


Applying AI: The elements of a comprehensive AI strategy



“In the past, a lot of S&P 500 CEOs wished they had started thinking sooner than they did about their Internet strategy. I think five years from now there will be a number of S&P 500 CEOs that will wish they’d started thinking earlier about their AI strategy.”

Andrew Ng

Travel back in time just 20 years and many industries look significantly different than they do today: To rent a movie for a relaxing night at home, you would have to go to a Blockbuster store rather than using Netflix or another streaming service. It was not Amazon but companies like Barnes & Noble, Toys “R” Us, and Sears who dominated retail. And Microsoft’s market position seemed unsailable until Google came into the world and began to dominate the Internet market. Many companies underestimated the potential of the Internet and had to watch their business models become irrelevant. Only a few learned to adapt; the others did not, and today nobody even remembers their names.

Today, just like the Internet 20 years ago, artificial intelligence (AI) is already here, even though you might not have noticed it yet. AI is often not obvious, forming part of other software or tools that companies already use. When you navigate with your smartphone, an AI-based algorithm will find the best route using real-time data from many other smartphones. When you book an Uber or Lyft, an AI-based algorithm will determine the price based on many factors including demand. When you want to watch a film on Netflix, you get recommendations for other films based on your previous film choices and on the choices of others – using an AI-based algorithm. And when you talk with your Echo speaker or your Google Assistant, AI-based algorithms are translating your voice into text, understanding your intention (or at least trying to), and delivering answers.

With more available data, greater computing power at a lower price, and by enabling technologies such as the Internet of things (IoT), AI has emerged to become a disruptive force for the years to come. **There is little doubt that AI will be relevant for all companies, regardless of their industry or size.** A recent McKinsey report suggests that AI could create an annual potential economic value of up to US\$ 5.8 trillion. This value potential is also reflected in the valuations of AI startups: For instance, Chinese AI startups SenseTime and Face++, both providing facial recognition technology, are valued at more than US\$ 3 billion each.

But where are you with respect to AI? Have you heard the starting pistol? Or are you still at the edge of the track, curiously or apprehensively watching what’s going on and not really knowing what to do?

You wouldn’t be alone: In a recent BCG survey, **only 20% of the surveyed companies said that they use a concrete AI application.** And those companies that apply AI do so to a varying degree – some apply AI within just some products or processes, some integrate AI into their strategy, and a few build their strategy around AI. So far, only a few companies already use AI at scale, predominantly the large technology companies. Google uses AI within virtually all its products, and in 2017 the company adopted its “AI first” strategy. Amazon’s enormous investments

in AI have made it the most valuable company worldwide and Facebook’s high market valuation is due to the same reason. And in China, companies like Tencent and Alibaba are aggressively investing in AI.

However, making a profit from AI takes more than simply connecting to an AI black box and feeding it terabytes of data. AI is certainly not a magic bullet that solves all problems. Therefore, companies need to first separate the hype from the actual capabilities and potential of AI, understanding how they define AI and how it might create value for them. Starting the process of moving an entire company towards the adoption of AI is a challenging task. It requires an understanding of the technology, as well as of its possibilities, and requires deep organizational and cultural changes. Furthermore, companies need to create the necessary enabling factors, including access to the right data, an adequate infrastructure, and the right talent – and not only in terms of programming skills.

That sounds overwhelmingly complex – but don’t get discouraged. AI certainly is not reserved for technology companies like Google and Facebook, especially as there is no need to reinvent the wheel. A lot of AI algorithms are readily available “off-the-shelf,” and there are many applications where relatively low-complexity AI applications can already outperform the current manual processes. The goal of this paper is to help you assess the potential of AI for your company and to provide guidance on how to develop an AI strategy. It is the first part of a multiple-piece series on AI strategy by the appliedAI initiative.

Getting started: What do we mean by artificial intelligence?

The term “artificial intelligence” was first coined in the 1950s. However, there has been a long, and still ongoing, debate on what is actually meant by artificial intelligence. Often AI is defined by referring to human intelligence – it is defined as computers acting “intelligently”¹ or when they “do things at which, at the moment, people are better.”² However, these definitions are not really helpful as human intelligence is not clearly defined, and also a simple calculator will most likely perform the multiplication of two ten-digit numbers better, and certainly faster, than a human.

Broadly defined, **AI refers to a machine that can learn, reason, and act independently when faced with new situations.** What we mean when we talk about AI today is artificial narrow intelligence, which includes more practical and readily available technologies. It has nothing to do with simulating real human intelligence but rather involves using math and algorithms to make machines predict answers – systems and algorithms that can perform one single task very well but can’t be reused for a different task or problem.

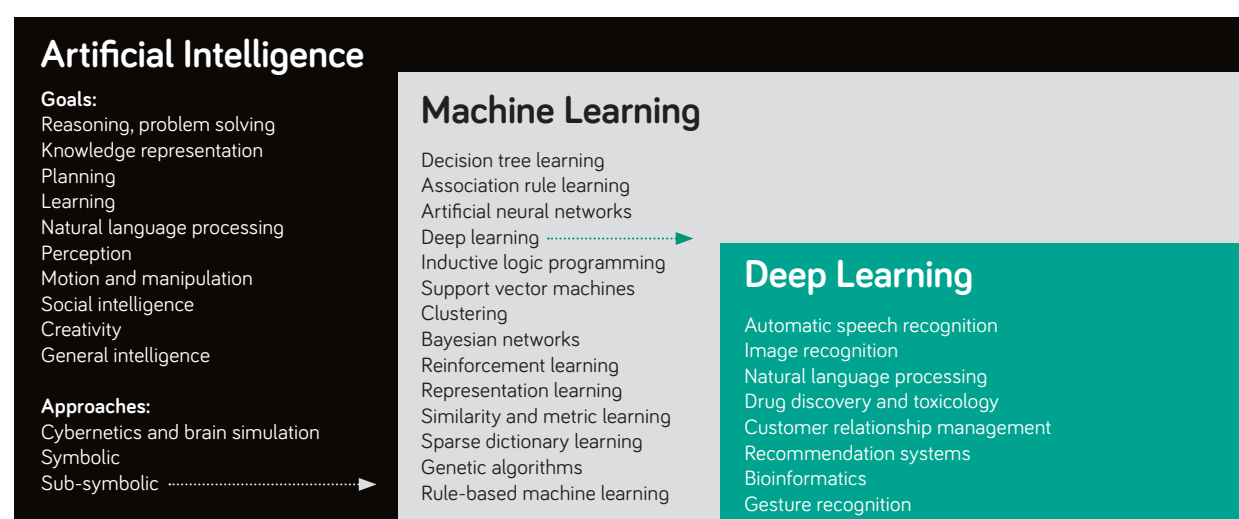
AI is an umbrella term that combines many different technologies and subfields. The bulk of the advances in AI are in machine learning, the most important subfield of AI. Machine learning uses statistics to find patterns in large amounts of data and uses these patterns to make predictions. The technical ideas behind most machine learning applications have been around for many deca-

des, but we have only recently been able to combine this form of AI with enough computational power and data to make it work really well.

At the heart of this technology is the idea of automating complex areas of knowledge work, using machines where we traditionally couldn’t due to the lack of “intelligent” capabilities. In particular, AI has advanced capabilities in six domains:

- Computer vision
- Computer audition
- Linguistics
- Mathematical/logical machines
- Interpersonally intelligent machines
- Motion and robotics

To understand the business value of AI, it is helpful to understand the ability of AI to make predictions more cheaply and accurately than ever before – an idea raised by Ajay Agrawal, Joshua Gans, and Avi Goldfarb in their book “Prediction Machines.” To understand this idea, consider the example of a translation model: A sentence is fed in in one language and the translation of it is returned in another language. In reality, the model is predicting the (most likely) sentence in the target language, given the input sentence. To decide if this translation is accurate still requires human judgement.



¹ John McCarthy, 1956

² Elaine Rich, 1983

You may well feel that this is thinking around a couple of corners. But in general, reframing problems that can potentially be solved by AI as predictions is quite a good way to approach possible use cases. And it is worthwhile: for classical predictions, such as finding indications of lung cancer on X-ray images, AI can help to complete the task more accurately, quickly and cheaply, which, in some cases, can even alter a business model. For example, when the value of human prediction decreases due to the fact that an AI algorithm is simply better at the job, employees and doctors will have more time to focus on the actions resulting from the predictions.

This way of thinking can also increase the use of AI, especially if one extends its usage to similar problems that until now have not been viewed or framed as prediction tasks.

At this point it should be clear that, for companies, the successful adoption of AI could enhance performance enormously and hence create a significant competitive advantage.

How does AI create value for companies?

AI can create value for companies in various ways: It can enhance productivity, optimize processes, create new

processes, and improve the customer experience.

This is reflected in the results of a recent survey: Companies are applying AI to pursue a wide range of goals. More than half of the surveyed companies' representatives said they wanted to use AI to make their existing products better. This is followed by optimizing internal business processes, freeing up workers, and making better decisions.


More systematically, the extent to which companies can use AI is defined by two dimensions – the field of application as well as the depth of internal AI knowledge:

The field of application. There are two AI-application fields:

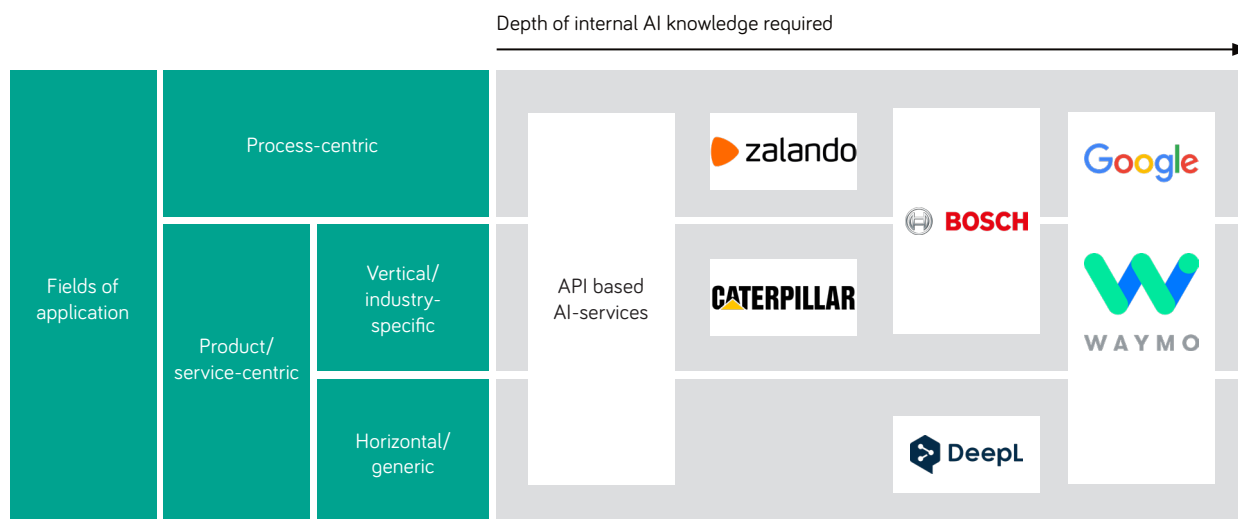
- Product/service-centric AI: Using AI to improve existing products or create new AI-driven products – either specific to one industry (vertical) or generic (horizontal)
- Process-centric AI: Using AI to improve or redesign internal processes

The depth of internal AI knowledge. Companies can use AI with a varying degree of internal knowledge – either doing their own research and development, taking part in joint developments with external developers, or using existing AI solutions.

Taken together, these two dimensions define a company's potential AI usage.



Source: Deloitte State of Cognitive Survey, August 2017



Let's take a closer look at a few examples: At one extreme, there is Google – a company that adopted an “AI-first” strategy in 2017. Google uses AI at scale, both for improving internal processes as well as within a wide range of products. These applications are fueled by heavy investment into AI research – Google employs more than 1,800 AI researchers in more than ten global locations. However, there are also examples of German companies using AI:

mentation, demographic predictions, pricing and price forecasting, invoice steering, and advanced fraud detection. Zalando has built up significant AI engineering and research teams, and employs more than 100 machine learning engineers who are embedded in different teams. There is also a dedicated Zalando research department with 15 researchers closely connected to universities and AI experts.

“Ten years from now, scarcely any Bosch product will be conceivable without artificial intelligence. Either it will possess that intelligence itself, or AI will have played a key role in its development or manufacture.”

Volkmar Denner, Bosch CEO

Bosch applies AI for improving processes, for example to optimize manufacturing or supply chain management, as well as to create new products. Bosch has also built up AI research teams with a staff of more than 300 in four global locations, investing more than € 300 million.

German fashion e-tailer Zalando uses AI to improve all their core processes across different business functions. Among them are product search and customer recom-

But, as mentioned previously, **AI is not restricted to the big players. Smaller companies can apply AI without starting from scratch – by using existing online machine learning infrastructures.**

Common pitfalls of AI applications

To create value from AI, companies need to approach its application strategically. Often, we observe the opposite: A company decides to “do AI” and begins by hiring a few data scientists and machine learning experts. They are then allocated to a newly created department, most likely on the third or fourth level of the hierarchy; they are then left to come up with some use cases.

Let’s say that these new employees, despite their lack of domain knowledge, come up with a couple of potential applications. Highly motivated, they go hunting for the appropriate data – and get nowhere. Other departments are unlikely to be willing to cooperate, either because of a “not invented here” approach, meaning not in that particular department, or due to other reasons. Maybe they fear that the AI project could be a threat to their authority, importance, existence, or even their jobs.

But let’s assume that the data scientists do manage to get some data; not the amount they wanted or in the quality they were hoping for, but enough to start a project. If they get this far, they will probably be confronted with the problem that AI projects are difficult to implement within existing budget processes; as AI projects contain uncertainty by design, it is difficult to estimate costs and quality before the project implementation has begun. If they give it a go, they will ask for the appropriate hardware: perhaps some GPUs or the all-clear to train their models in the cloud. And then they will be told by the IT department that all employees get the same computers, without exception. And that no data is to leave the company’s servers, without exception.

Does that scenario sound familiar to you? More systematically, we commonly observe five pitfalls when it comes to AI applications:

1.

Companies develop isolated AI use cases that do not create value for their business. A lot of companies start to implement AI use cases with neither a clear understanding of whether a particular use case is linked to the overall objectives of their organization nor a realistic estimate of the added value. After all, AI is not an end in itself.

2.

Difficulties in scaling use cases from pilot to production. In many cases, companies develop projects – but they have difficulties in scaling them up from a pilot project to being applied on a much larger scale. Developing pilot use cases in the lab is not always easy, but it is typically a lot simpler than moving those pilot applications into

production. One reason for this is data management, or the lack of it. In the lab, models are often developed with data from only one source, whereas in the real world the database has to be continuously updated, maintained, and optimized.

3.

Many companies lack the internal resources and capabilities to implement AI use cases – and they also have no idea how to bridge this gap. Which raises the question of why this is the case? The answer is quite simple: Just try and hire an experienced AI programmer in Cham, Berchtesgaden, Wedel, or Buxtehude. In fact, anywhere outside of Berlin, Munich, or Hamburg. Typically, non-IT companies lack both the talent and the infrastructure to develop use cases internally, which can develop into a vicious circle, for example if these companies invest in AI/data infrastructure without a clear understanding of the requirements of their potential applications and the suitability of their newly acquired hardware components for certain use cases. If they have underinvested, they might have further problems finding experienced programmers as the availability of decent GPUs is a determining factor for almost all experts seeking a job.

4.

Investment in AI/data infrastructure without a clear understanding of applications/use cases. However, overinvesting into infrastructure is not a magic bullet, at least not on its own. As well as having the hardware, you also have to get the related processes right. You need to govern access to data, and data protection must be taken into account – you also need a data strategy to apply AI.

5.

Implementation of use cases without a clear maintenance strategy. All use cases must be maintained carefully because every AI project requires an iterative approach. In general, AI project teams start off with a preliminary data set, build a first model, and evaluate the results. The evaluation provides hints on how to improve the model. Does the team need more or better data – or is a better model required? Depending on the outcome, the process is repeated and the model is evaluated again. And this cycle continues, even once the model has been put into production. There is always a need to update the data as, for example, customer behavior could change. And there is always a need to check whether the model is still providing the expected outcome.

The systematic approach towards AI

The very first step for applying AI has nothing to do with AI itself. You have to be clear about the overarching objectives or the purpose of the company: What is its goal? That sounds simple – but there are lots of examples where companies have applied AI projects that did not really support the overall purpose of the company but, for example, helped with a comparatively minor part of the production process. This could well be fine, but it could also turn out not to be economically feasible. Moreover, concentrating on side issues in the production process might lead to huge AI opportunities in the core business areas being missed.

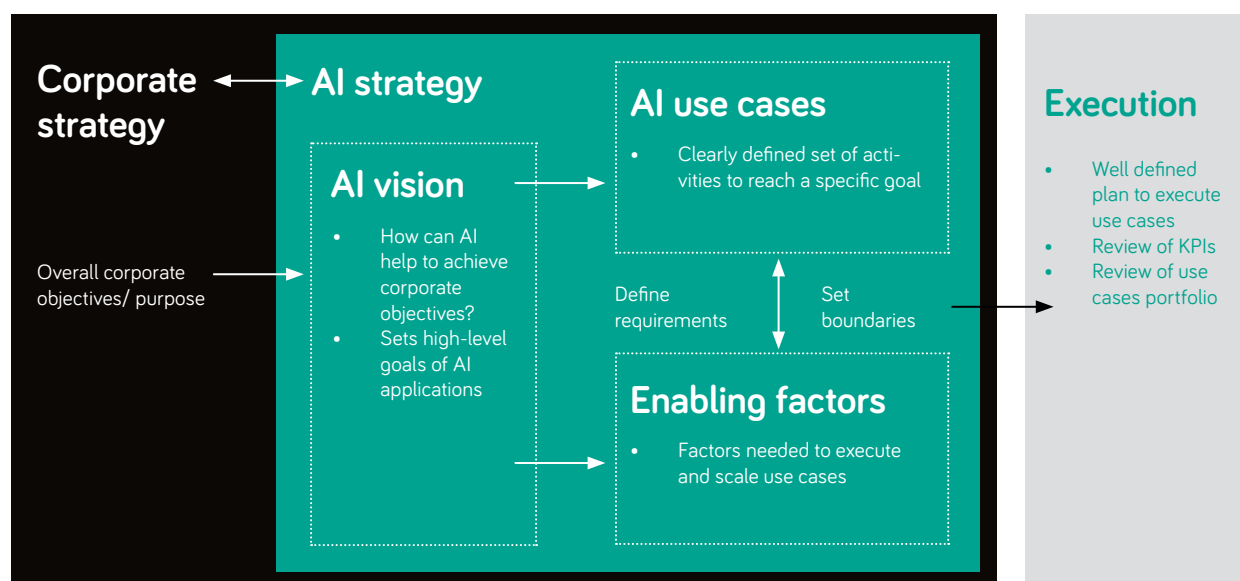
And this is only one side of the coin. The other is to understand how AI can help to achieve your objectives. Applying AI should not be focused on technology alone. Perhaps there is much greater potential for increase revenues via customer recommendations rather than improving the production process of the offered products? Finding out how to apply AI in the most effective way should not (only) be left to the data scientists. **It is a management task, and it requires executives to understand the principles and possibilities of AI.**

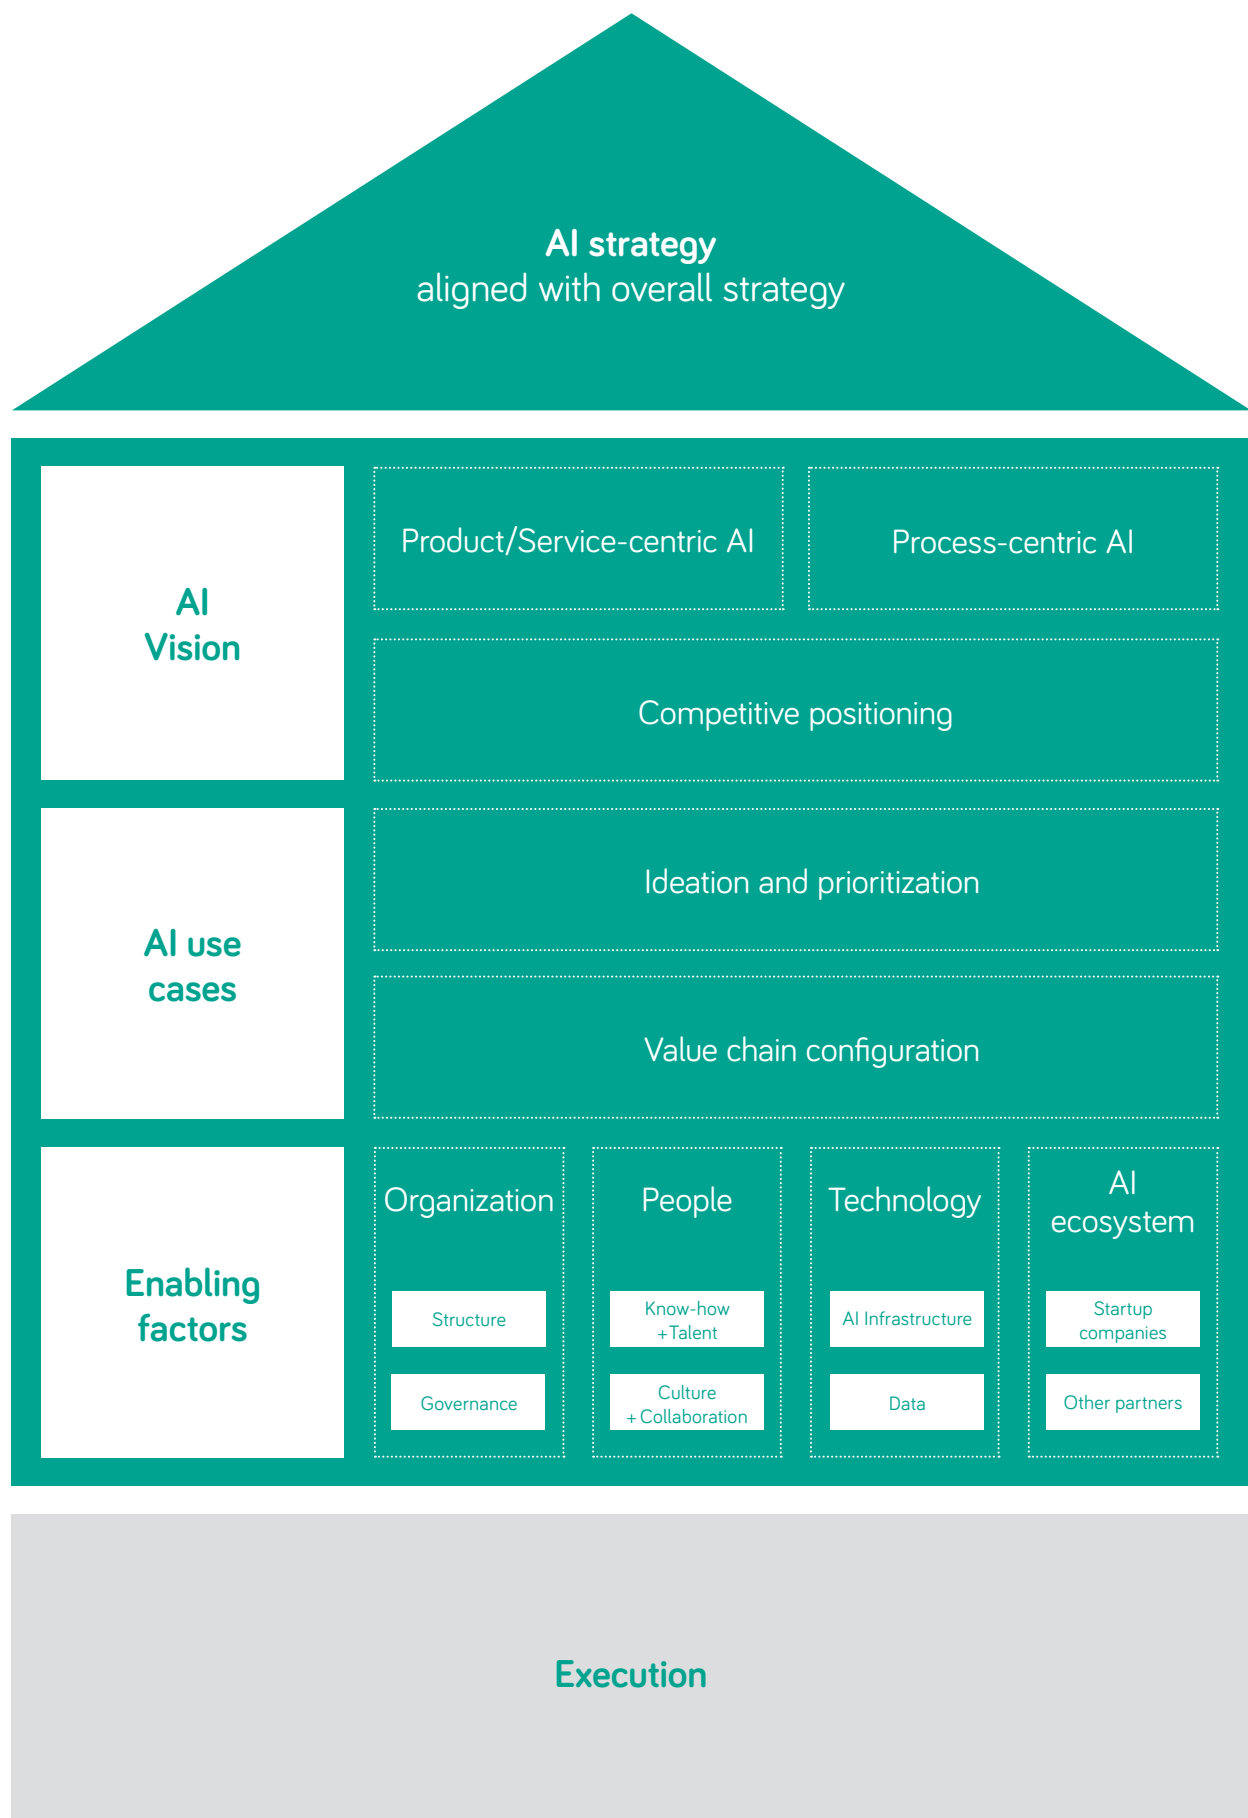
Once both of these prerequisites have been fulfilled, it is time to set up the company's AI vision, which sets the high-level goals of any AI application to be developed or deployed. Needless to say, this should also be in line with the overarching corporate strategy. The two should be closely linked – but also monitored. A company needs to understand how AI and “cheap predictions” might change its key sources of value. In quite a few cases, the successful implementation of AI projects or models has changed the overall corporate objectives.

At this point, it is time to get more specific. Based on its AI vision, a company needs to identify relevant use cases and also determine the enabling factors. This includes the organization, the people, the technology, and the AI ecosystem required to execute and scale the use cases. Again, this is not a linear process as the AI use cases and the enabling factors are closely linked: On the one hand, the use cases a company wants to implement determine the requirements for data, talent, and infrastructure. On the other, the available resources also determine which use cases can be implemented in the short term.

To handle and manage the iterative process of setting up and developing an AI project, a company should develop clear plans for how its AI projects are to be set up, run, and monitored, and also how the resulting AI-based solutions are to be further maintained in the future.

Putting the elements together leads to the “applied AI strategy house.” **Every company that wants to implement AI and avoid the aforementioned pitfalls needs to address all the elements comprehensively**, that is: Develop an AI vision, a portfolio of AI use cases, and a clear strategy for the required enabling factors.





How to get there, step by step

The following section aims to give some guidance on how to systematically approach all three pillars of the “strategy house” mentioned before.

1.

Define your AI vision

The first and central element of any AI strategy is the AI vision. It describes how a company wants to apply AI.

My company: The first step in drafting an AI vision is to understand the current position of the company:

- What are the overarching objectives of my company?
- What are my main products and process blocks?
In which of these do I have specific (secret) domain knowledge and assets?

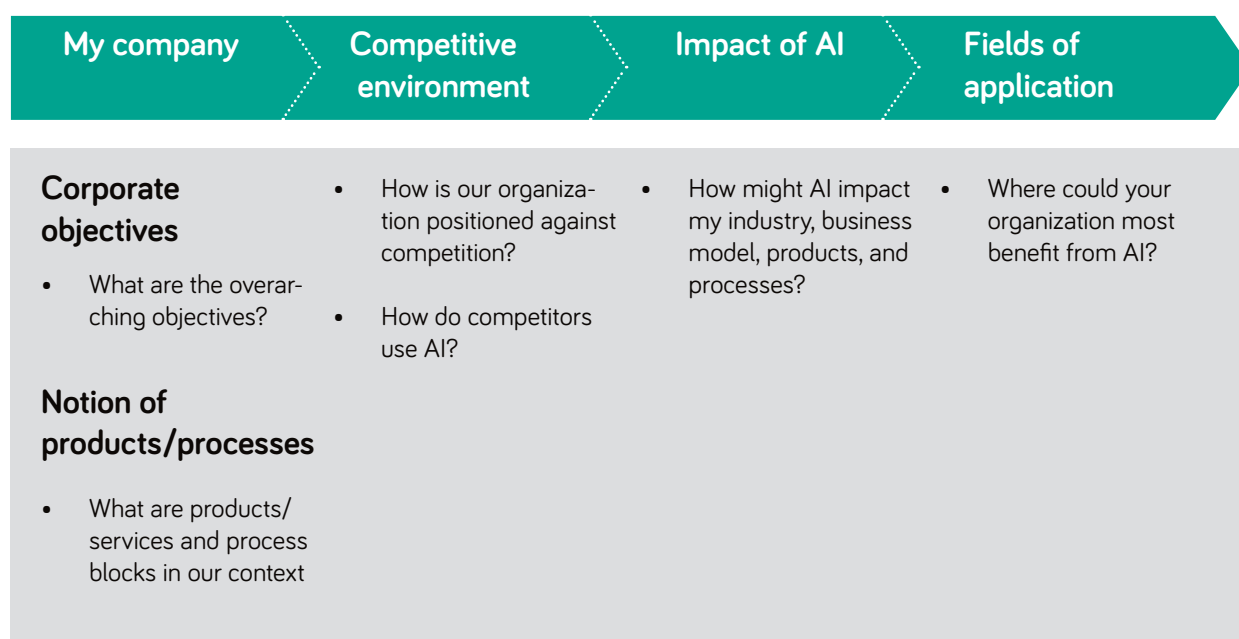
Competitive environment: The second step is to understand the competitive position and industry dynamics. Key questions to ask yourself are:

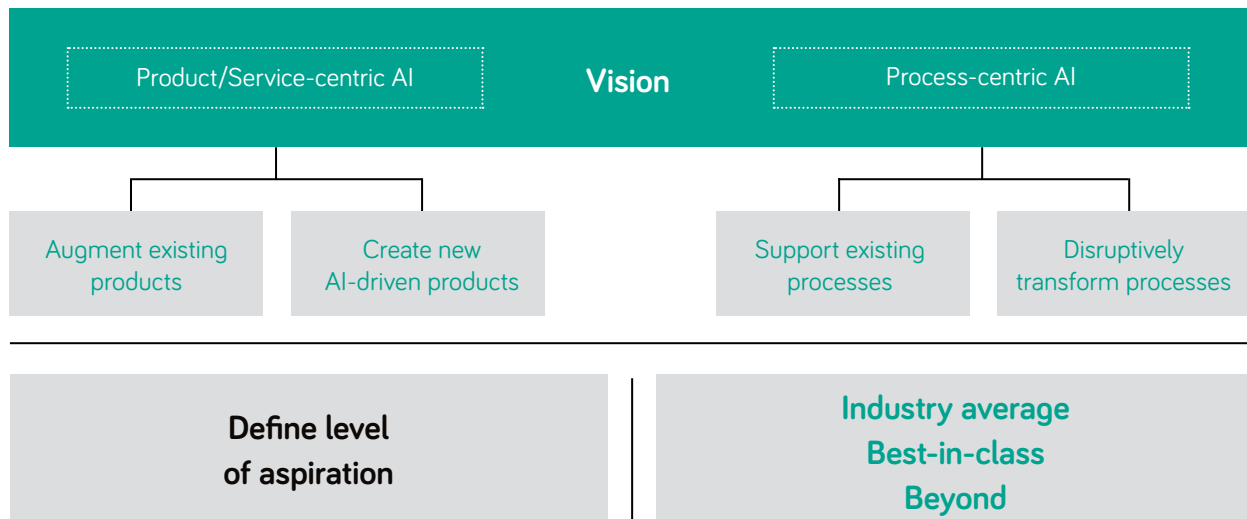
- How is AI currently used and how will AI be used by my company and my industry competitors?
- Which assets and capabilities are available at my company and at my competitors?
- Which new digital disruptors (might) move into my industry? How are they already using AI?

Impact of AI: Third, you need to understand how AI might change the fundamental source of value generation and hence the business models in your industry.

Fields of application: Fourth, based on the aforementioned analysis you need to decide where your organization could benefit most from AI – within a specific product or service and/or by improving processes. While a company focused on low prices might use AI to optimize processes, a company with an innovation focus should rather invest in new AI-augmented products. Understanding the two different ways to apply AI is paramount as they each require different stakeholders and different resources. Companies just getting started with applying AI should focus on the application where they expect the highest value as available resources are typically limited.

Using AI for products or services could involve augmenting existing products or services with AI, for example adding a voice-based assistant to a car. However, it could also involve creating new AI-driven products, for example an insurance company offering a service that can automatically assess damages by using an AI application. Process-centric AI applications could, for example, involve an algorithm that predicts raw-material prices based on various parameters such as time of year, some key economic indicators and the publicly available revenue or profit figures of major competitors. The introduction of such an algorithm would doubtless limit the power of purchasing managers within a company, who, as one can imagine, might not be too happy about this development. Or it could involve the implementation of AI-based pre-selection of potential job candidates in the HR department. SAP, for example, offers a model that does exactly that. This too might lead to apprehension within the HR department as some, or even several dozen, staff might feel that they will not be needed in the future.





If you think about it, there is another question underlying the entire process: How far do you want to go? Any company starting to move towards AI will have to define its level of aspiration at some point. And again, there are some landmark facts that might help in determining this. One of these pillars is the industry average. What level of AI do direct competitors already apply? Generally, a company needs to at least match this level to stay competitive. If all of a car manufacturer's competitors have implemented voice-based controls, this manufacturer also has to implement them.

In many industries, AI is not – yet – widely used. This offers a unique opportunity as, right now, a company can differentiate itself by building best-in-class AI solutions that are ahead of the competition within their own industry, or even beyond.

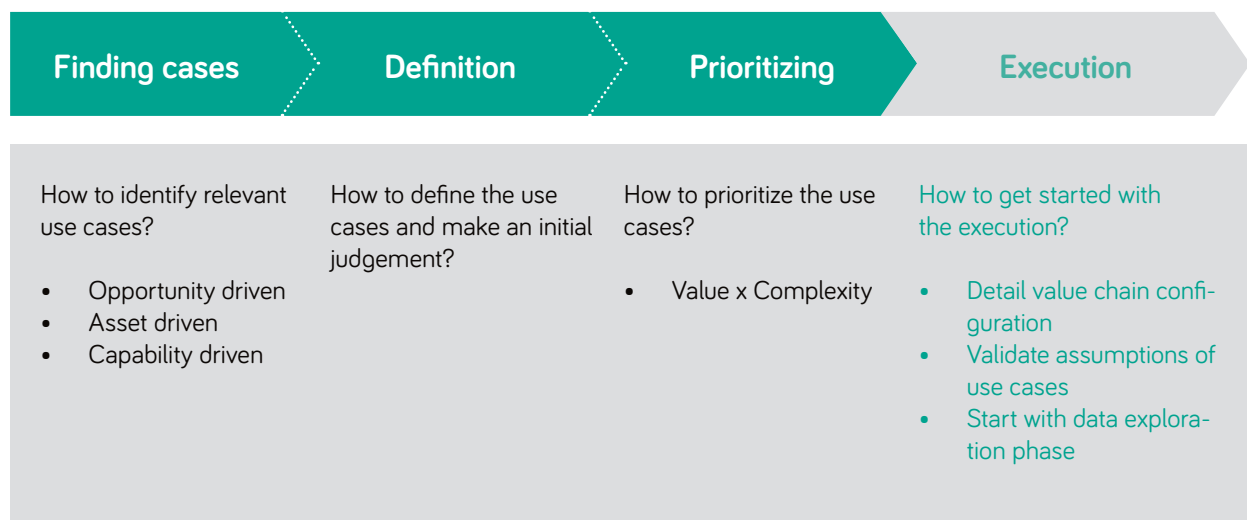
2.

Find and prioritize use cases

As a next step, and based on their overall AI vision, companies need to implement a portfolio of use cases. It is important to start thinking of use cases with a final solution in mind: How can AI solve a particular problem or improve a particular product? At this stage, don't forget that AI is not and cannot be the solution for everything.

To find the use cases, companies should follow a four-step approach:

- First of all, the use cases need to be identified. To do this, identify use cases that build upon strategic goals (opportunity driven) as well as existing strength (asset and capability driven). 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?



- Then, these use cases need to be defined in a clear and comprehensive way.
- Only after this can each use case be prioritized based on its potential value for the implementing company as well as the expected level of complexity. Companies should start with those use cases that are relatively easy to implement and provide high value. But they also have to think strategically about use cases that would provide high value but are not easy to implement yet.
- Finally, the prioritized use case needs to be executed. For this, companies should develop a clear plan.

3.

Creating the enabling factors

To be able to follow the AI vision and implement the use cases, a set of enabling factors is required: The organization, the people, the technology, and the AI ecosystem. However, the respective importance of the individual factors will differ depending on the aspired-to AI vision. First, companies need to set up the right AI organization. You have to create the right organizational setup – where is your AI team located? But you also have to create the right governance structure, and this also includes changes to the board roles.

Second, you need to get your employees ready for AI and recruit the necessary talent. New roles are emerging: AI engineers are required who can build learning systems from an engineering perspective. Getting the right people with the right skills is currently a major challenge. As hiring is hard, reskilling your existing data scientists or software engineers might be an important option.

But that on its own is not enough. Other employees, including the executives, need to have a basic understanding of what AI will enable and how it will change their working lives. You will have to bring everyone on board, as silent resistance at various points within a company can be detrimental for the success of an AI project due to the iterative approach that is required.

Companies not only need new AI skills but also have to adapt their culture: Using learning systems implies accepting failure. Employees' fears need to be addressed to create acceptance for the use of AI-based solutions. And the company needs a very high level of transparency in order to implement AI projects successfully. All this is not easily done within an organization.

Third, a company needs to build up the required technology. This includes the AI infrastructure as well as the data. Data certainly is the key element of applying AI, as training AI models requires a great deal of it. If a company does not already have well-defined data governance it is

often unlikely to have useful data. You'll need to identify data sources, build data pipelines, clean and prepare data, identify potential signals in your data, and measure your results.

But you also need the right IT infrastructure. Principally, you will have to decide whether to use your own servers and GPUs or rely on the cloud. This question is not only about data security but also about cost and economic feasibility.

Fourth, you need to address your ecosystem. At this point, no company has truly comprehensive experience when it comes to applying AI. Therefore, a company should exchange knowledge externally – with startups, academia, and other companies.

4.

Execution

Finally, the use cases need to be implemented. AI isn't like traditional software: an AI system learns continuously as new data is fed into the system. Thus, an AI system needs to be monitored to ensure that a model is still delivering the expected results. For this, a company needs to put in place the right processes.

About appliedAI

The appliedAI Initiative, Europe's largest non-profit initiative for the application of artificial intelligence technology, aims to bring Germany into the AI age and offers its wide ecosystem of established companies, both large and small, researchers, and startups neutral ground in which to learn about AI, implement the technology, and connect with each other. NVIDIA, Google, MunichRe, Siemens, Deutsche Telekom, and many more are partners of the initiative, which started in early 2018.

You can find more information about appliedAI at: www.appliedai.de

Authors:

Philipp Hartmann serves appliedAI as Head of Strategy. Prior to joining appliedAI, he spent four years at McKinsey & Company as a strategy consultant and was a researcher at the Technical University of Munich where he investigated factors of competitive advantage in Artificial Intelligence.

Andreas Liebl is Managing Director at UnternehmerTUM as well as appliedAI. Before joining UnternehmerTUM, he worked for McKinsey & Company for five years and did his PhD at the Entrepreneurship Research Institute at the Technical University of Munich.

Alexander Waldmann is Operative Director at appliedAI. Before that he served as Visionary Lead of [x], UnternehmerTUM's unit for disruptive technologies and helped create ~300 startup-teams during that time. Alex studied Computer Science at Technical University of Munich and MIT.

Hendrik Brakemeier is Senior AI Strategist at appliedAI. Before joining appliedAI Hendrik pursued a doctoral degree at the software and digital business group at TU Darmstadt during which he worked on the economics of data-based business models. Furthermore, he served as part of Accenture's analytics strategy and transformation advisory practice.

The authors thank Wolfgang Reuter for his support in writing this report and Henrike Großer for designing this publication.

