

AI Hallucinations: A Misnomer Worth Clarifying

Negar Maleki
University of South Florida
negarmaleki@usf.edu

Balaji Padmanabhan
University of Maryland
bpadmana@umd.edu

Kaushik Dutta
University of South Florida
duttak@usf.edu

Abstract—As large language models continue to advance in Artificial Intelligence (AI), text generation systems have been shown to suffer from a problematic phenomenon often termed as “hallucination.” However, with AI’s increasing presence across various domains, including medicine, concerns have arisen regarding the use of the term itself. In this study, we conducted a systematic review to identify papers defining “AI hallucination” across fourteen databases. We present and analyze definitions obtained from all databases, categorize them based on their applications, and extract key points within each category. Our results highlight a lack of consistency in how the term is used, but also help identify several alternative terms in the literature. We discuss the implications of these findings and call for a more unified effort to bring consistency to this important contemporary AI issue, which can significantly affect multiple domains.

Index Terms—Hallucination, Generative AI, Computer Science, Healthcare, Large Language Models (LLMs)

I. INTRODUCTION

One of the early uses of the term “hallucination” in the field of Artificial Intelligence (AI) was in computer vision, in 2000 [1], where it was associated with constructive implications such as super-resolution [1], image inpainting [2], and image synthesis [3]. Interestingly, in this context hallucination was regarded as a valuable asset in computer vision rather than an issue to be circumvented. For instance, an image with low resolution might have been rendered more useful with careful hallucination [1] that generated additional pixels specifically for this purpose.

Despite this (more positive) beginning, recent research has started to employ the term “hallucination” to describe a specific type of error in image captioning [4] and adversarial attack in object detection [5]. In this context, “hallucination” refers to instances where non-existent objects are erroneously detected or incorrectly localized at their anticipated positions. This latter (more negative) interpretation of “hallucination” in computer vision mirrors its analogous usage in language models. For instance, in 2017, researchers highlighted challenges in language models, such as “*the output of the Neural Machine Translation (NMT) system is often quite fluent but entirely unrelated to the input*” [6], or “*language models presume likelihood, but the generated content is ultimately incorrect and unsupported by any information*” [7], which is interpreted as a form of hallucination in AI.

To date, a precise and universally accepted definition of “hallucination” remains absent in the discussions related to this in the increasingly broader field of AI [8]. Diverse definitions, or implied interpretations, persist; sometimes even contradic-

tory, as previously highlighted within the field of computer vision where multiple, disparate interpretations coexist under the same term.

Beyond the AI context, and specifically in the medical domain, the term “hallucination” is a psychological concept denoting a specific form of sensory experience [9]. Ji et al. [10], from the computer science perspective (in ACM Computing Surveys), rationalized the use of the term “hallucination” as “*an unreal perception that feels real*” by drawing from Blom’s definition — “*a percept, experienced by a waking individual, in the absence of an appropriate stimulus from the extracorporeal world.*” On the other hand, Østergaard et al. [11], from the medical perspective (in Schizophrenia Bulletin, one of the leading journals in the discipline), raised critical concerns regarding even the adoption of the “hallucination” terminology in AI for two primary reasons: 1) The “hallucination” metaphor in AI from this perspective is a misnomer, as AI lacks sensory perceptions, and errors arise from data and prompts rather than the absence of stimuli, and 2) this metaphor is highly stigmatizing, as it associates negative issues in AI with a specific issue in mental illness, particularly schizophrenia, thereby possibly undermining many efforts to reduce stigma in psychiatry and mental health.

Given AI’s increasing presence across various domains, including the medical field, concerns have arisen regarding the multifaceted, possibly inappropriate and potentially even harmful use of the term “hallucination” [11], [12]. To address this issue effectively, two potential paths of work offer some promise: 1) The establishment of a consistent and universally applicable terminologies that can be uniformly adopted across all AI-impacted domains will help, particularly if such terminologies lead to the use of more specific and nuanced terms that actually describe the issues they highlight (as we will show later, such vocabulary does exist, but needs more consistent use) and 2) The formulation of a robust and formal definition of “AI hallucination” within the context of AI. These measures are essential to promote clarity and coherence in discussions and research related to “hallucination” in AI, and to mitigate potential confusion and ambiguity in cross-disciplinary applications.

Motivated by these issues, in this paper, we conduct a systematic review of the use of “AI hallucination” across 14 databases with a focus on identifying various definitions that have been used in the literature so far (our review covers more fields than just healthcare and computer science, including ethical and legal settings, and domains as diverse as physics,

sports, etc. in order to explore any broader issues). Recently, two papers ([10], [13]) explored the concept of hallucination in Natural Language Generation- (NLG-) specific tasks (e.g., text translation, text summarization, knowledge graph, etc.). Our work builds on these studies to also consider the application of NLG in diverse domains. The pervasive nature of AI extends beyond these specific tasks, affecting numerous domains and applications. Consequently, our broader review done here reveals the extensive utilization of Large Language Models (LLMs) across almost a much broader space of domains to date, and provides a comprehensive understanding of how the term has been leveraged across various fields. Generally we see that research attempting to define "AI hallucination" does so based on their individual understanding and the challenges encountered within their respective fields. The findings from our systematic and broad review underscore the challenge that the term "AI hallucination" lacks a precise, universally accepted definition, resulting in the observation of various characteristics associated with this term across different applications. We present a summary of these different interpretations and provide some guidance going forward.

II. METHODOLOGY

Our systematic review covered an extensive database search across various domains, including computer science and health, with a focus on the following databases: PubMed, MEDLINE, Scopus, PubMed Central, Web of Science, BioMed Central, Embase, PLOS, CINAHL, ACM, IEEEExplore, ScienceDirect, Google Scholar, and arXiv (no relevant documents were found in MedlinePlus, Cochrane Library, and UpToDate databases, so those were excluded).

Our search methodology was tailored to adapt to the volume of results as well as the relevance of the papers to our research objectives. We manually reviewed every paper that made it through this process in order to identify possible definitions/usage of the term "hallucination" in AI. Given this goal we had to adapt the search in some cases to identify papers most closely relevant to this objective as noted further below. Also, given differences in how search queries are interpreted across the different databases we had to iteratively modify the search term within each database as well in many cases. For clarity, we present all the specific details of this below (in order to be transparent about how we created the subset of papers from which to examine the definitions). However, the summary of these is provided in Table I, including details of the study period for each database.

In the PubMed Database, we initiated an advanced search employing the keywords "Artificial Intelligence" AND "Hallucination" within the "All field" category, yielding 103 papers within the last 10 years. However, the query "AI+hallucination" yielded only 3 papers. Conversely, within the Scopus database, searches for "AI hallucination" or "AI+hallucination" resulted in a total of 1445 records across all fields over the same 10-year period. To manage this extensive dataset, we refined our search criteria to focus on the Title, Abstract, or Introduction, which reduced the results to 483

relevant records. A detailed review of each abstract led us to download papers that appeared pertinent to AI hallucination. This approach significantly differed from searching within abstracts alone, which produced only 49 records and missed some relevant documents.

In PubMed Central (PMC), the query "AI+hallucination" yielded just 1 paper. Consequently, we conducted an advanced search using the keywords "Artificial Intelligence" AND "Hallucination" within the "Text Word" category, uncovering 371 records from the past decade. PMC does not provide abstracts, necessitating the manual examination of each paper to assess its relevance to AI hallucination.

In both the MEDLINE and Web of Science databases, we employed the term "AI hallucination" within the "All field" category, yielding 157 and 139 records, respectively, spanning the last 10 years. In the case of MEDLINE, each record underwent a thorough review, and records containing definitions for AI hallucination were downloaded. Conversely, Web of Science offered abstracts for the records, enabling us to screen them individually and select those relevant to AI hallucination.

Within the BioMed Central (BMC) database, a search for "AI hallucination" led to the retrieval of 76 papers published within the last 10 years. Subsequently, we accessed each paper individually and downloaded those containing pertinent definitions of AI hallucination. In the Embase database, our search for "AI hallucination" produced 80 records published within the last 10 years. These records underwent meticulous abstract review, and those relevant to AI hallucination were selectively downloaded for further analysis.

In the PLOS database, our initial search with "AI hallucination" resulted in a substantial 1064 records across all fields for the past 10 years. Given this large dataset, we refined our search to focus on the "Body" section, yielding 885 records. We proceeded to review the abstracts of each paper and downloaded those demonstrating relevance to AI hallucination. Within the CINAHL database, we conducted an advanced search utilizing the terms "AI" or "Artificial Intelligence" combined with "Hallucination" within the "All field" category, yielding 34 records published in the past 10 years.

In the ACM and IEEEExplore databases, we performed advanced searches using the terms "AI" AND "hallucination" within the "Full Text" category, resulting in 264 and 257 records, respectively, spanning the past 10 years. Each record underwent individual review, and those containing definitions related to AI hallucination, particularly within the field of LLMs, were downloaded.

In the ScienceDirect database, searches for "AI hallucination" or ("AI" AND "hallucination") yielded an identical number of records, specifically 769 English records, spanning the last decade. Each record underwent meticulous examination, and we selectively downloaded papers containing defined concepts of AI hallucination within the LLMs domain.

Searching for "AI hallucination" on Google Scholar yielded 17,000 records from the last 10 years, rendering a com-

TABLE I
METHOD SUMMARY

Source/ Database	Search Category	Query Terms	Num. of Papers	Study End Date*
PubMed	All Field	"Artificial Intelligence" AND "Hallucination"	103	09/27/2023
MEDLINE	All Field	"AI hallucination"	157	09/28/2023
Scopus	Title, Abstract, or Introduction	"AI+hallucination" OR "AI hallucination"	483	09/27/2023
PubMed Central	Text Word	"Artificial Intelligence" AND "Hallucination"	371	09/28/2023
Web of Science	All Field	"AI hallucination"	139	09/28/2023
BioMed Central	All Field	"AI hallucination"	76	09/29/2023
Embase	All Field	"AI hallucination"	80	09/29/2023
PLOS	Body	"AI hallucination"	885	09/29/2023
CINAHL	All Field	"Hallucination" AND ("AI" OR "Artificial Intelligence")	34	09/29/2023
ACM	Full Text	"AI" AND "hallucination"	264	09/30/2023
IEEEExplore	Full Text	"AI" AND "hallucination"	257	09/30/2023
ScienceDirect	All Field	"AI hallucination" OR ("AI" AND "hallucination")	769	09/30/2023
Google Scholar	All Field	"AI hallucination" AND "hallucination in AI"	89	10/01/2023
arXiv	All Field	"AI" AND "hallucination"	40	10/01/2023

* The start date is the same for all databases: 01/01/2013 (Date format: mm/dd/yyyy).

prehensive review unfeasible. To address this challenge, we employed Google Scholar's advanced search feature, identifying records containing the exact phrases "AI hallucination" and "hallucination in AI," which reduced the results to 89 records from the last decade. Subsequently, we conducted a meticulous screening of these records to identify those providing definitions for AI hallucination.

Similarly, within the arXiv database, we conducted an advanced search using the keywords "AI" AND "hallucination" within the "All field" category, resulting in the retrieval of 40 relevant papers. As with our prior search, we meticulously examined each paper and downloaded those containing definitions for AI hallucination.

The eligibility criteria encompassed any type of published scientific research or preprints, such as articles, reviews, communications, editorials, and opinions, that contained the following search terms: "AI hallucination," "AI" AND "hallucination," "Hallucination in AI," or ("AI" OR "Artificial Intelligence") AND "hallucination" in any part of the document, including the title, abstract, and full text. As explained for each database, we employed the most appropriate search terms.

Initially, our search yielded 3753 records, in total, matching these criteria. However, we refined our search to focus exclusively on records that offered a definition of "AI hallucination" within the context of LLMs. It is essential to clarify that we excluded other types of hallucination, such as face hallucination, auditory voice/verbal hallucination, etc., as they were not the primary focus of this review. Our search involved thorough examination of entire documents, and we collected any documents that indicated the presence of a definition for AI hallucination.

Our exclusion criterion was limited to non-English records. The precise database search strategy encompassed all available documents from January 1st, 2013, to October 1st, 2023. In

total, we identified 333 records that provide a definition either independently or by inference from a referenced paper. The summary of the methodology is provided in Table I, including details of the study period for each database. While our review of this work is one of the broadest to date, we acknowledge limitations that are implicit in the methodology above - particularly ones where we had to reduce the retrieved number of papers to focus on potentially more relevant ones due to the manual nature of our review (i.e. where we individually reviewed each paper to identify how the term was used and extract the relevant definition in the proper context). Therefore, while the definitions presented here are certainly those that were used it is possible we may have missed a few other definitions that may have newer connotations not identified in our work here. All the 333 definitions are provided in the supplementary material.

III. RESULT

We reviewed all retrieved papers and documented the definitions provided in each. One main takeaway was that a formal and consistent definition of hallucination simply does not currently exist. There is also little agreement on the specific characteristics of AI hallucination. Depending on the application, we observe varying characteristics, sometimes even contradictory ones.

For instance, in the context of text translation, Koehn and Knowles [6] described hallucination as "fluent but irrelevant," or Miao et al. [26] characterized it as "fluent but inadequate," while Lee et al. [27] attributed "abnormal and unrelated" characteristics to it, thus illustrating different attributes within the same context. In the text summarization context, hallucination refers to generated content that is inconsistent with the source document [28], [29], with some studies categorizing it into

TABLE II
ALTERNATIVE TERMS USED

Alternative Terms	Definitions	References
Confabulation	AI generated responses that sound plausible but are, in fact, incorrect.	[14]
	Definition was not provided.	[15]
Delusion	AI generated responses that are false.	[16]
Stochastic Parroting	The repetition of training data or its patterns, rather than actual understanding or reasoning.	[17]
	LLM model generates confident, specific, and fluent answers that are factually completely wrong.	[18]
	Definition was not provided.	[19]
Factual Errors	Inaccuracies in information or statements that are not in accordance with reality or the truth, often unintentional but resulting in incorrect or misleading information.	[20]
Fact Fabrication	The occurrence where inaccurate information is invented, not represented in the training dataset, and is presented lucidly.	[21]
Fabrication	The phenomenon where, as a generative AI, ChatGPT generates outputs based on statistical prediction of the text without human-like reasoning, potentially resulting in plausible-sounding but inaccurate responses.	[22]
	The phenomenon in ChatGPT output where the text is cogent but not necessarily true.	[23]
	Definition was not provided.	[24]
Falsification and Fabrication	Definition was not provided.	[12]
Mistakes, Blunders, Falsehoods	Answers that are fabricated when data are insufficient for an accurate response.	[25]
Hasty Generalizations, False Analogy, False Dilemma	AI models making inferences that do not follow from the premises; also "hasty generalizations," i.e., the fallacy of making (too) strong claims based on (too) limited data.	[11]

subtypes: "Intrinsic hallucination" and "Extrinsic hallucination" [30], [31], raising concerns, particularly regarding the latter.

Before the launch of ChatGPT on November 30, 2022, we hardly observed definitions for AI hallucination in fields other than computer science. However, with the advent of ChatGPT, researchers have recognized the urgent need for Large Language Models (LLMs) in various fields, including medicine. Therefore, over time, we have observed that the definition has changed and seems to have become a problem more relevant to ChatGPT, albeit with different characteristics under the same term across various applications.

In recent times, for reasons discussed earlier in this paper as well as broader concern about giving AI "human" characteristics inadvertently by using this term, researchers have made efforts to replace the term 'hallucination,' deeming it unsuitable and advocating for its renaming or for alternatives. We have compiled many of the suggested terms found in the literature in Table II, along with their definitions in the respective papers. This is a start in the right direction perhaps - in the search for specific definitions and specific characteristics that we want to model - but does illustrate the lack of consistency in the literature that we pointed out in this paper.

Based on the alternate terms we found, some "old" problems appear to re-surface: the terms confabulation and delusional for instance have connections to mental health conditions as well. However, fabrication, stochastic parroting and hasty generalization together suggest three viable alternatives. Fact fabrication captures many of the cases previously attributed to 'hallucination' without the negative connotations, while stochastic parroting appears to be an appropriate descriptive term for the reasons behind fact fabrication in Generative AI. While we need clarity in terms of distinguishing between stated facts (in the training data) and inferences, the reference

to hasty generalization does start to capture such a distinction.

Finally, since our focus here was on reviewing AI hallucinations across various applications, we grouped all the final papers examined by category, extracted definitions related to AI hallucination, and used ChatGPT 3.5 [32] to extract key points. The applications included chatbots, dialogue settings, generative AI, academia, health, legal and ethical settings, science, technology, text translation, question and answering, text summarization, and others. As shown in Table III, the extracted summaries share similar characteristics, but highlight different extents of inaccuracy, ranging from "deviating from established knowledge", "factual incorrectness", "fictional" to "nonsensical" – offering further considerations for a robust taxonomy that will be needed to bring out such nuances.

IV. DISCUSSION

"Hallucinate" secured its position as the word of 2023 ([38], [39]) and Dictionary.com noted a 46% surge in searches for the term over the past year. The popular press has also keyed in on this (Table IV highlights topics some recent articles discuss and the meaning of "AI hallucination" they convey). These articles primarily feature interviews with CEOs of big tech companies, who discuss future efforts to prevent "hallucination" in chatbots' outputs. Indeed, preventing such occurrences continues to be a key research goal, but few solutions have emerged so far. In the meanwhile, Generative AI continues to expand its applications into multiple domains, making the need for good solutions vital. As a precursor to even developing solutions, this article calls for more systematic, consistent and semantically nuanced terms that can replace "hallucinations" for the reasons noted here. As one step toward such a call, we presented a short summary from one of the broadest manual literature reviews on this topic to date. Our findings illustrate the current lack of consistency and consensus on this issue, but also bring to light some

TABLE III
KEY POINTS OF "HALLUCINATION" DEFINITIONS WITHIN EACH APPLICATION. THE CHARACTERISTICS OF DEFINITIONS ARE PRESENTED IN **BOLD**, ALTHOUGH THEY MAY BE SIMILAR ACROSS DIFFERENT APPLICATIONS.

Application	Number of Papers	LLM Generated Key Points of Definitions
Chatbot	34	The definitions collectively highlight the central theme of AI-generated content deviating from factual correctness , at times even leading to entirely fictional or erroneous information . In essence, AI hallucination underscores the ongoing challenge of maintaining accuracy and reliability in AI-generated content within the context of chatbot applications.
Dialogue Setting	8	The definitions collectively underscore the challenge of ensuring accuracy and reliability in dialogue systems, given the potential pitfalls associated with generating content that is unsupported, nonsensical, or factually incorrect . These issues are particularly pertinent when deploying large pre-trained language models in dialogue applications, as they struggle with maintaining fidelity to the source material while generating coherent and accurate responses.
Generative AI	50	The definitions collectively emphasize the complexity of ensuring factual accuracy and reliability in AI-generated content within generative AI applications, highlighting the potential pitfalls of deviating from adherence to factual correctness .
Academia	88	A common thread among these definitions is the generation of text or content by AI models that lacks fidelity to factual accuracy, reality, or the intended context .
Health	82	The key idea common to all the definitions is that "AI hallucination" occurs when AI systems generate information that deviates from factual accuracy, context, or established knowledge . In essence, AI hallucination manifests as the production of text that, though potentially plausible, deviates from established facts or knowledge in health applications.
Legal and Ethical Setting	16	The definitions collectively emphasize the multifaceted challenges posed by AI hallucination in the legal and ethical context. They highlight issues of accuracy, confidence, relevance, context, and potential misinformation , underscoring the critical importance of addressing these challenges to ensure the responsible and ethical use of AI systems.
Science	10	Across the definitions, the central theme is that AI hallucination involves the generation of text or information that deviates from factual accuracy, coherence, or faithfulness to the input or source content , with potential consequences for scientific accuracy and integrity.
Technology	8	The definitions reflect the multifaceted nature of AI hallucination in technology applications, encompassing accuracy, unpredictability, credibility, and the balance between reasonableness and correctness .
Text Translation	4	The definitions collectively emphasize the central theme of "AI hallucination" in text translation, which revolves around challenges related to maintaining fidelity, coherence, and relevance in the generated translations to ensure accurate and meaningful output.
Question and Answering	7	"AI hallucination" in question and answer applications raises concerns related to the accuracy, truthfulness, and potential spread of misinformation in AI-generated answers, emphasizing the need for improving the reliability of these systems.
Text Summarization	19	The definitions highlight the multifaceted challenges posed by "AI hallucination" in text summarization, encompassing issues related to fidelity, coherence, factual accuracy, and the preservation of the original meaning in generated summaries.
Others *	7	These diverse applications collectively emphasize the challenge of maintaining accuracy, coherence, and trustworthiness in AI-generated content, highlighting the need for tailored approaches to address domain-specific concerns.

* Including: Investment portfolio, Journalism, Reinforcement Learning, Retail, Sport, and Survey Setting.

TABLE IV
SOME POPULAR PRESS ARTICLES ON AI HALLUCINATION

	What the Press Article Discussed...	The Real Meaning the Press Article Conveys about "AI Hallucination"	Source
1	CNBC provided some examples where ChatGPT generated outputs that sounded correct but weren't actually true, such as a legal brief written by ChatGPT to a Manhattan federal judge	When an AI model "hallucinates," it generates fabricated information in response to a user's prompt, but presents it as if it's factual and correct	CNBC [33]
2	The New York Times asked ChatGPT, Google's Brad, and Microsoft's Bing: When did The New York Times first report on "artificial intelligence"?	Chatbots provide inaccurate answers to questions; although false, the responses appear plausible as they blur and conflate people, events, and ideas	The New York Times [34]
3	The New York Times traced the evolution of the term "hallucination" throughout the newspaper's history	-	The New York Times [35]
4	CNN addressed the major issue of "AI hallucination" and narrated on the responses of OpenAI's and Google's CEOs to the question: Can hallucination be prevented?	AI-powered tools like ChatGPT impress with their ability to provide human-like responses, but a growing concern is their tendency to just make things up	CNN [36]
5	Forbes narrated the history of artificial neural networks, which started around eight decades ago, when researchers sought to replicate the functioning of the brain	"AI hallucination" refers to unrealistic ideas about achieving "artificial general intelligence" (AGI), while understanding of how our brains work is limited	Forbes [37]

recent options that are good alternatives. More work is needed to develop a systematic taxonomy that can be more widely adopted as we discuss these issues in the context of AI applications.

V. ACKNOWLEDGMENT

The sections of the paper that use AI-generated text have a citation to the AI system used to generate the text.

REFERENCES

- [1] S. Baker and T. Kanade, "Hallucinating faces," in *Proceedings Fourth IEEE International Conference on Automatic Face and Gesture Recognition (Cat. No. PR00580)*, 2000, pp. 83–88.
- [2] H. Xiang, Q. Zou, M. A. Nawaz, X. Huang, F. Zhang, and H. Yu, "Deep learning for image inpainting: A survey," *Pattern Recognition*, vol. 134, p. 109046, 2023.
- [3] A. Pumarola, A. Agudo, A. Sanfeliu, and F. Moreno-Noguer, "Unsupervised person image synthesis in arbitrary poses," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 8620–8628.
- [4] A. F. Biten, L. Gómez, and D. Karatzas, "Let there be a clock on the beach: Reducing object hallucination in image captioning," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022, pp. 1381–1390.
- [5] A. Braunegg, A. Chakraborty, M. Krumdick, N. Lape, S. Leary, K. Manville, E. Merkhofer, L. Strickhart, and M. Walmer, "Apricot: A dataset of physical adversarial attacks on object detection," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*. Springer, 2020, pp. 35–50.
- [6] P. Koehn and R. Knowles, "Six challenges for neural machine translation," 2017.
- [7] S. Wiseman, S. M. Shieber, and A. M. Rush, "Challenges in data-to-document generation," 2017.
- [8] K. Filippova, "Controlled hallucinations: Learning to generate faithfully from noisy data," *arXiv preprint arXiv:2010.05873*, 2020.
- [9] T. R. Insel, "Rethinking schizophrenia," *Nature*, vol. 468, no. 7321, pp. 187–193, 2010.
- [10] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. J. Bang, A. Madotto, and P. Fung, "Survey of hallucination in natural language generation," *ACM Computing Surveys*, vol. 55, no. 12, pp. 1–38, 2023.
- [11] S. D. Østergaard and K. L. Nielbo, "False responses from artificial intelligence models are not hallucinations," *Schizophrenia Bulletin*, p. sbad068, 2023.
- [12] R. Emsley, "Chatgpt: these are not hallucinations—they're fabrications and falsifications," *Schizophrenia*, vol. 9, no. 1, p. 52, 2023.
- [13] H. Ye, T. Liu, A. Zhang, W. Hua, and W. Jia, "Cognitive mirage: A review of hallucinations in large language models," 2023.
- [14] C. Karakas, D. Brock, and A. Lakhotia, "Leveraging chatgpt in the pediatric neurology clinic: Practical considerations for use to improve efficiency and outcomes," *Available at SSRN 4475000*, 2023.
- [15] D. L. Rodgers, M. Needler, A. Robinson, R. Barnes, T. Brosche, J. Hernandez, J. Poore, P. VandeKoppel, and R. Ahmed, "Artificial intelligence and the simulationists," *Simulation in Healthcare: Journal of the Society for Simulation in Healthcare*, 2023.
- [16] M. G. Madden, B. A. McNicholas, and J. G. Laffey, "Assessing the usefulness of a large language model to query and summarize unstructured medical notes in intensive care," *Intensive Care Medicine*, pp. 1–3, 2023.
- [17] Z. Li, "The dark side of chatgpt: Legal and ethical challenges from stochastic parrots and hallucination," *arXiv preprint arXiv:2304.14347*, 2023.
- [18] J. Ge and J. Lai, "Artificial intelligence-based text generators in hepatology: Chatgpt is just the beginning," *Hepatology Communications*, vol. 7, no. 4, 2023.
- [19] N. Curtis *et al.*, "To chatgpt or not to chatgpt? the impact of artificial intelligence on academic publishing," *The Pediatric Infectious Disease Journal*, vol. 42, no. 4, p. 275, 2023.
- [20] A. Borji, "A categorical archive of chatgpt failures," *arXiv preprint arXiv:2302.03494*, 2023.
- [21] A. J. Thirunavukarasu, D. S. J. Ting, K. Elangovan, L. Gutierrez, T. F. Tan, and D. S. W. Ting, "Large language models in medicine," *Nature medicine*, pp. 1–11, 2023.
- [22] D. S. J. Ting, T. F. Tan, and D. S. W. Ting, "Chatgpt in ophthalmology: the dawn of a new era?" *Eye*, pp. 1–4, 2023.
- [23] A. Sriwastwa, P. Ravi, A. Emmert, S. Chokshi, S. Kondor, K. Dhal, P. Patel, L. L. Chepelev, F. J. Rybicki, and R. Gupta, "Generative ai for medical 3d printing: a comparison of chatgpt outputs to reference standard education," *3D Printing in Medicine*, vol. 9, no. 1, p. 21, 2023.
- [24] J. Gravel, M. D'Amours-Gravel, and E. Osmanliu, "Learning to fake it: limited responses and fabricated references provided by chatgpt for medical questions," *Mayo Clinic Proceedings: Digital Health*, vol. 1, no. 3, pp. 226–234, 2023.
- [25] A. Bryant, "Ai chatbots: Threat or opportunity?" *Informatics*, vol. 10, no. 2, 2023.
- [26] M. Miao, F. Meng, Y. Liu, X.-H. Zhou, and J. Zhou, "Prevent the language model from being overconfident in neural machine translation," *arXiv preprint arXiv:2105.11098*, 2021.
- [27] K. Lee, O. Firat, A. Agarwal, C. Fannjiang, and D. Sussillo, "Hallucinations in neural machine translation," 2019. [Online]. Available: <https://openreview.net/forum?id=SkxJ-309FQ>
- [28] Z. Zhao, S. B. Cohen, and B. Webber, "Reducing quantity hallucinations in abstractive summarization," *arXiv preprint arXiv:2009.13312*, 2020.
- [29] Y. Dong, S. Wang, Z. Gan, Y. Cheng, J. C. K. Cheung, and J. Liu, "Multi-fact correction in abstractive text summarization," *arXiv preprint arXiv:2010.02443*, 2020.
- [30] J. Maynez, S. Narayan, B. Bohnet, and R. McDonald, "On faithfulness and factuality in abstractive summarization," *arXiv preprint arXiv:2005.00661*, 2020.
- [31] Y. Lyu, C. Zhu, T. Xu, Z. Yin, and E. Chen, "Faithful abstractive summarization via fact-aware consistency-constrained transformer," in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 2022, pp. 1410–1419.
- [32] OpenAI. (2022) Chatgpt 3.5. <https://www.openai.com/>.
- [33] C. DeVon. (2023) Ai chatbots can 'hallucinate' and make things up—why it happens and how to spot it. Accessed on Jan 03, 2024. [Online]. Available: <https://www.cnbc.com/2023/12/22/why-ai-chatbot-s-hallucinate.html>
- [34] K. Weise and C. Metz. (2023) When a.i. chatbots hallucinate. Accessed on Jan 03, 2024. [Online]. Available: <https://www.nytimes.com/2023/05/01/business/ai-chatbots-hallucination.html>
- [35] S. Diamond. (2023) A.i. chatbots, hens and humans can all 'hallucinate'. Accessed on Jan 03, 2024. [Online]. Available: <https://www.nytimes.com/2023/12/17/insider/ai-chatbots-humans-hallucinate.html>
- [36] C. Thorbecke. (2023) Ai tools make things up a lot, and that's a huge problem. Accessed on Jan 03, 2024. [Online]. Available: <https://www.cnn.com/2023/08/29/tech/ai-chatbot-hallucinations/index.html>
- [37] G. Press. (2023) Celebrating 80 years of hallucinating about artificial intelligence. Accessed on Jan 03, 2024. [Online]. Available: <https://www.forbes.com/sites/gilpress/2023/12/30/celebrating-80-years-of-hallucinating-about-artificial-intelligence/?sh=6b6b090c2e96>
- [38] C. Dictionary. (2023) The cambridge dictionary word of the year 2023 is... Accessed on Jan 03, 2024. [Online]. Available: <https://dictionary.cambridge.org/editorial/woty>
- [39] Dictionary.com. (2023) Word of the year. Accessed on Jan 03, 2024. [Online]. Available: <https://content.dictionary.com/word-of-the-year-2023/>