

Implicit causality in GPT-2: A case study

Minh Hien Huynh and Tomas O. Lentz and Emiel van Miltenburg

Department of Communication and Cognition

Tilburg University

hienhuynh.tdn@gmail.com

{T.O.Lentz, C.W.J.vanMiltenburg}@tilburguniversity.edu

Abstract

This case study investigates the extent to which a language model (GPT-2) is able to capture native speakers' intuitions about *implicit causality* in a sentence completion task. Study 1 reproduces earlier results (showing that the model's surprisal values correlate with the implicit causality bias of the verb; Davis and van Schijndel 2021), and then examine the effects of gender and verb frequency on model performance. Study 2 examines the reasoning ability of GPT-2: Is the model able to produce more sensible motivations for why the subject VERBED the object if the verbs have stronger causality biases? For this study we took care to avoid human raters being biased by obscenities and disfluencies generated by the model.

1 Introduction

This paper is a case study, highlighting different ways to analyse the linguistic abilities of a language model, with respect to an established linguistic phenomenon, namely Implicit causality (IC) bias (Hartshorne, 2014). Speakers associate either the subject or the object of a verb with the cause of the state or event described by that verb. For example, the verb *frighten* is a *subject-biased* verb because native speakers of English tend to see the subject as the cause of the frightening event. Thus, given a main clause like in (1a), participants in a sentence completion task would tend to provide a reason referring to the subject, as in (1b).

- (1) a. [_{MAIN CLAUSE} John scared Mary because ...]
b. [_{REASON} he put on a Halloween costume.]

Earlier work by Upadhye et al. (2020) and Davis and van Schijndel (2020, 2021) investigated the extent to which language models are able to capture native speakers' IC biases. This paper aims to reproduce some of their earlier results, using GPT-2 (Radford et al., 2019) as an example. Using the

same sentence completion task as Davis and van Schijndel (2020), we further investigate how bias, subject gender, and verb frequency influence the behavior of this model (§3). Next, we will more thoroughly assess the quality of the completions generated by GPT-2 (§4), asking: Does the model's performance hold up to further scrutiny?

Why GPT-2? Although it is neither the most recent, nor the best performing open-source language model around (see Black et al. 2022 for alternatives), GPT-2 is still a very popular choice for many researchers and practitioners. (See <https://huggingface.co/gpt2> for statistics.) This popularity is at least partly due to its size, as the model can be run and fine-tuned on consumer hardware. The model's popularity means that studying its capabilities and limitations may be more impactful (at least in the short term) than studying the capabilities and limitations of larger but less accessible systems. For us, GPT-2 offers the right balance of complexity and efficiency; as shown by Upadhye et al. (2020), its outputs are good enough to have a meaningful discussion about the assessment of language model performance, without requiring a large (and expensive) computational infrastructure.

Contributions Next to the value of reproducing earlier work, and providing further details on the IC-related behaviour of GPT-2, the main innovation of this paper is the controlled assessment of output quality. We took great care to separate issues with the fluency and offensiveness of the output from the content of the generated continuations. (See Section 6 for the limitations of this study.) This not only makes the task less harmful to our participants, but it also increases their focus on our construct of interest: The reasoning abilities of language models. Our code and data from this paper is available at <https://github.com/hienhuynhtdn/GPT2andImplicitCausality/>.

2 Data

We present two studies. Study 1 investigates next-word surprisal values, and Study 2 looks at continuations that are generated based on a prompt. For both our studies, we provide the model with input sentences of the following form:

(2) SUBJECT VERB-*ed* OBJECT *because* . . .

The verbs are derived from a list of 246 IC verbs compiled by [Ferstl et al. \(2011\)](#), who also provide a *bias score* for each of these verbs. This score is derived from a human experiment, where participants were asked to complete sentences like (1a). The human bias scores range from -100 (i.e., all valid continuations produced by respondents in Ferstl et al.’s experiment uniquely referred to objects of the preceding clauses) to 100 (i.e., all valid continuations referred to subjects of the preceding clauses).

The subjects and objects are provided by [Davis and van Schijndel \(2020\)](#), who produced a list of 14 noun pairs that are grammatically male and female (e.g. *man, woman* or *brother, sister*). A combination of the nouns and verbs, our set of stimuli consists of 6888 examples (246 verbs \times 14 pairs of gender-mismatched nouns \times 2 subject genders).

3 Study 1: IC-bias and pronoun use

First, we investigate whether verbs in GPT-2 possess the same subject/object bias as in the human experiment described above.

Set-up To test the hypothesis, we use the approach from [Davis and van Schijndel \(2021\)](#) and checked for each prompt whether the model assigned a lower surprisal to a male or female pronoun (i.e. *he* or *she*).¹ If GPT-2 captures the IC bias, then subject-biased verbs with a female subject should lead the model to produce lower surprisal values for *she*. Since each noun pair is used in both orders (either a male or a female noun in subject position), we have a perfectly balanced dataset.

Results Figure 1 shows the results split by subject gender and bias scores from [Ferstl et al. \(2011\)](#). GPT-2 generally picks up on the subject or object bias of the verb. The gender produces more subject-based explanations if the verb’s IC bias is more subject-biased. There is only one exception: Performance for sentences with both subject-biased

¹Next-word surprisal is estimated for a target upcoming word by taking the inverse log of this word’s probability: $surprisal(w_t) = -\log P(w_t | w_1 \dots w_{t-1})$




Figure 1: Heatmap table showing the percentage of outputs matching the subject or object bias of the verb, with scores separated by subject gender.

verbs and female subjects is at chance level for all bias scores.

IC Bias and Gender. We used the `lme4` ([Bates et al., 2015](#)) and `lmerTest` package in R ([Kuznetsova et al., 2017](#)) to carry out a mixed effects regression analysis of the relationship between the bias scores and GPT-2’s subject-preference scores, corresponding to the difference between the surprisal for the object-congruent pronoun and the surprisal of the subject-congruent pronoun. Details about our statistical analyses are provided in Appendix B. There were significant fixed effects of human bias score ($b = 0.003$, $SE = 0.0001$, $p < .001$) and of the interaction between human bias score and subject gender ($b = 0.0014$, $SE = 0.0002$, $p < .001$). This confirms our observations from Figure 1. The effect of subject gender was found to be not significant ($b = -0.0103$, $SE = 0.012$, $p = .374$).²

Verb frequency and model performance. We then investigated whether verb frequency is positively correlated with the language model performance. For more frequent verbs, we hypothesize that the model more closely matches the human subject/object bias. To test this hypothesis, we regressed the squared errors from the previous mixed effects model to the log-transformed word frequencies of verbs used in the stimuli. As a proxy for verb frequency in the training corpus, we used the

²In addition, the coefficient of determination of the model was calculated based on the method developed by [Nakagawa and Schielzeth \(2013\)](#) using the `MuMIn` library in R ([Bartoń, 2022](#)). Approximately 14% of the variance in the GPT-2’s subject-preference score could be explained by the fixed effects alone (marginal $R^2 = .142$) while 23% of the variance in the subject-preference score could be explained by both fixed and random effects (conditional $R^2 = .230$).

SUBTLEX-US word frequencies from Brysbaert and New (2009). We found a significant fixed effect of log-transformed word frequencies of the verbs in our materials ($b = -0.049$, $SE = 0.012$, $p < .01$). This supports our hypothesis: More data leads to a better approximation of human IC bias, in terms of the preference for either male or female pronouns.

4 Study 2: Assessing continuations

We now turn to continuations generated by GPT-2 based on the prompts described in Section 2. Our results above suggest that GPT-2 can generate continuations according to the causality bias pattern, given enough data. Based on this result, we now hypothesize that such continuations are better when the IC pattern is clearly present in the data. The stronger the IC bias is, the clearer the pattern of continuations in the training set. Of course, this task is much more difficult than simply generating the right pronouns. We now want to see whether the continuations actually *make sense* in the eyes of human judges. To this end, we collected human ratings for a carefully controlled subset of the continuations generated by GPT-2 (see Appendix C for details). IRB approval was obtained prior to this study. (Ethical considerations in Appendix A.)

Participants We used the Prolific participant pool to recruit 75 participants for the sentence rating task. Our items were spread across 25 different lists, and each participant was only allowed to provide ratings for one list. We restricted potential participants to native speakers of English, either from the UK or from the USA. After assessing response quality, we recruited five additional participants, to obtain three reliable judgments per item.

Task and target construct Each continuation was assessed by three different participants, who judged whether the continuations were reasonable, given the prompt. We set up our experiment as a rating task, where each participant indicated for a list of 40 items whether they agreed with the statement that the continuation was ‘reasonable.’ Participants could indicate their agreement on a five-point Likert scale, ranging from ‘Strongly Agree’ to ‘Strongly Disagree.’ With the addition of some examples in our task description (see Appendix G), we targeted our participants’ intuitions for what makes a good reason to do something. As we will discuss below, we aimed to avoid any influence of

the form of the output as much as possible.³

Prompt selection Due to financial limitations, we were not able to obtain ratings for all 6888 continuations. Following the recommendations from van Miltenburg et al. (2021), we used a stratified sampling approach. We selected the 5 most frequent noun pairs, and the 10 most frequent verbs for each of the 10 different bias levels (as illustrated in Figure 1). This selection gives us a sense of the upper bound performance with respect to continuation quality. Frequency was again determined using the SUBTLEX-US data (Brysbaert and New, 2009). Prompts were constructed in the same way as before (see Ex. 2), with each noun pair being presented in both orders. This yields 10 verbs \times 10 bias levels \times 5 nouns \times 2 orders = 1000 prompts.

Data preparation We used GPT-2 to generate continuations for each prompt. As shown earlier, the problem with these continuations is that they may be offensive or contain disfluencies (most notably repetition, see Fu et al. 2021). This creates two problems: (i) Offensive output may cause psychological harm for our participants, and (ii) offensiveness and disfluencies may lower the reliability of the ratings, if participants consistently provide lower scores for offensive/disfluent outputs (even though they may be consistent with the prompt). To prevent harm, and to avoid noise in the ratings, we took the following approach:

1. If the output is offensive, select a different noun pair from the 9 remaining pairs to generate a non-offensive alternative continuation.
2. If the output contains repetition, manually remove repeated elements from the sentence, so that the core content of the continuation remains largely unchanged.

Reliability To assess annotator reliability, we used a leave-one-out approach to correlate each participant’s scores with the mean scores of the two other participants who rated the same responses. Initial correlations ranged between 0.17 and 0.84. Five annotators scored below our cutoff of 0.4, and thus we recruited five more participants. After re-computing the reliability scores, we kept only the

³Given the terminological confusion in the field of Natural Language Generation (Howcroft et al., 2020), Belz et al. (2020) developed a categorization system for evaluation criteria in NLG. In terms of their taxonomy, our informal notion of ‘reasonable continuation’ clearly focuses on the content of the output, but the frame of reference is harder to define. If we look at the completed sentence as a whole, the sentence is evaluated in its own right, and so it is a question of *Coherence*.

Rating	1	2	3	4	5
Raw counts	689	571	333	782	625
% of ratings	23	19	11	26	21
% avg rating \geq rating	100	79	55	32	7

Table 1: Distribution of the ratings. Ratings correspond to a Likert scale, where 1=Strongly disagree, 2=Somewhat disagree, 3=Neither agree nor disagree, 4=Somewhat agree, 5=Strongly agree.

three highest-scoring participants per task. This way, we obtained a mean score of 0.64, with a standard deviation of 0.11. This is a strong correlation, considering the subjective nature of the task.

Continuation quality Table 1 shows the ratings in three different ways. The top row shows the (bimodal) raw score distribution: Ratings tend to be either negative or positive, with relatively few ratings in the middle of the scale. The second row provides the same values as percentages, as a guide for the third row. The third row shows the percentage of continuations for which the average rating is greater than or equal to a given rating. For example: Only 32% of all the continuations have an average rating greater than or equal to 4.

A total of 72 sentences was only rated at the lowest level. Table 2 (next page) provides examples of low-quality categorisations, with a rough categorisation. The three error categories are:

1. Non-sensical: Continuations that do not make sense as a reason for anything.
2. Invalid reason: Continuations that provide a reason, but the reason is not applicable.
3. Subject-object reversal: Continuations that would have made sense if subject and object were reversed.

These categories are not mutually exclusive (and fairly subjective) because it is hard to pin down what makes for a good/bad continuation. Nevertheless, it is clear that the example continuations in Table 2 provide poor reasons indeed. Future research could carry out a more systematic error analysis (along the lines of van Miltenburg et al. 2021), and present the distribution of the different kinds of erroneous continuations.

There were 69 continuations that always received the highest rating seem to conform to the human IC bias. One random example is ‘The woman thanked the man because he was a good man.’; the verb ‘to thank’ is strongly object-biased (raw subject bias score -92).

Continuation diversity Besides continuation quality, we also observe low continuation diversity. Table 3 shows the five most frequent continuations for our prompts, split by subject gender. It is clear from the table that the continuations generated by GPT-2 are very repetitive, and tend to be generic without any specific details. So there are no examples like (3) in the generated continuations:

- (3) The woman thanked the man because he gave her a nice book for her thirty-seventh birthday.

Explaining model performance We again used a mixed effects model to analyze our results. Our aim is to explain GPT-2’s performance (i.e., how reasonable the continuation is) in terms of verb frequency and absolute IC bias (which only looks at strength of the bias).⁴ Full details about the model and model fitting are provided in Appendix D.

Subject gender and its interactions were considered, but did not significantly improve model fit, and were dropped. Absolute IC bias has a small but positive effect on the rating ($b = 0.003$, $SE = 0.001$, $p = 0.015$). The effect of frequency is negative ($b = -0.121$, $SE = 0.038$, $p = 0.002$). There is also a negative interaction of bias and frequency ($b = -0.006$, $SE = 0.001$, $p < .001$), indicating that the positive effect of bias diminishes for higher frequency verbs. Thus, while we concluded in Study 1 that the accuracy of GPT-2 with respect to subject/object bias increases as the verb becomes more frequent, we now find that higher frequency does not give us more reasonable continuations. The (small) positive effect of absolute IC bias on the ratings can be explained by the intuition that the IC bias pattern is likely more clearly present in the training data for verbs that are more strongly biased. This makes it easier to pick up the pattern.

5 Discussion

The difference between Study 1 and Study 2 indicates a qualitative difference between generating pronouns and providing explanations: The latter requires higher-level reasoning which may not be present in language models like GPT-2 (Bender and Koller, 2020). Though the fact that GPT-2

⁴IC bias and log-transformed frequency values were determined as for Study 1. IC bias scores were made absolute; the mean absolute IC bias is 49.3 (SD 27.5, range [2, 92]). The mean (untransformed) frequency per million is 45.4 (SD 90.7, range [0.02, 502.27]). The correlation between IC bias and frequency was not significant and low (Pearson’s $r = -0.02$, $p = .35$).

Error category	Examples
Non-sensical	The girl hit the boy because he was too young to be a boy. The girl chased the boy because he was wearing a black hoodie and a black hoodie with a black hoodie on.
Invalid reason	The mother approached the father because she was afraid of him. The mother scared the father because he was a little too big for her.
Subject-object reversal	The woman affected the man because she was afraid of him. The woman surprised the man because he was wearing a black suit and a black tie.

Table 2: Rough categorisation of examples of low-quality continuations

Continuations	Male subject	Frequency	Continuations	Female subject	Frequency
he was afraid of her		521	she was afraid of him.		476
she was a woman		102	he was a good man and he was a good man		166
he was a good man and he was a good man		68	he was a good man		124
she was a good girl		64	he was a good boy		92
she was a woman, and he was a man		60	she was a woman		57

Table 3: Most frequent continuations for the prompts in Study 2, split by subject.

follows the implicit causality bias in pronoun selection is at least compatible with knowledge of causality, the quality of the continuations suggests the pronoun selection is based on superficial heuristics rather than a deep understanding of language (also discussed as *fast* versus *slow*; see Choudhury et al. 2022; Kahneman 2011). Although existing suites for LM evaluation (e.g. Ettinger 2020) are useful, slower forms of assessment (such as human evaluation) are helpful to tease out this difference.

6 Limitations

The main limitation of our paper is that we focused on only one language model (GPT-2), and only in one language (English). So while our findings provide insights into the capacities of the English GPT-2, they cannot be generalised to other language models or other languages (which is also illustrated by Davis and van Schijndel (2021)). Our main contribution is methodological, namely exploring how to assess the linguistic capacities of language models. Because our approach treats GPT-2 (mostly) as a black box, our analysis can easily be applied to other models as well.

For Study 2, a negative effect of verb frequency on the quality of generated continuations was found. As the materials used for the rating study were not gathered to cover the full range of frequencies, this pattern should not be generalized and may reflect a hidden effect. The negative interaction with IC bias suggests the same. The model fit shows an imperfect fit, but without one clear deviation from lin-

earity. Hence, our findings may well be explained better with more appropriate independent variables.

7 Conclusion

This paper showed two different ways to assess the linguistic capacity of a language model (GPT-2), using implicit causality as a case study. The techniques used above can be applied in a black box setting, without the need to look at the internals of the model. We hope that this paper is useful for others aiming to assess the ability of other language models to capture different linguistic phenomena as well. Our findings also showed that automatic assessment methods may not be enough to determine whether semantic phenomena like implicit causality are learned by a language model. Human evaluation remains a necessary complement to automatic evaluation (van der Lee et al., 2021). Our paper shows one way to do this without participants being influenced by factors like grammaticality and offensiveness of the output.

Acknowledgments

We thank all participants of our crowdsourcing experiment. This study was approved by the Research Ethics and Data Management committee at the Tilburg School for Humanities and Digital Sciences (TSHD) (code: 2019.40ab). We thank the committee for their feedback. The second author acknowledges the support of the Dutch Research Council (NWO) as part of the InDeep consortium (NWA.1292.19.399).

References

- Kamil Bartoń. 2022. *MuMin: Multi-Model Inference*. R package version 1.46.0.
- Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. [Fitting linear mixed-effects models using lme4](#). *Journal of Statistical Software*, 67(1):1–48.
- Anya Belz, Simon Mille, and David M. Howcroft. 2020. [Disentangling the properties of human evaluation methods: A classification system to support comparability, meta-evaluation and reproducibility testing](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 183–194, Dublin, Ireland. Association for Computational Linguistics.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, Michael Pieler, Usven Sai Prashanth, Shivanshu Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. 2022. [GPT-NeoX-20B: An open-source autoregressive language model](#). In *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 95–136, virtual+Dublin. Association for Computational Linguistics.
- Marc Brysbaert and Boris New. 2009. Moving beyond kučera and francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for american english. *Behavior research methods*, 41(4):977–990.
- Sagnik Ray Choudhury, Anna Rogers, and Isabelle Augenstein. 2022. [Machine reading, fast and slow: When do models "understand" language?](#)
- Forrest Davis and Marten van Schijndel. 2020. [Discourse structure interacts with reference but not syntax in neural language models](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 396–407, Online. Association for Computational Linguistics.
- Forrest Davis and Marten van Schijndel. 2021. [Uncovers constraint-based behavior in neural models via targeted fine-tuning](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1159–1171, Online. Association for Computational Linguistics.
- Leon Derczynski, Hannah Rose Kirk, Abeba Birhane, and Bertie Vidgen. 2022. [Handling and Presenting Harmful Text](#). *arXiv e-prints*, page arXiv:2204.14256.
- Allyson Ettinger. 2020. [What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Evelyn C Ferstl, Alan Garnham, and Christina Manouilidou. 2011. [Implicit causality bias in english: A corpus of 300 verbs](#). *Behavior Research Methods*, 43(1):124–135.
- Karën Fort, Gilles Adda, and K. Bretonnel Cohen. 2011. [Last words: Amazon Mechanical Turk: Gold mine or coal mine?](#) *Computational Linguistics*, 37(2):413–420.
- Zihao Fu, Wai Lam, Anthony Man-Cho So, and Bei Shi. 2021. [A theoretical analysis of the repetition problem in text generation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12848–12856.
- Florian Hartig. 2022. *DHARMA: Residual Diagnostics for Hierarchical (Multi-Level / Mixed) Regression Models*. R package version 0.4.5.
- Joshua K. Hartshorne. 2014. [What is implicit causality?](#) *Language, Cognition and Neuroscience*, 29(7):804–824.
- Tatsunori B. Hashimoto, Hugh Zhang, and Percy Liang. 2019. [Unifying human and statistical evaluation for natural language generation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1689–1701, Minneapolis, Minnesota. Association for Computational Linguistics.
- David M. Howcroft, Anya Belz, Miruna-Adriana Clinciu, Dimitra Gkatzia, Sadid A. Hasan, Saad Mahamood, Simon Mille, Emiel van Miltenburg, Sashank Santhanam, and Verena Rieser. 2020. [Twenty years of confusion in human evaluation: NLG needs evaluation sheets and standardised definitions](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 169–182, Dublin, Ireland. Association for Computational Linguistics.
- Daniel Kahneman. 2011. *Thinking, fast and slow*. Farrar, Straus and Giroux, New York.

- Alexandra Kuznetsova, Per B. Brockhoff, and Rune H. B. Christensen. 2017. [Imertest package: Tests in linear mixed effects models](#). *Journal of Statistical Software*, 82(13):1–26.
- Emiel van Miltenburg, Miruna Clinciu, Ondřej Dušek, Dimitra Gkatzia, Stephanie Inglis, Leo Leppänen, Saad Mahamood, Emma Manning, Stephanie Schoch, Craig Thomson, and Luou Wen. 2021. [Underreporting of errors in NLG output, and what to do about it](#). In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 140–153, Aberdeen, Scotland, UK. Association for Computational Linguistics.
- Emiel van Miltenburg, Desmond Elliott, and Piek Vossen. 2018. [Measuring the diversity of automatic image descriptions](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1730–1741, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Shinichi Nakagawa and Holger Schielzeth. 2013. A general and simple method for obtaining r^2 from generalized linear mixed-effects models. *Methods in ecology and evolution*, 4(2):133–142.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). Technical report, OpenAI.
- Shiva Upadhye, Leon Bergen, and Andrew Kehler. 2020. [Predicting reference: What do language models learn about discourse models?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 977–982, Online. Association for Computational Linguistics.
- Chris van der Lee, Albert Gatt, Emiel van Miltenburg, and Emiel Krahmer. 2021. [Human evaluation of automatically generated text: Current trends and best practice guidelines](#). *Computer Speech & Language*, 67:101151.
- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, Zac Kenton, Sasha Brown, Will Hawkins, Tom Stepleton, Courtney Biles, Abeba Birhane, Julia Haas, Laura Rimell, Lisa Anne Hendricks, William S. Isaac, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2021. [Ethical and social risks of harm from language models](#). *CoRR*, abs/2112.04359.
- Hugh Zhang, Daniel Duckworth, Daphne Ippolito, and Arvind Neelakantan. 2021. [Trading off diversity and quality in natural language generation](#). In *Proceedings of the Workshop on Human Evaluation of NLP Systems (HumEval)*, pages 25–33, Online. Association for Computational Linguistics.

A Ethical considerations

Because our study deals with human subjects, we first obtained ethical approval from our IRB. We describe our considerations below.

A.1 Information letter and informed consent

Our IRB mandates the use of a separate information letter (Appendix E) and informed consent form (Appendix F). With the information letter, we give a general description of the study, and provide an indication of potential risks and benefits of the study. The informed consent form is provided separately to prevent information overload, and to ensure that participants know what they are agreeing to, if they decide to take part in our study.

A.2 Crowdsourcing and payment

Crowdsourcing has been criticized for its potentially exploitative nature (Fort et al., 2011). We explicitly frame our task as an *experiment* with *human participants*, rather than a *human intelligence task* with *crowdworkers*, and apply the same considerations and protections as for lab experiments. Nevertheless, it is still work, and work needs to be paid. Based on experience, we expected participants to spend roughly 15 minutes on the task, and set the compensation to £2.40, which amounts to £9.60/hour; 10 cents above the current UK minimum wage.⁵ In the end, the vast majority of our participants spent less time than expected on our task (time in mm:ss format: range: 03:06–16:18, median: 06:18, mean: 06:59, SD: 02:55.).

All participants were compensated for their time, including those providing low-quality responses.

A.3 Offensive material

We wanted to avoid confronting our participants with profanity or otherwise potentially harmful language. We manually identified potentially offensive continuations generated by the model, and replaced harmful outputs with alternative continuations generated for different prompts. We considered continuations potentially harmful if they contained profanity or made reference to religion, violence, or sexual acts. All originally generated sentences and their replacements can be found in the GitHub repository associated with our paper.

⁵See <https://www.gov.uk/government/publications/the-national-minimum-wage-in-2022>

A.4 Language model-related harms

Language models are associated with several different harms (Bender et al., 2021; Weidinger et al., 2021), but these harms also depend on the task at hand. For example, since we used a pretrained model, our study did not incur any additional training costs. And as described above, since no one other than the authors were directly exposed to the model’s output, we could prevent our participants from seeing harmful or toxic content. Thus we are mostly left with inference costs, which are relatively low, since GPT-2 can run on a personal computer.

A.5 Intended use of this work

This study serves two purposes: (i) To explore the ability of a language model (GPT-2) to capture native speakers’ intuitions about implicit causality, and (ii) to develop an evaluation methodology that isolates coherence of the responses from other factors like offensiveness and (un)wellformedness. We do not wish to make any claims about the cognitive capacity of language models in general, nor do we want to claim that GPT-2 can somehow reason about the world. We just want to see whether the model can generate outputs that follow earlier observations about implicit causality, and that are internally consistent. Follow-up studies in the spirit of this work are encouraged.

A.6 Licensing

All resources used for this study were developed for research purposes, but not all materials have a clearly indicated license. GPT-2 is provided under the MIT license.⁶ The work by Davis and van Schijndel (2020, 2021) is provided on GitHub without any license⁷ and both Brysbaert and New (2009) and Ferstl et al. (2011) published their work in the *Behavior Research Methods* journal without a clear license, but with a clear intention for their work to be used for research purposes. We thus conclude that academic use of these resources is warranted.

B Study 1: Statistical analysis

As noted in Section 3, we used the lme4 (Bates et al., 2015) and lmerTest packages in R (Kuznetsova et al., 2017) to carry out a mixed

⁶<https://huggingface.co/gpt2>

⁷<https://github.com/forrestdavis/ImplicitCausality>

effects regression analysis of the relationship between the model’s subject-preference score and human bias score of every IC verb, reducing the complexity by removing terms that do not significantly improve fit.

As fixed effects, human bias score, subject gender of the stimulus sentence-fragments (with two levels, namely male or female) and their interaction were included. Moreover, we included item, which indicated the pairs of antonymous nouns used in the stimulus sentences as subjects and objects, as a random effect, and added by-item random slopes for the effects of human bias score, subject gender and their interaction. However, the fitting of the full model including all fixed and random effects failed to converge. Therefore, the by-item random slopes for the effects of human bias, subject gender and their interaction were removed. Our final model is the following in R notation:

(4) subject-preference score ~ human IC bias score
* subject gender + (1|item)

We regressed the squared errors obtained from the linear mixed effects regression model we performed in the previous step on the log-transformed word frequencies of IC verbs used in the stimuli. As random effects, item was also entered into the model, and by-item random slope for the effect of log-transformed frequency was added.

C Study 2: Continuations

While Study 1 looked into the surprisal values for the generation of male/female pronouns, this study looks at continuations themselves.

C.1 Consistency check

We first checked whether the IC bias pattern continues to be present in the full continuations. In all but 176 cases, the model generated a gendered pronoun. These cases were coded manually, to determine whether the continuation referred to the subject or the object. Figure 2 shows the proportion of continuations referring to either the subject or the object of the prompt (Y-axis), split by the bias score (X-axis). We see the same trends as in the first study: Overall we see that references to the subject increase as the bias score increases. This trend also holds when we split the prompts by subject gender, but for the female subjects the proportion of references to the subject never exceeds 50%.

C.2 Patterns

We then inspected the frequencies of the continuations. We observe that more than 20% of the outputs is repeated more than 100 times. This lack of diversity is a common issue in (neural) Natural Language Generation (e.g. [van Miltenburg et al. 2018](#); [Hashimoto et al. 2019](#); [Zhang et al. 2021](#)). It is not necessarily a problem at the individual level (the continuation may be a bit generic but still appropriate for the given context), but at the corpus level these ‘one size fits all’ continuations are shortcuts that prevent more varied outputs.⁸

C.3 Offensive output

Looking over the generated outputs, there were several occasions where the model generated offensive continuations, including instances of sexism, racism, and misogyny. On top of this, the model outputs also contained sexually explicit words, and some continuations described acts of violence. Following the recommendations from [Derczynski et al. \(2022\)](#), we do not provide any examples in this paper. As discussed in Section 4, we also removed (potentially) offensive continuations from our human rating experiment. For transparency reasons, we do provide those sentences in our GitHub repository.

D Study 2: Statistical analysis

A full linear model was built with the factors subject gender, absolute IC bias (centered), log-transformed word frequency (z-scored) and their interactions as fixed effects, and random effects for participant (intercept, and slopes for the three main effects), sentence (intercept), and subject/object pair (intercept). The model was then subjected to the step function of `lmerTest`, which removes insignificant components. Though the full model and some of the first reductions of it did not converge, this procedure is still appropriate (a model that does not converge should not be used, and as more data is not available, the model should be simplified). The reduced final model, that did converge, is, in R syntax:

⁸We might also question whether it is possible at all to properly assess the cognitive capacity of a model that keeps using such shortcuts (which may be seen as cheating). As an alternative, we can imagine a two-stage process where researchers first generate an unrestricted set of continuations, and then force the model to avoid common continuations.

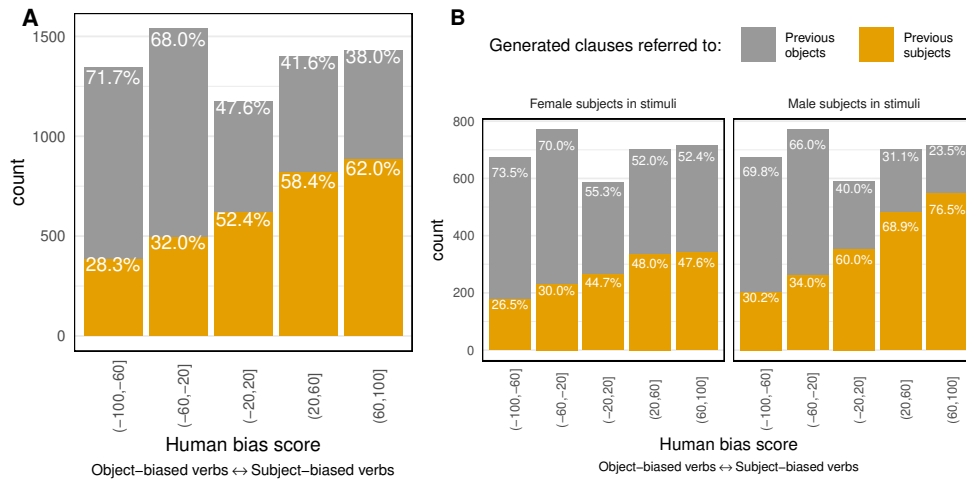


Figure 2: Percentage of continuations referring to subjects or objects by bias scores of verbs. Panel A shows overall proportion, Panel B shows proportion split by subject gender

$$(5) \text{ rating} \sim \text{abs(IC bias_abs_c)} * \text{frequency} + (1 | \text{sentence_ID}) + (1 | \text{participant}) + (1 | \text{item})$$

The DHARMA package (Hartig, 2022) was used to assess model fit. Though the residuals are not normally distributed, the deviations show no clear pattern. To avoid spurious conclusions, we corroborated the significant estimates by checking if their 95% confidence intervals included 0, which they did not. Hence, treating ratings as a continuous numeric variable was not problematic.




Figure 4: Ratings by absolute value of human IC bias, separated for four different levels of verb frequency.




Figure 3: Ratings by absolute value of human IC bias.

E Information letter

Title: Assessing computer-generated texts

Introduction

We invite you to take part in our study to assess computer-generated texts. This study is part of a larger project to see how good or bad computers are at producing or understanding human language, such as English. In this study you will be asked to rate the quality of

computer-generated texts/sentences. We will use this information to see what the computer is good at, and to see where it can still be improved.

What do I have to do?

As mentioned above, you will be asked to rate the quality of computer-generated texts. In this study, you will be asked to read 40 short sentences, and to provide your judgment. We are interested in your intuition as a native speaker of English, so you don't need to think too long about it.

Expected duration

We expect this study to take about fifteen minutes of your time. Other than this, we do not foresee any risks associated with this study. On the positive side, your participation will improve our understanding of the language capacity of modern computer models.

Ethics and rights

This study was approved by the Research Ethics and Data Management Committee (REDC) at Tilburg University.

Your participation is completely voluntary. Your consent to participate generally applies for the duration of this study. However, you have the right to decline to participate and withdraw from the research once participation has begun, without any negative consequences, and without providing any explanation.

Your participation is completely anonymous. We will not store any identifying information, so your answers cannot be traced back to you. We only see the demographic information that Prolific provides. Do let us know if you would like to have a copy of your responses, and we will try to obtain them based on the Prolific ID.

Use of data

Your responses will be used for the current study, and possible follow-up studies in the future. This means that the data will be presented in research articles, that are publicly available. For full transparency, we will also publicly share the anonymised individual responses. As such, they will be stored indefinitely.

Contact

If you have any questions, or if you would like to learn more about this study, please contact C.W.J.vanMiltenburg@tilburguniversity.edu for any of your questions.

Ethics approval

If you have any remarks or complaints regarding this research, you may also contact the “Research Ethics and Data Management Committee” of Tilburg School of Humanities and Digital Sciences via tshd.redc@tilburguniversity.edu.

F Informed consent

By agreeing to this consent form, you confirm that you have read the study description and that you have been offered the opportunity to ask questions (via email). Remember that your participation is voluntary, and that you have the right to decline to participate and withdraw from the research once participation has begun, without

any negative consequences, and without providing any explanation.

I hereby give permission to:

- Store my anonymised responses to this survey.
- Analyse the anonymised data (both manually and automatically through statistical software).
- Make the responses to this survey publicly available upon completion of the study.

Yes ⇒ continue to the survey.

No ⇒ continue to the end of the survey.

G Task instructions

Task instructions

All questions below are of the same form. You will see the start of a sentence on the first line, and a continuation generated by a computer model on the second line. Your job is to assess the quality of the continuation on the second line.

Example of a reasonable continuation:

For the following sentence: *The clown startled the girl because*

A reasonable continuation would be: *his make-up was scary.*

- Strongly agree
- Somewhat agree
- Neither agree nor disagree
- Somewhat disagree
- Strongly disagree

This is a reasonable continuation because scary make-up can cause someone to be startled. So here you would answer *Strongly agree*.

Example of a less reasonable continuation:

For the following sentence: *The clown startled the girl because*

A reasonable continuation would be: *she liked him.*

- Strongly agree
- Somewhat agree
- Neither agree nor disagree
- Somewhat disagree
- Strongly disagree

This is a less reasonable continuation, because being liked by someone is generally not a reason to startle them. So here you would answer one of the disagree options.