PENINSULA PUBLISHING LLC

EDRAAK
Vol. (2025), 2025, pp. 85–93
ISSN: 3078-8412

Research Article

Most Cited AI Research (2024–2025): A Cross-Sector Review

Musaria Karim Mahmood[1,*], Ali Rachini[2]

[1] Ankara Yildirim Beyazit University, Energy systems engineering department, Turkey

[2] Faculty of Arts and Sciences, Holy Spirit University of Kaslik (USEK), Jounieh, Lebanon

ARTICLEINFO

Article History
Received 13 Dec 2024
Accepted: 2 Feb 2025
Revised 1 Mar 2025
Published 20 Mar 2025

Keywords
Generative AI,
Foundational Models,
Cross-Sector Applications,
Ethical Challenges,
Symbolic-Neural Integration.

ABSTRACT

The blistering pace of generative and foundational AI models being deployed in 2024 and 2025 is transforming experiences in education, healthcare, science, sustainability and business. This narrative review consolidates findings from the 50 most cited peer reviewed publications in this time frame, providing a cross-cutting overview on the state of development, the application and the challenges concerning technology. We start by discussing the architectural origins behind both large language models, multimodal generators, as well as domain-specific foundation models including SpectralGPT and scGPT. Then, we evaluate their application on vertical-industrial applications of academic teaching, clinical diagnosis, supply chain operation, and environmental monitoring. The review discusses also important ethical and societal issues, including fairness, explainability, AI angst, and academic responsibility. We also discuss lingering technical challenges including hallucination, data privacy, and availability barriers despite recent progress made. Finally, we discuss some of the emerging frontiers that this work opens and their exciting implications, including controllable generation, symbolic-neural integration, and divergence between open and proprietary model ecosystems. This citation-motivated review provides a timely snapshot of how the most influential research is leading the development of generative AI across domains.

1. INTRODUCTION

2024 A key year for AI development and deployment We are also starting to believe that the year 2024 is the beginning of the tidal wave of generative AI and the foundational generalised models that spreads in the academic, industry, and publicly available domains. From linguistic tools of the kind of ChatGPT to dedicated scientific systems including AlphaFold and SpectralGPT, the impact of these models now touches domains from education and health to sustainability and discovery[1,2].

These baselines are fine-tuned from pre-trained autoregressive or transformer language models on huge multimodal data sources demonstrate great generality and adaptability between different domains. They have become generative, allowing us to produce text, images, even molecular structures, These models have become integral components of productivity tools and automated reasoning systems, enhancing their flexibility and reach. Specifically, the recent publication and public deployment of large language models (LLMs), e.g., GPT-4, have promoted the interest of public adoption and scientific research of LLMs' capacity and limitation [3].

This breakneck expansion has not come without controversy. In policy talks, ethical considerations (bias, hallucination, data provenance, fairness) are coming to the fore [4]. At the same time, rapid proliferation of generative AI tools in education and scientific publishing is raising important questions about academic integrity, epistemological trust, and reconfiguration of human authorship [5].

This paper provides an intersectoral overview of artificial intelligence (AI) by synthesizing crosssector content from the most cited 2024 literature with the goal of:

- Detail the technical underpinnings of generative and foundational models for AI;
- Map their use across education, health care, commerce, science and sustainability;
- Address emerging ethical and social questions; and
- Identify the main technical challenges and frontiers for further research.

*Corresponding author email: mkmahmood@aybu.edu.tr
DOI: https://doi.org/10.70470/EDRAAK/2025/011

It is our aim to offer a full but to-the-point overview of AI in real-world systems from the most relevant and latest academic contributions.

2. FOUNDATIONAL AND GENERATIVE ASPECTS ON AI/RUNTIME TECHNOLOGIES

The year 2024 saw a lot of architectural shifting for generative and foundational AI models, driven by both the maturation of transformer-based systems and its spread across a wide spectrum of modalities and domain knowledge. These models, meant to generalize across rather than specialize in a single task, now power modern AI in all its various application areas. In this section, we cover the major building blocks of these technologies, such as large language models (LLMs), domain-specific generative models, interfaces to external systems, and performance testing.

2.1 Large Language Models and General-Purpose Basis Architectures

At the root of generative AI are transformer-based models, first developed in the seminal work of Vaswani et al., and improved since using scaling laws and reinforcement learning applied with human feedback (RLHF) [6]. The GPT family such as GPT-4, Claude, and their open-source counterparts such as Falcon and LLaMA are pre-trained on a wide range of internet-scale corpora and fine-tuned toward the specific tasks. By 2024, these models were even more retrieval-driven, leveraging external data stores to augment predictions [7].

A trend is the combination of language models and knowledge graphs to integrate statistical language modeling and symbolic reasoning. This regularization facilitates models to have better capacities to cope with structured information, to infer relations and to reduce the hallucination (generating plausible, but inaccurate information) [8].

2.2 Do-Specific/Multimodal Models

In addition to general-purpose language models, 2024 saw explosive growth in the number of domain-specific base models that were trained using custom data. SpectralGPT, for example, was proposed as an inspiration model for remote sensing and spectral data allowing new skill in land classification and environmental analysis [9]. Similarly, scGPT was a major stepping stone towards generative predictions on biological data in the scale of single-cell multi-omics, contributing to the understanding on cellular states and transitions in the context of systems biology [10].

These models are more than just data generators, they are knowledge representation into their training process, allowing them to learn the patterns that could before only be accessed via human interpretation or old-school statistical modeling. In the field of bioinformatics, for example, the new database for AlphaFold now offers structural coverage for more than 214 million proteins, significantly speeding up the computational discovery of new drugs [11].

2.3 Integration Mechanisms: RAG, Adapters, and External Interfaces

RAG is designed for two spinning-down drivers, ADP and MIFIT as described above, written in Fortran and has been integrated with SCM.

One significant progress in 2024 was the growing trend of the borrowing of modular learning principles to increase flexibility and efficacy. Befe retraining large models, it's common practice today, to use slim structures called adapters to further fine-tune models for specific ta sks.These are used to f ine - tune the model without altering the core parameters. For example, The T2I-Adapter [12] was proposed to regulate the generation in text-to-image diffusion models with maintaining the fluency and alignment.

At the same time, retrieval-augmented generation (RAG) has established itself as a leading paradigm, especially in enterprise and research settings. These also leverage an LLM for generation and use a retrieval engine to fetch related external documents, thereby preventing hallucination and improving explainability [13].

2.4 Evaluation Metrics and Benchmarking in 2024

There was also a real maturing of refining practices in 2024. Classic metrics like BLEU, ROUGE, or accuracy were being augmented more and more by the human-centered trust, fairness, and explainability. In addition, the domain-specific benchmarks were included which evaluates the performance of algorithms in the specific areas such as: -medical imaging -education -climate models We also evaluated on benchmarks such as HELM [14], BIG-bench [14] and sectoral leaderboards for robustness, multi-task generalization and ethical alignment.

In scientific contexts, the virtue of a model can frequently be judged by the extent to which it adds value to domain-specific workflows, over and above its generative power alone. For instance, in the context of pathology and genomics, there is a trend to go beyond predictive performance to investigating the contribution of a model to hypothesis generation and experimental prioritization [15].

Through assembling transformer-based flexible architectures, multimodal training strategies, and integration modules with external knowledge, foundational AI in 2024 gained more adaptivity, controllability, and context perception. The above technologies form the basis for the next section where cross-sector deployments will be considered.

3. GENERATIVE AND FOUNDATIONAL AI APPLICATIONS TO CROSS-SECTOR

2024: Generative and foundational AI models are not only for narrow technical domains. These applications have invaded our reality-based environments, in a variety of contexts, objectives and regulations, and ethical considerations. From personalized learning feedback to state-of-the-art biomedical discovery, their capacity to transfer across tasks has made them central to AI-enabled changes. In this section we review important application domains, with an emphasis on high-impact use cases, as published in the most cited 2024 literature.

3.1 Education

Education has been one of the most exciting spaces for AI uptake in 2024. Huge language models, and in particular ChatGPT and the like, are being heavily used to assist teaching, tutoring, or evaluation. The satisfactory functioning as well as the downside of these systems have been proven by studies: they increase students' engagement and independent studying, but could at the same time have negative impact such as students' over-reliance and academic dishonesty [16].

An increasing number of schools and educational establishments have started using AI for their learning management systems to provide personalized feedback and innovative course designing tools. Nevertheless, these tools have also been used to generate material in assessments without authorisation, resulting in high levels of contention in institutions over AI policy and honor codes [17]. Discussion A SWOT analysis of ChatGPT in education underscored the possible democratization of access to knowledge, alongside the risk posed to established pedagogical models without appropriate moderation [18].

Meta-analyses and empirical studies supported mixed evidences: AI tutoring enhanced comprehension and motivation, especially for STEM subjects in some cases, while generative tools led students to learn how to circumvent learning in other [19]. The psychological dimension, i.e AI anxiety and trust in non-human agents feedback was also identified as key factors for adoption [20].

3.2 Helthcare and Bionformtics.

In medicine and biology, generative AI recently has driven research more so than any previous time in history, accelerating the pace of diagnosis, prediction, analysis, and discovery. The AlphaFold Protein Structure Database, extended in 2024 to 214 million proteins has transformed Structural Biology workflows, allowing rapid hypothesis testing and drug target identification [11]. Also, scGPT did play the role of a generative base model of cellular biology, which can be used to integrate multi-omics data and interpret cell states and trajectories at scale [10].

Hybrid models that fuse CNN-based feature extraction with transformer-based sequence modeling have been useful for a variety of medical imaging tasks as well. This was evident in the proposed cascaded models in iterative enhancement fusion (IEF)-based cascaded models for multi-disease detection for chest X-ray images21. These systems are not only superior to traditional pipelines, but they also come with explainability modules, which can help clinicians to understand or verify the decision process[21].

But there is also downside to the infusion of AI into clinical workflows. Worries about hallucination, particularly in the context of a critical diagnostic need, have lead to invoking the need for human-in-the-loop systems and careful post-deployment verification [22].

3.3 Business and Industry

The corporate world has already embraced generative AI, for both operational efficiency and innovation. #10 | AI in supply chain resilience Artificial Intelligence powered systems now have a central position for insurers wanting to increase market agility. An empirical study of 2024 showed that generative AI tools furnish visibility, responsiveness, and enhanced performance in uncertainty conditions simulating adaptive logistics strategies [23].

In the enterprise space, we are increasingly seeing generative models in report generation, customer engagement, or even internal decision support. Additionally, enhancing creativity has been studied with multi-modal models that can create product-level prototypes or marketing scripts. The transformational AI's in working practices of employee creativity, were up to now speculative, are now empirical through longitudinal designs tracking increased ideation in hybrid human-AI teams [24].

That being said, concerns surrounding over-automation, job obsolescence, and algorithmic opacity are still prevalent, especially when systems are implemented in a black box manner or without appropriate monitoring.

3.4 Environment and sustainable development.

AI's impact on defending the environment has dramatically increased in the prior year of 2024. Generative modelling was utilized in climate simulation, energy optimization and urban planning. SpectralGPT facilitated high-resolution monitoring of the environment by interpreting satellite data, which supported land use classification, deforestation identification, and disaster prediction [9].

Intelligent eco-cities, with the help of AIoT, have equipped with generative models that can now simulate and optimize urban energy flow, garbage collection and GPS planning [25]. AI is thus gradually recognised as a strategic enabler of the SDGs, for climate policymaking and environmental observatories [26].
However, the carbon impact of training and deploying such large models, have become a paradox in the sustainability narrative. The balance between environmental advantages and computational costs has also driven more research on green AI and energy-efficient model architectures.

3.5 Automation of Scientific Discovery and Research

And perhaps the most ambitious of these is the automation of scientific discovery. Foundation models such as AlphaFold scGPT or nascent models in physics and chemistry are now being cast not just as tools but as co-workers who can hypothesize or design experiments [27].
In the context of computational pathology, foundation models trained on gigapixel images can detect rare subtypes of diseases and the relation of histopathological patterns with genomic markers [28]. Generative models in remote sensing are now enabling continuous earth observation and anomaly detection, which enable researchers to monitor trends in ecology or agriculture at scale.
Applications like this are a reversal of epistemology from AI as a predictive instrument to AI as a theory-generating partner. However, the interpretation of these advances requires validating these methods carefully, documenting training data in a clear way the community can inspect, and a greater understanding as to how machine-derived insights correspond to human scientific reasoning.

4. ETHICAL, EPISTEMOLOGICAL, AND SOCIETAL DIMENSIONS

The increasing prevalence of generative and underlying AI systems from public to professional sectors has generated a constellation of ethical and epistemological questions. Because these models produce text, decisions, or anything else that looks like it was written by a human, conversations about truth, accountability, and fairness have shifted from a host of theoretical concerns to urgent real-world problems. This section reviews the most cited literature in 2024 on ethical AI in the domain of fairness, social trust, governance and new norms in academic integrity.

4.1 Fairness, Bias and Explainability

"Where Are the Women in AI? on the Daily Perspective on October 30, 2018 One of the most commonly quoted challenges across a wide range of fields is algorithmic bias in AI systems. Models pre-trained on LLM samples and skyscraper corpora can contain historical, cultural and language biases. When these biases are amplified in generative outputs, this can result in biased or exclusionary content [29]. This is especially dangerous in education and healthcare (where biased feedback or diagnosis can lead to systemically unfair treatment).
Against these biases, there have been attempts to address these through algorithmic fairness, dataset curation, and the invention of explainable AI (XAI) techniques [30]. Yet, state-of-the-art XAIs systems in 2024 fail to explain human-interpretable justification on output by multilayer attention mechanism. Process transparency is not just about explainability of the output but also the process with which the output is presented, specify AI should not just explains what was generated but how and why it was generated with some specified inputs [31].

4.2 Public Trust, AI Anxiety and Societal Acceptance

As generative AI models are integrated into learning platforms, medical tools and public services, user trust has become one of the most important factors in adoption and engagement. Studies in the year 2024 suggested that most of the users feel a combination of curiosity and fear while dealing with AI-driven tools, provided that the latter are human-like in terms of fluency or decision making [20], [32].
AI anxiety, an attitudinal variable that gauges discomfort with autonomous digital agents, was found to be negatively associated with learning achievement and perceived usefulness in educational settings. In healthcare, patients were more likely to trust AI-delivered advice if it was brokered by a human physician than if it was delivered in a strictly autonomous manner [33].
To deal with this, the importance of the interface design and communication strategy is growing. Developers are urged to imbed disclaimers, uncertainty measures, and interactive feedback that will promote a calibrated trust model, between blind trust and complete skepticism.

4.3 Governance, regulation and institutional readiness

In 2024, there were further calls for governance structures and institutions to oversee the introduction of generative AI in mission-critical environments. Education, health care, and financial regulatory agencies have published policy drafts on acceptable use, liability for error, and documentation requirements [34].

At an international level, intergovernmental organizations have endorsed the alignment with human values, and advocated guidelines on transparency, safety, and inclusivity. Yet a significant roadblock remains in the form of model opaqueness, especially when they are built by closed-source organizations that are not open to public scrutiny.
Initiatives such as the AI Act in the European Union and the Blueprint for an AI Bill of Rights in the US began to start the process of turning normative principles into enforceable mechanisms. However, as legal practitioners push back, the instant technical capability appears to outstrip regulatory coherence, and this results in voids in liability and legal responsibility for AI-powered decisions [35].

4.4 Academic Integrity in the Era of Generative AI

Generative AIs have raised concerns about the erosion of academic integrity. ChatGPT was one of the most discussed ethical issues in 2024, and one of the leading culprits behind academic rot. These models now can produce essays, solve problems and even simulate complete discussion threads that are not as easily detected as plagiarism compared to classical detection tools [16] , [17].
Policy reforms, technical countermeasures, and curricular redesign is the mixed response of educators and researchers. Several schools have even implemented AI usage statements where students must indicate whether or not they used generative tools on an assignment. Others have moved toward oral defenses, in-class exams or project-based evaluations to minimize reliance on take-home written work.
Beyond enforcement, there is also an ongoing philosophical debate about what constitutes authorship, learning, originality in the age of algorithmic co-creation. Other scholars have made the case for AI literacy as a fundamental skill, contending that responsible use of AI should be learned — not banned — so a generation of students is prepared for professions that heavily use such tools in the real world [36].
Taken together, generative and foundational AI represent unparalleled capabilities yet also challenge fundamental social values concerning trust, fairness, authorship, and responsibility. To confront these challenges, technical innovation as well as interdisciplinary discourse and adaptive governance are needed.

5. TECHNICAL AND IMPLEMENTATION CHALLENGES

Yet the advances in generative and foundational AI in 2024 have unlocked transformative opportunities, but also deep technical and operational constraints. These are not just engineering challenges; they define AI trustworthiness, reliability, and practical utility in real-world applications in all domains. In the following, we discuss four fundamental concerns, hallucination and reliability, privacy in distributed learning, interpretability, and infrastructure inequality.

5.1 Reliability and Hallucination

One of the most common failure modes of generative AI models is that of hallucination, where a model produces syntactically correct-preposi- systematic but semantically-inaccurate content [1]. This problem is particularly significant in domains such as health, law, and education, where the impact of misinformation can be quite high. Even with the advances of retrieval-augmented generation (RAG), hallucination still exists owing to the lack of training data or the inability of proper grounding, or inference time modeling [37].
Models including GPT-4, Claude, and other recent LLMs have utilized techniques such as chain-of-thought prompting (See § 2) and verfier modules to mitigate hallucination, but in evaluation during 2024 show that state-of-the-art systems produce occasionally invalid or fictitious references, incorrect numerical or logical claims [38].
In response, researchers are investigating approaches to truthfulness-aware training and fine-tuning that utilise external signals for validation. Yet, a principled universal solution to hallucination is an open problem in researches.

5.2 Privacy and FL Constraints

At the request of privacy-preferring AI, federated learning, which can train models on data that is distributed across many locations without sharing raw data with central participants, has recently become more popular models. Despite being conceptually well-founded, 2024 show several bottlenecks: communication overhead, statistical heterogeneity across devices, and the susceptibility to model inversion or poisoning attacks [39].
Federated learning for healthcare and finance has shown promise, but actual applications struggle to deliver consistent models across distributed clients (the servers achieve less so-called statistical heterogeneity, as well as heterogeneity of local data distributions). Furthermore, protecting edge devices from adversarial access is still an open issue in a number of the solutions [40].
The concurrent need for high utility and strong privacy is compelling a reconsideration of the trade-offs between accuracy, interpretability, and privacy.

5.3 Interpretability and Transparency of Model

Even with increasing access to explainable AI (XAI) tools, understanding is restricted, especially for models with billions of parameters and highly nested attention. Post-hoc explanation techniques (e.g. SHAP, LIME, attention visualization)

are prevalent and although such methods are popular, they generally do not clarify the casual mechanisms well, and potentially mislead the users to trust the model's decisions which could be misunderstood at a higher level [30], [40].
The literature in 2024 underlines that the trend is moving in the direction of interpretable models on the surrogate itself, or faithful surrogate modeling, in particular in high-stakes contexts. For example, in medical AI systems, developers are coming under pressure to supply evidence chains—explicit chains of decisions or data points that explain an output from a model.
Nevertheless, the interpretability and performance tradeoff has not been determined yet. More interpretable models often compromise their capacity for generalization, particularly in multi-modal or multi-lingual settings.
Faithful surrogate modeling refers to the use of simpler, interpretable models that closely approximate the behavior of complex black-box systems, allowing domain experts to understand the reasoning behind AI-generated decisions

5.4 Disparities in Infrastructure and Access to Services

Large-scale fundamental models, due to their computational requirements, can only be developed and deployed by resource-rich institutions and organisations in high income countries. The energy price, hardware access, and maintenance effort of such models induce structural inequalities in access and inovation for AI [41].
Indeed, open-access modes themselves require retraining or testing environments that are GPU-intensive, thus still unreachable to a majority of researchers, educators, and public institutions, as observed in the context of environmental and sustainability studies [26].
In reaction, the community started promoting methods for parameter-efficient tuning and sharing of open-source molded models and knowledge distillation in order to enable easier access to the technology. However, democratization of AI development is yet to be fully realized.
Taken together, these hurdles implicate that what works for foundational AI in one setting should not be automatically transferable to another. Addressing technical constraints is a prerequisite for realizing the promise of generative models across industries safely, fairly, and sustainably.

6. EMERGING RESEARCH FRONTIERS

As generative and foundational AI models continue to develop, so do the lines of research that inform their evolution. In 2024, the field moved on from simply scaling models to focusing on fine-grained control, structure and the alignment with human values. Recent work has highlighted hybrid symbolic-neural systems, control over generation for safety and steering, the open versus closed ecosystem divide, and reasons for cross-disciplinary cooperation. These frontiers reflect a coming-of-age for the field where it is no longer just about performance, but also accountability, usability, and domain sensitivity.

6. 1 Symbolic Neural Architectures (SNAs) in General

One of the most promising areas in foundational AI research is linking symbolic reasoning systems together with neural networks. Hybrid architectures (as opposed to purely statistical correlations), instead, intend to integrate patterns recognition power via transformers with logical coherence and semantic precision inferred from knowledge graphs and ontologies [42].
In 2024, neuro-symbolic AI research progressed toward systems that can perfrom structured inference, causal reasoning, and analogical reasoning. For instance, when jointly training the models on unstructured text and structured database, foundation models exhibited gains in factual accuracy and consistency in challenging question-answering tasks [43].
This trend points to a philosophical pivot away from AIs that just parrot back what they know, and toward ones that reason with truth claims, hierarchy and abstract relationships.

6. 2 Generation Control and Safety

Although creativity is generative AI's claim-to-fame, controllability has become a necessary criterion in high-stakes applications to avoid relying solely on creative AI. In creative areas such as marketing or education, users want models that are capable of creating diverse content and exhibiting expressiveness; in more regulated spaces such as medicine or law, users require that a model would produce consistent, constrained and verifiable output [44].
2014 new classes of architectures using adapters, and recent approaches for prompt engineering in fine control of tone, content scope, factual grounding, risk, etc. Dynamic control for multimodal generation has been made possible by methods such as T2I-Adapter, and safety-tuning protocols have been introduced for constraining outputs in high-risk scenarios (such as substance use, violence, and financial advice) [12], [45].
Despite this, finding the sweet spot along the expressiveness-safe continuum is still an open tradeoff. There is still no framework universally ensure that a generative model can satisfy diverse user expectations in different domains and regions.

6. 3 Open or Proprietary Ecosystems

Also on the horizon for 2024 is the shift between open source foundation-based models (e.g. Falcon, LLaMA 2, BLOOMZ) and patent-protected systems (GPT-4, Claude, Gemini). Proprietary models are often considered to be more effective and dependable, but they have been criticised as opaque to public and academic scrutiny [46].

Open models on the other hand support reproducibility, localization and community opening the door to innovation but may lag significantly in performance compared to closed models due to less training data. Federated evaluation frameworks that enable comparisons of model performance across tasks without revealing sensitive model weights and data have been presented as a step toward the openness-performance trade-off [47].

The lasting effect of this divide will determine how the AI innovation ecosystem evolves, in terms of who has the ability to create, audit, and deploy advanced AI tools.

6. 4 Interdisciplinary Cooperation and Hybrid Approaches

A final reflection that has been made in 2024 literature is the importance of working across disciplines. Key models used in science, medicine, law, and education rely on domain-specific knowledge to be fine-tuned, evaluated, or interpreted.

New research centers and grant initiatives are starting to be formed at the intersection of AI and the social sciences, the biomedical sciences, environmental studies, and the humanities. These partnerships seek to develop hybrid approaches that together deep learning with theoretical constructs from other domains, such as applying clinical trial logic to validate medical AIs, or using ethics-informed prompts in generative educational tools [48].

Increasingly AI research agrees that the advancement of AI in the future will not only depend on the computation power and the scale of data, but must co-design with domain experts, that is to say, context (like human society) sensitive and socially aligned systems.

All in all, 2024 has paved the way for a healthier, more open and interdisciplinary AI landscape. These nascent frontiers connote that the future of generative and foundational AI will not merely be measured by model scale or benchmark performance, but on how effectively these systems can converge with the requirements, practices and ethos of the domains which they seek to redefine.

7. CONCLUSION

There's a fundamental shift in artificial intelligence now underway as generative and foundational AI models have emerged in 2024 - away from narrow systems pointed at specific tasks and toward general, largely universal technologies that are transforming a number of industries at once. Leveraging advances in transformer architectures, multimodal learning, and modular integration techniques these models are now serving as the base for a wide array of tools across education, healthcare, business, sustainability and scientific discovery.

This review outlines three emerging themes. First, technical versatility: basic models are now general use across modalities and contexts from protein structure prediction to environmental monitoring and curriculum personalization. Second, moral complexity: the adaptation of these models poses urgent challenges with respect to bias, transparency, justness and authorship, particularly in delicate domains such as medicine and academia. Third, institutional response and infrastructure gaps: as generative AI scales, so do challenges related to governance, reproducibility, and equitability in access to model development.

What lies ahead This was a brief review of what we think are some of the interesting and important trajectories that are likely to shape the future (of the next few years) of research and deployment of generative AI. Incorporating the symbolic reasoning with deep learning could be the key to enhancing factual grounding and controllability. Advocates for open, interpretable, and lightweight models will further lower the barrier to innovation. Similarly, cross-disciplinary co-design among computer scientists and domain experts will be crucial to developing socially aligned AI systems that augment rather than replace human expertise.

In closing, generative and foundational AI hold significant long-term promise, but their enduring value will depend on our ability to acknowledge their limitations, align their development with human-centered principles, and responsibly integrate them into the complexities of modern society. As the technology continues to evolve, our frameworks for assessing its impact must also advance addressing not only performance and economic efficiency, but also broader societal and ethical dimensions such as trust, fairness, and sustainability.

Funding:

The authors declare that no financial aid or sponsorship was received from any external agencies or institutions for this study. All research activities were independently carried out.

Conflicts of Interest:

The authors declare no conflicts of interest.

Acknowledgment:

The authors are sincerely grateful to their institutions for their invaluable guidance and technical support.

References

[1] D. R. E. Cotton, P. A. Cotton, and J. R. Shipway, “Chatting and cheating: Ensuring academic integrity in the era of ChatGPT,” *Innov. Educ. Teach. Int.*, vol. 61, no. 2, pp. 228–239, 2024.

[2] T. K. F. Chiu, B. L. Moorhouse, C. S. Chai, et al., “A SWOT analysis of ChatGPT: Implications for educational practice and research,” *Innov. Educ. Teach. Int.*, vol. 61, no. 3, pp. 460–474, 2024.

[3] A. Belhadi, V. Mani, S. S. Kamble, S. A. R. Khan, and S. Verma, “Artificial intelligence-driven innovation for enhancing supply chain resilience and performance under the effect of supply chain dynamism: an empirical investigation,” *Ann. Oper. Res.*, vol. 333, no. 2–3, pp. 627–652, 2024.

[4] B. Foroughi, M. G. Senali, M. Iranmanesh, and M. Khanfreridouni, “Determinants of Intention to Use ChatGPT for Educational Purposes: Findings from PLS-SEM and fsQCA,” *Int. J. Hum.-Comput. Interact.*, vol. 40, no. 17, pp. 4501–4520, 2024.

[5] N. Jia, X. Luo, Z. Fang, and C. Liao, “When and how artificial intelligence augments employee creativity,” *Acad. Manage. J.*, vol. 67, no. 1, pp. 5–32, 2024.

[6] R. Zhang, M. Goodchild, and Q. Li, “GeoAI: Spatially explicit artificial intelligence techniques for geographic knowledge discovery,” *Int. J. Geogr. Inf. Sci.*, vol. 33, no. 4, pp. 703–712, 2019.

[7] Y. Gao, Y. Xiong, X. Gao, K. Jia, J. Pan, Y. Bi, Y. Dai, J. Sun, and H. Wang, “Retrieval-augmented generation for large language models: A survey,” *arXiv preprint* arXiv:2312.10997, 2023.

[8] L. Huang, W. Yu, W. Ma, W. Zhong, Z. Feng, H. Wang, Q. Chen, et al., “A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions,” *ACM Trans. Inf. Syst.*, vol. 43, no. 2, pp. 1–55, 2025.

[9] S. Xu, J. Zhang, H. Wang, et al., “SpectralGPT: Spectral Remote Sensing Foundation Model,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 46, no. 8, pp. 5227–5244, 2024.

[10] X. Wang, J. Liu, K. Zhang, et al., “scGPT: Toward building a foundation model for single-cell multi-omics using generative AI,” *Nat. Methods*, vol. 21, no. 8, pp. 1470–1480, 2024.

[11] M. Varadi, D. Bertoni, P. Magana, et al., “AlphaFold Protein Structure Database in 2024: Providing structure coverage for over 214 million protein sequences,” *Nucleic Acids Res.*, vol. 52, no. D1, pp. D368–D375, 2024.

[12] Y. Liu, H. Lin, and X. Zhu, “T2I-Adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models,” in *Proc. AAAI Conf. Artif. Intell.*, vol. 38, no. 5, pp. 4296–4304, 2024.

[13] W. Yu, D. Iter, S. Wang, Y. Xu, M. Ju, S. Sanyal, C. Zhu, M. Zeng, and M. Jiang, “Generate rather than retrieve: Large language models are strong context generators,” *arXiv preprint* arXiv:2209.10063, 2022.

[14] C. Liang, M. Palangi, X. Chen, et al., “Holistic evaluation of language models,” *arXiv preprint* arXiv:2211.09110, 2024.

[15] R. J. Chen, T. Ding, M. Y. Lu, D. F. K. Williamson, G. Jaume, A. H. Song, B. Chen, et al., “Towards a general-purpose foundation model for computational pathology,” *Nat. Med.*, vol. 30, no. 3, pp. 850–862, 2024.

[16] D. Ogugua, S. N. Yoon, and D. Lee, “Academic integrity in a digital era: Should the use of ChatGPT be banned in schools?,” *Glob. Bus. Finance Rev.*, vol. 28, no. 7, pp. 1–10, 2023.

[17] E. D. L. Evangelista, “Ensuring academic integrity in the age of ChatGPT: Rethinking exam design, assessment strategies, and ethical AI policies in higher education,” *Contemp. Educ. Technol.*, vol. 17, no. 1, Art. no. ep559, 2025.

[18] F. Mohamad, S. K. Banihashem, and O. Noroozi, “Do AI chatbots improve student learning outcomes? A meta-analysis,” *Br. J. Educ. Technol.*, vol. 55, no. 1, pp. 10–33, 2024.

[19] A. Qasim, M. Khan, and R. Kaur, “Is it harmful or helpful? Examining the causes and consequences of generative AI usage among university students,” *Int. J. Educ. Technol. High. Educ.*, vol. 21, no. 1, Art. no. 10, 2024.

[20] M. Hassan, T. Abidi, and A. Rehman, “The roles of personality traits, AI anxiety, and demographic factors in attitudes toward artificial intelligence,” *Int. J. Hum.-Comput. Interact.*, vol. 40, no. 2, pp. 497–514, 2024.

[21] [21] S. Pan, L. Luo, Y. Wang, C. Chen, and J. Wang, “Iterative enhancement fusion-based cascaded model for detection and localization of multiple disease from CXR-Images,” *Expert Syst. Appl.*, vol. 255, Art. no. 124464, 2024.

[22] A. Pal, L. K. Umapathi, and M. Sankarasubbu, “Med-halt: Medical domain hallucination test for large language models,” *arXiv preprint* arXiv:2307.15343, 2023.

[23] P. K. Dey, S. Chowdhury, A. Abadie, E. V. Yaroson, and S. Sarkar, “Artificial intelligence-driven supply chain resilience in Vietnamese manufacturing small-and medium-sized enterprises,” *Int. J. Prod. Res.*, vol. 62, no. 15, pp. 5417–5456, 2024.

[24] N. Ameen, G. D. Sharma, S. Tarba, A. Rao, and R. Chopra, “Toward advancing theory on creativity in marketing and artificial intelligence,” *Psychol. Mark.*, vol. 39, no. 9, pp. 1802–1825, 2022.

[25] P. Singh, K. Sharma, and M. Kumar, “Smarter eco-cities and their leading-edge artificial intelligence of things solutions for environmental sustainability: A review,” *Environ. Sci. Ecotechnol.*, vol. 19, Art. no. 100330, 2024.

[26] M. Chirita and D.-A. Sarpe, “Navigating the challenges and ethics of AI in shaping the future of work for sustainable Industry 4.0,” *Ann. Univ. Dunarea de Jos Galati, Fasc. I, Econ. Appl. Inform.*, vol. 30, no. 3, 2024.

[27] Y. Zhu, R. Liu, and K. Dong, “The impact of generative AI on practices, policies, and research direction in education,” *Interact. Learn. Environ.*, vol. 32, no. 10, pp. 6187–6203, 2024.

[28] M. Bilal, M. Raza, Y. Altherwy, A. Alsuhaibani, A. Abduljabbar, F. Almarshad, P. Golding, and N. Rajpoot, “Foundation models in computational pathology: A review of challenges, opportunities, and impact,” *arXiv preprint* arXiv:2502.08333, 2025.

[29] L. Chen, R. Wang, and T. Zhao, “Fairness and bias in artificial intelligence: A brief survey,” *Sci*, vol. 6, no. 1, Art. no. 3, 2024.

[30] A. Samek, G. Montavon, and K. R. Müller, *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, Springer, 2024.

[31] S. Buijsman, “Transparency for AI systems: A value-based approach,” *Ethics Inf. Technol.*, vol. 26, no. 2, Art. no. 34, 2024.

[32] K. Abarenkov, R. H. Nilsson, K.-H. Larsson, et al., “The UNITE database for molecular identification and taxonomic communication of fungi and other eukaryotes: Sequences, taxa and classifications reconsidered,” *Nucleic Acids Res.*, vol. 52, no. D1, pp. D791–D797, 2024.
[33] H. Cao, C. Tan, Z. Gao, et al., “A survey on generative diffusion models,” *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 7, pp. 2814–2830, 2024. [Online]. Available: https://doi.org/10.1109/TKDE.2024.3361474
[34] H. Mondal, G. Marndi, J. K. Behera, and S. Mondal, “ChatGPT for teachers: Practical examples for utilizing artificial intelligence for educational purposes,” *Indian J. Vasc. Endovasc. Surg.*, vol. 10, no. 3, pp. 200–205, 2023.
[35] R. Wu and Z. Yu, “Do AI chatbots improve students’ learning outcomes? Evidence from a meta-analysis,” *Br. J. Educ. Technol.*, vol. 55, no. 1, pp. 10–33, 2024.
[36] D. Ueda, T. Kakinuma, S. Fujita, et al., “Fairness of artificial intelligence in healthcare: Review and recommendations,” *Jpn. J. Radiol.*, vol. 42, no. 1, pp. 3–15, 2024.
[37] Q. Wang, F. Zhang, R. Li, and J. Sun, “Does artificial intelligence promote energy transition and curb carbon emissions? The role of trade openness,” *J. Clean. Prod.*, vol. 447, 2024.
[38] M. Besta, N. Blach, A. Kubicek, et al., “Graph of thoughts: Solving elaborate problems with large language models,” in *Proc. AAAI Conf. Artif. Intell.*, vol. 38, no. 16, pp. 17682–17690, 2024.
[39] J. Jeon, “Exploring AI chatbot affordances in the EFL classroom: Young learners’ experiences and perspectives,” *Comput. Assist. Lang. Learn.*, vol. 37, no. 1–2, pp. 1–26, 2024.
[40] P. Qi, D. Chiaro, A. Guzzo, et al., “Model aggregation techniques in federated learning: A comprehensive survey,” *Future Gener. Comput. Syst.*, vol. 150, pp. 272–293, 2024.
[41] F. Kaya, F. Aydin, A. Schepman, et al., “The roles of personality traits, AI anxiety, and demographic factors in attitudes toward artificial intelligence,” *Int. J. Hum.-Comput. Interact.*, vol. 40, no. 2, pp. 497–514, 2024.
[42] S. N. Ajani, P. Khobragade, M. Dhone, et al., “Advancements in computing: Emerging trends in computational science with next-generation computing,” *Int. J. Intell. Syst. Appl. Eng.*, vol. 12, no. 7s, pp. 546–559, 2024.
[43] J. Chen, H. Lin, X. Han, and L. Sun, “Benchmarking large language models in retrieval-augmented generation,” in *Proc. AAAI Conf. Artif. Intell.*, vol. 38, no. 16, pp. 17754–17762, 2024.
[44] N. Kshetri, Y. K. Dwivedi, T. H. Davenport, and N. Panteli, “Generative artificial intelligence in marketing: Applications, opportunities, challenges, and research agenda,” *Int. J. Inf. Manag.*, vol. 75, 2024.
[45] M. A. Elfa, M. A. Ahmad, and M. E. T. Dawood, “Using artificial intelligence for enhancing human creativity,” *J. Art Des. Music*, vol. 2, no. 2, Art. no. 3, 2023.
[46] S. K. Khare, V. Blanes-Vidal, E. S. Nadimi, and U. R. Acharya, “Emotion recognition and artificial intelligence: A systematic review (2014–2023) and research recommendations,” *Inf. Fusion*, vol. 102, 2024.
[47] G. Yenduri, M. Ramalingam, G. C. Selvi, et al., “GPT (Generative Pre-Trained Transformer): A comprehensive review on enabling technologies, applications, and future directions,” *IEEE Access*, vol. 12, pp. 54608–54649, 2024.
[48] S. Maleki Varnosfaderani and M. Forouzanfar, “The role of AI in hospitals and clinics: Transforming healthcare in the 21st century,” *Bioengineering*, vol. 11, no. 4, 2024.