



International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

Artificial Intelligence in Modern Computing: A Comprehensive Review of Applications, Challenges, and Future Trends

S. Jyothrimaye Reddy^{1*}, Radha Thangarajan², Mahadev, A. Gawas³, S. Poongavanam⁴, Ishwar Prasad Sahu⁵

^{1*}Associate Professor, Marwadi University, E-mail: sathi.reddy@marwadieducation.edu.in, ORCID: 0000-0002-1629-5777, sathi.reddy@marwadieducation.edu.in

²Associate Professor, School of Commerce and Management, St. Clare College Autonomous, Bangalore, India radhasaran26@gmail.com

³Assistant Professor, Department: State Higher Education Council, Institution : Directorate of Higher Education, Goa, India, Mail – gawas-dhe.goa@gov.in, ORCID: 0000-0002-5719-3830

⁴Professor, Amet Business School, Amet University, Chennai, Tamil Nadu, India, ORCID: 0000-0002-5798-7251, E-mail: s.poongavanam@gmail.com

⁵Assistant Professor, Department of Mechanical Engineering, Rungta International skills University, Bhilai, Rungta College of Engineering & Technology, 490024, Chhattisgarh, India, * ORCID ID: 0009-0006-6408-7514, Email: ishwarprasad.sahu@rungta.org

Abstract

Artificial intelligence (AI) has evolved from a novelty to a necessity in today's world. By consolidating the existing secondary quantitative indicators and analyzing the cross-cutting data, this study investigated the current state of AI, its possibilities, and problems. Differing from the conventional review, the paper utilized the latest practices, possibilities, investments, infrastructure, and regulations about AI as the arguments and extracted the comparative data from them. In particular, the article analyzed the Stanford AI Index 2026, International Energy Agency's report 'Energy and AI', NIST AI Risk Management Framework, EU Artificial Intelligence Act, UNESCO guidelines, and many studies related to AI in healthcare, cybersecurity, and computing. The descriptive statistics of the studies demonstrate that 88% of the organizations have adopted AI in 2025, and 70% of them have utilized generative AI technologies in at least one area of their operations, implying that the difference between the lines is 18%. Meanwhile, the utilization of AI agents is still below 10%, which reflects the difference of 78% between AI agents and overall AI applications. Moreover, the performance of models on SWE-bench Verified increased from 60% to almost 100% per annum, and the electric power consumption of data centers is forecasted to increase from 415 TWh in 2024 to 945 TWh in 2030, which is 127.7% higher and constitutes 14.7% average annual growth. The domain scoring analysis also identified cybersecurity, finance, healthcare, and industrial systems as the highest AI impact areas needing governance attention. Based on the findings, the paper proposed a five-lever Modern Computing AI Integration Framework, covering data, infrastructure, models, applications, and governance. Ultimately, the study concluded that the most pressing question in modern computer science was not whether to embrace artificial intelligence, but whether its rise and accompanying expansion and investment would serve the needs of reliability, safety, and energy, security, and human-centered governance, and whether it would overcome these challenges.

Keyword: artificial intelligence; modern computing; generative AI; AI adoption; cloud computing; edge AI; cybersecurity; responsible AI; agentic AI; sustainable computing

Introduction

The concept of artificial intelligence (AI) in present-day computers is generally the process by which they learn from a specified data set, automate, operate interactive, and adapt to user needs as well as changing circumstances. Older programs relied on deterministic algorithms where the programmers usually laid down the relationships between inputs and outputs. Contemporary AI has started to take over the process, enabling the systems to infer patterns, generate content, predict outcomes, and increasingly make selections based on the learned representations. Such a transition has been enabled by deep learning, transformer architecture, foundation models, generative AI, and specific hardware (LeCun, Bengio, & Hinton, 2015; Vaswani et al., 2017; Bommasani et al., 2021).

The scale of adoption makes AI an infrastructure issue rather than an algorithmic issue. 88% of companies adopted AI at a managerial level in 2025, according to the Stanford Institute for Human-Centered Artificial Intelligence. Generative AI was used in at least one business function in 70% of surveyed companies. More than 90% of notable frontier models were produced by industry, and the technical frontier was closely tied to corporate computing infrastructure (Stanford HAI, 2026). These figures indicate that AI is embedded within cloud servicing, corporate software, cyber defense, software-development, healthcare platforms, financial systems, educational tech and industrial computers.

The rapid proliferation is associated with a second transition. The AI systems are becoming capable of operating within the confines of software environments. On the SWE-bench Verified, a benchmark designed on real software-

engineering tasks, the frontier performance has improved from approximately 60% to near 100% in one year (Stanford HAI, 2026). This is a significant improvement because the modern computing relies increasingly on the AI for not only prediction but also code generation, debugging, information retrieval, execution of work flow and use of tools. Yet, the same report remarks that the AI agent deployment is still in the single digits for most of the business functions. The discrepancy between the ubiquitous adoption of general AI and the use of autonomous-agent is a significant gap of readiness.

A third transition is represented by the material requirements of large-scale AI. AI model training and inference require an extraordinary amount of compute capacity, which typically takes the form of data centres and accelerated servers. The International Energy Agency estimates global electricity consumption of data centres to amount to 415 TWh in 2024, and 945 TWh in 2030 in the International Energy Agency's base case scenario (IEA, 2025). The electricity implications of AI are of utmost importance and intersect AI development with questions of sustainable computing and grid capacity, hardware efficiency and capital investments

These interrelated phenomena raise a research question that cannot be solved by considering only the applications of new technologies. The development of modern computer science is inseparable from such factors as the level of its implementation, accompanying risks and threats, and the pressures on its infrastructure. The current study undertakes a secondary data empirical analysis and develops a conceptual framework that allows the researcher to evaluate the situation with the critical implementation of artificial intelligence in the IT field. In other words, the paper aims to provide a comprehensive assessment of the latest authoritative data related to the use of new technologies in computer science.

1.1 Problem Statement

Although AI is a common topic of discussion across fields of computing, much discussion separates out areas such as applications, technical performance, governance and infrastructure. This makes it difficult to assess the overall level of maturity of AI-enabled computing. High model capability does not necessarily imply safe deployment, adoption does not necessarily imply autonomous operation, and rapid AI diffusion can impose infrastructure and regulatory burdens that are not captured in conventional accuracy metrics. There is therefore a need for a research framework connecting adoption, capability, application intensity, risk and demand.

1.2 Research Objectives

- To quantify recent changes in AI adoption, model capability, investment and computing-infrastructure demand using verified secondary data.
- To compare the intensity of AI integration and governance burden across major modern-computing application domains.
- To identify the principal technical, operational, ethical and sustainability challenges that constrain reliable AI deployment.
- To develop a Modern Computing AI Integration Framework that connects data, infrastructure, models, applications and governance.
- To derive future research priorities from the gap between AI capability, adoption and deployment readiness.

1.3 Research Questions

1. What do recent quantitative indicators say about AI diffusion/capability rates in contemporary computing?
2. What are the computing domains with the highest combination of AI integration intensity and governance burden?
3. How fast is the infrastructure requirement for AI-enabled computing growing?
4. What technical and governance mechanisms are needed to transition from assistive AI to autonomous and agentic AI responsibly?

2. Background and Theoretical Foundation

2.1 From Machine Learning to Foundation Models

Machine learning was a paradigm shift in programming that allowed models to learn statistical regularities from data rather than relying only on hard-coded instructions (LeCun et al., 2015). The concept of deep learning builds on the previous paradigm by enabling the network to learn high-level abstractions of images, sound, and text using multi-layer neural networks (LeCun et al., 2015). The newly introduced transformer architecture allows to create flexible models that are able to perform a variety of tasks through the use of an innovative mechanism of attention (Vaswani et al., 2017).

Foundation models represent a class of large-scale systems, the main difference of which is that the model can be fine-tuned for different downstream tasks (Bommasani et al., 2021). However, modern generative AI can do more than that, namely, it can generate text, code, images, audio, and other data. This is a huge step for the modern computer science industry since artificial intelligence can become a universal tool for working with other programs and databases, which would not be limited to only one specific task.

2.2 AI as a Computing-System Layer

The use of artificial intelligence can be seen as comprising several areas, such as data acquisition, model creation, accelerator hardware, cloud or edge execution, application logic, human interaction, and governance. In other words, it can be asserted that the effectiveness of any AI application is determined by factors beyond the quality of the produced model. Those crucial elements are data, system security, latency, number of processors, system monitoring, access control, and organization's rules regarding the use of technology. Notably, the NIST AI Risk Management Framework highlights the need to perform activities of Govern, Map, Measure, and Manage, which implies that trustworthiness must be integral in every stage of an AI system's life cycle (Tabassi, 2023).

2.3 High-Impact Application Contexts

AI is currently being applied in healthcare, cybersecurity, finance, education, manufacturing, software engineering and scientific computing. In healthcare, applications range from imaging to risk prediction, workflow support and decision assistance in surgery (Varghese et al., 2024). For cybersecurity, typical applications are anomaly detection, malware classification, predictive analytics and automated response (Ofusori, Bokaba, & Mhlongo, 2024). Generative AI in education enables tutoring, content generation and feedback but also creates privacy, academic integrity and verification challenges (Miao & Holmes, 2023). The fields have dramatically different risk profiles, time-criticality, regulatory requirements and error tolerance, therefore an assessment across domains should take into account the degree of integration and governance overhead.

3. Materials and Methods

3.1 Research Design

The article employs the secondary-data empirical research design and the domain-analysis rubric. The rationale behind the former choice was that the research problem pertains to a phenomenon of systemic change, which has been already operationally measured by an international authority. Instead of resorting to a fictional population of survey respondents, the authors use the already published quantitative indicators as statistical entities, which they analyze and re-synthesize using descriptive calculations. The analytical unit in this study is the modern-day AI-centric computer system, while the application domains are the comparative cases.

3.2 Data Sources and Inclusion Criteria

Data were selected according to the following inclusion criteria: (1) source type: authoritative institutional report, official regulation, or peer-reviewed publication; (2) indicator relevance: directly measurable aspects of AI adoption, capability, infrastructure, risk, or application; (3) temporal relevance: underlying period sufficiently recent to capture generative and agentic AI developments; (4) interpretability: reported value could be understood without requiring reconstruction of inaccessible proprietary datasets. The quantitative foundation was provided by the Stanford AI Index 2026 and the IEA Energy and AI report. Governance frameworks were analyzed using NIST AI RMF 1.0, the EU Artificial Intelligence Act and UNESCO guidance. Peer-reviewed studies provided context for domain-specific application characteristics.

Table 1. Secondary indicators used in the empirical analysis

Indicator	Reported value	Period	Primary source
Organizational AI adoption	88%	2025	Stanford HAI (2026)
Generative AI used in ≥1 business function	70%	2025	Stanford HAI (2026)
Population-level generative AI adoption	≈53%	2026 reporting	Stanford HAI (2026)
University students using generative AI	≈80%	2026 reporting	Stanford HAI (2026)
Industry share of notable frontier models	>90%	2025	Stanford HAI (2026)
SWE-bench Verified performance	60% → near 100%	1-year change	Stanford HAI (2026)
Private AI investment growth	+127.5%	2025	Stanford HAI (2026)
AI agent deployment	Single digits in most functions	2025	Stanford HAI (2026)
Data-centre electricity consumption	415 TWh	2024	IEA (2025)
Projected data-centre electricity consumption	945 TWh	2030 base case	IEA (2025)

3.3 Derived Measures

Three estimates were calculated. First, the organizational implementation gap is defined as the difference between overall organizational AI adoption and organizational generative-AI use. Second, the autonomy-readiness gap was conservatively estimated by assuming that overall organizational AI adoption minus the upper bound of single-digit agent deployment (10%) equals the actual gap; the actual gap is therefore higher than the estimated 8%. Third, energy growth was calculated from the 2024 IEA estimate and the 2030 base-case projection and expressed as a compound annual growth rate (CAGR): $(945/415)^{(1/6)} - 1$.

3.4 Cross-Domain Scoring Procedure

To compare application contexts, each of the seven domains was scored on the basis of a five-point ordinal rubric for five dimensions. These dimensions were breadth of application of artificial intelligence, degree of autonomy, dependence on real-time processing, potential negative impact of errors, and governance burden. Each dimension was scored from 1 (lowest) to 5 (highest). The level of integration was defined as the average score across the first three dimensions, whereas governance burden reflected the average score across the last two dimensions. The scoring is intended to be an analytical rather than an empirical exercise and should be viewed as an approximation of the relative differences among domains.

Table 2. Cross-domain AI integration and governance scoring rubric

Domain	Breadth	Autonomy	Real-time	Error consequence	Governance burden	Integration intensity	Risk burden
Software engineering	5	4	2	3	3	3.67	3.00
Cybersecurity	5	4	5	5	5	4.67	5.00
Healthcare	4	3	4	5	5	3.67	5.00
Finance	5	4	4	5	5	4.33	5.00
Education	4	3	2	4	4	3.00	4.00
Manufacturing & robotics	4	4	5	5	4	4.33	4.50
Scientific computing	4	3	3	3	3	3.33	3.00

3.5 Analytical Approach

The study uses descriptive statistics, percentage-point differences, relative change, CAGR and structured comparison. It makes no claim of causal inference since the data is drawn on heterogeneous secondary sources rather than randomized or longitudinal primary dataset. It aims to identify patterns at the level of systems, pertinent to the computing research and then translating the patterns into an integration framework.

4. Results

4.1 AI Diffusion Is Broad, but Deployment Depth Is Uneven

Overall organizational AI adoption reached 88%, while generative AI was used in at least one business function by 70% of organizations. The resulting implementation gap is 18 percentage points. This suggests that many organizations identify themselves as AI users even though generative AI is not yet embedded across all operational functions. Population-level generative AI adoption of approximately 53% also implies that consumer diffusion is substantial but lagging in comparison to enterprise-level adoption reporting.

The more substantial gap appears to be in the area of autonomous operation. The deployment of AI agents remained in single digits across nearly all areas of business operations. Using 10% as a conservative ceiling, the autonomy-readiness gap between general AI adoption stands at a minimum of 78 percentage points (88% - 10%). This suggests that organizations are significantly more open to using AI as an auxiliary or embedded capability than to delegating multistep autonomous actions to AI systems.




Figure 1. Selected indicators of AI diffusion and capability. Source data: Stanford HAI (2026). ">90%" and "near 100%" are reported approximations.

4.2 Capability Growth Is Outpacing Operational Assurance

Performance on SWE-bench Verified increased from roughly 60% to nearly 100% within one year, an increase of roughly 40 points. This is extremely important because it is in the domain of software engineering, where the benchmarks are about real world repository-level problems, as opposed to just making up snippets of code. This means that AI systems are growing more capable of performing tasks that require understanding, modifying or reasoning about code in a multi-step process.

However, this does not mean that the capabilities are limitless, as benchmarks can saturate, the deployment environment can vary, and models can behave differently when interacting with external tools, credentials, data stores, and permissions. The gap between the capability of a model as demonstrated on a benchmark and its actual deployment as an agent is what leads to the interpretation that operational assurance, rather than capability, is the new constraint on the autonomy of AI systems.

4.3 Computing Infrastructure Pressure Is Increasing Rapidly

IEA statistics state that the use of electricity by data centers will grow from around 415 TWh in 2024 to almost 945 TWh by 2030. In such a way, the demand for electricity by data centers will increase by 127.7% over six years. The CAGR of such a rise will be 14.7%. According to the IEA, AI is the most significant driver of this increase along with other digital services, which has accelerated servers' electricity consumption (IEA, 2025).




Figure 2. Growth in global data-centre electricity consumption. Source: IEA (2025); CAGR calculated by the present study.

4.4 Application Domains Differ in Integration Intensity and Risk

Cross-domain rubric indicates that cybersecurity has the highest integration intensity score (4.67) and the maximum risk-burden score (5.00). These two factors reflect wide application of artificial intelligence, need in

day-to-day functioning, adversarial nature of the environment and potential negative impact on the security in case of error. Finance and manufacturing/robotics also have a high integration intensity (4.33), but finance is characterized by a higher governance burden due to ability of algorithms to make decisions affecting credit, trading, fraud detection and regulation.

Healthcare sector has a moderate-to-high rate of integration-intensity (3.67), but a maximum risk-burden score (5.00) suggesting that severe consequences are associated with clinical errors and professional liability. The education segment has a relatively low integration-intensity score (3.00) but involves a significant governance burden since AI impact on learners' privacy, testing practices and quality of educational information. Scientific computing has the lowest risk burden in the scoring matrix (3.00) of the assessed domains but reproducibility, data provenance and scientific validity are major concerns.

Table 3. Ranking of domains by integration intensity and risk burden

Rank	Domain	Integration intensity	Risk burden
1	Cybersecurity	4.67	5.00
2	Finance	4.33	5.00
3	Manufacturing & robotics	4.33	4.50
4	Software engineering	3.67	3.00
5	Healthcare	3.67	5.00
6	Scientific computing	3.33	3.00
7	Education	3.00	4.00

4.5 Capital Growth Reinforces Infrastructure Concentration

Stanford HAI noted that private AI investment grew 127.5 percent year-over-year in 2025, and that investment in generative-AI grew by over 200 percent. Industry also produced over 90 percent of noteworthy frontier models. Combined, these are signs that frontier AI development is becoming increasingly capital-intensive and concentrated in entities with the ability to fund compute, data, specialized talent and infrastructure. This does not mean that smaller models or open systems are not important, but rather that the frontier of capability training is strongly influenced by those with large-scale capital and computing power.

5. Proposed Modern Computing AI Integration Framework

In the light of the presented results, the current study put forward the Modern Computing AI Integration Framework (MCAIF). The framework suggests that successful AI application deployment should be considered as a five-level process. It emphasizes that AI model choice is not the key point in the system, while requirements to each level define the upper level's criteria. At the same time, governance should be regarded as a cross-cutting assurance dimension.




Figure 3. Modern Computing AI Integration Framework (MCAIF) proposed in this study.

5.1 Layer 1: Data and Knowledge

AI systems rely on structured data, unstructured data, images, code, sensor data, and enterprise data. The main aspects of research include data quality, provenance, representativeness, access control, and lifecycle governance. Faulty or low-quality data may lead to errors at all stages below.

5.2 Layer 2: Computing Infrastructure

This layer comprises of cloud platforms, edge devices, accelerators, data centres, communication networks and storage. An analysis of the energy use demonstrates that infrastructure efficiency should be considered as a primary research variable in AI, not as a background cost. Edge AI and smaller specialized models can reduce the latency and transfer of data, while the cloud infrastructure allows large-scale training and inference.

5.3 Layer 3: AI Models and Intelligence

This layer consists of machine learning, deep learning, foundation models, multimodal systems, and agentic models. The rapid rise in benchmarks reflects the need to focus on the model-level improvement, and the framework emphasizes the importance of evaluating reliability, robustness, security, and task-specific appropriateness in addition to standard metrics.

5.4 Layer 4: Applications and Agents

Application logic dictates the influence of AI outputs on real-world users and systems. An AI assistant that recommends code is different to an agent authorised to modify production software. Similarly, a clinical decision-support system is different to an autonomous medical decision-maker. The framework thus requires explicit permission boundaries, escalation paths and human-control points.

5.5 Layer 5: Governance and Assurance

Governance encompasses risk assessment, accountability, monitoring, audit, human factors, and regulatory compliance. NIST's AI RMF promotes ongoing risk management, whereas the EU AI Act imposes specific regulatory requirements on AI systems of a certain risk level. Governance is the top-level function but should be embedded in all other elements of the framework during design and implementation.

6. Challenges in AI-Enabled Modern Computing

Data quality and bias

AI performance is determined by the representativeness and accuracy of data. Historical bias, insufficient sampling, label errors and domain shift can cause unstable or unjust outputs. High-impact systems thus require documentation of data sets, bias testing and post-deployment monitoring.

Hallucination and factual reliability

Generative models may generate fluent but incorrect outputs. In software, healthcare, finance and research, plausible errors can be more insidious than apparent bugs. Retrieval grounding, tool verification, constrained generation and human validation help, but do not solve the problem

Cybersecurity and adversarial use

AI systems provide attack surfaces such as prompt injection, data poisoning, model extraction, adversarial examples and malicious tool use. Agentic systems exacerbate these risks since these models have access credentials or perform actions on behalf of an entity. Sandbox hardening, least-privilege access and audit logging are required.

Privacy and confidential information

The modern applications of artificial intelligence tend to involve processing private, organizational, financial, health or biometric data in one way or another. Thus, there is a privacy concern regarding possible data use during training, inference, logging, retrieval and third-party processing, which motivates the need for privacy-preserving learning, local inference, encryption and data minimization.

Explainability and accountability

Complex models may pose challenges in interpretation. In high risk contexts there is a need for users to understand the basis, limitations, and uncertainty of an AI supported decision. Explainability must be accompanied by accountability as an explanation on its own does not identify who is responsible for a harmful outcome.

Energy and infrastructure demand

The 127.7% surge in data centre electricity needs by 2030 (over 2024) is merely an example of a more substantial sustainability problem. Research on efficiency should focus on the consequences of model compression, quantization, specialized hardware, workload scheduling, or edge inference plus the energy source that will be utilized to power such operations.

Regulatory compliance

The EU Artificial Intelligence Act proposes a risk-based regulatory framework. At the same time, NIST provides a voluntary risk management framework. Entities operating in these jurisdictions will need to adopt technical documentation, risk assessment, logging, transparency, and governance practices that are applicable to multiple standards.

Concentration and access inequality

The development of frontier systems requires vast resources. Large foundational models and their derivatives are currently monopolizing industry resources. This undermines the value proposition of smaller-scale solutions, and

creates challenges for research accessibility, reproducibility and vertical integration. These issues are less acute for smaller, standardized models, and more efficient architectures.

Table 4. Challenge-to-control mapping for AI-enabled computing

Challenge	Primary impact	Recommended control direction
Data bias / poor quality	Unfair or unstable outputs	Data governance, bias testing, provenance, monitoring
Hallucination	Incorrect decisions or content	Grounding, verification, confidence checks, human review
Prompt injection / tool abuse	Unauthorized actions or data leakage	Sandboxing, least privilege, input filtering, audit logs
Privacy exposure	Disclosure of sensitive data	Data minimization, encryption, local processing, privacy-preserving learning
Opaque decisions	Low trust and weak accountability	Explainability, documentation, model cards, escalation rules
High energy demand	Cost and sustainability pressure	Efficient models, accelerators, edge inference, workload optimization
Regulatory mismatch	Compliance and liability risk	Risk classification, documentation, governance controls
Provider concentration	Dependency and unequal access	Open standards, portable architectures, smaller domain models

7. Discussion

The results demonstrate that modern computing has entered an era in which AI diffusion is no longer the central uncertainty. With 88% of organizations adopting AI, the key questions concern the degree of integration and the level of operational authority that it is granted. The 18-points gap between overall AI adoption versus generative-AI business use suggests that there is room for improvement, while the significantly bigger autonomy-readiness gap shows that firms remain hesitant to delegate actions to agents.

This is not without foundation, because a recommendation from a generative model can be subject to review before being implemented, but an AI agent embedded with enterprise tools can directly affect operations. The gap between extremely high standards in the system and inadequate agent deployment therefore indicates a governance and systems engineering problem. Future agent research must address permission boundaries, tool use safety, task verification, reversibility, auditing and escalation to humans.

The domain analysis lends support to this conclusion. Cybersecurity, finance, healthcare and industrial computing have a combination of high-application-potential and high-consequence-if-error. Domains like this are likely to be appropriate for AI, but are among the least appropriate for unbounded autonomy. In these domains the objective should not be to maximize the extent of automation, but to calibrate the degree of automation to the extent of reliability and the need for appropriate human oversight.

The energy analysis suggests another constraint. A data-centre electricity CAGR implied by the projection is in excess of 14.7%. This is a large number for infrastructure that must be supported by grids, cooling, facilities and capital investment. Algorithmic efficiency is therefore a strategic research priority. Advances in quantization, sparsity, model routing, efficient attention, local inference and domain-specific models can have a significant impact on not only the cost but also the environmental sustainability and accessibility of AI.

The proposed MCAIF connects the two findings by conceptualizing AI as a system consisting of data, infrastructure, models, applications, and governance. In this vein, this paper advances a layered perspective that can provide a better understanding of the ways in which successes in one domain can give rise to issues in others. For instance, an increase in the quality of the models can lead to an increase in the number of applications while also creating higher demands for infrastructure and governance. At the same time, problems in the domain of data or insecure infrastructure can have a detrimental effect on the performance of a sophisticated model.

8. Future Research Directions

Agentic AI assurance

Consider the application of formal methods to permission control, tool authentication, action verification, rollback, memory-security and human escalation in the context of multi-agent and autonomous workflows

Energy-aware model design

Report energy, latency, and hardware utilization in addition to accuracy. Research should compare the possibility of achieving acceptable task performance with a smaller and more specialized model that would require significantly less energy.

Edge and distributed intelligence

Build a research plan to investigate how cloud and edge models can work together to do cool stuff like privacy-preserving synchronization, personalization, and network-aware model selection.

AI security engineering

Standardize red-team practices for prompt injection, data poisoning, model extraction, adversarial inputs and agent-tool attacks. Security testing should be part of the development cycle

Domain-specific reliability metrics

Generic benchmarks are insufficient for high-risk applications. Healthcare, finance, cybersecurity and industrial systems require domain-specific metrics of uncertainty, error severity, and operational robustness.

Human-AI decision allocation

Study which of the tasks should be done by technology, by people, and by collaboration. Study should consider factors other than productivity, such as error detection, skills maintenance and accountability.

Governance interoperability

Create technical mappings between NIST AI RMF framework and the EU AI Act and other standards and regulations to allow organizations to implement controls once and demonstrate compliance to multiple frameworks.

Open and efficient AI ecosystems

Investigate open-weight models, modular architectures and interoperable AI services as means of reducing dependence on concentrated frontier infrastructure while preserving security and quality.

Table 5. Evidence-based future trends and associated research needs

Future trend	Evidence from results	Priority research need
Agentic AI	General AI adoption is high while agent deployment remains single digit	Secure autonomy, verification, permission boundaries
Efficient / green AI	Data-centre electricity projected to reach 945 TWh by 2030	Energy-aware training and inference
Smaller domain models	Frontier training is capital and compute intensive	Efficiency, specialization, local deployment
Edge AI	Need for low-latency and privacy-aware computing	Cloud-edge coordination and local adaptation
AI governance	High-risk applications face regulatory and accountability demands	Continuous risk assessment and interoperable controls
Human-AI collaboration	High capability does not guarantee autonomous readiness	Task allocation, oversight and skill preservation

9. Research Implications

9.1 Implications for Researchers

AI research needs to go further by exploring evaluation beyond accuracy and benchmark leadership. Such system-level metrics as energy, latency, calibration, security, explainability, operational failure rate and human-intervention frequency are becoming more and more indispensable for meaningful comparison. Research claims should also differentiate between model capability and deployment readiness.

9.2 Implications for Industry

Organizations should assess the application of artificial intelligence as a portfolio of use cases, not as a binary maturity category. Demands on high-risk applications require more extensive validation, access controls, monitoring, and documentation than less sophisticated applications of AI, such as productivity tools. The big gap between high-risk readiness and autonomy suggests that organizations need to establish governance before expanding the scope of agents' activities.

9.3 Implications for Policy

Policy frameworks should be designed to encourage innovation, while recognizing that AI risk is context dependent. Risk-based regulation is preferable to regulating all systems as if they're at equivalent risk. At the same time, policy should take into account infrastructure concentration and sustainability as the capacity to develop frontier AI is strongly coupled to access to capital, energy and advanced computer hardware.

10. Limitations

This study has four drawbacks. Firstly, the quantitative analysis employed aggregated secondary-indicator data that rendered the findings non-causal. Secondly, the analysis included approximated figures such as near 100% benchmark achievement and single-digit agent utilization. Thirdly, the cross-domain scoring rubric was developed for the purpose of the current analysis and, thus, its scoring criteria call for validation through expert-judgment panels and primary-indicator data collection. Fourthly, the IEA 2030 electricity-value estimation is a projected scenario value that is susceptible to hardware efficiency, AI demand, economic and grid-availability factors.

11. Conclusion

This research analyzed the phenomenon of artificial intelligence in the context of modern computer science, using secondary quantitative data and a domain analysis method. The findings indicate that the spread of AI technology

follows two trends. On the one hand, it is used in the corporate environment more frequently than ever before. On the other hand, the level of acceptance varies significantly depending on the type of technology. Thus, while overall adoption rates reached 88%, generative AI tools were adopted by 18% fewer organizations, and agent-type AI, in single digits. The gap between the introduction of new technologies and the readiness of companies to rely on them fully explains the discrepancy.

The study further unveils modern computing infrastructure as a critical bottleneck. Data-centre electricity consumption is predicted to surge from 415 TWh in 2024 to approximately 945 TWh in 2030 – representing a 127.7% growth and an implied CAGR of approximately 14.7%. This dynamic renders energy-efficiency, model-efficiency and infrastructure planning as key research areas for the future of artificial intelligence.

A cross-domain analysis reveals cybersecurity, finance, healthcare and industrial systems as prime domains for AI-integration, while simultaneously imposing the most stringent governance requirements. The proposed Modern Computing AI Integration Framework thus positions governance and assurance as inextricably linked to data, infrastructure, models and applications. The future of artificial intelligence in modern computing will increasingly be defined by the ability to deploy capabilities that are secure, efficient, transparent and appropriately accountable rather than solely focused on their ability to accomplish increasingly complex tasks.

References

1. Bommasani, R., Hudson, D. A., Adeli, E., et al. (2021). On the opportunities and risks of foundation models. arXiv:2108.07258.
2. Brown, T. B., Mann, B., Ryder, N., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877-1901.
3. European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, 2024/1689.
4. International Energy Agency. (2025). *Energy and AI*. Paris: IEA.
5. Jumper, J., Evans, R., Pritzel, A., et al. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596, 583-589. <https://doi.org/10.1038/s41586-021-03819-2>
6. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436-444. <https://doi.org/10.1038/nature14539>
7. Miao, F., & Holmes, W. (2023). *Guidance for generative AI in education and research*. Paris: UNESCO.
8. Ofusori, L., Bokaba, T., & Mhlongo, S. (2024). Artificial intelligence in cybersecurity: A comprehensive review and future direction. *Applied Artificial Intelligence*, 38(1), 2439609. <https://doi.org/10.1080/08839514.2024.2439609>
9. Stanford Institute for Human-Centered Artificial Intelligence. (2026). *The 2026 AI Index Report*. Stanford University.
10. Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
11. Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433-460.
12. Varghese, C., Harrison, E. M., O'Grady, G., et al. (2024). Artificial intelligence in surgery. *Nature Medicine*, 30, 1257-1268. <https://doi.org/10.1038/s41591-024-02970-3>
13. Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
14. Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K., & Zhang, J. (2019). Edge intelligence: Paving the last mile of artificial intelligence with edge computing. *Proceedings of the IEEE*, 107(8), 1738-1762. <https://doi.org/10.1109/JPROC.2019.2918951>