Solve business problems with AI

Each business problem asks for a specific method within Artificial Intelligence

May 2021

MAKING AN IMPACT THAT MATTERS since 1845

With deep learning techniques headlining today's news and commercial applications being powered by ever more complex models, organizations may be tempted to try to solve their use cases with state of the art AI models. However, whether you should use such complex methods or are likely to benefit more from simpler approaches depends on a variety of factors. So what method do you pick to REASON with AI?

By Sjors Broersen, Daniël Rennings & Naser Bakhshi

The Deloitte AI Loop (DAIL)

As humans, we sense our environment, we reason about it, use conclusions to take action and augment our interactions with our environment whilst continuously learning to improve this process over time. Likewise, the Deloitte AI Loop provides a framework that resembles this human approach in the space of artificial intelligence. Based on our experience in bringing cognitive solutions to our clients, we have lined out DAIL as a blueprint for all aspects that should be covered in a successful AI solution, as we explained in the introductory blog. Following a deep-dive on SENSE in our previous blog, this blog focuses on the reason component, consisting of algorithms based on (tacit) knowledge and patterns in data in order to reason.

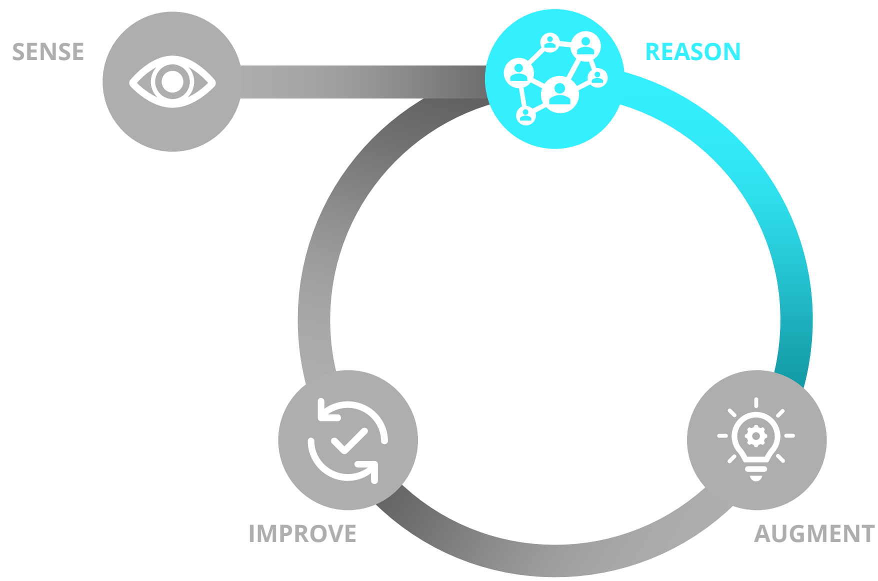

Algorithms based on (tacit) knowledge and patterns in data in order to reason and "understand"

May use **auxiliaries** like:
Pre-trained Resources
Ontologies

REASON as part of the complete DAIL

REASON in the age of AI – background

Reasoning might be the most elusive human skill for computers or artificial intelligence to learn. The ability to take into account different pieces of information, weigh them against each other, envisioning different scenarios and finally coming to a grounded decision is something that seems like a typical human skill. However, it very well sets us apart from computers.

Still, with new advancements in deep learning, AI is getting more intelligent and capable of making better decision based on more types of data. Recent examples include the language generating capabilities of GPT-3, the image generating capabilities of DALL-E or the protein folding solved by DeepMind. What these examples all have in common is that they take enormous amounts of data and through meticulous and lengthy training on vast computing resources, they learn to do something that makes sense. In doing so, they are able to generate tremendous (potential) value for business and society.

Due to the scientific breakthroughs and the ever-growing commercial applications of these advanced techniques over the last-decade, these deep learning techniques have become top of mind. However, if you want to obtain value from AI for your organization, this is not a silver bullet. Rather, you should follow a structured process, before you select your tool to reason.

How to REASON – in theory

Before you can start to REASON, there are two preconditions: data and knowledge. As discussed in our previous blog, having a proper data foundation in place is a strict precondition for enabling reasoning. In addition, knowledge in the form of domain expertise is required to properly reason about data and its context. Such knowledge can be located in the minds of business liaisons, but can also be expressed in machine readable formats, such as ontologies, embeddings or knowledge graphs. If data and knowledge are available, we can effectively start a reasoning process.

As displayed by the toy examples below, reasoning can come in different flavors: based on deductive, inductive and abductive logic.

Deduction	Induction	Abduction
All men are mortal	Socrates is a man	All men are mortal
Socrates is a man	Socrates is mortal	Socrates is mortal
Socrates is mortal	All men are mortal	Socrates is a man

Examples adapted from Bateson's Method: Double Description. What is It? How Does It Work? What Do We Learn?

Is AI really intelligent?
Even though the above mentioned capabilities are mind-blowing and better than humans could perform on these tasks, not many people would call the algorithms actually intelligent. Gary Marcus argues, for example, that GPT-3 is not grounded in reality, and does not actually think. While this may be true, for now, it does not need to stop organizations from generating value from them.

In deductive reasoning, one can draw precise conclusions from a given situation. As you can imagine, AI can easily outperform humans in crisp logic (think of computers crushing humans in a variety of board and digital games). However, in real-world applications, there is no generic rulebook that contains a section on "When does a player win". Instead, humans define what constitutes to a win for them, for instance via concrete KPI's. For example, in assortment optimization, one could identify KPI's based on a product margin and customer reach. In addition, one could define rules-of-thumb that should be considered on these KPI's, leading to so-called heuristics, such as "the margin of a product should be positive in each quarter".

With inductive logic, this changes as it considers reasoning by induction from a set of examples, which comes very close to the definition of machine learning. Based on a training set that should represent all possible cases as good as possible, a machine can learn to predict the outcome for new examples by induction. It does so by looking at specific features, which are defined by humans based on their domain knowledge, hereby they try to ensure computers have a proper overview of relevant parts of a real-world.

In abduction, we take this one step further, where we not only need to infer conclusions on new examples based on given examples, bust also have to come up with hypotheses. While this is beyond the scope of machine learning, it can be found in deep learning. For example, in the [deep question answering techniques that made IBM Watson win Jeopardy](). Now that we have familiarized ourselves with some background on reasoning, it is time to look at how to apply this theory in practice.

How to REASON – in practice

Having connected the different types of reasoning to different techniques of AI, we will now look at each of these techniques and when to apply them in organizations, following the framework below.

	1 HEURISTICS	2 MACHINE LEARNING	3 DEEP LEARNING
	Digitize explicit rules of thumb obtained from domain experts	*Assemble possibly relevant features and let ML find relationships to the value to predict*	*Assemble examples and let DL learn features itself through representation learning*
When?	• Analytics capabilities are not in place • Training data cannot be obtained	• Analytics capabilities are not advanced yet • Limited computing power and training data are available • Explainability of predictions is business critical	• Advanced analytics capabilities are available • Sufficient computing power and training data is available • Tailored DL models are readily available • Outperforming competition is business critical
Why?	• Establish a focus on concretely expressed rules-of-thumb and hypotheses • Validate hypotheses at scale • Establish a baseline model and move calculations to the digital space	• Identify patterns that humans may miss • Benefit from available data • Build trust in AI through easily-explainable models	• Identify patterns that humans can not see • Utilize data to it's full potential • Distant yourself from your competitors
Examples	• Assortment optimization based on KPI's, such as margin and penetration, while accounting for seasonality and locality • Extracting medical values from free text fields through regular expressions	• Forecast demand in various parts of an organization, such as incoming calls and appointments • Predict the language level of a speaker, based on voice input	• Track animals to identify well-being • Spot your brand in commercials • Identify relevant information from a vast corpus of documents

When and why to use different tools to REASON, distinguishing heuristics, machine learning (ML) and deep learning (DL)

Heuristics

Whilst employing heuristics does not sound sexy, heuristics are a viable option for solving certain business problems and should not be overlooked. These rules of thumb that reside in the heads of domain experts that are (potentially unconsciously) using them to make business decisions. By moving these ideas into a digital formula, you are empowered with concrete KPI's and you can validate to what extend rules-of-thumb (don't) hold to identify opportunities for improvement. This is also a valid strategy if you simply don't have any training data available and are unable to create such training data at a reasonable scale, or if your analytical capabilities are in a starting stage. While this may not bear the name Artificial Intelligence, it does bring the organization closer to it, as they can in identifying the first few data sources and features to take into account, whilst you also have set out a first digital footprint and overcome the challenges that come with it.

Client case

Retail

For a retailer we kickstarted their journey in analytics by digitizing the process of assortment optimization. Previously, they rationalized their assortment once a year, based on the margin gained on a product over the past year. By adopting a digital approach, we could move beyond an annual margin and provide calculations on a daily base, while also accounting for temporal and spatial patterns. As a result, we could confidently rationalize 10% of the assortment, but also identify products that where worth a wider distribution, leading to a potential annual revenue increase of over a million euros, based on nothing more than heuristics.

Machine Learning

Moving beyond rules of thumb, machine learning can provide you with a large improvement in predictive performance by identifying patterns that humans may miss. Humans, and especially experts in a field, can raise hypotheses and input features that may potentially be relevant, but no longer have to define their explicit relation to the value to predict. In addition, these relationships can be identified in white-box approaches or via additional tooling in some black-box approaches. Hence, this tool can also help you build trust in AI in your organization. As the requirements in analytical capabilities, training data and computing power are limited, ML can typically be applied to a wide range of problems in a variety of domains. Think about forecasting the demand for your products and services, detecting anomalies in your logistic processes, or doing a customer segmentation based on clustering. For this you can apply an arsenal of machine learning methods, which each come with their own pro's and con's. For example, you could use a Random Forest to benefit from less requirements on feature preprocessing or you could use a linear or logistic regression to benefit from clear explainability of your model.

Client case

Public

For a public organization we researched the possibility of predicting a language level of a speaker to personalize a voicebot conversation. While we all know that the uniqueness of words you use and the pace at which you speak may hint at your language level, we don't know their individual explicit impact and interaction effects. Additionally, due to the innovative nature and limited availability of training data, it made sense to adopt approaches that allow for clear explanations. Using a ML approach, we could accurately predict the language level within a certain range, demonstrating the feasibility of applying this in production.

Deep Learning

While machine learning is likely to provide the biggest performance increase over the heuristic baseline, deep learning may still provide a lot of additional value in specific cases. Cases where either a lot of data is available or pre-trained models/services exist that can be readily used. Especially in areas where typical human interpretation is needed, e.g. understanding what's in an image or determine the content of a text, deep learning shines. Deep learning excels in such cases, by learning a specific representation, i.e. we no longer have to define the features (as we did in ML), but we simply provide a large set of labeled examples and set some boundary conditions for the model. For some applications, a state of the art DL model is readily available - think of pre-trained models for coloring images or text-to-speech [footnote, Interested in concrete examples? Have a look at https://modelzoo.co/] that can readily be used when provided with proper input. Beyond that, DL is likely also valuable to adopt when outperforming your competition is critical, with obvious examples in the domain of trading. However, to use this tool successfully, it is required to have a vast amount of data and computing resources available, as well mature analytics capabilities.

Solving hurdles of AI with human reasoning

While the above provides guidance on which method to pick, you may still run into hurdles when applying AI, whether it is the limited availability of data or performance limitations of your model, there are still means to achieve success.

One of the main hurdles in AI is the limited scale of available training data. There is always a balance between the accuracy needed in a model to be successful and the data that is available for training and testing this model - while different methods require different scales of data. If too little data is available, it may be worth researching means to establish training data yourself. Can we devote resources to label examples or can we use the crowd to label data for us? Investing in generating training data can help to obtain an edge over your competition, as they would also have to take this hurdle to catch up with you, underlining the fact that a good domain specific dataset has actual business value.

In other cases, you may face a hard limit on the performance of your model: the model alone might not be good enough to practically solve your business problem. We here want to emphasize that there are different levels of automation, as displayed below: with initial models you may already be able to provide assistance to humans, without requiring full automation. Super-human performance may still be achieved if we utilize humans-in-the-loop to go over predictions on which your model may be bound to unsatisfactory performance. This enables your organization to focus on achieving a decent performance on the majority of predictions, whilst you and let AI and humans both do what they do best. Because after all, humans and computers are not the same.

Client case

Industrials

Deloitte teamed up with a major meat producer, animal welfare organizations and a major meat purchaser to create an AI solution that can automatically detect and label potential animal welfare deviations. Based on a deep learning algorithm that scans CCTV footage for deviations, QA employees can focus their efforts on reviewing a pre-selected fraction of the many hours of generated video footage. Since a pre-trained model was available, this solution could be realized without having to build the algorithm from scratch, but adjust an existing algorithm through fine tuning.

What to expect next?

This blog is part of a series in which we deep dive in the different components of DAIL and describe them in a more in-depth fashion. Next up, we will discuss AUGMENT, which will bring us from the models we employ to the value they bring to organizations.

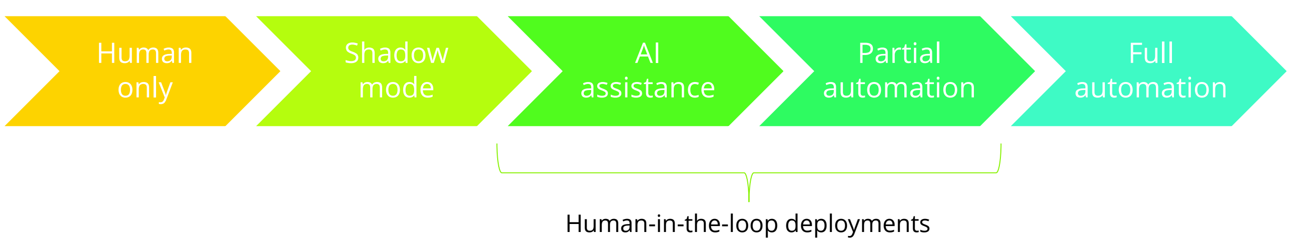

Image adapted from DeepLearning.AI: AI-enabled automation is often portrayed a binary "on or off", while in practice, automation is a spectrum and AI teams have to choose where on this spectrum to operate.

About the authors

Naser Bakhshi
Naser is a Director in the Deloitte Dutch Analytics service line. In the world of ever-growing data and continues new technological advances, his goal is to bring to bring the best, the newest and most innovative analytics solutions and unlock actionable insights for a fact-based decision making. Naser combines enabling technologies such as Machine Learning, Cognitive Analytics and intelligent self-learning systems to develop and embed advanced decisions support systems in the supply chain and in customer domain. Moreover, he's an experienced transition manager with a strong feeling for organizational change and sensitivity helping organizations to define analytics strategy and build analytical capabilities.

Sjors Broersen
Sjors is a Specialist Lead in the Dutch Deloitte Artifical Intelligence team. Sjors specializes in applications of AI, mainly in the fields of NLP and computer vision, driving innovation with the goal to bring state of the art solutions to clients. Sjors has over 5 years of experience in applying AI and analytics at clients in the Consumer and Energy and Resources industries. Before joining Deloitte, Sjors completed a PhD in astronomy.

Daniël Rennings
Daniël is a Consultant in the AI team within Deloitte NL. He has worked as a data scientist across various industries and has a background in Computer Science and Data Science. Daniël focuses on the delivery of NLP solutions while also educating clients on how to apply AI techniques within their organization to realize their opportunities in the world of analytics. Before joining Deloitte, Daan graduated on diagnosing deep neural networks in IR and was a visiting student at MIT, researching the application of IoT in urban environments.

References

Marcus, G., & Davis, E. (2020). GPT-3, bloviator: OpenAI's language generator has no idea what it's talking about. *Technology Review*.

Mindmatters.ai (2019). A type of reasoning AI can't replace.

Hui, J., Cashman, T., & Deacon, T. (2008). Bateson's Method: Double Description. What is It? How Does It Work? What Do We Learn?. *In A legacy for living systems* (pp. 77-92). Springer, Dordrecht.

Magazine, A. I. (2010). The ai behind watson—the technical article. *AI Magazine*.

Ng, A. (2021). The Batch - Issue 80, Deeplearning. AI.

Deloitte.