



Profitability study

Summary

This study aims to assess the profitability of modernizing a simple business use case by integrating artificial intelligence solutions. To do this, we will first analyze the cost of the business process in its current state, detailing its steps and inefficiencies. Next, we will compare three main categories of natural language understanding technologies: approaches based on machine learning, those based on pre-trained language models, and finally those using knowledge-based systems. This comparison will take into account criteria such as implementation costs, performance, scalability, and adaptability of each solution to identify the most relevant one for optimizing the use case under study.


Presentation of the use case

Let's take a simple use case that affects many companies: talent attrition. Bertrand is the CEO of an insurance group in Niort, France. Many employees at his company are unhappy and are resigning. Given the numbers and frequency, this situation worries him, and he decides to discuss it with Sophie, his Human Resources Director. After listening to Bertrand and sharing her observations with him, Sophie suggests conducting a study on job satisfaction in his company to better understand the reasons why employees are leaving. Bertrand thinks Sophie's idea is a good one and asks her to come back to him with areas for improvement within three months.

Sophie remembers hearing about a specific study on employee job satisfaction. After some research, she finds work carried out by Frederick Irving Herzberg, an American clinical psychologist. He has worked extensively on motivation in the workplace. She decides to create a questionnaire to gather information on fulfillment, recognition, the work itself, responsibility, and social advancement.

The employees' answers to these questions would give her a better understanding of which factors were out of balance. Once identified, she could propose actions to correct these discrepancies and ensure that her company's employees were satisfied.

This HR process, which concerns a job satisfaction study, can be broken down into six distinct activities.



The first task for Sophie and her team is to formalize the survey questions. They agree on eight questions and encounter their first problem. The company has 9,000 employees. So, conducting individual face-to-face interviews is not feasible. They therefore decide to use a software product available on the company's intranet to distribute a questionnaire to employees.

The second task is to ask the IT team to prepare the distribution of this questionnaire to employees.

The third task involves the employees who will have to answer the questions. Of course, the responses will be treated anonymously.

The fourth task is to collect all the employees' responses. These responses will be stored in a company database hosted on an internal server. Yes, this is a precautionary measure, as this data is sensitive and may contain personal information.

The fifth activity involves Sophie and her team reading the employees' responses to assess their satisfaction.

Finally, the sixth and last step involves summarizing all the information from the analysis and generating a report that will then be presented to Bertrand.


When humans perform work

Everything is falling into place: the questionnaire has been prepared, the employees will respond, and Sophie's team will analyze the results. But then, things don't go as planned.

To analyze job satisfaction criteria, Sophie and her team will rely on Herzberg's study. Herzberg identified 17 factors of job satisfaction and dissatisfaction. Each of these factors has a specific meaning that must be recognized when reading the employees' responses. In addition to these, it is necessary to capture all words related to emotional (8 primary emotions and 24 secondary emotions) and affective states (positive and negative) to better assess what employees appreciate and what to improve. This will take time, a lot of time, even if the employees' responses are short.

Due to time constraints, Sophie and her team cannot read all the responses provided and must sample them to meet Bertrand's deadline. They read a few answers to assess job satisfaction factors, emotions, and affective states. Then prepare a presentation for Bertrand to report on what is going well and what needs improvement.

These results do provide some clues, but they have one drawback. Due to time constraints, just a sample of responses were analyzed, so certain job satisfaction factors were overlooked and don't appear. The examination is therefore incomplete.



However, Bertrand is pleased with these initial results and decides on two things. First, he wants these employee surveys to be conducted twice a year. He asks Sophie to talk to the IT department to find solutions, particularly artificial intelligence solutions capable of understanding natural language, so that all employee responses can be analyzed and the time spent on analysis can be reduced. Sophie will present her use case to Julien, the IT Director.

When work is done by humans and machines

Assessing job satisfaction requires reading responses, identifying specific words, and understanding them. Today, these processes requiring human intelligence are performed by people because it is difficult to modernize them with traditional technologies.

But tomorrow, companies may combine human and machine capabilities to perform the same processes more efficiently. To make this possible, we need to use software solutions capable of replicating human capabilities for symbolic information processing. And that is one of the goals of artificial intelligence.

Julien and his team identify three types of solutions that could meet Sophie's needs. The first type uses artificial intelligence learning technologies, such as machine learning (ML), Microsoft Azure, IBM Watson, and AWS Comprehend. The second type of solution, which is also learning-based, uses language models such as ChatGPT, Grok, Mistral, Gemini, or CoPilot. Finally, there is a third type of solution that is knowledge-based, such as Aeteos' Percipion™ secure Symbolic Cognitive AI.

All of these solutions process or understand natural language and offer a wide range of features to meet Sophie's needs. But what is the cost of using them? Julien decides to estimate this. First, he will evaluate the cost of work done by the human resources team.

Labor cost estimate

Sophie works in the insurance industry. The average size of a large company in this sector is 11,000 employees, but Sophie's company only has 9,000 employees. Therefore, she will send her questionnaire to all of them, but only 70% respond to it. She and her team will have to process a total of 6,300 responses. In each questionnaire, each employee answered 8 questions. Sophie and her team will have to read each response and identify all possible semantics related to two affective states (negative and positive), eight primary emotions, twenty four secondary emotions, and 17 factors of job satisfaction and dissatisfaction, for a total of 51 different semantic fields. Then, they will report on any detected factors. As it takes 10 seconds to detect one factor in a response, Sophie's team will have a workload equivalent to 2.56 FTEs (full-time equivalents) per study, or 5.12 FTEs in total, as Bertrand has requested that this survey be conducted twice a year.

As Director of Human Resources, Sophie has many other responsibilities, which is why she only participates in this project one day per month. She has entrusted this analysis work to junior human resources experts who are supervised by a senior expert working part-time. As the company is based in Niort, not a large city as Lyon, Marseille, or Paris, the premises and remuneration of its teams are more reasonable. The total annual salary (including bonuses, health insurance, meals, PC, software, smartphone, and transportation) for a junior HR profile in Niort has been estimated at €65,425, that of a senior profile at €87,753, and that of Sophie at €125,307. In total, the cost of this analysis activity has been estimated at €385,880.

Modernization estimate

Julien will then compare the cost of Sophie's team's work to what it would cost him to use one of these software solutions that enable natural language understanding.

First, there is the cost of the license.

Whether using natural language processing solutions such as those from IBM, AWS, Google, or Microsoft, or a language model such as ChatGPT, the monthly cost per person is around €30. In some cases, depending on the publisher, additional software components or services must be added (storage services, model creation software workshop, database, etc.), and each user involved in the project must have access. For this study, we estimate that this cost represents €800 per month, all-inclusive. As for Percipion™, its monthly cost is €2,800, all-inclusive (unlimited number of users, unlimited use).

Then there is the cost per use.

Whether they are natural language processing solutions or language models, they all work on the same principle: pay-per-use. The more characters there are to process, the more you have to pay. The more users there are, the higher the license cost will be. You pay for the data you send and for the amount of data you receive.

In this specific case, what exactly are we dealing with? We are dealing with short answers to eight questions. The answer size is a few words, similar to a text message or a tweet. Let's say the response size is 213 characters. So we have to process 6,300 questionnaires, containing 8 responses, or 50,400 files of approximately 213 characters (a total of 10.24 MB of data) per study, or 100,800 files per year (20.48 MB of data).

These software solutions often charge per-character package, i.e., 100, 1,000, or 10,000 characters, costing a few cents or less. Suffice to say that given the volume to be processed, this cost will be negligible, probably a few dozen euros. For example, if we used ChatGPT model GPT4(8K), the token price would be €0.03, and we would need to use a total of 1,154 tokens. This represents a cost of €34.62. Another example would be the use of AWS Comprehend. In this case, we would need to analyze 107,352 tokens, and each token would cost €0.0001. In this case, we have a total cost of €10.73. Again, these estimates are for a single study, and to get the annual cost, we would need to multiply prices by two. As these solutions will be used both during the project phase and during the operation or maintenance phase, the cost of this consumption related to the

project or to the maintenance of the model should also be added to these estimates. All in all, we have estimated that this cost would be €200 per month.

As for Percipion™, its cost is not based on usage. This means that there can be an unlimited number of users processing an unlimited amount of data. This does not affect the final cost.

We are also dealing with a specific business use case. The software solution must correctly identify two affective states, eight primary emotions, 24 secondary emotions, and 17 factors of satisfaction and dissatisfaction at work, for a total of three reference frameworks containing 51 parent entities. The good news is that these technologies are learning technologies, so all we need to do is teach them these semantic contexts.

The first difficulty Julien's IT team will encounter is that they do not know all the words associated with this frame of reference, nor does Sophie. This list of words, correctly associated with each concept in these three frames of reference, must be provided.

A total of 1,045 words were found for job satisfaction and dissatisfaction factors. 5,163 words were found for emotions, and 4,190 were identified for positive and negative affective states.

Since these technologies are learning-based, we will need to provide these words to the machine. But it is not that simple. First, when collecting these words, we must ensure that they are of good quality and correctly classified. They must not contain spelling errors, must not be discriminatory or harmful, and must be consistent with the use case context. They must therefore be selected with care, which will take time.

Once this first phase is complete, several files must be created for each word to place it in a sentence with a grammatically positive form, a negative form (there are no fewer than seven negative forms in French), interrogative form, use its canonical form (without errors), then use variations of the same word, such as its masculine or feminine form, its plural form, its compound forms, its regional semantic variants, its slang or colloquial expressions, and finally represent it with common spelling mistakes or typos. Finally, if the word is a verb, it should be conjugated in common forms, and certain specific grammatical forms should be considered, such as its past participle. All of the words found must be placed in various sentences and contexts, each sentence must be represented by a file (or a specific entry), and this file must be correctly labeled and tagged. This approach provides a high-quality dataset for machine learning. Depending on the estimates provided and the desired quality of analysis, there may be between 20 and 50 or even 100 files (or entries) per word if they have multiple meanings

(polysemous). Finally, specific datasets will also need to be created to avoid unfortunate associations that would prevent the machine from correctly identifying the word. It should also be noted that it is possible to use computer tools or techniques to automatically generate certain variants, but not all of them, simplifying both the work to be done and the time spent building this repository.

This dataset must be prepared, fed into the machine, the result analyzed, the files corrected, and so on. Each iteration (learning cycle) will improve understanding quality until an acceptable result is obtained. Typically, it takes between 10 and 20 iterations or "epochs." These variations in the number of cycles are related to the pre-training of the model. The less trained it is, the greater the number of cycles.

Once complete, several cases must be tested, adjusted, and retrained if necessary, before deploying this solution and using it to help Sophie and her team.

In short, it is a real project!

Since Julien wants to work quickly, instead of generating 20, 50, or even 100 files for each word, he will only retain the most common variations and generate only 9 files (or specific entries) for each word. He therefore has a total of 93,582 files to generate, and each file generation, including labeling and all the necessary checks, is estimated to take between 1 and 3 minutes per file, depending on the difficulty of generating the desired variation, for a total of 285 days or 1.35 FTE.

Once this first phase of the project is complete, the model can be trained. Julien opts for 12 supervised training cycles, representing 238 days or 1.13 FTE.

To reduce the costs of his projects, Julien usually relies on a near-shore delivery center located in Tunisia. This destination is ideal for many reasons. Firstly, professionals in Tunisia are highly skilled, speak French, and labor costs are lower than in France. In addition, Tunisia is only 1 hour behind France, which makes it easy for his teams to work with these Tunisian service providers.

The junior data scientist will cost €320, the senior data scientist who will supervise them will cost €540, and the service provider project manager will cost €640. As this is a simple project, the project manager will be present one day per week. In Niort, a consultant will liaise with this team and will be accompanied by a project manager who will also be present one day a week.

In total, 653 days of services will be provided near-shore, 129 in Niort, and 31 will be entrusted to psychology experts to compile the list of words for these three frames of reference. In total, the cost of this project will be €378k (at an average

daily rate of €464). But Julien knows that if the target solution is a language model, it is more operational than traditional machine learning because it is pre-trained, and it will require less work. He estimates that he will only incur 40% of these costs and that the project will cost him a total of €151k. Today, many simpler, less pre-trained language models can be installed on a simple server and offer a slightly higher level of security. The use of these models would increase workload on the project side and has not been considered in this study on the profitability of artificial intelligence solutions.

Once this project is complete, this solution will be implemented to assist Sophie's team. However, new terms may appear, and it is also possible that the solution may have been poorly trained on certain terms, particularly those with double meanings, or on certain unidentified errors. To ensure this, Julien will need to mobilize data scientists at a delivery center in Tunisia. For the first six months, two engineers will be assigned to this project full-time. Then one engineer for the following year, and one engineer working part-time for the final year. This will ensure that any necessary adjustments can be made within a reasonable timeframe. This maintenance phase adds €169k to the project cost.

For information and completeness, companies such as OpenAI, Google Meta, DeepSeek, Mistral, xAI, Baidu, and Anthropic rely on service centers, often located in Africa and Southeast Asia, for annotation, correction, and improvement of AI model tasks. These centers play a key role in optimizing model performance. These service centers, sometimes called "data annotation centers," employ people to perform tasks essential to model development and maintenance. These tasks include labeling, classifying, or structuring data to train AI models. These individuals review model outputs to identify errors, biases, or problematic content (inappropriate or inaccurate responses) and correct them. This is called "reinforcement learning from human feedback" (RLHF). These companies are often located in countries where labor is less expensive (some employees are paid \$2 an hour), such as Kenya, Uganda, the Philippines, or India. In addition to the exploitation of this labor force, the processing of this data by these service providers also raises the issue of data sovereignty. To sum up, reinforcement learning based on human feedback is the only way to create a carefully annotated and corrected dataset that can be used to build high-quality models. What these companies do to tune their models must also be done by Julien's team, and this maintenance phase is not only essential but must be sustained over time.


What about Percipion™? Well, since Percipion™ has a "knowledge-based" approach, this work has already been done by the publisher, and the solution is, so to speak, ready to use. The job satisfaction reference framework is included in the subscription. There are no additional costs. We handle any updates to the frame of reference, no extra cost.

When estimating the costs associated with this project, we did not include transportation, accommodation, governance workload, or review these estimates based on the level of risk taken (near-shore component). Depending on the solution used, the result is provided as data that must still be exported to other software products to generate a report for the end user. This last step has not been quantified.

Now that all the costs are known, it's time to compare them and see how long it will take to see a return on investment (ROI).

Return on investment

Over two years, if this activity is handled by Sophie's team, the labor cost is €772k. If this activity is entrusted to learning artificial intelligence, its cost over two years would be €409k for a machine learning solution and €372k for a solution using language models. If this solution used a knowledge-based approach, such as Percipion™, the cost would be €67k.



The first conclusion we reach is that, in all cases, modernizing this activity through artificial intelligence will enable this insurer to reduce its operating costs.

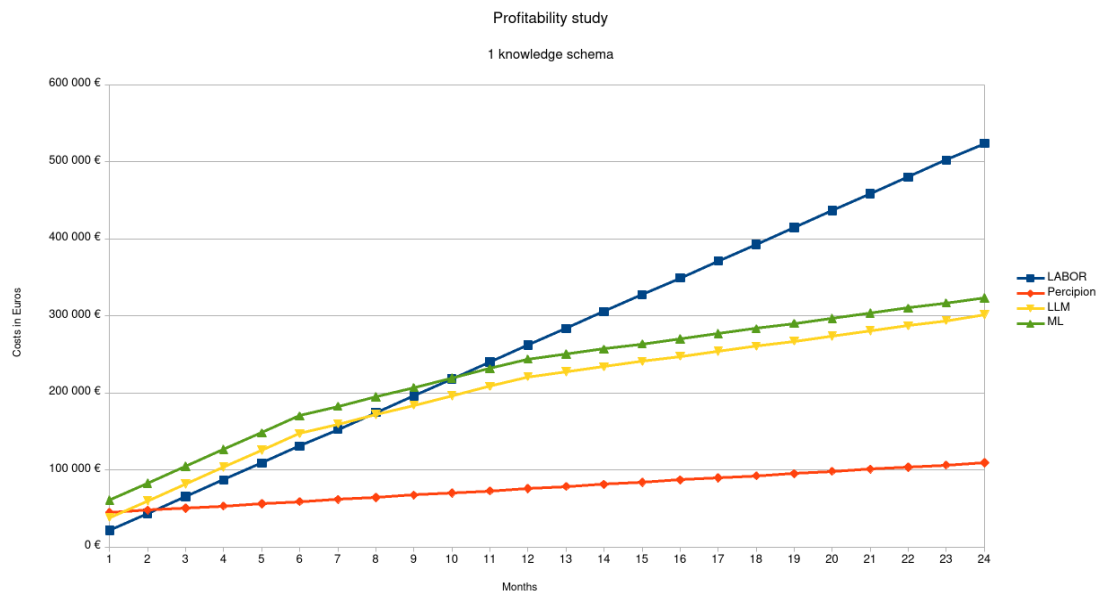
The second conclusion is that knowledge-based technologies have lower operating costs.

All of these solutions offer real benefits. But how long does it take for them to become profitable?

When evaluating the return on investment (ROI) of these solutions, machine learning technologies offer benefits from the 9th month onwards, language models from the 7th month onwards, and knowledge-based technologies offer immediate benefits.

This project is complex because it requires the creation of three specific reference frameworks. Julien therefore also asked himself what would happen if there were only one specific reference framework to create and, in this specific case, this new reference framework were not available on the knowledge-based technology side. Here, too, there would need to be a project phase, and these technologies could then be better compared.

The most significant workload for this type of project remains the creation of qualified data repository that presents each word in multiple ways in various contexts, as well as training the solution.



In terms of knowledge-based technologies, particularly Percipion™, some of this work is unnecessary because Percipion™ not only has an innovative data structure that represents data in the machine's memory as it would be represented in human long-term memory, but also a processing unit that can assess whether two words are similar or not, which significantly reduces the correct identification of the number of variations that this word could have in relation to its canonical form. For this evaluation, if we were to create a reference framework with the same complexity as the reference framework containing the factors of job satisfaction and dissatisfaction, the creation of this reference framework containing 1,045 objects would take one month at a projected cost of €20k. As this reference framework would be specific to this project, the cost of its use by Percipion™ would be zero.

For this second case, over a two years period, if this activity were to be carried out by Sophie's team, the cost of the work would be €524k. If this activity is entrusted to a learning artificial intelligence system, the cost over two years would be €323k for a machine learning solution and €301k for a solution using language models. If this solution used a knowledge-based approach, such as Percipion™, the cost would be €109k.

When evaluating the return on investment (ROI) of these solutions in the second case, machine learning technologies show gains within 10 months, technologies using language models show gains from the 8th month onwards, and those based on knowledge offer gains from the 2nd month onwards.

It is crucial to note that the high variability of learning cycles depending on the learning technologies used (machine learning versus language model) and the workload associated with these activities have not been taken into account in this article. According to Gartner, "Value generation is slow (more than two years)." The use of such solutions would be profitable after 24 months. Gartner adds that "by 2028, more than 50% of companies that have developed large-scale AI models from scratch will abandon their efforts due to the costs, complexity, and technical liabilities associated with their deployments." Profitable, yes, but not right away. Sustainable over time? Not so sure. Do they generate value? According to a recent MIT study (July 2025), it seems not: "Despite \$30-40 billion in corporate investment in generative AI (GenAI), this report reveals a surprising finding: 95% of organizations are not seeing a return on their investment."

Conclusion

To conclude, modernizing this activity through artificial intelligence technologies offers many benefits to the company in all cases. Technologies such as Percipion™, which are knowledge-based, are easier to implement and generate productivity and financial gains quickly, even immediately, than technologies based on machine learning.

In addition to these financial criteria, the choice of a software solution to modernize activities such as those of Sophie's team is also linked to other constraints.

For example, these solutions must comply with the European AI Act and be legally viable for production. Great care must be taken with the data that will be fed into learning solutions to ensure that these solutions do not discriminate or process incorrect information. As these technologies are "open," incorrect data input can quickly skew or invalidate the results of the processing. With this type of

technology, how can we guarantee that the results of an analysis performed yesterday will be identical to those obtained today? As the reference framework is constantly evolving and the statistical relationships are changing, it is impossible to guarantee this reproducibility. With Percipion™, this is not possible because the reference frameworks are highly qualified and protected from any data input.

These learning solutions are software solutions produced by a publisher who creates their own source code and relies either on other components or libraries that they have created, or on third-party components or libraries that do not belong to them. As these solutions are artificial intelligence solutions, these publishers must also have control over the origin of these third-party components and be able to identify them all. When using high-level languages, particularly scripting languages, a large number of libraries are required, and this requirement quickly becomes very complex. It is not the case with Percipion™. It has been entirely programmed in C and only uses standard libraries (C99 ANSI).

The vast majority of machine learning artificial intelligence solutions are black boxes. In most cases, the publisher is unable to produce a process that is both transparent and capable of explaining the semantic relationships between frames of reference concepts to provide the user with verifiable results. How can you trust this type of solution? Percipion™ allows all processing steps to be displayed during execution to offer a high level of transparency.

In the example given in this article, we focused primarily on the work that needed to be done to inject a new reference frame into a machine, taking into account only simple variations in the canonical form of a word. However, this very conservative approach does not allow us to account for the richness of natural language expression or the substitution of a word with an iconic representation (emojis). The use of such an iconic representation must also be interpreted and correctly associated with its parent concept. In addition, some sentences use words that introduce valences. These modify the meaning of the word and must also be taken into account. Taking this into account will significantly increase the amount of data that needs to be fed into the machine and will therefore significantly increase the number of possible variations of the same word. To clarify this point, there are 38 words in the French language that carry a positive semantic valence and 49 words that carry a negative semantic valence. In addition to these semantic valences, there are also psychological valences linked to our positive and negative affective states (for a total of 836 words) and those linked to our primary emotions (for a total of 729 words). Therefore, if these variations are not included in the data corpus that the machine must learn during training, the consideration of these valences will be compromised, and the machine will not be able to process these cases and will produce false analyses. Percipion™ not only transcribes the emojis present in a verbatim transcript into the corresponding word so that it can be taken into account by the language

comprehension solution, but it is also capable of processing semantic and psychological valences, giving it a much more nuanced understanding of natural language.

You should also ask yourself questions about the hosting of your natural language understanding solution, network usage, and any data that may be transferred from your servers to third-party servers. Is your data truly secure? Is it shared with other solutions? Is it being captured? How can you trust it? In our example, Sophie needs to analyze sensitive data that must be processed anonymously and securely. Can these solutions guarantee that? With Percipion™, the answer is simple. Percipion™ can be installed wherever you want, and this software solution does not use the network. With zero network activity, the processing performed by Percipion™ can be done offline and in real time.

Before concluding, let's also address the inability of many natural language processing solutions to handle large volumes of verbatim text. Based on APIs, these solutions only accept texts of a limited size. But in some cases, what can be done when it is necessary to analyze an entire document? How can we perform a semantic analysis of a text as a whole if, due to technical constraints, it is necessary to split the text into a set of files? In these cases, the analysis becomes more complex and requires the consolidation of the different results of these different processes at the end of the treatment. This adds difficulty. Percipion™ is capable of processing documents of any size in a single pass.

Finally, unlike learning artificial intelligence software solutions that consume a significant amount of energy, Percipion™ consumes only 2.37Wh per CPU core, or, by analogy, less than 1 centiliter of water per CPU core. Here again, Percipion™ is ahead of its competition and complies with European Union objectives.

Julien has identified technologies that can help Sophie in her work. The activity of reading and identifying specific semantics related to affective states, emotions, or factors related to job satisfaction or dissatisfaction can indeed be identified by a software solution, and the use of these solutions offers many benefits. A knowledge-based approach, such as that offered by Percipion™, is the least expensive and offers the fastest return on investment. In addition, numerous functional, regulatory, and security advantages make Percipion™ a relevant choice for modernizing many use cases within the company.

First humans, then machines

After eight years of research and development in neuroscience, cognitive and clinical psychology, linguistics, philosophy, and computer science, we created Percipion™, a Symbolic Cognitive AI platform that is explainable, secure, trustworthy and sustainable capable of assisting humans in their work by transforming textual data into strategic insights for informed and rapid decision-making.

It is as if on this platform you would have assistants capable of reading text, understanding it, and grasping its nuances, identifying words and associating them with concepts, extracting meaning from text through precise contextual analysis, anticipating, reasoning, and drawing conclusions, producing reports, communicating, and taking action.

With this new technology, Aeteos™ helps companies improve their decision-making and operational efficiency by enabling them to build their own secure Symbolic Cognitive AI platform and modernize numerous use cases.

Percipion™ enables daily analysis of data, whether internal or external to your organization. It can also be used to monitor the digital space and help collect essential information, proactively identify risks, and issue alerts. By providing accurate contextual analysis, Percipion™ helps you make informed decisions while giving users complete control.

Regardless of your field of activity, Percipion™ will speed up information analysis and facilitate decision-making in more than 40 use cases.

With its own dataset, it will complement the teams' diagnostics. Everything remains under their control.

It will provide expertise where no expert is available by supplying information on sensitive data, freeing up your teams to focus on higher-value-added activities.

If you would like to learn more, please don't hesitate to contact us!