A Comprehensive Study on The Adoption of Large Language Models Across Industries

MADHURI GUPTA, ROOBAN AGRAWAL, VERTIKA SHRIVASTAVA, DR ARUN KUMAR CHOUDHARY, DR. MUNISH KUMAR, SUMEDHA ARYA

Abstract: *Large language models (LLMs) have moved from experimental novelty to a standard component of enterprise technology stacks in under four years. This study synthesises survey data, industry reports, and peer-reviewed research to characterise how organisations across eight major sectors, healthcare, banking and financial services, legal services, retail and e-commerce, software development, customer service, education, and manufacturing, are deploying LLM-based tools, and with what measured effect. Drawing on McKinsey's 2025 global AI survey, the World Economic Forum's Future of Jobs Report, and sector-specific data from organisations including the International Legal Technology Association, GitHub, and peer-reviewed economics research, the study finds that adoption is now widespread, with 88% of organisations reporting regular use in at least one business function, but shallow: roughly two-thirds of adopters remain in pilot or experimentation phases, and few report measurable enterprise-level financial impact. Sector-level evidence shows the most rigorously documented productivity gains in customer service and software engineering, while legal, healthcare, and retail show rapid individual-level uptake that is outpacing organisational governance. The study concludes with an analysis of recurring barriers, including hallucination, data privacy, and regulatory uncertainty under frameworks such as the EU AI Act, and outlines priorities for closing the gap between adoption and measurable value.*

Keywords: Large Language Models; Generative Ai; Enterprise Adoption; Artificial Intelligence Governance; Industry Case Studies; Workforce Impact

Madhuri Gupta, Assistant Professor, Department of Management,
Rukmini Devi Institute of Advanced Studies, New Delhi, India
Email: madhuriguptardias@gmail.com
Rooban Agrawal, Assistant Professor, Department of Master of Computer Applications,
Meerut Institute of Engineering and Technology, Meerut, UP, India,
Email: rooban.agrawal@miet.ac.in
Vertika Shrivastava, Head, Department of Computer Application,
Maharaja Agrasen International College, Raipur, Chhattisgarh, India
Email: mail2vertika@gmail.com
Dr Arun Kumar Choudhary, Prof. & Dean, Venkateshwara Open University, Itanagar, AP, India
Email: choudharyarun@rediffmail.com
Dr. Munish Kumar, Business Strategy Manager (IT), Nebraska Department of Labor (NDOL), Dublin OH - 43016, USA
Email: munish2012@gmail.com
Sumedha Arya, IT Project Manager, Cardinal Health, Dublin OH - 43016, USA
Email: arya.sumedha@gmail.com

Introduction

Since the architecture underlying modern large language models was first described in 2017, and especially since general-purpose conversational interfaces reached mass consumer adoption in late 2022, organisations across virtually every sector have begun experimenting with LLM-based tools for tasks ranging from drafting and summarisation to coding, diagnostic support, and customer interaction (Vaswani et al., 2017). What began as isolated pilot projects has, within roughly three years, become a near-universal feature of enterprise technology strategy, even as the depth and financial return of that adoption vary enormously by sector and by firm.

This study draws together survey-based and case-study evidence to characterise the present state of LLM adoption across eight industries chosen for their economic weight and the maturity of available evidence: healthcare, banking and financial services, legal services, retail and e-commerce, software development and information technology, customer service, education, and manufacturing and supply chain. Rather than treating AI adoption as a single undifferentiated trend, the study distinguishes between breadth of use, the share of organisations reporting any use at all, and depth of use, the share that have moved beyond pilots to embed LLM tools in core workflows with measurable effect.

The remainder of the study proceeds as follows. Section 2 outlines scope and methodology. Section 3 summarises overall enterprise adoption trends. Section 4 examines adoption sector by sector. Section 5 compares patterns across sectors. Section 6 discusses recurring barriers and governance challenges. Section 7 considers economic and workforce implications, and Sections 8 and 9 discuss future directions and conclusions.

Scope and Methodology

This study synthesises evidence from primary industry surveys (McKinsey & Company, the World Economic Forum, the International Legal Technology Association, the Higher Education Policy Institute), peer-reviewed research (The Quarterly Journal of Economics and clinical-informatics literature indexed in PubMed Central), and verifiable industry and financial reporting published primarily between 2023 and early 2026. Industries were selected for the combination of available evidence quality and economic weight: the eight sectors examined together account for a substantial share of global employment and output. Because public reporting on AI adoption varies considerably in rigour, ranging from large quasi-experimental studies to vendor-published marketing statistics, this study prioritises figures attributable to a named survey methodology, a peer-reviewed study, or an official regulatory source wherever multiple estimates conflict..

The State of Enterprise Llm Adoption

McKinsey's 2025 global survey of nearly 2,000 executives across 105 countries found that 88% of organisations now report regular AI use in at least one business function, up from 78% the year before, with AI agent experimentation rising to 62% and 23% of organisations reporting they are scaling an agentic system in at least one function (McKinsey & Company, 2025). Yet breadth of adoption has outpaced depth: nearly two-thirds of adopting organisations remain in an experimentation or pilot phase, and only about one-third report having begun scaling AI programmes across the enterprise. Larger firms scale fastest; nearly half of companies with more than five billion dollars in annual revenue report having reached the scaling phase, compared with 29% of firms below 100 million dollars in revenue (McKinsey & Company, 2025).

Financial impact remains concentrated and modest relative to the breadth of adoption. Just 39% of organisations attribute any enterprise-level earnings impact to AI, and most of those report that less than 5% of earnings before interest and taxes is attributable to AI use. Cost benefits cluster in software engineering, manufacturing, and information technology, while revenue benefits concentrate in marketing and sales, corporate strategy, and product development (McKinsey & Company, 2025). A small cohort the survey terms "AI high performers," roughly 6% of respondents, are nearly three times more likely to have fundamentally redesigned workflows around AI rather than layering it onto existing processes, and they invest disproportionately more of their digital budgets in AI capabilities. This gap between broad experimentation and narrow, redesigned deployment recurs, in different forms, across each of the sectors examined in Section 4.

Industry-By-Industry Adoption

The remainder of this section examines adoption in eight sectors chosen for the strength of available evidence and their combined share of global employment and economic output.

4.1 Healthcare and Life Sciences

The most mature enterprise use of large language models in healthcare is ambient clinical documentation. AI scribes that listen to a clinician-patient encounter and draft a structured note have moved from pilot projects to mainstream deployment at large health systems; one widely cited account documented over 2.5 million uses across a single health system's first year of deployment (Tierney et al., 2025). Measured benefits are real but more modest than early marketing suggested: a large multi-site study of 1,800 clinicians across five academic medical centres found that AI-scribe users saved approximately 16 minutes of documentation time and 13 minutes of electronic-record time per eight-hour shift (STAT News, 2026), while a separate quality-improvement study across 45 clinicians in 17 specialties found a median reduction of 2.6 minutes per appointment and a 29% cut in after-hours record-keeping (Beyond Human Ears, 2025). Error rates for LLM-based scribes, roughly 1% to 3%, compare favourably with older dictation systems at 7% to 11%, though clinicians caution that adoption is outpacing independent validation, and that some vendors now position the same tools as revenue-cycle

and coding-optimisation systems rather than purely as burnout remedies (Beyond Human Ears, 2025).

4.2 Banking and Financial Services

Large banks have built some of the most extensive internal LLM deployments of any industry. JPMorgan Chase's internally developed LLM Suite, a model-agnostic platform integrating tools from multiple providers, had been rolled out to more than 230,000 employees by 2026, with an estimated $1.5 billion in annual business value and updates roughly every eight weeks (Tearsheet, 2025; The Digital Banker, 2026). Morgan Stanley's AI @ Morgan Stanley Assistant, built with OpenAI and launched in 2023 for its wealth-management division, lets advisors query roughly 70,000 annual research reports in natural language; by late 2023, 98% of advisor teams were using it regularly (Morgan Stanley, 2023). Consumer-facing deployment has also scaled quickly: Wells Fargo's virtual assistant logged 245 million interactions in 2024, more than double internal projections, with an 80% adoption rate for its Spanish-language version (Tearsheet, 2025). McKinsey has separately estimated that generative AI could add $200 to $340 billion in annual value to the banking sector through automation of knowledge work, with productivity gains of 30% to 50% in targeted functions (McKinsey & Company, 2023).

4.3 Legal Services

Legal services illustrate a widening gap between individual and organisational adoption. The International Legal Technology Association's 2025 survey of nearly 600 firms found 80% using or exploring generative AI, including all firms with 700 or more attorneys, with Microsoft 365 Copilot the most-used tool at 68% of firms, followed by Thomson Reuters' CoCounsel and the legal-AI platform Harvey (ILTA, 2025). A separate survey of more than 1,300 legal professionals found individual attorney use of general-purpose tools rising from 27% in 2024 to 69% two years later, while firm-wide organisational adoption lagged at only 21% (LawNext, 2026). Independent benchmarking lends some support to the optimism: a 2025 evaluation found CoCounsel scoring 79.5% across four legal task categories, well above a lawyer baseline near 50% to 54% on document summarisation and transcript analysis (Toppe Consulting, 2025). Vendor-side investment has followed: Harvey AI raised $760 million in 2025 alone, reaching an $8 billion valuation and more than 1,500 enterprise customers, with one London firm moving from a 70-lawyer pilot to 80% firm-wide daily use within two years (Harvey AI, 2026).

4.4 Retail and E-Commerce

Retail adoption is broad but unevenly documented; an NVIDIA survey of global retail and consumer-goods executives found 80% adopting or piloting generative AI projects, with brands including Sephora, Estee Lauder, and Home Depot using LLM-based tools for product recommendation, customer engagement, and employee knowledge support (NVIDIA, 2025). The clearest measured commercial signal comes from an analysis of 2025 holiday-season traffic, which found that shoppers arriving from generative-AI sources converted at rates 31% higher than those from traditional channels, with revenue per visit from AI-referred shoppers up 254% year over year (Adobe Analytics, 2026). Bain & Company separately reports that 95% of United States companies now use some form of generative AI, up from 83% the previous

year (Bain & Company, 2025), though much of this growth reflects experimentation with content generation and personalisation rather than fully re-engineered retail operations.

4.5 Software Development and Information Technology

Software engineering shows the clearest, most directly measured productivity effect of any sector studied. GitHub Copilot surpassed 20 million all-time users by mid-2025 and is used by 90% of Fortune 100 companies, with enterprise customer adoption growing 75% quarter over quarter (GitHub, 2025). In a controlled study with Accenture, developers using Copilot completed coding tasks approximately 55% faster than a control group, and average pull-request cycle time fell from 9.6 days to 2.4 days (GitHub & Accenture, 2024). Active users report that Copilot now generates close to half of the code they write, up from roughly a quarter at the tool's 2022 launch, illustrating how quickly AI-assisted coding has moved from novelty to default workflow.

4.6 Customer Service

Customer service offers the single most rigorously evidenced productivity study in this review. Brynjolfsson, Li, and Raymond's analysis of more than 5,100 customer-support agents at a Fortune 500 software firm, published in The Quarterly Journal of Economics, found that access to a GPT-based real-time chat assistant raised the number of issues resolved per hour by 14% on average, with a 34% improvement for novice and lower-skilled agents and little measurable effect on the most experienced agents (Brynjolfsson et al., 2025). The authors attribute this pattern to the tool disseminating the conversational patterns of top performers: an agent with only two months' tenure using the tool performed comparably to a six-month-tenure agent without it. The same study found improved customer sentiment and higher employee retention among tool users, making this one of the few sectors where LLM adoption has been linked to worker-level skill transfer rather than only headline productivity metrics.

4.7 Education

Adoption among students has outpaced both institutional policy and instructor readiness. The Higher Education Policy Institute found that the share of UK university students using generative AI tools for assessed work rose from 53% in 2024 to 88% in 2025, with any-tool usage reaching 92%; only 29% of students felt their institution actively encouraged AI use, and a digital divide persisted by gender, subject, and socioeconomic background (HEPI, 2025). In the United States, 69% of high-school students reported using ChatGPT for schoolwork by May 2025, while more than 85% of school administrators viewed AI literacy as a valuable part of the curriculum, even as most students and parents remained ambivalent about whether the benefits outweighed the risks (College Board, 2025). Teacher-side adoption and training have lagged furthest behind, with most K-12 teachers reporting no formal AI training despite widespread informal use.

4.8 Manufacturing and Supply Chain

Manufacturing adoption is real but less extensively documented than other sectors. EY estimates that roughly 40% of supply-chain organisations are investing in generative AI for demand forecasting, risk simulation, procurement negotiation, and predictive maintenance (EY, 2025). At the production level, automotive manufacturers including BMW and Siemens

have deployed LLM-powered systems that translate raw sensor and diagnostic data into natural-language maintenance reports, helping technicians prioritise interventions without manually parsing telemetry. Combined with longer-standing machine-learning-based predictive maintenance, which one academic synthesis associates with 30% to 50% reductions in unplanned downtime, these tools position language models as an interpretive layer atop existing industrial AI rather than a replacement for it (An Automated Survey of Generative AI, 2023). Table 1 summarises representative examples of LLM deployment across the eight sectors examined above.

Table 1. Representative Examples of LLM Adoption by Industry

Industry	**Representative Platform / Tool**	**Reported Adoption or Impact Metric**
Healthcare	Ambient AI scribes (Nuance DAX, Speke, Tandem Health)	Over 2.5 million documented uses in one health system's first year; documentation time cut by up to 16 minutes per 8-hour shift
Banking & Financial Services	JPMorgan Chase LLM Suite; Morgan Stanley AI Assistant	LLM Suite deployed to 230,000+ employees, ~$1.5B estimated annual value; 98% of Morgan Stanley advisor teams using their assistant regularly
Legal Services	Microsoft 365 Copilot, Harvey, CoCounsel	80% of law firms using or exploring GenAI in 2025; individual attorney use rose from 27% to 69% in two years
Retail & E-Commerce	NVIDIA-surveyed deployments (Sephora, Home Depot)	80% of retail/CPG executives adopting or piloting GenAI; 31% higher conversion from AI-referred shoppers
Software Development & IT	GitHub Copilot	20 million+ users by mid-2025; used by 90% of Fortune 100; tasks completed 55% faster in controlled study
Customer Service	GPT-based real-time chat assistants	14% average productivity gain; 34% gain for novice agents
Education	ChatGPT and similar general-purpose tools	88% of UK university students use GenAI for assessments; 69% of US high-schoolers use ChatGPT for schoolwork
Manufacturing & Supply Chain	LLM-based maintenance reporting (Siemens, BMW)	~40% of supply-chain organisations investing in GenAI for forecasting and maintenance

Cross-Industry Patterns

Read together, the eight sectors examined above reveal at least three recurring patterns. The first is a split between knowledge-intensive, text-heavy professions, including law, banking research, customer service, and software engineering, where LLMs operate directly on the firm's core product, whether documents, code, or conversation, and physical or production-oriented sectors, including manufacturing and retail operations, where LLMs more often function as an interpretive layer atop existing sensor data, inventory systems, or supply-chain software rather than as the primary tool of production. Measured productivity gains are correspondingly larger and easier to isolate in the first group: the customer-service and software-engineering studies cited in Sections 4.5 and 4.6 both rely on controlled or quasi-experimental designs precisely because the underlying work is digital and already instrumented, whereas manufacturing and retail studies more often report adoption intentions or investment levels rather than directly measured output changes.

A second pattern concerns the relationship between individual and organisational adoption. Legal services and education show the widest gap: individual professionals and students have adopted general-purpose tools such as ChatGPT far faster than their employing institutions have built governance, training, or formal deployment around that use, leaving informal use ahead of policy in both sectors. Banking and software development show the opposite pattern: adoption has been led top-down by large, well-resourced organisations building purpose-built internal platforms, with individual experimentation following organisational rollout rather than preceding it.

A third pattern is regulatory exposure. Healthcare, banking, and legal services, the three most heavily regulated sectors examined, all show evidence of adoption outpacing independent validation or oversight, whether in the form of clinical scribe deployments still being externally validated, AI-generated legal citations referencing non-existent cases, or new automated coding and compliance reviews introduced specifically in response to LLM-assisted documentation. This suggests that regulatory and professional-liability frameworks, rather than technical capability, are likely to be the binding constraint on how quickly the most heavily regulated sectors can move from individual experimentation to organisation-wide scaled deployment.

5.1 Cross-Cutting Challenges and Governance

Three challenges recur across every sector studied. The first is accuracy: McKinsey's 2025 survey found that 51% of AI-adopting organisations had experienced at least one negative consequence from AI use, with nearly a third citing inaccuracy or hallucination specifically as the cause (McKinsey & Company, 2025). This concern is most acute in regulated, high-stakes domains: healthcare commentators warn that ambient scribes can introduce critical omissions or misattributions even as their overall error rates improve on older dictation tools (Beyond Human Ears, 2025), and law firms have reported instances of AI-generated citations referencing non-existent cases.

The second is governance maturity, which lags far behind tool deployment. Only about one-third of organisations describe themselves as having moved beyond pilots to enterprise-wide

scaling, and the organisations that do scale successfully share specific practices, including defined human-validation checkpoints, redesigned rather than AI-overlaid workflows, and senior-leadership ownership, that most adopters have not yet implemented (McKinsey & Company, 2025).

The third is regulatory uncertainty. The European Union's AI Act, in force since August 2024, imposes transparency obligations on generative AI systems, including mandatory disclosure of AI-generated content and deepfakes, with penalties for the most serious violations reaching 7% of global annual turnover (European Commission, 2026). Implementation of the Act's high-risk-system rules, covering use cases such as employment screening, credit decisions, and education, was pushed from August 2026 to December 2027 under a "Digital Omnibus" package agreed in 2026, reflecting the difficulty regulators and industry alike have had in operationalising oversight for fast-moving generative AI deployments (European Commission, 2026). Outside the European Union, most jurisdictions still rely on sector-specific or voluntary guidance, leaving multinational organisations to navigate a fragmented and shifting compliance landscape even as deployment accelerates.

5.2 Economic and Workforce Impact

The clearest economy-wide projection comes from the World Economic Forum's Future of Jobs Report 2025, based on a survey of more than 1,000 employers across 22 industry clusters and 55 economies. It projects that all macro trends combined, including AI, will create 170 million jobs and displace 92 million by 2030, a net gain of 78 million but a churn rate equivalent to 22% of current employment (World Economic Forum, 2025). Isolating AI and information-processing technologies specifically, the same report attributes roughly 11 million created jobs against 9 million displaced, a narrower but still substantial reallocation concentrated in roles intensive in routine information processing (World Economic Forum, 2025).

Employer expectations diverge from this aggregate optimism in the near term. Forty-one percent of employers surveyed by the Forum expect to reduce headcount specifically where AI automates tasks, even as almost half expect to redeploy affected staff elsewhere in the business rather than eliminate roles outright (World Economic Forum, 2025). McKinsey's parallel data shows a similar pattern at the function level: a median of 17% of respondents reported AI-related workforce declines in the past year, rising to a median expectation of 30% for the year ahead (McKinsey & Company, 2025).

The clearest empirical evidence on how AI affects different workers, rather than how many jobs survive, comes from the customer-service study discussed in Section 4.6: productivity gains concentrated among less experienced workers, with the AI tool functioning as a channel for transferring the skills of top performers to newcomers rather than simply automating existing tasks (Brynjolfsson et al., 2025). Whether this skill-levelling pattern generalises beyond customer service, or whether other sectors will see AI primarily substitute for, rather than complement, less experienced workers, remains one of the most consequential open questions raised by the evidence reviewed here.

Conclusion

Large language models have achieved, in under four years, a breadth of enterprise adoption that took earlier general-purpose technologies a decade or more to reach. Yet the evidence surveyed here shows this breadth has consistently outpaced depth: most organisations remain in pilot mode, measurable financial impact is concentrated among a small minority of high performers, and the strongest documented productivity gains, in software engineering and customer service, derive from a handful of large, well-instrumented studies rather than broad-based industry data. Closing the gap between adoption and value will likely depend less on further model-capability gains than on the organisational practices, workflow redesign, governance maturity, and workforce-transition support, that this study finds concentrated among the minority of adopters already seeing returns.

Several trends are likely to shape the next phase of adoption. First, the shift from conversational assistants toward autonomous agents capable of executing multi-step workflows is already underway: McKinsey's 2025 survey found 62% of organisations experimenting with agentic systems, concentrated in IT and knowledge management, though scaled agentic deployment remains rare in any single function (McKinsey & Company, 2025). Second, regulatory frameworks such as the EU AI Act are likely to become a de facto global baseline as multinational organisations standardise compliance practices across markets rather than maintaining separate regional approaches (European Commission, 2026). Third, the gap this study repeatedly observes between broad experimentation and narrow, redesigned, value-generating deployment is likely to narrow only for organisations willing to redesign workflows around LLM capabilities rather than layering tools onto unchanged processes, the practice McKinsey's data most strongly associates with its small cohort of AI high performers.

References

Adobe Analytics. (2026). 2025 holiday season e-commerce report.

An automated survey of generative artificial intelligence: Large language models, architectures, protocols, and applications. (2023). arXiv:2306.02781.

Bain & Company. (2025). Generative AI adoption among United States companies.

Beyond human ears: Navigating the uncharted risks of AI scribes in clinical practice. (2025). PubMed Central.

Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. The Quarterly Journal of Economics, 140(2), 889-942.

College Board. (2025). New research: Majority of high school students use generative AI for schoolwork. College Board Newsroom.

European Commission. (2026). AI Act: Regulatory framework for artificial intelligence. Shaping Europe's Digital Future.

EY. (2025). Generative AI investment trends in supply-chain organisations.

GitHub. (2025). GitHub Copilot surpasses 20 million users.

GitHub & Accenture. (2024). Quantifying GitHub Copilot's impact on developer productivity and happiness.

Harvey AI. (2026). Ten legal AI adoption metrics for law firms.

HEPI (Higher Education Policy Institute). (2025). Student Generative AI Survey 2025.

ILTA (International Legal Technology Association). (2025). 2025 Legal Technology Survey.

LawNext. (2026). AI adoption among legal professionals has more than doubled in a year.

McKinsey & Company. (2023). The economic potential of generative AI: The next productivity frontier.

McKinsey & Company. (2025). The state of AI in 2025: Agents, innovation, and transformation.

Morgan Stanley. (2023). AI @ Morgan Stanley Assistant: Generative AI for wealth management advisors.

NVIDIA. (2025). Generative AI in retail and consumer packaged goods: Executive survey.

STAT News. (2026). Large AI scribe study finds modest time savings, inconsistent use.

Tearsheet. (2025). The AI reality check: Q1 2025 edition.

The Digital Banker. (2026). JPMorgan Chase's LLM Suite drives AI transformation across the enterprise.

Tierney, A. A., et al. (2025). Ambient artificial intelligence scribes: Learnings after 1 year and over 2.5 million uses. NEJM Catalyst, 6(CAT.25.0040).

Toppe Consulting. (2025). CoCounsel vs. Harvey AI vs. Lexis+: Which legal AI tool delivers the best ROI in 2025?

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30.

World Economic Forum. (2025). The Future of Jobs Report 2025.