

# Beyond the hype

How generative AI  
can unlock real value  
for Europe





## PREFACE

The vast majority of us will be familiar with sentiments such as “AI is an overrated technology with few use cases for the corporate sector. The AI market is a bubble.” Ironically, if you replace “AI” with “the Internet”, you find yourself back in a debate that raged 25 years ago. Although the dotcom bubble did indeed sober things up for a while, the prophecies of the time were nevertheless fulfilled. Alphabet, Amazon and Meta today rank among the most valuable companies in the world – except that they have now been overtaken by AI chip producer NVIDIA.

That why the hype surrounding (generative) AI has a feeling of déjà-vu to it: Yes, AI is already revolutionizing the world of business. But at the same time, the very real benefits and use cases are still hard for many companies to grasp.

This publication therefore takes stock of where Europe stands today in relation to AI and explains how companies and government agencies can make better use of the technology. Our focus is primarily on Europe’s frequently underestimated strengths: its data and the vast wealth of expertise that lies dormant in our companies. If we succeed in leveraging both to a greater degree, the EU can win back its seat at the global economic table.

This publication also marks a change within our organization. The Internet Economy Foundation (IE.F) was founded in 2016 to play a part in shaping Europe’s agenda for the digital economy and, in the process, to promote start-ups, scale-ups, bold investments, digital infrastructure and fair competition. Yet for all the considerable progress that has been made, Europe has fallen further behind. In response, we have renamed our foundation and will now position ourselves more clearly as an independent voice for European technological sovereignty.

The road ahead is clear: As we pass new legislation to open up the Internet economy once again, Europeans must learn to play more to their own strengths. We remain convinced that the EU has a bright future ahead of it. To realize this future, however, it must press ahead with innovation more quickly and more boldly. That is why our new name is Innovate Europe Foundation (IE.F). We invite all our readers to stand with us for a resilient and prosperous Europe.



**Clark Parsons**  
Managing Director  
Innovate Europe  
Foundation

# “Strengthen your strengths”

Three areas where action is needed to build an independent EU ecosystem

## 1. Transfer knowledge

We must leverage the domain knowledge possessed by European industry

## 2. Tear down dependencies

We must specifically reinforce sovereign technology solutions from Europe

## 3. Respond to the AI arms race


We must mobilize more private capital for European AI innovations

# CONTENTS

<b>PREFACE</b>	<b>3</b>
<b>1</b> <b>INTRODUCTION: THE GREAT MISUNDERSTANDING</b>	<b>6</b>
<b>2</b> <b>THE DOMINANCE OF BIG TECH AND A MISGUIDED FOCUS ON A SINGLE SUPER-LARGE LANGUAGE MODEL</b>	<b>9</b>
<b>3</b> <b>WHERE THE RUBBER HITS THE ROAD: WHY EVEN THE BEST TECHNOLOGY MISSES THE MARK WITHOUT DOMAIN KNOWLEDGE</b>	<b>18</b>
<b>4</b> <b>“STRENGTHENING YOUR STRENGTHS”: HOW TO SUCCESSFULLY BUILD AN INDEPENDENT ECOSYSTEM</b>	<b>26</b>

# 1

## INTRODUCTION: THE GREAT MISUNDERSTANDING



While the hype around generative AI (GenAI) has not yet abated, a degree of disillusionment is already doing the rounds on the shop floor. True, many companies and government agencies have begun to incorporate AI language models in their processes. Yet the hoped-for breakthrough – broad-based implementation in the corporate sector (B2B) and in government (B2G) – has (so far) failed to materialize. This is due less to the technology than to the misleading criteria used to measure its “success”.

All too often, the focus is on what at first glance appear to be the impressive capabilities of the major all-round models such as OpenAI. Applications in the B2C segment in particular have attracted considerable media attention. In contrast, however, the practical needs of companies and public administrations are increasingly slipping into the background. As a result, many professional users have not yet progressed beyond the experimental phase. Alternatively, they are indulging in expensive showcase projects that, quite clearly, will not deliver the targeted gains in productivity anytime soon.

Against this backdrop, this policy paper ventures an honest look at where we stand today. We address the issue of why, despite the justified hype, the technology has not (yet) revolutionized everyday corporate and government life. We also stake out the conditions that must be met in order at last to realize the vast potential of GenAI.

**The practical needs of companies and public administrations are increasingly slipping into the background.**

One key aspect of our explanation is that too many hopes are being pinned on the technology alone. In practice, properly embedding the technology in internal processes and linking it to employees' knowledge in specific domains is at least as important to the success of this venture. We understand this "domain knowledge" to be the specialist expertise accumulated by a skilled worker in a given industry or field over many years. Without this feedback loop to the world of practitioners, even the best AI will be unable to realize its full impact.

Chapter 2 explores the causes of this one-sided focus on technology. We trace how competition for the largest language model has accelerated and grabbed all the attention. As is so often the case, this competition is dominated by the big US tech companies who play in a league of their own as far as investment, partnerships and infrastructure are concerned. Their dominance has enticed pessimists to revive the time-honored lament that the Europeans lag hopelessly behind. As we see it, this one-sided perspective does not match the reality.

Chapter 3 draws on practical examples to demonstrate that the greatest benefits do not necessarily derive from what are reputed to be the most powerful large language models (LLMs) simply because they are the biggest. In many cases, smaller models tailored to a given application can add significantly more value in practice. Conversely, general benchmark tests can easily prove misleading because they pay little or no attention to specific user situations.

That said, Europe must definitely accelerate its endeavors. As our recommendations in the closing

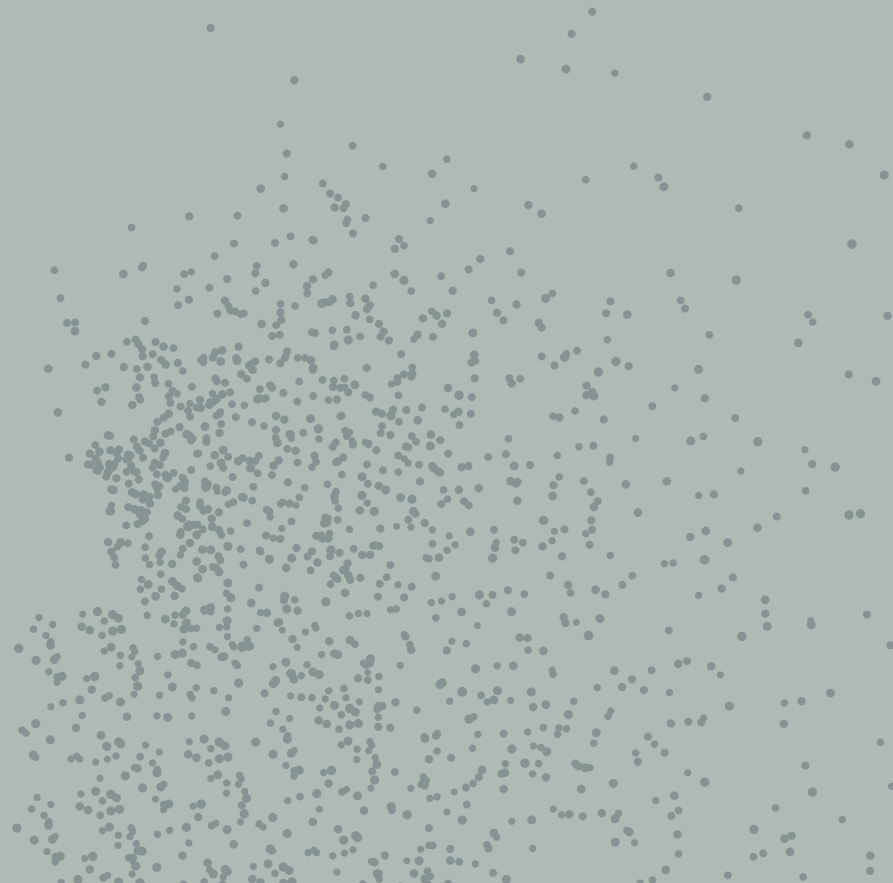
chapter point out, better use must be made of European industry's domain knowledge and unilateral infrastructure dependencies must be dismantled. The same goes for the shortfall on the capital markets: It is high time to close the gap. Only a largely independent ecosystem that brings together technology, domain knowledge and infrastructure can lay the foundation for AI-driven value creation in Europe and get the continent's economy in shape for the future.





# 2

## THE DOMINANCE OF BIG TECH AND A MISGUIDED FOCUS ON A SINGLE SUPER-LARGE LANGUAGE MODEL



Given the massive competition that exists for the most powerful generative AI, it is no wonder that all eyes are currently on huge language models. As is so often the case, the focus is on the tech giants in general and Alphabet, Amazon, Meta and Microsoft in particular. These behemoths are fighting a pitched battle for market supremacy, announcing new advances almost by the day: Together, these four corporations want to invest 200 billion US dollars in data centers, chips and other equipment to run and operate GenAI models in 2024 alone. That is a year-on-year increase of 45% and a new record. Over the next five years, this quartet's capital expenditure will likely rise to as much as a trillion US dollars. →A


However, these investments are not restricted solely to expanding their own infrastructure. Increasingly, big tech is also stepping up as a venture capitalist and, in the process, securing access to the latest technology from ambitious start-ups. In a departure from the logic of traditional financial investments, they instead aim to consolidate their market dominance. The partnerships between Microsoft and OpenAI on the one hand and Microsoft and French start-up Mistral on the other are only two developments that underscore this trend. Google and Amazon are likewise supporting Californian start-up Anthropic to the tune of billions of dollars.

---

## A The 200-billion-dollar gamble

---

Capital spending by the four AI titans [USD bn]



---

Source: Capital IQ, Forbes

---

2  
The dominance of big tech and a misguided focus on a single super-large language model

This new engagement is also showing up in the venture capital (VC) statistics. According to the analysts at Pitchbook, two thirds of the total of 27 billion US dollars that were invested in up-and-coming GenAI firms in 2023 came from Alphabet, Amazon and Microsoft. In contrast, traditional VC providers are increasingly falling behind. →B

The decision for a VC investment and against an acquisition is one of the key differences to previous waves of innovation. Major takeovers like Meta's acquisition of WhatsApp and Instagram a few years ago are not on the agenda right now – probably due in part to worries about the antitrust watchdogs. Notwithstanding, fears remain about the independence of GenAI start-ups that are funded by big tech.

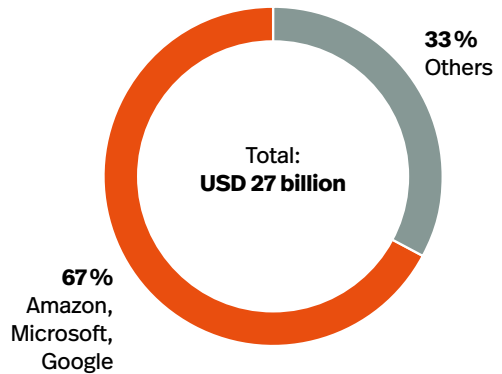
Thanks to their dominant market position, the tech titans are also having a field day recruiting the few AI talents on the market. Research firm Live Data Technologies says that Alphabet, Amazon, Apple, Meta and Microsoft have snapped up at least 30 leading AI experts from OpenAI, Anthropic and Cohere since 2022. And only a few months ago, in March 2024, Microsoft also announced that it was taking on almost the entire workforce of Inflection AI. Among the new recruits is Mustafa Suleyman, co-founder and former head of the applied AI department at DeepMind. When he left DeepMind in 2022, he and other partners teamed up to launch Inflection AI, a company that has specialized in machine learning and generative AI. In June 2023, Inflection won one of the biggest rounds of AI funding with a volume of 1.3 billion US dollars. One of the main investors was Microsoft, where Suleyman is taking

---

**B Start-up funding instead of start-up takeovers**

---

Big tech's share of US venture capital investments in generative AI, 2023



---


Source: Pitchbook

---

## C An expensive business

---

Estimated training costs for selected LLMs [USD m]



---

Source: Stanford University

charge of the new Microsoft AI division. Again, power in the international AI research and development space is being concentrated still further.

The antitrust authorities are not prepared to stand idly by and watch this development unfold. To take one example: In a recent interview with the *Financial Times*, Jonathan Kanter, the top expert on monopolies at the US Department of Justice, stated his intention to look more closely at what are known as “acqui-hires”. The term refers to the acquisition-like recruiting of employees with which dominant companies buy start-ups’ intellectual property. The same goes for access to hardware and other critical resources that could restrict free competition.

Nor is spending on human resources and equity investments the only activity devouring ever larger sums of money. Development costs for super-LLMs are likewise going through the roof. Individual companies rarely disclose information about the exact training costs for their models. However, the duration of training and the type, scope and utilization rate for the training hardware allow the volume of investment to at least be approximated. Working on this basis, Stanford University, in collaboration with AI research institute Epoch AI, came up with training costs of around 78 million US dollars for OpenAI’s GPT-4 and 191 million US dollars for Google’s Gemini Ultra. The training costs for the original transformer model that introduced the architecture and underpins virtually every modern LLM totaled around 900 US dollars in 2017. →C

It is worth noting that these are only the costs for successful training runs. The hundred or so failed

experiments that usually precede one successful one are not yet priced into this calculation. The need for capital thus quickly becomes gigantic, even though it is not clear whether the benefits are in proportion to the investment. Developing the next generation will probably be even more expensive.

**The development costs for super-LLMs are going through the roof, even though it is not clear whether the benefits are in proportion to the investment.**

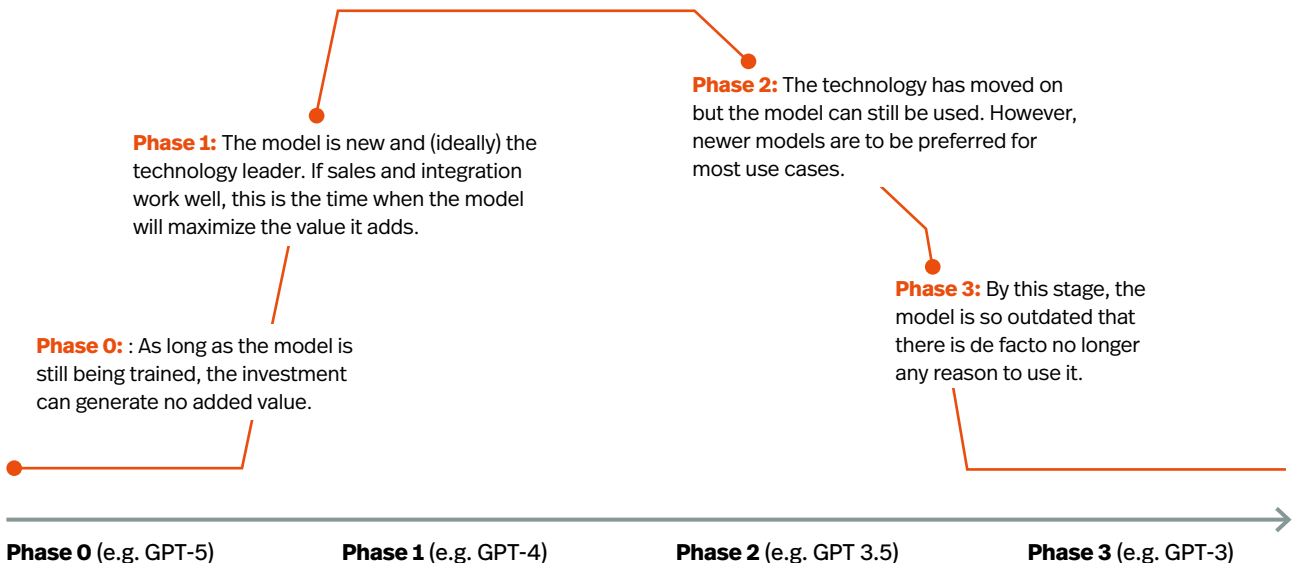
These astronomical training costs are problematic for another reason: Each AI language model only enjoys a relatively short half-life that has recently shrunk to about a year. After that, the LLM barely adds any more economic value. Why? Because every successive model that is based on a new software architecture has to be trained from the ground up. This fact marks a radical distinction between GenAI products and previous software innovations that remained fundamental to all subsequent product adaptations for years or even decades.

Similarly, the hardware needed for training has a relatively short lifecycle of its own. In extreme cases, little therefore remains in the medium term of the hundreds of millions of euros (dollars) invested in chips, infrastructure and operations – aside from reputational gains, the customer base and the user data (in models with no data sovereignty). →D

The exorbitant cost of training and running super-LLMs currently constitutes a sizable barrier to practical

## D A short half life

Added value generated by investments in AI language models



Source: Innovate Europe Foundation

2  
The dominance of big tech and  
a misguided focus on a single  
super-large language model


application. On the other hand, it should be noted that the willingness to pay is still limited on the demand side. While the big tech firms say nothing about how successful (or otherwise) the monetization of their models is, estimates indicate that expectations have not yet been fulfilled.

Despite this, the three dominant cloud providers Alphabet, Amazon and Microsoft are already profiting from the GenAI revolution. They control the infrastructure that is

required for both the training and use of language models. No other company can make the requisite computing power available on a comparable scale. Above and beyond their dominance in AI models, their primacy in cloud services also explains why the market value of the three big providers has increased by 2.5 trillion US dollars since November 2022. Right now, the capital markets in general likewise see the greatest potential for added value not in the industry but in the tech giants – witness their share price development in recent months. →E

E AI's current value-added potential from investors' perspective

Trend in the Shiller P/E index since 2023 for selected sectors




Source: GuruFocus

---

## F The USA as the undisputed leader

---

Number of newly funded AI firms by country, 2013-2023



---

Source: Stanford University



In view of the dominant position of US tech companies, it is hardly surprising that the capitalization gap between the USA and Europe has not narrowed. On the contrary, the imbalance is worsening. In 2023, the USA once again led the rankings for private investment in generative AI, followed by China and the UK. At 67.22 billion US dollars, the investment volume in the USA was 30 times higher than that of the European Union. Regarding the number of newly launched AI firms, a similar picture emerges: Here again, the USA leads the rest of the world by some distance. The familiar question is thus posed with increasing frequency: How can Europe close the gap on the USA? But is this even the right question to ask? Or is it not far more important to ask how the EU can strengthen its position in the international battle for talent and resources without submitting to any benchmark binges? →F

US companies' technological leadership is undisputed. Indeed, a glance at current investment volumes suggests that their dominance will likely be further reinforced in the years ahead. Yet despite this, there is no reason for excessive pessimism. Alongside the powerful all-round models, the big tech corporations have also ramped up their commitment to working on smaller and more application-oriented models. As the next chapter explains, the breakthrough in practice will largely depend on how these special (context-sensitive) models are embedded in the corporate routine.

**The breakthrough  
in practice will  
largely depend on  
how special (context-  
sensitive) models  
are embedded in the  
corporate routine.**



3

**WHERE THE RUBBER  
HITS THE ROAD: WHY EVEN  
THE BEST TECHNOLOGY  
MISSES THE MARK WITHOUT  
DOMAIN KNOWLEDGE**

Experience gained in previous phases of upheaval shows that it can take a while before the potential of a new technology impacts the whole of the economy. The average company and the average employee need time to progress through certain learning curves and experiential processes. Experts who are able to adapt new technologies, integrate them and redefine value creation processes must first be trained. And the past has shown that it takes time for these developments to filter down and permeate corporate practice. Before the PC era, it took at least a decade until computers became a common feature of everyday office life. The same will be true regarding the use of extensive generative AI applications: It will take a while before expectations of steep productivity gains come to practical fruition.

It is true that promising use cases already exist, and their number is growing almost by the day. According to a survey by the American Census Bureau, however, comparison of different sectors of the economy shows that it is mainly firms in the information industry that are currently experimenting with generative AI to an above-average degree. In contrast, implementation tends to be a slow process at manufacturing companies, even though very few doubt the potential of AI language models to add value.

The fact that companies are not yet incorporating generative AI in their processes to a far greater extent is due to a whole series of practical obstacles. The worst of these is the lack of expertise among user companies, according to a survey by IT group IBM. In second and third place we find problems that relate closely to this deficiency: Either the companies' data complexity is too

great, or the projects cannot be properly integrated and scaled in the context of internal processes. →G

The lack of knowledge and the difficulty of embedding this technology in existing processes also explains a further issue: As things stand, larger companies with 250 or more employees are substantially overrepresented among the early adopters. They are the ones most likely to have sufficient resources to secure the loyalty of AI experts and invest the requisite sums in processes and infrastructure – or to have enough capital to buy in such expertise as a service. For smaller firms, that is much harder.

Examining real-world practice also makes it clear that a larger and more powerful AI language model in itself is not enough to overcome the obstacles to use. On the contrary, trusting in super-LLMs alone could actually widen the gulf between technology partners on the one hand and industry on the other. Nor is training large AI language models the only expensive aspect: Operating them too requires vast computing capacity, and that is a deterrent for many users.


Yet size alone is not the only criterion. Analyses show that models with a smaller number of parameters can definitely compete with the logical capabilities of larger models and even outperform them in critical areas. Parameters are generally understood to be numeric values that are learned during training. They determine how a given model interprets input data and what forecasts it arrives at. Here again, though, more is not always better. The quality of the data sets also has an important part to play. For example, if repetitive or

---

## G A lack of expertise

---

Biggest obstacles to the successful use of GenAI applications in practice



Status: November 2023

Source: IBM

---

3  
Where the rubber hits the road:  
Why even the best technology  
misses the mark without  
domain knowledge

irrelevant data can be filtered out intelligently and carefully, smaller models can indeed achieve similar levels of performance.

Language models that are trained using carefully curated data sets are not only more cost effective but also present an added advantage: Customers may be able to run them on their own premises – i.e. in their own hardware environment, instead of beaming information

to the cloud – with less computing power and superior energy efficiency. Especially for privacy-conscious users, this could be a decisive consideration. →H

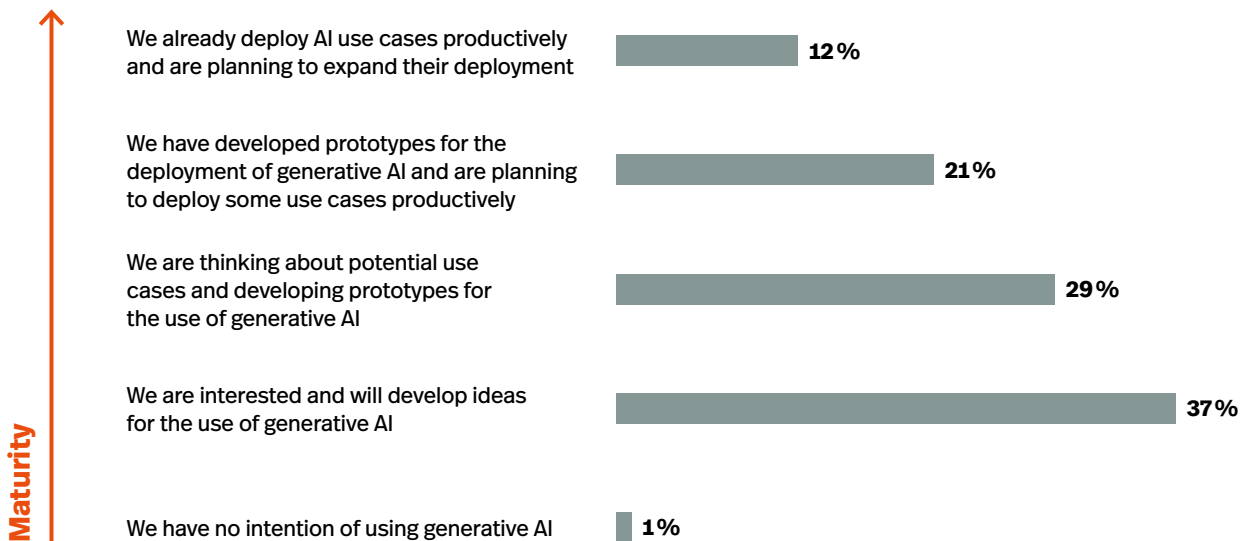
Existing use cases also provide evidence that companies are well advised to select their AI language model based on the requirements of specific use cases. For example, one model may perform far better than a rival model in general benchmarking tests but nevertheless produce

---

## H Tentative start

---

Current status of GenAI usage across all industries



---

Source: Coleman Parkes Research

inferior results in categories of relevance to a particular application. In such cases, choosing the nominally “better model” would be the wrong decision, because the selected model is less well suited to the use case for which it is needed. A sound understanding of real-world applications and a thorough evaluation of benchmarking outcomes will therefore determine the success or failure of an AI language model in practice.

Another relevant consideration: In the view of researchers at Stanford University, existing benchmark tests that measure the capabilities of a model in different areas of application have in any case already reached their useful limits. Accordingly, these tests are referenced ever less frequently to assess performance. Under certain circumstances, the models may even have been trained for a given benchmark test with special catalogues of questions, which would obviously undermine the validity of the results. It is therefore helpful to look very closely at those areas of relevance to a particular use case and not simply to go for an all-round model. Models that focus on “following visual illustrations”, “moral thinking” or “planning” would be three examples. Ultimately, though, it is up to the user to define what requirements are needed in which areas. Only then can the best model for this context be identified, while the technology partner can provide any necessary proprietary fine-tuning.

As we have already seen, the choice of language model is by no means the only or all-important success factor. Depending on the company and the specialist discipline, ever greater importance is now being attached to individual fine-tuning. In this way, every user can

essentially operate their own LLM – or multiple LLMs simultaneously – to which only their own staff have access. This is also a good way to protect full sovereignty and prevent the exodus of domain knowledge via interaction data. Experience with use cases to date suggests that precisely this combination of technology coupled with domain knowledge and a grasp of the company’s specific processes adds considerable value. In practice, this could play out along the lines of the following three use cases:

## CASE 1

# How a virtual AI assistant can actually make production more profitable

Radical transformation of the manufacturing landscape necessitates the continuous optimization of production processes to reduce downtime and cut costs. The use of AI-driven virtual assistants can provide valuable support along the way. These assistants enable routine tasks to be automated, and better decisions can be made based on real-time data synchronization. However, if this kind of virtual assistant is to genuinely make production more profitable in practice, it must be tailored perfectly to the given processes. Since factory workers often simply do not have the time to type questions into a computer interface by hand, for example, it helps if staff can interact directly with a chatbot.

In many cases, international experts are involved in the production process. In such constellations, the option of using the AI assistant to communicate in different languages is a critical factor, even if the user

manual is only available in, say, German or English. The same goes for communication via visuals. In this particular use case, this capability is more important than other main categories that are assessed in traditional benchmark tests. If it works well, a factory specialist could, for example, photograph the position of a robot and ask the AI assistant whether it is safe. The efficiency gains and reduced outages that can be achieved by closely intertwining this kind of special-purpose AI model with the production process are obvious for all to see.

## CASE 2

# Genuinely adding value – rather than short-lived “efficiency gains” – in customer service

Using generative AI to optimize customer service is a classic use case that has already been trialed for a long time. The basic idea is this: By accessing real-time decision aids based on a knowledge of every case the company has ever handled, customer service staff can resolve tickets better and more efficiently. Ideally, this kind of AI service agent could itself deal with many standard questions before they are passed on to humans. That would leave human employees more time to tackle exceptionally difficult cases or provide individual support.

In practice, however, integrating a service agent often proves to be far more complex. Witness the case of an AI-assisted breakdown service: Analysis of the customer data showed that different customer groups wanted accidents to be processed in different ways. Whereas Gen Z generally preferred self-help tools, other groups wanted to communicate directly with a

member of the customer service team after a certain point in the process. The challenge was therefore to train the AI service agent to make this important decision, and not to target short-lived “efficiency gains” by deploying it for everything. Only once the system had been optimized in this way was productivity genuinely improved. Customer satisfaction levels of 98% were then reached.



### CASE 3

## The vast potential of human-machine collaboration in compliance contexts

At the present time, compliance audits number among the most difficult use cases for AI language models, for several reasons: One is that what are usually strictly confidential contracts and documents have to be examined. Another is that the content is highly complex and very specialized. Moreover, any mistakes can do huge damage to a company’s reputation. Understandably, many companies therefore see it as too risky to outsource audits to AI in their entirety.

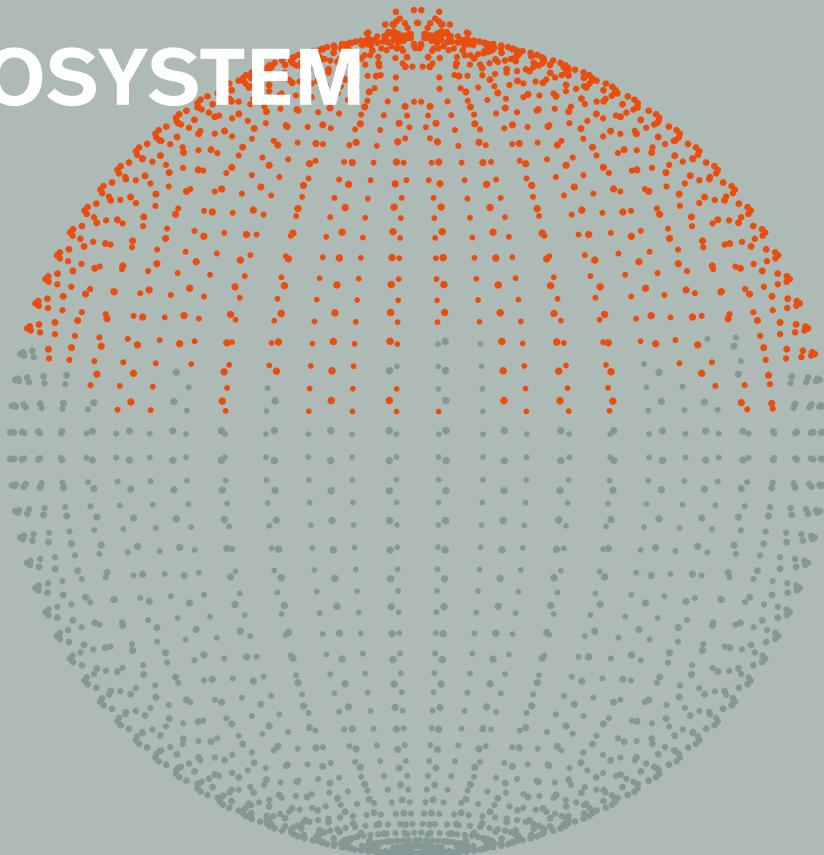
In practice, it nevertheless turns out that human-machine collaboration can add a great deal of value. The EU’s new Digital Operational Resilience Act (DORA) illustrates the point: The purpose of DORA was to establish a unified framework for the effective and comprehensive management of cybersecurity risks on the financial markets. Specifically, the companies affected have to prove that there is no danger to the security of information systems even in the event of a serious business interruption.

Until recently, this was a painstaking and, above all, cost-intensive process during which sometimes hundreds of contracts per company had to be audited. Initial practical tests are now showing that compliance audits can be made many times more efficient with the aid of AI. But to make that happen, a series of conditions must be met. First, a special-purpose model is needed that has been trained precisely to recognize the relevant contractual clauses. Conventional general-purpose models are not suitable for this requirement.

Interaction between the human specialist and the AI is equally important. The model must receive transparent instruction covering all room for interpretation and all inherent uncertainties. Only then can the human compliance expert make a final decision successfully without extra work and without decisions being made on a defective basis. This also means that feedback loops must be integrated in the workflow in such a way that, while it is in use, the model continually learns and is fine-tuned on the basis of existing domain knowledge. The success of this use case also depends on the availability of offerings that guarantee user sovereignty. Understandably, many companies insist on maintaining absolute sovereignty over their confidential data. In cases of doubt, this means that it must also be possible to run the AI language model from the company’s own data center.

# 4

## “STRENGTHENING YOUR STRENGTHS”: HOW TO SUCCESSFULLY BUILD AN INDEPENDENT ECOSYSTEM



The experimental phase we have seen to date shows that trusting in the potential of technology alone is not enough. Simply assuming that productivity gains will come about automatically is not a good idea. Experience shows that the best results are always achieved when the choice and application of the language model are guided by subject-specific and domain-specific considerations. And this is where technology partners, industry and the government must all step up to the plate: Their job is precisely to facilitate a harmonized triad of “technology, processes and people”. However, this can only succeed within a smooth-running ecosystem of potent technology firms, innovative industrial companies and brilliant minds.

Three basic conditions must be met to begin building such an ecosystem: First, industrial companies must utilize their people’s domain knowledge to apply AI-based technologies. Second, these technologies can only be operated if a data infrastructure shaped by competition, diversity and interoperability is in place. Lastly, investment and innovation require an attractive, modern and competitive capital market.

## 1. Leveraging European industry’s domain knowledge

Generative AI can drive superior productivity and hence greater prosperity only if the new technology is adopted by broad swathes of the economy. The successful implementation of artificial intelligence hinges not just on the technology itself, but on how it dovetails with corporate processes based on the domain knowledge of the workforce.

If that is to happen, both the company and its staff need to be aware of the importance of their knowledge. Yet we cannot simply assume that this is the case: Most of the knowledge possessed by very experienced experts in particular is essentially implicit. This “silent reserve of knowledge” is not formalized and is closely bound up with processes and activities. The successful use of AI-based technologies will always depend on this knowledge – which is why employees must be involved in developing concrete use cases for generative AI.

At the same time, it is important for them to have an adequate understanding of the AI technologies they will be using. In part, this is so that they can play an active role in developing the specific use case. But their involvement is also needed to ensure that friction-free processes exist at the interface between humans and AI. To put that another way: Very extensive context-specific expertise is no substitute for an understanding of digital

technologies. On the contrary: Especially those staff who possess extensive domain knowledge also need digital skills to ensure that generative AI can be deployed successfully. That is the reason why training and advanced education in the arena of digital technology and generative AI is imperative.

Subject-specific knowledge can be combined with digital skills and technologies in this way only if the right corporate culture prevails and a suitable organizational structure is in place. Regarding corporate culture, great importance must above all be attached to lifelong learning. Advanced education must be communicated and understood as an opportunity, not a burden. Beyond that, mistakes must be perceived as an incentive to improve. The organizational structure must then make sure that staff who possess specialist expertise collaborate seamlessly with the AI experts. To this end, fields of responsibility must be clearly defined, as must appropriate incentives.

## **2. No unilateral dependencies on infrastructure**

At a time when GenAI is booming, Europe is paying a double price for its tentative approach to ramping up home-grown cloud and data centers. Both the vendors and the users of language models need access to huge amounts of computing power. That is why the offerings of the three market-dominating American cloud services are often so attractive: These players are in a position to provide both the infrastructure and the AI applications within this infrastructure as well as the ancillary services. And they are increasingly exploiting their market power – to some extent because doing so will, for the foreseeable future, make it harder for their European counterparts to compete.

Governments must therefore master a tricky balancing act: If they restrict market access for US-based cloud providers, European companies could suffer a disadvantage due to lower productivity. Why? Because it would take years for Europe to build an equivalent cloud ecosystem of its own. Yet at the same time, it is vital to rein in the anticompetitive behavior of individual market players to a greater extent.

What are referred to as lock-in strategies, for example, have been known about for years but have never been satisfactorily prohibited. Even the EU's Digital Markets Act is likely to mark only a first step, especially as it must

now be rigorously enforced. Some cloud service providers still manage to prevent customers from defecting to other suppliers by a combination of contract design, technical barriers and high transfer charges. Governments and antitrust authorities must intervene more energetically on this score. They must prevent cloud services from demanding excessive fees for the outflow of data while making data inflow available free of charge. Discount tactics that seek to tie companies to one provider for longer periods should also be restricted, while strict interoperability prescriptions should become compulsory.

It should also be pointed out that, while European cloud firms do have a presence on the market, further growth is vital if they are to stand up to US competition. Governments could stimulate the growth of European cloud firms, for example by instructing the public sector to make greater use of their services. To make the use of resources as efficient as possible, precedence should be given to strengthening demand for sovereign offerings rather than funding GPU clusters with public money.

### 3. Closing the gap on the capital markets

The GenAI arms race has accelerated to a frantic pace. European firms in the growth phase will therefore need commensurate capital resources at their disposal to be able to compete. Even before the current AI boom, European companies had a far harder time stumping up the money they needed than American companies did.

If more private capital is not forthcoming, it will quite simply not be possible to fund the EU’s digital transformation. Right now, however, there is still no sign of a cross-border EU capital market (i.e. a cross-border market for capital instruments traded on and off the floor), despite the progress made of late. The obstacles – relating to insolvency and stock market law, for example – remain too daunting to leave any hope that these laws will be successfully harmonized any time soon. Governments should not be discouraged, however, but should continue to pursue the path of incremental improvements that they have already embarked upon.

A first step would be to strengthen institutional investors. By expanding supplementary funded pension provisions, for example, pension funds could accumulate more capital and then be placed under obligation to invest a reasonable proportion in the venture segment. Modernizing existing investment regulations would also be beneficial. The aim should be to get European pension funds to follow the Americans’ lead and invest in higher-risk instruments to the same extent.

Despite these huge challenges, it is important not to forget one good piece of news: With less competition between investors in Europe and in light of the lower influx of capital, returns in the European VC sector have lately exceeded those of the USA. Private investors too would therefore benefit from investment in this area. For that to happen, two conditions must be met: First, educational measures at the national level and across all age groups should improve people's financial literacy and promote a smart portfolio diversification strategy. Second, broader pan-European financial products should be launched and made available to private investors in particular. The vision behind this proposal is simple: to develop an ecosystem for European investments that is worthy of the name.

**Pan-European financial products are an important step toward an ecosystem for European investments that is worthy of the name.**

## **Publishers**

**Innovate Europe  
Foundation (IE.F)**  
Uhlandstraße 175  
10719 Berlin  
[www.ie.foundation](http://www.ie.foundation)

## **Authors**

**IE.F**  
**Clark Parsons**  
[c.parsons@ie.foundation](mailto:c.parsons@ie.foundation)

## **Contact**

**IE.F**  
**Clark Parsons**  
Managing Director  
Innovate Europe Foundation (IE.F)  
[c.parsons@ie.foundation](mailto:c.parsons@ie.foundation)  
Tel.: +49 30 8877 429-400

## **Publication date**

July 2024

## **Disclaimer**

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. IE.F shall not be liable for any damages resulting from any use of the information contained in the publication.

