


Applying AI: Building the organization for scaling AI




“A century ago, factories electrified without rethinking their production lines and therefore saw no productivity benefits. In much the same way, machine learning technology without management and organizational change will be ineffective.”

Erik Brynjolfsson
Director of the
MIT Center for Digital Business

Elements of a comprehensive AI strategy

There is little doubt that AI will become relevant for all companies, regardless of their industry or size. When it comes to creating value from AI, several pitfalls can be observed in practice – including the isolation of AI use cases, the lack of resources and capabilities, and a poor understanding of use cases and applications. To avoid this, a systematic approach towards AI is needed. Therefore, from the very beginning, you need to be clear on the overarching objectives or purpose of your company: What is its goal? Furthermore, it is necessary to understand how AI can help to achieve your objectives. A comprehensive AI strategy consists of three parts: an AI vision, a portfolio of AI use cases, and a clear strategy for the required enabling factors.

A company’s AI vision sets the high-level goals of any AI application to be developed or deployed. It includes an understanding of the current position of the company, its competitive position and industry dynamics, including potential changes to the industry’s business model. On this basis, it can be decided where the organization could benefit most from AI – within a specific product or service and/or by improving processes. The vision needs to be translated into a portfolio of AI use cases. To build this portfolio, you need to identify and prioritize relevant use cases. To execute the use cases, a set of enabling factors is required concerning the organization, the people, the technology, and the AI ecosystem. All of these aspects need to be taken into account when it comes to the development of a comprehensive AI strategy and are further detailed in our report “Elements of a comprehensive AI strategy.”



Introduction

Today, artificial intelligence (AI) is widely recognized as one of the key technologies fueling the digital revolution. The expected economic benefit of AI is huge. By 2030, McKinsey projects that AI will help deliver a 16 percent higher cumulative GDP¹. Companies feel enormous pressure to unlock the potential of AI and stay ahead of their competitors. This is particularly important since AI adoption tends to follow a “winner takes all” logic: Market leaders have access to more data, leading to superior trained algorithms and thus attracting more customers.

Leaders understand the importance and possibility of AI, but many are challenged to effectively implement AI initiatives within their own organizations. Various departments and regions run their own AI projects. Initially, this can be positive since a broad participation and commitment of the organization is a crucial ingredient for the AI transformation. However, many – mostly large – enterprises report they have reached a point where it has become difficult to coordinate different initiatives and pilots, establish standards and processes and to avoid duplicating their efforts. They are trying to understand how best to benefit from the scarce talent available and how to ensure the optimal output.

Many leaders struggle with how to effectively set up their organization for AI. Yet, the right organizational setup is essential for successfully applying AI. The initial arrangement is key for scaling AI applications from proof of concept (PoC) to actual, productive solutions.

¹ McKinsey Global Institute 2018: “Notes from the AI frontier: Modeling the impact of AI on the world economy”

Organizing for AI should not be treated independently from other transformation activities as there is a strong overlap within most enabling factors albeit one needs to reflect the distinct differences

Three principles for organizing AI

We have had numerous discussions with companies how to effectively set up their organization for AI. Clearly there is no one-size-fits-all solution for setting up and managing AI. Each company has its own nuances that make a cookie cutter solution impossible. However, there are three guiding principles that apply to all companies and enable leaders to find success with AI. The three principles are:

- Treat AI solutions as product, not project
- Avoid failure by setting the right balance between central coordination and decentralized ownership from the start
- Demonstrate powerful leadership and broad commitment from an AI-educated C-Level

Treat AI solution as product, not project

AI and machine learning (ML) solutions have some inherent characteristics that complicate traditional project management. With AI and ML, organizations can only judge the potential outcome of a product or solution after the development phase. Therefore, they can only partially define typical project parameters such as an a-priori specified outcome, predetermined end date and fixed budget.

Furthermore, AI solutions are never “finished”. They require continuous maintenance. As such, organizations need to put a fixed team in place, responsible for a specific AI product. This team is responsible for the product’s ongoing development and its complete integration and adoption within the organization. Difficulties in the handover usually stem less from changes in the lines of code but rather from an evolution of data, entailing the need to retrain the entire system. As such, organizations should approach AI activities from a

product, not project, framework.

Set the right balance between central coordination and decentralized ownership from the start

Among leaders and managers, there is wide agreement that certain aspects of AI are best managed centrally.


Central coordination prevents teams from duplicating efforts and ensures collaboration. For example, if one team has successfully solved a problem using AI, there should be a mechanism in place that enables that team to share that knowledge on a wide scale.

To support this collaboration, organizations must centralize certain types of AI expertise, such as machine vision, natural-language processing. Also, central coordination dramatically increases efficiency and effectiveness by:

- Consolidating best practices on AI tools (e.g., ML platforms or container solutions for deploying models)
- Providing benchmarking of solutions
- Offering central repositories of training and test data sets, and reference models
- Facilitating easy access to preferred tools

To demonstrate the importance of a centralized approach, we want to draw attention to an example from a large European automotive OEM. The company began hiring ML engineers decentrally and soon encountered problems. The engineers did not have enough power and influence to tackle large and cross-functional challenges. The set up created a lack of tangible outcomes and high dissatisfaction among the individuals. On a larger scale, centralization helps an

Overview of centralized vs decentralized responsibilities



organization align on direction. Through centralization, leaders enable the creation of an overarching AI strategy. They also set the general direction and give priority to the most pressing AI projects.

However, centralization does not offer a complete path to AI success. When organizations centralize AI responsibilities, they can inadvertently build isolated teams of “techies” that develop insular solutions without real business value. For example, a European company hired a large team of data scientists, many straight out of university. The team did not have access to relevant use cases or data from the business and had limited exposure to business domain experts. As such, they were unable to progress as a team, leading to massive frustration.

Therefore, the most successful AI initiatives set the right 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 be able to draw on centralized know-how and resources to effectively tackle identified challenges.

One best practice approach we often see in companies 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. The AI CoE can be complemented by an integration team responsible for assimilating and scaling PoCs into the organization’s structures. In some cases, this model is known as the “hub and spoke” model.

The hybrid approach ensures balance and promotes collaboration between the central teams and decentralized units, both key to the successful application of AI.

Demonstrate powerful leadership and broad commitment from an AI-educated C-Level

Historically, companies have held their analytical expertise in a specific and singular department. Sometimes it is a controlling department within the finance department that has built up analytical competencies and is now addressing “advanced” analytical topics like ML. In other cases, the know-how to manage AI resources is found in the IT department.

No one team or department is “best” for handling AI resources and responsibilities. We find organizations take a variety of approaches to setting up their AI CoE. Depending on the context, each CoE can work effectively. But in general, it makes sense to position the AI

CoE close to critical decision makers. The involvement of these individuals also highlights the importance of AI to the larger organization.

The specific C-suite member who leads the AI CoE can vary from company to company. In organizations where AI is mainly focused on process and efficiency improvements, the COO or CFO are often good leaders for AI. However, if the application of AI is part of a large and comprehensive digital transformation, the CEO or the CDO can be the right executive to lead the initiative.

On the other hand, if AI is largely used to drive product innovation, the CTO or CPO can be a suitable leader. Finally, the CIO can be an advantageous leader because they keep AI development close to the organization's IT infrastructure, ultimately helping with deployment.

But if choosing the CIO, the AI CoE needs to create a sufficiently strong link to business needs to avoid creating a group of isolated “techies.” A MITSloan/BCG study² indicates that the likelihood of capturing value from AI is twice as high when responsibility for AI is placed outside the CIO's department. Indicating that there is a strong risk that AI is treated as just another technology innovation rather than influencing entire value chains and business models.

The table shown below outlines potential AI CoE leadership assignments and the respective advantages and disadvantages of each. When choosing the right executive to align with the AI CoE, organizations should take a few factors into account. They should look at their current data repository, specifically the

location of data within the company. AI CoE and data responsibilities should be assigned near one another. Furthermore, the best assignment of AI resources also depends on the context in which AI is used and the company-specific historical developments.

Organizations must exercise great care in selecting the best leader to oversee the day-to-day activities of the AI CoE. The right person will understand AI from a technical standpoint but also have business savvy and expertise. They will serve as the chief evangelist for AI and be able to convince others to make changes to business models, product development plans and corporate culture. When working with other leaders, the AI CoE leader needs to have the gravity to defend resources and priorities when other (short-term) activities threaten to push AI downward on the agenda. Finally, this person needs to be an excellent strategist with the ability to manage the complexity of all AI activities within the organization.

Importantly, AI should be the responsibility of the whole management team, not just the person responsible for the AI CoE. Further, the CEO must clearly commit to spearheading the AI strategy as it will heavily influence and be influenced by the organization's overall corporate strategy. The AI CoE is only partially responsible for AI systems. The organization needs cross-divisional teams to implement AI processes or products.

In order to ensure true commitment, C-level executives need to understand the fundamentals of AI technology and the related implications. Thus, educating top management is paramount.

	Pros/cons	Applies if...
CEO	+ High authority - Lack of capacity	AI is expected to deeply transform business model
CFO	+ Owner of data warehouse & BI - Limited ties to products	CFO oversees large transformation budget
CIO/IT	+ AI expertise close to IT expertise - Often too far away from business	CIO is strongly business-oriented
COO	+ Deep understanding of internal processes - Limited ties to products	AI vision is mainly process-centric
CTO/CPO³	+ Close ties to products and innovation - Limited connection to internal processes	AI vision is mainly product-centric
CDO	+ Transversal and transformational leadership - Not directly involved in processes or products	AI is part of larger digitization journey requiring central leadership

² MITSloan/BCG 2019: “Winning with AI”

³ Chief Product Officer (CPO)

A prototypical organizational set up for AI

In the following, we outline a prototypical organizational setup for a successful AI transformation within larger enterprises. Of course, each organization has its own idiosyncrasies. But successful AI transformations do have commonalities. Below, we describe what we have seen in nearly all organizations with some form of an AI CoE.

As with most organizational initiatives, the AI CoE changes over time. Initially, the CoE must focus on defining and setting up appropriate processes, including various one-time tasks. Over time, however, the CoE shifts its focus toward more continuous tasks and forms central and outer core teams.

The AI CoE: A central team with decentralized extensions

The AI CoE is the main driver and point of contact for all AI topics in an organization. It bundles expertise, implements the most strategic projects and enables the adoption of AI. Oftentimes, the CoE consists of a core team linked to (potentially virtual) decentralized extension teams. It can also operate as a dedicated

team sitting physically together or as a remote or virtual team spread out across departments.

Depending on the size and complexity of an organization, as well as its AI-specific problems, there can be more than one core AI team. For example, an automotive manufacturer will have different challenges when working on AI embedded in its vehicles (e.g. autonomous driving) from AI in its production processes. Or, it may have different business units working in different continents. In these cases, multiple teams make sense.


However, when an organization requires more than one AI team, they must ensure central coordination between these teams. The AI CoE should take the lead in this effort.

The responsibilities of an AI CoE


The responsibilities of the AI CoE are broad. Successful CoEs are classified by three main tasks:

- Defining strategy and use case prioritization
- Developing AI solutions
- Enabling the organization

Visualization of the organizational structure of an AI CoE



What does an AI CoE do?



1. Defining strategy and use case prioritization

The AI CoE bundles the organization's AI expertise. Therefore, it needs to proactively engage in defining AI strategy and identifying use cases. The CoE advises senior management on AI trends in and out of the industry and judges the feasibility and complexity of ideas around AI. Each organization will determine the extent of the group's oversight. For example, the CoE may facilitate the overall process of defining AI strategy. Or, it may give input through an expert advisor role while another unit drives the AI strategy process.

Beyond strategy, the AI CoE advises the organization on concrete AI use cases. By maintaining a strong network in the organization, as well as observing, listening and facilitating workshops, the CoE identifies use case opportunities. Once ideas have been collected, the CoE supports business units with prioritization. In this case, the AI CoE is the expert in evaluating the feasibility and complexity of a given use case. It uses these two factors, as well as business value, as the main criteria for prioritization. In one instance, we found the central AI team of a large German utility company to regularly hold a series of use case workshops with dispersed business units. The team would aggregate their findings in their overall use case portfolio. By managing the use case portfolio centrally, the AI CoE also helps to keep the focus on creating value from AI.

Additionally, the AI CoE sets the overall process and governance framework for AI within an organization. In other words, the CoE establishes standards for the training of AI algorithms, creates the necessary docu-

mentation, provides insight on decision making, shares results and enables the reuse of solutions. The CoE is also responsible for creating central data repositories to benchmark AI solutions.

2. Developing AI solutions

Operationally, the AI CoE plays a role in implementing AI solutions. The extent of that involvement is specific to the organization and its resources. The CoE can drive the implementation of an AI project or give that responsibility to another team.

There are certain factors that advocate a central development approach. These include:

- **Low organizational maturity:** At the beginning of their AI journey, it makes sense for organizations to implement most AI solutions through central AI and implementation teams. As they mature with AI, they can gradually move toward decentralized AI teams that can run projects on their own.
- **Strategic relevance for the core business:** Frequently, the core team of the AI CoE is responsible for implementing the most strategic use cases for the overall organization. Local teams tend to work on use cases specific to their unit.
- **High technical/analytical complexity:** The core team of the AI CoE houses most of the organization's AI expertise. Therefore, complex solutions should be developed by this team.
- **AI solutions that are relevant for several units:** Solutions that are relevant for several business units should be developed centrally. A large insurance

company centrally developed AI solutions for OCR and image recognition that can be used by other BUs as a service to build specific products and solutions.

- Exploring new technologies and partners: A central team might create a PoC to test new technological approaches and/or new partners and external solutions. These learnings can then be used across the organization.

When it comes to operationalizing AI solutions, the role of the AI CoE varies. Typically, the respective decentralized units will oversee the operation of an AI solution. However, the central team may be accountable for heavy AI products or services that need continuous maintenance and retraining.

Remember, AI solutions should be treated as products. As such, the overarching responsibility for the solution should stay with the product owner from the respective business unit. Nonetheless, the AI engineer should also retain ownership for the maintenance of the solution over the entire life cycle, while remaining part of the central AI CoE.

The handover from the AI CoE to the decentralized unit or implementation team is particularly important. It can make or break a seamless integration of the product and the smooth transfer of responsibilities post-PoC.

“AI is the most important transformation topic in the history of our company”

CIO of an Automotive OEM

3. Enabling the organization

While less obvious than its involvement in use case ideation and creation, successful AI CoEs play a crucial role in enabling AI within the organization. With this, the AI CoE has partial responsibility for talent, culture, awareness and communication, collaboration, AI infrastructure, data and the AI ecosystem.

Finding and developing talent:

Organizations have only recently started to think of AI business applications as value drivers. Therefore, few organizations

understand fully how and where to deploy AI. Leaders cannot always draw the fine line between science fiction and feasible application. Also, business and domain experts struggle to identify where they can use AI to help their teams or customers.

Most companies that have established an AI CoE have tasked this group with training employees on AI, rather than leaving this responsibility to Human Resources or the Learning and Development team. In fact, we have seen organizations successfully cooperate with universities or other external partners to design and deliver such training programs. Some programs even culminate in some sort of certification.

Even with training, organizations continually find they have skill gaps. Therefore, in addition to educating the organization, the AI CoE supports the organization in recruiting AI talent. It does this by identifying specific job needs or assessing applicants' technical skills. A close collaboration with the HR department is paramount. Moreover, the AI CoE also needs to work with marketing and employer branding. Tech workers are continuously looking for challenging opportunities within the field of AI. They are especially interested in jobs that give them access to large data pools and the freedom to apply and extend their existing knowledge while also continuing to learn and grow on the topic. The AI CoE is best positioned to contribute language and specifics to the teams that are traditionally responsible for recruitment and hiring.

Also, we frequently see newly hired AI experts leave companies after only a short tenure as they do not mesh with traditional corporate culture and feel isolated. To avoid unwanted turnover, the AI CoE needs to take responsibility for developing suitable career paths for AI experts. With the CoE leading the charge, organizations can save time and money. However, our experience shows that when the AI CoE takes over this responsibility, it creates strong friction between the group and the business units. Typically, business units want to retain ownership of this task. As a result, the AI CoE needs strong backing from its executive leadership.

Fostering a culture that embraces AI:

Digitalization and AI can heavily influence how an organization operates. An AI transformation cannot be successful unless the

company's culture is prepared for this shift. Therefore, organizations need to move away from favoring deterministic knowledge to acknowledging experimentation and from long-term cascadic development processes to rapid iteration.

As part of the cultural shift, organizations must ensure they focus on the positive aspects of AI while a solid ethical foundation for the use of AI has to be set. For example, AI can eliminate boring, repetitive tasks and scale an organization's skills.

Leaders should also be transparent about situations where AI can lead to process automation and potentially, job reductions. Or, where business models get transformed from a product- to service-orientation. Even in cases where AI "only" enhances an existing product, it can be challenging to convince engineers to try new methods when their traditional approaches have served them well. Leaders also need to prepare employees that AI-based products create dynamic results that are probabilistic and not deterministic and therefore might not be fully reproducible.

The above poses a greater challenge to organizations than other digital transformations. Nonetheless, the AI CoE cannot and should not be responsible for an organization's overall cultural transformation. Instead, it should closely collaborate with other departments and support them in their charge.

An AI CoE should also help to create transparency around the AI culture. The AI CoE should regularly investigate the status of the cultural change to answer questions such as:

- To what degree do various parts of the organization understand AI opportunities?
- Are employees aware of landmark projects and overall strategic objectives of the organization?
- Do employees identify themselves with these objectives?
- What are the emotions regarding AI?

Creating awareness for AI:

The AI CoE should help create awareness of AI within the wider organization. Through showcases and introductory trainings, they can create curiosity and a sense of urgency around AI. As part of this, the CoE should ensure that managers develop a fundamental

understanding of how AI really works in practice.

Through a structured communication campaign, the AI CoE must also ensure that information and success stories are spread throughout the organization. This approach guarantees that the questions raised above are answered and that the general feeling toward AI is positive.

Facilitating collaboration and an AI knowledge exchange:

As the central AI unit, the CoE is a bracket around all AI activities in the organization and responsible for fostering collaboration. Thus, the CoE facilitates seamless cooperation on AI projects. It also creates a vibrant AI community that allows the constant exchange of ideas and products, a learning mechanism among AI experts. As such, the CoE establishes thematic groups around topics such as NLP, computer vision or deep learning. It organizes internal conferences dedicated to the expert community, but also half- or one-day events targeted to a broader audience with external speakers and show-cases.

Building the AI infrastructure:

The operation of AI is a relatively new activity. Therefore, the task of building the right infrastructure can be staggering for organizations. Further, the requirements for the infrastructure will mature with the AI CoE over time.

To build an AI infrastructure, the AI CoE has to ask questions that are aligned with the maturity of the software it is expected to produce or to integrate into the business environment. In the early stages, experimental flexibility is of the utmost importance. Platform wars related to hardware, software, and the cloud layer are ongoing. A young AI CoE needs to be flexible in its choice of the technical stack, be it a software architecture or a physical developer device.

Over time, however, as technical teams grow, their solutions move from PoC to production and users start to rely on them. This progression creates more responsibility for the AI CoE and the centralized decisions it has to make for developers. Historically, organizations with a multitude of unstandardized tools or unreliable libraries end up assigning engineering capabilities to individuals and ultimately create a so-called "technical debt".

An AI CoE must actively avoid this trap – the same one that many companies in the 2000s stepped into when code bases became unmanageable and crucial company know-how was lost with departing engineers.

Certain tool stacks must be publicly embraced and favored over others. Core activities should rely on a compatible code base that follows well-known best practices from the world of software engineering. This must not be confused with the enforcement of certain tools or libraries, crippling the speed the AI CoE has in the rapidly changing domain. Flexibility and maintainability are competing requirements but must be considered together.

Over time, the maturity of the AI CoE and the infrastructure lead the team to settle on favorable tech stacks on the hardware, cloud and software layer, in particular for existing solutions or challenges the AI engineers encounter regularly. At this stage, such choices might rely strongly on the AI CoE. Both hardware architectures and software frameworks for high-throughput computer vision might differ fundamentally from, say natural language processing. The same is true for AI engineering activities such as data versioning or reproducibility of experiments. So far we have seen well working “on-prem” strategies for large-scale problems with standardized infrastructure choices. But flexible application architectures, as well as strictly enforced “cloud first” strategies, where applications follow concrete architectural assumptions. This is still a case-by-case decision and should be taken with a partner that has experience in such questions. A crucial best practice is the inclusion of so-called “ML Ops” talent who address such questions. Ideally this talent is an experienced AI engineer who supports the developer’s lifecycle.

Managing data:

Many organizations have brilliant ideas on what to do with AI. However, lack of experience in collecting or consolidating relevant data make it difficult to turn those ideas into reality. Take a manufacturing company who wants to offer a predictive maintenance service to its customers. While most machines are already equipped with sensors to register the relevant data, many producers do not store or process data in a consolidated fashion in the cloud or the corporate repository. More importantly, some companies have

not yet collected their customer’s permission to store and use this data.

In other cases, organizations have plenty of data, but nothing is labelled for supervised-learning approaches. For example, a car insurer that has thousands of images of damaged cars would need detailed information on the type and cost of the damage to use this data for automated claims handling. However, generating this type of labeled data is tedious and time consuming.

In essence, this is a problem as the data itself is an integral part of the yet-to-be-created ML system. When data becomes code, it has to be documented and maintained with the same care as any digital asset. The AI CoE has to push and spread this central insight as one of their core tasks. Every individual who influences the design and development of an AI solution has a responsibility to think about its “data code.”

When it comes to data, two separate tasks have to be addressed:

1. Data curation and governance of “passive” data (typically in the hands of the CIO)
2. Acquiring and developing the right data as part of a concrete AI use case.

While the latter might be part of the AI CoE’s tasks, the former is often the responsibility of a separate central data team, frequently headed by a Chief Data Officer, as managing data is not only relevant for AI but also for other functions of the organization. This central data team is responsible for defining and implementing a coherent data governance structure to support central data integration. Their activities need to be closely aligned with the AI CoE. In some organizations, we have observed central data teams acting as suppliers for AI teams. The data team might manage a central data lake and operationally support the actual data pipeline for AI projects.

The task of acquiring and developing the right data for a specific use case should be handled in collaboration with business units and specific engineers in a decentralized fashion. Individual business units often generate ideas for the application of AI and use their available data and relevant domain knowledge to act on those ideas. However, they may need data beyond what is readily available to them. In these cases, the AI CoE

might help with finding or building up additional data sources.

Building the AI ecosystem:


An organization cannot excel in AI without leveraging an ecosystem of partners such as startups, academics and other industry players. By collaborating, sharing best practices and working with innovative startups, organizations avoid common pitfalls and accelerate the AI journey. While each part of an organization has the freedom to look for its own partners, the AI CoE should build a common bracket for this ecosystem and link the organization with suitable partners. The AI CoE needs to be proactive in this effort and dedicate some of its resources to building and maintaining the ecosystem.

Roles needed for an AI CoE


AI development teams will always be cross-functional, involving AI experts (strategist, product manager, engineer, data scientist), domain experts, UX designers and IT experts. AI expert roles can be filled by employees inside and outside the AI CoE. In any case, members of the AI CoE core team should assist the development team in an advisor role. Besides those mentioned, other roles that may be needed during the development of an AI project include devops, ML researchers, AI trainers or AI-trained lawyers.

Overview of role staffing depending on model

Pilots implemented by central CoE:



Pilots implemented by decentral CoE extensions:



To deliver on the diverse responsibilities around AI, we have found that AI CoEs are staffed with the four expert roles described below:

Role	Description
AI Strategists	AI strategists have a deep understanding of AI technology combined with expert business and domain knowledge. They act in an AI evangelist role and support the organization in finding its AI strategy and use cases. Their knowledge of AI and the business ensures that AI systems deliver value. Further, as AI enablers they ensure the AI CoE fulfills its duties to drive the AI transformation.
AI Product Managers	AI product managers are responsible for the AI product throughout its lifecycle, from PoC to the integration and application of the live product. The team developing the product might change during the lifecycle, but the product manager should be a constant. They develop the product and follow a clear vision and strategy. Furthermore, the product manager identifies requirements and potential risks in the AI system while ensuring that the focus on the target user is not lost.
AI Engineers	AI engineers implement the AI systems. They have technical responsibility for the overall solution and manage the scalability, reproducibility and overall quality of an AI product. Deploying skilled AI engineers is an important factor for a successful AI implementation, since developed models and algorithms need to be integrated into and aligned with existing processes, IT landscapes and customer requirements.
Data Scientists	Data scientists juggle data. More specifically, they derive insights from the vast amount of data available within the organization. Using their excellent statistical understanding they conduct analyses and discover root causes. They can validate and present data-driven use cases and hypothesis to be solved with AI.

Besides the experts, further roles are needed in a core AI team. These include:

Role	Description
Domain Experts	Not every business unit will have an educated AI strategist, so domain experts are needed to apply AI within their field. They develop use cases within their area of expertise in collaboration with the AI CoE.
UX / Designers	UX /designers are essential to make an AI product “useable” for customers through a human-centered approach. Trust in AI needs to be built up and maintained so that customers accept and use the product. Therefore, UX / designers need to be aware of and address all opportunities and challenges of a specific AI product.
Infrastructure Engineers	Infrastructure engineers manage the important intersection between existing technology and newly developed AI. They are responsible for building up an infrastructure solution that meets all stakeholders’ needs. They also manage existing computing resources required for AI and are responsible for the system administration to ensure a smooth operation of AI and ML products.

The above-mentioned roles make up the core of an AI development team. Some of them reside within the AI CoE while others do not. The AI CoE bundles the core competencies and skills regarding AI in an organization, irrespective of the individual AI products. The development teams, however, consist of employees from the AI CoE as well as employees from the rest of the organization who are needed to develop a specific product

This approach brings opportunities for employees but can create friction in organizations. Employees are taken out of their line functions and (temporarily) moved to a central team that tackles new and innovative topics. Employees experience change, both when they are relocated and when they return to their original team. However, such a shift resembles talent/high potential and trainee programs that have been executed in larger organizations for years.

Ensure a smooth interaction between the central team and the decentralized units

The interaction between the central AI team and decentralized unit within an organization is key to success. But how does an organization achieve the ideal interaction?

First, organizations need dedicated people to lead the cooperation, with the AI strategist reaching out from the AI CoE to their counterparts in each division, called “AI ambassadors”. Together, they ensure effective and


efficient cooperation between the different organizational units on AI topics.

The AI strategist is part of the AI CoE and acts as a direct point of contact for business units regarding AI topics. They support the ideation process for interested business units and conduct workshops with them.

The AI ambassador bridges the gap between business and AI technology within the organization’s business units, thereby acting as a bridgehead to the AI CoE. The AI ambassador may also serve as the domain expert in an AI project team by encapsulating domain knowledge.

Both roles are key to ensure a smooth interaction between technology and business teams. They ensure AI products are built with a dedicated business purpose in mind.

Second, organizations need to clearly allocate responsibilities. We found a good example of collaboration between an AI CoE and business units in a leading German direct bank. The bank’s business units are responsible and liable for the AI system, while the AI CoE manages the technological side including versioning and updating the AI. Within the responsible business unit, there is a dedicated product owner for the AI system throughout its lifecycle, while the developing AI engineer is located within the AI CoE. The engineer might hand over the responsibility for the technical support to an ops team within the AI CoE depending on the complexity and continuous effort for the AI system.



 **AI strategist** at the central AI CoE acts as a direct point of contact

 **AI ambassador** acts as a decentral counterpart who bridges the gap between business and AI technology

Defining the governance structure

Even with a solid organizational setup, leaders are challenged to define a common bracket around the various top-down and bottom-up AI initiatives. AI projects are often experimental, particularly as organizations learn how to work with the technology. Therefore, they often find it difficult to estimate a return on investment and generate quick results from AI initiatives.

To counteract this challenge, the C-suite and board, particularly the CEO, need to show a strong commitment to advance AI and continuously develop it within the company. In executive and board meetings, leaders should regularly review AI efforts, new product developments, implementations and progress toward the overall AI transformation. Leaders should understand the overall strategic importance of AI and show support for the various organizational AI efforts and changes.

Organizations also need to set a clear governance framework. AI initiatives should be a regular topic on the board meeting agenda. However, to keep up with the AI vision and strategic goals, leaders also need to define decision-making processes and responsibilities as well as effective escalation mechanisms.

The following RACI-matrix can help organizations define appropriate AI governance. In general, AI governance should be in line with existing board and decision-making structures as well as other transformation initiatives.

Ensuring PoCs are moved into production

For many organizations, the process of moving from PoC to operation of an actual AI product is daunting. The handover from the AI CoE to the decentralized units is crucial. The process must allow for a seamless integration of the product and the shift of responsibilities.

A one-size-fits-all approach for handover and collaboration between the AI CoE and the business units does not exist. However, we have observed some common patterns. The

extent of the handover varies and depends on the amount of ongoing development needed for the product. The responsibility for heavy AI-based products or services with ongoing development needs might remain with the central team. However, products that are developed within the AI CoE but only require moderate continuous development can be taken over by the decentralized extensions.

Often, the AI CoE uses a dedicated intermediate implementation team to move the PoC into production. This team is responsible for integrating the PoC into existing IT systems and structures.

The changing role of the AI CoE

As an organization matures with AI, the setup of the AI CoE can and should evolve. Commonly, experienced organizations take a layered approach to how they shape the AI CoE. They form a central core team and an outer core team. The central core team is responsible for capturing efficiency gains of reference models and developing new technologies. The outer core team is responsible for provisioning central APIs to other units; coordinating learning content, central tools and data repositories; and acting as a quality gate for potential new AI applications.

The time needed to set up and transition tasks within the AI CoE varies in each organization. Organizations must account for previously established tasks and processes as well as the overall ability of employees and teams to adapt to new technologies.

Exemplary completed RACI-matrix

Governance Bodies Involved	Decision Domain					
	AI Strategy	AI Infrastructure	AI Architecture	AI Use-Case Identification	AI Investment and Prioritization	...
Senior Business Executive (C-Level)	A	I	I	I	A	
IT Lead / CIO	C	A/R	C		C	
AI Lead	R	C	A/R	A/R	R	
Business or Process Leaders	C			C	C	
Individuals / Small Groups	I			C		
...						

R = Responsible
 A = Accountable
 C = Consulted
 I = Informed

Learning from the big players: AI at Amazon and Uber

Within the world of AI, there are clear frontrunners. These organizations have been able to fully integrate AI into their operations and culture so they can optimize the technology. Typically, this means they had to shift the organization's focus to wholly invest in AI. Two obvious examples of these frontrunners are Amazon and Uber. Their stories provide organizations with guidance on how to further their own AI agenda.

Amazon has worked with AI and ML from an early stage and within different areas of the company. For instance, they have used AI to enhance their recommendation engine and their shipping optimization.

As more opportunities to leverage AI became available, Amazon stepped up its game and began to use more AI technology throughout the company. All departments were asked to think of ways to use AI and ML in product improvement. The push led to the creation of new areas of business, including the Alexa voice-controlled assistant.

To become more efficient and profit from other AI use cases within the company, Amazon started to introduce what they call the “flywheel organization.” This organization aims to move away from several siloed AI efforts and toward a system where one achievement fuels and benefits the next. Thus, a central AI exchange between business units has become a crucial aspect of Amazon's success story. The original department that developed the product retains the specific AI knowledge as well as ownership for and responsibility of the AI system. However, they share their knowledge and products with the whole company through a central unit dedicated to supporting AI and an applied research unit.

In addition, Amazon has built a central unit dedicated to supporting and sharing AI. The company has also created an applied research unit to attract talent with an interest in scientific work and more specifically, in tackling difficult problems with new scientific methods.

By asking and enabling all departments to incorporate AI and ML into their products, Amazon keeps the flywheel going. They use their existing AI systems in other products, helping to gather more data and improve their products. For Amazon, the result is attracting both customers and talent. The successful effort is powered by a hybrid organizational approach. In it, decentralized teams develop specific use cases, build the internal system and maintain responsibility for the AI, while the central AI team facilitates the development and proliferations of the technology throughout the company.

Uber has a goal to democratize ML across the company. It uses ML in different units to predict consumer choices at Uber Eats, forecast driving capacity requirements and estimate arrival times during Uber rides.

However, Uber was facing a broad skills gap and talent shortage for solving their different ML problems. The company had to come up with the right set-up to achieve its goal of spreading ML throughout its teams. Like Amazon, they implemented a hybrid approach. Dedicated product teams build, deploy and own the ML models. Throughout the development process they are supported by specialist teams that provide expertise not found in the core team.

Furthermore, Uber has a central ML platform team that develops general purpose ML tools. These tools can be used by teams throughout the company to address use cases and products not covered by the general platform. Finally, Uber has a central team of AI researchers who provide the product team with cutting-edge developments from the scientific world to help improve Uber's services.

Authors

Philipp Gerbert is the Future Shaper of UnternehmerTUM and Director at appliedAI. Previous activities include many years as Senior Partner and Lead of Digital Strategy at BCG, Fellow for AI in Business at the BCG Henderson Institute, as well as a Partner at the McKenna Group in Silicon Valley. He holds a PhD in Physics from the Massachusetts Institute of Technology (MIT).

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.

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

Maria Schamberger is a Senior AI Strategist at appliedAI. She has a rich background within the financial services industry from her former role as Vice President at the Allianz Group as well as consulting and research experience from McKinsey & Company. Maria studied Corporate Innovation at Stanford University and Banking at the Frankfurt School of Finance and Management.

Alexander Waldmann is the Director of Operations at appliedAI. Before joining appliedAI 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 the Technical University of Munich and MIT.

The authors would like to thank Ulrich Schlickewei for his support in writing this report and Henrike Noack for designing this publication.

Contributors

This report has drawn on the experience of leading experts from appliedAI partner companies:

Martin Arend is General Manager of Automotive Security at BMW Group and part of BMW's projectAI. Previously he led the BMW Research E/E Architectures and Technologies and was responsible to set up an AI strategy and shape the initial AI CoE. Martin studied Electrical Engineering at Ostbayerische Technische Hochschule Amberg.

Gemma Garriga heads the AI and Advanced Analytics department at Allianz. Prior to joining Allianz, she led research in the area of machine learning at INRIA, France, where she has held a tenured position since 2010.


Gunnar Lange is Director of IT Innovation at Audi AG. Before joining Audi AG he worked for Kearney as a management consultant and for Daimler AG in various management positions in Germany, the USA and Japan. His focus is on generating value through artificial intelligence.

Andreas Nawroth is Head of AI at Munich Re. He is responsible for developing AI solutions and platforms in the domain of text, image and risk solutions. He holds a PhD in Physics and has 10 years of experience in the banking and insurance business.

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, researchers, and startups neutral ground on 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



appliedAI Initiative GmbH
August-Everding-Straße 25
81671 Munich
Germany
www.appliedai.de

A joint
initiative

**UNTER
NEHMER
TUM**

 **Ipai**