

# The new productivity booster

How companies can harness the potential  
of generative AI

### The new productivity booster

How companies can harness the potential of generative AI

Generative artificial intelligence will fundamentally change companies' everyday routine. Yet it will do so in a different way than many stakeholders expect. According to a new global Roland Berger survey of executives across all industries, it will significantly increase the productivity of knowledge workers. However, the survey also shows that – at least in the short term – this will not be accompanied by a tangible reduction in headcount. Only 16% of managers expect generative AI-related job cuts of more than 5% by the end of 2025.

At the same time, it is becoming increasingly clear that some industries are implementing the new generative AI-related tools much faster than others. Wholesale and Retail; Information and Communication; Professional, Scientific and Technical Activities; and Other Service Activities lead our rankings, with adoption rates of at least 57%. Nor is the picture any different at intracompany level. While the new speech AI models already play a routine role in Marketing/Communications and Client/Customer Services, they are used only very sporadically elsewhere.

Despite these differences and the technical limitations in relation to data protection, for example, no company should voluntarily forego the benefits of the new productivity booster. We recommend that companies adopt a multistep approach: First, they should create transparency about their own value chain with the help of the *Roland Berger AI Readiness Radar*. Then, the management of unstructured data must be improved so that personalized AI language models can leverage the organization's knowledge. All this must take place within an infrastructure where top priority is given to security and data protection.

# Contents

P 4	<b>1/ Introduction</b> A way out of the productivity crisis?
P 6	<b>2/ Strong competition</b> The market potential of generative AI and how it works
P 12	<b>3/ Already well established in the business routine</b> What top managers hope for from the new tools
P 18	<b>4/ Limitations</b> Technical and practical obstacles to application
P 20	<b>5/ Recommendations</b> Three steps to AI readiness

# Introduction

A way out of the productivity crisis?



New technologies spawn fresh hopes, and generative AI is no exception to this rule. It is expected to do nothing less than resolve the prevailing productivity crisis in advanced economies. The grim reality is that, despite the advance of digitalization, productivity has grown only very modestly over the last two decades. As in previous years, the forecast for the current year is little to no growth. ▶ **A**

For the first time, however, there are now growing signs that new AI language models can prompt a turnaround. This alone is reason enough to take a sober look at the current hype surrounding generative AI and focus on the issue of potential productivity gains from the new AI models.

To this end, Roland Berger conducted both a global survey of around 100 executives from various industries and a subsequent series of qualitative interviews in July 2023. These exercises allowed us to make a first solid, quantitative approach to several questions.


In chapter 2, we will first explain the mechanisms and levers used by generative AI to boost productivity. Embedded in the context of digitalization as a whole, they provide powerful arguments that generative AI is indeed a revolutionary technological breakthrough that will radically change the field of services in particular.

In chapter 3, we discuss the results of our survey with regard to the following key questions: In which areas and sectors does generative AI hold out the greatest potential? How can the associated productivity gains be quantified? And what are the associated labor market effects?

It is important to note that the positive answers to these questions emerging from our survey cannot be taken for granted. Nor can the corresponding solutions. Companies that want to make profitable use of generative AI face the specific challenges and limitations that are outlined in chapter 4.

Lastly, chapter 5 formulates practical recommendations on how to overcome these challenges and limitations.

**A Slump in productivity**  
GDP per hour worked [USD]



Source: OECD

# Strong competition

The market potential of generative AI and how it works

A large, bold, black number '2' is positioned to the left of a large, white number '1'. The '1' is partially cut off by the right edge of the page. The numbers are set against a bright yellow background.

All economic activity takes place in one of three categories: production, transactions and interaction. The same goes for the entire history of digitalization, with distinctions drawn as a function of the category in which it is especially disruptive at any given time.

For example, "traditional" AI in the field of robotics and automated production – i.e. in the days before generative AI – can be understood as the last innovation cycle in production for the time being. In the area of transactions – i.e. how products get to consumers – digitalization has likewise had an enormous impact in recent years. We see this, for instance, in the number of new sales opportunities that have vastly reduced transaction costs in the platform economy.

Which leaves us with interaction. Here, the digitalization of customer services and other services, for example, is clearly still lagging behind.

Yet generative AI has the potential to also place interaction, i.e. exchanges between two economic actors, on a completely new digital footing.

This is possible first and foremost via three levers – cost reduction, reskilling/upskilling and the expansion of the product portfolio – which appear as circles in the figure at right. ► **B** In other words, the new AI models can, at least in theory, play a part in drastically reducing interaction costs, improving the quality of interaction and opening up completely new forms of interaction.


Importantly, these three levers cannot be cleanly distinguished from one another but are in fact interdependent. A closer look at the dynamics behind them illustrates the point.

## 1. COST REDUCTION

At least in the public arena, the debate surrounding generative AI is dominated by the resultant savings potential. It is widely argued that companies could literally slash their wage bill<sup>1</sup> once the potential for automation has been leveraged and machines have taken over many of the activities previously performed by humans. This, of course,

## B New dynamics

How generative AI works



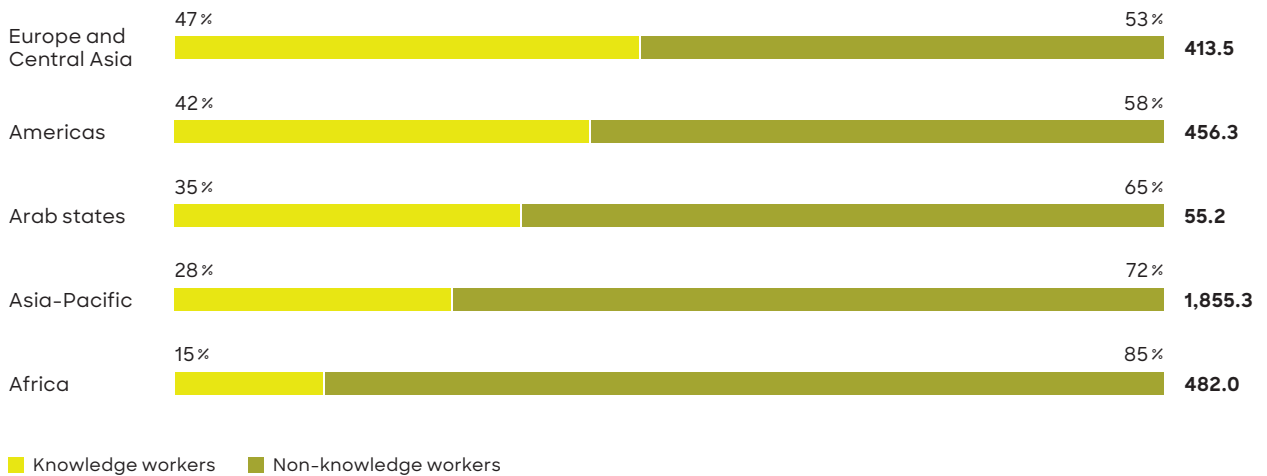
Source: Roland Berger Institute

is by no means a new discussion in the context of AI. What is new is that, for the first time, well-educated knowledge workers are at the center of substitution considerations. Why? Because, as never before, language AI tools can now be used to automate creative activities – such as writing more complex texts and programming codes – with relative ease. This category of knowledge workers accounts for a significant portion of employees in high-wage countries. ► **C**

<sup>1</sup> See (for example): Goldman Sachs, [The Potentially Large Effects of Artificial Intelligence on Economic Growth](#), March 2023; and the OECD Employment Outlook 2023, [Artificial Intelligence and the Labour Market](#), July 2023.

## C Focus on high-wage countries

Knowledge workers and jobs involving knowledge tasks [millions]



Source: ILO, Roland Berger Institute

## 2. UPSKILLING AND RESKILLING

Many studies show that the average knowledge worker today spends a significant amount of time on back-office activities. These include internal organizational arrangements, the search for, collection and dissemination of information, and assorted documentation obligations. Yet these are all routine tasks that could easily be handled by language AI in the future. The advantage for companies and employees is obvious: Knowledge workers would have more time to devote to more creative work that genuinely adds value for the company. That said, delegating routine tasks to AI is not the only way to realize efficiency gains: In complex work processes, too, interaction between humans and the speech AI model opens up completely new possibilities for increasing output – as attested by initial reports from the software industry.<sup>2</sup>

## 3. EXPANSION

There is much to suggest that the new AI tools will create entirely new products, services and professions. The list of possible use cases is already long. For example, an AI bot trained with proprietary knowledge could serve as a round-the-clock personal advisor for every investment decision made by bank clients. In retail settings, text-to-image generators could visualize the various purchase options and increase customer satisfaction many times over. Whole new lines of work are also already coming into view. One could be what are known as prompt engineers: highly trained specialists who can operate language AI in a targeted

<sup>2</sup> Peng et al., *The Impact of AI on Developer Productivity: Evidence from GitHub Copilot*, February 2023



manner so that the algorithms deliver optimal results. Other applications will be added, revolutionizing the business models of innovative companies in the process.

While it is still too early to precisely quantify the dynamics of these three levers, we expect the second scenario in particular to dominate. This means that generative AI will not lead to mass job losses, but will instead evolve into a new kind of digital co-worker that can be used at will for repetitive tasks. Companies will be more productive overall as a result.

Key players in the tech industry have long recognized the enormous economic potential of generative AI. Consider the initiatives of two major tech players – Microsoft and Google – for example.

OpenAI was originally founded as a non-profit company in 2015. In June 2018, it launched its first Generative Pre-Trained Transformer (GPT), ushering in the era of "large language models." The hotly debated ChatGPT-4 update followed in March 2023. This version is already said to deliver more accurate responses with a 40% higher probability than its predecessor, which was released just

## THE UNDERLYING TECHNOLOGY

### What is generative AI and what distinguishes it from other forms of AI?

To start with, generative AI differs from machine learning and deep learning in that it is no longer about automating learning processes but about generating new content – hence the name. However, there is a second specific feature in contrast to other forms of AI: Generative AI uses language and images to facilitate direct interaction on a user interface that can be operated by anyone, because it no longer requires the mastery of specific codes or a specific syntax. Instead, it is based on what are termed large language models (LLMs) and thus allows direct human-AI interaction. The advantages of such direct interaction between humans and AI are demonstrated by the fact that ChatGPT attracted 100 million users within two months, for example. An absolute record, one that no other app has achieved so far.

All of this would be unthinkable if it were not for recent exponential advances in the capacity of computing power, which has doubled every six months since deep learning approaches began in 2010. ChatGPT and its numerous rival products are now capable of performing a variety of human activities autonomously. The amazing thing is that everything is based purely on statistics: ChatGPT does not think like a human brain. It is not "creative": It merely collects data. The word-by-word (or, to be precise, token-by-token) text generation that takes place in models like GPT-4 can be thought of as similar to the autocomplete feature on smartphone keyboards, albeit with much larger volumes of data. The program simply uses statistics to assemble an answer from word modules to whose linguistic correctness the program's algorithms assign the greatest probability. Currently, AI models can be improved in two ways: either by enlarging the underlying computational pipeline that models use to establish relationships between and within datasets, or by improving the data input and/or data quality.

twelve months earlier. No wonder, then, that, earlier this year, Microsoft invested more than USD 10 billion in OpenAI, the maker of the ChatGPT language model.

Google is pursuing similarly ambitious goals. The tech company is currently striving to integrate generative AI in its search engines. To this end, Google invested USD 300 million in the AI startup Anthropic in February 2023. Two months later, the data giant launched its chatbot Bard. Other tech companies such as Amazon, ByteDance and Baidu are following this example and working on their own generative AI offerings. Nor should we overlook the wealth of other startups such as Stability AI, RunwayML and AI21, all of which are also working on promising products.

Further competitive pressure comes from a third direction: Open-source language models are increasingly entering the market and putting pressure on the established players. Facebook parent Meta has also opted for this approach. The hope is that it will be possible to develop generative AI applications even faster if the underlying data and programming codes are made available to the general public.

Which of these many models will ultimately prevail cannot be reliably predicted at the current stage of development. The technology is developing too rapidly for that. However, there is reason to believe that user companies will rely on various generative AI models for different use cases in the future.

Be that as it may, gold rush fever has broken out not only among the tech giants, but also among investors. While corporations like Meta, Amazon and Microsoft announced large-scale layoffs earlier this year and venture capital firms have been reluctant to invest, generative AI specialists are expanding right now and having little trouble finding backers. According to an analysis by venture capitalist Atomico, 35% of AI funding in Europe currently goes to generative AI projects – up from just 5% a year ago. ▶D


We have discussed the levers underlying the possible or expected potential of AI and how they operate. We have

also looked at the activities of companies and investors in response. But what does all this mean in concrete terms for those companies that want to realize this potential? Which parts of the company, which industries and sectors are the best suited? Where does the greatest potential for productivity gains lie? What labor market effects might be associated with this development? These questions are addressed in the next chapter based on the results of our survey.

There is reason to believe that user companies will rely on various generative AI models for different use cases in the future. ●●

## D Gold rush fever

Investment in generative AI startups  
(September 26, 2023)



Source: pitchbook

# Already well established in the business routine

What top managers hope for from  
the new tools



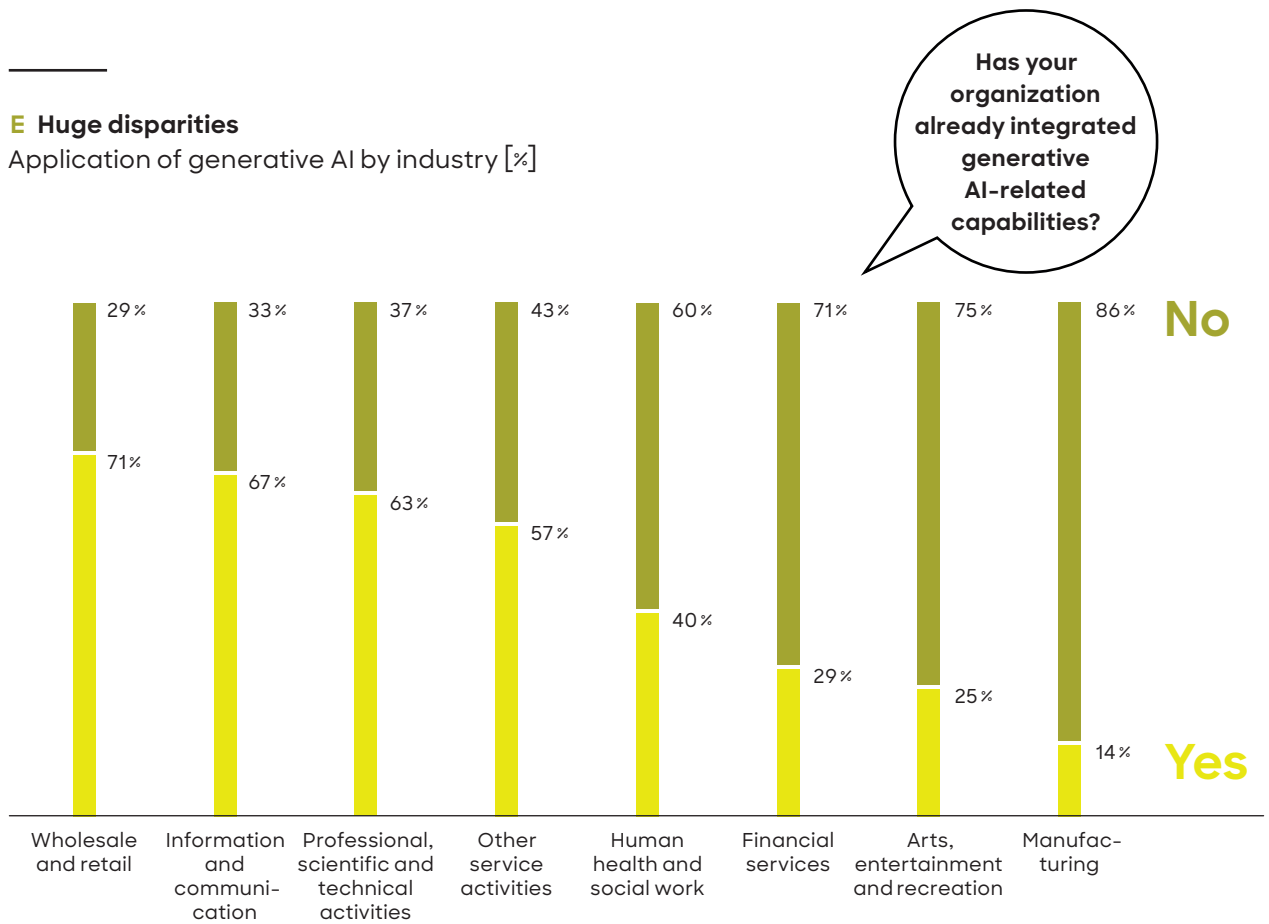
**O**ur survey results show that industries and sectors, but also individual functions and processes that are primarily concerned with economic interactions can gain particularly handsome benefits from generative AI. On the other hand, the new AI tools are of comparatively little significance in production contexts.

About half of our respondents (47%) have already deployed generative AI products such as ChatGPT and DALL-E 3 in their organizations.

However, there are some striking differences between different industries and sectors: While the new "tools" of generative AI are comparatively irrelevant in the area of traditional production, for example (only 14% of respondents from this area report already using such tools), the application rate in the four sectors Wholesale and Retail; Information and Communication; Professional, Scientific and Technical Activities; and Other Service Activities ranges from 57 to 71%. ► **E**

**E Huge disparities**

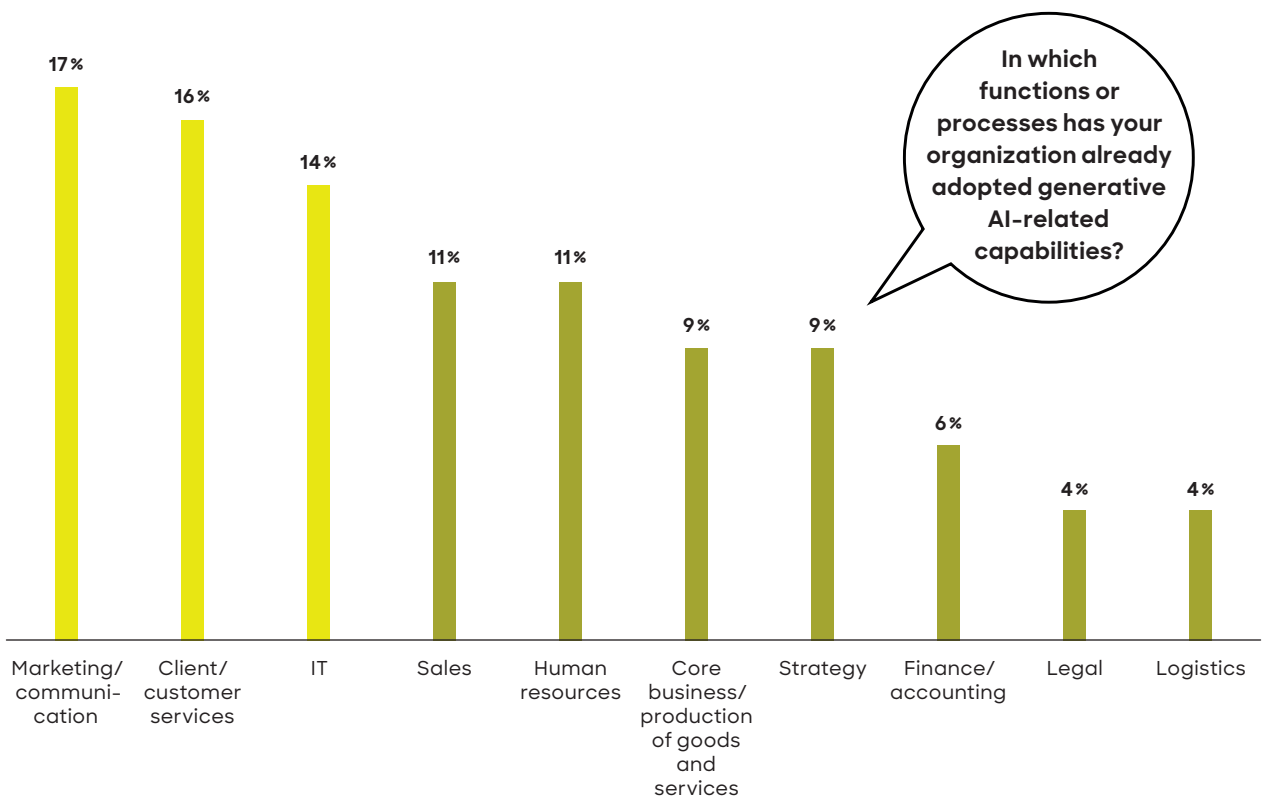
Application of generative AI by industry [%]



Source: Roland Berger Institute

## F The pioneers

Overview of functions and processes in which generative AI is currently applied



Source: Roland Berger Institute

The application of generative AI-related tools varies not only from industry to industry, however: Clear areas of focus can also be identified regarding those functions and processes in which generative AI is used to a greater or lesser extent.


In Marketing/Communication and Client/Customer Services, for instance, 17% and 16% of all companies

surveyed respectively have already introduced generative AI-related tools. In logistics and legal departments, on the other hand, these tools have so far only been used very sporadically. ►F

All in all, the experts surveyed believe that productivity can be significantly increased by 2025 thanks to generative AI. The vast majority (82%) consider an efficiency gain of

## G Optimistic outlook

Expected increase in productivity by 2025



Source: Roland Berger Institute


at least 6% to be realistic. Around one fifth of the survey participants even expect productivity gains of more than 20%. Only a tiny minority (3%) believe that the new options will not lead to any changes. ►G

Interestingly, the experts we interviewed do not expect these improvements in efficiency to lead to a drastic reduction in companies' headcount. Fully 57% expect to

see no change at all or even the creation of new jobs. Only 16% plan to cut jobs by 5% or more by the end of 2025. To put that another way: The executives surveyed believe that routine back-office tasks will increasingly be taken over by generative AI, and that employees will use the freed-up time for more creative and hence more productive tasks – a classic win-win situation. ►H

## H Job killer AI?

### Expected headcount change by 2025



Source: Roland Berger Institute



Two further arguments favoring such a scenario are described in our prior discussion of upskilling and reskilling. First, it has been shown historically that major technological advances create more new activities than they destroy in the long run. Second, unlike previous waves of innovation, generative AI, as we have seen, mainly affects knowledge workers in highly developed economies, i.e. in economic areas that are generally suffering from a shortage of skilled workers and from demographic change. This again makes the reskilling and upskilling variants more likely.

Initial findings in practice also suggest that generative AI will complement knowledge workers rather than substituting for them. One recent MIT study, for example, concludes that, with the aid of ChatGPT, the productivity of office workers in writing tasks (such as writing critical emails to the whole workforce and producing meaningful press releases) can be increased by an astonishing 40%.<sup>3</sup> Another study found that software developers complete tasks 55.8% faster when supported by Microsoft's GitHub Copilot.<sup>4</sup>

Without a doubt, these are impressive results. Yet they are possible only because humans play crucial roles in the value chain, feeding the AI and classifying its results. At least as important as the effect on overall employment levels, therefore, is the impact on job quality.<sup>5</sup> Properly implemented, there is significant potential to improve both employee satisfaction and job identification using the new language AI applications if they leave more time for "human" activities. Such activities could, for example, involve deepening customer relationships or structuring strategic issues of fundamental importance.

- Properly implemented,
- generative AI can improve both employee satisfaction and job identification if it leaves more time for "human" activities.

3 Shakked Noy & Whitney Zhang, [Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence](#), March 2023

4 Peng et al., [The Impact of AI on Developer Productivity: Evidence from GitHub Copilot](#), February 2023

5 ILO Working Paper, [Generative AI and jobs: A global analysis of potential effects on job quantity and quality](#), August 2023

# Limitations

Technical and practical obstacles  
to application

4

**A**s a rule, it is activities that are automated, not occupational groups. In many cases, doing the latter would entail too great a risk for companies. The limitations of the new AI generation are (still) too great to allow it to be integrated in everyday corporate life without human supervision. Five problems stand out in particular:

**1.**  
**PRIVACY**

Services such as ChatGPT use any user input to refine their own software. While this circumstance made the latest technological leap forward possible in the first place, it also drastically limits the scope of application: Copyrighted material, personal data, trade secrets and the like should therefore not be made accessible to language AI, as this could lead to massive liability risks for users. It is vital to prevent violations of the principles of intellectual property.

**2.**  
**TRANSPARENCY**

AI chatbots do not (yet) disclose how they work. Which facts are weighted and how? From what data basis was a result derived? All this remains the secret of the algorithm and is sometimes difficult to understand, even for its developers.

**3.**  
**MISINFORMATION**

On occasion, the new AI systems freely invent answers. In the tech industry, this inadequacy is commonly referred to as "hallucinations." These aberrations arise because the algorithm is designed to produce linguistically correct sentences that may not necessarily be correct in terms of content.

**4.**  
**BUSINESS INTEGRATION**

Deciding to use generative AI is one thing. But successfully integrating this new technology in practice is a far greater feat. The degree of complexity varies greatly from industry to industry. In some cases, embedding it in the organizational workflow is the biggest challenge. In others, the challenge lies in adapting the business model. Either way, the successful application of ChatGPT and similar solutions requires far more than just an investment decision.

**5.**  
**REGULATION**

How dynamically generative AI develops in the context of everyday business depends not least on the framework staked out for the new technology by governments. Various regulatory approaches are currently being discussed worldwide. While a consensus for stronger top-down regulation is emerging in the EU, for example, the USA seems to be taking the opposite path and is more open to risk. At the same time, the first signs of voluntary self-regulation in the form of new industry standards and best practice agreements are already appearing. They include collaboration in the USA between industry and the National Institute for Standards and Technology on the design, training and application of AI systems. Another example is work on what is known as "watermarking" systems, which are designed to verify whether content was created by a real person or a language AI system.

# Recommendations

Three steps to AI readiness



**T**he new language AI models have accomplished a feat that few new technologies can emulate: They have successfully established themselves on the market within a very short time. Moreover, they are developing so rapidly that their enormous potential is becoming ever more apparent in a seemingly endless stream of new use cases. No company should knowingly let this powerful productivity booster go to waste.

The initial experimental phase seems to be over for the time being. The relevant players have entered a new phase of maturity where it is no longer just about feeding more and more data into AI models. In addition to size and mass, the signs are now also pointing to quality.

It is therefore all the more important that companies identify the potential of this new technology for their own business model and develop a coherent AI strategy. Having done that, the new alignment should be tested in initial use cases and implemented step by step. Experience shows that companies can go from analysis (4-6 weeks) and strategy definition (4-6 weeks) to an implemented AI solution in four months. As a general rule, we recommend the following procedure:

## **1. ANALYZE THE IMPACT OF AI**

To develop an AI strategy, companies should first analyze the entire value chain to determine which stages are best and least suited to the three mechanisms described above. All further steps are contingent on this initial analysis. The *Roland Berger AI Readiness Radar* – the tool we recommend – reviews each link in a company's value chain to determine the extent to which the potential afforded by AI is already being exploited or is still being neglected. The findings from this research are then benchmarked against competitors' AI adoption and general industry trends. All this is flanked by analysis of the company's current IT infrastructure.

## **2. DEFINE YOUR AI STRATEGY**

After completing this inventorying exercise, we advise clients to define overarching goals and productivity benchmarks. The next step should then be to plan the implementation of initial use cases and define a framework of conditions to ensure that they succeed. Improving the management of all unstructured data sets is of critical importance here. To put it in a nutshell: Until you have genuinely clean data, not even the cleverest AI model can increase productivity. All process data must be prepared in such a way that established models can be fine-tuned to the given business purpose. Alternatively, it must be possible to develop special-purpose models from them for very specific tasks. The latter models require significantly less energy and computing power and are therefore far easier to set up.

## **3. ESTABLISH PROOF OF CONCEPT**

Once all the relevant technical details have been determined – such as selection of the AI model, the vendor and the integration plan – the implementation of the first use case can begin. This is followed by evaluation, the iteration of any improvements and, lastly, a company-wide rollout plan.

Experience has shown that the biggest obstacle is making the unstructured data that exists throughout the company accessible and analyzable in the first place. However, if a company does manage to make the whole of its organizational knowledge available via its own generative AI model or even a modified off-the-shelf model, the scaling potential is enormous. Agent-based systems can perform routine tasks for defined use cases, and a multitude of other apps and applications can be developed on this basis. Theoretically, this approach could put just about every interaction in a company on a completely new footing.

There is one important caveat, however: If sensitive company data flows into third-party models, this could also improve a competitor's value proposition and so turn into a tangible business disadvantage. To avoid such unintentional "second-order effects," security and privacy should therefore be given top priority. This means, for

example, that sensitive data must remain encrypted and in a secure environment at all times. Conversely, this also means that a clearly defined strategy must be used to decide what data goes into which model. The goal must always be to prevent any kind of IP conflicts and leaks from the word go.

**Experience shows that companies can go from analysis and strategy definition to an implemented AI solution in four months.**



# Credits

## AUTHORS

### **Jochen Ditsche**

Senior Partner, Global Head of Digital Practice  
+49 89 9230-8361  
jochen.ditsche@rolandberger.com

### **Hasmeet Kaur**

Partner, Global Managing Director  
+49 89 9230-8943  
hasmeet.kaur@rolandberger.com

### **Elias Aad**

Partner, Head of Digital Practice Middle East  
+971 4 446-4080  
elias.aad@rolandberger.com

### **Fabian Neuen**

Partner  
+49 40 37631-4344  
fabian.neuen@rolandberger.com



---


10.2023

**ROLANDBERGER.COM**

This publication has been prepared for general guidance only. The reader should not act according to any information provided in this publication without receiving specific professional advice. Roland Berger GmbH shall not be liable for any damages resulting from any use of the information contained in the publication.

© 2023 ROLAND BERGER GMBH. ALL RIGHTS RESERVED.

ROLAND BERGER is the only management consultancy of European heritage with a strong international footprint. As an independent firm, solely owned by our Partners, we operate 51 offices in all major markets. Our 3000 employees offer a unique combination of an analytical approach and an empathic attitude. Driven by our values of entrepreneurship, excellence and empathy, we at Roland Berger are convinced that the world needs a new sustainable paradigm that takes the entire value cycle into account. Working in cross-competence teams across all relevant industries and business functions, we provide the best expertise to meet the profound challenges of today and tomorrow.



## **Publisher**

### **Roland Berger GmbH**

Sederanger 1  
80538 Munich  
Germany  
+49 89 9230-0