

Review

Not peer-reviewed version

---

# Generative AI and Large Language Models: A Comprehensive Scientific Review

---

[Bogdan-Iulian Ciubotaru](#) \*

Posted Date: 10 April 2025

doi: 10.20944/preprints202504.0413.v2

Keywords: generative AI; large language models; LLM



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Review

# Generative AI and Large Language Models: A Comprehensive Scientific Review

Bogdan-Iulian Ciubotaru

Military Equipment and Technologies Research Agency (METRA), Ministry of National Defence, Clinceni, Romania; bciubotaru@acttm.ro

**Abstract:** This systematic review examines the evolution, technical architecture, applications, limitations, and future directions of generative artificial intelligence (AI) and large language models (LLMs). Through comprehensive analysis of scientific literature, it was traced the development of these technologies from early linguistic theories to modern transformer-based architectures. The findings presented in this review article reveal the transformative impact of LLMs across diverse domains including healthcare, education, software development, and creative industries. Significant technical limitations were identified, including hallucinations, context window constraints, and reasoning deficiencies, along-side ethical concerns regarding bias, privacy, and environmental impact. The review concludes by exploring emerging trends in model architecture, efficiency improvements, and ethical frameworks that will shape future development. This work provides researchers, practitioners, and policymakers with a comprehensive understanding of the current state and future trajectory of generative AI and LLMs.

**Keywords:** generative AI; large language models; LLM

---

## I. Introduction

Generative Artificial Intelligence (GenAI): is the hype real or it's another technology in its early stages, with a lot of promises? In recent years, the field of artificial intelligence has witnessed a remarkable transformation with the emergence of generative AI and Large Language Models (LLMs). These technologies have revolutionized how machines understand, process, and generate human language, marking a significant milestone in the evolution of AI capabilities [1]. Generative AI, particularly in the form of LLMs, has captured widespread attention not only within academic and research communities but also across industries, governments, and the general public [2]. The launch of GPT models by OpenAI in November 2022 can be considered as a turning point for artificial intelligence, which later can be referred as 0-day for GenAI.

Generative AI refers to artificial intelligence systems capable of creating new content, including text, images, audio, code, and other media formats, based on patterns learned from existing data [3]. Its ability depends on data which was used for training and the prompts provided by a user. The focal point of this technological revolution are Large Language Models—sophisticated neural network architectures trained on datasets of text data that can generate coherent, contextually relevant, and increasingly human-like responses to prompts [4]. These models have demonstrated unprecedented capabilities in understanding context, generating creative content, answering complex questions, and even proving some incipient reasoning abilities that were previously thought to be exclusively human domains [5].

LLMs are more than simple text generators. These models are transforming numerous fields, from healthcare and education to software development and creative industries [6]. In healthcare, LLMs are being utilized for clinical documentation, medical research synthesis, and patient communication [7]. Software developers leverage their coding capabilities using LLMs these models for code generation and debugging assistance [8], while creative professionals like artists and social

media influencers use them for content creation, and design [9]. Even in research LLMs can be used to analyze different data and provide insights which were not previously discovered. The versatility and adaptability of LLMs have positioned them as one of the most significant technological advancements of the 21st century [10].

Modern LLMs are basically the result of decades of work in areas like natural language processing, machine learning, and neural networks — all that research finally coming together [11]. The big game-changer came in 2017 when the transformer architecture was introduced — it completely changed the way AI handles texts [12]. This new setup, thanks to something called self-attention, allowed models to understand the bigger picture in language — like how words relate to each other even if they're far apart in a sentence — way better than anything before. [13]. Subsequent innovations in training methodologies, computational resources, and data availability have led to the rapid evolution of increasingly powerful models, from GPT (Generative Pre-trained Transformer) to BERT (Bidirectional Encoder Representations from Transformers), LaMDA, PaLM, and beyond [14]. While until recent days, LLMs required huge Graphical Processing Unit (GPU) resources, recent technological advances from the Chinese company DeepSeek now enable end users to deploy powerful, efficient language models using significantly fewer computational resources and even CPU usage. These advancements improve accessibility, allowing smaller organizations, researchers, and even individuals to leverage sophisticated LLM capabilities on more affordable hardware setups, significantly reducing operational costs and broadening the practical applicability of generative AI [15].

Even though LLMs are powerful, there are some important problems and limitations which development must overcome. These include tendencies to generate plausible-sounding but factually incorrect information (hallucinations) [16], perpetuate biases which can be detected in training data [17], and consume substantial high amounts of computational resources during training [18]. People have also started worrying about things like privacy, ownership of content, and how these models might be used in the wrong way — all of which have become really important issues.[19]. The main idea would be that if AI is doing something bad, who should be accounted for it. As these technologies continue to evolve and integrate into unimagined dimensions of modern society, addressing these challenges becomes increasingly important for responsible development and deployment [16].

This scientific review aims to provide a comprehensive examination of generative AI and Large Language Models, exploring their historical evolution, technical architecture, capabilities, applications, limitations, and future directions. By synthesizing insights from academic research, industry developments, and practical implementations, this review offers a holistic understanding of these transformative technologies. This article pulls together insights from trusted sources to give a clear and balanced look at where generative AI and LLMs stand today — and where they might be headed. While the article focuses on current status of LLMs, it is more than clear that in the future, the various applications of LLMs will interact with every aspect our lives, improving and providing new skills which will foster productivity in both personal and professional lives.

## II. Methods

### *A. Research Design*

This study employed a systematic review methodology to examine the current state of generative AI and large language models. The systematic approach was chosen to ensure thoroughness, minimize bias, and provide a structured framework for analyzing the diverse and rapidly evolving literature in this field [20].


### *B. Search Strategy*

A comprehensive search of scientific literature was conducted by using multiple databases including arXiv, IEEE Xplore, ACM Digital Library, Google Scholar, Web of Science, and MDPI. The

search was performed between January and March 2025, focusing on literature published between 2017 (marking the introduction of the transformer architecture) and March 2025.

The following search terms were used in various combinations, highlighted in Figure 1, Word cloud:

- “generative AI” OR “generative artificial intelligence”
- “large language model\*” OR “LLM\*”
- “transformer model\*” OR “attention mechanism\*”
- “GPT” OR “BERT” OR “LaMDA” OR “PaLM”
- “neural language model\*”, “self-attention”
- “LLM model\*”, “LLM training”
- “chain-of-thought prompting\*”, “prompt engineering”
- “tokenization”, “text-to-text transfer”



**Figure 1.** Word cloud for search strategy.

### C. Inclusion and Exclusion Criteria

#### Inclusion Criteria:

- Peer-reviewed journal articles, conference proceedings, and technical reports
- Literature focusing on generative AI and large language models
- Studies examining technical architecture, applications, limitations, or future directions
- Publications in English
- Literature published between 2017 and March 2025
- Websites for institutions (European Commission) and blogs for main LLM providers (OpenAI, DeepSeek, Google)

#### Exclusion Criteria:

- Non-English publications
- Opinion pieces without substantial technical or empirical content
- Studies focusing exclusively on other AI technologies without significant discussion of generative AI or LLMs
- Duplicate publications or multiple reports of the same study

### D. Data Extraction and Synthesis

Data extraction was performed using a standardized form that captured the following information:

- Publication details (authors, year, journal/conference)
- Study objectives and methodology
- Key findings and contributions
- Technical details of models or architectures discussed
- Applications and use cases
- Limitations and challenges identified

- Future research directions proposed

The extracted data was synthesized using a thematic analysis approach, identifying key themes and patterns across the literature. These themes were organized into categories corresponding to the major sections of this review: historical evolution, technical architecture, applications, limitations and challenges, and future directions.

#### *E. Quality Assessment and Limitations of Review Methodology*

The quality of included studies was assessed using criteria adapted from the Critical Appraisal Skills Programme (CASP) and the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines [21]. For technical papers, methodological rigor, clarity of reporting, and significance of contribution were evaluated. For empirical studies, study design, sample size, data collection methods, and analysis techniques were assessed.

Several limitations of selected methodology should be acknowledged. First, the rapid pace of development in generative AI means that some recent advancements may not be reflected in peer-reviewed literature. Second, proprietary details of commercial LLMs are often not fully disclosed in scientific publications, potentially limiting the understanding of state-of-the-art systems. Third, the interdisciplinary nature of the field necessitated searches across multiple databases, which may have resulted in some relevant studies being overlooked.

### **III. Results**

#### *A. Technical Architecture of Generative LLMs*

The analysis of the technical approaches of modern LLMs revealed several key architectural components and principles used for these technologies to achieve their capabilities.

##### **1). Transformer Architecture**

The transformer architecture consists of an encoder and a decoder, though many modern generative LLMs utilize only the decoder portion in a “decoder-only” architecture [4]. The defining innovation is the self-attention mechanism, which allows the model to consider all tokens in a sequence simultaneously, weighing their relevance to each other [12].

Self-attention computes a weighted sum of all token representations in a sequence, where the weights (attention scores) are determined by the relevance of each token to the current token being processed [22]. Multi-head attention splits the attention computation into multiple “heads”, each focusing on different aspects of the input sequence [23]. Positional encodings address the limitation that transformers process all tokens in parallel by providing numerical representations that encode the position of each token in the sequence [24].

##### **2). Components of Modern LLMs**

Modern LLMs convert tokens (words or subwords) into numerical representations called embeddings using subword tokenization methods such as Byte-Pair Encoding (BPE) or SentencePiece [25]. Each transformer block contains a feed-forward neural network that processes the output of the attention mechanism, introducing non-linearity and increasing representational capacity [26]. Layer normalization stabilizes training and improves convergence by normalizing the activations across features for each example [27]. Residual connections (skip connections) around each sub-layer facilitate gradient flow during training, especially in deep models [28].

##### **3). Architectural Variations**

The GPT family of models uses a decoder-only transformer architecture with masked self-attention, making the model autoregressive [29]. BERT uses only the encoder portion of the transformer, allowing bidirectional attention [1]. T5 (Text-to-Text Transfer Transformer) uses the



complete encoder-decoder architecture, framing all NLP tasks as text-to-text problems [30]. Google’s PaLM and Gemini models introduce architectural innovations for improved scaling and multimodal capabilities [31].

4). Scaling Properties, Emergent Capabilities and Efficiency Innovations

Research has identified several important scaling laws: performance on language tasks follows a power-law relationship with model size [13]; performance also scales with the amount of training data, though with diminishing returns [32]; and training compute (the product of model size and training tokens) is another key factor [33]. LLMs exhibit emergent abilities—capabilities not present in smaller models but appearing once models reach a certain scale—including in-context learning, chain-of-thought reasoning, and instruction following [34].

To address the computational demands of large models, researchers have developed various efficiency techniques: sparse attention mechanisms limit each token to attending only to a subset of other tokens [23]; parameter-efficient fine-tuning methods like LoRA allow for efficient adaptation with minimal parameter updates [35]; quantization reduces the precision of model weights to decrease memory requirements [36]; and knowledge distillation transfers knowledge from a large “teacher” model to a smaller “student” model [37].

B. Applications and Use Cases

Diverse applications of generative AI and LLMs across multiple domains were discovered, as illustrated in Figure 2, demonstrating their versatility and transformative potential.

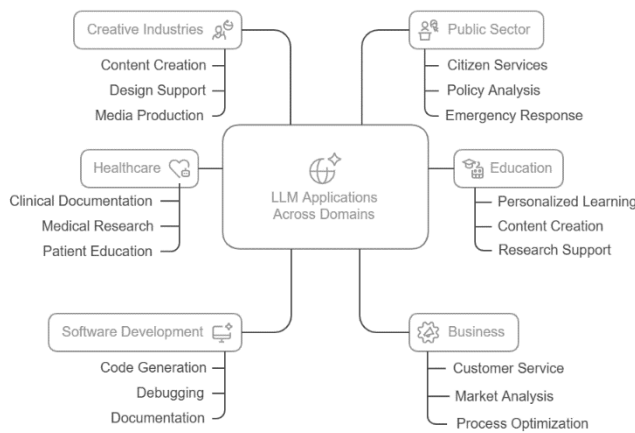


Figure 2. LLM applications in different domains.

LLMs excel at text summarization and content creation, distilling lengthy documents into concise summaries while preserving key information [38]. They have significantly advanced machine translation capabilities, supporting communication across language barriers with unprecedented fluency [39]. The conversational capabilities of LLMs have revolutionized virtual assistants and chatbots, enabling coherent, contextually appropriate dialogues [40].

LLMs offer unprecedented opportunities for personalized education through adaptive tutoring systems that explain concepts in multiple ways and provide tailored feedback [41]. Educators leverage LLMs to develop curriculum materials, lesson plans, and educational resources [42]. In academic contexts, LLMs assist with literature reviews by summarizing research papers and identifying connections between studies [43].

LLMs assist healthcare professionals with clinical documentation, generating notes, discharge summaries, and referral letters [7]. They synthesize findings across thousands of studies, identify emerging trends, and summarize evidence-based practices [44]. Artificial intelligence models can be used in combination with IoT devices to detect different pathologies or to detect early signs of future

disorders [45,46]. LLMs facilitate improved patient education by generating personalized health information that accounts for a patient's specific condition and health literacy level [23].

Programming-focused LLMs generate code snippets, complete partial code, and translate requirements into functional implementations [8]. They assist with debugging by identifying potential issues in code, suggesting fixes, and explaining the underlying problems [47]. LLMs can generate comprehensive documentation from code, including function descriptions, parameter explanations, and usage examples [48].

Enterprises deploy LLM-powered systems to handle customer inquiries, troubleshoot common issues, and provide product information [6]. LLMs analyze consumer feedback, social media conversations, and market trends to provide businesses with actionable insights [49]. In industrial settings, LLMs integrate with sensor data and operational metrics to optimize processes and predict maintenance needs [6].

LLMs support creative professionals in developing narratives, scripts, poetry, and other creative works [2]. Multimodal LLMs that combine text and image understanding support various design applications [50]. In music and audio production, specialized LLMs assist with composition, arrangement, and sound design [51].

Government agencies use LLMs to improve citizen services through more accessible information delivery and streamlined interactions [6]. LLMs assist policymakers by analyzing large volumes of data, simulating potential policy outcomes, and identifying unintended consequences [52]. During emergencies, LLMs help coordinate response efforts by processing incoming information and facilitating communication between different agencies [53].

### *C. Limitations and Challenges*

A range of critical limitations and challenges has been identified in generative AI and LLMs, all of which demand attention for responsible development and deployment. The following sections detail these issues, highlighting technical, ethical, regulatory, environmental, and integration aspects.

#### *1). Technical Limitations*

LLMs tend to generate content that appears plausible but contains factual errors or fabricated information—a phenomenon commonly referred to as “hallucination” [16]. Current LLMs operate within fixed context windows that constrain their ability to maintain coherence and consistency across long documents or conversations [23]. Despite their impressive language capabilities, LLMs struggle with complex reasoning, logical consistency, and mathematical accuracy [54]. State-of-the-art models require enormous computational resources for both training and inference, creating barriers to entry for smaller organizations and researchers [18].

#### *2). Ethical Concerns*

LLMs learn from vast corpora of human-generated text, inevitably absorbing and potentially amplifying biases present in that data [17]. The development and deployment of LLMs raise significant privacy concerns, particularly in domains like medicine where confidentiality is paramount [5]. LLMs trained on vast corpora of text raise complex intellectual property questions regarding copyright status of training data and ownership of AI-generated content [55]. The capabilities of advanced LLMs create potential for various forms of misuse, including generation of misinformation and automated production of harmful content [19].

#### *3). Regulatory and Compliance Challenges*

The rapid advancement of generative AI has outpaced regulatory frameworks, creating uncertainty and compliance challenges across different jurisdictions [56]. LLMs operate as “black boxes” with billions of parameters, making their decision-making processes opaque and difficult to

interpret [57]. The complexity and autonomy of LLMs complicate traditional accountability frameworks, with unclear responsibility distribution among developers, deployers, and users [52].

#### 4). Environmental Impact

The training of large language models requires enormous computational resources, resulting in significant energy consumption and carbon emissions [18]. Beyond training, the ongoing operation of LLMs for inference also consumes significant resources, particularly for high-traffic applications [58].

#### 5). Implementation and Integration Challenges

Adapting general-purpose LLMs to specific domains presents significant challenges, particularly for specialized fields with domain-specific terminology and knowledge [5]. Incorporating LLMs into existing technological ecosystems presents numerous integration challenges, including interfacing with legacy systems and ensuring consistent performance [6]. Evaluating LLM performance presents unique challenges compared to traditional software due to the subjective nature of many language tasks and the impossibility of comprehensive testing across all possible inputs [59].

#### *D. Future Directions and Emerging Trends*

Several promising research directions and emerging trends have been identified that are expected to influence the future development of generative AI and LLMs. The following sections outline key areas of ongoing innovation, highlighting advancements in model architecture, efficiency, domain specialization, reasoning capabilities, ethical AI, regulatory frameworks, and technological integration.

Researchers are exploring architectural innovations beyond transformers, including sparse attention mechanisms, mixture of experts approaches, retrieval-enhanced architectures, and neuro-symbolic approaches [22]. The integration of multiple modalities—text, images, audio, video, and structured data—represents a significant frontier in LLM development [50].

As model scaling faces economic and environmental constraints, research into parameter-efficient architectures is accelerating, including parameter-efficient fine-tuning (PEFT), knowledge distillation, quantization and pruning, and neural architecture search [35]. Innovations in training methodologies aim to reduce computational and data requirements through advances in self-supervised learning, curriculum learning, continual learning, and federated learning [60].

While general-purpose LLMs demonstrate broad capabilities, specialized models tailored to specific domains are likely to proliferate, including scientific, legal, financial, and healthcare-specific models [43]. Expanding beyond English-centric development to better serve global populations through multilingual models, cultural contextualization, and local knowledge integration is an important direction [61].

Addressing current limitations in structured reasoning represents a critical frontier, with improvements to chain-of-thought techniques, integration with external tools, and formal verification methods [34]. Developing stronger capabilities for understanding causality rather than mere correlation through causal inference, counterfactual reasoning, and temporal reasoning is an active area of research [62].

Ensuring that AI systems behave in accordance with human values and intentions through constitutional AI, interpretability research, red-teaming, and value alignment is increasingly important [63]. Addressing issues of bias and fairness through comprehensive frameworks for identifying and mitigating various forms of bias is a critical research direction [17].

The regulatory environment for AI is rapidly developing, with emerging international standards, risk-based regulation, certification and auditing systems, and industry-led self-regulation



initiatives [56]. Mechanisms to ensure responsible development and deployment through standardized documentation, explainability tools, and audit trails are being developed [64].

The integration of LLMs with broader technological ecosystems through AI-enabled infrastructure, autonomous agents, Internet of Things integration, and smart city applications represents a frontier of development [6]. Evolving paradigms for human-AI interaction through collaborative interfaces, augmented creativity, cognitive prosthetics, and personalized AI assistants are emerging.

## IV. Discussion

The presented systematic review reveals the remarkable trajectory of generative AI and LLMs from theoretical concepts to transformative technologies with wide-ranging applications. The evolution of these models has been characterized by several key developments: the breakthrough of the transformer architecture with its self-attention mechanism [12], the scaling of models to unprecedented sizes, and the emergence of capabilities not explicitly programmed [34]. These developments have enabled applications across diverse domains, from healthcare and education to creative industries and public services.

The technical architecture of modern LLMs, centered around self-attention mechanisms and deep neural networks, has proven remarkably effective at capturing the patterns and structures of human language [12]. The scaling properties of these models have revealed fascinating relationships between model size, training data, and performance, suggesting pathways for continued advancement through both scaling and architectural innovation [13].

However, the findings also highlight significant limitations and challenges that must be addressed. Technical constraints such as hallucinations, context window limitations [23], and reasoning deficiencies [54] impact the reliability and applicability of LLMs in critical domains. Ethical concerns regarding bias, privacy, intellectual property, and potential misuse needs careful mitigation strategies. Regulatory uncertainties (while some countries tend to liberate AI development, others want to regulate as strict as possible), environmental impacts, and implementation challenges further complicate the landscape of LLMs development [56].

### *A. Implications for Research and Practice*

The findings suggest several important directions for future research. First, addressing the technical limitations of current LLMs, particularly hallucinations and reasoning limits, requires fundamental advances in model architecture and training methodologies [10]. While there are new models arising daily, research must be continued, so new ways of LLM development approach should be proposed. Developing more efficient methods to train and deploy (in this particular case, inference) is essential for offering access to these technologies for new people and reducing their environmental impact [35]. Third, research into interpretability of different concepts and finding new ways how to explain why some models are behaving in a particulate way is critical for addressing the “black box” nature of current systems [57].

The interdisciplinary nature of challenges associated with LLMs necessitates collaboration across fields including computer science, linguistics, cognitive science, ethics, law, and social sciences [3].

For people working in the field, this review points out how important it is to have smart, well-planned strategies — ones that really take into account what today’s LLMs can and can’t do. Organizations deploying these technologies should implement robust evaluation frameworks, monitoring systems, and governance structures to ensure responsible use [64]. Domain-specific adaptation through fine-tuning or retrieval augmentation is essential for applications in specialized fields, like military and medicine.

The development of clear guidelines for human-AI collaboration can maximize the benefits of these technologies while maintaining appropriate human oversight. Educational initiatives to build

AI literacy among professionals and the general public are important for enabling informed engagement with these technologies [41,61].

### *B. Ethical Considerations*

The general adoption of generative AI and LLMs raises critical ethical questions that require consideration. The potential for these technologies to exacerbate existing inequalities through biased outputs or unequal access demands proactive approaches to fairness and accessibility [17]. The environmental impact of large-scale AI systems necessitates more sustainable approaches to development and deployment [18].

Questions of authorship, originality, and intellectual property in the context of AI-generated content challenge traditional legal and cultural frameworks [55]. The potential for automation to disrupt labor markets requires thoughtful approaches to workforce transition and the development of new roles that leverage human-AI collaboration [65].

### *C. Limitations of the Current Review*

Several limitations of this review should be acknowledged. First, the rapid pace of development in this field means that some recent advancements may not be fully reflected in the literature analyzed. Second, proprietary details of commercial LLMs are often not fully disclosed, potentially limiting understanding of state-of-the-art systems. Third, focus on English-language publications may have excluded valuable insights from non-English research communities.

## **V. Conclusions**

This systematic review has comprehensively examined the evolution, technical architecture, applications, limitations, and future directions of generative AI and Large Language Models. The findings demonstrate the transformative potential of these technologies across diverse domains, while also highlighting significant challenges that must be addressed for responsible development and deployment.

The technical architecture of modern LLMs, centered around self-attention mechanisms and deep neural networks, has enabled unprecedented capabilities in language understanding and generation. These technologies already demonstrate considerable promise across a wide variety of fields, including healthcare, education, software development, creative work, and public services. Their adaptability is clear from the way they're being used to optimize medical research, personalize learning experiences, generate innovative software solutions, spark new forms of art and media, and streamline government operations.

At the same time, there are important limitations which can't be overlooked. On the technical side, issues like model hallucinations and limited reasoning capabilities represent real challenges, often undermining the reliability of the systems and the adoption of such systems. On the ethics side, there are still big concerns about bias, data privacy, and making sure everyone has fair access. Figuring out how to regulate these technologies is still an open question, and on top of that, their environmental impact — since they use a lot of computing power — adds even more complexity. Putting these systems into real-world use isn't easy either, with challenges like high costs and the need to train people to work with them. Tackling all this will take teamwork between tech folks, policymakers, and others, along with ongoing research to keep pushing the limits of what these models can actually do.

New breakthroughs in how these models are built could make them a lot more accurate and efficient, while also using less energy. Improving how well they "reason" could help fix some of the current problems—like when they give wrong or confusing answers. It also makes sense to fine-tune models for specific industries or tasks, so the results they give are more accurate and useful. On top of that, combining LLMs with other types of AI—like computer vision or reinforcement learning—might open the door to totally new and unexpected applications.


These technologies aren't just about technical concepts — they're changing the way we live, work, and connect with each other. As LLMs become part of our everyday lives — helping students learn, supporting doctors, powering apps, inspiring creativity, and even helping write this review — their impact on society and culture keeps growing. To make sure that impact is a positive one, we need to build in ethical values, fairness, and accountability from the ground up. And we also need smart governance — laws, policies, and systems that put people first and protect things like privacy, equality, and responsible innovation.

Generative AI and Large Language Models represent a groundbreaking convergence of potential and complexity. If we make the most of what these tools can do—while keeping a close eye on their flaws and the bigger picture—we've got a real opportunity to shape them into something that truly helps people and benefits society as a whole. What is worth mentioning is that if we compare the evolution of smartphones—taking the launch of the first iPhone in 2007 as a reference point—with the launch of ChatGPT in 2022, and we extrapolate based on how far smartphone technology has come, it becomes clear that the next AI revolution is a matter of years, not decades.

**Acknowledgment:** During the preparation of this review, the author used Manus for the purposes of summarization and text-generation for some general statements. The author has reviewed and edited the output and takes full responsibility for the content of this publication.

## References

1. J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," May 24, 2019, *arXiv*: arXiv:1810.04805. doi: 10.48550/arXiv.1810.04805.
2. T. B. Brown et al., "Language Models are Few-Shot Learners," 2020, *arXiv*. doi: 10.48550/ARXIV.2005.14165.
3. R. Bommasani et al., "On the Opportunities and Risks of Foundation Models," 2021, *arXiv*. doi: 10.48550/ARXIV.2108.07258.
4. A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," OpenAI Blog. Accessed: Apr. 01, 2025. [Online]. Available: [https://cdn.openai.com/better-language-models/language\\_models\\_are\\_unsupervised\\_multitask\\_learners.pdf](https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf)
5. P. Yu, H. Xu, X. Hu, and C. Deng, "Leveraging Generative AI and Large Language Models: A Comprehensive Roadmap for Healthcare Integration," *Healthcare*, vol. 11, no. 20, p. 2776, Oct. 2023, doi: 10.3390/healthcare11202776.
6. G. Salierno, L. Leonardi, and G. Cabri, "Generative AI and Large Language Models in Industry 5.0: Shaping Smarter Sustainable Cities," *Encyclopedia*, vol. 5, no. 1, p. 30, Feb. 2025, doi: 10.3390/encyclopedia5010030.
7. M. Sallam, "ChatGPT Utility in Healthcare Education, Research, and Practice: Systematic Review on the Promising Perspectives and Valid Concerns," *Healthcare*, vol. 11, no. 6, p. 887, Mar. 2023, doi: 10.3390/healthcare11060887.
8. M. Chen et al., "Evaluating Large Language Models Trained on Code," 2021, *arXiv*. doi: 10.48550/ARXIV.2107.03374.
9. R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-Resolution Image Synthesis with Latent Diffusion Models," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, LA, USA: IEEE, Jun. 2022, pp. 10674–10685. doi: 10.1109/CVPR52688.2022.01042.
10. K. Shuster, S. Poff, M. Chen, D. Kiela, and J. Weston, "Retrieval Augmentation Reduces Hallucination in Conversation," in *Findings of the Association for Computational Linguistics: EMNLP 2021*, Punta Cana, Dominican Republic: Association for Computational Linguistics, 2021, pp. 3784–3803. doi: 10.18653/v1/2021.findings-emnlp.320.
11. Y. Bengio, R. Ducharme, and P. Vincent, "A neural probabilistic language model," in *Proceedings of the 14th International Conference on Neural Information Processing Systems*, in NIPS'00. Cambridge, MA, USA: MIT Press, 2000, pp. 893–899.
12. A. Vaswani et al., "Attention Is All You Need," 2017, *arXiv*. doi: 10.48550/ARXIV.1706.03762.

13. J. Kaplan et al., "Scaling Laws for Neural Language Models," Jan. 23, 2020, *arXiv*: arXiv:2001.08361. doi: 10.48550/arXiv.2001.08361.
14. H. Touvron et al., "LLaMA: Open and Efficient Foundation Language Models," 2023, *arXiv*. doi: 10.48550/ARXIV.2302.13971.
15. DeepSeek-AI et al., "DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning," 2025, *arXiv*. doi: 10.48550/ARXIV.2501.12948.
16. E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? , in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Virtual Event Canada: ACM, Mar. 2021, pp. 610–623. doi: 10.1145/3442188.3445922.
17. T. Bolukbasi, K.-W. Chang, J. Zou, V. Saligrama, and A. Kalai, "Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings," 2016, *arXiv*. doi: 10.48550/ARXIV.1607.06520.
18. D. Patterson et al., "Carbon Emissions and Large Neural Network Training," 2021, *arXiv*. doi: 10.48550/ARXIV.2104.10350.
19. A. Tamkin, M. Brundage, J. Clark, and D. Ganguli, "Understanding the Capabilities, Limitations, and Societal Impact of Large Language Models," 2021, *arXiv*. doi: 10.48550/ARXIV.2102.02503.
20. D. Moher, A. Liberati, J. Tetzlaff, D. G. Altman, and The PRISMA Group, "Preferred Reporting Items for Systematic Reviews and Meta-Analyses: The PRISMA Statement," *PLoS Med.*, vol. 6, no. 7, p. e1000097, Jul. 2009, doi: 10.1371/journal.pmed.1000097.
21. Critical Appraisal Skills Programme, *CASP Systematic Review Checklist*, 2018. Accessed: Apr. 01, 2025. [Online]. Available: <https://casp-uk.net/casp-tools-checklists/systematic-review-checklist/>
22. N. Kitaev, L. Kaiser, and A. Levskaya, "Reformer: The Efficient Transformer," 2020, *arXiv*. doi: 10.48550/ARXIV.2001.04451.
23. I. Beltagy, M. E. Peters, and A. Cohan, "Longformer: The Long-Document Transformer," Dec. 02, 2020, *arXiv*: arXiv:2004.05150. doi: 10.48550/arXiv.2004.05150.
24. P. Shaw, J. Uszkoreit, and A. Vaswani, "Self-Attention with Relative Position Representations," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, New Orleans, Louisiana: Association for Computational Linguistics, 2018, pp. 464–468. doi: 10.18653/v1/N18-2074.
25. R. Sennrich, B. Haddow, and A. Birch, "Neural Machine Translation of Rare Words with Subword Units," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany: Association for Computational Linguistics, 2016, pp. 1715–1725. doi: 10.18653/v1/P16-1162.
26. D. Hendrycks and K. Gimpel, "Gaussian Error Linear Units (GELUs)," 2016, *arXiv*. doi: 10.48550/ARXIV.1606.08415.
27. J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer Normalization," Jul. 21, 2016, *arXiv*: arXiv:1607.06450. doi: 10.48550/arXiv.1607.06450.
28. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
29. A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, and others, "Improving language understanding by generative pre-training," 2018.
30. C. Raffel et al., "Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer," 2019, *arXiv*. doi: 10.48550/ARXIV.1910.10683.
31. A. Chowdhery et al., "PaLM: Scaling Language Modeling with Pathways," 2022, *arXiv*. doi: 10.48550/ARXIV.2204.02311.
32. J. Hoffmann et al., "Training Compute-Optimal Large Language Models," Mar. 29, 2022, *arXiv*: arXiv:2203.15556. doi: 10.48550/arXiv.2203.15556.
33. T. Henighan et al., "Scaling Laws for Autoregressive Generative Modeling," Nov. 06, 2020, *arXiv*: arXiv:2010.14701. doi: 10.48550/arXiv.2010.14701.
34. J. Wei et al., "Chain-of-thought prompting elicits reasoning in large language models," in *Proceedings of the 36th International Conference on Neural Information Processing Systems*, in NIPS '22. Red Hook, NY, USA: Curran Associates Inc., 2022.

35. E. J. Hu et al., "LoRA: Low-Rank Adaptation of Large Language Models," Oct. 16, 2021, *arXiv*: arXiv:2106.09685. doi: 10.48550/arXiv.2106.09685.
36. S. Han, H. Mao, and W. J. Dally, "Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding," Feb. 15, 2016, *arXiv*: arXiv:1510.00149. doi: 10.48550/arXiv.1510.00149.
37. G. Hinton, O. Vinyals, and J. Dean, "Distilling the Knowledge in a Neural Network," Mar. 09, 2015, *arXiv*: arXiv:1503.02531. doi: 10.48550/arXiv.1503.02531.
38. J. Zhang, Y. Zhao, M. Saleh, and P. Liu, "Pegasus: Pre-training with extracted gap-sentences for abstractive summarization," in *International conference on machine learning*, PMLR, 2020, pp. 11328–11339.
39. M. Johnson et al., "Google's Multilingual Neural Machine Translation System: Enabling Zero-Shot Translation," *Trans. Assoc. Comput. Linguist.*, vol. 5, pp. 339–351, 2017, doi: 10.1162/tacl\_a\_00065.
40. S. Roller et al., "Recipes for building an open-domain chatbot," Apr. 30, 2020, *arXiv*: arXiv:2004.13637. doi: 10.48550/arXiv.2004.13637.
41. Z. Chu et al., "LLM Agents for Education: Advances and Applications," 2025, *arXiv*. doi: 10.48550/ARXIV.2503.11733.
42. X. Ouyang et al., "ERNIE-M: Enhanced Multilingual Representation by Aligning Cross-lingual Semantics with Monolingual Corpora," Sep. 17, 2021, *arXiv*: arXiv:2012.15674. doi: 10.48550/arXiv.2012.15674.
43. I. Beltagy, K. Lo, and A. Cohan, "SciBERT: A Pretrained Language Model for Scientific Text," Sep. 10, 2019, *arXiv*: arXiv:1903.10676. doi: 10.48550/arXiv.1903.10676.
44. A. Esteva et al., "CO-Search: COVID-19 Information Retrieval with Semantic Search, Question Answering, and Abstractive Summarization," Jun. 17, 2020, *arXiv*: arXiv:2006.09595. doi: 10.48550/arXiv.2006.09595.
45. J. C. L. Chow, V. Wong, and K. Li, "Generative Pre-Trained Transformer-Empowered Healthcare Conversations: Current Trends, Challenges, and Future Directions in Large Language Model-Enabled Medical Chatbots," *BioMedInformatics*, vol. 4, no. 1, pp. 837–852, Mar. 2024, doi: 10.3390/biomedinformatics4010047.
46. B.-I. Ciubotaru et al., "Frailty Insights Detection System (FIDS)—A Comprehensive and Intuitive Dashboard Using Artificial Intelligence and Web Technologies," *Appl. Sci.*, vol. 14, no. 16, Art. no. 16, Jan. 2024, doi: 10.3390/app14167180.
47. A. Svyatkovskiy, S. K. Deng, S. Fu, and N. Sundaresan, "IntelliCode compose: code generation using transformer," in *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, Virtual Event USA: ACM, Nov. 2020, pp. 1433–1443. doi: 10.1145/3368089.3417058.
48. S. Iyer, I. Konstas, A. Cheung, and L. Zettlemoyer, "Summarizing Source Code using a Neural Attention Model," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany: Association for Computational Linguistics, 2016, pp. 2073–2083. doi: 10.18653/v1/P16-1195.
49. M. Grootendorst, "BERTopic: Neural topic modeling with a class-based TF-IDF procedure," Mar. 11, 2022, *arXiv*: arXiv:2203.05794. doi: 10.48550/arXiv.2203.05794.
50. A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical Text-Conditional Image Generation with CLIP Latents," Apr. 13, 2022, *arXiv*: arXiv:2204.06125. doi: 10.48550/arXiv.2204.06125.
51. C. Hawthorne et al., "Enabling Factorized Piano Music Modeling and Generation with the MAESTRO Dataset," Jan. 17, 2019, *arXiv*: arXiv:1810.12247. doi: 10.48550/arXiv.1810.12247.
52. D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, "Concrete Problems in AI Safety," Jul. 25, 2016, *arXiv*: arXiv:1606.06565. doi: 10.48550/arXiv.1606.06565.
53. J. Wei, J. He, K. Chen, Y. Zhou, and Z. Tang, "Collaborative filtering and deep learning based recommendation system for cold start items," *Expert Syst. Appl.*, vol. 69, pp. 29–39, Mar. 2017, doi: 10.1016/j.eswa.2016.09.040.
54. D. Hendrycks et al., "Measuring Mathematical Problem Solving With the MATH Dataset," Nov. 08, 2021, *arXiv*: arXiv:2103.03874. doi: 10.48550/arXiv.2103.03874.
55. M. A. Lemley and B. Casey, "Fair Learning," *SSRN Electron. J.*, 2020, doi: 10.2139/ssrn.3528447.



56. European Commission, "Proposal for a Regulation laying down harmonised rules on artificial intelligence." Accessed: Jan. 04, 2025. [Online]. Available: <https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence>
57. F. Doshi-Velez and B. Kim, "Towards A Rigorous Science of Interpretable Machine Learning," Mar. 02, 2017, *arXiv*: arXiv:1702.08608. doi: 10.48550/arXiv.1702.08608.
58. E. Strubell, A. Ganesh, and A. McCallum, "Energy and Policy Considerations for Deep Learning in NLP," Jun. 05, 2019, *arXiv*: arXiv:1906.02243. doi: 10.48550/arXiv.1906.02243.
59. K. Guu, K. Lee, Z. Tung, P. Pasupat, and M.-W. Chang, "REALM: Retrieval-Augmented Language Model Pre-Training," 2020, *arXiv*. doi: 10.48550/ARXIV.2002.08909.
60. Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th Annual International Conference on Machine Learning*, Montreal Quebec Canada: ACM, Jun. 2009, pp. 41–48. doi: 10.1145/1553374.1553380.
61. N. T. Tuan, P. Moore, D. H. V. Thanh, and H. V. Pham, "A Generative Artificial Intelligence Using Multilingual Large Language Models for ChatGPT Applications," *Appl. Sci.*, vol. 14, no. 7, p. 3036, Apr. 2024, doi: 10.3390/app14073036.
62. Y. Bengio et al., "A Meta-Transfer Objective for Learning to Disentangle Causal Mechanisms," Feb. 04, 2019, *arXiv*: arXiv:1901.10912. doi: 10.48550/arXiv.1901.10912.
63. J. Leike, D. Krueger, T. Everitt, M. Martic, V. Maini, and S. Legg, "Scalable agent alignment via reward modeling: a research direction," Nov. 19, 2018, *arXiv*: arXiv:1811.07871. doi: 10.48550/arXiv.1811.07871.
64. M. Mitchell et al., "Model Cards for Model Reporting," 2018, doi: 10.48550/ARXIV.1810.03993.
65. D. Acemoglu and P. Restrepo, "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *J. Econ. Perspect.*, vol. 33, no. 2, pp. 3–30, May 2019, doi: 10.1257/jep.33.2.3.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.