A CyNam and HCR Law Insight Report

AI in practice

From foundational principles to future horizons

CYNAM

hcrlaw

Contents

The law firm with a passion for people

About this report

This report is the culmination of the Everyday AI programme, a series of forums co-hosted by CyNam and HCR Law across 2025 and early 2026. It draws on evidence from those forums, the CyNam and HCR Law Everyday AI Survey 2026 and UK research on artificial intelligence.

CyNam is the cyber security and emerging technology ecosystem anchored in Cheltenham. As a not-for-profit Community Interest Company, it sits between government, industry, academia and the public sector in a neutral position to provide honest insight and evidence-based analysis.

HCR Law is one of the UK's leading independent law firms, with deep experience in technology law, data protection, employment, intellectual property and the regulatory landscape. Its partnership with CyNam reflects a shared conviction that, to grow and succeed, organisations need practical insights grounded in expert evidence.

The forums brought together healthcare professionals, teachers, financial services practitioners, farmers, social care providers and professionals from technology and professional services. This report does not claim AI will solve every problem, nor that it poses existential risk. It tries to tell you what is known, what is uncertain and what your organisation should be thinking about now.

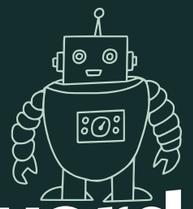

Foreword

When I agreed to chair the Everyday AI roundtables that form the foundation of this report, I expected to find organisations at different stages of their AI journey. Indeed we found this, with enthusiastic early adopters already embedding AI across their operations, to those approaching the technology with measured caution and some who remained genuinely sceptical about whether AI had any place in their work at all. All of these perspectives were valuable.

I am a technophile and a user of AI and I came to this roundtable programme with enthusiasm to see how others are using AI. But, as a technology lawyer, I also brought my own professional lens to assess the legal questions surrounding AI, such as data protection, bias, liability, intellectual property infringement and regulatory compliance.

The roundtables brought together healthcare professionals, teachers, farmers, financial services practitioners, and social care providers – people whose daily work could not be more different, yet who found themselves grappling with remarkably similar questions. What is AI? Where do we start? Should we use it? How can we use it to our advantage? Will we become more competitive by using it or, conversely, less competitive if we don't? Do we have budget for this? What are the governance and legal risks we should be worried about and how do we address them? How will using AI affect our workforce?

One observation stood out consistently across every session: while some attendees exhibited scepticism over use of AI or lacked the budget or governance mechanisms to implement it properly, all acknowledged that AI is a pivotal moment in our society. It is here to stay whether they actively embrace it or not. This report does not say you must use AI nor promise that using AI will solve every problem. Instead, it tries to tell you what is known, what remains uncertain, and what your organisation should be thinking about now and how it can address the issues and remain compliant.

What emerged most clearly from the programme is that the organisations making genuine progress with AI are not necessarily those with the most sophisticated tools. They are the ones that began with a specific, well-defined problem, have invested in the right AI tools, implemented them wisely to address the problem, and built governance structures that keep responsible humans genuinely in the loop and supervise AI outputs. I hope this report provides a foundation for others to do the same.

Frank Jennings
Partner, Technology and Innovation, HCR Law

Executive summary

Artificial intelligence started as a research discipline and has evolved in a practical set of interconnected technologies. However, its role in everyday organisational life remains uneven and often misunderstood. While some sectors are moving quickly, many organisations remain unsure where AI fits, what value it offers, what risks it introduces and how to maintain compliance with evolving legal and regulatory requirements. This report focuses on that reality, where ambition, uncertainty and capability often collide. It explores how organisations are navigating AI in practice. It also provides practical guidance on how organisations can identify, manage and address AI-related issues.

The evidence points to rapid but uneven adoption. UK business use of AI has roughly doubled over the past two years. Yet, adoption sits alongside widespread disengagement and limited understanding. 80% of UK businesses still see AI as irrelevant to their work. Only 17% of adults can explain what AI actually is, and just 11% of workers receive adequate training to use it safely and effectively.

CyNam and HCR Law's Everyday AI programme set out to understand this gap. Across forums in healthcare, education, financial services, agriculture, and social care, practitioners were trying to work out what AI actually is, whether the tools on offer were suited to their needs, what to do when something went wrong, and how to make reasonable decisions in the absence of clear guidance. Those conversations are the foundation of this report.

What follows works through the questions the forums raised: what AI is and what it already is in organisations that do not realise they are using it; whether the history of AI should make us cautious about current claims; what the international landscape means for UK organisations right now; how AI is being used across sectors; what responsible and lawful use looks like in practice; what the real risks are and how to mitigate them; what is happening to skills and employment; and how to get started on the AI adoption journey. This report also reviews what the regulatory landscape means in the context of compliance and offers actionable guidance.

Hollie Wakefield
General Manager, Cynam

Only 17% of adults can explain what AI actually is

Reflections on the forum series

The Everyday AI programme began with a commitment to listening to the experiences of those on the ground. CyNam and HCR Law convened professionals from healthcare, education, financial services, agriculture, and social care to hear what AI looks like from inside their organisations.

What emerged spanned vastly different sectors, scales, and levels of AI maturity. Yet, the same questions surfaced with remarkable consistency.

Key questions

What even is AI? Am I already using it without knowing?

In every forum, participants associated AI primarily with the tools they had chosen, such as ChatGPT, Copilot, and Gemini. Many did not recognise how AI was already embedded in their organisation.

An organisation that cannot identify its AI cannot assess whether it is working as intended, understand what might go wrong, or act responsibly toward the people, processes and planet whose decisions it affects. This is why the report begins with definitions.

Many participants had already adopted a wide range of new technologies. While these did change how people worked, they took longer to embed than expected, required significant organisational effort, and arrived amid much hype. Page 16 examines AI's own history with this in mind.

Technology is always advancing, but what makes this different?

What is happening globally, and what does it mean for us?

Participants repeatedly highlighted the international dimension of AI use. Their concerns focused on where physical infrastructure is located, how dependent organisations are on global supply chains and cloud services, and how resilient those systems are to disruption. Page 19 provides a more detailed overview of the geopolitical context.

An organisation that cannot identify its AI cannot assess whether it is working as intended

What is AI actually doing in organisations like ours?

Forum participants highlighted the difficulty of navigating a market saturated with AI products and services. Page 22 provides an account of what AI is actually doing across the UK sectors CyNam and HCR Law chose to focus on.

What does it mean to use AI responsibly?

Values questions surfaced in every forum in sector-specific forms. Page 34 addresses this, including the ethical principles that should shape AI use and what responsible use looks like in practice.

Forum participants were not primarily worried about superintelligent machines. They were worried about confident but wrong outputs in clinical decision-making; deepfake calls impersonating a CEO; AI recruitment tools quietly discriminating against underrepresented applicants; open-source models with safety controls removed; and cloud infrastructure physically targeted in a military conflict. Page 37 examines these risks in more detail.

What are the real risks, and which ones apply to us?

What are the skills and talent development implications?

Participants raised concerns about how AI affects skills, learning and workforce development. Issues raised ranged from individuals relying on AI rather than learning, to the loss of junior roles, pressures on under-resourced staff and digital skills gaps. Page 41 provides a more detailed review of the skills landscape.

Despite the uncertainty, every forum highlighted a consistent desire to find applications where AI could make a genuine difference. Page 52 outlines how to get started.

What is actually worth doing, and how do we start?

Despite the uncertainty, every forum highlighted a consistent desire to find applications where AI could make a genuine difference

Defining AI

If an organisation does not know what AI is, it cannot identify when it is using it, how to use it well, or act responsibly and compliantly toward the people, processes and data it affects. The Everyday AI forums found that most participants associated AI primarily with generative AI tools such as ChatGPT, Gemini, or Copilot. However, AI is much more than that. AI encompasses a wide range of interconnected algorithms, processes, and models which are used much more than people realise. The first step in any assessment is to conduct a thorough audit of existing systems to identify where AI is already present. The following diagram illustrates this point and some of the existing and emerging fields of AI.

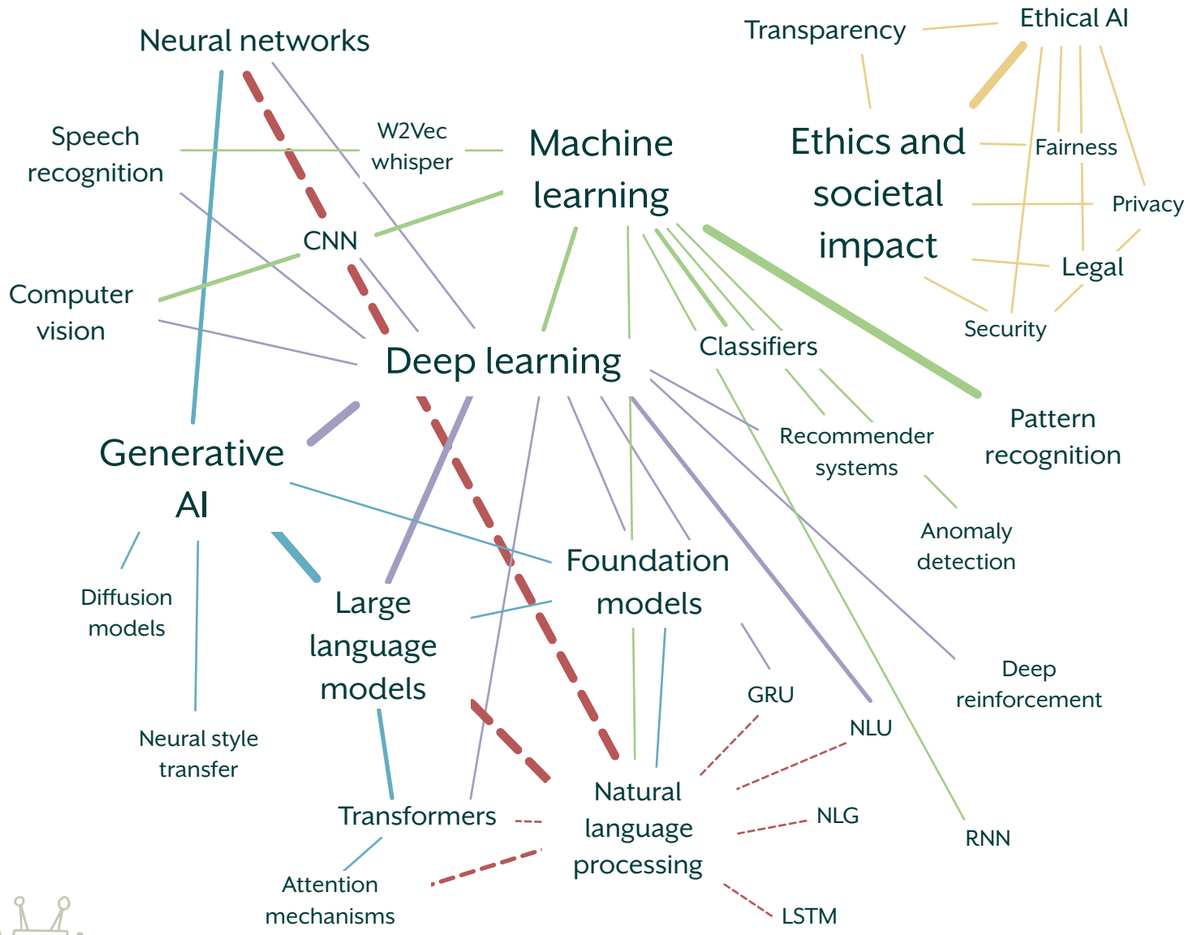

Reference: Artificial Intelligence Playbook for the UK Government, 2025

The UK Government adopts the internationally agreed OECD definition:

> A machine-based system that, for explicit or implicit objectives, infers from the input it receives how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments.

If an organisation does not know what AI is, it cannot identify when it is using it

The definition in the EU AI Act also mirrors this definition closely. Many other organisations adopt a broader view of Artificial Intelligence as the development of systems capable of performing tasks that typically require human intelligence and perhaps this accounts for the different understanding of AI.

For an overview of key definitions, please visit the glossary at the end of this report.

History of AI

Society has experienced repeated waves of transformative technology. Each provoked claims of fundamental changes to work, culture and power, yet none replicated core human cognitive abilities.

In what ways is AI different from earlier technological advances, and in what ways is it similar?

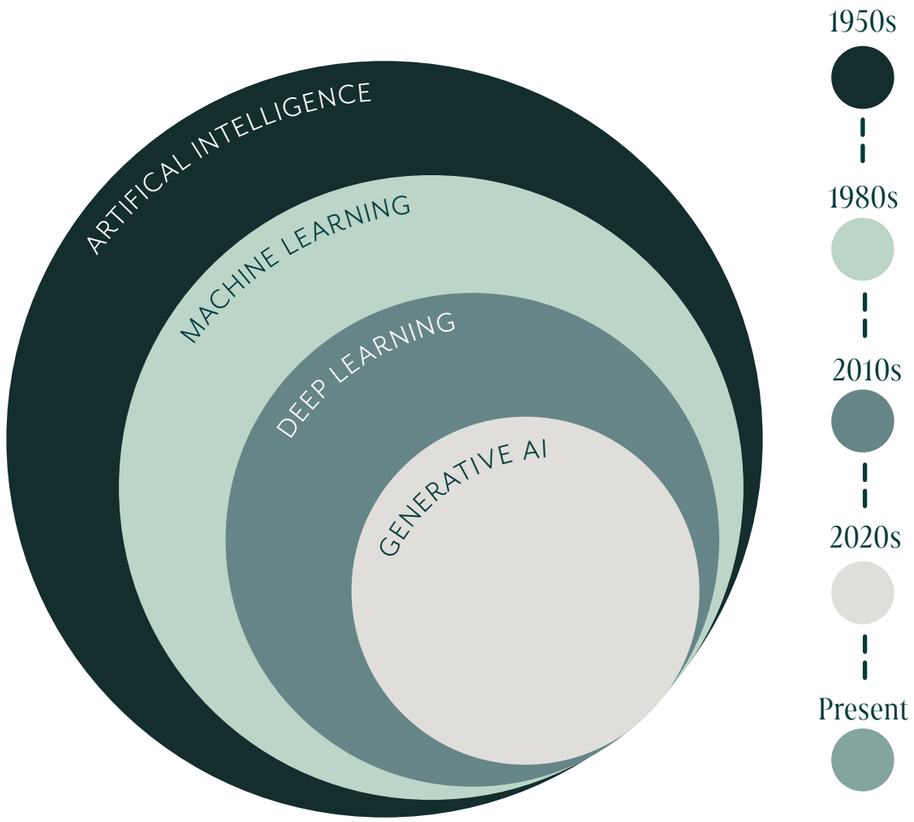

Source: A History of Artificial Intelligence, Oxford University

AI has developed over more than 70 years, progressing from academic conversations to a major field of science, engineering and industry. Its trajectory has not been smooth, with periods of optimism repeatedly met by disappointment, funding cuts and slower advancement than anticipated.

The foundations were laid in the 1930s–1950s by figures such as Alan Turing, who formalised the general-purpose computing machine and posed questions about machine intelligence, later known as the Turing Test. Early work in logic, cybernetics and information theory established concepts still used today, including artificial neurons, learning rules, and symbolic reasoning.

In 1956, the term "artificial intelligence" was coined at the Dartmouth Conference. Early programs appeared impressive in controlled settings, leading researchers to predict rapid progress. Languages like LISP, early neural networks, and systems like ELIZA and SHRDLU demonstrated narrow capabilities. However, these systems struggled with real-world complexity. Herbert Simon predicted in 1960 that machines would match human capabilities within twenty years; Minsky predicted in 1970 that general human-level intelligence would arrive within a decade. These predictions did not age well.

By the mid-1970s, limitations became clear. Government reviews (notably the UK's Lighthill Report) concluded AI had failed to meet its promises. Funding was sharply reduced, triggering the first "AI winter", a period when funding, corporate interest and enthusiasm collapse.

In the 1980s, AI revived through expert systems encoding specialist knowledge as explicit rules. Stanford's MYCIN diagnosed blood infections effectively, and Digital Equipment Corporation's XCON automated computer configuration, leading to a modest commercial industry.

Japan's Fifth Generation Computer Systems project aimed to develop logic-based, massively parallel machines. The project spurred international responses, including the UK's Alvey Programme and DARPA's Strategic Computing Initiative, but Japan's Prolog-based systems were overtaken by cheaper general-purpose PCs.

ChatGPT reached 100 million users within 2 months of its launch

A second downturn occurred in the late 1980s after expert systems proved expensive to maintain and fragile outside tightly defined tasks. Computers from Apple, IBM and Sun became fast enough to run rule-based software cheaply. By the 1990s, researchers renamed their work ‘data-mining’ or ‘analytics’ to avoid the stigmatisation of AI labels.

Despite this, key advances continued. Researchers developed backpropagation for training deep neural networks. This era’s milestone was IBM Deep Blue, which defeated the reigning world chess champion.

A major shift occurred in the late 2000s and 2010s, as vast data and compute power made deep learning possible. Systems such as AlexNet, AlphaGo, and AlphaFold showed that learning-based approaches could outperform earlier methods.

From 2018 onward, transformer architecture and vast collections of digital text led to large language models (LLMs), neural networks trained to predict the next token (a fragment of a word) after reading trillions of tokens.

Models such as BERT and GPT-series demonstrated that scaling improved performance. On 30 November 2022, OpenAI launched ChatGPT, reaching 100 million users in 2 months. Public deployment sharply increased awareness and commercial investment, alongside concerns about accuracy, misuse, and social impact.

From 2024–2025, global competition intensified. DeepSeek released its R1 reasoning model, assessed as comparable to OpenAI’s o1 at substantially lower cost. Other Chinese organisations, including Alibaba Cloud, Baidu, and 01.AI, released competitive models.

Several laboratories introduced models designed for multi-step reasoning, including OpenAI’s o-series, Google DeepMind’s Gemini 2.5, and Anthropic’s Claude 3.7. Parallel work focused on AI agents—systems combining language models with software tools to carry out semi-autonomous tasks.

AI’s history shows that progress depends on computing power, large-scale datasets and sustained engineering effort, often decades after core ideas were first proposed. Overall, though, the field’s long-term trajectory has been upward.

AI and geopolitics

Developments in artificial intelligence converge with global politics. From its earliest development, AI has been influenced by government funding, national security priorities and international competition.

In today’s increasingly fragmented world, organisations adopting AI are more than ever exposed to geopolitical risks embedded in supply chains, infrastructure, regulation, and data flows. These risks do not only affect governments or large corporations. They shape the reliability, cost, legality and sustainability of the AI tools used by organisations across the UK.

Understanding these risks means understanding that AI depends on physical components, physical infrastructure, and legal systems that are contested and changing.

Supply chain concentration

Most AI systems ultimately depend on the same small set of underlying components. The chips that power modern AI, known as graphics processing units (GPUs), are designed primarily by one US company (Nvidia) and manufactured almost entirely by one company in one place: Taiwan Semiconductor Manufacturing Company (TSMC) in Taiwan.

TSMC produces over 90% of the world’s most advanced semiconductors. Taiwan’s geopolitical position makes this supply chain unusually fragile. Military issues in Taiwan could lead to shortages in supply.

For organisations, this matters because chip shortages propagate quickly. When supply tightens, computing capacity would be constrained, and prices would rise, highlighting how businesses would be exposed to geopolitical disruption.

Physical infrastructure

AI depends on physical infrastructure including data centres, undersea cables and power grids.

Data centres are large, permanent buildings containing vast numbers of computers. They are resource-intensive, power-hungry and located in known places. Major cloud providers operate data centres across the UK, Ireland and mainland Europe and further afield including the Middle East.

AI adoption is not purely a technical or productivity decision

The UK's National Cyber Security Centre now classifies AI data centres and connectivity as critical national infrastructure. Disruption could come from severe weather, power failures, physical damage, or single-site dependency, leading to widespread service outages.

Undersea cables present a related vulnerability. The UK depends heavily on a small number of cable landing points. Sustained disruption would degrade internet connectivity, increase latency, and potentially render cloud-based AI systems unreliable or unusable for extended periods.

Power supply adds another layer. AI data centres consume far more electricity than conventional IT systems. As AI use grows, grid pressure increases. Power disruptions, whether from overload, weather, or an attack, could affect AI usage.

Data sovereignty and legal reach

When using AI tools to process data, that data might be transferred abroad. Organisations should map these data flows carefully and, where cross-border transfers are involved, ensure they comply with applicable data protection regimes, not just under UK law, but also the laws of the country where data is stored or processed.

Many widely used AI tools operate on infrastructure owned by US-headquartered companies.

Under the US CLOUD Act, American authorities can legally compel those companies to provide access to data under their control, even if the data is physically stored in the UK or the EU.

For most organisations, this remains a low-likelihood event. But for organisations handling sensitive, regulated or confidential information, it matters that compliance with UK GDPR does not override these obligations.

A more immediate data risk comes from AI training practices. Depending on the tool and its terms, information entered into AI systems may be used to improve or train future models. Without clear organisational rules, people may paste client data, internal documents or proprietary information into consumer-grade tools, unintentionally disclosing it under terms they have not read. Organisations should identify data-handling risks before deploying AI and establish clear acceptable use policies and review vendor terms of service.

Regulation fragmentation

The EU has taken a comprehensive approach to AI regulation through the EU AI Act, placing strict obligations on certain uses of AI, particularly in areas such as employment, law enforcement, public services, education and critical infrastructure. These rules apply based on where the AI is used, not where the organisation is based.

The UK has opted not to introduce an overarching AI law, instead relying on existing regulators to enforce AI-related obligations within their sectors.

The US has no single federal AI law, but a growing patchwork of state-level rules and expanding export controls affecting AI technologies and cloud computing.

For organisations operating across borders or using AI tools developed for multiple markets, this creates complexity. A system acceptable in one region or country may trigger obligations or restrictions in another. Organisations with cross-border operations or international client bases should map their regulatory exposure carefully to navigate overlapping or conflicting requirements.

Implications

These risks mean that AI adoption is not purely a technical or productivity decision.

AI systems depend on concentrated supply chains, contested infrastructure, evolving legal regimes and geopolitical dynamics that affect availability, cost, legality and trust. Organisations that understand this context are better placed to make resilient choices about tools, providers, data and skills.

AI systems depend on concentrated supply chains, contested infrastructure, evolving legal regimes and geopolitical dynamics

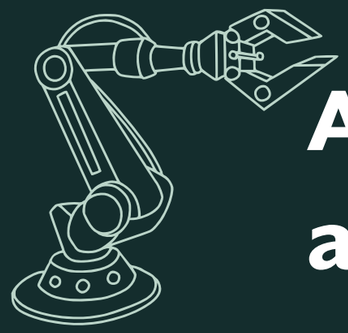

AI adoption and use cases

The state of adoption

The most comprehensive national picture of AI use in the UK private sector comes from DSIT-commissioned research, which draws on responses from 3,500 businesses. Its headline finding is that around one in six businesses (16%) are currently using at least one AI technology, while 80% have neither deployed nor planned to adopt AI. More telling still, over half of all businesses surveyed (51%) do not consider AI relevant to their organisation.

Against this national backdrop, the findings from CyNam and HCR Law's own Everyday AI Survey 2026 tell a markedly different story. All respondents were either already using AI or planning to introduce it within twelve months, 58% with AI actively deployed and in regular use, 25% in active piloting, and 12% in the planning phase. This divergence is not surprising: the survey's respondents skew towards businesses and individuals who have actively engaged with CyNam and HCR Law's AI work. What it does reflect is the views of those working through the practical realities of AI adoption in real time, which the wider business community can learn from.

Adoption in the survey cohort is also intensive. 25% of the respondents use AI tools daily or multiple times a day, and only 8% reported not currently using AI tools personally. This is a group for whom AI has become a routine part of working life, not an occasional pilot.

The national picture, however, provides important context. The gap between this cohort and the 51% of UK businesses that see AI as irrelevant is not simply a gap in access or tools. It reflects a deeper unevenness in awareness, confidence, and the internal capacity to navigate an increasingly complex landscape.

Who is adopting and where

At the national level, business size remains one of the most reliable predictors of AI adoption. Large businesses are more than twice as likely to use AI as micro businesses (36% compared to 14%), and mid-sized businesses sit between these at 23%. Larger organisations are more likely to have dedicated technical functions, clearer data infrastructure and the internal expertise needed to navigate an increasingly complex tooling landscape.

1 in 6
businesses are currently using at least one AI technology

Sector affiliation is equally significant. Information and communication businesses lead at 43% adoption, followed by business services and administration at 23% and finance and real estate at 21%. At the other end, construction, transport and storage, and hotel and catering all see adoption rates below 12%, reflecting a genuine mismatch between current AI capabilities and the physical, operational, or interpersonal nature of work in those sectors.

CyNam and HCR Law's Everyday AI forum series explored adoption in depth across four distinct sectors (financial services, health and social care, education, and agriculture), bringing together practitioners, sector leaders, legal experts, and researchers to examine what AI actually looks like in practice. The picture that emerged across all four was consistent with the national data in one important respect: AI is arriving unevenly, concentrated in specific functions and driven by motivated individuals rather than organisation-wide strategy. What varied was the nature of the use cases, the maturity of governance, and the specific barriers each sector must overcome.

What AI is being used for

Among businesses that have adopted AI nationally, use cases cluster heavily around generative AI. The majority (85%) are using natural language processing and text generation, which encompasses tools such as ChatGPT, Claude, and Gemini. The most common functional areas are marketing (72%), administration (72%), and IT (64%), with a focus largely on task automation and content generation.

The survey closely reflects this pattern. The two most common tasks reported were summarising and analysing information, and drafting documents, emails or reports, each cited by 83% of the respondents.

AI is being used to accelerate and augment professional work, not to replace the judgement and expertise of the person using it

Respondents described their use in ways that go beyond simple task completion. One described using AI as a "thinking assistant", deliberately instructing it to take an opposing view and argue back as a way of stress-testing their own reasoning. Another described AI as enabling tasks they "wouldn't have been able to otherwise." A technology professional valued it for "getting the ball rolling" with content, emphasising they never used AI output verbatim but adapted it to their own voice. These descriptions share a common thread: AI is being used to accelerate and augment professional work, not to replace the judgement and expertise of the person using it. The human remains firmly in the loop.

The sector forums revealed a considerably richer and more varied picture of what AI adoption looks like beyond the office, in fields, hospitals, classrooms, and trading floors.

Clare Day
Partner, Head of Financial Services, HCR Law

> Whilst it is clear that AI is already very much embedded in the Financial Services sector, particularly in areas such as automated valuations, compliance, and finance/credit applications, where its use is speeding up processes and reducing the need for as much manpower, it is also recognised that its use is not without risk, particularly with such sensitive information as financial and personal data being used.
>
> Incorrect use or overreliance does risk leading to inaccurate or misleading outputs. However, in an industry which relies so heavily on relationships and human interaction, if used cautiously, it could enhance relationships by streamlining processes and reducing costs. It is essential that those involved in the sector collaborate to develop a deeper understanding of how AI can best be used and to share best practice.
>
> **Clare Day**

The CyNam and HCR Law AI in Financial Services forum brought together professionals from across the sector and found AI already embedded across a wide range of applications, many of which go well beyond the language tools that dominate national statistics.

Machine learning models are being used to detect fraud and financial crime, assess credit risk, and monitor transactions in real time. In fraud detection specifically, forum participants noted that AI-driven systems have significantly reduced false positives, allowing human investigators to focus on genuinely suspicious activity, thereby improving outcomes and the efficiency of compliance teams. In lending, AI is enabling institutions to move beyond traditional credit scoring by analysing a broader range of data signals, resulting in faster, more accurate, and more personalised decisions for customers.

Customer service is another area of active deployment. Chatbots and voice assistants are providing around-the-clock support, reducing pressure on call centres and allowing staff to focus on complex or sensitive cases that genuinely require human judgement. Internally, tools such as Microsoft Copilot are being used to streamline routine tasks, analyse data, and support document creation, with forum participants reporting savings of several hours per employee per week.

One of the more sophisticated applications discussed was AI-enabled Know Your Customer (KYC) onboarding, where identity documents are verified, information extracted, and risk assessed through a combination of computer vision, natural language processing, and expert systems operating as a compound workflow. This kind of multi-model integration represents a step beyond single-tool adoption and is increasingly common in larger financial institutions.

Alongside these productivity gains, the forum was equally frank about the risks. Overreliance on AI outputs without applying critical thinking was flagged as a genuine concern in regulated environments where decisions carry real consequences for customers. Cyber security risks were also highlighted as AI increases the attack surface through system complexity and third-party integrations, and can itself be a vector for sophisticated attacks, including model manipulation and AI-enhanced phishing. A notable development discussed was the FCA's AI Live Testing programme, which aims to define what safe and responsible AI looks like in practice through real-world testing, which is an important signal of where regulatory expectations in the sector are heading.

> This culmination of the insights from CyNam and HCR Law's Everyday AI programme highlights that, while UK business adoption of artificial intelligence has doubled over the past two years, the successful implementation within organisations remains highly uneven. The constraints include widespread corporate disengagement and severe data-handling, legal, and skills gaps. The report shows putting AI to work needs to start with business outcomes or problem-first strategies before tool selection, and must include robust human oversight, and clear governance structures to mitigate the practical risks across sectors like healthcare, education, professional services, finance and agriculture.
>
> **David Terrar, CEO, Tech Industry Forum**

The CyNam and HCR Law AI in Health and Social Care forum opened with a caution that runs through the entire sector: a saturated market of AI products, many of which rely on generic off-the-shelf models beneath a polished surface, means that selection and governance matter enormously. As participants emphasised, AI works best when applied to specific, well-defined problems, supported by high-quality data and clear oversight frameworks.

Within those boundaries, the applications discussed were significant. AI is being used to help clinicians by surfacing the latest research and clinical evidence at the point of decision and by detecting patterns in large datasets to support diagnosis, treatment planning, and early intervention. Personalised care applications, including tailored exercise or lifestyle advice designed to reduce pressure on frontline services, represent a use case with substantial long-term implications for health system capacity.

Care-focused technologies were also discussed, with an emphasis on tools that promote independence without replacing the human relationships that underpin patient trust. This distinction was described as the shared principle across all examples considered.

The forum gave significant attention to the data challenges specific to this sector. Patient records are held across multiple legacy systems and organisations; cleaning, linking, and governing this data is not merely a technical task but a prerequisite for AI to function reliably and safely. Concerns about data ownership were raised regarding the growing volume of personal health data generated by consumer apps, wearables, and in-home devices, which raised live questions about consent, access, and third-party use. The work of the Oxford-led Digital Care Hub, which has published practical guidance on the responsible use of generative AI in social care, was highlighted as an important reference point for organisations navigating these questions.

> When I talk to health and care providers about AI, I encounter caution. Given the risks and hurdles, and the importance of confidentiality, safety and integration in these sectors, that's a sensible starting point.
>
> But the work I do with HealthTech suppliers who are harnessing AI, and the use cases in this report, show that, with human control and other appropriate guardrails in place, in some contexts AI offers major benefits across the health system, for patients, clinicians, staff and the 'back office'.
>
> AI is tricky to govern, evaluate, select, procure and prove, but there are routes through it all.
>
> **David Hall**

David Hall
Partner, Technology and Innovation

Kristine Scott
Partner, Head of Education

> The education sector findings reflect what we are seeing. We are aware that teachers are using AI to support lesson content and design – this differs from school to school and can certainly save preparation time. As with everything, the content still needs to be checked to ensure it is accurate.
>
> Aside from the challenges already highlighted, where AI implementation and management usually fall to the IT team within a school (which may be one person), we have seen issues arise when AI has been implemented without appropriate controls and training for those using it across the staff and governing body. This is particularly relevant as governors are volunteers and may not be in school regularly.
>
> For example, we now frequently see an AI note-taker (sometimes several – one for each participant) in online meetings, which can be helpful in keeping notes or a transcript. This is useful; however, schools have to consider what happens to those notes. If notes are automatically sent to all attendees, including those who only attend part of a meeting, is that appropriate, particularly where the meeting is sensitive, for example with a parent or to discuss a pupil?
>
> **Kristine Scott**

The AI in Education roundtable, attended at full capacity by representatives from schools and colleges across South West England, surfaced a sector experiencing AI adoption from the bottom up, faster and less predictably than institutional frameworks can currently accommodate.

Students are often more proficient with AI tools than their teachers. This inversion creates a particular challenge: schools are expected to integrate AI into the curriculum while simultaneously managing the risk that students will use it to complete assessed work rather than learn. As one roundtable participant put it, the point of school is to learn, not to complete tasks, and completed tasks are only evidence of learning if they reflect genuine engagement. The forum discussed this tension at length, noting that existing guidance or policy frameworks do not yet resolve the ethical and educational questions raised by AI in assessment.

Institutionally, adoption is fragmented. Some schools use AI to streamline administrative processes and save staff time; others have deployed AI tools for marking or content delivery; a smaller number have restricted AI use entirely to prevent plagiarism. These differences reflect access, cost, and senior leadership confidence as much as they reflect considered strategy.

The forum noted that AI adoption is beginning to affect school competitiveness, adding further pressure, particularly for under-resourced schools who may not be able to embrace AI as effectively as better-resourced schools.

A recurring structural concern was the tendency for AI governance to be delegated downward. Rather than being owned by senior leadership teams, digital strategy and AI policy in some schools has fallen to computing departments, placing undue pressure on subject heads to design and manage policies for which they may lack both authority and time. The forum was clear that this is unsustainable: AI governance in education, like safeguarding or health and safety, requires top-down ownership and strategic commitment from leadership.

The broader curriculum challenge also featured prominently. Many schools struggle to offer GCSE and A-level computing due to teacher shortages, and the curriculum is widely considered outdated compared to developments in AI. Schools that have engaged with programmes such as CyberFirst have seen the value of industry-connected, practical exposure, but these remain the exception rather than the norm.

Rory Hutchings
Partner, Head of Agriculture and Estates

Agriculture is one of the sectors least represented in national AI adoption statistics. Yet, the CyNam and HCR Law AI in Agriculture forum revealed a sector where specialist AI applications are already delivering measurable operational value, often in ways that bear little resemblance to the language tools dominating broader adoption surveys.

The forum exhibited the most scepticism of all our forums, with some attendees questioning whether a tool or application was truly AI and whether it would benefit them. Conversely the forum also documented a range of deployed applications. Wearable sensors using machine learning monitor each cow's temperature, movement, eating habits, and rumination patterns, identifying deviations from the norm that can signal illness hours or days before visible symptoms appear, enabling timely treatment and preventing production losses. Overhead cameras using deep learning and computer vision automatically identify early indicators of lameness in dairy herds by interpreting gait and movement patterns, providing consistent detection that manual observation cannot match at scale. Robotic milking systems integrate computer vision, sensor analytics, and machine learning to guide robotic arms, monitor yield and udder health, and flag early signs of conditions such as mastitis.

In arable farming, precision-spraying systems using computer vision distinguish weeds from crops in real time, activating individual nozzles to apply herbicide only where necessary, reducing chemical use, lowering costs, and improving input efficiency at the plant level. In poultry, AI-driven environmental and behavioural monitoring has demonstrated the ability to detect heat-stress-related crowding early enough to make environmental adjustments that prevent losses.

Beyond animal and crop management, AI is also being applied to regulatory compliance and environmental reporting, processing soil data, land boundaries, and environmental designations to generate compliant nutrient management plans, and analysing structured data on livestock and fertiliser use to produce greenhouse gas emission profiles aligned with national reporting standards.

The forum also surfaced a challenge that cuts across the sector: interoperability. Participants described using multiple farm management apps that do not communicate with each other, resulting in duplicated data entry and limited ability to derive integrated insight. One farmer described a situation in which none of their digital tools shared data, significantly undermining the potential value of the information being collected. This is a structural barrier to the kind of joined-up, data-informed decision-making that AI at its most useful could enable.

Trust, cost, and data ownership were identified as the three practical concerns most likely to determine the pace of adoption across the sector. Farmers need to understand how AI systems reach their conclusions before they will act on them with confidence. Return on investment must be demonstrable, particularly for small and medium-sized farms. And the question of who owns the sensitive operational data flowing through third-party platforms requires clear answers before many farmers will commit to further integration.

How businesses are adopting AI: trial, buy, build, or partner

Businesses approaching AI adoption face a fundamental strategic choice between four broad routes, each carrying distinct implications for cost, control, capability, and risk.

The first approach was to use an unpaid version of any off-the-shelf solution. This happens where the business does not have an AI strategy and budget, or has not endorsed a particular tool. This creates "Shadow AI" and allows some form of benefit but without addressing any of the risks of using AI.

Buying off-the-shelf is by far the most common approach and the natural entry point for most organisations. General-purpose tools, including subscriptions to large language models, AI features embedded in existing software suites, or specialist point solutions, require little technical expertise to deploy and can deliver value within days. The 'There's an AI for That' catalogue listed over 48,500 tools across more than 11,300 distinct task types as of April 2026. The challenge for businesses is not access but navigation: knowing which of thousands of available tools is the right fit for a given problem. This challenge was raised explicitly in the financial services forum, where participants noted that build, buy, or partner decisions involve real trade-offs between cost, speed, control, and risk that organisations need to think through carefully rather than defaulting to the most visible solution.

Building bespoke solutions remains the domain of larger organisations. Nationally, 46% of large businesses plan to build their own AI models, and 39% plan to develop bespoke AI systems. Building in-house offers the greatest degree of customisation and control, and allows organisations to train models on proprietary data. However, it requires significant investment in technical talent, data infrastructure, and ongoing maintenance. In agriculture, the interoperability problems surfaced in the forum suggest that even where individual tools are bought off the shelf, the absence of open integration standards can undermine the value of the overall system, pointing to a gap that only more deliberate platform-level thinking can address.

75% of AI-using businesses reported improved workforce productivity

The fourth approach – partnering with specialist providers – represents a middle path that is increasingly attractive for organisations that need more than a generic tool but lack internal capacity to build. The health and social care forum was particularly clear on this point: the market contains many AI products built on generic models, and sector-specific needs like clinical safety, data governance, regulatory compliance, and patient trust demand providers with genuine domain expertise rather than AI capability alone.

In practice, many organisations will combine these approaches, selecting different routes for different problems and different stages of their AI journey. The question of how to compose an effective AI capability, rather than simply which single tool to adopt, is becoming increasingly central to how forward-thinking organisations approach the technology.

The benefits

The productivity case for AI is well evidenced at the national level. Three-quarters (75%) of AI-using businesses reported improved workforce productivity, and over half (56%) reported increased employee productivity since adopting AI. Our survey reflects this: the average self-reported impact score among respondents was 7.5 out of 10, with 25% rating it 10 and none rating it below 5. All respondents expect AI to change their organisation significantly or fundamentally within five years.

The sector forums revealed the specific contours of these benefits more clearly than headline statistics can. In financial services, gains are concentrated in fraud detection accuracy, faster lending decisions, reduced call centre load, and increased internal productivity through tools like Copilot. In agriculture, benefits are most visible in animal welfare outcomes, reduced chemical use, and the early detection of health and compliance issues that would otherwise require manual observation at scale. In health and social care, the opportunity is described primarily in terms of better decision support, reduced pressure on frontline services through prevention and personalised care, and the ability to analyse datasets too large for human review. In education, the benefits are real but contested: AI is saving staff time on administration and supporting lesson preparation, while simultaneously creating new problems around academic integrity and equitable access.

The national data provides an important point alongside all of this. Over three-quarters (77%) of AI-using businesses have not yet seen a change in revenue, a reminder that efficiency gains at the task level do not automatically translate into commercial outcomes. This gap between productivity improvement and revenue growth reflects the current stage of adoption rather than a fundamental limitation of the technology, but it does mean that businesses should be clear-eyed about the near-term value proposition. AI is currently performing best as a tool for doing existing work more efficiently and more effectively.

From access to action

Across both national research and CyNam and HCR Law's survey and forum series, a consistent message emerges: the barrier to AI adoption is no longer access to tools, but knowing how to begin using them effectively.

Many organisations lack the confidence and internal capability to move forward, largely due to the absence of strategic direction. In the survey, 58% of respondents identified a lack of strategy or leadership buy-in as a key challenge, placing it alongside data privacy and security as the most significant barriers to AI adoption. Without senior ownership, organisations struggle to move beyond discussion into coordinated action.

Respondents were explicit about what is needed to address this gap. They highlighted the importance of "leadership, ownership of strategy, practical steps and starting points," alongside "a clear strategy and budget to implement tools more consistently across the business, rather than just talk about it," and "professional guidelines." These responses point to a shared need for clarity, structure, and direction rather than additional tools

In the absence of this, AI adoption tends to be fragmented and inconsistent, with isolated efforts that fail to build broader organisational capability. The issue is not a lack of interest or opportunity, but a lack of alignment and readiness: organisations often struggle to connect business problems to appropriate AI solutions, to assess trade-offs between buying, building, or partnering, and to develop the internal understanding required to use AI responsibly and effectively.

AI ethics

AI is often adopted with the intention of improving business outcomes, yet it can lead to unforeseen consequences, unfair outcomes, and potentially regulatory exposure. AI ethics are important to mitigate this. AI ethics relate to principles that steer how to optimise the benefits of artificial intelligence while reducing risks and adverse outcomes. Embedding ethical principles into AI governance is not merely good practice, it is increasingly becoming a compliance requirement.

The core principles

Multiple international frameworks converge on six common principles, including:

- **Fairness:** AI should not produce unjustifiably different outcomes for different groups.
- **Transparency:** the logic and limitations of AI should be understandable to those affected.
- **Accountability:** clear responsibility must exist for AI behaviour.
- **Privacy:** AI should not undermine individuals' reasonable control over their personal data.
- **Non-maleficence:** AI should not cause unnecessary harm.
- **Beneficence:** AI should actively contribute to human wellbeing.

The UNESCO recommendation

The UNESCO Recommendation on the Ethics of AI, adopted unanimously by all 193 Member States in November 2021, is the first global instrument specifically on AI ethics. It is built on four core values:

1. Human rights and human dignity. Respect, protection and promotion of human rights and fundamental freedoms and human dignity
2. Living in peaceful, just and interconnected societies
3. Ensuring diversity and inclusiveness
4. Environment and ecosystem flourishing.

AI ethics are principles that steer how to optimise the benefits of AI while reducing risks and adverse outcomes

It also lays out a ten-principle approach to the Ethics of AI:

Proportionality and do no harm

Transparency and explainability

Safety and security

Human oversight and determination

Right to privacy and data protection

Sustainability

Multi-stakeholder adaptive governance and collaboration

Awareness and literacy

Responsibility and accountability

Fairness and non-discrimination

These values and principles can be used as a practical framework for organisations. This provides a structured way to think through whether, where, and how AI should be adopted. Rather than slowing progress, principles such as fairness, transparency, accountability, privacy, beneficence, and non-maleficence act as structured decision-making prompts: they help organisations assess who is affected by an AI system, what could go wrong, how risks would be identified and who would be responsible for acting. Applied in this way, these principles support better procurement choices, clearer governance, and earlier issue identification and remediation. This is, by definition, responsible AI adoption.

The risks of AI adoption

Bias and discrimination

AI systems learn from historical data. When that data is inaccurate, incomplete or reflects historical discrimination, AI perpetuates and amplifies those patterns at scale.

The UK cases are documented. The DWP's Universal Credit algorithm produced statistically significant disparities across every protected characteristic analysed. Pulse oximeters produced systematic underestimation of hypoxaemia in darker-skinned patients throughout the pandemic. Amazon's AI recruitment tool, trained on a decade of predominantly male application data, penalised CVs that mentioned women's organisations. A 2024 University of Washington study found AI résumé screening tools favoured white-associated names 85% of the time. Organisations deploying AI in recruitment, credit decisions, or other high-impact contexts should conduct bias audits before deployment and establish ongoing monitoring processes. Where automated decision-making affects individuals' rights, they should take steps to comply with requirements for meaningful human involvement.

Cyber security

The relationship between AI and cyber security is a double-edged sword. As a weapon, AI enables attacks of greater sophistication, at greater scale, with less specialist expertise required. As a defence, AI enables threat detection, vulnerability scanning, and incident response beyond the capacity of humans alone.

The UK Government's April 2026 open letter to business leaders was direct: 'A new generation of AI models are becoming capable of doing work that previously required rare expertise: finding weaknesses in software, writing the code to exploit them… Frontier model capabilities are doubling every 4 months.' Anthropic's Project Glasswing, which is a frontier model that found high-severity vulnerabilities in every major operating system and web browser, was restricted and not released publicly.

As another example, within days of Google releasing Gemma 4 E4B, researchers published a version with safety guardrails removed, achieving 97.5% compliance with harmful prompts.

39% of UK workers admit uploading company information into a public AI tool

Data privacy and confidentiality

Data privacy risks arise not only from how systems are deployed, but from how they are trained and how people interact with them in everyday use. The ICO's £7.5 million fine of Clearview AI established that privacy obligations apply at the point of training-data collection, not only at deployment, making clear that AI can infringe data protection rights before producing any output at all. For most organisations, however, the more immediate risk stems from unmanaged employee use: EY's 2025 survey found that 32% of UK employees use AI tools their organisation has not approved, often inputting sensitive data into consumer-grade systems. Highlighting this, KPMG found that 39% of UK workers admit to uploading company information into public AI tools. To mitigate these risks, organisations should implement clear policies on data usage, conduct data protection impact assessments for AI deployments and ensure that AI procurement decisions involve appropriate privacy and legal review.

8.8%
of total UK electricity demand in 2030

Environmental harms

AI systems require vast amounts of computing power, concentrated in data centres that consume significant electricity and water for processing and cooling. Oxford Economics projects that UK data-centre electricity consumption will increase fivefold by 2030, reaching 26.2 TWh, which is equivalent to 8.8% of total UK electricity demand. At the model level, researchers estimate that training a single large natural language processing model can emit the equivalent of over 600,000 pounds of CO2. However, efficiency gains have reduced some harms, notably DeepMind's AI-driven optimisation of data-centre cooling systems, which cut cooling energy use by 40%. At the same time, AI has the potential to mitigate emissions elsewhere; the LSE Grantham Research Institute estimates that AI applications in power, transport, and food systems could reduce global emissions by 3.2–5.4 billion tonnes of CO2-equivalent annually by 2035.

Existential and catastrophic risks

Existential risks from AI refer to scenarios in which highly capable AI systems could cause large-scale, irreversible harm to humanity, including outcomes that threaten societal stability. In a 2024 survey of 2,778 AI researchers, between 37.8% and 51.4% estimated a probability of at least 10% that advanced AI systems could lead to such catastrophic consequences, often characterised in public discourse as scenarios in which humans lose meaningful control over critical systems or decision-making. These risks are the subject of active debate within the research community and remain highly uncertain in both likelihood and form.

Intellectual property infringement

This arises from uncertainty over how training data is sourced, how outputs are generated, and who bears responsibility when those outputs infringe existing rights. The High Court judgement in Getty

Images v Stability AI (November 2025) clarified that AI model weights themselves are not infringing copies of training data, but also confirmed that infringement can occur at the point of output, as seen where generated images contained visible Getty watermarks. This distinction matters for organisations using generative AI: even where training occurs lawfully, outputs may still reproduce protected material in ways that infringe copyright or trademarks. Organisations should review their use of AI-generated content carefully and, where outputs are used commercially or externally, assess and manage infringement risk.

Employment

This risk lies in the restructuring of labour markets that disproportionately affects junior, entry-level, and routine cognitive roles. Dr Bouke Klein Teeselink's 2025 analysis of UK employment data found that firms with high AI exposure reduced total employment by 4.5% on average, with junior positions falling by 5.8%, while highly exposed occupations experienced a 23.4% drop in job postings and a £2,951 reduction in advertised salaries.

At the same time, AI-skilled workers command an average global wage premium of 56%, indicating that AI is redistributing opportunity as much as eliminating it. The risk, therefore, is a widening labour-market divide: AI substitutes for the routine tasks through which junior workers traditionally build skills and experience, while disproportionately rewarding those with the capacity to design, deploy, or supervise AI systems. Over time, this hollowing out of entry-level roles threatens workforce pipelines, exacerbates inequality, and concentrates economic gains among already advantaged groups, particularly disadvantaging women, younger workers, and those from lower socioeconomic backgrounds. Organisations implementing AI should consider employment law implications and consultation obligations and the need to manage workforce change fairly including re-training or re-skilling staff.

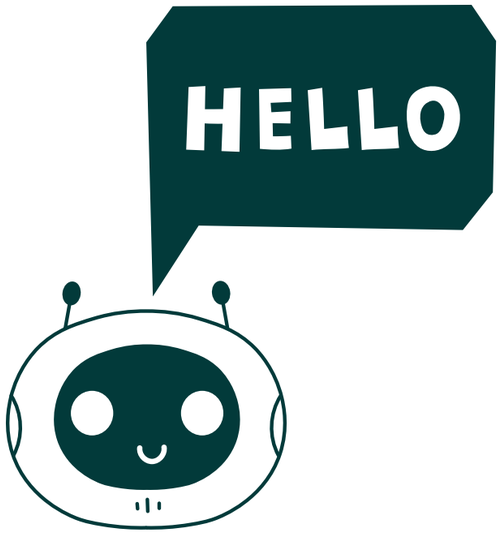

Accountability

This arises when responsibility for consequential decisions becomes unclear or fragmented across humans, organisations, and third-party technology providers. When AI systems influence legal, financial, clinical, or operational decisions, errors may emerge from a combination of model behaviour, human reliance, and platform design, making it difficult to determine who is answerable when things go wrong. The legal consequences of misplaced reliance are already visible. In Ayinde v Haringey (June 2025), the court issued wasted-costs orders against legal representatives who included AI-hallucinated case law in their submission to the court, with Dame Victoria Sharp DBE warning of “serious implications for the administration of justice.” Organisations should establish clear accountability structures for AI use and document decision-making processes to manage liability exposure.

Explainability and transparency

This risk arises because many AI systems operate as ‘black boxes’. While their inputs and outputs are observable, the internal processes by which they arrive at particular results are not readily interpretable by humans. This matters because AI systems are increasingly used to influence decisions, yet organisations may be unable to reconstruct, scrutinise, or justify how specific outcomes were produced. Post-hoc explainability techniques such as LIME and SHAP can help surface which inputs contributed to an individual output. Still, they explain what happened in a specific case, not why the system behaves as it does overall or whether its behaviour is appropriate, consistent, or reliable across contexts. The result is a gap between reliance and understanding: organisations may depend on AI outputs they cannot fully explain, making it harder to detect errors, identify bias, challenge flawed decisions, or maintain confidence when outcomes are questioned.

Misinformation and manipulation

This arises from its ability to generate persuasive, human-like content at scale, regardless of truth. Generative AI systems are designed to produce fluent and plausible outputs, not to verify facts, enabling the rapid creation of false, misleading, or fabricated text, images, audio, and video. Because this content often appears authoritative and lacks clear signals of inauthenticity, it can be used to impersonate individuals, distort information, or influence decisions before errors are detected.

One high-profile example is the UK engineering firm Arup, which lost $25 million to a deepfake video conference impersonating its CFO.

Organisations should consider taking steps to address this, such as Retrieval-Augmented Generation, to verify outputs.

Skills and the future of work

AI's impact on employment and skills is producing a measurable redistribution of opportunity. It is creating wage premiums and productivity gains for workers who develop AI skills, while reducing employment and entry-level opportunities in occupations where AI can substitute for routine cognitive work.

What UK workers are experiencing

DSIT's AI Skills for Life and Work survey found that while 73% of UK adults had used AI in the past month, most did so passively, using virtual assistants or predictive text. Only 28% of respondents felt confident in deliberately using AI tools. At work, 36% had used AI in the past month, and 21% said it had increased their productivity.

EY's 2025 Work Reimagined Survey found that 83% of UK employees use generative AI at work, but only 11% receive adequate training to do so safely and effectively. 32% are using tools their organisation has not approved, often inputting sensitive data. 43% worry that over-reliance on AI is eroding their skills.

The CIPD's Autumn 2025 Labour Market Outlook found that one in six employers expected AI to shrink their workforce, with 62% of those identifying clerical, administrative, and junior professional roles as most affected.

The impact on junior and entry-level roles

The most consistent finding across recent UK research is that AI's early impact is concentrated in junior and entry-level positions. Dr Bouke Klein Teeselink's 2025 analysis found that junior positions at AI-exposed firms fell by 5.8%, while total employment fell by 4.5%. Advertised salaries in high-exposure occupations fell by an average of £2,951.

AI is creating a skills divide as well as a jobs divide

This matters beyond its direct economic impact. Entry-level roles are where people develop the knowledge, judgement, and competence that cannot be acquired solely through formal training. If AI tools perform the tasks that junior professionals previously learned from, the pipeline through which people develop into more senior roles narrows.

The skills gap and the wage premium

DSIT's AI Labour Market Survey found 97% of AI-sector organisations reporting at least one skills gap, with 28% saying skills shortages had directly impacted business goals. Nash Squared found that 52% of UK tech leaders face AI skills shortages, a 114% increase from the previous year. The Government estimates the AI skills gap costs the UK economy £400 billion in unrealised potential.

PwC's Global AI Jobs Barometer found an average global wage premium of 56% for AI-skilled workers. In the UK, lawyers with AI skills command a 27% premium; database designers 58%. AI-exposed industries are seeing wages grow twice as fast for workers with AI skills as for those without. AI is creating a skills divide as well as a jobs divide.

Diversity and the AI skills divide

Women comprise only 22% of the UK's AI and technology workforce, a proportion that has been declining in the last few years. Female-founded AI companies receive six times less venture capital than their male-founded counterparts. Only 9% of technology employees come from lower socioeconomic backgrounds, versus 33% of the wider workforce.

These statistics are a risk that AI systems built by homogeneous teams will be inadequately tested for the needs and contexts of the majority of the population.

Addressing the six persistent barriers

Skills England has identified six systemic barriers to AI skills development in England, based on expert evidence from learners, employers, and training providers.

Inconsistent use of the term 'AI skills'

There is no shared understanding of what counts as an AI skill. In some workplaces, the term refers to using tools such as ChatGPT or Copilot to complete everyday tasks. In others, it refers to more complex knowledge such as understanding bias, data protection or ethical decision-making. This lack of clarity makes it difficult for training providers to design courses, for learners to explain what they have learned and for employers to know what training they need. As a result, people often complete AI courses without being able to describe their skills, and employers hesitate to invest in training because they are unsure what it should cover.

Lack of digital literacy

Many learners struggle with basic digital tasks such as typing, saving files, managing passwords or using a web browser. These challenges are especially common among older adults, people with learning differences and those with limited access to devices.

When AI courses assume these basic skills, learners often disengage early because they cannot complete simple steps such as logging in or navigating a platform. This gap between learner readiness and course design prevents people from progressing and reinforces digital exclusion.

Fragmentation across the training ecosystem

AI training is offered by many different organisations, including colleges, libraries, universities, online platforms, employers and community groups. Although this creates a wide range of options, the system is poorly coordinated and difficult for learners to navigate. In some areas, training is duplicated or disconnected from local job opportunities; in others, it is limited or absent. Promising local projects often fail to grow because they are not linked to wider systems or long-term funding. Learners frequently struggle to understand which courses are recognised, how to progress or where to go next.

Curriculum responsiveness to emerging AI needs

Training providers find it difficult to update their courses quickly enough to keep pace with rapid developments in AI. Annual review cycles, strict approval processes and limited staff time mean that new tools and examples take too long to be introduced. As a result, learners often encounter outdated content that does not reflect current practice. Educators are keen to update their teaching, but the systems around them, such as governance, validation processes and institutional structures, are not flexible enough to support rapid change. This creates a gap between real-world AI use and what learners are taught.

Costs and funding

Many AI training programmes rely on short-term grants or temporary funding. This leads to stop-start provision, making it difficult for providers to retain skilled staff and disrupting learner progress. Even when courses are free, people still face indirect costs such as travel, childcare, broadband, and device access. These practical barriers prevent many learners from taking part in or completing training. The lack of stable funding also makes it harder for organisations to plan long-term provision or build trust with communities.

Limited employer understanding of workforce AI skills needs

Many employers, particularly small and medium-sized businesses, are unsure what AI means for their organisation or what skills their staff need. Even when businesses use AI tools, they often do not know what type of training would help staff use them safely and effectively. This uncertainty leads to limited investment in training and an overreliance on vendor-led courses that focus on specific tools rather than broader skills. Employers often lack frameworks to assess their needs or plan how AI will affect roles and teams. Without clear guidance, organisations miss opportunities to improve productivity and support staff development.

If AI tools perform the tasks that junior professionals previously learned from, the pipeline narrows

> The findings throughout this report reflect a consistent theme we are seeing across both public and private sector organisations. The conversation is rapidly moving beyond AI experimentation towards the practical challenge of deploying AI safely, responsibly and at scale. As organisations increasingly adopt agentic AI, the focus shifts from individual models to the governance, assurance and resilience of the agents acting on behalf of the business. Maintaining an agnostic approach to AI will be critical, enabling organisations to retain flexibility and avoid dependency on any single technology provider, while robust agent assurance frameworks will be essential to ensure transparency, accountability and trust. As AI becomes embedded within business critical operations, resilience and continuity must be designed in from the outset, ensuring organisations can innovate with confidence while maintaining control.
>
> **Jon Bridges, CIO, Exponential-e**

Getting started

The most common mistake organisations make when approaching AI is starting with the technology rather than the problem. They ask 'how do we adopt AI?' when they should be asking 'what specific challenge are we trying to address, and could AI help us address it better than our current approach?' A structured, problem-first approach not only improves outcomes for the business and the success of the AI roll-out, but also supports compliance by ensuring that AI deployment is proportionate, justified, and subject to appropriate governance from the outset.

AI is a broad field with many different applications. Each field of AI has distinct capabilities, limitations, and implications for the people affected by its actions. Choosing whether to use AI, and which type, only makes sense once you understand the problem you are trying to solve, the outcomes you want to achieve, the data you have available, and the people who will be affected.

The problem-oriented starting questions

Before evaluating any AI tool, an organisation should be able to answer: What specific challenge are we facing, with enough specificity that we could measure whether a solution had worked? What does a good outcome look like, and how would we know if we had achieved it? What data do we have that is relevant, and is it representative, accurate, and usable? Who are the people affected, and what do their interests require? And what are the acceptable and unacceptable failure modes? If the system gets things wrong, what are the consequences, and for whom?

Buy, build, or partner: choosing how to adopt AI

Once an organisation has a clear problem to solve, the next decision is how to access AI capability. There is no universally right answer, and the choice depends on the complexity and sensitivity of the problem, the organisation's internal capability, available budget, and the degree of control and customisation required. Three broad routes are available.

Buying off the shelf means adopting an existing AI product or platform, whether a general-purpose tool such as Microsoft Copilot, a sector-specific application, or a specialist tool available via a marketplace. This is the fastest and most accessible route for most organisations, and it currently accounts for the majority of AI adoption across UK businesses. The trade-off is limited customisation, reliance on the vendor's data practices and update cycles, and the risk of over-reliance on tools not fully suited to specific organisational needs. Organisations should review vendor contracts, paying attention to data processing terms, liability provisions, and compliance with applicable regulations.

Building internally means developing a bespoke AI system or fine-tuning an existing foundation model on proprietary data to meet a specific organisational need. This offers the greatest control over data, design, and outcomes, and is increasingly the route taken by larger organisations with dedicated technical capability. It also carries the highest cost, the longest timelines, and the greatest internal demand for AI and data engineering expertise. It is rarely appropriate for smaller organisations without dedicated technical teams, and the absence of those skills is a recognised barrier.

Partnering means working with an external organisation, such as a technology provider, systems integrator, academic institution, or sector peer, to develop or deploy AI capability jointly. A partnership can provide access to skills, data, and infrastructure that an organisation could not easily obtain alone, while sharing risk and cost. It is particularly relevant when a use case is sector-specific and no off-the-shelf product exists, or when building internally would require disproportionate investment.

The choice between these routes is not fixed. Many organisations begin by buying off-the-shelf tools, develop internal expertise over time, and selectively partner where specialist capability is needed. What matters is that the decision is made deliberately, based on the nature of the problem, an honest assessment of internal capability, and a clear view of the risks each route introduces.

Principles for responsible AI adoption

Regardless of whether an organisation buys, builds, or partners, adoption must be guided by principles that actively mitigate the risks AI introduces and support ongoing compliance. This can be modelled on the UNESCO Recommendation outlined earlier in the report. Applying these principles systematically - at procurement, deployment, and throughout the AI lifecycle - provides both a practical governance framework and an evidence base for demonstrating responsible use to regulators, clients, and other stakeholders.

Organisations often start with technology instead of the problem

Fairness and non-discrimination

AI systems must not produce unjustifiably different outcomes for different groups of people. This is the primary mitigant for bias and discrimination risk. It also mitigates employment risk: the disproportionate impact of AI on junior roles, entry-level workers, women, and those from lower socioeconomic backgrounds requires deliberate intervention to prevent AI from widening existing inequalities. Fairness requires active testing across demographic groups and diverse teams in the design, testing, and governance of AI systems.

Safety and security

AI systems must be reliable, tested in context, and protected against attack. This is the primary mitigant against cyber security risks, including specific threats such as prompt injection, adversarial manipulation, deepfake impersonation, and open-source models with safety guardrails removed. It also addresses misinformation risk: AI tools used to generate or process information must be subject to quality controls sufficient to prevent harmful or false outputs from causing real-world damage. Before deployment, organisations should verify that AI tools have been tested in their specific environment, not just on benchmark datasets, and that cyber defences have been updated to address AI-specific threat vectors.

Right to privacy, data protection and IP

AI adoption must not compromise individuals' right to control their personal data. This principle directly mitigates data privacy risks, including unauthorised data use in AI training, shadow IT behaviours where employees upload sensitive data to unapproved tools, and the legal reach of cloud infrastructure under foreign jurisdictions such as the US CLOUD Act. It also addresses intellectual property risk: organisations must ensure that both the data used to train or fine-tune AI and the outputs generated by AI do not infringe existing rights. Data Processing Agreements, enterprise-tier tooling, and clear internal policies on data handling are the practical instruments through which this principle is implemented.

Sustainability

AI adoption must account for its environmental impact. This is the primary mitigant for environmental harm: the energy and water consumption of AI data centres is substantial and growing. Organisations making procurement decisions about AI tools should consider whether vendors provide environmental impact data, whether they have credible net-zero commitments, and whether the scale of contemplated AI use is proportionate to the environmental costs it imposes.

Proportionality and do no harm

AI should only be used when the benefits are proportionate to the risks introduced and when it does not cause unnecessary harm to individuals, communities, or the environment. In practice, this means asking whether AI is the right tool at all, not simply whether it is available. This principle directly mitigates risks of bias and discrimination, employment harm, and existential or catastrophic misuse: it requires organisations to weigh who bears the costs of AI adoption alongside who receives the benefits, and to avoid deploying AI in high-stakes contexts where the consequences of failure are severe and the case for AI is weak.

Human oversight and determination

AI systems must remain subject to meaningful human oversight, particularly in consequential decisions. This is not simply a matter of having a human in the loop, as rubber-stamping AI outputs without genuine review is not oversight. This principle directly mitigates accountability and existential risk, and is especially important as agentic AI systems take on longer, more autonomous tasks. It also mitigates explainability risk: a human reviewer who genuinely understands what an AI system is doing and why is better placed to identify errors before they cause harm. Organisations should ask not only whether a human reviews AI outputs, but whether that person has the knowledge, authority, and time to exercise real judgement.

Awareness and literacy

AI adoption requires that the people deploying and affected by AI systems understand what those systems are, what they can and cannot do, and what responsibilities they carry. This principle mitigates bias and discrimination risk as people without AI literacy are more susceptible to AI-generated false content, and employment risk as workers without AI skills are more vulnerable to displacement and less able to benefit from the productivity gains AI enables.

Multi-stakeholder adaptive governance and collaboration

AI governance should not be confined to technology teams or senior leadership. Effective oversight requires the involvement of those who use AI systems, those affected by their outputs, and those with expertise in the relevant regulatory and ethical context. This principle mitigates accountability risk and the fragmentation of responsibility that occurs when AI decisions affect people, but no single individual or body is answerable for them. It also mitigates employment risk by ensuring that workers' perspectives are included in decisions about AI deployment that affect their roles. Governance structures must be adaptive: the regulatory environment, the capability of AI systems, and the risks they introduce are all changing rapidly, and static governance frameworks will quickly become inadequate.

Responsibility and accountability

Someone must be clearly responsible for each AI system, including its design, its outputs, its effects on people, and what happens when it goes wrong. This principle is the primary mitigant for accountability risk. The legal consequences of misplaced reliance are already evident, and in clinical, financial, and public-sector contexts, accountability gaps expose organisations to regulatory, legal, and reputational risks. Responsibility must be assigned, visible, and backed by genuine authority to act, including the authority to pause or stop a system that is causing harm.

Transparency and explainability

People affected by AI decisions should know that AI is involved, and organisations should be able to explain how those decisions are reached. This mitigates the explainability and transparency risk: the black-box problem that makes it difficult to identify errors, detect bias, or challenge flawed decisions. It also mitigates the risk of misinformation by ensuring that AI-generated content is identifiable as such.

AI adoption must be guided by principles that actively mitigate risks

AI adoption: workforce and skills

AI adoption is a large-scale transformational decision which also affects the workforce. The evidence from across the Everyday AI forums and the UK research reviewed in this report points to a consistent gap between the pace at which organisations deploy AI tools and the pace at which they invest in the human capability needed to use them well, govern them responsibly, and adapt to the changes they bring.

Organisations adopting AI should treat workforce and skills as a first-order consideration, not an afterthought. This means addressing four interconnected questions from the outset.

Who needs to understand what?

Not everyone needs to be able to build AI. Still, everyone working alongside it needs to understand what it is, what it cannot reliably do, and what their responsibilities are when it produces an output they must act on. The Alan Turing Institute's AI Skills for Business framework usefully distinguishes four levels: AI citizen (awareness of what AI is and does), AI worker (using AI tools competently in a specific role), AI practitioner (configuring, deploying, or evaluating AI systems), and AI builder (designing and developing AI capability). Mapping these levels to existing roles is a practical starting point for identifying where skills gaps are most consequential.

What is the impact on roles and career pathways?

The most consistent finding across recent UK research is that AI's early workforce impact is concentrated in junior and entry-level positions: the tasks through which people traditionally develop professional judgement and progress into more senior roles. Organisations should be explicit about how AI changes role expectations, where tasks will be automated, and how they intend to support career development in an environment where some traditional pathways are narrowing. This is both an ethical and a practical obligation, as hollowing out entry-level roles without addressing the pipeline implications creates long-term capability risk for the organisation.

How will training be delivered, and will it reach those who need it most?

Training strategies must be designed to reach the staff who most need support, not only those already motivated to engage with AI independently. Generic online modules that certify completion rather than competence are not adequate in contexts where AI errors carry real consequences.

How will the organisation manage the human oversight required by AI?

Meaningful human oversight of AI isn't passive. It requires the knowledge to assess outputs critically, the authority to raise concerns, and the time to apply real judgement. As AI takes on more tasks, human roles shift towards review, exception handling and governance, work that still relies on expertise and sound judgement. Organisations must plan for this shift, ensuring that the staff responsible for oversight have the skills and capacity to exercise it effectively. The emergence of agentic AI, capable of pursuing multi-step goals autonomously, makes this planning more urgent: oversight of systems that act independently across longer time horizons demands a lot from the humans responsible for them.

In summary, the route to adopting AI is as follows: define a problem you can measure, choose an adoption approach (buy, build, or partner) that aligns with your capacity and requirements, and apply clear principles so that responsible adoption is at the core. Workforce readiness should be treated as part of delivery because the quality of oversight, training, and role design will ultimately determine whether AI creates value or creates avoidable harm.

In practical terms, organisations can move forward now by:

- nominating accountable owners (senior sponsor and operational lead)
- selecting one low-risk, high-frequency workflow and run a time-boxed pilot with clear success metrics and red lines
- setting minimum controls: approved tools, data handling rules, access permissions, logging, an escalation route for incidents, and clear guidance on when to seek compliance advice
- training the people who will supervise outputs, not just the people who will use the interface (including how to verify, challenge, and document decisions)
- reviewing and iterating: keep what works, retire what does not, and expand only when you can evidence value and manage the risk.

> AI isn't slowing down for anyone. What keeps me up at night about AI? The decisions being made around it are often made without the people most affected having any say. We still have time to build this properly. Not perfectly, but properly. This report asks the uncomfortable questions most organisations are avoiding, and that's exactly why it matters.
>
> **Kate Bennett, CEO, Disruptive Live**

Conclusions

The Everyday AI programme set out with a straightforward question: what does artificial intelligence adoption actually look like in organisations across the UK? The answer, drawn from forums in healthcare, education, financial services, agriculture, and social care, from the CyNam and HCR Law Everyday AI Survey 2026 and wider industry research, is more nuanced than either AI optimism or AI scepticism. What emerges clearly is that successful AI adoption requires not only technical capability but also robust governance, clear accountability, and proactive compliance management.

The adoption gap is real, and it matters

The national picture is stark. 80% of UK businesses still view AI as irrelevant to their work, yet among the organisations engaged with through this programme, usage is intensive, routine, and growing. The gap between these two realities is not primarily one of access or cost. It is a gap in awareness, confidence, and the internal capacity to navigate an increasingly complex landscape. Closing this gap requires sustained, practical, and evidence-based engagement of the kind the Everyday AI programme has sought to provide.

Understanding AI is a precondition for using it well

Across every forum, participants found themselves operating in environments where AI was already present without always recognising it as such. The report's opening focus on definitions is not academic. Organisations that cannot identify the AI they use cannot govern it, assess whether it is working as intended, or act responsibly toward the people it affects. AI literacy is not a technical specialism. It is a baseline organisational capability that must be deliberately built and maintained over time.

Sector evidence reveals the specificity of AI's value and risk

The sector forums confirmed what national statistics obscure: AI's value is context-specific. In financial services, the gains are real in fraud detection and compliance efficiency, but so are the risks of over-reliance in regulated decision-making. In health and social care, a saturated market demands careful selection and strong governance before any AI system touches patient outcomes. In education, AI is arriving faster than institutional frameworks can accommodate, with governance too often delegated downward rather than owned by senior leadership. In agriculture, measurable operational value is already present in animal welfare, precision spraying, and environmental compliance. Yet, interoperability, trust, and data ownership remain practical barriers to the full realisation of that value.

These differences matter. One of the clearest findings across the programme is that organisations making genuine progress with AI are typically those that began with a specific, well-defined problem, rather than with the technology itself.

Ethics is not an obstacle to adoption; it is a condition of it

The UNESCO framework and the principles common to responsible AI literature should be used as structured decision-making tools. Applied at the point of procurement, design, and deployment, principles such as fairness, transparency, accountability, and human oversight help organisations identify who bears the costs of AI alongside who receives the benefits, and ensure that consequential decisions remain understandable and challengeable.

The risks are real

Forum participants were not primarily worried about existential AI risk. They were worried about wrong clinical outputs, deepfake impersonation, discriminatory recruitment tools, and employees uploading sensitive data to consumer-grade systems they had never been trained to use. These risks are well evidenced, are already causing documented harm in UK organisations, and are addressable through practical governance. The legal consequences of inadequate oversight are real.

The skills challenge is systemic

AI is not simply changing which tasks organisations need people to perform. It is restructuring the pathways through which people develop professional competence. The hollowing out of junior and entry-level roles threatens the workforce pipelines through which organisations develop their own long-term capability. A 56% global wage premium for AI-skilled workers and an estimated £400 billion cost of the UK's AI skills gap describe conditions shaping the labour market today. Training strategies that reach only the already-motivated will not close this divide. The six structural barriers identified by Skills England require systemic responses.

Geopolitics is part of the AI landscape for every organisation

The concentration of AI capability in a small number of global supply chains, the legal reach of foreign jurisdictions over cloud-hosted data, and the physical vulnerability of the infrastructure on which AI depends are operational risks for any organisation using AI tools at scale. Understanding those dependencies and making deliberate choices about which tools, providers, and data-handling arrangements are appropriate is part of responsible AI adoption for organisations of all sizes.

What this programme has demonstrated

The consistent finding across all four sectors is that the organisations making genuine progress with AI are not necessarily those with the most sophisticated tools. They are the ones that began with a specific, well-defined problem; invested in the human capability needed to supervise AI outputs; and built governance structures that kept responsible humans genuinely in the loop. The technology is not the hard part. The organisational readiness, including the ability to ask the right questions, assign clear accountability, manage compliance obligations, and make considered decisions about when and how AI should be deployed, lies in the real work.

CyNam and HCR Law will continue to support organisations across the UK in navigating this landscape. The questions raised in these forums, about what AI is, what it can and cannot reliably do, and how to adopt it in ways that are genuinely good for the people it affects, do not have simple answers. Still, this report has provided some answers grounded in the experiences of practitioners working through these challenges in real time. That evidence base and the community of practice that generated it are themselves among the most valuable outcomes of the Everyday AI programme.

Successful AI adoption requires technical capability, governance, accountability, and proactive compliance

Glossary

The following terms are used throughout this report to cover the full landscape of AI, including all fields shown in the AI landscape diagram shared across the Everyday AI forums. All definitions are written for a professional audience without a specialist AI background.

A

Agentic AI

AI systems that go beyond responding to prompts to autonomously pursue goals: planning sequences of actions, using tools such as web browsers and databases, and adapting based on outcomes. Raises distinct questions about accountability and human oversight.

Artificial general intelligence (AGI)

A theoretical system capable of performing any intellectual task a human being can, with the same flexibility and breadth. No AGI system exists today.

Artificial narrow intelligence (ANI)

The only form of AI that currently exists. ANI systems are built for specific tasks and cannot transfer their ability to unrelated domains. A system that reads chest X-rays cannot write code. Every AI product on the market today is ANI.

Artificial intelligence (AI)

The broad term for systems that can perform tasks typically requiring human intelligence. An umbrella term covering many different technologies.

Artificial superintelligence (ASI)

A theoretical system that surpasses human intelligence in every domain. Purely speculative and not a near-term practical concern for organisations. The subject of long-range safety debates among AI researchers.

B

Bias (in AI)

Systematic errors in AI outputs that produce different results for different groups of people. Typically arises from training data that reflects historical inequalities, or from proxy variables that encode protected characteristics without naming them.

C

Computer vision

AI that enables machines to interpret visual information from images or video.

D

Deepfake

AI-generated synthetic media that depicts people saying or doing things they did not say or do.

Deep learning

A subset of machine learning that uses neural networks with many layers. Each layer learns increasingly abstract representations of the data. Underpins image recognition, speech recognition, and language generation.

E

Ethical AI

The design, development, and deployment of AI systems in accordance with ethical principles including fairness, transparency, accountability, and human wellbeing.

Explainability

The degree to which an AI system's outputs can be understood and explained in terms that a human can follow.

F

Fairness (in AI)

The requirement is that AI systems do not produce unjustifiably different outcomes for different groups. Multiple technical definitions exist that cannot always be simultaneously achieved.

Requires deliberate design, diverse development teams, and ongoing monitoring.

Foundation models

Large AI models trained on broad, diverse datasets can be adapted to many different tasks. LLMs are one type; others process images, audio, or multiple modalities.

G

Generative AI

AI systems that create new content in response to prompts. The most visible category of AI to most users today.

H

Hallucination

When an AI model generates an output that is statistically plausible but factually incorrect. Not a malfunction, as the model is doing exactly what it was trained to do (predict likely outputs).

L

Large language models (LLMs)

A category of foundation model trained on vast quantities of text. LLMs learn to predict the most likely next word fragment and, by doing so at an enormous scale, develop wide-ranging language capabilities. Examples: GPT-4, Claude, Gemini, Llama.

M

Machine learning)

A family of techniques in which AI systems learn patterns from data rather than following pre-programmed rules - the basis of most modern AI applications.

Multimodal AI

AI systems can process and generate multiple types of data within a single integrated system. Most frontier AI models are now multimodal.

N

Natural language generation (NLG)

The NLG component focused on production: generating coherent, contextually appropriate text or speech.

Natural language processing (NLP)

The broad field of AI enables computers to understand, interpret, and generate human language.

Natural language understanding (NLU)

The NLU component focused on comprehension: understanding the meaning, intent, and context of text or speech.

Neural networks

A machine learning architecture loosely inspired by the brain: interconnected layers of artificial nodes that process and transform data, adjusting connections during training until reliable at a task.

P

Privacy

In the AI context, the right of individuals to control how their personal data is collected, used, and processed, including in AI training and AI-driven decisions.

Prompt

The instruction, question, or input given to a generative AI system. The content, structure, and framing of a prompt significantly affect the quality and relevance of the output.

Prompt engineering

The practice of designing prompts to elicit more accurate, relevant, or useful outputs from AI systems. An emerging professional skill, particularly valuable in applications requiring consistent AI performance in constrained contexts.

R

Responsible AI

Using AI in a way that is genuinely honest, fair, and accountable, and building the practical habits that make that possible consistently. A commitment to doing right by the people AI affects, expressed through how decisions about AI are made and reviewed every day.

S

Security (in AI)

Protection of AI systems against attack, manipulation, and unauthorised access. AI-specific risks include adversarial inputs, prompt injection, model theft, and training data poisoning.

Speech recognition

AI that converts spoken language into text. Used in virtual assistants, transcription tools, clinical documentation AI, and accessibility applications.

T

Training

The process by which an AI model learns from data: large quantities of examples are fed through the system, and internal parameters are adjusted to minimise errors. Computationally intensive and energy-consuming. Once trained, a model is essentially static until retrained or fine-tuned.

Transparency (in AI)

The property of being open about an AI system's nature, operation, and limitations to users and to people affected by it. Distinct from explainability: a system can be transparent about its existence and purpose without being fully explainable in its internal workings.

“Frank is simply brilliant. He has been with us since the beginning and always takes time to understand what we need and how best to deliver for us. If you care about your business, you definitely want Frank by your side. He'll care about it too - and deliver for you.”

Rafah Knight, CEO and Founder, Secure AI

Scan here for further information and supporting materials.

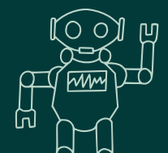

Printed on 100% recycled paper.

hcrlaw

Our Technology and Innovation team supports organisations navigating digital transformation, acting as an innovation partner to help turn ideas into practical outcomes with clear, commercially focused advice. Find out more and get in touch through our website.

www.hcrlaw.com