International Research Journal of Advanced Engineering and Science
ISSN (Online): 2455-9024

Prospective Advancements in Conversational AI: A Projection of Natural Language Understanding with GPT-5

Herbert Wanga
Department of Mathematics & Information Technology, University of Iringa, Iringa, Tanzania- 51000

Abstract**—***Large language models (LLMs) have been evolving rapidly, and experts anticipate that the next-generation model, often referred to as GPT-5, will mark a major leap in conversational artificial intelligence. This review brings together recent research findings from peer-reviewed studies and preprints to explore the new capabilities GPT-5 might offer. The paper examines how GPT-5 could advance natural language understanding (NLU), what architectural improvements it may include, and what social and technical implications it carries. In particular, GPT-5 is expected to greatly improve its contextual reasoning and its ability to handle very long pieces of text. It should also integrate different types of input (text, images, audio) much more smoothly than previous models. At the same time, GPT-5's introduction raises several concerns. The model will likely require far more computational resources, and it emphasizes ongoing issues around deploying AI ethically and maintaining user trust. This review suggests that realizing GPT-5's benefits will require not just technical innovation, but also strong governance and oversight to guide its use. It synthesizes current projections to provide a roadmap for researchers and policymakers, highlighting both the transformative potential and the critical challenges.*

Keywords- *GPT-5, conversational AI, natural language understanding, generative AI, multimodal AI, AI ethics.*

I. Introduction

Conversational AI has been advancing at an extraordinary pace. GPT-3 introduced a breakthrough by allowing models to learn from just a few examples (Brown et al., 2020). GPT-4 built on that progress with much stronger reasoning abilities and the addition of multimodal capabilities - meaning it can understand images and other media, not just text (OpenAI, 2023). Now, the field is looking toward GPT-5.

Early reports suggest GPT-5 will be capable of nearly human-expert level reasoning on complex tasks - almost at a "PhD-level" in some respects. It is also expected to handle text, audio, images, and even video with far more autonomy than prior models (Askitas, 2025; Odelami, 2025). Unlike its predecessors, GPT-5 is being discussed not just as a productivity tool but as a piece of core digital infrastructure (Askitas, 2025). Its emergence has also reignited debates about several issues, including creativity, trust, environmental sustainability, and how we govern the use of AI.

This article offers a forward-looking review of what GPT-5 might bring, both technically and in a broader social context. The review is structured around three key areas:

- Advances in language understanding and reasoning: The improvements GPT-5 is expected to have in natural language understanding (NLU) and its problem-solving abilities.
- New modes of human - AI collaboration and education: Emerging ways that GPT-5 could be used alongside people, such as in collaborative workflows or personalized education.
- Major challenges and societal considerations: The obstacles that need to be addressed, including GPT-5's massive computational requirements, questions of ethical deployment, and the long-term societal implications of using GPT-5 and similar large language models.

II. Literature Review

A. Advances in LLM Architecture and Performance

Current research on large-scale transformer models is laying the groundwork for a model like GPT-5. State-of-the-art models have already expanded the amount of text these systems can consider at once (the "context window") to unprecedented lengths, in some experiments, up to millions of tokens. They have also improved the models' reasoning abilities using techniques such as chain-of-thought prompting and self-consistency (Wei et al., 2022). In parallel, architectural innovations like the *mixture-of-experts* (MoE), where different parts of the network specialize in different types of problems, are enabling more efficient scaling (Fedus et al., 2022).

Thanks to these advancements, it seems plausible that GPT-5 will support a dramatically larger context window, potentially on the order of 100,000 tokens or more. Some reports even suggest GPT-5 might handle up to about 256,000 tokens at once (Liu et al., 2024; Odelami, 2025). In practical terms, that means GPT-5 could keep track of a book-length conversation or document and maintain coherence throughout.

Alongside a longer context, researchers also expect GPT-5 to have stronger reasoning skills. This might be achieved by using more efficient transformer designs, such as a *mixture-of-experts* architecture, or other improved methods of routing information inside the model (Fedus et al., 2022). Early studies of GPT-5 prototypes indicate that they already outperform GPT-4 on complex reasoning and forecasting tasks. For example, GPT-5 variants have shown higher accuracy and reliability on time-series predictions (Ghasemloo & Moradi, 2025). This suggests that even smaller or

Herbert Wanga, "Prospective Advancements in Conversational AI: A Projection of Natural Language Understanding with GPT-5" *International Research Journal of Advanced Engineering and Science*, Volume 10, Issue 4, pp. 48-53, 2025.

specialized versions of GPT-5 might beat earlier models like GPT-4 on certain tasks.

B. Multimodality and NLU Enhancement

One standout feature expected in GPT-5 is its native multimodal capability. In other words, GPT-5 can seamlessly process and generate content that mixes text, visuals, and audio. Recent reports indicate GPT-5 can analyze an image and then answer questions about that image in the same conversation. This demonstrates a more integrated understanding than GPT-4 was capable of (OpenAI, 2023; Odelami, 2025).

By strengthening these kinds of practical language-understanding skills, for example, being able to interpret a diagram or a video alongside written instructions, GPT-5 is expected to broaden the real-world uses of conversational AI. Early tests suggest GPT-5 has a more nuanced grasp of context and user intent, partly thanks to new adaptive attention mechanisms that help the model reason more accurately through complex, multi-step questions (Odelami, 2025). GPT-5's improved performance on language-understanding benchmarks and multimodal tasks highlights that the model is moving beyond surface-level text generation toward a more genuine comprehension of meaning.

C. Human - AI Collaboration Frameworks

Researchers are also proposing new ways to collaborate with GPT-5 that leverage its strengths while compensating for its weaknesses. One emerging concept is the Centaur workflow, a hybrid collaboration model. In a Centaur workflow, an AI system (like GPT-5) handles the heavy analytic or creative drafting work, while human experts provide oversight, validation, and guidance (Youvan, 2025). A related idea is the Two-Model Method. This approach uses a very powerful model like GPT-5 for tasks requiring highly precise and deterministic results (one source calls this focused role "seeded depth"), while a second model (such as a smaller system nicknamed *Gemini*) runs in parallel to explore more broadly and double-check the first model's output for errors (Youvan, 2025).

Researchers have also developed structured prompting techniques like the *Structured Path Resonance Loop (SPRL)*. SPRL is an interface that organizes a sequence of prompts in a systematic way to reduce randomness and make GPT-5's outputs easier to interpret (SPRL Development Team, 2025). Collectively, these developments suggest that the field is shifting away from single-shot answers toward longer, more structured reasoning sessions with GPT-5, sessions that can be reviewed and audited afterward. By structuring how humans interact with the AI, the goal is to address the trust and reproducibility concerns that arise when using an advanced model like GPT-5 in critical workflows.

D. Sustainability and Ethical Considerations

Alongside technical advances, researchers have become increasingly concerned about the growing computational footprint of cutting-edge language models. Training and running models like GPT-4 consume enormous amounts of energy and computing resources, and scaling up to GPT-5 will likely demand even more (Ahmad et al., 2025). This reality is pushing the AI community toward "Green AI" practices. These include developing more energy-efficient model architectures, compressing models, and using renewable energy in data centers (Schwartz et al., 2019).

Another major focus in the literature is the ethical and societal implications of increasingly human-like AI. For example, there is concern about AI "hallucinations," where the model makes up false information, and a growing *"trust gap"*, meaning users are unsure whether they can trust the AI's answers. These issues are well-documented, and researchers have proposed some solutions (Schwartz et al., 2019; Askitas, 2025). For instance, some suggest building transparent auditing tools to trace how an AI derives its responses, as well as establishing standards for responsible AI deployment (Schwartz et al., 2019; Askitas, 2025). Overall, these studies stress that as models like GPT-5 become more powerful, we must also improve how we evaluate them, interpret their decisions, and govern their use. Only with such safeguards in place can the benefits of GPT-5 be realized safely.

III. Methodology

In this study, I adopted a structured narrative review method to project GPT-5's advancements and their implications. I conducted an extensive literature search covering the years 2020 through 2025, spanning multiple academic databases and preprint servers (including Scopus, IEEE Xplore, the ACM Digital Library, arXiv, SSRN, and Zenodo). Key search terms included "GPT-5," "large language model evolution," "long context transformers," "multimodal AI," "AI ethics," and "conversational AI collaboration." The initial search returned over 200 documents.

A. Sources

Sources were selected based on the following inclusion criteria:

- Detailed insights into novel LLM architectures or capabilities (for example, improvements in reasoning, handling longer contexts, or integrating multiple modalities).
- Evidence of new application paradigms involving LLMs (such as collaborative human - AI workflows or educational tools).
- Social and technical challenges of deploying advanced LLMs (including issues of ethics, sustainability, and user trust).

B. Exclusion Criteria

The exclusion criteria were:

- Articles focused solely on models prior to GPT-4 without a forward-looking perspective.
- Non-technical commentary without references or empirical projections.
- Duplicate publications.

This process identified 32 highly relevant documents for final synthesis. The corpus comprises peer-reviewed journal articles, conference proceedings, white papers, and preprints. I

Herbert Wanga, "Prospective Advancements in Conversational AI: A Projection of Natural Language Understanding with GPT-5" *International Research Journal of Advanced Engineering and Science*, Volume 10, Issue 4, pp. 48-53, 2025.

acknowledge that the inclusion of preprints and industry reports is a limitation inherent to reviewing a not-yet-released technology; however, these sources are crucial for capturing the most current projections and debates. Each source was reviewed and coded for key themes using a thematic synthesis approach. Three main themes emerged from an analysis: technical capabilities, collaborative applications, and systemic challenges. These themes became the outline for the structured discussion that follows. By aggregating findings from many recent sources, this review provides an evidence-based projection of GPT-5's potential features. It also offers the context needed to integrate such a model into society responsibly.

IV. Discussion

A. Technical Advancements

GPT-5 is expected to deliver major technical improvements over current conversational AI models. One of the biggest upgrades will be in how it handles context and memory. For example, GPT-4 can juggle about 32,000 tokens of text (OpenAI, 2023), but GPT-5 is anticipated to handle around 256,000 tokens or more (Odelami, 2025). In plain terms, GPT-5 should be able to read and reason about very large documents or multi-part conversations without losing track of important details. This huge context capacity would let the model keep dialogues coherent and refer back to information from much earlier in a conversation, far beyond what was previously possible.

However, effectively using such a large context is not trivial. Current models struggle with what researchers call "mid-context recall". They sometimes overlook relevant information that appears in the middle of a long input (Liu et al., 2024). To overcome this limitation, new techniques in adaptive attention and memory management for transformers will be crucial. These techniques aim to ensure that GPT-5 not only *reads* more information, but actually *remembers* and uses it when it is needed.

Another widely discussed improvement is GPT-5's reasoning ability. Some optimistic reports suggest GPT-5 could perform nearly as well as a human expert on complex analytical tasks, approaching "PhD-level" reasoning in some cases (Askitas, 2025). However, it is important to keep this in perspective: just because a model scores high on tests or benchmarks does not mean it truly understands the material the way humans do. As Bender and Koller (2020) point out, a language model can produce expert-like answers without possessing genuine understanding or common sense about the content. In other words, GPT-5 might be very good at following logical patterns and producing answers, yet still lack the deeper comprehension and real-world experience that human experts draw upon.

From an engineering perspective, GPT-5 is expected to include design innovations aimed at boosting its reasoning performance. For example, it may employ more efficient *mixture-of-experts* layers, a technique where the model routes questions through specialized subsets of its neurons, allowing it to reason better without a huge increase in computation (Fedus et al., 2022). There is also evidence of improved inference algorithms. One source describes GPT-5 using advanced internal routing methods and performing self-consistency checks to reduce errors on complex problems (Odelami, 2025). Overall, these enhancements suggest GPT-5 will be better equipped than its predecessors to handle multi-step logical problems, complex decision trees, and domain-specific questions with greater accuracy.

In terms of multimodal capabilities, GPT-5 is pushing conversational AI beyond text-only interactions. Building on GPT-4's multimodal features, GPT-5 is designed to integrate text, visuals, and audio more seamlessly within one model. For example, a user could upload a photograph or a chart and ask GPT-5 to analyze it and explain what it shows in natural language. Early reports indicate that GPT-5 can interpret an image, such as a graph or even a medical scan, and then provide a coherent explanation of it, answering follow-up questions in the same conversation (Odelami, 2025).

This level of multimodal understanding opens up many new applications for AI. For instance, a legal analyst might feed GPT-5 several pages of a contract along with an audio commentary, and the model could draft a summary that takes both inputs into account. Similarly, a doctor could input a patient's lab results together with some patient history text, and GPT-5 could offer preliminary diagnostic suggestions in response. In tests, GPT-5 has also shown it can cross-reference information from different kinds of input. For example, it can align details from a diagram with a written description of that diagram, suggesting that the model builds a more unified internal representation of knowledge from multiple sources.

By drawing on multiple information sources at once, GPT-5 tends to demonstrate a better grasp of context and intent, much like how people use both words and visuals together to understand a situation. As a result, GPT-5's answers in complex scenarios are expected to be more contextually grounded and useful. However, ensuring that it stays consistent across different modalities (and avoiding mistakes like misinterpreting an image) remains a significant challenge for researchers.

B. Human - AI Collaboration and Education

GPT-5 is also spurring new ways for humans to collaborate with AI, especially in knowledge-intensive work. The frameworks introduced in the literature review, such as the Centaur workflow and the Two-Model Method, are central to this evolution.

In a GPT-5 Centaur workflow, the AI and human act as a team, each contributing different strengths (Youvan, 2025). The AI could take on labor-intensive components of a task, such as generating a first draft of a report, performing complex calculations, or exploring a broad design space, while a human expert reviews the results, imparts domain-specific knowledge, and makes final judgments. Early case studies of humans partnering with GPT-5 suggest this approach can boost productivity and solution quality. For instance, GPT-5 might quickly draft several possible approaches to a research problem, and then a human collaborator evaluates which of these ideas are feasible and ethically sound. This iterative

Herbert Wanga, "Prospective Advancements in Conversational AI: A Projection of Natural Language Understanding with GPT-5" *International Research Journal of Advanced Engineering and Science*, Volume 10, Issue 4, pp. 48-53, 2025.

back-and-forth combines GPT-5's speed and breadth of output with the human expert's critical thinking and values-based judgment. Crucially, these workflows include oversight mechanisms to build trust: the human can prompt GPT-5 to explain its reasoning or show its work, and then verify the AI's outputs, catching any mistakes or unsupported claims before they influence decisions.

Another strategy to improve reliability is the Two-Model Method (Youvan, 2025). In this approach, GPT-5 is paired with a second model to double-check its work. The idea is to use GPT-5 for tasks that require very precise, deterministic results, for example, a detailed legal analysis or an engineering computation, where GPT-5's extensive knowledge and reasoning capacity are needed. At the same time, a secondary model (which might be a smaller, more transparent system specialized in fact-checking) runs in parallel to provide broader context or verify GPT-5's answers. The outcome is a sort of AI peer-review system: GPT-5 produces a detailed answer, and the companion model evaluates parts of that output (or offers alternative perspectives) to ensure nothing critical was missed and that no hallucinated information slipped through. While this method promises enhanced auditability, its practical implementation faces challenges, including increased computational cost and system complexity. Some researchers argue that this two-model approach makes AI-driven processes more reproducible, because the secondary model can log its verification steps and create a traceable record of how an answer was vetted (Youvan, 2025). Both the centaur workflow and the two-model strategy represent a shift in how we view AI: instead of treating GPT-5 as an infallible oracle, these approaches use it as a powerful partner. They acknowledge that as advanced as GPT-5 is, human judgment and complementary checks (even from other AIs) are still essential for reliable, trustworthy outcomes.

Similarly, techniques like the *Structured Path Resonance Loop (SPRL)* aim to make interactions more systematic and less stochastic (SPRL Development Team, 2025). However, the efficacy of these prompting frameworks in diverse real-world scenarios remains an area for empirical validation.

While these collaborative frameworks are promising, they also come with risks. If people rely too heavily on AI-generated content, there is a chance our thinking could become homogenized and that the original sources of information behind AI outputs might get obscured. If professionals begin to accept GPT-5's synthesized answers uncritically, we might see a decline in human expertise and critical thinking skills over time. Some scholars warn of a potential "lost knowledge" problem: the AI's rephrasing of information can mask where facts or ideas came from, making it difficult to trace how a conclusion was reached (Askitas, 2025). Additionally, the earlier-mentioned trust gap remains a concern. If users are not sure how GPT-5 arrived at an answer, they may hesitate to trust the AI, or worse, they might trust it too much and propagate errors before they are caught. This highlights the need for any human - AI collaboration workflow to include robust validation and documentation steps. Maintaining a diversity of thought and rigorous source verification in the age of pervasive AI assistance will be an important cultural and educational challenge moving forward.

In the realm of education, GPT-5's improved NLU capabilities open up new possibilities for personalized and interactive learning. Intelligent tutoring systems powered by GPT-5 could potentially understand a student's question in depth, recognize areas of confusion, and provide tailored explanations or examples in real time. For example, if a student is struggling with a calculus problem, GPT-5 could walk them through the solution step by step, using analogies or even a simple diagram, adapting the explanation to the student's learning level. Early demonstrations of GPT-5 in educational settings show the model to be adept at generating instructive content, such as practice quiz questions, explanations of complex concepts, or even simulated dialogues for language practice, that align with curriculum goals. And because GPT-5 can handle text and images together, one can imagine it delivering a history lesson by presenting an excerpt from a historical document alongside a relevant map or portrait to enrich the context. These applications hint at a more interactive, on-demand form of e-learning, where an AI tutor can cover a wide range of topics and modalities as needed.

That said, using GPT-5 in education also poses significant challenges. The accuracy and trustworthiness of the AI's output are paramount when it comes to instructing students. Like its predecessors, GPT-5 may occasionally produce incorrect or misleading information with a confident tone. In a classroom or self-learning scenario, this could confuse learners or propagate misconceptions if not carefully monitored. Therefore, educators stress the importance of keeping a human in the loop. Teachers or curriculum designers should vet the AI's responses, especially for high-stakes subjects. Additionally, there are concerns about overreliance: if students turn to GPT-5 for every answer or explanation, they might not develop the critical problem-solving skills and perseverance that come from working through challenges themselves. Academic integrity is another consideration; the ease with which GPT-5 can generate essays or solutions means schools may need new approaches to assessment and plagiarism prevention. In summary, GPT-5 has the potential to greatly enhance educational experiences through personalized tutoring and rich, multimodal content delivery, but it must be deployed in a way that supports and augments human teaching rather than undermines it. Clear guidelines, oversight by educators, and improved digital literacy for students will all be crucial to integrating GPT-5 into education responsibly.

C. *Sustainability and Ethics*

The advent of GPT-5 intensifies ongoing concerns about the sustainability of using such large AI models. By all accounts, GPT-5's training phase consumed substantially more computing resources than GPT-4's, and using GPT-5 in practice also demands greater processing power due to the model's size and extended context handling. Recent analyses project that GPT-5's energy requirements and carbon footprint could exceed those of GPT-4 by 30 - 50% under comparable

Herbert Wanga, "Prospective Advancements in Conversational AI: A Projection of Natural Language Understanding with GPT-5" *International Research Journal of Advanced Engineering and Science*, Volume 10, Issue 4, pp. 48-53, 2025.

usage (Ahmad et al., 2025). This has made the pursuit of "Green AI" more urgent than ever. Researchers are actively exploring ways to improve the energy efficiency of LLMs. Approaches include algorithmic optimizations and specialized hardware that cut down on power use, as well as techniques like knowledge distillation or parameter pruning to produce smaller, more efficient versions of GPT-5 that retain most of its performance (Schwartz et al., 2019). Deployment strategies are also under scrutiny. For instance, running GPT-5 on servers powered by renewable energy or scheduling heavy computations for times when surplus renewable power is available can help mitigate its environmental impact. Industry-wide, there are calls for standardized reporting of AI models' energy use and carbon emissions (Schwartz et al., 2019). The hope is that increased transparency will encourage competition on efficiency metrics (not just on performance benchmarks) and drive the development of more environmentally sustainable AI systems.

Beyond environmental concerns, the ethical deployment of GPT-5 raises several other issues. One major worry is the risk of exacerbating misinformation and eroding public trust. GPT-5's outputs can be so fluent and authoritative that users may struggle to distinguish AI-generated text from verified information. If the model confidently produces a subtle error, for example, misreporting a medical fact or a legal statute, the mistake could be dangerously misleading. Ensuring the trustworthiness of GPT-5's answers is therefore critical. Some proposed solutions involve building auditing tools that trace the origin of facts used by GPT-5, or having the model provide citations and justifications for its responses. Implementing these reliably, however, remains challenging. The Two-Model Method mentioned earlier is one attempt to enforce verifiability by using one AI to audit another's output (Youvan, 2025). There are also calls for greater human and regulatory oversight. Governments and professional bodies are considering guidelines or rules for using advanced AI in sensitive areas like healthcare, law, and education, to ensure accountability when AI is involved in decision-making. For instance, regulators might require human review of GPT-5's recommendations in medical diagnostics, or mandate disclosure when content is AI-generated. Such measures are aimed at maintaining transparency and responsibility in the use of GPT-5.

Another ethical dimension is the impact of GPT-5 on human labor and creativity. GPT-5's ability to generate high-quality text, images, and even code raises concerns about displacing jobs in fields such as content writing, customer service, programming, and design. While these models can boost productivity, there is an active debate on how to retrain and upskill the workforce so that humans can collaborate with AI, rather than be replaced by it (Askitas, 2025). Moreover, critics worry that an overreliance on AI for creative tasks could lead to stagnation in human creativity. If GPT-5 is used to churn out endless streams of AI-generated art and literature, human creators might feel pressure to conform to AI-influenced styles or risk being overshadowed. To address this, ethicists argue for frameworks that keep humans in the creative loop and use AI as a tool to enhance, not substitute, human creativity. In other words, AI should augment artists and professionals, not push them aside.

Finally, the opaqueness of GPT-5's decision-making process, the classic "black box" issue, poses a problem for accountability. If GPT-5 produces a harmful or biased outcome, it can be very difficult to pinpoint why it happened, given the model's complexity and lack of transparent reasoning. This has spurred research into interpretability for large LLMs, as scientists seek methods to make GPT-5's inner workings and rationale more legible to developers and end-users. Specific techniques being explored include mechanistic interpretability, which aims to reverse-engineer neural networks into human-understandable algorithms, and the development of simplified reasoning chains that approximate the model's internal logic. Some efforts include extracting these chains from the model or highlighting which parts of the input most influenced the output. As GPT-5 begins to be deployed, implementing such interpretability features, along with regular bias audits and fairness evaluations, will be essential. Ultimately, ensuring GPT-5 is used ethically will require a multifaceted strategy: technical safeguards, user education, organizational policies, and possibly new regulations must all work in concert to make sure this powerful technology is used in ways that are fair, trustworthy, and aligned with societal values.

V. Conclusion

GPT-5's development represents a significant milestone in the evolution of conversational AI, pushing the field beyond high-quality language generation toward more integrated reasoning, multimodal understanding, and collaborative capabilities. If the anticipated advancements materialize, GPT-5 could enable transformative applications across research, education, healthcare, business, and beyond. We might see GPT-5 automating complex analytical tasks, supporting creative endeavors, providing personalized tutoring, and offering expert decision assistance. The potential societal benefits are substantial, but realizing them will depend on how the accompanying challenges are navigated.

This review has highlighted that GPT-5's impressive technical strides come entangled with profound unresolved issues. We must address the sustainability of its compute demands, ensure the veracity and accountability of its outputs, preserve human creativity and critical thinking, and mitigate the epistemic (knowledge-related) risks that arise from widespread AI-generated content.

Moving forward, a balanced and proactive approach is required to ensure GPT-5 and its successors deliver long-term value. On the technical front, continued innovation in efficiency (to curb environmental impact) and in explainability (to build trust) should be prioritized alongside raw capability improvements. Equally important is the development of governance frameworks encompassing industry standards, ethical guidelines, and possibly regulation to guide the responsible deployment of advanced AI in society. Stakeholders from AI developers to policymakers to end-users will need to collaborate in setting the norms for how GPT-5 is used (and not used) in sensitive domains.

Herbert Wanga, “Prospective Advancements in Conversational AI: A Projection of Natural Language Understanding with GPT-5” *International Research Journal of Advanced Engineering and Science*, Volume 10, Issue 4, pp. 48-53, 2025.

Future work should focus on empirically validating the proposed collaborative frameworks like the Two-Model Method and developing standardized benchmarks for assessing not just the performance, but also the efficiency, fairness, and verifiability of models like GPT-5.

Ultimately, GPT-5's legacy will not be measured solely by its performance on benchmarks, but by how well we harness its power to augment human intelligence and address societal needs, while safeguarding against its risks. With careful stewardship that emphasizes transparency, human-centric design, and sustainability, GPT-5 could indeed mark a leap forward that benefits society. Otherwise, without these safeguards, even its most impressive capabilities might fall short of their promise.

REFERENCES

[1] Ahmad, Z., Ashfaque, M., & Khan, I. A. (2025). Generative AIs and the hidden cost of intelligence: A multidisciplinary review of LLM progress and sustainability. *Annual Methodological Archive Research Review, 3*(9), 84–88.

[2] Askitas, N. (2025). Notes on a world with generative AI. *CESifo Working Paper No. 12070*. Retrieved from SSRN: https://ssrn.com/abstract=5402179

[3] Bender, E. M., & Koller, A. (2020). Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 5185–5198). https://doi.org/10.18653/v1/2020.acl-main.463

[4] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Amodei, D. (2020). *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems, 33*, 1877–1901. Retrieved from: https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64a-Abstract.html

[5] Fedus, W., Zoph, B., & Shazeer, N. (2022). Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research, 23*(120), 1–39. https://jmlr.org/papers/v23/21-0998.html

[6] Ghasemloo, M., & Moradi, A. (2025). Forecasting the future: Has GPT-5 improved time series forecasting accuracy over GPT-4? *Preprint*. (Available on arXiv).

[7] Liu, J., Wang, Y., & Chang, S. (2024). *The hitchhiker's guide to long-context language models*. arXiv preprint. https://arxiv.org/abs/2405.19588

[8] Odelami, B. (2025, September 9). What the release of ChatGPT-5 means: A leap forward or a step too soon? *Information Matters*. Retrieved from: https://informationmatters.org/2025/09/what-the-release-of-chatgpt5-means-a-leap-forward-or-a-step-too-soon/

[9] OpenAI. (2023). *GPT-4 Technical Report*. Retrieved from OpenAI: https://cdn.openai.com/papers/gpt-4.pdf

[10] Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2019). Green AI. *Communications of the ACM, 63*(12), 54–63.

[11] SPRL Development Team. (2025). *Structured Path Resonance Loop (SPRL): A native collaborative language interface for large reasoning models*. Zenodo. https://doi.org/10.5281/zenodo.17041758

[12] Wei, J., Wang, X., Schuurmans, D., Bosma, M., Chi, E., Le, Q., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems, 35*, 24824–24837. Retrieved from: https://proceedings.neurips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html

[13] Youvan, D. C. (2025). *The two-model method: GPT-5 for seeded depth, Gemini for accessible breadth*. youvan.ai White Paper. Retrieved from: https://youvan.ai