

Confabulation: The Surprising Value of Large Language Model Hallucinations

Peiqi Sui^{1*} Eamon Duede^{2†} Sophie Wu¹ Richard Jean So^{1‡}

¹McGill University, CA

²Harvard University, USA

*peiqi.sui@mail.mcgill.ca †eduede@g.harvard.edu ‡richard.so@mcgill.ca

Abstract

This paper presents a systematic defense of large language model (LLM) hallucinations or ‘confabulations’ as a potential resource instead of a categorically negative pitfall. The standard view is that confabulations are inherently problematic and AI research should eliminate this flaw. In this paper, we argue and empirically demonstrate that measurable semantic characteristics of LLM confabulations mirror a human propensity to utilize increased narrativity as a cognitive resource for sense-making and communication. In other words, it has potential value. Specifically, we analyze popular hallucination benchmarks and reveal that hallucinated outputs display increased levels of narrativity and semantic coherence relative to veridical outputs. This finding reveals a tension in our usually dismissive understandings of confabulation. It suggests, counter-intuitively, that the tendency for LLMs to confabulate may be intimately associated with a positive capacity for coherent narrative-text generation.

1 Background

Artificial intelligence is rapidly becoming ubiquitous in science and scholarship with autoregressive large language models (LLMs) leading the way (Duede et al., 2024). Since the concept’s integration into natural language processing (NLP) discourse, hallucination in LLMs has been widely viewed as a conspicuous social harm and a critical bottleneck to real-world applications of LLMs. Both popular, comprehensive academic surveys (Ji et al., 2023; Kaddour et al., 2023) and public-facing technical reports (OpenAI, 2023; Touvron et al., 2023; Google, 2023) position the problem of hallucination as one of LLM’s chief ethical and safety pitfalls, one that should be heavily mitigated in conjunction with other concerns like bias and toxicity. Yet, despite a wide range of such mitigation efforts (Huang et al., 2023), hallucinations endure, presenting an imminent risk to model trustworthiness

(Luo et al., 2024), with epistemic consequences in truth-sensitive domains like law (Curran et al., 2023), medicine (Pal et al., 2023), finance (Kang and Liu, 2023), science (Alkaissi and McFarlane, 2023; Duede, 2022, 2023), and education (Zhou et al., 2024).

With so much at stake, NLP researchers have naturally developed a **normative** stance against hallucinations, rightly extending epistemic worries to more far-reaching, ethical concerns like misinformation (Li, 2023), deception (Zhan et al., 2023), fairness (Wang et al., 2023), and the broader goals of alignment research (Ouyang et al., 2022). Consequently, the observed commitment to uniform reduction of hallucinations to negligible levels is not just seen as a technical challenge but also a critical component of the broader mission to mitigate the societal stigmas and systemic risks associated with the broad deployment and wide adoption of LLMs.

However, a smaller body of work advances the view that hallucinations are not inherently harmful. According to this view, hallucinated outputs have particular affordances. For instance, such outputs can be utilized for the exploration of domain-specific scenarios where properties of hallucination could be leveraged constructively (Duede and So, Forthcoming). Cao et al. (2022) demonstrate that certain types of hallucinations are factual and could positively contribute to text summarization. Yao et al. (2023) explore the potential of hallucinatory responses as adversarial examples for enhancing the robustness of models. Following an existing practice in computer vision to use hallucinated materials as synthetic training data, (Fei et al., 2023) devise a visual scene hallucination mechanism to dynamically fill in the missing visual modality for text-only inference with multimodal machine translation models (McKee et al., 2021; Wu et al., 2023; Shah et al., 2023).

This **explorative** view highlights the potential

Type	Definition of Confabulation
Clinical Def.	“A neuropsychiatric disorder wherein a patient generates a false memory without the intention of deceit” (Wiggins and Bunin, 2023)
NLP Def.	“...where an LLM seemingly fills gaps in the information contained in the model with plausibly-sounding words” (Nejjar et al., 2023)
(Clinically-Informed) NLP Def.	“The generation of narrative details that, while incorrect, are not recognized as such... mistaken reconstructions of information which are influenced by existing knowledge, experiences, expectations, and context” (Smith et al., 2023)

Table 1: Extant definitions of ‘confabulation’ from the literature.

affordances and plausible necessity of hallucinations. The possible necessity of hallucinations is supported by recent research positing that hallucinations are a statistical inevitability (Kalai and Vempala, 2023; Xu et al., 2024), and impossible to eliminate from LLMs as efforts to do so would be limited to a trade-off between generativity, creativity, and information accuracy (Lee, 2023; Sinha et al., 2023; Zhang, 2023). Moreover, in many domain-specific applications, achieving an optimized balance between creativity and factuality is more effective for maximizing the utility of LLMs than merely attempting to eliminate hallucinations. LLM use cases where hallucinations could be particularly valuable include discovering novel proteins (Anishchenko et al., 2021), providing inspiration for creative writing (Mukherjee and Chang, 2023), and formulating innovative legal analogies (Dahl et al., 2024).

In this paper, we attempt to broaden the concept of hallucination and argue that hallucination is closer in kind to the concept of ‘confabulation’, a term that has already gained popularity in public discourse on AI but has yet to become widespread in the academic literature. Reorienting our conception of what is typically conceived of as hallucination towards confabulation offers a more flexible way to characterize, measure, and analyze factually inaccurate outputs.

We first operationalize a common measure of narrativity following (Piper et al., 2021; Piper and Bagga, 2022; Antoniak et al., 2023) and then demonstrate empirically (see Section 3) that ‘hallucinations’ manifest significantly higher ‘narrativity’ than strictly factual outputs. Scholarship in cognitive science and cultural analytics suggests that narrativity provides beneficial cognitive and

social effects and that, by extension, confabulations that demonstrate such features are likely to offer similar affordances. Taking this broader perspective, in turn, allows researchers to move beyond a strictly normative approach that narrowly focuses on identifying and discarding hallucinations. As such, casting hallucination as confabulation offers a practical and tangible framework for adjusting to the nuance of hallucinations.

2 Related Work

2.1 Confabulation vs Hallucination

While both are anthropomorphizing analogies borrowed from psychiatry, ‘confabulation’ has recently emerged as a preferred alternative to ‘hallucination’ in the public discourse on AI for three reasons. First, confabulation does not pertain to precepts or stimuli, thus avoiding the thorny implication that LLMs have sensory experiences or even consciousness (Smith et al., 2023). Second, as a more neutral term, confabulation is more sensitive to groups dealing with mental health disorders and better at signaling a lack of malicious intent in that the output is not to be thought of as a fabrication (Emsley, 2023). Finally, the concept’s applicability to non-pathological scenarios is supported by research from psychiatry suggesting that everyday memory reconstruction often involves some degree of confabulation (French et al., 2009) and is equally common for healthy individuals to inadvertently fictionalize details of stories without the intention to deceive (Riesthuis et al., 2023) (points we return to in detail in Section 4).

Despite conceptual distinction and advantages over ‘hallucination’, ‘confabulation’ is still commonly used interchangeably with hallucination in

the NLP literature (Romera-Paredes et al., 2023), quite often trivially so (Liu et al., 2023b; Zhan et al., 2023; Suraworachet et al., 2024). From the perspective of model architecture, confabulation has been considered as a formal disparity in confidence level between a model’s queried output and probed internal representation of truthfulness (Liu et al., 2023a). More severe implications of the term have been discussed from the perspective of medicine. For instance (Hatem et al., 2023; Brenner, 2023) have debated the trade-off between the stigmatization that attaches to ‘hallucination’ and the semantic baggage that attends the anthropomorphism of LLM outputs characterized as confabulation. However, neither can offer a more substantive definition than their cited clinical baselines (Wiggins and Bunin, 2023).

2.2 Towards a Narrative-Centered Definition of Confabulation

However, extant definitions (see Table 1) in the literature do not sufficiently account for confabulation’s social and cognitive benefits in human communication. That is, current definitions attributed to LLMs fail to consider that, unlike hallucination, confabulation affords communicative and cognitive benefits to humans when endeavoring to fill in gaps in their knowledge of contextually relevant details to produce coherent linguistic communication. A conflation of confabulation with hallucination limits the potential affordances to NLP research that the concept of confabulation could play.

To address this, we define confabulation in a Narrative-Centric way as a latent narrative impulse to generate more substantive and coherent outputs—a characteristic of LLM textual outputs that closely mirrors the human predisposition to storytelling as a cognitive resource for sensemaking. Specifically, confabulation is a narrative impulse to schematize the information at hand into self-consistent stories, even if there might not be enough available details to do so, in which case would result in the generation of fictional yet plausible information. When humans do not have access to sufficient information to formulate coherent semantic meaning, they often confabulate to ‘fill in the blanks’ with self-consistent narratives that are not necessarily factual but bear close semantic verisimilitude to reality. Following existing conceptions of narrativity in narrative understanding (Piper et al., 2021), we further hypothesize that the extent of this compensatory storytelling can be

measured as a scalar construct that quantifies the degree of narrativity.

3 Data, Methods, and Results

3.1 Empirical Results for Higher Narrativity in Hallucinations

To validate the narrative-rich properties of confabulation, we compare the narrativity of hallucinations and their ground truth counterparts across three popular dialogic hallucination benchmarks:

- FaithDial (Dziri et al., 2022a) is a hallucination-free dialogue benchmark between an information-seeking user and a chatbot, adapted from the Wizard of Wikipedia (WoW) benchmark (Dinan et al., 2018). Mechanical Turk annotators labeled WoW’s human-generated responses as either ‘Hallucination’ or as truthful responses. Truthful responses are broken into three classes: ‘Entailment’, ‘Uncooperative’, and ‘Generic’, and produced faithful and knowledge-grounded edits to 21,445 original responses, all of which we sample for our study.
- BEGIN is a preliminary study of FaithDial conducted to select an existing benchmark for subsequent large-scale annotation and editing (Dziri et al., 2022b). As a smaller expert-curated set, it consists of information-seeking queries, as well as both human-written and model-generated (GPT-2, DoHA, and CTRL) responses, each labeled with a slightly different hallucination taxonomy than FaithDial (with the addition of ‘Partial Hallucination’ as a label) by expert annotators. We adopt BEGIN as a model and dataset agnostic validation of our findings on HaluEval, to confirm the consistency and robustness the narrative patterns across different datasets and models.
- HaluEval (Li et al., 2023) is a comprehensive dataset featuring plausible but hallucinated ChatGPT generations alongside their groundtruth counterparts. Instead of the more fine-grained hallucination labels of FaithDial and BEGIN, HaluEval only differentiates between hallucinated and ground truth responses. We only utilize the dialogue sections of HaluEval with 10,000 samples to maintain domain consistency with the other benchmarks.

	FaithDial			HaluEval		BEGIN		
	Hallucinated	Partial	Truth	Hallucinated	Truth	Hallucinated	Partial	Truth
Count	4485	14108	2852	10000	10000	1019	239	1139
Mean	0.620	0.606	0.518	0.655	0.638	0.658	0.612	0.561
Std	0.178	0.178	0.186	0.138	0.168	0.183	0.188	0.187
Min	0.041	0.052	0.061	0.167	0.029	0.064	0.073	0.059
25%	0.500	0.486	0.379	0.563	0.528	0.541	0.477	0.433
50%	0.640	0.624	0.524	0.669	0.664	0.695	0.642	0.567
75%	0.756	0.741	0.664	0.759	0.768	0.799	0.761	0.698
Max	0.975	0.974	0.959	0.952	0.985	0.965	0.956	0.965

Table 2: Summary statistics for Narrativity across our three benchmark datasets. We observe that, in all three datasets, outputs labeled as ‘hallucinated’ text have, on average, the highest narrativity.

For both the FaithDial and BEGIN datasets, we treat all outputs that do not contain a ‘hallucination’ label as ‘truth’ and all outputs that contain the ‘hallucination’ label as well as an additional truthful label as a ‘partial’ hallucination/truth. This aggregation allows for more direct comparison across datasets. In Table 2 and Figure 1, we present empirical support for the claim that confabulated texts, in general, exhibit increased levels of narrativity and, as such, can be viewed as a form of narrative-rich behavior.

Modeling Narrativity. The degree of narrativity is measured as the softmax probability output of an ELECTRA-large based (Clark et al., 2020) text-classification model fine-tuned on an expert-annotated narrative detection dataset recently curated from Reddit (Antoniak et al., 2023). Antoniak et al. observe that a fine-tuned RoBERTa model outperforms zero-shot and few-shot approaches with larger models like GPT-4 at distinguishing between texts that represent stories from texts that represent non-stories. Following this strategy, we fine-tune a series of encoder-only models for story detection as a text classification task, which perform similarly when evaluated against both Antoniak et al.’s bootstrapped test set (AUC=0.83-0.85) and Piper et al.’s larger (AUC=0.81-0.84) (Piper and Bagga, 2022). We then qualitatively evaluate each model’s classifications to select the model (the fine-tuned ELECTRA-large) that best generalizes to the dialogue contexts of our target benchmarks and further verify its robustness through the manual evaluation of its inference output.

Results. Across all 3 benchmark datasets, we find that hallucinated *qua* confabulated responses exhibit higher narrativity than both partial and

non-hallucination categories and their own ground truth counterparts (See Table 2). Additionally, we find significant correlations between the narrativity score and the hallucination label. In this study, we employed a binomial logistic regression model to investigate the predictive relationship between the independent variable narrativity x and a binary dependent variable y that indicates ‘hallucination’ or ‘not hallucination’ such that $P(y_i = 1|x_i) = \frac{e^{\beta_0 + \beta_1 x_i}}{1 + e^{\beta_0 + \beta_1 x_i}}$. For the purposes of this exercise, we class partial hallucinations as 1. Give the slightly lower mean ‘narrativity’ score for partial hallucination compared to pure hallucinated texts, we suspect that the results of our logistic regression, in fact, under-estimate the magnitude of the association. The logistic regression model allowed us to estimate the probability of y being 1 as a function of x . The model’s coefficients were estimated using maximum likelihood estimation (MLE). Table 3 indicates the results of fitting the model to predict whether an unedited response in FaithDial and BEGIN will likely be labeled as a hallucination. The positive coefficient and low p-value indicate that higher degrees of narrativity are a significantly predictive feature of the hallucination label in both benchmarks.

Variable	Coefficient	Std. Error
Narrativity	0.631 ***	0.059
Intercept	0.368 ***	0.038

Table 3: Regression results of logistic regression. Dependent variable is the hallucination classification (1 for hallucinated output, 0 for ground truth output). Observations: 43,842 | Log Likelihood: -27,397.6 | *** denotes significance at the 0.01 level.




Figure 1: The left panel illustrates distribution for narrative score of hallucinated outputs (blue) and the edited version of the output (gray) in the FaithDial dataset. The hallucinated texts are, in general more narrative rich than those that are edited to resolve inaccuracies. The right panel illustrates distribution for non-hallucinated texts from the FaithDial dataset.

4 In defense of confabulation

In this section, we present a two-part argument that calls for a more careful and nuanced examination of confabulation and possible utilization of its affordances. We argue that confabulation’s narrative-rich properties should not be viewed as a flaw but a hallmark for LLM alignment with a well-established human tendency to use narratives as a versatile tool for persuasion, identity construction, and social negotiation (Bruner, 1991). The normative view’s unreflective dismissal of confabulation, in turn, would risk eliminating critical behavioral and cognitive capacities central to communication and sense-making from the capabilities of LLMs. We also consider how aspects of narrative theory could complicate and enrich our understanding of newly emergent concepts in NLP, like faithfulness and factuality.

4.0.1 Empirical Support for Association of Narrativity and Coherence in Confabulated Texts

To empirically verify the correlation between narrativity and coherence, we extend our analyses to include an evaluation of dialogic coherence by implementing DEAM, a state-of-the-art metric for the task (Ghazarian et al., 2022). DEAM measures coherence in dialogues with a RoBERTa-large model fine-tuned on conversation-level semantic perturbations designed to generate more natural samples of incoherence, which achieves better generalizability from their training set than other metrics for coherence evaluation.

We find that higher narrativity is associated with higher coherence across all three benchmarks. Table 4 shows the results of a beta regression with

Variable	Coefficient	Std. Error
Narrativity	0.372 ***	0.029
Intercept	0.433 ***	0.018

Table 4: Regression results of beta regression. Dependent variable is Coherence and independent variable is output Narrativity. Observations: 65,287 | Pseudo R^2 : 0.004 | Log Likelihood: 109,935.2. | *** denotes significance at the 0.01 level.

narrativity score as an independent variable and coherence as a dependent variable. Beta regression was used because all values in our dataset range between 0 and 1. Beta regression is particularly suited for modeling continuous variables restricted to the interval (0, 1), making it ideal for proportions and rates (Ferrari and Cribari-Neto, 2004). It accounts for the heteroscedasticity commonly observed in such data and provides more accurate estimates than linear regression in these contexts. Our regression model is simple: $\text{Beta}(\text{Coherence}) = \alpha + \beta_1 \text{Narrativity}$ and we apply epsilon smoothing for cases where coherence scores are exactly 0 or 1. We observe that there is generally a highly statistically significant association between increasing narrativity and increasing coherence.

4.1 Narrative, Discourse, and Coherence

Storytelling is pivotal in maintaining discourse-level coherence, a cognitive pattern that text linguistics identifies as a standard default assumption about any text or conversation until proven otherwise (Brown and Yule, 1983). When coherence is not immediately apparent, people actively search for evidence of its otherwise latent presence. In conversation scenarios, for example, the search for

latent coherence can take the form of a cooperative effort between participants, where they collectively work towards making sense of the subject at hand (Grice, 1975).

This latent presence of coherence typically exists in textual spaces “where it is not evident in the surface lexical or propositional cohesion” (Stubbs, 1983, 179), such as implicit semantic relationships like the inclusion of cultural subtexts and schemata that underpin discursive structures (Brown and Yule, 1983), which are, in turn, often represented and conveyed in the form of narratives. Stories, therefore, serve as a *vital communicative and cognitive resource for sense-making*, assisting individuals in navigating complex social and cultural contexts by providing a heuristic framework for understanding implicit meanings and relationships within discourse.

4.1.1 Narratives Help Us Articulate and Understand Complex Arguments

The narrative paradigm (NP) is a communication framework that posits that meaningful communication primarily occurs via forms of storytelling rather than discursive argumentation (Fisher, 1984), based on the widely validated premise that the rhetorical logic of compelling narratives tends to be more persuasive than that of structured arguments (Dahlstrom, 2010; Schreiner et al., 2018; Oschatz and Marker, 2020), particularly in the context of healthcare (Chen et al., 2016; Ballard et al., 2021). NP construes humans as especially disposed to storytelling and whose communication channels and internal world models are inherently sensitive to and ordered by narratives. This notion is echoed by foundational works in psychology (Sacks, 1985; Mercier and Sperber, 2017), philosophy of identity (Taylor, 1992), and philosophy of literature (Bruner, 1987; Dennett, 1988). Moreover, NP introduces the binomial nomenclature “homo narrans” to highlight the importance of storytelling in the human condition, underscoring its essential role in ensuring meaningful communication, mutual intelligibility, and the formation of common sense (Fisher, 1984).

As a model of human behavior, NP measures the effectiveness of conveying and understanding complex information through the narrative merit of its underlying stories. The judgment of this merit is based on two criteria of narrative rationality: 1) narrative coherence, the internal consistency of a story that allows it to make sense in its own con-

text, and 2) narrative fidelity, the degree to which a story aligns or resonates with the receiver’s existing understanding of reality and experiences with other stories (Fisher, 2021). In other words, the coherence of a message depends less on the completeness of its discursive structures than the self-consistency of the stories around which its content is framed.

4.1.2 Narratives Maintain the Consistency of Our Own Internal World Models

Cognitive narratology is a branch of narrative theory that examines the “nexus of narrative and mind” (Herman, 2009, 137) through the cognitive underpinnings of storytelling practices. It considers narratives to be not merely cultural artifacts for literary interpretation but cognitive instruments that serve as an indispensable resource for navigating the world and scaffolding, compartmentalizing, and negotiating our experience (Dannenberg, 2008; Herman, 2013). The foundation of this approach rests on frameworks in cognitive linguistics positing that language and thought are fundamentally shaped by stories and so-called “parabolic projections” that map them onto experience (Turner, 1996).

Leveraging the semantics of possible worlds (Kripke, 1963) and theory of mind (Zunshine, 2006), cognitive narratologists argue that stories accomplish more than building worlds to ensure narrative fidelity with the interpreter’s mental models; they also enhance the self-consistency and robustness of such mental models through the simulation of counterfactual scenarios (Lewis, 1979) and alternative “storyworlds” (Ryan, 1991; Herman, 2013; Gerrig, 2018). This exercise of make-believe not only facilitates narrative understanding but also enriches the interpreter’s cognitive landscape by providing a dynamic platform for testing and refining mental models through imaginative engagement and adaptation (Doležel, 1998; Ryan, 2001), a process that has been linked to evolutionary mechanisms (Easterlin, 2012), especially a predisposition towards play in evolutionary psychology (Boyd, 2009).

Earlier periods of AI research took a particular interest in this mind-narrative nexus due to narrative’s capacity to facilitate complex interpretations with very limited discourse-level information available. Computer scientists have attempted to represent these capabilities formally as 1) cognitive mechanisms like story grammar that could parse stories into sequences and embeddings (Man-

der, 2014), 2) frame representations of common-sense knowledge that reduce interpretive complexity by schematizing expected sequences of events and storing them in memory (Schank and Abelson, 2013). In addition, cognitive narratologists have also suggested that the narrative semantics of possible worlds, in which the coherence of a narrative is complete as its ability to construct immersive story-worlds compellingly, could serve as a hermeneutic approach to probe into the neural processes of automatic story generation (Ryan, 1991).

4.1.3 Narratives Enable Patients to Negotiate the Coherence of Their Experiences

One prominent display of narratives' cognitive affordances is in the medical domain. Narratives have been documented to play a critical role in emotional recovery and rehabilitation in the aftermath of traumatic events that have disrupted the internal coherence of patients' understanding of the world. Clinically validated narrative interventions have improved patient care and outcomes, especially by helping patients grapple with cataclysmic ruptures in the fabric of their experience (Rosário et al., 2018; Coats et al., 2020). Therapeutically, narratives serve dual purposes: for patients, narratives provide a vital sense of agency, giving them control over how to tell their own stories, a process crucial for reconstructing a coherent identity and self-concept (Frank, 2013); for physicians, engaging with patient illness narratives allows them to enhance their empathy by gaining insights into the patient as a whole person, rather than viewing them merely as a collection of symptoms or statistics (Charon, 2001).

5 Limitations and Directions for Future Research

This paper advances a systematic defense of LLM confabulations as a potential resource instead of a categorically negative pitfall. We believe the mental picture that LLMs hallucinate because they are untrustworthy, unfaithful, and ultimately unhumanlike is an oversimplified one. Instead, they confabulate and exhibit narrative-rich behavioral patterns closely resembling the human urge of storytelling – perhaps hallucinations make them more like us than we would like to admit.

While our current findings reveal intriguing associations between increased narrativity and significant increases in coherence, we must refrain from asserting that narrativity drives coherence.

However, this perspective does garner substantial support from the interdisciplinary perspectives presented in our qualitative analysis. This paper serves as a preliminary exploration in this direction, but we require more robust methods for modeling narratives and more comprehensive human evaluations to elucidate the intricacies of this association. We are eager to implement these steps in the next phase of our study.

The textual behaviors and effects we observe in LLM hallucination benchmarks, i.e., higher degrees of narrativity and coherence and the significant correlation between them, are widely considered beneficial in human-to-human communication. But the extent to which these affordances generalize to human-AI interactions—whether humans indeed benefit more in terms of user experience from engaging with high-narrative confabulations versus low-narrative but factual generations—needs to be further validated with human-based evaluations. We plan to follow up this study with experiments with human participants to verify the benefits of narrative engagement as hypothesized.

If the effectiveness of narratively rich confabulations is robustly validated, it could lead to exciting avenues for future research, testing the efficacy of confabulations across diverse domains. Beyond the applications we have already outlined in Section 1, exploring the utility of confabulations in fields such as journalism and advertising could yield valuable insights. By adopting confabulations as a conceptual framework, we could inspire more cross-disciplinary explorations, opening up new possibilities for LLM applications that extend beyond purely factual text generations.

6 Acknowledgement

For the research and completion of this paper, we acknowledge the resources and support provided by the McGill Faculty of Arts, McGill Department of English, and the Division of Research and Faculty Development at the Harvard Business School. Partial funding for the project was provided by Siva Reddy through a FRQNT New Researchers award and a NSERC Discovery grant. We thank Maria Antoniak for her help with the paper. We also acknowledge the helpful feedback from audiences at the University of Chicago's Globus Labs, Data Science Institute, and Knowledge Lab, and Harvard Digital Data Design Institute. Finally, we sincerely thank our anonymous reviewers for their generous

time, attention, and feedback.

References

- Hussam Alkaiissi and Samy I McFarlane. 2023. Artificial hallucinations in chatgpt: implications in scientific writing. *Cureus*, 15(2).
- Ivan Anishchenko, Samuel J Pellock, Tamuka M Chidyausiku, Theresa A Ramelot, Sergey Ovchinnikov, Jingzhou Hao, Khushboo Bafna, Christoffer Norn, Alex Kang, Asim K Bera, et al. 2021. De novo protein design by deep network hallucination. *Nature*, 600(7889):547–552.
- Maria Antoniak, Joel Mire, Maarten Sap, Elliott Ash, and Andrew Piper. 2023. Where do people tell stories online? story detection across online communities. *arXiv preprint arXiv:2311.09675*.
- Anjuly M Ballard, Ashlee Davis, and Cynthia A Hoffner. 2021. The impact of health narratives on persuasion in african american women: A systematic review and meta-analysis. *Health Communication*, 36(5):560–571.
- Brian Boyd. 2009. *On the origin of stories: Evolution, cognition, and fiction*. Harvard University Press.
- Teva D Brender. 2023. Chatbot confabulations are not hallucinations—reply. *JAMA Internal Medicine*, 183(10):1177–1178.
- Gillian Brown and George Yule. 1983. *Discourse analysis*. Cambridge university press.
- Jerome Bruner. 1987. Life as narrative. *Social research*, pages 11–32.
- Jerome Bruner. 1991. The narrative construction of reality. *Critical inquiry*, 18(1):1–21.
- Meng Cao, Yue Dong, and Jackie Cheung. 2022. **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)*, pages 3340–3354, Dublin, Ireland. Association for Computational Linguistics.
- Rita Charon. 2001. Narrative medicine: a model for empathy, reflection, profession, and trust. *Jama*, 286(15):1897–1902.
- Meng Chen, Robert A Bell, and Laramie D Taylor. 2016. Narrator point of view and persuasion in health narratives: the role of protagonist–reader similarity, identification, and self-referencing. *Journal of health communication*, 21(8):908–918.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*.
- Heather Coats, Nadia Shive, Ardith Z Doorenbos, and Sarah J Schmiede. 2020. Integration of person-centered narratives into the electronic health record: study protocol. *Nursing research*, 69(6):483.
- Shawn Curran, Sam Lansley, and Oliver Bethell. 2023. Hallucination is the last thing you need. *arXiv preprint arXiv:2306.11520*.
- Matthew Dahl, Varun Magesh, Mirac Suzgun, and Daniel E Ho. 2024. Large legal fictions: Profiling legal hallucinations in large language models. *arXiv preprint arXiv:2401.01301*.
- Michael F Dahlstrom. 2010. The role of causality in information acceptance in narratives: An example from science communication. *Communication Research*, 37(6):857–875.
- Hilary P Dannenberg. 2008. *Coincidence and counterfactuality: Plotting time and space in narrative fiction*. U of Nebraska Press.
- Daniel C Dennett. 1988. Why everyone is a novelist. *TLS-THE TIMES LITERARY SUPPLEMENT*, 4459(1016).
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2018. Wizard of wikipedia: Knowledge-powered conversational agents. *arXiv preprint arXiv:1811.01241*.
- Lubomír Doležel. 1998. *Heterocosmica*. Johns Hopkins University Press.
- Eamon Duede. 2022. Instruments, agents, and artificial intelligence: novel epistemic categories of reliability. *Synthese*, 200(6):491.
- Eamon Duede. 2023. Deep learning opacity in scientific discovery. *Philosophy of Science*, 90(5):1089–1099.
- Eamon Duede, William Dolan, André Bauer, Ian Foster, and Karim Lakhani. 2024. Oil & water? diffusion of ai within and across scientific fields. *arXiv preprint arXiv:2405.15828*.
- Eamon Duede and Richard Jean So. Forthcoming. The humanistic case for AI optimism. *Poetics Today*.
- Nouha Dziri, Ehsan Kamaloo, Sivan Milton, Osmar Zaiane, Mo Yu, Edoardo M. Ponti, and Siva Reddy. 2022a. **FaithDial: A faithful benchmark for information-seeking dialogue**. *Transactions of the Association for Computational Linguistics*, 10:1473–1490.
- Nouha Dziri, Sivan Milton, Mo Yu, Osmar Zaiane, and Siva Reddy. 2022b. **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*, pages 5271–5285, Seattle, United States. Association for Computational Linguistics.

- Nancy Easterlin. 2012. *A biocultural approach to literary theory and interpretation*. JHU Press.
- Robin Emsley. 2023. Chatgpt: these are not hallucinations—they’re fabrications and falsifications. *Schizophrenia*, 9(1):52.
- Hao Fei, Qian Liu, Meishan Zhang, Min Zhang, and Tat-Seng Chua. 2023. [Scene graph as pivoting: Inference-time image-free unsupervised multimodal machine translation with visual scene hallucination](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5980–5994, Toronto, Canada. Association for Computational Linguistics.
- Silvia Ferrari and Francisco Cribari-Neto. 2004. Beta regression for modelling rates and proportions. *Journal of applied statistics*, 31(7):799–815.
- Walter R Fisher. 1984. Narration as a human communication paradigm: The case of public moral argument. *Communications Monographs*, 51(1):1–22.
- Walter R Fisher. 2021. *Human communication as narration: Toward a philosophy of reason, value, and action*. Univ of South Carolina Press.
- Arthur W Frank. 2013. *The wounded storyteller: Body, illness & ethics*. University of Chicago Press.
- Lauren French, Maryanne Garry, and Elizabeth Loftus. 2009. False memories: A kind of confabulation in non-clinical. *Confabulation: Views from neuroscience, psychiatry, psychology, and philosophy*.
- Richard Gerrig. 2018. *Experiencing narrative worlds*. Routledge.
- Sarik Ghazarian, Nuan Wen, Aram Galstyan, and Nanyun Peng. 2022. [DEAM: Dialogue coherence evaluation using AMR-based semantic manipulations](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 771–785, Dublin, Ireland. Association for Computational Linguistics.
- Google. 2023. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*.
- Herbert P Grice. 1975. Logic and conversation. In *Speech acts*, pages 41–58. Brill.
- Rami Hatem, Brianna Simmons, and Joseph E Thornton. 2023. Chatbot confabulations are not hallucinations. *JAMA Internal Medicine*, 183(10):1177–1177.
- David Herman. 2009. *Basic elements of narrative*. John Wiley & Sons.
- David Herman. 2013. *Storytelling and the Sciences of Mind*. MIT press.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2023. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.
- Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert McHardy. 2023. Challenges and applications of large language models. *arXiv preprint arXiv:2307.10169*.
- Adam Tauman Kalai and Santosh S Vempala. 2023. Calibrated language models must hallucinate. *arXiv preprint arXiv:2311.14648*.
- Haoqiang Kang and Xiao-Yang Liu. 2023. Deficiency of large language models in finance: An empirical examination of hallucination. *arXiv preprint arXiv:2311.15548*.
- Saul Kripke. 1963. Semantical considerations on modal logic. *Acta Philosophica Fennica*, 16:83–94.
- Minhyeok Lee. 2023. A mathematical investigation of hallucination and creativity in gpt models. *Mathematics*, 11(10):2320.
- David Lewis. 1979. Counterfactual dependence and time’s arrow. *Noûs*, pages 455–476.
- Junyi Li, Xiaoxue Cheng, Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023. [HaluEval: A large-scale hallucination evaluation benchmark for large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6449–6464, Singapore. Association for Computational Linguistics.
- Zihao Li. 2023. The dark side of chatgpt: Legal and ethical challenges from stochastic parrots and hallucination. *arXiv preprint arXiv:2304.14347*.
- Kevin Liu, Stephen Casper, Dylan Hadfield-Menell, and Jacob Andreas. 2023a. Cognitive dissonance: Why do language model outputs disagree with internal representations of truthfulness? *arXiv preprint arXiv:2312.03729*.
- Zhenzhen Liu, Chao Wan, Varsha Kishore, Jin Peng Zhou, Minmin Chen, and Kilian Q Weinberger. 2023b. Correction with backtracking reduces hallucination in summarization. *arXiv preprint arXiv:2310.16176*.
- Junliang Luo, Tianyu Li, Di Wu, Michael Jenkin, Steve Liu, and Gregory Dudek. 2024. Hallucination detection and hallucination mitigation: An investigation. *arXiv preprint arXiv:2401.08358*.
- Jean Matter Mandler. 2014. *Stories, scripts, and scenes: Aspects of schema theory*. Psychology Press.

- Daniel McKee, Bing Shuai, Andrew Berneshawi, Manchen Wang, Davide Modolo, Svetlana Lazebnik, and Joseph Tighe. 2021. Multi-object tracking with hallucinated and unlabeled videos. *arXiv preprint arXiv:2108.08836*.
- Hugo Mercier and Dan Sperber. 2017. *The enigma of reason*. Harvard University Press.
- Anirban Mukherjee and Hannah Chang. 2023. The creative frontier of generative ai: Managing the novelty-usefulness tradeoff. *arXiv preprint arXiv:2306.03601*.
- Mohamed Nejjar, Luca Zacharias, Fabian Stiehle, and Ingo Weber. 2023. Llms for science: Usage for code generation and data analysis. *arXiv preprint arXiv:2311.16733*.
- OpenAI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Corinna Oschatz and Caroline Marker. 2020. Long-term persuasive effects in narrative communication research: A meta-analysis. *Journal of Communication*, 70(4):473–496.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2023. **Med-HALT: Medical domain hallucination test for large language models**. In *Proceedings of the 27th Conference on Computational Natural Language Learning (CoNLL)*, pages 314–334, Singapore. Association for Computational Linguistics.
- Andrew Piper and Sunyam Bagga. 2022. Toward a data-driven theory of narrativity. *New Literary History*, 54(1):879–901.
- Andrew Piper, Richard Jean So, and David Bamman. 2021. Narrative theory for computational narrative understanding. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 298–311.
- Paul Riesthuis, Henry Otgaar, Glynis Bogaard, and Ivan Mangiulli. 2023. Factors affecting the forced confabulation effect: a meta-analysis of laboratory studies. *Memory*, 31(5):635–651.
- Bernardino Romera-Paredes, Mohammadamin Barekatin, Alexander Novikov, Matej Balog, M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang, Omar Fawzi, et al. 2023. Mathematical discoveries from program search with large language models. *Nature*, pages 1–3.
- FRANCISCO Rosário, ISABEL CORREIA, JORGE OLIVEIRA, and JOAO RAPOSO. 2018. Narrative medicine in diabetes—a randomized trial. *Diabetes*, 67(Supplement_1).
- Marie-Laure Ryan. 1991. *Possible worlds, artificial intelligence, and narrative theory*. Indiana University Press.
- Marie-Laure Ryan. 2001. Narrative as virtual reality. *Immersion and Interactivity in Literature*, pages 357–359.
- Oliver Sacks. 1985. The man who mistook his wife for a hat and other clinical tales.
- Roger C Schank and Robert P Abelson. 2013. *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Psychology Press.
- Constanze Schreiner, Markus Appel, Maj-Britt Isberner, and Tobias Richter. 2018. Argument strength and the persuasiveness of stories. *Discourse Processes*, 55(4):371–386.
- Anshul Shah, Aniket Roy, Ketul Shah, Shlok Mishra, David Jacobs, Anoop Cherian, and Rama Chellappa. 2023. Halp: Hallucinating latent positives for skeleton-based self-supervised learning of actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18846–18856.
- Ritwik Sinha, Zhao Song, and Tianyi Zhou. 2023. A mathematical abstraction for balancing the trade-off between creativity and reality in large language models. *arXiv preprint arXiv:2306.02295*.
- Andrew L Smith, Felix Greaves, and Trishan Panch. 2023. Hallucination or confabulation? neuroanatomy as metaphor in large language models. *PLOS Digital Health*, 2(11):e0000388.
- Michael Stubbs. 1983. *Discourse analysis: The sociolinguistic analysis of natural language*. Basil Blackwell.
- Wannapon Suraworachet, Jennifer Seon, and Mutlu Cukurova. 2024. Predicting challenge moments from students’ discourse: A comparison of gpt-4 to two traditional natural language processing approaches. *arXiv preprint arXiv:2401.01692*.
- Charles Taylor. 1992. *Sources of the Self*. Harvard University Press.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Mark Turner. 1996. *The literary mind: The origins of thought and language*. Oxford University Press.

- Skyler Wang, Ned Cooper, Margaret Eby, and Eun Seo Jo. 2023. From human-centered to social-centered artificial intelligence: Assessing chatgpt’s impact through disruptive events. *arXiv preprint arXiv:2306.00227*.
- Amanda Wiggins and Jessica L. Bunin. 2023. *Confabulation*. StatPearls Publishing, Treasure Island (FL).
- Jing Wu, Jennifer Hobbs, and Naira Hovakimyan. 2023. Hallucination improves the performance of unsupervised visual representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16132–16143.
- Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. 2024. Hallucination is inevitable: An innate limitation of large language models. *arXiv preprint arXiv:2401.11817*.
- Jia-Yu Yao, Kun-Peng Ning, Zhen-Hui Liu, Mu-Nan Ning, and Li Yuan. 2023. Llm lies: Hallucinations are not bugs, but features as adversarial examples. *arXiv preprint arXiv:2310.01469*.
- Xiao Zhan, Yifan Xu, and Stefan Sarkadi. 2023. *Deceptive ai ecosystems: The case of chatgpt*. In *Proceedings of the 5th International Conference on Conversational User Interfaces, CUI '23*, New York, NY, USA. Association for Computing Machinery.
- Chen Zhang. 2023. User-controlled knowledge fusion in large language models: Balancing creativity and hallucination. *arXiv preprint arXiv:2307.16139*.
- Kyrie Zhixuan Zhou, Zachary Kilhoffer, Madelyn Rose Sanfilippo, Ted Underwood, Ece Gumusel, Mengyi Wei, Abhinav Choudhry, and Jinjun Xiong. 2024. "the teachers are confused as well": A multiple-stakeholder ethics discussion on large language models in computing education. *arXiv e-prints*, pages arXiv–2401.
- Lisa Zunshine. 2006. *Why we read fiction: Theory of mind and the novel*. Ohio State University Press.