the data foundation and 'traditional' analytics since this is where they can reap the most benefit in the short term.



## Top-down-bottom-up

Top management typically knows that their company must become data driven. Translating this vision into a concrete strategy requires a "meet-in-the-middle" of business priorities and AI capabilities. To achieve this, the level of understanding must be increased on all levels of the organization.

What helps is when top leaders have a technical background. A new kind of business leaders with technical acumen is emerging. It helps them to grasp more easily the potential of this technology.

The sweet spot is where the vision meets the art of the possible. A mismatch on this level leads to disappointment. At UCB they look at business problems through the lens of new technological capabilities. And at Port of Antwerp, they regularly re-evaluate the options. What was impossible yesterday may have become possible today.

## Do something with data

Some boards have been mobilized by evangelists, marketeers or strategy consultants that paint a picture of doom on the one hand or utopia on the other.

On the upside this can get the company in motion and kickstart the transformation. However, when it is not generating a well-developed strategic plan, one that infuses data and AI in all aspects of the business, including milestones and associated budget, it leads to disappointment. After an evangelist session to the board one of our CIOs received the mission that they should "do something with data". It was not quite clear what was expected...

## Crystal ball

Several companies muse about an AI that can predict big economic shifts. There have been too many examples lately of unexpected events with a major macro-economic impact. But how can you predict that a ship will get stuck in the Suez Canal or that a virus will cause a global pandemic? AI will not come to the rescue here since even though they are very good at handling predictable events, not so much for these black swan events that appear out of the blue.



## Stay close to your core

Once the vision is launched, there will be an overflow of ideas, some valuable, others less so. One way to avoid distractions is to only down select projects that are aligned with the core mission. Generating proof points here will also help address the distrust that may live in your company.

At Port of Antwerp one of the key concerns is CO2 emission. They have set up a program to optimize the distance each of the 20 tugboats has to travel before arriving at the ship it has to guide into or out of port. Others prefer first to gain experience, like DHL which started experimenting outside of the core before now also touching their fine-tuned package delivery machine.



## It's a race

It is not enough to translate an AI vision into a strategy driven at direction committee level. A factory approach is necessary to keep it humming. It's a race. As always, the competitive advantage you can gain through technology has an expiry date. This is why execution is so important.



# EXECUTION

## Data doesn't tell you anything

While data may very well be "the new oil", you still need human experts to ask the right questions: by itself data doesn't spontaneously tell you anything.

Fortunately there are many ways to find use cases. Some are well known in the industry, for others you can tap into the combined knowledge of your employees, both business and IT combined.


Focus groups with a certain competitive element, for example in a shark-tank model, seem to work well. There are two advantages here:

1. The ideas are grounded in your own business and
2. The teams can regularly present cases to the ExCo, which brings the team a sense of pride and simultaneously keeps the ExCo engaged.

The challenge now becomes to prioritize. You will receive more proposals than you can handle so you need a structured approach to down select the most promising ones, not based on the "coolness" of the technology or whoever shouts the loudest, but by defining metrics that ensure alignment with the business value.

Another interesting avenue is to create an open innovation platform where different partners are invited to experiment. This is what Port of Antwerp has done with the "Plug and Play Maritime" initiative.

One dominating category of use cases has to do with planning and prioritizing human work. At Picanol they prioritize maintenance work on client machines, at Randstad the seller call-list, at Coca-Cola the bars to visit for cross-sell, at DHL the route to drive, at Kinopolis the number of interim workers needed for the next day.



## Expectation management

Realistic expectation management is critical to long term success. One multi-national refused to put an information extraction system in production because the recognition accuracy was not the 98% they thought they needed. Their inflated expectation made them miss out on real business value.

On a more realistic and very valuable level one company achieved a 20% reduction in manual work with an information extraction system that had only 50% accuracy. And with just 65% accuracy UZA reduced the duration of a migration project from 5 years to only 1.5. While manual validation is still necessary in these cases, a significant speedup is realized because the system helps the human data curators more quickly identify where to find the information. Since it would require super-human abilities, fully deciphering a doctor's handwriting is still squarely in the realm of science-fiction though.

## The feedback loop

AI models, in particular neural networks, are highly sensitive to correlations between the input parameters. This is, after all, what they were trained to detect. When changes in the real world occur, there will be inevitable entropy in the data. The inputs will slowly change which will, at some point, lead to unexpected predictions. This is the infamous model drift. Depending on your use case these may have a smaller or larger business impact, but the result stays the same: as the system starts to make more and more strange decisions, the confidence of the business in the ability of the AI solution to deliver value deteriorates.

Interestingly almost no interviewed company actively monitors their AI systems for model drift. If the topic is addressed at all, then this is as part of a regular project update lifecycle. That can be too slow in any case and what if the update is delayed or cancelled altogether? This is a disaster waiting to happen. At some point your AI model will no longer perform as expected. This is a certainty. To handle this risk, you must build the feedback loop into your core approach from the start.

If you don't have many models it doesn't have to be complicated. You can schedule a regular manual retraining. Quickly you will discover that this becomes unwieldy. BNPP and KBC have hundreds of models in production. They have no choice but to use MLOps to automate the whole process end to end, from data capture to model monitoring.. As KBC discovered, supporting such end-to-end processes has a significant impact on a more traditional release calendar that had fixed release windows during the year. This ability to deploy on-demand to production, whenever there is a need, sounds of course familiar in the context of the IT industry's move to agile and an associated continuous release management capability.

Even when you are not yet automating the process it is still necessary to plan the capture of new training data, the correction of errors and the retraining of the model. At one company a project initially failed because the project plan did not account for regular retraining which was needed to make the system "learn", i.e., to make it more accurate over time. The errors were corrected but the training data was not updated, and the model not retrained. If this part is not planned, budgeted, and communicated at the start, then disappointment is inevitable.

At Randstad the feedback loop is built in. They run benchmark simulations on historical data and have continuous monitoring at runtime.

Sometimes it is a bit more difficult to monitor the accuracy of an automated system. Think of postal services that use automated routing of letters. A wrongly routed letter or package is only discovered at its final destination, which can be days or weeks later. Setting up a system to capture this and link it back to the original routing decision is not easy.

At this time this end-to-end automation is at the pinnacle of maturity, with only few companies at this level. In the end the only way to professionally use AI is to control the full lifecycle, from data gathering, through model creating and training, up to the monitoring in production, including the retraining when needed. There is work to do.

# CAPABILITIES

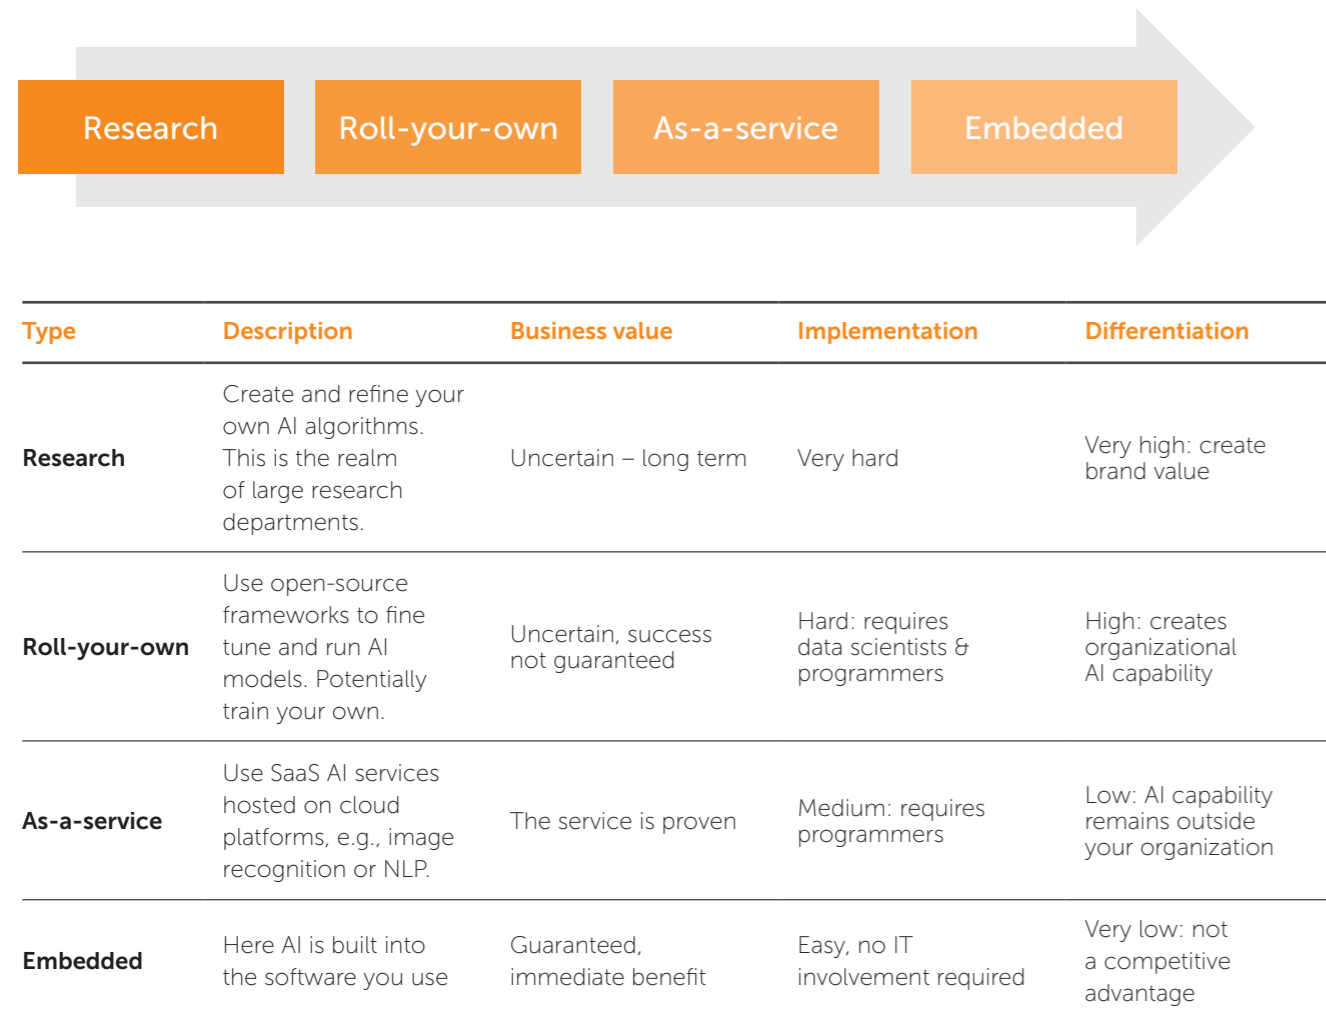
We have a broad spectrum of AI capabilities at our disposal, not only in terms of functionality but also with regards to consumability, i.e., the ease, or lack, of acquiring, developing, maintaining, and running an AI solution.

On one end of the spectrum there is fundamental research that is typically done by large corporations. These might be multi-year, multimillion Euro efforts where they aim to push the frontiers of the possible. Examples are IBM's Watson or DeepMind's Alpha Go. These programs require a sizable budget as well as world class expertise while the time to value is uncertain.

On the other side we see AI embedded in off-the-shelf software where it delivers business value without requiring any knowledge of its inner workings. Qualiphar, for example, uses a SaaS software with embedded NLP for customer feedback analysis. While this is very valuable, it is not an area for competitive differentiation since others can do the same.

Belgian CIOs often find themselves in the middle, where data-scientists and developers glue together open-source and SaaS components and fine tune those to their specific tasks and with their own data. You need to decide where you want to play for each domain in your business.

Figure 1 – Continuum of AI capabilities



## The first mover disadvantage

Over time AI capabilities will trickle down from research to off-the-shelf at which point they become commodity. The speed at which this happens seems to accelerate. Visual recognition on static images used to be the pinnacle of AI research and has now matured to a point where it is available as a service from any cloud provider. The same is now happening to information extraction from documents which was initially developed in research and academia, then moved to start-ups when the underlying Transformer technology became easier to use. Today it is being offered as a service by large cloud platforms.

When is a good time to get into the game? The answer to this question depends on the strategic intent you have with regards to the competitive advantage you want to get from AI.


What matters is not so much which functionality or use case you are aiming at, but rather which capability you want to build within your organisation in order to create not only business value, but also a competitive advantage which will last.

The long-term costs also need to be considered. While embedded or as-a-service solutions are bringing short term value at a low entry point, the long-term costs are often overlooked. An example is Internet of Things (IoT). When it became hot, several industrial companies developed an in-house solution to capture sensor data for downstream analysis. Because in that sector departments typically operate independently, some of them now have a separate solution per department. Today they face a migration cost to consolidate these.

## AI is the solution, what is your problem?

While AI gives us the ability to process data in a previously unthinkable way, we should not forget all the other tools that remain at our disposal. In a phenomenon typical for the IT industry, it seems that every new technology is, for a while, considered as the solution for everything. It can be worthwhile to address a problem first with "traditional" predictive models like linear regression or rules based expert systems. If anything, this approach provides a first accuracy baseline that more advanced models can be benchmarked against.





# SKILLS



To be successful there is a need to increase the level of understanding on all levels of the organization to create a common culture and vocabulary encompassing data, analytics, and AI. This starts with the top that must get a grip on the business potential. A recurring experience is that training sessions given by consultants go too deep into the technology, sometimes even diving into the mathematics. Although beautiful and fascinating this is not the right level. What we want to see here is a clarification, in plain language, of what AI is and what it can do for the business.

---

*“What is worse: a false positive or a false negative?”*

---

## Take back control

A key consideration here is how you can keep control of your strategic knowledge and the intellectual property you create when developing AI systems. Apply serious thought to what you are willing to keep internal and what you are ready to externalize. As you will often work with external partners, make sure to have clear agreement on what they can do with the outcome of the work.

An AI that is trained and finetuned on your data contains core company knowledge, your secret sauce. Do you agree that this knowledge will be shared, through your supplier, with your competitor? After years of outsourcing, we now see a trend to “take back control”, i.e., to become self-sufficient again for core IT capabilities, including AI because of its potential to deliver competitive advantage.

When Isabel automated its invoice processing through AI powered information extraction, it agreed on a clear separation between the IP of its vendor and its own, i.e., the in-house rules that determine what constitutes a valid invoice.

Consider this scenario: a large machine manufacturer deploys its machines on-premises at the factories of its clients. These clients use their own expertise to fine tune the configuration of the machines based on the quality of the inputs and the desired output. The machine manufacturer centrally captures all these settings from all

its clients and is now able to train an AI system to automate the machine configuration. In practice this mean there is a knowledge transfer from one client to another in the form of a trained AI model. In some highly competitive markets, the top performing clients will rightfully not accept this sharing of their core knowledge and expertise.

One interesting way technology can come to the rescue is through the use of containers. A sometimes-overlooked advantage of containerization is that it smoothens previously difficult collaboration by allowing you to securely bring third party algorithms onto your own platform. The principle “bring the analytics to the data” becomes practically feasible across companies. The reason is that containerized algorithms can process your data without giving the third-party access to it and the reverse is also true, the third party can deploy its intellectual property in your environment without exposing it. You only need to agree on the interface, how data is fed to the container and how the container delivers its result, that’s it.

Explaining, for example, important concepts such as false positives and negatives and the impact they have on the business case.

We need deep thinking on the transformational possibilities of AI. What does it mean when you can process a client request in real-time when it used to take weeks because a human had to make a decision? What does it mean when you can do personalized risk scoring which also allows you to determine a personalized price? A training program tailored to the ExCo is in order.

Businesspeople on all levels must get a grounding in the concepts so they can identify interesting use-cases in their domain and also spot AI washing, i.e., cases where the impact of AI is either grossly overinflated or non-existent altogether.

For IT people training is crucial as well, not all of them have had the opportunity to work with the new technologies, even though it will impact them in all aspects of their job. Think for example about the emerging AI-powered source code generators that will assist programmers.

The whole organization must develop a data/AI awareness, a common culture. Often the maturity is different between departments (e.g., a research group can have high maturity and while it may be lower in the sales team). Awareness must be raised at all levels or projects get stuck between vision and execution. At UZA awareness is already high because healthcare has been a focus domain for AI for many years now. In many companies there is now a targeted education program since insights seldom appear through spontaneous osmosis. UCB has a “digi-minds” & “data-minds” program. At Bekaert they hold “fireside chats” and Picanol has a DigiCube to make employees familiar with new technologies.



*“People join for the ambition but are confronted with reality”*



## Ambition vs reality

The war for talent is real. It poses a challenge in general in the IT sector and is especially fierce when it comes to data science. On top it is hard for companies to retain AI talent because they lack the volume of projects. Data scientists want to do data science work and if you cannot offer that they will leave.

To address this need there are AI service companies that have created talent pool of skilled consultants they can deploy. This is a solution for some of the companies interviewed, they either fully lean on their partners or create a healthy mix of internal and external skills. This approach also ensures a constant inflow of experience and ideas. Others find the domain too strategic to outtask. By leveraging AI as a catalyst for innovation and by treating it as a transformational project at corporate level, they are able to sustain a pipeline that keeps the data-science team occupied and challenged.

Internal people can also be reoriented towards data-science: at NMBS they have set up a learning & development program to retrain people to become data scientists. Maybe these people will not get to the level of the academically trained but considering the commoditization of AI this is also not always needed.

With regards to the mix between internal and external skills the consensus lies with at least a strong core internal team that is extended with external consultants. In some cases, like at BNPP, these externals seamlessly blend in with the internals to form one team. The distinction disappears, they are treated equally. This makes sense since they are also equally likely to stay or to leave the company. The exact ratio between internal and external tends to differ significantly and typically lies between 20/80 and 80/20. A minority of companies have no internal people and fully rely on outside parties and on the other end of the spectrum there is KBC who have no external people at all in their data-science team.

## The world is your playground

Startups in US have no trouble working remote, tapping into a talent pool infinitely larger than the one in Belgium. They also easily collaborate with experts in India or Eastern Europe. Our companies more often wait for the perfect candidate that speaks the same language and lives within driving distance. Depending on the area you are in this might be next to impossible leading to companies actively trying to hire from the competition. Some also struggle because they are not perceived as "hot" technology companies. The more reason to cast a wider net.

Many companies have set up collaboration with universities (Barco, KBC, BNPP, J&J, Port of Antwerp, ...). Some, like Renson, also sponsor a specialized chair. This brings you into contact with the brightest minds and can be an excellent pool for recruitment.



## The business of play

The PhD level data scientists at Mediahuis stay engaged because they have sufficient interesting projects, and they get room to experiment. Experimentation is particularly important in the AI space because the path to a working model is often not clear ahead of time. There are countless decisions to make in terms of data, model selection, tuning, etc. that it's easy to get lost in the rabbit hole. That is why the experimentation itself should be structured and timeboxed, for example by allocating one sprint (two weeks) to data and then a second sprint to model testing. After that a decision must be made whether to continue the chosen path. To avoid too much freewheeling the innovation team at VDAB is tasked with delivering at least three projects to production every year. This keeps the team focused on business results next to experimentation. Results also bring recognition of the team to the top.

## Nobody knows your business like you

A recurring frustration is that working with AI startups involves too much hand holding. They may have the technical knowledge but often lack the business knowledge necessary to solve the problem at hand. If anything, this challenge should look familiar, because IT has exactly the same, i.e., how to gain credibility at the business side. At Coca-Cola they handle this by moving business analysts with few years of business experience into IT where they infuse their domain knowledge and continue to serve as a bridge between the departments.

# DATA

It does not matter whether you are a small or a large company, everyone is struggling with data. Many companies are still working on the basics in this respect: what data do you have, can you access it and what does that data mean?

There are countless examples of misunderstandings on seemingly unambiguous concepts such as full-time equivalent (not the same for HR as for operations), a subscription to a magazine, a train (where is the front, where the back?), etc. If a data element is stored for one purpose it is not necessarily useable for another.

At UZA the iNNOCENS project aims at creating a digital assistant that can monitor and predict blood poisoning for premature babies in neonatology intensive care. Even though 10 years of data was captured only a subset turned out to be usable for training the predictive algorithm. At first the goal was to correlate sensor data with clinical observations. Matching automatically captured timestamps with manually registered ones turned out to be impossible. You cannot go back in time to fix these issues. In the case of physical devices, such as those built by Barco, you might not even be able to fix them for the future when timestamps are not captured, and it is hard or impossible to roll out a firmware update. In the end the sensor data at UZA was useful by itself for predictions, without the extra data. Still this experience led to a new protocol for capturing data and a hospital wide, centralized data warehouse.

## Let's build a lake

It might be tempting to start capturing all possible data into a huge data lake from which, at an indetermined point in the future, value will arise. It quickly becomes clear that the cost of doing this is prohibitive. It also ignores the fact that data has diminishing value over time. It's better to start from a solid understanding of the business value of each dataset, i.e., what are the core, differentiating data elements for your company, in your context. At UCB they have avoided this trap by consciously not building a huge, centralized data lake up front but rather by starting from the concrete needs of each project.

## Sharing is caring

Since weather has an impact on almost every industry it is an often-used data source. Carrefour uses it to determine how to stock the shelves (ice-cream!) and Kinopolis to estimate the number of movie enthusiasts coming to the cinema. Interestingly it is the change of weather that is important for them, rather than the fact if it is sunny or rainy. In any case, once clients started to order tickets online, which 70% of people now do, the prediction was pretty much a solved problem, reaching 85% accuracy. This helps Kinopolis tremendously with the staffing for the next day, a process that runs with an automated link to the interim agency.

Weather data can easily be purchased on the market. Other external data sources may be harder to come by. Especially when it comes to data available within an ecosystem there is reluctance to share. The big players like Facebook, Google or Apple tend to leverage their access to data as a way to force you onto their platform. On top it is not always easy to understand their data and correlate with your own. That is why companies find it easier to collaborate with local partners. This can be through a direct partnership, e.g., Renson (Healthbox) with Niko (Home Control) or sometimes in the context of a university project, which seems to increase the trust and willingness to share between partners.

In some cases, one company alone is unable to gather enough training data. For example: Port of Antwerp has a project to continuously assess the quality of quayside ladders through visual recognition. The challenge is to collect enough pictures of damaged or missing ladders, a situation which is luckily exceptional but at the same time important. Sharing such images between ports creates a win/win proposition.

## Show me the money


Many companies have the ambition to create a new revenue stream through selling data or insights derived from data. This external data monetization is not easy to achieve, particularly in B2C. How much will you pay for a monthly subscription that promises, but not guarantees, to save you some money on your heating bill? On the B2B side Carrefour has succeeded in monetizing their core data asset: your supermarket receipt. Through their Carrefour Links initiative, they sell anonymized and aggregated purchasing behaviour, which is a treasure trove for consumer goods companies and marketeers that are eager to better understand their customers.

## Accuracy is overrated

The existing data infrastructure has historically been created to serve a reporting need. AI puts different requirements in place that force us to rethink our approach. Higher data quality is better, but it does not have to be one hundred percent perfect. A neural network for example will find the statistically significant signal in the noise. A bit of noise can be acceptable, as long as everyone is aware that there are different data quality levels in the company, some very accurate, some less so. A data catalogue becomes crucial to avoid painful misunderstandings.

On the flip side, AI can be used to generate reliable information from incomplete master data, as NMBS is doing when they determine the exact composition and orientation of their trains. At some point, as it travels through the country, a train might split, or conversely be merged with another train. This changes the position of the first-class cabins, the restrooms, and the special carriage for bicycles, all of which is information NMBS wants to provide accurately and timely to the travellers. The train can also change orientation, where the front becomes the back, which is confusing to maintenance workers sent out to fix the "left door of the second carriage". Trains have a lifespan measured in decades. Retrofitting them with new sensors is not trivial. Through AI it has been possible to create the information necessary for real-time mobile use cases.

Don't let the data struggle stop you from starting with AI, there is so much that can be done without waiting for the perfect data platform to be in place, something that might never happen.



# GOVERNANCE

It is no longer acceptable today to just throw an AI model into production and see what happens, we need to make it manageable at scale, auditable, and above all, continuously valuable to the business. The only way to graduate from experiment to predictable, real-world value is by professionalizing the governance around AI.

In healthcare this is business as usual since they need go through strict compliance reviews in order to put their product into the market, for example when an AI would be used for clinical diagnostics. This is the case with the Medical Device Software regulation in EU.



## Love and hate

Regulation is a double-edged sword. Healthcare companies feel constrained because of the European privacy regulations. Hospitals like UZA ask themselves why they must comply to the same rules as commercial companies when their motivation is to provide better care for all. Researchers and pharma companies struggle for example to get access to enough genetic data and healthcare records which they need to combine in order to discover meaningful insights.

On the other hand, although compliance comes at a steep cost, it did force finance and healthcare companies to get their data house in order, well ahead of other sectors. They now benefit from the investments they made in the past because it forced them to build the foundation that today makes them leaders in the application of data and AI.

## Infiltrate the processes

What works best is to infuse the attention to data & AI into the existing governance to make sure it becomes the standard way of working. Rather than creating separate, new processes it is much easier to become part of existing processes, committees, and checklists. This is how companies like BNPP & Coca-Cola, that already have strong governance in place, have managed to break the code. And if the business is not so much involved yet in AI then you can do like Carrefour, that has embedded this as part of the architectural guidelines that every project has to comply to. A project will not be approved without a clear position on the capture and use of data and analytics. If the project team believes that this is not applicable to their project, then this will be seriously challenged.

## Risky business

Since the use of AI implies exposure to new risks and because it is in generally the right thing to do, there are also new governance requirements that can be handled by a new data privacy & ethics board. Since this is not technical but rather about business risk this board should be managed by the business in a multi-disciplinary team.

The ethics board at VDAB is comprised of 50% external experts. Would it not be valuable to infuse this outside-in point of view in every such board?

## Finding a home

Where to place the data scientists in the organization is a common challenge. In practice they can be found all over the organization, in finance, in HR, in the business but also, in most cases, they are housed within in IT. Wherever they are, they are often pooled together, for several reasons:

- Work distribution: there are only a few data scientists, and they are pooled to work for different departments.
- Retention: any one department does not have enough interesting work to offer, data scientists leave because of lack of growth perspective.
- Organization: when they are not pooled, they might be "claimed" by a department that does not want "their" data scientists to work with or move to another department.
- Maturity: in many cases the business is not yet considered to be ready to host data scientists. They may lack the understanding and language to communicate with the data scientists. In one company the data scientists were moved from the business back to IT after a first experiment to embed them failed. A second try, with more guidance, is now planned.
- AI is a multi-disciplinary undertaking that cannot be run by IT alone, it requires deep understanding of business, data, and algorithms. Most companies seem to settle on a centralized pool of experts that work closely with, but are not directly assigned to, a particular business department. At BNPP this happens under the umbrella of a Centre of Excellence.

# TRUST

---

In many cases an AI can make better predictions than the best human expert. It can simply pull from more data sources and a richer history of examples. This is not easy for those experts to accept. The problem is particularly visible with black-box AI models like neural networks which cannot easily be explained.

---

*“ The more you are an expert,  
the more you believe  
you are infallible ”*

---

The lack of explainability is a recurring roadblock for AI success: at one company a customer churn prediction algorithm was not believed and acted upon, at another company employees didn't even propose the next-best-offer to clients because they couldn't understand and explain to the client the reason the proposal was generated. The good news is that there are now techniques to understand to what extent certain input parameters contribute to an AI decision, for example with loan approval, where amount, salary, age, savings amount, etc. all can contribute positively or negatively in the decision to approve or decline the loan. As a side note, it seems like we put the bar higher for the AI than to human experts, who are not always asked to explain their own decisions either.

Coca-Cola successfully implemented an algorithm to optimize cross and upsell. It determines which bars and restaurants salespeople should visit and what to propose. This approach helps giving priority to promising clients. The key here was to first implement this approach in a region with a tech-savvy leader, i.e., one that has a better understanding of the underlying technology, leading to more trust and a willingness to give the new approach a chance. This use case, prioritizing seller call lists, was mentioned often during the interviews. It is also generally successful. Although at one company the program was abandoned after the sellers ignored the recommendations and management did not impose its use.

Even when the experts were involved in the ideation and the development of the system, their initial positive sentiment can change once the AI starts impacting their own work on a day-to-day basis. Work prioritization is a popular domain where this happens a lot. To ease the transition, it can help to initially position the AI as just an extra data point that can help the expert to decide. There is a big difference between "here is an AI tool that tells you what to do" and "here is some extra information that could be useful". The latter, more humble approach also allows you to learn as you go since in many cases an AI system will require fine tuning as it is tested out in the real world. At Port of Antwerp the AI prediction on the required number of tugboats is presented as an extra data point on the same dashboard that the planners were already using.

## Google doesn't know everything

You might discover that the experts still bring a lot of insight to the table. At Picanol they were at first baffled by a quality problem in the factory, until an expert pointed out that the salt level in ground water had gone up because of recent spreading of de-icing salt in wintertime. This knowledge can then lead to extra sensors and the inclusion of this data in the predictive model.

Drivers at DHL can modify their AI generated package delivery route. After all, the drivers know the ins and outs of the routes they travel daily. The Sunday morning market that blocks the only road through a tiny village may not be registered in a database and is therefore not available to the AI. And at UZA the queue in the emergency room is prioritized by an AI, but the doctors always have the final decision power.

## Legal & reputation

Depending on the business and the use cases, AI comes with board level concerns on trust and ethics which are often underestimated. Concerns like bias, explainability and accountability can have far reaching legal ramifications and impact on reputation. The attention that Europe is placing on the fair use of AI and the strict attention to privacy can be seen as a catalyst in this respect since it focusses the thinking and keeps the discussion alive at the top executive level. It might be best, although admittedly somewhat challenging, to view the regulator as an ally that helps push the company forward in the right direction. Ask yourself, would you move at the same speed without this external push?

One job matching service suffered from unexpected bias when it used NLP to extract information from curriculum vitae and job openings to match job seekers with jobs. They eventually addressed this challenge by working with an intermediate ontology that controlled the possible outcomes of the model, instead of letting it freely find associations.

All these concerns can be addressed if you include them into the project from the start. It is next to impossible to fix them as an afterthought.

# IN SUMMARY

---

AI can be a game changer for both business and IT. Sharing our experiences, the good, the bad and the ugly, helps us discover the path to success together.

While the technology is fascinating and sometimes feels like science-fiction, in a business context nobody really cares, what matters is the concrete value and the opportunity for transformation. Taking a strategic approach, driving the AI initiative from the very top of the organization, is critical. It is not just the CIO's responsibility; the whole board will have to sharpen its technical acumen so it can align its vision with the art of the possible. With a well-defined mission it then becomes clear which use cases to prioritize and which data to capture. The team that can navigate the tricky path to success requires skills

and experience, which you can nurture with a healthy mix between internal and external experts. It now becomes 'just a matter of' automating the flow end-to-end, from data to models in production. To be fair, this part is state-of-the-art today but it does become imperative as the number of AI use cases grows. Controlling and automating the process also allows you to structurally handle governance challenges such as bias and explainability which are instrumental to a successful business outcome.

AI has huge potential and when you do it right, your future will be bright...

# ABOUT THE AUTHOR THE RESEARCH

---



**Wouter Denayer**  
Innovation magician & founder

Wouter is a technology optimist, innovator, and keynote speaker. After 25 years of working with top enterprises, the last seven as CTO for IBM Belgium, he has now started his own company Briteflo where he uses all that experience to help companies on their digital transformation journey. His focus is on the efficient use of new technologies for true disruption while putting the human squarely at the center. He is also on a mission to strip away both hype and fear surrounding Artificial Intelligence, so that its true potential for business and society can be realized.



### About CIONET

CIONET is the leading community of IT executives in Europe and LATAM. With a membership of over 7000 CIOs, CTOs and IT Directors, CIONET has the mission to help IT executives achieve their aspirations. CIONET opens up a universe of new opportunities in IT management by developing, managing and moderating an integrated array of both offline and online tools and services designed to provide real support for IT executives, so they can do more than just keep up with change but ultimately define it.

[www.cionet.com](http://www.cionet.com)



### About briteflo

BRITEFLO helps private equity firms make better investment decisions in companies for which technology is core to their business. Our technical due diligence combines state-of-the-art technology expertise with deep business acumen. We work with the largest sustainable private equity firm in Europe, leveraging our expertise in data science, IT architecture and security to help them grow their portfolio. We believe that passion makes the difference, and that people naturally excel when they work for a goal they believe in.

[www.briteflo.com](http://www.briteflo.com)





**briteflo**