

AI for Austria

The importance of inclusive dialogue, a common governance, and the transformative potential of AI.

Imprint

© 2024 Microsoft.

All rights reserved. Names and products of other companies may be registered trademarks of their respective owners. This document is provided 'as is'. The views and information contained in this document (including URLs and other references to websites) are subject to change without prior notice. You bear the risk of use. This document does not constitute legal or regulatory advice and does not contain any representation or contractual obligation on the part of Microsoft. Please seek independent legal advice regarding your legal and regulatory obligations. This document does not grant you any intellectual property rights in any Microsoft product. You may copy and use this document for internal reference purposes.



Table of Contents

The Promise of AI for Austria	<u>5</u>
Our Investment in Austrian Digital Infrastructure	<u>7</u>
AI Explained	<u>13</u>
AI Impact on Austria	<u>21</u>
AI Case Studies	<u>26</u>
AI for Innovation	<u>28</u>
AI for Government	<u>43</u>
AI for Healthcare	<u>47</u>
AI for Energy	<u>53</u>
AI for Education	<u>59</u>
AI for Sustainability	<u>63</u>
Deployment of Responsible AI	<u>69</u>
How do we Govern AI?	<u>75</u>
Microsoft's Commitments	<u>81</u>
Sources and references	<u>85</u>
Appendix	<u>86</u>

The promise of AI for Austria



Hermann Erlach
General Manager
Microsoft Austria

Endless opportunities

AI is leading the way in cutting-edge technology, enabling organizations to improve their productivity and competitiveness while also enhancing the potential of individuals. A recent study on the use of AI and its economic effect on the Austrian GDP reveals a growth potential of 18 percent within the next 10 years, which is equal to the economic output of Vienna and Styria together.

AI will change the way we work and foster innovation at scale in various domains, including health care, education, the environment, business, and the wider economy. While AI presents exciting possibilities, it also brings forth challenges – ethical, legal, and social. It is essential to develop and deploy generative AI in a responsible and trustworthy manner, one that respects human values, rights, and dignity and ensures fairness, transparency, and accountability. As a global technology company and a pioneer in AI research and innovation, Microsoft is committed to advancing and supporting responsible AI in collaboration with industry and civil society stakeholders, as well as among policy makers.

With our extensive Austrian partner network and deep customer relationships across the region, Microsoft is well placed to help bring stakeholders together to advance the responsible use of AI. Our mission is to empower every person and every organization to use and trust our technology. It's both a privilege and responsibility that informs our approach to AI.

The opportunities are boundless, and the time to seize them is now. The industry is just starting to scratch the surface on ways in which AI can generate value. Together, we can shape the future and unlock the full power of AI in our communities and economies.



Our investment in an Austrian digital infrastructure

How our cloud region will provide access to cutting-edge technologies for all businesses and green growth in Austria.

The Cloud as a Catalyst: Growth through Digitalization



Florian Slezak
Cloud Region Lead
Microsoft Austria


Providing access to cutting-edge technologies for all businesses through Austria's data center region

At Microsoft, we are taking steps to ensure that Austria makes the most of our AI era. Our investment in the country's first full cloud data center region, a specific geographic area where Microsoft has multiple data centers to ensure redundancy, availability, and low-latency access to cloud services, represents not only a response to the increasing demand for advanced cloud and AI technologies, but also a clear commitment to digital sovereignty and sustainable growth in our country. Microsoft's on-premises cloud region goes beyond technology delivery - it's an innovation program that will serve the country, people, business, and politics. Cloud and modern data centers democratize innovation by granting companies of all sizes cost-effective access to cutting-edge technologies. This levels the playing field for small-to-medium enterprises, allowing them to compete with large corporations by creating efficient usage options without investment or entry barriers, regardless of where in Austria these companies are geographically located.

The cloud as catalyst: growth through digitalization

In today's connected world, digital advancements are critical to economic success. This became even more apparent during the pandemic and the subsequent supply chain crisis. Given that digital technologies have a clear advantage, we want to ensure that we are making the best possible use of this opportunity in Austria.

A [study by EcoAustria, Accenture, Avanade](#) and Microsoft Austria provides fascinating perspectives on this. An increase in digital competitiveness of just 1 percent could increase gross domestic product per capita by 0.16 percent by 2030 – an impressive potential of 100 billion euros in overall economic value. Most of this potential can be leveraged through increased use of cloud services.



Microsoft is actively shaping this digital growth in Austria by investing in on-premises cloud infrastructure, which not only creates added economic value but also achieves real, sustainable cost effects. [Spar](#), for example, uses Microsoft Azure's AI capabilities to make precise predictions about the optimal demand for goods in its stores, reducing food waste. [TietoEvry](#) has developed a system that provides winemakers and farmers with detailed 48-hour forecasts of late frost, allowing them to better protect their harvests.

Increasing skilling to make the most of digital opportunities

Cutting-edge technology is meaningless without people who know how to use this technology. It's vitally important that Austria's workers are trained in the latest available technology and ready to make the most of the

opportunities available in our current digital age. That's why Microsoft Austria is investing in the largest skilling initiative since our inception in Austria. Since 2021, we have trained over 100,000 Austrians in digital technologies and supported thousands of certifications. Our investment in Austria's cloud region is poised to become one of the largest digital training initiatives in the country.

Sustainability as a cornerstone of the cloud region

A crucial element that will distinguish our new cloud region and underline its forward-looking role in the Austrian tech sector and economy at large is its clear commitment to sustainability. We rely [on 100 percent renewable energy in our data centers](#) and design them with a [Power Usage Effectiveness \(PUE\) score of 1.12](#)—well below the global average of all data centers.

A large part of the electricity purchased from providers flows directly into the operation of the IT infrastructure, making it possible to continuously optimize energy consumption for cooling, lighting, and general building operations. This results in just over 10 percent of total energy consumption onsite.

This sustainable orientation is part of Microsoft's long-term, global strategy. We are committed to covering the entire energy needs of our data centers using renewable energies by 2025 and [achieving water neutrality by 2030](#).

In October 2023, Federal Minister Leonore Gewessler paid a [visit to the Microsoft data center near Vienna](#), which is currently under construction. The Federal Minister for Climate Action, Environment, Energy, Mobility,

Innovation and Technology emphasized the importance of the coming years in the fight against the climate crisis. Digital solutions have the potential to contribute to the protection of our climate while also supporting the climate-friendly orientation of the economy. As Minister Gewessler noted, this will require the appropriate digital infrastructure, where climate protection is prioritized, and the operation is energy efficient. She appealed to business and political leaders to jointly mobilize all available resources to effectively meet the current challenges.

Microsoft's local cloud region in Austria will not only promote technological growth, but also provides a sustainable path into the digital future.





AI explained

AI is changing how we learn, work, live and play.

Generative AI is transforming how we interact with digital technology on every level.

What exactly is AI?



Harald Leitenmüller
Chief Technology Officer
Microsoft Austria

Microsoft's AI systems provide billions of intelligent experiences every day to people using our business software, gaming on our Xbox platform, and relying on our services to keep their organizations secure.

Our AI tools and technologies are designed to benefit everyone at every level in every organization. They are already being used in workplaces, home offices, academic institutions, research labs, hospitals, manufacturing facilities, and emergency preparedness organizations around the world.

These tools and technologies are helping everyone from scientists and salespeople, to farmers, software developers, front-line healthcare workers, and security practitioners.


What exactly is AI?

Using math and logic, an artificially intelligent computer system simulates the reasoning that humans use to learn from data and make decisions. Algorithms, or sets of instructions

that the system follows to process and analyze data, are at the core of every AI system. Different algorithms are used for different tasks, such as processing images or language.

An AI model is trained on a dataset representative of the task or problem it is being developed to solve. The AI model uses the patterns and relationships it discovers in the data to adjust and optimize its performance, a process called machine learning.

As an example, training a computer program to recognize cats would involve feeding a machine learning algorithm many pictures, each labelled as to whether the image contained a cat. By analysing these labelled examples, an algorithm can learn to identify patterns such as shapes, colours, and textures associated with cats, and then create a model based on this learning. Once trained, the model can be used to analyze new, unseen images to provide a prediction of whether they contain a cat or not. The more examples the algorithm sees, the better the model becomes at recognizing cats accurately.



A subset of machine learning under the broader umbrella of AI is reinforcement learning, which involves directing an AI system to make a sequence of decisions to maximize a reward. The AI system learns through trial and error, receiving feedback in the form of rewards—or penalties—based on its choices. By learning which actions lead to higher rewards, the AI system improves its decision-making over time.

AI development has also been inspired by the structure and functioning of the human brain. Neural networks are computational models that consist of interconnected nodes called artificial neurons, organized in layers. Each artificial neuron receives input, performs a computation, and passes the result to the next layer. Neural networks excel at pattern recognition and train their models to demonstrate complex relationships in data.

AI is a wide umbrella under which sits a range of different but complementary computer science fields and sub-fields—many of which have

been the subject of research for decades. No matter what the approach used to develop an AI system, data and compute-power are fundamental building blocks.

AI-powered possibility


AI enables machines to carry out a wide range of tasks, many of which were not possible before its advent. Natural Language Processing (NLP), for example, extracts meaning from text or speech data.

NLP enables computers to process languages and perform tasks like sentiment analysis, translation, and text generation. AI-powered Computer Vision enables machines to “see” and decode the visual world involving techniques such as object detection and image recognition. Computer Vision has a wide range of potential applications, from medical imaging to autonomous vehicles.

These capabilities and more are transforming how we use, benefit from, and interact with computer technology. At Microsoft, we believe that AI can serve as a “copilot” or virtual assistant to augment our human ingenuity and creativity by automating tasks and providing new insights. Over the last decade, the AI field has made considerable progress on image, sound, and language tasks. This progress has been accelerated by recent advances in generative AI—a class of AI models that can generate new content such as text, images, code, and more.

Generative AI is underpinned by a class of large-scale models known as foundation models. Foundation models are trained on massive amounts of data and can perform a wide range of tasks. Many recent generative AI models are not only good at generating text but also at generating, explaining, and debugging code. GitHub Copilot, for example, leverages OpenAI’s Codex model to assist developers in writing code.

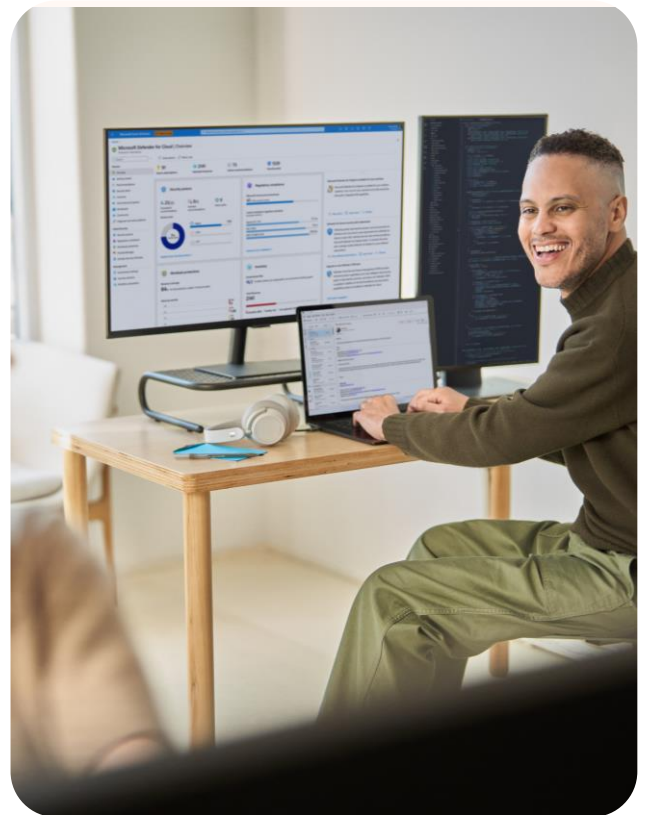
Another example of the cutting-edge capabilities of AI systems is the new AI-powered Bing search engine. The web search experience often involves the time-consuming task of reviewing and synthesizing information from a variety of sources identified from different search queries. The new Bing can do the heavy lifting, working behind the scenes to make the necessary queries, collect results, synthesize the information, and present a single complete answer.



Putting People at the Centre

From development to deployment to ongoing oversight, people play an essential role in every stage of the AI process. By prioritizing safety and trustworthiness throughout the process, AI developers and practitioners can put in place measures to mitigate potential harm, ensure fairness, and promote transparency. This hands-on human involvement is critical to shaping and guiding AI systems to achieve reliable and beneficial outcomes.

Recent months have put a spotlight on advancements in generative AI and how it might shape the future of industry. The pace and scale of technological innovation happening today is unprecedented. Generative AI has potential to provide a powerful boost to our efforts to address the major societal challenges we face in a range of areas.



Health

We are at a unique moment in history where medicine, biology, and technology are converging on a large scale. Health care organizations and life sciences firms are using AI to develop new treatments and medicines, as well as helping doctors and nurses improve patient care.

Sustainability


AI is helping accelerate and scale our efforts to transition to clean energy and mitigate the impact of global warming on people and communities. To curb the progression of climate change, humans urgently need to change the way we use the planet's resources and accelerate decarbonization.

Agriculture

Ensuring everyone has access to affordable food as the world's population grows and weather patterns become more extreme requires us to think differently about farming. AI is helping farmers improve yields even as they use fewer resources.

Manufacturing

More Europeans are employed in manufacturing than in any other sector. AI is helping to make the way we produce goods safer, more efficient, and more profitable.



Education

AI is transforming the way we learn and train, for the benefit of educators and students alike. By analyzing vast swathes of data to see which approaches work best, AI can help students with learning difficulties, support teachers in creating individualized learning, reduce school dropout rates, and enable differentiated education programs.



AI impact on Austria

Artificial Intelligence (AI) has the potential to significantly boost Austria's economic growth by enhancing labour productivity. AI could save a substantial amount of working hours, leading to a notable increase in economic output. However, challenges such as low digitalization levels, deficiencies of digital skills, insufficient data quality, and a heterogenous governance need to be addressed.

Recommendations include a governance model, improving digital skills and enhancing data access and quality for AI adoption. By tackling these issues, Austria can fully leverage AI to drive economic growth and productivity.

A potential equivalent to 2.24 billion working hours



Martin Hörmann
Director Government Affairs
Microsoft Austria

Since 2009, Austria's real economic growth has been trending significantly below the growth rate experienced from 1995 to 2008. AI is a broadly applicable technology that has the potential to affect not only manual but also intellectual labour across all sectors of the economy.

Demography and Productivity

Demographic development forecasts indicate that the number of people of working age will remain relatively constant, while the proportion of those over 65 will increase significantly. Additionally, there is a trend of reducing working hours per person. This will lead to a decrease in the total working hours (number of employed persons multiplied by the average working hours per person) in the long term. To maintain economic performance and the associated social benefits (and potentially provide for more people in the future), other components of macroeconomic growth factors must compensate. Artificial intelligence has the potential to do this as a technological advancement.

According to a [2024 study by Economia](#) considering the current digital intensity of companies, the full utilisation of the productivity increase through artificial intelligence could result **in a potential equivalent to 2.24 billion working hours.**

This corresponds to approximately one-third of the total labour or roughly the number of hours worked in Vienna and Styria together in 2023. If these hours are translated into value added, this results in **an increase of up to 18 percent compared to 2021 (+70.9 billion euros)**. This corresponds to a tenth federal state, which would have the second-highest economic output behind Vienna (which in 2021 was 91.9 billion euros). These efficiency gains can free up human resources for more strategic and creative tasks, enhancing overall workforce productivity.

Technological Advancement

AI can drive technological progress by enabling new capabilities and improving existing processes.

The continuous development and integration of AI technologies can lead to sustained improvements in labor productivity and overall economic performance. This technological advancement can also position Austria as a leader in AI innovation and application.

Sectoral Impact

AI has the potential to revolutionize various sectors, including healthcare, manufacturing, finance, and transportation. In healthcare, AI can improve diagnostic accuracy and patient outcomes. In manufacturing, AI can optimize production processes and reduce waste. In finance, AI can enhance risk management and customer service. In transportation, AI can improve traffic management and reduce emissions.

Recommendations

To fully harness the transformative potential of Artificial Intelligence (AI) and digitalization, it is essential for Austria to adopt a strategic and comprehensive approach.

1. State and Politics

- **Lead and partner in Digitalization and AI:** Austria is encouraged to take a pioneering role in specific areas of digitalization and AI application. By leveraging existing expertise, partnerships with industry, and setting positive examples, Austria can become a model for other nations.

- **Digitalisation and AI as a top priority:** Digitalisation and the use of AI must be declared a top priority. This requires clear and decisive leadership at the highest level to ensure that digital strategies and AI initiatives are implemented effectively. By prioritising these issues, companies and government institutions can mobilise the necessary resources and attention to make the digital transformation a success. Strong leadership can also promote collaboration between different sectors and ensure that Austria fully capitalises on the opportunities of the digital era.
- **Adapt Laws and Regulations:** It is recommended to adapt laws and regulations to meet the needs of AI development and deployment. This includes creating a legal framework that supports innovation while ensuring ethical standards and data protection.

2. Companies

- **Focus on Low Digital Intensity Sectors:** Companies are advised to concentrate efforts on sectors with low digital intensity. By doing so, they can unlock significant productivity gains and drive digital transformation across the economy.
- **Support SMEs:** Providing accessible solutions and support to small and medium-sized enterprises (SMEs) is crucial. This can include financial incentives, training programs, and access to digital tools that enable SMEs to adopt AI technologies effectively.


- **Encourage Experimentation:** Companies are encouraged to foster a culture of experimentation. This can be achieved through the use of sandboxes, where new technologies can be tested in a controlled environment, and by sharing best practice examples to inspire innovation.

3. Data and AI Applications

- **Improve Data Availability, Governance and Quality:** Efforts should be made to enhance the availability and quality of data. High-quality data is the foundation of effective AI applications and improving data infrastructure will support AI development.
- **Ensure Data Protection:** It is crucial to maintain robust data protection measures. This includes implementing policies and technologies that safeguard personal and sensitive information, building trust among users.
- **Break Down Data Silos:** Breaking down data silos is recommended to enable high-quality AI applications. Facilitating data sharing and integration across different sectors and organizations will enhance the effectiveness and reach of AI solutions.

4. Society

- **Increase Digital Skills:** It is essential to enhance digital skills among the population. This includes offering educational programs and resources that equip individuals with the knowledge and abilities needed to thrive in a digital world.



- **Ensure Data Literacy:** Promoting data literacy is vital. People should be educated on how to understand, interpret, and use data responsibly. This will empower them to make informed decisions and participate actively in the digital economy.
- **Promote Balanced Discourse:** Encouraging a balanced discourse on AI's opportunities and risks is important. This involves fostering open discussions that highlight both the benefits and potential challenges of AI, ensuring a well-informed public.



AI case studies

AI for innovation

Artificial Intelligence (AI) is redefining the landscape of innovation, empowering industries to push the boundaries of what's possible. By harnessing AI's advanced capabilities, organizations can accelerate research and development, create groundbreaking products, and revolutionize processes. AI-driven tools enable businesses to innovate faster, from predictive analytics that anticipate market trends to automation that frees up creative potential. As AI continues to evolve, it is unlocking unprecedented opportunities for innovation, driving forward a future where the limits of creativity and efficiency are continuously expanded.

CASE STUDY:

Strabag SE

Strabag SE builds a risk management solution to improve efficiency using Microsoft's Intelligent Data Platform.

As one of the oldest industries in the world, construction hasn't always kept pace with the modernization that has swept across other industries, particularly when it comes to driving operational efficiencies with data. Global construction company Strabag SE is determined to change that. Based in Austria, Strabag SE specializes in construction and related services in locations spanning the globe. The company has approximately 75,000 employees and its projects range from civil engineering to transportation infrastructure.

Strabag SE partnered with Microsoft to build a Data Science Hub to collect decentralized data and leverage it for insights, thereby enabling the organization to develop use cases to prove the value of data including its risk management project. The solution uses an algorithm to pinpoint at-risk construction projects, saving Strabag SE time and reducing financial losses. With data now at the helm, Strabag SE is

cementing its future in a traditional industry on the path to digitalization. "The construction industry is very traditional, and innovation can be slowed by its inefficiencies," said **Dr. Marco Xavier Bornschlegl, Head of Innovation and Digitalisation at Strabag SE**. "But for us, innovation isn't just about building incredible solutions and products; it's about making small changes that can have a huge impact. One area of the business where we saw we could make a difference was with data sharing within the company. We have a lot of data, but it was stored in single units and business fields, with no connection to another."

Building a Data Science Hub unlocked a myriad of possibilities, allowing the team to begin developing use cases to prove the value of data and explore its potential. One successful initiative was the AI-based risk management solution that assesses which construction projects may fail. "By using data and AI to pinpoint which projects are a potential risk for us," said Bornschlegl, "we can save the organization huge amounts of money and time. It's a perfect example of the power of data."



Laying the foundation for innovation

“Construction is one of the oldest industries,” explained Bornschlegl. “Even to this day, it’s a sector built by hand. While it’s important to honor this legacy, it’s vital that we bring innovation straight to the construction sites to improve efficiency. Within this framework, a key challenge we faced at Strabag SE was that each of our global sites were decentralized and employees needed multiple logins to access anything,” Bornschlegl continued. “This is where our digitalization journey began, and it marked the beginning of our partnership with Microsoft.”

With strong support from the internal IT infrastructure department, Strabag SE’s data science team started small, using Microsoft Azure Active Directory to set up logins and enable multifactor authentication. They adopted cloud technology and implemented tools and applications such as Microsoft 365, Microsoft Teams, and Azure Active Directory to support

employees with their day-to-day workloads. Now they had a foundation of technology to build upon. “Without the street, you can’t drive the car,” said Bornschlegl. “So, in the last three years, we’ve been constructing the street.”

Establishing a centralized data hub

With a bedrock of technology in place, it was time to start exploring ways to digitalize and leverage data. Initially, this proved challenging, as from both a cultural and technical perspective, the sector had fallen behind other more progressive industries. “One of the great things about the construction industry,” explains Bornschlegl, “is that it’s all about collaboration, where multiple companies with different skillsets come together to create something incredible. But this way of working wasn’t reflected in how companies used and managed data – no one was willing to share what they generated. We needed to revolutionize this culture and put data at the heart of it.”

Over the last 18 months, Strabag SE collaborated with the Microsoft Industry Solutions Delivery team to embrace Microsoft's Intelligent Data Platform. Implemented as part of a strategy to phase out costly and complex solutions, the technology also allowed Strabag SE to establish the Data Science Hub. This decentralized system uses Microsoft Purview for data governance, as well as various Microsoft products including Azure Synapse Analytics, Azure Databricks, Azure Machine learning, and Azure SQL.

This enabled the team to start building use cases which could then be scaled globally. This strategy is helping to provide Strabag SE with a more structured approach, while also generating key insights, business value, and shareable data.

"Throughout the whole process, Microsoft has been more than a technology partner," said Bornschlegl. "They advised us every step of the way. Together, we ran workshops to explore the organization's needs. We used their Intelligence Driven Organization (IDO) framework to create a strategy that would address our digital priorities. We chose two initiatives to work on that would have the biggest impact: improving the existing Data Science Hub infrastructure, and risk management."

Strabag leverages Microsoft's Intelligence Data Platform, adopting some data mesh principles and moving from soiled data to a federated AI data hub where all business units can be contributors as well as consumers of the data.

The solution follows the cloud-scale analytics scenario guidelines and best practices, with a data management landing zone that governs the platform and secures the data. The data management landing zone connects the unique data landing zones, which may have different ingestion capabilities depending on their sources, to serve the creation of data products.

Mitigating risk using the power of AI

The margins in construction are narrow, so it's important for companies to pick the right projects to work on. "If you choose the wrong one," explained Bornschlegl, "it can be financially detrimental to the company. This is why risk management is such an important part of what we do."

Strabag SE used Azure OpenAI Service to create a solution that allows it to assess the potential impacts of a project by comparing it against all its realized projects – therefore minimizing risk. Specialists then use a Power App to carry out further comparisons, review the projects, and reprioritize the work.

“With just three months of data, this algorithm can predict risk with 80% accuracy—it’s mind-blowing!” said Bornschlegl. “With this solution, enabled by Azure OpenAI Service, we know when the weather will prevent us using a crane on-site. Moreover, by the aid of generative design, our real estate development unit knows with a few clicks how much maximum gross floor area they can develop on a site. This kind of information is vital when you have projects worth two million euros with a margin of 4-5 percent.” “The algorithm can work more quickly and efficiently than humans, but the process still requires people to make the ultimate decisions. Even so, technology this powerful scares people, and right now our biggest challenge is addressing this anxiety among our employees.”

To tackle this and empower employees to embrace everything that innovation has to offer, Strabag SE is running initiatives to help transform how employees view and use technology. Given that the construction industry is in dire need of more people, the aim of digitalization is not to replace employees, but to relieve them of parts of their workload.

Constructing a future built on data

“In time,” says Bornschlegl, “STRABAG Innovation and Digitalisation hopes that employees see that in the same way you need different materials to work together to build a sturdy house, when you collect and connect information, you can see the bigger picture and build a stronger organization—and that’s data science.” With the help of Microsoft, the team are continuing to develop use cases to prove the value of being data-driven and have around 15 projects in the pipeline.



“We want to use data to improve efficiency, optimize operations and become more sustainable. That’s what sets apart the organizations that are successful and those that are not.”

For the team at Strabag, leveraging data isn’t just about supporting the construction industry of today. “Data science helps us prepare for the future,” says Bornschlegl. “Our vision is to build a better tomorrow, not just for Strabag, but for the construction industry. To achieve this, we must embrace innovation and unleash the power that data has to offer.”

CASE STUDY:

AVL

AVL: Secure, AI-based cloud solutions for efficient and sustainable development.

AVL and Microsoft are setting the stage for digital transformation in vehicle development. We are partnering with the world's leading mobility solutions provider to jointly support original equipment manufacturers (OEMs) in getting to market faster. Thanks to the power of data stored safely in the cloud and evaluated by AVL's intelligent algorithms, development cycles can be accelerated and made more sustainable.

AVL currently employs around 3,000 software experts to make the car of the future, a process that is more environmentally friendly because it's conducted virtually. AVL's software experts can work anytime, from anywhere. The data collected from simulation and testing flows together in the cloud and can be evaluated and used with maximum efficiency. This allows AVL to optimize its processes from the earliest stages of development and reduce its CO2 footprint. A high level of transparency and the evaluation of real-time data allows for a shorter time to market.

Cloud-based methodology: High flexibility, worldwide usability, short development times

The cloud-based approach enables broad integration of AVL's software tools which are available to development teams around the world thanks to the software-as-a-service model. This enables simultaneous access for everyone, which promotes the development of collaborative and innovative solutions. Even with limited storage space or computing power in engineering, simulation or testing, AVL doesn't need to invest in new hardware capacities. The integration of High-Performance Computing (HPC) complements this approach, making it possible to handle computationally intensive tasks without the limitations of physical hardware capacity. This not only contributes to the scalability and efficiency of software testing, but also enables parallel execution.

Jens Poggenburg, Executive Vice President

AVL: "Through the strategic partnership with Microsoft, we are creating a boost in digital transformation. The bundling of expertise allows us to further develop the offering for our customers in engineering and toolchains at the highest speed. We are proud to combine our specific application experience with Microsoft's digitalization expertise. Also, because we can derive new important business models from this."

Sebastian Jagsch, Global Head of Partner

Management AVL: "AVL combines its extensive expertise in automotive development with Microsoft's innovative cloud and AI technology. This results in revolutionary solutions that take the development and validation of vehicle systems to a new level. The groundbreaking cooperation marks a significant step towards

the future of mobility and opens up new opportunities for our customers worldwide."

Sustainability as a key goal: Together for greener mobility


The strategic partnership between AVL and Microsoft is a milestone on the way to greener and more sustainable mobility. It combines AVL's extensive expertise in automotive development with Microsoft's innovative cloud and AI technology to create sustainable solutions. But the efforts of AVL and Microsoft go beyond creating technological solutions. The partnership aims to revolutionize the automotive industry while addressing environmental challenges. This ambition is nourished by a common goal: to pave the way to more sustainable and greener mobility.



Joacim Damgard, President of Western Europe at Microsoft, sums up the essence of this partnership vision: "Our collaboration with AVL is not only a synergy of technology, but also a desire for greater sustainability. Together, we are using our resources and expertise to fundamentally change the way vehicles are developed and used, and to significantly improve their sustainability."

The seamless integration of artificial intelligence not only results in accelerated development processes and increased resource efficiency, but also makes a significant contribution to reducing AVL's CO2 footprint. This innovative approach is particularly evident in the field of Advanced Driver Assistance & Autonomous Driving Systems (ADAS/AD). Under the umbrella of Advanced Driver Assistance, a wide range of technologies support drivers in vehicle control, such as lane departure warning, automatic braking, and adaptive cruise control. Autonomous Driving Systems go one step further in pursuit of the goal of controlling the vehicle without human intervention. These technologies have the potential not only to increase passenger safety and comfort, but also to optimize traffic flow and reduce fuel consumption.

Generative AI is now coming to the fore of ADAS/AD technology by creating highly realistic simulations and scenarios for autonomous driving. Generative AI cleverly merges real-world data and conditions to create environments in which advanced driving technologies can be tested and refined.



This innovative approach revolutionizes the development, validation, and optimization of autonomous driving systems and paves the way for future-oriented mobility that is not only safer and more efficient, but also more sustainable.

"The use of artificial intelligence in vehicle development opens up a new chapter in mobility for us, which is characterized by efficiency and sustainability. We are proud to be at the forefront of this transformation together with AVL and to create groundbreaking solutions that drive positive change in the automotive industry," emphasizes **Hermann Erlach, General Manager at Microsoft Austria**.

CASE STUDY:

ANDRITZ

International technology Group
ANDRITZ is working with Microsoft
on the future of the process industry
by enabling fully autonomous
factories and creating a trusted data
ecosystem based on the ANDRITZ
Metris digital platform and Microsoft
Cloud for Manufacturing.

The collaboration will also promote modern working methods within ANDRITZ, thereby accelerating internal processes for rapid and seamless customer support. This initiative represents an important step in the digital transformation of the manufacturing industry.

Joachim Schönbeck, CEO of ANDRITZ, commented: "This collaboration is a milestone in our digital transformation. By combining our industry knowledge with Microsoft's digital solutions, we are enhancing our Metris digital platform and providing our employees with a streamlined work environment that enables breakthrough business models and improved customer service."

The ANDRITZ Metris digital platform forms the basis for the path to autonomous plant operation. The platform offers comprehensive support for industrial assets throughout their lifecycle, addressing customers' key challenges such as operating costs, process stability, asset reliability, and sustainability. A Copilot solution based on the Microsoft Azure OpenAI Service facilitates the detection of anomalies and optimizes communication between employees and AI. A chat interface provides real-time insights into the machines, aiding employee decision-making and thereby helping to increase plant efficiency, product quality, and resource utilization.

Digitization within the company has also received a significant boost. The use of Copilot for Microsoft 365 underlines ANDRITZ's commitment to modern working methods and product improvements. By automating repetitive tasks and providing actionable insights, Copilot helps make internal processes more efficient, allowing employees to focus more on strategic tasks and customer needs.



Rohit Pant, Chief Automation Officer at ANDRITZ, emphasizes the benefits for the customer: "The integration of Microsoft's digital know-how into our Metris platform enables our customers to make data-driven decisions and unlocks a number of benefits. Through qualification programs and specialized training, ANDRITZ's teams are constantly evolving to meet the demands of a digitally driven market."

"ANDRITZ is helping to shape the future by using Microsoft Azure to further develop its solutions for autonomous factories," says **Ralph Haupter, President, Microsoft EMEA**. "The deep integration of their products with Azure cloud services is the kind of technological innovation that drives sustainable and efficient change in the industry."

CASE STUDY:

BTV VIER LÄNDER BANK

BTV VIER LÄNDER BANK (BTV) recognized the potential of generative artificial intelligence early on, and recently put its own text generator into operation.

Founded in 1904 and headquartered in Innsbruck, BTV's market area stretches from Tyrol, Vorarlberg, and Vienna, through southern Germany and German-speaking Switzerland to South Tyrol. BTV VIER LÄNDER BANK combines personalized customer service with a strong will to innovate. In its digitalisation efforts, the bank is pursuing a sustainable differentiation strategy: "Our aim is to be more digital than our classic competitors, but at the same time to act more personally than pure neobanks, thus combining our strengths in personal advice with the opportunities of the digital world," explained **Markus Perschl, Member of the Board of Management for Digitalisation at BTV VIER LÄNDER BANK.** "Our progress in this area has also been confirmed by the award as "Austria's Most Innovative Bank 2023" by the financial magazine Börsianer."

Rather than seeking to replace the personal relationship at the core of their business philosophy, BTV is seeking to supplement and further expand its mission through AI technology.

From marketing to ESG analysis: The AI assistant in Microsoft Teams for all business units

Two years ago, the bank initiated an important shift: implementing Microsoft Teams as a central platform for all employees.

In 2023, BTV integrated ChatGPT from the Microsoft Azure Cloud into Teams as a chatbot and made accessible the AI assistant accessible to all employees. Trained on the bank's company data, the AI assistant is designed to facilitate and speed up the creative process of writing.




In marketing, the AI assistant supports the creation of advertising text and customer mail, among other things. In customer service, the assistant serves as an aid for formulating correspondence and translations. BTV is soon planning to use AI in balance sheet analysis, where it will be tested in a pilot project next year to evaluate complex financial data more efficiently. AI also provides valuable services in extracting and interpreting ESG figures in sustainability reports and quickly finding information in internal resources.

Overall, the technology allows BTV employees to move away from routine tasks and instead focus on more creative and value-adding activities, which both increases internal productivity and improves customer experience.

Artificial intelligence as a driver of innovation and development at BTV

The AI assistant also offers BTV employees the opportunity to gain practical experience with AI language models in their everyday work. By encouraging its teams to actively seek out ways to use artificial intelligence and implement it in their workspaces, BTV hopes to improve the employee experience and better compete for talent.

"BTV VIER LÄNDER BANK has deliberately prioritized the implementation of the AI assistant so that our employees can gain their own practical experience with generative AI models as quickly as possible," said **Herwig Petz, Head of Digital Unit at BTV VIER LÄNDER BANK**. "Not only have we implemented a new technology, but we've also established a new mindset in our teams that we use to actively shape innovation."



Data protection, "Do's & Dont's" for safe use in banking

BTV paid close attention to security measures when introducing ChatGPT in Teams. The closed system implemented by the bank guarantees that all chat traffic takes place only within the bank and is not accessible to outsiders—both external persons and other employees.

In addition, communication follows strict guidelines as well as automatic deletion routines to meet the high privacy and security standards of the banking sector. For this demanding task, the bank relies on an interdisciplinary team of data protection and information security specialists as well as banking experts.

BTV on the way to new horizons with artificial intelligence

BTV sees the integration of ChatGPT into Teams as an essential building block of its digital evolution. The positive response the move has received so far motivates individual employees in various departments to further develop and experiment with AI solutions.

Guided by an ambitious AI strategy and with an adequate investment budget, BTV's goal is to further integrate AI technologies into its business processes, continuing to enrich the customer experience without neglecting personal contact. In this ongoing digital transformation, BTV VIER LÄNDER BANK is facing the challenges of the future while preserving its traditional values.

AI for government

Artificial Intelligence (AI) is transforming the government and public sector by enhancing the efficiency, transparency, and responsiveness of public services. Governments around the world are leveraging AI to improve decision-making, streamline administrative processes, and deliver more personalized services to citizens. AI-driven tools enable better resource allocation, predictive analytics for public safety, and more effective management of infrastructure projects. Additionally, AI is helping to detect and prevent fraud, optimize emergency response systems, and improve regulatory compliance. By integrating AI, the public sector is not only improving service delivery but also fostering greater trust and accountability in governance, ultimately leading to more effective and inclusive public administration.

CASE STUDY:

AMS Berufsinfomat

The Austrian public administration has been actively integrating AI and Large Language Models (LLMs) across various sectors. A notable example is the AMS Berufsinfomat, which provides personalized career advice to job seekers.

Other implementations include digital assistants in ministries to handle citizen inquiries more efficiently, predictive analytics to forecast and optimize resource allocation, and automated document processing to analyze large volumes of documents swiftly.

The primary drivers for AI adoption in the public sector are efficiency improvements, enhanced citizen services, and fostering innovation, all of which are compounded by the upcoming retirement wave. Around 30 percent of public servants in Austria are expected to retire by 2034. This significant retirement wave presents an opportunity for AI to fill the gap by

automating routine tasks, improving efficiency, and ensuring continuity in public services.

AI use cases such as digital assistants, predictive analytics, and automated document processing will become increasingly vital.

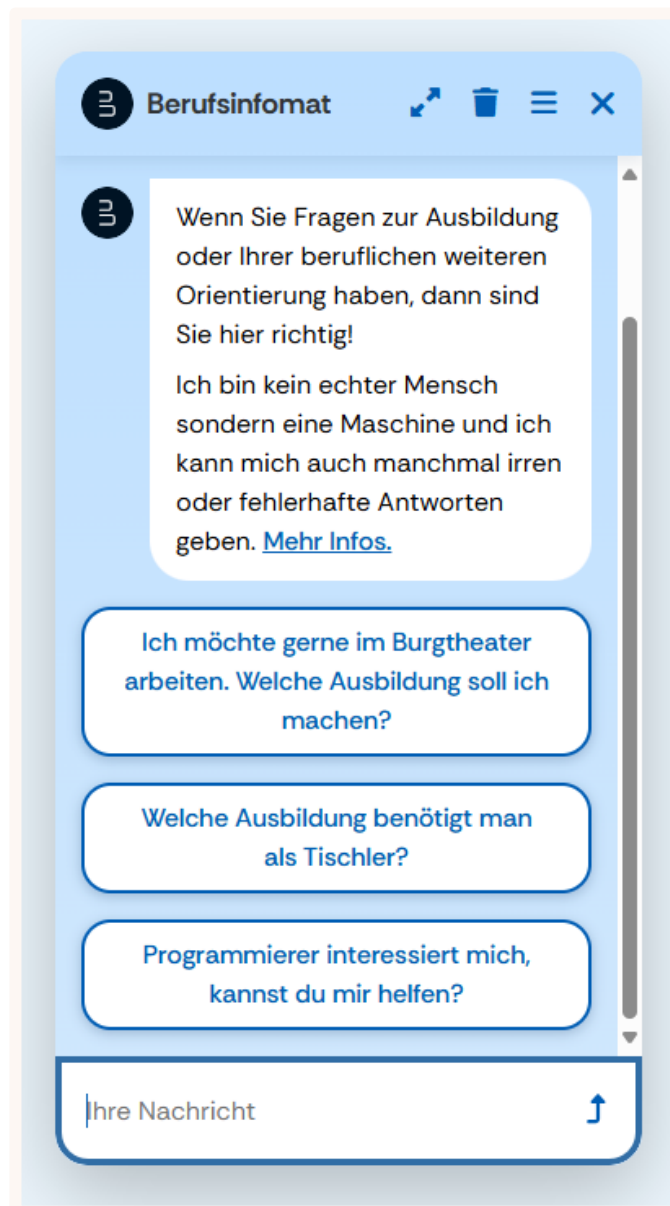
The Austrian AI strategy, "[AIM AT 2030](#)," underscores the country's commitment to harnessing AI for economic and societal transformation. The strategy focuses on using AI for the common good, establishing Austria as an internationally recognized center for AI research and innovation, and enhancing the country's competitiveness as a technology and economic hub.

Planned developments include expanding technological infrastructure, promoting educational programs to prepare the workforce for AI, developing ethical guidelines and legal frameworks for responsible AI use, and supporting research and innovation initiatives.

The opportunities presented by AI include improved decision-making through the analysis of large datasets, cost reduction through automation, and increased transparency and accountability in processes. However, there are also risks, such as data privacy concerns, potential biases in AI systems, and over-reliance on technology, which could lead to issues if systems fail or malfunction.

To maximize the benefits of AI while mitigating risks, it is recommended that all Austrian public administration employees receive training in AI and LLMs. Additionally, developing and implementing ethical guidelines and legal frameworks is crucial for ensuring the safe and responsible use of AI. Pilot projects should be conducted to test the benefits of AI in various areas, with successful approaches being scaled up.

By taking these steps, the Austrian public administration can effectively leverage AI's potential while addressing its challenges, ensuring that the technology serves the public good and supports the country's broader strategic goals.



The application interacts with the user in the form of a chatbot, enabling natural language input without prior knowledge.

AI for healthcare

Artificial Intelligence (AI) is revolutionizing the healthcare sector, and Microsoft is at the forefront of this transformation. By leveraging AI, Microsoft aims to enhance patient care, streamline clinical workflows, and accelerate medical research. Solutions like Microsoft Cloud for Healthcare integrate AI to provide predictive analytics, improve diagnostic accuracy, and personalize treatment plans. Additionally, Microsoft's AI-driven tools help in managing large datasets, enabling healthcare providers to make data-driven decisions swiftly. With a commitment to privacy and security, Microsoft ensures that AI applications in healthcare are both innovative and compliant with regulatory standards, ultimately aiming to improve health outcomes and patient experiences globally.

CASE STUDY:

Medical University of Vienna

A talented team of researchers at the Medical University of Vienna — or “MedUni Vienna” — recognized that tumor analysis was lagging other specialized medical fields when it came to AI.

While fields such as radiological research were benefitting from machine learning solutions for CT and MRI scanning technologies, tumor analysis was reliant on outdated processes, creating a painful and drawn-out patient experience. Since working alongside Microsoft, the university has found a partner and platform to pursue its vision of accurately categorizing tumors, potentially without biopsies.

MedUni Vienna is a world-class institution known for its excellence in research, education, and patient care. Researchers regularly publish industry-shaping work in top journals across the globe, cementing the institution as a true thought leader in medicine.

One such research team—a collaboration between the Center for Medical Physics and Biomedical Engineering, consisting of medical

physicist **Professor Thomas Beyer and Ph.D. candidate Laszlo Papp, and the Division of Nuclear Medicine, clinical director Professor Marcus Hacker**— has been working on pioneering new approaches to tumor characterization.

“Cancer is the second leading cause of death worldwide,” explained Papp. “So, what we want to do in principle is to generate predictive models, based on hybrid, anato-metabolic imaging, without needing to take biopsy samples from patients. This is because, on the one hand, a biopsy is not a pleasant experience, and on the other hand, it’s just not precise. Hybrid imaging, on the other hand, particularly PET/CT (Positron Emission Tomography/Computed Tomography), gives an overview of the whole tumor in 3D. We can learn crucial information about the biology of tumors.”

Beyer, Hacker, and Papp chose to work with Microsoft because of the company’s enthusiasm for the project’s potential. “When we approached Microsoft with our ideas and needs, they listened to us,” Papp said. “They believed in our project. They understood why it’s so important.”



Harnessing AI to analyse tumors better than ever before

The routine diagnosis of tumors is typically performed through invasive biopsy sampling, where medical imaging is incorporated to visually interpret and detect lesions and to guide biopsies. But since no form of cancer is the same as any other, this process is complex, labor-intensive, and not always completely accurate.

“Currently, cancer treatment is successful in half of the patients worldwide,” Papp explained. “The treatment is generally expensive, costing up to one hundred thousand Euro per year and patient.”

Through the Microsoft Azure infrastructure as a service (IaaS) platform and Microsoft Azure Cognitive Services, the team is now able to put the power of AI behind its hybrid imaging methods to understand tumors better than ever before.

“We have many ongoing projects that are built on the results of our Microsoft-powered research, focusing on tumors such as cervical, prostate, or breast cancer. The knowledge we’re gaining and the outcomes we’re generating are very exciting,” added Hacker.

Research is becoming faster, more accurate, less time-consuming and, most importantly, is being applied directly to improving patient care.



Using predictive modelling to save precious patient time

In current standard practice, biopsies can take up to two weeks to produce a result, in which time fast-growing tumors can get progressively worse. Waiting for news of biopsy results is a drawn-out and stressful experience for patients. Beyer, Hacker, and Papp believe that predictive models can help physicians outrun such aggressive types of cancer. Speeding up the process is not all that predictive modelling can potentially do for patients: the time saved can literally be lifesaving.

As Papp explained, "Imagine being confronted with your doctor telling you about a lesion that could be cancerous. You go to the hospital and get a hybrid imaging exam. Following that, instead of going through a painful, inaccurate

biopsy analysis, followed by a long decision-making process about your therapy, by the time the scan is finished, and you get off the imaging bed, artificial intelligence has already performed an in-depth analysis of your medical images. This is relayed almost instantly to your attending doctor, who can then put together a personalized therapy plan."

The accuracy, ease of use, and reduced stress associated with such an AI approach make for a far better patient experience. Beyer added, "Cancer is a scary prospect, so anything that can be done to make the diagnostic process a better experience for people affected has a huge impact on their mind and on healthcare in general." All three experts' reason that their approach to a clinical decision support system will help create better patient experiences and lead to dramatic cost reductions over the long term.

Providing computing power for seamless data analysis


In order to teach the models to predict patterns in hybrid images effectively, the research team recognized that gathering the “extreme amount of data” it needed to work with was going to be a challenge.

“For such high-level computation, typically we need a lot of CPUs, GPUs, and maybe hundreds of thousands of GPU cores, and this is just not available in a university environment,” Papp noted. The Azure platform is a versatile landscape that’s able to hold huge amounts of data impossible to manage with standard, in-house IT systems. This development of processing and infrastructure ability could unlock broader transformation across this project.

Transforming medicine to transform lives

The collaboration with Microsoft, Beyer, Hacker, and Papp can continue to unravel the complexities of tumor characterization, shortening the path from diagnosis to treatment for patients across the globe.

Their vision of a world where AI and clinicians work together to transform patient care is one that could change the face of medicine. As they put it, “We believe that our combination of AI and hybrid imaging for cancer care will help enhance the chances of survival and ultimately, the quality of life for our patients.



When we approached Microsoft with our ideas and needs, they listened to us. They believed in our project. They understood why it's so important.” – **Laszlo Papp, PhD Candidate, Medical Physics & Biomedical Engineering and Division of Nuclear Research team**

“Cancer is a scary prospect, so anything that can be done to make the diagnostic process a better experience for people affected has a huge impact on their minds and on healthcare in general.” – **Professor Thomas Beyer, Medical Physicist, Medical Physics & Biomedical Engineering and Division of Nuclear Medicine Research team.**

AI for energy

Artificial Intelligence (AI) is revolutionizing the energy sector, driving advancements in efficiency, sustainability, and innovation. By integrating AI, energy companies can optimize grid management, forecast energy demand with greater accuracy, and enhance the performance of renewable energy sources. AI-driven analytics enable predictive maintenance for energy infrastructure, reducing downtime and operational costs. Moreover, AI is instrumental in advancing smart grid technology, enabling more efficient energy distribution and consumption. As the energy sector shifts towards a more sustainable future, AI plays a critical role in accelerating the transition to renewable energy, optimizing resource use, and reducing environmental impact, ultimately leading to a cleaner and more resilient energy landscape.

CASE STUDY:

VERBUND

Austria's leading energy company, VERBUND, and Microsoft are entering into a technological partnership that will create a digital platform offering solutions for the energy sector. In addition, Microsoft will also focus on 100% sustainable energy and CO2 reduction in its new data center for the recently announced cloud region in Austria.

VERBUND will rely on Microsoft cloud technologies such as Azure, which offers scalability and ubiquity for data storage and processing. Over the next three years, VERBUND expects that a significant proportion of its processing capacities will be transferred to intelligent cloud platforms. VERBUND will also introduce Office 365 and Dynamics 365 for its employees as a productivity-enhancing solution for communication and collaboration. With this partnership, VERBUND can accelerate the acquisition of skills necessary for the digital age among its employees.

"Digital solutions are the catalyst for a successful energy future in Austria," said **Achim Kaspar, member of the Executive Board of VERBUND AG** responsible for generation.

"Digitalization is part of our VERBUND strategy, our digital roadmap includes use cases for the entire Group—from generation to trading to sales. Through the strategic partnership with Microsoft, we benefit from the development of innovative, digital technologies. We are looking forward to this groundbreaking Austrian cooperation between our companies. Together, we will implement projects and plans that focus on efficient generation and digital sales and trading at VERBUND."

As part of this partnership, both companies will aim to jointly identify innovative and sustainable solutions for the energy industry based on cloud computing technologies, artificial intelligence, the Internet of Things (IoT), big data and mixed reality, among others. Projects range from the implementation of digital twin projects for power plants to sustainability projects such as AI applications in fish bypasses for VERBUND's power plants.



"By introducing a completely new way of bringing people, data and processes together, VERBUND is a strong example of what it means to take on a leadership role in the digital age. We are looking forward to the joint partnership, in which we can support VERBUND on the path to digital transformation, which will help to shape and accelerate the next wave of innovation in the energy sector," said **Hermann Erlach, General Manager, Microsoft Austria.**

Climate-neutral data centre

Microsoft plans to meet the global energy requirements of its data centres entirely from sustainable renewable sources by 2030. VERBUND is supplying Microsoft's data centre in Austria with all renewable energy, a major first step towards the goal of a CO₂-free, climate-neutral power supply for data centres.

"With our efforts to ensure completely CO₂-free operation of all our data centres, we want to show that a CO₂-free economy is possible," notes Erlach.

VERBUND power plants in Mayrhofen/Tuxbach and Freudenau are already making the climate-neutral data center a reality. These data centers are CO₂-neutral and supplied by 100 percent Austrian hydropower. With VERBUND as a partner, customer data is secured, and e-mail and cloud platform services are reliably available to users.

Traceability is ensured by providing verified certificates of origin (in accordance with Article of EU Regulation 2018/2001) on an hourly basis. In this way, the VERBUND/Microsoft partnership is making an active contribution to the energy transition and the expansion of sustainable energy generation in Austria.

CASE STUDY:

Kelag Energie & Wärme GmbH

Winter tourist destinations have an enormous energy requirement that fluctuates wildly based on occupancy rate, weather, and other special events. For Kelag Energie & Wärme GmbH, an energy supply company operating throughout Austria, the question of how much district heating needs to be fed into the networks has historically posed a challenge.


Because the district heating network is temperature-controlled, sometimes more heat ends up in the networks than is needed. Thanks to an innovative new data-sharing program, Kelag Energie & Wärme can calculate heat demand more precisely in the future, operating the network more efficiently and saving delivery routes.

The well-known winter sports resort Obertauern is focusing on a data-driven revolution: Thanks to the innovative data-sharing program of the Viennese start-up nexyo, the resort's heat requirement is now calculated much more precisely not only based on temperature

conditions, but also on numerous other factors. This is ensured by an intelligent linking of different data sources.

"It makes a difference what day it is, what the booking situation is in the town and whether, for example, there is a big concert there. From now on, we will bring all this data together and thus ensure that we can forecast the amount of energy to be generated much more precisely and in line with demand. In this way, sustainable digitization brings an increase in the efficient use of resources and a much better fuel forecast, which can be of particular relevance to the exposed location of Obertauern," said **Michael Wagner, Managing Director of Kelag Energie & Wärme.**

These exact forecasts are made possible by a new form of data sharing developed by nexyo. As **Lisa Höllbacher, Managing Director of nexyo, explained:** "nexyo is a decentralised data infrastructure technology for secure data exchange for companies and research institutions, where you can determine in detail who can use which data and under what conditions. You never lose control of your own data—and you can still use its full potential."



Optimising district heating control in Austria's highest biomass heating plant

In Obertauern, this means that various project participants such as ZAMG, the Austrian National Tourist Office, or mobile phone provider Drei provides specifically requested, very heterogeneous data that only unfolds its full information power when combined. Microsoft provides the storage of the data and looks after the IT systems.

"Where in the past only large providers managed data and infrastructure, partner ecosystems can now work together beyond their own company boundaries and create new advantages. Keyword platform economy—we firmly believe in this model for the future," said **Stefan Nussbaumüller, Enterprise Account Executive at Microsoft**, about the strategic relevance of data ecosystems for the company. "However, this is not only a question of technology, but also of mentality and the will to innovate. With the cloud as an intelligent infrastructure, we want to make new business models possible and unleash unprecedented potential."

AI for education

Education is undergoing a profound transformation as Artificial Intelligence (AI) introduces more personalized, accessible, and effective learning experiences. Through AI, educators can customize instruction to meet the unique needs of each student, creating adaptive learning paths that foster deeper understanding and engagement. Advanced AI tools provide instant feedback, streamline administrative tasks, and offer valuable insights into student performance, allowing teachers to dedicate more time to teaching and mentoring. Moreover, AI expands access to quality education through intelligent tutoring systems and virtual classrooms, breaking down traditional barriers and ensuring that learning is inclusive and equitable. This integration of AI is reshaping education, equipping students with the skills and knowledge needed for a rapidly changing world.

CASE STUDY:

Austrian Federal Ministry for Education

The Austrian Federal Ministry of Education, Science and Research (BMBWF) views generative AI as a significant opportunity in the education sector. The ministry has published guidance for teachers and students on how best to use AI in the classroom.

The ministry sees AI as a tool that can increase productivity, improve quality of life, and help tackle global challenges while also acknowledging the challenges associated with AI, such as data protection and security.

The BMBWF has established an expert group to propose and oversee measures regarding the use of AI in school. Initiatives include AI pilot schools to support learning and teaching processes with AI tools. The ministry sees the professional use of generative AI tools as key to the future of learning in Austria. It also highlights the need to make AI socially equitable and ensure that AI tools are available and accessible to all.

Austrian higher education institutions also recognize the importance of generative AI and are actively engaged in research and teaching in this area.

Additionally, there are initiatives exploring the collaboration between human intelligence and generative AI, promoting the integration of AI into daily life and higher education. Austria also supports international efforts to bridge the digital divide and develop skills and technologies in the AI field to promote fairer and more environmentally friendly global progress. The Austrian government recognizes generative AI's disruptive nature and the need to embrace digital transformation.

By using Microsoft Azure and Microsoft AI products and services, the Austrian government can support the modernization of legacy systems and applications and enhance the user experience.

AI for sustainability

The landscape of environmental stewardship is evolving as Artificial Intelligence (AI) introduces groundbreaking approaches to addressing global sustainability challenges. AI technologies are enabling more efficient resource management, reducing waste, and cutting carbon emissions through advanced data analysis and predictive modeling. From optimizing energy use in smart grids to enhancing agricultural practices with precision farming, AI is driving significant improvements in environmental impact. It also supports more effective monitoring and response strategies for climate change and natural resource management. By integrating AI into sustainability efforts, industries and governments are better equipped to foster long-term environmental health and resilience, ensuring a balanced approach to economic and ecological well-being.

CASE STUDY:

SPAR

Global retail company SPAR sees AI not as merely a tool, but as a catalyst – a force capable of propelling the company towards inclusive growth, sustainable development, and unparalleled creativity. Thanks to its new intelligent supply chain, SPAR is promoting sustainability and minimizing food waste using the Microsoft Azure Cloud.

Knowing today what customers will buy tomorrow is the dream scenario of every retail company. This allows goods to be ordered sensibly and the supply chain to be designed efficiently. As a market leader in food retailing, SPAR is breaking new ground to best meet the needs of its customers. The benefits are manifold—not only for the company, suppliers, customers and employees—but also, and above all, for the environment. It's a balancing act for both sides: Customers expect everything they want to be available in the store. On the other hand, the food trade tries not to have too many goods in stock to prevent spoilage.

Together with SPAR ICS, the IT company of the SPAR Austria Group, the Austrian family-owned company headquartered in Salzburg relies on Azure, Microsoft's intelligent cloud technology.

Demand forecasting: Predicting demand and demand for goods

The vision of an intelligent supply chain takes process optimization to a new level. Thanks to precise assessments with the help of data and artificial intelligence, this is no longer pure theory. Based on Microsoft Azure, SPAR's IT unit developed a customer-oriented solution with its Microsoft partner paiqo to enable accurate demand forecasts for all supermarket branches. The solution uses AI to analyse and process data on sales volumes, weather conditions, special offers, marketing promotions, seasonality, and other factors, making it possible to predict the optimal amount of fruit and vegetables per store with over 90 percent accuracy. The solution is a win-win: Customers can find what they are looking for in any given store and SPAR's predictability and sales improve. Fewer perishable goods are disposed of and food waste is kept to a minimum.



"The major challenges of our time, such as climate change, cannot be addressed today without technology overcoming them is only possible with the right digital solutions. With this project, we are using the potential of technology and AI to meet the needs of our customers and save resources - this also benefits our employees and the environment," said **Stefan Edelmayer, Head of Data & Analytics at SPAR ICS**, about the project.

"We are very pleased to be able to support the SPAR ICS team as a strategic cooperation partner in setting up the data and AI platform. There are still many innovative and exciting topics that will create a new shopping experience and optimize further processes," said **Markus Nemeth, Managing Director at paiqo**.

"The topic of sustainability will be of particular importance here. We are proud to have already brought about noticeable positive changes within the framework of the Demand Forecast solution together with the SPAR ICS team and have set ourselves the goal of continuing to

ensure a long-term-oriented and responsible use of all resources in the future."

From best practice to platform change

AI-based predictions are a valuable support for employees involved in the procurement process. Rather than replacing the previous processes, AI complements the team by optimising deliveries to more than 1,500 stores.

"The added value of the smart supply chain for retail is immense: it is obvious that efficiency is increased, and costs are reduced. But we were also able to learn from it: When the cloud and AI suddenly come into play, then you realize what else is possible in process optimization," said Edelmayer.

Microsoft Cloud can manage the large amounts of data and variable demand for computing power required to create an intelligent supply chain with advanced analytics. Microsoft's solutions provide the necessary prerequisites to scale with the requisite amount of data, both horizontally and according to workload. One challenge in the field of AI is that while enormous computing power is needed, demand fluctuates.

This requirement can be met cost-effectively, especially via a highly scalable cloud. The same flexibility can also be found in the selection and use of software: Microsoft itself covers a large part through innovative and robust services, such as Azure Synapse Analytics, Azure Databricks, Azure Purview, and Power BI, but also allows for the use of third-party and open-source solutions. This means that the requirements of users with a wide variety of expertise can be served to simplify strategic processes, but also to use innovation without restriction.

Building on its good experience with the flexibility of the Microsoft platform for the Demand Forecast solution, SPAR began to explore whether all data should run on this platform in the future: including the business warehouse. Standardizing all data on an open platform has the decisive advantage of data and algorithms interacting in an integrated way, creating synergies and thus enabling the large number of planned innovation projects to be implemented more quickly. For end users, this means faster data availability and ultimately more data-based decisions in all areas. This enables SPAR to achieve significant cost savings of 60 percent per year.

Against this background, the SPAR Group for Central Eastern Europe decided to replace the existing solution with Microsoft Azure Synapse Analytics in all areas of the company. Group-wide reporting for executives and branch managers is currently being rolled out not only in Austria, but also in Italy, Slovenia, Croatia, and Hungary. SPAR is in the process of migrating one of Europe's largest traditional business warehouse installations of more than 100 terabytes to the Microsoft Cloud to put the innovation pipeline of the next few years on a solid footing.

The project includes development of the Snack Away app, which allows customers to order their snack in advance and pick it up during their lunch break without waiting, to the digitization of the SPAR stores.

Allowing fresh ingredients to make their way from the shelf onto the breakfast table without unnecessary storage time means less food is disposed of, both in retail stores and at the customer's home. Everyone benefits from an efficient supply chain—including and above all the environments.



Deployment of AI Responsibly

Helping customers develop and deploy AI responsibly



Anthony Cook
Corporate Vice President
Microsoft

At Microsoft, all of our AI systems undergo a responsible AI process that includes reviews by a multidisciplinary team of experts to help us understand potential harms and mitigate these harms. Examples of mitigations include refining the dataset used to train models, deploying filters to limit the generation of harmful content, integrating techniques like query blocking on sensitive topics that helps to prevent misuse by bad actors, and applying technology that can return more helpful, representative, and diverse results.

Another very important safeguard is intentional and iterative deployment. Microsoft takes a measured approach to rolling out new AI-powered experiences. We often start with a limited preview with a select number of customers with well-defined use cases. Collaborations with these early customers help us ensure the responsible AI safeguards are working in practice before we scale adoption.

One of our most important responsible AI commitments is sharing our learnings with customers. We provide transparency documentation for our platform AI services. Transparency Notes communicate in clear, everyday language the purposes, capabilities,

and limitations of AI systems so our customers can understand when and how to deploy our platform technologies. They also identify use cases that fall outside the solution's capabilities and our Responsible AI Standard. Transparency Notes fill the gap between marketing and technical documentation, proactively communicating information that our customers need to know to deploy AI responsibly.

Customers also need practical tools to operationalize responsible AI practices. Over the years, responsible AI research at Microsoft has led to the incubation of tools such as Fairlearn and InterpretML. Our collection of tools has grown in capability, spanning many facets of responsible AI practice including the ability to identify, diagnose, and mitigate potential errors and limitations of AI systems. Since their original conception within Microsoft, these tools continue to improve and evolve through the contributions of active open-source communities. Our latest tool, which is currently in preview, is Azure Content Safety. This tool helps businesses create safer online environments and communities through models that are designed to detect hate, violence, sexual, and self-harm content across languages in both images and text.

Building on the Responsible AI Toolbox, Microsoft's responsible AI program has invested in integrating some of the more mature responsible AI tools directly into Azure Machine Learning for our customers' benefit. This collection of capabilities, known as the Responsible AI Dashboard, offers a single pane of glass for machine learning practitioners and business stakeholders to debug models and make informed, responsible decisions as they build AI systems or customize existing ones.

The community involved in developing, evaluating, and using AI extends beyond our direct customers. To serve this broad ecosystem, we publicly share key artefacts from our responsible AI program, including our Responsible AI Standard, Impact Assessment Template and Guide, detailed primers on the implementation of our responsible AI by design approach, and collections of cutting-edge research. Our digital learning paths further empower leaders to craft an effective AI strategy, foster an AI-ready culture, innovate responsibly, and more.

Every organization that creates or uses AI systems will need to develop and implement its own governance systems. To help them do so, we will continue to share our knowledge and tools. We will continue to share the work we are doing to build a practice and culture of responsible AI at Microsoft, including key parts of the curriculum that we use to train Microsoft employees. And we will continue to invest in dedicated resources and expertise in regions around the world to respond to customer questions about deploying and using AI responsibly.

Furthermore, we created an AI Assurance Program to help customers ensure that the AI applications they deploy on our platforms meet the legal and regulatory requirements for responsible AI.

This program is open to our customers around the world.

Regulator engagement support

We have extensive experience helping customers in the public sector and highly regulated industries manage the spectrum of regulatory issues that arise when dealing with the use of information technology. For example, in the global financial services industry, we worked closely for a number of years with both customers and regulators to ensure that this industry could pursue digital transformation on the cloud while complying with its regulatory obligations.

One learning from this experience has been the industry's requirement that financial institutions verify customer identities, establish risk profiles, and monitor transactions to help detect suspicious activity, collectively known as the "know your customer" requirements. We believe that this approach can apply to AI in what we are calling "KY3C", an approach that creates certain obligations to know one's cloud, one's customers, and one's content. We want to work with our customers to apply KY3C as part of our AI Assurance Program.

Customer councils

We are committed to bringing customers together in customer councils to hear their views on how we can deliver the most relevant and compliant AI technology and tools.

Regulatory advocacy

Finally, we'll play an active role in engaging with governments to promote effective and interoperable AI regulation. Our Microsoft Blueprint for AI governance (see next section) presents our proposals to governments and other stakeholders for appropriate regulatory frameworks for AI.

We are committed to supporting customers in Europe and around the world in implementing their own AI systems responsibly, including by developing responsible AI programs for our

partner ecosystem. Many of our partners have already created comprehensive practices to help customers evaluate, test, adopt, and commercialize AI solutions, including creating their own responsible AI systems.

We are launching a program with selected partners to leverage this expertise to assist our mutual customers in deploying their own responsible AI systems. PwC and EY are our initial launch partners, and we aim to add European partners to the program in the future.

Ultimately, we know that these commitments are only the start, and we will need to build on them as the technology and regulatory conditions evolve. We are excited by this opportunity to partner more closely with our customers as we continue our responsible AI journey together.

KY3C

Applying to AI services the "know your customer" concept developed for financial services

1

Know your cloud



2

Know your customer



3

Know your content





How do we
best govern AI?

Around the world, governments and organizations are looking for guidance on how to best govern artificial intelligence (AI). How can we use this new technology to solve our problems? How do we avoid or manage new problems it might create? There is, of course, no single or right approach. However, after several years of working to ensure our own AI systems are safe, secure, and trustworthy, we offer our ideas and suggestions as a company to empower other organizations on their responsible AI journey.

In May 2023, we published [Governing AI: A Blueprint for the Future](#), which presents a blueprint to address several current and emerging AI issues through public policy, law, and regulation. We hope this blueprint continues to drive important conversations on the necessary governance of AI. Here are some key takeaways in the paper:

A five-point blueprint for public governance of AI

1. Implement and build upon new government-led AI safety frameworks.

We believe one of most effective ways to move quickly is to build on recent advances in governmental work that advance AI safety. A new AI Risk Management Framework introduced by the U.S. Department of Commerce's [U.S. National Institute of Standards and Technology](#) is a great example of this.

2. Require effective safety brakes for AI systems that control critical infrastructure.

We propose new AI safety requirements for controlling the operation of designated critical infrastructure; part of a comprehensive approach to system safety that would keep effective human oversight, resilience, and robustness top of mind.

3. Develop a broad legal and regulatory framework based on the technology architecture for AI.

We believe that regulatory responsibilities should be placed on organizations based upon their role in managing different aspects of AI technology, such as at the applications layer, the model layer, and the infrastructure layer of AI technology.

4. Promote transparency and ensure academic and nonprofit access to AI.

We believe a critical public goal is to advance transparency and broaden access to AI resources. Our commitment is to inform the public about our AI policies, systems, performance, and progress through an annual transparency report and take other steps to expand transparency for our AI services.

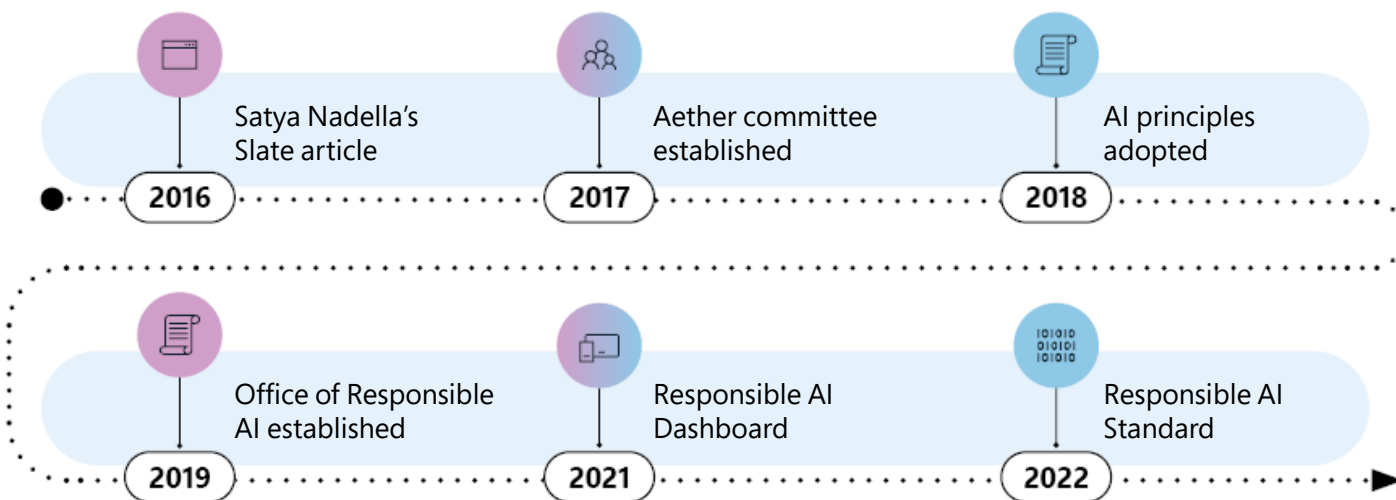
5. Pursue new public-private partnerships to use AI as an effective tool to address the inevitable societal challenges that comes with new technology.

We have witnessed what democratic societies can accomplish when they harness the power of technology and bring the public and private sectors together. We need to work together to use AI to protect democracy and fundamental rights, provide broad access to the AI skills that will promote inclusive growth, and use the power of AI to advance the planet's sustainability needs.

Responsible by design: building AI systems that benefit society

In addition to public policy, companies have an important role to play in advancing AI responsibly. For the past seven years, we have worked to advance responsible AI that is grounded in strong ethical principles. We have approached our work with a humble recognition that trust is not given but earned through action, and a deep understanding of our responsibility not just to Microsoft but our community more broadly. To put these responsible AI principles into practice, Microsoft has taken the following steps:

Our responsible AI journey



Foundational governance structures

Microsoft's Office of Responsible AI developed a governance system that incorporates many diverse teams and functions across the company.

Standardization

While our responsible AI principles state the enduring values we seek to uphold, more specific guidance can help to build and deploy AI systems responsibly. This is why we developed our Responsible AI Standard, a practical guide that memorializes a set of rules for our engineering teams to uphold AI principles in their daily practice.

Checks and balances

Our Sensitive Uses program provides an additional layer of oversight for teams working on higher-risk use cases of our AI systems, including ongoing review and oversight for high-impact issues. Think of the Sensitive Uses program as a reporting, review, and guidance framework. It begins with a mandatory reporting requirement, followed by a hands-on responsible AI project review and consulting process, culminating in project-specific guidance and requirements that are tailored to the specific project at hand.

Our Responsible AI Standard provides actionable guidance to secure beneficial uses of our products and guard against potential harms.



Principles

Which enduring values guide our responsible AI work?



Goals

What are the outcomes we need to secure?



Requirements

What are the steps we must take to secure these goals?



Tools and practices

Which aids can help us meet the requirements?



Microsoft's
commitment to
responsible AI

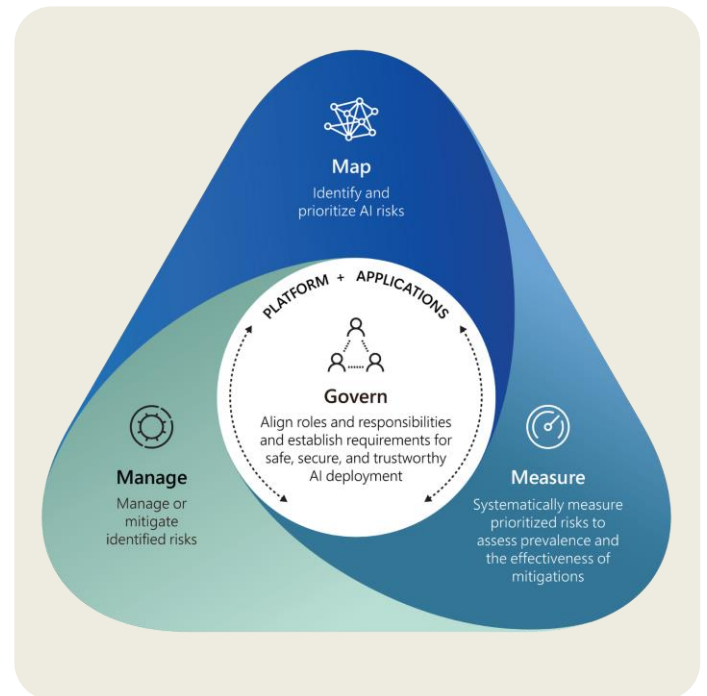
How we build, support our customers & grow

AI can be a powerful tool with the potential to transform how we work and learn—when it’s designed with people in mind.

At Microsoft, six principles guide every decision we make and product we release: transparency, accountability, fairness, inclusiveness, reliability and safety, and privacy and security.

How we build generative AI applications responsibly

Our approach to responsible AI (RAI) is guided by our Responsible AI Standard, which integrates our AI principles directly into requirements for the AI development lifecycle and tooling processes. This standard allows us to implement AI technologies with the necessary safeguards, facilitating ongoing experimentation, learning, and the establishment of cutting-edge best practices for deploying safe, secure, and trustworthy AI systems. We’ve also formalized a set of generative AI requirements building on our RAI Standard and the [NIST AI Risk Management Framework](#) to proactively map, measure, and manage risks associated with generative AI applications.



How we make decisions about releasing generative AI applications

At each stage of the map, measure, and manage process for generative AI releases, we’ve built best practices, guidelines, and tools that reflect our learnings from the last year of releasing generative applications. We use tools like responsible AI impact assessments, privacy reviews, security reviews, and red teaming to map risks associated with generative AI products. We’ve developed metrics that enable us to systematically measure known risks like ungrounded output and content risks. We’ve developed a range of technical and policy mitigations to manage identified risks at the platform- and application-level, such as content filters, appropriate human oversight, transparency documentation, and ongoing monitoring.

How we support our customers in building AI responsibly

Our goal is not only to develop AI responsibly ourselves but also to enable our customers to do the same. We invest in our customers' responsible AI goals in three ways:

- We stand behind our customers' deployment and use of AI through our AI Customer Commitments.
- We've released 30 responsible AI tools with over 100 features to support our customers' responsible AI journeys.
- We provide transparency documentation, including 33 Transparency Notes, to customers to share important information about our AI platforms and applications.


How we learn, evolve, and grow

As we've prioritized our company-wide investments in responsible AI over the last eight years, people remain at the center of our progress. From our growing internal community to the global responsible AI ecosystem, the individuals and communities involved continue to push forward what's possible in developing AI applications responsibly.

At Microsoft, each team member is pivotal in driving the responsible development of AI applications. Our commitment to innovation, collaboration, and adaptation is relentless as we strive to advance our goals and lead the industry in scaling responsible AI practices across the globe.

There is no finish line for responsible AI. Our process to govern, map, measure, and manage risks are iterative and constantly evolving.

Read the Microsoft Responsible AI Transparency Report



Learn more about the Responsible AI Transparency Report

As a company at the forefront of AI research and technology, we are committed to sharing our responsible AI practices as they evolve.

aka.ms/RAITransparencyReport2024 | aka.ms/rai



Sources and references

Sources and references

- AI Impact Story: <https://news.microsoft.com/de-at/osterreichs-wachstumspotenzial-18-prozent-mehr-wertschopfung-durch-kunstliche-intelligenz/>
- GitHub Copilot: GitHub <https://github.com/features/copilot>
- Microsoft Bing Blog: Azure OpenAI Service https://blogs.bing.com/search/2020_05/AI-at-Scale-in-Bing
- Power Platform: Microsoft <https://www.microsoft.com/en-us/power-platform>
- Microsoft Designer: Microsoft <https://designer.microsoft.com/>
- Bing and Edge: Yusuf Mehdi, Microsoft, February 2023, Reinventing search with a new AI-powered Microsoft Bing and Edge, your Copilot for the web
[Reinventing search with a new AI-powered Microsoft Bing and Edge, your copilot for the web - The Official Microsoft Blog](#)
- Microsoft AI Skills Initiative [AI Skills - Training and Resources | Microsoft CSR](#)
- Microsoft Skills for Jobs <https://www.microsoft.com/en-ie/skills-for-jobs/>
- Microsoft AI for Good Lab <https://www.microsoft.com/en-us/research/group/ai-for-good-research-lab>
- Microsoft, Anthropic, Google, and OpenAI launch Frontier Model Forum - Microsoft On the Issues
[Microsoft, Anthropic, Google, and OpenAI launch Frontier Model Forum - Microsoft On the Issues](#)
- Microsoft announces new supercomputer, lays out vision for future AI work
[Microsoft announces new supercomputer, lays out vision for future AI work – Source](#)
- OECD AI Principles <https://oecd.ai/en/ai-principles>
- OECD Framework for the Classification of AI systems [OECD Framework for the Classification of AI systems | OECD](#)
- UNESCO Recommendation on the Ethics of Artificial Intelligence [Ethics of Artificial Intelligence | UNESCO](#)
- Microsoft's AI Access Principles: Our commitments to promote innovation and competition in the new AI economy [Microsoft's AI Access Principles: Our commitments to promote innovation and competition in the new AI economy - Microsoft On the Issues](#)
- Microsoft announces new Copilot Copyright Commitment for customers
[Microsoft announces new Copilot Copyright Commitment for customers - Microsoft On the Issues](#)



Appendix



A comprehensive approach to advance safe, secure, and trustworthy AI


With AI technology rapidly evolving, there is an emerging consensus that advancing safe, secure, and trustworthy AI will require a multi-stakeholder, global approach. Much like managing air traffic and aviation standards, global bodies can shape a shared approach to high-risk AI. This will require participation from industry leaders, domestic policymakers, and global governance organizations.

Applying our voluntary commitments globally

In the United States, the White House initiated a new, collaborative process to create its own Voluntary AI Commitments. It recognized that safety, security, and trustworthiness were the core issues that needed to be addressed quickly. But rather than developing a set of proposals, it engaged the most advanced AI companies with helping define the highest risks and finding collaborative ways to address them.

This approach helped develop the clearest and most actionable set of voluntary commitments in the current ecosystem, with the necessary level of transparency and accountability.

Microsoft has endorsed all these voluntary commitments and added several of our own to help ensure the promise of AI stays ahead of its risks. We will apply these commitments globally and look forward to their broad adoption by others in the industry.



Microsoft AI Policy Commitments

Safe

Build and deploy AI that is safe to use, functions as expected, and accounts for diverse users.

Microsoft will:

- ✔ Test our AI systems using red-teaming and systematic measurements
- ✔ Contribute to industry efforts to develop evaluation standards for emerging safety and security issues
- ✔ Build and implement tools to help people identify AI-generated audio or visual content
- ✔ Implement globally recognized AI risk management framework developed by the U.S. National Institute of Science and Technology
- ✔ Implement robust reliability and safety practices for high-risk models and applications

Secure

Fortify our AI systems from cyber threats to preserve integrity.

Microsoft will:

- ✔ Ensure that the cybersecurity risks of our AI products and services are identified and mitigated
- ✔ Participate in global multistakeholder exchange of threat information
- ✔ Support the development of a global governance regime for highly capable models, including international standards and a licensing framework
- ✔ Support the development of an expanded “know-your-customer” concept for AI services

Trustworthy

Develop and deploy AI systems in a way that earns people’s trust.

Microsoft will:

- ✔ Release an annual transparency report on the governance of our responsible AI program
- ✔ Design our AI systems so that people know when they are interacting with an AI system and be transparent about system capabilities and limitations
- ✔ Increase investment in our academic research programs, including the Accelerating Foundation Model Research program
- ✔ Explore pilot projects to stand up National AI Research Resources globally
- ✔ Support the development of national-level registries of high-risk AI systems

Public-private collaboration

In addition to the adoption of voluntary commitments, leading AI companies have an obligation to engage with and provide clear information to policymakers. Microsoft is working with policymakers globally to share information about the latest advancements, both in terms of the technology itself, as well as around the risk-mitigation approaches being developed. We are also studying the regulatory models built for other technologies, including civil aviation, maritime shipping, financial transactions, and more to determine how those can be applied to AI.

Industry leaders in AI, including Microsoft, must be willing to learn from policymakers and answer questions in a timely manner. This technology is moving quickly, and Microsoft is committed to being part of this public-private collaboration to help ensure it is applied responsibly.



Tech stack explainer: *At a glance*

- A technology stack is a collection of hardware, software, and services that include tools, programming languages, and frameworks that are used to build and run applications. In the context of AI, a technology stack would include tools and frameworks that are specifically designed for developing, deploying, and using AI applications.
- We propose that a legal and regulatory architecture should consider each layer of the AI technology stack – or AI tech stack- how it applies to different stakeholders based on their role.
- In many cases, existing laws and regulations can be applied and built upon, but new regulatory approaches will be needed in some certain instances such as for the development of new powerful AI models, and for the deployment and use of these models in advanced datacenters.


Regulatory architecture for the AI tech stack

When you enter a search query into Bing Copilot, you are interacting with a sophisticated AI application. That search query will be processed by a generative AI model, but between you and that model are layers of AI technology, and the model itself runs on high-powered infrastructure. Together, these layers are known as the 'AI tech stack.'

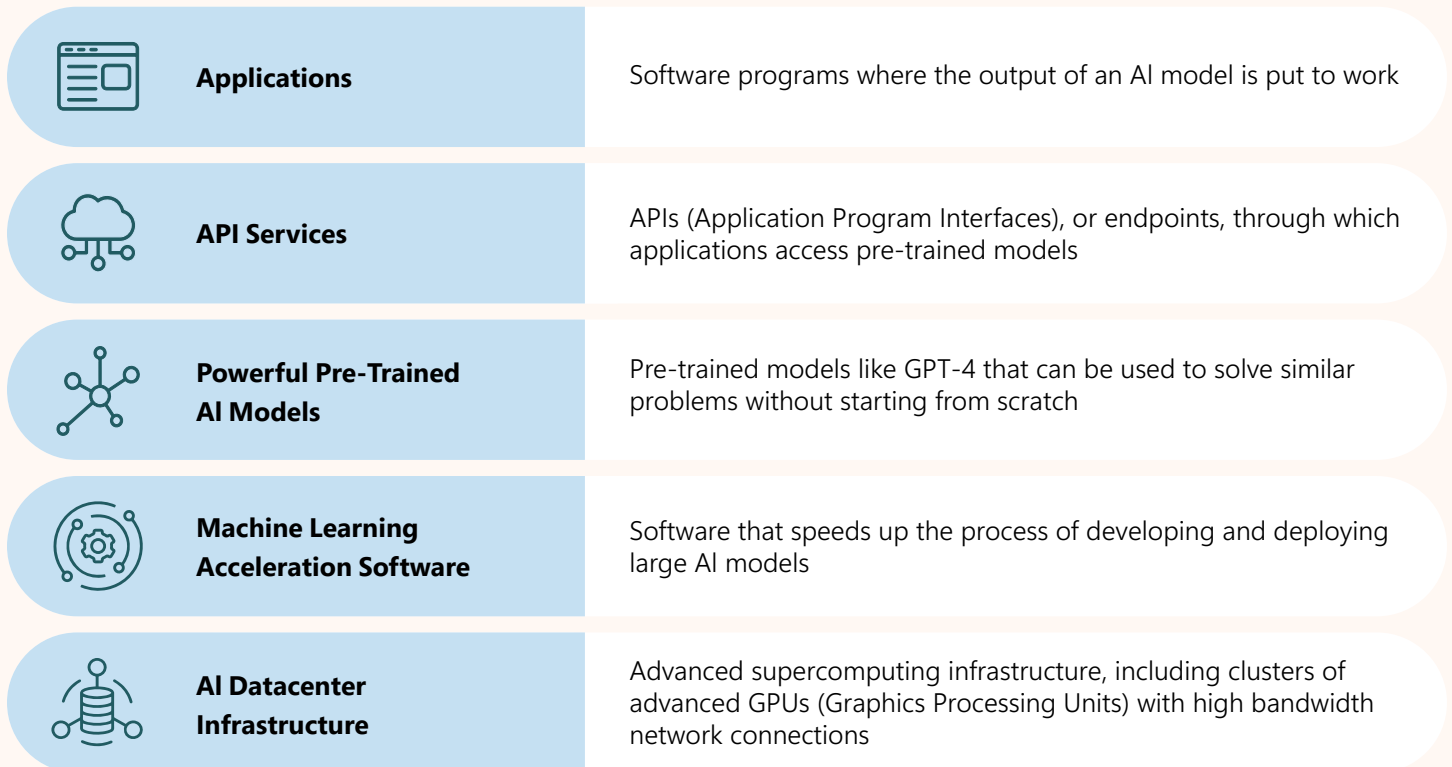
Understanding the AI tech stack is important when considering how to create policies aimed at ensuring AI is safe, secure, and trustworthy.

The diagram below shows a model of an AI tech stack. On the top layer, users interact with a model like GPT-4 through **applications** like ChatGPT, Bing Chat, and GitHub Copilot. The power of new generative AI models has prompted companies to create new or modify existing applications to incorporate AI-enhanced features.

The next layer is **API (Application Programming Interface) Services**, which provide access to models and capabilities that are hosted on Microsoft's infrastructure. This means that Microsoft Azure OpenAI customers can build applications by calling the API to submit prompts and receive outputs from GPT-4. Powerful pretrained AI models like GPT-4 are created by developers and research scientists at companies like OpenAI. In the case of GPT-4, OpenAI developed their models by using the AI supercomputing infrastructure that Microsoft built for them.



The technology stack for AI foundation models



Source: [Governing AI: A Blueprint for the Future](#)

Create a regulatory architecture that reflects AI's technology stack

With the rapid advances in innovation, we need laws and regulations that apply to applications that use AI model outputs to deliver services to individuals and organizations. Fortunately, we have long had laws in place to protect the public and, repeatedly, our courts and agencies have upheld the values we regard as timeless amidst constant technological change. Rapid advances in AI mean they will need to do so again.

For example, it is unlawful for a bank to discriminate based on race or gender when deciding who to approve for a mortgage. If a bank now wants to use AI to help it make lending decisions, it will need to ensure that this does not lead to unlawful discrimination. And what's true for banks and mortgages is true in every field. Existing laws will continue to apply to the decisions and actions of organizations and individuals alike.



We propose that a legal and regulatory architecture should be in place that considers each layer of the AI technology stack, and the role that different stakeholders in the AI technology stack have. For most individuals and organizations, the application layer is where their safety and rights will be most directly impacted. As a result, we need to ensure that existing laws and regulations that govern conduct and societal impact apply to applications.

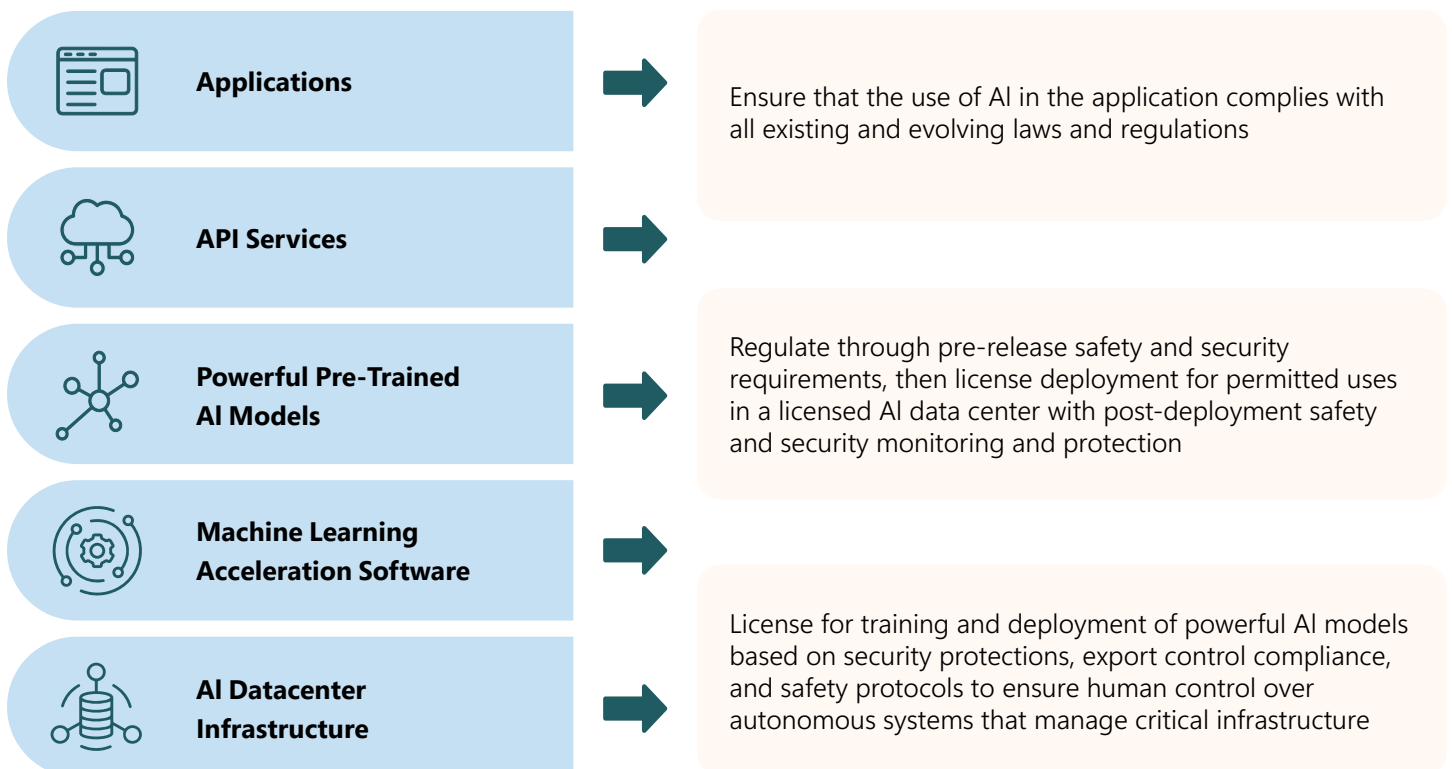
No one will be able to stand up in court and proclaim, “but Your Honor, a machine made me do it.”

While existing laws and regulations can be applied and built upon for the application layer of the AI tech stack, we believe that new approaches are needed for two additional layers beneath that. The first of these is for the development of the most powerful new AI models, and the second is for the deployment and use of these models in advanced datacenters. From our work on the frontiers of AI, we have seen a new class of model emerge. Highly capable foundation models that are trained on internet-

scale datasets and are effective out-of-the-box at new tasks—a model like GPT-4 allows you to create a never-seen-before image using words in one prompt.

At the cutting-edge, the capabilities of foundation models are at once very impressive and can be harder to predict. As the models have been scaled up, we have seen anticipated advances in capabilities, as well as surprising ones that we and others did not predict and could not observe on a smaller scale.

A proposed AI regulatory architecture



Source: [Governing AI: A Blueprint for the Future](#)

We shouldn't leave these issues of societal importance to good judgment and self-restraint alone. We need regulatory frameworks that anticipate and get ahead of the risks.

The work needed to harness the power of these models and align them to the law and societal values is complex and evolving. These characteristics of highly capable models present risk surfaces that need to be addressed. To date, we have benefited from the high safety standards self-imposed by the developers who have been working at the frontiers of AI model development. We propose that a legal and regulatory architecture should be established that considers each layer of the AI tech stack so that all stakeholders are held accountable to produce safe, secure, and trustworthy AI.



