10 Ethical AI Issues in 2026: From Principles to Dated Rules

Edition: 2026/27 · English · v1.0 · 22 October 2026

The ten issues we described in 2024 are still with us, but most now come with a legal date, a measurable control or published evidence. Here is what has changed and what to do about each.

From ethics checklist to schedule

In 2024 this article described AI ethics as a set of principles: security, privacy, explainability, bias mitigation and calibration. Those principles still hold. What has changed is that many of them now carry application dates. The EU AI Act, as amended by the Digital Omnibus (Regulation (EU) 2026/1744, in force since 27 July 2026), sets a clear timetable:

- **Since 2 February 2025:** prohibited practices apply, including social scoring and emotion recognition at work and in education (with medical and safety exceptions).
- **Since 27 July 2026:** a revised AI-literacy duty. Providers and deployers must take measures to support AI literacy; no specific level is required of any individual.
- **Since 2 August 2026:** Article 50 transparency duties, such as telling people they are interacting with an AI system.
- **2 December 2026:** machine-readable marking of synthetic content for generative systems already on the market, and a new prohibition on AI systems that generate non-consensual intimate imagery or child sexual abuse material.
- **2 December 2027:** high-risk duties for Annex III uses, including employment, credit, insurance, education and public services.
- **2 August 2028:** high-risk duties for AI in regulated products (Annex I).

Outside the EU the dates differ, but the direction is the same. Korea's AI Basic Act has applied since 22 January 2026, and Colorado and California rules on automated decisions apply from 1 January 2027. Labelling of AI-generated content is the one rule most regimes share.

The technology underneath has changed too

Ethical controls now have to keep pace with a faster stack. Enterprises are multi-model by default: 81% of large companies use three or more model families (a16z, January 2026). Models also change quickly. OpenAI retires the GPT-4o and GPT-4.1 snapshots on 14 January 2027, and Anthropic retires Claude Sonnet 4.5 on 30 November 2026, about 14 months after its release. Every model swap is a moment when a bias test, an accuracy check or a documented decision can silently go out of date.

The ten issues, updated

1. Copyright and training data. The legal picture is clearer, though not settled. General-purpose AI model obligations under the AI Act have applied since 2 August 2025, and the Commission can fine providers from 2 August 2026. The GPAI Code of Practice (July 2025), which includes a copyright chapter, has signatories including Amazon, Anthropic, Google, IBM, Microsoft, Mistral AI and OpenAI. In the UK, the High Court ruled in *Getty Images v Stability AI* on 4 November 2025; Getty had dropped its primary infringement claims during trial and won only "extremely limited" trademark points. *Practice:* ask suppliers whether they signed the Code, and check rights and licence terms for each tool that generates content.

2. Biometric identification. Real-time remote biometric identification in public spaces for law enforcement is prohibited in the EU, with narrow exceptions, as is untargeted scraping of facial images. *Practice:* keep any biometric

use in your AI inventory and confirm its legal basis before deployment.

3. Deepfakes and synthetic media. Machine-readable marking becomes mandatory in the EU for generative systems already on the market from 2 December 2026. China has required explicit and metadata labels since 1 September 2025, and India since 20 February 2026. Our companion article on synthetic faces and voices covers this in detail.

4. AI-enabled fraud and cyber attacks. Cyber incidents are the top business risk for the fifth year running (42% of responses), and AI jumped from tenth to second place (32%) in the Allianz Risk Barometer 2026. Shadow AI is involved in breaches at about one firm in five (IBM, via Help Net Security). *Practice:* make shadow AI visible through an inventory and give staff a sanctioned alternative.

5. Accountability when systems change. Fast model cycles mean responsibility must sit with a named owner, not with the original project team. *Practice:* record which model version each system uses and re-run your quality and fairness checks before every switch.

6. Work and workers. Measured effects are real but specific. In customer support, a generative AI assistant raised issues resolved per hour by 14% on average and 34% for less experienced agents (NBER). In a field experiment with about 70,000 applicants, AI-led interviews raised job starts by 18% (Erasmus). Recruitment and worker management are Annex III uses from 2 December 2027, and Italy's AI law already requires employers to inform staff when AI is used at work. *Practice:* involve employees and works councils from the start, and decide in advance where released capacity goes.

7. Accuracy in legal and justice contexts. Studies report hallucination rates of 17–33% in legal research tools (LLRX, February 2026). Administration of justice is an Annex III area. *Practice:* require source citations and human review for any output that affects rights.

8. Healthcare. 81% of US physicians now use AI, mostly for documentation (AMA, March 2026). AI in medical devices falls under Annex I, with AI Act duties from 2 August 2028. *Practice:* choose tools on documented validation and data-handling facts, and measure clinician time saved.

9. Automated decisions in public services and credit. US federal agencies reported 3,611 AI use cases in 2025, 445 of them high-impact, up from 1,757 in 2024 (OMB inventory). Creditworthiness assessment and access to public benefits are Annex III uses. Colorado's SB 26-189 requires notice, an explanation within 30 days of an adverse decision, and human review where commercially reasonable, from 1 January 2027. *Practice:* map every automated decision about people and prepare notices and an appeal route.

10. Surveillance and social scoring. Social scoring is banned in the EU, as is emotion recognition at work and in education, both since February 2025. *Practice:* check monitoring and analytics tools against these prohibitions, including features added in supplier updates.

What good practice looks like now

Three practices make the ten issues manageable:

- **One AI inventory.** Record every system, its model version, its owner, its intended use and the earliest deadline that applies to it in each jurisdiction where you operate.
- **Documented supplier facts.** Ask for hosting region, training use of your data, retention and exit terms in writing. In the Good AI Prism study, only 10 of 107 customer-service AI offerings publicly documented both an EU hosting option and limits on training with customer data: 8 state no training and 2 train only with opt-in consent (Wave 1; provider information reviewed 19–20 October 2026).
- **Measured pilots.** Agree a baseline, a success threshold and a stop condition before the pilot starts. Standards help here: 271 ISO/IEC 42001 AI management-system certificates had been issued worldwide by the end of 2025 (ISO Survey, via DQS).

> **What to decide now** - Build or update your AI inventory and map each system to its earliest deadline: 2 December 2026, 1 January 2027 or 2 December 2027. - Name an owner for each system who re-tests quality and fairness whenever the model version changes. - Write the four supplier facts (hosting, training use, retention, exit) into every new contract and renewal.

The full evidence base, including the regulatory timetable for more than 30 economies, is in the Good AI Prism Global Market Study 2026/27 at ethority.ai/ai-prism-study.

Sources

- Regulation (EU) 2026/1744 (Digital Omnibus on AI): https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng
- Lewis Silkin, The Digital Omnibus on AI enters into force today (27 Jul 2026): https://www.lewissilkin.com/en/insights/2026/07/27/the-digital-omnibus-on-ai-enters-into-force-today-102nedo
- nicfab, Digital Omnibus on AI in the Official Journal: https://www.nicfab.eu/en/posts/digital-omnibus-ai-official-journal/
- Regulation (EU) 2024/1689 (AI Act): https://eur-lex.europa.eu/eli/reg/2024/1689/oj
- European Commission, GPAI Code of Practice: https://digital-strategy.ec.europa.eu/policies/contents-code-gpai
- Cleary Gottlieb, Getty Images v Stability AI (Nov 2025): https://www.clearygottlieb.com/news-and-insights/publication-listing/uk-high-court-issues-landmark-ruling-in-getty-images-v-stability-ai
- a16z, enterprise AI survey (Jan 2026): https://a16z.com/leaders-gainers-and-unexpected-winners-in-the-enterprise-ai-arms-race/
- OpenAI, deprecations: https://developers.openai.com/api/docs/deprecations
- Anthropic, model deprecations: https://platform.claude.com/docs/en/about-claude/model-deprecations
- Allianz Risk Barometer 2026: https://commercial.allianz.com/news-and-insights/news/allianz-risk-barometer-2026.html
- Help Net Security, IBM Cost of a Data Breach 2025: https://helpnetsecurity.com/2025/08/04/ibm-cost-data-breach-report-2025
- NBER Working Paper 31161: https://www.nber.org/papers/w31161
- Erasmus University, AI job interviews field experiment: https://www.eur.nl/en/erim/events/ai-automation-job-interviews-natural-field-experiment-2025-10-28
- Norton Rose Fulbright, Italy Law No. 132/2025: https://www.nortonrosefulbright.com/en/knowledge/publications/9bfedfea/italy-enacts-law-no-132-2025-on-artificial-intelligence-sector-rules-and-next-steps
- LLRX, hallucinations in legal research (Feb 2026): https://www.llrx.com/2026/02/what-the-science-says-about-hallucinations-in-legal-research/
- AMA, AI use among physicians (Mar 2026): https://ama-assn.org/press-center/ama-press-releases/ama-ai-usage-among-doctors-doubles-confidence-technology-grows
- OMB, 2025 Federal Agency AI Use Case Inventory: https://github.com/ombegov/2025-Federal-Agency-AI-Use-Case-Inventory
- Colorado SB 26-189: https://leg.colorado.gov/bills/sb26-189
- ITA, South Korea AI Basic Act: https://www.trade.gov/market-intelligence/south-korea-ai-basic-act
- Loeb & Loeb, China AI labelling measures: https://www.loeb.com/en/insights/publications/2025/03/chinas-ai-labeling-measures-and-mandatory-national-standards-take-effect-september-1
- Hogan Lovells, India synthetic-content labelling (Feb 2026): https://www.hlc.com/en/publications/india-introduces-mandatory-labelling-for-ai-and-3hour-takedown-for-illegal-content
- DQS, ISO Survey 2025: https://www.dqsglobal.com/es/acerca-de/sala-de-prensa/iso-survey-2025-las-certificaciones-iso-crecen-un-11-en-el-mundo-y-un-28-en-espana
- ETHORITY, Good AI Prism Global Market Study 2026/27: https://ethority.ai/ai-prism-study

This article is not legal advice. Dates are application dates as published by October 2026.

Updated edition of the article first published on ethority.net in 2024; the original remains archived.

How to Navigate 961 AI Offerings: A Buyer's Guide to the 2026 AI Landscape

Edition: 2026/27 · English · v1.0 · 22 October 2026

A list of 100 tools no longer helps anyone choose; a map of where AI does work, what data it handles and what technology it rests on does.

Why we replaced the list

In 2024 we published a guide to more than 100 AI tools, sorted into categories such as natural language processing, computer vision, machine-learning frameworks and chatbots. Many of the names in it were developer libraries; several have since been renamed, merged or retired. The market has also grown too large for a list. For the Good AI Prism Global Market Study 2026/27 we mapped 998 market entries, 961 of them AI offerings, each recorded with a dated source from the provider's own pages.

The market also changes faster than any list. Within a few weeks of research we found 34 renames and 26 ownership changes, for example Acrolinx becoming Markup AI and LivePerson moving to SoundHound AI. So this edition is a method, not a catalogue: how to move from a business need to a defensible shortlist.

The map: three rings

The Good AI Prism maps the AI market from the outside in.

Outer ring: 24 business areas, where AI does work. - *14 functions:* Customer Service; Assistants & Copilots; Knowledge & Search; Documents, Workflow & Agents; Sales & CRM; Marketing & Content; HR & Recruiting; Finance & Accounting; Legal & Compliance; IT & Security Operations; Software Development; Data & Analytics; Design, Image & Video; Meetings, Voice & Language. - *10 industries:* Health & Life Sciences; Financial Services & Insurance; Manufacturing & Industrial; Retail & Consumer; Energy & Utilities; Mobility & Logistics; Public Sector; Education; Media & Entertainment; Science & Research.

Middle ring: 8 data types, the kind of data an offering handles: language, voice, audio, image, video, tables, code and multimodal.

Inner ring: 6 technology layers, the technology underneath: models, chips and compute, clouds, data, agents and trust.

An offering can sit in more than one area; the 998 entries hold 1,037 placements on the map.

Step 1: Start with the job, not the tool

Pick one process in one business area and measure it before you look at suppliers: volume, cost per case, quality and cycle time. Measured gains are real but task-specific. In customer support, a generative AI assistant raised issues resolved per hour by 14% on average and 34% for less experienced agents (NBER). In software work, measured effects range from 19% slower (METR) to 21% faster (Google RCT). A supplier's productivity figure is not your baseline.

Step 2: Name the data type

The middle ring narrows the field quickly. A voice agent, a document workflow and a coding assistant handle different data, carry different risks and need different checks. Voice and meeting tools, for example, raise consent

questions: 65% of US workers who met with an AI notetaker said consent was inconsistent or missing (survey via KRDO).

Step 3: Understand the layers underneath

Most business buyers purchase applications, but each application rests on layers that change independently.

- **Models.** The frontier now has three closed flagships: OpenAI GPT-6 Astra, Anthropic Claude Fable 5.1 and Google Gemini 4 Argon. Open-weight options include Mistral Large 4, previewed in October 2026 with European deployment and on-premises options. Prices span roughly 100 times, from USD 0.10 to USD 10 per million input tokens, and 81% of large enterprises already use three or more model families (a16z). Models also retire: OpenAI shuts down its GPT-4o and GPT-4.1 snapshots on 14 January 2027.
- **Chips and compute.** Capacity stays tight. ASML expects its planned 2027–28 supply "will not fully meet customer demand", and big-tech capital spending is set to rise a further 75% in 2026 (IEA). At the same time, inference prices per unit keep falling steeply (Epoch AI).
- **Clouds and data.** Hosting region, subprocessors and data use decide whether an offering is usable for a European buyer at all.
- **Agents.** Spending on AI agents and assistants is forecast to grow from USD 29.2bn in 2026 to USD 65.5bn in 2027 (Gartner), yet Gartner also expects more than 40% of agentic AI projects to be cancelled by the end of 2027.
- **Trust.** Documentation, certifications and their scope, oversight and incident handling.

Verified relations between offerings show applications resting on a few hubs: Zendesk, Salesforce and the Model Context Protocol on the application side; OpenAI, Anthropic, Amazon Bedrock and Hugging Face underneath. A single model or platform change can move many products at once, so ask every supplier which model and cloud they depend on.

Step 4: Ask for the deciding facts in writing

The facts a careful buyer needs are mostly not on public pages. In the Customer Service area, 10 of 107 offerings publicly documented both an EU hosting option and limits on training with customer data: 8 state no training and 2 train only with opt-in consent (Good AI Prism Market Study, Wave 1; provider information reviewed 19–20 October 2026). They include BOTfriends X, Dynamics 365 Customer Service, PolyAI, sipgate AI Agents and Typewise. For the other 97, at least one fact was not documented. That is not the same as unavailable: it means you have to ask.

Two further patterns run across the market. Four in five offerings we added in 2026 are described only on marketing pages, not in technical documentation. And prices come in at least seven units: per resolution, per seat, per conversation, per minute, credits, plan tier and quote only.

Step 5: Compare on cost per accepted outcome

Because prices come in different units, convert every option to one measure: total monthly cost (usage or licence, amortised integration, monitoring and training, human review, handling of failed cases) divided by the outcomes your business accepts. In an illustrative support example with 10,000 cases a month, €11,000 of total cost and a 40% acceptance rate give €2.75 per accepted outcome, which you can set against today's cost per resolved case.

Step 6: Check the rules for your intended use

Whether an offering is high-risk depends on what you use it for. In the EU, Article 50 transparency duties apply since 2 August 2026; high-risk duties for Annex III uses such as recruitment, credit and education apply from 2 December 2027 under Regulation (EU) 2026/1744. Colorado and California duties on automated decisions apply from 1 January 2027.

> **What to decide now** - Choose one process, place it on the map (business area, data type, layer) and measure its baseline before any demo. - Send every shortlisted supplier the same written questions: hosting region, training use, billing unit, model and cloud dependencies, exit terms. - Compare options on cost per accepted outcome at your own volume, with a stop condition agreed in advance.

The Good AI Prism, the full taxonomy and the Wave 1 findings are available at ethority.ai/ai-prism-study, where providers and readers can also submit corrections.

Sources

- ETHORITY, Good AI Prism Global Market Study 2026/27 (Wave 1, research edition 20 Oct 2026): https://ethority.ai/ai-prism-study
- NBER Working Paper 31161: https://www.nber.org/papers/w31161
- METR study, via Simon Willison: https://simonwillison.net/2025/Jul/12/ai-open-source-productivity/
- Google RCT on developer productivity: https://arxiv.org/html/2410.12944v2
- KRDO / Stacker, AI notetaker survey (Jul 2026): https://krdo.com/stacker-business-economy/2026/07/16/ai-notetakers-have-sat-in-on-1-in-3-us-workers-meetings-but-only-a-third-say-they-were-asked-first/
- OpenAI, GPT-6 Astra: https://openai.com/index/gpt-6-astra/
- Anthropic, models overview: https://platform.claude.com/docs/en/about-claude/models/overview
- Google Cloud, Gemini 4 Argon: https://cloud.google.com/blog/products/ai-machine-learning/gemini-4-argon-available-on-google-cloud
- Mistral AI, Mistral Large 4: https://mistral.ai/news/mistral-large-4/
- OpenAI, GPT-6.1 Sol and Luna pricing: https://openai.com/index/introducing-gpt-6-1-sol/
- a16z, enterprise AI survey (Jan 2026): https://a16z.com/leaders-gainers-and-unexpected-winners-in-the-enterprise-ai-arms-race/
- OpenAI, deprecations: https://developers.openai.com/api/docs/deprecations
- ASML, Q3 2026 results (6-K): https://www.sec.gov/Archives/edgar/data/0000937966/000162828026065813/pressreleasefinancialresul.htm
- IEA, data-centre electricity 2025: https://www.iea.org/news/data-centre-electricity-use-surged-in-2025-even-with-tightening-bottlenecks-driving-a-scramble-for-solutions
- Epoch AI, inference price trends: https://epoch.ai/data-insights/ai-inference-price-trends
- Gartner, worldwide AI spending forecast (16 Sep 2026): https://gartner.com/en/newsroom/press-releases/2026-09-16-gartner-forecasts-worldwide-ai-spending-to-grow-49-point-5-percent-in-2026
- Gartner, agentic AI project cancellations: https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027
- Regulation (EU) 2026/1744: https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng
- Colorado SB 26-189: https://leg.colorado.gov/bills/sb26-189
- California Privacy Protection Agency, ADMT regulations: https://privacy.ca.gov/2025/09/california-finalizes-regulations-to-strengthen-consumers-privacy/

Updated edition of the article first published on ethority.net in 2024; the original remains archived.

AI Growth Facts 2026: What the Evidence Shows Across Industries

Edition: 2026/27 · English · v1.0 · 22 October 2026

Spending on AI grows by half in 2026, use is broad and measured returns are thinner; these are the sourced figures that matter for planning 2027.

A new rule for this edition

The 2024 version of this article collected market-size projections from secondary aggregators: trillions added to GDP by 2030, millions of new jobs, market sizes per industry. Many of those figures could not be traced to a primary study, and several described years that have now passed. This edition uses only figures with a named publisher, a date and, where relevant, a sample. Survey results are shown separately; they are not stages of one funnel.

Spending: up by half, mostly into infrastructure

- Worldwide AI spending is forecast at **US$2.67 trillion in 2026, up 49.5%** from US$1.79 trillion in 2025, and US$3.64 trillion in 2027. AI infrastructure accounts for US$1.48 trillion of the 2026 total (Gartner, September 2026).
- Within that total, spending on **AI agents and assistants** grows from US$29.2 billion in 2026 to US$65.5 billion in 2027 (Gartner).
- Enterprise generative AI spending reached **US$37 billion in 2025**, 3.2 times the 2024 level (Menlo Ventures, December 2025).
- European AI spending is forecast to reach **US$290 billion by 2029**, growing 33.7% a year (IDC, April 2026).
- Germany's AI market is forecast at **€28.7 billion in 2026 (+48%)** (Bitkom/IDC).

Investment and capacity

- US private AI investment reached **US$285.9 billion in 2025**, against US$20.9 billion for Europe and US$12.4 billion for China (Stanford AI Index 2026). HAI notes that state-backed Chinese funds are not captured.
- Global AI compute capacity reached about **17 million H100-equivalents** in Q4 2025 (Stanford AI Index 2026).
- NVIDIA reported **US$89.0 billion in data-centre revenue** for Q2 of fiscal 2027, up 117% (August 2026).
- Big-tech capital spending exceeded US$400 billion in 2025 and is set to rise **a further 75% in 2026** (IEA, April 2026).

Adoption: broad, but measured differently

- In a global survey, **nearly nine in ten** organisations report regular AI use in at least one function; **44%** report scaling it and **37%** attribute some EBIT effect to AI (McKinsey, August 2026; n=1,719, 97 countries).
- Official statistics give a soberer picture: **20% of EU enterprises** with ten or more staff used AI in 2025, up from 13.5% (Eurostat). Denmark leads at 42%.
- In Germany, **57% of companies** with 20 or more employees use AI in 2026, up from 36% in 2025 and 20% in 2024 (Bitkom, n=603).
- In the US, **19.8% of businesses** used AI to produce goods or services in May 2026 (Census Bureau).
- Among people, **18.8% of the working-age population** worldwide used generative AI in Q2 2026 (Microsoft, telemetry).

- In every national study we compared, large firms use AI more often than small ones, by **1.6 to 4.3 times** (EU: 55% vs 17%).

Sectors: where it stands

Sector	Latest figure	Source
Health	81% of US physicians use AI, mostly for documentation	AMA, Mar 2026
Banking	92% of EU banks deploy AI	EBA
Manufacturing	88% of US mid-market manufacturers have integrated AI	RSM, 2026
Retail	About 14% of US retail firms use AI	US Census, May 2026
Utilities	81% of North American utilities use AI	Itron
Logistics	Active AI use rose from 16% to 43% in a year	Logistics Management, Sep 2026
Research	84% of researchers use AI, up from 57%	Wiley
Education	60% of US teachers used AI; weekly users save 5.9 hours a week	Gallup-Walton

Models and prices

The model layer has moved fast. At the top of the leaderboard, the best US and Chinese models are 2.7% apart, and the best closed model leads the best open-weight model by 3.3% (Stanford AI Index 2026). Flagship API prices now span about 100 times, from USD 0.10 to USD 10 per million input tokens (OpenAI, Anthropic, Google list prices, October 2026). Per-unit inference prices have fallen by 9 to 900 times a year depending on the capability threshold (Epoch AI). Cheaper units do not mean smaller bills: volume grows faster.

Energy and physical AI

- Data centres used about **485 TWh in 2025**, projected to reach about **950 TWh by 2030**, around 3% of global electricity demand; use by AI-focused data centres rose 50% in 2025 (IEA, 2026).
- In Ireland, data centres used **22% of metered electricity in 2024** (CSO).
- More than **600,000 industrial robots** were installed worldwide in 2025 (+11%), and the operational stock reached a record **5 million** (IFR, September 2026).
- Almost **250,000 professional service robots** were sold in 2025 (+24%); transport and logistics accounted for 47% (IFR).

Value: real, specific and still rare

- **37%** of organisations attribute some EBIT effect to AI (McKinsey, August 2026); **56% of CEOs** report no significant financial benefit so far (PwC, January 2026).
- Where measured, gains are task-specific: **+14%** issues resolved per hour in customer support, +34% for less experienced agents (NBER).
- In finance, **59%** of finance functions use AI, flat year on year, and 91% report low or moderate impact (Gartner).
- Know-how is the barrier named most often: **71%** of EU firms that considered AI and **85%** of German non-users (Eurostat; Bitkom).

What the figures mean

Spending and use are rising fast; proven value lags behind. The distance between trying AI and running it in a core process is where value is won. For most organisations, the useful question for 2027 is not whether to use AI but which process to change first, how to measure it and who owns the result. The Good AI Prism study turns these figures into a buyer worksheet and a pilot charter for exactly that step.

> **What to decide now** - Name the one process where AI in production would matter most to your results, and its accountable owner. - Measure that process today (volume, cost, quality, cycle time) so that any future gain is your own number, not a market average. - Budget for the full cost: in Germany, firms name cloud and compute (51%), data preparation (50%) and integration (41%) as cost drivers ahead of licences (21%).

All figures, with sample sizes and definitions, are in the Good AI Prism Global Market Study 2026/27 at ethority.ai/ai-prism-study.

Sources

- Gartner, worldwide AI spending forecast (16 Sep 2026): https://gartner.com/en/newsroom/press-releases/2026-09-16-gartner-forecasts-worldwide-ai-spending-to-grow-49-point-5-percent-in-2026
- Menlo Ventures, State of Generative AI in the Enterprise (Dec 2025): https://menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise/
- IDC, European AI spending to 2029 (Apr 2026): https://www.idc.com/resource-center/press-releases/european-ai-spending-to-reach-290-billion-by-2029-driven-by-banking-software-and-agentic-ai-adoption/
- Bitkom/IDC, German AI market: https://www.bitkom.org/index.php/Presse/Presseinformation/Deutscher-KI-Markt-waechst-um-48-Prozent
- Stanford HAI, AI Index Report 2026: https://hai.stanford.edu/ai-index/2026-ai-index-report
- Stanford HAI, AI Index 2026, technical performance: https://hai.stanford.edu/ai-index/2026-ai-index-report/technical-performance
- Stanford HAI, AI Index 2026, research and development: https://hai.stanford.edu/ai-index/2026-ai-index-report/research-and-development
- NVIDIA, Q2 FY2027 results: https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Announces-Financial-Results-for-Second-Quarter-Fiscal-2027/
- IEA, data-centre electricity 2025 (Apr 2026): https://www.iea.org/news/data-centre-electricity-use-surged-in-2025-even-with-tightening-bottlenecks-driving-a-scramble-for-solutions
- IEA, Key Questions on Energy and AI: https://www.iea.org/reports/key-questions-on-energy-and-ai/executive-summary
- McKinsey, The state of AI (Aug 2026): https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
- Eurostat, 20% of EU enterprises used AI in 2025: https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20251211-2
- Bitkom, majority of companies use AI (Sep 2026): https://www.bitkom.org/Presse/Presseinformation/Erstmals-nutzt-die-Mehrheit-der-Unternehmen-KI
- US Census Bureau, AI use by businesses (May 2026): https://census.gov/library/stories/2026/05/ai-use-businesses.html
- Microsoft, The continued state of global AI diffusion in 2026: https://blogs.microsoft.com/on-the-issues/2026/09/21/the-continued-state-of-global-ai-diffusion-in-2026/
- AMA (Mar 2026): https://ama-assn.org/press-center/ama-press-releases/ama-ai-usage-among-doctors-doubles-confidence-technology-grows
- EBA, AI in EU banking: https://www.eba.europa.eu/sites/default/files/2025-09/146b3558-d026-47bf-a872-f05e93ed30d2/Rising%20application%20of%20AI%20in%20EU%20banking%20and%20payments%20sector.pdf
- RSM, manufacturers using AI 2026: https://rsmus.com/insights/industries/manufacturing/manufacturers-using-ai-2026.html
- Itron: https://investors.itron.com/node/25786
- Logistics Management (Sep 2026): https://www.supplychain247.com/article/logistics-ai-use-adoption-trust-study
- Wiley: https://newsroom.wiley.com/press-releases/press-release-details/2025/AI-Adoption-Jumps-to-84-Among-Researchers-as-Expectations-Undergo-Significant-Reality-Check/

- Gallup-Walton, teachers and AI: https://nextgeninsights.waltonfamilyfoundation.org/wp-content/uploads/2025/06/Teachers_AI_Report.pdf
- Epoch AI, inference price trends: https://epoch.ai/data-insights/ai-inference-price-trends
- Anthropic, models overview (prices): https://platform.claude.com/docs/en/about-claude/models/overview
- OpenAI, GPT-6.1 Sol and Luna pricing: https://openai.com/index/introducing-gpt-6-1-sol/
- CSO Ireland, data centres metered electricity 2024: https://www.cso.ie/en/releasesandpublications/ep/p-dcmec/datacentresmeteredelectricityconsumption2024/keyfindings
- IFR, five million robots (Sep 2026): https://ifr.org/ifr-press-releases/five-million-robots-now-operate-in-factories-globally
- IFR, World Robotics 2026 Service Robots (Sep 2026): https://kyodonewsprwire.jp/release/202610017079
- PwC, 29th Global CEO Survey (Jan 2026): https://www.pwc.com/gx/en/news-room/press-releases/2026/pwc-2026-global-ceo-survey.html
- NBER Working Paper 31161: https://www.nber.org/papers/w31161
- Gartner, finance AI adoption (Nov 2025): https://www.gartner.com/en/newsroom/press-releases/2025-11-18-gartner-survey-shows-finance-ai-adoption-remains-steady-in-2025
- ETHORITY, Good AI Prism Global Market Study 2026/27: https://ethority.ai/ai-prism-study

Updated edition of the article first published on ethority.net in 2024; the original remains archived.

AI in Inventory and Logistics 2026: What the Evidence Measures

Edition: 2026/27 · English · v1.0 · 22 October 2026

Half of German logistics firms now use AI and logistics robots lead the service-robot market; the gains that count are the ones you measure against your own baseline.

A note on the original headline

Our 2024 article looked at how AI improves inventory accuracy in retail. This edition builds on what has been measured and published since, and adds a method to measure the effect in your own operation.

Where adoption stands

Logistics has moved quickly from pilots to use:

- **51% of German logistics firms** with 20 or more staff use AI in 2026, up from 22% in 2022 (Bitkom, September 2026, n=502). Across the German economy the figure is 57%.
- In the same survey, 47% use warehouse management systems, 39% big data and analytics, 23% digital twins, 17% robotics or autonomous transport and 15% smart shelves.
- In a US logistics survey, active AI use rose from **16% to 43%** in one year, yet only **7%** of respondents highly trust AI output, and 65% say staff use AI with or without approval (Logistics Management, September 2026).
- About **14% of US retail firms** use AI (US Census Bureau, May 2026), well below information (39.7%) and finance and insurance (33.9%).
- Only **17%** of supply-chain leaders pursue a transformational redesign with AI; 83% scale it incrementally (Gartner, May 2026, n=140).

Read together, these figures say that use is spreading, but trust and depth of use are the next constraints.

What AI does in inventory and the supply chain

The core techniques described in 2024 remain valid: demand forecasting from sales history, seasonality and external signals; automated replenishment; dynamic safety stock; and route and warehouse optimisation. Three things have changed in the technology underneath:

- **Language models handle the documents.** Delivery notes, customs papers, supplier emails and claims can now be read and checked by general-purpose models. Prices for small models start at USD 0.10 per million input tokens (OpenAI GPT-6 Luna, Anthropic Claude Haiku 5.5), which makes high-volume document work affordable. Document fraud is one of the risks logistics leaders name.
- **Agents act across systems.** Spending on AI agents and assistants is forecast to more than double to USD 65.5 billion in 2027 (Gartner). 55% of supply-chain leaders expect agentic AI to reduce entry-level hiring needs (Gartner, February 2026, n=509).
- **Models change often.** OpenAI retires its GPT-4o and GPT-4.1 snapshots on 14 January 2027. Any agent or document workflow needs a regression test before a model switch.

Expected benefits are clear in the German data: 66% of logistics firms expect better supply-chain transparency, 60% a lower environmental impact and 49% lower costs (Bitkom).

Physical AI: robots in the warehouse

Logistics is where robots earn their keep today.

- Almost **250,000 professional service robots** were sold worldwide in 2025 (+24%). **Transport and logistics** accounted for 117,500 units, or **47%**, the largest application (IFR, September 2026).
- Robots-as-a-service is the dominant business model in the US (IFR).
- About **7,000 full-size humanoid robots** were sold globally in 2025, the first IFR count. For warehouses, humanoids are pilots; mobile robots and automated storage are where volume deployments sit.
- Industrial robot stock reached a record **5 million** units; Germany has 449 robots per 10,000 manufacturing employees, against an EU-27 average of 231 (IFR).

Two rules matter for anyone buying robots for the EU. The Machinery Regulation (EU) 2023/1230 applies from **20 January 2027**, including requirements for AI-enabled safety functions. Updated robot safety standards ISO 10218-1 and -2 were published in 2025.

People and constraints

AI and robotics in logistics respond to a labour gap. 87% of German logistics firms cite a shortage of skilled staff, including 48% for warehouse staff (Bitkom). Globally, about 2.9 million truck-driver jobs are unfilled (IRU). Energy matters too: every German logistics firm surveyed named energy prices as a challenge.

Where AI touches workers, rules apply. Worker management and performance monitoring are Annex III uses under the EU AI Act, with high-risk duties from **2 December 2027** (Regulation (EU) 2026/1744). Emotion recognition at work has been prohibited since February 2025.

How to measure the effect yourself

A credible number comes from your own operation, measured before and after:

1. **Set the baseline.** Inventory accuracy (physical vs system), forecast error, stock turns, stock-out rate, on-time rate and cost per shipment, measured for at least one comparable period.
2. **Check data readiness first.** In manufacturing, data quality is behind 51% of failed AI pilots (RSM, 2026).
3. **Count accepted outcomes.** For document or replenishment workflows, divide total cost (usage, integration, review, rework) by the cases your team accepts, not by the cases processed.
4. **Agree a stop condition.** Decide in advance which result scales the pilot and which ends it.
5. **Make shadow AI visible.** With 65% reporting AI use with or without approval, an inventory of tools in use is the starting point for trust.

The Good AI Prism maps offerings in the Mobility & Logistics and Retail & Consumer areas by data type and technology layer, which helps narrow a shortlist before any demo.

> **What to decide now** - Pick one inventory or logistics process and record its baseline: accuracy, forecast error, cost per shipment. - For any robot purchase from 2027, write Machinery Regulation conformity and ISO 10218:2025 into the specification. - If AI monitors warehouse staff, prepare for the 2 December 2027 high-risk duties now, with works-council involvement.

Evidence and buyer tools for logistics and retail are in the Good AI Prism Global Market Study 2026/27 at ethority.ai/ai-prism-study.

Sources

- Bitkom, logistics and AI 2026 (Sep 2026): https://www.bitkom.org/Presse/Presseinformation/Die-Logistik-investiert-in-KI-doch-es-fehlt-Knowhow

- Bitkom, majority of companies use AI (Sep 2026): https://www.bitkom.org/Presse/Presseinformation/Erstmals-nutzt-die-Mehrheit-der-Unternehmen-KI
- Logistics Management, AI use and trust study (Sep 2026): https://www.supplychain247.com/article/logistics-ai-use-adoption-trust-study
- US Census Bureau, AI use by businesses (May 2026): https://census.gov/library/stories/2026/05/ai-use-businesses.html
- Gartner, AI and supply-chain operating models (May 2026): https://www.gartner.com/en/newsroom/press-releases/2026-05-06-gartner-survey-shows-ai-is-not-driving-supply-chain-operating-model-transformation
- Gartner, agentic AI and entry-level hiring (Feb 2026): https://www.gartner.com/en/newsroom/press-releases/2026-02-25-gartner-survey-shows-55-percent-of-supply-chain-leaders-expect-agentic-ai-to-reduce-entry-level-hiring-needs
- Gartner, worldwide AI spending forecast (Sep 2026): https://gartner.com/en/newsroom/press-releases/2026-09-16-gartner-forecasts-worldwide-ai-spending-to-grow-49-point-5-percent-in-2026
- OpenAI, GPT-6.1 Sol and Luna pricing: https://openai.com/index/introducing-gpt-6-1-sol/
- Anthropic, models overview: https://platform.claude.com/docs/en/about-claude/models/overview
- OpenAI, deprecations: https://developers.openai.com/api/docs/deprecations
- IFR, World Robotics 2026 Service Robots (Sep 2026): https://kyodonewsprwire.jp/release/202610017079
- IFR, five million robots (Sep 2026): https://ifr.org/ifr-press-releases/five-million-robots-now-operate-in-factories-globally
- IFR, robot density (Apr 2026): https://ifr.org/downloads/press_docs/DE-2026-APRIL-_8-IFR_PRESSMELDUNG_Roboterdichte-deutsch.pdf
- Regulation (EU) 2023/1230 (Machinery Regulation): https://eur-lex.europa.eu/eli/reg/2023/1230/oj
- ISO 10218-2:2025: https://www.iso.org/standard/73934.html
- IRU driver shortage report (via DC Velocity): https://www.dcvelocity.com/supply-chain/other-services/labor/iru-report-finds-global-truck-driver-shortage
- Regulation (EU) 2026/1744: https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng
- RSM, manufacturers using AI 2026: https://rsmus.com/insights/industries/manufacturing/manufacturers-using-ai-2026.html
- ETHORITY, Good AI Prism Global Market Study 2026/27: https://ethority.ai/ai-prism-study

Updated edition of the article first published on ethority.net in 2024; the original remains archived.

Retail AI in 2026: Dynamic Pricing and Personalisation, From Promise to Proof

Edition: 2026/27 · English · v1.0 · 22 October 2026

AI-driven discovery and agentic shopping are changing how customers find and buy, so retailers who price and personalise with AI now need consumer trust and measured results as much as algorithms.

What has changed since 2024

Our 2024 article described how AI adjusts prices to demand, stock and competition, and how it tailors offers to each customer. The mechanics still hold. What has changed is the setting. Customers now research products with AI assistants, AI agents are starting to buy on their behalf, and regulators have set dates for transparency. The 2024 version also leaned on a widely quoted consumer statistic and an unnamed retailer case. This edition replaces both with dated, sourced figures.

Where retail stands

- About **14% of US retail firms** use AI, and about 17% expect to within six months (US Census Bureau, May 2026). That is below the national average of 19.8% and well below information (39.7%).
- **30% to 45% of US consumers** already use generative AI for product research and comparison (Bain & Company, November 2025).
- AI accounts for **up to 25% of referral traffic** for some retailers, but still less than 1% of total traffic (Bain).
- **Half of consumers** are not comfortable letting AI handle end-to-end transactions on their own (Bain).
- Consumers trust a retailer's own on-site agent **three times more** than a third-party agent, though Bain expects that gap could narrow.

In other words, retail AI use inside the firm is still modest, while customers are moving faster. The gap is an opportunity for retailers who build trusted, owned channels.

How AI pricing and personalisation work in 2026

The core loop is unchanged: models combine sales history, inventory, competitor prices, seasonality and local signals to forecast demand and recommend a price or an offer; feedback from results retrains the model. Three developments in the technology stack matter for retailers now:

- **Language models are affordable at retail volume.** Small models from OpenAI (GPT-6 Luna) and Anthropic (Claude Haiku 5.5) start at USD 0.10 per million input tokens, and flagship prices span about 100 times. Personalised product descriptions, search and service answers can be generated at scale, with the cost per interaction measurable.
- **Retailers run several models.** 81% of large enterprises use three or more model families (a16z). Models also retire: OpenAI shuts down its GPT-4o and GPT-4.1 snapshots on 14 January 2027, so every customer-facing workflow needs a test before a model switch.
- **Agents enter commerce.** Spending on AI agents and assistants is forecast to grow from USD 29.2 billion in 2026 to USD 65.5 billion in 2027 (Gartner). Gartner also expects more than 40% of agentic AI projects to be cancelled by the end of 2027. Scope agent permissions tightly, especially for anything that changes a price or completes a payment.

Trust is now a design requirement

Personalisation creates value only when customers accept it. Three practical points follow.

Own the relationship. Bain advises retailers to "dig the customer moat" with exclusive products and loyalty benefits, rethink retail media for natural-language queries, and retain ownership of first-party data and checkout wherever possible.

Explain the AI. In the EU, Article 50 of the AI Act requires since 2 August 2026 that people are told when they are interacting with an AI system, such as a shopping assistant. From 2 December 2026, generative systems already on the market must mark synthetic outputs, including product images, in a machine-readable way (Regulation (EU) 2026/1744).

Know which decisions are regulated. Creditworthiness assessment of natural persons, for example for instalment payments at checkout, is an Annex III high-risk use with duties from 2 December 2027. In the US, Colorado and California rules on automated decisions with significant effects apply from 1 January 2027. Personalised prices are also a consumer-protection topic in several markets; check the disclosure duties in each market where you sell before you launch.

Choosing tools on facts

Retail AI often starts with customer service and search. Here the facts that decide a European shortlist are rarely public. In the Good AI Prism study, 10 of 107 customer-service AI offerings publicly documented both an EU hosting option and limits on training with customer data: 8 state no training and 2 train only with opt-in consent (Wave 1; provider information reviewed 19–20 October 2026). Prices come in at least seven units, from per resolution to per seat to quote only. Ask for hosting region, training use and billing unit in writing before any demo.

Measured gains exist: in customer support, a generative AI assistant raised issues resolved per hour by 14% on average and 34% for less experienced agents (NBER).

How to measure pricing and personalisation

The evaluation methods from 2024 remain the right ones, with sharper rules:

- **Use a control group.** Run AI pricing or personalisation against a holdout of stores, products or customers, so that seasonality and promotions do not inflate the result.
- **Measure margin, not only sales.** Track conversion, gross margin, markdown rate and return rate together.
- **Track AI-driven discovery.** Measure AI-referred traffic and its conversion separately; it is small today but growing.
- **Watch trust signals.** Complaint rate, price-related service contacts and opt-outs show whether customers accept the personalisation.
- **Agree a stop condition.** Decide in advance which result scales the pilot and which ends it.

> **What to decide now** - Choose one pricing or personalisation use case and run it against a control group with margin and complaint rate as outcome measures. - Decide which AI-driven channels you want to own (on-site assistant, loyalty, checkout) and which third-party agents you will serve, and on what data terms. - Map every customer-facing AI system to its disclosure duties: Article 50 now, machine-readable marking from 2 December 2026, credit decisions from 2 December 2027.

Sector evidence for Retail & Consumer and the full buyer worksheet are in the Good AI Prism Global Market Study 2026/27 at ethority.ai/ai-prism-study.

Sources

- US Census Bureau, AI use by businesses (May 2026): https://census.gov/library/stories/2026/05/ai-use-businesses.html

- Bain & Company, Agentic AI poised to disrupt retail (13 Nov 2025): https://www.bain.com/about/media-center/press-releases/20252/agentic-ai-poised-to-disrupt-retail-even-with-50-of-consumers-cautious-of-fully-autonomous-purchasesbain--company/
- OpenAI, GPT-6.1 Sol and Luna pricing: https://openai.com/index/introducing-gpt-6-1-sol/
- Anthropic, models overview: https://platform.claude.com/docs/en/about-claude/models/overview
- a16z, enterprise AI survey (Jan 2026): https://a16z.com/leaders-gainers-and-unexpected-winners-in-the-enterprise-ai-arms-race/
- OpenAI, deprecations: https://developers.openai.com/api/docs/deprecations
- Gartner, worldwide AI spending forecast (Sep 2026): https://gartner.com/en/newsroom/press-releases/2026-09-16-gartner-forecasts-worldwide-ai-spending-to-grow-49-point-5-percent-in-2026
- Gartner, agentic AI project cancellations: https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027
- Regulation (EU) 2026/1744: https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng
- Lewis Silkin, Digital Omnibus on AI (Jul 2026): https://www.lewissilkin.com/en/insights/2026/07/27/the-digital-omnibus-on-ai-enters-into-force-today-102nedo
- Colorado SB 26-189: https://leg.colorado.gov/bills/sb26-189
- California Privacy Protection Agency, ADMT regulations: https://privacy.ca.gov/2025/09/california-finalizes-regulations-to-strengthen-consumers-privacy/
- NBER Working Paper 31161: https://www.nber.org/papers/w31161
- ETHORITY, Good AI Prism Global Market Study 2026/27: https://ethority.ai/ai-prism-study

Updated edition of the article first published on ethority.net in 2024; the original remains archived.

Synthetic Faces and Voices in 2026: From VASA-1 to Deepfake Rules

Edition: 2026/27 · English · v1.0 · 22 October 2026

In 2024 a research demo showed a lifelike talking face made from one photo and a voice clip; in 2026 the question is no longer whether this is possible, but how synthetic media is labelled, governed and used well.

By Sten Franke, Founder, Principal Architect & Lead Creator, ETHORITY

Where we started: VASA-1

In April 2024 we wrote about Microsoft Research's VASA-1. From a single portrait and an audio clip, it generated a talking face with synchronised lips, natural head movement and expressions. The researchers reported 512×512 video frames at 45 frames per second in offline batch mode and up to 40 frames per second in online streaming with 170 milliseconds of starting latency, on a desktop PC with a single NVIDIA RTX 4090 GPU.

The team handled the release with care. Their project page states that they have "no plans to release an online demo, API, product, additional implementation details" until they are confident the technology will be used responsibly. They noted it "could still potentially be misused for impersonating humans" and that generated videos "still contain identifiable artifacts."

Our questions then were: how do we tell AI-created content from authentic recordings, what happens to identity and privacy, and how does such a tool fit into an information space already under pressure? Two and a half years later, regulators have given concrete answers to the first question, and businesses have learned a lot about the second and third.

What has changed in the technology

Real-time synthesis is no longer a research exception. The capability that needed a high-end GPU in 2024 now sits close to everyday hardware: Apple's M5 chip offers over four times the peak GPU compute for AI of its predecessor, and Qualcomm's Snapdragon X2 Elite Extreme carries an 80 TOPS neural processor. Counterpoint expected AI-capable PCs to exceed half of global shipments in 2026.

Legitimate uses have grown alongside. Voice agents are now a standard part of customer service; the Good AI Prism maps them in the Customer Service and the Meetings, Voice & Language areas. Synthetic presenters are used in training and accessibility. The same building blocks also lower the cost of impersonation, which is why cyber incidents remain the top business risk in the Allianz Risk Barometer 2026 (42%), with AI climbing to second place (32%).

The rules in 2026: labelling is the common denominator

Most AI regimes differ in scope, but almost all now require synthetic content to be identifiable.

European Union. The AI Act's Article 50 transparency duties apply from 2 August 2026, as confirmed by the Digital Omnibus (Regulation (EU) 2026/1744). Providers of systems that generate synthetic audio, image, video or text must mark outputs in a machine-readable format; systems already on the market before 2 August 2026 have until **2 December 2026**. Deployers who publish deepfakes must disclose that the content has been artificially generated or manipulated. From **2 December 2026**, the Omnibus also adds a prohibition on AI systems that generate non-consensual intimate imagery of identifiable people or child sexual abuse material.

China. Measures for labelling AI-generated synthetic content have applied since **1 September 2025**. They require explicit labels on text, images, audio, video and virtual scenes, plus implicit labels in file metadata, and oblige app stores to check labelling when reviewing apps.

South Korea. The AI Basic Act has applied since **22 January 2026**. Providers must tell users in advance that AI is in use and may need to disclose that images, video or audio are AI-generated where they are hard to distinguish from reality. A grace period of at least one year applies to fines.

India. Amended IT Rules have applied since **20 February 2026**. Permitted synthetic content must be prominently labelled and carry permanent metadata or provenance markers that platforms may not remove; takedown orders must be acted on within three hours.

What this means for businesses

For most organisations, synthetic media is now a governance topic with three sides:

- **As a creator.** If you produce AI-generated video, voice or images for marketing, training or service, you need labelling and marking that work across the EU, China, Korea and India. Build them into the production workflow rather than adding them at publication.
- **As a buyer.** Ask suppliers of generative tools how outputs are marked, which metadata survives export, and whether voice or face cloning requires documented consent from the person depicted.
- **As a target.** Voice and video impersonation of executives is a fraud vector. Payment approvals, password resets and supplier bank-detail changes should never rely on a voice or video call alone.

Consent deserves special attention in meetings. In a US survey, 33% of workers had met with an AI notetaker, and 65% of them said consent was inconsistent or missing.

A long view

Twenty years ago, at my company Gridpatrol, we monitored the web for brands and once intercepted a compromising image of a retail manager before it spread. Then the problem was static images and text. Today it is real-time video and cloned voices. What has stayed the same is the principle: detect early, label honestly, and keep a human accountable for what is published.

Rules alone will not make synthetic media trustworthy. Clear labels, documented consent and verification steps inside business processes will. The opportunity for 2027 is to use synthetic voices and faces where they genuinely help, in accessibility, training and service, while making their origin visible.

> **What to decide now** - Inventory every tool that generates images, video or voice and confirm how it meets EU machine-readable marking by 2 December 2026. - Require documented consent for any cloned voice or likeness, including employees and presenters. - Add a second, non-voice verification step to payment approvals and supplier bank-detail changes.

Labelling rules for more than 30 economies are compared in the Good AI Prism Global Market Study 2026/27 at ethority.ai/ai-prism-study.

Sources

- Microsoft Research, VASA-1 project page: https://www.microsoft.com/en-us/research/project/vasa-1/
- VASA-1 project site (release and responsible-AI statement): https://vasavatar.github.io/VASA-1/
- Apple, M5 announcement (Oct 2025): https://www.apple.com/newsroom/2025/10/apple-unleashes-m5-the-next-big-leap-in-ai-performance-for-apple-silicon

- Futurum, Snapdragon X2 Elite (Oct 2025): https://futurumgroup.com/insights/snapdragon-x2-elite-pushes-ai-pc-performance-to-new-heights
- Counterpoint, AI PCs in 2026: https://counterpointresearch.com/ko/reports/ai-advanced-pcs-to-surpass-half-of-global-shipments-in-2026
- Allianz Risk Barometer 2026: https://commercial.allianz.com/news-and-insights/news/allianz-risk-barometer-2026.html
- Regulation (EU) 2026/1744: https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng
- Regulation (EU) 2024/1689 (AI Act): https://eur-lex.europa.eu/eli/reg/2024/1689/oj
- Lewis Silkin, Digital Omnibus on AI (Jul 2026): https://www.lewissilkin.com/en/insights/2026/07/27/the-digital-omnibus-on-ai-enters-into-force-today-102nedo
- nicfab, Digital Omnibus on AI in the Official Journal: https://www.nicfab.eu/en/posts/digital-omnibus-ai-official-journal/
- Loeb & Loeb, China AI labelling measures: https://www.loeb.com/en/insights/publications/2025/03/chinas-ai-labeling-measures-and-mandatory-national-standards-take-effect-september-1
- ITA, South Korea AI Basic Act: https://www.trade.gov/market-intelligence/south-korea-ai-basic-act
- Hogan Lovells, India synthetic-content labelling (Feb 2026): https://www.hlc.com/en/publications/india-introduces-mandatory-labelling-for-ai-and-3hour-takedown-for-illegal-content
- KRDO / Stacker, AI notetaker survey (Jul 2026): https://krdo.com/stacker-business-economy/2026/07/16/ai-notetakers-have-sat-in-on-1-in-3-us-workers-meetings-but-only-a-third-say-they-were-asked-first/
- ETHORITY, Good AI Prism Global Market Study 2026/27: https://ethority.ai/ai-prism-study

This article is not legal advice.

Updated edition of the article first published on ethority.net in 2024; the original remains archived.