

ACKNOWLEDGEMENTS

We would like to thank Mina Lee, Nick Lourie, Naomi Saphra, Richard Pang, and Nitish Joshi for their input at various stages of the project. We would also like to thank Abby Rabinowitz and Alexander Landfair from the NYU Expository Writing Program for valuable discussions during the planning stages. We would finally like to acknowledge all the writers recruited for this project without whom this work would not have been possible. This work is supported by Open Philanthropy, AWS AI, and the National Science Foundation under Grant No. 1922658.

REPRODUCIBILITY STATEMENT

All data collection experiments were done between March and July 2023. For the sake of reproducibility, the essays with character-level logs as recorded by the interface along with all model suggestions presented to users will be released after the review period. This allows readers to replay entire writing sessions using the code released by Lee et al. (2022a). We also make public all the scripts to process the raw logs to perform the analyses run in this paper.

ETHICS STATEMENT

We detail all the procedures for user recruitment and remuneration in Appendix A.1. The plan for our user study was approved by the Institutional Review Board of our university. We obtained consent from all participants to share the essays publicly post the completion of the study. We detail further limitations of our study in Appendix D.

REFERENCES

- Nader Akoury, Shufan Wang, Josh Whiting, Stephen Hood, Nanyun Peng, and Mohit Iyyer. 2020. Storium: A dataset and evaluation platform for machine-in-the-loop story generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6470–6484.
- Kenneth C Arnold, Krysta Chauncey, and Krzysztof Z Gajos. 2020. Predictive text encourages predictable writing. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*, pages 128–138.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- 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*, pages 610–623.
- Advait Bhat, Saaket Agashe, Parth Oberoi, Niharika Mohile, Ravi Jangir, and Anirudha Joshi. 2023. Interacting with next-phrase suggestions: How suggestion systems aid and influence the cognitive processes of writing. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*, pages 436–452.
- Rishi Bommasani, Kathleen A Creel, Ananya Kumar, Dan Jurafsky, and Percy S Liang. 2022. Picking on the same person: Does algorithmic monoculture lead to outcome homogenization? *Advances in Neural Information Processing Systems*, 35:3663–3678.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dorottya Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter

- Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Nieves, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Daniel Buschek, Martin Zürn, and Malin Eiband. 2021. The impact of multiple parallel phrase suggestions on email input and composition behaviour of native and non-native english writers. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–13.
- Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, et al. 2023. Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint arXiv:2307.15217*.
- Tuhin Chakrabarty, Vishakh Padmakumar, and He He. 2022. [Help me write a poem: Instruction tuning as a vehicle for collaborative poetry writing](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6848–6863, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Quanze Chen, Chenyang Lei, Wei Xu, Ellie Pavlick, and Chris Callison-Burch. 2014. Poetry of the crowd: A human computation algorithm to convert prose into rhyming verse. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 2, pages 10–11.
- Elizabeth Clark, Anne Spencer Ross, Chenhao Tan, Yangfeng Ji, and Noah A Smith. 2018. Creative writing with a machine in the loop: Case studies on slogans and stories. In *23rd International Conference on Intelligent User Interfaces*, pages 329–340.
- Kathleen Creel and Deborah Hellman. 2022. The algorithmic leviathan: Arbitrariness, fairness, and opportunity in algorithmic decision-making systems. *Canadian Journal of Philosophy*, 52(1):26–43.
- E. Durmus, K. Nyugen, T. I. Liao, N. Schiefer, A. Askeel, A. Bakhtin, C. Chen, Z. Hatfield-Dodds, D. Hernandez, N. Joseph, L. Lovitt, S. McCandlish, O. Sikder, A. Tamkin, J. Thamkul, J. Kaplan, J. Clark, and D. Ganguli. 2023. Towards measuring the representation of subjective global opinions in language models. *arXiv preprint arXiv:2306.16388*.
- Sebastian Gehrmann, Elizabeth Clark, and Thibault Sellam. 2023. Repairing the cracked foundation: A survey of obstacles in evaluation practices for generated text. *Journal of Artificial Intelligence Research*, 77:103–166.
- Katy Ilonka Gero and Lydia B Chilton. 2019. Metaphoria: An algorithmic companion for metaphor creation. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pages 1–12.
- Marjan Ghazvininejad, Xing Shi, Jay Priyadarshi, and Kevin Knight. 2017. Hafez: an interactive poetry generation system. In *Proceedings of ACL 2017, System Demonstrations*, pages 43–48.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. News summarization and evaluation in the era of gpt-3. *arXiv e-prints*, pages arXiv–2209.

- Jeffrey T Hancock, Mor Naaman, and Karen Levy. 2020. Ai-mediated communication: Definition, research agenda, and ethical considerations. *Journal of Computer-Mediated Communication*, 25(1):89–100.
- Daphne Ippolito, Ann Yuan, Andy Coenen, and Sehmon Burnam. 2022. Creative writing with an ai-powered writing assistant: Perspectives from professional writers. *arXiv preprint arXiv:2211.05030*.
- Maurice Jakesch, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman. 2023. Co-writing with opinionated language models affects users’ views. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–15.
- J. Kleinberg and M. Raghavan. 2021. Algorithmic monoculture and social welfare. In *National Academy of Sciences*.
- Mina Lee, Percy Liang, and Qian Yang. 2022a. Coauthor: Designing a human-ai collaborative writing dataset for exploring language model capabilities. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–19.
- Mina Lee, Megha Srivastava, Amelia Hardy, John Thickstun, Esin Durmus, Ashwin Paranjape, Ines Gerard-Ursin, Xiang Lisa Li, Faisal Ladhak, Frieda Rong, et al. 2022b. Evaluating human-language model interaction. *arXiv preprint arXiv:2212.09746*.
- X. L. Li, A. Holtzman, D. Fried, P. Liang, J. Eisner, T. Hashimoto, L. Zettlemoyer, and M. Lewis. 2023. Contrastive decoding: Open-ended text generation as optimization. In *Association for Computational Linguistics (ACL)*.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2022. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*.
- Q Vera Liao and Ziang Xiao. 2023. Rethinking model evaluation as narrowing the socio-technical gap. *arXiv preprint arXiv:2306.03100*.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- C. Meister, T. Pimentel, G. Wiher, and R. Cotterell. 2023. Locally typical sampling. *Transactions of the Association for Computational Linguistics (TACL)*, 11.
- Piotr Mirowski, Kory W Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-writing screenplays and theatre scripts with language models: Evaluation by industry professionals. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–34.
- OpenAI. 2023. [Gpt-4 technical report](#).
- L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano, J. Leike, and R. Lowe. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Vishakh Padmakumar and He He. 2022. Machine-in-the-loop rewriting for creative image captioning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 573–586.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: A method for automatic evaluation of machine translation. In *Association for Computational Linguistics (ACL)*.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

- Melissa Roemmele and Andrew S Gordon. 2015. Creative help: A story writing assistant. In *Interactive Storytelling: 8th International Conference on Interactive Digital Storytelling, ICIDS 2015, Copenhagen, Denmark, November 30-December 4, 2015, Proceedings 8*, pages 81–92. Springer.
- Melissa Roemmele and Andrew S Gordon. 2018. Automated assistance for creative writing with an rnn language model. In *Proceedings of the 23rd international conference on intelligent user interfaces companion*, pages 1–2.
- Melissa Roemmele, Andrew S Gordon, and Reid Swanson. 2017. Evaluating story generation systems using automated linguistic analyses. In *SIGKDD 2017 Workshop on Machine Learning for Creativity*, pages 13–17.
- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cino Lee, Percy Liang, and Tatsunori Hashimoto. 2023. Whose opinions do language models reflect? *arXiv preprint arXiv:2303.17548*.
- Abigail See, Aneesh Pappu, Rohun Saxena, Akhila Yerukola, and Christopher D. Manning. 2019. [Do massively pretrained language models make better storytellers?](#) In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 843–861, Hong Kong, China. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur P Parikh. 2020. Bleurt: Learning robust metrics for text generation. *arXiv preprint arXiv:2004.04696*.
- Irene Solaiman, Zeerak Talat, William Agnew, Lama Ahmad, Dylan Baker, Su Lin Blodgett, Hal Daumé III, Jesse Dodge, Ellie Evans, Sara Hooker, et al. 2023. Evaluating the social impact of generative ai systems in systems and society. *arXiv preprint arXiv:2306.05949*.
- Z. Song, T. Cai, J. D. Lee, and W. J. Su. 2023. Reward collapse in aligning large language models. *arXiv preprint arXiv:2305.17608*.
- Simeng Sun, Ori Shapira, Ido Dagan, and Ani Nenkova. 2019. [How to compare summarizers without target length? pitfalls, solutions and re-examination of the neural summarization literature](#). In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 21–29, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ben Swanson, Kory Mathewson, Ben Pietrzak, Sherol Chen, and Monica Dinalescu. 2021. Story centaur: Large language model few shot learning as a creative writing tool. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 244–256.
- Reid Swanson and Andrew S Gordon. 2008. Say anything: A massively collaborative open domain story writing companion. In *Interactive Storytelling: First Joint International Conference on Interactive Digital Storytelling, ICIDS 2008 Erfurt, Germany, November 26-29, 2008 Proceedings 1*, pages 32–40. Springer.
- Reid Swanson and Andrew S Gordon. 2012. Say anything: Using textual case-based reasoning to enable open-domain interactive storytelling. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 2(3):1–35.
- R. Taori and T. B. Hashimoto. 2023. Data feedback loops: Model-driven amplification of dataset biases. In *International Conference on Machine Learning (ICML)*.
- Guy Tevet and Jonathan Berant. 2021. [Evaluating the evaluation of diversity in natural language generation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 326–346, Online. Association for Computational Linguistics.
- Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. [Wordcraft: Story writing with large language models](#). In *27th International Conference on Intelligent User Interfaces, IUI ’22*, page 841–852, New York, NY, USA. Association for Computing Machinery.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#).

Y. Zhu, S. Lu, L. Zheng, J. Guo, W. Zhang, J. Wang, and Y. Yu. 2018. Taxygen: A benchmarking platform for text generation models. *arXiv preprint arXiv:1802.01886*.

Jacob Ziv and Abraham Lempel. 1977. A universal algorithm for sequential data compression. *IEEE Transactions on information theory*, 23(3):337–343.

Jacob Ziv and Abraham Lempel. 1978. Compression of individual sequences via variable-rate coding. *IEEE transactions on Information Theory*, 24(5):530–536.

A TASK DETAILS

A.1 USER RECRUITMENT ON UPWORK

We recruit a total of 38 participants from Upwork, all of whom are native English speakers and have prior experience in writing or copyediting on the platform. Their prior experience varied from 5 to over 200 previously completed tasks on the platform. Our participants are exclusively from the United States and encompassing diverse racial backgrounds with an approximately equal distribution between genders. Each participant is asked to complete a session consisting of three essay writing tasks—one in each of the three settings and each on a different topic. In each session, the order of the three settings and the assignment of topics are randomized. We don’t enforce that users include a certain amount of machine-generated text as we want to study how collaborative writing in its most natural form differs from solo writing without model help. The average machine-written fraction is 0.32 for InstructGPT and 0.35 for GPT (Table 9) indicating that users obtain useful help from the models without over-relying on the same. The vast majority of essays have machine-written fractions ranging from 0.2 to 0.5 with a few outliers being as low as 0.1 and as high as 0.8 (Figure 13 in Appendix C). Users are allowed to edit their essays post-completion. Their objective is to convey their opinions effectively. We acknowledge that this might affect word usage patterns which again motivates the contribution of our proposed key point-based evaluation as these key points are less likely to be changed during post-editing. Each essay takes 15 minutes on average and hence a session usually takes the participants under one hour. Compensation was provided through hourly contracts, prorated to \$20 per hour. After each participant completed one session, their essays were reviewed manually by the authors of this work for the relevance to the essay prompt as well as coherence of writing. This process does not involve policing the actual content of the essays and we encourage our participants to express their opinions in detail. This was used to filter out a few participants whose sessions were reassigned. The remaining participants were then invited to complete two additional sessions, each with different topics. As a result, each of the 38 participants completed between 1 and 3 sessions. In total, we obtain 10 essays on each of the 10 topics for each setting, resulting in 100 essays per setting.

A.2 DECODING PARAMETERS ON OPENAI API

To avoid reduced diversity due to decoding, we use a decoding strategy biased towards higher diversity. Specifically, we sample continuations from both models following the “high randomness” setting from [Lee et al. \(2022a\)](#). The decoding parameters we use are:

- Engine: davinci for the base language model and text-da-vinci-003 for the feedback-tuned language model
- Response length (word piece): 30
- Temperature: 0.9
- Top P: 1
- Frequency penalty: 0.5
- Presence penalty: 0.5

Prompt Code	Prompt (Source URL)
school	What Are the Most Important Things Students Should Learn in School? In your opinion, what are the most important things students should learn in school? What is the most important thing you have learned in school? How has this knowledge affected your life? How do you think it will help your success in the future? (Link)
stereotype	What Stereotypical Characters Make You Cringe? What stereotypical characters in books, movies or television shows make you cringe and why? Would you ever not watch or read something because of its offensive portrayal of someone? (Link)
audiobook	Is Listening to a Book Just as Good as Reading It? Do you listen to audiobooks? What are the benefits, in your opinion, of listening instead of reading? Are there advantages to reading that cannot be gained by listening? Which method do you prefer? Why? (Link)
athletes	Should College Athletes Be Paid? Do you think college athletes should be paid? Or is a college scholarship and other non-monetary perks like the opportunity to play in front of cheering fans enough? What possible difficulties or downsides might there be in providing monetary compensation to players? (Link)
extremesports	Is It Selfish to Pursue Risky Sports Like Extreme Mountain Climbing? Some sports, like extreme mountain climbing, are dangerous. Since there are varying degrees of risk in most, if not all, sports (such as the possibility of concussions, broken bones and even death), how does one decide where the line might be drawn between what is reasonable and what is not? Are some sports simply too dangerous to be called a sport? (Link)
animal	Is It Wrong to Focus on Animal Welfare When Humans Are Suffering? Would you be surprised to hear that a study found that research subjects were more upset by stories of a dog beaten by a baseball bat than of an adult similarly beaten? Or that other researchers found that if forced to choose, 40 percent of people would save their pet dog over a foreign tourist. Why do you think many people are more empathetic toward the suffering of animals than that of people? In your opinion, is it wrong to focus on animal welfare when humans are suffering? Why do you think so? (Link)
news	Are We Being Bad Citizens If We Don't Keep Up With the News? In your opinion, are we being bad citizens if we don't keep up with the news? Do you think all people have some responsibility to know what is going on in the world? Does engaging with current events actually do anything at all? Why do you think the way you do? (Link)
mindfulness	Should Schools Teach Mindfulness? Have you ever tried mindfulness or meditation, practices that focus on the present moment and being aware of your thoughts, feelings and bodily sensations? If so, what was it like for you? If not, does it sound like something you'd like to try? Do you think that such practices have a place in schools? Why or why not? (Link)
screen	How Worried Should We Be About Screen Time During the Pandemic? The coronavirus pandemic ended the screen time debate: Screens won. We all now find ourselves on our screens for school, for work and for connecting with family and friends during this time of social distancing and increased isolation. But should we be worried about this excessive screen use right now? Or should we finally get over it and embrace the benefits of our digital devices? (Link)
dating	How Do You Think Technology Affects Dating? Have you had any experience with dating? Have you ever used dating apps? If so, what has it been like for you? If not, why not? How do you think technology — like apps, Netflix, social media and texting — affects dating and relationships? In your opinion, does it improve or worsen romantic interactions? (Link)

Table 5: List of prompts used in the experiments.

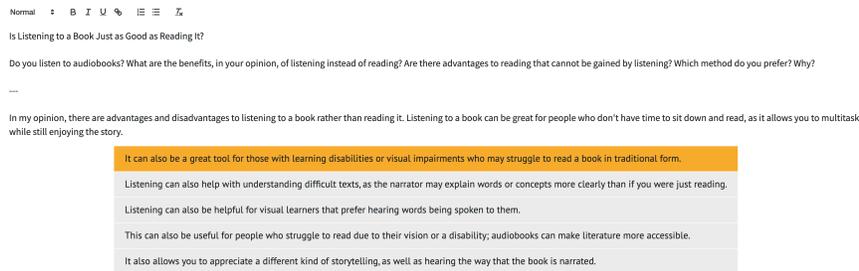


Figure 5: Text editor interface showing suggestions when requested by the user.

A.3 USER INSTRUCTIONS

Essay Requirements To complete an essay assignment successfully, you would need to write a short 3-passage essay given the prompt presented to you. The essay should reflect your opinion of the assigned topic and you’re writing an argumentative piece for why you feel that way. The expected length of the essays is around 300 words (3-4 short paragraphs of 3-4 sentences each). Each piece should take you between 10 and 15 minutes (as calculated from prior experiments). When writing with AI help, you are required to obtain model suggestions at least 5 times (we encourage you to do this more if you find it helpful, more the better). Some example responses are provided below (though these were written by non-experts so I’m sure you can identify flaws and improve on the style)

Instructions in Detail

- When you are completing a task, make sure you have a Session ID. If you don’t have this already, please send us a message on Upwork.
- View the spreadsheet of assignments and navigate to your assigned session. Each session will correspond to three rows in the spreadsheet, each with an accompanying URL.
- Each of these links corresponds to the three assignment essays you would have to complete in order to finish the task.
- If you click one of these, you will see the text editor pop up which looks like this which shows the prompt for the assigned essay.
- Out of the three assignments, please note that one is the essay that you will have to write without the assistance of the AI. Please complete the three assignments in the order provided.
- When you write with the AI, you will have the option to hit the TAB key on your keyboard and obtain suggestions from the model in the form of a dropdown like this. You are free to continue the writing process as you normally would, however, we encourage you to make use of model help when you are looking for ideas when you are stuck or just want to take a look at some possible continuations :)
- When you write with the AI, you need to request at least 5 suggestions from the model. You do NOT need to accept all of the suggestions, you are free to accept one then edit it, ask for suggestions again at the same time, or even just reject all of the model suggestions. We encourage you to use the model as much as possible :) The objective of the study is to understand the effect of the model, so more interaction is better.
- Once you complete the essay, make sure you hit ‘Save’ at the bottom of the screen to obtain a verification code on completion. We ask that you complete the writing in one sitting. We aim to record the entire writing process so please refrain from copy-pasting any text into the editor.

Essay

While I believe the concerns regarding children’s screen time are valid, I believe it is somewhat biased to not take this problem, which is a genuine issue right now, as an everyone problem, and not a student one. With that in mind, I do believe that for the majority of our generation it would not be as hard to disconnect with devices after the pandemic.

I know that I, along with many other teenagers, would like nothing more than to go back to school, play sports outside, meet new people and such. I doubt I could say the same for our older generations, because while parents ban teenagers from using their phones, they often spend many more hours on their devices. Cranky grandparents grumble about how ‘back in their day they lived without a mobile phone’ usually forget the point of their stories. Back in their day indeed.

We live in a different era, and to think that the same rules should apply to teenagers, especially during a pandemic, isn’t logical. Pinterest, Youtube, Twitter aren’t just places where teens while away their time. They are also places for new ideas, watching college lectures, and political discourse. Limiting screen time, while in the short run is helpful, is doubtful to do anything in the long run. Banning technology is not stopping kids from using it. It’s just stopping them from telling their parents about it.

I would like to end by saying, parents, please trust your teenager just a little bit more, and don’t be too worried about their screen time!

Keypoints

- The concerns regarding children’s screen time are valid
- The problem of screen time should be considered an everyone problem, not just a student one
- For most young people, it would not be hard to disconnect from devices after the pandemic
- Older generations often spend more time on their devices than teens, despite their complaints
- Social media can be used for educational and constructive purposes
- Limiting screen time may not be effective in the long run
- Parents should trust their teenager more and not worry too much about their screen time

Table 6: Example of summarization into keypoints

A.4 SUMMARIZATION INTO KEY POINTS

Motivated by recent encouraging results of using LLMs for zero-shot summarization (Goyal et al., 2022), we summarize each essay into a list of key points by prompting gpt-3.5-turbo. Table 6 contains a full example of an essay with the generated key points. We use a simple prompt, “Summarize this essay into a set of simple and distinct bullet points. Make sure that the bullet points cover all the information from the essay.”

B SIGNIFICANCE TESTING

B.1 PERMUTATION TEST FOR SIGNIFICANCE OF DIVERSITY RESULTS

To test for the significance of the results on diversity (Section 5.2), since we do not have multiple instances of essay corpora from each setting (Solo, GPT3 and InstructGPT), we employ three permutation tests between each pair of settings. A walkthrough example of the permutation test setup is as follows. We wish to evaluate the significance of the difference between the diversity scores of Solo and InstructGPT, as measured by the clustering of key points (Table 3) or by traditional lossless

compression algorithms (Table 13). We first calculate the statistic (difference between the diversity scores) on the collected corpora of essays. We then take the union of Solo and InstructGPT essays, randomly partition this into two equal sets, and recalculate the statistic on these two permuted sets of essays. This process is then repeated for 1000 different permutations. We then obtain the p-value of this two-tailed permutation test as the proportion of times the absolute value of the statistic calculated on the permuted data is greater than the statistic calculated on the observed data from the user study. This is then repeated for all pairs out of Solo, GPT3, and InstructGPT. Bold values in the results tables (Table 3, Table 13) indicate those instances where the p-value on the permutation test for a setup (InstructGPT) was significant at the 5% level over both other setups (GPT3 and Solo).

B.2 CHI-SQUARE TEST FOR SIGNIFICANCE FOR n -GRAM DISTRIBUTIONS

To confirm the significance of the difference between the categorical distributions of n -grams in the different setups, we employ a chi-square test on the count of occurrences. We evaluate if writing with InstructGPT results in a change of 5-gram usage in Figure 3 and Figure 4(b). To perform this test, we first identify and take the union of the 50 most frequently occurring 5-grams in each setup. We then obtain the counts of occurrences of each of these 5-grams in both corpora and perform a chi-square test for significance on these frequencies. A p-value < 0.05 results in a rejection of the null hypothesis and the observation of a significant difference in the categorical distributions. We perform these for all pairs out of Solo, GPT3, and InstructGPT and find that the repetition of n -grams when writing with InstructGPT is more similar by a significant margin compared to both other setups (Figure 3).

B.3 SIGNIFICANCE TESTING FOR RESULTS ON HOMOGENIZATION

As noted in Section 2.2 and Appendix A.1 we recruit a fixed set of users who wrote essays in each of the three setups in a randomized order. We ensure that they write essays on different topics in each setup to prevent a repetition of their opinions. To test for the significance of the change in homogenization across all topics between the three setups, we conducted independent samples t-tests in Figure 2. To further test the significance of this result, we observe that within each topic, each writer only submits an essay to one of the three setups. This matches the between-group experimental setup, albeit with less power as we only collected 10 essays per setup per topic. We compute the significance of the difference in homogenization scores per topic via an independent samples t-test. We report the p-values comparing each pair of setups within each topic at both the key point (Table 7) and essay level (Table 8). Comparisons with significance at the 5% level are marked with an asterisk. We observe that the InstructGPT setup exhibits higher homogenization over Solo writers with significance at the 5% level on 8 out of 10 topics at the raw essay level and 6 out of 10 at the key point level using both Rouge-L and BertScore. The same trend holds on GPT3 on 7 topics at the raw essay level and 5 at the key point level. We believe that this confirms the overall trend that writing with InstructGPT results in higher homogenization in the essays created.

C ADDITIONAL RESULTS

Reporting more basic statistics about the collected essays We aim to improve the comprehensiveness of our analysis by reporting additional metrics on the collected essays in Table 9.

- We compute basic statistics such as perplexity (via GPT2), average sentence length, and essay length on the essays collected. Predictably, writing with model-help reduces the perplexity of the essays which now incorporate suggestions sampled from an LM distribution. The total length of essays, in terms of word count, is roughly similar, with writers writing slightly shorter sentences when writing with InstructGPT (albeit with a slightly high standard error).
- Additionally, by calculating the average height of the dependency parse trees of the sentences in the essays, we observe that writing with InstructGPT results in sentences with fewer nested structures, indicative of slightly lower complexity.
- Writing with model help also results in fewer unique POS-Ngrams again providing some evidence that collaborative writing could result in more homogenized usage of language. The effect here is even between GPT3 and InstructGPT, a slight departure from the results

Significance Testing - Key point level							
Rouge-L				BertScore			
Topic	Solo vs InstructGPT	InstructGPT vs GPT3	GPT3 vs Solo	Topic	Solo vs InstructGPT	InstructGPT vs GPT3	GPT3 vs Solo
Stereotype	0.02*	0.34	0.01*	Stereotype	0.13	0.18	0.81
School	0.03*	0.05*	0.86	School	0.78	0.15	0.17
Mindfulness	0.00*	0.00*	0.03*	Mindfulness	0.00*	0.11	0.64
Athletes	0.24	0.14	0.53	Athletes	0.09	0.05*	0.67
Audiobook	0.73	0.03*	0.25	Audiobook	0.00*	0.06	0.01*
Animal Welfare	0.67	0.13	0.01*	Animal Welfare	0.31	0.02*	0.00*
Extreme Sports	0.03*	0.02*	0.93	Extreme Sports	0.43	0.11	0.38
Screen Time	0.00*	0.09	0.00*	Screen Time	0.00*	0.00*	0.06*
Dating	0.17*	0.00*	0.04*	Dating	0.01*	0.01*	0.41
Citizens	0.51	0.08	0.16	Citizens	0.16	0.11	0.87

Table 7: Evaluating the significance of the difference of homogenization between the three setups across each topic (Corresponding to the box plots in Figure 9). We conduct an independent samples t-test to compare the document homogenization values between the setups pairwise. Pairs with a significant difference at the 5% level are highlighted with an asterisk. The overall trend is that InstructGPT has significantly higher homogenization across a majority of topics.

Significance Testing - Raw essay level							
Rouge-L				BertScore			
Topic	Solo vs InstructGPT	InstructGPT vs GPT3	GPT3 vs Solo	Topic	Solo vs InstructGPT	InstructGPT vs GPT3	GPT3 vs Solo
Stereotype	0.02*	0.01*	0.00*	Stereotype	0.02*	0.01*	0.03*
School	0.17	0.12	0.71	School	0.14	0.22	0.62
Mindfulness	0.00*	0.01*	0.42	Mindfulness	0.02*	0.22	0.81
Athletes	0.00*	0.00*	0.52	Athletes	0.00*	0.00*	0.90
Audiobook	0.01*	0.15	0.69	Audiobook	0.00*	0.00*	0.76
Animal Welfare	0.49	0.11	0.15	Animal Welfare	0.13	0.00*	0.00*
Extreme Sports	0.35	0.00*	0.00*	Extreme Sports	0.04*	0.00*	0.09
Screen Time	0.00*	0.00*	0.00*	Screen Time	0.00*	0.00*	0.70
Dating	0.00*	0.61	0.00*	Dating	0.88	0.01*	0.02*
Citizens	0.00*	0.14	0.11	Citizens	0.00*	0.04*	0.11

Table 8: Evaluating the significance of the difference of homogenization between the three setups across each topic (Corresponding to the box plots in Figure 8). We conduct an independent samples t-test to compare the document homogenization values between the setups pairwise. Pairs with a significant difference at the 5% level are highlighted with an asterisk. The overall trend is that InstructGPT has significantly higher homogenization across a majority of topics.

	Solo	GPT3	InstructGPT	
Perplexity of Essays (via GPT2)	25.067	22.10	20.26	
Sentence Length (in words)	23.51 (0.25)	23.66 (0.24)	22.51 (0.24)	
Height of Syntax Tree	5.93 (0.06)	5.98 (0.06)	5.71 (0.05)	
Essay Length (in words)	376.44 (8.03)	380.87 (9.30)	368.39 (8.43)	
Human Ratings	Relevance to Prompt	4.30	4.15	4.10
	Grammaticality	4.00	4.10	4.00
	Depth of Discussion	3.95	3.85	3.80
Unique POS Ngrams	1	73	56	58
	2	487	437	426
	3	1297	1261	1235
	4	2044	1975	1988
	5	2423	2338	2414

Table 9: Reporting basic statistics about the collected essays. Writing with the help of a model reduces the perplexity (measured via GPT2) of the essays in InstructGPT and GPT3. We also observe that users write essays of similar length across all setups with slightly shorter sentences in InstructGPT. We also examine the height of the dependency parse trees of the sentences in each setup to find that writing with InstructGPT also results in sentences with less complexity on average. Along with these average numbers, we report the standard errors in brackets. We also find that writing with model help results in fewer unique POS- n -grams. Finally, we rate the quality of one essay per topic per setup (selected randomly) along three different dimensions (grammaticality, relevance to the prompt, and depth of discussion) to find that there is no significance in essay quality across setups.

on diversity which is more focused on content. We finally plot the distributions of the 50 most frequent POS-Ngrams (Figure 6) and find that on 4 and 5-grams, writing with model help leads to higher repetition of frequent POS-Ngrams.

- To ensure that the quality of the essays is not a confounder in the result, we present human ratings for a random sample of 10 essays (one per topic) from each setup that we conducted as part of our quality control checks. Here we judged the essays on 3 different axes—relevance to the topic, grammaticality, and the depth of discussion on their opinions. This was per the advice from faculty at the Expository Writing program at our university. This analysis showed no clear differences between the setups, explained in part due to the observation that the machine-written fraction in the collaborative essays was on average around 32%-35% (Table 9).

Writing with InstructGPT results in more similar content at the key point and raw essay level

In addition to the results presented in Section 4.2 we also report the corpus homogenization scores at the key point and essay level in Table 10. Across both Rouge-L and BertScore, the corpus of essays written with InstructGPT exhibits higher corpus homogenization than both other groups by a statistically significant margin (p -value < 0.05) on both levels. We also plot the homogenization scores of all the essays calculated at the raw essay and key point levels in Figure 8 and Figure 9. We choose to report the homogenization at the key point level mainly because string similarity metrics such as Rouge-L and BertScore are less reliable on longer text documents (Sun et al., 2019; Gehrmann et al., 2023).

Essays written with InstructGPT repeat higher- order n -grams more frequently Building on from the results in Section 5.2, we plot the n -gram distributions of the raw essays varying n from 1 to 5 in Figure 10. Here we truncate the distribution to the most common 50 n -grams from each setup. While the distributions for 1, 2 and 3-grams are almost identical across the setups, on 4 and 5-grams we see a concentration of probability mass at the head of the distribution for the essays written with InstructGPT. The reduction in lexical diversity (Table 3a) is manifested in this increased repetition of higher-order n -grams.

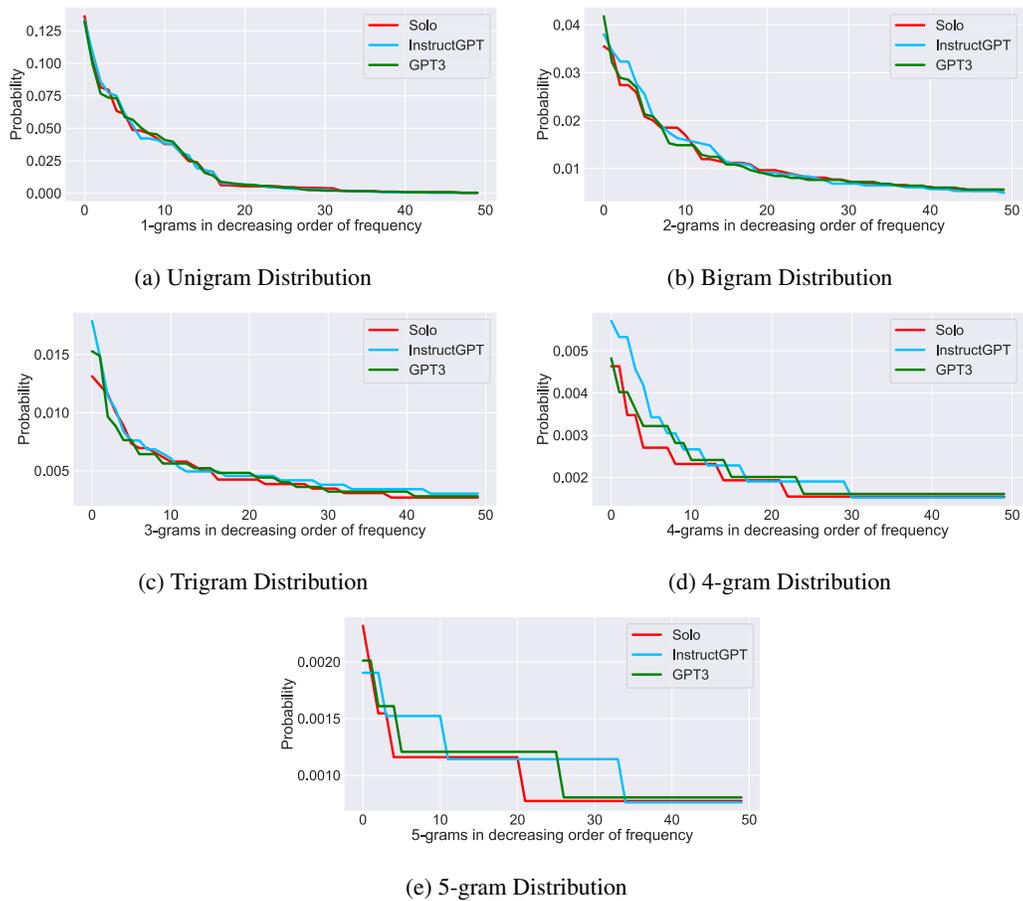


Figure 6: Distribution of the top-50 POS- n -grams in the essays from the various setups varying n from 1 to 5. While the lower order n -gram usage patterns are similar across groups, essays written with InstructGPT exhibit higher repetition of common 4- and 5-POS-ngrams.

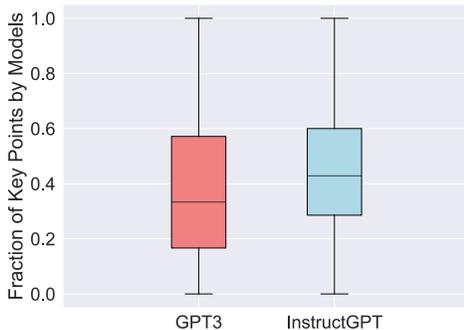


Figure 7: Boxplot of the fraction of key points contributed by the model in each essay. Both models contribute a considerable amount of key points. However, the high variance suggests varied reliance on the model by different users.

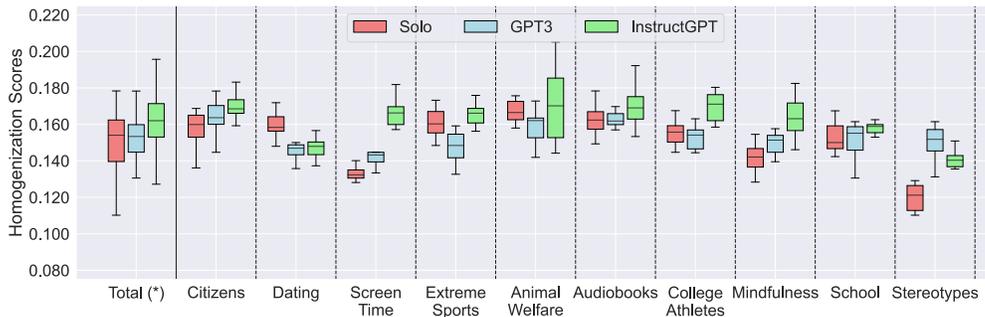
		Solo	GPT3	Instruct-GPT
Key point Level	Rouge-L	0.1536	0.1578	0.1660
	BertScore	0.6305	0.6284	0.6437
Essay Level	Rouge-L	0.1498	0.1523	0.1621
	BertScore	0.6044	0.6130	0.6318

Table 10: Corpus homogenization scores comparing essays from the three setups at the key point and essay level using both Rouge-L and BertScore to calculate homogenization. Essays written with InstructGPT exhibit higher homogenization than Solo and GPT3 across both levels and similarity metrics by a statistically significant margin (p -value < 0.05).

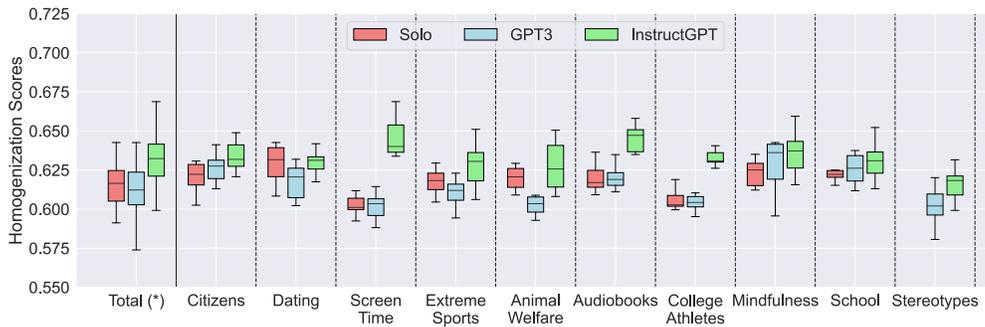
Essays written with InstructGPT are more compressible. We also show that the reduction of diversity measured by linguistic units such as n -grams and key points also correlates with less diversity in an information-theoretic sense. Specifically, we measure the compression ratio of essays written in the three settings using various lossless compression algorithms: LZMA, ZLIB, and GZIP, each being a variation of the Lempel-Ziv algorithm (Ziv and Lempel, 1977; 1978). These algorithms compress data by identifying repeated patterns and using dictionary mappings for encoding. To compute the compression ratio, we concatenate all essays from a setting into a single text file. We then run the compression algorithm on it and calculate the ratio of the compressed file size to the original file size. From Table 13, we see that the InstructGPT essays are consistently more compressible across all three compression algorithms, which is statistically significant with $p < 0.05$ according to permutation tests (Appendix B.1).

InstructGPT presents less diverse suggestions to users than GPT3 To test the hypothesis that adapting a model with human feedback reduces the diversity of the presented suggestions, we calculate the average similarity of all pairs of the five suggestions presented by InstructGPT and GPT3 upon each user query. In addition to the similarity plotted with Rouge-L in Figure 12 we also plot the average similarity scores of suggestions computed using BertScore in Figure 12(b).

Higher correlation between AI written fraction of the document and homogenization, on InstructGPT than GPT3 To study the relationship between increased model intervention and homogenization, we plot the fraction of the essay written by the model as a function of the document homogenization with BertScore and Rouge-L in Figure 13. We observe a weak correlation when users write with InstructGPT, particularly in the case of BertScore.

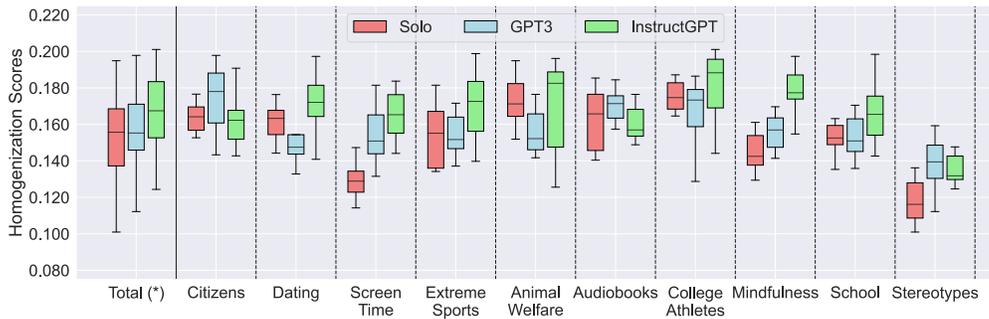


(a) Homogenization calculated with Rouge-L

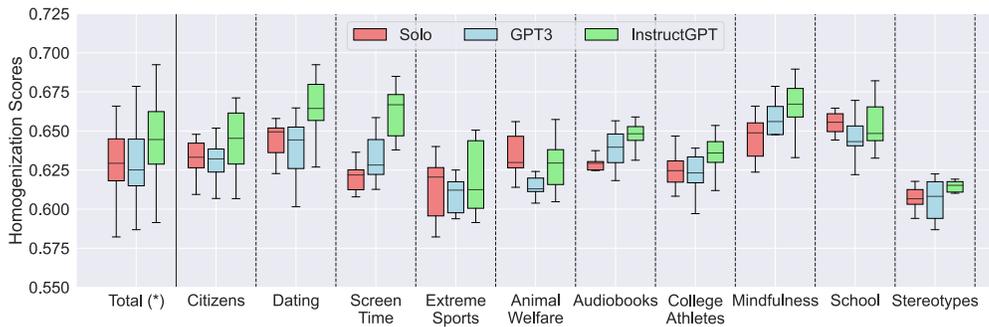


(b) Homogenization calculated with BertScore

Figure 8: Boxplots of homogenization scores for all three groups (Solo, InstructGPT, GPT3) when comparing the raw essays and calculated using (a) Rouge-L and (b) BertScore as measures of similarity. The left-most column (Total) shows essay homogenization scores for all topics and the other columns show essay homogenization scores by topic. Essays written with InstructGPT exhibit higher corpus homogenization by a statistically significant margin (Section 4.2).



(a) Homogenization calculated with Rouge-L



(b) Homogenization calculated with BertScore

Figure 9: Boxplots of homogenization scores for all three groups (Solo, InstructGPT, GPT3) when comparing essays at the key point level and calculated using (a) Rouge-L and (b) BertScore as measures of similarity. The left-most column (Total) shows essay homogenization scores for all topics and the other columns show essay homogenization scores by topic. Essays written with InstructGPT exhibit higher corpus homogenization by a statistically significant margin (Section 4.2).

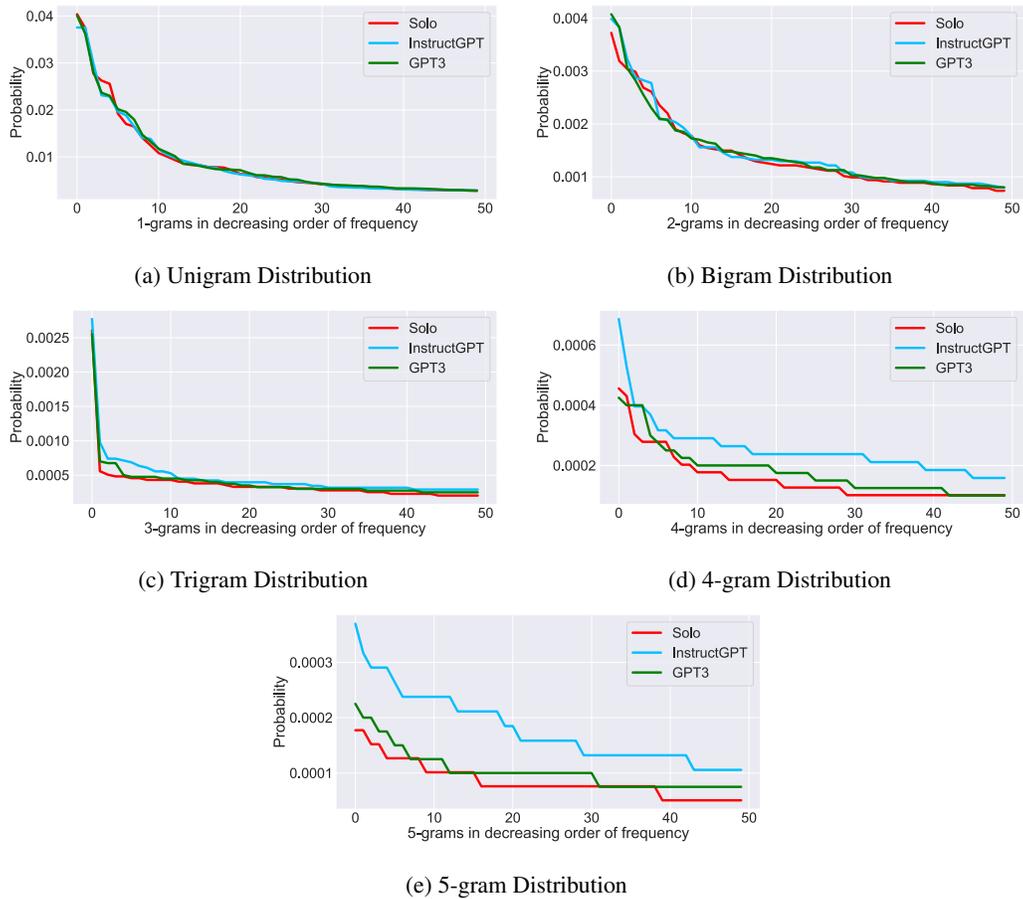


Figure 10: Distribution of the top-50 n -grams in the essays from the various setups varying n from 1 to 5. While the lower order n -gram usage patterns are similar across groups, essays written with InstructGPT exhibit higher repetition of common 4-ngrams and 5-ngrams.

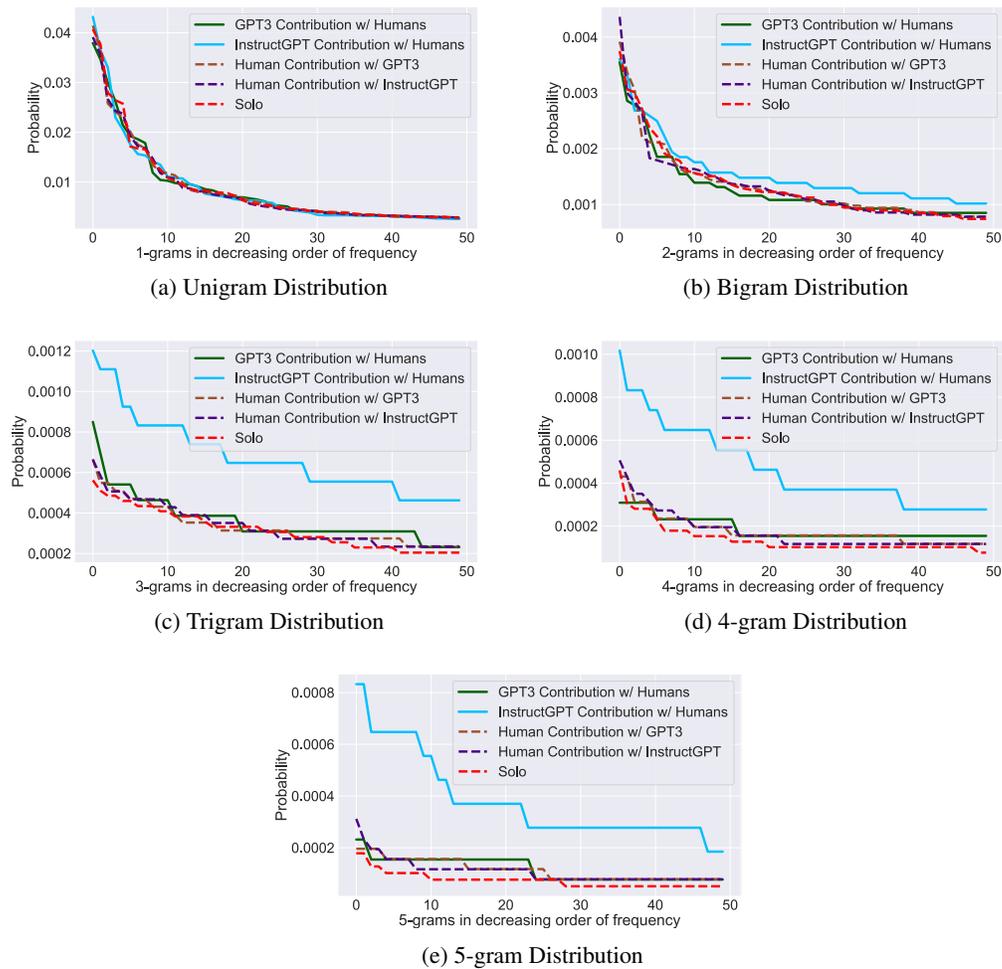


Figure 11: n -gram distribution of tokens introduced by the user and the model during the two co-writing setup with InstructGPT and GPT3 respectively. The distribution of human-written n -grams from Solo essays is also provided as a reference. The distribution of user-written text in all settings is similar to each other regardless of model assistance, whereas InstructGPT contributed text notably has larger probability mass on common 3, 4, and 5-grams. This indicates that the reduced lexical diversity in the InstructGPT group is mainly due to model introduced text.

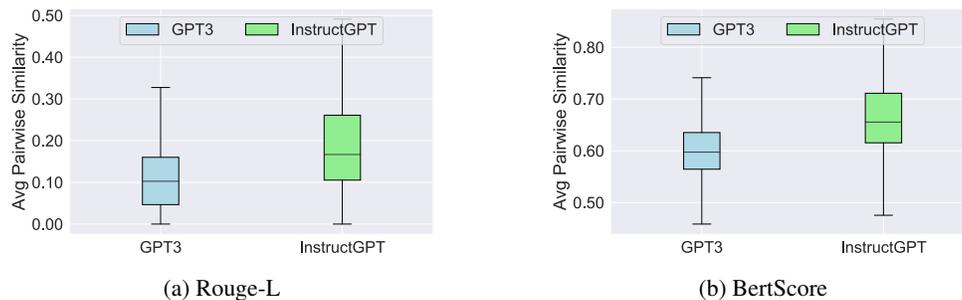


Figure 12: Boxplot of average pairwise similarity of suggestions provided by InstructGPT and GPT3 when calculated using (a) RougeL and (b) BertScore as a measure of similarity. InstructGPT presents more similar suggestions on average to the user by a statistically significant margin.

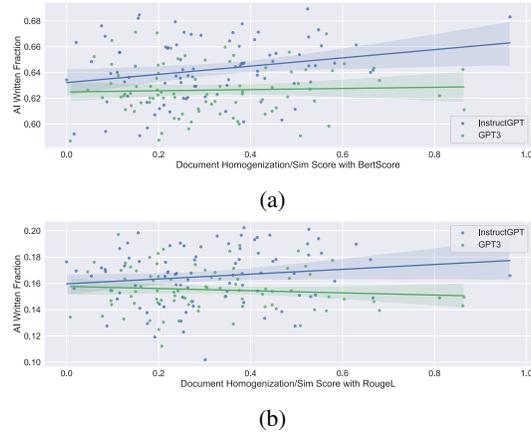


Figure 13

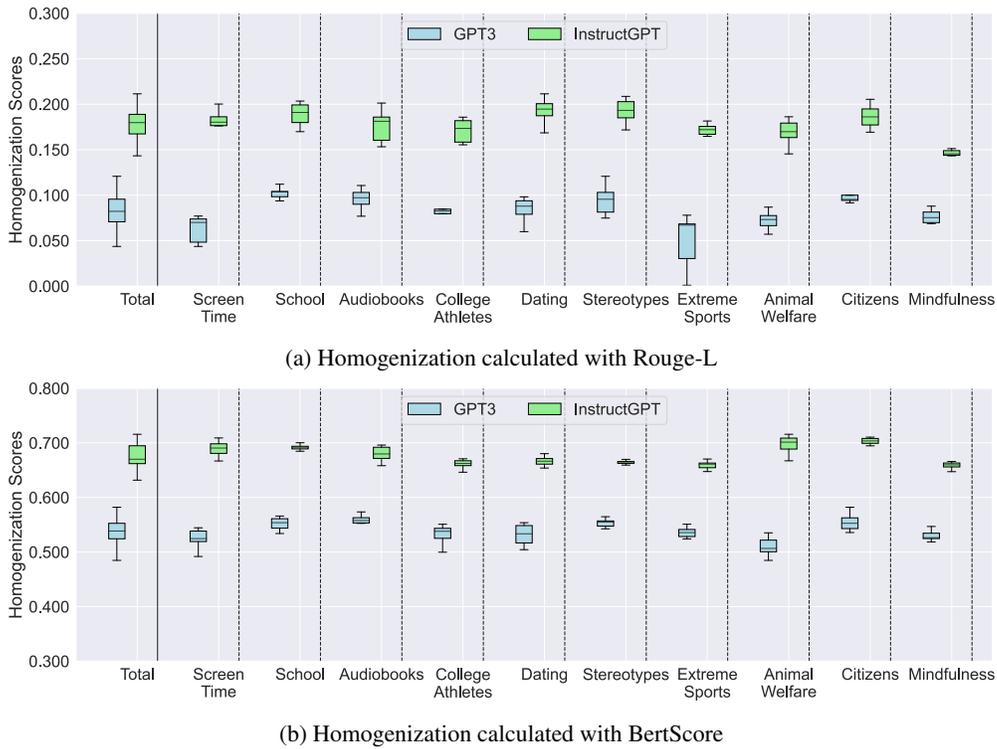


Figure 14: Boxplots of homogenization scores comparing essays generated solely via InstructGPT and GPT3 models at the raw essay level and calculated using (a) Rouge-L and (b) BertScore as measures of similarity. The left-most column (Total) shows essay homogenization scores for all topics and the other columns show essay homogenization scores by topic. Essays written with InstructGPT exhibit higher corpus homogenization by a statistically significant margin.

<i>n</i> -gram size	Solo	GPT3	InstructGPT
1	0.119	0.116	0.115
2	0.602	0.585	0.579
3	0.898	0.886	0.869
4	0.973	0.967	0.953
5	0.991	0.988	0.977

Table 11: *N*-gram diversity in each setting. Bold values indicate the lowest diversity score as measured by the fraction of unique *n*-grams. Essays written with InstructGPT are the least diverse across *n*-gram sizes.

Thresholds	Solo	GPT3	InstructGPT	Thresholds	Solo	GPT3	InstructGPT
0.5	0.982	0.971	0.950	0.1	0.998	0.997	0.992
0.6	0.941	0.927	0.877	0.2	0.981	0.976	0.941
0.7	0.792	0.779	0.738	0.3	0.805	0.787	0.730
0.8	0.543	0.514	0.494	0.4	0.321	0.338	0.292

(a) RougeL

(b) BertScore

Table 12: Diversity measured by agglomerative clustering of the key points of essays with (a) RougeL and (b) BertScore as the metric. We report the scores using different distance thresholds for clustering. Bold values are different to other columns by a statistically significant margin ($p < 0.05$) using a permutation test. InstructGPT consistently exhibits lower content diversity across both similarity metrics and all selected thresholds.

D LIMITATIONS

Interaction interface. Our interface provides suggestions to users in the form of continuations of the current text in the draft. Further investigation is needed to evaluate if the reduction in content diversity with feedback-tuned models can be mitigated with prompt engineering or richer forms of interaction (e.g., through a dialogue).

User selection. The outcome of the experiments can be affected by the specific group of participants. We try to ensure a diverse user group and detail our recruitment procedures in Appendix [A.1](#). However, it is unclear whether these results will generalize to other groups such as students learning to write or second language speakers, who have different goals and incentives for using writing assistants.

LLM access. Our experiments are conducted using two limited-access models from OpenAI. While we believe they are representative of properties of current LLMs, it is possible that the other models may exhibit different behavior, especially given that the RLHF pipeline is highly customized.

Algorithm	Solo	GPT3	InstructGPT
LZMA	0.305	0.302	0.293
ZLIB	0.353	0.351	0.342
GZIP	0.352	0.341	0.350

Table 13: Compression ratios of the essays using various lossless compression algorithms. Bold values are different from each of the other columns with statistical significance ($p < 0.05$) using permutation tests. InstructGPT exhibits lower diversity, i.e. lower compression ratios, than GPT3 and InstructGPT across all three compression algorithms.