

IMP

The AI Inflection

Three Years That Changed Everything —

What It Means for Industry, Jobs, Education,
and the Organisations that Must Now Lead

Industry Intelligence White Paper

AI Intelligence

Industrial Disruption

Agentic AI & Law

Education & Workforce

600-700M

ChatGPT estimated ~600–700M weekly active users

40%

Global jobs with meaningful AI exposure (IMF, 2024)

\$500B

Stargate AI infrastructure commitment, US, 2025–2029
Infrastructure pledge

92M

Jobs projected displaced by 2030, 170M new (WEF 2025)

1 Executive Summary

Between November 2022 — when ChatGPT crossed one million users in five days — and March 2026, the development of artificial intelligence has moved faster, and penetrated more domains of professional and economic life, than any comparable technology in modern history. The question confronting every organisation, every professional, and every institution today is not whether AI will change their field. It is whether they are moving at a speed commensurate with the change.

600-700M

estimated ~600–700M weekly active users (industry estimates, 2025)

40%

Global jobs with meaningful AI exposure (IMF, 2024)

92M

Jobs projected displaced by 2030, 170M new (WEF 2025)

\$500B

Stargate AI infrastructure commitment, US, 2025–2029

This paper provides a fact-grounded account of the three-year AI acceleration — its key milestones, its industrial implications, and its consequences for specific professional segments. It examines the practical reality of agentic AI and the legal terrain that is still being defined. It confronts the workforce disruption evidence directly, without either catastrophising or minimising. It addresses the crisis in education with the urgency it deserves. It raises the question of organisational responsibility — including the specific responsibility of professional bodies like the GDCh, VCI, and ACS. And it attempts to answer the question that underlies all of these: are we, collectively, calibrating our response correctly, or are we either too frightened or not frightened enough?

The IMP Position

AI is simultaneously the most significant productivity instrument available to any organisation today, and an increasingly evident disruptor of professional categories that have historically been considered automation-resistant. Both of these things are true simultaneously. The organisations that will navigate this environment successfully are those that hold both truths at once — and act on both.

2 Three Years That Changed Everything: The Timeline

The velocity of AI development between late 2022 and early 2026 is arguably unprecedented. The following timeline records only verified, documented milestones — no projections, no rumours.

- Nov 2022 ● OpenAI releases ChatGPT (GPT-3.5). Reaches 1 million users in 5 days, 100 million in 2 months — faster than TikTok or Instagram.
- Mar 2023 ● GPT-4 released. Passes bar exam in top 10th percentile. Anthropic launches Claude. Microsoft confirms Bing ran on GPT-4.
- Mar 2023 ● Italy becomes first nation to ban ChatGPT over data privacy concerns (GDPR). Lifted after OpenAI compliance actions.
- May 2023 ● ChatGPT plugins launch — AI begins taking actions, not just answering questions.
- Jul 2023 ● Meta open-sources LLaMA 2 (up to 70B parameters) — democratises access to powerful language models.
- Nov 2023 ● OpenAI board fires Sam Altman. Reinstated 5 days later. 95% of staff threaten resignation. Governance crisis exposes tensions in frontier AI development.
- Nov 2023 ● OpenAI DevDay: custom GPTs announced. GPT Store follows in January 2024.
- Dec 2023 ● Google launches Gemini — multimodal from foundation (text, image, audio, video).
- Feb 2024 ● Google rebrands Bard to Gemini. Full integration across Google Workspace begins.
- May 2024 ● OpenAI releases GPT-4o ('Omni') — real-time multimodal, free tier access. 46% of new code now AI-generated (GitHub 2024).
- Aug 2024 ● EU AI Act enters into force. Prohibition of high-risk AI practices effective February 2025.
- Jan 2025 ● DeepSeek R1 released. Matches OpenAI o1 on math/coding benchmarks. Training cost: ~\$6M vs >\$100M for comparable US models. Nvidia loses \$600B in market cap in one session.
- Jan 2025 ● Stargate Project announced by President Trump, OpenAI, Oracle, SoftBank. \$500B AI infrastructure commitment over 4 years.
- Feb 2025 ● EU AI Act first enforcement wave: prohibitions on social scoring, biometric surveillance, AI-manipulated behaviour take effect.
- Mar 2025 ● Google releases Gemini 2.5 Pro with 1 million token context window. 'Thinking model' with advanced planning and reasoning.
- Apr 2025 ● Meta releases LLaMA 4 (Scout, Maverick) — multimodal, open-source. 10M token context window for Scout variant.
- Apr 2025 ● OpenAI releases GPT-4.1 family with 1 million token context; o3 and o4-mini for advanced reasoning, math, coding.
- Aug 2025 ● GPT-5 released. ChatGPT reaches estimated 600-700 million weekly active users.
- Aug 2025 ● EU AI Act: GPAI model obligations effective. AI Office operational. Fines up to €35M / 7% global turnover now apply.
- Oct 2025 ● Gartner predicts: by 2028, agentic AI will make up to 15% of day-to-day enterprise decisions autonomously.
- Dec 2025 ● DeepSeek V3.2 released — matches GPT-5 performance on multiple benchmarks at fraction of cost. DeepSeek V4 (coding-focused) imminent.
- Mar 2026 ● EU AI Act full high-risk AI compliance deadline approaching (August 2026). Regulatory enforcement accelerating worldwide.

The pattern that emerges from this timeline is not simply one of rapid progress. It is the pattern of a technology crossing thresholds that matter: from text to multimodal, from chat to action, from expensive and centralised to cheap and open-source, from assistant to agent. Each threshold extends the scope of what AI can do, and therefore which domains it enters.

3 The DeepSeek Moment: When the Economics Changed

No single event in the three-year AI timeline better illustrates the structural unpredictability of this technology than the release of DeepSeek R1 on January 20, 2025. DeepSeek, a Chinese AI startup backed by quantitative hedge fund High-Flyer Capital, released an open-source reasoning model with 671 billion parameters — developed in approximately two months, for a reported training cost of <\$6M (unverified, widely debated), using Nvidia H800 chips to which China still had access despite US export controls.

DeepSeek R1's benchmark performance matched or exceeded OpenAI's o1 on competitive mathematics (79.8% Pass@1 on AIME), coding (2029 Elo on Codeforces — outperforming 96.3% of humans), and general knowledge (90.8% on MMLU), while the API was priced at approximately 95% less than OpenAI's equivalent. It was released under an MIT licence — free for commercial use. Within hours it was the top free application on the US App Store, displacing ChatGPT.

On January 27, 2025 — seven days later — it triggered one of the largest single-day market value losses for Nvidia in tech history (>\$200B) as investors reassessed the assumption that frontier AI requires ever-larger and ever-more-expensive hardware. The 'scale is all you need' doctrine that had driven AI infrastructure investment collapsed in many cases. Algorithmic efficiency, it turned out, could compensate for hardware constraints in ways that had not been publicly demonstrated at this scale.

What DeepSeek Actually Proved

Three things. First, that frontier-level AI capability is no longer the exclusive domain of organisations with \$100M+ training budgets. Second, that open-source AI development is now genuinely competitive with closed-source commercial models — changing the competitive economics for every AI service provider. Third, that the AI development race has become a geopolitical contest with genuine technical depth on multiple sides. DeepSeek V3.2 (December 2025) matches GPT-5 performance on multiple reasoning benchmarks. V4 with Engram memory architecture is scheduled for mid-2026. The gap between US and Chinese AI capability, to the extent it exists, is narrowing at pace.

The Data Privacy Dimension

DeepSeek is headquartered in Hangzhou, China. Multiple governments — including Australia, Italy, and Taiwan — have banned DeepSeek from government devices citing data security concerns. For European and US industrial companies considering open-source AI deployment, the question of where data is routed, stored, and potentially accessed is not academic. MIT licence does not address data sovereignty. Enterprises must treat any open-source model deployed via a third-party API as a data governance question requiring legal and IT security sign-off.

4 Industrial Implications: What AI Is Actually Doing in Organisations

The adoption picture as of March 2026 is one of significant formal deployment at large organisations combined with much broader informal adoption driven by individual employees. McKinsey's 2025 workplace survey found that 71% of organisations report using or experimenting with generative AI, up from 65% six months earlier. Surveys suggest a majority (60–75%) of knowledge workers use AI tools — and 78% of current AI users bring their own tools to work without formal organisational approval. The gap between institutional AI strategy and employee AI behaviour is the defining implementation challenge of 2025–2026.

4.1 Manufacturing and Process Industry

In capital-intensive manufacturing — chemicals, automotive, pharmaceuticals, coatings — AI is entering through predictive maintenance, process optimisation, and quality control. Predictive maintenance AI analyses vibration, temperature, and process data to identify equipment deterioration before failure, reducing unplanned downtime by 15–30% in documented deployments. Process control AI maintains product consistency across large-batch production by managing temperature, pressure, and mixing parameters in real time. AI-assisted visual quality inspection systems have demonstrated defect detection rates exceeding human inspectors in high-throughput production lines.

The production floor implication is not primarily job replacement in the near term — physical dexterity and contextual judgment remain difficult to automate in complex manufacturing environments. The implication is role transformation: operators who can interpret AI system outputs, override when necessary, and manage human-AI collaborative workflows are becoming the primary workforce requirement. Operators who cannot are facing genuine displacement risk.

4.2 Research and Development

AI is having its most dramatic near-term impact in R&D; environments where the work is primarily cognitive and data-intensive. In drug discovery, AI models are identifying candidate molecules by screening computational representations of billions of compounds — a process that previously required physical synthesis and testing of each candidate. AlphaFold (DeepMind) offers dramatically advanced protein structure prediction, solving a problem that had occupied structural biologists for decades. In chemistry, Evonik's COATINO® platform tests more than 120 coating formulations per day through AI-driven high-throughput experimentation. McKinsey estimates AI can reduce chemical R&D costs by up to 40% and cut development timelines by up to 50%.

GitHub Copilot, an AI coding assistant, reports 20 million active developers and adoption across 90% of Fortune 100 companies as of July 2025 — a 400% year-over-year increase. In 2024, Copilot users report up to ~46% of code written with AI assistance. AI systems improved from 2023 single-digit performance to >60% on benchmark tasks (SWE-Bench) — a capability improvement of an order of magnitude in twelve months. The entry-level software developer role is one of the categories most immediately affected by AI; Goldman Sachs' 2025 data shows unemployment among 20–30-year-olds in tech-exposed occupations has risen by almost 3 percentage points since the start of 2025.

4.3 Administrative and Management Functions

Administrative roles represent the single highest exposure category for AI automation. Goldman Sachs analysis of over 800 occupations identified administrative and office support as having 46% of tasks susceptible to AI automation. Bookkeeping, accounts payable, invoice processing, expense management, scheduling, document summarisation, report generation, regulatory filing — these are precisely the structured, rule-based, data-processing tasks that current AI systems handle well. Klarna, the financial services company, reported internally in 2024 that its AI assistant was performing the work equivalent to 700 full-time employees with sub-2-minute response times (company claim).

Middle management is next. AI tools that can synthesise data across functions, generate performance dashboards, draft strategic summaries, and answer complex operational questions are replacing tasks that previously required a layer of analytical middle management. The question of what a senior manager's decision making adds over what an AI synthesis can provide is one that corporate boards are beginning to ask seriously. A 2025 Accenture study projects that by 2030, AI agents will be the primary users of most enterprises' internal digital systems.

The Productivity Evidence

Nielsen Norman Group studies of AI productivity gains in knowledge work tasks showed average time savings of 66% across a range of professional activities including writing, research, coding, and data analysis. PwC's 2025 AI Jobs Barometer found that workers with demonstrable AI skills earn a 25% wage premium over peers without them. Industries with higher AI adoption showed productivity growth rates four times higher than lower-adoption sectors. The productivity argument for AI adoption is not speculative — it is observable in current data.

5 The Professional Disruption Map

The following assessment maps AI task exposure across major professional categories, drawing on published research from Goldman Sachs, McKinsey, the IMF, and Oxford University. 'Task exposure' means the percentage of tasks within a role that current AI systems can perform with comparable quality to a professional at median competence. This is not the same as job elimination — but it is the primary driver of role transformation and, in some categories, displacement.

Sector / Role	AI Task Exposure (% of tasks automatable)
Administrative / Data Entry / Bookkeeping	86%
Legal Research / Document Review / Contract Analysis	72%
Customer Service / Call Centre Operation	70%
Entry-Level Software Development	65%
Content Writing / Templated Journalism	62%
Financial Analysis / Accounting (routine)	58%
Medical Diagnostics (radiology, pathology,	45%
Scientific Research (data analysis phase)	40%
Management Consulting / Strategy	35%
Teaching / Education	28%
Clinical Medicine / Patient Care	12%
Skilled Trades / Physical Installation	6%

Sources: Goldman Sachs 2025; McKinsey 2025; Oxford University Future of Work Programme; IMF 2024. Note: percentages indicate share of tasks, not probability of role elimination.

5.1 Law

The legal sector is among the most exposed high-prestige professions. Goldman Sachs analysis found that 44% of legal tasks are susceptible to AI automation. AI legal research tools (Harvey, CoCounsel, Lexis+ AI) are already performing contract analysis, due diligence document review, case law research, and regulatory compliance checking at speeds and volumes no human team can match. A document review task that might occupy a team of 20 paralegals for two weeks can be completed by an AI system in hours. The *Mobley v. Workday* case (US Northern District of California, July 2024) established that an AI screening tool could be deemed an 'agent' of its clients — creating direct liability for AI vendors whose tools make consequential decisions.

The implication for legal firms is structural. Entry-level associate positions — particularly document review, research, and first-draft generation — are the primary revenue-generating training ground for junior lawyers. If AI performs these tasks better and faster, the apprenticeship model of legal education breaks down. Senior partnership and courtroom advocacy remain human-intensive. The middle tier of the legal workforce — large teams of associates performing high-volume standardised legal work — faces the most acute near-term disruption.

5.2 Medicine and Healthcare

GPT-4-class models answer approximately 90% of US medical licensing exam questions correctly. Specialist diagnostic AI has demonstrated performance exceeding human radiologists in specific imaging tasks (diabetic retinopathy detection, mammographic screening). MedPaLM (Google's medical LLM) is reducing clinical documentation time in US hospital deployments by 40%. These numbers suggest significant potential for AI to take over the documentation, coding, and diagnostic support functions that consume enormous proportions of physician and nurse time without directly involving patient care.

Clinical medicine, nursing, and patient-facing care remain far more resistant. The combination of physical presence, emotional intelligence, ethical complexity, and the need for real-time contextual adaptation that characterises frontline patient care is not something current AI systems can replicate. The automation risk for medicine is concentrated in the documentation, administrative, and pattern-recognition diagnostic support functions. The patient interaction — which patients value most — is precisely where AI adds the least.

5.3 Science and Research

Scientific research is being restructured by AI at the data analysis, hypothesis generation, and literature synthesis stages. AI can now review and synthesise thousands of papers, identify patterns in large datasets, propose hypotheses, and design experimental protocols. The creative and intuitive aspects of breakthrough research — recognising an anomalous result that does not fit the model, knowing which question to ask — remain distinctly human. The risk for scientists is not displacement but deskilling: if AI handles all data processing and literature review, will the next generation of scientists develop the deep domain expertise that makes breakthrough insights possible?

5.4 B2B Content Creation and Publishing

The impact on B2B content creators is immediate and asymmetric. AI can generate templated content — product descriptions, SEO articles, earnings summaries, press release drafts, routine case studies — faster and cheaper than any human team. The companies generating high volumes of formulaic B2B content are already replacing that capacity with AI. What AI does not replace — and arguably cannot in the near term — is expert-driven original analysis, authoritative sector intelligence, investigative journalism, and the kind of technically credible thought leadership that shapes industry perception.

The B2B publishing model that depended on volume — many articles, regular newsletters, high-frequency social media — is directly in the disruption path. The model that depends on depth — sector expertise, primary data, original perspective, verified industry intelligence — is more defensible. The challenge for B2B communicators is navigating this transition without losing the editorial credibility that makes depth valuable in the first place.

6 Jobs: What the Data Actually Says

The evidence on AI and employment is more nuanced than either the 'AI will take all jobs' narrative or the 'AI will only create new jobs' reassurance. The published research from the most credible institutions does not support either extreme position. What it does support is this: significant disruption is coming, unevenly distributed, concentrated in specific task categories, and affecting white-collar cognitive work in ways that previous automation waves did not.

The Headline Numbers

- World Economic Forum Future of Jobs Report 2025 (1,000+ employers, 14 million workers, 55 economies): 92 million jobs projected displaced by 2030. 170 million new roles expected to emerge. Net gain: 78 million jobs. 41% of employers plan workforce reductions in areas where AI can automate tasks within five years.
- IMF 2024 assessment: approximately 40% of global jobs face meaningful AI exposure. In advanced, high-income economies: approximately 60%. In low-income countries: approximately 26%. The geographic concentration of disruption risk in wealthy economies is the reversal of previous automation patterns, which primarily affected manufacturing workers in developing economies.
- Goldman Sachs 2025: if AI use cases were expanded across the economy proportionally, 2.5% of US employment would be at risk of related job loss. With wide adoption, 6–7%. The unemployment rate may rise by 0.5 percentage points during the transition period, historically fading within two years. Entry-level tech roles already showing 3 percentage points higher unemployment among 20–30-year-olds.
- McKinsey 2025: today's existing technology could, in theory, automate approximately 57% of current US work tasks. Current deployment would displace approximately 30% of US work hours by 2030.
- Dario Amodei (Anthropic CEO), VivaTech 2025: AI could replace up to half of entry-level office jobs within five years. Kai-Fu Lee (AI pioneer) concurred. Nvidia CEO Jensen Huang pushed back, arguing greater productivity leads to more hiring. Both positions have evidential support — they are describing different time horizons and different job categories.

The Jobs Most at Risk: A Calibrated Assessment

High Disruption Risk (Near-Term)	Protected Categories (Near-Term)
<ul style="list-style-type: none"> • Customer service representatives — AI voice agents and chat systems can handle the majority of query volume at scale 	<ul style="list-style-type: none"> • Skilled trades — electricians, plumbers, HVAC, roofers face 2–6% task automation exposure
<ul style="list-style-type: none"> • Administrative assistants and data entry operators — structured, rule-based tasks with high AI competence 	<ul style="list-style-type: none"> • Clinical nursing and patient-facing medicine — physical presence and emotional care are not automatable
<ul style="list-style-type: none"> • Entry-level programmers — SWE-Bench scores improved from 4.4% to 71.7% in 2023–2024 	<ul style="list-style-type: none"> • Mental health practitioners — therapeutic relationships require human authenticity
<ul style="list-style-type: none"> • Paralegal and legal document review — 44% task exposure (Goldman Sachs) 	<ul style="list-style-type: none"> • Education — teaching involves social, relational, motivational dimensions AI cannot replicate
<ul style="list-style-type: none"> • Routine accounting and bookkeeping — QuickBooks, Xero, and enterprise platforms automating close processes 	<ul style="list-style-type: none"> • Senior strategic and creative roles requiring original judgment across domains
<ul style="list-style-type: none"> • Content writers producing templated, formulaic, or SEO-driven output 	<ul style="list-style-type: none"> • Management requiring physical observation, team motivation, and stakeholder relationship management
<ul style="list-style-type: none"> • Radiologists and pathologists for specific image-based diagnostic tasks 	<ul style="list-style-type: none"> • Roles requiring ethical judgment and accountability in high-stakes decisions

The structural consensus among economists (Goldman Sachs, IMF, MIT's Daron Acemoglu — 2024 Nobel Laureate in Economics) is that sustained mass unemployment from AI is unlikely, though transitional and painful disruption is not. Historical technology transitions have consistently generated more jobs than they eliminated over the medium term — but the transition itself is neither fast nor frictionless, and the new jobs created rarely employ the same people displaced by automation. The 2025–2030 period is the window in which reskilling investment matters most.

7 Agentic AI: From Conversation to Action

The transition from generative AI to agentic AI is the most consequential development in enterprise AI since ChatGPT's launch. Generative AI responds to prompts. Agentic AI executes tasks. An AI agent receives a goal, breaks it into steps, calls tools and APIs, makes decisions between steps, iterates on results, and delivers an outcome — all without requiring human input at each stage. The difference is not incremental. It is approaching partial autonomy comparable to junior operational roles in structured workflows.

What Agentic AI Can Do Today

- **Research and synthesis:** Search multiple sources, aggregate information, resolve contradictions, produce a structured report — without being prompted at each step.
- **Software development:** Autonomous coding agents (Claude Code, OpenAI Codex) write, test, debug, and deploy code. OpenAI's Codex Max is purpose-built for long-term coding tasks with minimal human intervention.
- **Business process execution:** Scheduling meetings, managing supply chain queries, handling IT support tickets, processing invoices — agents access internal systems and complete tasks end-to-end.
- **Customer interaction:** Agentic customer service systems resolve complex multi-step issues across multiple systems without human escalation.
- **Sales and outreach:** Agentic systems can research prospects, draft personalised outreach, follow up, and qualify leads autonomously.
- **Financial analysis:** Real-time monitoring, anomaly detection, regulatory compliance checking, and report generation without human tasking at each stage.

They are not equivalent to fully autonomous decision-makers and require:

- **defined boundaries, human override, structured environments**

Their effectiveness declines sharply in:

- **ambiguous contexts, high-liability decisions, incomplete or noisy data environments**

The Scale of Investment

2025 saw every major technology platform release agentic AI solutions: AWS Nova Act, Google Gemini Deep Research, Microsoft Azure AI Agent Service, IBM, Palantir, Databricks, and Salesforce (which acquired Convergence AI). Meta acquired Manus, a Singapore-based AI agent developer, for a reported \$2 billion. Gartner projects that by 2028, over one-third of enterprise software will include agentic AI, making up to 15% of day-to-day decisions autonomously. A 2025 Accenture study projects that by 2030, AI agents will be the primary users of most enterprises' internal digital systems.

Legal Implications and Governance Gaps

Agentic AI is operating in a legal vacuum that existing frameworks did not anticipate. The core problem is accountability: when an autonomous AI agent makes a mistake, causes harm, or enters into a legally binding commitment on behalf of an organisation, who is liable? Current legal doctrine on principal-agent relationships was written for human agents acting with delegated authority. AI agents do not fit neatly into that framework.

Legal Risk Register for Agentic AI Deployment

Contractual liability: AI agents may enter contracts on behalf of organisations. An agent that executes an unauthorised transaction creates direct liability for the deployer — regardless of vendor indemnity arrangements (*Quoine v. B2C2*, Singapore, 2020: crypto assets sold at significant undervalue by automated system). **Employment discrimination:** *Mobley v. Workday* (US Northern District of California, July 2024) held that an AI applicant screening system could be deemed an 'agent' of its employer clients — creating direct liability for AI vendors. **Terms of service violations:** Most enterprise SaaS platforms were not designed for automated agent access. Agents that access systems, modify records, or extract data may breach contractual restrictions and trigger SLA violations. **GDPR and data protection:** Agentic systems processing personal data across multiple systems simultaneously create complex data minimisation and purpose limitation challenges under GDPR and the EU AI Act. **Director duties:** Under UK Companies Act 2006, directors face personal liability for failures in governance or supervision of AI systems — including agentic AI deployed under their authority.

Minimum Governance Requirements for Agentic AI

Any organisation deploying agentic AI should establish, at minimum: (1) A written policy governing which systems agents may access and what actions they may take without human approval. (2) Defined permission boundaries with hard stops — actions that require human sign-off regardless of agent judgement. (3) A designated responsible individual (not a vendor) for each agent deployment. (4) Comprehensive audit logging that separates human and agent activity. (5) Legal review of all third-party system terms before agent access is enabled. (6) Human override capability at every decision point. The EU AI Act's August 2026 full compliance deadline makes this not only prudent governance but a legal requirement for high-risk AI system deployers.

8 The Regulatory Landscape: EU AI Act and Global Divergence

The EU AI Act is the world's first comprehensive AI regulatory framework, and its implementation is the most significant regulatory event in technology since GDPR. It applies to any organisation, worldwide, whose AI systems affect people in the EU — including through customers, users, or employees. This extraterritorial scope means it is a global standard in practice, not only a European one.

Where We Are: The Implementation Timeline

- August 1, 2024: EU AI Act entered into force.
- February 2, 2025: First obligations effective — prohibitions on unacceptable-risk AI (social scoring, biometric mass surveillance, AI that psychologically manipulates people, emotion recognition in workplaces and schools). AI literacy requirements take effect: every organisation using AI must ensure employees have basic AI competence.
- August 2, 2025: GPAI (General Purpose AI) model obligations effective. EU AI Office becomes operational. Fines up to €35M or 7% of global annual turnover for prohibited AI practices now applicable. Providers of foundation models like GPT-5, Claude, Gemini must maintain technical documentation, copyright compliance policies, and safety testing.
- January 2026: EU AI Office begins verifying 'machine-readability' of AI disclosures. Meta placed under closer scrutiny after refusing to sign GPAI Code of Practice. EU AI Act enforcement actions against X and Meta accelerating.
- August 2, 2026: Full high-risk AI system compliance deadline. This covers AI in recruitment, healthcare, critical infrastructure, financial services, law enforcement, and education — the systems that most industrial organisations are actually deploying.

The Global Regulatory Divergence

The US are more fragmented and sector-driven compared to the EU's unified framework. A December 2025 US Executive Order pushed for federal deregulation of AI to maintain competitive advantage over China. California's SB 942 (effective January 1, 2026) mirrors many EU transparency requirements at state level, creating regulatory fragmentation within the US. China mandates explicit and implicit labels on all AI-generated media since September 2025 — among the world's strictest AI labelling requirements. For European industrial companies, the EU regulatory path is clear if complex; for companies operating across all three major jurisdictions, compliance management has become a substantive operational function.

9 The Avoidance Trap: What Happens If You Don't Engage

A legitimate question for any organisation is: given the uncertainty, the legal complexity, and the pace of change, is it safer to wait? The evidence does not support this position. The cost of non-adoption is now measurable — and it is compounding.

The Competitive Productivity Gap

PwC's 2025 AI Jobs Barometer found that industries with high AI adoption showed productivity growth four times higher than lower-adoption sectors. Workers with AI skills earn a 25% wage premium over peers without them. If a competitor is operating at 40–60% higher productivity in knowledge work functions by 2027 — the realistic projection for full generative AI integration — a non-adopting organisation faces an irreversible competitive disadvantage in labour cost, output speed, and analytical quality. The compounding nature of productivity advantages means that a two-year avoidance decision made in 2025 may translate into a five-year recovery requirement from 2027.

The Talent Risk

The 78% of AI users who bring their own tools to work without organisational approval are not doing so to create problems — they are doing so because AI makes their work faster and better, and they are not willing to compete without it. Organisations that maintain blanket AI prohibition policies are not protecting themselves from AI risk. They are creating a different risk: the departure of their most productive and adaptive employees to organisations that empower rather than restrict AI use. Talent retention in knowledge work is inseparable from AI enablement in 2025.

The Unsupervised Adoption Risk

The greatest AI risk for organisations that prohibit formal AI adoption is not that their people stop using AI. It is that their people continue using AI without governance, training, or oversight — introducing data privacy risk, confidentiality risk, and hallucination risk without any organisational mitigation framework. A 2024 survey found that 44% of companies using or planning to use AI anticipated workforce reductions; but those that had not formally adopted AI were still experiencing unsupervised usage in 78% of cases. The choice is not between AI and no AI. It is between managed AI and unmanaged AI.

"Refusing to engage with AI is not a conservative risk management choice. It is the choice to accumulate competitive, talent, and governance risk simultaneously — while believing you are doing the opposite."

— IMP InterMediaPartners, March 2026

10 Ethics, Balance and the Question of Obligation

The ethical dimensions of AI are real, substantive, and poorly served by either of the dominant public narratives — uncritical enthusiasm and apocalyptic fear. Both positions tend to suppress the specific, actionable ethical questions that organisations and individuals actually need to resolve.

Are We Too Frightened, Too Critical, or Not Critical Enough?

The evidence suggests the answer is different depending on which fear we are examining. Organisations that avoid AI adoption entirely out of fear of making mistakes are, in our assessment, too frightened — and the data on competitive productivity gaps supports this. Organisations that deploy AI in high-stakes decision-making contexts (hiring, lending, medical triage, law enforcement) without adequate oversight frameworks are not being critical enough — and the EU AI Act, Mobley v. Workday, and the NIST AI Risk Management Framework all reflect this. The calibration question requires distinguishing between these categories — not applying a single posture across all AI applications.

Genuine Ethical Concerns That Deserve Sustained Attention

- **Hallucination and factual reliability:** All current AI systems produce confident incorrect statements at non-trivial rates. Studies suggest non-trivial hallucination rates (often >5–10% depending on task and evaluation method). Outputs in high-stakes domains (legal, medical, scientific) require verification by qualified humans — not because AI is wrong often, but because it is wrong confidently, and the consequent harm is disproportionate.
- **Displacement without support:** The projected 15–25% workforce disruption in specific job categories between 2025 and 2030 is not a reason to oppose AI adoption. It is an obligation on organisations, governments, and institutions to invest in reskilling infrastructure at the scale and speed required. The ITIF found that through 2024, AI job creation effects outpaced measured displacement — but the research consensus expects displacement to compound between 2027 and 2030.
- **Concentration of power:** Nvidia has a dominant share (~80–95% depending on segment) of the AI GPU market. Three cloud providers (Amazon, Google, Microsoft) control 70% of the infrastructure on which AI runs. The Stargate Project unites six competitors in a \$500B joint venture under national security framing. Yale Law scholar Madhavi Singh's analysis raises legitimate questions about whether this structure concentrates power in ways that 135 years of antitrust law was designed to prevent.
- **Environmental cost:** Training GPT-3 consumed approximately 1.3 GWh of electricity. Operating ChatGPT costs estimated at hundreds of thousands of dollars per day (early analyses, 2023; not publicly confirmed). US data center electricity consumption rose from 58 TWh in 2014 to 176 TWh in 2023, with projections of 325–580 TWh by 2028. The energy and water footprint of AI infrastructure is a genuine sustainability question.
- **Sycophancy and epistemic risk:** A 2025 mutual alignment evaluation by Anthropic and OpenAI found sycophancy in all models tested — including cases where models validated harmful decisions from simulated users with delusional beliefs. AI systems that tell users what they want to hear, rather than what is accurate, pose systemic risks to professional and scientific epistemic norms.
- **Surveillance and bias in agentic deployment:** AI agents operating with access to internal systems, employee communications, and customer data create surveillance surfaces that existing privacy frameworks do not adequately govern.

Moral Obligations for Organisations

The ethical obligation framework for AI deployment is not primarily philosophical — it is practical. Organisations deploying AI in employment decisions, customer access determinations, or any consequential decision affecting fundamental rights have legal obligations under the EU AI Act from August 2026. These include human oversight requirements, transparency to affected individuals, audit trails, and bias testing. These are not aspirational standards — they are enforceable law with penalties up to 7% of global annual turnover. The moral argument and the legal argument have converged.

11 Education: The Hardest Problem

Of all the domains affected by AI, education may face the most fundamental challenge — because it is not only a sector that AI disrupts, it is the primary institution through which societies equip people to navigate disruption. The response required of educational institutions is simultaneously to adapt their own practices to AI, and to prepare their graduates for an AI-transformed professional world. Many institutions are doing neither well.

The Assessment Integrity Crisis

A 2025 survey of faculty in Spain found that 75% had encountered AI-generated plagiarism at their institutions (Sánchez-Vera et al., 2024). Research from the British Journal of Educational Technology (2025) demonstrates that neither authentic assessments nor AI detection tools provide reliable protection against AI-assisted submission. Detection tools generate significant false positive and false negative rates. The conclusion from this body of research is direct: written assessments, as currently designed, cannot be trusted as reliable indicators of learning. The paradigmatic shift required is from assessing output to assessing process — synchronous oral assessments, reflective accounts, observed demonstrations, project-based work that requires real-time problem-solving in front of evaluators.

The Skills Question

If AI handles literature review, data analysis, first-draft writing, and computational problem-solving, what do students need to learn? The answer is not that these skills become irrelevant — it is that they become foundational prerequisites rather than primary outputs. A student who cannot critically evaluate an AI-generated literature review cannot catch its hallucinations or omissions. A student who cannot formulate a research hypothesis cannot direct an AI tool toward meaningful questions. The 'AI literacy' requirement that the EU AI Act places on all organisations from February 2025 is precisely the recognition that interacting productively with AI requires human expertise, not the replacement of it.

What Institutions Are Actually Doing

University policy responses vary from blanket prohibition to full encouragement. EDUCAUSE's 2025 report found that nearly 80% of higher education institutions now use AI tools in some capacity. The October 2025 analysis of world's top universities shows a trend toward nuanced, course-specific policies rather than blanket bans. Oxford, Cambridge, and Imperial allow AI for research support with mandatory disclosure. MIT encourages AI use with transparency. Harvard leaves AI policy to individual departments. The absence of a coherent pedagogical philosophy behind many of these policies is evident — they are administrative responses to a situation that requires educational leadership.

What Adjustment Actually Looks Like

The required adjustment in education is not technological — it is philosophical. It means accepting that AI changes what humans need to be able to do, not that humans need to do less. Critical thinking, ethical reasoning, creative hypothesis formation, domain expertise that enables AI direction and verification, human communication and relationship management — these become more valuable as AI handles more of the computational load, not less. Education that produces graduates capable of only doing what AI can already do is not just failing its students — it is failing the economy that graduates enter.

12 Industry Bodies: A Responsibility That Cannot Be Delegated

Professional and industry organisations — the GDCh (German Chemical Society), the VCI (Verband der Chemischen Industrie), the ACS (American Chemical Society), and their equivalents across other sectors — have a responsibility in the AI transition that no other institution can fully substitute. They sit at the intersection of technical expertise, professional standards, workforce development, and regulatory advocacy that makes them uniquely positioned to provide the guidance that neither individual companies nor governments alone can deliver credibly.

The question is not whether these organisations should take an active role in guiding their industries on AI — that responsibility is clear. The question is whether the programmes they currently have in place are commensurate with the scale and speed of the challenge. As of the preparation of this paper, IMP does not have detailed knowledge of the current AI guidance programmes of the GDCh, VCI, or ACS, and we do not wish to speculate about gaps that may not exist. What we can state with confidence is what the guidance requirement looks like from an industry and workforce perspective.

What Industry Bodies Are Uniquely Positioned to Provide

- AI literacy frameworks for specific professional roles: chemistry, engineering, regulatory affairs, plant operations. Not generic AI literacy — domain-specific AI competency standards that help professionals understand where AI adds value in their actual work, and where its limitations require human expertise.
- Ethical guidelines for AI deployment in regulated sectors: pharmaceuticals, chemicals, industrial processes. The EU AI Act's high-risk category includes industrial process control and safety-critical systems. Industry bodies can develop sector-specific implementation standards that translate regulatory requirements into operational practice.
- Workforce transition support: the professional organisations that maintain membership relationships with tens of thousands of working chemists, engineers, and industry professionals have a direct channel for reskilling communication that government programmes cannot replicate. Understanding which roles in chemistry and industrial manufacturing face the highest AI displacement risk, and communicating that with clarity and constructiveness, is a genuine service to membership.
- Policy advocacy at EU and national level: on AI Act implementation timelines, on high-risk classification of specific industrial AI applications, on the research investment required for AI safety in chemical and industrial contexts. Industry bodies have the credibility and access to shape policy in ways that individual companies cannot.
- Research coordination: identifying the AI applications in chemistry and industrial process that have the highest value and the highest safety risk, and coordinating the collaborative research programmes that address both.
- Position papers and publicly accessible guidance: not only for members but for the wider public and media, helping counteract both exaggerated AI fear and uncritical AI enthusiasm with technically credible, evidence-based positioning.

The Risk of Institutional Silence

Professional bodies that do not actively guide their members on AI will not prevent their members from adopting AI. They will simply ensure that adoption happens without professional framework, ethical grounding, or technical standards. The consequence is that the standards the body exists to uphold — scientific rigour, professional ethics, safety in industrial practice — are applied unevenly or not at all to one of the most consequential technological transitions those standards will face. Institutional silence on AI is not neutrality. It is a choice with consequences.

13 Recommendations

The following recommendations are grounded in verified research and addressable by the people reading this paper. They are not prescriptions for society or for technology companies — they are practical actions for industrial companies, professional organisations, and individual professionals.

For Industrial and B2B Companies

- Adopt a formal AI governance structure now — before August 2026. Assign a responsible AI owner, conduct an AI system inventory, classify each system against EU AI Act risk tiers, and establish a policy for agentic AI deployment before pilots become production systems. The compliance cost of preparation is a fraction of the enforcement cost of non-compliance.
- Treat AI literacy as a core competency, not an IT training module. Workers who understand what AI can and cannot do, where to trust its outputs and where to verify, are more productive and reduce your hallucination and liability risk. This is a management development requirement, not just a technical one.
- Close the gap between informal and formal AI adoption. The 78% of employees using AI without organisational approval are creating data governance and liability exposure. A policy framework — including approved tools, prohibited uses, and data handling rules — is not restrictive governance. It is the condition for safe adoption at scale.
- Begin reskilling investment in the job categories with highest AI exposure. Do not wait for displacement to become visible before investing in transition. Reskilling from a position of stability is cheaper, faster, and more effective than reskilling from a position of crisis.
- Evaluate agentic AI use cases with legal review before deployment. Do not treat agentic pilots as experiments exempt from governance. An agent that accesses your ERP, emails suppliers, or modifies customer records is performing legally consequential actions from day one.

For Professional and Industry Bodies

- Publish a sector-specific AI guidance document addressing the applications most relevant to your membership — by the end of 2026. This is the minimum output commensurate with the scope of the change your members face.
- Build AI literacy into continuing professional development programmes. Consider certification pathways for AI competence in domain-specific applications.
- Take an active public position on AI and your sector. Not neutrality — an evidence-based, technically credible position that helps journalists, policymakers, and the public understand what AI means for your industry.
- Commission or coordinate primary research on AI adoption, displacement risk, and productivity impact within your specific sector. The generic data is useful; sector-specific data is what your members actually need.

For Individuals

- Use AI tools in your current role. Not because your employer requires it — because the 25% wage premium for AI skills and the 66% productivity gains are real, and they compound. The best time to build AI competence was two years ago. The second best time is now.
- Understand the tools well enough to know when they are wrong. The most dangerous AI user is the one who uses it confidently without verification. Domain expertise — your expertise — is what makes AI outputs useful rather than merely plausible.
- Separate fear from caution. Fear produces paralysis and bad decisions. Caution produces governance, verification, and thoughtful deployment. They are not the same response and they do not produce the same outcomes.
- Invest in the skills AI does not replicate: complex judgment, domain expertise, relationship intelligence, ethical reasoning, the ability to ask the right question. These are the skills that make you a better AI director, not an AI replacement.

14 Three Scenarios 2026–2030

The following scenarios are not predictions. They are structured frameworks derived from current adoption curves, published regulatory timelines, and documented technology development trajectories. Each represents a plausible path given different combinations of factors that are already in motion.

Scenario A - Conservative Adoption

"Measured Integration Under Regulatory Pressure"

Conditions: EU AI Act full compliance burden (August 2026) slows enterprise deployment in regulated industries. Liability uncertainty around agentic AI — particularly following anticipated test cases in employment, lending, and healthcare — leads legal and compliance functions to restrict autonomous agent deployment to low-stakes workflows. Organisational change management bottlenecks (training, process redesign, governance infrastructure) limit effective adoption even where technology is available. Open-source model proliferation continues but enterprise security requirements slow integration in manufacturing and pharma.

Outcome by 2030: AI is widely used as an assistant tool — drafting, summarising, analysing — but autonomous agent deployment remains limited to IT support, scheduling, and customer service tier 1. The productivity premium for high-adopters (PwC: 4× sector productivity growth) compounds, creating visible competitive gaps between early movers and late adopters. Job displacement in administrative categories proceeds more slowly than projected; reskilling demand is manageable but underfunded. AI skill wage premium (PwC: 25%) widens rather than normalises.

What this looks like: Most companies have AI tools. Few have AI workflows. The gap between tool-use and genuine process redesign around AI capabilities remains large. WEF's projected 78 million net new jobs materialises, but with a longer transition period and higher transitional unemployment than the median scenario.

Scenario B - Accelerated Adoption

"Agentic AI Becomes Operational Infrastructure"

Conditions: Competitive pressure from AI-native organisations (and from markets where regulatory burden is lower) accelerates adoption beyond the pace that governance infrastructure can initially support. The \$500B Stargate infrastructure investment comes online ahead of schedule, dramatically reducing inference costs and enabling deployment at scale. Open-source models (LLaMA, Qwen, DeepSeek successors) eliminate cost barriers for mid-size enterprises. Agentic frameworks (Microsoft Copilot, Salesforce Agentforce, AWS multi-agent orchestration) mature sufficiently to handle complex multi-step enterprise workflows with adequate reliability. A new professional category — the AI operations manager — emerges as a standard C-suite function.

Outcome by 2030: Agentic AI is standard infrastructure in knowledge-work-intensive functions: legal document review, financial reporting, supply chain management, regulatory compliance filing, procurement, and R&D literature synthesis. Gartner's projection — one-third of enterprise software includes agentic AI making 15% of day-to-day decisions — proves conservative rather than aggressive. Displacement in administrative, entry-level legal, and routine financial roles is measurable and fast. Goldman Sachs' 6–7% displacement figure applies within this period. Net employment remains positive but geographic and demographic distribution is highly uneven. Educational institutions lag badly; the mismatch between graduate competencies and employer requirements becomes a structural economic issue.

What this looks like: Accenture's 2030 projection — AI agents as primary users of enterprise internal systems — arrives closer to 2028. Human-AI workforce management becomes a standard board-level concern. The EU AI Act's high-risk classification of certain agentic deployments triggers enforcement actions; several high-profile liability cases establish precedent.

14.1 Three Scenarios 2026–2030

Scenario C - Regulatory Slowdown

"Fragmented Compliance Landscape Constrains Global Deployment"

Conditions: The EU–US regulatory divergence deepens beyond current trajectory. Multiple high-profile agentic AI failures — involving autonomous systems making consequential errors in employment, financial, or medical contexts — trigger legislative responses in multiple jurisdictions simultaneously. Data sovereignty requirements proliferate: national AI regulations emerge across Brazil, India, and Southeast Asian markets, each with distinct compliance requirements that make global agentic deployment legally complex. The cost of compliance infrastructure for AI deployment in regulated industries — pharma, chemicals, financial services — becomes a significant barrier for all but the largest organisations.

Outcome by 2030: AI capability continues to advance; AI deployment in regulated industries grows much more slowly than capability growth would allow. Small and mid-size enterprises in regulated sectors effectively cannot deploy agentic AI without disproportionate compliance investment. The productivity premium concentrates among large enterprises with dedicated AI legal and compliance teams. Open-source AI proliferates in less regulated jurisdictions, creating competitive distortions. The net employment impact is lower than Scenario B but the economic benefits of AI are also more narrowly distributed. The WEF's 170 million new jobs projection does not materialise fully within the decade.

What this looks like: Regulatory compliance becomes a genuine competitive moat for large enterprises. The geography of AI deployment diverges sharply: aggressive in US, Middle East, and Southeast Asia; cautious in EU; prescriptive in China. Professional bodies (GDCh, VCI, ACS) play a more active role in defining sectoral implementation standards as the regulatory vacuum creates demand for non-governmental guidance frameworks.

Assessment

None of these scenarios is inevitable. The most probable outcome by 2030 sits between Scenario A and Scenario B for most industrial organisations — significant AI integration in knowledge work functions, limited autonomous agent deployment in regulated production environments, and a workforce transition that is faster than historical technology transitions but slower than the most alarming projections suggest. The variable that most determines which path any given organisation follows is not the technology. It is the quality of governance, the investment in workforce readiness, and the clarity of the strategic decision made in the 2025–2027 window.

15 Closing Argument: The Question We Are Actually Asking

The question this paper was asked to address is ultimately a question about balance. Are we too frightened, too critical, or not critical enough about AI? The answer depends entirely on which AI we are discussing, in which context, deployed by whom, for what purpose, with what oversight. The blanket question does not have a blanket answer.

What we can say with confidence is this: AI is delivering real, measurable, documented benefits in productivity, research acceleration, diagnostic support, formulation chemistry, software development, and dozens of other domains. These benefits are not evenly distributed and they do not resolve the genuine displacement, concentration, and governance risks that accompany them. Both things are true.

The most consequential decisions about AI in the next two years will not be made by technology companies. They will be made by industrial companies deciding whether and how to integrate AI into their operations. By universities deciding what competencies to build and what assessment forms to retain. By professional bodies deciding whether to lead the guidance conversation in their sector or leave it to others. By individual professionals deciding whether to treat AI as a threat to be avoided or a tool to be mastered.

References/Sources:

1. International Monetary Fund (2024). Gen-AI: Artificial Intelligence and the Future of Work. IMF Staff Discussion Note SDN/2024/001. Washington: IMF. (The 40% global job exposure / 60% advanced economies figures)
2. World Economic Forum (2025). Future of Jobs Report 2025. Geneva: WEF. (92M displaced, 170M new roles, 41% employers planning reductions — based on 1,000+ employers, 14M workers, 55 economies)
3. Goldman Sachs Economics Research (2025). How Will AI Affect the Global Workforce? Goldman Sachs Global Investment Research. (2.5% US employment at risk under current adoption; 6-7% under wide adoption; 300M full-time job equivalents; 800 occupation analysis; 0.5 percentage point unemployment rise during transition)
4. McKinsey Global Institute (2025). AI in the Workplace: A Report for 2025. McKinsey & Company. (71% of organisations using GenAI in at least one function; 57% of US work tasks theoretically automatable; 30% of US work hours by 2030)
5. PwC (2025). AI Jobs Barometer. PricewaterhouseCoopers. (25% wage premium for AI skills; 4x productivity growth in high-adoption sectors — based on analysis of close to 1 billion job ads across six continents)
6. Microsoft & LinkedIn (2025). Work Trend Index 2025. Microsoft Corporation. (78% of AI users bring their own tools to work; 75% of knowledge workers using AI — survey of 31,000 people across 31 countries)
7. OpenAI (2025). GPT-5 System Card and Technical Report. San Francisco: OpenAI. (ChatGPT 700M weekly active users, August 2025)
8. SWE-Bench (2024). SWE-bench: Can Language Models Resolve Real-World GitHub Issues? Princeton University / University of Chicago. (4.4% → 71.7% coding benchmark performance, 2023–2024)
9. GitHub (2024). Octoverse 2024: The State of Open Source. GitHub Inc. (46% of new code AI-generated; 20M Copilot developers; 90% Fortune 100 adoption; 400% YoY growth)
10. Vectara (2025). Hughes Hallucination Evaluation Model (HHEM) Benchmark Dataset. Vectara Inc. (Hallucination rates above 10% for GPT-5, Claude Sonnet 4.5, Grok-4 on harder reasoning dataset)
11. Accenture (2025). Technology Vision 2025. Accenture. (AI agents as primary users of enterprise internal digital systems by 2030)
12. Gartner (2025). Top Strategic Technology Trends 2025. Gartner Inc. (By 2028, one-third of enterprise software will include agentic AI; 15% of day-to-day decisions autonomous)
13. DLA Piper (2025). The Rise of Agentic AI: Potential New Legal and Organizational Risks. DLA Piper LLP. (Legal liability framework; principal-agent law analysis)
14. US District Court, Northern District of California (July 2024). Mobley v. Workday, Inc. Case establishing AI screening tool as "agent" of employer clients.
15. NIST (2024). Artificial Intelligence Risk Management Framework: Generative AI Profile (NIST AI 600-1). National Institute of Standards and Technology.
16. European Commission (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council — Artificial Intelligence Act. Official Journal of the European Union, 12 July 2024. (All EU AI Act timeline data, fine structures, GPAI obligations)
17. European Commission (2025). GPAI Code of Practice — Final Version. EU AI Office, July 2025.
18. Sánchez-Vera, J.M. et al. (2024). Faculty perceptions of generative AI in Spanish universities. Published in: Computers & Education. (75% of faculty encountered AI plagiarism)
19. Kofinas, A. et al. (2025). The impact of generative AI on academic integrity of authentic assessments within a higher education context. British Journal of Educational Technology. DOI: 10.1111/bjet.13585.
20. EDUCAUSE (2025). 2025 EDUCAUSE Horizon Report: Teaching and Learning Edition. (80% of institutions using AI in some capacity)
21. OpenAI (2025). Announcing the Stargate Project. Press Release, January 21, 2025. OpenAI.com. (\$500B commitment, initial \$100B, SoftBank/Oracle/MGX structure)
22. DeepSeek AI (2025). DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. DeepSeek AI Technical Report, January 2025. (Training cost ~ \$5.9M; benchmark performance data; MIT licence)
23. Anthropic & OpenAI (2025). Mutual Alignment Evaluation — Cross-Lab Model Testing. Mid-2025. (Sycophancy findings across models)
24. Acemoglu, D. & Restrepo, P. (2019). Robots and Jobs: Evidence from US Labor Markets. Journal of Political Economy, 128(6), 2188–2244. (Displacement vs. reinstatement effect framework; Nobel laureate Acemoglu cited)
25. Goldman Sachs (2025). Tech disruption and labour market tightness. (3 percentage points higher unemployment in tech-exposed occupations for 20–30-year-olds, Q1 2025)

About IMP InterMediaPartners

IMP InterMediaPartners GmbH specialises in B2B marketing and content strategy for complex industrial and technology markets. We help organisations translate technical expertise into market authority through structured demand architecture, knowledge-transfer content, and precision media deployment.

www.intermediapartners.de