Comprehending and Reducing LLM Hallucinations

Harsh[1] ; Dr. Shobha T[2]
[1] Student, BMS College of Engineering, Bengaluru
[2] Assistant Professor, BMS College of Engineering, Bengaluru

Abstract:- The integration of large language models (LLM) into many artificial intelligence applications shows the best performance in tasks such as text mining, typing, question answering. Despite his success, his LL.M. The biggest concern is the emergence of so-called "hallucinations", especially in text-based systems and Q&As that rely on LL M. These hearings may lead to the spread of misinformation or fraud. This article explains the basics of AI illusions and highlights their importance in AI. Work involves deploying visualizations to a variety of tasks, including machine translation, surveys, interviews, content writing, LLM maps, and visualization questions. Additionally, this article explores potential strategies to reduce negative perceptions in order to increase the overall credibility of the LL.M.

Keywords:- *LLMs, Hallucination, Artificial Intelligence, Hallucination Mitigation, Factualness.*

I. INTRODUCTION

The large field of language models (LLM) includes GPT-3 (21), IntroductGPT (22), FLAN (23), PaLM (24), LLaMA (25), etc. continues to evolve with new developments such as important collaborations. While LL.M is good at many things, he also displays a flaw that affects his self-confidence and self-confidence: skepticism. Citing Berrios and Dening (30), vision is thought to be slightly different from actual perception, the main difference being the lack of evidence. This allows for a nuanced assessment of the connection between perception and perception. In the context of many cognitive concepts that focus on the analysis of human gestures, revealing vision in cognitive skills needs to be done. Hallucinations, defined as the creation of concepts that arise but are inappropriate information or false facts, cause serious problems in important areas such as medicine (8), finance (22) and other sensitive areas of necessity. The question at hand is: Why do large language models (LLMs) gain insight? Factors such as lack of real-world knowledge, bias, or misinformation can push the model to produce positive but uncertain results. The real problem is the incomplete understanding of ideas that leads to abnormal production.

The outcome of daydreaming in this study involves the production of written content, such as text or response, that reveals reality, relationship, and reality but deviates from or is distorted at the pace of the original source interpretation. The real truth (23). Investigating efforts based on large language models (LLM) is important to avoid biases that can influence decision-making strategies and lead to negative outcomes (24).

LL.M. Identification and reduction of visual impairment developed by. Since the launch of ChatGPT in 2022, the world has seen exponential growth in LLM-based operations and tools. Recently, much interest in science and industry has been directed towards exploring side effects of LLM, such as insight. In a previous study (23), elements of auditory processing in functional studies were identified and linked to early development of natural language. Techniques for writing effective guidelines for implementing LLMs, including the use of NLP criteria, human reasoning, and dynamic LLMs, are discussed in (25). Another study (16) investigates voice therapy in which LL.M.s are guided or inspired to correct their vision. In contrast to these trainings, our contribution is to provide a comprehensive review of the vision of the LL.M., the inclusion of various methods and access to their advantages and disadvantages. The main contribution of this article is an in-depth analysis of the field of research available in LL.M. To achieve this goal, we review and categorize relevant studies across various fields and disciplines. We also discuss visualization and reduction methods in LL.M. Evaluate the advantages and disadvantages of these reductions by demonstrating the principles behind the results presented. The final chapter, “Unborn Perspectives,” suggests future directions and raises questions about current interests.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24JUL882

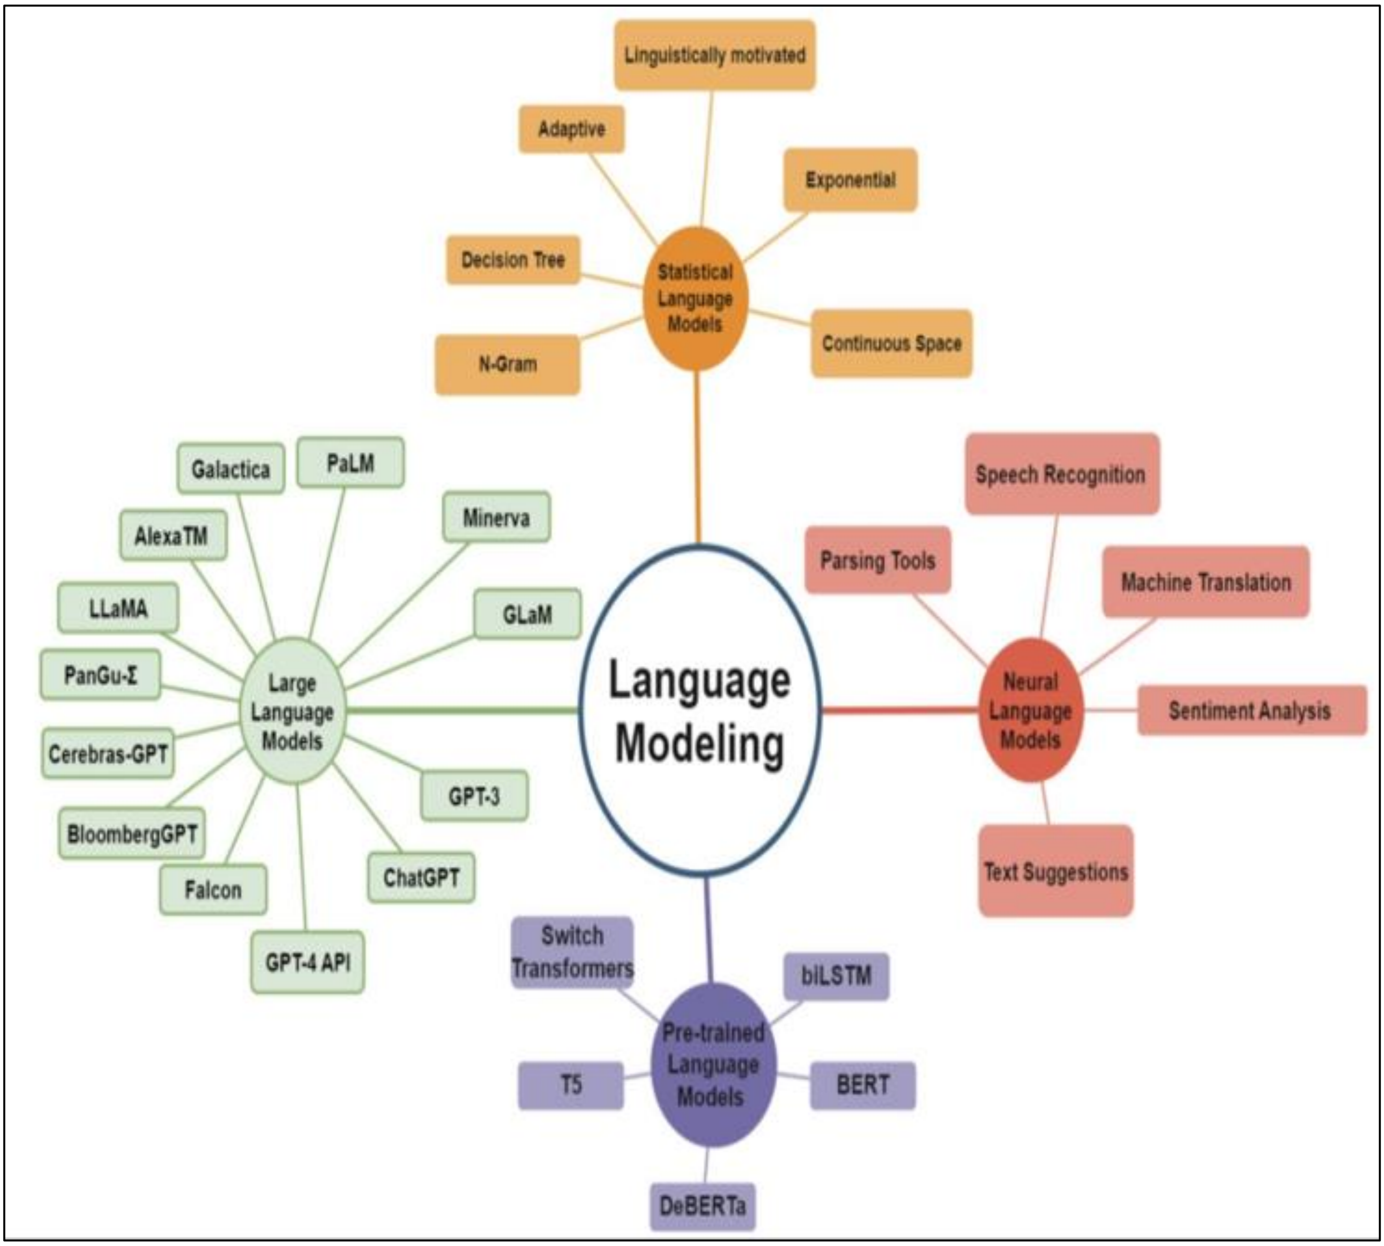

Fig. 1. Types of Language Models (1)

II. RELATED WORK

Surprising opening has been discussed in the context of GPT-4 by Bubeck and colleagues (16 , 3). (3, Runner 82) shows the challenges these visionaries faced due to their risks and impact. Blind eye opener is a comprehensive investigation that involves collected data and tasks beyond the current discussion and is considered to be more difficult. This study shows that it is possible to at least partially resolve the apparent illusion without the need for external resources. The word "illusion" used in this work refers to information that is not based on knowledge. In fact, there are two types of errors: Errors that may arise from errors in knowledge (for example, the misconception that people use only 10% of their brain) and from unlawful behavior. The two types may require different treatments. Less error training programs or the use of methods such as RLHF (17) may help reduce errors. However, illegal crimes, which are the focus of our research, pose a great challenge for the smart and are difficult to solve by developing information that indicates each other. This difference is explained in more detail by Evans et al. (6). Previous studies specifically examining open source concepts similar to ours are limited. Some projects, such as (8), aim to understand what training is most appropriate in a field. In a recent independent study in the field of healthcare, Athaluri et al. (1) Evaluates hearing-related information empirically. Similar to our method, they use Google search with real match strings for sales evaluation. Our auditory processing analysis allows us to predict visual perception for different models, and as discussed in previous studies, the auditory issue will be important since users will give more weight to what they believe if the activation model is correct (16). A recent workshop discussed a black-box technique for trustworthiness testing built on linguistic models (LMs). Although these tutorials focus on real trust, their approach is consistent with our work. For example, Kadavas et al. (10)

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24JUL882

Predictions using LMs can be directly connected to form only LMs, such as ChatGPT using slices. Lin et al. (12) Show that LM can represent the approximation by generating numbers or words that represent the three ways. Finally, Manakul et al. (13) Do a volume check when gathering information. These workshops all used direct research, which directly influenced the design of our research. Due to space limitations, we do not delve into the study of unlimited space illusions (e.g., paraphrasing or summarizing), but instead refer to the discussion of recent work by Ji et al. (9).

Various approaches have been adopted to develop large language models (LLMs), including strategies such as using human feedback or using grammar for optimization (Bakker et al., 2022 ; Ouyang et al., 2022). Ouyang et al. Plan improved LLM-developed content by supporting learning with human feedback. Their recommendations include improving the LL.M. However, it is known that fine-tuning can often lead to poor performance patterns in other tasks (Kirkpatrick et al., 2017). In this study, we take an unbiased approach by assuming that the model cannot be accessed without modification or modification.

Another approach applicable to our context was proposed by Burns et al. (2022) called differential consistency research (CCS). However, CCS needs to turn the statements into questions, evaluate the LLM in two different parts of the statement, and request information from the same data (content language) as an experiment. These limitations LL.M. It makes it impossible to implement the statements made by. Furthermore, while CCS only increases accuracy by 4% over 0-shot LLM queries, our method improves accuracy by almost 20% over 0-shot LLM queries.

III. HALLUCINATIONS DETECTION

Various methods have been proposed to verify accuracy and reliability in large language models (LLMs). Some methods rely on the central process (e.g., recording the result) to define the uncertainty of the written sequence (18), (19). However, external APIs of standards such as ChatGPT do not provide users with access to important information, making this system ineffective in decision-making processes. LLM's fact-checking system can use external repositories and organizations such as Wikipedia (20) to verify visual information. However, there are concerns about the reliability of content on Wikipedia. Azaria and Mitchell (21) proposed a method to evaluate the accuracy of messages using latent representations of LL.M.s for use by multiple layers. The system is based on trained supervision and relies on registration information and the internal state of the LLM, which may not be accessible via API. In their system, LL.Ms are reminded to evaluate the accuracy of their previous guesses, the likelihood of the answer, or the answer indicating that they were correct. Kadawas et al. (22) introduced a nonvisual detection system called sonometry. This study investigates how language models can evaluate the validity of their own responses and predict their accuracy. Larger models that show good predictability for different questions can make predictions in open-ended tasks, estimating the probability of the answer being correct (“P(True)”). They also estimated their confidence in good knowledge (“P(IK)”) along with partial understanding (IK stands for “I know”). Various visual discovery methods have been developed for "zero source" where external data are not available to verify the authenticity of the LL.M. These methods can be divided into gray column methods and black box methods (23). The first assumes knowledge of the internal distribution of the model. The latter is designed for LL.Ms with limited API access and no access to relevant resources. Different techniques are used to eliminate gray box and black box illusions. Pre-graduate training knowledge, which includes training the future language of generic scripts to capture real-world knowledge and contextual context, is required for gray face detection. Figure 2 shows how the query and verification process work together. Varshny et al. (24) developed an improved technique and detected visual defects in GPT-3.5. Their involvement key spotting points, keyword initiation and "guidelines". They use the LLM feature to extract important information from the generated text. Comparison of the three methods shows that the "standard orientation" is better than the truth and initial words in determining the main meaning. They introduced the probability as the minimum value of the probability of the token and improved the method with the question recognition design step based on the answer-aware design model and search the website to answer the authentication questions. This approach achieved an impressive recall of 88 on GPT-3.5.

IV. MITIGATING LLM HALLUCINATIONS

Handling surprises in large language models (LLMs) has become a significant challenge, especially with the worldwide proliferation of LLM-based virtual chatbot agents and question answering systems. Although many methods have been published recently to solve this problem, some of them are only part of the effective use of vaccines because they can cause more blindness in LL.M. Varshny et al. (24) proposed an effective method that could reduce the reflection in GPT3.5 by 33%. The way to check for these artifacts is to provide a model that will correct them in the output. This process involves removing or changing inaccurate information based on the information collected. Although hearing LL.M. has emerged recently, many methods based on different standards have been proposed. These methods can be divided into the following categories:

- Fine-tuning
- Knowledge Graphs
- Memory augmentation
- Context Prompts
- Preemptive Strategies

Figure 3 provides a visual representation of mitigation methods along with their respective pros and cons. Fine-tuning, a widely adopted technique in machine learning to specialize a pre-trained model with a limited dataset [15], has been employed to mitigate hallucinations in Large Language Models (LLMs), as demonstrated by Lee et al. [16]. However, the high parameter count in LLMs, often in the millions, makes fine-tuning a resource-intensive solution. Knowledge graph methods offer the integration of structured and unstructured knowledge, providing LLMs with a broader

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24JUL882

foundation for various tasks [17]. However, the challenge lies in the time-consuming process of designing a well-curated knowledge base and the labor-intensive effort required to maintain up-to-date knowledge. Wu et al. [18] proposed an augmented transformer for knowledge-intensive Natural Language Processing (NLP) tasks to address the need for deep learning methods to expand their capabilities based on new knowledge. Although memory augmentation has benefited NLP models, its applicability to LLMs remains untested. Prompt-based solutions have recently emerged as a means to "de-hallucinate" LLMs. Jha et al. [19] introduced a self-monitoring prompting framework leveraging formal methods to autonomously identify errors in LLM responses. This framework utilizes the conversational abilities of LLMs for response alignment with specified correctness criteria through iterative refinement. Luo et al. [10] proposed Self-Familiarity, a method challenging existing State-of-the-Art (SOTA) techniques by introducing a zero-resource, pre-detection approach to mitigate the risk of LLMs producing inaccurate information. This method extracts and processes conceptual entities from the instruction, employing prompt engineering to derive a familiarity score for each concept.

Low instruction-level familiarity scores indicate a higher likelihood of the LLM generating erroneous information, prompting it to refrain from generating a response. Feldman et al. [11] developed a method based on context-tagged prompts. They formulated a set of questions and created context prompts to assist LLMs in providing more accurate answers. Validation of the context prompts and questions ensured their intended functionality. Experiments with various GPT models were conducted to evaluate the impact of context prompts on the accuracy of LLM responses.

V. FUTURE PERSPECTIVE

This section outlines considerations regarding Large Language Models' (LLMs) hallucinations and mitigation strategies. Current developments in zero-resource hallucination detection are in their nascent stages, suggesting potential avenues for future exploration to enhance the accuracy and reliability of these techniques across a broader spectrum of scenarios. Black-box hallucination detection poses additional challenges due to the absence of access to the LLM's internal states. Future research in this area could focus on devising novel black-box hallucination detection methods or optimizing existing approaches for greater effectiveness. Another aspect to explore is hallucination detection tailored for specific tasks. While current techniques are generally applicable, task-specific customization may yield more effective results. For example, designing hallucination detection methods for factual question answering could leverage the understanding that factually accurate responses are more likely to be grounded in real-world knowledge. Multimodal LLMs, a novel category capable of handling text, images, and other media types, present a unique challenge for hallucination detection. Despite the complexity, addressing hallucination detection in multimodal LLMs is crucial due to their increasing popularity.

VI. CONCLUSION

In summary, this paper provides a comprehensive review of the phenomena observed in large language models (LLMs). It classifies different types of hearing and explores their root causes, which arise from limitations in knowledge, models and inference methods. The authors recommend several mitigation strategies, including improving data quality, improving the design model, and incorporating a robust verification process. They also highlight the need to develop reliable evaluation measures to assess the effectiveness of these strategies. Through a combination of theoretical insights and empirical experiments, this article demonstrates the potential of advanced technologies to reduce LLM thinking and thus make them more reliable and effective outcomes.

REFERENCES

[1]. V. Raunak, A. Menezes, M. Junczys-Dowmunt, The curious case of hallucinations in neural machine translation, in: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, Online, 2021, pp. 1172–1183. URL: https://aclanthology.org/2021.naacl-main.92. doi:10.18653/v1/2021.naacl-main.92.

[2]. N. M. Guerreiro, D. Alves, J. Waldendorf, B. Haddow, A. Birch, P. Colombo, A. Martins, Hallucinations in large multilingual translation models, ArXiv abs/2303.16104 (2023). URL: https://api.semanticscholar.org/CorpusID:257771892.

[3]. D. Dale, E. Voita, J. Lam, P. Hansanti, C. Ropers, E. Kalbassi, C. Gao, L. Barrault, M. R. Costa-jussà, Halomi: A manually annotated benchmark for multilingual hallucination and omission detection in machine translation, ArXiv abs/2305.11746 (2023). URL: https://api.semanticscholar.org/CorpusID:258823059.

[4]. J. Pfeiffer, F. Piccinno, M. Nicosia, X. Wang, M. Reid, S. Ruder, mmt5: Modular multilingual pre-training solves source language hallucinations, ArXiv abs/2305.14224 (2023). URL: https://api.semanticscholar.org/CorpusID:258841429.

[5]. S. Lin, J. Hilton, O. Evans, TruthfulQA: Measuring how models mimic human falsehoods, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 3214–3252. URL: https://aclanthology.org/2022.acl-long.229. doi:10.18653/v1/2022. acl-long.229.

[6]. L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. P. Xing, H. Zhang, J. Gonzalez, I. C. Stoica, Judging llm-as-a-judge with mt-bench and chatbot arena, ArXiv abs/2306.05685 (2023). URL: https://api.semanticscholar.org/CorpusID:259129398.

[7]. V. Adlakha, P. BehnamGhader, X. H. Lu, N. Meade, S. Reddy, Evaluating correctness and faithfulness of instruction-following models for question answering, ArXiv abs/2307.16877 (2023). URL: https://api.semanticscholar.org/CorpusID:260334056.

[8]. L. K. Umapathi, A. Pal, M. Sankarasubbu, Med-halt: Medical domain hallucination test for large language models, ArXiv abs/2307.15343 (2023). URL: https://api.semanticscholar.org/ CorpusID:260316324.

[9]. N. Dziri, S. Milton, M. Yu, O. Zaiane, S. Reddy, On the origin of hallucinations in conversational models: Is it the datasets or the models?, in: Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, Seattle, United States, 2022, pp. 5271–5285. URL: https://aclanthology.org/2022.naacl-main.387. doi:10.18653/v1/2022.naacl-main.387.

[10]. S. Das, S. Saha, R. Srihari, Diving deep into modes of fact hallucinations in dialogue systems, in: Findings of the Association for Computational Linguistics: EMNLP 2022, Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 2022, pp. 684–699. URL: https://aclanthology.org/2022.findings-emnlp.48. doi:10.18653/v1/2022. findings-emnlp.48

[11]. N. M. Guerreiro, D. Alves, J. Waldendorf, B. Haddow, A. Birch, P. Colombo, A. Martins, Hallucinations in large multilingual translation models, ArXiv abs/2303.16104 (2023). URL: https://api.semanticscholar.org/CorpusID:257771892.

[12]. D. Dale, E. Voita, J. Lam, P. Hansanti, C. Ropers, E. Kalbassi, C. Gao, L. Barrault, M. R. Costa-jussà, Halomi: A manually annotated benchmark for multilingual hallucination and omission detection in machine translation, ArXiv abs/2305.11746 (2023). URL: https://api.semanticscholar.org/CorpusID:258823059.

[13]. J. Pfeiffer, F. Piccinno, M. Nicosia, X. Wang, M. Reid, S. Ruder, mmt5: Modular multilingual pre-training solves source language hallucinations, ArXiv abs/2305.14224 (2023). URL: https://api.semanticscholar.org/CorpusID:258841429.

[14]. S. Lin, J. Hilton, O. Evans, TruthfulQA: Measuring how models mimic human falsehoods, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 3214–3252. URL: https://aclanthology.org/2022.acl-long.229. doi:10.18653/v1/2022. acl-long.229.

[15]. L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. P. Xing, H. Zhang, J. Gonzalez, I. C. Stoica, Judging llm-as-a-judge with mt-bench and chatbot arena, ArXiv abs/2306.05685 (2023). URL: https://api.semanticscholar.org/CorpusID:259129398.

[16]. V. Adlakha, P. BehnamGhader, X. H. Lu, N. Meade, S. Reddy, Evaluating correctness and faithfulness of instruction-following models for question answering, ArXiv abs/2307.16877 (2023). URL: https://api.semanticscholar.org/CorpusID:260334056.

[17]. L. K. Umapathi, A. Pal, M. Sankarasubbu, Med-halt: Medical domain hallucination test for large language models, ArXiv abs/2307.15343 (2023). URL: https://api.semanticscholar.org/ CorpusID:260316324.

[18]. N. Dziri, S. Milton, M. Yu, O. Zaiane, S. Reddy, On the origin of hallucinations in conversational models: Is it the datasets or the models?, in: Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, Seattle, United States, 2022, pp. 5271–5285. URL: https://aclanthology.org/2022.naacl-main.387. doi:10.18653/v1/2022.naacl-main.387.

[19]. S. Das, S. Saha, R. Srihari, Diving deep into modes of fact hallucinations in dialogue systems, in: Findings of the Association for Computational Linguistics: EMNLP 2022, Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 2022, pp. 684–699. URL: https://aclanthology.org/2022.findings-emnlp.48. doi:10.18653/v1/2022. findings-emnlp.48.

[20]. N. Dziri, E. Kamalloo, S. Milton, O. Zaiane, M. Yu, E. M. Ponti, S. Reddy, FaithDial: A Faithful Benchmark for Information-Seeking Dialogue, Transactions of the Association for Computational Linguistics 10 (2022) 1473–1490. URL: https://doi.org/10.1162/tacl_a_00529. doi:10.1162/tacl_a_00529.

[21]. N. Dziri, H. Rashkin, T. Linzen, D. Reitter, Evaluating attribution in dialogue systems: The begin benchmark, Transactions of the Association for Computational Linguistics 10 (2021) 1066–1083. URL: https://api.semanticscholar.org/CorpusID:233481654.

[22]. W. Sun, Z. Shi, S. Gao, P. Ren, M. de Rijke, Z. Ren, Contrastive learning reduces hallucination in conversations, Proceedings of the AAAI Conference on Artificial Intelligence 37 (2023) 13618–13626. URL: https://ojs.aaai.org/index.php/AAAI/article/view/26596 . doi:10.1609/ aaai.v37i11.26596.

[23]. D. Tam, A. Mascarenhas, S. Zhang, S. Kwan, M. Bansal, C. Raffel, Evaluating the factual consistency of large language models through news summarization, in: Findings of the Association for Computational Linguistics: ACL 2023, Association for Computational Linguistics, Toronto, Canada, 2023, pp. 5220–5255. URL: https://aclanthology.org/2023. findings-acl.322. doi:10.18653/v1/2023.findings-acl.322.

[24]. M. Cao, Y. Dong, J. Cheung, Hallucinated but factual! inspecting the factuality of hallucinations in abstractive summarization, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 3340–3354. URL: https://aclanthology.org/2022.acl-long.236. doi:10.18653/v1/2022.acl-long.236.

[25]. J. Shen, J. Liu, D. Finnie, N. Rahmati, M. Bendersky, M. Najork, "why is this misleading?": Detecting news headline hallucinations with explanations, in: Proceedings of the ACM Web Conference 2023, WWW '23, Association for Computing Machinery, New York, NY, USA, 2023, p. 1662–1672. URL: https://doi.org/10.1145/3543507.3583375. doi:10.1145/ 3543507.3583375.

[26]. Y. Qiu, Y. Ziser, A. Korhonen, E. Ponti, S. B. Cohen, Detecting and mitigating hallucinations in multilingual summarisation, ArXiv abs/2305.13632 (2023). URL: https://api.semanticscholar.org/CorpusID:258841008.

[27]. J. Yu, X. Wang, S. Tu, S. Cao, D. Zhang-li, X. Lv, H. Peng, Z. Yao, X. Zhang, H. Li, C. yan Li, Z. Zhang, Y. Bai, Y.-T. Liu, A. Xin, N. Lin, K. Yun, L. Gong, J. Chen, Z. Wu, Y. P. Qi, W. Li, Y. Guan, K. Zeng, J. Qi, H. Jin, J. Liu, Y. Gu, Y. Gu, Y. Yao, N. Ding, L. Hou, Z. Liu, B. Xu, J. Tang, J. Li, Kola: Carefully benchmarking world knowledge of large language models, ArXiv abs/2306.09296 (2023). URL: https://api.semanticscholar.org/CorpusID:259165244.

[28]. N. Mihindukulasooriya, S. M. Tiwari, C. F. Enguix, K. Lata, Text2kgbench: A benchmark for ontology-driven knowledge graph generation from text, ArXiv abs/2308.02357 (2023). URL: https://api.semanticscholar.org/CorpusID:260611736.

[29]. Y. Li, Y. Du, K. Zhou, J. Wang, W. X. Zhao, J. rong Wen, Evaluating object hallucination in large vision-language models, ArXiv abs/2305.10355 (2023). URL: https: //api.semanticscholar.org/CorpusID:258740697.

[30]. T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry

[31]. L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. F. Christiano, J. Leike, R. Lowe, Training language models to follow instructions with human feedback, in: S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Eds.), Advances in Neural Information Processing Systems, volume 35, Curran Associates, Inc., 2022, pp. 27730–27744. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/ b1efde53be364a73914f58805a001731-Paper-Conference.pd f.

[32]. J. Wei, M. Bosma, V. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, Q. V. Le, Finetuned language models are zero-shot learners, in: International Conference on Learning Representations, 2022. URL: https://openreview.net/forum?id=gEZrGCozdqR.

[33]. A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, P. Schuh, K. Shi, S. Tsvyashchenko, J. Maynez, A. Rao, P. Barnes, Y. Tay, N. M. Shazeer, V. Prabhakaran, E. Reif, N. Du, B. C. Hutchinson, R. Pope, J. Bradbury, J. Austin, M. Isard, G. Gur-Ari, P. Yin, T. Duke, A. Levskaya, S. Ghemawat, S. Dev, H. Michalewski, X. García, V. Misra, K. Robinson, L. Fedus, D. Zhou, D. Ippolito, D. Luan, H. Lim, B. Zoph, A. Spiridonov, R. Sepassi, D. Dohan, S. Agrawal, M. Omernick, A. M. Dai, T. S. Pillai, M. Pellat, A. Lewkowycz, E. Moreira, R. Child, O. Polozov, K. Lee, Z. Zhou, X. Wang, B. Saeta, M. Díaz, O. Firat, M. Catasta, J. Wei, K. S. Meier-Hellstern, D. Eck, J. Dean, S. Petrov, N. Fiedel, Palm: Scaling language modeling with pathways, ArXiv abs/2204.02311 (2022). URL: https://api.semanticscholar.org/CorpusID:247951931.

[34]. H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, G. Lamp