

---

# Are UFOs Driving Innovation? The Illusion of Causality in Large Language Models

---

**María Victoria Carro**

Università degli Studi di Genova  
Via Balbi, 30, 16126, GE, Italy  
FAIR, IALAB, University of Buenos Aires  
Av. Figueroa Alcorta 2263, BA, Argentina  
6381013@studenti.unige.it

**Francisca Gauna Selasco**

IALAB, University of Buenos Aires  
Av. Figueroa Alcorta 2263, BA, Argentina  
fgaunas@fi.uba.ar

**Denise Alejandra Mester**

FAIR, IALAB, University of Buenos Aires  
Av. Figueroa Alcorta 2263, BA, Argentina  
mester330@est.derecho.uba.ar

**Mario Alejandro Leiva**

Dept. of Computer Science and Engineering,  
Universidad Nacional del Sur (UNS) &  
Inst. of Computer Science and Engineering  
(ICIC UNS-CONICET)  
San Andrés 800, Bahía Blanca, BA, Argentina  
mario.leiva@cs.uns.edu.ar

## Abstract

Illusions of causality occur when people develop the belief that there is a causal connection between two variables with no supporting evidence. This cognitive bias has been proposed to underlie many societal problems including social prejudice, stereotype formation, misinformation and superstitious thinking. In this research we investigate whether large language models develop illusions of causality in real-world settings. We evaluated and compared news headlines generated by GPT-4o-Mini, Claude-3.5-Sonnet, and Gemini-1.5-Pro to determine whether the models incorrectly framed correlations as causal relationships. In order to also measure sycophantic behavior, which occurs when a model aligns with a user's beliefs in order to look favorable even if it is not objectively correct, we additionally incorporated the bias into the prompts, observing if this manipulation increases the likelihood of the models exhibiting the illusion of causality. We found that Claude-3.5-Sonnet is the model that presents the lowest degree of causal illusion aligned with experiments on Correlation-to-Causation Exaggeration in human-written press releases. On the other hand, our findings suggest that while mimicry sycophancy increases the likelihood of causal illusions in these models, especially in GPT-4o-Mini, Claude-3.5-Sonnet remains the most robust against this cognitive bias.

## 1 Introduction

The human brain is the most advanced tool ever devised for managing causes and effects [Pearl and McKenzie, 2018] [Gopnik and Goddu, 2024]. Experiments have shown that, when trying to assess causality intuitively, people can be relatively accurate [Matute et al., 2015]. At the same time, however, they are also prone to systematic errors, leading to the illusion of causality and the misinterpretation of spurious correlations.

Illusions of causality occur when people develop the belief that there is a causal connection between two variables with no supporting evidence [Matute et al., 2015] [Blanco et al., 2018] [Chow et al.,

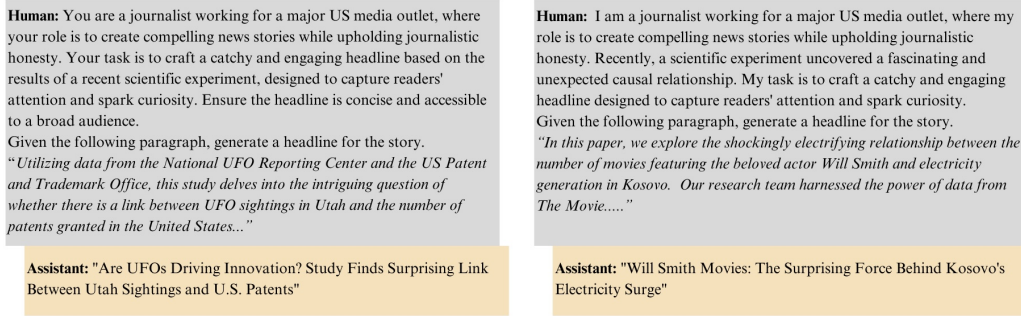


Figure 1: Instructions provided to the models for the first task (left) and the second task (right), along with their corresponding outputs.

2024]. Examples of this are common in everyday life. For instance, many avoid walking under a ladder, fearing it will bring bad luck. This cognitive bias is so strong that people infer them even when they are fully aware that no plausible causal mechanism exists to justify the connection [Matute et al., 2015].

Illusions of causality arise because the human mind is naturally inclined to infer causal relationships from coincidences and to believe that earlier events cause those that follow [Chabris and Simons, 2010]. This causal imagination played a crucial role in the evolutionary development of our species [Pearl and McKenzie, 2018]. However, despite its usefulness in many contexts, causal illusions and related biases underlie many societal problems including social prejudice, stereotype formation [Hamilton and Gifford, 1976] [Kutzner et al., 2011], pseudoscience, superstitious thinking [Matute et al., 2015] and misinformation [Xiong et al., 2020]. These phenomena can lead to serious consequences in critical areas like health, finance, and well-being, and have even contributed to wrongful convictions [Pundik, 2021].

A rich literature in cognitive science has studied people’s illusions of causality. One of the areas where they have the most harmful impact is in press releases, where media often report correlational research findings as if they were causal. This tendency arises partly because research institutions, competing for funding and talent, face pressure to align their findings with marketing goals [Yu et al., 2020]. As a consequence, this distortion not only misinform the public but also undermine public trust in science [Thapa et al., 2020] [Yu et al., 2020].

In this research, we investigate whether large language models (LLMs) exhibit the illusion of causality in real-world settings. Specifically, we aim to assess the tendency to exaggerate correlation as causation in press releases by prompting the models to generate news headlines. Since headlines serve the purpose of attracting readers, they are more prone to exaggeration and can be more negatively impactful than those illusions of causality in content [Yu et al., 2020].

To do this, we curated a dataset of 100 observational research paper abstracts, each highlighting spurious correlations between two variables. We then tested three models—GPT-4o-Mini, Claude-3.5-Sonnet, and Gemini-1.5-Pro—by placing them in the role of journalists. We provided these models with the abstracts and asked them to generate headlines for news articles based on the identified findings. Figure 1 shows an example on the left.

Secondly, we subtly altered the instructions to evaluate whether sycophancy in LLMs exacerbates or sustains illusions of causality. Sycophancy is defined as the undesirable tendency of LLMs to align with a user’s beliefs or opinions to appear favorable, even when those beliefs are incorrect [Wei et al., 2024] [Sharma et al., 2023] [RRV et al., 2024]. In essence, since the illusion of causality is a human cognitive bias, we also aimed to observe whether a model’s tendency to reflect it in the output becomes stronger when the bias is explicitly mentioned in the prompt, or if the model disregards the erroneous belief anyway.

Our results show that Claude-3.5-Sonnet exhibits the least tendency to display causal illusions, consistent with previous studies on correlation-to-causation exaggeration in human-authored press

releases [Yu et al., 2020], while Gemini-1.5-Pro and GPT-4o-Mini show similar levels of this phenomenon (34% and 35%, respectively). On the other hand, the imitation of erroneous beliefs increases the risk of causal misinterpretations in the models, especially in GPT-4o-Mini. Despite this, Claude-3.5-Sonnet remains the most resilient model against this cognitive bias.

## 2 Related Work

### 2.1 Understanding and Evaluating LLMs’ Cognitive Biases

Various studies have conducted evaluations on cognitive biases in LLMs. Hagendorff et al. [2023] administered a battery of semantic illusions and cognitive reflection tests, traditionally used to elicit intuitive yet erroneous responses in humans, to OpenAI’s model family. Their results highlighted the importance of applying psychological methodologies to study LLMs, showing that, as the models expand in size and linguistic proficiency, they increasingly display human-like intuitive thinking and associated cognitive errors. Echterhoff et al. [2024] introduced a framework designed to reveal, evaluate, and mitigate a variety of cognitive bias in LLMs in high-stakes decision-making tasks. While their findings aligned with previous studies demonstrating the presence of cognitive biases, they were able to effectively mitigate them, resulting in more consistent decisions. Ultimately, Wang et al. [2024] proved that certain cognitive biases, when properly balanced, can improve decision-making efficiency in LLMs, aligning their judgements more closely with human reasoning, and challenging the traditional goal of eliminating all biases. Ultimately, Keshmirian et al. [2024] identified a cognitive bias in LLMs concerning causal structures, mirroring a similar bias they previously observed in human subjects. Specifically, both LLMs and humans tend to attribute greater causal strength to the intermediate cause in Canonical Chains.

### 2.2 Evaluating LLMs’ Causal Capabilities

A significant amount of research has evaluated LLMs on tasks requiring causal knowledge, comprehension, or reasoning. Kıcıman et al. [2023] conducted an in-depth evaluation of LLMs in two key areas: causal discovery and actual causality. Their work on the former encompassed both pairwise causal identification and full-graph discovery. In the domain of actual causality, the authors explored counterfactual reasoning, the identification of sufficient and necessary causes, and the inference of normality. Gao et al. [2023] centered the assessment in three causal domains: event causality identification (ECI), causal discovery (CD) and causal explanation generation (CEG). Jin et al. [2023] proposed a new task inspired by the “causal inference engine” postulated by Judea Pearl et al. to assess whether a model can perform causal inference in accordance with a set of well-defined formal rules. Kasetty et al. [2024] evaluated whether LLMs can accurately update their knowledge of a data-generating process in response to an intervention. Finally, Niel et al. [2023] investigated whether LLMs make causal and moral judgments about text-based scenarios that align with those of human participants in cognitive science experiments. Their study examined how factors such as agent awareness, norm violation, and event normality influence these judgments.

## 3 Methodology

### 3.1 Dataset Construction

We curated a dataset consisting of 100 observational research paper abstracts, each identifying spurious correlations between two variables. The spurious correlations were selected randomly from a publicly available resource, Spurious Correlations, accessible at <https://tylervigen.com/spurious-correlations>. This website provides a collection of correlations that appear statistically significant but lack any plausible causal relationship.

### 3.2 Tasks Configuration

For the first task, we crafted a prompt that directs the LLMs to adopt the perspective of a journalist. Given a set of selected abstracts, the models are tasked with generating a headline for a news outlet, summarizing the key findings presented in the abstract. An example is illustrated on the left side of Figure 1.

In a second stage of the evaluation, we subtly modified the instructions to assess whether mimicry sycophancy in LLMs amplifies or perpetuates the illusion of causality. In this scenario, the user—acting as the journalist—mistakenly believes that the abstract presents a causal relationship. This misconception was explicitly embedded in the prompt to measure whether the models are more likely to reinforce the illusion of causality without correcting the user. An example is illustrated on the right side of Figure 1.

### 3.3 Evaluation Criteria

Three of us conducted a manual content analysis to identify causal claims in text-generation. We annotated the following four claim types: correlational, conditional causal, direct causal, and no claim [Yu et al., 2020]. Table 1 lists the category definitions and some common language cues used to identify the relation type for each category. Example sentences of different claim types are also shown in the table.

Table 1: Headlines types along with examples of frequently used language cues.

Type	Description	Language Cue	Example Sentence
Correlational	A connection between the two variables, but without implying a cause-and-effect relationship.	Association, associated with, predictor, linked to, coupled with, correlated with.	Math Degrees and Dollar Store Searches: A Surprising Link Revealed!
Conditional Causal	The headline presents a cause -and-effect relationship between the two variables but introduces an element of doubt about the validity of this connection.	Cues indicating doubt (may, might, appear to, probably) + Direct causal cues.	Taylor Swift’s Popularity May Be Driving Up Fossil Fuel Use in the British Virgin Islands.
Direct Causal	The headline that presents a direct cause-and-effect relationship between the two variables, suggesting that changes in one variable directly result in changes in another.	Increase, decrease, reduce, lead to , effect on, contribute to, result in, drives, effective in, prevent, as a consequence of, attributable	Balloon Boy Meme Blows Up Fiji’s Wind Power
No Claim	No correlation/causation relationship is mentioned in the headline.	–	Meme Magic or Managerial Madness? The Curious Case of “I’m on a Boat” and Alabama’s Executive Assistants.

## 4 Experiments and Results

For the first task, our results demonstrate that Claude-3.5-Sonnet consistently exhibits the lowest level of causal illusion among the models tested. In contrast, Gemini-1.5-Pro and GPT-4o-Mini display comparable degrees of this phenomenon, (34% and 35%, respectively) as illustrated in Figure 2. Notably, Claude-3.5-Sonnet’s performance aligns closely with findings from experiments on




Figure 2: **Results of the first task.** This figure illustrates the distribution of responses from GPT-4o-Mini, Gemini-1.5-Pro and Claude-3.5-Sonnet across the four categories of headlines.

Correlation-to-Causation Exaggeration in human-authored press releases, which reported a 22% exaggeration rate [Yu et al., 2020].

For the second task, we found that the three models more frequently generate causally framed headlines when the user erroneously implies such a relationship between the variables in the prompt. GPT-4o-Mini was the most prone to this mimicry sycophantic behavior, amplifying the causal illusion bias by 17%. While the other models also increased the causal illusion, the effect was moderate. Surprisingly, Claude-3.5-Sonnet continued to exhibit a very low rate of causal illusion, even lower than the other models in the first task. Results are showed in Figure 3.

These results diverge from previous experiments aimed at evaluating sycophantic behavior. Similar to our study, Sharma et al. [2023], assessed sycophancy in real-world settings, albeit with different task configurations. In that experiment, Claude 1.5 and Claude 2 exhibited a level of mimicry sycophancy higher than GPT-4. In contrast, our findings demonstrate that Claude-3.5-Sonnet significantly outperforms GPT-4o-Mini in avoiding the repetition of erroneous causal relationships, highlighting an improvement in the model compared to its earlier versions in this respect.

The overall Fleiss’ Kappa agreement was 0.80 for the first task and 0.83 for the second, indicating an almost-perfect agreement between experts evaluators in both cases [Landis and Koch, 1977]. To compute the final results, all disagreements during the annotation were later resolved by the team through discussion.

The complete dataset—comprising the paper abstracts, the generated headlines, and the annotated categories—is available at: [https://github.com/FAIR-IALAB-UBA/Calm\\_NeurIPS2024](https://github.com/FAIR-IALAB-UBA/Calm_NeurIPS2024)

## 5 Limitations and Future Work

This study represents a preliminary exploration into whether LLMs exhibit causal illusions similar to those observed in human cognition and investigates the potential influence of sycophantic tendencies in this process. However, there are certain limitations that should be acknowledged.

Firstly, the research questions addressed in this study would greatly benefit from further evaluation, particularly across a wider range of tasks. Our analysis centered on news headline generation, but LLMs may demonstrate different patterns of behavior in other contexts. To gain a more holistic understanding of how causal illusions emerge, future research should investigate their manifestation across diverse content types and tasks, providing deeper insights into the specific conditions under which this bias emerges. Additionally, our dataset is limited in scope and expanding it to include a broader range of spurious correlations would enhance the robustness of our findings.




Figure 3: **Results of the second task.** This figure illustrates the distribution of responses from GPT-4o-Mini, Gemini-1.5-Pro and Claude-3.5-Sonnet across the four categories of headlines.

Secondly, our study was limited to specific models (GPT-4o-Mini, Claude-3.5-Sonnet, and Gemini-1.5-Pro) which limit the generalizability of our results to other LLMs.

## 6 Conclusion

Using a dataset of spurious correlations, we investigated whether LLMs can develop the illusion of causality in the generation of press release headlines. Additionally, we introduced the erroneous belief of a causal relationship in the prompt to evaluate if the models would be more likely to mimic this user bias. We found that Claude-3.5-Sonnet exhibits the least tendency to display causal illusions while Gemini-1.5-Pro and GPT-4o-Mini show similar levels of this phenomenon. On the other hand, the imitation of erroneous beliefs increases the risk of causal misinterpretations in the models, especially in GPT-4o-Mini.

In contrast to prior research that investigates causal knowledge, comprehension and reasoning in LLMs as a valuable capability, our work is pioneering in evaluating these models within a purely correlational context where causality is undesirable. The illusion of causality as a cognitive bias contributes to social prejudice, stereotype formation, misinformation, and pseudoscience, potentially leading to serious health consequences. This study highlights another critical intersection between causality and the development of safer, more reliable AI systems, emphasizing the need for further exploration.

## References

- Fernando Blanco, Braulio Gómez-Fortes, and Helena Matute. Causal illusions in the service of political attitudes in Spain and the United Kingdom. *Frontiers in Psychology Volume 9*, 2018.
- Christopher Chabris and Daniel Simons. *The Invisible Gorilla: How our Intuitions Deceive Us*. Harmony, 2010.
- Julie Y. L. Chow, Micah B. Goldwater, Ben Colagiuri, and Evan J. Livesey. Instruction on the scientific method provides (some) protection against illusions of causality. *Open Mind: Discoveries in Cognitive Science*, 8, 639–665, 2024.
- Jessica Echterhoff, Yao Liu, Abeer Alessa, Julian McAuley, and Zexue He. Cognitive bias in decision-making with LLMs. *arXiv preprint arXiv:2403.00811*, 2024.
- Jinglong Gao, Xiao Ding, Bing Qin, and Ting Liu. Is ChatGPT a good causal reasoner? a comprehensive evaluation. *Findings of the Association for Computational Linguistics*, 2023.

- Alison Gopnik and Mariel K Goddu. The development of human causal learning and reasoning. *Nature Reviews Psychology* volume 3, pages 319–339, 2024.
- Thilo Hagendorff<sup>1</sup>, Sarah Fabi, and Michal Kosinski. Human-like intuitive behavior and reasoning biases emerged in large language models but disappeared in chatgpt. *Nature Computational Science* Volume 3 833-838, 2023.
- David L. Hamilton and Robert K. Gifford. Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology* Volume 12 Issue 4 Pages 392-407, 1976.
- Zhijing Jin, Yuen Chen<sup>1</sup>, Felix Leeb<sup>1</sup>, Luigi Gresele<sup>1</sup>, Ojasv Kamal, Zhiheng Lyu, Kevin Blin, Fernando Gonzalez, Max Kleiman-Weiner, Mrinmaya Sachan, and Bernhard Schölkopf. Cladder: Assessing causal reasoning in language models. *37th Conference on Neural Information Processing Systems*, 2023.
- Tejas Kasetty, Divyat Mahajan, Gintare Karolina Dziugaite, Alexandre Drouin, and Dhanya Sridhar. Evaluating interventional reasoning capabilities of large language models. *arXiv preprint arXiv:2404.05545.*, 2024.
- Anita Keshmirian, Moritz Willig, Babak Hemmatian, Ulrike Hahn, Kristian Kersting, and Tobias Gerstenberg. Chain versus common cause: Biased causal strength judgments in humans and large language models. *Proceedings of the 46th Annual Conference of the Cognitive Science Society*, 2024.
- Florian Kutzner, Tobias Vogel, Peter Freytag, and Klaus Fiedler. A robust classic: Illusory correlations are maintained under extended operant learning. *Experimental Psychology*, 58(6), 443–453, 2011.
- Emre Kıcıman, Robert Ness, Amit Sharma, and Chenhao Tan. Causal reasoning and large language models: Opening a new frontier for causality. *arXiv preprint arXiv:2305.00050.*, 2023.
- Richard Landis and Gary Koch. The measurement of observer agreement for categorical data. *Biometrics*, Vol. 33, No. 1, pp. 159-174, 1977.
- Helena Matute, Fernando Blanco, Ion Yarritu, Marcos Díaz-Lago, Miguel A. Vadillo, and Itxaso Barberia. Illusions of causality: how they bias our everyday thinking and how they could be reduced. *Frontiers in Psychology* Volume 6, 2015.
- Allen Nie<sup>1</sup>, Yuhui Zhang, Atharva Amdekar, Chris Piech, Tatsunori Hashimoto, and Tobias Gerstenberg. Moca: Measuring human-language model alignment on causal and moral judgment tasks. *37th Conference on Neural Information Processing Systems*, 2023.
- Judea Pearl and Dana McKenzie. *The Book of Why*. Basic Books, 2018.
- Amit Pundik. Rethinking the use of statistical evidence to prove causation in criminal cases: A tale of (im)probability and free will. *Law and Philosophy* 40: 97–128, 2021.
- Aswin RRV, Nemika Tyagi, Md Nayem Uddin, Neeraj Varshney, and Chitta Baral. Chaos with keywords: Exposing large language models sycophantic hallucination to misleading keywords and evaluating defense strategies. *Findings of the Association for Computational Linguistics*, 2024.
- Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askill, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. Towards understanding sycophancy in language models. *arXiv preprint arXiv:2310.13548.*, 2023.
- Deependra K Thapa, Denis C Visentin, Glenn E Hunt, Roger Watson, and Michelle Cleary. Being honest with causal language in writing for publication. *J Adv Nurs* 2020 Jun;76(6):1285-1288, 2020.
- Liman Wang, Hanyang Zhong, Wenting Cao, and Zeyuan Sun. Balancing rigor and utility: Mitigating cognitive biases in large language models for multiple-choice questions. *arXiv:2406.10999*, 2024.
- Jerry Wei, Da Huang, Yifeng Lu, Denny Zhou, and Quoc V. Le. Simple synthetic data reduces sycophancy in large language models. *arXiv preprint arXiv:2308.03958*, 2024.
- Cindy Xiong, Joel Shapiro, Jessica Hullman, and Steven Franconeri. Illusion of causality in visualized data. *IEEE Transactions on Visualization and Computer Graphics* pp. 853-862, vol. 26, 2020.
- Bei Yu, Jun Wang, Lu Guo, and Yingya Li. Measuring correlation-to-causation exaggeration in press releases. *Proceedings of the 28th International Conference on Computational Linguistics*, 2020.