2978-5669

AI DECISIONS

Human–AI Decision-Making Systems

Bridging Algorithms, Humans and Governance in Decision-Making

Volume 2, Issue 1 — March 2026

Published by AI Decisions Ltd
London, United Kingdom
journal.aidecisions.ai

AI DECISIONS HUMAN–AI DECISION-MAKING SYSTEMS

VOLUME 2 — ISSUE 1
MARCH 2026
QUARTERLY OPEN ACCESS ACADEMIC JOURNAL

ISSN: 2978-5669
DOI PREFIX: 10.65114

Contact: journal@aidecisions.ai
Frequency: quarterly (March, June, September, December)

Published By AI Decisions Ltd, London, United Kingdom

OPEN ACCESS

INDEX

TABLE OF CONTENT

FROM THE EDITOR-IN-CHIEF

Ai DECISIONS

The conviction that founded this journal — that the decision, not the model, is what matters — feels more alive today than when we first wrote it. The technology has moved. The questions have deepened. And the stakes, for organisations and for the people within them, have grown considerably.

We stand at a moment when AI is no longer something organisations experiment with at the edges. It is moving into the centre — into the workflows, the judgments, the accountabilities that define how institutions actually function. That shift makes the human side of human–AI decision-making not less important, but more. Who decides, on what basis, with what oversight, and with what recourse: these questions are becoming urgent in ways they were not before.

The year ahead promises to be the most consequential yet for our field. We expect it to challenge comfortable assumptions, surface new tensions, and demand more rigorous thinking about what good decision-making looks like when AI is a genuine participant rather than a tool.

We are grateful to our authors, reviewers, and readers for the quality of engagement that makes this work worth doing. We look forward to the year ahead — with curiosity, and with appropriate caution.

Professor Adil Khan

Received: 10 February 2026	Revised: 20 February 2026	Accepted: 25 February 2026

DOI: https://doi.org/10.65114/aide.gmry6g42

Journal of Human–AI Decision-Making Systems

AI DECISIONS

Artificial intelligence in 2026: an analytical review of key trends

Yury Korolev

AI Decisions, London, UK

Correspondence

Yury Korolev, CEO, AI Decisions,
107B Cottenham Park Road,
London, SW20 0DS, UK
Email: yury@aidecisions.ai

Abstract

The years 2025–2026 mark a turning point in the development of artificial intelligence. Corporate AI investment reached a record $252.3 billion in 2024, and the share of organisations using AI rose from 55% in 2023 to 78% in 2024. Yet behind these figures lies a paradox: despite mass adoption, only 5–7% of companies generate significant value from AI, whilst 60–80% report no material impact on financial results. This article presents a systematic analysis of 15 authoritative industry reports published in 2025 – early 2026 by academic institutions, consulting firms, technology companies, analyst agencies, and venture funds. The analysis identifies 12 key conclusions and maps areas of high expert consensus (high expert consensus) as well as notable divergences. Agentic AI (AI agents capable of autonomously executing complex multi-step tasks) emerges as the dominant technological trend of 2026, already accounting for 17% of total AI value and projected to reach 29% by 2028. The article further examines AI maturity and the "Gen AI Paradox," infrastructure and compute trends, the economics of AI, regulation and governance, AI in science and medicine, and industry-specific developments in finance, healthcare, manufacturing, and retail. Twenty verifiable checkpoint questions for end-of-2026 assessment are proposed, together with quantitative thresholds for evaluating forecast accuracy. The overall direction can be summarised as "from experiments to transformation through agents."

Keywords

Artificial intelligence, agentic AI, AI adoption, AI maturity, generative AI, enterprise AI, hybrid cloud, AI regulation, AI in medicine, sovereign AI, workforce transformation, industry analysis.

1. Introduction

1.1 Why this analysis matters

The years 2025–2026 have marked a turning point in the development of artificial intelligence. As the Stanford AI Index Report 2025 observes, "AI is no longer just a story of what's possible — it's a story of what's happening now and how we are collectively shaping the future of humanity" [1, p. 2].

Corporate investment in AI reached a record $252.3 billion in 2024, an increase of 25.5% over the previous year [1, p. 247]. The share of organisations using AI rose from 55% in 2023 to 78% in 2024 [1, p. 3]. At the same time, according to McKinsey, 88% of organisations already use AI regularly in at least one function [2, p. 2].

Behind these impressive figures, however, lies a paradox: despite mass adoption, most companies do not see a material impact on their financial results. According to BCG, only 5% of companies generate significant value from AI, whilst 60% see no material return on investment [4, pp. 3–4]. McKinsey confirms that more than 80% of companies using generative AI report no significant impact on EBIT [3, p. 4].

This difference between adoption and value is the central theme of the present analysis. The year 2026 will be the year in which companies must transition from experimentation to the systematic extraction of value, and agentic AI will become the key instrument of that transition.

1.2 Methodology

Sources of analysis

This report is based on a systematic analysis of 15 authoritative industry documents published in 2025 – early 2026. Table 1 presents the full corpus, organised by source category, together with each document's thematic focus and data-collection methodology.

Table 1. Corpus of analysed sources: focus and methodology.

#	Source	Category	Focus	Methodology
1	Stanford AI Index Report 2025 [1]	Academic	Comprehensive review of the state of AI globally	Analysis of >100 databases, academic publications, patents, public surveys
2	McKinsey / QuantumBlack — The State of AI: Global Survey 2025 [2]	Consulting	State of corporate AI adoption and value	Survey of 1,993 respondents across industries and geographies, June–July 2025

#	Source	Category	Focus	Methodology
3	McKinsey / QuantumBlack — Seizing the Agentic AI Advantage [3]	Consulting	Agentic AI strategy and scaling	Qualitative analysis, executive interviews, case studies
4	BCG — The Widening AI Value Gap [4]	Consulting	AI maturity and the adoption–value gap	Survey of 1,250 companies globally
5	Deloitte — Tech Trends 2026 [5]	Consulting	Technology trends for enterprise	Survey of 500 US technology leaders, June–July 2025
6	PwC — 2026 AI Business Predictions [6]	Consulting	Corporate AI forecasts and responsible AI	PwC Responsible AI Survey data, February 2026
7	Accenture — Sovereign AI: Own Your AI Future [7]	Consulting	Sovereign AI and geopolitical dynamics	Global executive survey, July–August 2025
8	Accenture — Banking Top Trends FY26 [8]	Consulting	AI in banking and financial services	Industry analysis, transaction data
9	Google Cloud — AI Agent Trends 2026 [9]	Technology	Agentic AI adoption and architecture	Survey of 3,466 executives across industries
10	Google Cloud — State of AI Infrastructure [10]	Technology	Enterprise AI infrastructure	Global survey of 2,500+ IT and business respondents
11	Microsoft — The Strategic CIO's Generative AI Playbook [11]	Technology	GenAI strategy for CIOs	Microsoft Work Trend Index data, Gartner benchmarks
12	IBM Institute for Business Value — The AI Multiplier Effect [12]	Technology	Role of data and the CDO in AI transformation	Survey of 1,700 Chief Data Officers, Q3 2025
13	Gartner — Top 10 Strategic Technology Trends 2026 [13]	Analyst	Strategic technology trends	Gartner research methodology, client enquiry analysis
14	Andreessen Horowitz (a16z) — Big Ideas 2026 [14]	Venture	Emerging technology and market forecasts	Expert opinions of fund general partners
15	Alan Turing Institute / PAIF — GenAI Model Risk Management [15]	Research institute	AI governance in financial services	Collaborative research with 5 industry partners (Accenture, Aviva, BNY, HSBC, Mastercard)

Source selection criteria

Sources were selected on the basis of the following criteria:

- Published in 2025 – early 2026 with an explicit focus on 2025–2026+ trends
- Availability of quantitative data and/or large-scale executive surveys
- Global or regional coverage (no fewer than 500 respondents for survey-based sources)
- Publisher reputation in technology analytics and/or independent research standing

Aggregate survey coverage

The combined corpus represents a substantial evidence base. Table 2 summarises the survey-based sources by sample size and collection period.

Table 2. Survey-based sources: sample characteristics.

Source	Sample size	Collection period
Stanford AI Index [1]	>100 databases; global academic and industrial data	Continuous; report published April 2025
McKinsey State of AI [2]	1,993 respondents	June–July 2025
BCG [4]	1,250 companies	2025
Deloitte [5]	500 US technology leaders	June–July 2025
Accenture Sovereign AI [7]	Global executive survey	July–August 2025
Google Cloud AI Agents [9]	3,466 executives	2025
Google Cloud AI Infrastructure [10]	2,500+ respondents	2025
IBM CDO Study [12]	1,700 CDOs	Q3 2025

Category distribution

The 15 sources span six categories, ensuring methodological and perspectival diversity: academic (1), consulting firms (7), technology companies (4), analyst agencies (1), venture funds (1), and research institutes (1).

1.3 What is excluded and why

The following are excluded from the analysis:

- **Vendor marketing materials** without independent data verification
- **Forecasts beyond a 5-year horizon** owing to high uncertainty
- **Narrowly specialised technical benchmarks** without business relevance
- **Regional reports** with a sample of fewer than 500 respondents
- **Publications without stated data-collection methodology**

2. Executive summary

12 key conclusions:

1. **AI adoption has reached critical mass, but value extraction lags.** 88% of organisations use AI, yet only 5–7% generate significant value (McKinsey, BCG — high consensus, 4+ sources) [2; 4].
2. **Agentic AI is the dominant technological trend of 2026.** AI agents already account for 17% of all AI value and are projected to reach 29% by 2028 [4, p. 10]. Interest in multi-agent systems has grown by 1,445% over 15 months [13, p. 12].
3. **Workflow redesign is the critical success factor.** Companies that fundamentally redesign processes for AI are 2.8× more likely to achieve meaningful results [2, p. 11]. Technology delivers only 20% of value; 80% comes from redesigning work [6, p. 6].
4. **The US maintains leadership by a factor of 12×.** Private AI investment in the US totalled $109.1 billion — ~12 times that of China ($9.3 billion) and 24 times that of the UK ($4.5 billion) [1, p. 3].
5. **The gap between the US and China in model quality has nearly closed.** On the MMLU benchmark the difference narrowed from 17.5 p.p. to 0.3 p.p. within one year [1, p. 14].
6. **Inference costs have fallen by a factor of ~280 in 18 months.** The cost of generation at the GPT-3.5 level dropped from $20 to $0.07 per million tokens [1, pp. 4, 12].
7. **Data quality is the primary bottleneck for scaling AI.** 70% of organisations experience difficulties with data governance [10, p. 20]. 80% of corporate knowledge resides in unstructured data [14, p. 1].
8. **Hybrid cloud dominates for AI workloads.** 74% of organisations prefer a hybrid approach [10, p. 43]. Only 2% use a fully self-managed infrastructure.
9. **Regulatory pressure is growing exponentially.** The number of AI laws at the US state level rose from 1 (2016) to 131 (2024) — a 2.7× increase in a single year [1, p. 342]. 59 federal AI regulations were adopted in 2024 — more than doubling year-on-year [1, p. 4].
10. **AI in medicine is approaching an inflection point.** FDA-approved AI devices grew from 6 (2015) to 223 (2023) [1, p. 3]. The MedQA benchmark reached 96% (o1), approaching saturation [1, p. 301].
11. **Sovereign AI is becoming a strategic priority.** 61% of organisations are more inclined towards sovereign technologies owing to geopolitical turbulence [7, p. 5]. The sovereign AI infrastructure market is estimated at $1.5 trillion [7].
12. **Workforce transformation is accelerating.** 32% of companies expect headcount reductions of ≥3% owing to AI [2, pp. 17–18]. At the same time, AI is equalising productivity: junior specialists see gains of 21–40%, whilst senior specialists see only 7–16% [1, p. 269].

3. Key trends for 2026

3.1 Agentic AI and autonomous systems

Definition and current state

What is agentic AI?

Sundar Pichai (CEO, Google) defines AI agents as "systems that combine the intelligence of advanced AI models with access to tools to take actions on your behalf, under your control" [9, p. 3].

A more detailed definition: agentic AI is AI that "moves beyond answering questions to understanding goals, making plans, and taking actions through applications to achieve an outcome with significant human guidance and control" [9, p. 3].

Key characteristics of AI agents:

- Autonomous decision-making within defined boundaries
- Ability to plan multi-step tasks
- Memory and context from prior interactions
- Integration with external systems and tools
- Adaptability to changing conditions

Agent spectrum by complexity [11, p. 28]:

Table 3. Agent spectrum by complexity level.

Type	Description	Examples
Retrieval agents	Rule-based search and information retrieval	RAG systems, search bots
Task agents	Automation of defined processes	Invoice processing, scheduling
Autonomous agents	Self-directed planning and learning	Complex workflow orchestration

Current adoption level

Source consensus (3+ reports): Agentic AI is at an early stage of mass adoption with high interest but a low level of production deployment.

Table 4. Current adoption levels for agentic AI across sources.

Source	Metric	Value
McKinsey State of AI [2]	Organisations experimenting with agents	62%
McKinsey State of AI [2]	Organisations scaling agents in at least 1 function	23%
BCG [4]	Companies experimenting with agents	46%

Source	Metric	Value
Google Cloud AI Agents [9]	Executives with AI agents in production	52%
Deloitte [5]	Organisations with AI agents in production	**only 11%**
Gartner [13]	Growth in multi-agent system enquiries (Q1'24→Q2'25)	**1,445%**

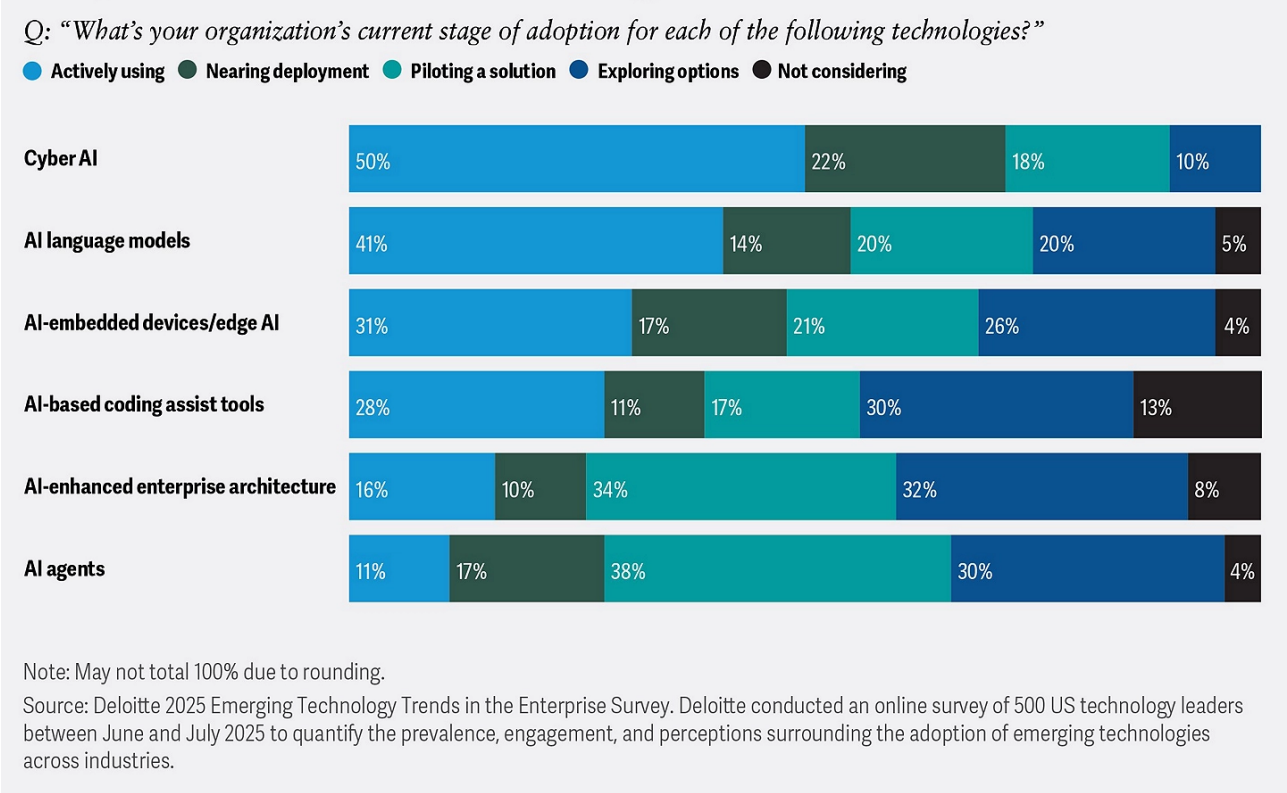

Figure 1. AI technology adoption stages. Adapted from: Deloitte [5, p. 15].

Barriers to agent adoption

According to McKinsey [3, p. 9], six key barriers impede scaling:

1. **Fragmented initiatives** — bottom-up approach; fewer than 30% have CEO sponsorship
2. **Lack of mature packaged solutions** — need for custom development
3. **Technological limitations of LLMs** — inaccuracy, passivity, limited memory
4. **Siloed AI teams** — disconnection from IT, data, and business functions
5. **Data accessibility and quality gaps** — particularly unstructured data
6. **Cultural apprehension and organisational inertia** — fear of disruption

Deloitte stands out: "The problem is not the technology, but the attempt to automate existing processes designed for humans" [5, p. 8].

Forecasts across firms

Table 5. Industry forecasts for agentic AI adoption.

Source	Forecast	Horizon
BCG [4]	Agentic AI's share of total AI value reaches 29%	2028
Gartner [13]	70% of MAS will use narrowly specialised agents	2027
Gartner [13]	60% of MAS will have multivendor interoperability	2028
Microsoft [11]	81% of leaders expect agent integration into AI strategy	12–18 months
Google Cloud [9]	Agents will manage complex multi-step workflows across systems	2026

Multi-agent system evolution [13, p. 12]:

- **Phase 1 (current):** Single platform — agents within one ecosystem
- **Phase 2 (2026–2027):** Cross-platform — interaction between platforms
- **Phase 3 (2028+):** "Internet of agents" — an open ecosystem of agents

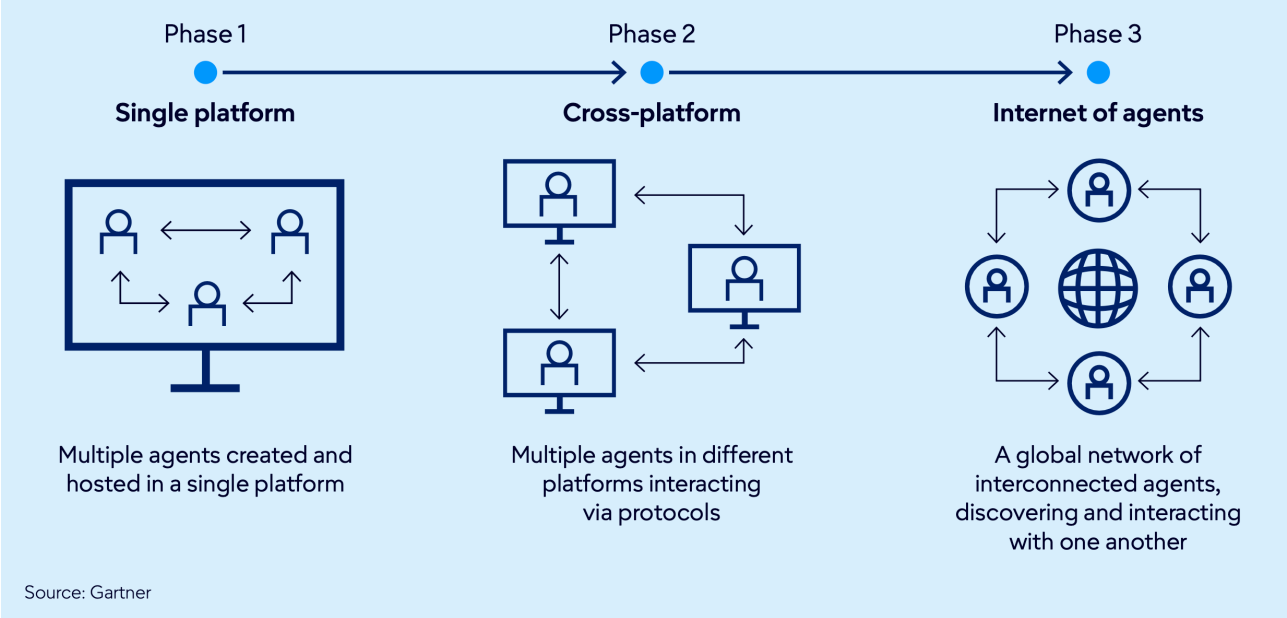

Figure 2. The evolution of multiagent systems. Adapted from: Gartner [13, p. 12].

Technological protocols for interoperability:

- **Agent2Agent (A2A)** — Google's open protocol for inter-organisational agent communication [9, p. 18]
- **Model Context Protocol (MCP)** — a standard for connecting LLMs to data and tools [9, p. 18]
- **Agent Payments Protocol (AP2)** — a protocol for agentic e-commerce [9, pp. 21–22]

Consensus and divergences

Consensus (mentioned in 5+ sources):

- Agentic AI is the key trend for 2025–2026
- Workflow redesign is more critical than technology
- Human-in-the-loop remains necessary
- Governance for agents is the main unresolved question

Divergences:

Table 6. Areas of divergence on agentic AI.

Question	Optimists	Sceptics
Speed of mass adoption	Google: 52% already in production [9]	Deloitte: only 11% in production [5]
ROI timeline	BCG: 17%→29% of value by 2028 [4]	McKinsey: >80% without EBIT impact currently [3]
Level of autonomy	a16z: "Internet of agents" by 2028 [14]	PwC: agents handle ~50% of human tasks [6]

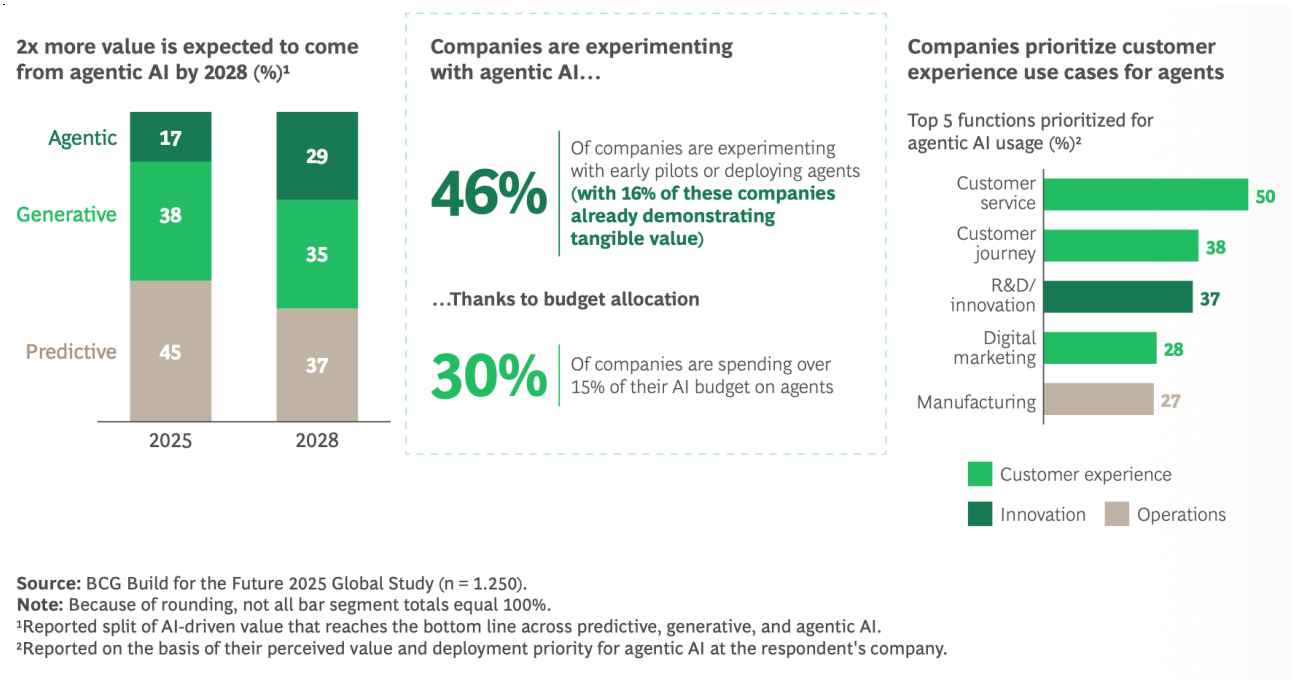

Figure 3. Agentic AI value trajectory, 2025–2028, and top five functions for agents. Adapted from: BCG [4, p. 10].

3.2 From pilots to scaling (AI maturity)

Current maturity level

"Gen AI Paradox" — this is how McKinsey describes the situation in which "nearly 8 in 10 companies use generative AI, but just as many see no significant impact on financial results" [3, p. 4].

Distribution of organisations by AI deployment phase [2, p. 2]:

Table 7. Distribution of organisations by AI deployment phase (McKinsey).

Phase	Share of organisations	Characteristic
Fully scaled	7%	Enterprise-wide deployment
Scaling	31%	Scaling beyond pilots
Piloting	30%	Testing at limited scale
Experimenting	32%	Exploration without production

BCG data yield a similar picture [4, p. 4]:

Table 8. AI maturity distribution (BCG).

BCG category	Share	Characteristic
Future-built	5%	Generating value at scale
Scaling	35%	Active scaling
Emerging	46%	Early stages
Stagnating	14%	Stagnation or no progress

The data show that ~60% of organisations derive no material value from AI despite their investments.

Maturity by industry

Production deployment by industry [10, p. 7]:

Table 9. GenAI production deployment and perceived importance by industry.

Industry	GenAI in production	Importance of GenAI (very/extremely)
IT consulting	58%	91%
Hardware/software	53%	88%
Manufacturing	48%	75%

Industry	GenAI in production	Importance of GenAI (very/extremely)
Healthcare	45%	75%
Retail	39%	72%
Financial services	**7%**	**51%**

Financial services exhibit the lowest level of production deployment (7%) despite significant investment, a finding explained by stringent regulatory requirements.

Leaders in AI maturity [4, p. 8]:

- Software
- Telecom
- Payments/Fintech

Laggards:

- Fashion & Luxury
- Chemicals
- Real Estate

Characteristics of AI leaders

What distinguishes "AI high performers" [2, pp. 10–15]:

Table 10. Characteristics of AI high performers vs other organisations.

Characteristic	High performers	Others	Gap
Intent for transformational change	50%	14%	**3.6×**
Fundamental workflow redesign	55%	20%	**2.8×**
>20% of digital budget on AI	35%	7%	**4.9×**
Senior leadership commitment (strongly agree)	48%	16%	**3.0×**

BCG further reports that future-built companies outperform laggards by a factor of **1.7** on revenue growth, **1.6** on EBIT margin, and **3.6** on three-year total shareholder return [4, pp. 4–5].

The top barriers to GenAI adoption, according to Google Cloud, are security risks (cited by 39% of organisations), data privacy concerns (36%), regulatory concerns (29%), and cost (26%) [10, p. 18].

Data challenges [10, p. 20]:

- 70% of organisations experience difficulties with data governance
- Integration of data into AI models
- Shortage of training data

Organisational barriers [1, p. 179]:

Table 11. Organisational obstacles to AI adoption.

Obstacle	% of respondents
Knowledge and training gaps	51%
Resource/budget constraints	45%
Regulatory uncertainty	40%
Technical limitations	32%
Organisational resistance	22%
Lack of executive support	16%

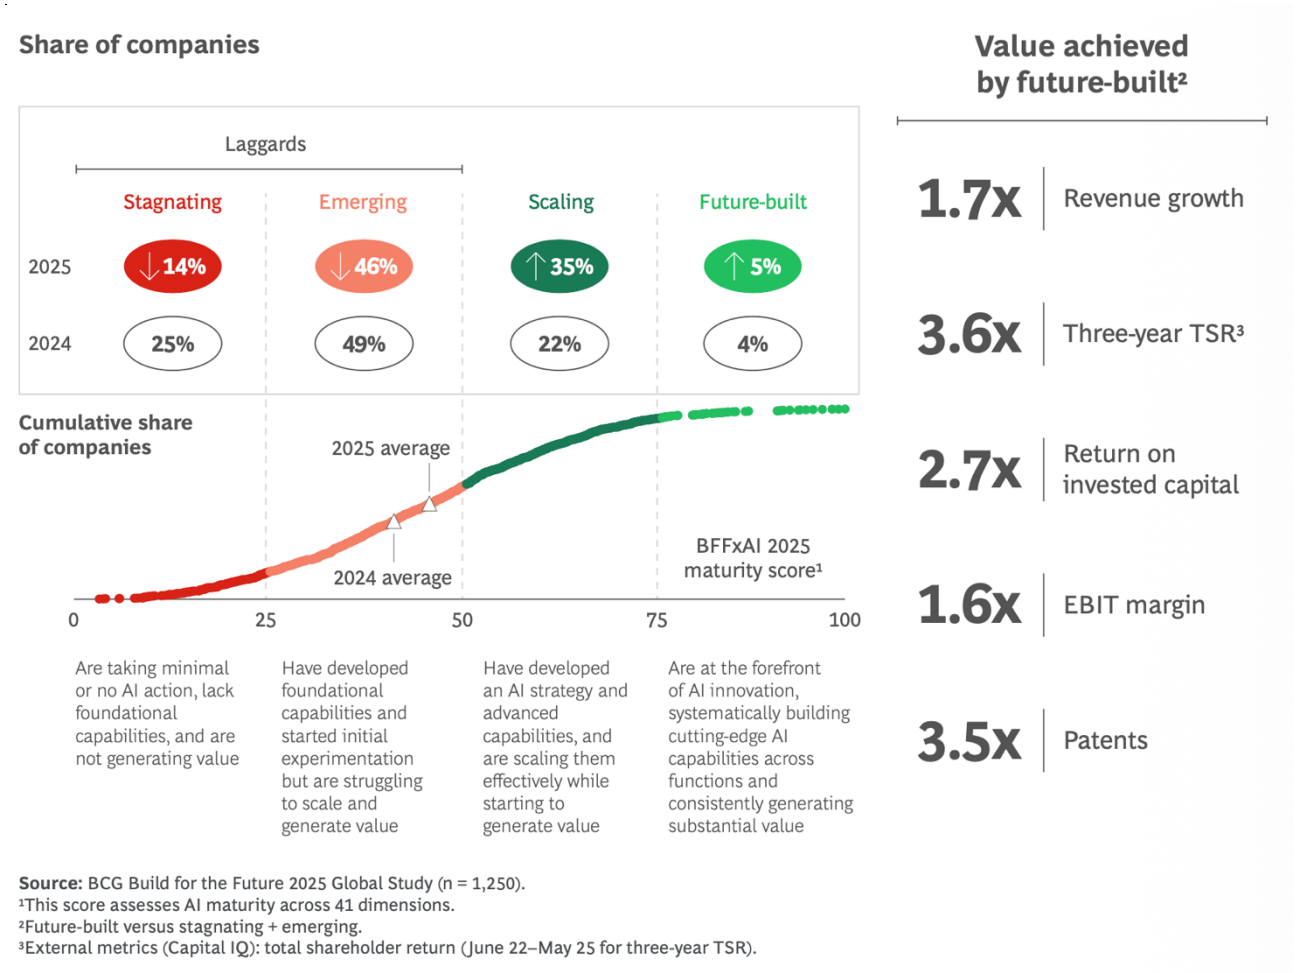

Figure 4. AI maturity distribution: 5% future-built, 35% scaling, 46% emerging, 14% stagnating + value multipliers. Adapted from: BCG [4, p. 4].

Adoption forecasts

Table 12. Consensus adoption forecasts, 2026–2030.

Source	Forecast	Horizon
Microsoft/Gartner [11; 13]	>80% of enterprises will use GenAI APIs/models in production	2026
Gartner [13]	80% of organisations with AI-augmented "tiny teams"	2030
Gartner [13]	40% of enterprise apps on AI-native platforms	2030
Google Cloud [10]	Enterprise infrastructure adoption growth >30%	2026

3.3 Infrastructure and compute

Hardware trends

Key technological shifts:

1. **Exponential growth of training compute**
   - Doubling period: ~5 months [1, p. 12]
   - Dataset sizes (LLM): doubling approximately every ~8 months [1, p. 12]
2. **Dramatic decline in inference costs**
   - Cost at the GPT-3.5 level: $20 → $0.07 per million tokens (~280× reduction in 18 months) [1, pp. 4, 12]
   - Hardware performance: +43% annually [1, p. 13]
   - Energy efficiency: +40% annually [1, p. 4]
3. **Test-time compute revolution**
   - OpenAI o1 on the IMO qualifying exam: 74.4% vs GPT-4o 9.3% [1, p. 14]
   - Cost: o1 is 6× more expensive and 30× slower than GPT-4o [1, p. 14]
   - Latency: o1 is approximately 40× slower than GPT-4o (29.7 s vs 0.72 s) [1, p. 112]
4. **Reduction in model size whilst maintaining quality**
   - 2022: PaLM (540B parameters) for 60% on MMLU
   - 2024: Phi-3-mini (3.8B parameters) achieves the same
   - **A 142× reduction in two years** [1, p. 99]

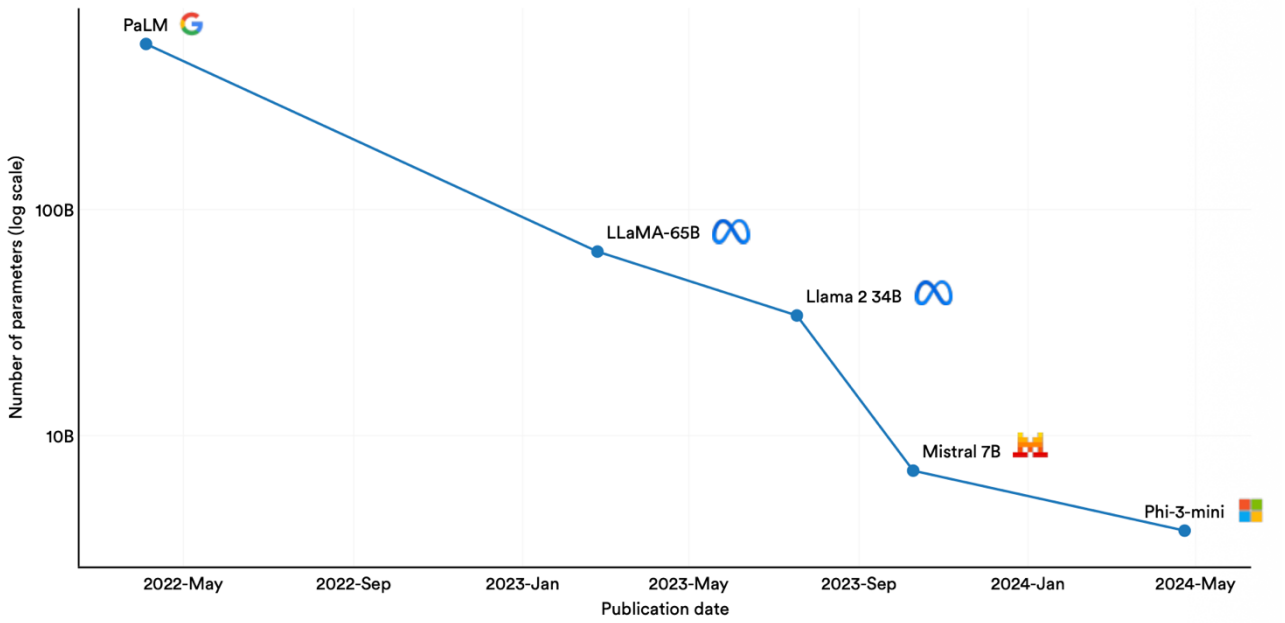

Figure 5. Smallest models scoring 60%+ on MMLU. Adapted from: Stanford AI Index [1, p. 99].

Heterogeneous computing [13, pp. 7–8]:

- CPUs, GPUs, AI ASICs, quantum processors, neuromorphic chips
- Unified APIs: CUDA, NCCL, ROCm, RCCL, QuantumSDKs
- Forecast: 40% of enterprises with hybrid computing architectures by 2028 (up from 8%)

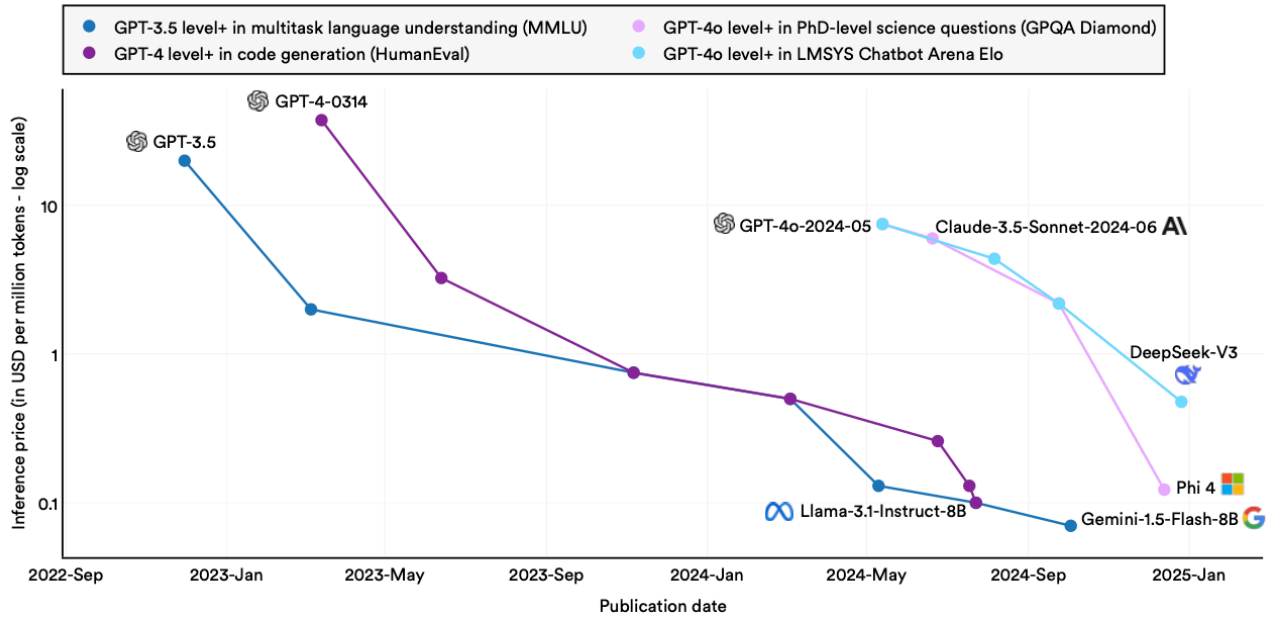

Figure 6. Decline in inference costs by benchmark. Adapted from: Stanford AI Index [1, p. 64].

Cloud infrastructure

Dominance of hybrid cloud [10, p. 43]:

Table 13. Enterprise AI infrastructure approaches by share of organisations.

Approach	Share of organisations
Hybrid cloud (on-prem + cloud)	**74%**
— On-prem + multicloud	41%
— On-prem + single cloud	33%
Multicloud only	19%
Single cloud only	4%
On-premises only	4%

Three-tier hybrid architecture [5, pp. 8–9]:

- **Cloud** — for elasticity
- **On-premises** — for consistency
- **Edge** — for immediacy

The majority of organisations (57%) take a hybrid approach to infrastructure management, blending managed and self-managed services in roughly equal proportions. A further 21% rely primarily on managed services and 17% primarily on self-management, whilst fully managed (4%) and completely self-managed (**2%**) approaches remain rare [10, p. 47].

Edge computing [10, p. 38]:

- 73% consider deployment of GenAI on edge devices important or critically important
- Gartner forecast: 80% of warehouses with robotics/automation by 2028 [13]

Energy consumption and sustainability

The carbon cost of training frontier models is rising sharply. Training GPT-4 emitted an estimated 5,184 tonnes of CO_2, while Llama 3.1 405B produced roughly 8,930 tonnes — equivalent to the annual footprint of approximately 500 average Americans (18 tonnes each) [1, pp. 13, 73].

Energy consumption of training continues to rise, doubling approximately every 12 months [1, p. 12].

Investments in energy for AI [1, pp. 219, 222]:

- Microsoft: $1.6 billion to restart the Three Mile Island reactor
- Google and Amazon: small modular reactor (SMR) agreements

The efficiency paradox [6, p. 11]:

- AI is rapidly becoming more energy-efficient, yet usage grows even faster
- Rising efficiency makes AI cheaper → accelerates usage → affects emissions, water, energy prices

PwC recommendations [6]:

- Approve token usage only where significant value is present
- Carbon scheduling to reduce emissions and costs
- Diversify energy sources; build own renewables

Infrastructure forecasts [10, p. 48]:

- Demand for AI-ready data-centre capacity growth: 33% annually until 2030
- Growth in data-centre spending over 5 years: 2× (doubling)

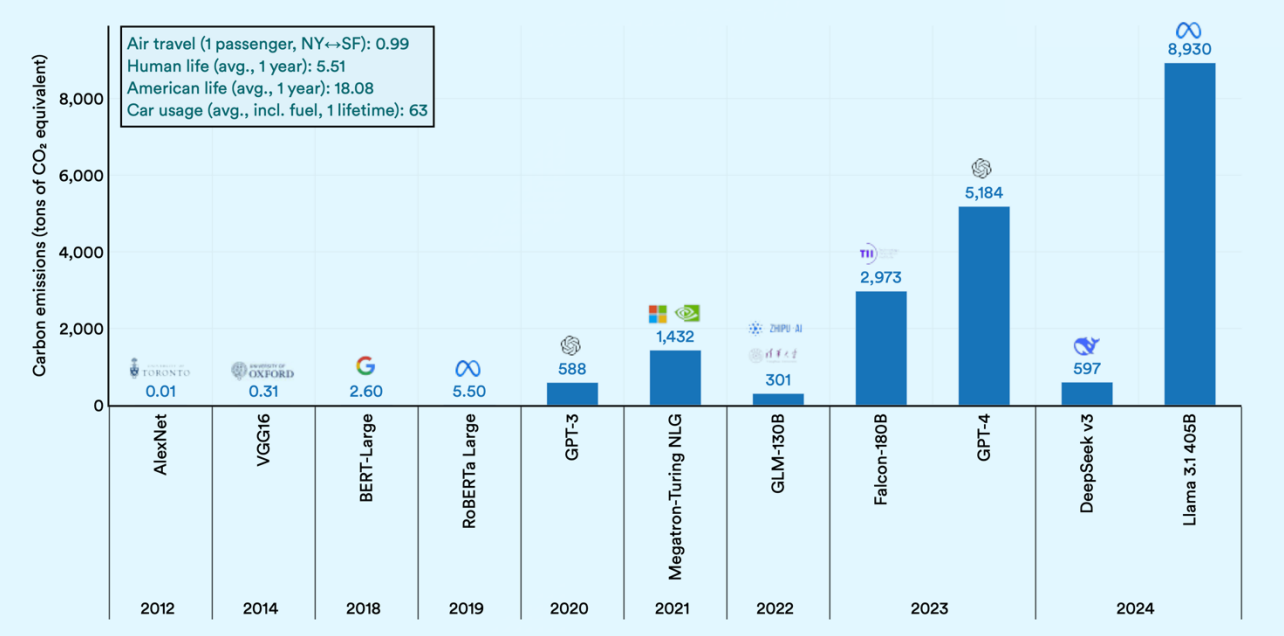

Figure 7. Carbon emissions from AI model training. Adapted from: Stanford AI Index [1, p. 73].

3.4 The economics of AI

Investment by region

In 2024, corporate AI investment reached a record **$252.3 billion** (+25.5% year-on-year), private AI investment totalled **$150.79 billion** (+44.5%) (over the past decade, private AI investment has grown by a factor of 13), and the generative-AI segment alone attracted **$33.9 billion** in private funding (+18.7%) [1, pp. 218, 247–260].

Regional distribution of private investment [1, pp. 252–255]:

Table 14. Private AI investment by country, 2024.

Country	Investment	Share relative to US
USA	**$109.1 bn**	100%
China	$9.3 bn	8.5%
UK	$4.5 bn	4.1%
Sweden	$4.3 bn	3.9%
Canada	$2.9 bn	2.7%
France	$2.6 bn	2.4%
Germany	$2.0 bn	1.8%

These figures underscore the scale of US dominance: private AI investment in the United States exceeds that of China by a factor of roughly 12 and that of the United Kingdom by a factor of 24.

Major deals in 2024 [1, p. 18]:

Table 15. Largest AI deals and partnerships, 2024.

Date	Company	Type	Amount / Valuation
Oct 2024	OpenAI	Funding	$6.6 bn / $157 bn
Dec 2024	Databricks	Funding	$10 bn / $62 bn
Dec 2024	xAI	Funding	$6 bn / $40+ bn
Nov 2024	Anthropic + Amazon	Partnership	+$4 bn (total $8 bn)
Jan 2024	Synopsys + Ansys	Acquisition	$35 bn

Government investment [1, pp. 4, 20]:

Table 16. Government AI investment programmes.

Country	Investment	Purpose
France	€109 bn	AI infrastructure
Saudi Arabia	$100 bn	Project Transcendence
China	$47.5 bn	Big Fund III (semiconductors)
Canada	$2.4 bn	AI advantage package
India	$1.25 bn	IndiaAI Mission

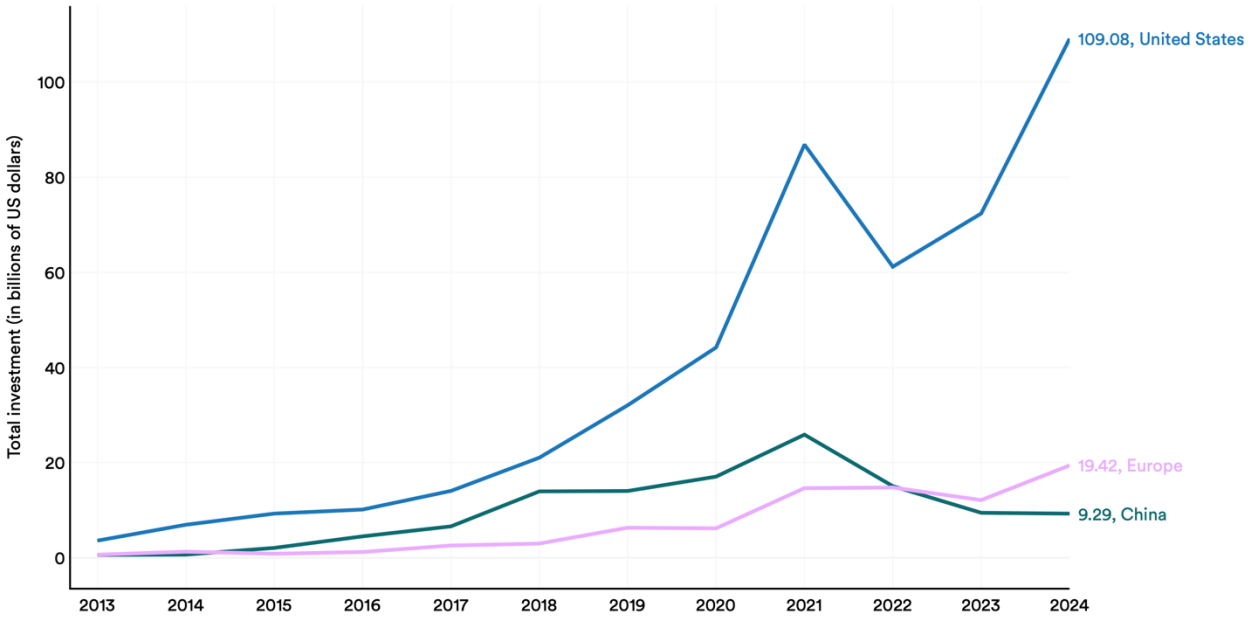

Figure 8. Global private investment in AI by geographic area, 2013–2024. Adapted from: Stanford AI Index [1, p. 254].

ROI and business value

The aggregate economic potential of generative AI is substantial. McKinsey estimates the total addressable value at **$2.6–4.4 trillion annually** across all industries and functions — equivalent to adding an economy the size of the United Kingdom [3, p. 4]. However, realising this potential requires the transition from horizontal use cases (copilots, chatbots) to vertical, function-specific solutions with direct P&L impact, and ultimately to agentic workflows that redesign entire business processes [3, pp. 8–9].

Financial impact of AI by function [1, p. 265]:

Table 17. Financial impact of AI by business function.

Function	Cost savings	Revenue gains
Service operations	49%	57%
Supply chain	43%	63%
Software engineering	41%	44%
Marketing & sales	34%	**71%**

When asked about the primary business benefits they expect from GenAI, respondents most frequently cited increased employee productivity (22%), improved customer satisfaction (21%), and better workflows (20%) [10, p. 27].

Potential for the banking sector [8, p. 38]:

- Potential benefit for the top-200 banks from scaled GenAI (3 years): **$289 billion**
- Expected revenue increase: +5%
- Reduction in operational costs: −8%
- Reduction in loan-loss provisions: −16%

ROI among early adopters [9, p. 16]:

- 88% of early adopters report positive ROI on at least one use case

The labour market

AI and productivity — the equalising effect [1, pp. 268–270]:

Table 18. AI productivity gains by skill level across studies.

Study	Task	Low-skill gain	High-skill gain
Brynjolfsson et al.	Customer support	34%	~0%
Dell'Acqua et al.	Consulting	42.96%	16.5%
Cui et al.	Software engineering	21–40%	7–16%
Hoffman et al.	Software engineering	12–27%	5–10%

Across the four studies, AI boosts productivity by 10–45%, with the largest gains accruing to less experienced workers — effectively narrowing the gap between junior and senior specialists.

AI skills demand in the US — growth 2023→2024 [1, p. 226]:

Table 19. Fastest-growing AI skills in US job vacancies, 2023–2024.

Skill	Vacancy growth
RAG	**+2,047%**
Microsoft Copilot	**+2,238%**
Multimodal models	**+1,566%**
Prompt engineering	+350%
Generative AI	+323%
Large language modelling	+295%

Looking ahead, 32% of companies expect AI-driven headcount reductions of 3% or more over the next 12 months, 43% anticipate no change, and 13% foresee workforce growth of 3% or more [2, pp. 17–18].

Note: Companies consistently overestimate the scale of AI-driven workforce reductions. When asked about the **coming** year, a median of 30% of respondents anticipated headcount cuts of ≥3%; however, when the **same question** was posed retrospectively about the **past** year, only 17% reported that such cuts had actually materialised [2, pp. 17–18].

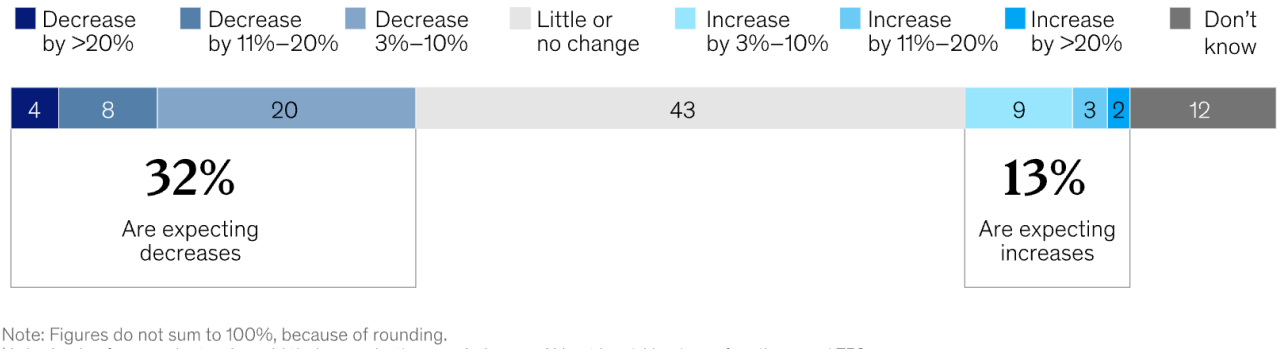

Figure 9. Expected workforce change. Adapted from: McKinsey [2, p. 18].

New workforce structure [6, pp. 6–7]:

- **Knowledge work:** an "hourglass" shape — junior + senior, fewer mid-tier
- **Front-line task work:** a "diamond" shape — more mid-level for agent orchestration, fewer entry-level

New roles:

- AI workflow designers
- Agent supervisors / orchestrators
- Governance leads
- "Chief of Staff for AI" [9, p. 41]

3.5 Regulation and governance

Key initiatives by region

Global AI legislation has grown rapidly: from a single law in 2016 to a peak of 44 in 2022, followed by 30 in 2023 and 40 in 2024, bringing the cumulative total to **204 laws** across **39 countries** [1, pp. 337–349].

Top countries by AI legislation (2016–2024):

1. USA: 27 laws
2. Portugal: 20
3. Russia: 20
4. Belgium: 18
5. South Korea: 13

US state-level AI legislation has exploded: from a single law in 2016 to 49 in 2023 and **131** in 2024 — a **2.7-fold increase** in one year [1, pp. 340–343]. California led with 22 laws in 2024, followed by Utah (12), Maryland (8), and Virginia (6).

US federal AI regulations more than doubled, from 25 (issued by 21 agencies) in 2023 to **59** (issued by **42** agencies) in 2024 — growth of **136%** in the number of regulations and 100% in the number of issuing agencies [1, pp. 349–352]. The most active agencies in 2024 were the Department of Health and Human Services (14 regulations), the Centers for Medicare and Medicaid (7), the Commerce Department (7), and the Executive Office of the President (5).

EU AI Act: Set to enter full force in 2026. Requires watermarking for AI-generated content [13, pp. 21–22].

Deepfake regulation in the US [1, pp. 344–345]:

- Before 2024: 5 states with regulation of election deepfakes
- By end of 2024: **20 states**
- Intimate deepfakes: 25 states with laws for all, 5 — only for minors

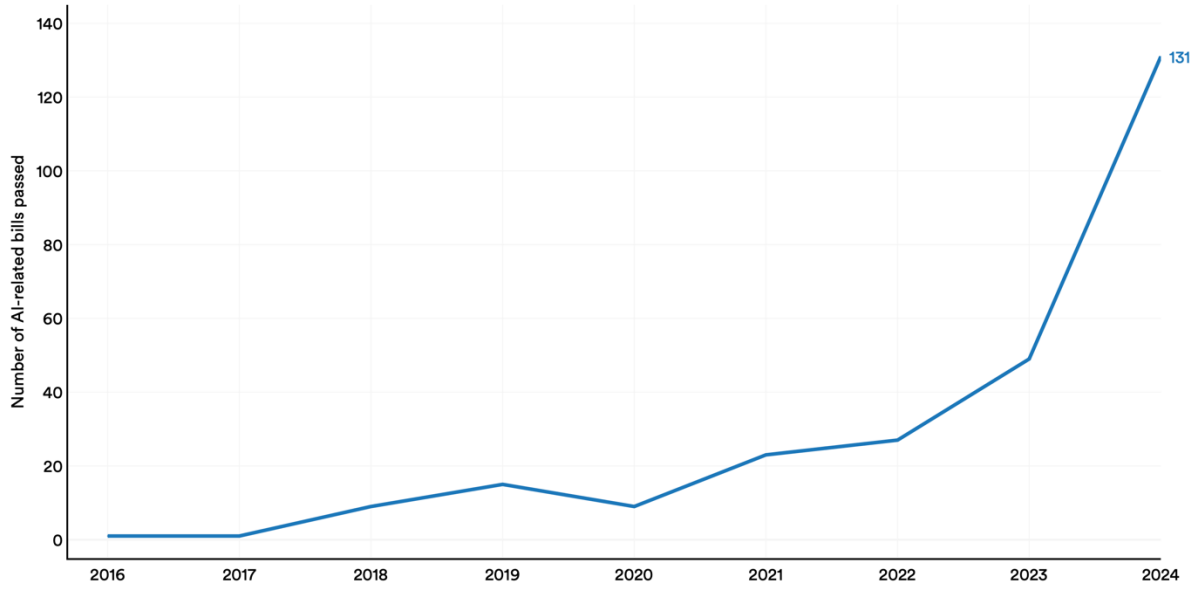

Figure 10. State-level AI bills in the US, 2016–2024. Adapted from: Stanford AI Index [1, p. 342].

Responsible AI

AI incidents — a sharp rise [1, p. 167]:

- 233 incidents in 2024 — a record high
- Growth of 56.4% compared with 2023

RAI risks — awareness vs mitigation gap [1, p. 176]:

Table 20. Responsible AI risks: awareness vs active mitigation.

Risk	Consider relevant	Actively mitigating	Gap
Cybersecurity	66%	53%	−13 p.p.
Regulatory compliance	63%	50%	−13 p.p.

Risk	Consider relevant	Actively mitigating	Gap
Privacy	60%	46%	−14 p.p.
IP infringement	**57%**	**38%**	**−19 p.p.**
Inaccuracy	60%	55%	−5 p.p.

The pattern is consistent across all five risk categories: recognition outpaces action, with mitigation lagging awareness by 5 to 19 percentage points.

Hallucination rates among the leading models have converged in a narrow band: GLM-4-9b-Chat and Gemini-2.0-Flash-Exp share the lowest rate at **1.3%**, followed by o1-mini (**1.4%**), GPT-4o (**1.5%**), and GPT-4 (**1.7%**) [1, p. 171]. On the HELM safety benchmark, Claude 3.5 Sonnet leads with a mean score of **0.977**, closely followed by o1 (**0.976**) and Claude 3 Opus (**0.964**). GPT-4 Turbo scores **0.957**, while DeepSeek R1 trails at **0.863** [1, p. 202].

PwC Responsible AI Survey 2025 [6, p. 8]:

- 60% of executives: RAI improves ROI and efficiency
- 55%: improves customer experience and innovation
- ~50%: difficult to translate RAI principles into operational processes

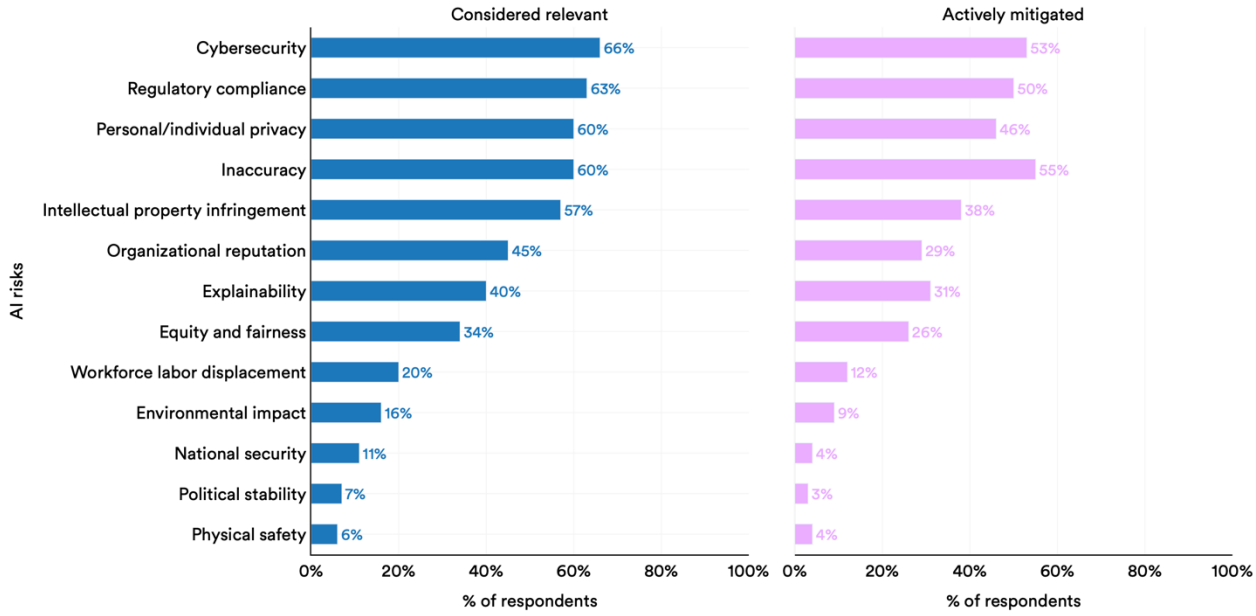

Figure 11. RAI risks: considered relevant vs actively mitigated. Adapted from: Stanford AI Index [1, p. 176].

Regulatory pressure forecasts

Table 21. Regulatory and sovereignty forecasts by source.

Source	Forecast	Horizon
Gartner [13]	65% of governments with technology sovereignty rules	2028
Gartner [13]	75% of enterprises with geopatriation workloads	2030
Gartner [13]	80% of AI incidents from internal policy violations	—
Accenture [7]	35% of countries tied to region-specific AI platforms	2027

AI governance within organisations:

- Only 14% have dedicated AI governance roles [1]
- Foundation Model Transparency Index: mean score rose from 37% to 58% (October 2023 → May 2024) [1, pp. 199–201]

3.6 AI in science and medicine

Breakthroughs 2024–2025

Key achievements in biomedical AI [1, pp. 285–291]:

Table 22. Key biomedical AI achievements, 2024–2025.

Model / Project	Organisation	Achievement
ESM3	EvolutionaryScale	98 billion parameters; created esmGFP — a protein whose evolution would have required nature 500 million years
AlphaFold 3	Google / Isomorphic Labs	Modelling of protein–DNA–RNA–ligand interactions; 93.2% accuracy given a defined binding pocket
AlphaProteo	Google DeepMind	Binding proteins up to 300× more effective than SOTA; 88% successful binders for BHRF1
GluFormer	Nvidia / Weizmann	Diabetes prediction 4 years in advance: 66% of new cases, 69% of cardiac deaths
Virtual AI Lab	Stanford	Autonomous AI laboratory created 92 nanobodies; >90% successfully bound SARS-CoV-2

Growth of protein models [1, pp. 291–296]:

- ProGen (2020): 1.2B parameters
- ESM2 (2023): 15B parameters
- ESM3 (2024): 98B parameters
- **Growth of 82× in 4 years**

Nobel Prizes 2024: Physics and Chemistry awarded for AI research [1, pp. 5, 20].

Clinical application of AI

Progress on the MedQA benchmark has been rapid. The best score stood at 67.6% in late 2022; by 2023 it had risen to 90.2%, and in 2024 OpenAI's o1 model reached **96.0%** — an improvement of 28.4 percentage points in two years and a level approaching saturation [1, p. 301].

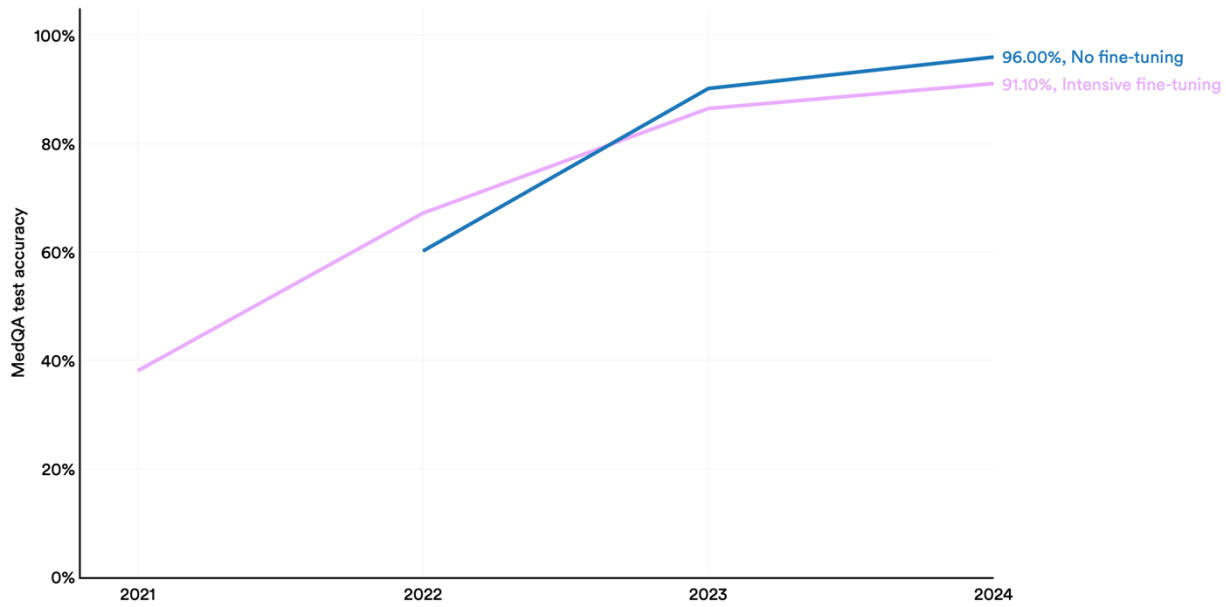

Figure 12. MedQA test accuracy timeline. Adapted from: Stanford AI Index [1, p. 301].

In a controlled diagnostic study, **GPT-4** operating alone achieved **92%** accuracy, whereas physicians assisted by GPT-4 scored only 76%, barely above the 74% achieved by physicians using conventional resources [1, pp. 305–306]. The implication is counter-intuitive: simply providing clinicians with access to an LLM did not improve their performance — a finding that underscores the need for workflow redesign rather than tool overlay.

The number of FDA-approved AI medical devices has grown exponentially: from a single device in 1995 to six in 2015 and **223** by 2023 [1, p. 309].

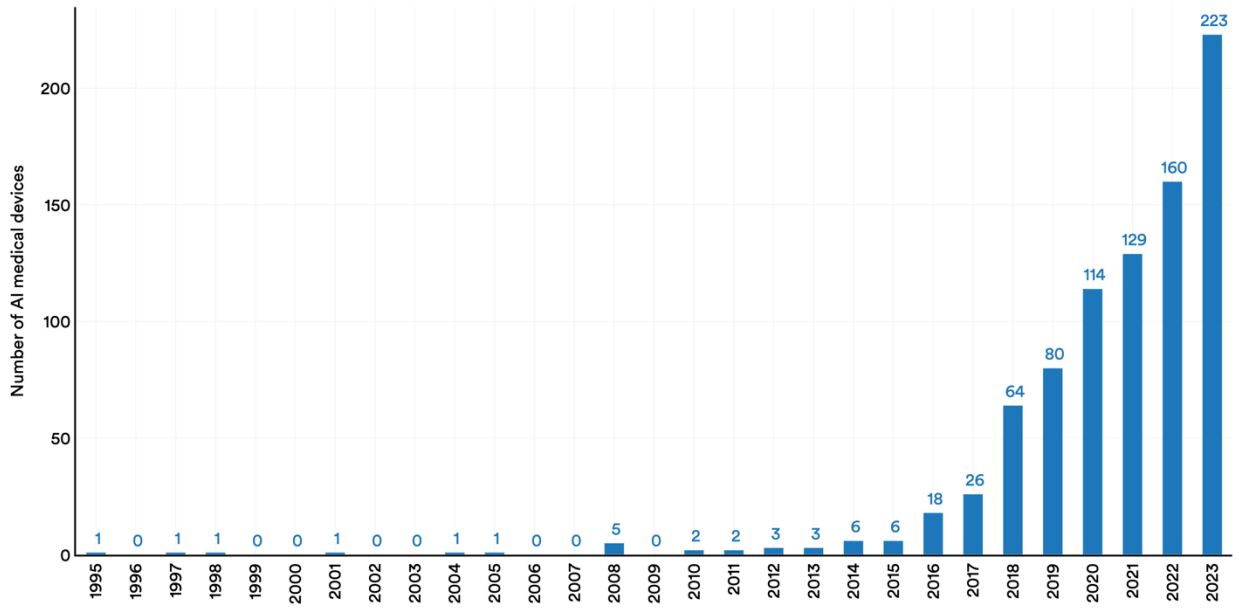

Figure 13. FDA-approved AI medical devices, 1995–2023. Adapted from: Stanford AI Index [1, p. 309].

Ambient AI scribes [1, pp. 307–308]:

- Kaiser Permanente: 3,442 physicians, 303,266 visits in 2 months
- Stanford: 55% adoption, saving ~20 minutes/day on EHR
- Burnout reduction: 26%; burden reduction: 35%
- Investment in the technology: ~$300M in 2024

AI foundation models in science

Models of 2024 [1, pp. 320–323]:

Table 23. AI foundation models for science, 2024.

Date	Model	Domain	Significance
Feb 2024	CrystalLLM	Materials science	49% vs 28% metastability rate
Apr 2024	ORBIT	Climate	113B parameters, 1.6 exaFLOPS
May 2024	Aurora	Earth systems	SOTA forecasts for air quality, waves, cyclones
Jul 2024	NeuralGCM	Weather/climate	Physics+ML hybrid; decadal forecasts
Sep 2024	FireSat	Fires	Detection of fires at 5×5 m within 20 minutes
Dec 2024	GenCast	Weather	15-day forecasts in minutes (vs hours)
Dec 2024	AlphaQubit	Quantum	SOTA quantum error detection

Academic interest in the ethics of medical AI is surging. The number of relevant publications rose from 288 in 2020 to 1,031 in 2024 — an increase of **258%** [1, p. 318].

NIH grants for the ethics of medical AI:

- 2023: 25 grants, $16.3M
- 2024: 337 grants, $276M
- **Growth of 17× in one year**

4. Trends by industry

4.1 Finance

Current state [10, p. 7]:

- GenAI in production: **only 7%** (the lowest among all industries)
- Importance of GenAI (very/extremely): 51%

Why the lag?

- High regulatory requirements (SR 11-7, SS1/23)
- Need for continuous monitoring vs point-in-time validation [15, p. 26]
- Vendor risk as a critical factor [15, pp. 9–11]

Trends [8]:

1. **Digital currencies entering the mainstream:**
    - Stablecoin transaction volume (adjusted, October 2025): $10.7 trillion annualised (+88% YoY)
    - 87% of financial institutions exploring tokenisation
    - 135 countries researching CBDC
    - Up to $13 trillion of transactions may shift to alternative payment methods by 2030
2. **Agentic payments and "smart money":**
    - 35% of consumers willing to entrust AI with final purchase selection
    - 57% of business leaders expect mainstream agentic commerce within 3 years
3. **Tech debt as a critical risk:**
    - ~70% of bank IT budgets spent maintaining legacy systems
    - 2 in 3 banks plan to reduce manual coding by 10–50% with GenAI
4. **Competition for the balance sheet:**
    - Stablecoins attacking deposits ($110 trillion globally)
    - Private credit attacking loans ($106 trillion globally)
    - Risk to US bank pre-tax income: 22%
    - ~40% of US deposits at risk from stablecoins (Treasury estimate)

The concept of the "10× bank": "One person manages a team of AI co-workers for exponentially greater impact" [8, p. 32].

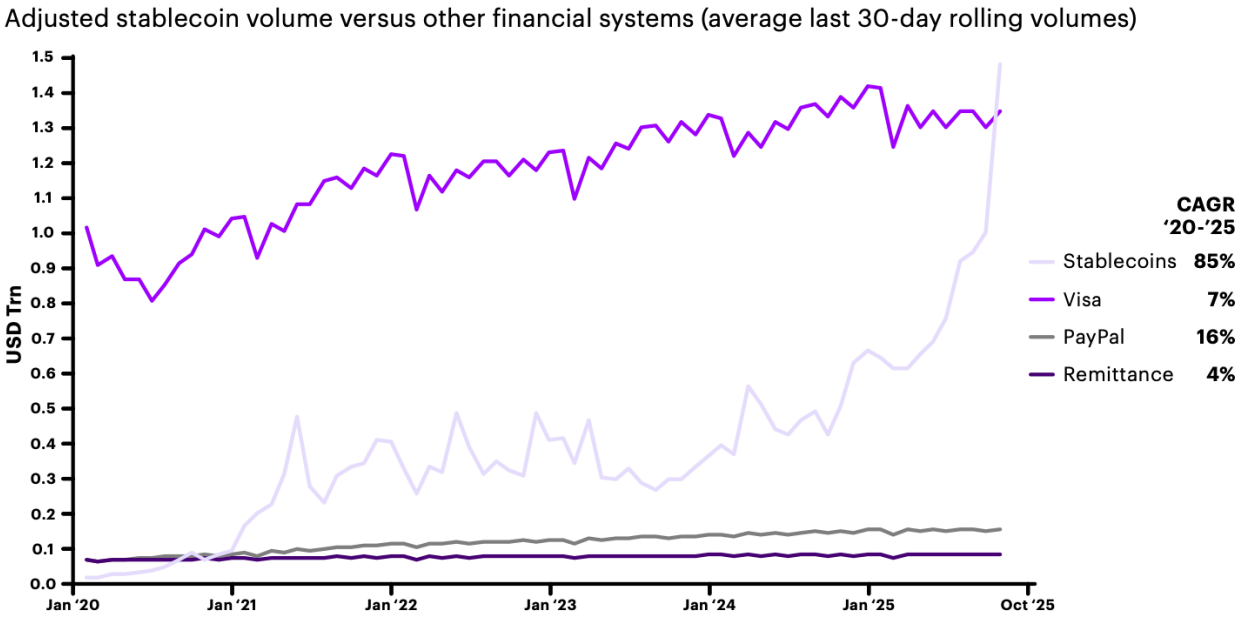

Figure 14. Adjusted stablecoin transaction volume (30-day rolling average) vs Visa and PayPal. Adapted from: Accenture [8, p. 7].

4.2 Healthcare

Current state [10, p. 7]:

- GenAI in production: 45%
- Importance of GenAI (very/extremely): 75%

Key applications:

1. **Clinical decision support:**
   - MedQA: o1 reached 96.0% — approaching saturation
   - More complex assessments required
2. **Ambient AI scribes:**
   - 55% adoption at Stanford
   - Saving ~20 minutes/day on EHR
   - Burnout reduction: 26%
   - ~$300M investment in 2024
3. **Drug discovery:**
   - AlphaFold 3: 93.2% docking accuracy
   - AlphaProteo: binders up to 300× more effective than SOTA
   - ESM3: creation of proteins requiring 500M years of evolution
4. **Predictive health:**
   - GluFormer: diabetes prediction 4 years in advance (66% of new cases)
   - "Healthy MAUs" — a new segment: consumers regularly monitoring their health [14, p. 11]

Healthcare stands out as the industry where AI's transformative potential and its practical barriers are most sharply juxtaposed. On the one hand, model performance on medical benchmarks has reached physician-level accuracy, ambient scribes are saving clinicians meaningful amounts of time, and protein-modelling tools are compressing decades of drug-discovery research into months. On the other hand, production deployment sits at 45% — lower than IT consulting, hardware, or manufacturing — and the sector's unique constraints (patient safety, regulatory oversight, data sensitivity) mean that scaling requires more than technical capability. The mismatch between what AI can do in controlled settings and what it is permitted to do in clinical practice may prove to be the defining tension of 2026 in this industry.

Regulatory context:

- 223 FDA-approved AI devices as of 2023
- Exponential growth continues

Ethical concerns [1, p. 318]:

1. Bias (~28%)
2. Privacy (~25%)
3. Equity (~14%)

4.3 Manufacturing

Current state [10, p. 7]:

- GenAI in production: 48%
- Importance of GenAI (very/extremely): 75%

Key applications:

1. **Robotics and automation:**
    - Gartner forecast: 80% of warehouses with robotics/automation by 2028
    - China: 276,300 industrial robots (51.1% of the global market) — 6× more than Japan
    - Collaborative robots: 2.8% (2017) → 10.5% (2023)
2. **Physical AI:**
    - Humanoid robotics: Figure 02, Tesla Optimus, Boston Dynamics Atlas
    - AI-native industrial base: simulation, automated design, AI-driven operations from day one [14, pp. 15–16]
3. **Quality control and predictive maintenance:**
    - Toyota: 20% reduction in model-creation time, >10,000 hours of savings per year [10, p. 45]

"Factory mindset" [14, pp. 16–17]:

- Modular deployment of AI + autonomy + skilled workers
- Mass-producing nuclear reactors, housing, data centres
- "The factory is the product" (Elon Musk)

4.4 Retail

Current state [10, p. 7]:

- GenAI in production: 39%
- Importance of GenAI (very/extremely): 72%

Key applications:

1. **Hyper-personalisation:**
   - "The year of me" — products are ceasing to be mass-produced and becoming made for you [14, p. 13]
   - AI learns from camera rolls, messaging patterns, routines
2. **Customer service automation:**
   - Home Depot Magic Apron: AI agent for expert consultations 24/7 [9, p. 26]
   - Danfoss: 80% automation of resolutions; response time from 42 hours to near real-time [9, p. 29]
3. **Demand sensing and forecasting:**
   - A priority area for agents [6, p. 3]
4. **Agentic commerce:**
   - Agent Payments Protocol (AP2): PayPal adoption [9, p. 22]

4.5 Other industries

IT consulting [10, p. 7]:

- GenAI in production: **58%** (leader)
- Importance of GenAI: **91%** (leader)

Hardware/software:

- GenAI in production: 53%
- Importance of GenAI: 88%

Cybersecurity [14, p. 2]:

- 2013–2021: unfilled jobs grew from <1M to 3M
- AI automates repetitive level-1 security work
- Agentic SOC: 82% of analysts concerned about missing real threats owing to "alert fatigue" [9, p. 32]
- Torq: 90% automation of tier-1 SOC tasks, 10× faster response

Education:

- AI-native university: an adaptive academic organism in real time [14, p. 14]
- ASU + OpenAI partnership; SUNY AI literacy requirements
- 86% of students use AI in their studies; 61% of instructors in their teaching [1, p. 393]

5. Where experts agree (consensus)

Trends mentioned in 3+ sources (high consensus)

Table 24. High-consensus trends across sources.

#	Trend	Sources	Key metric
1	**Agentic AI is the key technological trend**	McKinsey (3), BCG, Google (2), Deloitte, PwC, Gartner, Microsoft, a16z, Accenture	62% experimenting (McKinsey)
2	**Adoption–value gap**	McKinsey (2), BCG, Stanford, PwC, Deloitte, Google, Microsoft, Accenture	60–80% with no material impact
3	**Workflow redesign > technology**	McKinsey, BCG, PwC, Deloitte, Google	2.8× for high performers
4	**Data quality is the primary bottleneck**	Google (2), IBM, a16z, Stanford, McKinsey, Alan Turing	70% with difficulties
5	**Hybrid cloud dominates**	Google (2), Deloitte, Gartner, Accenture, Alan Turing	74% prefer
6	**AI security as a separate discipline**	Gartner, Google, Deloitte, IBM, Stanford	39% — top barrier
7	**Regulatory pressure growing**	Stanford, Gartner, Accenture (2), Alan Turing	131 state laws in 2024
8	**CEO/leadership engagement is critical**	McKinsey (2), BCG, PwC, Microsoft	3–12× for leaders
9	**Sovereign AI gaining momentum**	Accenture, Gartner, Stanford, a16z	61% more likely
10	**AI equalises productivity**	Stanford, McKinsey, BCG	Junior +21–40%, Senior +7–16%

Consensus forecasts for 2026–2028

Table 25. Consensus forecasts for 2026–2028.

Forecast	Source consensus	Horizon
>80% of enterprises with GenAI in production	Microsoft/Gartner, Google	2026
Agentic AI 29% of AI value	BCG	2028
Multi-agent systems as standard architecture	Gartner, Google, McKinsey	2027–2028
Hybrid cloud remains dominant	Google, Deloitte, Gartner	2026–2028
Regulatory pressure continues to grow	Stanford, Accenture, Gartner	2026–2030

6. Where experts disagree (debates)

Key divergences

Table 26. Key areas of expert divergence.

Question	Optimists	Sceptics
Agentic AI adoption speed	Google: 52% already in production [9]	Deloitte: only 11% in production [5]
Timeline to mass value	Google: 88% of early adopters with ROI [9]	BCG/McKinsey: 60–80% without impact [3; 4]
Agent autonomy level	a16z: "Internet of agents" by 2028 [14]	PwC: agents do ~50% of human tasks [6]
Open-source vs proprietary	Stanford: gap narrowed to 1.7% [1]	Google Infra: 39% proprietary, 17% open-source [10]
Sovereign AI — defensive vs value	Accenture: $1.5T potential [7]	Accenture: 46% motivated by compliance, 12% — monetisation [7]
Workforce impact	Stanford: productivity +10–45% [1]	McKinsey: 32% expect ≥3% cuts [2]

Detailed analysis of divergences

1. Agentic AI adoption rates

Google Cloud AI Agents [9, p. 7]: 52% of executives with AI agents in production. **Deloitte** [5, pp. 8, 12]: Only 11% of organisations have successfully deployed AI agents in production.

The gap likely reflects differences in both sample composition and terminology. Google surveyed organisations already using GenAI, whereas Deloitte's broader sample included firms at earlier stages, and the two reports appear to define "production deployment" differently.

2. China vs USA gap

Stanford [1, pp. 14, 97–98]: The gap in model quality has nearly closed.

- MMLU: 17.5 p.p. → 0.3 p.p.
- MATH: 24.3 p.p. → 1.6 p.p.
- HumanEval: 31.6 p.p. → 3.7 p.p.

Stanford [1, pp. 252–255]: The investment gap remains enormous.

- USA: $109.1 billion
- China: $9.3 billion (12× smaller)

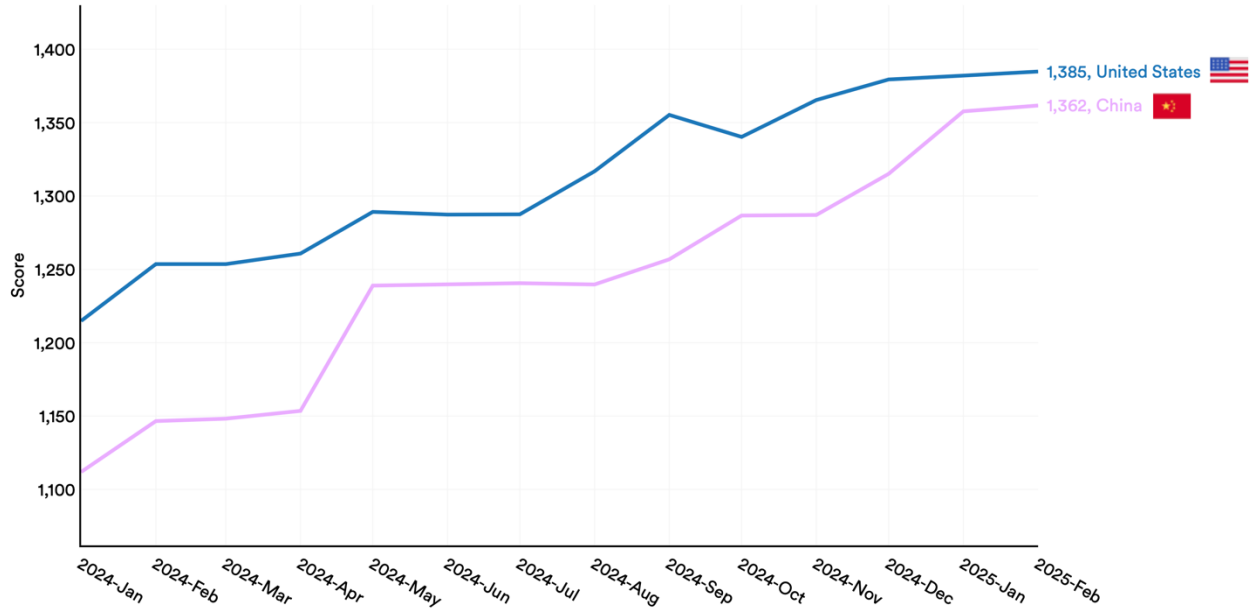

Figure 15. US vs China models on Chatbot Arena. Adapted from: Stanford AI Index [1, p. 97].

Technological parity amid a 12× investment gap is a potentially unstable equilibrium, and one that bears close monitoring through 2026.

3. GenAI project success rate

Microsoft/Gartner [11, p. 20]: "30% of GenAI projects will be abandoned after PoC by 2025 due to poor data quality, inadequate risk controls, escalating costs, or unclear business value."

Google Cloud [9, p. 16]: 88% of early adopters report positive ROI on at least one use case.

The upshot is that not all projects succeed, but organisations find at least one successful use case.

4. AI–human collaboration

Stanford [1, pp. 305–306]: GPT-4 alone (92%) outperforms Physician + GPT-4 (76%).

McKinsey Agentic AI [3, p. 15]: "AI is rarely a standalone solution. Companies extract value when they effectively equip employees with real-world domain experience to interact with AI solutions at the right points."

Taken together, these data suggest the contradiction between theoretical AI effectiveness in isolation and the practical necessity of human–AI collaboration requires further research into workflow design.

7. Metrics and forecasts

7.1 Summary table of quantitative forecasts

Table 27. Summary of quantitative forecasts by domain.

Source	Metric	Current value	Forecast	Horizon
Adoption				
Microsoft/Gartner [11; 13]	Enterprises with GenAI in production	<5% (2023)	>80%	2026
BCG [4]	Companies generating AI value at scale	5%	TBD	2025
Gartner [13]	Organisations with AI-augmented "tiny teams"	—	80%	2030
Gartner [13]	Enterprise apps on AI-native platforms	2%	40%	2030
Agentic AI				
BCG [4]	Agentic AI share of total AI value	17%	29%	2028
Gartner [13]	MAS with narrowly specialised agents	—	70%	2027
Gartner [13]	MAS with multivendor interoperability	—	60%	2028
Infrastructure				
Gartner [13]	Enterprises with hybrid computing architectures	8%	40%	2028
Gartner [13]	Processing protected by confidential computing	—	75%	2029
Google Cloud [10]	Enterprise infrastructure adoption	—	+30%	2026
Google Cloud [10]	Demand for AI-ready data centres (annually)	—	33%	2030
Regulation				
Gartner [13]	Governments with technology sovereignty rules	—	65%	2028
Gartner [13]	Enterprises with geopatriation workloads	—	75%	2030
Accenture/Gartner [7; 13]	Countries with region-specific AI platforms	—	35%	2027

Source	Metric	Current value	Forecast	Horizon
Security				
Gartner [13]	Enterprises with AI security platforms	—	50%+	2028
Gartner [13]	Spending on pre-emptive security solutions	—	50%	2030
Gartner [13]	AI incidents from internal policy violations	—	80%	—
Robotics				
Gartner [13]	Warehouses with robotics/automation	—	80%	2028
Gartner [13]	Top 10 AI vendors with physical AI products	—	5 of 10	2028
Domain-specific				
Gartner [13]	Enterprise GenAI models that are domain-specific	—	60%+	2028
Gartner [13]	GenAI workloads on DSLMs on-prem/on-device	—	30%	2028

7.2 Current key metrics (baseline 2024–2025)

Table 28. Baseline AI metrics, 2024–2025.

Category	Metric	Value	Source
Investment	Corporate AI investment (2024)	$252.3 bn	Stanford [1]
	US private AI investment (2024)	$109.1 bn	Stanford [1]
	GenAI private investment (2024)	$33.9 bn	Stanford [1]
Adoption	Organisations using AI	78–88%	Stanford [1], McKinsey [2]
	Organisations using GenAI	71–79%	Stanford [1], McKinsey [2]
	Fully scaled AI	7%	McKinsey [2]
	AI high performers	5–7%	BCG [4], McKinsey [2]

Category	Metric	Value	Source
Model performance	MMLU (best result)	92.3% (o1-preview)	Stanford [1]
	MATH (best)	97.9% (o3-mini)	Stanford [1]
	SWE-bench Verified	71.7% (o3)	Stanford [1]
	MedQA	96.0% (o1)	Stanford [1]
Cost	Inference costs (GPT-3.5 level)	$0.07/M tokens	Stanford [1]
	Reduction over 18 months	~280×	Stanford [1]
Regulation	US state AI laws (2024)	131	Stanford [1]
	Federal AI regulations (2024)	59	Stanford [1]
	AI incidents (2024)	233	Stanford [1]
Medicine	FDA AI medical devices (2023)	223	Stanford [1]
	NIH grants on AI ethics (2024)	$276M	Stanford [1]

8. Checkpoint questions for end of 2026

This section contains **20 key verifiable hypotheses**, selected from the analysis as the most significant and verifiable. Each question will enable an assessment of the accuracy of 2025 forecasts and will form the basis for a follow-up report in 2027.

Question selection criteria:

- **Verifiability:** availability of publicly accessible data for verification
- **Significance:** reflection of the key trends described in this article
- **Source consensus:** priority given to forecasts with high consensus (3+ sources)
- **Quantitative precision:** ability to provide an unambiguous answer

8.1 Checkpoint questions

Table 29. Checkpoint questions for end-of-2026 assessment.

#	Question	Forecast source	How to verify	Why it matters
Agentic AI — the main trend of the year				
1	Has the share of organisations with AI agents in	Google Cloud AI Agents [9, p. 7] (baseline: 52% in	Industry surveys from McKinsey,	Key indicator of agentic AI's transition

#	Question	Forecast source	How to verify	Why it matters
	production reached 70%?	2025); Deloitte [5, p. 12] (baseline: 11%)	Gartner, Google Cloud 2026	from hype to mainstream
2	What share of AI value is generated by agentic solutions (~23% by interpolation)?	BCG [4, p. 10] (17% → 29% by 2028)	BCG/McKinsey surveys 2026	Validation of the agent value growth trajectory
3	Have mass cross-organisational A2A integrations appeared in production?	Google Cloud AI Agents [9, pp. 18–19]	Public integration announcements, GitHub adoption of A2A protocol, case studies	Indicator of transition to the "cross-platform" phase of multi-agent systems
4	What share of e-commerce transactions is initiated by AI agents?	a16z [14, p. 22]; Accenture Banking [8, p. 11] (35% of consumers willing to delegate choice to AI)	Payment system reports, PayPal data on AP2 protocol	Indicator of consumer trust in autonomous decisions
Overcoming the "Gen AI Paradox"				
5	Has the share of companies with material EBIT impact from AI risen above 50%?	McKinsey [2, p. 6] (baseline: 39% in 2025); McKinsey [3, p. 4] (>80% without impact)	McKinsey Global Survey 2026	The principal indicator of bridging the adoption/value gap
6	Has the 80% threshold of enterprises with GenAI in production been reached?	Microsoft/Gartner [11, p. 5] (forecast for 2026)	Gartner surveys, industry reports	Validation of a consensus forecast (high consensus: 4+ sources)
7	Has the share of "future-built" / AI high performers risen above 10%?	BCG [4, pp. 3–4] (baseline: 5%); McKinsey [2, pp. 9–10] (7%)	BCG/McKinsey surveys 2026	Indicator of the scaling of leadership practices
8	What actual percentage of GenAI projects has been abandoned after PoC?	Microsoft/Gartner [11, p. 20] (forecast: 30% by 2025)	Gartner research, industry surveys	Validation of the "pilot purgatory" hypothesis

#	Question	Forecast source	How to verify	Why it matters
Technological progress and geopolitics				
9	Has China achieved sustained parity with the US on major benchmarks (MMLU, MATH, HumanEval)?	Stanford [1, pp. 97–98] (baseline: 0.3–3.7 p.p. gap in 2025)	LMSYS Chatbot Arena, public leaderboards	Critical geopolitical indicator
10	Does the US vs China gap in private AI investment remain >10×?	Stanford [1, pp. 218, 252–255] (baseline: ~12× in 2024)	Stanford AI Index 2026, CB Insights	Sustainability of technological parity amid an investment gap
11	Has test-time compute (o1-style) become the standard for frontier models?	Stanford [1, pp. 110–112]	Announcements by major AI labs (OpenAI, Anthropic, Google, DeepSeek)	Identification of the dominant paradigm for model development
12	Has the hallucination rate of leading models dropped below 1%?	Stanford [1, p. 171] (baseline: 1.3% in 2025)	HHEM Leaderboard, SimpleQA results	Indicator of readiness for high-stakes applications
Infrastructure and sovereign AI				
13	Does hybrid cloud remain the dominant approach (>70%)?	Google Cloud [10, p. 43] (baseline: 74%)	Cloud adoption surveys, IDC/Gartner	Validation of infrastructure strategy
14	What percentage of organisations apply sovereignty to AI models?	Accenture [7, p. 19] (baseline: 22%); Gartner forecast: 35% of countries with region-specific platforms by 2027 [13]	Accenture survey 2026, Gartner research	Indicator of fragmentation of the global AI landscape
15	Has the forecast of >30% growth in enterprise AI infrastructure adoption been confirmed?	Google Cloud [10, p. 48]	Cloud-provider reports, IDC/Gartner	Verification of the scale of infrastructure transformation

#	Question	Forecast source	How to verify	Why it matters
Regulation — the year of the EU AI Act				
16	Have the main provisions of the EU AI Act entered full force?	Stanford [1, p. 20]; Gartner [13, pp. 21–22]	Official EC publications, enforcement actions	The largest regulatory experiment in AI
17	How many US states have adopted comprehensive AI regulation?	Stanford [1, pp. 340–345] (baseline: 131 laws in 2024, 20 states with deepfake regulation)	NCSL database, state legislature records	Pace of regulatory pressure in the US absent a federal law
Risks and safety				
18	Has the number of AI incidents exceeded 350?	Stanford [1, p. 167] (baseline: 233 in 2024, +56.4% YoY)	AIAAIC Repository	Indicator of the risks of scaling AI
19	Is the 80/20 ratio (internal vs external) for AI incidents confirmed?	Gartner [13, p. 23]	Cybersecurity incident reports, AIAAIC classification	Validation of focus on internal governance vs external threats
Workforce and society				
20	Have expectations of ≥3% workforce reductions in companies using AI been confirmed?	McKinsey [2, pp. 17–18] (32% expected ≥3% reductions)	Bureau of Labor Statistics, industry HR reports, McKinsey survey 2026	Social impact of AI transformation; validation of the "Jevons Paradox"

8.2 Checklist for a follow-up report, Q1 2027

Quantitative thresholds for assessment:

Table 30. Quantitative thresholds for end-of-2026 assessment.

Metric	"Forecast confirmed"	"Partially"	"Not confirmed"
AI agents in production	>60% of organisations	40–60%	<40%
EBIT impact from AI	>50% with material impact	40–50%	<40%
GenAI in production	>75% of enterprises	65–75%	<65%
AI incidents (2026)	>300	250–300	<250
US–China benchmark gap	<2 p.p. across all major benchmarks	2–5 p.p.	>5 p.p.

Metric	"Forecast confirmed"	"Partially"	"Not confirmed"
Hybrid cloud dominance	>70%	60–70%	<60%
Hallucination rate	<1%	1–1.3%	>1.3%

Sources for verification:

Table 31. Primary verification sources by category.

Category	Primary sources	Frequency
Adoption & maturity	McKinsey Global Survey, BCG AI Report, Gartner	Annually (Q2–Q3)
Models & benchmarks	LMSYS Chatbot Arena, Stanford AI Index, Papers with Code	Continuously / Annually
Investment	CB Insights, PitchBook, Stanford AI Index	Quarterly / Annually
Regulation	NCSL, EUR-Lex, Congress.gov	Continuously
Incidents	AIAAIC Repository	Continuously
Workforce	BLS, LinkedIn Economic Graph, McKinsey	Monthly / Annually

9. Conclusion

9.1 Overall trajectory

Analysis of 15 reports from leading research organisations and consulting firms reveals a principal trajectory for AI development in 2026, which can be characterised by the formula: **"From experiments to transformation through agents."**

Three phases of AI evolution

Phase 1 (2022–2024): Discovery and experimentation

- Mass familiarisation with GenAI through ChatGPT and analogues
- A wave of pilot projects and proof-of-concept initiatives
- Focus on horizontal use cases (copilots, chatbots)
- Adoption reached 78% of organisations [1, p. 3]

Phase 2 (2025–2026): Scaling and disillusionment

- "Gen AI Paradox": ~78% use it, but >80% see no material P&L impact [3, p. 4]
- Only 5–7% of organisations attain "future-built" / "AI high performer" status [4, pp. 3–4; 2, pp. 9–10]
- The critical transition from pilots to production is the principal barrier
- Recognition that process reinvention, not task automation, is required

Phase 3 (2026+): Transformation through agentic AI

- AI agents as the catalyst for the transition from experimentation to value
- Agentic AI's share of total AI value: 17% (2025) → 29% (2028) [4, p. 10]
- Transition from "human + tools" to "AI agents + human orchestration"
- Formation of "Frontier Firms" — organisations built around hybrid teams

Key paradigm shifts

Table 32. Key paradigm shifts: 2024 vs 2026.

Aspect	Previous (2024)	Emerging (2026)
Approach to AI	Add-on to existing processes	AI-first process redesign
Employee role	User of AI tools	Orchestrator of agent teams
Unit of work	Task	Business process (workflow)
Success metric	Adoption rate	P&L impact
Infrastructure	Cloud-first	Hybrid (cloud + on-prem + edge)
Governance	Declarative	Operationalised
Competition	Technological	Organisational (speed of transformation)

9.2 Ten consensus conclusions

The analysis identified **high consensus** (mention in 3+ sources) on the following points:

1. **Agentic AI is the key trend of 2026** — 12 of 15 sources confirm the central role of AI agents.
2. **The adoption–value gap is critical** — BCG, McKinsey (both reports), PwC, Deloitte, Google, Microsoft, Stanford, Accenture are unanimous: 60–80% of organisations see no material impact.
3. **Workflow/process redesign matters more than technology** — BCG (the 10-20-70 rule), McKinsey (2.8× for leaders), PwC (80% of value from work redesign; the difference with BCG's 10-20-70 split reflects classification of data, not substance).
4. **CEO/C-level engagement is a necessary condition for success** — BCG (12× gap), McKinsey (3×), PwC, Microsoft point to the criticality of leadership buy-in.
5. **Data quality is the primary bottleneck for scaling** — Google Infrastructure (70% experience difficulties), IBM, Stanford, McKinsey, Accenture, Alan Turing, a16z (80% of data unstructured).
6. **Hybrid cloud is becoming the dominant architecture** — Google (74%), Deloitte, Gartner confirm the transition from cloud-first to hybrid.
7. **AI security requires a separate discipline** — Gartner, Google, Deloitte, IBM, Stanford identify AI safety as an emerging field.

8. **Regulatory pressure is growing exponentially** — Stanford, Gartner, Accenture (both reports), and the Alan Turing Institute document a sharp acceleration: US state-level AI laws rose from 1 (2016) to 131 (2024), and 59 federal regulations were adopted in 2024 alone.
9. **Sovereign AI is becoming a strategic priority** — Accenture, Gartner, Stanford, and a16z concur: 61% of organisations are more inclined towards sovereign technologies owing to geopolitical turbulence, and the sovereign AI infrastructure market is estimated at $1.5 trillion.
10. **AI equalises productivity across skill levels** — Stanford, McKinsey, and BCG present converging evidence that AI disproportionately benefits less experienced workers (productivity gains of 21–40% for junior specialists vs 7–16% for senior specialists), effectively narrowing the expertise gap.

9.3 What this means for business

For company leaders (CEO/Board)

Imperative for action: The time for experimentation is ending. Companies that do not transition from pilots to scaling in 2026 risk finding themselves among the "60% laggards" with a growing deficit.

"The most important point is that time is short. The technology is advancing fast, making catching up more difficult with each passing week." — BCG [4, p. 17]

Key decisions for 2026:

- Selecting 3–5 processes for wholesale transformation (not optimisation)
- Creating a centralised AI hub / AI studio
- Transitioning from adoption metrics to P&L impact metrics
- Investing in upskilling: 50% of the workforce among leaders vs 20% among laggards [4, p. 15]

For CIO/CTO

Infrastructure challenge: Traditional cloud infrastructure is approaching an inflection point at which it will be unable to cope with the demands of AI workloads [10, p. 48].

Priorities:

1. Three-tier hybrid architecture (cloud + on-prem + edge)
2. Agent-native infrastructure for a new type of workload
3. Data readiness and governance as the foundation
4. AI security platforms as a separate layer

For CDO/Chief Data Officers

New role: From data management to enabling AI transformation.

"Enterprise AI at scale is finally within reach. The technology is ready — as long as organisations can feed it the right data." — IBM CDO Study [12, p. 2]

Focus:

- Resolving data silos (each AI project = 6–12 months of data cleansing)
- "Bring AI to data" rather than centralising data
- Proprietary data as a competitive advantage (78% of CDOs) [12]
- AI agents as data consumers (83% of CDOs see benefits) [12]

For HR/CHRO

Workforce transformation: Building hybrid teams of "people + agents."

New realities:

- "Hourglass" structure for knowledge work: junior + senior, fewer mid-tier [6, pp. 6–7]
- New roles: agent orchestrators, AI workflow designers, governance leads
- Half-life of a skill = 4 years (in tech = 2 years) — the necessity of continuous learning
- 84% of employees want a greater focus on AI, but only 29% feel supported [9, p. 40]

9.4 What this means for society

Positive prospects

1. **Democratisation of expertise:** AI is narrowing the gap between junior and senior specialists (21–40% vs 7–16% productivity gain) [1, pp. 219, 268–270]
2. **Breakthroughs in medicine:**
   - MedQA benchmark: 96% (o1) — expert-physician level
   - FDA-approved AI devices: 6 (2015) → 223 (2023) — exponential growth
   - Ambient AI scribes: −20 min/day in documentation, −26% physician burnout
3. **Scientific discoveries:**
   - AlphaFold 3: modelling of protein–DNA–RNA–ligand interactions
   - ESM3: creation of a protein whose evolution would have required nature 500 million years
   - Nobel Prizes 2024 for AI research
4. **Safety:**
   - Waymo: 81% fewer airbag deployments, 77% fewer injuries
   - AI for cybersecurity: 50% of organisations already using it [5, p. 12]

Challenges and risks

1. **Labour market:**
   - 32% of companies expect ≥3% reductions [2, pp. 17–18]
   - However: "Jevons Paradox" — automation may increase demand for specialists
   - 60% of people believe AI will change their work within 5 years
2. **Inequality in attitudes to AI:**
   - China: 83% optimism vs USA: 39% vs Netherlands: 36% [1, p. 401]
   - 61% of Americans fear driverless cars
3. **AI incidents:**
   - 233 incidents in 2024 (+56.4% YoY) — a record [1, p. 167]
   - The upward trend is set to continue
4. **Contraction of the data commons:**
   - Full robots.txt restrictions: 10% (2017) → 48% (2024) [1, pp. 194–195]

    - Epoch AI forecast: the current data stock will be exhausted between 2026 and 2032
5. **Geopolitical tensions:**
    - ~70% of leading AI models originate in the US, ~25% in China
    - 61% of organisations more inclined towards sovereign solutions owing to geopolitics [7, p. 5]

9.5 Final forecast for end of 2026

Based on the synthesis of all sources, by end of 2026 with high probability:

Will happen (high confidence, consensus of 5+ sources):

- ✅ Agentic AI becomes mainstream (>50% of enterprises in production)
- ✅ The gap between AI leaders and laggards widens
- ✅ >80% of enterprises use GenAI APIs/models in production
- ✅ Hybrid cloud becomes the dominant architecture for AI
- ✅ EU AI Act enters full force
- ✅ The number of AI incidents exceeds 300

Likely to happen (moderate confidence):

- 🔶 AI models achieve US–China parity across all major benchmarks
- 🔶 Test-time compute becomes the standard for frontier models
- 🔶 The first mass A2A integrations emerge
- 🔶 "Agent orchestrator" roles are formalised in corporate structures

Unlikely (low confidence or source divergence):

- ❌ A federal AI law is passed in the US
- ❌ Trust in self-driving exceeds 30% in the US
- ❌ >50% of organisations close the adoption/value gap

"AI is no longer just a story of what's possible — it's a story of what's happening now and how we are collectively shaping the future of humanity." — Stanford AI Index 2025 [1, p. 2]

A follow-up assessment using the checkpoint questions in Section 8 is recommended for Q1 2027.

10. Appendices

10.1 Summary table: Source → Topics → Key figures

Table 33. Cross-source coverage matrix.

Source	Agentic AI	AI maturity	Investment	Infrastructure	Regulation	Medicine / Science
Stanford AI Index [1]	RE-Bench: AI 4× better than humans (2h), 2× worse (32h)	78% adoption	$109.1B US, $33.9B GenAI	Inference cost −280×	59 federal regulations, 131 state laws	MedQA 96%, FDA 223 devices
BCG [4]	17%→29% agentic value	5% future-built, 60% laggards	+120% leader investment	—	—	—
McKinsey State of AI[2]	62% experimenting, 23% scaling	38% scaling/scaled, 7% fully scaled	—	—	—	—
McKinsey Agentic [3]	~90% vertical use cases in pilot	>80% without EBIT impact	$2.6–4.4T GenAI potential	—	—	—
a16z [14]	5,000+ subtasks per 1 goal	—	—	80% of data unstructured	KYA (Know Your Agent)	—
PwC [6]	~50% of tasks agents can do	—	—	—	60% RAI boosts ROI	—
Deloitte [5]	11% agents in production	—	—	Three-tier hybrid	AI for cyber + cyber for AI	—
IBM CDO Study [12]	83% benefits > risks	13% IT budget on data	—	75% data platform cross-silo	—	—
Google AI Agents [9]	52% agents in production	88% positive ROI for early adopters	—	—	82% SOC alert fatigue	—

Source	Agentic AI	AI maturity	Investment	Infrastructure	Regulation	Medicine / Science
Microsoft CIO [11]	81% agent integration 12–18 months	24% organisation-wide	—	—	—	—
Accenture Sovereign[7]	—	—	$1.5T sovereign AI market	55% mix global+local	61% lean sovereign	—
Accenture Banking[8]	10× bank concept	—	$289B benefit top-200 banks	~70% IT on legacy	—	—
Gartner [13]	1,445% MAS enquiry growth	80% tiny teams by 2030	—	40% hybrid by 2028	80% AI incidents = internal	—
Google Infrastructure[10]	—	98% use GenAI	—	74% hybrid cloud	62% security/privacy concerns	—
Alan Turing Institute [15]	—	—	—	—	GenAI MRM framework; SR 11-7 / SS1/23	—

10.2 Glossary of key terms

Table 34. Glossary of key terms.

Term	Definition	Source
Future-built companies	Companies achieving value from AI at scale (5% per BCG)	BCG [4, pp. 3–4]
Frontier Firm	An organisation built around hybrid teams of humans and AI agents	Microsoft [11, pp. 6–7]
Agentic AI Mesh	A composable, distributed, vendor-agnostic architecture for managing AI agents	McKinsey [3, pp. 17–19]
10-20-70 rule	10% technology, 20% data, 70% people/organisation/processes — the distribution of barriers	BCG [4, p. 18]
Tiny teams	Teams of 2 people + AI agents replacing large dev teams	Gartner [13, p. 5]
10× bank	A concept in which one employee manages a team of AI co-workers	Accenture [8, p. 32]
Sovereign AI	The ability of countries to independently develop and deploy AI on their own infrastructure, data, models, and talent	Accenture [7, p. 4]
Geopatriation	Transfer of workloads from global hyperscale clouds to sovereign/local environments	Gartner [13, pp. 25–26]
Neoclouds	AI-native infrastructure providers (Nebius, CoreWeave, Lambda)	Accenture [7, pp. 15–16]
Agent2Agent (A2A)	An open protocol for inter-organisational agent communication	Google Cloud [9, p. 18]
Model Context Protocol (MCP)	A standard for connecting LLMs to databases and tools	Google Cloud [9, p. 18]
Test-time compute	Inference-time reasoning (o1-style), in which the model "thinks" longer for a better result	Stanford [1, pp. 110–112]
DSLMs (Domain-specific LMs)	Language models for specific industries (finance, healthcare, HR)	Gartner [13, pp. 14–15]
Data commons	Data on the internet available for AI training	Stanford [1, pp. 194–195]
Vibe coding	Software creation by people without technical expertise via AI	Gartner [13, p. 5]; PwC [6, pp. 9–10]

Conflict of interest: The author declares no conflict of interest.

Declaration of generative AI and AI-assisted technologies in the writing process: During the preparation of this work, the author used an AI tool, namely Claude, in order to correct grammatical mistakes and edit the language professionally. After using this tool/service, the author reviewed and edited the content as needed.

ORCID

Yury Korolev https://orcid.org/0000-0001-8316-0058

Decision Impact Summary

This analytical review addresses a question at the heart of human–AI decision-making: how should organisations decide when the technology they rely on for decisions is itself changing faster than the governance structures around it? By synthesising 15 industry reports (combined survey base: >13,000 respondents), the article maps the current state of enterprise AI adoption and surfaces a central paradox — mass adoption without proportionate value — that directly concerns any practitioner designing human–AI workflows.

For the readership of this journal, three implications stand out. First, the finding that workflow redesign, not technology capability, explains the 2.8× performance gap between AI leaders and laggards reinforces the journal's premise that the design of the human–AI interface is the binding constraint on decision quality. Second, the emergence of agentic AI — systems that plan, act, and learn with limited human oversight — fundamentally alters the locus of decision authority in organisations; the article's evidence on adoption rates, governance gaps, and divergent expert forecasts provides the empirical grounding that governance frameworks will need. Third, the 20 checkpoint questions and quantitative thresholds in Section 8 offer a reusable methodology for evaluating forecast accuracy against real-world outcomes — a contribution to the broader challenge of calibrating trust in AI-generated predictions.

Practitioners engaged in AI governance, organisational transformation, or human–AI teaming will find in this review both a data-rich baseline for 2025–2026 and a structured framework for revisiting its claims in 2027.

References

[1] Stanford University HAI. Stanford AI Index Report 2025. (April 2025).
URL: https://hai.stanford.edu/ai-index/2025-ai-index-report

[2] McKinsey & Company / QuantumBlack. The State of AI — Global Survey 2025. (November 2025). URL: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai/

[3] McKinsey & Company / QuantumBlack. Seizing the Agentic AI Advantage. (June 2025).
URL: https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage

[4] Boston Consulting Group. The Widening AI Value Gap. (September 2025).
URL: https://www.bcg.com/publications/2025/are-you-generating-value-from-ai-the-widening-gap

[5] Deloitte. Tech Trends 2026. (December 2025).
URL: https://www.deloitte.com/us/en/about/press-room/deloitte-tech-trends-2026.html

[6] PwC. 2026 AI Business Predictions. (February 2026).
URL: https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-predictions.html

[7] Accenture. Sovereign AI — Own Your AI Future. (July–August 2025).
URL: https://www.accenture.com/content/dam/accenture/final/accenture-com/document-4/Sovereign-AI-Report.pdf

[8] Accenture. Banking Top Trends FY26. (2026).
URL: https://www.accenture.com/content/dam/accenture/final/industry/banking/document/Banking-Top-Trends-FY26-Report-Final.pdf

[9] Google Cloud. AI Agent Trends 2026. (2025).
URL: https://cloud.google.com/resources/content/ai-agent-trends-2026

[10] Google Cloud. State of AI Infrastructure. (2025).
URL: https://cloud.google.com/resources/content/state-of-ai-infrastructure

[11] Microsoft. The Strategic CIO's Generative AI Playbook. (2025).
URL: http://aka.ms/Copilot/CIOplaybook

[12] IBM Institute for Business Value. The AI Multiplier Effect (CDO Study). (Q3 2025).
URL: https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/2025-cdo

[13] Gartner. Top 10 Strategic Technology Trends 2026. (2025).
URL: https://www.gartner.com/en/articles/gartner-top-strategic-technology-trends-2026

[14] Andreessen Horowitz (a16z). Big Ideas 2026. (December 2025).
URL: https://a16z.com/newsletter/big-ideas-2026-part-1/#agent-native-infrastructure-becomes-table-stakes

[15] Alan Turing Institute / PAIF. GenAI Model Risk Management and Governance in Financial Services: Principles and Practice. (2025).
URL: https://www.turing.ac.uk/news/publications/genai-model-risk-management-and-governance-financial-services-principles-practice

AI DECISIONS HUMAN–AI DECISION-MAKING SYSTEMS

VOLUME 2 — ISSUE 1
MARCH 2026
QUARTERLY OPEN ACCESS ACADEMIC JOURNAL

ISSN: 2978-5669
DOI PREFIX: 10.65114

Published by AI Decisions Ltd,
London, United Kingdom
journal@aidecisions.ai
journal.aidecisions.ai