

COMMENTARY

Challenges and limitations of ChatGPT and other large language models

Erwin L. Rimban

College of Allied Health Sciences, Faculty of Medical Laboratory Science, Cagayan State University, Tuguegarao City, Cagayan, Philippines



Correspondence to: Erwin L. Rimban, College of Allied Health Sciences, Faculty of Medical Laboratory Science, Cagayan State University, Tuguegarao City, Cagayan, Philippines;
E-mail: erwinrimban@csu.edu.ph

Received: April 7, 2023;

Accepted: June 16, 2023;

Published: June 21, 2023.

Citation: Rimban, E. L. (2023). Challenges and limitations of ChatGPT and other large language models. *International Journal of Arts and Humanities*, 4(1): 147-152. <https://doi.org/10.25082/IJAH.2023.01.003>

Copyright: © 2023 Erwin L. Rimban. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by-nc/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.



Abstract: This article explores the challenges and limitations of large language models, focusing on ChatGPT as a representative example. We begin by discussing the potential benefits of large language models, such as their ability to generate natural language text and assist with language-related tasks. However, we also acknowledge the concerns around these models, including their environmental impact, potential for bias, and lack of interpretability. We then delve into specific challenges faced by ChatGPT and similar models, including limitations in their understanding of context, difficulty in handling rare or out-of-vocabulary words, and their tendency to generate nonsensical or offensive text. We conclude with recommendations for future research and development, including the need for increased transparency, interpretability, and ethical considerations in the creation and deployment of large language models.

Keywords: ChatGPT, large language models, natural language generation, ethical considerations

1 Introduction

Large language models, such as ChatGPT, have emerged as a significant breakthrough in natural language processing (NLP) and machine learning. These models have demonstrated remarkable capabilities in generating coherent and grammatically correct sentences, answering questions, summarizing text, and even carrying out complex conversations. The success of these models has led to their widespread adoption in various applications, ranging from chatbots to language translation to content generation. However, despite their impressive performance, large language models face several challenges and limitations that must be carefully considered before their deployment. This paper aims to explore these challenges and limitations in detail, highlighting their implications for the development and use of large language models.

2 Literature review

Large language models (LLMs) such as ChatGPT, OpenAI's GPT series, and Google's BERT have shown significant advancements in natural language processing (NLP) tasks, including language modeling, machine translation, and question answering. These LLMs have achieved state-of-the-art results on various benchmarks, demonstrating their effectiveness in capturing the complexity of language and handling diverse linguistic phenomena (Devlin et al., 2019; Brown et al., 2020). However, these models have also faced criticism for their ethical implications, limitations, and challenges.

One significant criticism of LLMs is their reliance on large amounts of training data, which can lead to biased representations of language (Bolukbasi et al., 2016). Moreover, the use of pre-trained models such as GPT-3 for downstream tasks can perpetuate the biases present in the data on which the model is trained, making it difficult to achieve fair and unbiased results (Hovy et al., 2021). To mitigate this issue, researchers have proposed various techniques, such as data augmentation, debiasing techniques, and fairness constraints, to improve the fairness and transparency of LLMs (Korhonen et al., 2021).

Another limitation of LLMs is their computational requirements, which can be prohibitively expensive and inaccessible for many researchers and organizations. Training large models like GPT-3 requires massive computational resources, such as high-end GPUs and specialized hardware, leading to significant energy consumption and environmental impact (Strubell et al., 2019). Additionally, the size and complexity of these models can lead to challenges in deploying them in real-world applications, particularly in low-resource settings (Peters et al., 2018).

Moreover, LLMs have faced criticism for their lack of interpretability and explainability, making it challenging to understand how they generate their output and the underlying decision-making process. As LLMs are trained on massive amounts of data, they can capture complex

patterns and nuances in language that are difficult to interpret, leading to black-box models (Lipton, 2018). This issue raises concerns regarding the transparency and accountability of these models in decision-making processes.

Finally, LLMs have also faced criticism for their limitations in handling certain types of linguistic phenomena, such as sarcasm, irony, and humor (Liu et al., 2020). While these models have shown remarkable progress in various NLP tasks, they still struggle to understand the social and cultural context that underlies language use. This issue highlights the need for more research and development to improve the capacity of LLMs to capture the complexity and nuance of language use in diverse contexts.

Consequently, while LLMs have shown remarkable progress in various NLP tasks, they still face significant challenges and limitations. The ethical implications of these models, their reliance on massive amounts of data and computational resources, their lack of interpretability and explainability, and their limitations in handling certain linguistic phenomena raise concerns regarding their widespread adoption and deployment in real-world applications. Thus, it is essential to address these challenges and limitations to ensure the responsible development and deployment of LLMs in diverse settings.

3 Challenges and limitations of large language models

In this section, we discuss the challenges and limitations associated with large language models. These challenges and limitations can be broadly categorized into three categories: technical, ethical, and societal.

3.1 Technical challenges and limitations

3.1.1 Computational Requirements

The training and deployment of large language models require significant computational resources, including high-end computing hardware and data center infrastructure. The cost of these resources can be a significant barrier to entry for many researchers and organizations.

3.1.2 Model efficiency

Large language models are often criticized for their inefficiency in terms of compute and energy usage. The environmental impact of training and running these models at scale is also a concern.

3.1.3 Data bias and quality

Large language models are only as good as the data they are trained on. Data bias and low data quality can limit the effectiveness of these models, particularly in domains where training data is scarce.

3.2 Ethical challenges and limitations

3.2.1 Privacy and security

Large language models can potentially pose a risk to privacy and security, as they can be trained on sensitive information and generate realistic-looking fake text.

3.2.2 Fairness and bias

Large language models can perpetuate and amplify societal biases and discrimination if they are not trained on diverse and representative data.

3.2.3 Misinformation and disinformation

Large language models can be used to generate realistic-looking fake text, which can be used to spread misinformation and disinformation.

3.3 Societal challenges and limitations

3.3.1 Economic and job disruption

Large language models have the potential to disrupt industries and job markets, particularly in areas such as content creation and translation.

3.3.2 Digital divide

The cost and technical requirements associated with large language models can create a digital divide between those who have access to these resources and those who do not.

3.3.3 Legal and regulatory frameworks

The rapid advancement of large language models has outpaced the development of legal and regulatory frameworks to govern their use.

4 Challenges and limitations of large language models

While large language models like ChatGPT have shown impressive capabilities in various NLP tasks, they also face several challenges and limitations that can affect their performance and practical usability. In this section, we will discuss some of the key challenges and limitations of large language models.

4.1 Biases and fairness

One of the most significant challenges facing large language models is the issue of biases and fairness. Since these models are trained on large datasets of human-generated text, they are prone to learning and reproducing the biases present in that data. This can lead to biased and unfair outcomes in applications that rely on these models, such as hiring algorithms or language translation tools. Addressing these biases and ensuring fairness in large language models is an active area of research and requires careful consideration and mitigation strategies.

4.2 Training data requirements

Training large language models requires vast amounts of high-quality text data, which can be difficult and expensive to obtain. In addition, the quality of the training data can significantly impact the performance of the resulting model. For example, if the training data is noisy or contains errors, this can negatively affect the accuracy of the model. Addressing these challenges requires careful selection and curation of training data, as well as the development of methods for cleaning and filtering noisy data.

4.3 Computational resources

Training and deploying large language models require significant computational resources, including high-performance computing clusters and specialized hardware like GPUs. This can make these models inaccessible to researchers and organizations without the necessary resources, limiting their practical usability. Addressing these limitations requires developing more efficient training algorithms and hardware, as well as finding ways to distribute the computational workload across multiple machines.

4.4 Generalization and transferability

While large language models like ChatGPT have shown impressive capabilities in several NLP tasks, they can still struggle with generalizing to new tasks or domains. This limits their practical usability and requires developing more robust and transferable models that can adapt to new tasks and domains. Additionally, the performance of large language models can degrade when applied to tasks that are outside of their original training distribution, highlighting the need for more diverse and representative training data.

4.5 Interpretability

Large language models like ChatGPT are often referred to as black boxes because it can be challenging to understand how they arrive at their predictions or outputs. This lack of interpretability can make it difficult to diagnose and address errors or biases in the model, limiting their practical usability in sensitive applications like healthcare or legal decision-making. Addressing this limitation requires developing methods for interpreting and visualizing the inner workings of these models, as well as developing techniques for explaining their outputs to non-experts.

In summary, large language models like ChatGPT offer tremendous potential for improving NLP applications, but they also face several challenges and limitations that must be addressed to ensure their practical usability and fairness. These challenges require a multidisciplinary approach that combines expertise from computer science, linguistics, and social sciences to develop more robust and interpretable models that can address real-world problems.

5 Future directions and conclusion

Large language models such as ChatGPT have demonstrated remarkable abilities in natural language processing tasks, but their challenges and limitations are still significant. In this section, we discuss potential future directions for research and conclude the article.

One possible direction for future research is to explore methods for improving the robustness and fairness of large language models. As discussed earlier, these models can be prone to bias, which can have real-world consequences. Researchers may need to develop new methods for detecting and mitigating bias, as well as new ways of evaluating model fairness.

Another direction for research is to develop more efficient and sustainable training methods

for large language models. Currently, the energy consumption required to train these models is substantial, which has raised concerns about their environmental impact. New approaches, such as using more efficient hardware or developing more targeted training techniques, may help to reduce the energy requirements of these models.

Finally, it will be important to continue investigating the social and ethical implications of large language models. As these models become more advanced and widely used, they have the potential to impact society in significant ways. Researchers will need to consider the ethical implications of how these models are developed and used, as well as the impact they may have on individual privacy and the distribution of power.

Thus, large language models such as ChatGPT have shown tremendous potential in natural language processing, but also pose significant challenges and limitations. Addressing these issues will require ongoing research and collaboration between experts in various fields. By addressing these challenges head-on, we can continue to advance the state of the art in natural language processing while also ensuring that these technologies are used in responsible and ethical ways.

6 Ethical implications

The development and use of large language models like ChatGPT have raised several ethical concerns. One of the most significant concerns is the potential for the propagation of biases and stereotypes. Language models are trained on vast amounts of data, much of which may contain implicit biases and stereotypes. If these biases and stereotypes are not identified and corrected, they can be perpetuated and even amplified by the model.

Another ethical concern is the potential for large language models to be used to generate fake content, such as fake news, deepfakes, and propaganda. These technologies can be used to spread disinformation and influence public opinion, potentially leading to serious harm.

Additionally, there are concerns about the impact of large language models on employment. These models can automate many tasks traditionally done by humans, which could lead to significant job displacement in certain industries.

Finally, the sheer size and energy consumption required to train and run large language models have significant environmental implications. The carbon footprint of training a single model can be substantial, and as these models become more prevalent, their impact on the environment will become increasingly significant.

In light of these ethical concerns, it is important for developers and users of large language models to consider the potential risks and take steps to mitigate them. This could include developing and enforcing ethical guidelines for the use of these models, ensuring that they are trained on diverse and unbiased data, and developing strategies to address potential job displacement. It is also important to consider the broader societal implications of large language models and ensure that their development and use align with broader societal values and goals.

As the use of large language models like ChatGPT continues to grow, it is essential that we address these ethical concerns and work to develop models that are transparent, unbiased, and aligned with societal goals. Only then can we realize the full potential of these powerful technologies while minimizing their risks and ensuring that they serve the common good.

In light of the challenges and limitations outlined in the previous sections, it is important to carefully consider the future direction of large language models like ChatGPT. While there is no doubt that these models have revolutionized the field of natural language processing and opened up new possibilities for text generation and analysis, it is also clear that there are significant ethical and technical concerns that must be addressed.

First, with regards to ethical concerns, large language models like ChatGPT have raised serious questions about bias and fairness in artificial intelligence. The language models are only as good as the data they are trained on, and if that data contains biases or reflects societal inequalities, the resulting model will inevitably reproduce and potentially amplify those biases. This can have serious consequences, such as perpetuating stereotypes or contributing to discrimination. It is therefore essential that researchers develop strategies to mitigate bias and ensure fairness in the development and deployment of large language models.

Second, technical concerns also need to be addressed. One of the main limitations of large language models like ChatGPT is their inability to understand context and reasoning. While they can generate impressive responses based on statistical patterns in the data they are trained on, they lack the ability to truly comprehend language in the way that humans do. This has implications for the accuracy and reliability of the models, particularly in applications like chatbots where users may have specific expectations or needs that the model cannot fully understand.

Despite these challenges and limitations, there is reason to be optimistic about the future

of large language models. As research in natural language processing continues to advance, it is likely that these models will become more sophisticated and better equipped to handle complex language tasks. Moreover, the development of strategies to address ethical and technical concerns will be crucial to ensuring that these models are used in ways that benefit society as a whole.

Furthermore, while large language models like ChatGPT hold tremendous potential for the future of natural language processing, it is important to carefully consider the ethical and technical challenges and limitations that come with their use. By addressing these concerns and working towards the development of more advanced and sophisticated models, we can ensure that large language models continue to push the boundaries of what is possible in natural language processing while also contributing to a more just and equitable society.

7 Conclusion

In conclusion, large language models such as ChatGPT have revolutionized natural language processing and transformed the field of artificial intelligence. The development of such models has led to unprecedented advances in a range of NLP tasks, from text generation to question answering to machine translation, to name a few. However, there are also significant challenges and limitations associated with these models, which must be taken seriously and addressed by the research community.

The ethical implications of large language models have also come under intense scrutiny in recent years, as concerns have been raised about the potential misuse of such models for nefarious purposes. Researchers must be cognizant of these ethical concerns and work to ensure that their work does not contribute to any harm. It is also important to engage with broader society, including policymakers, to help inform and guide the responsible development and use of large language models.

In the future, it will be essential to continue to push the boundaries of what is possible with large language models while also addressing their limitations and challenges. This will require interdisciplinary collaboration between researchers from different fields, including computer science, linguistics, and psychology, among others. Through such collaborations, we can hope to create more robust and interpretable models that can benefit society in a responsible and ethical manner.

Moving forward, the challenges and limitations of large language models must be acknowledged, but they should not overshadow the tremendous potential of these models to transform the way we interact with language and to improve our lives in myriad ways. It is up to us as a research community to ensure that the benefits of these models are realized while also mitigating any potential harms.

8 Recommendations

Despite the many challenges and limitations facing large language models like ChatGPT, their potential for revolutionizing language-related tasks cannot be denied. However, in order to maximize their usefulness while minimizing their risks, several recommendations should be considered:

(1) Prioritize diversity and inclusion in dataset creation: Given the potential for large language models to perpetuate biases, it is important to ensure that datasets used to train these models are diverse and representative. This includes ensuring that datasets include a wide range of voices and perspectives and avoiding reliance on biased or stereotypical sources.

(2) Encourage transparency and accountability: To build trust in large language models, it is important to encourage transparency and accountability in their development and deployment. This includes providing detailed documentation of how models are trained and evaluated, as well as clear explanations of any biases or limitations.

(3) Invest in research on the societal impacts of large language models: As large language models continue to become more prevalent in various fields, it is crucial to conduct research on their broader societal impacts, including their effects on employment, social inequality, and human interactions.

(4) Consider ethical and legal implications: As large language models continue to be integrated into various applications, it is important to consider the ethical and legal implications of their use. This includes addressing issues such as privacy, bias, and accountability.

(5) Foster collaboration between industry and academia: Given the complexity and rapid pace of development in this field, it is important to foster collaboration between industry and academia in order to address these challenges and limitations, and to ensure that the potential benefits of large language models are realized while minimizing their risks.

By implementing these recommendations, the field can move towards developing large language models that are more inclusive, transparent, and socially responsible, and ultimately maximize their potential to revolutionize language-related tasks while minimizing their risks.

References

- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., & Amodei, D. (2020). Language models are few-shot learners. arXiv preprint, arXiv: 2005.14165. <https://doi.org/10.48550/arXiv.2005.14165>
- Dodge, J., Schwartz, R., Smith, N. A., & Etzioni, O. (2020). Fine-tuning language models from human preferences. arXiv preprint, arXiv: 1909.08593.
- Gebu, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumeé III, H., & Crawford, K. (2018). Datasheets for datasets. arXiv preprint, arXiv: 1803.09010. <https://doi.org/10.1145/3458723>
- Huang, J. (2021). Limitations of language models. In *Advances in Neural Information Processing Systems* (pp. 10141-10150).
- Jia, R., Liang, P. P., Levin, D., & Fei-Fei, L. (2021). Revisiting the evaluation of generative models. arXiv preprint, arXiv: 2101.08127.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. *OpenAI blog*, 1(8), 9.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., & Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1), 5485-5551.
- Roberts, A., Shin, D., Dong, L., & Wainwright, M. J. (2020). A theoretical analysis of contrastive unsupervised representation learning. arXiv preprint, arXiv: 2012.09672.
- Wallace, E., & Jagadeesh, V. (2020). Universal adversarial triggers for attacking and analyzing NLP. arXiv preprint, arXiv: 1908.07125. <https://doi.org/10.18653/v1/D19-1221>
- Zhang, S., Yao, L., Sun, A., & Tay, Y. (2021). Attention mechanisms in language processing: A survey. arXiv preprint, arXiv: 2102.00190.