



# AI in Europe

Meeting the opportunity  
across the EU





# AI in Europe

Meeting the opportunity  
across the EU

## Contents

### 04 The promise of AI

05 Foreword by Ralph Haupter

### 08 AI explained

09 What exactly is AI?

10 AI-powered possibility

11 AI at work – Q&A with Kevin Scott

12 A brief history of AI: From Turing to ChatGPT

14 Developing and scaling AI

15 Putting people at the centre

### 16 AI in action

18 AI for a healthier future

21 AI for a more sustainable future

22 AI for a growing future in agriculture

25 AI for smarter manufacturing

27 AI for better education

29 The value of AI

### 30 Responsible development and use of AI

32 Responsible AI – Q&A with Natasha Crampton

35 Helping our European customers develop  
and deploy AI responsibly – Antony Cook

37 Advancing trustworthy AI

### 38 AI governance

39 Advancing AI governance in Europe  
and internationally – Brad Smith



# The promise of AI



**Ralph Haupter is Microsoft's President for Europe, Middle East and Africa, leading an organisation with over 11,000 employees across the region.**

# Foreword by Ralph Haupter

## The age of AI has arrived

We use and experience it every day. In the office, for research, when we're gaming, and so much more. In the workplace, it helps us focus on what we as humans are best at: creative problem solving, empathy and collaboration.

Across every industry AI is helping organisations be more efficient and competitive. And it's playing an important role in boosting productivity and addressing the pressure Europe's demographic challenge is putting on many sectors, including health care.

It's hard to overstate AI's potential and its importance.

Europe has much to give and to gain from AI. The EU is the world's largest exporter of manufactured goods and services. It ranks first for both outbound and inbound international investment and is the top trading partner for 80 countries. It is deeply connected to the fabric of the global economy. Harnessing AI provides an opportunity for Europe to build on these advantages – to innovate, compete and grow as business around the world becomes ever more digital-first.

Europe is home to almost 50% more software developers than the United States. Imagine each and

every developer using an AI copilot for coding to boost their productivity. Just think of the powerful turbo charge that gives Europe's digital economy and our competitiveness in the global marketplace.

I already see many great examples of how European firms are innovating with AI and creating value for themselves and their customers, such as ABB, Vodafone, Commerzbank, Erasmus Medical Center, the Flemish government and so many more. In health care, for example, Novartis is using algorithms to develop next-generation medicines; to manufacturing, with Denmark's Danfoss turning factory floor data into insights that help improve operational efficiency while reducing energy consumption by up to 40% at the same time.

In sustainability, where Europe is showing strong global leadership, AI is playing a major role in the transition to green energy. For example, Vestas is using AI to generate more energy from wind turbines. And Spanish energy firm Repsol is using AI to plan the majority of its green efforts as it looks to become the first company in the sector to achieve zero net emissions by 2050.

AI innovation is gaining momentum. Analyst group IDC forecasts annual spending on AI in Europe will more than double in the next three years from €31 billion today to reach €65 billion in 2026.

The key question for me, though, is not whether AI will transform the way we Europeans work, or drive innovation at scale in business. Rather, we should be asking, how are we going to benefit from it?

It's up to us – Europe – to decide how we want to leverage AI and on what terms. How we set up and empower European companies and organisations to take advantage of AI, and how we use it to advance our welfare, wellbeing and generate equitable wealth.

Answering these questions requires us to look holistically at AI across government, commerce, the public sector, and civil society. With our extensive European partner network and deep customer relationships across the region, Microsoft is well placed to help bring stakeholders together. Our mission is to empower every person, every organisation, and millions of Europeans today to use and trust our technology. It's both a privilege and responsibility that informs our approach to AI.

Of course, AI is not new. The computer science behind the technology has been in research and development for many decades. Today, it has stepped into the limelight with the advent of new generative AI applications like ChatGPT. Popular consciousness of the power and potential of AI is growing rapidly. And, with it, the need to have clarity on how, here in Europe, we want to and will benefit from AI both today and into the future.

## Our commitment to be a responsible steward of AI and partner to Europe

There's no doubt AI is a powerful tool. Not since the printing press have we seen an innovation, it has been said, that could have as much economic and social impact. Like the printing press, it is a tool that can be used to help advance society. But it can also be used by bad actors.

Since 2017, Microsoft has invested to ensure AI systems are responsible by design.

Today, the company has a large and growing team responsible for implementing best practices around building safe, secure, and transparent AI

systems designed to benefit society. Microsoft's Office of Responsible AI developed a governance structure that pulls on the expertise of core teams within engineering, research and policy to advance responsible AI priorities.

To oversee progress, drive accountability and leadership alignment on responsible AI commitments Microsoft established a Responsible AI Council. The Council is chaired by senior staff including Vice Chair and President, Brad Smith, and Chief Technology Officer, Kevin Scott. The council meets regularly, convening representatives of Microsoft's core research, policy, and engineering teams driving Microsoft's responsible AI efforts. This includes evaluating potential uses that might pose a conflict of interest with Microsoft's approach to responsible AI – and working with customers to find responsible and appropriate approaches.

Ensuring AI benefits everyone goes far beyond our four walls, however. It will take commitment and collaboration across a diverse set of stakeholders including tech firms and their customers, policymakers, academics, and civil leaders. For example, in the EU, Microsoft has been actively calling for AI regulation since 2019, particularly around AI systems like facial recognition technology, that could have consequential impact on privacy and civil liberties and added [safeguards for responsible use](#). We need strong regulatory guardrails to protect citizens, governments,

“What really matters is a commitment to asking the right questions”

and businesses across Europe, which we discuss in our recently released *Governing AI: A Blueprint for the Future* whitepaper.

Further, the private and public sectors have a shared responsibility to make certain those in the workforce are equipped with the skills they need to flourish. That is why EU initiatives such as the European Skills Agenda and the Digital Education Action Plan are so important. It's also why Microsoft continues to invest in our global skills initiative, which has already benefited several million people across Europe.

No one group or individual has all the answers. That's OK.

What really matters is a commitment to asking the right questions – questions around fairness, transparency, inclusivity, and, ultimately, how any new innovation will benefit people – and working collectively to secure the benefits of AI while mitigating the potential risks.

Europe has what it takes to play a leading role in AI globally. We have the technology ecosystem. We have the talent. Most importantly, we have a set of shared values that will see AI shape more prosperous, inclusive, and sustainable communities.

*For the purposes of this document, references to Europe include the EU, and the EFTA.*

## Microsoft's responsible AI principles



### Fairness:

AI systems should treat all people fairly



### Reliability & safety:

AI systems should perform reliably and safely



### Privacy & security:

AI systems should be secure and respect privacy



### Inclusiveness:

AI systems should empower everyone and engage people



### Transparency:

AI systems should be understandable



### Accountability:

People should be accountable for AI systems



# AI explained

AI is changing how we work, live, and play. Microsoft AI provides billions of intelligent experiences every day to people using our business software and services, gaming on our Xbox platform, and in keeping organisations secure.

Our AI tools and technologies are designed to benefit everyone at every level in every organisation. They're used in workplaces, home offices, academic institutions, research labs, and manufacturing facilities around the world. They're helping everyone from scientists and salespeople to farmers, software developers, and security practitioners.

# What exactly is AI?

Using mathematics and logic, an artificially intelligent computer system simulates the reasoning that humans use to learn from data and make decisions.

Algorithms are at the core of every AI system. They are sets of instructions or rules that the system follows to process and analyse data. Different algorithms are used for different tasks, such as recognising images or processing language, for example.

An AI model is trained on a dataset representative of the task or problem it's being developed to solve.

Data could be in the form of unstructured data, such as text, images, or audio, or structured data that can be stored in a standardised format.

The AI model uses the patterns and relationships it discovers in the data to adjust and optimise its performance resulting in a learned representation of the training data, which can be used to interpret and make inferences from new input or data.

One of the techniques used to train AI is machine learning. It's akin to teaching a computer program to recognise patterns by showing it examples. Just as we learn from experience, machine learning algorithms learn from data.

“AI development has also been inspired by the structure and functioning of the human brain”

They look for similarities and relationships in the data to perform tasks such as making predictions, recommendations, classifications and more.

Imagine you want to train a computer program to recognise cats. You would feed a machine learning algorithm a large number of pictures, each labelled as to whether the image contains a cat or not. By analysing these labelled examples, the algorithm can identify patterns such as shapes, colours, and textures associated with cats, and create a model based on this learning. Once trained, the model can then be used to analyse new, unseen images and predict whether they contain a cat or not. The more examples

the algorithm sees, the better the model becomes at recognising cats accurately.

A subset of machine learning under the broader umbrella of AI is reinforcement learning. It involves asking an AI system to make a sequence of decisions to maximise 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 neurons, organised in layers. Each neuron receives input, performs a computation, and passes the result to the next layer. Neural networks excel at pattern recognition and can learn complex relationships in data.

As you can see, AI is a wide umbrella under which sit a range of different but complementary computer science fields and sub-fields – many of which have been in research for decades. But whichever approach is used to develop an AI system, having access to data and computing power to train your AI model are fundamental build blocks.

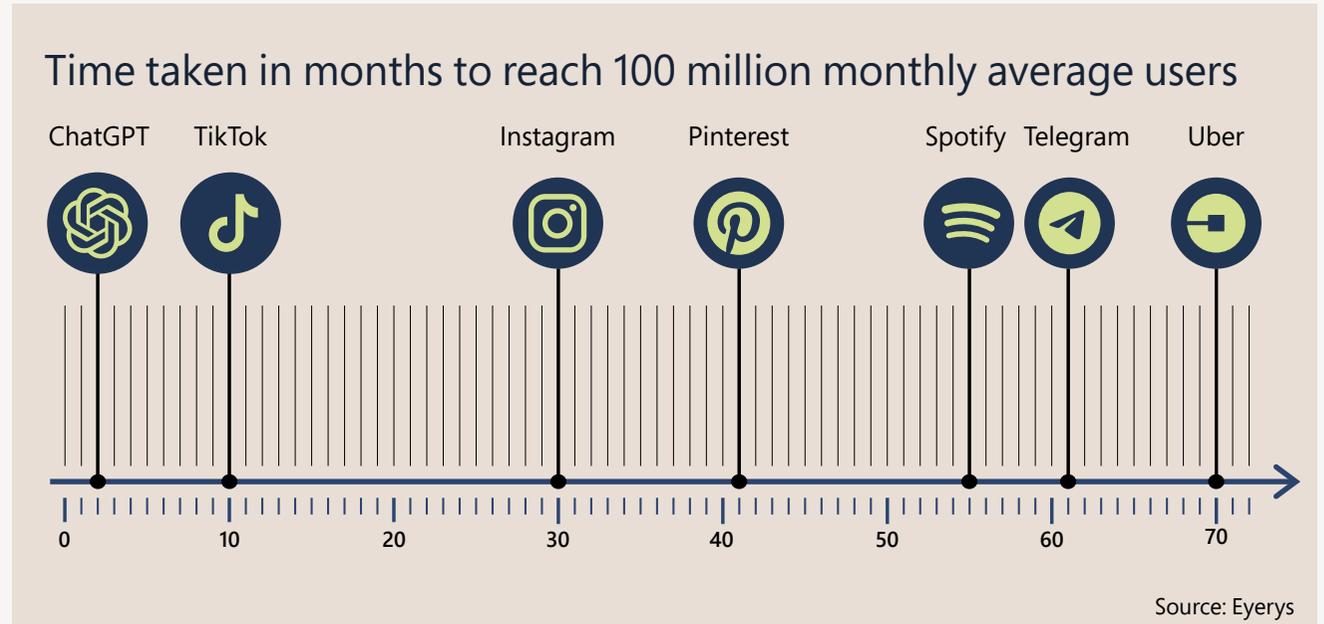
# AI-powered possibility

AI enables machines to carry out a wide range of tasks, many of which weren't possible before its advent.

Natural language processing (NLP), for example, extracts meaning from text or speech data. It enables computers to interpret language and perform tasks like sentiment analysis, language translation, and text generation. AI-powered Computer Vision similarly enables machines to "see" and decode the visual world around us. It involves techniques such as object detection and image recognition, and has a wide range of applications from medical imaging to autonomous vehicles.

These capabilities and more are transforming how we use, benefit from, and interact with computer technology. AI can serve as a "copilot" to augment our human ingenuity and creativity through automating tasks and providing new insight.

Over the last decade, the AI field has made significant progress on perception and language tasks. There have been recent advances in the form of 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 are capable of performing a wide range of tasks. With a simple prompt like "describe a scene of the sun rising over the beach", generative AI models can output a detailed description or produce an image, which can then be animated or even turned into video.

Many recent generative AI models such as language models are not only good at generating text but also generating, explaining, and debugging code. [GitHub](#)

[Copilot](#), for example, leverages OpenAI's Codex model to assist developers in writing code.

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

# AI at work – Q&A with Kevin Scott



**Kevin Scott is Executive Vice President of Technology & Research, and the Chief Technology Officer of Microsoft. Scott's**

**20-year career in technology spans both academia and industry as researcher, engineer and leader.**

## How do you see AI benefiting us at work?

The entire knowledge economy is going to see a transformation in how AI helps out with repetitive aspects of your work and makes it generally more pleasant and fulfilling. Not only will it help with tasks like writing, editing or creating reports, it is going to apply to almost anything – designing new molecules to create medicine, or making manufacturing “recipes” from 3D models.

This is the “copilot for everything” dream – that you would have a copilot that could sit alongside you as you’re doing any kind of cognitive work, helping you not just get more done, but also enhancing your creativity in new and exciting ways.

## Will AI help make work more rewarding?

In many cases, people now have new, interesting, and fundamentally more effective tools than they’ve had before. For some workers, it’s literally enhancing that core flow that you get into when you’re doing the work; it speeds you up. It’s like having a better set of running shoes to go run a race or marathon.

This is exactly what we’re seeing with the experiences developers are having with Copilot; they are reporting that Copilot helps them stay in the flow and keeps their minds sharper during what used to be boring

“[AI]... speeds you up. It’s like having a better set of running shoes to go run a race or marathon”

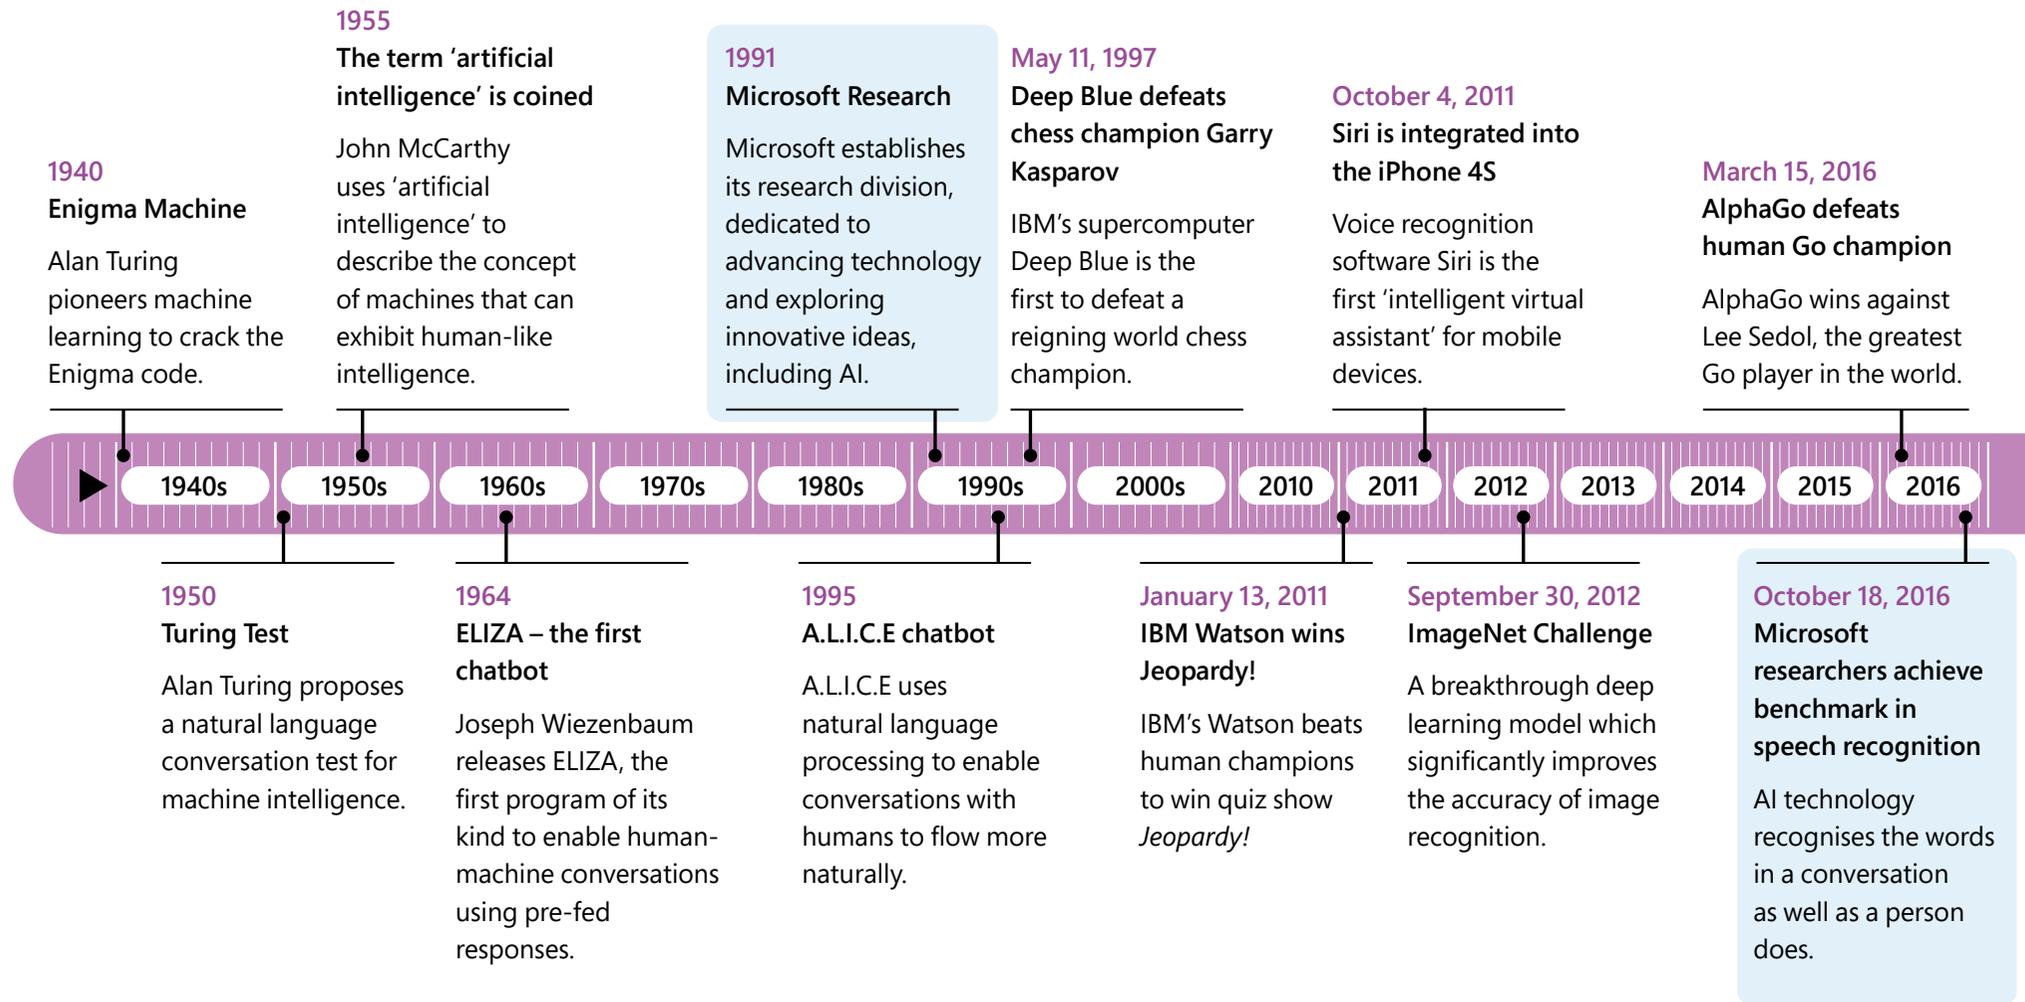
and repetitive tasks. And when AI tools can help to eliminate drudgery from a job, something that is super repetitive or annoying or that was getting in their way of getting to the thing that they really enjoy, it unsurprisingly improves satisfaction.

## Will AI transform the way we work?

We’ve gone from machine learning in a few places to being integrated into a wide range of different products and services: Everything from how your Outlook email client works, your predictive text in Word, your Bing search experience, to what your feed looks like in Xbox Cloud Gaming and LinkedIn. There’s AI all over the place making these products better.

As all these AI systems continue to grow and evolve, I think we can expect that these advances are going to fundamentally change the nature of work, in some places more than others, and in some cases create a whole spate of new jobs that didn’t exist before. You can look back and see the same thing happen adjacent to all kinds of famous paradigm shifts in technology over history: the telephone, the automobile, the internet. And, just like with those examples, we’re going to need new ways to think about work and new ways to think about skills.

# A brief history of AI: From Turing to ChatGPT



Microsoft projects and collaborations are highlighted in a blue box

**July 12, 2017**

### Microsoft AI for Earth

Microsoft launches the AI for Earth programme, to harness AI to address environmental challenges such as climate change, biodiversity loss, and water conservation.

**February 14, 2019**

### GPT-2 language model

OpenAI's GPT-2 language model generates impressive text content, translates, answers questions, summarises passages.

**May 19, 2020**

### Microsoft announces its first AI supercomputer

Microsoft announces it has built one of the top five publicly disclosed supercomputers in the world, making new infrastructure available to train extremely large AI models.

**October 14, 2020**

### Azure Cognitive Services achieves benchmark in image captioning

Microsoft researchers build an AI system that can generate captions for images with high accuracy.

**November 30, 2022**

### ChatGPT

ChatGPT marks a step-change in chat bots, with capability to provide detailed responses and articulate answers across many domains of knowledge.

**March 14, 2023**

### GPT-4

GPT-4's multimodal capabilities allow it to accept and produce text and image inputs and outputs, differentiating it from the text-only capabilities of GPT-3 and 3.5.

2017

2018

2019

2020

2021

2022

2023

**December 3, 2018**

### Microsoft real-time translation achieves benchmark.

Real-time voice translation and captions is integrated into Office 365, based on deep learning work from the Microsoft Research team.

**February 13, 2020**

### Turing-NLG language model

Microsoft Project Turing introduces Turing Natural Language Generation (T-NLG), the largest model ever published at 17bn parameters, which outperforms a variety of language modelling benchmarks.

**June 11, 2020**

### GPT-3

The GPT-3 deep learning model is launched, trained on large text datasets of hundreds of billions of words.

**June 21, 2022**

### GitHub Copilot

GitHub Copilot is made generally available to individual developers. For the first time in the history of software, AI can be broadly harnessed by developers to write and complete code.

**February 7, 2023**

### AI comes to Bing and Edge

Microsoft integrates AI into the new Bing and Edge, creating a 'new era of search' with added conversational tools.

**March 16, 2023**

### Microsoft 365 Copilot

Microsoft announces a new AI copilot feature for Microsoft 365, with the ability to create documents, read and summarise emails, craft presentations and more.

# Developing and scaling AI

More than a decade ago, Microsoft forecast an exponential growth in demand for AI systems and started to build special computing infrastructure to handle it. Inspired by early research developments, we've integrated large-scale language models in services ranging from [Microsoft Bing](#) to Microsoft 365, and have enabled other companies to take advantage of these technologies through Azure AI and the [Azure OpenAI Service](#).

We believe that every organisation in the world should benefit from the power of large-scale AI models. We've developed platforms, tools and a supercomputing infrastructure that would allow any developer to build and scale their own AI innovation.

Through Azure OpenAI Service, we provide businesses and developers with access to high-performance AI

models, such as GPT-4, Codex and DALL-E 2.

The service is based on the same platform we use to power AI models in our own products, including [GitHub Copilot](#), [Power Platform](#), [Microsoft Designer](#), [Bing and Edge](#).

We continue to evolve our AI infrastructure based on feedback and insights from training and serving AI models at scale, and our teams are working in lockstep with industry partners on the design of processors, networks and datacentres optimised for AI.

Microsoft is committed to building Azure into an AI supercomputer for the world, serving as the foundation of our vision to democratise AI as a platform. We pushed the frontier of cloud supercomputing technology, announcing our first top-five supercomputer in 2020, and have subsequently constructed multiple AI supercomputing systems at massive scale.

The convergence of innovations in infrastructure, machine learning acceleration software, platform service, and modelling powered by cloud technology has created the perfect conditions to accelerate innovation in AI and enable every company to become an AI technology company.

## Why the cloud is the critical infrastructure behind AI



**Training AI models requires both large datasets and AI-optimised cloud computing infrastructure**



**Cloud-based AI platforms make it easier for organisations to develop their own AI applications**



**Cloud computing resources give organisations of all sizes a cost-effective way to run and scale AI innovation**

## The technology stack for AI foundation models



### Applications

Software programs where the output of an AI model is put to work



### API services

APIs (Application Program Interfaces), or endpoints, through which applications access pre-trained models



### Powerful pre-trained AI models

Pre-trained models like GPT-4 that can be used to solve similar problems without starting from scratch



### Machine learning acceleration software

Software that speeds up the process of developing and deploying large AI models



### AI datacentre infrastructure

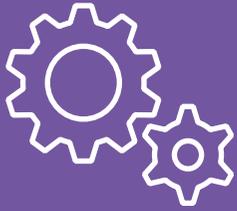
Advanced supercomputing infrastructure, including clusters of advanced GPUs (graphics processing units) with high bandwidth network connections



## Putting people at the centre

At every stage from development to deployment and ongoing supervision, people play an essential role in AI: collecting, preparing, and curating data to train AI systems; evaluating which algorithms or models to use; validating system performance in real-world scenarios to assess its behaviour and robustness before deployment; gathering and analysing ongoing user feedback to identify errors, biases, or limitations to make improvements to the data, algorithms, or models; and many more roles.

By prioritising safeguarding throughout, AI developers and practitioners can put in place measures to mitigate potential harm, ensure fairness, and promote transparency. This hands-on involvement is critical to shaping and guiding AI systems to achieve reliable and beneficial outcomes.



# AI in action

Recent months have put a spotlight on advancements in generative AI, and the possibilities of how it can shape the future of industry. The pace and scale of technological innovation happening today is unprecedented and breathes optimism into a demanding economy. It's also providing a powerful boost to our efforts to address the major societal challenges we face.

And that matters to all of us, as it's ultimately about improving people's lives.



## Health

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



## Sustainability

The climate challenge is a race against the clock to decarbonise. AI is helping accelerate and scale our efforts to transition to clean energy and mitigate the impact of global warming on people and communities.



## Agriculture

Ensuring we have access to affordable food even 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, especially water.



## Manufacturing

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



## Education

AI is transforming the way we learn and train, making it more accessible, personalised and impactful, benefiting educators and students alike.



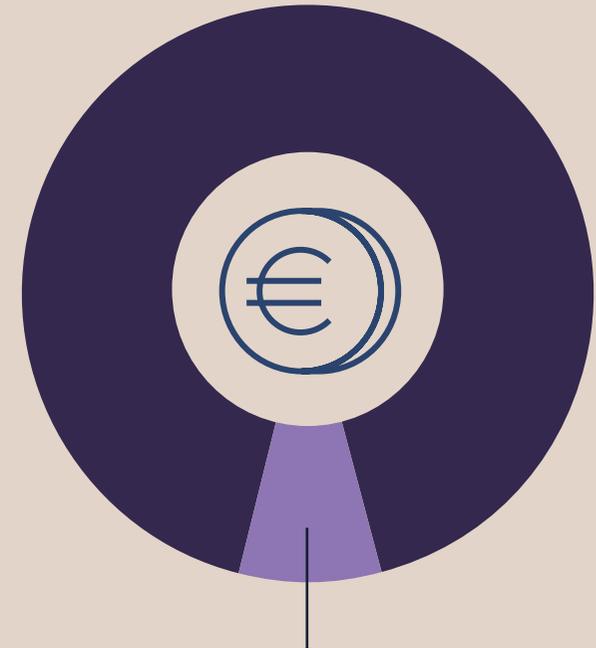
## AI for a healthier future

Europe is at a health crossroads. By 2050, one in four Europeans will be over 65. Health systems will need to shift from episodic care to an approach that is much more proactive and focused on long-term care management, according to a report by [McKinsey and EIT Health](#). The report surveyed health care professionals, who said AI solutions can ease administrative tasks, facilitate home-based care, and improve clinical practice.

According to Eurostat, health care expenditure is rising across Europe. In 2021, EU governments spent 8.1% of their GDP on health. And more than half of health care expenditure in the EU in 2020 was for curative and rehabilitative care. Prevention is better than a cure, and AI can help accelerate public health and clinical research to identify solutions to some of the most pressing health care challenges.

On top of this, the sector faces a staffing crisis. Europe faces an ageing health workforce, with many doctors approaching retirement age and a surge in chronic illnesses. Dr Hans Henri P. Kluge, WHO Regional Director for Europe stated in March 2023: “Health providers and workers across our region are clamouring for help and support.” AI and Cloud technologies can help burned-out health care staff tackle this workload.

### Proportion of expenditure by EU governments spent on health in 2021



“In 2021, EU governments spent 8.1% of their GDP on health”

Source: Eurostat



## Listening out for heart attacks

### Denmark

---

Quick ambulance dispatch in the case of cardiac arrest can mean the difference between life and death. But the people making the call aren't medically trained and often struggle to communicate clearly.

The Emergency Medical Services (EMS) in Denmark's Capital Region is pioneering the use of AI to identify cardiac arrests during emergency calls. In collaboration with Microsoft and technology partner Corti, EMS has developed a virtual assistant that can listen in and flag a likely cardiac arrest using a database of some 1,400 recorded calls to the emergency number.

Freddy Lippert, CEO of EMS, says: "The first 90 seconds of that [emergency] call are crucial to the safety of the patient... once that diagnosis is made, we initiate a cascade of high-priority response measures." He adds: "Copenhagen alone receives 130,000 emergency calls per year, 1% of which usually is a cardiac arrest ... Identifying and correctly diagnosing that 1% is often difficult, and that's where AI and technology can really make a difference."

The solution's success has paved the way for many more use-cases – playing a crucial role in EMS's response to time-critical cases.

“A virtual assistant that can listen in and flag a likely cardiac arrest”



## Teaching Oxford's hospitals

### UK

---

Doctors have been taking notes since Ancient Egypt, when they used papyrus to note medical histories. However, even with digital tools, writing them up takes up precious time, making it a perfect job for AI support. Nuance, which develops AI and voice recognition for medical applications, is used by more than 10,000 health care organisations worldwide.

Oxford University Hospitals NHS Foundation Trust (OUH) is a world-renowned centre of clinical excellence and one of the largest National Health Service teaching trusts in the UK. It relied on writing and sending printed letters to patients in many departments.

In 2017, one department at the OUH was averaging a 12-day turnaround of letters from frontline doctors. Their target was five days. The issue was a chronic shortage of administration staff and the complex, costly work of transcription, printing, and mailing.

Nuance's Dragon Medical One is a tool that uses AI to input voice-generated content directly into clinical systems to develop structured clinical letter templates. Giving this mundane, but indispensable, job to AI cut turnaround for letters from an average of 12 days to three. Where no blood results are required, letters can be sent in real-time.



## Innovating new cures Belgium

---

Developing new drugs has traditionally been a time-consuming, expensive, and uncertain process.

By leveraging the power of AI, researchers can shorten that process to weeks or even days. Tools that use AI can sift through results from decades of laboratory experiments and suggest new solutions – like reworking classic recipes from an old cookbook. These can then be fast-tracked for testing and, if proven safe and effective, be developed into new medications. This AI-augmented process could cut out years of trial-and-error experimentation.

Our partnership with Belgian pharmaceutical group UCB enables its scientists to discover new medicines in a more efficient way. As part of the open-science [COVID Moonshot project](#), UCB's medicinal and computational chemists contributed highly effective designs for an orally bioavailable anti-viral for COVID-19. Building on this work, UCB and Microsoft are [collaborating](#) to apply a similar approach to speed up the process of identifying candidate drugs for clinical trials that target diseases in immunology and neurology.

“Tools that use AI can sift through results from decades of laboratory experiments and suggest new solutions”



## AI for a healthier future UK

---

In the UK, [Our Future Health](#) is a collaboration between the private, charity, and public sectors to address some of the leading causes of death and serious illness in the UK, including dementia, cancer, diabetes, heart disease, arthritis, and strokes.

The programme is recruiting five million volunteers to share health information about themselves, which AI-powered tools will turn into health insights. Our Future Health will use Microsoft Cloud to securely store the data in a Trusted Research Environment (TRE) provided by DNAnexus, which will sit in a UK Azure region.

Cloud services “will be an integral part of Our Future Health, underpinning so many important systems that are essential to the running of the programme and ultimately helping to create one of the most detailed pictures we’ve ever had of people’s health,” says Andrew Roddam, Chief Executive of Our Future Health.



# AI for a more sustainable future

Sustainability is the challenge of the century. According to the World Meteorological Organization (WMO), there is a 98% likelihood that at least one of the next five years, and the five-year period as a whole, will be the warmest on record. To curb the progression of climate change, humans urgently need to change the way we use the planet's resources and accelerate decarbonisation.



## Ancient forces, future solutions Germany

A sustainable future means switching to clean, renewable energy sources. AI can play a vital role in making them more efficient. German start-up Solytic used Microsoft Azure to create monitoring and analytics software that helps solar plant operators lower costs, improve performance, and achieve greater photovoltaic efficiency. Photovoltaics are a key technology in the energy transition, and it is essential solar installations are configured to work as efficiently as possible. But this is not always the case if a panel is chipped, dirty, or in shadow.

This is where Solytic comes in: its software monitors installations, collects their data, and derives insights to inform decision-making. "Modules don't have to be visibly damaged to underperform. We identify the reason behind drops in performance," says Konrad Perényi, managing director and co-founder of Solytic. If an installation is not working as efficiently as it could be, the solution examines data from the installation and those nearby to work out what is amiss.

Solytic now monitors over 200,000 photovoltaic plants across 60 countries, and because the future is interoperable, the hardware-agnostic software connects with any type of Photovoltaic. Solytic receives thousands of readings from its customers every single second; analysing them would be impossible without the cloud and AI.

“AI can play a vital role in making clean energy more efficient”





## A breath of fresh AI Denmark

Wind power is one of the fastest-growing renewable energy technologies, according to the IEA. Abundant, free, and entirely without pollution, wind energy has many advantages. But one challenge is that the amount of energy that is harnessed can change daily based on weather conditions. Wind turbines cast a wake, or a “shadow effect” that can slow other turbines that are located downstream.

Danish manufacturer [Vestas Wind turbines](#) is using AI to counteract the effect by turning turbine rotors to point away from oncoming wind and deflect the “shadow effect.” The solution uses Microsoft Azure high-performance computing, Azure Machine Learning and reinforcement learning-based controller design platform from Microsoft partner [minds.ai](#). Reinforcement learning is a type of machine learning in which AI agents can interact and learn from their environment in real-time, and largely by trial and error. Reinforcement learning tests out different actions in either a real or simulated world.

“You can use AI to optimise the construction, siting and the operations of a wind farm, but more importantly, you can use AI to optimise across different systems, when it comes to consumption but also production,” says Espen Mehlum, head of the energy and materials programme on benchmarking for the WEF. “That’s where the huge untapped potential is for AI – we’re just scratching the surface and seeing the first use-cases.”

“AI agents  
can interact  
and learn  
from their  
environment  
in real-time”



## AI for a growing future in agriculture

Humans have been farming crops for millennia, but this ancient practice can stand to gain a great deal from modern technologies as the planet adjusts to feeding 8 billion people.

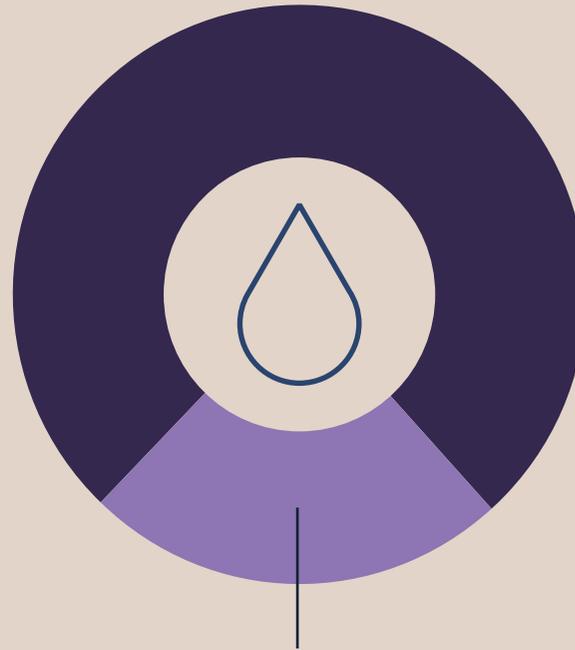
AI, Cloud, and Internet of Things (IoT) technologies have the power to transform farming and agribusinesses. And transformation is needed.

According to the [OECD-FAO Agricultural Outlook to 2031](#), to achieve the UN Sustainable Development Goals Zero Hunger target while reducing direct greenhouse gas emissions by 6%, overall agricultural productivity would need to increase by 28% over the next decade. At the same time, nearly one in three people worldwide lack regular access to adequate food. Optimising harvests while conserving precious resources, including water, is becoming essential.

In the EU alone, the agricultural sector accounts for 24% of water use and is now looking towards technology, powered by data and AI, to reduce its footprint.

“The soil is not getting any richer; the water levels are receding; there is climate change – these make the farmer’s life much harder,” says Microsoft’s Chief Technology Officer for Agri-Food, Ranveer Chandra. “One approach that can help is that of data-driven agriculture, where our goal is not to replace the farmer but to augment the farmer’s knowledge with data and AI.”

## Proportion of water use in the EU that is agricultural



“In the EU alone, the agricultural sector is accounting for 24% of water use”

Source: EEA



## Seeing the bigger picture International

Farmers know their land through decades, if not generations, of experience. Still, many decisions, such as when to sow seeds or fertilise soil, are based on rough estimates and gut feeling. AI can remove the guesswork.

[Project FarmVibes.AI](#), which runs on Microsoft Azure, makes it possible to predict and plot the ideal amounts of fertiliser and herbicide required based on the level of soil nutrients, forecast temperatures, and wind speeds across their fields. The technology can also help determine the ideal depth to plant seeds or irrigation needs based on soil moisture, and guide how different crops and practices can keep carbon sequestered in the soil.

Through these digital tools, farmers can augment their capabilities and knowledge about their farm with data and AI, helping them to make the best choices for optimising harvests with minimal resources.



## Outstanding in their field Germany

Founded in the 1860s, Bayer has been helping farmers grow crops more effectively ever since. It turned to Microsoft for a fruitful collaboration: Bayer AgPowered Services is built on top of the [Microsoft Azure Data Manager for Agriculture](#) platform. AgPowered Services are a set of capabilities and tools that turn data into insights with AI; they can track disease, insects and weed pressure, estimate heat stress impact and identify crop growth and production patterns.

Satellites, field sensors, drones, field equipment, and soil sensors generate hundreds of variables and attributes, but until recently, it's been difficult and expensive to turn that information into insights. "Only innovation can ensure global food security while protecting the planet," says Dr. Robert Reiter, Head of R&D for Bayer's Crop Science Division. "Customers can use the infrastructure and capabilities to build their own digital solutions and products on top of the most robust collection of ag data in the world."



## A sustainability dashboard Germany

[Eckes-Granini](#) is Europe's biggest fruit juice supplier. For Bernd Neufert, its procurement manager and the key area lead for sustainable sourcing, the idea of a "sustainable supply chain"

“The idea of a ‘sustainable supply chain’ is about every hand that touches that fruit, from tree to bottle”



is about every hand that touches the fruit, from tree to bottle, and about the land on which it grows. And it doesn't matter if that land is in Germany, Brazil, Peru, Poland, or Spain.

In Europe, companies are increasingly obliged to comply with corporate due diligence laws covering environmental, social, and human rights parameters. Eckes-Granini has set a goal to have 100% of its juice sustainably sourced by 2030. But keeping track of where the ingredients came from and the conditions under which they were produced and processed requires mountains of data.

Neufert turned to Microsoft for help to design a unique system, with a custom data model on the Microsoft Intelligent Data Platform using the AI-powered predictive insights the tool offers. It uses Azure Data Factory Flows to automate the process of inputting data from reporting agencies, sustainability ratings, suppliers, and their own records into a central database.

A customised Microsoft Power BI dashboard helps analyse the data. Power BI is a cloud-based platform that gives a visual overview of information. In this case, it's the risk levels associated with an individual country or all the data about one type of fruit. Even more critical is that it's easy to drill down into specifics about a particular vendor or supplier. For Neufert, the idea of "sustainable supply chain" benefits everyone involved – guaranteeing a steady supply of quality fruit for his company, as well as steady work for farm families and long-term stability and profits for all the players. On top of that, it's better for the environment. And it wouldn't be possible without big data.



## AI for smarter manufacturing

AI, Cloud and Big Data are already powering Industry 4.0, the next evolution of industry. Intelligent machines are already in place in smart factories, enabling manufacturing companies to be safer, more efficient, more productive, and more profitable.

AI can support simulations that prepare for real-world scenarios (digital twins), support predictive maintenance, and optimise processes to save energy. These smarter factories, in turn, can produce smarter devices that can send data back to manufacturers, who can update the software, improve the device, and design new features for the next iteration.

We are proud that European companies including [Siemens](#), [Mercedes-Benz](#), and [Bayer](#) have trusted Microsoft to support them in their own innovation journeys – using our AI, sustainability, and cloud offerings to digitally transform their organisations, while improving service and experience for their customers.



### What goes around... Germany

Factories are like organisms, and the heart is often a turbine or electric motor. Known in expert circles as rotating machines, typical examples include turbines in wind power generators, or electrical motors in car factories. It's not uncommon for these kinds of rotating machines to have an operational life of 10 to 20 years or more. If the *Mittelstand* is the backbone of Europe's economy, the rotating machine is the backbone of the *Mittelstand*.

What if the heart of the factory could talk? With a digital twin, it can. [Bosch](#) has developed an integrated asset performance management (IAPM) solution powered by a digital twin that

“AI can support simulations that prepare for real world scenarios”



runs on Microsoft Azure. This solution lets rotating machines indicate when they need maintenance, helping them to run with optimal costs and maximum efficiency. Repairing them when necessary, rather than as a precaution, plays a vital part in making the circular economy a reality.

“We know from our colleagues who work directly on these systems that the stress level rises when, for example, their machines are subjected to unusual loads,” says Bhuvan Shetty, Sales & Global Strategic Partnerships at Bosch. “Our digital twin enables them to make data-driven decisions to avoid breakdowns from the outset. As a digital assistant, it really lightens the day-to-day workload.” Bosch’s solution simplifies complex interrelationships, automates processes, boosts efficiency, and helps save resources.



## Rejuvenating a classic

Denmark

Denmark’s Novo Nordisk has been improving the lives of people with diabetes for almost a century and is the world’s largest manufacturer of insulin. The company’s production facilities are a whirlwind of activity, with lines constantly processing batches of medicine to serve more than 34 million diabetics in 170 countries.

Moving to new production formats and improving processes to meet demand requires careful planning and faultless implementation to ensure there are no production

“As a digital assistant, it really lightens the day-to-day workload”



interruptions. It’s essential that expert help is readily on tap for shop floor workers tasked with making those changes.

Instead of having to rely on being on-site, Novo Nordisk can now connect technicians virtually with the people they’re assisting through using Microsoft’s HoloLens 2 mixed reality technology, which leverages AI and Microsoft Dynamics 365. They can go through complex processes together with the help of real-time, 3D instructions. This helps to guarantee that every step has been completed successfully while allowing operators to visualise their full work environment.

“Depending on the specific use-case or instruction, you will be guided step by step,” Corporate Vice President of Product Supply Technology Transformation at Novo Nordisk, Ole Feddersen says. “It will say, ‘Loosen this bolt.’ ‘Turn this lever left.’ ‘Take this out and put a new one in.’ All the time there will be arrows pointing at what you need to do and where you need to go next, where to check, and so forth. You get a very intuitive execution throughout the operation while having both hands free.”

It’s helping engineers be more efficient with their time and support more colleagues, even those who may be thousands of miles away. And it’s helping Novo Nordisk scale its global operations. “The growth we’ve been seeing recently is requiring us to expand our manufacturing capacity, but it’s also a great opportunity to apply Industry 4.0 technology to some key parts of our work,” says Feddersen.



## AI for better education

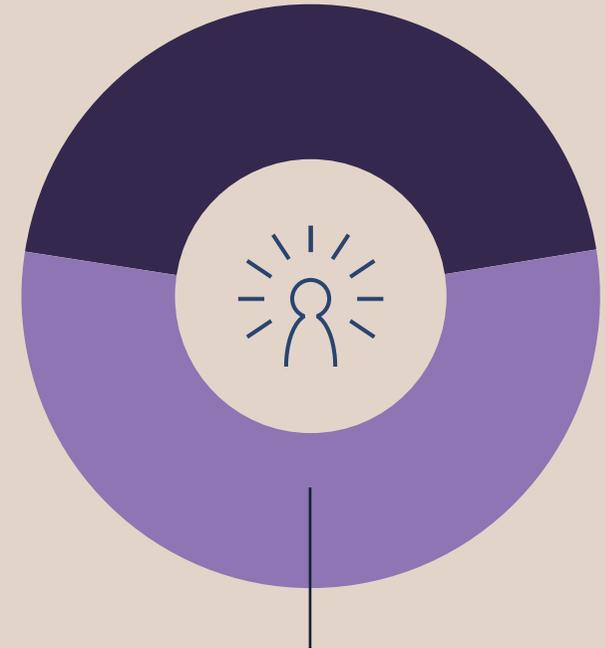
AI has immense potential to transform education and training for students, teachers, and school staff. It can help students with learning difficulties and support teachers in creating individualised learning. It could reduce school dropout rates and enable differentiated education programmes. It can help analyse vast swathes of data to see which approaches work best.

There is an urgent need to transform education: as the world is changing, so are careers. The next generation will be working closely with AI, yet there is a desperate shortage of digital skills. According to the latest DESI report, 55% of EU enterprises reported difficulties in filling ICT specialist vacancies.

UNESCO wants countries to harness the potential of AI technologies for learning “while ensuring that its application in educational contexts is guided by the core principles of inclusion and equity”. It notes in the Beijing Consensus, which sets out a common position on AI and learning, that the promise of “AI for all” must be that everyone can take advantage of the technological revolution under way and access its benefits, notably in terms of innovation and knowledge.

And for this to happen, there is a growing need for researchers, educators, and students to have a foundational understanding of AI and data use to be able to engage positively, critically, and responsibly with this technology

## Proportion of EU enterprises reporting difficulties filling ICT specialist vacancies



“55% of EU enterprises reported difficulties in filling ICT specialist vacancies”

Source: DESI



## Finns can only get better

Finland

In 2016, Helsinki was the first capital to have a digitalization strategy for education. “We are aiming for systemic change,” says Pasi Silander, Head of ICT Development Programmes in the City of Helsinki Education Division.

“It isn’t just about the technology or the operating culture or the pedagogy – it is all of those things plus the leadership strategy and the structure of the school. Everything has to be tackled all at once if the pieces are going to fit together and we are going to see a different picture.”

The team explored the ways in which AI and data analytics could enhance education.

The first step on the city’s journey with data and AI was to identify the pedagogical use-cases where technology can support their “well-learning” vision. AI and learning analytics have enabled teachers to personalise programmes.

“Teachers have maybe 25 to 30 students in a classroom and all the students are individuals with individual needs,” says Silander. “Even though our teachers are Super Ladies and Super Men, they can’t really personalise each of the student’s learning processes.” By collecting and analysing data, they can make better decisions for students’ future while prioritising responsible AI considerations specific to education.

“By analysing data, they can make better decisions for students”



## Tackling the digital skills gap

France

Europe faces a dramatic shortfall of ICT professionals, as previously noted. Furthermore, there is an imperative for greater diversity: 81% of ICT specialists working in the EU are men. To help address this issue, Microsoft partnered with Simplon, a social enterprise focused on digital training that has reached 20,000 learners from diverse backgrounds across several countries; about 40% of whom are women.

Participants in the Microsoft AI School by Simplon spend four to seven months in the classroom, seven hours per day. Then they spend a year in an apprenticeship at a company, with one week in class and three weeks on the job. Companies select candidates for apprenticeships from Simplon from the start of the programme. In this way, it supports companies in a non-traditional recruitment process – producing strong recruits outside of the normal channels, such as level of education, professional background, and previous experience.

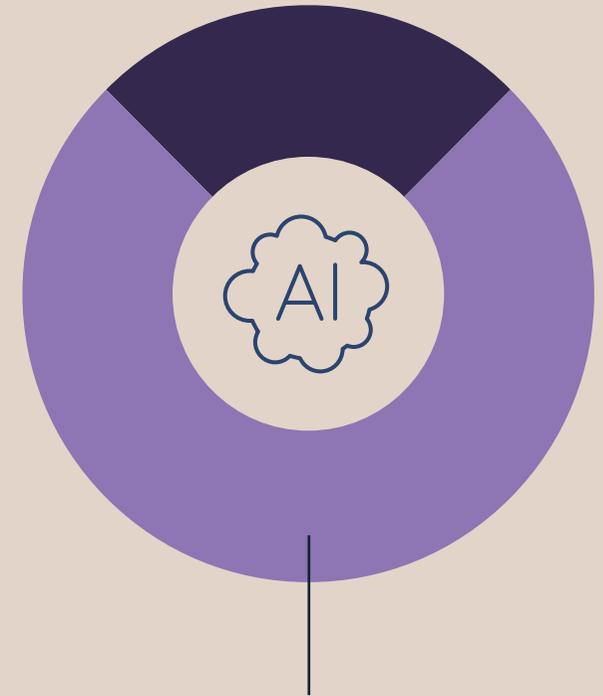
Microsoft and Simplon have collaborated on another training programme that focuses on cybersecurity. The Cybersecurity School is at the heart of Microsoft’s Cybersecurity Skills Plan in France, which aims to support the training of 10,000 new cybersecurity professionals by the end of 2025. This is part of Microsoft’s broader cybersecurity skills initiative in Europe and recent pledge to train 100,000 learners in Europe with cybersecurity skills over a two-year period.

# The value of AI

The scale of value that AI could help European organisations and entrepreneurs create for European economies and citizens is truly exciting. According to McKinsey, it could add €2.7 trillion, or 20%, to the region's economic output by 2030. If Europe can accelerate the development of its AI competencies and digital innovation capacity, that figure could rise by a further €900 billion.

Key to this will be fostering a thriving 'AI ecosystem' and ensuring AI is an affordable and accessible platform for European organisations of all kinds and sizes. By 2030, the EU wants 75% of businesses across the block to be using AI, cloud, and big data solutions. The true value of AI will only be realised through diffusing its use throughout the economy, and through the development of AI applications.

## 'AI ecosystem' goals



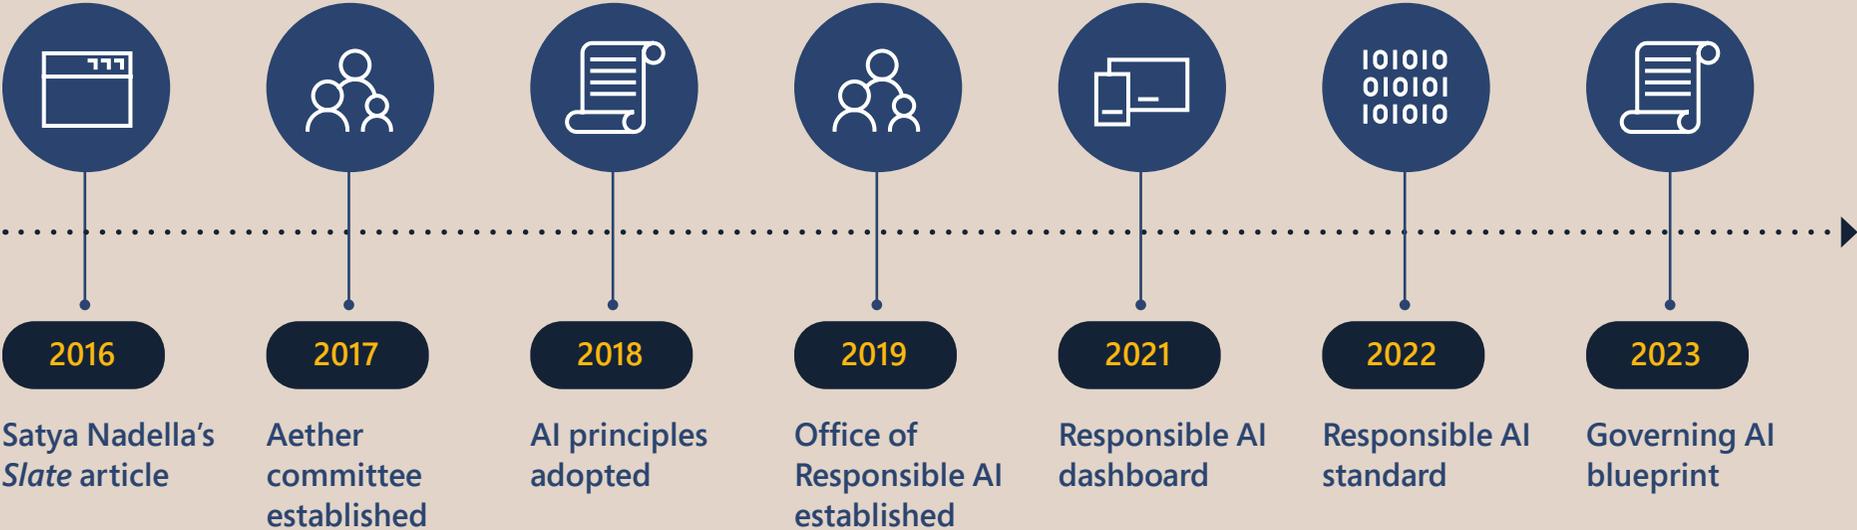
“By 2030, the EU wants 75% of businesses across the block to be using AI, cloud, and big data solutions”

Source: EU



# Responsible development and use of AI

# Our responsible AI journey



# Responsible AI – Q&A with Natasha Crampton



Natasha Crampton leads Microsoft's Office of Responsible AI, as the company's first Chief Responsible AI Officer. The Office of Responsible AI puts Microsoft's AI principles into practice by defining, enabling, and governing the company's approach to responsible AI. The Office of Responsible AI also collaborates with stakeholders within and outside the company to shape new laws, norms, and standards to help ensure that the promise of AI technology is realised for the benefit of all.

## Why is responsible AI so important?

AI offers tremendous opportunities to improve our world – stimulating economic growth, making work more enjoyable and helping address pressing societal issues like sustainability. However, like every technology before it, some people, communities and countries will turn this advancement into both a tool and a weapon. Some people will use this technology to exploit the flaws in human nature, deliberately target people with false information, undermine democracy and explore new ways of pursuing criminal endeavours.

This creates a profound sense of responsibility, both in terms of AI development and its deployment. Since 2017, Microsoft has invested in a cross-company programme to ensure that our AI systems are responsible by design.

## How does Microsoft approach responsible AI?

For Microsoft, responsible AI development is about both a practice and a culture.

Practice is how we formally operationalise responsible AI across the company, through governance processes, policy requirements, and tools and training to support implementation. Culture is how we empower our employees to actively champion responsible AI.

We see three essential tenants of upholding the right culture and practices:

### 1. Leadership from the top.

Microsoft's responsible AI journey began in 2016 when Microsoft's Chairman and CEO, Satya Nadella, penned an [article](#) articulating his vision for humanity powered

by AI. Satya's vision set in motion the beginnings of our core AI principles. Microsoft's responsible AI efforts continue to be championed by Microsoft's senior leadership team today. Microsoft's Responsible AI Council is chaired by Microsoft's Vice Chair and President, Brad Smith and by Chief Technology Officer Kevin Scott and convenes senior staff from research, policy, and engineering teams who lead responsible AI efforts across the company. The council serves as a forum for leadership alignment and oversees the progress we are making on our responsible AI commitments.

### 2. Inclusive governance models.

Since starting our work on responsible AI seven years ago, we learned the need to create a governance model that was inclusive and encouraged multi-disciplinary collaboration across engineering, policy, and research. The governance model established by Microsoft's Office of Responsible AI ensures that there are senior leaders tasked with spearheading responsible AI across every business group. These leaders are supported by a large network of trained "responsible AI champions". Our continually updated Responsible AI Standard, which is publicly [available](#), offers this community guidance for building AI systems responsibly. Taking lessons from long-standing, cross-company

commitments to privacy, security, and accessibility, we realised that responsible AI must be supported by the highest levels of leadership in the company and championed at every level across Microsoft. Responsible AI is a company-wide mandate.

### 3. Invest in people:

Microsoft currently has nearly 350 people working on responsible AI, with 129 dedicated to it full-time. These team members hold positions across the company in policy, engineering, research, sales, and other core functions. We are committed to continue investing in hiring diverse talent, assigning additional talent to focus on responsible AI full time, and upskilling more people throughout the company.

### How does the Responsible AI Standard help teams put Microsoft's principles into practice?

From crafting an AI system's purpose to designing how people interact with it, we must keep people at the centre of all AI decisions. While our responsible AI principles state the enduring values we seek to uphold, we needed more specific guidance on how to build and deploy AI systems responsibly. This is why we developed our Responsible AI Standard, a more practical guide that memorialises a set of "rules of the road" for our engineering teams so that upholding our AI principles is a daily practice.

## Microsoft's responsible AI principles



### Fairness:

AI systems should treat all people fairly



### Reliability & safety:

AI systems should perform reliably and safely



### Privacy & security:

AI systems should be secure and respect privacy



### Inclusiveness:

AI systems should empower everyone and engage people



### Transparency:

AI systems should be understandable



### Accountability:

People should be accountable for AI systems

The Responsible AI Standard was the result of a multi-year, cross-company effort that reflected a vast array of input from researchers, engineers, lawyers, designers, and policy experts. We consider it to be a significant step forward for our practice of responsible AI because it sets out concrete, practical guidance on how to identify, measure, and mitigate harms ahead of time. It also requires teams to adopt tools and controls to secure beneficial uses while guarding against potential misuses of their products.

When building and updating the Responsible AI Standard, we recognised early on that it is impossible to reduce all the complex sociotechnical considerations – for many different use-cases – into an exhaustive set of pre-defined rules. This led us to create a programme and process for ongoing review and oversight of high-impact cases and rising issues and questions, which we call Sensitive Uses. The Sensitive Uses review process is triggered when Microsoft personnel are involved in developing or deploying an AI system and the foreseeable use or misuse of that AI system could have a consequential impact on a user’s legal status or life opportunities, present the risk of significant physical or psychological injury, or restrict, infringe upon, or undermine the ability to realise an individual’s human rights.

Our Sensitive Uses programme provides a critically important additional layer of oversight for teams working on higher-risk use-cases of our AI systems.

## The Rome Call for AI Ethics

The Rome Call for AI Ethics is an important global effort to promote an ethical approach to artificial intelligence. Microsoft has been involved in this initiative since 2020, which stands for three key principles:



Importance of ethical safeguards to ensure AI technology is transparent, inclusive, and governed with human interests in mind.



Shared responsibility among tech companies and developers to create and implement AI with a broad perspective on ensuring a positive societal impact.



The importance of AI advancing the global community’s sustainability efforts.

# Helping our European customers develop and deploy AI responsibly – Antony Cook



Antony Cook leads Microsoft Customer and Partner Solutions group, a team comprised of more than 350 legal and government affairs professionals located in more than 50 countries around the world. They focus on all legal support for the global sales organisation including pioneering new legal and corporate affairs solutions that advance business goals to benefit consumers, industry partners, and communities.

Microsoft has a responsible AI process that our AI systems go through, which includes reviews by a multidisciplinary team of experts, that helps us understand potential harms and find mitigations. 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, or applying technology that can return more helpful, representative, and diverse responses and results.

Another very important safeguard is intentional and iterative deployment. We take a measured approach to rolling out many of these new AI-powered experiences. We may start with a limited preview with a select number of customers with well-defined use-cases in

mind. Collaborations with these early customers will help us make sure the responsible AI safeguards are working in practice so we can 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 the 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 operationalise responsible AI practices. Over the years, responsible AI research at Microsoft has led to the incubation of tools such as Fairlearn and InterpretML. The 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 in preview, is Azure Content Safety. It helps businesses create safer online environments and communities through models that are designed to detect hate, violent, sexual, and self-harm content across languages in both images and text.

Building on the Responsible AI Toolbox, Microsoft's responsible AI programme has invested in integrating some of the more mature responsible AI tools directly into Azure Machine Learning so our customers will also benefit from the development of engineering systems and tools. The 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 customise 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 programme, 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.

The bottom line is that every organisation 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 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 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 have [announced](#) that we are creating an AI Assurance Program to help customers ensure that the AI applications they deploy on our

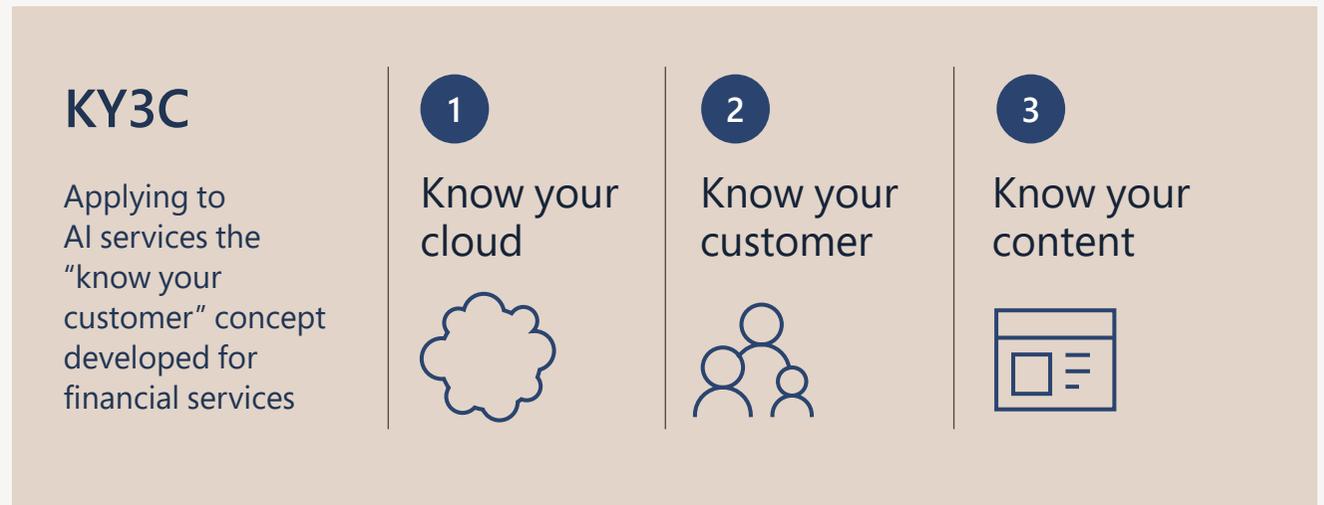
platforms meet the legal and regulatory requirements for responsible AI. This programme is open to our customers around the world.

## This programme will include:

**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, 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 will bring 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. The recently launched 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 implement their own AI systems responsibly, and we will develop responsible AI programmes for our partner ecosystem.

Many of our partners have already created comprehensive practices to help customers evaluate, test, adopt, and commercialise AI solutions, including creating their own responsible AI systems. We are launching a programme 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 will be looking to add European partners to the programme.

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

## Advancing trustworthy AI

For more than 30 years, [Microsoft Research](#) has been advancing the foundations of computing and translating new scientific understanding into innovative technologies to create value for our customers and broad benefit to society.

Our researchers collaborate across disciplines, institutions, and geographies to deliver cutting-edge advances in vision, speech, language, decision-making, and machine learning. They have pioneered AI breakthroughs in [conversational speech recognition](#), [machine translation](#), [image captioning](#), [natural language understanding](#), and [commonsense question answering](#).

As part of our ongoing commitment to [building AI responsibly](#), research scientists and engineers at Microsoft are also pursuing methods and technologies aimed at helping builders of AI systems cultivate appropriate trust – that is, building trustworthy models with reliable behaviours and clear communication that set proper expectations.

When AI builders plan for failures, work to understand the nature of the failures, and implement ways to effectively mitigate potential harms, they help engender trust that can lead to a greater realization of AI's benefits.

Facilitating trustworthy measurement, improving human-AI collaboration, designing for natural language processing (NLP), advancing transparency and interpretability, and exploring the open questions around AI safety, security, and privacy are key considerations for developing AI responsibly.

The goal of trustworthy AI requires a shift in perspective at every stage of the AI development and deployment life cycle. We're actively developing a growing number of [best practices and tools](#) to help with the shift to make responsible AI more available to a broader base of users. Many open questions remain, but as innovators, we are committed to tackling these challenges with curiosity, enthusiasm, and humility.



# AI governance



As Microsoft's Vice Chair and President, Brad Smith is responsible for spearheading the company's work and representing it publicly on a wide variety of critical issues involving the intersection of technology and society, including artificial intelligence, cybersecurity, privacy, environmental sustainability, human rights, digital safety, immigration, philanthropy, and products and business for non-profit customers. He leads a team of roughly 2,000 business, legal, and corporate affairs professionals located in 54 countries and operating in more than 120 nations.

# Advancing AI governance in Europe and internationally – Brad Smith

As this paper highlights, there are enormous opportunities to harness the power of AI to contribute to European growth and values. But another dimension is equally clear. It's not enough to focus only on the many opportunities to use AI to improve people's lives. We need to focus with equal determination on the challenges and risks that AI can create, and we need to manage them effectively.

This is perhaps one of the most important lessons from the role of social media. Little more than a decade ago, technologists and political commentators alike

“We are committed and determined as a company to develop and deploy AI in a safe and responsible way”

gushed about the role of social media in spreading democracy during the Arab Spring. Yet, five years after that, we learned that social media, like so many other technologies before it, would become both a weapon and a tool – in this case aimed at democracy itself.

Today, we are 10 years older and wiser, and we need to put that wisdom to work. We need to think early on and in a clear-eyed way about the problems that could lie ahead. As technology moves forward, it's just as important to ensure proper control over AI as it is to pursue its benefits. We are committed and determined as a company to develop and deploy AI in a safe and responsible way. We also recognise, however, that the guardrails needed for AI require a broadly shared sense of responsibility and should not be left to technology companies alone. In short, tech companies will need to step up, and governments will need to move faster.

When we at Microsoft adopted our six ethical principles for AI in 2018, we noted that one principle was the bedrock for everything else – accountability. This is the fundamental need to ensure that machines remain subject to effective oversight by people, and the people who design and operate machines remain accountable to everyone else. In short, we must always

ensure that AI remains under human control. This must be a first-order priority for technology companies and governments alike.

This connects directly with another essential concept. In a democratic society, one of our foundational principles is that no person is above the law. No government is above the law. No company is above the law, and no product or technology should be above the law. This leads to a critical conclusion: People who design and operate AI systems cannot be accountable unless their decisions and actions are subject to the rule of law.

In May, Microsoft released a whitepaper, *Governing AI: A Blueprint for the Future*, which sought to address

“Europe is now at the forefront of establishing a model for guiding and regulating AI technology”

the question of how do we best govern AI and set out Microsoft’s [five-point blueprint](#).

The blueprint builds on lessons learned from many years of work, investment, and input. Here we draw some of that thinking together with a specific focus on Europe, its leadership on AI regulation, and what a viable path to advancing AI governance internationally may look like.

With the European Parliament’s recent vote on the European Union’s Artificial Intelligence (AI) Act and ongoing trilogue discussions, Europe is now at the forefront of establishing a model for guiding and regulating AI technology.

From early on, we’ve been supportive of a regulatory regime in Europe that effectively addresses safety and upholds fundamental rights while continuing to enable innovations that will ensure that Europe remains globally competitive. Our intention is to offer constructive contributions to help inform the work ahead. Collaboration with leaders and policymakers across Europe and around the world is both important and essential.

In this spirit, here we want to expand upon our five-point blueprint, highlight how it aligns with EU AI Act discussions, and provide some thoughts on the opportunities to build on this regulatory foundation.



## A five-point blueprint for governing AI



1

Implement and build upon new government-led AI safety frameworks



2

Require safety brakes for AI systems that control critical infrastructure



3

Develop a broader legal and regulatory framework based on the technology architecture for AI



4

Promote transparency and ensure academic and public access to AI



5

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



## Implement and build upon new government-led AI safety frameworks

One of the most effective ways to accelerate government action is to build on existing or emerging governmental frameworks to advance AI safety.

A key element to ensuring the safer use of this technology is a risk-based approach, with defined processes around risk identification and mitigation as well as testing systems before deployment. The AI Act sets out such a framework and this will be an important benchmark for the future. In other parts of the world, other institutions have advanced similar work, such as the AI Risk Management Framework developed by the U.S. National Institute of Standards and Technology, or NIST, and the new international standard ISO/IEC 42001 on AI Management Systems, which is expected to be published in the fall of 2023.

Microsoft has committed to implementing the NIST AI risk management framework and we will implement future relevant international standards, including those which will emerge following the AI Act. Opportunities to align such frameworks internationally should

continue to be an important part of the ongoing EU-US dialogue.

As the EU finalises the AI Act, the EU could consider using procurement rules to promote the use of relevant trustworthy AI frameworks. For instance, when procuring high-risk AI systems, EU procurement authorities could require suppliers to certify via third-party audits that they comply with relevant international standards.

The pace of AI advances raises new questions and issues related to safety and security, and we are committed to working with others to develop actionable standards to help evaluate and address those important questions. This includes new and additional standards relating to highly capable foundation models.

“...a comprehensive approach to system safety that would keep effective human oversight...”



## Require effective safety brakes for AI systems that control critical infrastructure

Increasingly, the public is debating questions around the control of AI as it becomes more powerful. Similarly, concerns exist regarding AI control of critical infrastructure like the electrical grid, water system, and city traffic flows. Now is the time to discuss these issues – a debate already underway in Europe, where developers of high-risk AI systems will become responsible for warranting that the necessary safety measures are in place.

Our blueprint proposes new safety requirements that, in effect, would create safety brakes for AI systems that control the operation of designated critical infrastructure. These fail-safe systems would be part of a comprehensive approach to system safety that would keep effective human oversight, resilience, and robustness top of mind.

They would be akin to the braking systems engineers have long built into other tech such as elevators, school buses, and high-speed trains, to safely manage not just everyday scenarios, but emergencies as well.

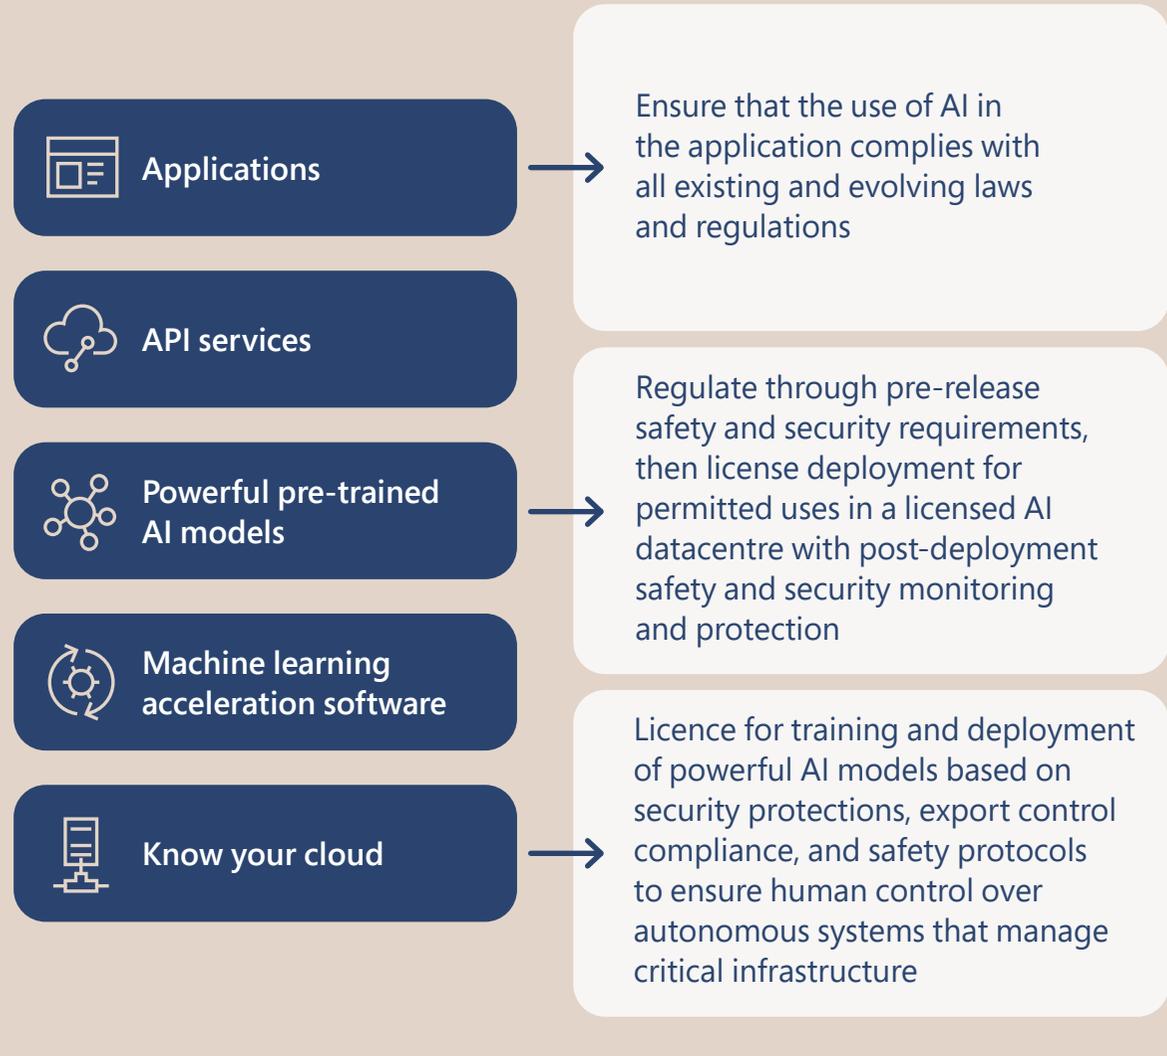
In this approach, the government would define the class of high-risk AI systems that control critical infrastructure and warrant such safety measures as part of a comprehensive approach to system management. New laws would require operators of these systems to build safety brakes into high-risk AI systems by design. The government would then oblige operators to test high-risk systems regularly. And these systems would be deployed only in licensed AI datacentres that would provide a second layer of protection and ensure security.



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

As we've worked the past year with AI models at the frontier of this new technology, we've concluded that it's critical to develop a legal and regulatory architecture for AI that reflects the technology architecture for AI itself. Regulatory responsibilities need to be placed upon different actors based on their role in managing different aspects of AI technology. Those closest to relevant decisions on design, deployment, and use are best placed to comply

## A proposed AI regulatory architecture



with corresponding responsibilities and mitigate the respective risks, as they understand best the specific context and use-case. This sounds straightforward, but as discussions in the EU have demonstrated, it's not always easy.

The AI Act acknowledges the challenges to regulating complex architecture through its risk-based approach for establishing requirements for high-risk systems. At the application layer, this means applying and enforcing existing regulations while being responsible for any new AI-specific deployment or use considerations.

It's also important to make sure that obligations are attached to powerful AI models, with a focus on a defined class of highly capable foundation models and calibrated to model-level risk. This will impact two layers of the technology stack. The first will require new regulations for these models themselves. And the second will involve obligations for the AI infrastructure operators on which these models are developed and deployed. The blueprint we developed offers suggested goals for each of these layers.

The different roles and responsibilities require joint support. We are committed to helping our customers apply "Know Your Customer" (KYC) principles through our recently announced AI Assurance Program. Financial institutions use this framework to verify customer identities, establish risk profiles and monitor transactions to help detect suspicious activity.

As Antony Cook [details](#), we believe this approach can apply to AI in what we are calling "KY3C": know one's cloud, one's customers and one's content. The AI Act does not expressly include KYC requirements, yet we believe such an approach will be key to meeting both the spirit and obligations of the Act.

The AI Act requires developers of high-risk systems to put in place a risk management system to ensure that systems are tested, to mitigate risks to the extent possible, including through responsible design and development, and to engage in post-market monitoring. We fully support this. The fact that the AI Act hasn't yet been finalised shouldn't prevent us from applying some of these practices today. Even before the AI Act is implemented, we will test all our AI systems prior to release, and use red teaming to do so for high-risk systems.

**"Even before the AI Act is implemented, we will test all our AI systems prior to release"**



## Promote transparency and ensure academic and nonprofit access to AI

It's also critical to advance the transparency of AI systems and broaden access to AI resources. While there are some inherent tensions between transparency and the need for security, there exist many opportunities to make AI systems more transparent. That's why Microsoft has committed to an annual AI transparency report and other steps to expand transparency for our AI services.

The AI Act will require that AI providers make it clear to users that they are interacting with an AI system. Similarly, whenever an AI system is used to create artificially generated content, this should be easy to identify. We fully support this. Deceptive, AI-generated content or "deepfakes" – especially audio-visual content impersonating political figures – are of particular concern in terms of their potential harm to society and the democratic process.

In tackling this issue, we can start with building blocks that exist already. One of these is the Coalition for Content Provenance Authenticity, or C2PA, a global

standards body with more than 60 members including Adobe, the BBC, Intel, Microsoft, Publicis Groupe Sony, and Truepic. The group is dedicated to bolstering trust and transparency of online information including releasing the world's first technical specification for certifying digital content in 2022, which now includes support for Generative AI.

As Microsoft's Chief Scientific Officer, Eric Horvitz, said last year, "I believe that content provenance will have an important role to play in fostering transparency and fortifying trust in what we see and hear online."

There will be opportunities in the coming months to take important steps together on both sides of the Atlantic and globally to advance these objectives. Microsoft will soon deploy new state-of-the-art provenance tools to help the public identify AI-generated audio-visual content and understand its origin. At Build 2023, our annual developer conference, we announced the development of a new media provenance service.

The service, implementing the C2PA specification, will mark and sign AI-generated videos and images with metadata about their origin, enabling users to verify that a piece of content is AI-generated. Microsoft will initially support major image and video formats and release the service for use with two of Microsoft's new AI products, Microsoft Designer and Bing Image Creator.

Transparency requirements in the AI Act, and potentially several of the to-be-developed standards related to the Act, present an opportunity to leverage such industry initiatives towards a shared goal.

We also believe it is critical to expand access to AI resources for academic research and the nonprofit community. Unless academic researchers can obtain access to substantially more computing resources, there is a real risk that scientific and technological inquiry will suffer, including that relating to AI itself. Our blueprint calls for new steps, including those we will take across Microsoft, to address these priorities.



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

One lesson from recent years is that democratic societies often can accomplish the most when they harness the power of technology and bring the public and private sectors together. It's a lesson we need to build upon to address the impact of AI on society.

“Democratic societies often can accomplish the most when they harness the power of technology”

AI is an extraordinary tool. But, like other technologies, it too can become a powerful weapon, and there will be some around the world who will seek to use it that way. We need to work together to develop defensive AI technologies that will create a shield that can withstand and defeat the actions of any bad actor on the planet.

Important work is needed now 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. Perhaps more than anything, a wave of new AI technology provides an occasion for thinking big and acting boldly. In each area, the key to success will be to develop concrete initiatives and bring governments, companies, and NGOs together to advance them. Microsoft will do its part in each area.

# International partnership to advance AI governance

Europe's early start towards AI regulation offers an opportunity to establish an effective legal framework, grounded in the rule of law. But beyond legislative frameworks at the level of nation states, multilateral public-private partnership is needed to ensure AI governance can have an impact today, not just a few years from now, and at the international level.

This is important to serve as an interim solution before regulations such as the AI Act come into effect, but, perhaps more importantly, it will help us work towards a common set of shared principles that can guide both nation states and companies alike.

In parallel to the EU's focus on the AI Act, there is an opportunity for the European Union, the United States, the other members of the G7 as well as India and Indonesia, to move forward together on a set of shared values and principles. If we can work with others on a voluntary basis, then we'll all move faster and with greater care and focus. That's not just good news for the world of technology, it is good news for the world as a whole.

Working towards a globally coherent approach is important, recognising that AI – like many technologies – is and will be developed and used

across borders. And it will enable everyone, with the proper controls in place, to access the best tools and solutions for their needs.

We are very encouraged by recent international steps, including by the EU-US Trade & Tech Council (TTC) at the end of May, announcing a new initiative to develop a voluntary AI Code of Conduct. This can bring together private and public partners to

“Multilateral public-private partnership is needed to ensure AI governance can have an impact today, not just a few years from now”

implement non-binding international standards on AI transparency, risk management, and other technical requirements for firms developing AI systems.

Similarly, at the annual G7 Summit in Hiroshima in May 2023, leaders committed to “advance international discussions on inclusive artificial intelligence (AI) governance and interoperability to achieve our common vision and goal of trustworthy AI, in line with our shared democratic values.”

Microsoft fully supports and endorses international efforts to develop such a voluntary code. Technology development and the public interest will benefit from the creation of principle-level guardrails, even if initially they are non-binding.

To make the many different aspects of AI governance work on an international level, we will need a multilateral framework that connects various national rules and ensures that an AI system certified as safe in one jurisdiction can also qualify as safe in another. There are many effective precedents for this, such as common safety standards set by the International Civil Aviation Organization, which means an airplane does not need to be refitted midflight from Brussels to New York.



## In our view, an international code should:



Build on the work already done at the OECD to develop principles for trustworthy AI;



Provide a means for regulated AI developers to attest to the safety of these systems against internationally agreed standards;



Promote innovation and access by providing a means for mutual recognition of compliance and safety across borders.

Before the AI Act and other formal regulations come into force, it's important that we take steps today to implement safety brakes for AI systems that control critical infrastructure. The concept of safety brakes, along with licensing for highly capable foundation models and AI infrastructure obligations, should be key elements of the voluntary, internationally coordinated G7 code that signatory nation states agree to incorporate into their national systems. High-risk AI systems, relating to critical infrastructure (e.g., transport, electrical grids, water systems) or systems that can lead to serious fundamental rights violations or other significant harms, could require additional international regulatory agencies, based on the model of the International Civil Aviation Organization, for example.

The AI Act creates a provision for an EU database for high-risk systems. We believe it's an important approach, and it should be a consideration for the international code. Developing a coherent, joint, global approach makes immeasurable sense for all those involved in developing, using and regulating AI.

Lastly, we must ensure academic researchers have access to study AI systems in depth. There are important open research questions around AI systems, including how one evaluates them properly across responsible AI dimensions, how one best makes them explainable, and how they best align with human

“AI governance is a journey, not a destination. No one has all the answers, and it's important we listen, learn and collaborate”

values. The work the OECD is doing on the evaluation of AI systems is seeing good progress. But there's an opportunity to go further and faster by fostering international research collaboration and boosting the efforts of the academic communities by feeding into that process. The EU is well placed to lead on this, partnering with the US.

AI governance is a journey, not a destination. No one has all the answers, and it's important we listen, learn and collaborate. Strong, healthy dialogue between the tech industry, governments, businesses, academia and civil society organisations is vital to make sure governance keeps pace with the speed at which AI is developing. Together, we can help realise AI's potential to be a positive force for good.

# Sources

## The promise of AI

### Cited sources:

[Atomico, June 2022, State of European Tech 2022](#)

[International Data Corporation \(IDC\), March 2023, Worldwide Artificial Intelligence Spending Guide \(V1 2023\)](#)

### References:

[Novartis: Microsoft, November 2021, Novartis empowers scientists with AI to speed the discovery and development of breakthrough medicines](#)

[Danfoss & Repsol: Microsoft, Sustainability. Good for Business. 2021 and beyond](#)

[Vestas: Microsoft, March 2022, How one of the world's largest wind companies is using AI to capture more energy](#)

[Safeguards for facial recognition: Microsoft, June 2022, Responsible AI investments and safeguards for facial recognition](#)

[Governing AI: A Blueprint for the Future whitepaper: Microsoft, May 2023, Governing AI: A Blueprint for the Future](#)

### General sources:

[Natasha Crampton, Microsoft, May 2023, Reflecting on our responsible AI program: Three critical elements for progress](#)

[John Frank, Microsoft, September 2019, The stakes are too high to let facial recognition regulation slide](#)

[Microsoft, July 2021, Mind the digital skills gap: A closer look at the global skills initiative in Europe](#)

[Sylvie Laffarge, Microsoft, March 2021, Boosting digital skills for all: An update on the European impact of our global skills initiative](#)

[Microsoft](#)

## AI explained

### References:

[GitHub Copilot: GitHub](#)

[Copilot report: GitHub, September 2022, Research: quantifying GitHub Copilot's impact on developer productivity and happiness](#)

[Power Platform: Microsoft](#)

[Microsoft Designer: Microsoft](#)

[Bing and Edge: Yusuf Mehdi, Microsoft, February 2023, Reinventing search with a new AI-powered Microsoft Bing and Edge, your copilot for the web](#)

[Azure AI supercomputer: Microsoft, May 2020, Microsoft announces new supercomputer, lays out vision for future AI work](#)

### General sources:

[Microsoft, February 2023, What is Microsoft's approach to AI?](#)

[Microsoft, March 2023, Explaining AI: The fundamentals and the frontiers](#)

[Microsoft, A conversation with Kevin Scott: What's next in AI](#)

[Microsoft, February 2022, The innovation behind AI at Scale](#)

## AI in action

### Cited sources:

[McKinsey & Company, February 2019, Tackling Europe's gap in digital and AI](#)

[McKinsey & Company and EIT Health, March 2020, Transforming healthcare with AI](#)

[Eurostat, February 2023, Government expenditure on health](#)

[World Health Organization, March 2023, The health workforce crisis in Europe is no longer a looming threat – it is here and now. The Bucharest Declaration charts a way forward](#)

[World Meteorological Organization, May 2023, Global temperatures set to reach new records in next five years](#)

[OECD, June 2023, OECD-FAO Agricultural Outlook to 2031](#)

[European Commission, The Digital Economy and Society Index \(DESI\)](#)

[UNESCO, Artificial intelligence in education](#)

[Eurostat, May 2023, ICT specialist in employment](#)

### References:

[Emergency Medical Services \(EMS\): Microsoft, August 2021, Capital Region of Denmark: Using AI to improve patient care on the national healthcare system](#)

[Cova da Beira Hospital: Microsoft, June 2020, How technology is revolutionizing care for patients with chronic conditions](#)

[Oxford University Hospitals NHS Foundation Trust \(OUH\): Nuance, Oxford University Hospitals boost digitisation and accelerate uptake of the EPR](#)

[Covid Moonshot project: DNDi, February 2023, COVID Moonshot](#)

[UCB: UCB, February 2021, UCB and Microsoft Expand Collaboration to Accelerate Drug Discovery and Development](#)

[Our Future Health: Microsoft, September 2022, UK's largest health research program is aiming to transform the prevention, detection and treatment of some of the most common diseases – and the Microsoft Cloud will help](#)

[Metroscope: Microsoft, May 2022, Metroscope enables energy producers to improve efficiency with digital twin solution supported by Azure Kubernetes Service \(AKS\) and Microsoft Security solutions](#)

[Solytic: Microsoft, June 2021, Driving the energy transition with smart monitoring: How start-up Solytic is making solar energy more efficient with Azure](#)

[Vestas: Microsoft, March 2022, How one of the world's largest wind companies is using AI to capture more energy](#)

[Water usage: Microsoft, March 2023, How AI is helping farmers keep up with sustainable food production](#)

[Project FarmVibes.AI: Microsoft, October 2022, Extracting intelligence from farm data and remote sensing sources](#)

[Bayer AgPowered Services: Microsoft](#)

[Eckes-Granini: Microsoft, April 2022, From orange farmer in Brazil to juice maker in Germany, making sustainability a winning game](#)

[Siemens: Microsoft, April 2023, Siemens and Microsoft drive industrial productivity with generative artificial intelligence](#)

[Mercedes-Benz: Microsoft, October 2022, Mercedes-Benz and Microsoft collaborate to boost efficiency, resilience and sustainability in car production](#)

[Bayer: Bayer, Scaling Sustainability Through Bayer & Microsoft Partnership](#)

[Bosch: Microsoft, May 2022, Machines raise their voice when they need help: Bosch's Digital Twin – IAPM takes human-machine collaboration thanks to Microsoft Azure to the next level](#)

[Novo Nordisk: Microsoft, June 2022, HoloLens 2 helps Novo Nordisk employees see work in new ways](#)

[City of Helsinki: Microsoft, March 2022, Empowering Learning and well-being in the City of Helsinki using data and AI](#)

[Microsoft AI School by Simplon: Microsoft, March 2023, AI training program helps close gender gap, bringing fresh faces to Europe's digital workforce](#)

### General sources:

[Judson Althoff, Microsoft, April 2023, The era of AI: How the Microsoft Cloud is accelerating AI transformation across industries](#)

[Simplon, Parcours Opérateur-trice de solutions de sécurité Cloud et Hybride](#)

[Microsoft, March 2022, The urgency of tackling Europe's cybersecurity skills shortage](#)

[Microsoft, May 2022, Microsoft lance son Plan Compétences Cybersécurité pour former 10 000 nouveaux professionnels en France en 3 ans](#)

[Microsoft, April 2023, Expanding Europe's cybersecurity talent pool](#)

[Microsoft, April 2023, Microsoft pledges support to EU Cybersecurity Skills Academy](#)

## Responsible development and use of AI

### References:

[Article by Satya Nadella: Satya Nadella, Microsoft, June 2016, The Partnership of the Future](#)

[Responsible AI Standard: Microsoft](#)

[Microsoft Research: Microsoft](#)

[Conversational speech recognition: Microsoft, August 2017, Microsoft researchers achieve new conversational speech recognition milestone](#)

[Machine translation: Microsoft, June 2019, Neural Machine Translation Enabling Human Parity Innovations In the Cloud](#)

[Image captioning](#)

[Natural language understanding: Microsoft, October 2020, Novel object captioning surpasses human performance on benchmarks](#)

[Commonsense question answering: Microsoft, December 2021, Azure AI milestone: Microsoft KEAR surpasses human performance on CommonsenseQA benchmark](#)

### General sources:

[Natasha Crampton, Microsoft, May 2023, Reflecting on our responsible AI program: Three critical elements for progress](#)

[New York Times, May 2023, Microsoft call for A.I. rules to Minimize the Technology's Risks](#)

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

[Microsoft, May 2023, Governing AI: A Blueprint for the Future](#)

[Brad Smith, Microsoft, January 2023, Putting AI in the service of the common good](#)

[Antony Cook, Microsoft, June 2023, Announcing Microsoft's AI Customer Commitments](#)

[Microsoft, February 2022, Advancing AI trustworthiness: Updates on responsible AI research](#)

[Eric Horvitz, Microsoft, December 2021, Building AI responsibly from research to practice](#)

## AI governance

### References:

[Governing AI: A Blueprint for the Future whitepaper: Microsoft, May 2023, Governing AI: A Blueprint for the Future](#)

### General sources:

[Brad Smith, Microsoft, May 2023, How do we best govern AI?](#)

[Eric Horvitz, Microsoft, January 2022, A Milestone Reached in Digital Content Provenance](#)

