2nd edition

initiative for applied artificial intelligence

Elements of a Comprehensive AI Strategy

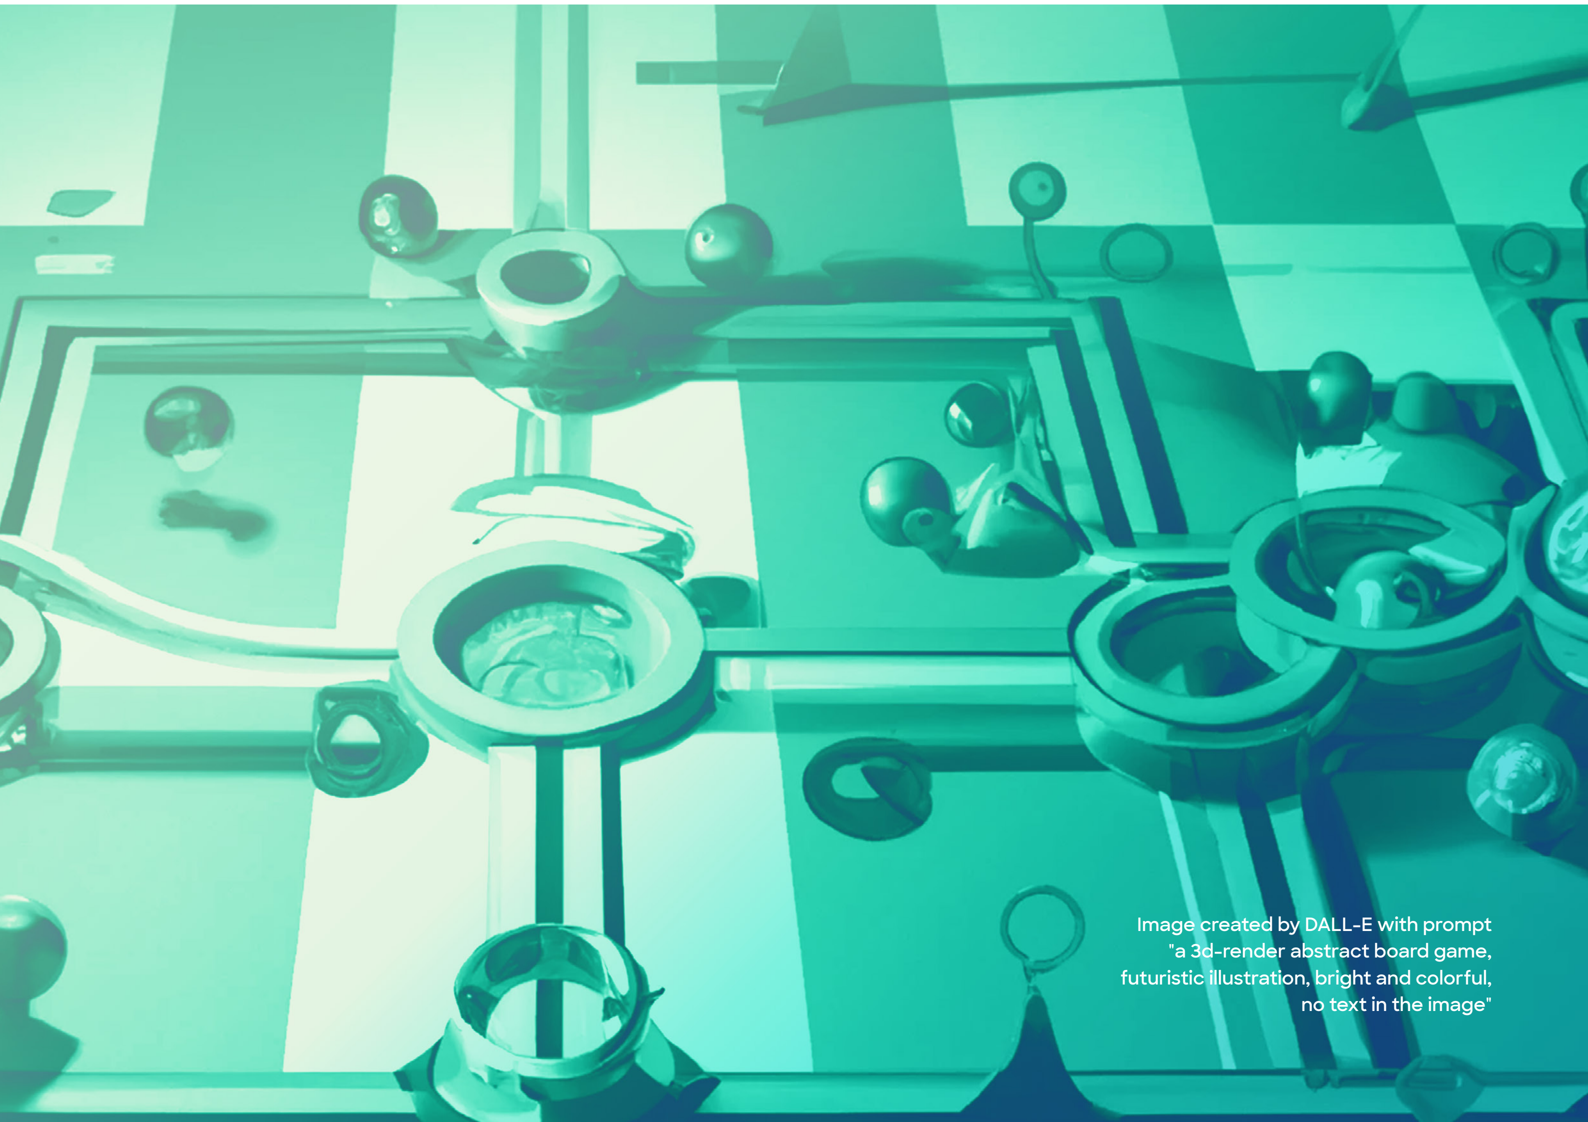

Image created by DALL-E with prompt "a 3d-render abstract board game, futuristic illustration, bright and colorful, no text in the image"

Contents

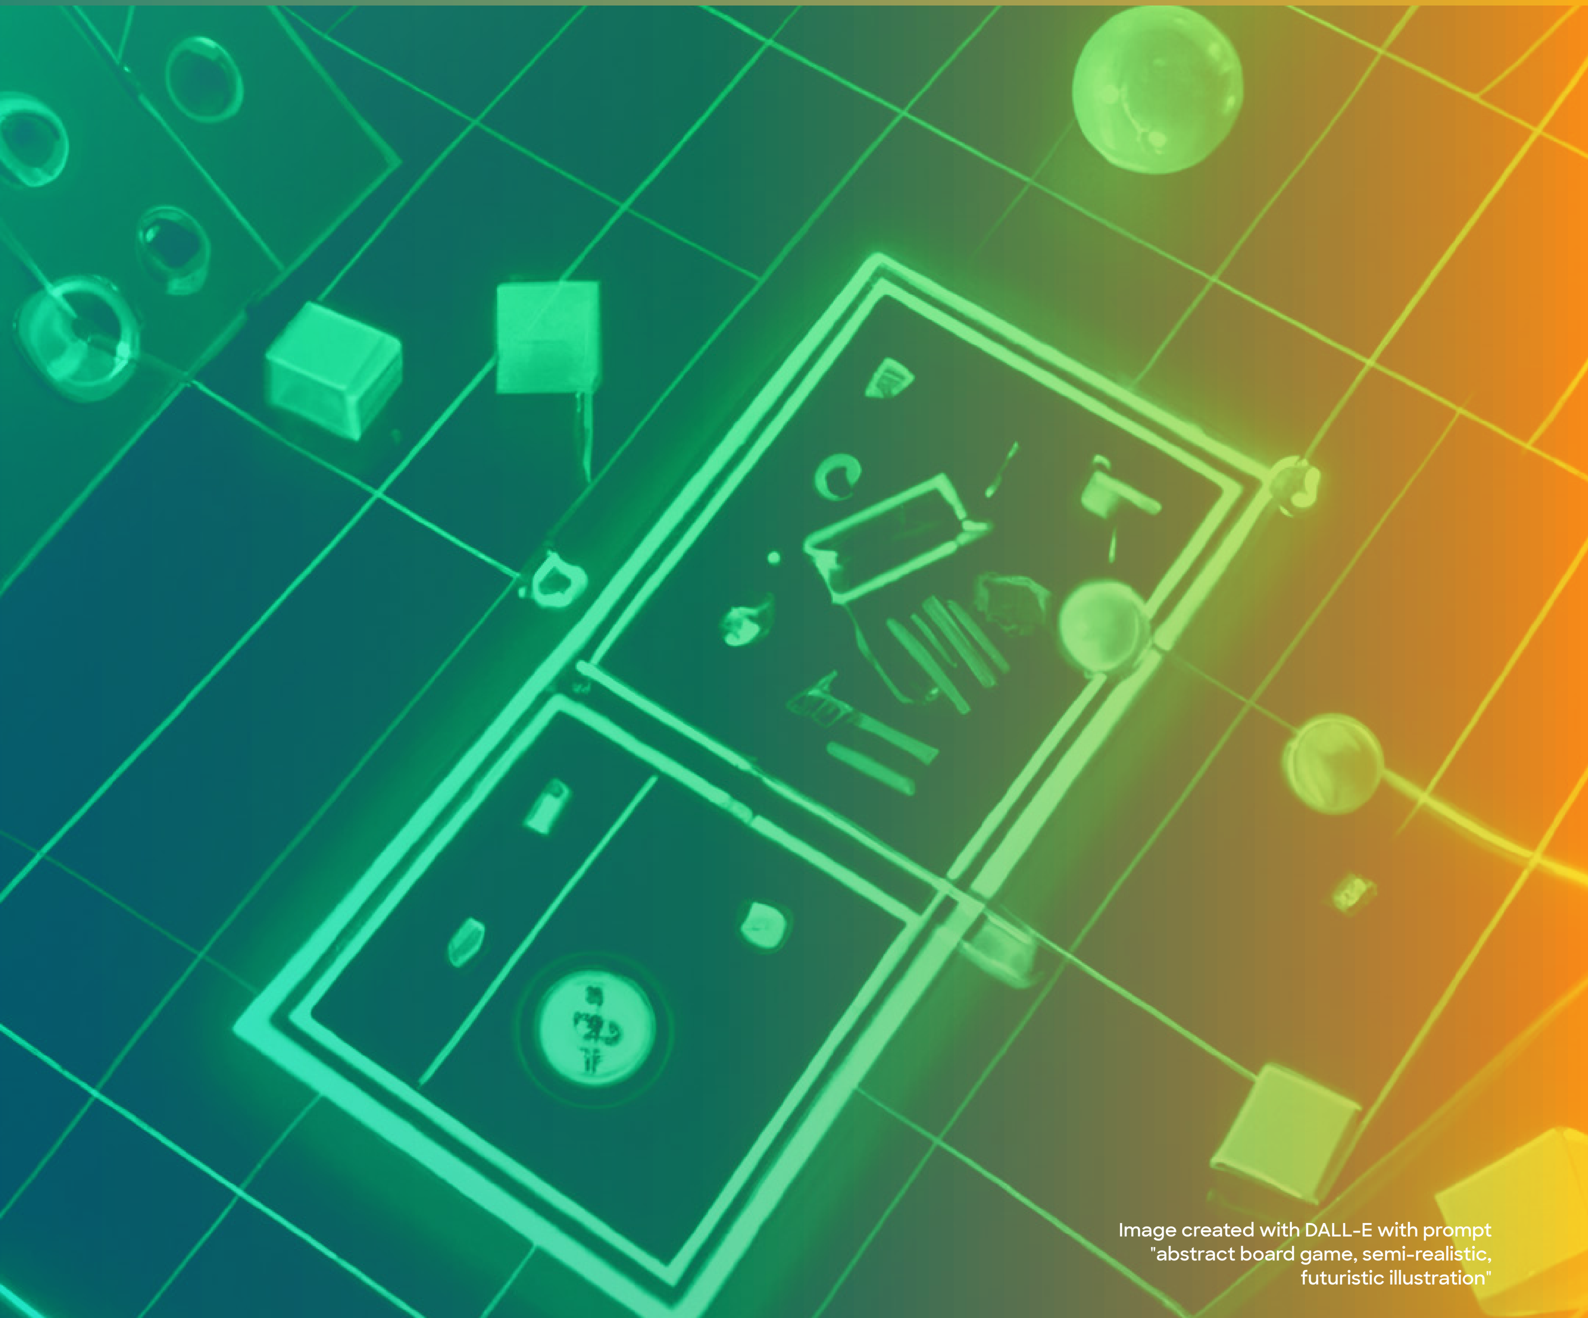

Image created with DALL-E with prompt "abstract board game, semi-realistic, futuristic illustration"

Introduction

Just 20 years ago, many industries looked significantly different than they do today. To rent a movie for a relaxing evening at home, you might have gone to a Blockbuster store rather than browse Netflix or another streaming service. Retail was dominated by companies like Barnes & Noble, Toys "R" Us, and Sears rather than Amazon. And if you wanted to book a hotel for your next holiday, you would have probably used a travel agency, not Airbnb. Many companies that underestimated the potential of the Internet during this period ended up seeing their business models become severely disrupted, and in many cases even obsolete.

Today, just like the Internet did 20 years ago, artificial intelligence (AI) is increasingly transforming our daily lives and with recent advancements of generative AI, this trend has accelerated even further. Some have suggested that AI has a stronger potential to change society than the Internet, fire, or electricity. AI can facilitate cancer diagnosis, robots can increasingly carry out sophisticated operations, and language models such as OpenAI's ChatGPT have startled established industries such as software engineering and academia. In recent years, we have seen significant advancements in AI development, leading to a vast and rising number of sectors being transformed. One recent study showed that the percentage of organizational revenue influenced by AI more than doubled over the last three years, with a prediction that it will triple by 2024[1].

Many companies are investing in AI technologies to enhance their products, services, and processes, all aiming to leverage the technology to gain competitive advantage. Yet there is a major gap in AI-related value creation between companies. While those like Google and Amazon are already making massive profits leveraging AI technology, the majority of firms have not yet been able to see any such positive impact. There are many reasons for this, but most can be traced back to the same basic problems. For many companies, there is no clear understanding of why AI is needed, which use cases should be implemented, what exactly is needed to do so, and in which order challenges should be tackled. In short, there is no comprehensive AI strategy.

Companies that have so far been successful with AI started to think about adopting the technology in its early stages and sought to understand the challenges that would arise with implementing AI at scale. As such, they realized that a holistic and comprehensive view of AI was needed. If we consider AI adoption as a journey that spans from a proof-of-concept (PoC) to routinely applying AI in daily operations, the companies generating value out of AI are those at later stages in the journey. A recent study found the average difference in return on AI investments between companies in the PoC phase and companies strategically applying AI to be 110 million US-Dollar[2]. Reaching a high level of AI maturity requires a comprehensive AI strategy with a structured understanding of the dimensions that need to be addressed. Such an understanding must go beyond discussions around use cases and address questions concerning AI ambition, enabling factors within organizations, and execution. It is also necessary to know how to develop these dimensions systematically over time.

This journey will differ depending on each organization's status quo and level of AI maturity, but a clearly defined path will help to navigate an ever-evolving AI landscape. AI is not just for big tech companies. This whitepaper will help you assess the potential of AI for your company and provide guidance on how to develop a coherent AI strategy.

1 Accenture (2022). The Art of AI Maturity. Retrieved 14 June 2023, from https://www.accenture.com/us-en/insights/artificial-intelligence/ai-maturity-and-transformation
2 Accenture (2019). AI: Built to Scale. Retrieved 14 June 2023, from https://www.accenture.com/us-en/insights/artificial-intelligence/ai-investments

What is Artificial Intelligence?

The term artificial intelligence was first coined in the 1950s, but there is an ongoing debate on what the term actually means. Often, AI is defined as machines acting intelligently, doing tasks that would require intelligence if done by a human. However, this definition is limited as the capabilities that AI brings are in many ways different from those possible with human intelligence. Some tasks that are easy for humans are hard for AI, and vice versa.

All of the AI applications we see today can be considered 'narrow AI', a term that refers to AI programmed to solve a single task, such as filtering out spam mails, recommending movies, playing chess, or translating text. This has very little to do with simulating real human intelligence, and instead involves using math and algorithms to make machines predict the best action given an input.

Most AI applications and virtually all major advances in AI in recent decades are in machine learning. Compared to traditional software, machine learning uses search and pattern recognition functions to derive rules and learn on its own. Traditionally, programmers automated tasks by writing programs. As such, programmers of conventional software must specify and know the rules in advance. In machine learning, the programmers (data scientists) feed the computer data and specify the output, and the computer writes its own program to fit those data. In this way, AI can tackle problems that were previously unfeasible or too complicated for humans to address. This includes circumstances where exact rules cannot be formulated, that deal with uncertainty, or where software needs to adapt to frequent changes in its environment.

To better understand AI, it's helpful to understand the areas for which it brings new capabilities. There are five main areas:

- **Perception:** Recognising structure and patterns in sensory data
- **Analysis:** Identifying trends and associations and making predictions
- **Language:** Understanding language and computer code
- **Planning:** Making tactical or strategic decisions
- **Generation:** Synthesising new data samples

Getting Started: How Will AI Impact Your Business?

Similar to electricity or the internet, AI is a transformative, general-purpose technology that has the potential to substantially change business opportunities, company risk, investment logic, human resource and skills requirements, processes, structures, and more. Existing business models may quickly become obsolete as new competitors emerge, often from adjacent industries.

These predictions are no longer about a distant future, as in certain industries AI is already a major source of competitive advantage. In just seven years, TikTok has scaled its service to 834 million users, relying on recommendation algorithms based on machine learning that drive user interaction. When you order a ride from Uber, an AI algorithm determines the optimal price, alerts the next available driver, and gives you an estimate of the arrival time. AI is also transforming traditional industries such as manufacturing, where companies like Boeing or General Motors are benefiting from using AI for intelligent maintenance, product quality inspection, and demand planning. The food manufacturer Danone Group uses AI to enhance demand forecast accuracy and in doing so has achieved reductions of 20% in forecasting errors, 30% in lost sales, and 50% in demand planner workload.

AI's impact on businesses has so far been significant and far-reaching, with some industries experiencing earlier and stronger effects than others. As AI continues to advance, we can see the emergence of a competitive winner-takes-all landscape in industries like e-commerce, online travel, and social media. It is therefore vital for companies to be proactive in embracing and leveraging AI in their business models. Waiting too long could easily result in being left behind by competitors who have already gained an advantage through AI implementation. Yet before considering how the transformation to AI can be best applied, it is important to first understand where AI's transformative power comes from.

Pushing the frontiers of automation

AI pushes the frontiers of automation by enabling machines to perform tasks that were previously only possible for humans. By leveraging advanced machine learning algorithms, AI systems can identify complex patterns and relationships within data, recognize objects and speech, make predictions and decisions, and even learn to perform new tasks through reinforcement learning. This has led to the automation of tasks across industries, from manufacturing and logistics to healthcare and finance. As AI systems continue to improve and become more sophisticated, they are expected to take on ever more complex tasks, further expanding automation's frontiers. Once a product or process is automated with AI, it can be scaled at low marginal costs.

A prime example is the Asian financial services company AntGroup, which uses AI to drive all core operations from customer services to loan approval. For instance, their 3-1-0 digital lending product, where borrowers can complete an online loan application in 3 minutes and obtain approval in 1 second with 0 humans involved, relies heavily on AI. This allows the company to serve more than ten times as many customers as the largest U.S. banks with less than a tenth of the employees.

Shifting the value of knowledge assets

As a general-purpose technology, AI shifts value towards those that master it, decreasing the importance of domain expertise and challenging existing competitive advantages. This enables organizations to compete in totally new areas. AI is increasingly being used to automate tasks that were traditionally reliant on domain knowledge. An example is DeepMind's AlphaFold, an AI-based prediction system that uses deep learning to predict the 3D structures of proteins, a traditionally time-consuming task that required significant domain expertise. AlphaFold is highly accurate, outperforming methods that rely heavily on domain knowledge.

Creating a virtuous cycle of AI

Placing AI at the core of a business model can help to create a virtuous cycle between user engagement, data collection, prediction, and improvement. Products and services are continuously improved through use, leading to even more use and a better competitive position relative to other products performing the same function – a 'flywheel' that can produce a powerful strategic advantage. Having data on how users or customers interact with your services and/or products can be incredibly beneficial, as it can be used to shape the development process, in turn creating a better user experience. As your product and service offerings improve, their use will generate even more data, which then further fuels the virtuous cycle by informing development of new products and services.

An illustration of this flywheel effect can be seen in how OpenAI improves their language models. As more users interact with the models, they provide feedback and corrections that are used to refine and improve the models. These improvements attract even more users, who provide more feedback, creating a virtuous cycle that continuously improves model quality. The cycle is a self-sustaining loop of improvement that results in language models that are increasingly sophisticated and capable of handling a huge range of tasks.

A Systematic Approach to AI

To successfully adopt AI, companies need to follow a systematic approach consisting of four elements. First, companies need to define their ambition regarding AI and ask themselves how AI can help to achieve corporate objectives, setting clear goals for AI application. Second, feasible use cases need to be determined, evaluated, and prioritized. Here, conscious decisions must be made about make-or-buy with regards to AI use cases. Third, companies have to make sure that enabling factors are in place to successfully implement the AI strategy and execute the use cases. Finally, the use cases have to be realized, i.e., researched, developed, or acquired and then deployed.

AI STRATEGY HOUSE

<table>
<tr><td>Ambition</td><td colspan="2">Future competitive advantage</td><td colspan="2">Fields of action</td><td colspan="2">Commitment</td></tr>
<tr><td>AI use cases</td><td colspan="2">Discovery & specification</td><td colspan="2">Make or buy</td><td colspan="2">Portfolio management</td></tr>
<tr><td>Enabling factors</td><td>Organization</td><td>Expertise</td><td>Culture</td><td>Data</td><td>Technology</td><td>Ecosystem</td></tr>
<tr><td>Execution</td><td colspan="2">Research & exploration</td><td colspan="2">Development & validation</td><td colspan="2">Operationalization & maintenance</td></tr>
</table>

Figure 1. The elements of an AI strategy

Define Your AI Ambition: Where to Play

The AI ambition sets high-level goals for AI application. Ambition should be aligned with overall corporate strategy. Fundamentally, you need to determine the area(s) in which AI will be able to give you a competitive advantage in the future, the fields of action where your company could benefit the most from AI, and the level of commitment needed to successfully implement the strategy.

Yet before defining the company's AI ambition, it is important to assess the status quo. Where lies the company's competitive advantage? What are its core products and services? How do its competitors use AI? How do AI developments on the market change customer expectations and needs?

An understanding of current competitive advantage will help to make predictions about your industry's future dynamics. The external environment will greatly influence your AI ambition. It is crucial you understand the **future competitive advantage** that AI can bring and how it will impact your products and processes, your business model, and the industry more widely. As part of this process, it is important to consider not only the incremental changes possible through AI but also the opportunities for completely rethinking established processes or products. An anecdote from history provides a useful illustration. In 1819, the SS Savannah became the first steam-powered vessel. The engineers at the time equipped the traditional sailing ship with steam-powered wheels on its sides. However, merely applying this new technological advancement (steam-power) to a traditional vessel turned out to be neither profitable nor functional. The vessel was hard to steer so the steam power was hardly used. The situation required a dramatic rethink about the engineering of the vessel, and it took another 20 years until regular transatlantic steam service was established.

Similar to the engineering process of the SS Savanna, AI offers an opportunity to take an established business model and improve it through incremental enhancements. However, much higher business value might be possible by exploring more dramatic shifts in value pools and disruptive AI. For instance, the customer journey of a typical insurance company might include claim submission, document processing, claim adjudication, claim evaluation, and fraud detection. When thinking in terms of incremental improvements, the insurer would most likely come up with AI use cases such as intelligent document generation, chatbots, or ML-based fraud detection. However, considering more disruptive AI-facilitated changes in the business model might lead to a realization that document processing and claim adjudication could be fully automated and integrated, and hence no longer be steps in the customer journey. This is not a hypothetical example, as emerging AI-first insurance companies have already made such realizations.

Key questions to ask yourself at this stage are:

- **What value drivers are enabled through AI and how can they be exploited?**
- **How do competitors use AI? How do organizations in adjacent industries use AI?**
- **What core sources of value could go AI-first or be fully automated?**
- **What defendable value-adds (beyond assets) exist without AI? Engineering? Production? Marketing/Pricing? Sales? Support?**
- **What transformational dynamics might AI trigger?**
- **What can be learned and expected from existing players?**
  - **What are the most disruptive business models you have seen in related fields?**
  - **What are the most dangerous lateral attacks you could imagine?**
  - **What are the most interesting start-ups you have come across?**

The next step is to define **fields of action**, which involves exploring where your organization would benefit most from AI and where it is likely to create business value in the future. Is it through selling a specific product or service or by improving your processes, or perhaps both? Understanding the distinct ways that AI could transform processes, products, and services within the company is crucial, as each may involve different stakeholders and require different resources. Companies just getting started with AI should focus on the application they expect to have the highest value as available resources are typically limited.

One field of action might involve augmenting existing products or services with AI, for example by adding a voice-based assistant to a car. It might also involve creating new AI-powered products, like an insurance company offering a service that uses the technology to automatically assess damage from images.

Process-centric AI applications could, for example, involve algorithms that predict demand for goods based on parameters such as time of year, weather, cost of complementary and substitute goods, and market movements, or an algorithm that pre-screens potential job candidates.

Additionally, companies should be conscious of the different horizons of AI application. These can be broadly categorized into two – those that are already mature but essentially incremental, and those that are still experimental but have the potential to have a disruptive impact. For instance, capabilities like computer vision, natural language processing, and speech recognition are already mature and can be readily applied in various domains. These capabilities have been refined through years of research and development and have already demonstrated their effectiveness in many real-world applications. Other capabilities, however, such as reinforcement learning, are still in experimental stages and require more research and development yet have the potential to have a significant impact on industries where there are complex and uncertain environments.

When considering the overall 'playing fields' of action in AI for a company, it can be helpful to view this as four distinct quadrants, as shown in **Figure 2.**

Finally, it is important to agree on the level of **commitment.** Set clear and measurable objectives and commit and allocate resources towards AI implementation.

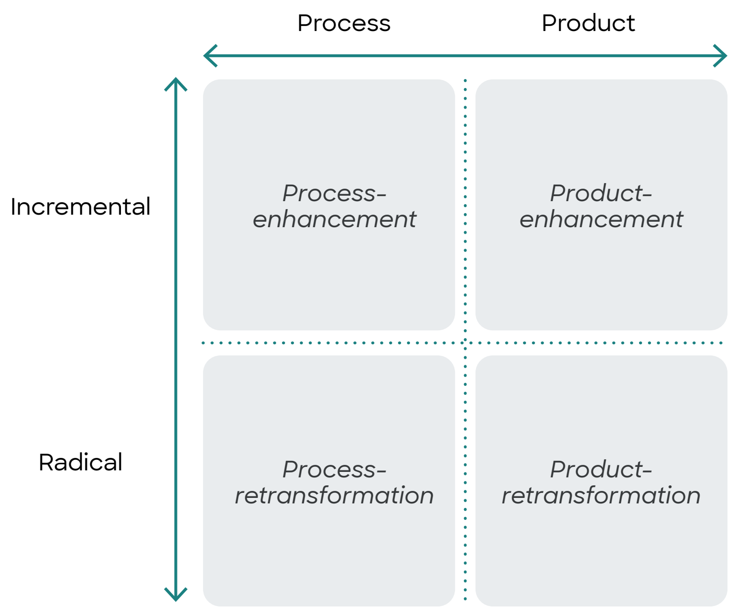

Figure 2. Playing fields of action in AI

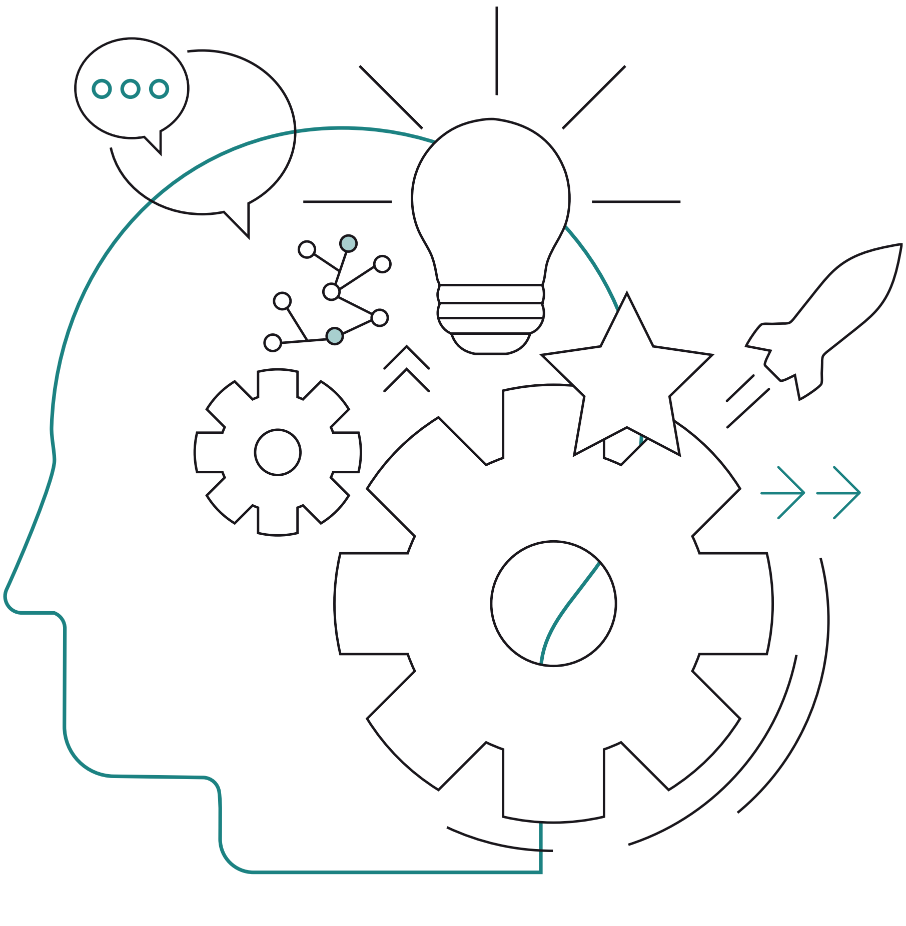

Finding and Managing AI Use Cases

A company's AI ambition sets high-level goals for AI application. To generate value, this needs to be translated into a portfolio of AI use cases. Building this portfolio requires identifying and prioritizing relevant use cases, developing a clear understanding of when to make-or-buy a use case[1], and maintaining a portfolio of use cases.

Using a Systematic Approach for Finding, Assessing, and Prioritising Use Cases

To identify and evaluate relevant use cases, companies can follow a four-step approach:

1. First, potential use cases are identified, either building upon strategic goals (opportunity-driven) or existing strengths (asset and capability-driven). Guiding questions are:

- Which use cases do you need to implement to execute your AI vision?
- What data do you have?
- Which AI capabilities have you already developed?

In this step, it is important to know the issue you want to solve. This may sound obvious, but many companies, often driven by a fear of missing out, start out using machine learning and then identify problems as they come, rather than the more beneficial method of identifying problems before starting. You should always ask yourself whether using ML is the right and easiest solution for a specific problem. Adding AI to a bad process will not make it a good process.

2. Subsequently, all potential use cases need to be defined clearly and comprehensively. They must be assessed in terms of potential economic value and ease of implementation. Consider potential AI red flags, i.e., implementation aspects that would substantially increase the effort (or decrease the value) of the use case. Red flags include regulatory and ethical concerns, as well as cybersecurity or black swan events that could impact the use case.

3. After this step, potential use cases should be prioritized based on their potential value and opportunity. Start with those that are relatively easy to implement and would provide high value. It may not be necessary to disregard high-value use cases that are difficult to

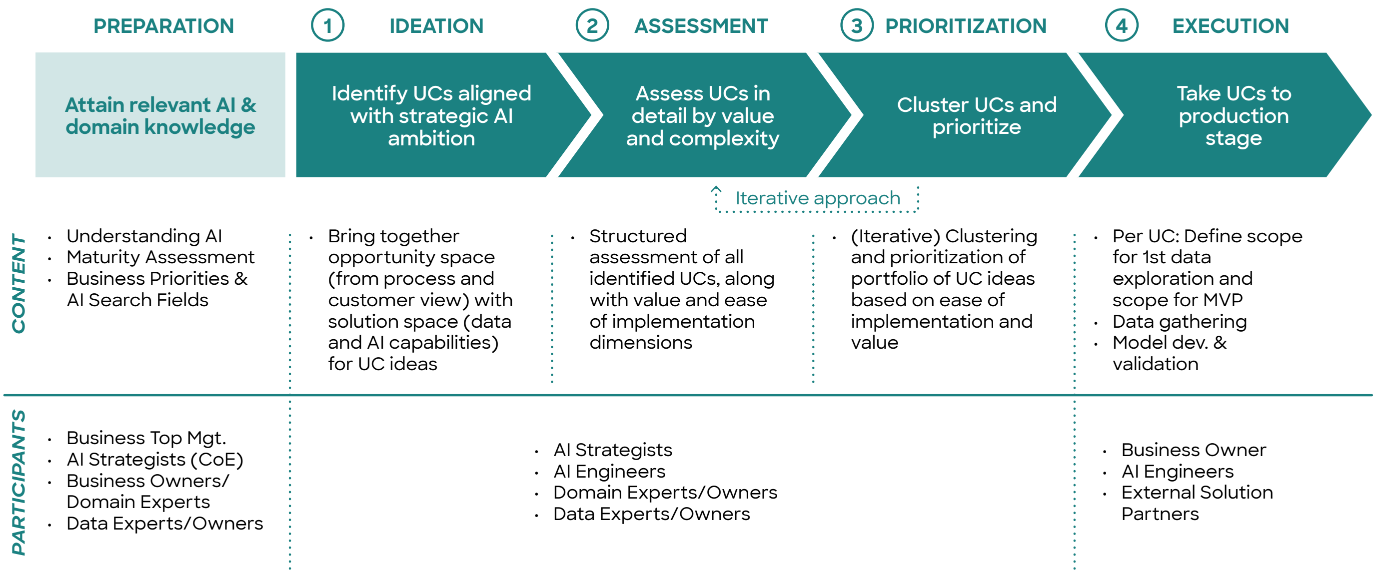

Figure 3. A four-step approach to identify and evaluate relevant AI use cases

1 An AI use case is a set of activities required to reach a specific goal from a business or customer perspective, which involve a substantial application of artificial intelligence.

implement. Use cases can sometimes be clustered by interdependencies or broken up into smaller use cases that are easier to execute.

4. Finally, selected use cases can be executed. Companies should develop an execution strategy, taking an explorative and iterative approach to building an initial proof of concept based on the use case's potential value and internal capabilities. Once the use case has been validated, it needs to be brought into production and maintained to deliver continuous value. **Figure 3** provides further information about use case execution.

Developing a clear understanding of when to make-or-buy an AI use case

Deciding to apply AI does not mean that you should build everything yourself. Rather, you should make a conscious decision about which cases you want to implement. Consider what it would take to come up with a proof of concept and to build a minimal viable product. When thinking about working with an external partner or purchasing a tool, note that AI is rarely plug-and-play and often requires substantial (re)training of algorithms. Concerns may arise when suppliers require access to internal data and the resulting algorithm contains critical business knowledge. In such situations, deciding on contracting and managing the relationship can have complex and unique issues. When balancing the pros and cons of make-or-buy, you should ask the following questions:

- **For building and managing the AI application, do you want to do it yourself or partner with others?**
- **Who should you partner with?**
- **How should a partnership or contract be structured?**

If you want to dive deeper into the topic, please refer to appliedAI's whitepapers "Enterprise Guide for Make-or-Buy Decisions" and "A Guide for Large Language Model Make-or-Buy Strategies: Business and Technical Insights".

Managing a Use Case Portfolio

Once possible use cases have been identified, a use case portfolio should be managed effectively. This involves ongoing evaluation and prioritization of use cases based on their potential impact, feasibility, and alignment with business objectives. It is also important to regularly assess external and internal factors that might impact viability and adjust the portfolio accordingly. This approach ensures that AI investments are focused on the most promising opportunities and that resources are allocated effectively. Managing the use case portfolio can also help organizations identify potential risks or dependencies between use cases and take proactive steps to mitigate them. In doing so, organizations can stay ahead of the curve and maximize the potential of AI to drive growth and innovation.

For further guidance on identifying and prioritizing AI use cases, please refer to appliedAI's whitepaper "How to Find and Prioritise AI Use Cases".

Enabling Factors

To implement the company's AI ambition and execute AI use cases, several enabling factors are needed. These include organizational set-up, expertise, culture for AI, technology, data, and ecosystem. The following sections provide guidance for approaching each dimension and suggest additional resources.

Organization

When electricity superseded steam engines and water wheels in factories, the initial impact on productivity was very small. It was only when the outlay of the factory changed and the rigid requirements of mechanical links were removed that productivity really took off. The way employees worked together also changed fundamentally. AI presents a similar situation. To successfully leverage AI, organizations must create a suitable environment or structure for it to thrive. There are three guiding principles to do this:

Manage AI solutions as a product, not project. Traditional project management approaches are not sufficient for AI and machine learning solutions, as these technologies have unique characteristics that make them more complex to manage. One such characteristic is that the potential outcome of an AI or ML solution can only be judged after the development phase, making it difficult to define project parameters such as a predetermined outcome, end date, and fixed budget. Moreover, AI solutions are never really finished as they require continuous maintenance and improvement. To manage this, organizations need to establish a dedicated team responsible for the ongoing development and integration of the AI product, rather than treating it as a one-off project. This team should be responsible for the entire lifecycle of the product, including its maintenance and retraining as necessary. Because changes in data can necessitate system retraining, difficulties in the handover process usually result from this retraining rather than changes in the lines of code. By adopting a product-oriented framework, organizations can better manage the complexity of AI solutions and ensure their continued success.

Set the right balance between central coordination and decentralized ownership from the start. Certain aspects of AI are best managed centrally to prevent teams from duplicating efforts and promote collaboration. On a larger scale, centralization helps an organization align in terms of direction, with leaders enabling generation of an overarching AI strategy. Leaders also assign priority to the most pressing AI projects. However, when organizations centralize AI responsibilities, they risk building isolated solutions without real business value. As such, the most successful AI initiatives set a balance between centralized and decentralized activities. Use cases should develop in a bottom-up way from business units where people understand customer needs and process pain points. Leaders need to create enough room for local, decentralized initiatives, but these initiatives should draw on centralized knowledge and resources to tackle identified challenges. Central coordination dramatically increases efficiency and effectiveness by:

- Consolidating best practices on AI tools
- Providing benchmarking of solutions
- Offering central repositories of training and test datasets, and reference models
- Facilitating easy access to preferred tools

One best practice approach we often see is a hybrid approach. Here, a central AI team, often called the Center of Excellence (CoE), bundles certain functions and expertise while maintaining strong links to decentralized units and the rest of the organization.

Demonstrate powerful leadership and broad commitment from an AI-educated C-level. Organizations should identify a leader who can balance the various interests and requirements involved. This individual should possess technical expertise in AI, as well as business savvy and strategic thinking skills. They will be able to act as chief evangelist for AI, convincing others within the organization to embrace new approaches and change existing business models, product development plans, and corporate culture. In working with other leaders, the AI leader should be able to effectively defend resources and priorities when competing demands threaten to push AI down the agenda.

It should also be recognized that AI is the responsibility of the entire management team, not just the individual in charge of the AI initiative. By demonstrating a commitment to AI at the organization's highest levels, C-level executives can set the tone for a culture of innovation and collaboration where AI is seen as a critical driver of business success. This commitment must be broad-based, encompassing the CTO, CIO, CEO, and other key decision-makers. With a strong AI-educated C-level leadership team in place, organizations can drive successful AI initiatives that deliver tangible business results and create value for stakeholders.

For further information please refer to appliedAI's whitepaper "Building the Organization for Scaling AI".

Expertise

Organizations need the right mix of talent to drive AI, translate business needs into AI use cases, and build, deploy, maintain, and scale AI to support the overall AI strategy. So, recruiting, attracting, retaining, and building talent is essential. To nurture AI expertise in the organization, it is necessary to hire and develop talent, and to support the organizational impact of AI by fostering culture and collaboration.

AI talent is still rare. Only a small fraction of college graduates is educated to work in AI or adjacent fields. The question, then, is how your company can go about finding and hiring this talent. In the first instance, you should ask yourself which kind of technical and non-technical roles are required for your AI activities. **Figure 4** gives an overview of what the distribution of AI roles in a company might look like. Roles can encompass multiple functions and skills, ranging from data scientists, data engineers, or ML engineers to business consultants and domain experts. Note that roles that bridge gaps between business and technology, such as AI strategists or AI product managers, are crucial.

Not all of these roles are necessary, especially at the beginning. However, you should make sure to identify the skills needed to successfully realize your AI projects, thinking beyond the technical expertise needed to also consider relevant strategic skills.

Roles can be staffed in three ways: built, acquired, or rented. You could internally develop AI talent by offering interested and qualified employees the necessary time and resources to upskill. It is now relatively straightforward to learn about AI with courses from universities such as Stanford being freely accessible online, while platforms such as Coursera and Datacamp offer sophisticated AI coding courses.

Culture

Successful AI implementation requires close collaboration between AI and humans. Even the most advanced AI solutions will be ineffective if they are not adopted by users or embraced within organizations. Culture determines the speed at which new directions or strategies like the use of AI take hold. In a recent study of Fortune 1000 C-suite executives, over 92% identified culture as the biggest hurdle to deriving value from AI[1]. With this in mind, organizations need to answer a series of organizational and cultural questions as well as building technical and infrastructure capabilities.

1 Bean (2021). Why Is It So Hard to Become a Data-Driven Company? Harvard Business Review. https://hbr.org/2021/02/why-is-it-so-hard-to-become-a-data-driven-company

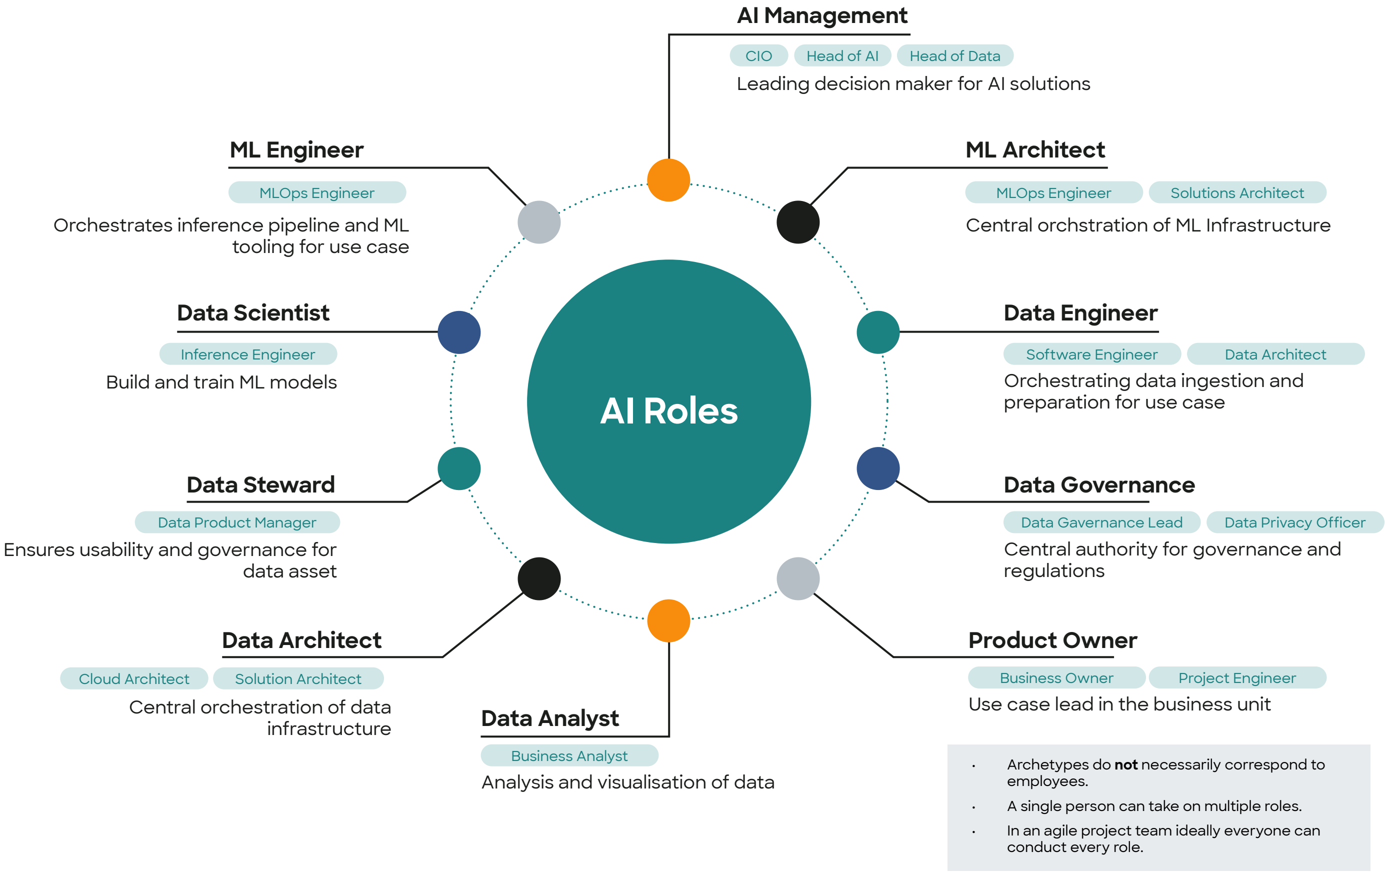

Figure 4. Overview for allocation of AI roles : Exemplary role profiles, alternative titles and task descriptions. These role profiles are not exhaustive or mutually exclusive.

AI can potentially impact organizations at a much deeper level than any other transformation. A common mistake is to think that AI transformation is just another digital transformation. There are many reasons why this is not the case. First, AI can impact the entire value chain and fundamentally shift employee role profiles. Second, with probabilistic systems, issues of trust arise and need to be met by providing transparency wherever possible. Third, AI is not limited to IT. Ideally, experts in a central, dedicated team work with decentralized extensions in business units to implement. Fourth, with AI roles spreading across different skill sets and roles, cross-functional collaboration is a key success factor. Open communication across departments is needed to foster organizational learning and efficient knowledge exchange. Fifth, AI solutions should be treated as products rather than projects. Lastly, AI and the complexity that comes with it require specific talent and expertise.

A two-folded approach for overcoming people-related challenges

Because AI transformation is fundamentally different from earlier digital transformations, specific measures are needed to meet AI-specific challenges and foster culture and collaboration. Organizations are recommended to follow a systematic approach to overcome cultural obstacles to AI adoption that places humans at the centre. We suggest a two-fold approach. On the one hand, development of AI applications should follow a human-centered approach accompanied by a change plan for implementation. On the other, the entire organization must be enabled.

Build your AI application with people in focus. Supporting AI implementation with suitable change measures requires adaptation, whereby processes affected by the AI application are adapted so that new procedures make sense, training, so that personnel working with the AI application can understand and use it, and incentives and management for correct use. This last point might be the least obvious, especially for decision-makers who are convinced of a particular solution. This, however, can be a make-or-break factor.

Empower your organization. To sustainably establish AI in your business, the organization should be empowered by following five steps: (1) helping people to understand the direction, (2) changing management processes, (3) generating momentum, (4) building skills, and (5) setting up the right structure.

First, your leaders and employees need direction regarding AI ambition and what it means for the organization. They need to understand the rationale and embrace the why. Tools are rarely the problem, as most problems arise from a lack of understanding. A new emphasis on discovery, creativity, cross-functional collaboration, and data-driven business needs to be broadly embraced.

Second, because AI deals in exploration space, there should be changes in how budgets are allocated and renewed. Rewards cannot be given on KPIs such as milestones or revenue. They need to be shifted to innovation and exploration measures.

Third, many organizations become sceptical about technological change as many have previously come and gone, failing to deliver on their promises. Responding to this challenge of morale requires building momentum and focusing on easy early successes, as Google did with their first AI use case in speech recognition outlined above. Such successes need to be communicated well and staff need to be supported in seeing the path.

Fourth, a learning approach needs to be put in place as AI can only be scaled with an understanding of its general working and the unique demands it places on the organization. AI literacy is particularly important for management, for whom a deep understanding is key for guiding the organization and deciding on an approach for value creation. Many AI pioneers have commented on the need for so-called ‘translators’ in organizations who can bring AI, IT, and domain experts and leaders together.

For a deeper explanation including case studies and step-by-step instructions on cultural topics please refer to appliedAI’s whitepaper “Culture, Change, Communication”.

Technology

AI needs proper IT infrastructure and technology, which requires critical hardware and software decisions. These include whether to use proprietary servers and GPUs or rely on the cloud, and which software solutions to use for building and scaling AI. These decisions will depend on organization-specific factors such as budget or security concerns. For highly sensitive projects, working in the cloud might not be an option. While cloud solutions offer convenience and enable rapid development with out-of-the-box solutions, this comes at the price of vendor lock-in and potential price increases. Such decisions need to be carefully evaluated with consideration of the organization’s resources and requirements.

Lack of the right resources will severely impair the ability to deploy AI. While these are standard decisions in enterprise IT, the dependency of data is special for machine learning. In machine learning, the world is modeled not by developing complex algorithms, but by utilizing insights hidden in data. A machine learning model is a way to organize and understand these insights. Machine learning needs a common framework to prevent it from becoming a random collection of broken or outdated experiments. This framework is Machine Learning Operations (MLOps). It's a structured process throughout the machine learning lifecycle. And it serves to obtain and ingest relevant data of sufficient quality while adhering to governance principles, keeping track of different experiments, model architectures and their respective performances in the modelling stage, as well as guiding through the deployment phase and implementing continuous monitoring of production performance.

Data

AI thrives on data. Data is needed to train algorithms, as input for the algorithms to function, and as feedback to improve AI model accuracy. Data availability is of course important, but quality also matters. A well-defined data infrastructure can make it easy to access and extract value from data. To create one, you will need to identify data sources, build data pipelines, clean and prepare data, identify patterns in your data, version the data, and measure results. For scaling AI, it is important to generate data catalogues and make sure that everyone in the organization knows what data is available and how to access it.

These are not trivial tasks and it is important to face some potentially inconvenient truths about data for AI projects. Creating value from AI requires getting data in the appropriate form for exploration, modeling, and operations. This can be difficult for organizations of all sizes and across industries. Also, data is not generally good or bad, but rather can only be good enough for specific use cases. Beyond these considerations, an efficient data strategy needs alignment of appropriate tools, processes, and people. Frustration can often arise from the fact that much of the value of data investment is not measurable in terms of direct revenue attribution, such as through mitigated risk, better product, or operational decisions.

These obstacles can be overcome. Organizations should identify key data based on a solid understanding of high-value use cases and AI ambition, with a use case portfolio that identifies common data needs. It should be noted that data quality comes at a price and there should be a clear business case to justify data quality initiatives. Again, it's advisable to start with use cases and have a tiered system, remembering not to treat all data equally. Further, data should be managed as a product, with a clear data owner and established agreements for data and its quality. Similar to working with code, a systematic review process for schema changes can ensure traceability and structure. As in the product development process, you should start with building and deploying a data Minimum Viable Product (MVP) that delivers a specific data component for a desired use case. These capabilities can later be scaled to other use cases.

Ecosystem

AI at scale is a comparatively young discipline that nevertheless involves many actors. An ecosystem is developing around AI, composed of private and public research laboratories, communities and interest groups, startups, universities, big tech companies, and various other organizations and investors. To successfully work with AI, cooperation with others in this ecosystem is key. Not everything needs to be done in-house. To define its AI ecosystem strategy, an organization needs to address the questions of which players its AI ecosystem should consist of, which partners can help in which situations, and how best to work with these partners.

There are generally four types of ecosystem partners. The first are service providers, which might include analytics or data service providers but also joint venture programs with government organizations to provide data and capabilities. The second are startups or startup accelerators, who can be approached to access innovative solutions and talent. Third are universities and research laboratories, which can provide access to the latest research developments as well as talent. The fourth are competitors, who can be collaborated with in co-opetition. By working with a competitor, your organization might get data access much more efficiently. Examples might be a company in another industry interested in the same AI technology, or a third-party interface that multiple players in an industry contribute to.

When deciding on a cooperation partner, your priorities regarding make versus buy should always be kept in mind.

Image created with DALL-E with prompt "abstract illustration on elements of AI strategy, semi-realistic abstract style, muted color palette"

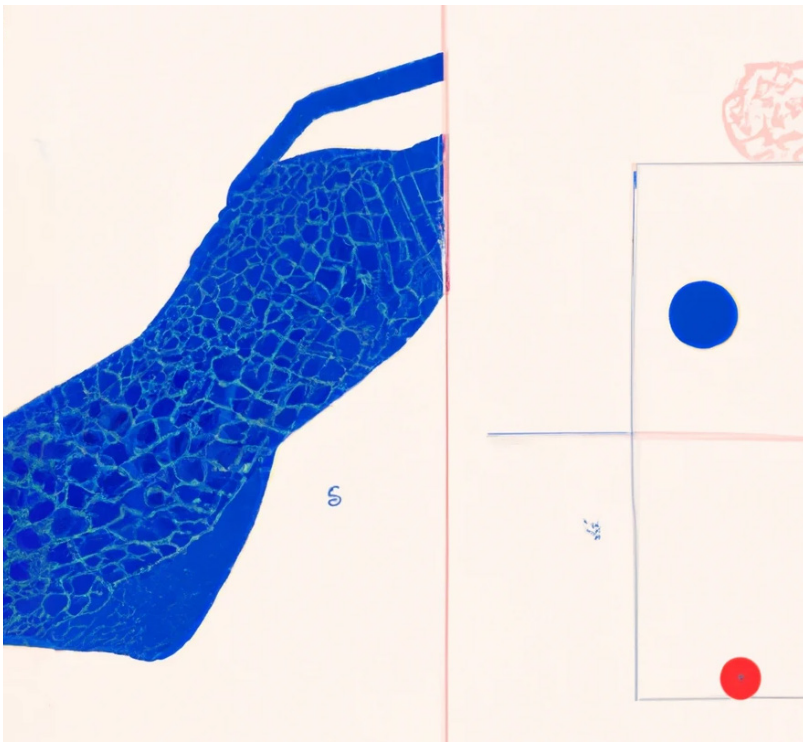

Execution

When suitable AI use cases and projects have been determined, they need to be executed and brought into production. The execution process can be complex and involve many steps, including data collection and preparation, feature engineering, model training and evaluation, model deployment, serving and monitoring, and model maintenance. A range of people need to be involved, from data analysts, data scientists, data engineers, software engineers, infrastructure engineers, and user interface designers to business consultants, AI strategists, and domain experts.

The process of developing, training, and deploying a machine learning model using huge amounts of data is summarised in the ML Lifecycle shown in **Figure 5.**

The goal of the ML lifecycle is to support planning and mitigate challenges during the development of an AI project. Consideration of this lifecycle during project definition and execution will support planning of the steps each organization needs to take. These steps guide the process of making a system work effectively while preventing unintended outcomes.

For a comprehensive explanation of the ML Lifecycle, the different stages and requirements please refer to appliedAI's "The Enterprise Guide to ML".

Execution can be roughly separated into three phases: Research and Exploration, Development and Validation, and Operationalization and Maintenance.

Research and Exploration. This phase deals with the clear definition of goals and benchmarks against which the model will be assessed and the forming of assumptions and hypotheses to create a proof of concept. Desired business objectives and outcomes, as well as specific user requirements, need to be translated into metrics and KPIs for machine learning. Training data is fundamental in how well a model performs. This phase should, therefore, include gathering and preparing data used to teach the algorithm, which can be a laborious process. Additionally, it involves researching any known implementations of the use case, so any relevant knowledge can be transferred to the project. This phase mainly involves domain experts, business analysts, data engineers, and data scientists to show the general feasibility of the project using initially available data. The research and exploration phase is time-boxed, ending after a set duration with a decision of whether or not to pursue the use case. This avoids putting too many resources into infeasible use cases.

Development and Validation. The goal of this phase is to develop a minimal viable product to validate the value of the use case. It involves using data to train, build, and validate the AI model. Data scientists look to see how well the models predict outcomes based on gathered data, and so will need to experiment with various architectures. They will then seek to retrain and improve models based on the predicted outcome. This is an iterative process and may require substantial time and be computationally expensive, depending on the complexity of model training and retraining. The aim should be to demonstrate a working product that can gain feedback from stakeholders.

Operationalization and Maintenance. Once the model has been trained and has reached a certain level of accuracy, it can be transferred from the lab to the real world. Writing code for AI models is just a small part of the ecosystem. Deployment is a large process that involves software engineers, infrastructure engineers, and user interface designers. For many engineers, the project really starts here. MLOps best practices should be heavily considered during this phase. After the model has been deployed, it needs to be monitored and maintained as its

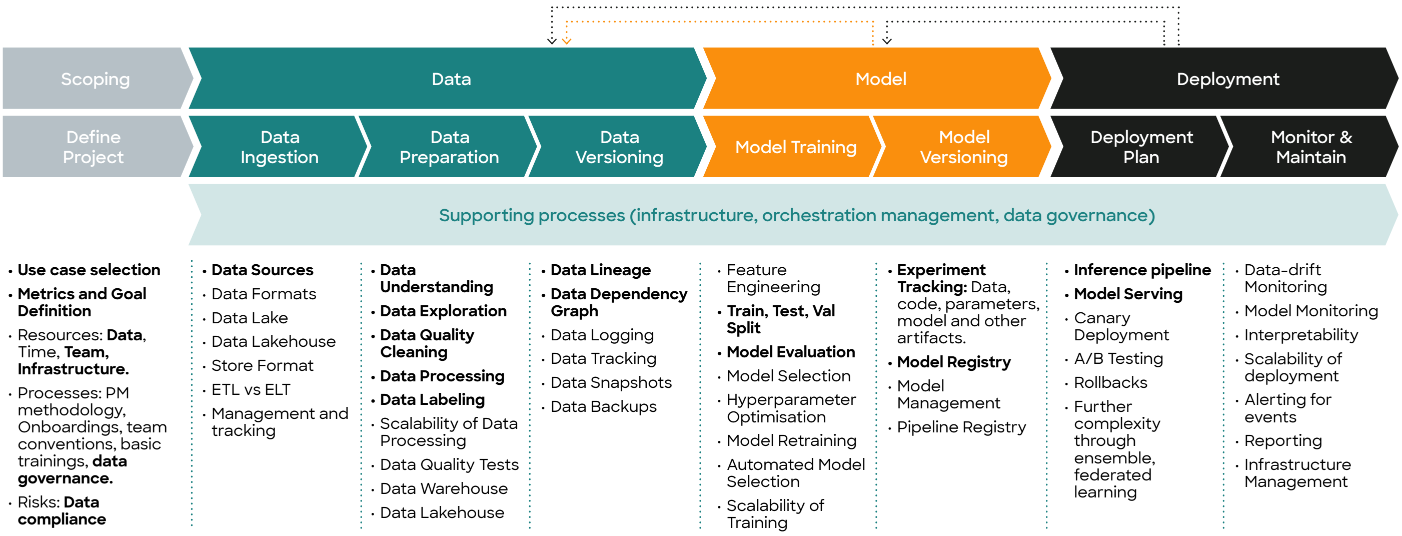

Figure 5. ML Lifecycle

performance could degrade over time due to factors such as data drift or environment changes. Systems that deal with constantly changing data or environments never really 'go into' maintenance but instead create a continuous cost and need for adjustment. Such instances can subvert the viability of each AI use case. Aspects to monitor include service health, data quality and integrity, data drift, general model performance and performance by segment, bias, and fairness, and how the system handles outliers.

The nature of the project changes throughout the execution process. Uncertainty is reduced through progress over these phases and necessary activities turn from being experimental and research-focused to more traditional software engineering.

Getting There Step by Step: The AI Journey

Applying AI is a journey. You and your organization start at one point and develop over time. Numerous companies have already started to embrace AI and are reaping its benefits. However, applying AI at scale has turned out to be a complex endeavor that requires long-term planning and engagement. Application can be a transformative process. Based on experiences with partner companies, appliedAI has mapped out the stages that companies go through, from first experiments with the new technology (what we call Experimenter Level) to enterprise-wide application of artificial intelligence at scale (Professional Level) and even shaping the overall market and AI ecosystem (Shaper Level). We call these stages the AI Journey.

The challenges and relevant questions related to each stage are dynamic and will develop as companies reach AI maturity. Organizations just starting with AI face different challenges regarding strategic dimensions than those that have been working with AI for some time. Furthermore, there should be appropriate development at each stage. Investing in compute infrastructure or building an ML team will likely not create value if there is a lack of understanding about the right use cases or required data.

Figure 6. Stages of AI maturity

There are four stages of AI maturity:

The AI Experimenter

Experimentation is the first step of every AI journey. In companies at this level, interest in AI has been sparked, and some degree of lobbying for AI occurs across the organization. The organization realizes the need for a vision that embeds AI and for a strategic approach to pursue that vision. These first steps will involve selecting and experimenting with the first use cases, eventually solidifying the need for AI technology. The goal for experimenters is to establish common understanding, interpretation, and expectation of AI. Common challenges at this stage will be a lack of management backing for AI initiatives, a lack of expert knowledge to drive AI vision development, and unclear responsibilities for steering AI activities. It is also typical that required skill sets are not yet available within the organization, with the company's data scientists not yet understanding AI sufficiently. For a case study of an AI Experimenter, see this article about appliedAI's work with the World Food Programme.

The AI Practitioner

The next stage in the AI journey is the Practitioner level. Operationalization of AI is a major challenge faced with every AI strategy and implementation. Practitioners begin to systematically develop and implement their AI strategy, focusing on optimizing processes or reinventing products. This should involve a principled approach to identifying opportunities, setting up central AI teams, and establishing first collaborations. Practitioners also set up training programs for the organization and contribute to forming a data strategy. At this stage, change management helps to reduce barriers to AI adoption, technical prototypes are taken from proof of concept to a running service, and processes are adapted to allow for the new development. Challenges for the AI Practitioner include no clear view of what will become the organization's competitive advantage when it comes to AI. There might also be too many use cases for the central AI team to deal with. A lack of alignment between teams could result in work being carried out redundantly. Further, the company might not yet be viewed as an attractive employer for AI talent, making it difficult to hire the required expertise. For a case study of an AI Practitioner, see this article about appliedAI's work with energy provider EnbW.

The AI Professional

Once first activities and systems are operationalized, raising effectiveness and efficiency of AI-related operations across the organization is the central goal of the AI Professional. Focus should now begin to shift. AI vision and strategy are underlined by longer-term resources, cooperation, and allocation decisions. For AI Professionals, activities of opportunity, discovery, and innovation are adjusted to incorporate AI across the organization. Governance structures are enriched with liability and protection strategies, while internal talent and culture development is rolled out at scale. Central processes to maintain, monitor, and operate the AI solutions that drive the business are established. Professionals start to set the standard for AI. The challenges facing AI Professionals revolve around changing company culture. There might also be a slow maturation of systems, platforms, and tools to develop AI technology. A further challenge is the transition from a centralized to a more decentralized structure for AI activities. Professionals need to figure out how to efficiently serve a large number of business units. Finally, there might be increased computational requirements due to the scaling of AI within the organization.

The AI Shaper

Becoming a Shaper is the last stage in the AI journey. It's important to note that not every organization needs to become a shaper, so this will not be a goal for every company. For AI Shapers, AI is a natural part of the conversation and daily working life at both management and employee levels. The company DNA is transformed as AI Shapers truly scale AI, both internally and externally, which requires different adoption approaches. At this stage, stimulating the ecosystem and driving AI both horizontally and vertically become fundamental issues. From having a data strategy that is considered an active step of development to longer-term acquisition decisions and open source activities, Shapers face a distinct set of challenges. Conducting research, attracting talent, and generating ideas for monetization are key activities of AI shapers. They also develop new markets and customer segments with their innovative solutions and create and publish new AI development processes and tools. Their challenges lie in needing to guarantee AI applications are ethical, and in shaping regulatory discussions. Exploration of new solutions will likely lead to never-before-seen challenges. The management of AI shapers might face the challenge that they have to actively manage public opinions on AI activities.

Outlook

Successfully applying AI can seem like a daunting task. The field is developing almost exponentially, and it is hard to keep track of the burgeoning algorithmic advancements, use cases, initiatives, and startups. Fear of technological displacement, irrational expectations, and a general lack of understanding about the capabilities and limitations of AI can be huge challenges, and finding talent and partners can be difficult and expensive. However, in the mid and long term, the benefits outweigh the costs. There are already plenty of examples of organizations successfully applying AI to create and capture value. AI has tremendous potential to drive economic growth and fundamentally impact business models. To harness the power of AI, a clear and comprehensive strategy is needed. Even if organizations are only at the beginning of an AI journey, it is important to start acting now and create the right environment to successfully apply AI.

Authors

Philipp Hartmann serves appliedAI as Director of AI Strategy. Prior to joining appliedAI, he spent four years at McKinsey&Company as a strategy consultant. Philipp holds a PhD from Technical University of Munich where he investigated factors of competitive advantage in Artificial Intelligence.

Desiree Modic is Head of AI Strategy & Transformation at appliedAI, with experience in applying AI in an industrial context gained during her 6-year tenure at thyssenkrupp's internal management consultancy. She holds a degree in industrial engineering with a focus on mechanical engineering and informatics from the Karlsruhe Institute of Technology.

Susanne Klausing is a Senior AI Strategist at appliedAI with experience in AI strategy and use case development as well as AI regulation. Prior to joining appliedAI, Susanne studied Technology Management at the LMU Munich, the Center for Digital Technology and Management (CDTM) and the University of Oxford.

About appliedAI

appliedAI is Europe's largest initiative for the application of cutting edge AI. Our vision is to shape the European AI ecosystem as a trusted enabler and innovator.

With partners such as NVIDIA, Google, BMW, Siemens, Deutsche Telekom and many more, we have been strengthening and building the next champions in AI since 2018.

You can find more information about appliedAI at:

www.appliedai.de

Acknowledgements

The authors thank Alexander Machado and Max Mynther for their contributions.

initiative for applied artificial intelligence

Elements of a Comprehensive AI Strategy

appliedAI Initiative GmbH

August-Everding-Straße 25
81671 München
Germany
www.appliedai.de