



AI Across Industries & Sectors

(AIStatus 2)

Where and How AI Is Being Applied Across Economic and Societal Sectors

AI Sourced Facts (AISF) Pte. Ltd.

© AI Sourced Facts (AISF)

All rights reserved.

Edition: Version 1.0

Publication Date: September 2026

Preface

Artificial intelligence is now being applied across a remarkably wide range of economic and societal activity.

It is assisting clinicians and researchers, teachers and students, bankers and insurers, manufacturers and engineers, retailers and travellers, farmers and scientists, public servants and emergency responders. It is being incorporated into telecommunications networks, professional services, creative production, buildings, transport systems, energy infrastructure and many other environments.

But the extent and nature of that use differ substantially.

Some applications have been operating for years and form part of established systems. Others are expanding rapidly as generative AI becomes incorporated into professional, commercial and public-service activity. Some have reached genuine operational deployment but only in particular organisations, countries or circumstances. Others remain pilots, demonstrations, research projects or developing capabilities.

These differences can make the present position difficult to interpret.

A technically impressive system may not yet be routinely deployed. A major organisation's implementation may not represent its sector. A technology available to millions of users may not be used extensively by them. A pilot may demonstrate capability without establishing effectiveness at scale. And an application that is established in a highly digitised economy may face very different conditions elsewhere.

AIStatus has been developed to provide a structured view of what is happening at a defined point in time.

This publication examines artificial intelligence through the **Industries & Sectors lens**. It considers where and how AI is currently being applied across sixteen important areas of economic and societal activity.

The purpose is not to predict what these sectors will become, recommend how organisations should adopt AI, or determine whether particular applications are desirable.

It is to establish, as coherently as the available evidence permits, **what is currently happening**.

The Status presented in this publication is assessed **as at 30 August 2026**.

Purpose of This AIStatus

AIStatus — Industries & Sectors is designed to help readers **DISCOVER where and how artificial intelligence is currently being applied across economic and societal sectors.**

The publication examines sixteen Industries & Sectors:

1. Healthcare & Life Sciences
2. Education & Training
3. Financial Services
4. Manufacturing & Engineering
5. Retail & E-commerce
6. Hospitality & Tourism
7. Transport & Logistics
8. Construction & Real Estate
9. Energy & Utilities
10. Agriculture & Food
11. Government & Public Services
12. Science & Research
13. Media, Entertainment & Creative Industries
14. Professional Services
15. Telecommunications
16. Security & Emergency Services

Each is examined independently but through a common set of questions. This allows readers to explore a particular sector on its own while also comparing developments across the wider economic and societal landscape.

The publication concentrates on the **current state**.

It distinguishes, where the evidence permits, between what AI systems can demonstrably do, what is actually being deployed in sector environments, how established adoption appears to be, and what remains experimental, limited or uncertain.

This is not a forecast of the future of industries or employment. It is not a technology-selection guide, implementation manual, consultancy report, investment analysis or set of recommendations.

It is a point-in-time examination of the position as it can reasonably be understood from available evidence.

Global Perspective & Neutrality Statement

Artificial intelligence is being developed and applied across countries with very different economic structures, infrastructure, institutions, regulatory environments, skills, languages and technological capabilities.

Its adoption is consequently not uniform.

A large digital retailer does not begin from the same position as a small independent retailer. A highly instrumented factory differs from a manufacturing operation with limited digital infrastructure. A national education system with extensive connectivity differs from one where many schools remain offline. A sophisticated electricity network, hospital system, research institution or telecommunications operator may have capabilities unavailable in lower-resource environments.

This publication therefore does not assume that developments in any single country or region represent the global position.

International evidence and examples are included where they materially improve understanding of the Status. The sixteen chapters contain evidence from multiple regions and from economies at different stages of technological and institutional development.

At the same time, the available public evidence is itself uneven. Some countries, industries and large organisations are documented much more extensively than others. Lack of comparable evidence from a region should not automatically be interpreted as absence of AI activity there.

A genuinely global perspective therefore requires two disciplines: **actively seeking meaningful international evidence, while not manufacturing geographical symmetry where comparable evidence does not exist.**

AIStatus is neutral, educational and non-prescriptive.

It does not rank sectors according to whether they are adopting AI quickly enough. It does not advocate automation. It does not recommend particular technologies, providers or approaches.

Its purpose is to help readers understand the current position and the qualifications necessary to interpret it.

Editorial Standard

The publication has been developed around several recurring evidential distinctions.

Technical capability is not the same as real-world effectiveness. A system may perform well on a defined task without demonstrating that it improves the wider activity in which the task sits.

Availability is not the same as adoption. A capability incorporated into a widely available platform may potentially reach millions of organisations or users without being actively used by all of them.

A pilot, trial or demonstration is not the same as routine deployment. Experimental evidence can establish that something is possible without establishing that it has become normal practice.

A prominent implementation is not evidence of sector-wide adoption. Leading organisations often possess exceptional data, capital, infrastructure and technical capability.

Regulatory approval, policy support or institutional investment is not the same as adoption. These may demonstrate readiness, legitimacy, direction or development without showing how extensively a technology is actually used.

Use is not the same as impact. A system can have many users without independently established evidence that it improves productivity, learning, safety, service quality, profitability, scientific progress or other outcomes.

AI is not synonymous with digitisation or automation. Robotics, sensors, digital twins, connected equipment, smart infrastructure and automated systems may incorporate AI, but they are not necessarily AI themselves. The distinction is particularly important when interpreting adoption claims.

These distinctions are applied throughout the publication.

How to Use This AIStatus

Readers seeking an overall picture of AI across economic and societal sectors may read the sixteen chapters sequentially.

Readers interested in a particular sector may go directly to the relevant chapter. Each has been designed to stand independently and provide sufficient context without requiring the reader to consult the other sectors.

To provide consistency and comparability, every sector is examined through the same seven questions:

1. Where is AI currently being applied?

Identifies important current areas of application and use.

2. What problems or needs is AI being used to address?

Examines the practical pressures, problems, needs or opportunities towards which AI is being directed.

3. What can AI currently do?

Describes demonstrated current capabilities while distinguishing technical capability from wider real-world effectiveness where necessary.

4. How established is adoption?

Considers actual adoption and distinguishes, where useful, between established deployment, growing adoption, emerging use, pilots, trials and experimentation.

5. What appears to be driving adoption?

Examines significant evidenced factors contributing to development, testing and use.

6. What are the present limitations?

Identifies technological, organisational, data, infrastructure, reliability, integration, skills, governance, regulatory, economic, evidential and other constraints evident at the Status date.

7. What you should be aware of

Provides important context and qualifications needed to interpret the Status, including differences between capability and deployment, the role of people, accountability, geographical variation and limitations in the available evidence.

Each sector concludes with **Evidence Notes** identifying principal sources supporting its Status assessment.

The common questions make comparison possible. They do not imply that AI is developing at the same rate, through the same technologies or under the same conditions in every sector.

Those differences are part of the Status.

Version & Governance

AIStatus 2 — AI Across Industries & Sectors

Where and How AI Is Being Applied Across Economic and Societal Sectors

Edition: Version 1.0 / 2026 Edition

Status Reference Date: 30 August 2026

Publication Date: September 2026

Series: AISF AIStatus

Published by: AI Sourced Facts (AISF) Pte. Ltd.

AIStatus publications provide point-in-time examinations of a changing artificial intelligence landscape.

The Status contained in this edition has been assessed using information available up to and including **30 August 2026**. Developments occurring after that date do not form part of the Status assessment for this edition.

Future editions may reassess the subject at a later reference date. Such editions may reflect changes in technologies, applications, sector deployment, adoption, evidence and the structure required to describe the position adequately.

An earlier edition remains a record of the Status assessed at its stated reference date.

Editorial oversight of the AISF AIStatus series is maintained by AI Sourced Facts (AISF).

Table of Contents

Preface**Purpose of This AIStatus****Global Perspective & Neutrality Statement****Editorial Standard****How to Use This AIStatus****Version & Governance****Introduction — Understanding the Industries & Sectors Status****Industries & Sectors**

1 — Healthcare & Life Sciences

2 — Education & Training

3 — Financial Services

4 — Manufacturing & Engineering

5 — Retail & E-commerce

6 — Hospitality & Tourism

7 — Transport & Logistics

8 — Construction & Real Estate

9 — Energy & Utilities

10 — Agriculture & Food

11 — Government & Public Services

12 — Science & Research

13 — Media, Entertainment & Creative Industries

14 — Professional Services

15 — Telecommunications

16 — Security & Emergency Services

The Status Across Industries & Sectors**AIStatus — Continuing the Exploration****Future Status Updates & User Submissions****Acknowledgements****AISF Universal Disclaimer****About AI Sourced Facts (AISF) Pte. Ltd.****Publication & Governance Information****Back Page**

Introduction — Understanding the Industries & Sectors Status

AIStatus — Industries & Sectors examines where and how artificial intelligence is currently being applied across economic and societal sectors.

It is designed as a **point-in-time Status**: a structured view of what can reasonably be evidenced as at 30 August 2026, rather than a prediction of where industries are heading or a prescription for how organisations should respond.

The sixteen sectors cover a broad range of human activity—from healthcare, education, finance and manufacturing to agriculture, science, telecommunications and emergency services.

They should not, however, be understood as sealed compartments.

The same AI capability may legitimately appear in several sectors. Computer vision can inspect manufactured products, monitor construction sites, grade food, analyse medical images and support security operations. Forecasting can assist retailers, electricity networks, transport operators, farmers and hospitality businesses. Generative AI can support teachers, public servants, scientists, lawyers, engineers and creative professionals.

What changes is the **sector context in which the capability is being applied**.

This publication therefore does not attempt to allocate every application exclusively to one chapter. Where overlap is necessary to understand what is happening within a sector, the application is examined from that sector's perspective.

The common seven-question structure provides consistency without implying uniformity.

That distinction matters because one of the clearest findings across the sixteen chapters is that **there is no single maturity level for AI across industries and sectors**.

Some applications are already established.

Machine learning has long supported recommendation and personalisation in digital commerce and media; fraud and risk analysis in financial services; forecasting and optimisation in networks and operations; computer vision in selected industrial settings; and specialised analysis across parts of science and healthcare.

Other applications are moving rapidly from experimentation into operational use.

Generative AI is increasingly supporting document-intensive, information-intensive and communication-intensive work across professional services, government, education, healthcare, telecommunications, media, retail and other sectors.

Still other applications remain substantially earlier.

More autonomous systems that undertake consequential actions, operate in difficult physical environments, control complex infrastructure or perform extended professional or scientific processes are increasingly visible, but their maturity is generally lower and more conditional.

Telecommunications provides examples of operational closed-loop and agentic systems, yet much AI-native networking and AI-enabled 6G remains in development or trial. Security & Emergency Services has genuine operational deployments, but the strongest present role remains specialised sensing, analysis and decision support inside human-led systems rather than autonomous crisis management.

The evidence also shows that **adoption itself has several dimensions**.

Education provides a particularly clear example. Individual learner use can be extremely high while institutionally provided, educationally validated and consistently governed adoption remains considerably less mature.

Retail demonstrates another distinction. A platform may make AI-enabled commerce infrastructure available at enormous scale, but technical availability does not establish equivalent merchant utilisation or transaction volume.

Government provides a third. AI activity has been reported across most OECD governments, but geographical breadth does not demonstrate uniform depth of use across their agencies or activities.

Consequently, **reach, availability, individual use, organisational deployment, institutional adoption, demonstrated effectiveness and sector prevalence are different measures**. They should not be collapsed into a single concept of “AI adoption”.

Sector conditions also matter.

AI generally develops more readily where relevant activity already produces accessible digital information. Large digital retailers possess extensive transaction and behavioural data. Telecommunications operators continuously generate network data. Modern factories may contain sensors and digitally controlled equipment. Scientific disciplines with machine-readable datasets and computationally evaluable outcomes can provide particularly favourable environments for AI.

Where those foundations are weaker, adoption can be more difficult.

Infrastructure, connectivity, computing resources, data availability, institutional capability, capital and skills differ markedly across and within countries. Even education illustrates the underlying constraint: available UNESCO evidence indicates that only about 40% of primary schools and 50% of lower-secondary schools globally are connected to the internet.

The physical environment creates another important distinction.

Generating text or analysing a digital record is different from controlling machinery, navigating roads, manipulating irregular objects, responding to emergencies or operating critical infrastructure. Physical conditions change. Sensors can fail. Weather, people, unusual objects and rare events introduce uncertainty.

Manufacturing evidence illustrates why adaptable physical AI remains harder than bounded repetitive automation. The same underlying qualification appears in transport, agriculture, construction and emergency response.

Human involvement therefore remains a significant feature of the current Status.

That involvement takes different forms.

Clinicians review medical information. Engineers validate designs and recommendations. Teachers determine educational objectives. Financial professionals retain accountability for regulated decisions. Scientists establish whether results constitute reliable knowledge. Creative professionals select and edit material. Public institutions retain legal responsibility for governmental decisions. Emergency personnel act on information supplied by analytical systems.

The fact that a person remains formally responsible does not mean the AI is inconsequential. A model may materially influence what evidence is seen, which cases are prioritised, what risks are identified, what recommendation is made or what action is considered.

But neither does substantial AI involvement mean that the wider professional, institutional or operational activity has become autonomous.

The chapters repeatedly demonstrate the importance of separating **the task AI performs from the larger activity in which that task sits**.

An AI system can identify a medical image feature without practising medicine. It can estimate construction progress without determining contractual entitlement. It can generate a scientific candidate without establishing scientific knowledge. It can detect a person in search-and-rescue imagery without conducting the rescue. It can produce creative material without independently directing an entire professional production.

This distinction is central to understanding current AI capability.

The evidence itself also requires interpretation.

Provider evidence can demonstrate that a system exists and has been deployed. Organisational evidence can establish real implementation. Surveys can provide useful information about reported use and attitudes. Academic studies can demonstrate capability or effectiveness under defined conditions. Regulators and governments can establish institutional and legal developments.

None of these sources, alone, necessarily establishes global sector-wide prevalence or effectiveness.

For example, advanced factories can demonstrate the industrial frontier without representing the average manufacturer. Large retail platforms can demonstrate enormous operational scale without representing smaller retailers. Scientific researcher surveys can demonstrate substantial individual use without measuring global institutional adoption.

The appropriate interpretation of AI across Industries & Sectors is therefore neither that AI remains principally experimental nor that industries have broadly become AI-operated.

The current position is more differentiated.

AI is materially present across economic and societal activity. Established specialised systems coexist with rapidly expanding generative applications, emerging agentic systems and much earlier forms of physical and operational autonomy. The depth and consequences of that use depend substantially on the sector, task, organisation, geography and surrounding human and institutional system.

That unevenness is not a problem to be removed from the assessment.

It is part of the Status.

1. Healthcare & Life Sciences

Status as at 30 August 2026

Artificial intelligence is now being applied across a substantial part of Healthcare & Life Sciences, but its maturity differs markedly from one application to another. Some relatively narrow AI systems are already part of routine clinical or laboratory workflows. Other applications are being deployed at increasing scale but remain under close human supervision. Still others—notably some uses of generative AI for clinical reasoning and parts of AI-enabled drug development—are scientifically significant but remain developmental.

The clearest current picture is therefore not one of AI replacing healthcare professionals or scientists. It is one of AI increasingly performing particular analytical, interpretive, administrative and research tasks within larger human-led processes.

1. Where is AI currently being applied?

One of the most established clinical areas is **medical imaging**. AI systems analyse X-rays, CT scans, MRI scans, mammograms, ultrasound and other images to detect, highlight, measure or classify abnormalities. Applications include cancer detection, cardiovascular disease, neurological conditions, fractures, lung disease and image reconstruction or enhancement. The US Food and Drug Administration's list of authorised AI-enabled medical devices illustrates both the scale and concentration of present development: radiology accounts for a large proportion of authorised systems, alongside a growing number of cardiovascular, neurological and other applications. Regulatory authorisation, however, does not establish that every authorised system is widely used.

AI-assisted screening is also becoming established in particular settings. An important global-health example is tuberculosis. The World Health Organization has recommended computer-aided detection of pulmonary tuberculosis from digital chest X-rays for adults since 2021, and in 2025 identified six products that met its performance standards. Such systems can substitute for human interpretation of X-rays for the defined task of TB screening, an especially significant capability where trained readers are scarce.

A second rapidly developing application is **clinical documentation**. Generative AI-based ambient systems can listen to a clinician-patient consultation and generate draft notes, summaries, letters or structured information for electronic records. NHS England has issued implementation guidance specifically for such technology. Singapore reported in April 2026 that its Note Buddy system had already supported healthcare workers across its public healthcare institutions and generated more than 80,000 medical and administrative notes. Singapore nevertheless requires healthcare professionals to review AI-generated notes before adding them to clinical records.

AI is also being applied to **clinical decision support and risk prediction**. Systems can combine information from health records, laboratory results, images and patient histories to estimate risks, identify patients who may require attention and suggest possible clinical pathways. Singapore, for example, is deploying or developing AI applications to identify

people at elevated risk of diabetes, hyperlipidaemia and chronic kidney disease and to support medication-related decisions. These systems illustrate a broader movement from retrospective analysis towards earlier identification of risk, although many remain subject to further clinical validation.

Other current applications include remote patient monitoring, analysis of signals from wearable or home-based devices, ECG interpretation, sleep monitoring, pathology and laboratory image analysis, treatment planning, patient scheduling, coding, forecasting attendance and operational resource management. The FDA's authorised device lists include AI-supported cardiovascular monitoring, wearable technologies, imaging systems and other sensor-based applications intended for clinical and home environments.

Within **life sciences**, AI is being used to analyse biological data, predict molecular and protein structures, identify potential drug targets, design molecules, prioritise compounds, analyse experimental results and support clinical development. AlphaFold provides perhaps the best-known example of a research capability that has moved beyond demonstration: the AlphaFold database provides hundreds of millions of predicted protein models for research, with major expansions into protein-complex prediction during 2026.

AI-assisted drug discovery has also progressed from computational experimentation into human clinical development. Rentosertib, whose target identification and molecular design involved generative AI, produced published Phase IIa trial results in idiopathic pulmonary fibrosis and entered a Phase III trial in China in July 2026. This is an important milestone, but it should not be interpreted as evidence that AI-designed medicines generally succeed clinically: drug candidates discovered using AI must still pass the same demanding clinical tests of safety and effectiveness as other medicines.

2. What problems or needs is AI being used to address?

Several pressures explain the concentration of present applications.

Healthcare generates quantities of information that are difficult for people to examine consistently and quickly. Imaging departments produce large numbers of scans; medical records accumulate substantial amounts of text and structured data; monitoring devices generate continuing streams of signals; and biological research produces enormous molecular and genomic datasets. AI is being used principally where computational analysis can reduce this information burden.

In diagnostics and screening, the problem is frequently **capacity combined with consistency**. An AI system can inspect images repeatedly, flag suspected abnormalities and prioritise cases without suffering fatigue. The practical attraction is particularly strong where specialist personnel are limited. WHO's use of AI-assisted chest-X-ray interpretation in TB screening illustrates how this can have relevance well beyond highly digitised hospital systems.

Documentation addresses a different problem: **clinician time**. Record keeping, summarisation and administrative work can consume substantial parts of a healthcare professional's working day. Ambient AI is being introduced to draft documentation so clinicians can spend less time typing or reconstructing encounters after consultations. A 2025 systematic review found the

emerging evidence centred particularly on efficiency, document quality and stakeholder experience, while also showing that the evidence base was still relatively small.

Risk-prediction systems address the problem that important deterioration or future disease may be difficult to identify from large, fragmented patient records. AI can search combinations of variables that would be cumbersome to process manually and help identify groups warranting attention.

Healthcare organisations are also applying AI to non-clinical bottlenecks: scheduling, predicting attendance, coding, administrative correspondence and workflow management. These may attract less public attention than diagnostic AI but potentially affect large numbers of staff and patients because they sit within everyday healthcare operations.

In life sciences, the underlying problem is different again. Drug discovery and development involves searching extremely large biological and chemical spaces while experimental testing is expensive and slow. AI is being used to prioritise where researchers should look, generate candidate structures, identify patterns across scientific information and sometimes eliminate unpromising options earlier. Regulators in Europe and the United States now explicitly recognise AI use across non-clinical research, clinical development, manufacturing and post-market monitoring.

3. What can AI currently do?

Current AI systems can perform some specialised tasks extremely well. They can classify images, identify patterns in physiological signals, measure anatomical structures, estimate probabilities, extract information from clinical records and generate structured text from conversations.

In restricted circumstances, they can perform a task previously requiring a specialist reader. WHO's guidance allowing appropriately validated computer-aided detection systems to interpret adult chest X-rays for pulmonary-TB screening is an unusually clear example. Importantly, this does not mean that the same AI can diagnose all conditions visible on the image or replace general radiological expertise.

Generative AI can now produce credible clinical drafts, summaries and correspondence and can retrieve or reorganise medical information. This makes it useful as an assistant. Its fluent output, however, should not be confused with dependable clinical judgement. WHO continues to identify inaccurate, incomplete or fabricated responses, bias and automation bias as material risks of large multimodal models used in health.

Research AI can search scientific relationships at a scale that would be impossible manually. Protein-structure systems can predict structural information useful to researchers; drug-discovery systems can help identify biological targets and propose molecules; and machine-learning models can analyse real-world or clinical-trial data. FDA reported that it had experience with more than 500 drug and biological-product submissions containing AI components between 2016 and 2023, showing that AI in pharmaceutical development is no longer merely conceptual.

What AI generally cannot do reliably is independently manage the full complexity of healthcare. A high-performing system normally operates within a defined task, population, data environment and workflow. Performance in one setting does not automatically establish performance in another.

4. How established is adoption?

Adoption should be described as **layered rather than uniform**.

Imaging represents one of the most established areas. Large numbers of regulated products exist, and health systems are deploying AI for image analysis and workflow support. Yet adoption varies by clinical speciality, institution and country. A US health-system survey published in 2025 found imaging and radiology to be the most widely deployed clinical AI category among respondents, while reported success was less uniform than deployment itself.

Ambient documentation is undergoing particularly rapid expansion. In the same US study, every responding health system reported some activity involving ambient notes, although that included development and piloting as well as deployment. England now has formal implementation guidance, while Singapore has moved beyond isolated pilots in parts of its public system. This supports describing ambient documentation as a **rapidly growing application**, rather than universally established practice.

Some specialised screening uses are more mature. AI-assisted adult TB screening has international WHO backing and relevance in high-burden settings. Other uses remain under large-scale evaluation. England's EDITH breast-screening programme, for example, was designed to test AI involving nearly 700,000 women. Its scale is important, but it remains a trial and therefore should not be presented as proof of routine national deployment.

European evidence similarly suggests movement beyond experimentation without uniform maturity. A European Commission study published in March 2026 described growing embedding of AI across clinical, operational and strategic healthcare domains. At the same time, the Commission was still consulting stakeholders during 2026 specifically about real-world impact, scalability, enabling factors and barriers.

Life-science adoption is established in research workflows but much less established if the test is whether AI has produced approved new medicines. AI is routinely useful for analysis and prioritisation, while AI-assisted candidates are only gradually progressing through clinical pipelines. Rentosertib reaching Phase III is therefore significant precisely because such clinical progression remains unusual rather than commonplace.

Globally, adoption is also constrained by differences in digitisation, infrastructure, financing, workforce availability and regulatory capacity. The current Status cannot therefore be represented accurately by extrapolating from a few highly digitised health systems.

5. What appears to be driving adoption?

The strongest driver is the combination of **rising demand and constrained human capacity**. Ageing populations, chronic disease, workforce shortages and increasing diagnostic and administrative workloads create pressure for technologies that can help existing staff process more information.

Productivity matters, but healthcare adoption is not simply a cost-reduction exercise. Clinical applications also promise earlier detection, more consistent screening, greater access to specialised analytical capability and improved ability to identify high-risk patients.

Documentation AI has an especially straightforward value proposition because reducing clerical work potentially benefits clinicians without requiring AI to make the final clinical decision. This partially explains its rapid uptake.

Data availability is another major driver. Electronic medical records, digital imaging, genomics and connected monitoring systems create the datasets on which AI depends. Singapore's experience illustrates the relationship clearly: its health ministry has explicitly linked successful AI adoption to integrated digital systems, good-quality data and a supporting governance structure.

In life sciences, scientific productivity is central. AI offers mechanisms for narrowing enormous search spaces and helping researchers decide which targets, compounds or hypotheses deserve expensive experimental work. Regulators themselves are responding to expanding AI use. In January 2026 the FDA and EMA jointly issued principles covering AI throughout medicine development, including research, trials, manufacturing and safety monitoring.

6. What are the present limitations?

The most important limitation is that **performance is contextual**. An AI model validated using one population, hospital, scanner, laboratory protocol or clinical practice environment may perform differently elsewhere. Local data quality, prevalence of disease and workflow design all matter.

Healthcare data are often incomplete, inconsistent and fragmented across incompatible systems. Even technically capable AI may therefore be difficult to integrate into actual clinical practice. WHO emphasises representative, high-quality and appropriately governed health data as fundamental to reliable AI.

Reliability remains particularly important for generative AI. Such systems can produce plausible but inaccurate statements. In clinical contexts, fluency can make errors harder rather than easier to recognise because an incorrect answer may appear authoritative.

Bias is another persistent limitation. If training data inadequately represent particular populations, performance may differ according to age, sex, ethnicity, geography, disease prevalence or socioeconomic circumstances. This is one reason locally contextualised models

are being developed in some countries rather than assuming that systems trained elsewhere will transfer unchanged.

Integration imposes practical limits. Hospitals must connect AI safely to existing records, devices and workflows; determine who sees an alert; prevent excessive false alarms; train staff; monitor performance; manage cybersecurity; and establish responsibility when outputs are wrong. Economic value also remains uncertain for many applications. In the US health-system survey, immature technology, financial concerns and regulatory uncertainty were among the barriers reported by participating systems.

Clinical evidence presents another constraint. Retrospective accuracy studies, prospective trials, regulatory authorisation and routine deployment answer different questions. A model can perform impressively in a controlled study without demonstrating that it improves patient outcomes in routine practice.

Drug development demonstrates this particularly clearly. AI may accelerate target identification or molecule design, but biological uncertainty remains. Clinical trials are not bypassed. Regulators continue to require evidence of quality, safety and efficacy regardless of whether AI contributed to developing the medicine.

7. What you should be aware of

The most important qualification is that **“AI in healthcare” is not one technology or one stage of maturity**. AI that enhances an MRI image, software that screens an adult chest X-ray for tuberculosis, a system drafting consultation notes and a general-purpose model suggesting diagnoses present very different evidential and safety questions.

Readers should therefore be cautious about headline claims that AI has “matched doctors” or “discovered a drug”. Such statements may be valid for a defined experiment or component of a process without showing that AI can replace the wider professional activity.

Human involvement remains substantial. Even where systems undertake meaningful analytical work, healthcare professionals commonly review outputs, interpret them in context and retain responsibility for patient care. Singapore's explicit requirement for clinician review of AI-generated notes is representative of this broader principle.

Regulation is also evolving and differs between jurisdictions. Some AI functions qualify as medical devices; others do not. The United States, European Union, United Kingdom, Australia, Singapore and other jurisdictions apply different legal and regulatory structures. International alignment is increasing in areas such as AI use in medicine development, but there is not a single global regulatory regime.

Privacy and cybersecurity deserve particular attention because AI often depends on highly sensitive health information. Questions extend beyond whether data can legally be processed to where models operate, whether information can be reconstructed or disclosed and whether large datasets adequately represent the populations affected.

Access may also become an important dividing line. AI could broaden access to analytical capabilities in places lacking specialists—as TB screening demonstrates—but sophisticated systems also require digital infrastructure, computing resources, suitable devices and reliable data. AI therefore has the potential both to reduce and to reinforce existing health inequalities.

Finally, the evidence base itself is uneven. Healthcare contains unusually strong incentives to publicise technical breakthroughs, promising trials and regulatory milestones. Those developments matter, but they should not be mistaken for widespread clinical use or established patient benefit.

Status at 30 August 2026: AI has moved materially into Healthcare & Life Sciences. Its strongest current position is as a specialised analytical and workflow technology embedded within human-led systems. Imaging, selected screening applications and increasingly documentation demonstrate meaningful real-world use. Risk prediction, clinical decision support and broader generative-AI applications are developing but vary considerably in maturity. In life sciences, AI has become an important research tool and is beginning to demonstrate clinical consequences through AI-assisted drug candidates, but it has not removed the fundamental scientific and clinical uncertainties involved in developing medicines.

Evidence Notes

1. **US Food and Drug Administration — AI-Enabled Medical Devices**, current 2026 list. Demonstrates substantial regulatory activity and the concentration of authorised AI applications in radiology and other diagnostic fields.
2. **FDA — Sensor-Based Digital Health Technologies**, 2026. Evidence of authorised wearable and home-monitoring applications.
3. **NHS England — Guidance on AI-enabled ambient scribing**, originally April 2025, updated 29 July 2026. Evidence that generative-AI documentation has reached implementation-policy stage within England.
4. **Singapore Ministry of Health**, 17 April, 9 July and 4 August 2026. Primary evidence concerning Note Buddy deployment, national AI use cases, healthcare infrastructure and mandatory clinician review of AI-generated notes.
5. **WHO — Computer-aided detection for tuberculosis**, June 2025 and subsequent 2026 work. Strong example of a clinically defined AI application with international relevance beyond wealthy healthcare systems.
6. **UK Department of Health and Social Care/NHS — EDITH breast-screening trial**, February 2025. Useful evidence for distinguishing very large prospective evaluation from routine deployment.
7. **European Commission — Study on AI in healthcare**, 17 March 2026, and stakeholder consultation, June 2026. Evidence of growing but uneven European adoption and continuing scalability questions.
8. **JAMA Network/PubMed — US health-system AI adoption survey**, published 2025. Provides useful adoption evidence but is geographically restricted and based on self-report from participating health systems; it should not be generalised globally.

9. **WHO — Ethics and governance of large multimodal models for health**, 2025 edition. Principal evidence concerning generative-AI reliability, bias, privacy, automation bias and governance limitations.
 10. **WHO Europe — Health-data governance and AI**, March 2025, and regional readiness assessment, November 2025. Evidence concerning data quality, governance and cross-country variation.
 11. **EMBL-EBI/Google DeepMind — AlphaFold Database**, current through May 2026. Primary evidence for the scale at which AI-generated protein-structure information is now available to biological researchers.
 12. **Nature Medicine — Rentosertib Phase IIa trial**, 2025, together with July 2026 Phase III initiation reporting. Provides an unusually advanced example of AI-assisted target and molecule discovery progressing through human trials; continued clinical testing means ultimate therapeutic success remains undetermined.
 13. **FDA and EMA — Guiding Principles of Good AI Practice in Drug Development**, January 2026, together with FDA regulatory-decision guidance. Evidence that AI is sufficiently embedded in medicine development to require specific regulatory frameworks while remaining subject to conventional evidential requirements.
 14. **Nature Reviews Drug Discovery**, April and August 2026. Useful independent scientific context on AI-supported target identification and the present state of AI drug discovery; these reviews reinforce the distinction between computational advances and demonstrated clinical productivity.
-

2. Education & Training

Status as at 30 August 2026

Artificial intelligence is now materially present across Education & Training, but its adoption is unusually uneven because much of the technology can be used directly by teachers and learners without being formally introduced by an educational institution.

Generative AI has become commonplace among many university students and increasingly familiar to teachers, trainers and school students. At the same time, education systems are introducing more purpose-built AI for adaptive learning, tutoring, feedback, lesson preparation, student support and administration. Some countries have moved further towards institutionally governed integration, while others remain primarily at the stage of guidance, experimentation and locally determined practice.

The resulting picture is not simply one of educational institutions “adopting AI”. It consists of several overlapping developments: learners independently using general-purpose AI; educators using it as an individual productivity tool; institutions providing approved AI services; purpose-built educational systems being integrated into teaching and learning; and education authorities beginning to redesign policy, assessment and AI literacy around a technology that is already widely accessible.

Evidence that AI can improve learning is also becoming stronger, but remains highly dependent on design and context. Controlled studies have found meaningful gains from particular AI-supported tutoring approaches, including secondary-school learning in Nigeria, mathematics tutoring in the United States and university science instruction. Other evidence demonstrates that AI assistance can improve performance while the technology is available without necessarily improving subsequent independent performance. These findings demonstrate both educational capability and the importance of how AI is designed and used; they do not establish that generic AI use, or AI tutoring generally, produces better educational outcomes.

1. Where is AI currently being applied?

AI is being applied across teaching, learning, assessment, educational support and administration.

One of the most visible applications is learner assistance. Students use general-purpose generative AI to explain concepts, answer questions, summarise material, develop ideas, provide examples, assist with writing, translate content, generate practice questions and support revision. In UK higher education, the 2026 HEPI survey found that 95% of surveyed full-time undergraduates reported using AI in at least one way and 94% reported using generative AI in connection with assessed work. These figures demonstrate extremely high use within the surveyed UK population, not a global student adoption rate.

More purpose-built AI tutoring and adaptive-learning systems are also operating. These systems can adjust explanations, questions, content or progression according to a learner’s responses. Singapore provides an unusually developed national example. Its Student Learning

Space is available throughout the national school system and includes AI-enabled features for adaptive learning, learning assistance, feedback and teacher support. Its Adaptive Learning System recommends learning resources, practice questions and feedback according to student readiness.

AI is increasingly being used by teachers and instructors for lesson planning, developing teaching materials, generating examples and exercises, adapting material to different abilities, summarising subject matter and preparing communications. OECD's TALIS 2024 data indicate that 37% of lower-secondary teachers across participating education systems reported using AI for their work. Among teachers using AI, 68% reported using it to learn about or summarise a topic and 64% to generate lesson plans. Use for assessment or grading was substantially lower, at 26%.

There is also evidence that some of these uses can produce measurable workload effects. An Education Endowment Foundation/NFER trial involving 68 English schools found that teachers using ChatGPT with implementation guidance for Key Stage 3 science lesson preparation spent an average 56.2 minutes per week on preparation compared with 81.5 minutes in the non-generative-AI comparison group—a 31% reduction—without an apparent reduction in the quality of lesson resources as judged by a blinded expert panel. This is useful evidence of a defined workload application, not evidence that generative AI reduces all teacher workload by 31%.

Feedback and assessment constitute another important area. AI can generate formative feedback, propose marking comments, construct questions, analyse student responses and assist educators in evaluating work. New Zealand's Ministry of Education, for example, states that AI tools may support marking but should not replace teachers' professional judgement. Its guidance was developed with the New Zealand Qualifications Authority, which separately addresses authenticity and acceptable AI use in assessment.

AI is simultaneously changing assessment design because students themselves can generate sophisticated text, code, images and analysis. Some institutions are therefore changing assignments, oral examinations, supervised assessment, process evidence and expectations about disclosure of AI assistance. In the 2026 HEPI survey, 65% of UK undergraduate respondents said assessment had changed significantly in response to AI.

Applications extend to student services and educational administration. AI systems can support information enquiries, advising, course or resource discovery, communications, scheduling, document processing and analysis of institutional information. Higher-education institutions are also experimenting with AI across staff work rather than teaching alone. An EDUCAUSE survey conducted in 2025 received 1,960 qualifying responses and found widespread institutional activity around AI strategies, pilots, policies and workforce development, although its respondents should not be interpreted as representing every institution internationally.

AI is also becoming an explicit subject of education. Schools, universities, vocational providers and workplace-learning programmes increasingly teach AI literacy, responsible use, prompting, critical evaluation and, at more advanced levels, AI development itself. UNESCO has produced separate AI competency frameworks for students and teachers. China went further in April 2026 by issuing a national "AI + Education" action plan that envisages AI education across educational stages and wider society, alongside deeper integration of AI into teaching and educational administration.

Professional development and workplace learning form another expanding area. Generative AI can create explanations, simulations, practice scenarios and tailored learning material; assist employees in locating organisational knowledge; and support the rapid production or modification of training content. Here, however, it can become difficult to distinguish an educational application from ordinary workplace AI assistance. Its defining educational role is the development or support of learning rather than completion of the work task itself.

2. What problems or needs is AI being used to address?

The largest underlying educational problem is the difficulty of providing individual attention at scale.

A teacher working with many learners cannot continually provide each person with individually paced explanations, practice questions and feedback. AI systems can respond repeatedly, vary explanations and supply assistance at the point when a learner requests it. This makes personalisation one of AI's strongest educational attractions.

Evidence increasingly suggests that this capability can produce genuine learning gains when the system and learning activity are appropriately designed. A World Bank randomised controlled trial in Nigeria used Microsoft Copilot, supported by teachers, with senior-secondary students studying English over six weeks and reported a 0.31 standard-deviation improvement on the relevant assessment. This is significant evidence from an African educational setting, while remaining evidence about a particular intervention rather than general-purpose AI use across education.

A separate US study of Tutor CoPilot provided tutors with real-time AI assistance. In a randomised trial involving more than 700 tutors and approximately 1,000 students from underserved communities, students whose tutors had access to the system were four percentage points more likely to master mathematics topics, with larger gains among students taught by lower-rated tutors. Here AI assisted the human tutor rather than substituting for one.

AI is also being directed towards the time pressures facing educators. Lesson preparation, marking and administrative work consume substantial teacher time. Generative AI can draft and transform material rapidly, and the English lesson-preparation trial provides evidence that such capability can translate into measurable time savings in a defined activity. This helps explain the technology's productivity attraction even while its effects on other teaching activities and student outcomes remain less certain.

A third need is timely feedback. Conventional feedback may arrive days or weeks after work is completed. AI can provide an immediate response, allowing students to revise, retry or request another explanation. Singapore's AI-enabled Student Learning Space illustrates institutional efforts to use this capability within curriculum-aligned and governed environments rather than relying entirely on public chatbots.

Education systems are also responding to differences between learners. Students differ in prior knowledge, language, pace, disability, confidence and access to specialist support. Translation, text simplification, speech and language technologies, different forms of explanation and adaptive practice can make some educational material more accessible. Whether this improves

educational equity depends heavily upon who can access the technology, the quality of the systems and whether different populations are represented adequately.

Another increasingly important need is preparation for an AI-influenced society and labour market. Education systems are no longer concerned only with whether students should use AI to complete existing educational tasks. They are increasingly considering what people need to understand about AI itself. UNESCO's competency frameworks and national initiatives in countries including China and Singapore reflect this movement towards AI literacy and capability as an educational objective.

Finally, AI is addressing the operational requirements of very large educational organisations. Universities, school systems and training providers manage substantial information flows, student enquiries and administrative processes. AI offers mechanisms for searching, categorising, summarising and responding to this information, although many such applications remain less visible publicly than classroom uses.

3. What can AI currently do?

AI can already perform a substantial range of educational tasks, but its ability to perform a task should not be confused with evidence that using it improves learning.

Generative systems can provide explanations, examples and practice exercises; answer follow-up questions; summarise and transform learning material; translate between languages; draft lesson plans; produce assessment questions; role-play scenarios; give preliminary feedback on student work; and support educators in adapting resources to different levels.

Purpose-built systems can go further by combining AI with curriculum structures, student-progress information or defined instructional methods. Adaptive systems can recommend the next exercise or resource. AI tutoring systems can ask questions rather than simply provide answers. Teacher-assistance systems can suggest pedagogical responses during a tutoring interaction.

Controlled studies demonstrate that well-designed systems can sometimes produce measurable learning gains. A 2025 randomised controlled university study reported that students using a specially designed AI tutor learned more in less time than students receiving the corresponding active-learning classroom lesson. The system was explicitly constructed around research-based pedagogical principles. This distinction matters: the finding does not demonstrate that providing students with unrestricted access to a general chatbot produces the same result.

Evidence from Türkiye illustrates the distinction particularly clearly. A field experiment involving nearly 1,000 high-school mathematics students found that access to a general-purpose GPT-4 interface substantially improved performance during assisted practice, but students subsequently performed 17% worse than the control group when AI access was removed. A tutor version incorporating educational safeguards largely removed the negative effect.

AI can therefore help people perform educational tasks, and appropriately designed systems can help people learn. Those propositions are related but are not equivalent.

AI can also help teachers analyse information about learners and potentially identify where additional support is required. Yet decisions about why a learner is struggling, how an explanation should be framed and what wider personal or educational circumstances matter remain substantially human activities.

AI cannot presently be assumed to provide consistently correct educational content. General-purpose generative systems can fabricate information, misinterpret prompts, oversimplify difficult material or provide answers that appear convincing despite being wrong. Their outputs can also vary between prompts and model versions.

Nor can AI independently reproduce the full role of a teacher, lecturer or trainer. Education involves motivation, relationships, classroom management, professional judgement, pastoral understanding, cultural context, evaluation of individual development and decisions about what should be learned—not merely generation of information.

4. How established is adoption?

Adoption is best understood as several layers progressing at different speeds.

Individual learner use is highly established in some environments. UK undergraduate evidence shows near-universal reported AI use, while wider international surveys also demonstrate substantial student uptake. But available surveys differ considerably in geography, population and definitions and therefore cannot be combined into a single global adoption percentage. A 2024 Digital Education Council survey, for example, covered 3,839 higher-education students across 16 countries and demonstrated widespread engagement, but was not a representative census of global higher education.

Teacher use is established but far from universal. TALIS 2024 found 37% of lower-secondary teachers reporting AI use for their work across participating education systems. There were also substantial differences between countries and between activities.

Institutional adoption trails individual use in many settings. In the 2026 HEPI survey, 95% of surveyed UK undergraduates reported some AI use, but only 38% reported that their institution provided AI tools and 36% felt encouraged by their institution to use them. This is a particularly clear illustration of the difference between technology being widely used by members of an institution and the institution itself having mature AI provision.

In some systems, however, AI has reached much more substantial institutional deployment. Singapore's Student Learning Space is available to teachers and students throughout its national school system and incorporates multiple AI-enabled educational functions. This represents system-level infrastructure rather than an isolated classroom trial, although individual AI capabilities within the platform continue to be developed and expanded.

National policy is also moving from initial reaction towards structured integration. Australia maintains a national framework for generative AI in schools, endorsed across education jurisdictions and reviewed to respond to changing technology. New Zealand provides national guidance covering educational and assessment use. European guidance for educators was updated in May 2026. China's April 2026 action plan envisages much deeper system-wide

integration over the remainder of the decade. Policy commitment, however, is evidence of institutional direction rather than completed adoption.

The current Status is therefore unusual: in some parts of Education & Training, use of AI by individuals can reasonably be described as mainstream, while educationally validated, institutionally integrated and consistently governed use remains substantially less mature.

5. What appears to be driving adoption?

The first driver is accessibility. Earlier educational technologies normally required institutions to purchase, install or integrate them. General-purpose generative AI can be accessed directly by individual teachers and learners. OECD identifies this availability outside institutional control as an important difference between generative AI and previous educational technology waves.

A second driver is immediate usefulness. Students can obtain an explanation or summary within seconds. Teachers can create a first draft of a lesson plan or exercise rapidly. In the UK 2026 undergraduate survey, students commonly associated AI with saving time, improving understanding and receiving immediate support.

Teacher workload creates similar incentives. Lesson preparation, marking and administrative activity occupy significant time, making tools that can accelerate drafting, adaptation and routine information handling attractive even where institutions remain cautious about AI undertaking educational judgements. Evidence from the English lesson-preparation trial suggests that, in at least some defined activities, the expected productivity benefit can translate into measurable time savings.

The growing evidence that some forms of AI-assisted instruction can improve learning provides a further driver. The Nigerian, US tutoring and university learning trials move the discussion beyond technological possibility by demonstrating measurable educational outcomes under defined conditions.

Governments are also increasingly treating AI capability as an educational requirement. UNESCO's competency frameworks, Singapore's EdTech strategy and China's 2026 national action plan all illustrate the shift from managing a new technology towards preparing learners and educators to function in environments where AI is increasingly normal.

Finally, the technology itself is becoming integrated into existing educational platforms and productivity systems. This reduces the distinction between deliberately "using AI" and simply using a contemporary learning, writing, search or administrative system that contains AI-enabled functions.

6. What are the present limitations?

The most fundamental limitation is that successful task completion is not equivalent to learning.

A student can use AI to produce an excellent answer without acquiring the knowledge or reasoning required to produce it independently. The developing evidence therefore points in two directions at once: appropriately designed AI-supported learning can be effective, while transferring too much cognitive work to AI can improve immediate performance without increasing—and in some circumstances can impair—subsequent unaided performance. The Türkiye field experiment provides a particularly clear example of this distinction.

Reliability remains another major limitation. Generative AI may provide incorrect references, fabricated facts or superficially plausible explanations. Learners with the least knowledge of a subject may also be the least able to recognise the error.

Assessment authenticity has consequently become a major challenge. Educational institutions need to determine what work demonstrates the learner's own knowledge when AI can generate increasingly sophisticated responses. At the same time, students may be uncertain about acceptable use. Jisc reported in 2025 that although 86% of surveyed UK universities and 49% of colleges had student AI guidance, students continued to report uncertainty about what was permissible.

There is no simple technical solution. AI-generated-text detection does not convert the authorship problem into a reliably automated judgement, making broader assessment design, process evidence and human judgement important parts of the emerging response.

Teacher capability is uneven. Using AI personally does not necessarily mean being able to judge its pedagogical value, identify bias, design an effective AI-supported learning activity or understand its data implications. UNESCO notes that teacher competency development remains insufficiently established internationally and has developed a 15-competency framework spanning human-centred thinking, ethics, AI foundations, pedagogy and professional development.

Privacy and data governance are particularly important because education involves children and extensive personal information. Questions arise about what data students enter into public systems, whether interactions are retained, how systems profile learners and whether automated recommendations affect educational opportunity. UNICEF's 2025 child-centred AI guidance highlights safety, privacy, non-discrimination, transparency and accountability as significant considerations.

Equity presents a double-edged limitation. AI can expand access to tutoring, translation and educational support where specialist resources are scarce. The Nigeria study illustrates meaningful educational results in a setting very different from wealthy Western universities. Yet access still depends on connectivity, devices, digital capability, appropriate language support and, increasingly, whether users can afford more capable AI services.

This constraint is substantial at global level. Available UNESCO evidence indicates that only about 40% of primary schools and 50% of lower-secondary schools globally are connected to the internet. AI may therefore expand educational access for some learners while remaining unavailable to others.

Regulation also varies. Under the European Union's AI Act, particular applications in education and vocational training—such as systems determining access or admission, evaluating learning outcomes in specified circumstances or monitoring prohibited behaviour

during tests—can fall within high-risk classifications. This illustrates the difference between relatively low-stakes assistance such as drafting teaching material and AI that materially affects a person’s educational trajectory.

Finally, the evidence base remains young. Some prominent studies involve a single course, limited duration, particular subjects or specifically engineered systems. Longer-term questions—including retention of knowledge, effects on independent reasoning, changes in study habits and how continuing AI assistance affects expertise development—remain much less settled.

7. What you should be aware of

“AI in education” does not describe one form of use.

A university student asking a public chatbot to explain an article, a teacher generating a lesson plan, an adaptive mathematics platform, a national education system providing curriculum-aligned AI, an automated admissions system and a workplace learner practising with an AI simulation all raise different educational, evidential and governance questions.

High student usage should therefore not be interpreted as evidence that educational institutions have completed an AI transformation. In some environments the opposite is closer to the current reality: individuals have adopted AI faster than institutions have determined how it should be integrated, taught, assessed and governed. The UK undergraduate evidence provides a particularly clear example.

Nor should AI performance be equated automatically with educational effectiveness. Controlled trials now demonstrate that appropriately designed AI-supported learning can produce meaningful benefits. These are important developments. But they coexist with evidence that learners can use AI to complete tasks while doing less of the intellectual work from which learning itself arises. The Türkiye experiment reinforces why the design of the AI-supported learning process matters, rather than merely whether AI is available.

Human involvement therefore remains substantial. Even highly developed institutional examples such as Singapore explicitly place teacher agency, human involvement and professional judgement within their AI-in-education framework. New Zealand similarly states that AI-supported marking should support rather than replace teachers’ professional judgement.

Assessment is one of the clearest places where the consequences are already visible. Generative AI has weakened the assumption that a polished submitted product necessarily demonstrates the learner’s unaided capability. Educational institutions are consequently experimenting with different combinations of permitted AI use, disclosure, supervised work, oral assessment and attention to the process by which work was produced.

International differences are substantial. Singapore has nationally integrated AI-enabled educational infrastructure; China has adopted an ambitious national integration programme; Australia has a nationally agreed school framework; European institutions operate within an increasingly consequential AI regulatory environment; and practices elsewhere can differ

between individual institutions, provinces, states or classrooms. These differences make it inappropriate to extrapolate a single country's experience into a global adoption position.

Access also matters. AI may make individualised explanation and tutoring available to people who previously could not obtain it. Conversely, differences in infrastructure, language coverage, institutional support and access to capable paid systems may generate new educational disparities. Globally uneven school connectivity demonstrates that the starting conditions for educational AI remain profoundly different.

Status at 30 August 2026: AI has moved decisively into Education & Training, but its strongest current position is as an increasingly pervasive learning and educator-assistance technology operating within predominantly human-led educational systems. Individual use—particularly among higher-education students in some countries—is already extremely high. Teacher use is substantial but uneven. Purpose-built tutoring, adaptive learning and feedback systems have demonstrated meaningful capability, and selected implementations have produced measurable learning gains. Institutionally governed adoption is advancing but generally remains less mature than individual use, with significant differences between countries and educational settings. The central unresolved educational question is therefore no longer whether AI can produce answers or assist educational work, but under what circumstances that assistance produces better learning rather than merely easier task completion.

Evidence Notes

1. **OECD — Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education, January 2026.** Major international synthesis of emerging evidence concerning generative AI in education. Particularly important for distinguishing improved task performance from genuine learning and for TALIS evidence showing substantial but non-universal teacher use. OECD evidence covers participating systems rather than every country and should not be interpreted as globally representative.
2. **OECD — TALIS 2024 results, published 2025.** Comparative evidence concerning lower-secondary teachers' AI use and purposes of use. Reports 37% of teachers across participating systems using AI for their work, with lesson planning and information summarisation more common than grading. Self-reported use and participating-country coverage limit generalisation.
3. **Higher Education Policy Institute/Kortext — Student Generative AI Survey 2026, March 2026.** Survey of 1,054 full-time UK undergraduates conducted by Savanta in December 2025. Provides unusually current evidence concerning widespread student AI use, assessment change and the gap between student usage and institutionally provided AI. It is UK-specific and should not be treated as a worldwide university adoption measure.
4. **World Bank — From Chalkboards to Chatbots: Evaluating the Impact of Generative AI on Learning Outcomes in Nigeria, May 2025.** Randomised controlled trial of a six-week AI-supported educational intervention with secondary students in Nigeria. Provides important evidence of measurable learning improvement and materially broadens the geographical evidence beyond North American and European settings. The results concern a specific supported intervention rather than unrestricted chatbot use.
5. **Wang et al./Stanford National Student Support Accelerator — Tutor CoPilot, 2025.** Randomised evidence involving more than 700 tutors and approximately 1,000 students. Demonstrates a human-AI model in which real-time AI assistance to tutors produced improved mathematics mastery, particularly for students

of lower-rated tutors. It concerns a defined tutoring environment and should not be generalised to classroom education overall.

6. **Kestin et al. — AI tutoring randomised controlled trial, Scientific Reports, 2025.** University-level study comparing a pedagogically designed AI tutor with active-learning classroom instruction. Provides evidence that carefully designed generative-AI tutoring can improve learning outcomes and reduce learning time in a defined educational context. A specialised experimental system is not equivalent to ordinary use of a general-purpose chatbot.
7. **Bastani et al., Proceedings of the National Academy of Sciences — Generative AI without guardrails can harm learning, 2025.** Field experiment involving nearly 1,000 high-school mathematics students in Türkiye. Demonstrated large improvements in assisted practice performance but poorer subsequent unaided performance from unrestricted GPT-4 use, with educational safeguards substantially mitigating the effect. Particularly important evidence for distinguishing performance from learning.
8. **Education Endowment Foundation/NFER — ChatGPT in Lesson Preparation — Teacher Choices Trial, completed May 2026.** Trial across 68 English schools found participating teachers using ChatGPT spent 31% less time on lesson and resource preparation without an apparent reduction in resource quality. Strong evidence for a defined teacher-workload use case; it should not be generalised to all teaching tasks or AI applications.
9. **Singapore Ministry of Education/GovTech — Student Learning Space and AI-in-Education materials, current through 25 August 2026.** Primary evidence of system-level educational AI deployment including adaptive learning, learning assistance, feedback and teacher-support functions within Singapore’s national school infrastructure. Particularly useful as evidence that AI integration has moved beyond isolated pilots in at least some national systems. Government descriptions establish deployment and design but are not independent evidence of educational effectiveness.
10. **Ministry of Education of the People’s Republic of China — “AI + Education” Action Plan, April 2026.** Primary evidence of national policy directed towards AI integration across educational stages, teaching, governance and wider AI literacy. It establishes policy direction and planned integration, not completion of the stated 2030 objectives.
11. **Australian Government Department of Education — Australian Framework for Generative AI in Schools, reviewed and endorsed June 2025.** Evidence of nationally coordinated governance principles covering teaching and learning, wellbeing, transparency, fairness, accountability, privacy, security and safety. A governance framework does not itself establish uniform classroom adoption.
12. **New Zealand Ministry of Education and NZQA — Generative AI and assessment guidance, current through 2026.** Evidence of national guidance emphasising teacher responsibility, professional judgement, human oversight, appropriate educational use and assessment authenticity. Useful for illustrating the distinction between AI assistance and delegation of professional assessment responsibility.
13. **UNESCO — Guidance for Generative AI in Education and Research and AI Competency Frameworks for Teachers and Students, updated through January 2026.** Principal global policy and competency evidence concerning human-centred use, teacher capability, learner capability, ethics and institutional preparedness. UNESCO’s material provides normative and policy guidance rather than evidence that the recommended practices have been universally adopted.
14. **Jisc — Student Perceptions of AI 2025.** Evidence drawn from UK further- and higher-education student discussions and institutional surveys. Useful for understanding academic-integrity uncertainty, student experience, perceived over-reliance, access and employability concerns. The evidence is principally UK based and includes qualitative findings rather than a globally representative survey.
15. **European Union — AI Act and 2026 educator guidance.** Establishes the developing European governance environment. Certain education and vocational-training applications affecting admission, learning outcomes or examination monitoring are classified as high-risk under specified circumstances. These provisions should not be interpreted as classifying all educational AI as high-risk.

16. **UNICEF — Guidance on AI and Children, Version 3.0, December 2025.** Internationally informed guidance concerning children’s privacy, safety, fairness, transparency, accountability and developmental interests. Particularly relevant where educational AI involves minors and student profiling.
 17. **Digital Education Council — Global AI Student Survey 2024.** Survey of 3,839 higher-education students across 16 countries. Useful as internationally distributed evidence that student engagement with AI extends beyond a single national system. It is not a statistically representative estimate of worldwide student adoption and should not be combined arithmetically with national surveys.
 18. **EDUCAUSE — The Impact of AI on Work in Higher Education, 2026.** Based on 1,960 qualifying survey responses gathered in late 2025. Provides evidence concerning institutional strategies, policy development, pilots and staff capability-building in higher education. Its institutional and respondent population limits its value as a global adoption measure but it provides useful evidence that AI activity extends substantially beyond student-facing teaching applications.
 19. **UNESCO — global school-connectivity evidence, 2025–26.** Evidence concerning persistent global inequalities in school internet access and the possibility that educational AI may amplify existing digital divides. It is important contextual evidence because the infrastructure required for AI-enabled education remains far from universally available. Evidence suggests connectivity of approximately 40% of primary schools and 50% of lower-secondary schools globally.
-

3. Financial Services

Status as at 30 August 2026

Artificial intelligence is already embedded across significant parts of financial services, but the maturity of its use varies considerably by activity. Machine-learning systems have become established in fraud detection, financial-crime monitoring, credit assessment, risk modelling, customer analytics and parts of trading and insurance. Generative AI is spreading quickly into document processing, information retrieval, customer interaction, coding, research and employee-assistance tasks, while more consequential uses involving autonomous decisions remain more controlled.

The sector is particularly suited to AI because financial institutions process very large quantities of structured and unstructured information, make repeated risk assessments and operate highly digitised transactional systems. At the same time, financial decisions can materially affect access to credit, insurance, investments and payments. AI use therefore sits within unusually extensive regulatory, model-risk, consumer-protection, operational-resilience and accountability frameworks.

The current Status is not one of AI independently running banks, insurers or investment organisations. It is principally one of AI performing analytical, detection, prediction, classification, information-processing and increasingly generative tasks within larger financial processes. International regulators are simultaneously documenting accelerating adoption and drawing attention to governance, data, cyber, third-party and systemic risks.

1. Where is AI currently being applied?

One of the most established areas is **fraud detection and financial-crime prevention**. Banks, payment providers and other financial institutions use machine-learning models to analyse transaction patterns, identify anomalies, assess fraud probabilities, detect suspicious accounts and prioritise cases for investigation. These systems supplement traditional rules-based controls by looking for combinations and network patterns that are difficult to define in advance. AI is also used in anti-money-laundering, sanctions screening, identity verification and cybersecurity.

AI is extensively applied in **lending and credit risk**. Models can combine conventional credit information with account, transaction, business and, in some settings, alternative digital data to estimate default risk, rank applications, set or inform credit limits and support ongoing borrower monitoring. In emerging markets, alternative-data models are also being used to assess borrowers with limited conventional credit histories. IFC reported in 2026 that lenders and fintech firms were deploying such approaches using sources including mobile-money, digital-payment, business and platform records.

Within **insurance**, AI is being applied across pricing, underwriting, claims processing, fraud detection, customer interaction and document handling. Computer vision can assist with damage assessment; language models and document-processing systems can extract information from claims and policies; predictive models estimate risks and possible losses; and

customer-service systems respond to routine enquiries. European supervisory evidence shows conventional AI already established in substantial parts of the insurance sector, while generative AI is expanding rapidly but remains considerably more experimental.

AI is also used within **investment management, securities and financial markets**. Machine learning can analyse market prices, economic information, company disclosures, news, earnings calls and alternative datasets; generate trading signals; assist portfolio construction; support order execution; monitor risk; and identify potentially abnormal market activity. Generative AI adds capabilities for summarising filings and research, retrieving information, generating code and supporting investment analysts. IOSCO's work confirms that AI use spans trading, investment research, advisory activities, market surveillance and operational processes, although maturity differs substantially among firms and applications.

A further major area is **customer service and personalisation**. Chatbots and virtual assistants answer questions, retrieve account information, help customers navigate products and increasingly produce conversational responses rather than selecting from predefined scripts. AI is also used to segment customers, identify likely needs and personalise communications. Automated and AI-assisted financial advice is developing alongside these applications, although consequential recommendations attract greater regulatory and conduct scrutiny.

Generative AI is becoming particularly visible in **internal knowledge and operational work**. Financial institutions use it to summarise documents, search internal policies, draft communications, assist software development, analyse regulatory material and support employees handling complex information. The Bank of England found UK firms' AI use concentrated substantially in internal-process optimisation, cybersecurity and fraud detection. European insurers similarly report substantial generative-AI activity in document processing, software development and other back-office functions.

There is also substantial experimentation. Hong Kong's GenA.I. Sandbox, for example, has included proposed applications involving credit-portfolio management, trade finance, market-risk alerts, fraud detection, investment insights and conversational customer service. These demonstrate active development by financial institutions but should be distinguished from evidence of routine deployment.

2. What problems or needs is AI being used to address?

Financial institutions operate in an unusually **information-intensive environment**. Banks must analyse enormous numbers of transactions, insurers process claims and risk information, investment firms absorb continuous market data, and regulated institutions must examine customers and transactions for financial crime. AI is attractive where the number of observations or possible relationships exceeds what people can examine manually.

Fraud illustrates this particularly clearly. Fixed rules can identify known patterns, but criminal methods change and often involve networks of apparently ordinary transactions. Machine learning can compare behaviours across accounts and transactions and identify combinations warranting investigation. BIS Project Hertha demonstrated this capability experimentally: adding payment-system analytics enabled participating models to identify 12% more illicit accounts in its synthetic environment and produced a 26% improvement when detecting

previously unseen behaviours. The project nevertheless used simulated data and did not establish equivalent effectiveness in a live national payment system.

Speed and operating cost are another major need. Lending applications, insurance claims, compliance reviews, customer queries and investment information frequently involve repeated processing of large quantities of material. AI can classify documents, extract relevant information, prepare summaries or predictions and prioritise cases so that people concentrate on exceptions or higher-risk decisions.

In credit, institutions face the additional problem of **incomplete information about borrowers**. Traditional credit histories may provide a poor basis for assessing individuals or small businesses that have limited previous borrowing. Alternative-data and AI-based scoring can potentially identify economically relevant patterns from digital activity that conventional credit files do not contain, which helps explain its particular relevance in parts of Asia, Africa, Latin America and other emerging markets.

Financial firms are also attempting to improve **risk differentiation**. Insurers want more accurate estimates of claims probabilities and severity; lenders seek better default predictions; asset managers search for signals relevant to expected returns and portfolio risk; and payment providers want to distinguish fraudulent transactions from legitimate ones without unnecessarily blocking customers.

A further pressure is the **regulatory and compliance burden**. Financial institutions must monitor transactions, maintain documentation, perform customer due diligence, assess risks, respond to regulatory requirements and detect suspicious conduct. AI is being directed towards information retrieval, sanctions screening, transaction monitoring, regulatory reporting and investigation support partly because these activities combine high information volumes with repeated analytical work.

Generative AI addresses another problem: substantial amounts of financial-sector knowledge are contained in **text rather than structured databases**. Policies, research, contracts, market commentary, customer correspondence and regulatory material can now be searched, summarised and reorganised through natural-language interfaces.

3. What can AI currently do?

AI can already perform many **defined financial analytical tasks at considerable scale**. Models can estimate probabilities of default or fraud, classify customers and transactions, identify unusual behaviour, recognise relationships across networks, extract information from documents, analyse images, forecast defined variables and rank cases according to risk.

These capabilities are especially useful where historical data contain recurring patterns. A fraud model, for example, can compare a transaction with previous legitimate and fraudulent behaviour far more quickly than a human investigator could examine the underlying records individually.

AI can also **combine more variables than traditional manual assessment generally permits**. Credit-scoring systems can analyse transaction histories alongside conventional

borrower characteristics. Investment systems can process market prices, filings, text, news and other datasets simultaneously. Insurance models can combine policy, claims and behavioural information.

Generative AI has added a different class of capability. It can produce fluent text, summarise lengthy material, retrieve relevant information from large document collections, generate draft reports or correspondence, interpret natural-language questions and assist with computer code. Within financial institutions these capabilities make AI increasingly useful as an employee assistant.

Machine learning can also reveal relationships that are difficult to specify through fixed rules. Network analysis is particularly relevant to fraud and money laundering because apparently unrelated accounts may form identifiable transactional structures when considered collectively.

What AI does **not** generally demonstrate is dependable independent competence across an entire financial activity. A model that predicts credit risk is not thereby capable of managing a lending relationship. A trading model that detects a market signal does not establish the ability to manage a portfolio through every market condition. A language model that explains a financial product fluently may still produce an incorrect statement.

Performance also depends on circumstances. Market relationships change; fraudsters adapt; economic conditions shift; customer populations differ; and training data may cease to represent current conditions. AI therefore does not remove the longstanding financial-sector problem of model risk—it introduces new forms alongside existing ones.

4. How established is adoption?

Adoption is substantial but **layered**.

Traditional machine learning is already established in many large financial institutions, particularly for fraud detection, credit risk, customer analytics, cybersecurity, trading and operational processes. Generative AI is expanding much faster from a more recent starting point and has generally progressed furthest where errors can be reviewed before affecting customers or financial decisions.

The Bank of England and Financial Conduct Authority's 2024 survey provides one of the clearest regulated-market snapshots. Seventy-five per cent of responding UK financial firms reported already using AI and another 10% planned to do so within three years. Insurance respondents reported 95% use and international banks 94%. Yet adoption was highly uneven: large banks reported many more applications than the median institution, and 62% of reported use cases were classified as lower materiality. Foundation models represented 17% of use cases. The results are important evidence of substantial adoption but concern responding UK-regulated firms and should not be extrapolated globally.

European insurance demonstrates the distinction between **use and maturity** particularly well. EIOPA's 2026 generative-AI survey covered 347 undertakings across 25 countries and found nearly two-thirds actively using generative AI. However, most were still at proof-of-concept

stage. Widespread experimentation and widespread mature deployment are therefore not equivalent.

Dubai provides another useful regional perspective. A 2025 DFSA survey of authorised firms in the Dubai International Financial Centre found 52% reporting AI use, up from 33% in 2024. The regulator nevertheless described adoption as nascent for many firms and reported that activity remained concentrated substantially in internal processes rather than customer-facing deployment.

South African regulators similarly describe fast adoption but considerable variability between institutions, with uses including chatbots, claims processing, forecasting, compliance and cybersecurity.

In India, AI/ML is already being used in fraud and credit applications, while some important public initiatives remain experimental. RBI's MuleHunter.AI, intended to help detect mule accounts, was reported as a pilot involving two large public-sector banks rather than an established nationwide banking capability.

Hong Kong's generative-AI sandbox likewise demonstrates a sizeable institutional development pipeline rather than universal deployment.

The overall Status is consequently stronger than “early experimentation”, but weaker than uniform transformation. **AI is established technology within financial services; individual AI applications range from routine production systems to proofs of concept.**

5. What appears to be driving adoption?

The first driver is **economics**. Financial services contain large quantities of repeated analytical and information-processing work. Even incremental improvements in fraud losses, claims handling, software development, customer servicing or staff productivity can matter when applied across millions of transactions.

A second driver is **digitisation and data availability**. Electronic payments, mobile banking, digital insurance, online investment platforms and increasingly structured customer interactions continuously generate machine-readable data. AI becomes more useful as these datasets increase in scale and quality.

Third is the escalating complexity of **fraud, cyber threats and financial crime**. Criminal activity itself increasingly uses sophisticated digital technology. Financial institutions and regulators are consequently seeking models able to recognise changing patterns more quickly than static detection rules. Singapore, for example, reported in August 2026 that MAS, government agencies and five banks were testing whether AI trained on cross-bank and public-private information could improve scam detection. Because this was a proof-of-value exercise, it represented development rather than an established industry utility.

Competition also matters. Digital banks, fintech firms, insurers, payment companies and established institutions compete on processing speed, convenience, personalisation and cost. AI can strengthen each of these dimensions.

Generative AI has materially lowered the technical threshold for using AI on text-intensive work. Employees can interact with complex systems in ordinary language rather than requiring every application to be separately programmed.

Finally, regulatory engagement is reducing some uncertainty while increasing expectations around control. India has developed its FREE-AI framework work; the UAE issued consumer-protection guidance for AI in financial institutions in February 2026; European authorities have clarified insurance and wider financial-sector expectations; Hong Kong operates regulatory sandboxes; and the FSB proposed international sound practices in June 2026. These developments indicate that regulators increasingly regard AI as a present operational issue rather than a hypothetical technology.

6. What are the present limitations?

The central technological limitation is **reliability outside the conditions represented by the model's data and design**. Financial relationships change. Credit models trained under one economic environment may behave differently during a severe downturn; trading relationships can disappear once markets change; and fraud models face adversaries deliberately attempting to evade them.

Data themselves present substantial problems. Financial information may be incomplete, inaccurate, historically biased, distributed across incompatible systems or restricted by privacy and confidentiality requirements. BIS's Financial Stability Institute identified privacy, data quality and security as major constraints on advanced AI adoption in financial services.

Explainability becomes particularly important when an AI-supported decision affects an individual. A consumer denied credit, charged an insurance premium or subjected to a fraud restriction may reasonably require an intelligible basis for the outcome. More complex models can make that harder. The UAE's 2026 guidance consequently gives specific attention to transparency, bias, accountability and explainability in high-impact financial decisions.

Bias and fairness are related limitations. Historical financial data may embody previous disparities, while alternative data can introduce relationships whose relevance or fairness is difficult to establish. A statistically predictive attribute is not automatically an appropriate basis for deciding access to financial services.

Generative models add the problem of **fabricated or inaccurate output**. Fluent language is valuable for drafting and information retrieval but does not establish factual reliability. This helps explain why regulated financial institutions have tended initially to favour internal, assisted and reviewable uses.

A substantial operational issue is **third-party dependence**. The 2024 UK survey found that one-third of AI use cases relied on third-party implementations, while considerable concentration existed among cloud, model and data providers. Financial firms may therefore depend on technologies they do not fully control or understand; 46% of firms using or planning AI reported only partial understanding of their implemented technologies.

Cybersecurity cuts in both directions. AI can strengthen threat detection, but the same technology increases attackers' ability to discover vulnerabilities, automate attacks, create convincing impersonation and accelerate fraud. IMF analysis in 2026 identified this interaction as potentially relevant to financial stability because institutions share infrastructure and technology providers.

Governance is also struggling to keep pace in some institutions. Australia's APRA reported in April 2026 that its supervisory work across selected banks, insurers and superannuation trustees found governance, assurance, risk-management and operational-resilience practices lagging the speed and complexity of AI adoption.

Finally, AI may create **system-level rather than merely individual-firm risks**. If many institutions use similar models, data, infrastructure or external providers, their behaviour may become more correlated. AI-driven market activity could also operate at speeds that reduce the time available for human intervention during stress. These remain partly prospective concerns rather than evidence that AI has already caused a systemic financial crisis, but the FSB and IMF regard them as sufficiently plausible to require active monitoring.

7. What you should be aware of

“AI in financial services” covers applications with very different consequences. A model detecting suspicious card payments, a chatbot explaining an account balance, an underwriting model affecting an insurance premium and an algorithm generating trading signals should not be treated as equivalent.

The distinction between **technical capability and decision authority** is especially important. AI may produce a credit score without making the final lending decision; identify a possible fraudulent transaction without determining guilt; propose an investment insight without controlling a portfolio; or draft customer correspondence that is reviewed before release.

Reported adoption percentages also need careful interpretation. Surveys differ in the institutions sampled, definition of AI and threshold for counting a use case. A firm with one internal machine-learning system may qualify as an AI user alongside a large international bank operating dozens of models. The UK's 75%, Dubai's 52% and European insurance statistics therefore describe different populations and should not be combined into a global adoption rate.

Geographical differences are material. Large financial centres and highly digitised banking systems have produced much of the detailed adoption evidence. Other markets may use AI through fintech, mobile-money or alternative-data models that do not resemble large-bank deployments. Evidence from IFC's work in emerging markets shows that AI's significance there can include extending credit assessment to customers inadequately represented by traditional credit files.

Regulation is increasingly consequential but is not globally uniform. Financial institutions are typically subject simultaneously to existing requirements covering credit, insurance, securities, privacy, model risk, consumer protection, operational resilience and financial crime, while some jurisdictions now also apply AI-specific frameworks. The European Union, United

Kingdom, India, UAE, Hong Kong, Australia and other jurisdictions are taking differing approaches. International bodies including the FSB, BIS, IOSCO and OECD are working towards greater common understanding rather than operating a single global AI regime.

Human accountability therefore remains substantial. Financial institutions cannot generally transfer responsibility for regulated decisions to a model or technology supplier merely because the AI produced the underlying analysis.

There is also a growing distinction between **AI used by financial institutions** and **AI used against them or their customers**. Generative AI assists financial employees, but it also enables more convincing impersonation, synthetic identities, investment scams, phishing and account-takeover attempts. AI is therefore simultaneously becoming part of financial-sector capability and part of the threat environment.

Status at 30 August 2026: AI is materially embedded in global Financial Services. Its most established position is in specialised prediction, classification, detection, risk analysis and process automation, particularly in fraud, credit, financial crime, insurance, markets and customer analytics. Generative AI is expanding rapidly into knowledge, documentation, software, customer-service and analytical assistance, but much of its more consequential deployment remains controlled, assisted or experimental. The sector's extensive digital data and repeated decision processes make AI particularly useful, while the financial consequences of errors make explainability, governance, resilience, human accountability and regulatory control unusually important.

Evidence Notes

1. **Financial Stability Board — *The Financial Stability Implications of Artificial Intelligence*, November 2024; *Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector*, October 2025; and *Sound Practices for Responsible Adoption of Artificial Intelligence*, consultation report, June 2026.** Principal cross-sector international evidence concerning present applications, adoption, governance, third-party concentration, model risk and possible financial-stability effects. The June 2026 sound-practices document remained a consultation report at the Status date; the final report was scheduled for later in 2026 and is therefore not treated as completed guidance.
2. **Bank of England and Financial Conduct Authority — *Artificial Intelligence in UK Financial Services — 2024*, November 2024; Bank of England, *Financial Stability in Focus: Artificial Intelligence in the Financial System*, April 2025.** Particularly useful quantitative adoption evidence across banking, insurance, investment and market-infrastructure firms, including materiality, foundation-model and third-party information. Results concern responding UK-regulated firms and cannot be treated as a global adoption rate.
3. **International Monetary Fund — *Artificial Intelligence and Cybersecurity in the Financial Sector*, June 2026, and Tobias Adrian, *How Central Banks Can Contain Financial Stability Risks as AI Accelerates Change*, July 2026.** Evidence concerning AI-enabled cybersecurity threats, shared infrastructure, trading, lending and potential systemic effects. Some systemic consequences remain risk assessments rather than realised outcomes.
4. **Bank for International Settlements / Financial Stability Institute — *In Data We Trust? Emerging Policy and Supervisory Approaches to AI Data Use in Financial Services*, March 2026.** Principal evidence concerning data quality, privacy, security, third-party dependence and supervisory challenges.

5. **BIS Innovation Hub — Project Hertha, June 2025.** Demonstrates technically meaningful application of AI and network analytics to payment-system financial-crime detection. Results were produced using a synthetic dataset representing 1.8 million accounts and 308 million transactions; they demonstrate capability rather than routine real-world deployment.
6. **BIS — Gambacorta, Sabatini and Schiaffi, *Artificial Intelligence and Relationship Lending*, February 2025.** Empirical evidence from Italian banking concerning AI-supported credit scoring and its interaction with relationship lending. Its defined dataset and historical period mean findings should not automatically be generalised to every credit market.
7. **International Finance Corporation — *Cracking the Credit Code: Alternative Data and AI for Financial Inclusion*, 2026.** Provides important emerging-market evidence concerning AI-supported credit scoring using mobile-money, payment, business and platform data, including applications to thin-file borrowers.
8. **EIOPA — *Generative AI Market Survey*, February 2026; *Opinion on AI Governance and Risk Management*, August 2025; and related insurance-sector material.** The survey, covering 347 undertakings across 25 European countries, is particularly useful for distinguishing active generative-AI use from proof-of-concept maturity. EIOPA's wider evidence establishes current applications across underwriting, pricing, claims, fraud and customer interaction. The evidence is European rather than global.
9. **IOSCO — *Artificial Intelligence in Capital Markets: Use Cases, Risks, and Challenges*, March 2025.** International securities-market evidence concerning AI in investment research, trading, advisory, surveillance and operational processes and associated investor-protection, market-integrity and stability risks. The publication was issued as a consultation report rather than evidence that every identified application is routinely deployed.
10. **Reserve Bank of India — *FREE-AI Committee materials*, August 2025; *RBI material concerning MuleHunter.AI and financial-sector AI*.** Provides evidence of India's institutional response to AI and a concrete fraud-detection application. MuleHunter.AI was reported in pilot use with two large public-sector banks and should not be interpreted as nationwide deployment.
11. **Hong Kong Monetary Authority — *Generative AI Sandbox materials*, 2025–26.** Provides Asian evidence of substantial bank experimentation involving risk, credit, trade finance, fraud, customer service and investment use cases. Sandbox participation demonstrates development and testing rather than routine production deployment.
12. **Monetary Authority of Singapore — *Annual Report remarks*, August 2026.** Primary evidence concerning the use and testing of AI against scams, including a proof-of-value involving government agencies and five banks. The cross-bank approach remained under evaluation at the Status date.
13. **Dubai Financial Services Authority — *AI Survey 2025, published November 2025, together with 2025–26 regulatory material*.** Provides Middle Eastern evidence across banking, capital markets, wealth management and fintech: 52% of participating DIFC-authorized firms reported AI use. The evidence concerns one major financial centre rather than the UAE or Middle East generally.
14. **Central Bank of the UAE — *Guidance Note on the Consumer Protection and Responsible Adoption and Use of AI and ML by Licensed Financial Institutions*, February 2026.** Evidence that AI in consequential financial decisions has progressed sufficiently to require explicit regulatory attention to transparency, bias, accountability, explainability and consumer outcomes. Regulatory guidance does not itself establish the extent of adoption.
15. **South African Reserve Bank / Prudential Authority — *South African financial-sector AI material*, 2025–26.** Provides useful African evidence of varied adoption across financial institutions and applications including chatbots, claims, forecasting, compliance and cybersecurity. It supports global breadth but should not be extrapolated to Africa generally.

16. **Australian Prudential Regulation Authority — AI industry letter and supervisory findings, April 2026.**
Primary evidence that AI adoption across selected Australian banks, insurers and superannuation trustees was accelerating while governance, assurance, risk-management and operational-resilience practices were not always keeping pace. The underlying engagement focused on selected large institutions rather than every regulated entity.
 17. **OECD — *Supervision of Artificial Intelligence in Finance*, January 2026, and OECD AI-in-finance material.**
Useful cross-jurisdictional context concerning supervisory approaches and applications across fraud, credit, risk management, customer service, compliance, investment and portfolio management.
-

4. Manufacturing & Engineering

Status as at 30 August 2026

Artificial intelligence is now materially embedded in parts of manufacturing and engineering, but adoption remains highly uneven between companies, factories, countries and applications. The strongest current position is in relatively bounded industrial tasks where AI can analyse large volumes of sensor, image or process data: quality inspection, predictive maintenance, production optimisation, anomaly detection, scheduling and selected forms of industrial automation. At technologically advanced factories, AI increasingly operates as part of integrated production systems rather than as an isolated experimental tool.

Generative AI is broadening the application map. Manufacturers and engineers are using or testing it for technical-information retrieval, maintenance assistance, engineering documentation, software and control-code support, design exploration, troubleshooting and access to accumulated organisational knowledge. AI-enabled robotics is also advancing as machine vision, learning systems and increasingly foundation-model approaches give machines greater ability to perceive and respond to changing environments. These developments are important, but their maturity is generally lower than that of established machine-learning and computer-vision applications.

The sector should not therefore be characterised simply as becoming “AI-powered”. Leading factories demonstrate significant real-world capability, but they are not representative of all manufacturing. In the European Union, for example, 18% of manufacturing enterprises with at least 10 employees reported using AI technologies in 2025. Australia reported AI use among 9% of manufacturing businesses in 2024–25. China reported that by the end of 2025 more than 30% of manufacturing enterprises above a specified revenue threshold had adopted AI technologies. Definitions and populations differ, so these figures should not be directly compared, but collectively they illustrate substantial geographical, firm-size and maturity differences.

1. Where is AI currently being applied?

One of the most established areas is **quality inspection and process monitoring**. Computer-vision systems inspect components, welds, surfaces, packaging, electronic assemblies and finished products for defects or deviations. Unlike traditional rule-based inspection, machine-learning systems can recognise complex visual patterns and classify defects using examples from production data. Industrial AI is also used to monitor process variables and detect abnormal operating conditions before they result in rejected products or equipment damage.

Predictive and condition-based maintenance is another important application. Data from vibration, temperature, pressure, acoustics, electrical systems and other sensors can be analysed to identify patterns associated with machine degradation. Rather than maintaining equipment solely at fixed intervals or after breakdown, organisations can use AI-generated estimates or warnings to focus attention on equipment showing signs of deterioration.

AI is increasingly applied to **production and process optimisation**. Models analyse combinations of process variables, machine states, energy consumption, material characteristics and historical production results to identify operating settings associated with higher yield, lower energy consumption, greater throughput or more consistent quality. In continuous-process sectors such as chemicals, metals, energy-intensive manufacturing and some food or pharmaceutical production, small improvements in process conditions can have significant operational value.

Planning, scheduling and factory orchestration form another important application area. AI can help sequence jobs, allocate machines and labour, anticipate bottlenecks and adjust production schedules when demand, materials or equipment availability change. China Merchants Heavy Industry's Haimen shipyard, for example, reported using AI-driven planning and digital-twin assembly among more than 25 digital solutions, while Rockwell Automation's Singapore site reported AI-enabled quality control and intelligent maintenance within a high-mix production environment. These are examples of advanced deployments, not evidence that equivalent systems are commonplace across manufacturing generally.

Engineering design and product development are increasingly affected. Machine learning can assist with simulation, optimisation and prediction of product or material behaviour. Generative and optimisation techniques can propose design alternatives subject to defined constraints such as weight, strength, manufacturability, cost or thermal performance. AI can also assist engineers in searching technical information, drafting documentation, developing or checking code and exploring design alternatives. In materials engineering, research has demonstrated AI applications for predicting performance and helping optimise material composition, although research performance should not automatically be equated with routine engineering practice.

AI is also being combined with **digital twins**—digital representations of machines, production lines or processes. Data from physical equipment can update the digital model while analytical systems simulate operating conditions, predict outcomes or identify potential changes. Australia's CSIRO identifies digital twins, real-time process optimisation, intelligent IoT, machine vision and robotic automation among its active manufacturing-AI areas.

Finally, AI is becoming more closely connected with **industrial robots and autonomous equipment**. Machine vision already allows robots to identify, position and manipulate objects under defined conditions. Newer systems are being developed to improve adaptability, including robotics foundation models trained across broader sets of tasks. Japan's METI and NEDO were funding dedicated robotics-foundation-model development during 2026, which demonstrates substantial developmental activity but not yet routine factory-wide deployment of such models.

2. What problems or needs is AI being used to address?

Manufacturing generates large quantities of operational data that people cannot continuously interpret unaided. Modern factories may contain thousands of sensors, machines, quality measurements and production events. AI is attractive where useful patterns are buried within these data and decisions have to be made repeatedly or quickly.

A central need is **higher productivity from expensive industrial assets**. Unplanned machine downtime can interrupt entire production lines. Predictive maintenance and anomaly detection seek to identify emerging problems earlier, while process optimisation attempts to obtain greater output, quality or energy efficiency from existing equipment.

Quality and consistency create another major demand. Human inspection can be highly skilled but is difficult to apply continuously to every item at high production speed. Computer vision can inspect large numbers of products according to repeatable criteria, while statistical and machine-learning systems can detect process drift that may lead to quality problems.

Manufacturers are also dealing with **greater production complexity**. Product variants, smaller batches, customised orders and volatile demand create scheduling and changeover challenges. AI-supported planning and optimisation can help evaluate large numbers of possible production arrangements that would be difficult to examine manually.

Skills shortages and loss of tacit knowledge are becoming significant drivers of newer AI applications. Experienced technicians frequently possess knowledge about machines and production processes that is poorly documented. Generative AI and knowledge-retrieval systems are being developed to make manuals, maintenance records, engineering documents and accumulated operating experience more accessible. The World Economic Forum's 2026 industrial evidence also emphasises workforce capability and human-machine collaboration as increasingly important components of advanced manufacturing transformation.

Engineering faces a related problem: the **size of the design space**. Engineers often have to choose among large numbers of possible geometries, materials, operating parameters or system configurations while meeting multiple constraints. AI-supported optimisation and simulation can narrow this search, identify promising alternatives and reduce some repetitive analytical work.

Energy, materials and sustainability pressures also create demand. AI can optimise heating, cooling, motors, process settings, material usage and production sequencing. World Economic Forum Lighthouse examples include factories using AI-enabled systems in energy-intensive processes and sustainability applications, although the reported performance improvements arise from broader technological transformations rather than AI alone.

3. What can AI currently do?

Current AI can perform **pattern-recognition and prediction tasks** very effectively where suitable data and clearly defined objectives exist. It can classify images, identify abnormal sensor patterns, forecast the likelihood of equipment failure, estimate product quality from process data and identify relationships among production variables.

Machine-vision systems can inspect products at production-line speed and identify defined categories of defects. AI-assisted robots can use visual information to locate or manipulate objects and perform tasks requiring more variability than conventional fixed automation. China reported, for example, a TCL production system using 3D vision and AI to automate connection of television interfaces, with a reported 99.8% product pass rate after introduction.

This establishes a real industrial deployment, but the reported performance of one system does not establish general capability across unrelated manufacturing tasks.

AI can also optimise complex processes by learning relationships between operating variables and outcomes. At advanced sites, these models can provide recommended settings or form part of closed-loop control systems. China's reported “Lighthouse” factories provide particularly advanced examples: by early 2026, the Ministry of Industry and Information Technology stated that AI had penetrated more than 70% of business scenarios in designated leading factories and that more than 6,000 vertical models had been developed. Because these are explicitly selected leading factories, the evidence demonstrates technical and organisational possibility rather than national manufacturing-wide maturity.

Generative AI can retrieve and reorganise technical information, summarise engineering or maintenance documentation, explain procedures, generate draft instructions and code, and support troubleshooting. It can make large repositories of industrial information easier to interrogate conversationally.

In design, AI can help generate candidate configurations, approximate results that would otherwise require lengthy simulation, optimise material or geometrical parameters and identify relationships within experimental or engineering datasets. These capabilities can significantly accelerate parts of engineering analysis.

What AI generally cannot do reliably is **independently perform the entire engineering or manufacturing function**. Industrial environments contain unusual events, changing materials, equipment wear, safety constraints and physical interactions that may not be represented in training data. A model that works well on one production line or product can perform differently following a process change, equipment modification or transfer to another factory.

Generative systems add another limitation: plausible language does not guarantee technically correct engineering content. Consequently, the ability to generate an apparently credible maintenance instruction, design explanation or control-code suggestion is different from demonstrating that the output is safe and correct in the relevant industrial environment.

4. How established is adoption?

Adoption is best described as **substantial but layered**.

Conventional machine learning, optimisation, computer vision and predictive analytics are established technologies in parts of large-scale manufacturing. Industrial robotics itself is highly established: 542,000 industrial robots were installed worldwide in 2024, with Asia accounting for 74% of new installations. Not all industrial robots use AI, but the scale of robotics creates a substantial installed base into which machine vision, learning and adaptive control can increasingly be integrated.

Leading factories have moved considerably beyond experimentation. The World Economic Forum's Global Lighthouse Network contained 238 advanced manufacturing and supply-chain sites across more than 30 countries by mid-2026. Within that selected population, analytical AI and machine learning accounted for approximately 62% of Lighthouse solutions in 2025

and generative AI 23%. The network includes sites in China, India, Saudi Arabia, Singapore, Europe and the United States showing scaled AI applications in quality, maintenance, process control, planning and other operations. This is compelling evidence of real deployment at the technological frontier, but the network deliberately represents advanced sites and should not be treated as a representative survey of ordinary factories.

Broader evidence shows a different picture. Eurostat reported AI use among approximately 18% of EU manufacturing enterprises with at least 10 employees in 2025. The Australian Bureau of Statistics reported 9% among manufacturing businesses in 2024–25. China reported more than 30% adoption among manufacturing enterprises above its specified annual-turnover threshold by the end of 2025. Different definitions and business populations prevent direct ranking, but all point to diffusion that is significant without being universal.

Survey evidence from large manufacturers tends to show higher adoption. Deloitte's 2025 survey of 600 executives from large manufacturers with headquarters or operations in the United States found 29% using AI/machine learning at facility or network level and a further 23% piloting it. For generative AI, 24% reported facility- or network-level deployment while 38% reported pilots. Because the respondents were from large companies with at least US\$500 million revenue and more than 1,000 employees, these results should not be extrapolated to smaller manufacturers or the world as a whole.

Regional maturity also varies. Evidence on Latin America remains comparatively limited. An Inter-American Development Bank study using Colombian manufacturing data found only 6% adoption in 2019–20 and strong concentration among larger firms; because the data pre-date the recent generative-AI expansion, they are useful mainly in demonstrating persistent diffusion and firm-size differences rather than the 2026 level of adoption.

Evidence for Africa is less comprehensive still. African Development Bank reporting identifies infrastructure, finance, technology and human-capital gaps as continuing constraints on industrialisation itself. This makes any assumption that adoption patterns seen in the most digitally advanced Asian, European or North American factories apply across Africa particularly unsafe.

Generative AI and more adaptable AI-enabled robotics therefore sit at an earlier maturity stage than computer vision, predictive analytics and established industrial automation. They are increasingly visible in deployments and pilots, but the current evidence does not support describing them as routine across global manufacturing.

5. What appears to be driving adoption?

The strongest driver is **productivity and competitiveness**. Manufacturers face continuing pressure to increase throughput, improve asset utilisation, reduce waste and maintain quality while controlling costs. Deloitte's surveyed manufacturers reported production-output, employee-productivity and capacity improvements from broader smart-manufacturing programmes, although those improvements cannot be attributed solely to AI because the programmes combine multiple technologies.

A second driver is the increasing availability of **industrial data**. Sensors, manufacturing execution systems, industrial IoT, machine histories, quality systems and connected equipment provide the data on which industrial AI depends. Factories with mature digital infrastructure are consequently better positioned to adopt AI than plants where operational information remains fragmented or largely manual.

Labour and skills pressures are also significant. Manufacturers in many economies report difficulty recruiting or retaining production, engineering and technical personnel. AI is being used both to automate selected tasks and to make existing workers more productive. Australia's Reserve Bank found that firms expected AI, machine learning, robotics and automation to attract increasing investment, while also emphasising the need for complementary skills, workflow changes and management capability.

The rapid improvement of general-purpose AI models has created an additional driver. Industrial information that was previously difficult to access can increasingly be searched and interpreted through natural-language interfaces. This makes AI relevant to technical workers who do not themselves build machine-learning models.

Government industrial strategy is also materially influencing adoption. China issued a dedicated "AI + Manufacturing" action programme in January 2026 and has been developing graded smart-factory programmes. Japan is supporting AI-ready manufacturing data and robotics foundation models through GENIAC. Australia is supporting AI and autonomous systems within its national AI and manufacturing initiatives. These programmes show that industrial AI is increasingly treated as an economic and technological capability rather than merely an individual-company software decision.

6. What are the present limitations?

The first limitation is **data quality and context**. Industrial AI depends on reliable sensor readings, machine histories and correctly labelled production data. Many factories contain equipment from different generations, proprietary control systems and inconsistent data structures. Information may be missing, inaccessible or collected for operational purposes rather than machine learning.

Models are also vulnerable to **changing production conditions**. Tool wear, new suppliers, altered raw materials, software updates or new product variants can change the relationship between inputs and outputs. A defect-detection model trained on yesterday's production may therefore require monitoring or retraining as conditions change.

Integration with operational technology is difficult. Factory systems operate continuously and often involve safety-critical machinery. Connecting AI to existing programmable logic controllers, manufacturing execution systems, robots and enterprise platforms may require substantial engineering work. Downtime during implementation can itself be expensive.

Cybersecurity becomes more significant as previously isolated industrial equipment becomes networked and data flows between operational technology, cloud services and AI systems. The consequences of an industrial cyber incident can include physical production disruption rather than merely loss of information.

Economics also constrain adoption. Large factories producing high volumes may be able to justify specialised AI systems because small performance improvements produce substantial returns. Smaller manufacturers frequently have fewer technical personnel, smaller datasets and less capital to fund integration. Official Australian and European evidence showing higher AI adoption among larger firms is consistent with this structural difference.

Skills remain a material constraint. Industrial AI requires a combination of data science, software, automation, engineering and domain knowledge that is not always available within one team. The problem is therefore not simply acquiring an AI model; organisations must connect digital expertise with people who understand the actual production process.

Generative AI introduces additional **reliability and accountability problems**. Incorrect technical information may appear convincing, while generated code or engineering suggestions may contain errors that are difficult to identify without appropriate expertise. Industrial safety and product-liability considerations therefore distinguish manufacturing applications from low-consequence uses of generative AI.

Physical AI and more general-purpose robotics face further limitations. Manipulating irregular objects, operating safely around people and responding robustly to unfamiliar physical situations remain much harder than performing a well-defined repetitive robotic task. The continuing public R&D investment in robotics foundation models in Japan and elsewhere itself indicates that these capabilities remain under active development rather than being solved industrial technologies.

Finally, evidence of economic impact remains difficult to interpret. Manufacturers frequently implement AI alongside sensors, robotics, cloud systems, digital twins, workflow redesign and workforce training. Reported productivity or quality gains therefore often represent the effect of an integrated transformation rather than AI operating independently.

7. What you should be aware of

“AI in manufacturing” covers technologies at very different stages of maturity. A computer-vision system identifying a known defect on a controlled production line, a model predicting bearing failure, a generative assistant answering maintenance questions and a humanoid robot learning multiple physical tasks present very different technological and evidential questions.

High-profile smart factories are important because they demonstrate what is already technically and organisationally possible. They do not establish typical sector adoption. The World Economic Forum's Lighthouse Network and China's designated leading factories are deliberately selected examples of advanced industrial transformation. Their achievements should therefore be interpreted as evidence of the frontier, not the manufacturing-sector average.

Likewise, industrial automation and AI should not automatically be treated as synonymous. A conventional robot following predetermined instructions may contain little or no AI. Conversely, an AI system analysing process data may never directly control a machine. The growing integration of AI with robotics makes the boundary increasingly blurred, but the distinction remains important when interpreting adoption statistics.

Human involvement remains substantial. Engineers, operators and maintenance specialists define objectives, verify recommendations, respond to unexpected conditions and remain important where safety, product quality or engineering judgement is involved. The International Labour Organization's 2026 tripartite work on AI in manufacturing reflects the significance of this human-machine relationship in a sector employing almost 500 million people worldwide.

Geography also matters. Industrial AI adoption depends on the existing manufacturing base, digital infrastructure, electricity and connectivity, capital availability, skills and access to technology. UNIDO notes major technological and industrial gaps between high-income and developing economies. Consequently, a global Status assessment cannot reasonably infer conditions in lower-income manufacturing economies from deployments in China, Singapore, Germany, the United States or other highly industrialised environments.

Engineering capability should also be separated from engineering responsibility. AI can generate alternatives, approximate simulations and assist analysis, but technically plausible output does not itself constitute a validated engineering design. Physical products remain subject to material behaviour, manufacturing tolerances, safety requirements and applicable technical standards.

Status at 30 August 2026: AI has moved materially into Manufacturing & Engineering. Its strongest present position is in specialised analytical, optimisation, inspection and maintenance tasks connected to increasingly digital factories. Leading manufacturers are demonstrating scaled integration of AI into production and decision-making, while broader statistics show that sector-wide adoption remains far from universal and strongly influenced by company size and regional digital maturity. Generative AI is rapidly expanding into engineering knowledge, maintenance and design workflows, and more adaptable AI-enabled robotics is progressing, but both remain less mature than established industrial machine learning, computer vision and automation. The current picture is therefore one of progressively more intelligent industrial systems operating within—and increasingly reshaping—human-designed production and engineering processes, rather than autonomous AI replacing manufacturing or engineering as a whole.

Evidence Notes

1. **Eurostat — Use of artificial intelligence in enterprises, 2025 / 2026 statistical reporting.** Official EU business statistics show that approximately 18% of manufacturing enterprises with at least 10 employees used AI technologies in 2025, compared with 20% across the covered business economy. Particularly valuable as representative adoption evidence rather than selected case-study evidence. Definitions encompass several types of AI and do not indicate depth of deployment.
2. **World Economic Forum — Global Lighthouse Network, January–July 2026.** Provides extensive evidence of real advanced-manufacturing deployments across more than 30 countries, including AI-enabled quality, maintenance, process optimisation, planning, sustainability and workforce applications. By mid-2026 the network included 238 advanced industrial sites. Because participating factories are deliberately selected for advanced performance, the evidence demonstrates the technological frontier rather than typical manufacturing adoption.

3. **China Ministry of Industry and Information Technology — “AI + Manufacturing” programme and smart-factory reporting, 2025–26.** Primary evidence of China's large-scale industrial-AI policy effort and reported penetration of AI across designated leading factories. Ministry reporting states that AI had entered more than 70% of business scenarios in leading factories and that more than 6,000 vertical models had been accumulated. These are selected frontier factories and should not be generalised across Chinese manufacturing.
4. **Chinese Government / Ministry reporting — national manufacturing adoption, March 2026.** Reported that by the end of 2025 more than 30% of manufacturing enterprises with annual main-business turnover of at least RMB20 million had adopted AI technologies. Useful national evidence but applies to a defined population of comparatively substantial enterprises and is not directly comparable with Eurostat or other surveys.
5. **Deloitte — 2025 Smart Manufacturing and Operations Survey.** Survey of 600 executives from large manufacturing organisations found facility- or network-level use of AI/ML by 29% and generative AI by 24%, alongside substantial pilot activity. Useful for distinguishing deployment from experimentation but restricted to large companies with headquarters or operations in the United States and therefore not globally representative.
6. **International Federation of Robotics — World Robotics 2025 and April 2026 regional analysis.** Reports 542,000 industrial-robot installations worldwide in 2024, with 74% of new installations in Asia. Establishes the substantial global base of industrial robotics into which AI capabilities are increasingly incorporated. Industrial-robot counts should not themselves be interpreted as AI-adoption counts.
7. **Japan METI and NEDO — GENIAC manufacturing-data and robotics-foundation-model programme, May 2026.** Primary evidence that Japan is funding development of AI-ready manufacturing datasets and robotics foundation models. Demonstrates strategic and developmental activity rather than established widespread deployment of foundation-model robotics.
8. **Australian Bureau of Statistics — Characteristics of Australian Business 2024–25, released June 2026.** Official data report AI use by 9% of manufacturing businesses compared with 12% of all businesses. Provides useful Oceania evidence and demonstrates that AI use remains far from universal outside highly publicised advanced deployments.
9. **CSIRO — AI for Manufacturing / Future Digital Manufacturing, current 2026 material.** Provides authoritative Australian evidence concerning active work in machine vision, robotics, intelligent IoT, digital twins, additive manufacturing, materials design, quality control and real-time process optimisation. Much of the material describes research and technology-development capability and should not be interpreted as prevalence data.
10. **Inter-American Development Bank — AI diffusion in Colombian manufacturing and wider Latin American digital-technology studies.** Firm-level research found 6% AI adoption among Colombian manufacturing firms in 2019–20 and significant concentration among larger firms. Particularly useful for illustrating unequal diffusion and the shortage of comprehensive Latin American adoption evidence, but the underlying adoption data pre-date the generative-AI expansion and cannot represent the 2026 rate.
11. **African Development Bank — Africa Industrialisation Index 2025, published May 2026.** Provides current context concerning the continent's comparatively small share of global manufacturing and continuing infrastructure, finance, technology and human-capital constraints. It does not provide a continent-wide manufacturing-AI adoption rate and is therefore principally evidence of the structural conditions affecting diffusion.
12. **UNIDO — Industrial Development Report 2026 and digitalisation/AI work.** Provides global context on widening technological capability differences, the rapid rise of AI and advanced digital production technologies and the risk of digital divides affecting industrial development. Useful for qualifying claims of worldwide convergence in industrial AI adoption.
13. **International Labour Organization — AI in the manufacturing industry, April 2026.** Tripartite conclusions involving governments, employers and workers provide evidence that AI's effects on productivity, work organisation, skills and human-machine interaction have become material enough to

require sector-level international consideration. The ILO notes that manufacturing employs almost 500 million workers globally.

14. **Reserve Bank of Australia — Technology Investment and AI: What Are Firms Telling Us?, November 2025.** Survey of more than 100 medium and large Australian firms, including 24 manufacturers, found expectations of greater investment in AI/ML, robotics and automation and emphasised complementary requirements such as skills, workflow change and managerial capacity. The sample was not designed to reproduce the exact Australian business population and under-represented small firms.
-

5. Retail & E-commerce

Status as at 30 August 2026

Artificial intelligence is now embedded across important parts of retail and e-commerce, but its maturity varies substantially by application and by type of retailer. Recommendation systems, search ranking, demand forecasting, fraud detection, inventory optimisation and aspects of automated pricing have been used for years and form part of the operating infrastructure of many large digital retailers. More recent generative AI is being added to customer service, product discovery, marketing, merchandising and merchant support, while conversational and agentic systems are beginning to undertake parts of the shopping process itself.

The sector is therefore moving through two developments simultaneously. Established predictive and optimisation systems continue to spread and improve, while generative and agentic AI is altering how consumers can search for, compare and increasingly purchase products. Major platforms in the United States and China have already moved some of these newer capabilities into customer-facing deployment. Amazon reported in 2026 that more than 250 million customers had used its AI shopping assistant during the year, while Alibaba connected its Qwen application to Taobao's catalogue of more than four billion products, allowing users in China to browse, compare, order and manage delivery conversationally.

This does not mean that retail has become uniformly AI-driven. Large e-commerce platforms possess unusually rich transaction data, sophisticated infrastructure and substantial technical resources. Smaller physical retailers may still be progressing through basic digitalisation. Singapore's 2026 Retail Industry Digital Plan, for example, reported that more than 75% of retail SMEs had adopted entry-level digital solutions and 45% intermediate solutions, while adoption of more advanced technologies remained limited. The appropriate Status is consequently one of significant real deployment combined with wide differences in sophistication, scale and organisational readiness.

1. Where is AI currently being applied?

One of the most established applications is **product discovery and recommendation**. Retailers use machine-learning systems to analyse browsing, purchase and product data and decide which products, offers or content to display. Recommendation engines support cross-selling, upselling and personalisation across websites, applications, email, loyalty programmes and increasingly physical-and-digital retail journeys. Singapore's IMDA, for example, provides a retail recommendation engine designed to combine data from point-of-sale systems, customer-management systems and web or mobile applications.

Search itself is changing. Conventional retail search depends heavily on product keywords and structured filters. Generative-AI shopping assistants can instead respond to requests expressed in everyday language, compare products, summarise reviews and incorporate the shopper's circumstances or preferences. Amazon's shopping assistant can answer product questions, recommend alternatives, track prices and use images or previous shopping activity; Walmart's

Sparky performs functions including review summarisation, comparison and occasion-based recommendation.

By 2026, some applications had moved beyond recommendation towards **agentic commerce**. Amazon reported capabilities including automatically adding products to a cart, reordering previous purchases and buying an item when its price reaches a customer's specified level. Alibaba's Qwen-Taobao integration enables conversational browsing, purchasing, order management and logistics-related actions. Shopify has meanwhile developed infrastructure allowing merchants' products to be discovered and purchased through external AI environments and has opened agentic-commerce capabilities to millions of merchants.

AI is also widely applied behind the storefront. **Demand forecasting and inventory management** systems estimate what will sell, where and when. The forecasts inform purchasing, replenishment, allocation, fulfilment and markdown decisions. Walmart describes AI-enabled systems that identify inventory imbalances and redirect products, provide unified inventory visibility and coordinate warehouse and transportation activity. Singapore's 2026 retail digitalisation programme similarly identifies AI-supported demand forecasting as a practical application for SME retailers.

Merchandising and product development are increasingly AI-assisted. Retail teams can analyse sales, search trends, customer behaviour and external signals to identify emerging demand. Generative AI can assist with product descriptions, imagery, campaign material and aspects of design. Walmart has described a system combining trend signals with generative AI to shorten parts of its trend-to-product process, while Alibaba reported AI tools being used across large numbers of merchants.

Customer interaction and service is another substantial area. Chatbots and conversational systems answer product questions, provide recommendations, check order status, support returns and undertake routine service interactions. Singapore's refreshed Retail Industry Digital Plan specifically includes in-store or online AI concierges and generative-AI customer-engagement chatbots.

AI also contributes to **pricing and promotions**. Retail and e-commerce organisations have long used algorithms to forecast demand, monitor competitor prices, optimise markdowns and adjust prices or discounts. Contemporary AI increases the amount of behavioural, transaction, inventory and contextual information that can be incorporated into such decisions. Academic reviews identify dynamic pricing as an established area of machine-learning research and application in e-commerce, although practices differ substantially and personalised pricing raises particular consumer and regulatory questions.

Other applications extend across fulfilment, logistics, payment and transaction-risk analysis, product-catalogue management, advertising, seller assistance, employee support and retail operations. China's June 2026 national policy on "AI + consumption", for example, explicitly identified e-commerce operations, customer service, design, marketing, livestreaming, recommendation, quality inspection and logistics as areas for deeper AI application.

2. What problems or needs is AI being used to address?

Retail presents an unusually large information problem. Large retailers may manage millions or billions of product records, enormous volumes of customer interactions, rapidly changing inventories and large numbers of transactions. Consumers face the inverse difficulty: there may be too many products to compare efficiently.

Recommendation, search and conversational systems address this mismatch. They try to reduce the effort required to find a suitable product while allowing retailers to make enormous catalogues more navigable. Amazon's and Alibaba's conversational shopping systems illustrate this movement from searching a database towards expressing an intention and allowing AI to interpret it.

Retailers also confront persistent uncertainty over demand. Too little stock causes lost sales; too much ties up capital and may lead to markdowns, spoilage or waste. Demand varies across locations, channels, seasons, promotions and external events. AI forecasting is used to identify patterns within these variables and support replenishment, distribution and fulfilment decisions.

Labour and operating-cost pressures create another demand for automation. Retail involves large volumes of repetitive analytical, administrative, marketing and customer-service work. Generative AI can draft product descriptions and promotional content, summarise commercial information, respond to routine enquiries and help staff interrogate business data. Singapore explicitly linked its 2026 expansion of AI tools for retailers to manpower constraints, rising operating costs and intensifying e-commerce competition.

E-commerce also creates a need for continuous catalogue management. Product information must be classified, enriched, translated, matched and kept consistent across channels. This becomes increasingly important as commerce moves into AI assistants because AI systems themselves need reliable information about product attributes, prices and availability. Shopify's 2026 development of structured catalogue infrastructure for agentic shopping illustrates this requirement.

Finally, retailers seek greater relevance from customer communication. AI enables segmentation and personalisation at a scale difficult to achieve manually. That can make recommendations and marketing more useful, but it also increases dependence on customer data and creates questions about privacy, manipulation, discrimination and the boundary between useful personalisation and excessive behavioural targeting.

3. What can AI currently do?

AI can already perform several retail tasks at large scale. Predictive models can rank products, forecast demand, identify customer and product patterns, estimate probabilities of purchase, support replenishment decisions and optimise aspects of pricing and promotion. A systematic review of 66 machine-learning studies in e-commerce identified personalisation, sales forecasting and pricing optimisation among the principal areas in which these methods are being applied.

Computer-vision and multimodal systems can identify visually similar products and allow shoppers to search using photographs rather than product names. Generative models can summarise large volumes of reviews and product information and answer questions that would previously have required the shopper to inspect multiple pages.

Large language models also make retail systems easier to interrogate. Instead of requiring staff to construct reports manually, an employee can increasingly ask questions about sales, inventory or merchandising in ordinary language and receive summaries or suggested actions. Walmart, for example, describes agentic tools through which associates can investigate inventory shortages and obtain recommended next steps.

Generative AI can create marketing text, product descriptions, training content and other commercial material. Singapore reported an example in which a retailer estimated a 50% reduction in content-production time after adopting a generative-AI marketing and sales-content system. This is evidence of a particular deployment rather than proof that equivalent gains occur across retail generally.

The newest systems can carry out actions. An AI shopping assistant can maintain context across several questions, identify products matching multiple conditions, compare alternatives, populate a shopping basket and, in some implementations, execute a purchase or reorder. These are genuine deployed capabilities by the Status date, rather than merely research demonstrations.

Their scope nevertheless remains bounded. A system that is effective at selecting products from a well-structured catalogue does not thereby possess general understanding of product quality or a consumer's wider interests. Generative systems can misinterpret requirements, produce incorrect product information or recommend an unsuitable item. Autonomous transaction capability therefore represents a material advance in what AI can *do*, but it does not establish that consumers or retailers are ready to delegate every purchasing decision to AI.

4. How established is adoption?

Adoption is best understood as layered.

Established deployment includes recommendation engines, search ranking, forecasting, fraud and risk analytics, advertising optimisation and numerous forms of inventory and pricing analytics. These are not new consequences of generative AI; machine learning has been part of major e-commerce operations for many years.

Growing deployment includes generative-AI customer service, review summarisation, conversational shopping, marketing-content generation and employee or merchant assistants. Deployment evidence is now substantial enough that these applications should no longer be described simply as experiments, although their penetration varies greatly.

At the largest platforms, conversational commerce has reached meaningful scale. Amazon stated that more than 250 million customers had used its shopping assistant during 2026 and reported strong growth in usage and interaction. This is provider-reported evidence and does

not establish sector-wide adoption, but it demonstrates that conversational shopping can operate at very large real-world scale.

China provides separate evidence of substantial movement. Alibaba's May 2026 integration of Qwen with the full Taobao catalogue moved its shopping agent beyond earlier limited-category trials, while China's national e-commerce report published on 29 August 2026 described AI as already being applied at scale across content generation, targeted marketing, intelligent customer service and supply-chain coordination.

Agentic commerce is emerging rather than universally established. Systems able to undertake transactions now exist in production, and commerce platforms have begun creating infrastructure enabling merchants to participate. Shopify reported in March 2026 that millions of merchants could become accessible through AI-chat commerce infrastructure, but platform availability is not equivalent to widespread merchant transactions or consumer adoption.

The SME position is materially different. Singapore's evidence shows why generalisation from major digital platforms is unsafe: although basic retail digitalisation is widespread among its SMEs, advanced solution adoption remains limited. The government's response has therefore moved towards practical AI guidance covering customer experience, forecasting and generative-AI applications.

Available adoption evidence is geographically uneven. Large deployments are particularly visible in the United States and China, with significant platform and policy activity elsewhere in Asia and increasing AI-commerce infrastructure available internationally. There is much less comparable evidence establishing the depth of AI adoption among small and medium-sized retailers across Africa, Latin America, the Middle East and many other markets. Absence of comparable evidence should not be interpreted as absence of AI use.

5. What appears to be driving adoption?

The most straightforward driver is commercial relevance. Better product discovery can increase the probability that a shopper finds something suitable; better forecasting can reduce stockouts and excess inventory; better customer service can reduce response times; and automation can reduce the effort required to maintain catalogues, create content and analyse operations.

Digital retail also produces the data on which many AI systems depend. Search queries, purchases, product characteristics, inventory positions, returns, browsing behaviour and service interactions create continuous feedback. Large e-commerce platforms consequently possess particularly favourable conditions for developing and refining AI.

Competitive pressure matters as well. Consumers increasingly encounter AI-assisted product discovery through general-purpose AI services as well as retailers' own applications. A July 2026 New Zealand survey of 1,000 consumers who had used AI in their shopping journey found that shoppers were using AI to discover products, compare alternatives and make purchasing decisions. A separate Singapore survey commissioned by a technology provider found substantial recent consumer use of AI for product discovery and decision-making. These

are not estimates of all consumers in either country, but they illustrate a changing route into commerce.

The technology has also become easier to access. AI capabilities that once required a large internal data-science organisation are increasingly supplied through commerce platforms and packaged services. Singapore's programmes explicitly seek to make AI-enabled retail tools available to SMEs, while Shopify's infrastructure can expose merchants' catalogues to multiple AI channels without each merchant developing its own agentic-commerce system.

Government policy is an additional driver in some markets. China in June 2026 called for deeper integration of AI throughout e-commerce, intelligent retail and logistics, while Singapore expanded its retail digitalisation framework to emphasise practical AI adoption. These policies do not themselves establish that deployment has occurred, but they create institutional and financial support for further adoption.

6. What are the present limitations?

AI depends heavily on the quality of retail data. Incorrect product attributes, stale inventory information, poorly matched catalogue records or fragmented customer data can produce poor recommendations irrespective of model sophistication. The problem becomes particularly significant when an AI system acts rather than merely advises: an incorrect size, price, delivery promise or compatibility assessment can propagate directly into a transaction.

Forecasting remains uncertain. Retail demand changes because of promotions, competitors, weather, economic conditions, fashion, local events and unexpected shocks. AI can improve analysis without eliminating uncertainty. Models trained on historical relationships can perform poorly when customer behaviour or market conditions shift.

Generative AI introduces additional reliability problems. A fluent shopping assistant may provide an answer that appears authoritative even when product details are missing or incorrectly inferred. Retail systems can constrain this problem by grounding answers in product catalogues and transaction data, but grounding does not guarantee complete accuracy.

Integration is a major organisational limitation. Inventory, point-of-sale, e-commerce, customer, pricing, logistics and supplier systems may have developed independently. A retailer cannot readily optimise across the business if the relevant data are incompatible, incomplete or delayed. Smaller retailers may additionally lack specialised staff, suitable infrastructure or the financial capacity to experiment with advanced systems.

The economic case is not equally clear for every application. Deloitte's 2026 survey of 200 retail and consumer-products executives characterised the position as one in which strategic interest was running ahead of execution, with uncertainty around returns, governance and the ability to move pilots into production. The survey is useful evidence of executive experience but is not a representative measure of global retail adoption.

Consumer trust and data use impose further limits. Personalised recommendations depend partly on observing customer behaviour, while more sophisticated personalisation can involve combining numerous data points. US regulators have been examining “surveillance pricing”,

including whether personal information may be used to influence individual prices or discounts. This investigation demonstrates the regulatory concern but should not be interpreted as evidence that personalised AI pricing is routine across retail.

Algorithmic pricing also raises competition questions. The UK Competition and Markets Authority has noted that increasingly sophisticated pricing systems may create risks ranging from conventional collusion implemented through algorithms to more complex forms of algorithmic coordination. These are potential competition effects requiring scrutiny, not evidence that AI pricing generally produces collusion.

Agentic commerce adds responsibility and control questions. When an AI can interact with customers, make representations, process refunds or initiate transactions, errors become operational actions rather than merely inaccurate text. The UK CMA's March 2026 guidance makes clear that businesses remain responsible for unlawful actions undertaken by AI agents they deploy.

7. What you should be aware of

“AI in retail” covers technologies at very different stages of maturity. A recommendation algorithm used millions of times each day, a demand-forecasting model, a generative product-description tool and an AI agent purchasing goods autonomously should not be treated as equivalent evidence of adoption.

Announcements also require careful interpretation. A commerce platform making an AI capability available to millions of merchants demonstrates infrastructure and technical availability; it does not show that millions of merchants are actively using it or generating significant sales through it. Conversely, very large usage figures from one major platform demonstrate genuine deployment but cannot automatically be extrapolated to retail generally.

AI is increasingly shaping not only what happens after a customer enters a retailer's website but **how the customer reaches commerce in the first place**. Product discovery is beginning to move into general AI assistants, conversational interfaces and agent-mediated environments. The resulting shift is important because the AI system may increasingly decide which products are surfaced, how alternatives are described and sometimes which actions are taken. By 30 August 2026 this transition is real, but still developing.

Personalisation has a similar dual character. It can reduce information overload and make product discovery more relevant, while simultaneously increasing the importance of questions about data collection, ranking criteria, commercial influence and whether different consumers receive materially different offers or choices.

Human involvement remains significant throughout the sector. Merchandisers determine ranges and strategy; operational teams manage supply chains and exceptions; customer-service staff handle difficult cases; and customers usually retain final purchasing authority. Even where autonomous capabilities exist, the presence of an AI agent does not mean the surrounding retail process has become autonomous.

Geographical differences also matter. The frontier visible through Amazon, Walmart, Alibaba and sophisticated commerce platforms reflects organisations with exceptional data, scale and technical resources. Retail sectors containing large numbers of small, informal or only partly digitised enterprises begin from a different position. Singapore's experience—where policymakers still distinguish basic digitalisation from advanced AI adoption even within a highly digitised economy—is a useful reminder of this distinction.

Regulation is developing across several overlapping domains rather than through a single global retail-AI regime. Data protection, consumer protection, competition law, advertising rules, pricing transparency and AI-specific regulation can all become relevant depending on the application. In Europe, the AI regulatory framework was itself continuing to evolve during 2026; the United Kingdom was simultaneously developing specific guidance on agentic AI and consumers.

Finally, much of the most detailed public evidence on frontier retail AI comes from technology providers and large retailers themselves. Such evidence is valuable for establishing that a capability exists and that particular deployments have reached scale, but commercial performance claims require appropriate caution. Independent, internationally comparable evidence on adoption, productivity, customer outcomes and return on investment remains much less complete.

Status at 30 August 2026: AI is already materially embedded in Retail & E-commerce. Its most established position is in specialised prediction, ranking, recommendation, forecasting and optimisation within large digital and omnichannel retail operations. Generative AI is increasingly being deployed for product discovery, customer interaction, content creation and staff or merchant assistance. The newest development is the emergence of agentic commerce, in which AI can move beyond providing information towards undertaking parts of the transaction itself. That transition has reached genuine production deployment but remains substantially less mature and less uniformly adopted than the predictive AI systems that already underpin modern retail.

Evidence Notes

1. **Singapore EnterpriseSG and Infocomm Media Development Authority — Refreshed Retail Industry Digital Plan, 26 May 2026.** Particularly useful primary evidence concerning SME retail digital maturity, demand forecasting, AI concierges, generative-AI marketing and training applications, and the continuing gap between basic and advanced technology adoption. It is specific to Singapore and should not be generalised globally.
2. **Singapore IMDA — Retail Recommendation Engine, updated April 2026.** Demonstrates practical application of AI recommendations using point-of-sale, CRM and online behavioural data. Performance figures are programme/provider evidence and may vary by implementation.
3. **Amazon — Alexa for Shopping/Rufus deployment information, 2026.** Primary evidence for large-scale conversational shopping, recommendations, product comparison, price tracking, reordering and agentic functions. Amazon reported more than 250 million users during 2026. Usage and conversion information is company-reported and should not be treated as independent sector-wide evidence.

4. **Alibaba Group — Qwen/Taobao integration, 11 May 2026.** Important evidence from China of an AI assistant connected to Taobao's catalogue of more than four billion products and able to support browsing, comparison, ordering, logistics and after-sales interaction. It demonstrates deployment within Alibaba's ecosystem rather than adoption across Chinese retail generally.
5. **Cyberspace Administration of China — China New E-commerce Development Report 2026, released 29 August 2026.** Particularly timely official evidence describing scaled AI application across content creation, targeted marketing, intelligent customer service and supply-chain coordination in Chinese e-commerce. National reporting does not provide directly comparable adoption metrics for other countries.
6. **China Ministry of Commerce and other ministries — Implementation Opinions on accelerating “AI + Consumption”, June 2026.** Primary policy evidence covering intelligent retail, e-commerce operations, customer service, design, marketing, livestreaming, recommendation, quality inspection and logistics. Policy direction establishes government support and intended application areas, not completed adoption.
7. **Walmart — AI-enabled global supply-chain and merchandising systems, July 2025, together with 2025–26 customer-facing AI material.** Provides concrete organisational examples of inventory rebalancing, demand and logistics optimisation, merchant assistance and conversational shopping. As company evidence, it demonstrates particular deployments rather than retail-sector prevalence.
8. **Shopify — Agentic commerce announcements and platform releases, January–June 2026.** Evidence that infrastructure allowing large numbers of merchants to expose products to AI assistants and support AI-mediated checkout had moved into production availability. Merchant eligibility or platform enablement should not be confused with active utilisation or transaction volumes.
9. **Gamboa-Cruzado et al. — systematic and bibliometric review of machine learning in e-commerce, 2025.** Review of 66 studies from 2018–2024 providing academic support for established application territories including personalisation, dynamic pricing and sales forecasting. Academic literature demonstrates capability and research concentration more readily than present global deployment levels.
10. **Ozay, Jahanbakht and Wang — AI-Based Dynamic Pricing, August 2026.** Large bibliometric synthesis covering 1,301 publications and documenting the development of AI-based pricing methods and the prominence of e-commerce as an application environment. It principally establishes the development of the research and technology field, rather than measuring retailer adoption.
11. **UK Competition and Markets Authority — AI and collusion, March 2026; agentic-AI consumer guidance, March 2026.** Authoritative evidence concerning emerging competition, consumer-protection and accountability issues associated with pricing algorithms and AI agents. Discussion of risks does not establish that prohibited behaviour is prevalent.
12. **US Federal Trade Commission — Surveillance Pricing inquiry and research, 2024–25, current at the Status date.** Important regulatory evidence concerning the potential use of detailed behavioural and personal data in algorithmic pricing and discounting. The FTC itself describes parts of this field as still relatively nascent, so the material should not be read as evidence of widespread personalised pricing.
13. **Deloitte — State of AI Adoption in Retail and Consumer Products, June 2026.** Survey of 200 executives providing useful evidence that organisational enthusiasm was exceeding execution in areas such as governance, investment and production scaling. Its sample and methodology mean the findings should not be treated as a global adoption rate.
14. **PwC New Zealand — The AI-powered consumer, July 2026.** Survey of 1,000 New Zealand consumers who had already used AI in their shopping journey. It provides useful evidence about how AI users employ these tools for discovery and comparison, but because respondents were selected on prior AI-shopping use it cannot establish prevalence among the New Zealand population generally.

6. Hospitality & Tourism

Status as at 30 August 2026

Artificial intelligence is now being applied across much of the Hospitality & Tourism value chain, from travel discovery and hotel pricing to guest communications, operational planning and destination management. Its maturity, however, differs substantially by application. Some uses are well established. Hotels, online travel platforms and other operators have used machine learning for years to forecast demand, recommend products, optimise prices, analyse customer behaviour and automate routine interactions. Newer generative-AI applications are expanding rapidly into conversational travel search, itinerary creation, review summarisation, marketing content, multilingual communication and digital concierge services. By 2026, several major travel and hotel groups had moved such tools from demonstrations into customer-facing deployment, although sometimes still explicitly in beta.

The wider global picture remains uneven. Large hotel groups, travel platforms and highly digitised destinations have greater access to data, integrated systems, specialist capability and capital than many independent hotels, restaurants and small tourism businesses. OECD evidence describes AI as increasingly enabling personalised services, operational optimisation, improved forecasting and visitor-flow management, while also noting that adoption remains uneven and many initiatives are still emerging.

The present Status is therefore not that hospitality and tourism has become AI-operated. Rather, AI is increasingly embedded in particular commercial, analytical and service activities within a sector that continues to depend heavily on people, physical places, local knowledge and human service.

1. Where is AI currently being applied?

One of the most visible areas is **travel discovery, search and trip planning**. Recommendation systems have long helped travellers rank accommodation, destinations, flights and activities. Generative AI now enables travellers to describe requirements conversationally, ask follow-up questions and obtain suggested destinations or itineraries. Booking.com has deployed an AI Trip Planner and AI-supported property search, question answering and review summaries in multiple markets. Hilton introduced a generative-AI planner on Hilton.com in March 2026, while IHG launched AI conversational search in the United States in July 2026 after introducing an IHG application within ChatGPT. Expedia has similarly expanded natural-language planning tools and in July 2026 acquired AI-native trip-planning platform Layla.

AI is also important in **revenue management and commercial optimisation**. Hotels use predictive models to forecast occupancy, assess booking patterns, recommend room rates, manage inventory and help determine when and through which channel rooms should be offered. IHG, for example, describes AI-supported revenue management as part of its core operating-platform development. These activities build on comparatively mature forms of predictive analytics rather than the recent generative-AI wave.

A related field is **personalisation and customer relationship management**. AI can use booking history, loyalty information, preferences and behavioural data to decide which accommodation, offer, ancillary product or communication may be most relevant to a particular traveller. Online travel businesses have been doing versions of this for years; hotel groups are increasingly seeking to extend personalisation from marketing into the stay itself.

Customer service and guest communication are another significant area. Chatbots and virtual assistants can respond to routine questions about reservations, facilities, arrival times, policies and local services, often continuously and in multiple languages. Generative models can produce or translate customer communications, summarise conversations and assist service employees with suggested responses. Japan's Tourism Agency has documented use cases including AI concierge functions, while its accommodation-sector productivity work includes AI concierge systems alongside other labour-saving technologies.

Within hotel operations, AI is being used or tested for **staff scheduling, housekeeping and maintenance planning, procurement, energy optimisation, demand prediction, food-and-beverage operations and management reporting**. Some applications combine AI with sensors or automation. The distinction matters: a cleaning or delivery robot is not necessarily an AI system, although AI may contribute to navigation, scheduling, recognition or interaction.

AI is also increasingly involved in **marketing and reputation management**. Systems generate draft advertising copy, destination descriptions, social-media content and images; translate material; analyse reviews and customer sentiment; group complaints; and identify recurring service issues. Review summarisation is becoming customer-facing as well: rather than requiring travellers to read hundreds of reviews, AI can extract recurring themes and answer questions about a property.

Beyond individual businesses, **destinations and tourism authorities** are applying AI and machine learning to tourism intelligence. Potential and demonstrated uses include combining booking, mobility, payment and other data to estimate demand, understand visitor movements, forecast pressures and improve marketing or destination management. Saudi Arabia's tourism data system, for example, combines traditional tourism statistics with mobile-positioning, payment and accommodation-booking information and uses machine-learning and AI modelling. OECD evidence indicates that real-time visitor-flow management and improved tourism forecasting are becoming significant areas of digital development.

AI is also changing how destinations become **discoverable through third-party AI systems**. In July 2026, New Zealand announced investment to organise tourism information so that AI search tools could access it more effectively and to expand Tourism New Zealand's AI travel assistant with regional content. This represents an emerging application distinct from conventional digital advertising: destinations are beginning to consider how accurate information about their businesses and attractions appears within AI-mediated travel discovery.

2. What problems or needs is AI being used to address?

Hospitality and tourism contains an unusually large number of choices. Travellers may need to compare thousands of properties, destinations, activities, restaurants and combinations of dates, prices and preferences. Conventional search requires people to translate complex

intentions into filters and browse large volumes of material. Conversational and recommendation AI is being used to reduce this **search and information burden**, turning requests such as a family holiday with particular amenities, budget and location requirements into narrower sets of possibilities.

The sector also faces **volatile and perishable demand**. An unsold hotel room tonight cannot be stored and sold tomorrow. Demand changes with seasonality, events, weather, air capacity, economic conditions and competitor behaviour. Forecasting and revenue-management systems therefore address a commercially important need: estimating likely demand and allocating limited inventory and prices accordingly.

A third pressure is **labour availability and productivity**. Accommodation and food-service businesses frequently operate long hours, require substantial frontline staffing and experience seasonal demand. Japan's Tourism Agency explicitly links automation, AI concierge systems and related technologies to accommodation-sector labour shortages. New Zealand's 2026 tourism policy similarly identifies AI as a potential source of productivity improvements across tourism and hospitality.

AI is also directed at **service responsiveness and multilingual communication**. International visitors may require assistance at any time and in many languages. Automated translation, conversational assistants and employee-support tools can expand the range and speed of responses without requiring every location to maintain specialist language capability continuously.

Personalisation addresses another characteristic of the sector: travellers do not value the same experience in the same way. A business traveller, family, backpacker and luxury visitor may respond very differently to the same property or destination. AI is consequently used to interpret preferences and make product selection, marketing and guest communication more individually relevant.

Destinations face a different set of pressures. Tourism authorities need to understand **where visitors are going, when congestion is developing, which regions receive tourism expenditure and how changing demand may affect infrastructure and communities**. OECD analysis increasingly connects AI and improved tourism data with real-time management of visitor flows and evidence-based destination decisions.

Finally, generative AI is addressing a simpler but widespread problem: the amount of **routine text and content work** performed throughout the sector. Drafting property descriptions, answering common enquiries, translating information, summarising reviews, producing marketing variants and preparing internal reports can now be partially automated.

3. What can AI currently do?

Current AI can forecast demand from historical and current data, identify booking and customer patterns, rank accommodation or activities, recommend prices, classify customer feedback, detect sentiment and help identify which offers or communications are likely to be relevant to particular customers.

Generative AI adds a substantially more natural interface. Travellers can now describe a desired trip in ordinary language, revise requirements conversationally and receive proposed destinations, hotels or itineraries. Systems connected to live commercial inventories can increasingly combine that conversation with actual availability, pricing, amenities and booking pathways. IHG's 2026 conversational-search deployment, for example, connects natural-language discovery with hotel information and booking channels rather than operating purely as a general-purpose chatbot.

AI can summarise large numbers of reviews, answer questions about property characteristics, translate guest communications, draft personalised messages and provide automated answers to routine service questions. It can also analyse operating information faster than staff could manually, enabling managers to identify demand patterns, staffing requirements or recurring guest concerns.

At destination level, machine learning can combine multiple data streams and identify patterns that conventional periodic tourism statistics may reveal only slowly. Saudi Arabia's use of administrative, mobile-positioning, point-of-sale and accommodation-booking information demonstrates the movement toward more granular and timely tourism intelligence.

What these capabilities do **not** establish is that AI can independently deliver the complete hospitality or travel experience. An itinerary generator may organise useful suggestions but can still recommend closed attractions, unrealistic travel times or inaccurate information. A pricing system can estimate demand but cannot eliminate sudden disruption. A chatbot can resolve routine questions yet may struggle with an unusual complaint requiring discretion and empathy.

Research published by August 2026 reinforces this distinction. Studies of travellers identify information quality, reliability and trust as central to continued generative-AI use, while experimental research shows that hallucinated information can reduce willingness to continue using AI travel services.

4. How established is adoption?

Adoption is best understood as **layered and highly uneven**.

Predictive analytics, recommendation systems, revenue management and automated customer interactions are among the more established uses, particularly among large hotel groups, online travel agencies and digitally mature operators. Their presence predates the current generative-AI cycle.

Generative AI is moving rapidly from experimentation into actual use but remains less mature. Major platforms now provide customer-facing conversational travel tools, and hotel groups are integrating similar capabilities into owned channels. Yet some important deployments remain explicitly described as beta products. Availability from Booking.com, Hilton, IHG or Expedia demonstrates genuine deployment; it does not demonstrate that conversational AI has become the normal travel-planning method globally.

Traveller adoption is growing but similarly far from universal. Deloitte's 2026 US travel research reported that nearly one-quarter of surveyed travellers were using generative AI for trip planning in late 2025. Its earlier 2025 summer survey found 15% usage among surveyed travellers, with materially higher use among younger groups. These surveys show rapid growth within their defined populations but should not be interpreted as global traveller adoption rates.

Hotel-industry surveys also produce apparently high but methodologically different adoption estimates. A global hotel-chain study reported 78% of participating chains already using AI, but only 6% operating with a comprehensive AI strategy. A Middle East survey of 35 senior tourism and hospitality leaders found 91% piloting or using AI but only 3% reporting enterprise-wide implementation. The contrast between “some AI use”, formal strategy and organisation-wide deployment is important.

Evidence from smaller or broader workforces gives a more restrained picture. New Zealand's 2025 hospitality and tourism workforce survey reported that about one in three employees used AI at work and that 51% said managers provided no resources for AI initiatives. Use was higher in larger urban organisations.

Geography matters as well. Japan is supporting tourism-sector generative-AI experimentation and implementation; Middle Eastern destinations are investing heavily in digitally enabled tourism; Asia-Pacific governments have placed AI and tourism on the APEC agenda; and New Zealand is explicitly funding AI-readiness and tourism discovery. OECD's 2026 assessment across 53 member and partner economies nevertheless concludes that adoption levels remain uneven and many destination initiatives are emerging rather than mature.

5. What appears to be driving adoption?

A principal driver is the economics of **better matching demand with capacity**. Accommodation inventory is perishable, margins can be sensitive to occupancy and rate decisions, and tourism demand is volatile. AI that improves forecasting, pricing or marketing therefore addresses established commercial priorities.

A second driver is **labour pressure**. Automation of routine enquiries, administrative work, scheduling and selected service activities is attractive where organisations face recruitment difficulties, seasonal peaks or high labour costs.

The rapid improvement of generative AI has created another driver: capabilities that previously required purpose-built interfaces can increasingly be accessed through natural conversation. This matters particularly in travel because planning is inherently iterative. Travellers often begin with uncertain preferences, compare alternatives and refine their intentions. Conversational AI maps naturally onto that behaviour.

Competitive pressure is also increasing. Major online travel businesses and hotel groups are introducing AI-powered discovery into their own customer channels. This potentially changes how accommodation and destinations are found, compared and selected. New Zealand's decision to make tourism information more readable by AI systems illustrates that AI-mediated discoverability is becoming a policy and commercial consideration in its own right.

Growing digital data provides the underlying fuel. Property-management systems, reservation platforms, loyalty programmes, online reviews, payments, mobile devices and destination data systems generate information from which AI can identify patterns. Organisations with integrated, well-governed data are consequently better positioned than businesses whose information remains fragmented.

Governments and destination authorities are also encouraging adoption for broader reasons: productivity, regional visitor dispersal, tourism intelligence and management of destination pressures. OECD, APEC, Japan and New Zealand all provide evidence that AI has moved onto formal tourism-policy agendas rather than remaining solely a private technology initiative.

6. What are the present limitations?

The first limitation is **data and system fragmentation**. Hospitality businesses commonly operate property-management, reservation, customer, payment, loyalty and distribution systems that were not designed as a single data environment. AI can only personalise or optimise reliably to the extent that relevant information is available, accurate and usable. Integration is repeatedly identified in hotel-industry evidence as a significant barrier.

Generative AI introduces a different limitation: **factual reliability**. Travel information changes constantly. Prices, opening hours, visa requirements, attractions, transport schedules and local conditions can become outdated. A fluent model may produce plausible information without establishing that it is current or correct. Research conducted specifically in tourism finds that hallucination and perceived inaccuracy can undermine trust and continued use.

Personalisation also depends on information about individuals, creating **privacy, security and governance questions**. Travel records can disclose movements, preferences, purchases and sometimes sensitive personal circumstances. OECD's G7 tourism work identifies data and consumer protection as important policy issues as AI-enabled personalisation expands.

Smaller businesses face particular constraints. Many lack specialist technology staff, integrated datasets or the investment capacity available to global hotel groups and online platforms. Skills are another limitation: employees and managers need sufficient understanding to use outputs appropriately and recognise where AI-generated information requires checking.

Economic evidence remains uneven. A system may reduce response time or help produce marketing content without demonstrating that it improves profitability, guest satisfaction or destination outcomes. Provider case studies establish that deployments exist but cannot by themselves demonstrate sector-wide effectiveness.

There is also a service-design limitation peculiar to hospitality: **efficiency is not the only objective**. Human welcome, empathy, cultural knowledge, judgement during service failure and interpersonal interaction can form part of the product travellers have purchased. A 2026 hotelier survey, despite reporting extensive AI use among its respondents, found 59% believed the front-desk welcome and check-in should remain human-led. The precise percentage is survey-dependent, but the underlying distinction is important: technologically automatable does not necessarily mean commercially or socially desirable to automate.

Finally, the industry remains vulnerable to wider physical circumstances that AI cannot remove. Weather, geopolitical events, natural disasters, transport disruption, local infrastructure, staffing availability and sudden demand changes continue to determine tourism outcomes.

7. What you should be aware of

“AI in hospitality and tourism” covers materially different technologies and levels of responsibility. An algorithm forecasting tomorrow's occupancy, a chatbot answering a question about breakfast hours, a recommendation system ranking hotels and an AI agent assembling an entire trip should not be treated as equivalent.

Headline adoption figures require similar care. Surveys may count an employee using a general-purpose generative-AI tool, a hotel piloting one application or a chain operating AI across multiple core systems under the same broad description of “AI use”. The Healthcare & Life Sciences pilot similarly distinguishes availability, trials and real deployment rather than treating them as interchangeable measures of maturity.

The distinction between **recommendation and transaction** is also changing. Earlier generative-AI travel tools primarily provided inspiration or text-based itineraries. By 2026, systems were increasingly being connected to real inventory, pricing, maps and booking pathways. This moves AI closer to influencing—and potentially executing—commercial decisions, but fully autonomous end-to-end travel management should not be inferred from individual announced capabilities.

Human involvement remains important. Employees commonly handle exceptions, complaints, safety issues and situations requiring judgement; travellers themselves continue to verify or reject recommendations. Experimental evidence published in 2026 indicates that human recommendations can remain more persuasive than generative-AI recommendations in some near-term travel decisions, illustrating that technical capability and user preference are separate questions.

Geographical differences are substantial. Large digital travel platforms operate internationally, but evidence concerning organisational deployment is concentrated disproportionately in technologically advanced markets and large businesses. SME-heavy tourism economies may experience a very different adoption pathway. Even within one country, large urban hotels may have significantly greater capability than small regional accommodation providers.

AI may also affect which destinations and businesses become visible to travellers. As conversational systems increasingly mediate discovery, the information they can access, the sources they rely on and the way they rank or describe alternatives could influence visitor flows. New Zealand's 2026 decision to make official tourism data more accessible to AI systems provides an early practical example of governments responding to this development.

Finally, evidence concerning generative AI remains young. A 2025 review described tourism as still being in relatively early stages of understanding and applying generative AI, while a 2026 systematic review found that focused research on consumer-facing AI-generated tourism content had emerged principally since 2023. Rapid expansion of research and deployment since

then is important, but long-term evidence about commercial outcomes, traveller behaviour, labour consequences and destination effects is necessarily much thinner than the level of current publicity may suggest.

Status at 30 August 2026: AI has moved materially into Hospitality & Tourism, but not uniformly. Predictive analytics, recommendations, revenue management and selected forms of customer-service automation are relatively established, especially among large digital platforms and hotel groups. Generative AI is advancing rapidly into travel discovery, conversational search, guest communication, marketing and service assistance, with several substantial deployments now operating but some important applications still in beta or early-scale use. Destination-level AI is developing for tourism intelligence, visitor management and AI-mediated discoverability. Across the sector, the strongest current position is AI embedded within wider human and digital processes rather than independently delivering the complete travel or hospitality experience.

Evidence Notes

1. **OECD — *Tourism Trends and Policies 2026, June 2026*.** Provides the broadest recent multi-country policy assessment, covering 53 OECD member and partner economies. It identifies AI and digitalisation as increasingly important to personalised tourism services, operational efficiency, forecasting and visitor management while explicitly noting uneven adoption. It is stronger for system-level Status than for measuring individual business deployments.
2. **OECD — *Artificial Intelligence and Tourism: G7/OECD Policy Paper, December 2024*.** Principal international sector-specific framework used for AI opportunities, consumer/data protection, employment implications and SME constraints. Its publication predates some of the rapid generative-AI deployments seen in 2025–26.
3. **Booking.com / Booking Holdings — *AI Trip Planner and generative-AI travel features, 2023–26*.** Primary evidence that conversational planning, natural-language property search, property question answering and review summarisation have moved into actual consumer-facing travel-platform deployment. Provider evidence establishes availability, not independent effectiveness or general adoption.
4. **Hilton — *Hilton AI Planner, March 2026*.** Primary evidence of a major global hotel company deploying a customer-facing generative-AI travel planner. The launch was described as beta, which limits conclusions about maturity.
5. **IHG Hotels & Resorts — *conversational AI search, June and July 2026*.** Primary evidence of movement from AI discovery through ChatGPT into conversational hotel search integrated with IHG's own digital channels. The July US deployment was explicitly a beta launch.
6. **Expedia Group — *AI travel developments, May and July 2026*.** Primary evidence of natural-language planning and continued investment in AI-native trip planning, including the acquisition of Layla. Announced strategy and acquisition activity should not be equated with sector-wide adoption.
7. **Japan Tourism Agency — *generative-AI tourism guidance and accommodation productivity material, 2025–26*.** Government evidence from Asia concerning practical generative-AI experimentation, tourism DX, AI concierge applications and responses to accommodation-sector labour shortages.
8. **New Zealand Ministry of Business, Innovation & Employment — *Tourism Policy Statement and AI tourism investment, June–July 2026*.** Primary evidence from Oceania concerning AI-enabled tourism productivity, destination discoverability, AI-ready tourism information and expansion of Tourism New

Zealand's AI travel assistant. These initiatives include policy and investment activity rather than evidence that all tourism businesses have adopted AI.

9. **New Zealand MBIE — *Hospitality and Tourism Employment: Strengthening Futures 2025***. Workforce survey evidence reporting AI use by approximately one-third of respondents and limited organisational support. Valuable counterweight to high adoption figures from technology-industry surveys, but geographically restricted to New Zealand and based on employee self-report.
10. **OECD — Saudi Arabia Data and Decision Support Deputyship case study, June 2025**. Useful evidence of destination-level integration of traditional statistics, mobile-positioning, payments, accommodation-booking data and machine-learning/AI modelling. It is a country case study rather than evidence of common global practice.
11. **PwC Middle East/Future Hospitality Summit — *AI at the Heart of Tourism and Hospitality*, December 2025**. Survey of 35 senior GCC tourism and hospitality leaders. Reports extensive piloting/use but only 3% enterprise-wide implementation, illustrating the difference between adoption activity and scale. The sample is small, senior and regionally concentrated.
12. **Cloudbeds/h2c Research — *AI and automation hotel-chain study, 2025***. Survey evidence covering 171 hotel chains and more than 11,000 properties, reporting significant AI use but few comprehensive AI strategies and substantial expertise/integration barriers. Vendor-associated research and respondent composition limit generalisation to the whole accommodation sector.
13. **Deloitte — US travel surveys and 2026 Travel Industry Outlook**. Provides evidence that traveller use of generative AI for planning is increasing and that some users act on AI recommendations. These are US consumer surveys and should not be treated as global adoption measures.
14. **Peer-reviewed tourism and hospitality research, 2025–August 2026**. Studies in the *International Journal of Hospitality Management*, *Tourism Management*, *Journal of Hospitality and Tourism Technology* and related journals provide evidence concerning traveller trust, privacy, information quality, hallucination, recommendation acceptance and the still-developing generative-AI evidence base. Much of this research measures attitudes, experiments or behavioural intention rather than routine operational outcomes.
15. **APEC Tourism Working Group — AI and Tourism initiative, 2025**. Evidence that AI's implications for tourism innovation, sustainability and inclusion have become a formal Asia-Pacific policy issue. The programme establishes policy attention rather than measuring business adoption across APEC economies.

7. Transport & Logistics

Status as at 30 August 2026

Artificial intelligence is already embedded in important parts of transport and logistics, but its maturity varies substantially by activity and operating environment. The most established applications are generally not fully autonomous vehicles. They are narrower systems used to predict demand and journey times, optimise routes and loads, manage fleets and assets, detect anomalies, support warehouse and port operations, improve traffic management and help organisations interpret large volumes of operational information.

Higher levels of physical autonomy have nevertheless moved beyond laboratory demonstration. Commercial robotaxi services now operate in selected cities; autonomous trucks are carrying commercial freight on defined US corridors; autonomous shuttles are carrying passengers in Singapore; and remotely operated and autonomous vessels are being tested internationally. China has developed particularly substantial robotaxi operations: Baidu reported more than 22 million cumulative Apollo Go public rides by April 2026. These developments demonstrate real capability, but they remain geographically and operationally bounded rather than representative of transport generally.

The present Status is therefore best understood as one in which AI increasingly assists the movement, scheduling and management of people, vehicles and goods, while progressively assuming control of selected transport tasks under defined conditions. Fully autonomous transport across unrestricted environments remains much less established.

1. Where is AI currently being applied?

A large and relatively mature area is **transport and logistics optimisation**. AI and machine-learning systems analyse historical and real-time information concerning orders, traffic, weather, vehicle location, capacity and delivery requirements to support routing, dispatch, estimated arrival times, load consolidation, network planning and demand forecasting. Logistics providers increasingly combine these capabilities with telematics and tracking systems so that plans can be recalculated as conditions change. DHL's Logistics Trend Radar identifies advanced analytics, computer vision and generative AI among prominent AI developments affecting logistics, including applications in route planning, demand forecasting and asset management.

AI is similarly applied to **fleet and asset management**. Data from vehicles, engines, tyres, handling equipment and infrastructure can be analysed to detect abnormal behaviour, anticipate maintenance requirements and support decisions about when equipment should be serviced. This extends across road fleets, rail, ports, shipping and aviation. In maritime transport, Singapore's Maritime and Port Authority describes uses including detecting equipment failures earlier, improving vessel performance and assisting ship managers with operational and regulatory decisions.

Warehousing and logistics facilities provide another major application environment. Computer vision can identify, count and track packages or pallets, detect damaged goods,

monitor loading processes and assist automated sorting and picking. Machine-learning systems support inventory positioning, labour and equipment allocation, demand prediction and the orchestration of increasingly automated facilities. Robotics may incorporate AI for perception and movement, although not every warehouse robot or automated conveyor should itself be classified as AI. DHL identifies computer vision as an increasingly significant logistics technology but also notes practical requirements such as adequate lighting and suitable operating environments.

Within **ports and maritime logistics**, AI is being used for vessel and berth planning, equipment management, document processing, predictive maintenance and decision support. Singapore, for example, has developed AI applications for processing shipping documentation and has been expanding sector initiatives covering ship management, chartering, shipping and bunkering operations. Its DocuMind and DocuMatch systems illustrate a relatively mundane but consequential form of AI: extracting and checking information from maritime documents rather than navigating a vessel autonomously.

AI also contributes to **traffic and transport-system management**. Transport authorities can use predictive analytics to estimate congestion, passenger flows and disruptions; optimise traffic signals or service scheduling; analyse incidents; and support planning and enforcement. The OECD's International Transport Forum identifies predictive analysis and automated systems as practical applications through which transport authorities are already seeking to improve transport operation, service delivery and inclusivity.

In **aviation**, AI applications include flight and disruption planning, aircraft maintenance, passenger and baggage processes and increasingly sophisticated assistance for air-traffic and airspace management. In June 2026 the US Federal Aviation Administration selected new technologies for national air-traffic flow management intended to coordinate schedules and trajectories and reduce congestion and delays. Europe is simultaneously establishing aviation-specific assurance requirements for AI, reflecting movement from experimentation towards regulated operational applications while more autonomous aviation remains considerably less mature.

The most publicly visible application is **automated driving**. Driver-assistance technologies are already present in production vehicles, while more capable automated-driving systems operate in geographically restricted commercial or trial environments. Baidu's Apollo Go reported 3.2 million fully driverless operational rides during the first quarter of 2026 and operation or testing across numerous international locations. Singapore had opened fixed-route autonomous shuttles in Punggol to public riders by April 2026, and more than 11,500 unique riders had used the services by 12 July.

Autonomous freight has also moved into real operation. In the United States, Aurora-powered trucks were carrying commercial freight for logistics customers on defined routes during 2026. Aurora reported nearly 440,000 driverless miles between commercial launch and the end of June, while Volvo Autonomous Solutions began autonomous commercial freight operations for DSV between Dallas and Houston. These examples establish genuine deployment, but on particular corridors and with fleets that remain very small compared with conventional trucking.

At sea, autonomous and remotely operated shipping remains earlier in development. The International Maritime Organization reported in May 2026 that fully crewless and remote-

operated ships remained limited, although trials were increasing internationally. Its adoption of the Maritime Autonomous Surface Ships Code provides the first comprehensive international safety framework for such operations.

2. What problems or needs is AI being used to address?

Transport and logistics operate under a fundamental optimisation problem: large numbers of people, vehicles and goods must move through networks whose capacity, demand and operating conditions continuously change. Human dispatchers and planners can manage this complexity, but the combinations become increasingly difficult as networks grow. AI is consequently attractive where large datasets must be converted quickly into routing, scheduling, allocation or exception-management decisions.

A related problem is **under-used capacity**. Trucks may travel partially loaded or empty, warehouses may position stock inefficiently, and vessels or aircraft may lose time waiting for facilities or routes. Predictive and optimisation systems are used to match capacity more closely with demand and to improve vehicle, container, warehouse and infrastructure utilisation. The World Economic Forum's work on freight logistics identifies route optimisation, capacity utilisation and modal choices as areas where AI can potentially improve both efficiency and environmental performance.

Uncertainty and disruption create another strong need. Congestion, weather, accidents, industrial action, equipment failures and geopolitical events can invalidate previously efficient plans. AI can continually reanalyse information and help operators identify an emerging problem, estimate its consequences and evaluate alternative routes or schedules. This does not make supply chains predictable, but it can shorten the time between disruption and operational response.

Transport is also equipment-intensive. Unexpected failure of an aircraft component, truck, locomotive, crane or shipboard system can interrupt entire chains of movement. Predictive-maintenance applications seek to identify patterns preceding failure so that intervention can occur before disruption becomes critical.

Workforce availability and operating cost are material drivers in several modes. Trucking, warehousing, shipping and public transport face varying degrees of recruitment difficulty, shift constraints or repetitive work. Singapore explicitly identifies manpower constraints as one reason for deploying autonomous systems at airports, seaports and on roads. Autonomous freight developers similarly emphasise the possibility of increasing long-haul utilisation while reallocating human drivers to other parts of transport networks.

Safety provides a further motivation. AI-supported perception systems can monitor roads, equipment and infrastructure continuously, while driver-assistance technologies can intervene in particular hazardous situations. US government-industry research found substantial reductions in front-to-rear crashes associated with automatic emergency braking in production vehicles. That evidence concerns a defined assistance technology, however, and should not be translated into proof that unrestricted autonomous driving is safer in all circumstances.

Finally, transport produces very large administrative information flows. Freight documentation, customs information, certificates, manifests, maintenance records and regulatory requirements remain labour intensive. Generative AI, document-recognition systems and language technologies are increasingly being used to extract, summarise, translate and check this information. Singapore estimated that an AI-enabled port-clearance system under development could eventually reduce some clearance processing from days to minutes, although that estimate concerns a developing system rather than demonstrated sector-wide performance.

3. What can AI currently do?

Current AI can perform defined perception, prediction and optimisation tasks at operationally useful levels.

It can estimate demand, traffic, arrival times and equipment condition from complex data; recommend or recalculate routes; allocate loads and resources; detect anomalies; identify objects in warehouse or traffic imagery; extract information from transport documents; and generate summaries or operational assistance from large bodies of text. These capabilities are particularly valuable because they can operate continuously across data volumes that would be impractical for individual staff to process manually.

AI can also exercise considerably more direct physical control than these examples imply. Automated-driving systems can perceive surrounding traffic, plan movement, steer, accelerate and brake without continuous human control inside defined operational domains. The scale achieved by Apollo Go demonstrates that this capability is no longer merely experimental: millions of driverless passenger journeys have occurred. Commercial autonomous trucks similarly demonstrate the technical ability to move freight without a person behind the wheel along approved routes.

The capability remains conditional. Weather, road layout, construction, unusual objects, emergency vehicles, connectivity, maps and interactions with unpredictable road users can all affect performance. Aurora reported that adverse weather had constrained its Texas driverless operations around 40% of the time during 2025 before a subsequent software expansion of its operating domain, illustrating how apparently simple statements that a vehicle “can drive itself” conceal important operating qualifications.

AI can therefore exceed people on some narrow transport calculations or perception tasks without possessing a general human understanding of the whole transport environment. Generative systems present an additional distinction: they can produce useful explanations, summaries, instructions and documentation, but plausible language is not equivalent to verified operational information.

4. How established is adoption?

Adoption is best described as **layered**.

Predictive analytics, optimisation, routing, estimated-arrival-time calculation, telematics-based fleet analysis and various forms of computer vision are established or increasingly established in digitally advanced logistics operations. They typically augment existing transport-management, warehouse-management and fleet-management systems rather than replace the organisation's operational structure.

Generative AI is growing rapidly but is less mature. Logistics providers are exploring it for customer interaction, document handling, employee assistance, knowledge retrieval and planning. Provider and industry reports demonstrate substantial interest and real deployments, but they do not establish that generative AI has become uniform across a globally fragmented sector containing multinational carriers, small trucking businesses, ports, postal systems, public operators and informal transport.

Autonomous road transport occupies a different maturity level. China now provides strong evidence of scaled robotaxi operation in selected cities, while US commercial operations and Singaporean public shuttle deployments confirm that autonomous passenger mobility is not merely a research proposition. Yet the geographical boundaries remain striking. Singapore's transport ministry stated in 2026 that there were fewer than 7,000 autonomous passenger vehicles of the relevant type globally and that mass deployment on its roads would still take time.

The regulatory environment reinforces this interpretation. In June 2026 UNECE adopted a global regulatory framework for Automated Driving Systems capable of operating without driver supervision. This is an important maturity milestone, but regulatory provision for deployment is not evidence that such vehicles have already become routine worldwide.

Autonomous trucking is commercially real but earlier in scale. The US examples involve defined routes and limited fleets. Europe's position was still substantially pre-deployment: in June 2026, 18 EU Member States agreed to establish cross-border testbeds aimed at coordinated testing and commercial pre-deployment of autonomous passenger and freight vehicles.

Maritime autonomy is earlier still. IMO's non-mandatory MASS Code took effect on 1 July 2026, but the organisation explicitly notes that the number of fully crewless or remotely operated vessels remains limited.

The global picture is therefore uneven. China and parts of the United States lead some high-autonomy road deployments; Singapore provides significant tightly governed passenger, port and logistics experimentation and use; Europe has substantial intelligent-transport deployment while proceeding cautiously towards higher automation; and evidence on advanced deployment is less extensive for many low- and middle-income transport systems. That difference should not be interpreted simply as different enthusiasm for AI: digital infrastructure, capital, vehicle fleets, road environments, regulation and data availability differ materially.

5. What appears to be driving adoption?

The strongest general driver is the economic importance of **better utilisation**. Vehicles, aircraft, vessels, warehouses and transport infrastructure are expensive assets. Even modest improvements in routing, turnaround, fuel consumption, loading or downtime can therefore have substantial value when repeated across large networks.

Data availability has made such optimisation increasingly feasible. GPS, telematics, sensors, electronic freight records, digital tickets, warehouse systems, cameras and connected infrastructure generate detailed operational information. AI becomes more useful as this digital foundation improves.

Competitive and service pressures also matter. Customers increasingly expect accurate arrival information, rapid delivery and visibility of goods. Operators facing these expectations are using forecasting and optimisation technologies to manage increasingly complicated service commitments.

Labour pressures encourage automation of both administrative and physical work. The attraction is particularly evident in repetitive warehouse activities, port operations and long-haul transport, although actual workforce consequences depend heavily on how technology is deployed.

Safety, congestion and sustainability provide public-policy drivers. Transport authorities are interested in AI for network management because the objective is not merely commercial productivity: traffic flow, accessibility, emissions and road safety are also public concerns.

Finally, improving AI capability and an emerging regulatory framework are making previously experimental applications more deployable. The 2026 UNECE automated-driving regulation and IMO MASS Code illustrate a broader transition from asking whether autonomous transport is technically imaginable to specifying how it should be validated, monitored and governed.

6. What are the present limitations?

Transport exposes AI to unusually difficult real-world conditions. Roads, ports, airports and logistics networks are open environments containing people, weather, damaged infrastructure, unusual objects, equipment failures and events that may be poorly represented in training data. A system that performs reliably during normal operations can therefore encounter rare circumstances in which its previous experience provides insufficient guidance.

The difficulty is especially visible in autonomous driving. In 2026 NHTSA recalled 3,791 Waymo automated-driving systems because software could permit vehicles to proceed into standing water on higher-speed roads. NHTSA also continued collecting mandatory crash reports for automated and Level 2 driver-assistance systems and warned that its data require careful interpretation because reports can contain duplicates, missing information and classification differences. These examples show both that operational learning is occurring and that significant edge cases remain.

Interactions with people create another problem. In July 2026 NHTSA publicly raised concerns about incidents involving driverless vehicles and first responders. A transport system cannot be judged only by whether a vehicle follows a mapped route: it must also respond appropriately to police, emergency personnel, pedestrians and unconventional road situations.

Data quality and interoperability constrain less visible AI applications. Transport information is distributed among carriers, shippers, ports, customs authorities, infrastructure operators and customers. Systems may use incompatible formats, incomplete records or commercially sensitive information. The EU's continuing electronic freight-transport-information programme demonstrates that creating interoperable digital information flows remains a substantial project even before sophisticated AI is applied to them.

Integration with legacy infrastructure can also be expensive. AI recommendations have limited value if dispatch, warehouse, vehicle or regulatory systems cannot act on them. Smaller operators may lack sufficient data, specialist staff or capital to adopt the same systems as large logistics groups.

Cybersecurity becomes more consequential as operational control becomes digital. Manipulation of vehicle, navigation, warehouse, traffic-management or shipping systems could have physical effects. Autonomous-transport regulation therefore increasingly treats cybersecurity, connectivity, system design and operational monitoring as safety issues rather than conventional information-technology concerns.

Generative AI adds familiar reliability problems: incorrect outputs, weak understanding of specialised rules and uncertainty over confidential information. Computer vision and worker-monitoring applications also create privacy and workplace-governance questions. DHL specifically identifies privacy, compliance, environmental requirements and specialist skills among constraints affecting AI logistics applications.

Finally, economic evidence remains uneven. A technically successful pilot may not prove that an application is cheaper or more reliable than existing practice after hardware, software, integration, supervision, insurance and maintenance costs are included.

7. What you should be aware of

The most important qualification is that **“AI in transport” encompasses very different technologies and evidential questions**. A system recommending a truck route, computer vision counting parcels, an aircraft-maintenance predictor, a driver-assistance system and a fully driverless robotaxi should not be treated as equivalent examples of adoption.

Automation levels also matter. Many production vehicles contain systems that assist with braking, steering or lane keeping while the driver remains responsible. NHTSA explicitly distinguishes these Level 0–2 assistance technologies from automated-driving systems capable of performing the driving task without continuous human supervision. Claims that a vehicle contains “AI” or “autonomous driving” therefore require careful interpretation.

Real deployment should likewise be separated from scale. Millions of robotaxi journeys in China are meaningful evidence of commercial maturity, but they do not establish that

autonomous vehicles can operate economically in every city or road environment. Conversely, Europe's comparatively cautious deployment should not be interpreted as lack of technical development: regulation, liability and public-transport structures influence how quickly capability becomes operational service.

Human involvement remains significant even in highly automated systems. Remote operations centres, fleet supervisors, safety operators, maintenance staff, dispatchers and regulators continue to form part of many autonomous operations. IMO's MASS framework, for example, preserves responsibility for the ship's master even where that person is not physically aboard.

Safety comparisons also require care. Autonomous systems generate extensive operational data, but differences in road type, weather, mileage, reporting rules and human intervention make simple comparisons between autonomous and human driving difficult. A provider's reported safety record can establish its experience; independent evidence and comparable exposure measures are required before broader conclusions are drawn.

The geographical evidence base is uneven. China, the United States, Singapore and Europe currently provide unusually visible evidence because they combine substantial investment, digitised infrastructure and active regulatory programmes. Reliable public evidence on AI adoption across much of Africa, Latin America, South and Southeast Asia, the Middle East and smaller transport markets is less systematic. Global Status should therefore not be inferred from a small number of highly digitised jurisdictions.

Status at 30 August 2026: AI has become a meaningful operational technology across Transport & Logistics. Its strongest current position lies in analysing, predicting and optimising movement: routing, scheduling, forecasting, fleet and asset management, traffic operations, document processing and increasingly warehouse and port activities. Physical autonomy has also crossed an important threshold into genuine passenger and freight deployment, most visibly in China, the United States and selected Asian locations. Yet autonomous road, maritime and other transport remains constrained by operating domains, regulation, infrastructure and unresolved safety and integration challenges. The sector is therefore not moving uniformly from human transport to autonomous transport; it is progressively inserting AI into particular decisions and physical tasks within still predominantly human-governed transport systems.

Evidence Notes

1. **International Transport Forum/OECD — *AI for Transport Authorities: Principles and Practical Guidance*, September 2025.** Provides authoritative cross-modal evidence concerning predictive analytics, automated systems and AI applications by transport authorities, together with governance and implementation limitations. Its focus is primarily public transport policy and administration rather than the full private logistics market.
2. **UNECE — *Automated/Autonomous Vehicle regulatory work and June 2026 Automated Driving Systems framework*.** Principal evidence that internationally harmonised regulation has reached systems capable of operating without driver supervision. Regulatory readiness should not be interpreted as widespread operational adoption.

3. **Baidu — Q1 and Q2 2026 corporate results concerning Apollo Go.** Provides primary operating evidence for one of the world's largest robotaxi deployments, including millions of quarterly rides, more than 22 million cumulative public rides by April 2026 and international expansion. These are company-reported figures and should not by themselves be treated as independent comparative evidence of safety or economic performance.
4. **Singapore Ministry of Transport and Land Transport Authority — autonomous-vehicle publications, March–July 2026.** Primary evidence concerning Punggol public autonomous shuttle operations, passenger numbers, regulatory assessment, safety-operator requirements and Singapore's distinction between trials, public rides and future commercial deployment. Particularly useful for demonstrating staged adoption rather than equating public availability with unrestricted autonomy.
5. **US National Highway Traffic Safety Administration — automated-vehicle safety, crash reporting and 2026 regulatory actions.** Provides authoritative evidence concerning distinctions between driver assistance and automated driving, continuing incident monitoring, recalls and emerging operational problems including interactions with first responders. Crash-report datasets contain methodological limitations explicitly identified by NHTSA and should not be used for simplistic human-versus-AV safety comparisons.
6. **NHTSA — Waymo ADS recall, May 2026.** Specific evidence that deployed high-autonomy systems can encounter materially important edge cases; approximately 3,791 automated-driving systems were covered by the standing-water-related recall. This is evidence concerning one technical issue and not a general assessment of autonomous-vehicle reliability.
7. **Aurora Innovation/Volvo Autonomous Solutions/DSV — 2026 autonomous trucking operational reporting.** Primary evidence that driverless and autonomous freight has entered genuine commercial logistics operations on selected US corridors, including Dallas–Houston and other Sun Belt routes. Company reporting establishes deployment and mileage but should not be generalised to the trucking sector or treated as independent proof of comparative safety or economics.
8. **International Maritime Organization — Maritime Autonomous Surface Ships Code, adopted May 2026 and effective 1 July 2026.** Principal international evidence concerning autonomous maritime transport. Particularly important because IMO simultaneously records a major regulatory milestone and acknowledges that fully crewless and remotely operated ships remain limited in number.
9. **Maritime and Port Authority of Singapore — maritime AI initiatives, 2025–July 2026.** Evidence concerning document processing, predictive maintenance, port and vessel optimisation, workforce capability and broader maritime-sector AI adoption. Some applications remain pilot or development programmes and are identified accordingly.
10. **European Commission — autonomous-vehicle cross-border testbeds, June 2026.** Evidence that substantial European activity remained at coordinated testing and commercial pre-deployment stage for important autonomous passenger and freight applications, illustrating regional differences in maturity.
11. **European Commission — Electronic Freight Transport Information Regulation implementation.** Provides useful context concerning the continuing digitisation and interoperability requirements underlying AI-enabled freight systems. It concerns digital freight infrastructure rather than AI adoption itself.
12. **European Union Aviation Safety Agency — aviation AI programme and 2025–26 regulatory work.** Evidence that aviation AI has progressed sufficiently to require specific assurance guidance covering machine-learning assistance and human-AI teaming, while more advanced autonomy remains under development and regulatory examination.
13. **US Federal Aviation Administration — air-traffic-management modernisation announcement, June 2026.** Primary evidence of advanced computational decision-support being introduced into national airspace-flow management. The announcement establishes procurement and intended operational capability rather than measured system-wide outcomes.
14. **DHL — Logistics Trend Radar 7.0, September 2024 and related AI analysis.** Provides logistics-industry evidence concerning advanced analytics, computer vision, generative AI and their application to routing, forecasting, asset management and warehouse processes. It is a provider-produced trend assessment and is

therefore used to establish use cases and industry direction rather than independently quantify global adoption.

15. **World Economic Forum with McKinsey & Company — *Intelligent Transport, Greener Future*, January 2025.** Useful cross-sector evidence concerning freight optimisation, capacity utilisation and the relationship between AI, operating efficiency and transport emissions. Its modelled potential benefits should not be interpreted as realised sector-wide outcomes.
-

8. Construction & Real Estate

Status as at 30 August 2026

Artificial intelligence is being applied across much of the built-environment lifecycle—from design and estimating through construction-site monitoring to property valuation, building operation and portfolio analysis—but its maturity is highly uneven. The strongest current position is not one of autonomous design, construction or property decision-making. It is AI performing defined analytical, predictive, visual and documentation tasks within processes still substantially directed and checked by people.

Construction remains comparatively early in organisational adoption. A 2025 Royal Institution of Chartered Surveyors (RICS) survey of more than 2,200 professionals found that 45% reported no AI implementation in their organisations and 34% were in early pilot phases. Just under 12% reported regular use in specific processes, while fewer than 1% described organisation-wide embedded use. The evidence nevertheless shows genuine deployments in progress monitoring, document analysis, scheduling, estimating and some site operations.

Real estate presents a different picture. Automated valuation, property-data analysis, predictive building maintenance and energy optimisation build on technologies and datasets that pre-date the recent generative-AI wave. At the same time, generative AI has prompted a large expansion of experimentation in lease and document analysis, portfolio intelligence, research and professional workflows. JLL's 2025 global technology research found that 88% of surveyed real-estate investors and 92% of corporate-real-estate occupiers had begun or planned AI pilots, but only a small minority reported achieving most programme objectives. That makes real estate a particularly clear example of why widespread experimentation should not be equated with mature scaled deployment.

1. Where is AI currently being applied?

Design and planning are important areas of application. AI-enabled tools can generate and compare design alternatives, analyse spatial arrangements, assist modelling and documentation, test options against constraints and help interrogate building information models (BIM). Generative systems are also being used to create draft specifications, summaries, reports and visual concepts. The significance lies less in AI independently designing complete buildings than in accelerating particular stages of an iterative professional process. RICS' global construction survey found design optioneering to be the area respondents most frequently expected AI to influence strongly, while Singapore's Building and Construction Authority (BCA) identifies AI-enhanced BIM, model checking, design review and documentation among present built-environment use cases.

Estimating, procurement and project management form another substantial cluster. Current systems analyse tenders and contracts, assist quantity take-offs and cost estimates, search project documents, summarise correspondence and support scheduling and risk analysis. Machine-learning approaches can compare current schedules with historical project patterns to identify activities that may be susceptible to delay. Generative AI is increasingly incorporated

into the document-heavy administrative layer surrounding projects, although the reliability of outputs varies with the underlying data and task.

Construction-site monitoring provides some of the clearest evidence of real deployment. Computer vision can process photographs, 360-degree imagery, video, drones and scans; locate imagery spatially against drawings or BIM; compare observed conditions with planned work; estimate progress; and flag discrepancies. Provider evidence shows deployments across substantial live-project portfolios. For example, Sir Robert McAlpine reported use of Buildots across five major projects covering more than 260,000 square metres, while OpenSpace reported in June 2026 that its visual-intelligence technology had been used on more than 1,000 data-centre construction projects. These examples establish real operational use, although vendor case studies should not be generalised into sector-wide adoption or treated as independent proof of claimed productivity benefits.

AI is also being applied to **safety and site operations**. Image analysis can assist recognition of protective equipment, people, equipment and potentially unsafe situations; AI can support safety planning and analyse incident information. Builders including Skanska, Turner Construction and Balfour Beatty were reported in 2026 to be using AI in safety-related activities. More automated physical operations remain less mature. Japan's Obayashi, for example, was in 2025 verifying a digitally integrated crane system incorporating simulation, machine guidance and remote or automated operation—significant technological development, but explicitly still a verification activity rather than evidence that autonomous cranes are normal construction practice.

Materials and construction logistics can also incorporate AI. Singapore's BCA documented use of Pan-United's AI-supported digital platform during a major 35-hour concrete foundation pour in January 2025, including dynamic scheduling, material replenishment and vehicle queuing. This illustrates AI embedded in a defined operational system rather than being used as a general-purpose assistant.

Within **real estate and property**, AI is being applied to valuation and market analysis, property search and matching, investment research, lease and document extraction, portfolio analysis, forecasting, customer interaction and property-management workflows. Automated valuation models (AVMs)—which may use statistical or machine-learning methods—are already sufficiently established for RICS to maintain dedicated professional material on their use. In the public sector, the UK Valuation Office Agency reported a production AVM supporting valuation work concerning 1.5 million domestic properties for the Welsh Council Tax revaluation.

Finally, **facilities and building operations** are an important bridge between construction and long-term asset management. AI can analyse sensor and building-management data to forecast energy consumption, identify anomalies, optimise heating, ventilation and air-conditioning, predict equipment problems and support maintenance planning. AI combined with BIM and digital twins is being developed and deployed for these purposes, although the sophistication and prevalence of such systems vary substantially between buildings and markets.

2. What problems or needs is AI being used to address?

Construction has longstanding problems of **productivity, coordination and uncertainty**. Large projects combine designers, engineers, contractors, subcontractors, suppliers and owners; generate thousands of documents and design changes; and take place in physical environments where actual progress may diverge from plans. Delayed or subjective information can make problems visible only after they have affected cost or schedule. AI is therefore particularly attractive where it can process project information faster, produce more current views of site conditions or identify patterns in data that would otherwise require extensive manual review.

A related problem is **fragmented information**. Project knowledge may sit across BIM models, drawings, schedules, photographs, contracts, emails, cost systems and site reports. Generative and analytical AI is increasingly being directed at extracting, comparing and retrieving information from these sources. Singapore's BCA, for example, identifies querying regulations and project documents, contract review, cost estimation, inspections and automated reporting as practical current built-environment use cases.

Construction also faces **skills and capacity constraints**. Experienced planners, surveyors, engineers and site managers must interpret growing amounts of project information while maintaining oversight of work conducted by multiple trades. AI can reduce repetitive reporting and checking or direct attention towards anomalies. RICS' 2025 survey found lack of skilled personnel itself to be the most frequently reported AI-adoption barrier, showing that workforce constraints simultaneously encourage interest in automation and make implementation more difficult.

Safety is another important need. Construction sites change continuously and contain moving workers, vehicles, equipment and temporary hazards. Computer vision and analytics can extend observation beyond periodic human site walks, but detecting a possible hazard and preventing an injury are different achievements. Present evidence therefore supports AI as an additional source of safety information rather than as a substitute for site safety management.

Real estate has parallel information problems. Valuers and investors must interpret transaction evidence, property characteristics, local conditions, rental information and other market data. Property managers deal with leases, maintenance records, tenant requests and equipment data. Large occupiers may manage thousands of spaces across portfolios. AI is being used to make these datasets more searchable, standardised and analysable. JLL's 2025 survey found data-related workflows, portfolio optimisation and energy management among the leading areas selected for corporate-real-estate AI pilots.

Energy and sustainability pressures create an additional demand. Buildings generate continuing data about occupancy, temperatures and equipment use, while owners face energy costs and increasingly demanding environmental targets. AI can help identify patterns and predict demand or equipment behaviour. European research and demonstration programmes are consequently combining AI, BIM and digital twins to compare design or renovation alternatives and optimise operational energy and comfort.

3. What can AI currently do?

Present AI can perform a growing range of **specialised built-environment tasks**. It can classify and retrieve documents, summarise contracts and reports, generate draft text, extract structured information, estimate quantities, assist cost and schedule analysis, compare alternatives and identify relationships within large project datasets.

Computer vision can recognise selected objects and conditions in construction imagery and connect visual information to its physical location. More specialised platforms can compare recorded site conditions with project models or programmes and provide progress information. These systems materially improve the ability to process visual records at scale; they do not understand every aspect of a complex site. Academic research continues to identify limitations arising from changing backgrounds, different object sizes, image quality and insufficiently representative construction datasets. A 2025 synthesis found substantial variation in available construction visual datasets in scale, modality, annotation and representativeness.

AI can also **estimate and forecast**. Models can analyse historical schedules, costs or building-performance data and calculate likely outcomes or anomalies. AVMs can estimate property values from available market and property data. Predictive-maintenance models can flag patterns associated with equipment deterioration. Building systems can adjust operations in response to sensor inputs.

Generative AI adds a broad language and reasoning interface. Professionals can ask questions of documents, generate draft reports, reorganise information and create preliminary analyses. This makes general-purpose AI potentially useful to small as well as large organisations because it does not always require construction-specific software development.

What AI cannot presently do reliably is independently manage the complete chain of judgement involved in designing, building, valuing or operating property. A system that identifies apparent construction progress does not thereby determine contractual entitlement to payment. An AVM estimate is not automatically equivalent to a professionally supportable valuation. A generated design still requires technical, regulatory, safety and client assessment. An AI summary of a contract may miss a qualification that materially changes its meaning.

The practical capability is therefore usually **task-level intelligence embedded in a broader human and organisational process**.

4. How established is adoption?

Construction adoption remains **early and uneven despite a growing number of genuine deployments**. The strongest global quantitative evidence identified for this assessment is RICS' 2025 survey. Among more than 2,200 respondents, 45% reported no AI implementation and 34% early pilots. Approximately 12% reported regular use in particular processes; only 1.5% reported use across multiple processes and fewer than 1% organisation-wide embedded use. Around three-quarters described their organisations as having either no AI capability or only limited exploratory capability.

This does not mean construction AI is merely theoretical. Visual progress monitoring has moved beyond isolated demonstrations at a number of major contractors, and AI-assisted document, scheduling and design functionality is entering mainstream construction software. Singapore provides evidence of public-sector support for practical AI adoption across design, tendering, inspections and reporting, while Japanese construction groups continue to develop robotics and automation. The resulting picture is one of **pockets of operational maturity inside a sector whose overall organisational maturity remains relatively low**.

Real estate is harder to express as a single maturity level. AVMs and some algorithmic property-analysis systems are established. Predictive maintenance and building-energy analytics have real operational use, particularly in sophisticated commercial portfolios. Generative AI and more ambitious portfolio optimisation are spreading very rapidly but remain disproportionately represented by pilots.

JLL's 2025 survey illustrates the distinction. Among large organisations in 15–16 markets, 88% of surveyed investors and 92% of occupiers had started or intended to start AI pilots. Yet only about 5% reported having achieved most programme goals, and more than 60% of surveyed real-estate investors remained strategically, organisationally or technically unprepared for scaled implementation. These data concern relatively large organisations and should not be read as adoption rates for the entire global property sector, particularly small agencies, landlords or firms in less digitised markets.

Geography matters substantially. Singapore's government is actively promoting AI, integrated digital delivery, robotics and other built-environment technologies. European programmes are advancing AI-enabled digital twins and renovation tools. Major Japanese contractors have extensive construction-technology programmes. New Zealand evidence indicates meaningful commercial-property experimentation and implementation. These examples show that activity is genuinely international, but available evidence is considerably stronger in technologically advanced markets than in much of Africa, Latin America and lower-income parts of Asia.

5. What appears to be driving adoption?

The first driver is the sector's continuing search for **greater productivity and predictability**. Delays, rework, coordination failures and incomplete information can have large financial consequences on major projects. Technologies that provide earlier visibility of divergence from programme or automate recurring information work therefore have a relatively direct economic rationale.

A second driver is the **digitisation that preceded AI**. BIM, cloud project platforms, digital drawings, connected equipment, 360-degree imagery, drones, building sensors and property databases have created information on which AI can operate. AI adoption is consequently easier where firms already have structured data and mature digital processes. JLL's real-estate research similarly found stronger AI outcomes among organisations that already had successful technology programmes.

A third driver is **workforce pressure**. Construction must deliver complex projects while experienced personnel are scarce in many markets. AI that reduces documentation, assists technical retrieval or gives managers visibility over multiple sites has obvious attraction.

Singapore's current construction policy explicitly connects robotics, automation and other technology with improving manpower productivity amid high construction demand.

Generative AI has also lowered the entry barrier. Earlier construction AI often required specialised systems, datasets and computer-vision expertise. General-purpose language and multimodal models can now be used for document interrogation, drafting and knowledge retrieval through familiar interfaces. This has broadened experimentation well beyond specialist digital teams.

Finally, **energy, carbon and asset-performance pressures** are increasing demand within real estate. AI-supported building controls, energy analytics and digital twins can connect operational and environmental objectives, making building operations one of the areas where AI potentially has both financial and sustainability value.

6. What are the present limitations?

The most pervasive limitation is **data**. Construction information is fragmented across organisations and projects; terminology and coding differ; historical records may be incomplete; and site conditions are inherently variable. Computer-vision datasets may not adequately represent different project types, climates, work practices or equipment. Real-estate datasets similarly vary in coverage, accuracy and timeliness. Poor inputs can limit an otherwise sophisticated model.

Integration is another major constraint. AI insight is useful only if it fits actual design, procurement, site, valuation or property-management workflows. RICS' construction survey identified integration with existing systems among the leading barriers to adoption, alongside skills shortages, data quality, implementation cost and uncertain returns.

Reliability and explainability are particularly important where outputs affect professional or contractual decisions. Generative AI may produce incorrect material with convincing language. Predictive systems can create false positives or miss unusual conditions. Valuation models may perform poorly where transactions are sparse or properties are atypical. Models trained in one market or project environment may not transfer cleanly to another.

These constraints become more consequential because construction and real estate involve large financial values, physical safety and legal responsibility. RICS' first global professional standard on responsible AI use in surveying came into effect on 9 March 2026. It requires governance, attention to data and system risks, professional judgement, reliability assessment and appropriate transparency where AI materially affects surveying services. The existence of such a standard itself reflects how AI has progressed from a purely experimental issue into a professional-practice issue.

Privacy, confidentiality and surveillance also matter. Construction computer vision can record workers; building systems can generate occupancy information; property firms hold financial, identity and contractual data. AI therefore introduces questions about what information is collected, who can access it, how long it is retained and whether third-party AI systems may process confidential information.

Physical automation has additional limitations. Construction sites are less controlled than factories. Layouts change, people and machines interact, weather intrudes and one project may differ substantially from the next. This helps explain why sophisticated construction robotics and autonomous machinery continue to include substantial trial, verification and supervised-operation activity.

Finally, the economic evidence remains incomplete. Vendors and customers report time savings, avoided delays and improved performance, but comparable independent studies across many projects are less common. Technical capability, successful customer deployments and industry-wide productivity improvement therefore remain different evidential propositions.

7. What you should be aware of

“AI in construction and real estate” covers applications with profoundly different levels of risk and maturity. Automatically locating construction photographs, producing a draft meeting report, estimating a property's value, adjusting a building's ventilation and controlling construction machinery are not equivalent uses of AI.

Headline adoption figures require particular care. A survey reporting that an organisation “uses AI” may include everything from staff using a general-purpose assistant to AI integrated into core property operations. Equally, the very high pilot rates reported among large global real-estate organisations do not mean that most real-estate decisions are now AI-driven. Construction evidence provides the opposite warning: impressive deployments at leading contractors coexist with low overall adoption in the wider industry.

Professional judgement remains substantial. RICS' 2026 standard explicitly preserves responsibility for assessing AI outputs where they materially affect surveying services. Similar human involvement remains evident in design approval, construction supervision, safety, contractual decisions and property investment.

Geographical differences should also be kept in view. AI adoption depends partly on BIM penetration, digital records, connectivity, property-data availability, capital, skills and the size and sophistication of organisations. Evidence from Singapore, major European programmes or large international property companies does not establish equivalent conditions everywhere.

Nor should all “smart building”, BIM, digital-twin, robotics or automation activity automatically be described as AI. These technologies can use AI but do not inherently do so. Conversely, AI can be embedded inside ordinary construction and property software without users necessarily recognising it as a separate technology.

The overall Status at 30 August 2026 is therefore one of **material but uneven transition**. AI has established useful positions in property valuation, building analytics, visual construction monitoring and increasingly in project and professional information work. Generative AI has greatly widened experimentation, while computer vision, predictive analytics and digitally connected systems are enabling more specialised operational applications. Yet construction remains early in organisation-wide adoption, much real-estate activity remains at pilot stage, and autonomous built-environment decision-making or construction is far from routine. The strongest current evidence supports AI as an increasingly capable analytical and workflow

layer within a still predominantly human-led sector.

Evidence Notes

1. **Royal Institution of Chartered Surveyors (RICS) — Artificial Intelligence in Construction Report, 2025.** Global survey of more than 2,200 professionals and the principal quantitative evidence used here for construction adoption, preparedness, perceived applications and barriers. The sample was geographically broad but disproportionately UK-based (48% of responses), and results are self-reported rather than independently measured adoption.
2. **RICS — Construction Productivity Report 2026.** Provides contemporary professional context for productivity pressures and applications such as scheduling, cost estimation, quality monitoring and resource allocation. Its discussion of AI is principally sector analysis rather than measured evidence that particular applications are widely deployed.
3. **Singapore Building and Construction Authority — Artificial Intelligence for the Built Environment, updated 5 August 2026.** Primary public-sector evidence of current recognised use cases spanning knowledge management, BIM, tendering, estimating, inspection, progress analytics and reporting. It establishes active policy and application areas in Singapore but does not quantify adoption across firms.
4. **Singapore BCA — Pan-United AiR Digital case study, June 2025.** Evidence of a defined operational AI application involving scheduling, materials replenishment and vehicle movements during a major concrete pour. It is a case example rather than evidence of industry-wide use.
5. **OpenSpace and Buildots deployment material, 2025–26.** Provides evidence that AI-assisted visual capture and progress analysis are being used across genuine construction projects and by major contractors. Provider and customer case-study claims establish deployment but should be treated cautiously when assessing performance or general prevalence.
6. **Construction Dive — AI and construction safety, May 2026.** Reputable trade-journalism evidence of safety-related AI use by major contractors including Skanska, Turner and Balfour Beatty. Useful for demonstrating current application, but not a quantitative measure of sector-wide safety adoption or effectiveness.
7. **Obayashi Corporation — ORCISM crane-management and automation programme, May 2025.** Primary evidence of Japanese development in digitally integrated and increasingly automated construction equipment. Obayashi described the activity as verification and technology development, so it is treated here as developmental rather than established autonomous construction.
8. **JLL — Global Real Estate Technology Survey 2025.** Survey evidence involving more than 1,500 senior investor and occupier decision-makers across multiple markets. Particularly valuable for distinguishing extraordinarily widespread AI piloting from low maturity and limited achievement of programme objectives. It primarily represents larger corporate and institutional organisations and should not be generalised to the full real-estate sector.
9. **UK Valuation Office Agency — Automated Valuation Model Algorithmic Transparency Record, August 2025.** Primary governmental evidence of an AVM in production supporting valuation work concerning approximately 1.5 million domestic properties in Wales. This demonstrates operational automated valuation at scale in one defined setting, not the global prevalence or accuracy of AVMs generally.
10. **RICS — Valuation standards and AI/AVM programme, current 2026.** Evidence that AVMs are an established professional issue and that AI is increasingly being incorporated into valuation workflows. RICS

was still developing dedicated global AI-in-real-estate-valuation guidance during 2026, illustrating continuing evolution in practice and governance.

11. **European Commission BUILD UP / Horizon Europe built-environment material, 2025–26.** Evidence concerning AI, BIM and digital twins for energy forecasting, renovation, operational optimisation and design comparison in Europe. Some initiatives remain research, development or demonstration programmes and are therefore not treated as evidence of universal commercial deployment.
 12. **Xiong et al. — OpenConstruction, 2025.** Academic synthesis of 51 publicly available construction visual datasets, used principally to establish limitations in dataset scale, modality, annotation and representativeness affecting computer-vision applications. It assesses the research-data landscape rather than commercial adoption.
 13. **RICS — Responsible use of artificial intelligence in surveying practice, first edition, published 2025 and effective 9 March 2026.** Global professional standard providing evidence of current governance requirements concerning AI systems that materially affect surveying services, including system governance, data management, risk, professional judgement, output reliability and transparency.
 14. **JLL New Zealand — AI at the property frontier in New Zealand, October 2025.** Supplementary Oceania evidence indicating real-estate AI implementation and research activity, including building operations and predictive maintenance. The underlying adoption figures derive from industry survey evidence and should not be treated as representative of all New Zealand property organisations or of Oceania generally.
-

9. Energy & Utilities

Status as at 30 August 2026

Artificial intelligence is already being used across significant parts of the energy and utilities sector, but the present position is uneven. Its most established applications are generally relatively bounded analytical tasks: forecasting electricity demand and renewable generation, detecting anomalies and equipment faults, analysing inspection imagery, predicting maintenance needs, optimising selected generating or industrial processes, and assisting operators to interpret large volumes of network data. Generative AI is also entering engineering, reporting, customer-service and knowledge-management workflows.

AI is becoming particularly relevant because electricity systems themselves are becoming more complex. Growing quantities of variable wind and solar generation, batteries, electric vehicles, distributed generation and increasingly large electricity-consuming data centres create a system in which supply and demand must be balanced across more assets, locations and timescales. The International Energy Agency (IEA) describes AI as already being used to optimise energy systems, improve production, reduce costs, raise efficiency and uptime, integrate renewable electricity and support maintenance, while cautioning that sector-wide benefits remain difficult to quantify beyond individual deployments.

This does not mean that power grids are becoming autonomously operated by AI. Most operational applications remain embedded within conventional engineering, control and safety arrangements. Deployment differs substantially between utilities and countries, reflecting variations in digital infrastructure, smart-meter coverage, sensor availability, data quality, investment capacity, regulation, technical skills and the condition of existing networks.

1. Where is AI currently being applied?

One of the most important areas is **forecasting**. Electricity system operators, generators and utilities use machine-learning techniques to predict electricity demand, wind and solar output, weather-dependent system conditions and increasingly localised network loads. These forecasts help operators schedule generation, storage and reserves and anticipate where networks may become constrained. In Britain, for example, an AI-based solar forecasting system developed by Open Climate Fix was embedded in National Energy System Operator control-room operations in 2025. ECMWF also made its Artificial Intelligence Forecasting System operational in 2025; among its potential energy applications are improved forecasts of wind speeds and solar radiation.

A second major application is **asset inspection and predictive maintenance**. Machine-learning models analyse sensor readings, operating histories, thermal images, photographs, acoustic signals and other condition-monitoring data to identify equipment that may be deteriorating or behaving abnormally. Computer vision can inspect transmission lines, poles, substations, wind turbines and other infrastructure from drone or vehicle imagery, directing engineers towards suspected defects rather than requiring every image or asset to receive identical manual attention.

New Zealand provides a useful example of real deployment. Northpower, Orion, Unison Networks and WEL Networks announced in November 2025 that they were deploying Tapestry's GridAware platform, combining drones and AI/computer vision for grid inspection across electricity distribution businesses serving more than 510,000 customers. Separately, New Zealand's Commerce Commission records regulated innovation projects involving AI/ML demand forecasting and AI-supported grid inspection.

AI is also being applied to **network monitoring, fault detection and outage management**. Models can examine streams of operational data to recognise abnormal patterns, help locate faults, estimate equipment-health risks and prioritise maintenance or field response. Brazil's Eletrobras moved during 2025 from an initial deployment across ten substations towards wider use of AI-based grid intelligence across its transmission assets for real-time fault monitoring and resolution. This establishes meaningful operational deployment, although claims about performance originate substantially from the participating organisations and technology provider.

Within **generation and industrial energy operations**, AI is being used for plant optimisation, equipment monitoring, process control and predictive maintenance. Dubai Electricity and Water Authority (DEWA), for example, reports AI-supported controllers and digital twins within power generation, alongside AI applications in transmission, distribution and water operations. Saudi Aramco has similarly deployed machine learning and predictive modelling in major oil and gas facilities and permanently adopted a reinforcement-learning system at its Fadhili Gas Plant following a pilot of acid-gas-removal operations. These examples demonstrate operational capability but should not automatically be generalised across generating plants or the wider global energy industry.

Another developing area is **grid planning and management**. AI can analyse scenarios involving demand growth, distributed generation, storage, electric-vehicle charging and network constraints. India's Ministry of Power reported in February 2025 that AI-related work supported under its distribution-sector programme included distribution-transformer health management, IoT monitoring of high-voltage switchgear and drone-based asset management, alongside smart-meter and advanced-metering infrastructure intended to improve the data foundation for advanced analytics.

Customer and organisational operations form a further category. Utilities are using machine learning and generative AI for billing analytics, customer interaction, document preparation, technical knowledge retrieval, regulatory or engineering searches, operational reporting and employee assistance. DEWA, for example, has deployed AI-based tools across engineering support, customer operations and internal functions. Such applications may resemble AI use in other industries, but within utilities they increasingly sit alongside sector-specific operational systems.

2. What problems or needs is AI being used to address?

The fundamental operational problem is **complexity combined with the need for continual balancing and reliability**. Electricity has to be generated and delivered at the time it is required. Wind and solar generation vary with weather; rooftop solar changes the amount of demand visible to the system; batteries can act as either load or supply; and electric vehicles

create new and sometimes concentrated demand. AI is attractive because it can detect relationships across large, fast-changing datasets that are difficult to process manually.

Forecasting therefore addresses uncertainty. Better estimates of demand, solar production and wind output allow operators to plan generation and reserves more accurately and potentially reduce unnecessary balancing interventions. The growth of variable renewable energy makes this increasingly important. The IEA identifies forecasting and system optimisation as central areas in which AI could assist electricity networks as they become more decentralised, electrified and connected.

Asset-management applications address a different problem: **large, ageing and geographically dispersed infrastructure**. Utilities may operate thousands of kilometres of lines and very large numbers of transformers, poles, switches and other components. Inspecting every asset manually is expensive and sometimes hazardous. AI-assisted image and sensor analysis allows attention to be directed towards assets displaying unusual characteristics or higher predicted risk.

Utilities are also trying to reduce **unplanned outages and maintenance costs**. A model that recognises deterioration before equipment fails can support condition-based maintenance rather than relying solely on fixed schedules or responding after failure. The practical value, however, depends on whether sufficient historical failure and condition data exist and whether the model's warnings are accurate enough to change maintenance decisions.

A further need is **better visibility of increasingly distributed electricity systems**. Network operators may not know precisely how much rooftop solar, flexible demand or future electric-vehicle charging will appear at every location. UK Power Networks, for example, ran an AI project from April 2025 to July 2026 aimed at improving visibility and forecasting of solar generation so that planning decisions would not rely unnecessarily on network reinforcement or flexibility procurement. This remained an innovation project rather than evidence that the approach had become universal utility practice.

AI is also being applied because utilities face **productivity and workforce pressures**. Engineers and operators increasingly have more telemetry, inspection imagery, maintenance records, technical documents and regulatory material to interpret. Generative AI can summarise or retrieve information and draft reports, allowing skilled personnel to spend less time on information processing. This is one reason generative AI may spread faster in supporting workflows than in safety-critical autonomous control.

Finally, AI itself is creating a new energy-system problem. Data-centre electricity consumption rose sharply during 2025, with AI-focused facilities growing particularly rapidly. The IEA reported in April 2026 that data-centre electricity demand had increased by 17% in 2025 and warned that expanding AI infrastructure was encountering power-system, equipment-supply and grid-connection bottlenecks. AI is therefore simultaneously a technology being used by the energy industry and an important new source of electricity demand that utilities must accommodate.

3. What can AI currently do?

Current AI systems can **make increasingly sophisticated predictions from historical and real-time data**. They can forecast demand and renewable generation, identify statistical anomalies in network measurements, estimate probabilities of equipment failure and rank assets according to predicted maintenance risk.

Computer-vision systems can examine large quantities of inspection imagery and **identify objects or visual characteristics associated with defects, vegetation interference, corrosion or other abnormalities**. Combined with drones, this can substantially change the inspection process: AI does not repair the infrastructure, but it can reduce the quantity of imagery engineers need to examine manually and help determine where physical inspection or maintenance is warranted.

AI can also **optimise a defined process within specified operating limits**. In generating plants and industrial energy facilities, models can learn relationships among temperature, pressure, fuel use, equipment settings and output, then recommend or implement adjustments intended to improve efficiency. Aramco reported that its Fadhili reinforcement-learning pilot reduced consumption of steam and amine and lowered power usage before the application was permanently adopted at the plant. The result is meaningful evidence for that particular process, not proof that equivalent gains will occur in other facilities.

In electricity networks, AI can help operators **interpret a changing system state and prioritise response**. Fault-monitoring applications can combine numerous operational signals, associate an abnormal event with relevant assets and provide operators with contextual information more rapidly than conventional manual investigation.

Generative AI adds a different capability. It can **summarise technical material, retrieve organisational knowledge, draft reports and requirements documents, answer natural-language questions and assist customer interactions**. These are already practical utility applications. They are less evidence that AI possesses independent engineering judgement than that general-purpose language technology can make large bodies of organisational information easier to access.

What AI generally cannot yet demonstrate is reliable, independent management of the full electricity system across unexpected conditions. A forecasting model does not operate a grid; a defect-recognition model does not determine every maintenance priority; and a generative system capable of explaining operating procedures does not thereby acquire authority to control critical infrastructure. Ofgem's 2026 guidance specifically addresses explainability in grid management and warns about reliance on opaque systems, illustrating the continuing importance of understanding and governing AI-supported decisions.

4. How established is adoption?

Adoption is best described as **established in selected analytical applications, growing across asset and operational management, and emerging or experimental in more autonomous and generative uses**.

Forecasting is among the more mature categories because utilities have used statistical and machine-learning techniques for many years. What is changing is the sophistication, resolution and range of inputs. Operational adoption of AI-based solar forecasting by Britain's electricity system operator demonstrates movement beyond experimental forecasting into routine control-room use.

Asset analytics and predictive maintenance are also established enough to have multiple real deployments, although maturity varies greatly between organisations and asset classes. The IEA's global assessment identifies predictive maintenance, leak detection, remote operations and related optimisation as existing energy applications, while also emphasising that overall benefits are difficult to quantify beyond particular cases.

Grid-monitoring and inspection applications are **growing rather than universal**. Deployments in Brazil and New Zealand show that utilities are moving from isolated proofs of concept towards larger operational implementations, but they do not establish global sector penetration.

Government programmes provide further evidence of movement into implementation. India's distribution-modernisation programme explicitly supports AI and machine-learning applications, while Singapore's SP Group is expanding sensors, digital twins, AI-assisted inspections and predictive grid management. Again, some stated outcomes remain targets extending to 2030 rather than completed deployment as at the Status date.

Generative AI is spreading through internal and customer-facing utility functions but remains comparatively new. Engineering assistants, reporting systems and customer-service agents are already being introduced, while autonomous “agentic” systems controlling operational infrastructure should be treated much more cautiously. A utility announcing an “AI-native” strategy, for example, indicates organisational ambition and development direction; it does not establish that all of its operations have already become AI-controlled.

Europe's position illustrates this mixed maturity. In June 2026 the European Commission launched a Strategic Roadmap for Digitalisation and AI in Energy and an AI.grids initiative involving grid operators and research organisations. This demonstrates substantial institutional commitment, but AI.grids itself was newly launched and therefore represents development and intended deployment rather than an established pan-European operating system.

Global adoption is consequently highly uneven. Highly instrumented utilities with smart meters, modern supervisory systems, extensive sensors and strong engineering-data capabilities can implement applications that are difficult to reproduce in networks with poor connectivity, limited digital records or weak financial capacity. Across parts of Africa, where more than 600 million people still lacked electricity access in 2025, basic infrastructure expansion and utility financial viability remain major priorities.

5. What appears to be driving adoption?

One major driver is the **rapid growth in system complexity**. Renewable generation, batteries, electric vehicles, distributed energy and increasingly flexible consumers create more variables for operators to predict and coordinate. AI's ability to process large multidimensional datasets

therefore becomes more useful as electricity systems become harder to manage using simple historical patterns.

A second driver is **reliability and resilience**. Power-system failures have substantial economic and social consequences. Utilities consequently have strong incentives to detect incipient equipment problems, improve vegetation and line inspection, locate faults quickly and understand network conditions before failures propagate.

Asset cost and ageing infrastructure also matter. Transmission and distribution equipment is expensive and often has long replacement cycles. Better information about asset condition can help utilities decide when maintenance or replacement is actually required rather than treating all equipment identically.

Another driver is the expanding **digital data foundation**. Smart meters, connected substations, drones, satellite observations, weather data, sensors and advanced distribution-management systems create information that did not previously exist at comparable resolution. India's Ministry of Power explicitly linked smart-meter and advanced-metering infrastructure to the reliable datasets required for AI models.

The energy transition provides a further stimulus. AI may help operators accommodate greater amounts of variable renewable electricity without simply adding infrastructure for every possible constraint. The IEA has modelled potentially substantial additional use of existing transmission capacity under widespread AI adoption, although such modelling represents potential rather than demonstrated present-day global gains.

There is also increasing **government and regulatory encouragement**. The European Commission's 2026 roadmap seeks to accelerate AI and digitalisation in energy, while Great Britain's Ofgem has established an AI regulatory laboratory and decided to proceed with an AI technical-sandbox pilot. These initiatives indicate that regulators now regard sector AI adoption as sufficiently material to require dedicated governance and experimentation frameworks.

Finally, the rapid expansion of AI computing is itself intensifying incentives to improve energy-system planning. Data centres are becoming exceptionally large, concentrated loads whose grid connections may require new generation and network investment. The relationship is therefore reciprocal: the energy system enables AI, while expanding AI demand increases the need for improved forecasting, planning and optimisation within the energy system.

6. What are the present limitations?

The most fundamental limitation is **data**. AI models depend on sufficiently complete, accurate and relevant records. Many utilities operate equipment installed over several decades, with inconsistent asset registers, incomplete failure histories and multiple information systems. Sensor coverage can also be uneven. A technically capable model cannot compensate fully for unreliable knowledge of the physical system it is intended to analyse.

Performance is also **context-dependent**. A fault model trained on one type of transformer, climate, network configuration or maintenance regime may not perform equally well

elsewhere. Rare but consequential events present a particular difficulty because historical examples may be insufficient for reliable model training.

A second limitation is **integration with safety-critical infrastructure**. Electricity and gas systems depend on specialised operational technology designed around reliability, deterministic controls and established engineering procedures. Connecting newer AI systems to operational environments creates questions about latency, cybersecurity, failure modes, fallback arrangements and authority over decisions.

Cybersecurity is especially material. More digital connectivity can increase the potential attack surface, while AI systems themselves can be manipulated or behave unexpectedly. In April 2026 the US Department of Energy and Lawrence Livermore National Laboratory launched the Mjölnir AI Testbed specifically to assess the security and reliability of AI models used in energy operations, planning and grid management. The existence of such a testbed is evidence both of AI's growing operational relevance and of unresolved security concerns.

Explainability and operator trust remain important where outputs could influence consequential decisions. A model may identify a component as high risk without giving an engineer a readily understandable causal explanation. This is less problematic for low-consequence assistance than for actions affecting system security, customer supply or infrastructure investment. Ofgem's updated 2026 ethical-AI guidance therefore includes explicit treatment of explainable AI in grid management, predictions and forecasting.

Generative AI introduces familiar but consequential **reliability problems**. A system may produce plausible but incorrect technical statements, misinterpret organisational documents or omit an important qualification. In an administrative task this may create inconvenience; in operational engineering the same behaviour could be materially more serious. Human verification therefore remains central to most consequential applications.

Economics are another constraint. Sensors, communications networks, cloud or edge computing, integration work, cybersecurity and workforce training all cost money. The benefits of avoiding failures can also be difficult to measure because the counterfactual—what would have happened without the AI system—is often uncertain. The IEA consequently cautions that claims about broader cost, reliability and resilience benefits remain difficult to quantify outside individual case studies.

There is finally a substantial **geographical infrastructure divide**. Sophisticated AI optimisation assumes an underlying level of electrification, connectivity, instrumentation and institutional capacity that does not exist uniformly worldwide. In countries still focused on expanding basic electricity access and strengthening financially weak utilities, AI may have useful specific applications, but it is not a substitute for generation capacity, transmission lines, distribution networks, meters, maintenance personnel or sound utility institutions.

7. What you should be aware of

“AI in energy” encompasses very different technologies and levels of consequence. A model predicting tomorrow's solar production, computer vision identifying damaged insulators, an algorithm adjusting a defined plant process, a generative assistant drafting an engineering

document and an autonomous system controlling network equipment should not be treated as equivalent applications.

Readers should also distinguish **optimisation potential from realised performance**. Studies and modelling can estimate large theoretical gains from AI-assisted grid operation, renewable integration or maintenance. Those findings show technical and economic possibilities; they do not establish that the gains have already been achieved across the world's electricity systems.

Similarly, individual deployments by prominent utilities demonstrate that real applications exist but not that equivalent use is widespread. Eletrobras, DEWA, British system operators or New Zealand distribution businesses can provide valuable evidence of present capability and deployment while remaining particular organisational examples.

Human expertise continues to play a central role. Electricity networks and other critical utilities are engineered systems governed by operating rules, physical limits and safety obligations. AI normally provides forecasts, classifications, recommendations or defined control functions within these systems rather than assuming unrestricted responsibility for their operation.

The regulatory position is also evolving. Great Britain has developed sector-specific ethical-AI guidance and regulatory experimentation, while the European Union has incorporated AI into a broader energy-digitalisation roadmap. Different jurisdictions nevertheless apply different combinations of energy regulation, cybersecurity requirements, data law and general AI regulation. There is no single global utility-AI governance regime.

AI's environmental relationship with the sector is unusually two-sided. AI may help improve efficiency, forecasting and renewable integration, but the computing infrastructure required for AI consumes substantial electricity. The IEA's 2026 assessment found rapidly growing data-centre demand and significant pressures on grids, transformers, generation equipment and connection processes. Assessments of AI's contribution to energy systems therefore need to recognise both sides of this relationship.

Finally, evidence quality is uneven. Utilities and technology providers understandably publicise successful deployments, while unsuccessful pilots, marginal improvements and abandoned projects receive less attention. Provider case studies establish that particular capabilities and implementations exist; independent, comparable evidence of long-term reliability, economic benefit and sector-wide penetration remains much less comprehensive.

The overall Status as at **30 August 2026** is therefore one of meaningful but uneven operational adoption. AI has become a practical tool in forecasting, asset analysis, inspection, maintenance, selected process optimisation and growing numbers of utility-support workflows. More ambitious applications in dynamic grid optimisation and autonomous control are developing, but their maturity should not be inferred from demonstrations, modelling studies or strategic announcements. The strongest present role of AI remains as an analytical and optimisation layer within human-engineered and human-accountable energy systems.

Evidence Notes

1. **International Energy Agency — *Energy and AI*, 10 April 2025, and *Key Questions on Energy and AI*, 16 April 2026.** Principal global evidence for the relationship between AI and energy, including present applications, data-centre electricity demand, system optimisation and limitations in quantifying sector-wide benefits. The IEA also models possible future gains; those scenarios have been treated as potential rather than current deployment.
2. **IEA — “AI for energy optimisation and innovation”, 2025.** Provides the principal cross-sector application map covering electricity systems, predictive maintenance, remote operations, leak detection, production optimisation and renewable integration. Particularly useful for distinguishing observed applications from modelled widespread-adoption outcomes.
3. **European Commission — *Strategic Roadmap for Digitalisation and Artificial Intelligence in Energy*, 3 June 2026.** Primary evidence of European policy development, including intended acceleration of AI-enabled grid technologies, energy-data sharing, cybersecurity, skills and AI.grids. The roadmap demonstrates institutional commitment rather than completed EU-wide deployment.
4. **European Commission — *AI.grids launch*, 4 June 2026.** Establishes the launch of a 48-partner initiative involving grid operators and research institutions to develop AI models for grid planning and management. Because the project had only just begun at the reference date, it is treated as developmental activity.
5. **Ofgem — *Ethical AI use in the energy sector*, May 2025, updated May 2026; associated *AI regulatory-laboratory and sandbox material*.** Evidence that AI has become sufficiently relevant to British energy regulation to require sector-specific treatment of accountability, explainability, forecasting, grid management, privacy and black-box systems.
6. **US Department of Energy/Lawrence Livermore National Laboratory — *Mjölnir AI Testbed*, 14 April 2026.** Primary evidence concerning cybersecurity and reliability testing for AI used in energy operations, planning and grid management. The testbed demonstrates active development and risk-management activity rather than widespread deployment of the models being evaluated.
7. **Government of India, Ministry of Power — *Lok Sabha response on implementation of AI technology in power distribution*, 13 February 2025.** Primary evidence of government-supported AI/ML work involving smart-meter data, distribution-transformer health, switchgear monitoring and drone-based asset management. It identifies programmes and work being undertaken but does not quantify national operational penetration.
8. **Northpower/Orion and New Zealand Commerce Commission, 2025–26.** Evidence of AI/computer-vision grid inspection involving four New Zealand electricity distributors, together with regulated innovation initiatives in AI/ML demand forecasting and grid analytics. Organisational announcements establish deployment but should not be extrapolated automatically to all New Zealand networks.
9. **National Energy System Operator/Open Climate Fix — *Quartz Solar*, November 2025.** Evidence that an AI solar-generation forecasting system moved into British control-room operations. Reported savings and emissions benefits originate substantially from the developer and should be interpreted accordingly; the important Status evidence is operational deployment.
10. **ECMWF — *operational Artificial Intelligence Forecasting System*, February 2025.** Primary scientific and operational evidence that machine-learning weather prediction has moved into routine forecasting infrastructure, with explicit relevance to wind and solar energy forecasting.
11. **Eletrobras/C3 AI and Reuters, August 2025.** Evidence that Brazil's major transmission utility moved from an initial ten-substation implementation towards broader deployment of AI-supported real-time fault monitoring. Provider claims concerning benefits are treated more cautiously than the fact of deployment itself.
12. **Dubai Electricity and Water Authority, 2025–26.** Primary organisational evidence of AI applications spanning generation, transmission, distribution, water operations, customer service and engineering support.

DEWA provides an unusually extensive utility example, but its “AI-native” roadmap contains both existing systems and strategic intentions and has therefore not been treated as evidence of universal AI operation.

13. **Saudi Aramco — AI and Big Data operational material, current through 2026.** Provides evidence of machine learning, predictive modelling and reinforcement-learning applications in oil and gas operations, including permanent adoption following a process-optimisation pilot at Fadhili. The reported performance improvements are company-reported and specific to the relevant facilities.
 14. **Singapore SP Group — FY2025/26 sustainability reporting.** Evidence of continuing work in predictive grid management, digital twins, sensors and AI-assisted asset inspection. Several stated productivity and deployment outcomes are 2030 targets, so they have not been represented as achievements already realised by August 2026.
 15. **World Bank/African Development Bank Mission 300 material, 2025.** Used principally to contextualise geographical differences. More than 600 million people in Africa remained without electricity access, making clear why sophisticated AI-grid deployment cannot be assumed to represent the priorities or infrastructure conditions of utilities worldwide.
-

10. Agriculture & Food

Status as at 30 August 2026

Artificial intelligence is now being applied across a substantial range of agricultural and food activities, but the extent of real-world adoption differs sharply by application, farm type, production system and geography.

Some uses have moved well beyond research. AI-assisted pest surveillance, satellite and image analysis, localised forecasting, optical food grading and selected precision-agriculture systems are operating in real production environments. India, for example, has deployed an AI-supported national pest-surveillance system used by more than 10,000 extension workers, while an AI-based local monsoon-forecasting initiative sent forecasts to approximately 38.8 million farmers across 13 states. Commercial machine-vision systems are also performing high-speed sorting and grading in food and fresh-produce facilities.

Other applications remain more uneven. AI-enabled field robotics, precision spraying and weeding, livestock monitoring and farm-management decision support are commercially available and deployed in some settings, but they are not normal practice across agriculture globally. Generative-AI agricultural advisers are developing rapidly, particularly because language and voice interfaces may make technical information accessible to farmers, yet many remain developmental, pilot-based or locally bounded.

The clearest current picture is therefore not of an autonomous AI-run food system. It is of AI increasingly performing particular sensing, prediction, classification, advisory and automation tasks within agricultural and food-production systems that remain strongly dependent on farmers, agronomists, veterinarians, food technologists, operators and other people.

1. Where is AI currently being applied?

One of the most important application groups is **crop observation and diagnosis**. Computer-vision systems analyse photographs of plants, pests and diseases; satellite, drone and sensor data are used to assess crop condition, soil moisture and field variation; and machine-learning systems can identify patterns warranting farmer or extension-worker attention. India's National Pest Surveillance System illustrates deployment at substantial scale: by early 2026 it covered more than 400 pest types and dozens of crops and was being used by more than 10,000 extension workers. China is similarly promoting combined satellite, aerial, sensor and AI-based monitoring within its national smart-agriculture programme.

A related field is **precision application and field automation**. AI can help machinery distinguish crop plants from weeds, determine where treatment is required and apply herbicide, mechanical action or other interventions selectively rather than uniformly across a field. Commercial computer-vision and laser-weeding systems demonstrate that this capability is no longer only experimental. Carbon Robotics, for example, reports commercial LaserWeeder ownership among growers in North America, Europe and Australia; this is useful evidence that deployment exists, although as provider evidence it does not establish general adoption.

European agricultural programmes are also testing and deploying AI-assisted robotic weed-control systems.

AI is being applied to **weather, yield and production decision support**. Models combine weather observations and forecasts with satellite, soil, crop and historical data to estimate events or conditions relevant to farming. India conducted a particularly large AI-based local monsoon-onset initiative during the 2025 Kharif season. Forecasts were sent to approximately 38.8 million farmers, and government follow-up surveys in Madhya Pradesh and Bihar reported that 31–52% of surveyed recipients changed aspects of land preparation, sowing or related decisions. This provides unusually strong evidence that an AI-generated output reached farmers and influenced some real decisions, although the survey covered only part of the programme's geography.

Agricultural advice and information access is becoming another significant application. India's Kisan e-Mitra voice-based AI service answers questions concerning agricultural support programmes in 11 regional languages and had processed more than nine million queries by late 2025. Generative-AI agricultural advisers are also being developed for agronomic questions. CGIAR and partners are working on agriculture-specific models for the Global South, while IFPRI and African agricultural-advisory organisations began testing generative-AI approaches in Kenya, Liberia, Nigeria and Uganda during 2026. These latter examples are important developments but remain testing and development rather than evidence of established African-wide use.

Within **livestock farming**, AI is applied to information from cameras, wearable sensors, feeding systems and other monitoring technologies. Current research and deployments cover animal identification and tracking, activity and behaviour classification, reproductive monitoring, detection of lameness or possible illness, feeding management and production analysis. Reviews of precision livestock farming show genuine operational potential but also large differences between intensive, highly digitised production systems and smaller or lower-resource farming environments.

AI is also entering **agricultural research, breeding and production planning**. New Zealand's AgResearch, for example, reported more than 50 AI-focused research projects in 2025, including AI analysis of livestock CT scans to extract traits relevant to feed efficiency and methane emissions. That demonstrates a research capability rather than routine farm deployment, but it illustrates how AI is being used to extract agricultural information from datasets that would otherwise require expensive manual analysis.

Downstream, **food production, processing and packing** provide another substantial application domain. Machine vision and deep learning are used to classify produce, identify visual defects, distinguish foreign material, grade products and automate portions of quality control. Commercial systems operate across apples, blueberries, cherries, citrus, kiwifruit and other foods. TOMRA reported more than 3,500 lanes of its deep-learning LUCAi sorting technology sold globally by 2025 and documents operational use by a New Zealand apple packer. These figures are provider-supplied, but they establish that AI-based grading has moved into commercial processing environments.

AI is additionally being investigated and used for **food safety, process control, predictive maintenance, supply-chain risk analysis and product development**. Academic reviews document applications combining AI with cameras, spectroscopy and other sensors to detect

defects, contamination, adulteration and quality characteristics. Nestlé, for example, reports using AI and real-time data to identify potential food-safety risks across its supply chain, while separately developing generative-AI tools for packaging research. These establish organisational activity, but not how widespread equivalent use is across the global food industry.

2. What problems or needs is AI being used to address?

Agriculture must continually make decisions under uncertainty. Weather changes, pest and disease outbreaks, soil differences, crop development, animal condition and market requirements create large quantities of information that are difficult to observe continuously or interpret consistently. AI is principally being directed towards circumstances where faster sensing, pattern recognition or prediction can improve the timing and precision of decisions.

Resource efficiency is a major need. Water, fertiliser, pesticides, energy, feed and labour are economically significant and can also carry environmental consequences. Precision systems attempt to determine what input is needed, where and when, rather than treating an entire field, herd or production line uniformly. China's smart-agriculture policy explicitly connects digital monitoring, precision water and fertiliser application, intelligent machinery and AI-supported decision systems with improved agricultural productivity and resource efficiency.

Labour availability and workload are also important. Japan's Ministry of Agriculture identifies a declining and ageing agricultural workforce as a central reason for encouraging smart agriculture: the country's core agricultural workforce fell from about 2.4 million in 2000 to approximately 1.16 million in 2023, with people aged 65 or over comprising a large majority. Robotics, autonomous machinery and automated monitoring are therefore being developed partly to perform repetitive or labour-intensive work and to enable fewer people to manage larger or more complex operations.

Climate and production risk provide another strong driver. Farmers must decide when to sow, irrigate, protect crops or move livestock despite uncertainty about rainfall, temperature, disease and other conditions. AI-based forecasting and predictive analytics seek to turn large meteorological and agricultural datasets into decisions relevant at farm level. India's monsoon-onset programme is a particularly clear example because its forecast was designed around the specific decision of when to begin Kharif agricultural operations.

Access to expertise is especially important for smallholders and locations where extension services are thin. Voice and conversational AI may make agricultural information available without requiring a farmer to search technical documents or communicate in a dominant international language. Work in India and Africa consequently places substantial emphasis on local language, speech recognition, locally grounded agronomic knowledge and connections to human advisers.

In livestock systems, the underlying problem is partly one of **continuous observation**. A person cannot watch every animal continuously. Sensors and computer vision can detect changes in movement, feeding or behaviour that may indicate reproductive state, illness, lameness or welfare problems and direct human attention towards particular animals.

Food processors face different pressures: **speed, consistency, quality, food safety and waste reduction**. Production lines may process enormous numbers of visually similar items. AI-assisted inspection can classify products repeatedly at line speed, allowing staff to concentrate on exceptions and process control. Research also shows substantial interest in combining AI with non-destructive sensing for rapid quality, contamination and authenticity assessment.

3. What can AI currently do?

Current agricultural AI can perform some narrowly defined tasks very effectively.

Computer vision can recognise specified weeds, pests, crop symptoms, fruit characteristics or animal features in images. When connected to machinery, the recognition can trigger a physical response—such as directing a treatment towards an identified weed. Commercial systems demonstrate that AI-guided recognition and action can operate at normal field or packing-line speeds, although performance inevitably depends on crop, environment, training data, machinery and operating conditions.

AI can also combine large datasets that would be cumbersome for an individual farmer to interpret. Satellite imagery, weather records, soil information, crop observations and historical patterns can be integrated to classify land, monitor crop condition or calculate probabilities relevant to production. India is using such approaches for crop monitoring, pest surveillance and agricultural decision systems.

In livestock systems, machine-learning models can identify and track animals and classify aspects of posture, movement and behaviour. They can therefore help detect patterns associated with heat, calving, feeding, lameness or possible health problems. These are primarily monitoring and early-warning capabilities: identifying a pattern is not equivalent to making a complete veterinary diagnosis.

Generative AI can retrieve and reorganise agricultural information and provide conversational responses in natural language. Voice interfaces extend this capability to farmers who may prefer spoken communication. Yet agricultural advice is unusually context-dependent: the correct response can vary with location, crop variety, soil, weather, available products, regulation and farming practice. Research programmes consequently emphasise grounding agricultural models in validated specialist knowledge and local conditions rather than relying solely on general-purpose models.

In food processing, computer vision and deep learning can inspect large numbers of products, classify visible defects and sort items according to predefined grades. AI combined with spectral and other sensing technologies can also distinguish chemical or physical patterns associated with freshness, adulteration, contamination or quality. The technical research base is substantial, but laboratory or controlled-testing performance should not automatically be interpreted as routine industrial deployment.

What AI generally cannot do reliably is manage the full complexity of a farm or food-production system without human involvement. A system that identifies weeds cannot necessarily diagnose crop disease; a yield model does not determine every management decision; and an agricultural chatbot does not possess the full contextual knowledge of an

experienced local agronomist.

4. How established is adoption?

Adoption is best described as **layered and highly uneven**.

Some specialised applications have reached meaningful operational scale. India's pest-surveillance, farmer-query and weather-advisory systems demonstrate public-sector deployments reaching thousands of extension staff and, in some cases, millions of farmers. Commercial optical sorting and grading are also established in parts of the fruit, vegetable, nut and processed-food industries.

Precision farming and field automation occupy an intermediate position. AI-powered machinery, robotic weeding, vision-guided spraying and autonomous capabilities are commercially real, and some growers are using them routinely. Their use, however, is concentrated in production systems where scale, crop value, labour economics, capital availability and suitable farm layouts justify the investment. A commercial deployment by dozens or hundreds of farms establishes real adoption but does not mean that a technology is standard across the world's farms.

European evidence demonstrates this distinction particularly clearly. The European Commission reported in 2026 that AI already supports precision farming, automation and agricultural decision-making, while simultaneously describing uptake of digital technologies in European agriculture as relatively slow. Reported obstacles include cost-effectiveness, interoperability, lock-in, trust, skills and data availability.

China is pursuing wide-scale smart-agriculture development through national policy covering crop production, protected agriculture, livestock and aquaculture. Its programme explicitly incorporates AI alongside sensors, robotics, remote sensing and intelligent machinery. This shows substantial state-supported development and deployment activity, although "smart agriculture" figures cannot automatically be treated as measures of AI adoption alone.

Japan likewise has an active smart-agriculture programme supported by legislation and public policy, but its continuing emphasis on demonstration, diffusion and technologies adapted to Japanese farm structures indicates an ongoing transition rather than universal adoption.

Across much of Africa, evidence is more strongly characterised by research, development programmes, digital-advisory initiatives and selected commercial applications. CGIAR's AgWise work and the 2026 IFPRI/AFAAS generative-AI trials demonstrate significant activity, but they should not be interpreted as evidence that AI farm systems are widely deployed among African smallholders. Reviews of precision dairy farming similarly identify substantial potential but continuing infrastructure, affordability and skills barriers.

Overall, AI has entered Agriculture & Food materially, but there is no credible basis for describing the sector as uniformly AI-enabled.

5. What appears to be driving adoption?

The strongest driver is the need to produce reliably despite **resource, labour and environmental constraints**. Technologies that reduce unnecessary chemical or water use, automate repetitive work, detect problems sooner or make scarce expertise available across more farms have relatively clear practical purposes.

Economics can be especially compelling for labour-intensive tasks. Commercial AI weeding and sorting systems are marketed around reductions in manual inspection or field labour. Customer and provider case studies indicate that such savings can be material in particular operations, although reported returns should not be treated as independent evidence that similar economics apply to every farm.

A second driver is the expanding **digital information base** surrounding agriculture. Satellites, weather systems, connected machinery, farm-management software, sensors, cameras and digital production records create the data from which AI systems can derive predictions or classifications. China's strategy explicitly seeks an integrated satellite-air-ground monitoring network, while India's Digital Agriculture Mission is building large farmer, crop and geospatial datasets.

Government policy is unusually important in this sector. China, India, Japan and the European Union are all funding or organising digital and AI-related agricultural development. Such activity reflects concerns extending beyond firm-level productivity to food security, rural labour, environmental sustainability and climate resilience.

Generative AI introduces another driver: the prospect of reducing the cost of **personalised agricultural information**. A single digital service can potentially communicate in many languages and respond to individual questions at far greater scale than conventional extension systems. India's Kisan e-Mitra provides real evidence of high-volume use, while current African and CGIAR programmes are testing whether similar approaches can provide sufficiently trustworthy agronomic advice.

6. What are the present limitations?

Agriculture is an unusually difficult environment in which to generalise AI performance. Fields, animals and food materials vary continuously. Lighting, weather, soil, crop variety, disease prevalence, farm practice and equipment differ between locations. A vision model trained successfully on one crop or region may therefore perform less reliably elsewhere.

Data availability and quality remain central constraints. European Commission work identifies fragmented data, insufficient high-quality datasets and weak interoperability as obstacles to agricultural AI. Agricultural data also come from many sources—machinery manufacturers, satellites, governments, sensors, processors and farmers—which can create questions about standards, access and ownership.

Economics and farm structure strongly influence adoption. Expensive cameras, robotics, connectivity and automated machinery may be viable on large farms or high-value horticultural crops while remaining uneconomic for small, fragmented or low-margin farms. Research on

agricultural technology adoption repeatedly identifies high capital requirements, credit constraints and uncertain returns as important barriers.

Connectivity and digital infrastructure are another dividing line. AI can sometimes operate locally on machinery or smartphones, but many systems depend on reliable communications, cloud services, geospatial data or continuing software support. These conditions cannot be assumed across rural regions globally.

Generative AI adds familiar **reliability problems** in an especially context-sensitive domain. Fluent advice can still be incorrect, insufficiently local or based on outdated information. Local languages introduce additional difficulties involving speech recognition, dialect, agricultural vocabulary and scarce training data. IFPRI's Indian work emphasises that vernacular terminology, local expertise, farmer feedback and human-centred design are essential to useful agricultural voice systems.

Physical automation has further limitations. Robots must operate safely in dust, mud, rain, uneven ground, foliage and changing light and around people and animals. Research prototypes can demonstrate excellent laboratory performance without establishing robust year-round field economics. A 2026 AI plant-protection robot study, for example, explicitly described its results as laboratory-based proof of concept requiring future field development.

Food-processing AI faces related questions concerning model generalisation, calibration, traceability and integration with existing quality systems. Reviews continue to identify interpretability, dataset diversity and industrial implementation as significant challenges despite strong technical results.

Finally, measurements of adoption are themselves limited. "Smart agriculture", "digital agriculture", "precision agriculture" and "AI" frequently overlap in surveys and policy material but are not synonymous. A farm using GPS steering or connected sensors is not necessarily using AI, and a national smart-agriculture programme may bundle AI together with IoT, robotics, remote sensing and conventional automation.

7. What you should be aware of

The most important qualification is that **AI in Agriculture & Food is not one technology or one maturity level**. An image system identifying a pest, an autonomous weeder, an AI weather model, a livestock-monitoring algorithm, a farmer chatbot and a food-grading system address substantially different tasks and operate under different evidential conditions.

Human expertise remains important even where AI performs significant work. Farmers still interpret recommendations in relation to their land and circumstances; agronomists and extension workers assess crop problems; veterinarians diagnose and treat animals; and food-production personnel determine acceptable specifications and manage safety systems. AI frequently changes what information those people can observe or how quickly they can act rather than removing them from the process.

Scale also needs careful interpretation. India's delivery of an AI-supported monsoon forecast to nearly 39 million farmers demonstrates exceptional distribution scale; it does not establish

that every recipient used the advice or benefited from it. Conversely, a robotic system operating successfully on a relatively small number of commercial farms may demonstrate strong real-world technical capability without constituting widespread agricultural adoption.

Geography matters unusually strongly. European Commission evidence indicates relatively limited European farm uptake despite extensive technological development, while countries including China, India and Japan are pursuing different forms of state-supported agricultural digitalisation. Smallholder agriculture in Africa and parts of Asia faces materially different economics and infrastructure from large mechanised operations.

Readers should also distinguish between **AI and the wider smart-agriculture system around it**. Sensors collect information; satellites provide imagery; GPS locates equipment; robots perform physical work; and conventional software stores records. AI may interpret or act on those inputs, but attributing the whole outcome to AI can overstate its independent role.

Finally, evidence is uneven. Agricultural technology suppliers have strong incentives to publicise performance and customer results. Such evidence is useful for establishing that a product and deployment exist, but does not by itself demonstrate sector-wide effectiveness. Academic research, meanwhile, contains many successful models and prototypes whose performance has not yet been established under diverse commercial conditions.

Status at 30 August 2026: AI has moved materially into Agriculture & Food, with some substantial real-world deployments. Its strongest current role is in sensing and classification, precision decision support, targeted automation, agricultural forecasting and food sorting or inspection. Large-scale public applications in countries such as India show that AI can reach farmers at population scale, while commercial systems demonstrate practical field and processing use. At the same time, adoption remains highly uneven. Generative agricultural advice, more autonomous farming and many sophisticated livestock and food-safety applications are developing rapidly but vary considerably in validation and deployment. The sector remains fundamentally human-operated, and the usefulness of AI continues to depend heavily on local data, physical conditions, infrastructure, economics and agricultural expertise.

Evidence Notes

1. **Government of India, Ministry of Agriculture & Farmers Welfare — Digital Agriculture and AI releases, 2025–2026.** Primary evidence for Kisan e-Mitra, the National Pest Surveillance System, crop monitoring and AI-assisted agricultural forecasting. The March 2026 release documents the local monsoon-onset pilot reaching approximately 38.8 million farmers across 13 states. Follow-up survey results provide unusual evidence of reported behavioural response, although surveys were conducted only in selected states and should not be generalised to all recipients.
2. **Ministry of Agriculture and Rural Affairs, China — Guidance on Developing Smart Agriculture, 2024/2025 implementation material.** Primary policy evidence for integrated remote sensing, field monitoring, precision application, intelligent machinery, livestock and aquaculture digitalisation and the incorporation of AI models into Chinese agricultural modernisation. Because the strategy combines AI with IoT, robotics and other digital technologies, it does not provide a clean measure of AI-specific adoption.

3. **European Commission — *Unlocking the potential of artificial intelligence for sustainable agriculture*, July 2025; *Apply AI agrifood work*, April 2026; *Apply AI Strategy*.** Important evidence that agricultural AI is technically advancing while European farm adoption remains uneven or limited. Identifies data quality, interoperability, cost-effectiveness, trust, skills, platform lock-in and regulatory uncertainty as material constraints.
4. **Japan Ministry of Agriculture, Forestry and Fisheries — *Smart Agriculture material*, April 2025.** Primary Japanese evidence on the relationship between agricultural workforce decline, ageing and promotion of robotics and AI-enabled smart agriculture. It illustrates a major adoption driver but does not establish universal use among Japanese farms.
5. **FAO — *Digital Agriculture and AI programme and 2025–2026 agrifood AI work*.** Provides global institutional context concerning AI, geospatial analysis, food security, agricultural data and responsible deployment. Particularly useful for showing that international work now treats AI as part of wider agrifood digital infrastructure rather than solely as individual farm technology.
6. **CGIAR/IFPRI/AFAAS — *agricultural advisory initiatives, 2025–2026*.** Evidence concerning AgriLLM, AI-enabled advisory development for the Global South, Indian voice-based agricultural advice and 2026 generative-AI trials with agricultural advisory organisations in Kenya, Liberia, Nigeria and Uganda. Much of this activity remains developmental or trial-based and should not be represented as widespread farmer adoption.
7. **CGIAR/IITA — *AgWise*, June 2025.** Provides evidence of AI- and data-supported agronomic decision tools directed towards African smallholders, including crop, soil and fertiliser decisions. Useful for illustrating activity outside wealthy mechanised farming systems; evidence of broad adoption remains limited.
8. **Frontiers in Animal Science/Veterinary Science and FAO AGRIS — *precision-livestock reviews, 2025–2026*.** Peer-reviewed evidence covering cameras, wearables, machine learning, animal tracking, behaviour analysis, reproductive and health monitoring and the barriers to livestock-AI deployment. Reviews also emphasise data, investment, skills and generalisation constraints, particularly in African and heterogeneous production environments.
9. **AgResearch New Zealand — *Annual Report 2025*.** Primary organisational evidence of more than 50 AI-focused research projects and AI analysis of livestock CT datasets for traits associated with feed efficiency and methane emissions. This principally establishes research capability rather than widespread commercial farm use.
10. **Carbon Robotics — *LaserWeeder deployment information, current through 2026*.** Provider evidence demonstrating commercial AI-enabled robotic/laser weeding on farms in several regions. Customer examples provide evidence of real deployment but performance and economic claims are vendor or customer case-study evidence and should not be generalised to agricultural operations overall.
11. **TOMRA Food — *LUCAi and AI sorting systems, current through 2025–2026*.** Provider evidence of commercial deep-learning fruit and food grading, including reported global sales and a New Zealand deployment. Strong evidence that AI-powered sorting is operational in commercial food environments; provider figures do not independently establish the overall prevalence of AI across food processors.
12. **Food Bioscience, Journal of Food Protection, Trends in Food Science & Technology and Current Research in Food Science — *reviews published 2025–2026*.** Peer-reviewed evidence for machine learning, computer vision, spectroscopy and deep learning in food quality, safety, contamination detection, adulteration assessment and process control. The literature demonstrates substantial technical capability but repeatedly notes data, interpretability, generalisation and industrial-deployment limitations.
13. **Nestlé — *food-safety and R&D material, current through 2025–2026*.** Organisational evidence that a major multinational food producer is using AI and real-time data for supply-chain food-safety risk analysis and developing generative-AI approaches to packaging research. Useful evidence of corporate activity, but it should not be generalised to the food-manufacturing sector as a whole.

14. **Scientific Reports — AI-driven plant-protection robot, 2026.** Peer-reviewed example demonstrating the distinction between impressive technical performance and commercial field deployment: the system achieved defined laboratory performance but was expressly presented as proof-of-concept work requiring future field development.
-

11. Government & Public Services

Status as at 30 August 2026

Artificial intelligence is now being used across a substantial range of government and public-service activities, but the depth of adoption varies considerably. The clearest areas of present use are citizen information and service delivery, internal administrative work, tax and revenue administration, public procurement and audit, document processing, employment services and selected forms of administrative decision support. Generative AI has widened experimentation particularly quickly because public servants can use it for drafting, summarisation, translation, search and knowledge retrieval without necessarily giving the system authority to make a public decision.

The overall Status is nevertheless more cautious than the volume of government AI initiatives may suggest. OECD evidence published in 2026 indicates that 35 of 36 reporting OECD countries used AI in at least one area of government, but use was much more common in internal processes and public services than in policymaking or oversight. The same evidence shows increasing generative-AI use but a continuing concentration in comparatively low-risk assistance tasks.

Adoption also extends well beyond OECD countries. Brazil has AI applications in tax administration, fiscal statistics, procurement oversight and public-service delivery; Singapore has deployed an in-house generative-AI assistant across a substantial part of its public workforce; Latin American public-sector databases record hundreds of systems; and other examples range from citizen-service systems in Chile and Portugal to government AI programmes in Australia and New Zealand. The resulting global picture is not one of an autonomous “AI government”. It is principally one of AI being inserted into particular components of larger administrative systems that remain governed, operated and accountable through public institutions and people.

National defence, military operations, national intelligence and national-security applications are outside the scope of this element.

1. Where is AI currently being applied?

One of the most visible areas is **citizen information and public-service access**. Governments use chatbots, virtual assistants and AI-supported search to answer questions, navigate government websites, direct people towards services and help users complete administrative processes. OECD case evidence includes applications in Portugal, Brazil, Chile and Singapore. Recife in Brazil, for example, has combined its Conecta Recife digital platform, covering hundreds of services, with chatbot-based assistance. Some systems principally retrieve existing information; newer generative systems can interpret natural-language questions and construct more tailored answers.

AI is also used inside **social, employment and other citizen-service programmes**. Public employment services are an important example. OECD research found that around half of public employment services in OECD countries had deployed AI-powered solutions, including

systems for matching jobseekers with vacancies, occupational classification, profiling unemployment risk, providing career guidance and answering client questions. Israel's employment service, for example, has incorporated AI into job matching.

A second major field is **internal government administration**. Machine learning, automation and increasingly generative AI are being applied to document classification, correspondence, meeting summaries, information retrieval, translation, drafting, HR processes, coding, procurement documentation and analysis of large administrative records.

Australia's whole-of-government Microsoft 365 Copilot trial illustrates this development. More than 7,600 public servants across over 60 agencies were given access during the 2024 trial, with prominent uses including summarisation, preparing first drafts and information search. The programme was explicitly an evaluation rather than proof of permanent government-wide adoption.

Singapore represents a more advanced example of routine generative-AI access. Its government-developed Pair assistant supports common public-service tasks and allows agencies and public officers to create specialised assistants. By the third quarter of 2025 it reportedly had 78,511 monthly active users, against a public-service workforce of approximately 154,000. That indicates substantial actual use within one highly digitised administration, although it does not establish equivalent adoption elsewhere.

Tax and revenue administration represents a comparatively mature form of government AI. Tax authorities have long used rules-based analytics, anomaly detection and machine learning to identify unusual transactions, assess compliance risks, prioritise investigations, classify documents and support taxpayer services. Brazil has used supervised machine learning to group similar tax-appeal files so officials can process them more efficiently.

AI is likewise being applied to **public financial management, procurement, audit and integrity activities**. Brazil's National Treasury has used machine learning to classify government expenditure, while its Alice system analyses public procurement documents and identifies potential irregularities for official examination. According to OECD evidence, Alice analysed nearly 191,000 acquisitions in 2023 and triggered 203 audits involving contracts worth several billion euros. These systems illustrate AI functioning primarily as a screening and prioritisation mechanism rather than independently determining wrongdoing.

Other applications include **administrative justice and case management**, regulatory analysis, processing applications and permits, fraud detection, inspection prioritisation, public consultation analysis and policy research. Generative AI is beginning to assist officials in searching legislation, summarising public submissions, drafting documents and synthesising evidence, but policy formulation and evaluation remain considerably less mature than administrative uses. OECD's 2026 Digital Government Outlook found AI supporting policymaking in only 13 of 36 reporting OECD countries, compared with 31 reporting use in internal processes.

2. What problems or needs is AI being used to address?

A fundamental pressure is the **scale and complexity of public administration**. Governments process large quantities of applications, correspondence, legislation, financial records, procurement documents, tax data and service enquiries. Much of this material is repetitive enough for computational assistance but too extensive for detailed manual examination at every stage.

Citizen services face a related problem: people must often navigate fragmented institutions, complex eligibility rules and large government websites. AI search and conversational interfaces are being used to make this information easier to locate and interpret and to reduce pressure on call centres and service staff.

Governments are also under continuing **productivity and workforce pressure**. Administrative processes can consume substantial amounts of professional time. Generative AI is therefore attractive for tasks such as summarising long documents, drafting routine text, extracting information and finding material scattered across organisational repositories.

Australia's Copilot trial provides unusually useful evidence of this practical motivation. Participants most consistently reported benefits in summarisation, drafting and information search. Sixty-nine per cent of post-use respondents said the tool improved the speed of completing tasks, although the study relied substantially on self-reported productivity and involved a non-randomised volunteer population, so the figures should not be interpreted as measured government-wide productivity gains.

Another need is **prioritisation**. Tax authorities, auditors, regulators and service agencies frequently confront more cases than can receive equal manual scrutiny. AI can rank records by apparent risk, identify anomalies, cluster similar cases or direct staff attention towards potentially significant transactions.

There is also demand for **greater personalisation and accessibility**. Employment services can match individuals to vacancies; citizen-service systems can tailor information to particular circumstances; language technologies can translate or simplify government content; and conversational systems can make information accessible without requiring users to know how a government agency is organised.

Finally, AI is being explored because governments need to analyse increasingly large and interconnected bodies of evidence when developing policy. Models can classify submissions, identify themes, search research and summarise information. This is an important emerging capability, but the use of AI in policymaking remains less established than its use in administrative support.

3. What can AI currently do?

Current AI can perform a number of bounded governmental tasks effectively.

Machine-learning systems can **classify, rank and identify patterns** across large administrative datasets. They can group similar cases, detect unusual expenditure or transactions, predict predefined risks and direct cases towards appropriate workflows.

Natural-language systems can **extract information from documents**, classify correspondence, search large text collections and summarise material. Generative AI extends this capability by creating draft emails, reports, consultation summaries, briefing material and other text.

AI can also operate conversational interfaces capable of **answering questions about public services** or guiding users towards appropriate information. Their practical usefulness improves when they are connected to controlled government information rather than relying purely on the general knowledge of a foundation model.

Government staff assistants can now operate across everyday office work. Australia's evaluation found perceived efficiency and quality benefits particularly in summarisation, first drafts and information searches. Singapore's Pair demonstrates that comparable capabilities can be made available through a government-controlled environment at substantial scale.

Some systems can **support decisions without making the final decision**. An AI system might score an application, flag a suspected irregularity or generate a summary for an official. Canada's Directive on Automated Decision-Making is notable because it explicitly treats such partial automation as consequential: systems that generate assessments, scores, classifications, recommendations or summaries that influence administrative decisions can fall within the directive even when an officer retains final authority.

What AI cannot generally do reliably is substitute for the full institutional context in which government decisions occur. Administrative decisions can involve law, policy, discretion, exceptional circumstances, procedural fairness and consequences for individual rights or entitlements. Generative models can additionally produce inaccurate or fabricated information. Technical fluency therefore does not establish suitability for autonomous decision-making.

4. How established is adoption?

Adoption is **substantial but highly uneven**.

At the broadest level, AI has clearly moved beyond isolated experimentation. OECD's 2026 Digital Government Outlook reported some government AI use in 35 of 36 reporting OECD countries. AI was used for internal processes in 31 of 36 countries and public services in 27 of 36. Generative AI was reportedly used for at least one purpose in 64% of OECD countries.

These statistics establish breadth, not depth. A country counts as using AI in an area even if only part of its government does so. They therefore should not be read as showing that most public servants, government transactions or citizen services are AI-enabled.

Some **analytical and rules-based applications are relatively established**, especially in taxation, fraud or anomaly detection, document processing and operational administration. These areas pre-date the recent generative-AI wave.

Generative AI for staff productivity is growing rapidly, but maturity varies. Singapore provides evidence of broad recurring access and use. Australia conducted a large whole-of-government trial but deliberately evaluated the technology before wider decisions; only about one-third of trial participants used Copilot daily. New Zealand's government guidance similarly describes agencies as being at different stages of implementation.

Citizen-facing conversational AI is established in selected services but not uniformly mature. Governments have operational chatbots and digital assistants, yet many applications still perform comparatively narrow information-navigation functions.

AI-supported administrative decisions require greater qualification. Such systems exist, but the sensitivity of benefits, immigration, employment, licensing and other decisions creates greater governance and evidential requirements. Canada's federal framework, for example, requires impact assessment and scaled safeguards for automated systems contributing to administrative decisions.

Policy analysis and evaluation remain less mature. In the 2026 OECD evidence, only 36% of reporting OECD countries used AI to support policymaking and 33% reported uses relating to oversight and accountability. Current government AI is therefore considerably more established as an administrative and service technology than as an autonomous policy technology.

5. What appears to be driving adoption?

The most immediate driver is **administrative productivity**. Governments face continuing demand for public services while operating within workforce, fiscal and processing constraints. Tasks involving documents, information search and repetitive analysis offer relatively straightforward opportunities for assistance.

A second driver is the rapid improvement and accessibility of **generative AI**. Public servants no longer require specialised machine-learning systems for every use case; general-purpose tools can perform summarisation, drafting and search within familiar office environments.

Existing digitisation and government data infrastructure also matter. Agencies with digitised records, integrated identity systems, accessible administrative data and mature digital services are better positioned to use AI. This helps explain why adoption differs significantly between jurisdictions.

There is also a **service-quality motivation**. AI can potentially reduce waiting times, make government information easier to navigate and help agencies identify people or cases needing attention.

Successful pilots reinforce interest. Australia's public-service evaluation, for example, found that 86% of participating respondents wished to continue using Copilot, even though actual usage was moderate and inaccuracies reduced some of the productivity benefits.

Governments are simultaneously creating frameworks that make experimentation institutionally easier. Canada has an AI Strategy for the Federal Public Service and an

established automated-decision regime; New Zealand has a Public Service AI Framework and responsible-AI guidance; Australia has developed government AI policies and assurance arrangements. The expansion of such institutional infrastructure is itself evidence that AI use has moved from occasional experimentation towards a recognised public-administration capability.

6. What are the present limitations?

The first limitation is **reliability**. Generative AI can produce plausible but incorrect information. In government, an erroneous answer about eligibility, legal requirements, taxation or a public service can have consequences beyond an ordinary factual mistake.

Australia's Copilot evaluation found that unpredictability, insufficient contextual knowledge and the need to verify outputs reduced productivity benefits, with some participants reporting that the tool added time to tasks.

A second limitation is **data quality and interoperability**. Government information is frequently divided among legacy systems, recorded using inconsistent standards or governed by access restrictions. An AI system cannot reliably integrate information to which it has poor access or whose underlying quality is inadequate.

Privacy, security and confidentiality are especially important. Public agencies hold tax records, employment data, health information, benefit records and other sensitive material. Public-service use therefore places constraints on which external models can receive data, where information is processed and how outputs and logs are retained.

Bias and fairness become material when AI contributes to decisions affecting people. Historical administrative data can reflect previous inequalities or institutional practices. A technically accurate model can also produce unequal error rates across groups.

Explainability and procedural fairness matter because government decisions may need to be justified and challenged. Canada's automated-decision framework demonstrates how substantially these requirements can affect deployment: relevant systems may require algorithmic impact assessment, explanation, testing, peer review, disclosure and mechanisms for recourse. Canadian privacy oversight has nevertheless reported that implementation of the requirements has not always been consistent across federal institutions.

There are also significant **skills and organisational limitations**. Effective use requires public servants who understand not only prompting or data science but the legal, ethical, operational and administrative implications of AI. Australia's evaluation found substantially greater confidence among participants receiving multiple forms of training.

Cost and integration remain uncertain. Licence fees, computing infrastructure, system integration, security controls, training, monitoring and human review can offset apparent efficiency gains. Australia's whole-of-government evaluation explicitly did not determine whether Copilot's benefits exceeded its costs.

Finally, the **evidence of real-world outcomes is uneven**. Governments can count systems, pilots and users comparatively easily. It is harder to establish whether AI materially improves service quality, fairness, public outcomes or whole-of-government productivity.

7. What you should be aware of

“AI in government” encompasses systems with very different consequences. An assistant summarising an internal meeting, a chatbot explaining a government service, an algorithm flagging a suspicious transaction and a model influencing eligibility for a public benefit should not be treated as equivalent applications.

Human involvement consequently differs. Some administrative automation can operate with limited intervention. Generative drafting normally requires verification. Higher-impact administrative decisions may require substantial human oversight, explanation and avenues for review.

Reported adoption figures also require care. OECD evidence that AI exists in 35 of 36 reporting countries is strong evidence of geographical spread, but not of uniform maturity or intensive use across each government. Individual jurisdictions themselves show variation between agencies.

High-profile demonstrations should likewise be distinguished from routine deployment. Australia's thousands of Copilot users constituted a deliberately controlled trial. Singapore's Pair usage provides stronger evidence of recurring operational adoption within a public service. These are different maturity signals.

The regulatory and governance position varies internationally. Canada places formal requirements around federal automated administrative decisions. New Zealand uses a principles-based public-service framework emphasising transparency, human oversight, safety and accountability. Other jurisdictions use different combinations of legislation, administrative law, procurement rules, data-protection requirements and AI-specific guidance.

Geography matters as well. Much systematic adoption evidence comes from OECD members and highly digitised governments. Nevertheless, databases and examples from Latin America—particularly Brazil—show that significant governmental AI use is not confined to North America, Europe or East Asia. The evidence base for many lower-income administrations remains less comprehensive, making global comparison difficult rather than demonstrating absence of activity.

Finally, government creates a distinctive accountability context. An inaccurate commercial recommendation and an inaccurate government determination can have very different consequences. The legitimacy of public administration depends not simply on whether an AI system is technically capable, but on whether decisions remain lawful, contestable, appropriately transparent and attributable to accountable institutions.

Status at 30 August 2026: AI has moved materially into Government & Public Services, particularly in internal administration, citizen-service support, tax and financial analysis, procurement oversight, information processing and staff productivity. Traditional machine-

learning and analytical applications are comparatively established in several functions, while generative AI is spreading quickly as a drafting, summarisation, search and knowledge-assistance technology. Citizen-facing use is increasingly visible, but high-impact administrative decision-making and AI-supported policymaking remain more constrained and uneven. The present Status is therefore best characterised as growing integration of specialised AI capabilities into human-governed public institutions, rather than autonomous administration by AI.

Evidence Notes

1. **OECD — Digital Government Outlook 2026, “Adopting and governing AI in government”, 2026.** Principal comparative evidence for adoption across OECD governments. Reports AI use in at least one governmental area in 35 of 36 reporting countries, with internal processes and public services substantially ahead of policymaking and accountability. The figures measure whether countries report use, not intensity or government-wide penetration.
2. **OECD — Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions, 2025.** Analysis of 200 governmental AI use cases across 11 functions, supplemented by European and Latin American datasets. Provides the principal cross-country application map and evidence that adoption remains fragmented and uneven. Individual examples demonstrate deployment but should not be generalised into sector-wide maturity.
3. **OECD — “AI in public service design and delivery”, 2025.** Evidence concerning AI-supported public employment services, citizen interfaces, personalisation and government service design. Useful for distinguishing established service applications from continuing experimentation.
4. **Australian Digital Transformation Agency — Evaluation of the whole-of-government Microsoft 365 Copilot trial, 2024.** Large real-world government evaluation covering more than 60 agencies and thousands of public servants. Provides unusually detailed evidence concerning use, perceived productivity, training, inaccuracies and adoption barriers. The trial was voluntary and non-randomised, and productivity results were substantially self-reported; it therefore does not establish whole-of-APS productivity impact.
5. **Open Government Products/GovTech Singapore and reporting on Pair, 2026.** Evidence that Singapore's government-developed generative-AI assistant had reached 78,511 monthly active users by the third quarter of 2025. Provides a strong example of substantial recurring generative-AI use within a national public service, but remains a single-country case.
6. **Government of Canada — Directive on Automated Decision-Making and associated scope and peer-review guidance, current through 2026.** Important evidence concerning actual governance of AI and other automated systems contributing to administrative decisions. Demonstrates requirements covering partial as well as full automation, including impact assessment, transparency, explanation, testing and review.
7. **Treasury Board of Canada Secretariat — Responsible use of artificial intelligence in government and AI Strategy material, 2025–2026.** Evidence of federal institutionalisation of AI policy, including generative-AI guidance, an AI register and the 2025–27 Federal Public Service AI Strategy. Government policy and governance infrastructure establish readiness and direction but do not by themselves prove widespread operational deployment.

8. **Office of the Privacy Commissioner of Canada — Parliamentary material on automated decision-making, February 2026.** Provides an important qualification to formal governance arrangements: the Commissioner reported inconsistent implementation of Algorithmic Impact Assessments across federal institutions.
 9. **New Zealand Government Chief Digital Officer — Public Service AI Framework and Responsible AI Guidance for the Public Service: GenAI, 2025.** Primary evidence concerning New Zealand public-service AI governance and reported cross-agency experimentation. The guidance states that agencies are at differing stages of implementation and emphasises human oversight, transparency, privacy, security and accountability. Guidance and identified use cases do not establish uniform government-wide adoption.
 10. **OECD evidence concerning Brazil, 2025.** Includes the AI Litigation Project for tax appeals, National Treasury expenditure classification, Alice procurement oversight and citizen-service applications. Particularly useful in demonstrating substantive Latin American deployment across several government functions rather than treating governmental AI as predominantly a US or European phenomenon.
 11. **OECD comparative public-sector datasets, 2025.** The OECD notes that the European Commission Public Sector Tech Watch repository contained nearly 1,500 governmental AI uses as at March 2025 and that a Latin America and Caribbean database contained approximately 700 systems. Such repositories establish breadth of experimentation and deployment but combine heterogeneous applications and levels of maturity, so case counts should not be interpreted as equivalent installations or outcomes.
-

12. Science & Research

Status as at 30 August 2026

Artificial intelligence has moved materially into scientific research. It is already an established analytical and computational tool in a number of data-intensive disciplines, and increasingly supports literature discovery, scientific data analysis, modelling, simulation, prediction, coding and the design of candidate molecules or materials. More recent generative-AI and agentic systems are extending this role towards hypothesis generation, experiment planning and research assistance, while AI combined with laboratory automation is beginning to create closed-loop experimental systems.

Maturity nevertheless varies sharply. Machine learning used to classify astronomical observations, analyse biological sequences or approximate computationally expensive physical calculations is substantially more established than systems intended to conduct large portions of the scientific process autonomously. AI can now generate scientifically useful candidates, uncover patterns in enormous datasets and accelerate some simulations and experiments, but these achievements do not establish that AI can independently produce dependable scientific knowledge across disciplines.

The current Status is therefore best understood as **AI becoming part of the scientific instrument set, while progressively reaching further into the research process**. Its strongest position is in bounded tasks where data, objectives and evaluation methods are comparatively well defined. More ambitious systems that connect literature, reasoning, simulation, experimentation and scientific writing are advancing quickly, but remain much less mature.

This development is international but uneven. Major programmes in the United States, China, Europe, Japan and Australia are building scientific models, computing infrastructure and automated laboratories. Other regions are participating, but differences in computing resources, research data, specialised skills and funding mean that the ability to exploit AI for science is not globally uniform.

1. Where is AI currently being applied?

One of the longest-established applications is **analysis of large and complex scientific datasets**. Machine-learning systems classify images and signals, recognise structures, identify anomalies and extract relationships that would be difficult to identify manually. Such methods are used in astronomy, particle physics, genomics, ecology, chemistry, earth observation, climate science and many other disciplines. US National Science Foundation AI institutes, for example, apply AI to data from facilities including the Large Hadron Collider, LIGO and IceCube and to problems in astrophysics, chemistry and materials science.

AI is also increasingly used for **scientific modelling and simulation**. Learned models can approximate processes that traditionally require computationally expensive numerical calculations, allowing researchers to investigate many more possible conditions or candidates. In chemistry and materials science, AI models predict molecular or material properties and

increasingly assist inverse design: researchers specify desired properties and a model proposes candidate structures. MatterGen, for example, demonstrated generative design of inorganic materials against defined property constraints, while other models approximate atomic interactions across ranges of temperature and pressure. These are significant research capabilities, but computationally generated candidates normally still require further calculation and ultimately experimental validation.

In biology, AI-supported protein and molecular modelling has become an important research activity. Protein-structure prediction is the best-known example, but applications extend to genomics, molecular interactions, drug targets and biological sequence analysis. This overlaps Healthcare & Life Sciences when the research concerns medicines or clinical applications, but within this element the important point is methodological: AI is helping scientists navigate biological spaces too large for exhaustive experimental investigation.

Another developing area is **environmental and earth-system science**. AI models can learn relationships within large meteorological, oceanographic and climate datasets and produce predictions much faster than traditional numerical models for particular tasks. GenCast, for example, demonstrated probabilistic global weather forecasting at substantially lower computational time than conventional numerical ensemble forecasting and achieved stronger performance than the comparator on most of the targets used in its published evaluation. This demonstrates scientific capability; it does not mean numerical modelling has generally been displaced.

AI is increasingly entering the **earlier stages of research**. Generative systems can search and summarise literature, extract information from papers and datasets, generate code, suggest hypotheses, help design experiments and identify relationships across scientific information. China's ScienceOne platform, released in 2025, included a literature assistant and scientific-tool orchestration, while the Chinese Academy of Sciences' broader Pangu/Panshi model programme had by April 2026 developed a scientific foundation-model architecture supporting literature extraction, scientific knowledge representation, reasoning and tool planning.

Scientific data discovery itself is becoming AI-supported. The Chinese Academy of Sciences' ScienceDB AI, launched in 2025, provides conversational discovery of scientific datasets across multiple disciplines and is designed to return traceable and citable data resources. This illustrates an important but less conspicuous application: AI can improve access to the inputs of research, not only analyse them.

A further frontier is **automated experimentation**. Machine learning can select experimental conditions, robotics can execute experiments, instruments can return measurements and algorithms can decide what experiment to perform next. Berkeley Lab's AutoBot, for example, demonstrated an automated iterative system for optimising synthesis conditions for perovskite materials. In the United States, the Department of Energy and National Science Foundation are investing in AI-enabled autonomous and remotely accessible laboratories, including a July 2026 NSF investment of approximately \$380 million across 20 teams. These initiatives are substantial evidence of institutional movement towards automated science, but many such facilities remain research infrastructure and test beds rather than routine laboratory practice.

At the most experimental end are systems attempting to connect several stages of the research cycle. The AI Scientist research published in *Nature* in 2026 demonstrated a system able, within machine-learning research, to generate ideas, write and execute code, analyse results

and produce a paper. The work is notable precisely because end-to-end automation remains difficult; human filtering was still involved, and performance in computational machine-learning experiments does not establish equivalent autonomy in experimental sciences.

2. What problems or needs is AI being used to address?

A fundamental pressure is **scientific scale**. Modern instruments, simulations, sensors and experiments can generate quantities of information that cannot realistically be inspected manually. Astronomy, genomics, particle physics, climate science and earth observation are obvious examples. AI provides mechanisms for classifying, filtering and extracting useful signals from this volume of data.

A related problem is the size of possible **search spaces**. The number of plausible molecules, materials, biological configurations or experimental parameter combinations can be enormous. Researchers cannot physically test every possibility. AI is therefore being used to prioritise candidates, predict likely properties and direct scarce experimental resources towards more promising regions of the search space.

Computational cost creates another bottleneck. Some scientific simulations and quantum-mechanical calculations are highly resource-intensive. AI-based surrogate or emulator models can sometimes approximate these processes far more rapidly, allowing scientists to explore more possibilities before applying slower, higher-fidelity methods. The attraction is not simply speed: increasing the number of feasible computational experiments can change what questions researchers are practically able to investigate.

Researchers also face an expanding **knowledge burden**. Scientific literature and research datasets are growing across thousands of specialised outlets and repositories. Literature-mining and generative-AI systems are being used to locate papers, extract findings, summarise bodies of work and connect information across fields. A 2025 survey cited by *Nature Sensors* found 61% of surveyed researchers using AI tools to discover and summarise scientific papers. The figure should not be interpreted as a global institutional adoption rate, but it indicates that research-information tasks have become an important practical use of AI.

Laboratory work creates further constraints. Repetitive experiments can be slow and expensive, while some instruments are scarce and costly to operate. Automated laboratories seek to increase experimental throughput, improve repeatability and shorten the feedback cycle between hypothesis, experiment and analysis.

More broadly, AI is being directed at the long-standing problem of **scientific productivity**: how to turn growing research expenditure, data and technical capability into useful discoveries. The OECD has described accelerating research productivity as potentially one of AI's most economically and socially valuable applications, while also emphasising that AI's full potential in science remains far from realised.

3. What can AI currently do?

AI can currently perform a substantial range of specialised scientific tasks. It can classify observations, recognise complex patterns, predict quantities from data, generate candidate structures, approximate simulations, extract information from scientific text, write and debug research code, and assist with statistical and computational analysis.

In sufficiently structured domains, AI can also **propose solutions rather than merely classify existing information**. Generative materials models can suggest previously unspecified structures against desired properties. Algorithm-discovery systems can search for improved algorithms. [Google DeepMind's](#) AlphaEvolve, for example, demonstrated an AI-assisted evolutionary approach to algorithm discovery and optimisation, including mathematical and computing problems. Such systems broaden AI's scientific role from prediction towards searching for candidate solutions.

AI can substantially accelerate some calculations. Scientific machine learning increasingly acts as an emulator of more computationally demanding physical models, while specialised weather models demonstrate that learned systems can produce sophisticated forecasts rapidly. The scientific value depends on the task: speed is valuable only where the approximation retains sufficient accuracy, robustness and physical validity.

AI can also operate within **iterative physical experiments**. When models are linked to robotics and measurement systems, experimental results can be fed back into an algorithm that selects subsequent conditions. This creates a closed loop in which humans define the research objectives and infrastructure but some experimental decisions are made algorithmically. Demonstrations such as AutoBot show that this can compress particular optimisation exercises dramatically.

Generative AI can now support much of the informational work surrounding science: literature searching, summarisation, translation, code generation, data exploration, drafting and research administration. More advanced research agents can link some of these functions. Australia's CSIRO, for example, has been developing a scientific AI environment integrating models, data, workflows and scientific agents, illustrating movement from general-purpose conversational tools towards research-specific systems.

What AI cannot presently be assumed to do is **reliably replace the scientific process itself**. Scientific validity involves more than generating a plausible hypothesis or obtaining a statistically interesting result. It requires appropriate methods, reliable data, interpretation, controls, uncertainty assessment, replication and connection to existing knowledge. The 2026 AI Scientist work demonstrates that considerably more of a computational research cycle can now be automated than was previously practical, but it also illustrates why an experimental system should not be equated with autonomous, generally dependable scientific discovery.

4. How established is adoption?

Adoption is **established but highly uneven by task, discipline and type of AI**.

Conventional machine learning is already embedded in many computational and data-intensive research fields. The Royal Society described AI by 2024 as becoming an established tool across scientific fields, and OECD analysis documents applications throughout numerous stages of research. Bibliometric research published in *Nature* in 2026, analysing 41.3 million natural-science papers, also found accelerating prevalence of AI-augmented research.

Generative-AI research assistance is newer but spreading rapidly. Literature discovery, summarisation, coding and drafting are increasingly normal individual-researcher activities. This does not mean research institutions have uniformly integrated such systems into controlled scientific workflows. Personal use, institutionally supported tools and validated scientific systems represent different levels of adoption.

Specialised scientific foundation models and AI-for-science platforms are in a **growing deployment and infrastructure-building phase**. China has developed coordinated scientific-model and research-platform initiatives; Japan's RIKEN is building dedicated AI-for-science models and computing infrastructure; Europe launched the RAISE pilot in November 2025 to pool computing, data, talent and funding; and the United States is investing heavily in AI research infrastructure and programmable laboratories.

Automated laboratories sit at a less mature level. Real systems exist and can perform useful experiments, but the scale of new public investment itself indicates that broadly accessible AI-driven laboratory infrastructure is still being built rather than representing standard scientific practice. The NSF's 2026 automated-laboratory awards were explicitly intended to test, scale and demonstrate such methods.

End-to-end autonomous research is earlier still. Systems capable of chaining hypothesis generation, experimentation, analysis and writing are important demonstrations, but remain experimental and tend to work in unusually machine-accessible environments. They should not be presented as evidence that autonomous AI scientists are routinely operating across research institutions.

Geographical maturity also differs substantially. The United States, China, Europe and Japan have large institutional programmes and access to advanced computing. Australia is developing national-science-agency capabilities. Africa has growing capacity-building initiatives, including the African Institute for Mathematical Sciences' AI for Science programme, but infrastructure and specialist-resource constraints remain material. International Science Council case studies similarly identify infrastructure, skills, funding and data gaps in countries including Chile and Colombia.

The global picture is consequently not one of uniform transformation. It is a combination of **mature task-specific scientific machine learning, rapidly growing generative-AI assistance, substantial investment in AI-for-science platforms, emerging automated experimentation and still-experimental autonomous research**.

5. What appears to be driving adoption?

The most fundamental driver is the growing **computational and informational intensity of science**. Larger instruments, richer sensors, expanding research databases and more powerful simulations create scientific problems that increasingly require computational assistance.

Improved AI capability is another driver. Deep learning expanded pattern recognition in scientific data; foundation models broadened the kinds of data and tasks that can be addressed; and generative and agentic systems are making natural-language interaction, tool use, coding and multi-stage workflows more accessible. Scientific AI is consequently becoming relevant to researchers who are not themselves machine-learning specialists.

The availability of digital scientific data is equally important. AI is particularly effective where disciplines possess large, machine-readable datasets, standardised representations or computationally evaluable outcomes. This helps explain strong progress in areas such as molecular science, genomics, astronomy, materials and weather modelling.

Compute and infrastructure investment is reinforcing the trend. The NSF reported that more than 600 research teams had gained access through the US National Artificial Intelligence Research Resource pilot by 2025. Europe is developing RAISE around pooled compute, data, talent and funding. Japan is combining dedicated AI-for-science computing with Fugaku, while China is developing scientific foundation models and platforms across Chinese Academy of Sciences institutes.

Governments and research organisations increasingly regard AI-enabled science as strategically important. The scale of US autonomous-laboratory funding, the European AI in Science Strategy, China's coordinated AI-for-science activity and Japan's AGIS programme show that AI for science has moved beyond isolated academic projects into national and regional research policy.

There is also a practical attraction to **shortening research cycles**. If AI can reduce a simulation from hours to seconds, narrow thousands of experimental possibilities to dozens, automate repetitive laboratory work or help a scientist interrogate a large literature, the potential productivity gain is readily understandable even without assuming that AI makes the eventual discovery.

6. What are the present limitations?

Scientific AI depends heavily on **data quality and coverage**. Data may be sparse, inconsistent, proprietary, biased towards well-studied phenomena or generated under incompatible experimental conditions. Some scientific questions have abundant historical data; genuinely novel phenomena may have little. AI consequently tends to be strongest where suitable machine-readable evidence already exists.

Generalisation is another problem. A model that performs well within the conditions represented in its training data may fail under unfamiliar physical, biological or experimental conditions. Scientific applications therefore face a particularly demanding version of the broader AI problem of distribution shift.

Scientific validity is not the same as predictive accuracy. A model can make accurate predictions without providing an explanation corresponding to the underlying causal mechanism. In some applications prediction alone is valuable; in others, understanding why something happens is the scientific objective. Interpretability, physical consistency and uncertainty therefore remain important limitations of AI-for-science systems. A 2026 review identified high-quality data, interpretability and physical consistency among continuing challenges.

Generative AI creates additional reliability problems. Systems can fabricate facts or citations, misrepresent papers, produce incorrect code and express uncertainty poorly. In research, such errors can propagate into analyses or publications if plausible output is accepted without verification. The European Commission's updated 2026 research guidelines explicitly retain human responsibility and identify hallucinations, inaccuracies, bias, confidentiality, intellectual-property and transparency issues.

Reproducibility can also be difficult. Results may depend on proprietary models, undocumented model versions, inaccessible training data, stochastic outputs or computing environments that other researchers cannot reproduce. This creates tension with scientific expectations that methods and results should be independently inspectable and replicable.

Infrastructure is a major constraint. Frontier scientific models may require specialised accelerators, supercomputers, large datasets and engineering expertise. National initiatives in the United States, Europe and Japan explicitly focus on improving access to these resources, indicating that access remains a bottleneck rather than a solved problem.

Automation of physical science faces additional difficulties. Robots must manipulate real materials; instruments require calibration; unexpected physical events occur; and experiments may take hours, months or years. Closed-loop laboratories are consequently easier to implement for highly standardised, high-throughput experiments than for many forms of fieldwork or bespoke experimental science.

Finally, productivity effects should not automatically be equated with better science. The 2026 *Nature* bibliometric study found AI-augmented researchers associated with substantially higher publication and citation outcomes, but also found a contraction in the collective range of topics studied and greater concentration around data-rich areas. Because the analysis is observational, it does not establish that AI alone caused these outcomes, but it raises an important qualification: optimisation of individual research productivity and expansion of scientific knowledge are not necessarily the same thing.

7. What you should be aware of

“AI for science” covers activities at very different levels of maturity. A neural network classifying astronomical observations, an AI model predicting molecular properties, a generative assistant summarising papers, an automated materials laboratory and an experimental system writing machine-learning papers should not be treated as equivalent evidence of scientific automation.

Claims that AI has “discovered” something therefore require interpretation. AI may have identified a candidate, generated a hypothesis, predicted a structure or selected an experiment while human researchers designed the overall investigation, supplied the data, determined validation criteria, performed physical verification and interpreted the result. The scientific achievement can still be significant without attributing the entire discovery process to AI.

Human responsibility remains central. The European Commission's May 2026 research guidelines state that researchers remain ultimately responsible for scientific outputs produced with AI support and that AI systems are not authors or co-authors. This reflects a wider distinction between computational contribution and scientific accountability.

The evidence base is also unusually susceptible to **demonstration bias**. Successful AI-generated molecules, algorithms or experiments are more visible than unsuccessful searches, integration difficulties or models that fail outside their benchmark conditions. Provider demonstrations establish that a capability exists; they do not by themselves establish reproducibility, broad scientific value or widespread adoption.

Research access may become an important geographical and institutional dividing line. Advanced AI-enabled science increasingly depends on combinations of compute, data, specialised personnel and sophisticated laboratory or digital infrastructure. Large national laboratories and well-funded universities can therefore operate under materially different conditions from smaller institutions and research systems. Initiatives in Africa and Latin America demonstrate growing engagement, but international assessments continue to identify unequal infrastructure, skills and funding.

AI may also influence **which science gets done**. Data-rich, readily measurable and computationally evaluable problems are particularly suitable for machine learning. Areas with limited datasets, difficult-to-quantify phenomena or expensive physical validation may progress differently. The resulting distribution of research attention matters when interpreting broad claims that AI is accelerating “science” as a whole.

Scientific research itself remains the validation mechanism. An AI-generated hypothesis is not a scientific finding merely because it is novel; a generated molecule is not a useful material merely because it is computationally plausible; and an automatically produced paper is not reliable knowledge merely because it resembles scientific writing. Experimental confirmation, replication, critical interpretation and accumulated evidence remain essential.

Status at 30 August 2026: AI has become a meaningful component of contemporary scientific research, with its strongest established position in data analysis, prediction, modelling, simulation and specialised computational tasks. It is expanding rapidly into scientific information retrieval, generative design, research assistance and integrated research platforms. Automated laboratories demonstrate that AI can participate directly in iterative experimentation, while early autonomous-research systems show that several stages of computational research can now be linked. These developments are significant, but adoption and capability remain highly uneven across disciplines and countries. The current position is therefore not autonomous science replacing scientists, but a progressively deeper integration of AI into human-led scientific investigation.

Evidence Notes

1. **OECD — *Artificial Intelligence in Science: Challenges, Opportunities and the Future of Research*, June 2023.** Broad international assessment of AI across scientific disciplines and stages of research, including data analysis, simulation, robotics, scientific communication, reproducibility and developing-country participation. Particularly useful as a conceptual baseline. Its publication predates the most recent generative and agentic systems, so it is used principally for established patterns and structural constraints rather than the 2026 frontier.
2. **Royal Society — *Science in the Age of AI*, 2024.** Independent scientific assessment of how AI is affecting scientific methods, research integrity, skills and research practice. Supports the conclusion that AI had already become established in multiple research fields before the latest generative-AI wave.
3. **US National Science Foundation — National AI Research Institutes, NAIRR and 2025–26 AI-for-science programmes.** Primary evidence of large-scale institutional investment and actual research infrastructure spanning materials, physics, astronomy, drug development and other areas. NSF reported more than 600 teams connected through the NAIRR pilot and announced approximately \$380 million for 20 AI-enabled programmable-laboratory teams in July 2026. Funding and infrastructure establish institutional commitment and capability-building, not completed scientific outcomes.
4. **US Department of Energy — AI-Driven Autonomous Laboratories/Genesis Mission, 2026.** Primary evidence concerning integration of AI, robotics, real-time analysis, experimental control and scientific infrastructure. Much of the programme concerns capabilities being built and therefore should not be interpreted as evidence that autonomous laboratories are already routine.
5. **Lawrence Berkeley National Laboratory — AutoBot, September 2025.** Concrete experimental demonstration of machine-learning-directed robotic optimisation of materials synthesis. Useful evidence that closed-loop experimentation has moved beyond theory, while remaining a specific materials-science application rather than evidence of broad laboratory automation.
6. **European Commission — European AI in Science Strategy and RAISE, 2025–26.** Primary evidence of Europe's coordinated programme to pool scientific computing, data, talent and funding and accelerate AI adoption across scientific fields. RAISE remained in pilot development at the reference date, making it stronger evidence of strategic institutionalisation than of widespread researcher adoption.
7. **European Commission/European Research Area Forum — *Living Guidelines on the Responsible Use of Generative AI in Research*, updated May 2026.** Principal evidence concerning research integrity, accountability, hallucination, bias, confidentiality, intellectual property, transparency and human responsibility for AI-assisted scientific outputs.
8. **Chinese Academy of Sciences — ScienceOne, ScienceDB AI and Panshi/Pangu scientific-model initiatives, 2025–26.** Primary evidence that China is developing AI-supported literature discovery, scientific data access, scientific foundation models, reasoning and research-tool orchestration at institutional scale. Provider/institutional reporting establishes existence and intended functionality but does not independently demonstrate widespread use across Chinese science.
9. **RIKEN — Advanced General Intelligence for Science programme and AI for Science infrastructure, Japan, current through July 2026.** Primary evidence of Japan's development of scientific foundation models, research agents and dedicated computing integrated with major scientific-computing infrastructure. The programme demonstrates substantial institutional commitment, while several capabilities remain developmental.
10. **CSIRO — Science Digital, Australia, current through 2026.** Evidence of an Australian national-science-agency programme integrating AI tools, workflows, scientific agents, simulation and research assurance. Useful for demonstrating Oceania participation and the movement towards governed scientific-AI platforms; parts of the programme remained under development at the reference date.
11. **Nature — Lu et al., *Towards end-to-end automation of AI research*, 2026.** Peer-reviewed evidence of an experimental system connecting ideation, coding, experimentation, analysis, writing and review within

machine-learning research. Important frontier evidence, but restricted to a highly computational domain and involving human filtering; it should not be generalised into autonomous science across disciplines.

12. **Nature — GenCast research, 2024.** Peer-reviewed demonstration that machine learning can perform sophisticated probabilistic global weather prediction rapidly and, on the study's defined evaluation, outperform the comparator ensemble on most targets. Evidence of technical capability rather than proof that conventional numerical weather science has been displaced.
 13. **Microsoft Research — MatterGen, MatterSim and AI for Science materials research, 2024–25.** Provider and associated published research demonstrating generative materials design and AI-supported atomistic simulation. Useful evidence of technical capability and scientific tooling; provider descriptions should not be treated as independent evidence of widespread materials-science adoption.
 14. **Google DeepMind — AlphaEvolve, May 2025.** Primary provider evidence of an AI agent applying large-language-model generation and automated evaluation to algorithm discovery and optimisation, including mathematical problems. Demonstrates an expanding role for AI in candidate-solution discovery but does not establish general autonomous mathematical reasoning.
 15. **Nature — *Artificial intelligence tools expand scientists' impact but contract science's focus*, 2026.** Large bibliometric analysis covering 41.3 million natural-science papers. Provides unusually broad evidence of increasing AI-augmented research and associations with publication, citation and topic-selection patterns. Its observational methodology supports associations rather than simple causal claims about AI's effects.
 16. **Nature Sensors — editorial on AI-mediated scientific knowledge, July 2026.** Reports a 2025 survey finding that 61% of surveyed researchers used AI tools to discover and summarise papers. Useful evidence of substantial researcher-level uptake in literature discovery, but survey definitions and population mean the figure should not be interpreted as a global institutional adoption rate.
 17. **International Science Council — *Preparing National Research Ecosystems for AI*, 2025 edition.** Cross-country evidence showing materially different national readiness, including infrastructure, data, skills and funding constraints in countries such as Chile and Colombia and more extensive coordinated activity in China. Particularly useful in avoiding extrapolation from leading AI research systems to the global position.
 18. **African Institute for Mathematical Sciences — AI for Science Master's programme, South Africa, current through 2026.** Evidence of dedicated African capacity-building in AI-enabled science, including researchers from across Africa and links to cosmology, epidemiology and ecology. It demonstrates meaningful ecosystem development but does not establish continent-wide research adoption.
 19. **The Innovation — *AI for science: Progress, challenges, and perspectives*, available August 2026.** Contemporary review spanning mathematics, physics, chemistry, materials, medicine, life sciences, astronomy, geoscience and other fields. Useful synthesis of movement towards connected AI-supported research workflows and of continuing limitations in data quality, interpretability and physical consistency. As a review, it consolidates underlying research rather than independently demonstrating deployment.
-

13. Media, Entertainment & Creative Industries

Status as at 30 August 2026

Artificial intelligence is now present across much of the media, entertainment and creative economy, but its position differs considerably according to the activity being examined. Recommendation algorithms, audience analysis, automated transcription, search, tagging, advertising optimisation and other forms of machine learning have been embedded in parts of digital media for years. Generative AI has added a newer layer: systems can now create or modify text, images, music, voices and video, assist with coding and design, translate and dub content, and help creative workers develop ideas and drafts.

This does not mean that complete AI-generated films, games, books, news services or music catalogues have become the normal model of professional creative production. Much of the more established use is narrower: assisting human creators, automating production tasks, adapting existing material, managing large content libraries, personalising distribution or reducing the time required to create variants. Fully synthetic content is technologically possible and increasingly abundant in some digital channels, but its commercial consumption, professional acceptance and economic significance can be quite different from the volume being generated.

The sector is distinctive because AI affects not only productivity but the material from which AI itself is developed. Books, journalism, music, photography, film, illustration, performances and other cultural works have been used in training or can be reproduced, imitated or transformed by generative systems. Questions concerning copyright, licensing, attribution, performers' voices and likenesses, cultural rights and the authenticity of content therefore form part of the current operating environment rather than being peripheral legal issues.

The global picture is also uneven. Evidence from North America and Europe is extensive, but rapid development is evident in China and other Asian markets, while UNESCO work in Central America, Africa, the Arab region and Asia-Pacific shows both growing experimentation and substantial differences in skills, infrastructure and governance capacity. The appropriate Status is consequently one of **substantial and expanding AI use, with mature applications coexisting with rapidly developing generative applications and significant unresolved questions about rights, reliability, economics and creative labour.**

1. Where is AI currently being applied?

AI is being applied across the creative lifecycle: development, production, post-production, distribution, audience engagement and commercialisation.

In **publishing and journalism**, established applications include transcription, translation, tagging, search, recommendation, automated metadata, archive retrieval, data analysis and assistance with editing. Generative systems are increasingly used to summarise documents, generate possible headlines, assist research, reorganise information, produce drafts and adapt

material for different formats or audiences. Some publishers are also experimenting with AI interfaces through which readers can ask questions about their content. The Reuters Institute's 2026 international survey of 280 media leaders in 51 countries and territories found back-end automation to be the most widely prioritised AI area, while newsgathering and product development were also rising in importance.

AI is simultaneously becoming part of **news distribution and consumption**. Recommendation systems already shape what users encounter on platforms, while generative chatbots are emerging as an additional gateway to news. Reuters Institute data published in June 2026 found that weekly use of AI chatbots for news had risen to 10% across the surveyed markets, although only 1% of respondents described AI as their main news source. Growth was particularly visible in parts of Asia, Africa and Latin America as well as Southern and Eastern Europe. This is meaningful movement in distribution, but not evidence that AI chatbots have replaced conventional news channels.

In **film, television, online video and audiovisual production**, AI supports script and concept development, storyboarding, previsualisation, image and video generation, visual effects, background replacement, restoration, colour and image enhancement, voice processing, translation, dubbing and localisation. China's National Radio and Television Administration in July 2026 identified content planning, character development, scripting, storyboarding, image and voice generation, dynamic previsualisation and clip synthesis among areas of applied AI innovation in broadcasting and online audiovisual production. It also identified restoration, super-resolution, noise reduction, colour correction and frame enhancement, showing that AI's role extends well beyond generative content.

Distribution platforms are deploying some applications at substantial scale. YouTube's automatic dubbing system, for example, had expanded to 27 languages by February 2026; the company reported that more than six million viewers per day had watched at least ten minutes of automatically dubbed content during December. That is evidence of real deployment of AI-supported localisation, although the company itself warns that automated dubbing remains susceptible to mistranslation, pronunciation errors and varying quality between languages.

In **music**, AI is used for mastering, audio restoration, source separation, recommendation, composition assistance, synthetic voices and instruments and increasingly for generating complete recordings. The volume of fully generated music has risen especially rapidly. Deezer reported in July 2026 that AI-generated tracks had exceeded half of new daily uploads at their June peak, averaging about 90,000 such tracks a day. Yet fully AI-generated music accounted for only around 1–3% of listening on Deezer, demonstrating that production volume is not equivalent to audience adoption.

In **games**, generative AI is being used for brainstorming, coding assistance, prototyping, text, concept development and some asset production. The 2026 GDC State of the Game Industry survey, covering more than 2,300 industry professionals, reported that 36% used generative AI as part of their work. Use was substantially higher among publishing, support, marketing and business respondents than among game-studio respondents. Research and brainstorming were much more common than player-facing AI features.

In **advertising and commercial creative production**, AI is increasingly involved in generating or modifying copy, imagery, video and variations of advertisements, audience segmentation, campaign planning, media buying, optimisation and measurement. IAB research

describes a shift in the United States from experimentation towards AI becoming part of campaign infrastructure, although its evidence is predominantly American and cannot establish equivalent adoption globally.

Across these fields, AI also supports **discovery and personalisation**. Streaming services, video platforms, music services, news feeds and advertising systems use machine learning to rank, recommend and target content. These applications pre-date the recent generative-AI wave and remain among the sector's most established forms of AI.

2. What problems or needs is AI being used to address?

A central need is the **cost and time involved in producing and handling large volumes of content**. Creative businesses routinely generate transcripts, edits, translations, metadata, promotional material, visual variants, audio versions and different formats for multiple platforms. AI can automate or accelerate parts of this work, allowing one underlying piece of content to be adapted more quickly.

A closely related pressure is the **fragmentation of audiences across platforms, markets and languages**. Publishers, broadcasters, creators, game companies and advertisers increasingly distribute content through multiple channels. Translation, dubbing and automated adaptation can make content accessible to audiences that would previously have required separate localisation teams or budgets. YouTube's deployment of automatic dubbing illustrates this demand particularly clearly.

AI is also being used to manage **information abundance**. News organisations must examine documents, recordings and datasets; publishers manage extensive archives and catalogues; streaming services must classify enormous libraries; advertisers process large quantities of audience and campaign information. Search, classification, recommendation and summarisation systems help narrow this material into something people can work with.

For creators themselves, AI can reduce the cost of **experimentation and iteration**. A writer can explore alternative outlines, an illustrator alternative compositions, a filmmaker previsualisations, a game developer prototype assets or code, and an advertiser many variations of a campaign. The significance is not necessarily that the AI produces the final work. It can increase the number of ideas that can be tested before expensive human production begins.

In advertising and commercial media, another need is **performance optimisation**. Organisations want to determine which audiences to reach, which creative variants perform best, where advertising should be placed and how budgets should be adjusted. Machine learning has long been used for parts of this process, with generative and agentic systems now extending automation further into campaign creation and execution.

There is also a demand for **lower barriers to production**. Generative tools allow individuals and small organisations to undertake activities that previously required specialist software, technical expertise or larger production teams. That can broaden creative participation, but it also intensifies competition because the quantity of material that can be produced rises substantially.

Finally, some uses respond not to a production problem but to a **verification problem created partly by AI itself**. Newsrooms and platforms increasingly need to determine whether images, video, voices or documents are authentic. Reuters, for example, describes combining conventional verification techniques with AI-detection tools when checking user-generated imagery, while stressing that detection systems are not infallible.

3. What can AI currently do?

Current systems can generate plausible text, illustrations, photographs, voices, music and increasingly sophisticated video from instructions or reference material. They can alter existing images and audiovisual sequences; mimic styles and voices; generate translations and synthetic dubbing; summarise and classify documents; create metadata; retrieve information from large collections; write and explain code; recommend content; and produce many variations of a creative asset quickly.

This makes AI particularly capable at **component tasks**. A system can generate storyboard frames without directing a film, draft advertising copy without designing the entire campaign, suggest dialogue without writing a finished game, clean an audio track without producing the underlying performance, or translate and dub an existing video without originating its content.

Generative video has advanced enough for AI-produced sequences to form part of professional or semi-professional production workflows, particularly in concept generation, previsualisation, animation, effects, short-form content and advertising. China's broadcasting regulator now treats multiple generative audiovisual tasks as sufficiently practical to seek applications that are capable of implementation, verification and wider replication. That demonstrates technological maturation, but it does not establish that wholly generated television or cinema is normal industry practice.

AI can also produce **complete digital works at very high volume**. Music provides unusually measurable evidence. Deezer's tens of thousands of detected AI-generated tracks per day demonstrate that generating finished recordings is no longer technically exceptional. The much smaller share of listening nevertheless shows why capability, supply and market acceptance must be distinguished.

In games, AI can help generate concepts, code and prototypes, but the 2026 GDC evidence indicates that relatively little professional use has moved into player-facing generative systems. Of surveyed professionals, the predominant applications were brainstorming and work assistance rather than finished game content or live AI interaction.

AI can increasingly imitate characteristics of identifiable performers, including voice and visual appearance. This capability has become sufficiently material that collective agreements in film, television and games specifically address digital replicas, synthetics and consent. The 2026 SAG-AFTRA television and theatrical agreement strengthened restrictions surrounding both digital replicas of identifiable performers and synthetic performers, while the video-game agreement contains consent, disclosure and compensation provisions for specified digital-replica uses.

What present systems do **not** demonstrate reliably is autonomous mastery of an entire professional creative process. Long-form narrative coherence, factual reliability, culturally appropriate judgement, originality, sustained character consistency, artistic intention and the selection of what is worth creating remain materially different problems from generating technically plausible outputs. The gap between impressive individual clips or passages and a successful film, game, book, news product or advertising campaign therefore remains important.

4. How established is adoption?

Adoption is best described as **layered**.

At the most established level are recommendation and personalisation systems, advertising optimisation, transcription, content classification, search, tagging, moderation assistance, analytics and other relatively narrow machine-learning applications. These have been embedded in large digital platforms and media businesses for years.

A second layer comprises rapidly expanding generative assistance in professional workflows. Journalism provides relatively strong evidence. In the Reuters Institute's 2026 international media-leader survey, 97% considered back-end automation important, 82% identified newsgathering uses as important and 81% pointed to coding and product development. Yet assessments of results remained mixed: 44% said their newsroom AI initiatives were showing promising results, while 42% described results as limited. The survey is strategically sampled rather than representative of every newsroom, but it demonstrates that meaningful organisational experimentation and implementation are occurring across many countries without establishing uniform maturity.

Advertising appears similarly advanced in some highly digitised markets. IAB evidence from US brands and agencies describes movement towards scaled execution and AI-supported campaign infrastructure. Its geographical concentration means that it should be interpreted as evidence of significant adoption in the US advertising ecosystem rather than a global adoption rate.

Games show a more mixed position. The GDC survey's 36% individual-use figure demonstrates substantial use, but only 30% of respondents at game studios reported generative-AI use compared with 58% among publishing companies, support teams and marketing or PR organisations. Only a minority of reported uses concerned asset generation or player-facing features. Generative AI is therefore present in the industry without yet constituting a universal core production technology.

Music illustrates yet another form of adoption. Synthetic content production is already extremely easy and widespread at the upload stage, but listening remains comparatively limited. Deezer's data therefore supports describing AI music **supply** as highly developed while consumer adoption of fully generated music on that service remains low.

Film and television are harder to quantify globally. AI-assisted editing, VFX, restoration, localisation and production tools are real, and generative content is increasingly entering workflows. Collective bargaining agreements and regulatory initiatives confirm that the

technology is sufficiently consequential to affect actual working arrangements. Nevertheless, individual demonstrations and studio announcements should not be treated as evidence that fully generative production has become the prevailing professional model.

Geographically, maturity differs sharply. China's regulator is actively promoting deployable AI across broadcast and online audiovisual production, while UNESCO's July 2026 Central American assessment describes adoption together with continuing needs for technical skills, specialised training and governance. Evidence from Australia similarly indicates substantial public experimentation: Creative Australia reported that two in five surveyed Australians had used AI tools to create art or generate ideas. These different measures should not be combined into a single global percentage because they examine different populations and forms of use.

5. What appears to be driving adoption?

The most immediate driver is **productivity**. AI can reduce the time required for transcription, localisation, coding, editing, metadata creation, image variation, research and other repetitive or preparatory work. In industries where a large proportion of expenditure is associated with specialist human time, even partial automation can attract considerable interest.

A second driver is **falling production cost**. Generative systems make sophisticated text, imagery, audio and video capabilities available through widely accessible software. Tasks that once required dedicated studios, production equipment or specialist technical staff can sometimes be attempted by individuals and smaller companies.

A third is **competition for audience attention**. Digital media businesses operate in environments containing extremely large volumes of content. AI supports recommendation, personalisation, rapid experimentation and the creation of more formats and variants. These capabilities are attractive where businesses need to engage different audiences across multiple platforms.

The growth of **creator-led media** is another factor. Individuals producing video, podcasts, music, illustrations, games or written material can use AI as a form of inexpensive production assistance. This widens the potential user base beyond traditional studios and publishing organisations.

Improving technical capability is itself a driver. The convergence of large language models with image, audio and video generation means that workflows that previously required several specialist AI systems can increasingly be undertaken through integrated tools.

In some regions, public policy is actively encouraging adoption. China's 2026 MediaAIAC initiative explicitly seeks AI applications that can be implemented, verified and promoted across broadcasting and online audiovisual services. UNESCO initiatives across Asia, Africa, Central America and other regions likewise show that AI and culture have become matters of active institutional development, although these initiatives should not themselves be confused with commercial deployment.

Finally, organisations are responding defensively as well as offensively. The spread of synthetic content creates demand for detection, provenance, rights management and

verification technologies. AI adoption therefore generates additional AI-related activity elsewhere in the media ecosystem.

6. What are the present limitations?

The most prominent limitation is **reliability**. Generative systems can produce convincing material that is factually wrong, internally inconsistent or inappropriate to context. This is especially consequential in journalism, documentary production and advertising claims, but it also affects fiction and entertainment where continuity, character identity and precise creative control matter.

Technical quality remains uneven. Generated video can struggle with temporal and physical consistency; synthetic speech may mishandle pronunciation or emotion; translation can miss idioms or cultural meaning; image systems can fail to maintain detailed continuity; and long-form text can lose narrative coherence. YouTube's own guidance on automated dubbing, for example, identifies difficulties with accents, dialects, names and background noise.

A second major limitation is **copyright and rights uncertainty**. Generative models may have been trained on copyrighted books, journalism, images, recordings and audiovisual works. Whether particular training practices are lawful differs between jurisdictions and remains the subject of litigation, regulation and policy development. The European Union now requires providers of general-purpose AI models to maintain a copyright-compliance policy and publish a sufficiently detailed summary of training content. New Zealand, by contrast, stated in 2026 that international approaches remain unsettled and that its government would continue examining a possible GenAI copyright framework.

Performers' identity rights present a related issue. AI can reproduce or simulate voices, faces and performances at increasingly low cost. The detailed digital-replica provisions negotiated by SAG-AFTRA in film, television and games demonstrate that consent, control and compensation have become practical employment questions rather than hypothetical future concerns.

There is also uncertainty over **ownership and commercial status of generated outputs**. Copyright treatment of works containing AI-generated material differs between legal systems, and contractual arrangements between creators, publishers, studios, platforms and AI providers may impose further restrictions.

Another limitation is **audience acceptance and trust**. IAB research found a gap between advertisers' expectations and younger consumers' attitudes towards AI-generated advertising, while Creative Australia reported that 82% of surveyed Australians thought AI use in creative works should be disclosed. Such evidence does not establish universal consumer opposition, but it suggests that technically effective generation does not automatically produce social acceptance.

The economics are also unsettled. Generative tools lower some production costs while potentially increasing the volume of competing content. Deezer's experience illustrates the resulting problem: huge quantities of AI-generated music can enter a platform without

attracting proportionate genuine listening, and the service reported unusually high levels of fraudulent streaming associated with detected AI-generated tracks.

Skills and infrastructure create further differences between regions and organisations. UNESCO's 2026 work on creative-industry employment identifies substantial existing precarity and rapidly changing skills requirements, while its Central American assessment highlights needs for specialist capability and governance. Smaller creative organisations may have less capacity to evaluate models, negotiate licences or implement provenance and security systems than large platforms and studios.

Finally, AI systems themselves can reflect **linguistic and cultural imbalances** in their training material. UNESCO's 2026 work on AI, culture and intellectual property has drawn particular attention to Global South cultural resources and to questions of cultural rights and value distribution. The ability to generate content about a culture does not necessarily mean that a system represents its language, traditions or knowledge accurately or that the originating communities share in the resulting value.

7. What you should be aware of

“AI in the creative industries” covers technologies with very different levels of maturity. A recommendation engine selecting songs, an automated transcription system, an AI-generated advertising image, a synthetic voice, a newsroom research assistant and a completely generated video are not equivalent applications.

The distinction between **AI-assisted and AI-generated** work is particularly important. A film containing an AI-assisted visual effect may remain overwhelmingly human-created. A journalist may use AI to transcribe an interview without delegating editorial judgement. A musician may use AI for mastering rather than composition. Conversely, some digital material can now be generated almost entirely by AI. Treating all of these activities as a single category can materially distort the Status.

The volume of generated content should also not be confused with its economic or cultural importance. Deezer's experience—very high AI-generated upload volumes but only a small listening share—is a clear example. Similarly, availability of powerful text-to-video systems does not demonstrate widespread adoption by established film producers.

Human involvement remains substantial across professional applications. People select ideas, provide source material, direct systems, edit outputs, establish factual accuracy, manage intellectual property and determine whether generated material meets artistic, editorial or commercial objectives. Collective agreements covering AI likewise demonstrate that human performers and creators remain central participants whose relationship with the technology is being renegotiated rather than simply removed.

The evidence on employment requires particular care. AI can reduce the human effort required for specific tasks and may change demand for some kinds of creative work. It can also allow individuals or small teams to undertake work previously beyond their resources. Current evidence does not support reducing those competing effects to a single global conclusion about creative employment. UNESCO notes that many cultural and creative workers were already

operating under precarious conditions, making technological change especially significant but also difficult to isolate from broader industry pressures.

Rights and regulation are changing rapidly and differ substantially between jurisdictions. EU transparency obligations concerning some AI-generated or manipulated content began applying on 2 August 2026, while copyright approaches internationally remain divergent. Industry rules also matter: contractual protections for performers can affect permissible AI use even where general law would otherwise allow greater freedom.

Synthetic media also changes the information environment. The same technologies used legitimately for entertainment, translation and creative experimentation can produce misleading images, voices and video. Verification therefore becomes increasingly important for journalism, platforms and audiences. AI-detection tools can assist, but they do not provide infallible proof of authenticity.

Finally, evidence remains uneven geographically and methodologically. Surveys of American advertisers, international media executives, game-industry professionals and general populations measure different things. Their percentages should not be combined as though they describe one global rate of AI adoption. Evidence from Central America, Africa, Asia-Pacific, the Arab region and Oceania confirms that AI is a genuinely global creative-industry issue, but the depth of deployment, access to technology, language support, regulatory structures and economic conditions vary substantially.

Status at 30 August 2026: AI is materially embedded in Media, Entertainment & Creative Industries, particularly in recommendation, advertising technology, analytics, transcription, search, localisation and workflow assistance. Generative AI is rapidly extending that position into writing, imagery, audio, video and interactive production, but maturity remains uneven. Some uses have reached substantial real-world scale; others remain experimental or supplementary. The defining current development is therefore not simply the arrival of machine-generated creativity, but the integration of increasingly capable generative systems into human-led creative, production and distribution processes while the sector simultaneously works through unresolved questions of reliability, authorship, rights, remuneration, authenticity and cultural value.

Evidence Notes

1. **Reuters Institute for the Study of Journalism — *Journalism, Media, and Technology Trends and Predictions 2026, January 2026*.** International survey of 280 digital media leaders from 51 countries and territories. Provides unusually useful evidence on organisational AI priorities, newsroom applications and reported results. The sample is strategic rather than statistically representative of all journalism organisations and therefore should not be treated as a global adoption rate.
2. **Reuters Institute — *Digital News Report 2026: Emerging uses of AI chatbots for news, June 2026*.** Provides cross-market evidence on the emerging use of generative-AI chatbots as a route to news, including greater growth in parts of Asia, Africa and Latin America. Measures audience behaviour rather than newsroom adoption.
3. **National Radio and Television Administration of China — *Fifth MediaAIAC announcement, 21 July 2026*.** Primary evidence that Chinese broadcasting and online audiovisual policy now treats AI applications

across content generation, scripts, storyboards, image and voice generation, enhancement and restoration as significant areas for implementable industry development. The competition demonstrates policy and industry activity, not adoption by every broadcaster or producer.

4. **YouTube — automatic dubbing documentation and February 2026 product update.** Primary evidence for real large-platform deployment of AI-supported audiovisual localisation, including 27-language support and reported viewing of automatically dubbed content. Provider evidence establishes existence and scale of its own deployment but not industry-wide effectiveness; YouTube itself identifies continuing translation and speech-quality limitations.
5. **Deezer — AI-generated music updates, April and July 2026.** Primary platform evidence showing the exceptionally rapid increase in fully AI-generated music uploads and, importantly, the much smaller share of actual listening. Especially useful for distinguishing content generation from consumer adoption. It represents one music service and should not be extrapolated mechanically to the entire global music market.
6. **GDC Festival of Gaming — 2026 State of the Game Industry.** Survey of more than 2,300 game-industry professionals. Provides evidence that generative AI is being used materially in games-related work but that adoption varies strongly by role and is concentrated in brainstorming, routine work, coding and prototyping rather than player-facing systems. The survey population was not geographically representative of the entire global games workforce.
7. **Interactive Advertising Bureau — 2026 Outlook and related AI research, January 2026.** Evidence of substantial AI adoption among US brands and agencies across creative development, targeting, optimisation and campaign execution, together with evidence of consumer concerns about AI-generated advertising. These studies are useful indicators for a major advertising market but are geographically restricted and based on defined survey samples.
8. **European Commission — AI Act general-purpose AI and transparency guidance, current through July/August 2026.** Authoritative evidence concerning copyright policies, summaries of model training content and transparency obligations relevant to synthetic and manipulated content. These are EU legal requirements and should not be presented as global rules.
9. **SAG-AFTRA — 2025 Interactive Media Agreement and 2026 TV/Theatrical Agreements.** Primary evidence that digital replicas, synthetic performers, consent, disclosure and compensation have become operational labour issues in film, television and games. These agreements principally concern covered US productions and performers and do not represent global employment law.
10. **UNESCO — Skills and employment in the culture and creative industries, May 2026, and ReShaping Policies for Creativity 2026.** Provides global institutional context on creative-industry employment, skills, precarity and the effects of digital and AI change. Useful particularly in avoiding an interpretation based solely on large technology and entertainment companies.
11. **UNESCO/CECC-SICA — exploratory assessment of AI and cultural and creative industries in Central America, July 2026.** Valuable regional evidence concerning adoption, skills, intellectual-property concerns and governance needs in the SICA region. It is exploratory rather than a globally comparable quantitative adoption survey.
12. **UNESCO — AI, Culture and Intellectual Property initiative, June 2026.** Draws on consultations across Africa, Asia-Pacific and the Ibero-American region and provides important context concerning cultural rights, cultural resources, knowledge systems and value distribution, particularly in Global South settings.
13. **Creative Australia — Generative AI and Creative Work principles and 2026 audience research.** Provides Oceania evidence that generative AI has entered creative practice and that disclosure and authenticity matter to many surveyed audiences. Public-use findings should not be interpreted as equivalent to professional-industry adoption.
14. **New Zealand Ministry of Business, Innovation & Employment — Copyright Act updates, August 2026.** Primary evidence of continuing uncertainty concerning the treatment of GenAI training and AI-generated content and of the differing approaches being taken internationally.

- 15. CISAC/PMP Strategy — global study on generative AI in music and audiovisual industries, December 2024.** Provides international economic modelling concerning potential effects of generative AI on creators and illustrates the scale of concern within rights-holder communities. Its future revenue estimates are projections rather than current Status evidence and therefore have not been used as measures of realised economic displacement.
-

14. Professional Services

Status as at 30 August 2026

Artificial intelligence is becoming materially embedded in Professional Services, but not uniformly across firms, professions or types of work. The clearest current pattern is AI being incorporated into particular stages of professional work—research, information retrieval, document analysis, drafting, data analysis, audit procedures, knowledge management and workflow administration—while qualified professionals continue to exercise judgement, review outputs and retain responsibility for the service delivered.

Large accounting, audit, consulting and legal firms have moved beyond isolated experimentation in some areas. Major audit networks are incorporating AI into global technology platforms; large law firms are deploying specialist generative-AI systems across multiple practices; and consulting firms have developed internal AI systems that search and synthesise organisational knowledge and support client work. At the same time, adoption outside the largest firms is more uneven. UK evidence, for example, shows substantial AI activity among mid-tier accountancy firms while regulatory monitoring indicates that use within some smaller and non-public-interest audit practices remains limited. Legal-sector surveys likewise show rapidly growing professional use but significant variation between individual experimentation, approved organisational deployment and fully redesigned service delivery.

Generative AI has accelerated this development because much professional work involves language and documents. Systems can summarise extensive material, locate information, compare contracts, prepare drafts, interrogate knowledge repositories and assist analysis. More recent agentic systems are beginning to orchestrate sequences of tasks rather than simply generating a single response. These capabilities are important, but they do not establish that AI can independently perform the wider role of an auditor, lawyer, accountant or consultant. Professional judgement, contextual understanding, accountability, confidentiality and verification remain central.

AI is also beginning to affect the economics of professional services. Where clients believe AI reduces the time required for research, review or drafting, traditional hourly or effort-based charging models can come under pressure. Evidence from legal services in particular shows clients expecting greater AI-enabled efficiency and changes in pricing. In consulting, AI simultaneously creates a substantial new advisory and implementation market while giving clients tools with which they can perform some analysis and knowledge work themselves. The resulting change in service models is real, but its eventual scale remained unsettled at the reference date.

1. Where is AI currently being applied?

AI is being applied across the information-intensive parts of professional-service production.

In **accounting and audit**, established data analytics already allow firms to analyse large populations of transactions rather than relying exclusively on samples. AI is increasingly being added to these technology environments for anomaly identification, risk assessment, document

interrogation, accounting and auditing research, preparation and review of workpapers, contract analysis, checking financial statements and drafting elements of audit documentation. Generative AI can also help professionals locate relevant accounting or auditing guidance and synthesise information from large collections of internal material.

The scale can be substantial within major international networks. EY announced in April 2026 that agentic AI was being integrated into its global Assurance platform, used across approximately 160,000 audit engagements in more than 150 countries and territories. Its existing platform processes more than 1.4 trillion lines of journal-entry data annually. This establishes that AI-enabled technology is being incorporated into a very large real audit infrastructure, although it does not mean that every procedure within those engagements is performed by AI.

In **accountancy practices more broadly**, AI is being used for technical research, client management, tax and accounting work, correspondence, document processing and internal knowledge management. ICAEW's 2026 study of UK mid-tier firms found AI use highest in internal knowledge management/technical research and client management, while substantial proportions reported some use in audit, accounting and tax delivery. Most use was nevertheless described as moderate rather than extensive.

In **legal services**, generative AI is being used for legal research, summarising judgments and documents, precedent searches, drafting and revising contracts and correspondence, due-diligence review, document comparison, litigation preparation, knowledge management, transcription and administrative work. Singapore's Ministry of Law reported examples from law practices and in-house teams covering legal research, drafting and knowledge management. The Solicitors Regulation Authority in England and Wales described AI in August 2026 as increasingly being used for research, drafting, document review and administration.

Organisational deployment is no longer confined to experimental laboratories. Harneys, for example, announced a firm-wide deployment of Harvey in January 2026 following pilots across practices and jurisdictions, using it for navigating large precedent and document collections, drafting support and issue identification. Major international firms were also collaborating with increasingly specialised legal-AI platforms by August 2026. Such cases prove substantive deployment by particular firms, but not universal adoption across the legal profession.

In **consulting and advisory services**, AI is being used both internally and in client delivery. Applications include searching firm knowledge, conducting preliminary research, synthesising information, analysing data, preparing working drafts and presentations, generating code, developing scenarios and accelerating aspects of technology implementation. McKinsey's Lilli platform, for example, combines generative AI with the firm's internal knowledge and has been incorporated into employee workflows; the firm also reports using experience gained from the platform in work with clients.

Professional firms have simultaneously become suppliers of **AI-related professional services**. Consulting, accounting and multidisciplinary firms provide AI strategy, implementation, governance, cybersecurity, risk, controls and assurance services. This is distinct from merely using AI internally: AI itself has become an object of professional advice and assurance.

2. What problems or needs is AI being used to address?

A defining characteristic of professional services is the conversion of large quantities of information into specialised judgement. AI is therefore being directed primarily at the information-processing burden surrounding that judgement.

Professionals routinely search regulations, legislation, precedents, accounting standards, contracts, financial records, corporate documents, research, client information and internal knowledge. Finding, comparing and synthesising this material can consume substantial professional time. Generative AI provides a conversational mechanism for retrieving and reorganising such information and producing an initial synthesis.

Professional firms also face pressure to process increasingly large datasets and document collections. Auditors may examine millions of transactions; lawyers may review thousands of documents in litigation or due diligence; consultants may combine internal knowledge with extensive external material. AI can classify, extract, compare, summarise and identify patterns at a scale difficult to reproduce manually.

A second need is **productivity**. Drafting routine correspondence, preparing first versions of documents, summarising meetings, reviewing lengthy materials and assembling research are significant components of professional workload. AI can reduce the time required for parts of these activities, allowing professionals theoretically to allocate more attention to interpretation, client interaction and difficult judgement.

A third issue is **consistency and access to organisational knowledge**. Large professional firms have enormous repositories of previous work, methodologies and specialist expertise. Systems such as McKinsey's Lilli illustrate the use of AI to make that accumulated knowledge more readily searchable and reusable.

Professional firms are also responding to **client expectations**. Thomson Reuters' 2026 legal-sector evidence, based on respondents across 46 countries, found a substantial gap between the AI-enabled improvements clients expected from law firms and what clients believed most providers were actually delivering. Seventy-one per cent of surveyed in-house legal professionals expected outside firms to alter their charging models as AI use increased, while only 28% of responding law firms said they had changed pricing in response. These figures concern the survey population and should not be generalised mechanically to every legal market, but they demonstrate that AI is becoming a client-service and commercial issue rather than solely an internal technology issue.

In audit, another objective is **quality and risk identification**, not simply faster completion. Large-scale analytics can examine entire transaction populations, identify unusual relationships and direct auditor attention to potential risks. UK Financial Reporting Council guidance published in March 2026 explicitly recognised potential improvements in both audit quality and efficiency from generative and agentic AI while maintaining professional accountability.

Finally, professional firms are responding to demand from clients that need help adopting AI themselves. AI governance, implementation, controls, assurance, regulatory interpretation and

organisational redesign have become areas of professional-service demand.

3. What can AI currently do?

AI can already perform a meaningful collection of bounded professional-service tasks.

It can retrieve and summarise information; extract clauses and data from documents; compare versions; categorise large document sets; draft emails, reports, contracts and other text; translate or reformat material; generate and explain code; search organisational knowledge; identify patterns in financial datasets; and provide first-pass analyses for professional review.

In legal work, specialist systems can combine language-model capabilities with authoritative or firm-specific sources, allowing lawyers to interrogate documents, compare precedents, generate draft wording and perform initial research considerably faster than manual searching in some circumstances.

In audit and accounting, AI can complement conventional analytics by highlighting anomalies, interrogating contracts and documentation, retrieving technical guidance and automating parts of document checking. Agentic systems can increasingly organise multi-step activities by selecting or coordinating tools rather than responding to only one prompt. EY's 2026 deployment demonstrates that such capabilities can be integrated into global audit workflow infrastructure rather than existing only as standalone demonstrations.

In consulting, AI can give professionals rapid access to codified organisational knowledge and provide starting points for research, analysis and problem solving. McKinsey reported that practical use of Lilli led to revised workflows and numerous additional use cases, while also revealing trade-offs between model accuracy, speed and operating cost—an indication that technical capability does not automatically translate into economically optimal deployment.

What AI cannot presently be assumed to do is perform an entire professional engagement with the reliability, responsibility and contextual judgement expected of a qualified professional.

A plausible legal analysis may contain fictitious authorities. An audit analysis may identify a statistical anomaly without understanding its commercial explanation. A consulting synthesis may reproduce weaknesses in underlying sources or omit organisational context. A draft accounting conclusion may overlook an unusual contractual or regulatory fact.

The distinction is therefore between **performing components of professional work** and **being the accountable professional**. At 30 August 2026 the evidence strongly supports the former; it does not establish the latter.

4. How established is adoption?

Adoption is significant but layered.

Large international professional firms are among the more advanced organisational adopters. Firm-wide legal-AI deployments, global audit platforms and consulting knowledge

systems demonstrate movement beyond individual experimentation. For some research, knowledge-management, analytics and drafting activities, AI is becoming part of normal professional infrastructure.

Legal services appear to be moving rapidly from experimentation into organised deployment, particularly among large firms. Thomson Reuters' 2026 research covered 736 law-firm respondents across 46 countries, although 421 were in the United States. Its results indicate widespread organisational attention to AI but also substantial differences in strategy, implementation and perceived value. The same research found unauthorised or “shadow” AI use among professionals, illustrating an important distinction: high individual use does not necessarily mean controlled institutional adoption.

Accounting and audit show an especially wide maturity range. ICAEW found that 86% of surveyed UK mid-tier firms had technology strategies explicitly including AI, and its underlying study reported AI use across knowledge management, client management, audit, accounting and tax. Yet ICAEW's audit-monitoring experience reported in May 2026 suggested that actual AI use within the non-public-interest audit firms it monitored remained very limited. Large-firm deployment and smaller-firm practice therefore should not be conflated.

Globally, professional accountants increasingly encounter AI. ACCA's 2026 global talent research reported high confidence among finance professionals in their ability to learn and apply AI skills, alongside continuing concern about job impacts. This indicates substantial awareness and use, but workforce surveys do not by themselves establish that AI has been deeply embedded into professional-service processes.

Audit regulators are themselves responding to adoption. The UK FRC issued specific generative- and agentic-AI guidance in March 2026 and opened an Audit Tech & AI Sandbox in June. Internationally, IAASB's technology work included contributions from more than 240 stakeholders across six continents and August 2026 proposed revisions to auditing standards addressing the increased use of technology. These developments demonstrate that AI use is material enough to require regulatory and standard-setting attention, although guidance and consultations are not evidence of universal deployment.

Consulting is harder to quantify consistently. Major firms clearly use AI internally and sell substantial AI-related services, but comparable global evidence showing the depth with which AI has transformed ordinary consulting engagements is limited. Public announcements often describe platforms, capabilities or client offerings more readily than actual utilisation across projects.

Geographical differences are also material. Strong public evidence is concentrated in large firms and in markets including North America, the United Kingdom, Europe, Singapore, Australia and New Zealand. This should not be interpreted as evidence that other regions lack adoption; rather, comparable public evidence about firm-level deployment in much of Africa, Latin America, the Middle East and parts of Asia is thinner.

5. What appears to be driving adoption?

The first driver is the unusually high proportion of professional work devoted to information processing. Modern language models are well matched to documents, research and structured professional knowledge, making professional services a natural field for experimentation and deployment.

The second is **competitive and client pressure**. Clients increasingly know that AI can reduce the time required for certain research, drafting, analysis and review tasks. This creates pressure for providers to demonstrate faster delivery, improved quality or lower cost rather than retaining the full economic benefit of productivity improvements. Legal-sector evidence shows this beginning to influence discussion of pricing and client relationships.

The third driver is **professional capacity**. AI can provide a first draft or initial analysis that would previously have required additional manual hours. For firms with large volumes of recurring work, even modest savings per matter can accumulate.

A fourth is **knowledge leverage**. Professional firms depend heavily on accumulated intellectual capital. AI offers a means of making previous work, methodologies and expertise more readily accessible across large organisations.

A fifth is the development of increasingly profession-specific AI products. Legal platforms, audit tools and systems connected to authoritative professional sources reduce some of the shortcomings of completely general-purpose models. Google's August 2026 expansion of Gemini Enterprise for Legal, for example, reflected increasing competition to embed AI within professional legal workflows and software environments. Provider availability does not demonstrate adoption, but expanding specialised supply lowers practical barriers to deployment.

Finally, AI is itself creating revenue opportunities. Organisations require external expertise in technology selection and implementation, AI governance, cybersecurity, regulatory compliance, risk, control design, operating-model changes and assurance. Professional firms can therefore have both an efficiency incentive to use AI internally and a commercial incentive to build AI expertise.

6. What are the present limitations?

The most visible technical limitation is **reliability**. Generative systems can produce statements, calculations, citations or interpretations that appear credible but are wrong. This is particularly consequential in professional services because users and clients may attach substantial weight to authoritative-sounding professional output.

Legal practice provides especially clear evidence. The Solicitors Regulation Authority warned in August 2026 of inaccurate or hallucinated information appearing in legal research, advice, analysis and court submissions. Australia's Federal Court similarly emphasised that generative AI may generate inaccurate, fictitious or plainly wrong material.

Confidentiality and privilege impose another constraint. Professional firms routinely hold commercially sensitive, legally privileged, personal and regulatory information. Entering that material into unsuitable external AI systems can create confidentiality, privacy, security or contractual risks. Singapore's legal-sector guidance therefore emphasises professional ethics, confidentiality and transparency, while New Zealand legal guidance similarly stresses human verification and protection of client information.

Professional accountability does not transfer to the model. Audit regulators continue to hold firms and responsible individuals accountable for audit quality. Lawyers remain responsible for submissions and advice. Accountants remain bound by professional and ethical duties. This limits the extent to which output can simply be accepted because it came from an advanced system.

Data and integration matter. A general model may know little about a firm's private knowledge, client history or current authoritative sources. Connecting AI securely to those resources requires technology, permissions, data management and governance. McKinsey's experience with Lilli also illustrates operational constraints: higher usage revealed differing trade-offs between models in accuracy, speed and cost.

Economics are not automatically favourable. AI systems require software expenditure, computing resources, integration, training, supervision and output verification. Time saved generating a draft may partly be offset if an experienced professional must extensively check or reconstruct it.

There are also **workforce and capability questions**. Some activities most amenable to AI—research, document review, routine drafting and basic analysis—have traditionally helped junior professionals acquire knowledge and judgement. If the task composition of early-career work changes materially, professional development models may also have to change. Thomson Reuters' 2026 law-firm research found that many respondents considered early-career lawyers dependent on experienced mentorship for skills affected by AI, but the longer-term consequences were not yet established.

Finally, evidence about productivity and quality is still uneven. Vendor demonstrations and firm announcements establish that systems exist and are deployed. They do not establish how much time they save across all engagements, whether savings persist after review costs, or whether professional outcomes consistently improve.

7. What you should be aware of

“AI in Professional Services” encompasses very different technologies and levels of responsibility. An audit system identifying unusual journal entries, a lawyer using AI to summarise a contract, a consultant searching a firm's internal knowledge and an autonomous agent coordinating several workflow steps should not be treated as equivalent forms of adoption.

Firm size also matters. Highly visible global networks have resources to develop proprietary systems, negotiate enterprise technology partnerships and establish dedicated governance

teams. Smaller practices may depend on general-purpose or third-party products and face different cost, expertise and integration constraints.

Individual professional use must also be distinguished from organisational deployment. Survey evidence showing that lawyers or accountants use generative AI can include personal experimentation, general productivity tools, authorised specialist systems and unauthorised “shadow AI”. These have very different implications for institutional maturity.

Human involvement remains fundamental. Regulators and professional bodies across several jurisdictions continue to place responsibility on the professional rather than the AI system. Singapore's legal guidance, the English and Welsh SRA warning, Australia's Federal Court practice note, New Zealand Law Society guidance and UK audit regulation all illustrate substantially the same underlying position despite differences in their legal frameworks: AI assistance does not remove existing professional obligations.

Readers should also separate **productivity claims from service quality**. Producing a document more quickly is demonstrable technical productivity; whether it produces better advice, better audits or better client outcomes is a different question.

The effect on professional-service economics is consequently still developing. AI can reduce effort in some tasks, but professional services are sold partly on judgement, accountability, trust, specialised knowledge and the capacity to deal with ambiguity. Client pressure on hourly billing and traditional staffing structures is therefore significant evidence of change, but not proof that a new dominant business model had emerged by August 2026.

The evidence is geographically uneven. Global professional-service firms provide genuine multinational examples, and surveys such as Thomson Reuters and IAASB's stakeholder work extend across many countries. Nevertheless, published evidence remains disproportionately rich for the United States, United Kingdom, Australia, New Zealand, Singapore and major international firms. The Status should not therefore be interpreted as showing uniform adoption globally.

Status at 30 August 2026: AI has moved materially into Professional Services. Its strongest present position is as an analytical, research, drafting, knowledge and workflow technology operating inside human-led professional processes. Large audit, legal and consulting organisations provide clear examples of production-scale deployment, while adoption among smaller firms and across countries remains much less uniform. Generative AI is also beginning to change client expectations, staffing assumptions and charging models, but evidence that AI can independently undertake the full judgement, responsibility and contextual reasoning of the professions remains absent. The current transformation is therefore substantial, but principally one of **professional work being reorganised around AI-assisted tasks rather than professional responsibility being transferred to AI**.

Evidence Notes

1. **Financial Reporting Council — Generative and Agentic AI Guidance, 30 March 2026.** UK regulatory evidence that generative and agentic AI were entering audit sufficiently rapidly to require specific guidance. Particularly useful for distinguishing potential efficiency and audit-quality benefits from continuing human accountability. UK-specific rather than global adoption evidence.

2. **Financial Reporting Council — Audit Tech & AI Sandbox, June 2026.** Shows active regulatory engagement with large, mid-tier and smaller audit firms experimenting with or scaling AI. Participation and regulatory interest establish development activity, not widespread adoption.
3. **IAASB — Technology Quality Management Roundtables, February 2026, and proposed revisions to ISA 330, ISA 500 and ISA 520, August 2026.** International evidence involving more than 240 stakeholders from six continents and showing that technology use in audit has become material to international standard setting. Consultation and stakeholder participation do not measure adoption rates.
4. **ICAEW — Evolution of Mid-Tier Accountancy Firms 2026.** Survey evidence that 86% of participating UK mid-tier firms included AI in their technology strategies and that AI was being used across knowledge management, client management, audit, accounting and tax, typically at moderate rather than extensive levels. Restricted to the participating UK mid-tier market.
5. **ICAEW Professional Standards Department — Artificial intelligence in audit work: managing the risks, 6 May 2026.** Important counterweight to high-level adoption claims: ICAEW's monitoring experience indicated very limited current AI use in the non-public-interest audit firms reviewed. The observation concerns a particular regulated population and should not be extrapolated to major international networks.
6. **EY — Global roll-out of enterprise-scale agentic AI in Assurance, April 2026, together with EY Audit Technology information.** Primary organisational evidence of AI integration into a global audit platform supporting approximately 160,000 engagements in more than 150 countries and territories. Strong evidence that deployment at large-firm scale exists, but provider evidence cannot independently establish resulting improvements in audit quality or sector-wide adoption.
7. **Thomson Reuters — Future of Professionals 2026, including Legal Report, June 2026.** Global survey evidence covering approximately 1,800 professionals across 62 countries overall and 736 law-firm respondents across 46 countries in the legal analysis. Provides evidence on professional use, client expectations, shadow AI, pricing pressure and organisational strategy. The legal sample is weighted substantially towards the United States and responses are self-reported.
8. **Singapore Ministry of Law — Guide for Using Generative Artificial Intelligence in the Legal Sector, 6 March 2026.** Government evidence on real legal-sector use cases and professional expectations covering research, drafting and knowledge management. Particularly useful for confidentiality, professional responsibility and transparency considerations and for providing substantive Asian evidence.
9. **Solicitors Regulation Authority, England and Wales — warning notice on misuse of AI, 17 August 2026.** Current regulatory evidence that AI had become part of day-to-day legal practice while generating concerns about hallucinated authorities, confidentiality and inadequate oversight. Regulatory warning evidence establishes material use and risk, not the percentage of firms adopting AI.
10. **Federal Court of Australia — Use of Generative Artificial Intelligence Practice Note, 16 April 2026.** Authoritative evidence recognising potential efficiency, cost and access benefits while explicitly identifying fabrication and accuracy risks in legal proceedings. Demonstrates institutional response rather than law-firm adoption rates.
11. **New Zealand Law Society — Generative AI guidance for lawyers, current 2026 guidance.** Oceania evidence emphasising independent checking, professional responsibility, confidentiality, privacy and cybersecurity. Useful for demonstrating that these constraints are not unique to the largest legal markets.
12. **Harneys — firm-wide Harvey deployment, 21 January 2026.** Organisational evidence of a multinational law firm moving from pilot activity to firm-wide deployment for drafting, precedent/document review and issue identification. Customer evidence establishes that real deployment exists but should not be generalised to the legal sector as a whole.
13. **Google — Gemini Enterprise for Legal launch reporting, 25 August 2026.** Evidence of increasingly specialised AI infrastructure for legal-service workflows and collaboration with several major international law firms. Product availability and collaboration do not establish completed or widespread adoption.

14. **McKinsey & Company — Lilli case material and June 2026 account of platform evolution.** Primary evidence of an international consulting firm incorporating generative AI into internal knowledge and professional workflows and learning from actual usage about model accuracy, speed and cost. Firm-supplied evidence does not independently quantify productivity or represent consulting firms generally.
 15. **ACCA — Global Talent Trends 2026.** Broad international workforce evidence showing considerable exposure to and confidence concerning AI among finance and accountancy professionals while recording continuing employment concerns. Workforce perceptions are informative about diffusion but do not directly measure firm-level process deployment.
-

15. Telecommunications

Status as at 30 August 2026

Artificial intelligence is now being applied across a substantial part of the telecommunications sector, from radio-network optimisation and fault management to customer service, fraud detection, energy management and network construction. Its maturity, however, differs considerably by application and by operator.

The most established uses generally involve relatively bounded analytical or automation tasks: forecasting network demand, detecting anomalies, optimising radio parameters, identifying fraud or spam, assisting customer-service staff and operating conversational service channels. Some large telecommunications groups now report AI-supported processes at substantial operational scale. Vodafone, for example, reported in May 2026 that its generative-AI customer assistant was live across all its European markets and that more than 70% of relevant conversations were being resolved end-to-end; it also reported AI-supported network operations and energy optimisation. These are company-reported outcomes rather than independent evidence of industry-wide performance, but they demonstrate that significant operational deployment exists.

A newer stage is also becoming visible: AI systems that do not merely recommend an action but participate in closed-loop network control, automated network construction or multi-step operational decision-making. NTT DOCOMO began commercial use of a generative-AI network-maintenance agent in February 2026, introduced an AI-driven mobile-network optimisation system in June, and used agentic AI in constructing a commercial 5G core. Mobily and Ericsson have reported closed-loop autonomous network management during the 2025 Hajj. These cases show that higher levels of autonomy are no longer purely theoretical. They do not, however, establish that telecommunications networks generally operate autonomously.

The clearest current Status is therefore one of **AI becoming embedded throughout telecommunications operations while remaining uneven in depth**. Established machine learning and automation coexist with rapidly expanding generative AI and a smaller number of operational agentic or highly autonomous systems. At the same time, much of what is described as “AI-native networking”, AI-RAN and AI-enabled 6G remains under development, standardisation, proof-of-concept or trial rather than routine commercial deployment.

1. Where is AI currently being applied?

AI is being applied first and foremost to the **planning, operation and optimisation of telecommunications networks**. Mobile operators collect very large quantities of information about traffic, signal quality, equipment performance, faults and customer experience. Machine-learning systems analyse these data to predict congestion, adjust network parameters, balance traffic between cells, identify coverage problems and prioritise interventions.

Radio access networks are an important area. AI can optimise radio-resource allocation, signal parameters, spectral use and traffic distribution. In August 2026 SoftBank and Ericsson

reported a trial of an AI-based scheduler on SoftBank's commercial Japanese 5G network, measuring improvements in spectral efficiency and user throughput. Because this was explicitly a validation exercise, it is evidence of capability on a commercial network rather than evidence of routine nationwide operation. By contrast, DOCOMO reported putting Nokia's AI-driven MantaRay AutoPilot into operation in June 2026 to continuously optimise mobile-network quality, illustrating movement from testing into operational use.

AI is also increasingly involved in **fault detection, service assurance and maintenance**. Models identify unusual patterns across network alarms and performance data, predict equipment problems, correlate apparently separate incidents and suggest likely causes. Generative-AI systems can search technical documentation, summarise incidents and support engineers troubleshooting complex multi-vendor networks. DOCOMO's network-maintenance AI agent entered commercial operation in February 2026 specifically to help diagnose faults across increasingly complex 4G and 5G environments.

Some operators are moving beyond decision support towards **closed-loop or self-healing operations**. During the Hajj, Mobily and Ericsson reported AI systems detecting coverage anomalies, dynamically balancing network loads and automatically carrying out hundreds of optimisation actions. Telstra demonstrated an agentic system in a live telco-cloud environment that detected an infrastructure outage and shifted network applications to healthy hardware, although Telstra described this as a proof of concept rather than production deployment.

A related application is **network design, configuration and deployment**. DOCOMO and its partners reported using agentic AI and GitOps to automate parts of 5G-core design and construction, reducing the construction period by approximately 80% compared with their conventional process. The resulting AWS-based 5G core entered commercial service in February 2026, although DOCOMO continues to describe complete AI-based network-operation automation as a future objective.

AI is being used for **energy and resource management**. Telecommunications networks contain large numbers of energy-consuming radio sites, network facilities and computing systems whose demand varies with traffic. AI can forecast demand and adjust capacity or equipment operation accordingly. Vodafone reports using AI for radio-access-network energy optimisation and energy-consumption forecasting. GSMA analysis nevertheless emphasises a two-sided effect: AI may reduce network energy consumption through optimisation while AI workloads, data centres and additional computation create new energy demand.

Another established area is **customer management**. AI assistants answer questions, troubleshoot connections, recommend products, summarise customer histories and assist contact-centre personnel. Generative AI has expanded the range and fluency of these systems. Vodafone's SuperTOBi deployment across its European markets is one large-scale example. Industry surveys also consistently identify customer service as one of the leading telecommunications AI applications.

AI is also used for **fraud, spam, scam and security detection**. Telecommunications networks provide valuable real-time information about calls, messages, identities and behavioural patterns. AI can identify unusual traffic, suspected scam calls, spam campaigns or account behaviour much faster than fixed blacklists alone. Australia's communications regulator reported in April 2026 that telecommunications operators were using AI to detect scam activity

in real time. AI-based network security products are also being deployed across operators in multiple regions.

Finally, telecommunications providers are beginning to incorporate AI into the **communications services they sell**, rather than using it only internally. e& UAE, for example, announced in August 2026 that agentic-AI capabilities were being embedded in certain business fixed and mobile connectivity propositions. Such offerings are emerging and should not yet be equated with mainstream communications services.

2. What problems or needs is AI being used to address?

The underlying telecommunications problem is one of **scale and complexity**. Modern operators manage millions of customers, thousands or tens of thousands of network sites, several generations of mobile technology, fixed and mobile infrastructure, cloud systems and equipment from multiple vendors. Network conditions also change continually as people move, devices connect and traffic shifts.

Traditional rule-based automation remains useful, but increasingly complex networks create combinations of conditions that are difficult to anticipate manually. AI is therefore being used to recognise patterns, predict events and choose responses from volumes of operational data that human teams could not continuously examine themselves. GSMA and TM Forum both identify fragmented systems, siloed data and multi-vendor complexity as central reasons why operators are seeking more intelligent forms of automation.

Reliability creates an additional pressure. Telecommunications is critical infrastructure: outages can interrupt personal communications, commercial activity, emergency access and public services. This creates demand for systems capable of detecting deterioration before users complain, isolating probable causes rapidly and restoring services with less manual intervention.

Capacity management presents a related problem. Traffic demand can shift quickly between locations and services, especially during major events. AI-based optimisation can allocate resources dynamically rather than relying solely on fixed engineering rules. The Mobily Hajj deployment illustrates the attraction in an unusually demanding environment involving millions of users concentrated in a limited area.

Operators also face substantial **cost and productivity pressures**. Large networks require continuous monitoring, field maintenance, software configuration and customer support while traditional connectivity markets in many countries offer limited revenue growth. The GSMA's 2026 telco-AI assessment describes an industry moving beyond experimentation but still seeking ways to turn AI from principally an efficiency mechanism into a source of sustainable revenue.

Customer-service applications address very large volumes of repetitive contacts while attempting to improve response time and consistency. Fraud applications respond to a different scaling problem: scam and malicious traffic can operate at machine speed and change rapidly, making static rules increasingly inadequate.

Energy use has also become material. Mobile access networks consume substantial electricity, while 5G expansion, traffic growth and emerging AI workloads can add demand. This gives operators an incentive to vary network resources more closely with actual traffic rather than operating all infrastructure at maximum capacity continuously.

3. What can AI currently do?

Current AI can already perform important specialised telecommunications tasks.

It can analyse streams of network measurements and recognise abnormal behaviour; forecast traffic and demand; identify likely faults; classify alarms; optimise selected radio parameters; balance traffic; recommend or initiate remedial actions; estimate energy consumption; analyse customer histories; conduct conversational customer interactions; and identify patterns associated with spam, scams or fraud.

Increasingly, AI can also connect several steps. Agentic systems can interpret an operational objective, gather relevant information, diagnose a situation and execute defined actions through software interfaces. DOCOMO's 5G-core work shows that AI can participate in network design and configuration workflows, while the Mobily deployment demonstrates that closed-loop systems can make repeated optimisation decisions during live network operations.

Generative AI adds a different capability. It can translate natural-language instructions into technical queries, summarise alarm histories, retrieve engineering knowledge, produce suggested procedures and allow operational staff to interact with complex telecommunications systems through conversational interfaces.

This does **not** mean that current AI can dependably operate every aspect of a national telecommunications network without human involvement. Telecommunications networks combine physical infrastructure, software, radio propagation, cybersecurity, regulatory requirements and unpredictable user behaviour. Errors in autonomous control can affect very large numbers of users.

The distinction is especially important for newer AI-RAN and 6G work. DOCOMO and VIAVI demonstrated AI-driven network control with digital twins in 2026, and DOCOMO and Samsung validated user-level AI-RAN optimisation. These demonstrate technically important capabilities, but the former concerned prospective 6G technology and the latter was explicitly a validation rather than established mass deployment.

4. How established is adoption?

AI adoption in telecommunications is **broad but strongly layered**.

Machine learning used for forecasting, network analytics, fraud detection, customer management and selected optimisation tasks is relatively established. Generative AI has moved rapidly into customer service, employee assistance and operational support. A 2025 NVIDIA survey of more than 450 telecommunications professionals found that almost all responding organisations were either deploying or assessing AI projects, with network planning, network

operations and field operations prominent use cases. Such vendor-sponsored surveys are useful indicators of activity but should not be interpreted as precise measures of worldwide deployment.

GSMA material reported that 65% of operators had an AI strategy by late 2024 and that, among operators covered by its research, substantial generative-AI deployment was already occurring alongside a large body of pilots and exploration. Again, definitions of “deployment” differ and operators vary enormously in size, geography and digital maturity.

Autonomous networking is less mature. TM Forum's model ranges from Level 0 manual operation to Level 5 full autonomy. Its current assessment is that most operators remain around Levels 2 and 3, although individual operators have now validated Level 4 capability in specific network domains. By June 2026 TM Forum reported 40 Level-4-related validations involving 17 companies. This is meaningful evidence of operational progress, but a validation in a defined domain is not equivalent to an entire operator achieving autonomous operation.

Geographical development is not uniform. Operational evidence is particularly visible in China and East Asia, Europe, the Gulf, North America and other highly digitised markets, but meaningful activity also exists more broadly. TM Forum's international work covers operators across regions and its surveys indicate substantially different autonomy strategies and maturity by geography. Infrastructure age, capital resources, spectrum arrangements, skills, vendor mix, cloud adoption and data architecture all influence how quickly AI can be incorporated.

The appropriate Status is therefore not that telecommunications has become an autonomous industry. It is that **AI is already established across numerous telecommunications activities, generative AI is scaling rapidly, and high-autonomy networking has begun to reach production in particular domains while remaining far from universal.**

5. What appears to be driving adoption?

One major driver is the economic need to operate increasingly complex networks without increasing human workload at the same rate. Automation can allow engineers to focus on difficult exceptions rather than repeatedly handling routine diagnosis, configuration or optimisation.

Network quality is equally important. Operators compete partly on coverage, reliability, speed and customer experience. AI that detects deterioration earlier, allocates scarce radio resources more efficiently or reduces restoration time can therefore affect both operating cost and service quality.

The transition towards cloud-based and software-defined telecommunications infrastructure is another enabling factor. Network functions that once depended primarily on specialised physical equipment are increasingly controlled through software and APIs. That makes more parts of a network observable and potentially adjustable by intelligent automation.

Large quantities of operational data provide another driver. Telecommunications companies generate extensive real-time information about network performance, traffic and service

behaviour. Where that information can be governed and integrated effectively, it provides unusually rich material for machine learning.

Competitive pressure around generative and agentic AI is accelerating experimentation. GSMA's 2026 assessment describes AI as a board-level telecommunications priority, while operators are increasingly examining AI not only for cost savings but also for differentiated enterprise services, network APIs, managed AI services and AI-ready connectivity.

Industry standardisation is also reducing some barriers. TM Forum, GSMA, NGMN, ITU and other bodies are developing common definitions, architectures, evaluation methods and operating frameworks. Their activity does not establish deployment, but it indicates that AI integration is becoming part of mainstream telecommunications engineering rather than remaining a collection of isolated experiments.

6. What are the present limitations?

The first major limitation is **data fragmentation**. Telecommunications operators frequently have hundreds of operational and business systems created at different times and supplied by different vendors. Data definitions, interfaces and quality may differ across radio, transport, core, cloud, billing and customer domains. AI cannot easily optimise an end-to-end service when it lacks reliable visibility across those components.

A second limitation is reliability. Telecommunications infrastructure has very little tolerance for damaging automated decisions. A customer-service hallucination can inconvenience an individual customer; an incorrect network-control action can potentially affect thousands or millions of connections. Greater autonomy therefore requires dependable models, constrained action spaces, testing, rollback mechanisms and governance.

Generative AI adds familiar problems of fabricated or incorrect output. In engineering environments, an apparently plausible but wrong diagnosis or configuration is particularly problematic. Agentic systems introduce the further issue that erroneous output may lead to an action rather than merely a recommendation.

Legacy technology and multi-vendor integration are substantial practical constraints. Operators cannot normally replace entire networks merely to accommodate AI. New systems must coexist with older equipment, established OSS/BSS platforms, regulatory controls and long-lived physical infrastructure.

Skills and organisational design also matter. Autonomous operation changes the role of network engineers from direct execution towards supervising automated systems, setting intent and constraints, handling exceptions and validating performance. NGMN's 2026 framework explicitly treats technology, people, skills and organisational culture as interdependent elements in movement towards agentic network operations.

Security creates a two-sided challenge. AI can improve detection of malicious traffic, but AI systems themselves introduce new software dependencies, data-access paths and attack surfaces. Automated systems with permission to change network configurations require particularly strong controls.

Energy effects are similarly mixed. AI can reduce energy consumption through better network optimisation while training and running AI models, expanding edge computing and supporting AI-intensive traffic can increase electricity requirements. The net effect therefore depends on the particular architecture and workload rather than AI inherently reducing telecommunications energy consumption.

Finally, evidence of economic return remains uneven. Operator case studies frequently report significant improvements, but results depend on baseline processes, network architecture and implementation scope. Provider or operator announcements demonstrate that a system exists; they do not establish that the same performance would occur across different networks.

7. What you should be aware of

“AI in telecommunications” encompasses technologies at very different stages of maturity. A machine-learning system predicting traffic, an AI chatbot answering billing questions, a closed-loop radio optimiser and an experimental AI-RAN architecture for 6G should not be treated as equivalent forms of adoption.

The term **autonomous network** also needs careful interpretation. Autonomy is normally assessed for defined operational scenarios or domains. An operator may reach a high autonomy level in fault management or radio optimisation while people continue to manage many other parts of its network. Statements that an operator has achieved “Level 4” therefore require examination of what precisely was assessed.

Similar caution applies to “AI-native”. International telecommunications bodies were still working during 2026 on concepts, definitions, use cases and proofs of concept for AI-native telecommunications networks. The term increasingly describes an architectural direction, but its use in announcements does not itself establish that a network has been rebuilt around mature AI-based operation.

Human involvement remains significant. Even highly automated telecommunications systems operate within engineering policies, predefined objectives, safety limits and accountability structures established by people. Operators also retain responsibility for service reliability, cybersecurity, emergency communications and regulatory obligations.

Customer-facing AI introduces additional questions around privacy, transparency and accuracy because telecommunications providers possess substantial information about customers and communications activity. Australia's ACMA, for example, recorded consumer concerns during 2026 about disclosure of AI use and the collection and storage of information.

AI is also changing the threat environment against which telecommunications companies operate. Generative systems can support fraud, impersonation and sophisticated scam activity at scale while telecommunications companies themselves increasingly deploy AI to identify that activity. AI is consequently both part of the security problem and part of the industry's defensive response.

Geographical differences remain substantial. Operators with modern cloud-based networks, extensive digital data and strong investment capacity can generally implement AI more readily

than providers operating older infrastructure or serving markets where capital, skills, connectivity or computing resources are constrained. Evidence from prominent global operators should therefore not be treated as representative of every telecommunications provider.

Finally, telecommunications has become central to a second AI relationship: **AI for networks and networks for AI**. Operators are using AI to manage connectivity, while growing AI workloads are simultaneously changing the traffic, computing and infrastructure that telecommunications networks must support. The latter is strategically important but should not be confused with evidence that AI itself is operating the telecommunications network.

Status at 30 August 2026: AI is materially embedded in telecommunications. Network analytics and optimisation, customer automation, fraud detection and operational decision support have achieved meaningful real-world scale. Generative AI is extending those applications, and a smaller but growing set of deployments now allows AI to carry out multi-step or closed-loop network actions. Fully autonomous or comprehensively AI-native telecommunications networks, however, are not yet the normal industry condition. Current deployment remains a mixture of established AI-assisted operations, expanding generative-AI use, domain-specific high autonomy and substantial continuing experimentation.

Evidence Notes

1. **GSMA — *Charting AI Monetisation: The Telco Landscape Report 2026*, January 2026.** Provides industry-level assessment that operators are already embedding AI while substantial structural barriers to scaling and monetisation remain. GSMA represents the mobile industry and its evidence should be interpreted partly as industry evidence rather than wholly independent assessment.
2. **GSMA and TM Forum — “Scaling Telco AI from Ambition to Reality”, June 2026.** Useful evidence concerning the sector's multi-vendor complexity, fragmented data and difficulties moving AI from individual applications to operational scale.
3. **TM Forum — Autonomous Networks material and 2026 Level 4 reporting.** Establishes the industry's Levels 0–5 maturity framework and provides evidence that most operators remain below comprehensive Level 4 autonomy while an increasing number have validated Level 4 in specific scenarios. TM Forum is an industry organisation and individual validations should not be interpreted as independent evidence of entire-network autonomy.
4. **TM Forum — worldwide autonomous-network research, 2025–2026.** Survey and operator evidence demonstrates geographical variation and gradual rather than uniform progress toward high network autonomy. Survey responses represent participating organisations and should not be treated as a census of all global operators.
5. **NTT DOCOMO — network-maintenance AI agent, February 2026.** Primary operator evidence that a generative-AI maintenance system moved into commercial mobile-network operation in Japan. As company-reported evidence, it establishes deployment but not independent performance validation.
6. **NTT DOCOMO/NEC/NTT DOCOMO BUSINESS — commercial 5G core and agentic-AI network construction, March 2026.** Important evidence that agentic AI has participated in building infrastructure subsequently used commercially and of a reported approximately 80% reduction in construction time. Results are operator-reported and specific to this implementation.

7. **NTT DOCOMO — MantaRay AutoPilot deployment, June 2026.** Primary evidence of operational AI-based continuous mobile-network optimisation rather than merely laboratory testing.
 8. **Mobily/Ericsson — Hajj 2025 autonomous-network deployment, reported 2026.** Provides a strong real-world example of AI-driven closed-loop optimisation under unusually demanding live traffic conditions. Performance figures are reported by the participating companies and concern a particular event and network.
 9. **Telstra — AI-enabled self-healing demonstration, March 2026.** Useful evidence of agentic self-healing capability in a live telco-cloud environment. Telstra explicitly characterised the work as a proof of concept, so it has not been treated as routine production deployment.
 10. **SoftBank/Ericsson — AI-in-RAN commercial-network trial, August 2026.** Demonstrates measurable AI-based radio optimisation on an operating Japanese 5G network. It remained a validation exercise at the Status date and therefore supports capability rather than broad deployment.
 11. **Vodafone Group — FY2026 Results, May 2026.** Provides operator-reported evidence of scaled customer-service AI, network automation, energy optimisation and other AI use across a multinational telecommunications group. Useful for demonstrating deployment scale but not independent evidence of sector-wide outcomes.
 12. **Australian Communications and Media Authority — AI sector research, April 2026.** Independent regulatory evidence that Australian telecommunications providers are using AI for operational efficiency, resilient networks, predictive fault management and real-time scam detection, while governance concerns remain. Geographically limited to Australia.
 13. **NVIDIA — State of AI in Telecommunications: 2025 Trends.** Survey of more than 450 telecommunications professionals provides international evidence of widespread AI assessment and deployment and significant interest in network planning, operations, field operations and generative AI. Vendor sponsorship, respondent selection and broad definitions of “AI activity” mean percentages should be treated as indicators rather than global prevalence estimates.
 14. **GSMA — Energy Implications of AI for Mobile Operators, February 2026.** Principal evidence for the mixed energy consequences of telecommunications AI: network optimisation can reduce consumption while AI computing and new workloads increase demand.
 15. **NGMN — Cloud-Native Next Chapter: Agentic AI-Based Operating Models, March 2026.** Operator-led evidence that movement towards agentic operations requires changes in technology, processes, skills and organisational culture and remains a staged maturity process.
 16. **ITU-T — Artificial Intelligence for Networks work programme and AI-native-network studies, current through August 2026.** Authoritative evidence that generative-AI assessment, AI-native telecommunications concepts, use cases and proofs of concept are active areas of international standardisation. The continuing work programme itself demonstrates that important aspects remain developmental rather than settled operational practice.
 17. **DOCOMO/Samsung and DOCOMO/VIAVI — AI-RAN and prospective 6G demonstrations, March–August 2026.** Useful evidence of advanced network-control capabilities while illustrating the need to distinguish technical validation from current routine deployment.
 18. **GSMA Foundry — Mobile Innovation Report 2026 and associated operator programmes.** Demonstrates substantial industry activity moving AI projects toward live operator environments across multiple countries, while the Foundry's innovation role means that much of the evidence concerns emerging rather than mature deployment.
-

16. Security & Emergency Services

Status as at 30 August 2026

Artificial intelligence is already being applied across a significant range of civilian security, safety and emergency-service activities. Its strongest present uses are generally **analytical and assistive**: interpreting video and sensor feeds, detecting fires and other hazards, analysing emergency calls and operational data, forecasting or identifying disaster risks, prioritising inspections, assisting search and rescue, supporting resource allocation and extracting useful information from large volumes of incident data.

The maturity of these applications differs considerably. AI-enhanced video analytics are already embedded in parts of organisational physical security. Some emergency services have moved AI-supported detection and operational analytics into real-world service: New South Wales reported approximately 4,000 fire detections by high-definition camera systems over a year, while Fire and Rescue NSW documented an operational search in June 2026 in which AI software on a drone helped identify two missing hikers. Singapore's Civil Defence Force is using predictive AI to target fire-safety inspections and data-driven systems to position emergency resources.

Disaster management presents a more mixed picture. AI is increasingly capable of combining satellite, meteorological, geographic, sensor and other information to improve hazard detection, forecasting and situational awareness, and a July 2026 UN-led report documented 27 case studies across the early-warning value chain. Yet European scientific advisers concluded in late 2025 that AI use in emergency and crisis management remained limited and that systems continued to struggle with interpretation of complex situations.

The present Status is therefore not one of autonomous AI emergency management or the removal of human responders and security professionals. It is predominantly one of **AI being inserted into particular detection, prediction, monitoring, communication and coordination tasks within human-led systems**.

1. Where is AI currently being applied?

One of the most established areas is **physical-security monitoring**. Modern camera systems increasingly incorporate machine-learning analytics capable of detecting people, vehicles, objects, movement and selected events rather than simply recording video for later review. Applications include intrusion and perimeter monitoring, unusual-event alerts, vehicle recognition, crowd or occupancy monitoring, smoke or fire detection and rapid searching of recorded footage. Industry evidence describes video analytics as one of the areas in which AI has already achieved substantial physical-security adoption.

Related systems extend into **workplace and public safety**. Cameras, wearable devices, industrial sensors and other smart monitoring tools can identify unsafe conditions or behaviours, detect falls, monitor equipment or environmental conditions and warn when predefined risk conditions occur. The International Labour Organization's 2025 global review

describes AI-enabled monitoring, advanced robotics and other digital systems as increasingly relevant to occupational safety, including reducing exposure to hazardous work.

In **fire and rescue**, AI is being used for earlier detection, risk assessment and operational intelligence. New South Wales uses AI-powered camera systems alongside thermal drones and satellite connectivity for bushfire management; government reporting stated that cameras detected nearly 4,000 fires in the preceding year and in some instances identified fires before emergency calls were received. Singapore's Civil Defence Force introduced its Targeted On-site Inspection Tool in April 2025, using predictive AI and historical enforcement and fire data to identify premises more likely to have fire-safety violations.

AI is also entering **search, rescue and dangerous-environment operations**. In June 2026 an AI-equipped Fire and Rescue NSW drone searched a defined area of the Snowy Mountains; its software identified two missing hikers in vegetation and thermal imaging was then used to confirm the detection. This is a useful example of actual field deployment, rather than a laboratory demonstration, but it remains a specialised human-operated system rather than autonomous rescue.

Within **emergency communications and dispatch**, AI and machine learning can classify or transcribe calls, distinguish some emergency from non-emergency contacts, extract location or incident information and help allocate resources. The US National 911 Program has documented Arlington County's use of AI and machine learning for diverting non-emergency calls, while Singapore reported by 2026 that SCDF was using AI-supported systems to analyse emergency calls more rapidly and accurately.

A closely related application is **resource planning and operational coordination**. Historical incident patterns, current demand, location and time data can be analysed to anticipate where ambulances or other resources may be required. Singapore's Dynamic Resource Optimisation platform analyses historical and real-time emergency-medical incidents to identify potential hotspots and assist deployment.

In **disaster risk reduction and early warning**, AI applications span risk mapping, hazard forecasting and monitoring, analysing Earth-observation data, estimating damage, interpreting social-media information and helping determine how warnings should be communicated. The July 2026 UNDRR/WMO/ITU/IFRC work on multi-hazard early warning systems describes applications across risk knowledge, forecasting, warning dissemination and preparedness.

Brazil illustrates developing capability outside the most highly digitised emergency systems. Research led by Cemaden, INPE and Serpro has applied machine learning to flash-flood forecasting while explicitly modelling uncertainty in hydrometeorological data. The work tested forecasts from 30 minutes to six hours ahead using more than 73,000 records, although it remained research intended to improve operational warning systems rather than evidence of completed national deployment.

Finally, generative AI is beginning to appear in **information and administrative support**: querying operational manuals, preparing reports, summarising material and assisting personnel in finding procedures or organisational knowledge. Singapore's SCDF AI Assistant, for example, provides staff with natural-language access to official internal material. These applications are operationally different from allowing a generative model to make emergency decisions.

2. What problems or needs is AI being used to address?

The common underlying problem is **too much information arriving too quickly for continuous human examination**. Security cameras, smoke and heat sensors, access-control systems, weather feeds, satellite imagery, emergency calls and field reports can produce continuous data streams. AI is attractive where it can identify a potentially important signal within those streams and direct human attention towards it.

This is particularly important for detection. A person cannot continuously inspect thousands of camera views or satellite images with equal attention. Automated visual analytics can screen large amounts of material and flag defined events. In physical security this helps address surveillance fatigue; in wildfire monitoring it can reduce the interval between ignition becoming visible and responders being notified.

Emergency services face a second pressure: **time-critical decisions under constrained resources**. Dispatch centres must interpret calls and determine urgency; ambulance and fire services need to decide where resources should be positioned; rescue teams must search large geographical areas; and incident commanders need changing situational information. AI is being applied to shorten parts of these processes rather than to assume overall command.

Disaster management adds the problem of **anticipating low-frequency but potentially severe events from incomplete and heterogeneous data**. Earth observation, weather models, hydrological sensors, maps, infrastructure information and vulnerability data may all be relevant. European Commission analysis notes that AI can process more Earth-observation data, assist damage estimation and extract useful information from social media, potentially increasing preparation time.

There is also an important **capacity and geographical-access problem**. Many communities lack dense monitoring networks, specialist forecasters, reliable communications or large emergency-response workforces. AI may help make greater analytical use of limited observations, but the same places are often those least able to provide the data, connectivity and computing infrastructure on which sophisticated systems depend.

Organisational security has a comparable problem at smaller scale. Security teams must monitor premises, investigate alerts, control access and respond to incidents while avoiding overwhelming staff with false alarms. The attraction of AI is therefore often less about replacing guards or operators than filtering routine information so that people can concentrate on exceptions requiring judgement.

Workplace safety presents another need: **reducing human exposure to dangerous environments**. The ILO identifies automation, robotics and smart monitoring as mechanisms capable of removing workers from some hazardous activities or identifying unsafe conditions earlier. That benefit coexists with new risks created by monitoring, automation and algorithmic management.

3. What can AI currently do?

Current AI can perform a number of specialised security and emergency tasks convincingly.

Computer-vision systems can recognise defined objects and patterns, distinguish people or vehicles from background movement, search footage, detect smoke or other visual anomalies and generate alerts. Thermal and visual AI can also support aerial searches. The NSW hikers incident demonstrates that such a system can contribute directly to finding people in a real rescue operation.

Predictive models can calculate **risk scores and probabilities** from historical and current information. Singapore's fire-safety system uses previous inspection and incident records to estimate which premises have a greater propensity for violations. Models can similarly identify possible incident hotspots or estimate how hazards such as floods and fires may develop.

AI can process **geospatial and Earth-observation information** faster than would generally be practical manually. It can classify satellite imagery, infer exposure or damage, combine multiple environmental variables and contribute to weather and hazard forecasts. Social-media or textual information can also be filtered and geolocated where sufficient data and appropriate permissions exist.

Language and speech systems can transcribe, classify and summarise communications and retrieve relevant information from large document stores. These capabilities make them potentially useful in dispatch, incident documentation and staff assistance.

What current AI generally cannot do reliably is comprehend the entire evolving context of a major emergency in the way an experienced multidisciplinary command team can. The European Commission's Scientific Advice Mechanism concluded that AI can improve aspects of emergency and crisis management but still struggles to interpret complex situations, while operational application remains limited.

Technical capability should therefore be distinguished from **operational authority**. Detecting an image pattern does not establish that an AI system can decide whether to evacuate a city. Predicting a fire-risk area does not mean it can determine all firefighting priorities. Summarising an emergency call does not demonstrate that it can safely replace a trained dispatcher across unusual, ambiguous or emotionally complex calls.

4. How established is adoption?

Adoption is best characterised as **uneven but materially real**.

AI-enabled video analytics are among the more established uses. ASIS International's 2025 survey found that 57% of roughly 445 responding security professionals reported using at least one of a broad set of AI tools in security. Among AI users, 46% described it as helpful and another 9% said it exceeded expectations; 20% considered it too early to judge. The survey included a wide definition of AI, was conducted primarily among ASIS members and customers and should therefore not be interpreted as a global population estimate.

Within that research, reported operational effects included automated incident response and expanded perimeter-patrolling capability. Yet relatively few respondents reported overall staffing reductions, reinforcing the picture of AI augmenting rather than simply eliminating security work.

Fire detection and emergency-service analytics have moved beyond technical demonstration in selected jurisdictions. New South Wales provides evidence of actual bushfire-camera operations, while Singapore provides evidence of implemented inspection-risk modelling and operational emergency-resource analytics. These cases establish deployment, not universal adoption by fire and rescue services.

Disaster early-warning AI is less uniformly mature. The July 2026 international report contained 27 real-world case studies, demonstrating that application is no longer purely hypothetical. But international organisations continue to describe AI as an enabling component of early-warning systems rather than a self-contained solution, and operational readiness differs significantly by country, hazard and infrastructure.

Some applications remain clearly developmental. Brazil's flash-flood work, for example, represents serious applied research designed to strengthen operational alerts but should not be presented as an already deployed nationwide AI warning system. Similarly, autonomous or coordinated firefighting robots remain substantially more developmental than ordinary sensors, camera analytics and human-operated drones.

The overall position is therefore a continuum: **established use in selected monitoring and analytics; growing deployment in detection, dispatch support and resource planning; substantial operational experimentation in disaster intelligence; and more limited maturity for autonomous emergency decision-making and robotics.**

5. What appears to be driving adoption?

The most immediate driver is the value of **speed**. Earlier recognition of a fire, intrusion, medical-demand hotspot or deteriorating hazard can create additional time for human intervention. In emergencies, even comparatively small reductions in detection or assessment time may matter.

A second driver is the rapid growth of **digital data and connected sensors**. Cameras, drones, satellites, environmental sensors, digital dispatch platforms and mobile communications provide increasingly machine-readable information. AI makes more of that data analysable at operational speed.

A third is **workforce and attention scarcity**. Security and emergency organisations cannot expand human monitoring in proportion to every new sensor or data source. Systems able to filter routine information and direct personnel towards higher-priority cases therefore have an evident operational attraction.

Technological improvement is also making applications more practical. Better computer vision, edge computing, thermal cameras, drone platforms, communications networks and machine-learning models allow analytics to occur closer to where incidents are happening rather than only in central computing environments.

Disaster risk is providing a particularly strong institutional driver. The international Early Warnings for All initiative seeks broader global warning coverage, and UNDRR, WMO, ITU and IFRC are explicitly examining AI as one means of closing existing gaps. Their 2026

assessment nevertheless emphasises infrastructure, partnerships and human expertise rather than treating AI alone as the solution.

Finally, organisational security is increasingly expected to demonstrate broader operational value. Cameras and access systems originally installed for security can also produce data relevant to occupancy, traffic, safety and facility operations. This encourages investment in analytics that can serve several organisational purposes.

6. What are the present limitations?

Reliability is the central technical limitation. Emergency and safety systems operate in environments containing poor visibility, smoke, noise, unusual behaviour, damaged infrastructure, incomplete information and rapidly changing conditions. A model that performs well on normal data may fail precisely when circumstances become abnormal.

False positives and false negatives both matter. Excessive alerts can overwhelm operators and reduce trust; missed detections can have serious consequences. Performance also varies with camera placement, lighting, weather, sensor quality, geography and the data on which a model was trained.

Prediction introduces a further difficulty: uncertainty cannot be removed simply by using AI. Brazil's flash-flood research is noteworthy partly because it explicitly models uncertainty and found trade-offs between predictive accuracy, robustness and computational cost. No single tested approach was optimal across all circumstances.

Emergency management is also deeply contextual. An incident may involve infrastructure failure, human behaviour, weather, conflicting reports, vulnerable populations and changing priorities simultaneously. The European scientific review's finding that AI struggles with complex situational interpretation is therefore particularly important.

Infrastructure is another major constraint. AI-assisted warning and coordination can depend on electricity, communications networks, cloud or edge computing and current sensor feeds—exactly the systems that disasters may disrupt. ITU's emergency-telecommunications work continues to emphasise restoring connectivity because communications remain fundamental to coordination.

Data quality and representativeness also vary geographically. Areas with sparse sensors, incomplete building records or limited historical incident data may be less well served by data-hungry models. This creates a risk that regions with strong digital infrastructure benefit first while communities facing high disaster vulnerability remain comparatively underserved.

Privacy and civil-liberties constraints are especially important for physical security. AI-enabled cameras can analyse identifiable people, movement and potentially biometric characteristics. The EU AI Act treats several biometric uses as high-risk and restricts certain applications; other jurisdictions apply different surveillance, privacy and employment frameworks.

Workplace safety illustrates the same tension. Automated monitoring may identify hazards but can also become intensive worker surveillance or algorithmic management. The ILO therefore describes digitalisation as creating both safety opportunities and new occupational risks.

Cost and integration remain practical constraints. Continuous video analysis can demand significant computing, storage, network and power resources. Existing dispatch, radio, access-control, sensor and incident-management systems may be old or fragmented. A technically effective pilot may therefore be difficult or expensive to scale. Industry analysis in 2026 specifically warned that video-heavy physical-AI deployments can encounter substantial infrastructure and data-transfer costs.

7. What you should be aware of

“AI in security and emergency services” covers applications with very different consequences. An algorithm searching archived camera footage, software suggesting where ambulances might be positioned, a model forecasting flood risk and a system influencing an evacuation decision should not be treated as equivalent merely because all use AI.

Human involvement remains extensive. Even the more impressive current deployments usually provide **information, prioritisation or detection to people who retain operational responsibility**. The NSW search-and-rescue example involved AI detecting possible people, thermal imaging confirming them, a trained drone operator controlling the platform and ground responders completing the rescue.

Evidence of effectiveness should likewise be interpreted carefully. A provider can demonstrate that a camera recognises a person or vehicle, but that does not establish lower incident rates. A forecasting model can improve a technical metric without demonstrating that communities receive or act upon warnings. A successful rescue or early fire detection establishes that an application can be operationally useful, not that the same result will occur in all environments.

Security uses also raise distinctive questions of proportionality and rights. AI may make surveillance much more searchable and scalable than conventional CCTV. The consequences therefore depend not only on recognition accuracy but on what information is collected, how long it is retained, who can access it and what decisions follow. Regulatory treatment differs substantially across countries.

Geographical differences are particularly important. High-income, highly connected jurisdictions provide many of the clearest documented operational examples, while disaster risk is often greatest in places with weaker communications, mapping, sensor and institutional capacity. The global Status should therefore not be inferred from leading emergency-service deployments alone.

There is also a distinction between **digitalisation and AI**. Modern emergency communications, drones, sensor networks, geographic information systems and cell-broadcast alerts may be important to AI-supported operations without themselves being AI. Singapore's SG Alert and the international expansion of emergency telecommunications, for example, strengthen the infrastructure around emergency response but should not automatically be counted as AI adoption.

Status at 30 August 2026: AI has become a meaningful component of Security & Emergency Services, particularly for monitoring, video analysis, risk identification, hazard detection, operational analytics and selected search-and-rescue and dispatch-support tasks. Real deployments exist across organisational security, workplace safety, fire and rescue and disaster-management settings. Adoption, however, remains uneven. The strongest present role is as a specialised sensing, analytical and decision-support layer inside human-led security and emergency systems; evidence does not support describing autonomous AI crisis management or autonomous emergency response as established sector practice.

Evidence Notes

1. **UNDRR, WMO, ITU and IFRC — *Leveraging AI to Enhance Multi-Hazard Early Warning Systems*, July 2026.** Principal international evidence concerning AI across risk understanding, forecasting and monitoring, warning dissemination and preparedness. Includes 27 case studies demonstrating real applications but deliberately treats AI as one component of wider warning systems rather than a standalone substitute for institutional capacity.
2. **European Commission Scientific Advice Mechanism/SAPEA — *Artificial Intelligence in Emergency and Crisis Management*, December 2025.** Important independent assessment of current capability and maturity. Finds useful applications but limited operational use and continuing difficulty interpreting complex crisis situations. European focus and strong emphasis on natural hazards limit direct generalisation to all emergency services.
3. **NSW Government — bushfire technology announcement, 23 March 2026.** Primary evidence of operational use of AI-powered camera systems in a large civilian fire-response environment, reporting nearly 4,000 fire detections during the preceding year. Government operational reporting establishes deployment but is not an independent comparative evaluation of effectiveness.
4. **Fire and Rescue NSW — AI-assisted missing-hiker rescue, 25 June 2026.** Detailed primary evidence of an actual field operation in which AI-equipped drone software contributed to detecting two missing people. Highly useful for distinguishing demonstrated real-world capability from general claims about autonomous search and rescue; it is a single operational example.
5. **Singapore Civil Defence Force/Ministry of Home Affairs — *AI and data-analytics deployments*, 2025–26.** Evidence concerning the Targeted On-site Inspection Tool, emergency-resource optimisation and associated operational systems. Demonstrates implementation within a highly digitised national emergency service, but Singapore's infrastructure and institutional setting should not be assumed representative globally.
6. **Singapore Ministry of Digital Development and Information — *Career Forward* remarks, 2026.** Government evidence that AI-supported emergency-call analysis was being used by SCDF. Useful confirmation of application but provides limited independent performance data.
7. **Brazil MCTI/Cemaden/INPE/Serpro — *flash-flood machine-learning research*, May 2026.** Provides a Latin American example of sophisticated AI work directed towards operational early warning and illustrates uncertainty, robustness and computational trade-offs. It remains research intended to improve warning systems and should not be treated as proof of completed routine deployment.
8. **International Labour Organization — *Revolutionizing Health and Safety: The Role of AI and Digitalization at Work*, April 2025.** Global evidence concerning AI, robotics, smart monitoring and other digital technologies in occupational safety and health. Particularly valuable for showing that safety benefits coexist with worker-monitoring and new technological risks. Its scope is wider than AI alone.

9. **ASIS International — *Understanding the Evolving Role of Security: 2025 Security Trends Research*.** Provides rare quantitative evidence concerning AI use among security professionals: 57% of approximately 445 respondents to the AI questions reported some use. The definition of AI was broad and respondents were primarily ASIS members and customers, so the percentage should not be interpreted as a global physical-security adoption rate.
 10. **Security Industry Association — video-surveillance and AI material, 2025–26.** Useful industry evidence that AI video analytics have become materially embedded in physical-security technology and that AI can automate detection and search. As industry material, it is stronger evidence of available capability and deployment patterns than of independent sector-wide outcomes.
 11. **US National 911 Program — AI for non-emergency call diversion, February 2025, and NG911 material.** Evidence that machine learning and AI are entering defined emergency-communications workflows. The documented example is specific rather than evidence of widespread US or global 911 adoption.
 12. **European Commission Joint Research Centre/Knowledge4Policy — AI in disaster-risk management, September 2025, updated August 2026.** Evidence concerning Earth-observation analysis, damage estimation, increased anticipation time and analysis of crisis-related social-media information. Supports technical capability rather than uniform operational deployment.
 13. **International Telecommunication Union — emergency telecommunications and Early Warnings for All material, current through 2026.** Establishes the continuing importance of resilient communications, emergency connectivity and infrastructure to digital and AI-supported emergency management. Particularly useful for understanding why AI capability cannot be separated from functioning communications systems.
 14. **European Union — Artificial Intelligence Act, Regulation (EU) 2024/1689.** Principal European regulatory evidence relevant to biometric identification, surveillance, workplace applications and safety components of critical infrastructure. Important for demonstrating regulatory constraints and jurisdictional differences; it does not establish the prevalence of the affected technologies.
-

The Status Across Industries & Sectors

Viewed together, the sixteen sectors reveal a broader current position that is difficult to see from individual applications or technology announcements alone.

The first and most important finding is straightforward:

AI has moved materially into economic and societal activity.

This is no longer principally a landscape of laboratory demonstrations or isolated experimentation.

AI systems are analysing medical images, detecting financial fraud, forecasting demand, inspecting industrial products, recommending goods and media, assisting travellers, monitoring construction, forecasting renewable generation, identifying agricultural pests, supporting public administration, analysing scientific data, assisting professional work, optimising telecommunications networks and helping emergency services interpret operational information.

The evidence supporting those applications differs in strength and scale, but collectively it establishes that AI has entered real operating environments across a wide range of sectors.

At the same time, the second finding is equally important:

AI adoption is highly uneven.

The unevenness occurs between sectors, within sectors, between countries, between large and small organisations and between different applications of AI.

Healthcare contains established specialised applications alongside much earlier generative and decision-support uses. Its strongest current position remains specialised analytical and workflow technology within human-led systems.

Financial Services has comparatively mature prediction, classification, detection and risk-analysis applications, while more consequential generative applications remain more controlled.

Manufacturing has established analytical, optimisation, inspection and maintenance applications in increasingly digital factories, while more adaptable AI-enabled robotics remains less mature.

Retail combines established recommendation, forecasting and optimisation with rapidly developing generative and agentic commerce.

Energy & Utilities demonstrates meaningful operational adoption in forecasting, asset analysis, inspection and optimisation, while more ambitious autonomous control remains developing.

Telecommunications has moved further into selected closed-loop and agentic operations than many sectors, but those examples still coexist with developmental AI-native network concepts rather than establishing autonomous networks as normal practice.

This pattern suggests that asking whether a sector “uses AI” has become increasingly uninformative.

A more useful question is **what kind of AI is being used, for what task, at what level of operational integration, with what degree of authority, and how widely?**

Specialised AI remains the most established layer

Across the sectors, some of the most mature applications are comparatively narrow.

They predict, classify, rank, detect, forecast, optimise or identify patterns within defined information environments.

That includes fraud detection, recommendation systems, demand forecasting, equipment-failure prediction, computer-vision inspection, anomaly detection, network optimisation, scientific classification, risk analysis and selected medical imaging.

These applications may be less visible in public discussion than generative AI, but in many sectors they represent the deeper operational layer of current AI adoption.

Their maturity also demonstrates an important principle: **AI does not have to resemble human conversation or general intelligence to have significant economic or societal consequences.**

A model making millions of recommendations, continuously analysing network conditions or screening large quantities of operational data may be deeply embedded even though users rarely encounter it directly.

Generative AI is broadening the application surface

Generative AI represents the most conspicuous expansion across the sixteen sectors.

Its importance lies partly in accessibility.

Many earlier AI systems required specialised datasets, technical teams and integration. Generative systems can be used directly through natural language and can operate on the documents, communications, images, code and other unstructured information already present in many forms of work.

This helps explain its rapid spread across education, professional services, government, media, healthcare administration, retail, hospitality, science and telecommunications.

The strongest present uses are often assistive: retrieving information, summarising, drafting, translating, generating alternatives, answering questions, supporting coding and organising knowledge.

In Media, Entertainment & Creative Industries, generative AI is expanding into text, imagery, audio, video and interactive production while established recommendation and analytical systems continue underneath it.

In Professional Services, it is being incorporated into research, information retrieval, document analysis, drafting, data analysis and knowledge work while qualified professionals continue to exercise judgement and retain responsibility.

The resulting development is significant, but it should not be confused with autonomous performance of entire professions or sectors.

AI is beginning to move from assisting towards acting

A newer development appears across several chapters.

AI systems are increasingly able not merely to produce information but to participate in multi-stage workflows, interact with systems and undertake bounded actions.

Agentic commerce can progress parts of a transaction. Telecommunications systems can undertake network optimisation actions. Some software can participate in professional workflows. Scientific systems can connect computational reasoning with automated experimentation.

This movement from **answering towards acting** is real.

But its maturity remains substantially below that of AI assistance overall.

The distinction becomes particularly important as consequences increase. Generating a draft and altering an operational system are not equivalent. Recommending a product and purchasing it are not equivalent. Suggesting a maintenance action and controlling infrastructure are not equivalent.

The more authority transferred to AI, the more questions of reliability, permissions, monitoring, reversibility, security and accountability become operational rather than theoretical.

Physical-world AI remains especially conditional

The sixteen sectors also show a divide between information environments and physical environments.

AI can perform increasingly capable tasks in factories, farms, transport systems, buildings, energy infrastructure and emergency operations. But physical environments contain variability that digital information tasks do not.

Weather changes. Equipment deteriorates. Objects vary. Roads and workplaces contain people. Farms differ by soil, crop and season. Emergencies are inherently unusual. Sensors provide incomplete information.

This means that impressive physical-world capability is frequently **conditional on the operating environment**.

A system that performs successfully within a defined route, factory process, crop, building, network or emergency scenario should not automatically be assumed capable of operating reliably outside those conditions.

The current Status is therefore one of increasing physical AI capability without evidence that general physical autonomy has become established across these sectors.

Digital foundations strongly influence adoption

Another recurring pattern is that AI frequently builds upon earlier digitalisation.

Data-rich platforms, connected equipment, digital records, sensors, cloud infrastructure, electronic transactions and established software systems create environments in which AI can be integrated more readily.

Retail demonstrates this through continuous product, transaction and behavioural data.

Manufacturing illustrates the complementary requirements of sensors, operational systems, integration and domain expertise.

Energy depends on instrumentation, forecasting data and digitally managed networks. Agriculture depends heavily on connectivity, sensors, machinery, data and production conditions.

This produces an important qualification when comparing sectors and countries:

AI adoption cannot be separated completely from the maturity of the underlying digital and physical system into which AI is being introduced.

A sector with extensive digital infrastructure may be able to adopt sophisticated AI while another remains occupied with more fundamental connectivity, digitisation or infrastructure requirements.

Scale has several meanings

The evidence repeatedly demonstrates that “scale” needs qualification.

A system can have large **potential reach** because it is available through a major platform.

It can have large **actual use** because millions of people interact with it.

It can have demonstrated **effectiveness** in a defined trial.

It can be deeply **institutionalised** within an organisation or national system.

Or it can be **prevalent across a sector**.

These are not equivalent.

Education illustrates the difference between widespread individual use and institutional adoption.

Retail demonstrates the difference between platform availability and merchant utilisation.

Government demonstrates the difference between geographical spread and depth of deployment.

Creative industries provide another striking example: large volumes of AI-generated music can coexist with a much smaller share of actual listening, demonstrating that generation volume, consumption and economic significance are different measures.

This distinction is essential when interpreting claims that AI has reached “scale”.

Human involvement remains central, but its form is changing

Across virtually all sixteen sectors, AI is operating within wider systems involving people.

That does not mean nothing important is changing.

AI can alter what people see, how quickly information is processed, which options are considered, which cases receive attention and how work is distributed.

A person may remain responsible for the final decision while relying substantially on AI-generated analysis.

Human involvement therefore should not be reduced to a binary question of whether a “human is in the loop”.

The more useful questions concern **what the AI does, what the person does, where judgement occurs, who possesses authority, and who remains accountable for the outcome.**

Those answers differ by sector.

Healthcare places particular weight on clinical judgement and patient safety. Finance operates within extensive regulatory and model-risk structures. Government must consider legality, contestability and public accountability. Professional Services retains professional responsibility. Science ultimately depends upon validation and reproducibility. Creative industries raise questions of authorship, consent and rights. Emergency services retain operational responsibility for decisions affecting safety.

The evidence therefore points towards changing **human–AI configurations**, rather than a simple transition from human activity to machine activity.

The value of automation is also sector-dependent

The sectors demonstrate that technological capability alone does not determine whether automation is desirable or adopted.

Hospitality provides a clear example. Human welcome, empathy and cultural knowledge can themselves form part of the service being purchased; an activity being technically automatable does not mean organisations or customers necessarily want it automated.

In government, legal and procedural legitimacy can constrain technically possible automation. In healthcare, clinical accountability and patient safety matter. In professional services, trust and professional responsibility remain important. In creative industries, cultural value and human authorship can matter independently of production efficiency.

AI adoption is therefore shaped not only by **what machines can do**, but by the economic, institutional, legal, professional and social character of the activity in which they operate.

AI can become both a solution and a new source of demand or risk

Several sectors reveal a two-sided relationship with AI.

Financial institutions use AI against fraud while confronting AI-enabled fraud and impersonation.

Creative industries use generative systems while also needing new approaches to provenance, authenticity and synthetic-media verification.

Energy systems use AI for forecasting and optimisation while simultaneously facing growing electricity demand from AI computing.

Telecommunications operators use AI to optimise networks while carrying rapidly increasing AI-related traffic and workloads.

This means AI adoption can itself create additional AI-related work, infrastructure requirements, controls and risks elsewhere in the same sector.

Geography remains fundamental to the Status

AI is unquestionably a global development.

The sixteen chapters contain meaningful activity from North America, Europe, China and wider Asia, Latin America, Africa, the Middle East and Oceania.

But the evidence does not support describing global adoption as uniform.

Infrastructure, capital, skills, regulation, language support, data availability, organisational scale and institutional capability differ greatly.

The most visible frontier deployments frequently come from organisations with unusually favourable conditions. Manufacturing evidence explicitly warns against treating advanced factories as representative of the sector average. Retail similarly shows why large digital platforms should not be treated as representative of smaller enterprises. Security & Emergency Services highlights the particular problem that places facing substantial disaster risk may have weaker communications, mapping, sensor and institutional capacity than jurisdictions producing many of the best-documented deployments.

The absence of comparable evidence from a country or region should therefore not be converted into a claim of absence of AI use.

But neither should documented frontier deployments be converted into an assumption of global diffusion.

Evidence remains stronger for activity than for ultimate outcomes

Across the sectors, it is generally easier to establish that AI is being developed, purchased, piloted or used than to establish its long-term consequences.

There is substantial evidence of applications and deployments.

There is less consistent independent evidence establishing sustained productivity gains, economic returns, learning outcomes, improved safety, better public services, employment effects or other sector-wide consequences.

This is partly because many applications are new.

It is also because AI is frequently introduced alongside other changes—new software, sensors, workflow redesign, cloud systems, automation, training and organisational transformation. Separating the effect of AI from the wider change can therefore be difficult. Manufacturing explicitly identifies this problem when interpreting productivity and quality gains.

Provider and organisational case studies remain valuable evidence that real deployment exists. They should not automatically be treated as independent evidence of general effectiveness.

The overall Status

Taken together, the sixteen sectors support a position that is neither one of marginal experimentation nor comprehensive AI transformation.

AI is now a meaningful component of economic and societal activity across a broad range of sectors.

Its most established role remains in specialised prediction, classification, detection, recommendation, forecasting, optimisation and information-processing tasks.

Generative AI is rapidly broadening that role into language, knowledge, communication, design, coding and content-intensive activities.

Agentic systems are beginning to move AI from assistance towards bounded action.

Physical AI and more autonomous operational systems are progressing, but remain more dependent on defined environments, infrastructure and human oversight.

The strongest present pattern is therefore not autonomous AI replacing industries or institutions.

It is AI becoming progressively embedded within them.

People, organisations and institutions continue to define objectives, provide data and context, establish permissions, handle exceptions, exercise judgement and remain responsible for many of the consequences that matter most.

The depth of AI integration will look very different in a hospital, school, bank, factory, farm, research laboratory, hotel, government agency, telecommunications network or emergency service.

Those differences are fundamental.

The Status as at 30 August 2026 is one of broad but uneven integration: established specialised AI, rapidly expanding generative AI, emerging agentic and autonomous capabilities, and continuing dependence on sector-specific data, infrastructure, expertise, governance and human accountability.

AIStatus — Continuing the Exploration

Artificial intelligence can be examined from several perspectives.

This publication considers AI through the **Industries & Sectors lens**: where and how AI is currently being applied across economic and societal sectors.

The wider AISF AIStatus series also examines the current AI landscape through two complementary lenses.

Business Functions considers where and how AI is currently being applied across the functions and activities of organisations.

Technology Landscape considers the significant AI technologies, system forms and enabling developments that currently constitute the applied AI landscape.

The three perspectives form the **AIStatus Tripod**:

Business Functions — Industries & Sectors — Technology Landscape

Each lens produces a standalone publication.

Readers do not need to read one AIStatus publication in order to understand another. The same development may legitimately appear through more than one lens because each asks a different question of the current landscape.

Together, the lenses provide complementary ways of discovering what is happening in artificial intelligence at a defined point in time.

AIStatus forms one part of the wider AISF educational architecture.

AIStatus helps readers DISCOVER what is happening.

AIGuides help readers NAVIGATE AI and the questions surrounding its use.

AISF Insight Books help readers UNDERSTAND important AI developments and their implications in greater depth.

Each serves a different primary reader purpose while contributing to the wider AISF objective of supporting clearer understanding of artificial intelligence.

Future Status Updates & User Submissions

Artificial intelligence continues to change rapidly.

New capabilities emerge. Existing technologies improve. Applications move from research to trials, from trials to operational deployment, or sometimes fail to progress as expected. Regulation, infrastructure, evidence and patterns of adoption also change.

For that reason, AIStatus is intentionally a **point-in-time publication**.

Future editions may reassess AI Across Industries & Sectors using a later reference date. Individual sectors may develop at different rates, and the sectors required to describe the landscape adequately may themselves change over time.

Earlier editions will remain records of the Status assessed at their stated reference dates.

AISF also welcomes information that may strengthen future Status assessments, including credible evidence of significant deployments, research, adoption patterns, limitations or developments from countries and sectors that may be less extensively represented in publicly available evidence.

Submissions do not guarantee inclusion. Material considered for future editions remains subject to AISF's editorial and evidential standards.

Acknowledgements

AI Sourced Facts acknowledges the work of the researchers, public institutions, international organisations, regulators, universities, professional and industry bodies, companies and other organisations whose publicly available information contributes to understanding the development and application of artificial intelligence.

Particular recognition is due to those organisations that publish not only successful applications and technological advances, but also limitations, evaluation results, implementation experience and evidence that helps distinguish technical possibility from real-world deployment.

AISF also acknowledges the continuing contribution of readers and professionals who identify emerging developments, question assumptions and help broaden understanding of how artificial intelligence is being applied across different parts of the world.

AISF Universal Disclaimer

This publication is provided for general educational and informational purposes only.

It is intended to support understanding of artificial intelligence and related developments and should not be interpreted as legal, regulatory, compliance, financial, investment, medical, technical, operational, strategic or other professional advice.

The publication does not recommend or endorse any particular artificial intelligence system, provider, product, platform, organisation, technology, approach or course of action.

Artificial intelligence technologies, services, capabilities, terminology, practices, laws, regulations and patterns of use can change rapidly. Their operation, suitability and adoption may differ between providers, models, versions, organisations, industries, jurisdictions and circumstances. Information and observations contained in this publication may therefore change over time and may not reflect subsequent developments.

Examples, studies, organisations, technologies and deployments are included to assist understanding of broader developments. Their inclusion does not constitute endorsement, recommendation or certification by AI Sourced Facts (AISF), and they should not be assumed to represent every organisation, jurisdiction, technology, sector or circumstance.

Evidence concerning artificial intelligence may be incomplete or subject to different definitions, samples, methodologies, assumptions and measures. Technical capability, availability, deployment, adoption and demonstrated effectiveness should not be assumed to be equivalent unless the underlying evidence supports that conclusion.

AISF seeks to maintain neutrality, clarity and appropriate recognition of uncertainty. Reasonable efforts are made to provide accurate and appropriately qualified information, but no representation or warranty is made as to its completeness, continuing accuracy or suitability for any specific purpose. No publication can eliminate all possibility of error, omission, subsequent change or differing interpretation.

Where decisions carry material consequences, responsibility for interpreting relevant information and making those decisions remains with the individual or organisation concerned and, where appropriate, their relevant qualified professionals or authorities.

Readers remain responsible for the use they make of the information contained in this publication.

About AI Sourced Facts (AISF) Pte. Ltd.

AISF is a Singapore-headquartered institution dedicated to structured reasoning, responsible AI navigation, and governance-informed adoption of artificial intelligence systems.

AISF operates with a capability-first, vendor-neutral posture. Its publications do not rank platforms, endorse providers, or promote specific technologies. Instead, AISF develops structured frameworks that help individuals, professionals, and institutions reason clearly before integrating AI into operational, strategic, or educational environments.

AISF's work spans whitepapers, applied insight books, education instruments, governance architectures, and structured research initiatives. These outputs are informed by cross-system AI research methodologies and reflect globally observed usage patterns at the time of publication. Human accountability remains central across all AISF frameworks.

AISF does not provide regulatory, legal, financial, investment, or compliance advice. Its publications are designed to support structured thinking, proportionate governance, and disciplined evaluation of AI capabilities prior to deployment or reliance.

As artificial intelligence systems continue to evolve, AISF's focus remains constant: clarity before integration, governance proportionate to capability, and long-term institutional resilience in the age of AI.

AI Sourced Facts (AISF)

Truth, Clarity & Guidance in the Age of AI.

Publication & Governance Information

Title: AI Across Industries & Sectors

Publication: AIStatus 2

Subtitle: Where and How AI Is Being Applied Across Economic and Societal Sectors

Series: AISF AIStatus

Edition: Version 1.0 / 2026 Edition

Status Reference Date: 30 August 2026

Publication Date: September 2026

Published by: AI Sourced Facts (AISF) Pte. Ltd.

Website: www.aisourcedfacts.com

© 2026 AI Sourced Facts (AISF). All rights reserved.

No part of this publication may be reproduced, distributed or transmitted in any form or by any means except as permitted by applicable law or with the prior permission of the copyright holder.

AISF maintains editorial oversight of this publication.

AIStatus is a point-in-time publication. Future editions may reassess the subject at a later reference date where material developments affect the Status of artificial intelligence across Industries & Sectors.

Any future edition remains subject to AISF editorial standards and governance.

Back Page

AI Across Industries & Sectors

Artificial intelligence is moving into industries and societal sectors around the world.

But where is it actually being applied — and how established is that use?

AI Across Industries & Sectors examines the Status of AI across sixteen areas of economic and societal activity, from Healthcare & Life Sciences, Education and Financial Services to Manufacturing, Agriculture, Government, Science, Telecommunications and Emergency Services.

Across each sector, the same seven questions examine current applications, the needs AI is addressing, demonstrated capability, adoption, drivers, limitations and what readers should be aware of.

The publication distinguishes technological capability from real-world deployment, adoption from experimentation, and current evidence from expectations about the future.

It provides a global, neutral and evidence-informed snapshot of the position:

Status as at 30 August 2026

AI Sourced Facts (AISF)

Truth, Clarity & Guidance in the Age of AI.

Website: www.aisourcedfacts.com

Quick Access: AISF.tech
