AI in Practice series

10 commandments

to create value with artificial intelligence

Insights from the field

About this publication

Welcome to our first booklet in the AI in Practice series. We aim to provide you with insights and actionable steps for establishing human(ity)-centered AI initiatives that demonstrate true business value.

This series is written for business leaders, product people, engineers, researchers, designers, and everyone else interested in applying AI responsibly, to create value for business, people and society.

For our inaugural edition, we compiled the ground rules for establishing successful AI initiatives in your organisation, based on the experience of AI leaders in The Netherlands. Through an evaluation of their experiences and challenges, we hope to provide you with practical considerations to apply artificial intelligence in your business.

How this booklet was made

This booklet was written based on interviews with 30+ AI business and technology leaders in the Netherlands. Some of these are featured in the deep dive case studies, other interview insights are summarised in the findings and recommended practices.

Want to contribute?

Our AI in Practice research initiative is ongoing, and we're looking forward to provide you with more insights in the near future. Want to contribute to our research? Don't hold back, and send us a message at **human@deus.ai**. We'd love to hear from you.

Introduction

Creating value with artificial intelligence for business, people & society

It is almost impossible to overstate the impact that the applications of artificial intelligence will have on our futures. In the words of Yuval Noah Harari, **"...AI is changing the world more than anything in the history of humankind. More than electricity."**

AI is on the radar of nearly every organisation and government, and it is prominently featured in news and literature. There is great excitement over the benefits of AI applications, but there is also a lot of noise, and this can make it difficult to separate real potential from hype. How can we distinguish practical reality from speculative science fiction?

As a society we're at a critical point: not only can AI play an important role for organisations, it can also be impactful on a political level, aiding humanitarian efforts, or combatting the climate crisis. Machine learning can help to significantly reduce all forms of waste, to improve health and safety issues, remove inefficiencies, optimise time usage, increase manufacturing accuracies and minimise carbon footprints to name just a few examples.

Beyond the speculations of what the future may bring, what are the practical considerations for establishing AI initiatives right now? As an industry, we are in the process of establishing best practices of how to deal with those considerations.

In this publication you will find our 10 Commandments (or in less 'biblical' terms - 10 important insights from the field), supported by practical case studies, to assist you with establishing or expanding your AI initiatives. We hope they help you create value for your business, and more importantly, people and society.

In this booklet

Our 10 commandments for establishing artificial intelligence initiatives

1. **Establish building blocks**
   Your future organisation needs a foundation
2. **Connect your data sets**
   Continuously improve while projects are ongoing
3. **Bridge AI and business**
   This is the era of the 'Analytics Translator'
4. **Decentralise!**
   Establish an interdisciplinary environment
5. **Centralise!**
   Create common practices and training
6. **Test & learn**
   Implement processes for rapid validation
7. **Create ethical AI**
   Anticipate the effect beyond the algorithm
8. **Explain AI reasoning**
   Provide transparency in decision making
9. **Be better together**
   Partner up to stay ahead
10. **Design your own future**
    Create value for society, people & business

01 Establish building blocks.

Your future organisation needs a foundation.

The challenge

There is a debate among industry leaders about the best way to get started with AI in an organisation. Does an organisation require a ‘top-down’ AI strategy that drives all initiatives? Or is it better to allow AI initiatives to arise organically, bottom-up, as a solution to business problems, without a specific strategy in place?

The response

Whether you have a centralised AI strategy or a bottom-up approach, to succeed with AI, the necessary foundation for AI initiatives should be established.

Opinions on the need for defining an AI strategy strongly vary. Some argue that it is a needed foundation for success, on the other end is argued that AI should only be considered as a potential solution to a specific business problem, and therefore should not require a separate strategy.

Regardless of the debate, one thing is clear. To get started and advance with AI, organisations should focus on establishing the building blocks that are required to successfully establish and scale AI initiatives.

A first building block could be value drivers: in which business area or activity can AI add most value? Secondly, processes: defining ways of working. Finally, a third and fourth block could cover establishing the needed data and technology infrastructure, and the resources needed to succeed.

Quote

"Artificial Intelligence is the glue between business and technology. Every business should build up an AI strategy the same way as they build a business and technology strategy."

Daniel Gebler, Chief Technology Officer, Picnic

Visualisation

Building blocks to consider when establishing AI initiatives

Value	**Define value** In which areas could AI make the most significant impact?		
Processes	**Team processes** How should teams work together on AI use cases?	**Organisational processes** How should teams collaborate with the rest of the organisation on AI use cases?	
Technology & Data	**Technology** What are the technology and infrastructure requirements?	**Data** What data is required, is it available and accessible?	
People	**Skills & People** What are the necessary skills? Will you hire, train or both?	**Partners** Which external parties could contribute to your efforts?	**Stakeholders** Which stakeholders are affected or should be involved?

Case study: ABN AMRO

Defining the strategy of an AI department

Dennis de Reus, Head of Artificial Intelligence at ABN AMRO, explains the building blocks that shape the strategy of his AI department.

The AI department develops use cases, based on the input of the business teams in different departments of the bank, as well as conducting research in specific AI related challenges, such as Explainable AI. This means that AI specialists in the bank function as subject matter experts working with business stakeholders on AI projects, as well as in the AI department to develop overall AI expertise.

At ABN AMRO, it is the AI department's mission to help the bank become better and more efficient, by utilising and creating value with artificial intelligence.

Building blocks of ABN AMRO's AI strategy in Innovation

Team	Defining how to work inside the team, together with the business, and hiring team members with broad profiles to successfully bridge business and technology.
Applications	Researching relevant topics, publications or new technological developments.
Use cases	Defining how to develop use cases within Group Innovation and the bank.
Capabilities	Enabling internal stakeholders to direct AI initiatives, and helping data specialists to gain a better business understanding.
Training	Ensuring that everyone handling data within the bank is trained in the same way.
Partnerships	Partnering with coalitions, academia, and potentially also with other banks on certain matters.

Consider how to...

Establish the building blocks to drive AI initiatives in your organisation.

Focus on the enabling capabilities

Regardless of whether you choose to define an AI strategy, evaluate which capabilities and infrastructure you currently have in place, and which need to be enhanced in order to establish AI initiatives.

A good starting point, is to evaluate your capabilities in terms of the needed: people and skills, data, technology, processes and partners.

Benchmark against leaders in your industry and AI leaders in other industries. Engage with specialists and peers and remember that Europe is behind in progress, when compared to the US and China.

Evaluate the organisational structure

Different organisational structures require different strategies for AI adoption. Most organisations go through an experimentation phase, in which they launch a few proof of concepts, before further defining the optimal operating model in the organisation.

Should AI become part of the data analytics group, should you establish a central AI department, or should AI expertise be decentralised across units? There is not one 'best' choice of operating model, however, dependent on your organisational structure, certain operating models will be a better fit.

Prioritise business value

Artificial intelligence initiatives should always be in line with the integral organisational strategy. The technologies in the field of artificial intelligence are a means to an end, and should never be considered independent from the business ecosystem. Not every problem is an AI problem: always consider with care what the best solution might entail.

Business strategy, challenges and opportunities should be driving AI initiatives. To ensure value, establish processes to evaluate the potential business value of new use cases up front, before investing significant resources.

02 Connect your data sets.

Continuous improvement while projects are ongoing

The challenge

Data is at the core of every AI project. Many organisations are facing data related challenges: transitioning to the cloud, data from different sources are not connected, and data governance, to name a few.

A common side effect of these challenges, is that data science and AI experts end up spending a significant amount of time accessing, gathering and cleaning data, instead of creating business value with data that is already available. This can undermine the perceived value of what AI can deliver, and limit the investment in an area that will be instrumental to the competitiveness and success of organisations.

The response

Establishing a data lake and cloud technologies can greatly benefit data science and AI efforts, but even during the required transition phase, organisations should aim to advance with AI projects.

Many companies have now gone through the process of establishing data lakes and moving towards cloud technologies. Valuable learnings are available when it comes to the implementation process.

Organisations that are initiating, or are in the middle of such a transitioning phase, should try and benefit from these learnings. They should not hold back efforts to make progress while the transition is taking place. Ongoing projects will help inform future data related decisions and help prioritise the data areas and AI initiatives that could bring most business value.

Quote

“Data is a big challenge: the access to data, data cleaning, and making it available across the organisation. In some projects, up to 80% of the time is spent on data preparation.”

Görkem Köseoğlu, Chief Analytics Officer, ING

Visualisation

Data science activities, proportionally

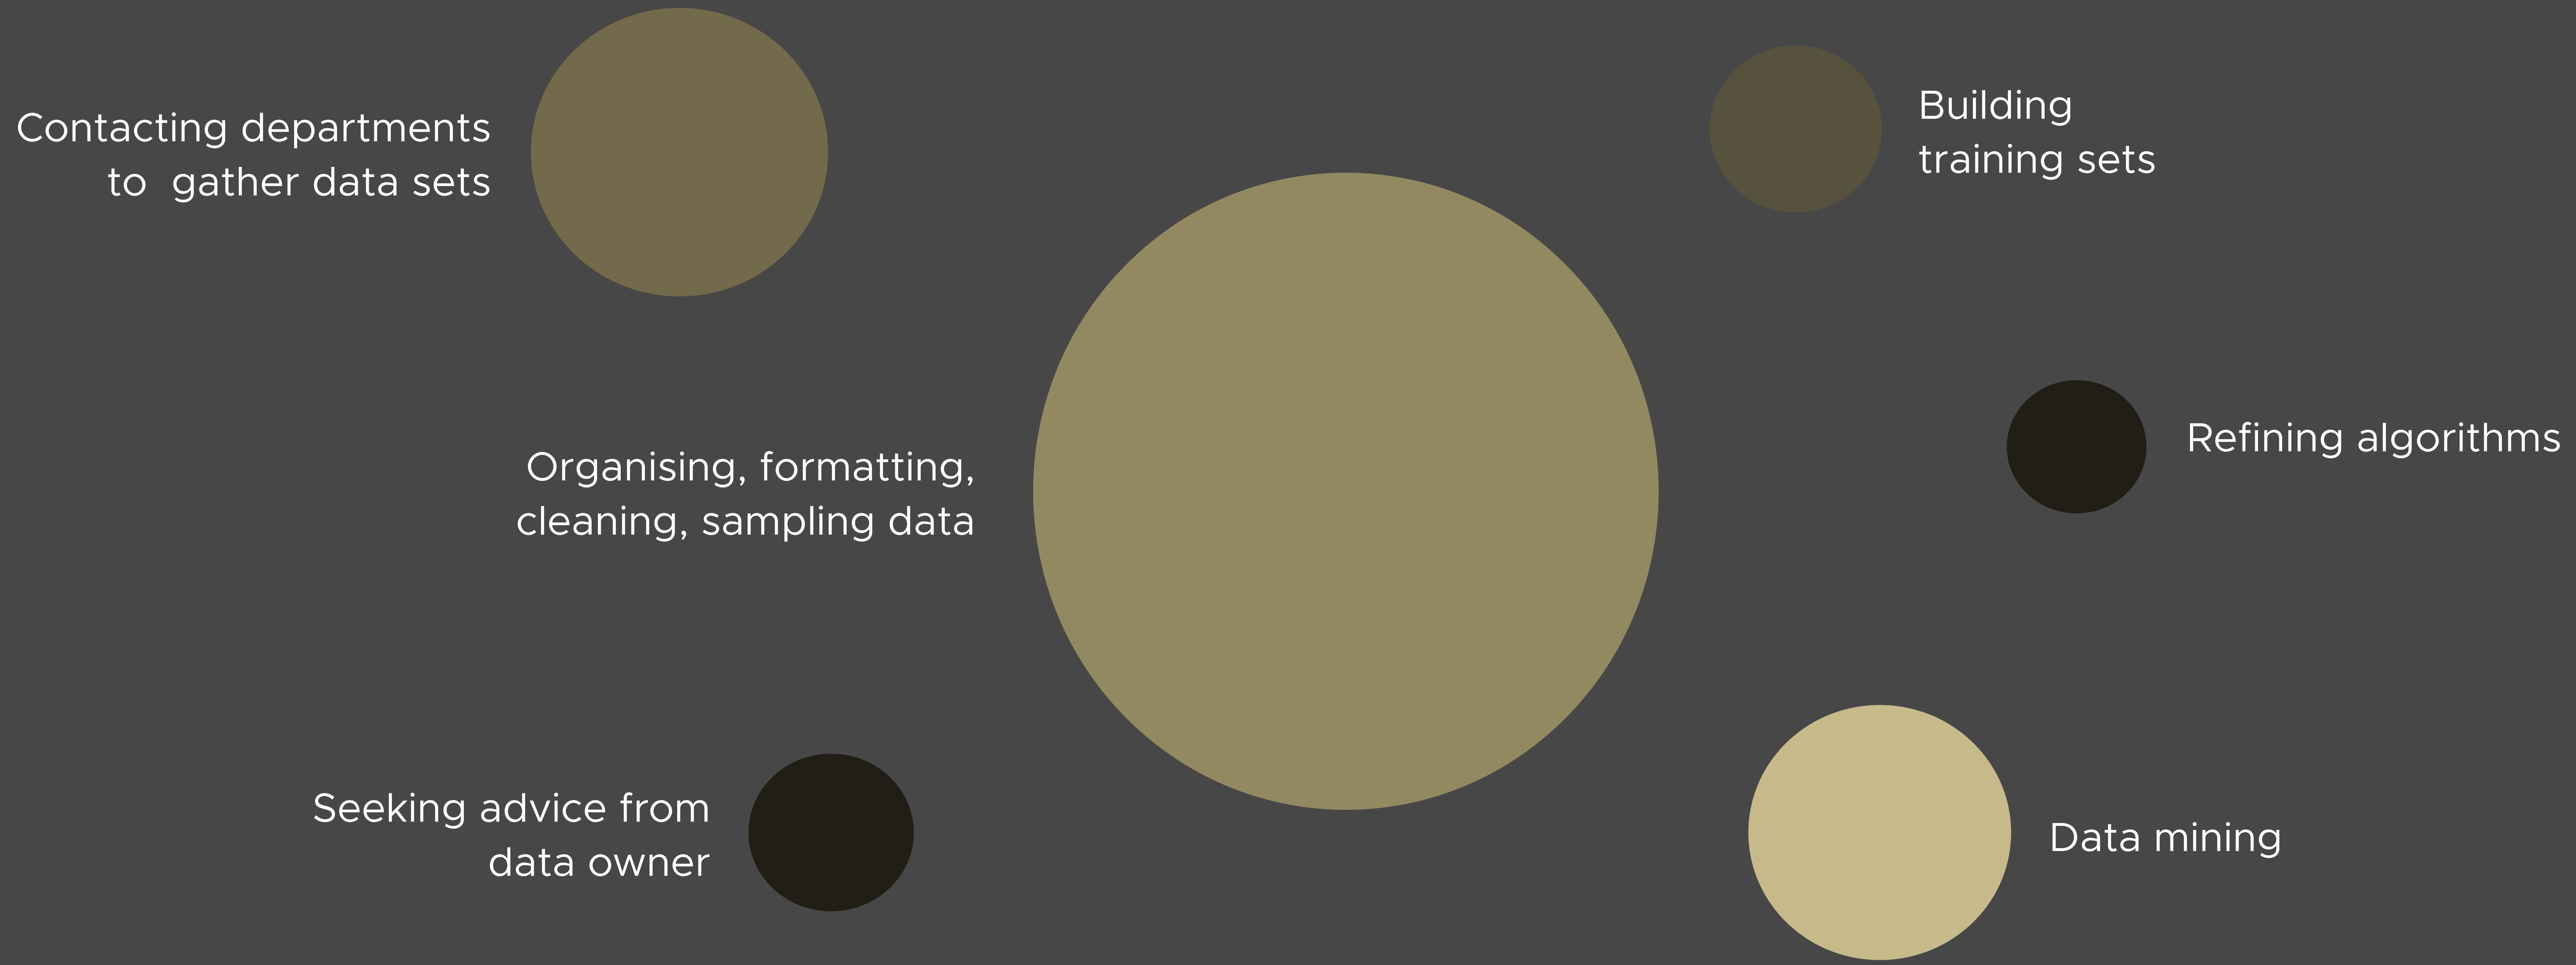

Case study: Arcadis

An advanced analytics transformation

Christoph Wollersheim, Global Director Data Analytics & Insights at Arcadis, explains their current efforts in their global advanced analytics transformation.

Arcadis is a global design, engineering and management consulting company, with over 28.000 employees in 40 countries. The organisation has developed “Vision 2030”, which serves as the “North Star” and guides their digital transformation. Both the digitalisation of the current business, as well as the invention of new business models, are heavily dependent on data and analytics.

Christoph states: “Data analytics is the cornerstone of our Digital Transformation and front and centre of everything we will be doing at Arcadis”.

The four phases of the analytics roadmap at Arcadis

1. Establishing the foundation

Setting up basic data analytics capabilities, supporting client development in bringing analytics to their customers, setting up a data infrastructure and data governance to better collect and leverage data, develop analytics skills (by training and hiring), developing new analytics use cases, and developing a vibrant analytics community within Arcadis.

2. Automate

Arcadis offers a range of global solutions, vertically (e.g. airports, train stations), and horizontally (e.g. business advisory services). The analytics team is actively involved in the automation of the current business, for instance by applying Design Automation, Natural Language Processing or Robotic Process Automation to improve margins, win more work by coming in at a lower price point, and to generate new revenue streams.

3. Advance

Pushing the boundaries towards even more impactful use cases by building out more advanced capabilities, in areas such as Computer Vision, nonlinear optimisation or by increasingly moving towards Data Exploration.

4. Repeat

As soon as analytics capabilities have been set up and established, they are being productised into reusable building blocks. For instance the company’s Computer Vision capabilities have been developed into such a tool, that enables Arcadians to conduct future projects more easily.

Consider how to...

Gradually evolve your data lake

Take an agile approach

Data lakes can be incredibly valuable for organisations. Through combining data sources from different parts of the business, they enable a more holistic view on business and its consumers. This in turn can speed up new analytics and AI initiatives that can drive more business value. However, building a data lake for all organisational data can take years. A phased approach should be favoured, in which data is gradually added to the data lake, according to prioritised use cases and the needed data for those use cases. AI efforts don't need to be delayed until all data has been centralised.

Ensure data governance

Simply moving an entire organisation's data into a data lake won't do much for an organisation's analytics efforts. Proper data governance is essential to ensure that the data is high quality, useful, trustworthy and well documented. The platform in which the data resides should allow for users with the right permissions to easily access, manage and protect the data.

The data scientists' challenge

Data science requires a lot more than hiring data scientists. Simply adding a data scientist to a team is not enough. Many organisations have experienced first hand that the data science role consists of many more aspects than solely analysing data and building models.

A large part of the data science role often is accessing, gathering and cleaning data, as data can be scattered across the organisation with little governance. In this context, soft skills are also crucial: the ability for a data scientist to effectively communicate with the business and understand business needs.

Centralising organisational data in data lakes, and setting up clear data governance practices will likely increase the overall effectiveness of data scientists in an organisation.

03 Bridge AI & business.

This is the era of the 'Analytics Translator'.

The challenge

Organisations are facing the challenge of finding employees who understand core concepts of data science and artificial intelligence, and also have the business acumen needed to identify valuable AI opportunities. This dual understanding is critical in shaping AI use cases that result in significant business value.

The response

New roles and training programmes are emerging to bridge the gap between business and artificial intelligence.

There are different descriptions for roles that fill this gap. The most commonly used description is ‘Analytics translator’. Other common roles are ‘AI Product Manager’, ‘AI Road Manager’ and ‘Business translator’.

The skill of understanding both a specific business domain, as well as data science and AI, is seen as critical, yet extremely difficult to find. Most organisations are therefore focussed on training existing employees, often those who have extensive business experience, to understand core concepts of AI, machine learning and data science. Employees with a dual understanding of business and AI can help teams identify impactful use cases, and help all involved disciplines to collaborate more effectively.

Quote

“Data scientists & engineers are scarce, but at least you can find them. However, people who understand the data science world and are also able to consider strategic business value, are very rare.”

Sarah Oey, Head of Data Analytics Global Commercial, Shell

Visualisation

AI and data analytics roles in organisations

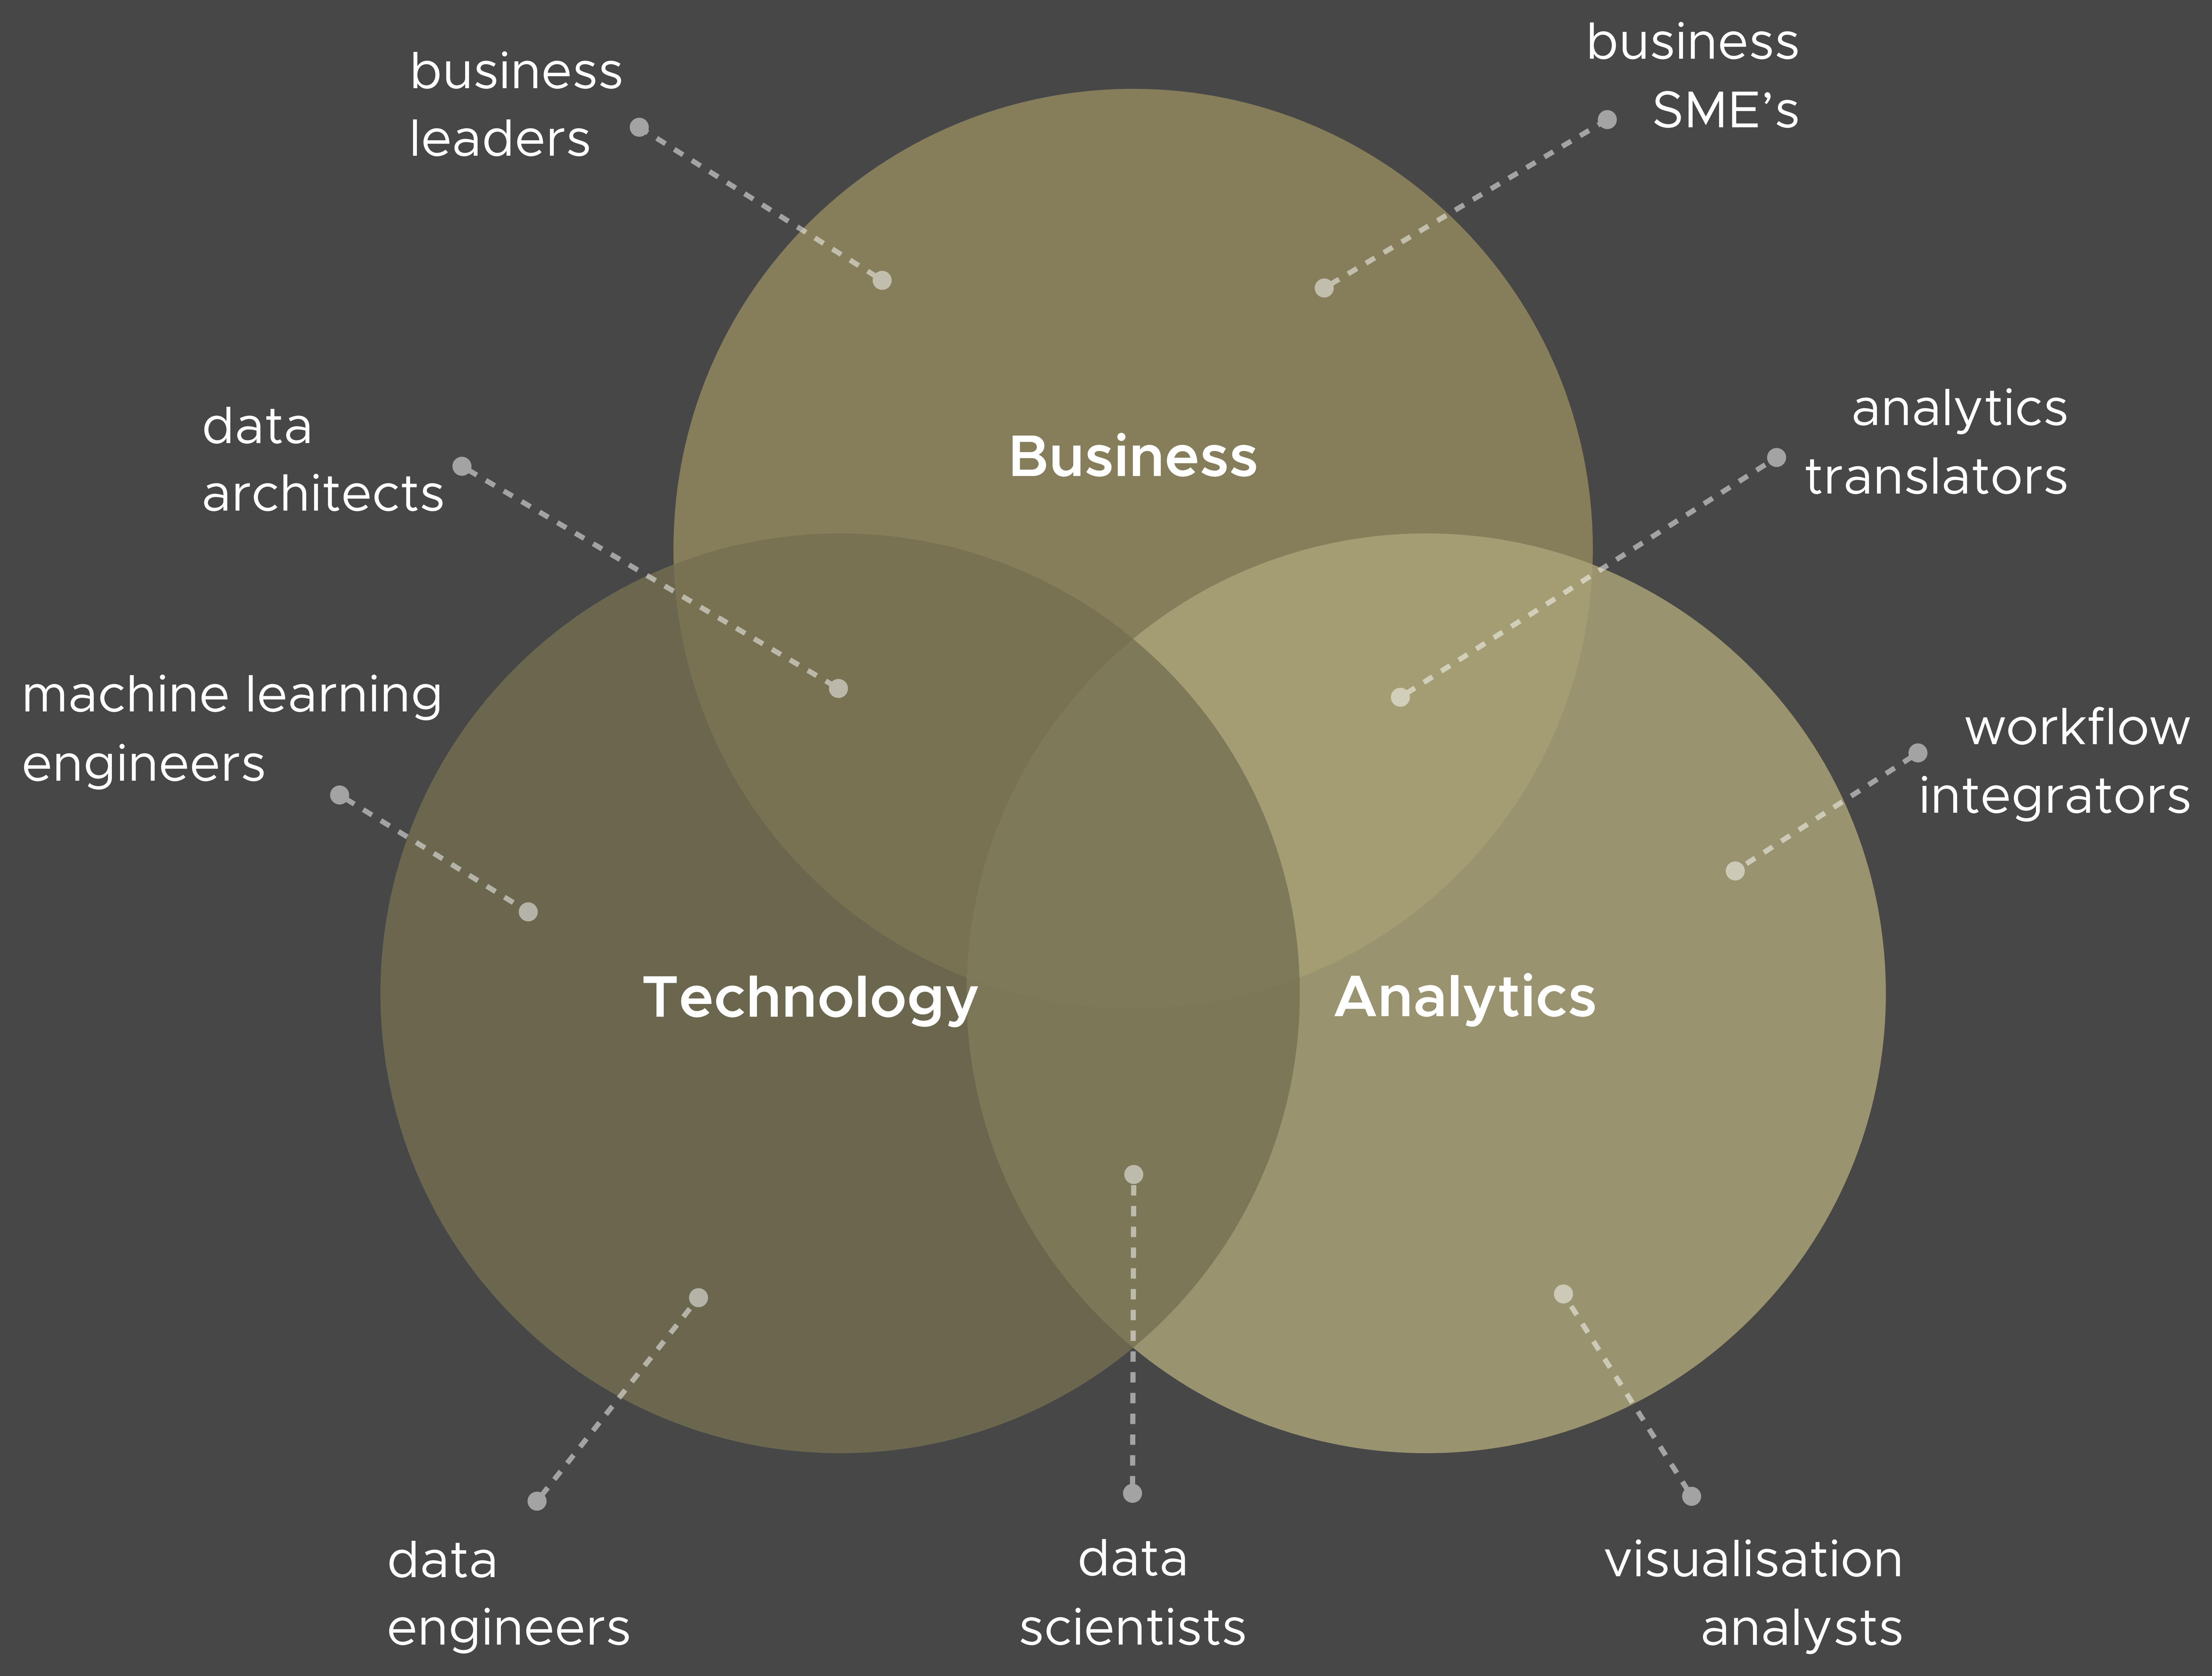

Visualisation adapted from McKinsey

Case study: Shell

Building analytics translator capabilities across the organisation

Sarah Oey, Head of Data Analytics Global Commercial at Shell, explains the importance of the business translator skill, within Shell's analytics efforts.

At Shell, the analytics translator skill (or 'business translator' in their terminology), is considered to be a critical factor in succeeding with AI and data analytics.

The importance of the business translator skill at Shell

Cross-functional analytics hubs, beyond silos

Shell established analytics hubs, with dedicated analytics teams combining technical expertise with business SME's. On top of these hubs, end-to-end roadmaps across departments were defined, such as "customer 360": creating a transparent customer view across marketing, sales and operations.

Scaling? Business translator skill is key

Shell started by establishing an analytics centre in London. Right now, the focus is on building up business translator capabilities, before establishing additional centres in other locations.

Business translator training

It is a strong belief at Shell that analytics should be done co-located. Therefore, Shell facilitates 2-day business translator trainings in classrooms, to help business employees understand core concepts of data science. However, a 2 day training only scratches the surface: the skill is only developed and proven after delivering a few use cases. Therefore, Shell is exploring other ways to enhance the translator skill, such as business translator Pro trainings for suitable candidates.

Consider how to...

Bridge the data science and business gap

Understand the business and AI gap

Whilst business leaders have a deep understanding of their area, they sometimes lack an understanding of the possibilities of AI and data science, and vice versa.

Simply adding an AI specialist or data scientist to a team, and expecting them to generate valuable output, can result in frustration. They need to understand which business problems exist and could benefit from an AI / data science solution. This is the area in which analytics translator capabilities can create great value.

Understand the needed skillset

Analytics translators bring business acumen and an understanding of data science, its possibilities and challenges. Therefore, the role calls for an ability to effectively communicate with both sides. Additionally, strong project management skills are required to successfully bring projects to fruition.

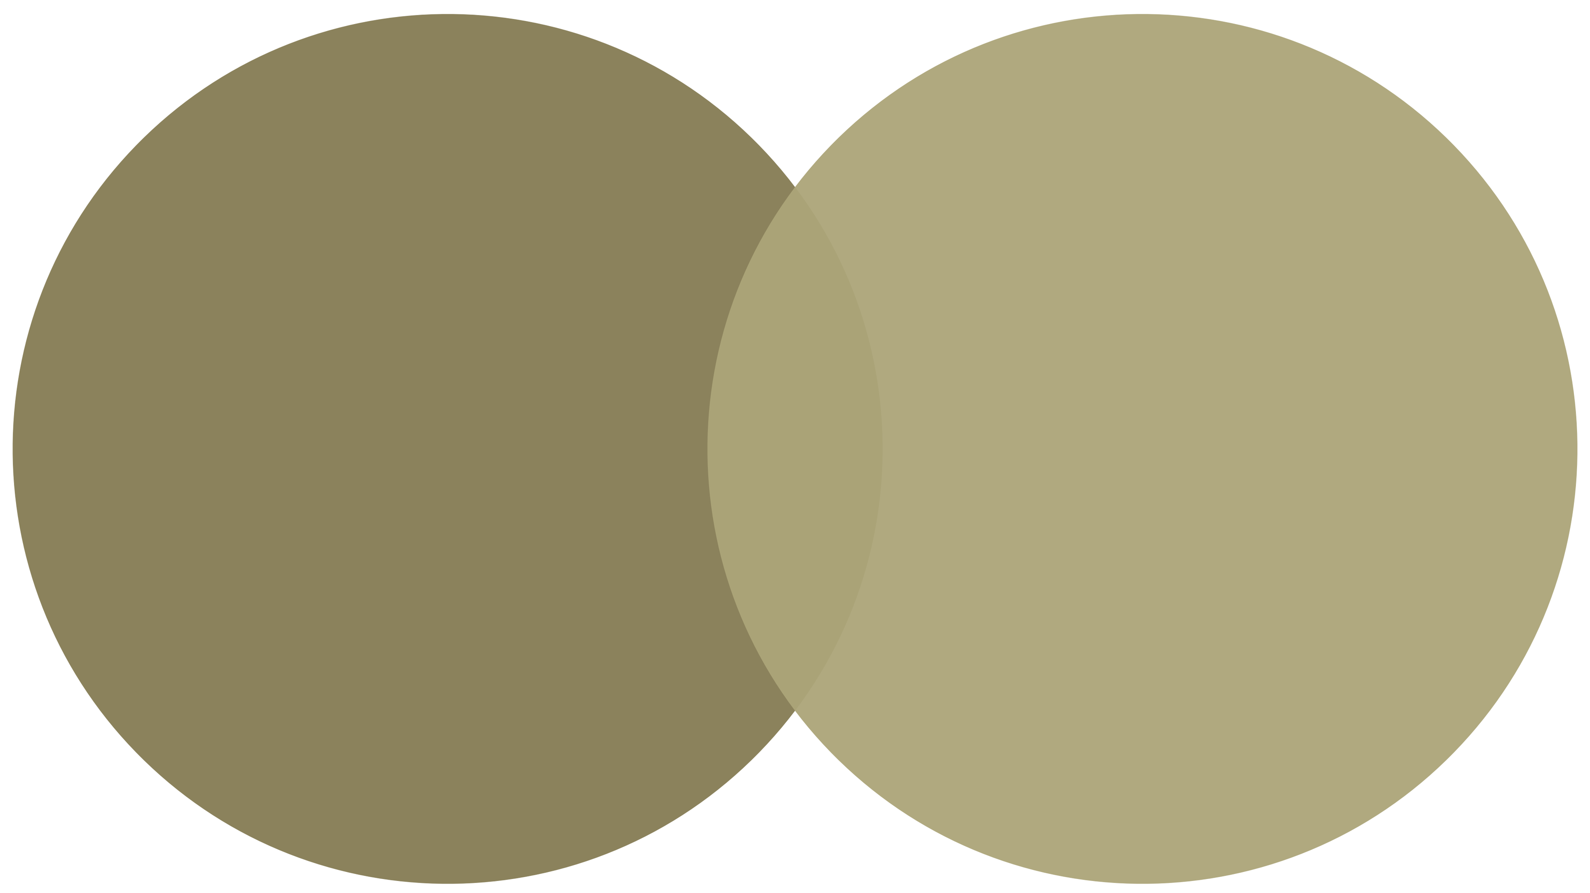

This combination of skills is the foundation that allows analytics translators to identify valuable use cases, and act as a liaison between business and data scientists.

Choose whether to train, hire, or do both

Organisations are increasingly facilitating training programmes to up-skill business leaders to become more data literate. A two day training programme however, isn't sufficient to become an analytics translator.

Hands on experience is required to gain experience, and learning from experienced peers is important to become familiar with the role and the challenges it brings.

04
De centralise!

Establish an interdisciplinary environment

The challenge

Should AI be a separate department in the organisation, or is it better to integrate expertise in business and product teams? What is the most effective organisational structure for effective AI adoption? These are common questions that arise when organisations are facing the challenge of identifying a suitable operating model for AI adoption.

The response

Companies have experimented with different structures and operating models for AI. Decentralising teams to a certain degree is key to success, allowing AI experts to closely collaborate with business units and product teams.

The right structure and operating model is highly dependent on the existing organisational structure, size, maturity and technological/ data infrastructure.

Regardless of the structure and operating model you adopt, it should allow for teams to work together cross-functionally. Optimal team structures typically include a business stakeholder, analytics translator, data science and engineering specialists and (UX) designers. Teams should be dedicated to projects and assigned to new challenges on a regular basis.

Quote

"Analytics and AI need to be as closely aligned to the business as possible. Because that is normally the biggest gap, and the hardest part to learn is the domain expertise."

Tad Slaff, Product Owner Management Information, Suitsupply

Visualisation

Different team structures of AI in organisations

Centralised model

Decentralised model

Hybrid model

Case study: Picnic

Full stack profiles and decentralising the data science team

Often times, organisations start out with a more centralised data science team, before transitioning to a decentralised model. Daniel Gebler, CTO at Picnic, explains the road to their current team set up.

Picnic is an online supermarket, that launched in 2015. From the start, Picnic focused on hiring engineers with a background in machine learning, as they anticipated this to be an area of opportunities. In their journey, Picnic transitioned from a centralised data science team, to a set up in which every product team has dedicated machine learning and data science capabilities.

Transitioning away from a centralised data science team

The risks of a centralised data science team

Picnic started with a separate Data Science team, as most organisations do. However, Daniel Gebler argues this to be an anti-pattern, as data science risks to become a single functional team that works on all kinds of different projects. This can result in difficulties in prioritising projects, as they run independently from the rest of the business.

Teams with dedicated data science capabilities

Picnic decided to take a more decentralised approach and explicitly dedicate team members to product teams. Each product team therefore has a product owner, front/back-end developers, data and machine learning engineers, and data scientists.

Full stack profiles

The importance of full-stack profiles is also emphasised in the team structure. Picnic tries to mitigate the risk of overspecialising, in which a team ends up waiting for a specific specialist to finish his or her task. At Picnic, successful data scientists should be familiar with the basics of machine learning and data engineering. Even though each role has a different focus, they should be able to work independent and "full stack”.

Consider how to...

Choose the structure & operating model for your organisation

Decentralised model

In a fully decentralised model, AI & analytics specialists are scattered throughout the organisation without any coordination, or overarching organisational governance. Each business unit is responsible for their own resources, operations and projects.

Centralised model

In a fully centralised model, there is a central team which manages all analytics or AI initiatives. They provide the expertise and prioritise projects based on the overall organisational strategy. As a result, the team implements the majority of AI and analytics projects, and are responsible for governance and coordination across the enterprise.

Hybrid model

There are a number of different hybrid variants, in which elements from the decentralised and centralised models are combined.

Frequently, a Center of Excellence (CoE) model is used. The CoE provides AI or data science services to the rest of the organisation, such as consultancy, implementation and support.

The CoE also ensures companywide training programs and governance on data and AI technology infrastructure. Meanwhile, the business units have their own capabilities and are responsible for prioritisation and assessment.

The best operating model?

There is not one clear winner in terms of operating model. However, most organisations that are significantly advancing with analytics and AI, have a hybrid operating model which combines a degree of centralisation with decentralisation.

Facilitate close collaboration

The main reason these models seem to be more successful than fully decentralised or centralised models, is because they allow for close collaboration between the business units and the data science experts, whilst maintaining a degree of organisation-wide governance.

05 Centralise!

Establish common practices and training initiatives

The challenge

Organisations are facing the challenge of ensuring that certain standards (such as data governance) are shared across the organisation, and that common practices are adhered to. Additionally, organisations that decentralise experts across teams, need to ensure that knowledge is shared across these team members.

The response

Whilst decentralisation is important to facilitate interdisciplinary collaboration, a degree of centralisation is necessary to ensure common practices.

Most organisations choose, either implicitly or explicitly, a hybrid operating model, in which decentralised and centralised practices are combined.

The main reason for a degree of centralisation, is to ensure that data governance and other standards are adhered to across the organisation, and to facilitate knowledge sharing among a community of peers.

Quote

“We have one community, a lot of events, and ‘knowledge sharing Fridays’. We also organise global events, and we have our own Slack. We can do more, and that is coming soon.”

Görkem Köseoğlu, Chief Analytics Officer, ING

Visualisation

A hybrid approach of decentralisation & centralisation

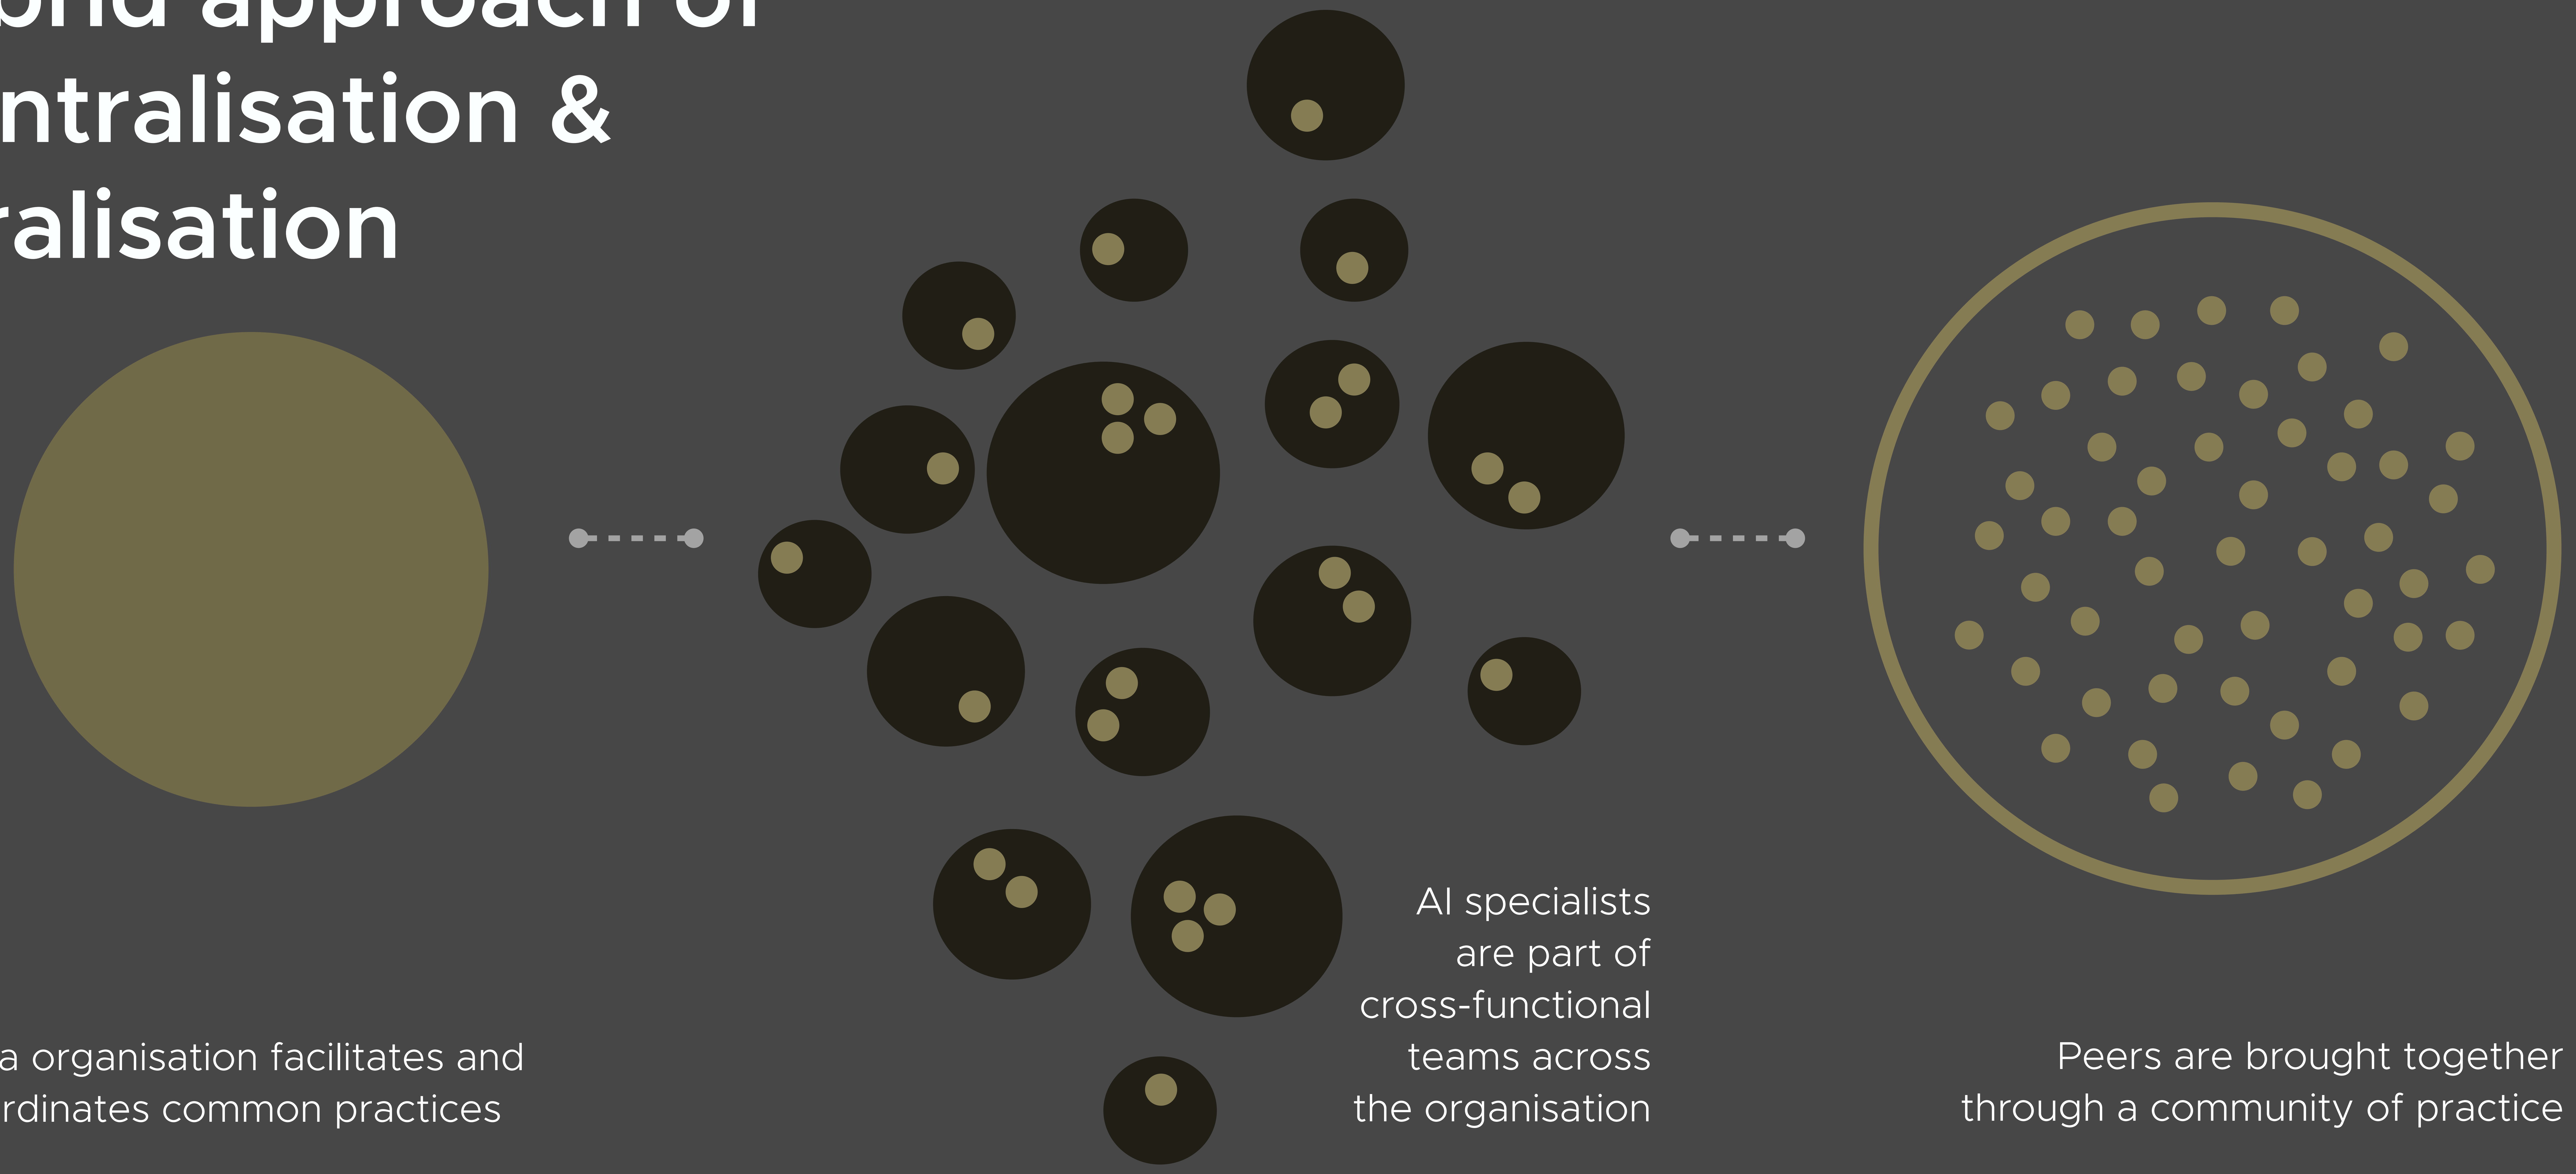

Case study: TomTom

Knowledge sharing across a decentralised organisation

In a decentralised organisation, it is essential to create a multitude of platforms for knowledge sharing. This holds especially true for newly established specialisation areas such as data science and artificial intelligence.

Pierluigi Casale, Group Data Scientist at TomTom, explains TomTom's big data strategy. They created a strategy that suits the decentralised nature of the organisation, allowing every business unit to take full advantage of TomTom's data and advanced analytics capabilities.

Each unit developed their own data strategy based on the general advice of the company's big data strategy, as every unit is working with different types of data on different types of problems. What ties them all together, is their community of practice on data and AI.

TomTom's community of practice on data and AI

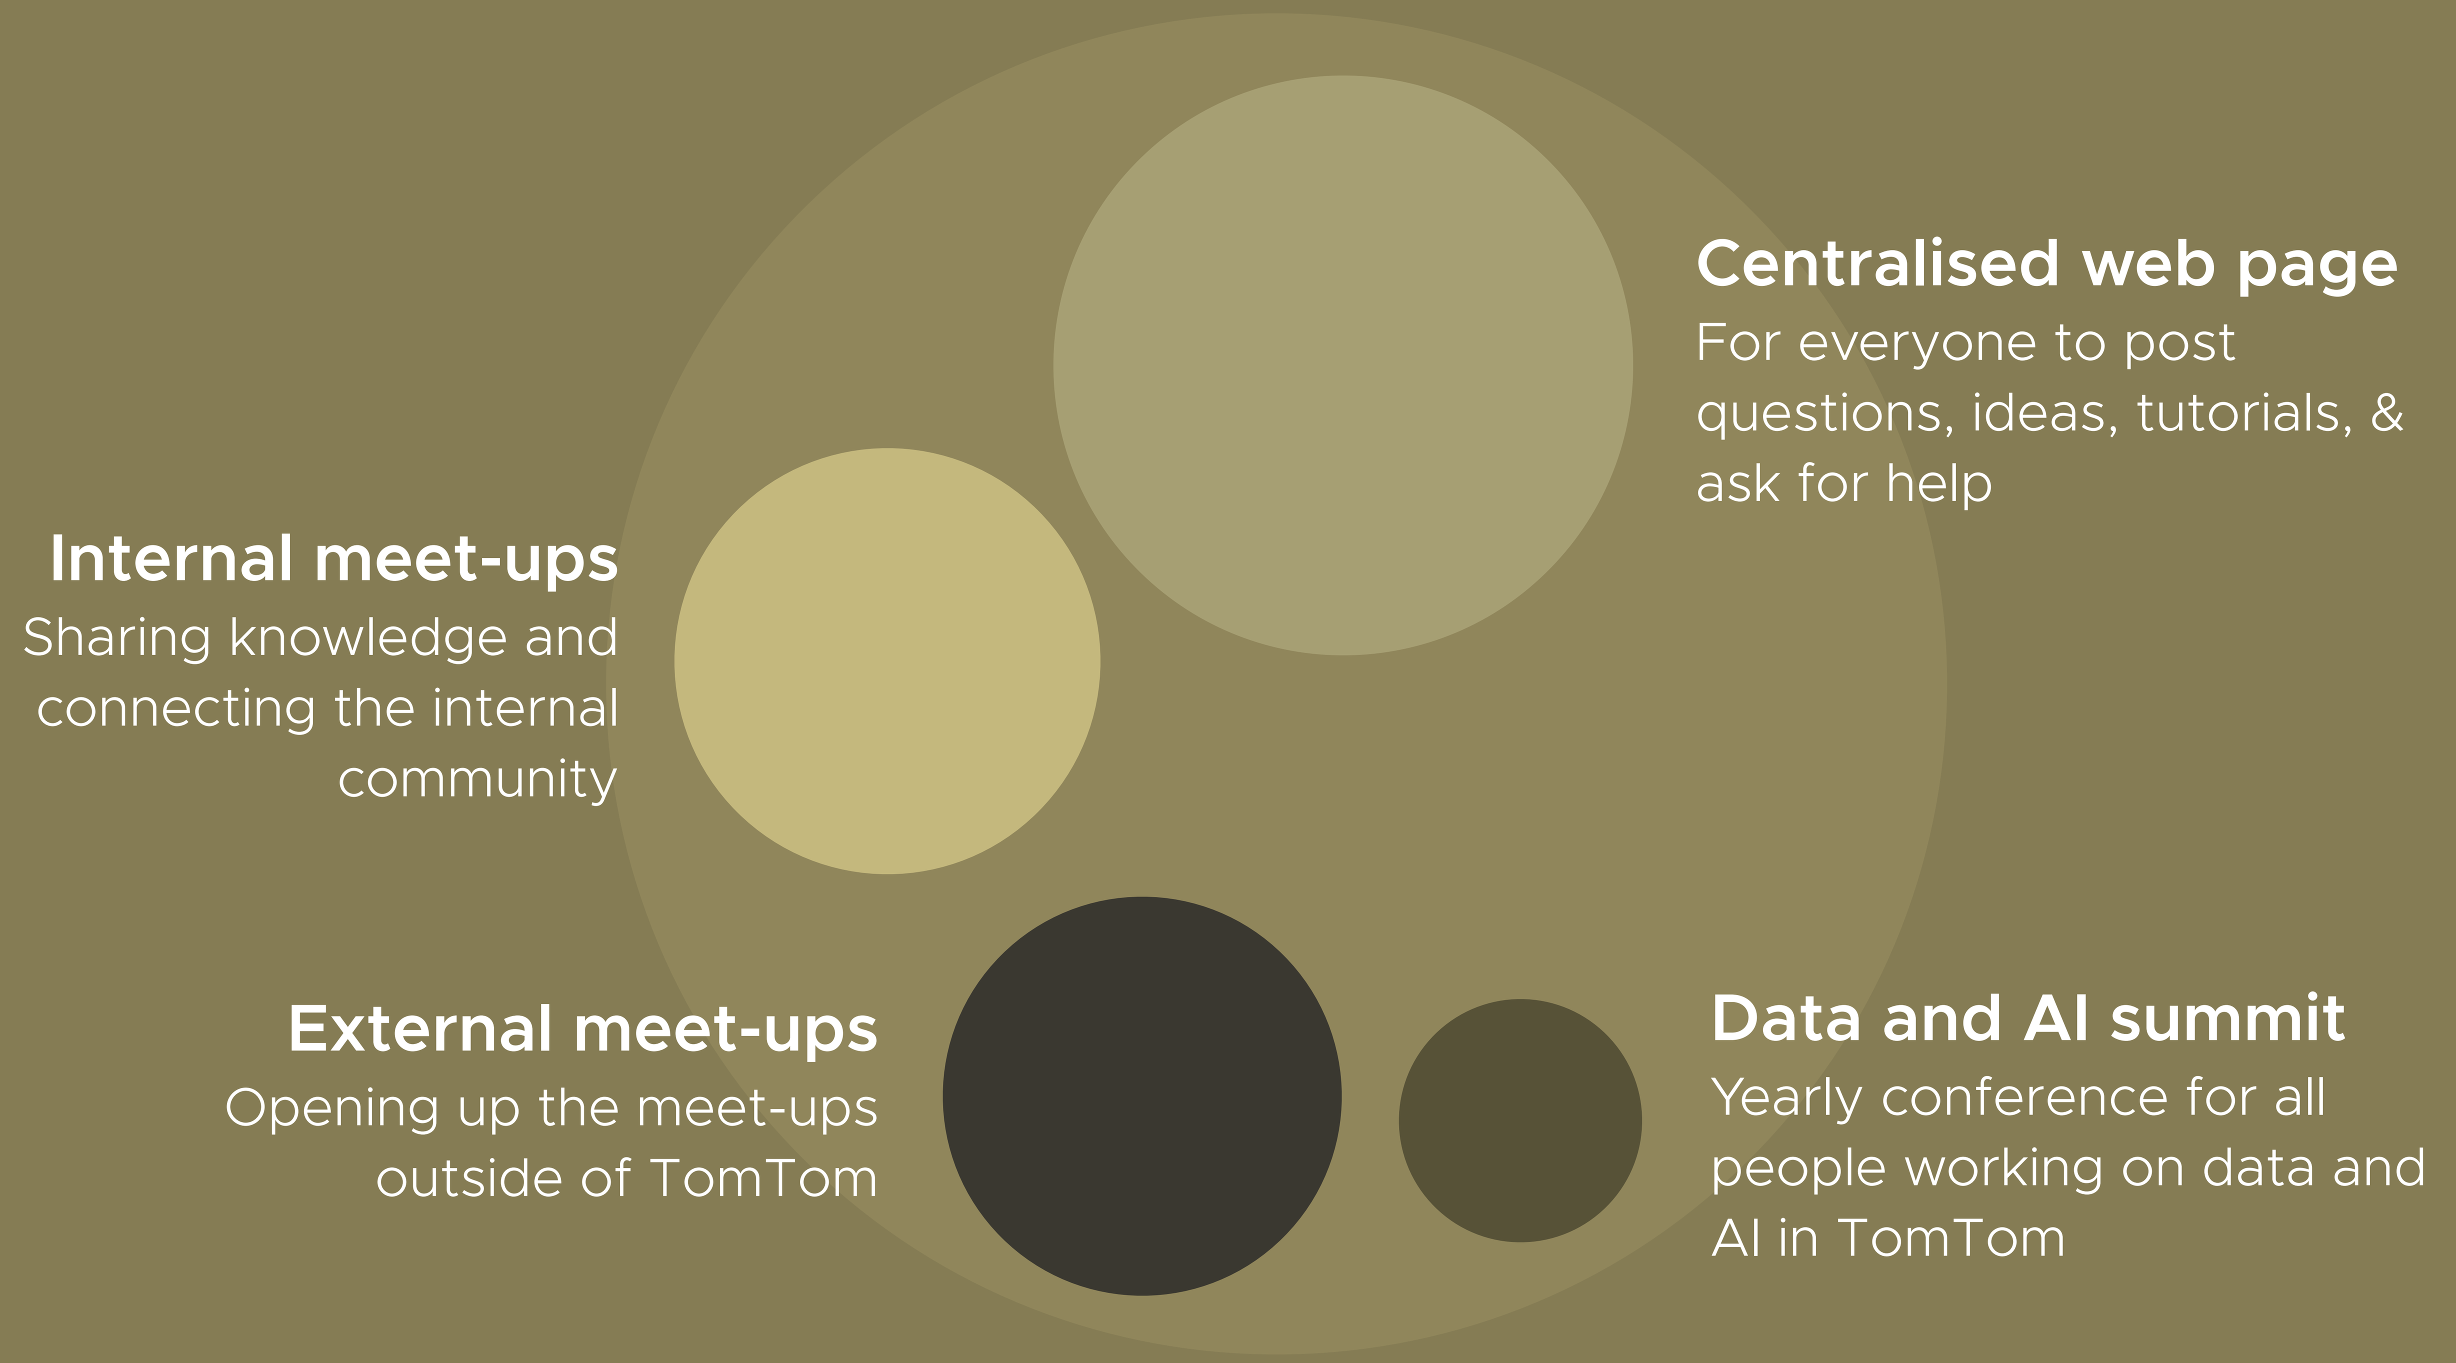

Consider how to...

Centralise common practices across the organisation

Centralise data governance practices and core models

Most successful organisations have a central party that focuses on data governance, how data should be used, the platforms and standard procedures for handling data and hosting models. Additionally, for some organisations it can be relevant to have a number of core models that can be used organisation-wide, to ensure consistency and prevent unnecessarily duplication.

Establish a community of practice

Regardless of whether you call it a community of practice, a chapter or something else, it is key to bring employees in the same area together regularly, especially when they are decentralised and working closely to business units across the organisation. A community of practice allows employees to connect, share knowledge, exchange ideas, get help on challenges and discuss emerging technologies.

Change the level of centralisation over time

Most organisations start out with a centralised analytics / AI team, as this allows for the team to gain experience, monitor quality, and build common practices.

However, over time, successful organisations tend to lower the degree of centralisation and integrate their analytics teams closer to the business units, while maintaining the established practices.

06
Test & Learn.

Implement processes for rapid validation

The challenge

Better use of data and AI can bring tremendous business value. But where to start? In a big organisation there are hundreds of opportunities. How do you prioritise?

Especially when organisations are still in the early stages of AI adoption, teams often struggle to identify a process that allows them to evaluate ideas at an early stage, and establish a continuous validation funnel.

The response

Most organisations are still in the process of establishing a structured practice that allows them to demonstrate the business value of AI consistently.

To prove this value exists, it is important to adopt a way of working that allows for fast and early stage validation of a large number of use cases. Successful validation methods will allow companies to screen more use cases rapidly and then identify, scale and capitalise on the successful ones.

Instead of creating long term strategies for unproven use cases, organisations that are successfully adopting AI are experimenting with pilot projects and validation methods that allow them to quickly evaluate whether or not a use case has the potential to create significant business value.

Quote

“We have an idea intake, followed by a BOOST-week to shape the idea. This way we try to offer our capabilities to the rest of the organisation. Low investment, and with a quick result. ”

Finbar Hage, Head of Data and Analytics, Rabobank

Visualisation

Short cycle innovation funnel for validating AI initiatives

Case study: ABN AMRO

4 weeks to get to a demonstrable outcome

Dennis de Reus, Head of Artificial Intelligence at ABN AMRO, explains their approach in establishing new AI initiatives. The approach focuses on quick value demonstration, with significant business involvement.

The first step in establishing a new use case is an initial conversation, where a use case template is filled out with:

- The value proposition
- Needed investment
- Needed data
- Ownership
- Outcome that is being optimised for
- What the solution will do
- Risks
- Assumptions

After refining the template and having a clear understanding of the use case, a 4 week cycle is started to get to the first demonstrable version of the use case.

The 4 week approach

Week 1

Designing the solution, requesting access for the data, setting up the development environments, so that everyone is ready in week 2 to work with the data.

Week 2-3

Try and train different models, initially more explorative: testing, understanding what the trade offs are, and what kind of feature engineering needs to be done.

Week 4

Decide on the suitable models and train those. Additionally, work on the presentation for the 4 week milestone: showing the model works, that the outcomes are better than current practices, a clear benefit, the road ahead. At the presentation (with a small steering committee) the decision to continue (or not) should be made.

The road ahead

2 additional months should be allowed to develop a POC and to demonstrate the business value of the use case. After a go / no go moment, 3 months can be added to bring it into production.

Consider how to...

Get started & quickly validate AI pilot projects

Ensure a strong connection with business goals

Most organisations that are succeeding with AI initiatives, initially focus on their core business areas. This allows them to demonstrate to internal stakeholders what value AI can bring. Only after initial value is demonstrated, it can make sense to also pick up more experimental projects that are not in an organisation's core business areas, and start using AI to improve upon secondary areas or processes.

Focus on desirability, feasibility and viability

Before starting a pilot project, it should be clear how the potential value of the project will be measured. Additionally, technical feasibility of the project should be ensured. Conducting a due diligence check can take a few days to a few weeks, to assess the suitability of the available data, the potentially needed additional data, and carry out experiments with the data to validate initial hypotheses.

Establish a process with short cycles

When it comes to getting started with AI, rather than going for the highest value projects, pick projects that allow you to easily get started with the technology and show measurable results in a short amount of time. Most pilot projects should be able to be deployed within 3-6 months. Starting with small projects, you can gain familiarity with the technology and the process.

Set metrics and measure results systemically

A successful AI project is not established the moment it is deployed; it is only successful once success can be measured. Before starting with your AI pilot project, determine how success will be measured by setting KPIs. Having clearly defined metrics will aid when explaining the value of AI projects to internal stakeholders, and thereby scaling AI initiatives across the organisation.

07
Create ethical AI.

Anticipate the effect beyond the algorithm

The challenge

While artificial intelligence can contribute in identifying human bias, it can also increase the impact of human bias and scale it, with tremendous consequences.

Organisations that are using AI, are facing the challenge of preventing bias and its consequences.

The response

Although the importance of ethics in AI is receiving a lot of public attention, only a few organisations have established guidelines for bias prevention.

While establishing guidelines is a good step forward, the practical implementation of these guidelines is often falling behind. Most organisations are still exploring a solid process to ensure bias is prevented when developing AI use cases. Currently, proper evaluation, approval and governance is not taking place in most organisations.

A growing number of examples highlight the need for regulation and guidelines for ethical and responsible AI: Cambridge Analytica's abuse of data, DeepFakes (and DeepNudes), racial bias in predicting future criminals, and many more, exemplify the growing need for industry-wide governance.

Organisations need to carefully consider how they will develop AI applications, as the backlash of irresponsible applications of AI can have a tremendous impact on individuals and society.

Quote

"Because humans can be biased sometimes, the models built by them, can be so as well."

Nic Huiskes, Head of Marketing Intelligence, Randstad

Visualisation

Consider how you will you develop AI ethically, in a way that is...

Case study: Rabobank

How to deal with the potential drawbacks of AI

Finbar Hage, Head of Data and Analytics at Rabobank, explains how they are dealing with neutralising the potential drawbacks of artificial intelligence.

Throughout the organisation, generally the perception of AI is very positive. However, worries also exist, especially on a supervisory level. The potential negative effects of how models are constructed and applied is something they are concerned about, and that causes doubt. Therefore, the potential drawbacks of AI are carefully considered in the implementation of new models.

Mitigating the potential drawbacks of AI

The business data committee

In order to mitigate the potential drawbacks of AI, Rabobank has established a business data committee. This committee is focused on evaluating the machine learning models, the way they were built, how they work and how data is used.

The business data committee determines what is allowed, what not, and evaluates whether an application could be prone to potential bias. The committee governs to ensure necessary documentation and responsible application of the data used, the artificial intelligence applied and the impact on the requested use case.

Ensuring data is only used for the purpose

Rabobank values privacy and transparency in the use of data. An important role of the committee, is ensuring that data is used within those strict boundaries. When data is used for activities that it was not originally collected for, the business data committee evaluates the proposal before allowing the data to be used. Thereby, they also ensure that privacy goals are met. For example, they prevent transaction data from being used to tailor commercial marketing, as that was not the agreement in place with the customer.

Consider how to...

Design & develop ethical AI

Explainable

People need to be able to understand the decision process and rationale behind the output of AI applications.

Consensual and secure

Usage of data should be consensual and not disregard the data rights of individuals. Data should be handled by using secure and reliable systems and software.

Accountable

Organisations need to ensure mechanisms are in place for the responsibility and accountability of AI systems. Critical systems need to be carefully audited by internal or external parties.

Unbiased

AI applications should be inclusive and minimise bias, and representing diversity both in the data collection as well as the teams defining the applications.

An important note

There is not one established framework for developing ethical AI applications. Organisations tend to adopt existing guidelines or develop their own frameworks.

Ethics in AI deserves much more than a one page explanation; the four principles summarised here only scratch the surface of the considerations needed to be made, and are merely a starting point for thinking about developing ethical AI applications.

For continued reading to familiarise yourself with the subject, explore the following directories:

- Crowdsourced directory AI ethics: bit.ly/ethics-directory
- AI guidelines and frameworks: bit.ly/ai-guidelines-frameworks
- Ethics in AI research paper collection: bit.ly/ethics-ai-research

08 Explain AI reasoning.

Provide transparency in decision making

The challenge

Artificial intelligence is fantastic for solving problems in which large quantities of data need to be processed.

However, understanding its decision making, can be incredibly challenging, if not impossible in some cases. This is a key challenge for organisations, as for certain use cases, it is critical to be able to trace back why a machine learning model provided a certain output.

The response

There is an increased interest in developing Explainable AI: AI that allows for interpretability of the outcome of AI models.

Explainable AI is an expanding field that aims to increase interpretability of machine learning outcomes. In short, it answers the reason why a model provided a certain outcome. This can be incredibly challenging, as it isn't uncommon for a machine learning model to have millions of parameters that influence the outcome of a model.

For now, explainability is largely a tradeoff: certain models are inherently more explainable than others. However, advancements in this field and technical design decisions will play a key role in users' ability to trust and work with AI systems and outcomes.

Quote

"The theme is new. Not for whoever is in the field, but for business people: now they need to care about what happens under the hood, and that it is explainable. That came at light with GDPR."

Patrick Attallah, Global VP Data & Analytics, DSM

Visualisation

Accuracy versus explainability: the tradeoff in creating explainable AI

Case study: Google Explainable AI

A set of tools and frameworks to deploy interpretable models.

Google launched a toolkit for explainable AI, to help users deploy and govern machine learning models that are interpretable. It provides users with an understanding of how each feature in the dataset contributed to the predicted result.

Google is one of the frontrunners in the field of Explainable AI. Whilst the academic literature on Explainable AI is rapidly expanding, the majority is theoretical. Therefore, practical initiatives, such as this toolkit, are much needed to support organisations in developing responsible and explainable AI applications. It is expected that in the near future there will be a great increase in the available AI explainability tools, with increasingly more advanced capabilities.

Features of the Explainable AI toolkit

AI explanations

The AI explanations tool provides developers a score that explains how each factor contributed to determining a machine learning model's prediction. With these insights, developers can improve the datasets or model architecture, and debug the performance of machine learning models.

What-If tool

The What-If tool allows users investigate a machine learning model's behaviour, without having to write code. Datasets can be automatically visualised, and users can manually edit examples from their dataset, and see how these changes affect the model's output. These exploration and visualisation tools allow users to understand the model's functioning better.

Continuous evaluation

The continuous evaluation tool samples predictions from machine learning models that have been trained and deployed. The tool assigns human reviewers, so they can provide ground truth labels (the output they determine to be correct) for the prediction input.

Consider how to...

Include explainability in your AI initiatives

Determine when explainability is needed

Consider whether or not explainability is required in the systems you are developing. For example, if a deep neural network is used for a healthcare diagnosis, a binary output would not be sufficient; explainability is required to establish trust, and thereby usefulness in the first place. On the other hand, in case of a shopping recommendation engine, a recommendation lacking explainability, will likely do less harm.

Consider explainability in advance

The earlier explainability is considered in the development process, the better. Therefore, it is recommended to evaluate at the pre-modelling stage whether explainability is required. While explainability can be considered during and after model development as well, it should be considered as a technical design decision: creating an AI system that is transparent to the degree that is necessary.

Evaluating the needed type of explainability

Explainable AI exists at different levels, geared towards different stakeholders. Consider the following stakeholders as a starting point for evaluating the necessary degree of explainability.

Data scientists, AI engineers, AI researchers and other technical users

This group requires an understanding of how the model works, whether the outputs make sense and whether the explanations make sense. They need to be able to identify potential bias and discriminative pitfalls.

Business, policy makers, and other non-technical users

This group needs to be able to trust the models and their outcomes in order to use them, but more importantly, ensure that the models are used fairly and responsibly, and be certain that they abide by regulations (such as GDPR).

End-users / consumers

Explainability aids greatly in establishing trust with end-users, by providing transparency through informing them on how certain actions were taken based on the model's output.

09 Be better together.

Partner up to stay ahead

The challenge

Organisations are short for AI talent. Not only is there a shortage of data science and data engineering talent, there is also a lack of people with an external perspective, to generate ideas and breakthrough innovations.

The field is rapidly evolving, which poses the challenge: how can organisations benefit the most from the talent and knowledge that is available?

The response

Organisations are actively teaming up with external parties in establishing their AI initiatives.

Significant value can be gained from partnering with government, academia and business, in terms of sharing learnings and enhancing capabilities.

Silos should be avoided, and organisations should actively reach out to partner up with other people and organisations in the field, to enhance expertise, learn and share knowledge.

Quote

“We work with a strategic consultancy, but also with start-ups, to become faster at validating. We do small pilots in specific markets, to see what value they might bring.”

Sarah Oey, Head of Data Analytics Global Commercial, Shell

Visualisation

Consider partnering with...

Academia
MSc, PhD positions, internships, labs

Data providers
Acquiring data for specific interest areas

Training partners
Organisation wide AI education programmes

Industry
Collaborating with industry partners on AI for Good initiatives

Consultancies
Strategy, organisational transformation, extra capacity

Start up incubators
Monitoring innovative initiatives

Start-ups / scale ups
Agile ways of working, strategic acquisitions

Fit-for-purpose solutions
Purchasing software for specific use cases

Case study: ING

Partnering to enhance available AI capabilities

Görkem Köseoğlu, Chief Analytics Officer at ING, explains ING's current partnering efforts. While ING has many AI and analytics capabilities in-house, they also work together with academia and other organisations to strengthen certain areas.

According to Görkem, there is room for more externals, as there is a lot of demand and it is not always easy to scale up. He considers their current partnerships with fintechs, specialist players and big tech important, because if they are successful, some of their capacity can be used on projects. In Know Your Customer, ING is actively working on about 25 initiatives, some of them are models and others are analytics platforms. There is room for scaling even further; in this domain only, there is room for 10's of more initiatives.

Görkem: "There is so much to do. It's not just data scientists that you need; you need subject matter experts, project management, and more. So it's not just scaling by adding more data scientists."

ING's partnerships in the field of Artificial Intelligence

Organisation wide trainings

ING works together with a specialised data company, that provides trainings throughout ING, to further educate staff on artificial intelligence subjects.

Strengthen specific AI capabilities

In some areas, ING works with externals to strengthen specific capabilities. One of these areas is natural language processing, for which ING collaborates with big tech. For certain targeted solutions, ING works with startups; for example, through licensing text extraction software for legal documents.

Providing PhD sponsorships in academia

ING partnered with ICAI, for a number of PhD sponsorships. Additionally, they partnered with the MIT CS labs, which provide access to research and recruiting. Partnerships with academic institutes often concern PhD sponsorships, and sometimes assistant professor sponsorships. ING invests in partnerships with academia for ideas, inspiration, but also for building ties and for potential future recruiting.

Consider how to...

Combine in-house and external expertise in AI initiatives

Partner to get started with AI

Most organisations that are in the early phases of AI adoption, start out with small experiments and often choose to partner with consultancies or specialised data companies. This allows organisations to experiment and learn, without having to bring all necessary capabilities in-house from the start; which can be quite the challenge, as the demand for AI capabilities is skyrocketing.

Especially in the early stages, specialised consultancies can be valuable, as they can provide support in up-skilling existing staff, and thereby smoothening the overall adoption of the technology.

Bring in specific expertise

Most commonly, as an organisation's AI capabilities grow, it becomes increasingly important to partner in order to strengthen the expertise needed to develop specific use cases. To do so, it can be interesting to work together with start-ups/scale-ups that can provide specific tooling, which saves the effort of developing the exact same thing in-house.

Additionally, it can be valuable to bring on board individuals with specific expertise to enhance capabilities on a project basis.

Partner with academia

Many organisations choose to partner with academia. These partnerships allow organisations to make significant advances in terms of deep diving into industry specific research areas. Academic partnerships are also considered to be a crucial element in the hiring funnel, as it allows them to build relationships with students, in the hope that the talent will choose to continue to work at the organisation.

A notable initiative is Kick-start AI, in which some of the biggest Dutch corporates (Ahold, Philips, KLM, ING and NS) are providing funding to stimulate private-public collaboration through joint appointments of academic staff.

10
Design your own future.

Create value for society, people & business

The challenge

The recent surge of artificial intelligence is only the beginning of the developments that are to come. There will be more and more AI around us, whether we are conscious of it or not.

It is our responsibility to explore the frontier of AI to create new and elegant human(ity)-centered processes that make it easier and cheaper and better to do the right thing.

How will you...

Design the future that you want, and create a positive impact on people, business and society?

To get you started, on the next page you will find a design thinking canvas for developing human-centered AI applications.

You can use this canvas to explore value drivers and to assess how artificial intelligence can aid in developing desirable, feasible and viable solutions.

Consider how to...

Design your own future.

Jobs to be done

- What are you trying to accomplish?

AI Promise

- How might AI help you with the problem you are trying to solve?
- Which benefits the use of AI might bring to your organisation?
- How might the use of AI help you in achieving your goal or completing your task?

Benefits for humans

- How might the human agents' part of your process benefit from using AI?
- What AI is freeing up in your human agents' capabilities?
- Which repetitive tasks that your human agents are performing can be eliminated?

Machine activities

- Which tasks in your process are best suited for machines rather than humans?
- How might you make the best use of machine precision and accuracy?
- Which limitations of human abilities could be taken over by machines in your process?

Collaboration

- What is the role of human agents in overseeing the performance of the machine and refining the model?
- How might you use the input from human agents to improve the machine's effectiveness?

Human activities

- What AI technology allows your human agents to focus on?
- How might you make best use of human agents' judgment, flexibility and creativity?
- Which activities where AI can't be used should be overtaken by your human agents?

Human reinforcement

- How might AI to influence the work of your human agents?
- How might AI enhance your human agents by adding to their capabilities and improving their decision-making process?
- What feedback loop will you build between the machine and your human agents to augment their intelligence and abilities?

Critical thinking and biases

- What are the limits to what you can do with AI in your process?
- Which major or complex decisions can't be taken over by machines in your process?
- How might you apply critical thought to your machine performance and overcome its potential biases?

Considerations and implications

- How might you ensure that your AI is developed in a responsible, safe and reliable manner?
- Which morals, ethical or philosophical norms should be taken into consideration?

Change management

- How will you manage the anxieties of your employees that might be worried about their role being disrupted?
- How will you re-train your workforce to ensure a smooth transition of your company to the age of artificial intelligence?

Conclusion

The time to start experimenting is now.

We hope we've inspired you to think of how *you* can start AI initiatives, responsibly, and create value for your business, people, and society.

Most of the commandments in this publication deserve a significant deep-dive independently (or even an entire dedicated library section); with these commandments, we are barely scratching the surface. However, we hope that these insights help you point your efforts in the right direction, and ask the right questions to drive progress with your AI initiatives.

1. **Establish building blocks**
   Your future organisation needs a foundation
2. **Connect your data sets**
   Continuously improve while projects are ongoing
3. **Bridge AI and business**
   This is the era of the 'Analytics Translator'
4. **Decentralise!**
   Establish an interdisciplinary environment
5. **Centralise!**
   Create common practices and training
6. **Test & learn**
   Implement processes for rapid validation
7. **Create ethical AI**
   Anticipate the effect beyond the algorithm
8. **Explain AI reasoning**
   Provide transparency in decision making
9. **Be better together**
   Partner up to stay ahead
10. **Design your own future**
    Create value for society, people & business

This publication was brought to you by the DEUS initiative

Our mission is to assist business, academia and government institutes in exchanging ideas and experiences in developing valuable human(ity)-centred AI applications.

We want to help shape our futures, and we believe that through data, engineering, user-experience and strategy, we are unlocking services that will transform the way we live.

Author: Nathalie Post
Editors: Ron Vrijmoet, Nick Mueller,
Danny Verloop, Richard Olyerhoek

Want to have a chat about AI, learn more, collaborate, have a coffee, or find out how we can help you?

We'd love to hear from you!

📱 **Give us a call at:**
+316 50 46 26 37

📧 **Mail us at:**
human@deus.ai

#️⃣ **Follow us on:**
LinkedIn: linkedin.com/company/deus-ai
Facebook: facebook.com/deus.initiative
Instagram: @deus_initiative

Want to read more?

Stay posted on our AI in Practice series or check out our podcast!

This publication is the first release in our AI in Practice series. More practical insights, case studies and tools will be released soon!

Want to receive the next edition straight into your inbox?
Sign up for our newsletter at **www.deus.ai**

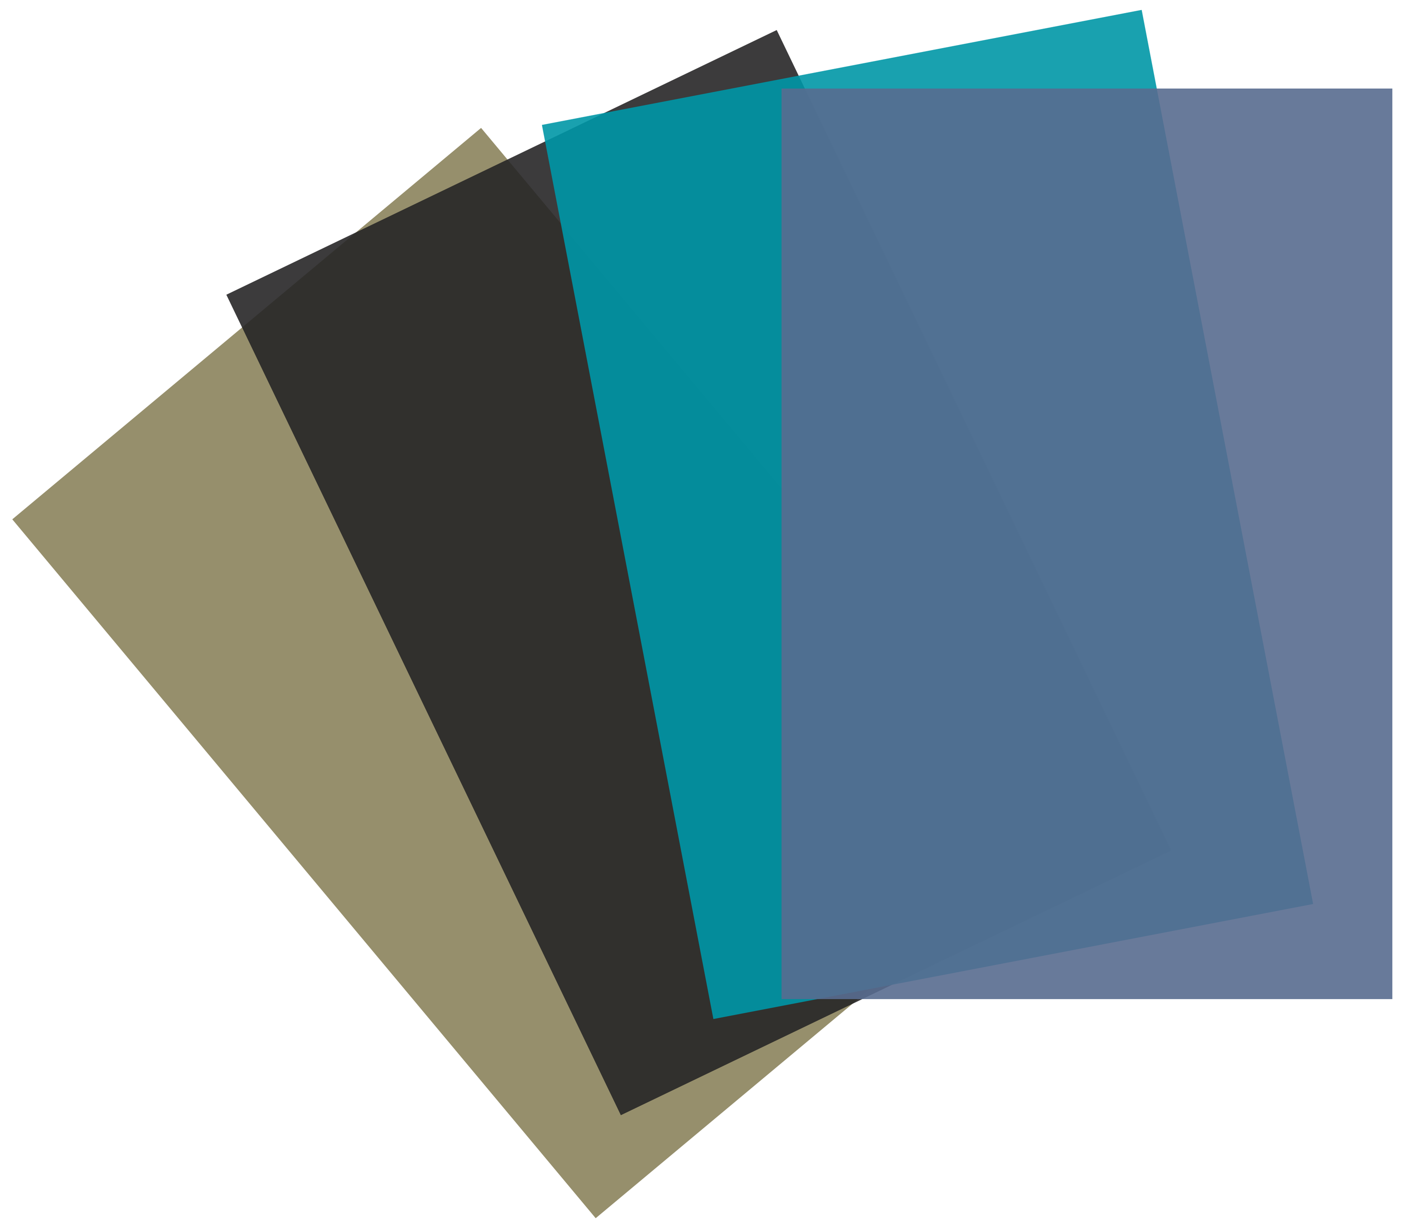

DE
US