

A Little Human Data Goes A Long Way

Dhananjay Ashok[†], Jonathan May[†]

[†]Information Sciences Institute, University of Southern California
{ashokd, jonmay}@isi.edu

Abstract

Faced with an expensive human annotation process, creators of NLP systems increasingly turn to synthetic data generation. While this method shows promise, the extent to which synthetic data can replace human annotation is poorly understood. We investigate the use of synthetic data in Fact Verification (FV) and Evidence-based Question Answering (QA) by incrementally replacing human-generated data with synthetic points on eight diverse datasets. Strikingly, replacing up to 90% of the training data only marginally decreases performance, but replacing the final 10% leads to severe declines. We find that models trained on purely synthetic data can be improved by including as few as **125** human generated data points. We show that matching the performance gain of a little human data requires an order of magnitude more synthetic data, and then estimate price ratios at which human annotation would be a more cost-effective solution. Our results suggest that even when human annotation at scale is infeasible, there is great value to having a small proportion of the dataset being human-generated.

1 Introduction

From BERT (Devlin et al., 2019) to GPT-4 (Achiam et al., 2023), the explosive growth of language models (LMs) has been underpinned by exponential increases in the size of available training data. However, the more complex and specialized the task, the more expensive and challenging it is to collect human-generated data at scale (Wang et al., 2021). Combined with growing concerns that LMs may soon exhaust the stock of publicly available training data (Villalobos et al., 2024), many turn to synthetic data generation, hoping to eliminate their reliance on human annotation.

Synthetic data generation has long been used to increase the amount of training data available (Simard et al., 2002; Krizhevsky et al.,

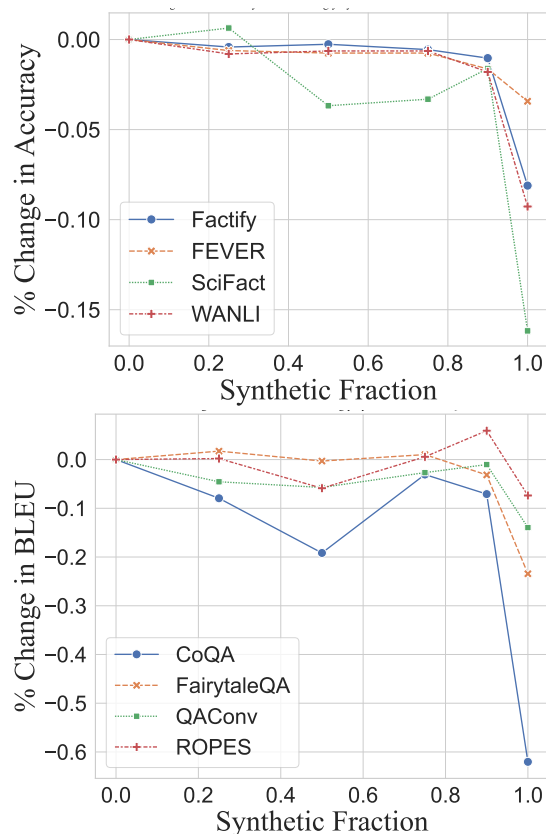


Figure 1: Change in model performance as the proportion of synthetic points in the training data is increased. Across datasets, the performance decrease when moving the synthetic proportion from 0 to 0.90 is often less than that of moving from 0.9 to purely synthetic data.

2012). Early NLP approaches use rule based methods (De Gispert et al., 2005; Chen et al., 2012), paraphrasing (Wang and Yang, 2015; Kobayashi, 2018), noising (Xie et al., 2017; Wang et al., 2018), and backtranslation (Sennrich et al., 2016; Yu et al., 2018), but are limited in their capability.

Modern LMs demonstrate the capability to solve myriad NLP tasks with minimal task specific data (Brown et al., 2020; Wei et al., 2022), making them more powerful synthetic data generators. Leveraging this, synthetic data approaches have

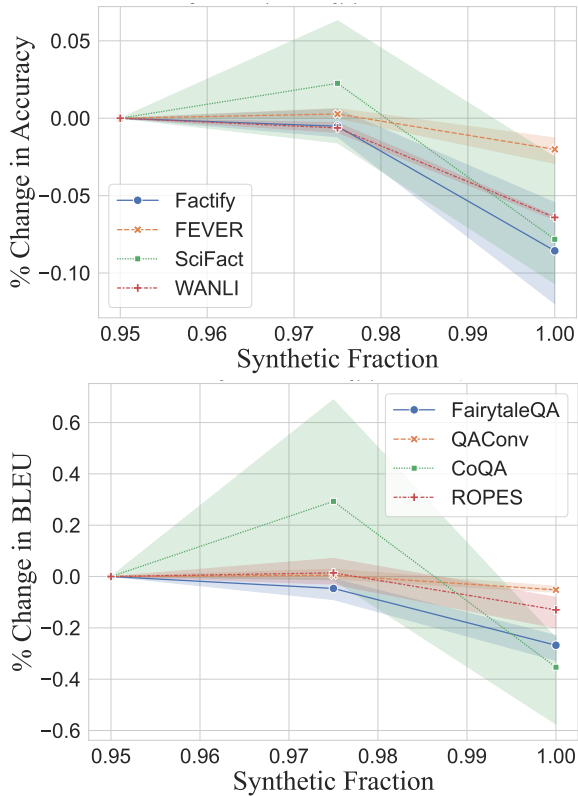


Figure 2: Model performance as the synthetic proportion of the training data varies from 0.95 to 1. Having **just 2.5%** of the training dataset being human-generated boosts performance.

seen increased use in tasks (Tan et al., 2024) such as QA (Wu et al., 2022), natural language inference (NLI) (Meng et al., 2022), text classification (Ye et al., 2022), instruction tuning (Li et al., 2024), evaluation (Dubois et al., 2024), and more (Tang et al., 2023).

The adoption has been particularly enthusiastic for tasks that require the model to ‘understand’ knowledge contained in an ‘evidence text’ e.g., FV (Tang et al., 2024), factual error correction (Ashok et al., 2023), NLI (Hosseini et al., 2024) and evidence-based QA (Schimanski et al., 2024). Such tasks are of vital importance in fake news detection (Sharma et al., 2023), retrieval-augmented generation (Gao et al., 2023) and dialogue systems (Weston et al., 2015). Recent datasets (Wu et al., 2022) and methods (Ye et al., 2022) exploit plentiful evidence texts (scientific journals, news articles, books, etc.), using synthetic generation to avoid being bottlenecked by the expensive annotation procedure (Liu et al., 2022).

Varying results across ML tasks suggest that whether completely replacing humans with synthetic data shows promise (Fan et al., 2024; Ham-

moud et al., 2024) or leads to failures (Bisbee et al., 2024; Guo et al., 2024) is task dependent (Li et al., 2023). In this work, we focus on **FV** and **Evidence-based QA**, performing the first investigation into the trade-offs presented by the use of synthetic data generation in these fundamental tasks.

We study eight diverse FV and QA datasets, using their ‘evidence texts’ to generate synthetic datasets. By holding the number of data points constant but increasing the percentage of the training data that is synthetic, we can compare the utility of synthetic data to the original human generated data points. Across multiple models, prompt models, and prompting strategies, we find (Figure 1) that while increasing the proportion of synthetic data typically causes only minor degradations in model performance, a significant decline occurs at the extremes; i.e., when the percentage of synthetic data exceeds 90%. Focusing on the extremes, we show that purely synthetically trained FV and QA systems can be meaningfully improved by including as few as **125** human-generated datapoints.

Our observations have actionable implications for researchers hoping to use synthetic data for FV and QA. The results (Figure 2, Figure 4) suggest that even when human annotation at scale is infeasible, there is great value to having a small proportion of the dataset being generated by humans.

To help guide this choice, we quantify the performance-cost tradeoff between human and synthetic data. We find (Figure 4) that matching the performance gain of just a little additional human data (only 200 data points) requires an order of magnitude more synthetic data points, empirically showing the per-data point price ratio at which human annotation is the more cost-effective solution. Finally, we conduct an analysis on the differing properties of synthetic vs. human data. Among other findings, we see that synthetic generations are longer and more extractive from the evidence texts than their human-produced counterparts.

2 Synthetic Data Generation from Evidence Texts

We study a synthetic generation pipeline representative of the methods used in the FV (Ni et al., 2024; He et al., 2023) and QA (Schimanski et al., 2024; Wan et al., 2024) literature. Using Few-Shot In-Context Learning (Brown et al., 2020), we generate synthetic (claim, label) pairs from an input evidence text. The prompt model is given exam-




Figure 3: Change in accuracy when the test set (shown in key) is not seen during training, and the training set is a mixture of other FV datasets. Increasing the synthetic proportion of the dataset leads to performance declines even in the OOD setting, showing that human data offers genuine performance increases.

ples of (evidence text, claim, label) from the human training data, and is then queried with the evidence text we seek to generate data for. QA synthetic data is generated analogously, see details in Appendix B.

In total, we use four FV/NLI datasets: FEVER (Thorne et al., 2018), SciFact (Wadden et al., 2020), WANLI (Liu et al., 2022) and FAC-TIFY1.0 (Mishra et al., 2022), as well as four QA datasets: ROPES (Lin et al., 2019), CoQA (Reddy et al., 2019), QAConv (Wu et al., 2022) and Fairy-TaleQA (Xu et al., 2022). Together, the datasets span a variety of domains (science, news, social media, reasoning, conversation, fiction). We confirm that the generations are of high quality by verifying that the diversity of the synthetic data is comparable to the human generated samples. For more details, including a discussion on data leakage, see Appendix C.

FV performance is measured by test accuracy, while QA is measured using BLEU (Papineni et al., 2002); we show robustness to choice of metric in Appendix A. Evaluation is always conducted on the (human-generated) test split of each dataset.

We use GPT-3.5-Turbo (Brown et al., 2020) for prompting and LoRA (Hu et al., 2022) on Llama3-8B (Dubey et al., 2024) for fine-tuning. Implementation details are provided in Appendix E, and our code is publicly available ¹

¹<https://github.com/DhananjayAshok/LittleHumanData/>

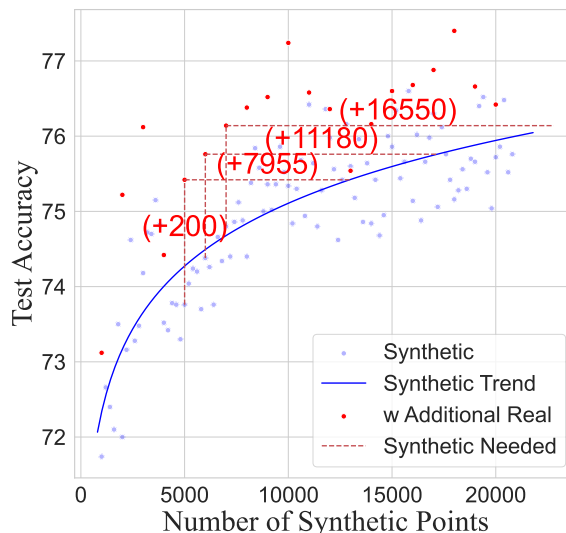


Figure 4: On the WANLI dataset, adding 200 real data points is as effective as adding an order of magnitude more synthetic data points.

3 Can Synthetic Data Replace Humans?

We investigate the potential of synthetic data to replace human annotation by holding the number of training data points fixed, incrementally increasing the proportion of the data that is synthetic, and fine-tuning a model on each training set.

Results: Across all datasets, using purely synthetic data leads to worse performance than the same amount of human data (Figure 1). We consider the possibility that this result could be caused by a spurious correlation between the human training and testing splits (e.g., annotation artifacts that are correlated with the label but not fundamental to the task). We conduct an out-of-distribution experiment, using different datasets for training and testing (e.g., training on FEVER + SciFACT and testing on WANLI). Increasing the synthetic proportion leads to performance declines even in the OOD setting (Figure 3), showing that human data offers genuine performance increases, and the results cannot be explained by a spurious correlation between the human test and human training samples (further discussion in Appendix A).

The performance decline is not uniform as we increase the synthetic proportion. On almost all datasets, there is only a minor degradation up until 90% replacement, after which the performance drops considerably. We zoom in on the 90%-100% interval, fixing the amount of training data at $n = 5000$ (500 for SciFact) and training on datasets with 95%, 97.5% and 100% synthetic

data (Figure 2). Surprisingly, the results show that there is a significant difference between the performance of models on 97.5% and 100% synthetic data; the addition of just **125** (2.5% of 5000) human generated datapoints reliably improves the performance of synthetically trained FV and QA models. These trends hold robustly over different languages (Arabic, Georgian, Indonesian), choice of fine-tuning model (Mistral, MPT), prompt model (GPT4 and Claude-3.5-Sonnet), prompting strategy (Chain-of-Thought), model size and dataset size (Appendix A).

4 When Should We Use Human Data?

Having observed the disproportionate value added by human data, we ask what the relative cost between human and synthetic data generation must be for us to prefer one over the other. We fine-tune models on purely synthetic datasets of varying sizes, and establish the synthetic baseline by fitting a curve of the form $y = a_0 + a_1 \log(x)$ where x is the size of the synthetic dataset and y is the performance. We then take the synthetic training sets with $\{1000, 2000 \dots\}$ points and observe the performance (y^*) when we add 200 human data points. $\exp(\frac{y^* - a_0}{a_1})$ is then the size of the purely synthetic dataset that achieves equivalent performance.

Results: Across all datasets, adding **200** human data points is usually comparable to adding at least an order of magnitude (often multiple orders of magnitude) more synthetic data points. On WANLI (Figure 4), more than 17,000 additional synthetic points are needed to achieve the performance gains of 200 human points. If the price of a synthetic point for WANLI exceeds 73 times the price of a human generated point, then an incremental amount of human annotation would be a more cost-effective way to achieve the same increase in accuracy. In the extreme case, the equation learned on FairyTaleQA suggests that it takes $2e5$ additional synthetic points to match the performance gain of 200 additional human data points. Rather than interpret these numbers literally, we take them to suggest that human data could have unique value in some settings, enabling performance levels that are impossible with purely synthetic datasets. See Appendix A for more results and details.

5 Discussion

The synthetic generations are as diverse as human data (Table 2), with comparable duplicate rates on

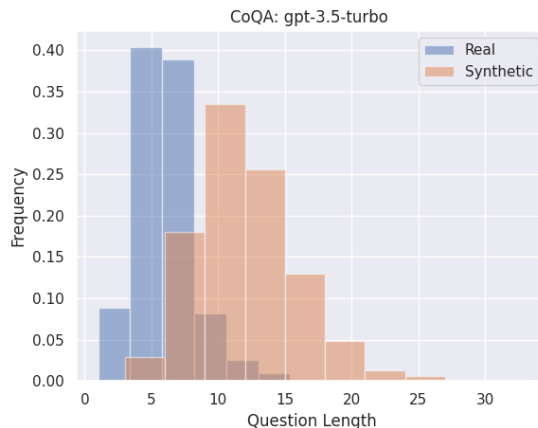


Figure 5: Synthetic questions are longer than human generated ones, a trend also seen in answers.

QA datasets, and markedly fewer duplicates on most FV datasets. This is evidence that our synthetic generations are of good quality, however, even on the datasets where the synthetic data is significantly more diverse, the synthetic data does not perform as well. This suggests that diversity is an insufficient measure of quality when evaluating how good the generated data is. Our analysis shows that synthetic data generation produces claims of comparable length to the human datasets, however synthetic questions and answers tend to be longer than human-generated counterparts for all QA datasets (Figure 5). We find that synthetic generations have a higher n-gram overlap with the evidence sentences. This suggests that synthetic data generation produces data points that are more directly taken from the evidence texts, while humans are more likely to employ rephrasing or different vocabulary than the evidence texts. Surprisingly, we find that synthetic data generation chooses more varied parts of the input text as sources for the question and answer content, with human annotation overwhelmingly more likely to create questions whose answers lie at the start of the evidence texts. We include a detailed discussion in Appendix D.

6 Related Work

The replacement of human annotation with synthetic data is extensively studied in the pretraining stage of LMs, where results consistently show (Shumailov et al., 2023; Seddik et al., 2024; Guo et al., 2024; Briesch et al., 2023) catastrophic forgetting, mode collapse, and performance deterioration.

In our setting, relying only on synthetic data still achieves reasonable performance across all

Dataset	Mean	Median	25th Percentile	75th Percentile
WANLI	17,671	16,905	9,711	22,931
ROPES	17,333	6,006	3,623	21,944
FairyTaleQA	281,951	36,901	15,129	813,135
FEVER	1,155	237	-1,400	7,073

Table 1: Additional synthetic data points needed to match the performance gain of 200 human data points. High values for FairyTaleQA suggest that human-generated data may unlock performance that purely synthetic data cannot achieve. Negative values for FEVER are due to a saturation of the performance gains, however, human data points reach the saturation point much faster (Appendix A)

Dataset	Synthetic	Human
FEVER	5.50	20.27
WANLI	0.23	1.22
SCIFACT	0.00	9.93
FACTIFY	0.27	14.93
NarrativeQA	3.85	1.42
CoQA	0.54	5.49
FairyTaleQA	2.26	0.18
ROPES	2.40	1.35

Table 2: Percentage of duplicated claims/questions for synthetic v.s. human data. Rates are comparable across datasets, but for fact verification datasets, synthetic datasets have fewer duplicates.

tasks. This suggests that the usage of exclusively synthetic data poses fewer risks when generations are grounded in diverse, natural ‘evidence texts.’

Interestingly, conclusions which confirm our findings are found more in the image and multi-modal domains, where recent work (Singh et al., 2024; He et al., 2023; Fan et al., 2024) finds that synthetic data holds promise, but must be used in conjunction with human data to mitigate its harms.

There is limited work on understanding whether synthetic data can replace human annotation in a task-specific setting for the language domain. Li et al. (2023) categorize text classification tasks by subjectivity, showing that synthetic data is less useful when tasks are more subjective. This draws them to focus on different tasks (sentiment classification, relation extraction and spam detection), and they do not study using a mixture of real and synthetic data. Bisbee et al. (2024) demonstrate that replacing political survey respondents with LMs produces unreliable results, while Ahmed et al. (2024) find that there are specific software engineering subtasks where synthetic data approaches human performance. Chen et al. (2024) show that

instruction-following capabilities are diminished when using synthetic data and present a machine unlearning approach to mitigate this. The diversity of results when evaluating the impact of using purely synthetic data confirms that the feasibility of replacing human annotation with synthetic data is highly task dependant. This work deepens our understanding of the problem by being the first to study whether synthetic data can replace human annotation on the fundamental tasks of fact verification and evidence-based question answering.

7 Conclusion

Showing impressive performance when human data is scarce, synthetic data generation seems poised to remain a key method in FV and QA. Our work sheds light on how the best way to use this method is in conjunction with human data. We show that a little human data goes a long way, with just **125** points being enough to see reliable gains on all datasets studied. With practical considerations in mind, we show that the alternative to small amounts of additional human data can be an order of magnitude more of synthetic data, suggesting that at times human annotation can be cost-effective relative to synthetic generation. We hope these results better inform design decisions on datasets and methods for fact verification and question answering.

8 Limitations

While we include results on multilingual Fact Verification datasets, the primary focus of our work is limited to the English language. Additionally, our results on multilingual datasets suggest that while similar claims can be made regarding the impact of replacing human annotation with synthetic data across different languages, the amount of human data needed to observe a meaningful performance increase may vary across languages. We also have a limited ability to control for dataset leakage, with

only one dataset from each of the tasks that is surely not leaked to GPT-3.5 (and, even these two datasets may have been seen by GPT-4). This can potentially bias the results in favor of synthetic data. Due to the scarcity of suitable available datasets (i.e., ones that have not been exposed to the prompt models) we are prevented from studying the problem more rigorously. Another limitation is that while we are able to identify clear differences between synthetic vs. real data distributions, our analysis of the errors made by models trained on 0% vs. 100% synthetic data failed to yield any generalizable insights that could inform modelling approaches. A more fine-grained study of the effect of using synthetic data on the behaviour of the downstream model is hence left as a subject of future research.

9 Ethical Considerations

The usage of synthetic data has several important ethical considerations. In the era of LMs trained on internet-wide corpora having poor documentation as to their exact data sources, it becomes challenging to ensure the privacy of individuals whose data may be obtainable via a public crawl (Yao et al., 2024). Additionally, models trained on massive internet-based data sources may contain implicit biases, as well as illegal and/or highly offensive material that is hard to audit and clean (Bender et al., 2021). This data affects the synthetic data obtained from prompt models, and could unknowingly impose cultural or ethical viewpoints that are unintended or not well aligned with the use case in mind. Specifically, prior work has shown that one of the prompt models studied in this work, GPT-3.5, often disagrees with humans on key ethical questions (Felkner et al., 2024). The endeavour to completely replace human annotation with synthetic data generation also has key implications on the extent to which the field of NLP employs human annotators. It is possible that an increasing reliance on purely synthetic data reduces the demand for human annotation, which would place a downward pressure on the working standards and compensation awarded to the remaining human annotators (Weidinger et al., 2022). We argue in this work that we should not try to eliminate human annotation from our dataset and method design, showing that their work contributes uniquely helpful data points.

10 Acknowledgements

This work was funded by the Defense Advanced Research Projects Agency with award HR00112220046. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the view of our sponsors.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Toufique Ahmed, Prem Devanbu, Christoph Treude, and Michael Pradel. 2024. [Can llms replace manual annotation of software engineering artifacts?](#) *ArXiv*, abs/2408.05534.
- Dhananjay Ashok, Atharva Kulkarni, Hai Pham, and Barnabas Poczos. 2023. [The student becomes the master: Outperforming GPT3 on scientific factual error correction](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6762–6778, Singapore. 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*, pages 610–623.
- James Bisbee, Joshua D. Clinton, Cassy Dorff, Brenton Kenkel, and Jennifer M. Larson. 2024. [Synthetic replacements for human survey data? the perils of large language models](#). *Political Analysis*.
- Martin Briesch, Dominik Sobania, and Franz Rothlauf. 2023. [Large language models suffer from their own output: An analysis of the self-consuming training loop](#). *ArXiv*, abs/2311.16822.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

- Jie Chen, Yupeng Zhang, Bingning Wang, Xin Zhao, Ji-Rong Wen, and Weipeng Chen. 2024. [Unveiling the flaws: Exploring imperfections in synthetic data and mitigation strategies for large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14855–14865, Miami, Florida, USA. Association for Computational Linguistics.
- Mei-Hua Chen, Shih-Ting Huang, Chung-Chi Huang, Hsien-Chin Liou, and Jason S Chang. 2012. Prefer: using a graph-based approach to generate paraphrases for language learning. In *Proceedings of the Seventh Workshop on Building Educational Applications Using NLP*, pages 80–85.
- Adria De Gispert, José B Mariño, and Josep Maria Crego. 2005. Improving statistical machine translation by classifying and generalizing inflected verb forms. In *INTERSPEECH*, pages 3193–3196.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Yann Dubois, Chen Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy S Liang, and Tatsunori B Hashimoto. 2024. AlpacaFarm: A simulation framework for methods that learn from human feedback. *Advances in Neural Information Processing Systems*, 36.
- Lijie Fan, Kaifeng Chen, Dilip Krishnan, Dina Katabi, Phillip Isola, and Yonglong Tian. 2024. Scaling laws of synthetic images for model training... for now. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7382–7392.
- Virginia Felkner, Jennifer Thompson, and Jonathan May. 2024. [GPT is not an annotator: The necessity of human annotation in fairness benchmark construction](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14104–14115, Bangkok, Thailand. Association for Computational Linguistics.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Qianyu Guo, Meng Wang, and Haofen Wang. 2023. [Retrieval-augmented generation for large language models: A survey](#). *ArXiv*, abs/2312.10997.
- Yanzhu Guo, Guokan Shang, Michalis Vazirgiannis, and Chloé Clavel. 2024. [The curious decline of linguistic diversity: Training language models on synthetic text](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3589–3604, Mexico City, Mexico. Association for Computational Linguistics.
- Hasan Hammoud, Hani Itani, Fabio Pizzati, Philip H. S. Torr, Adel Bibi, and Bernard Ghanem. 2024. [Synthclip: Are we ready for a fully synthetic clip training?](#) *ArXiv*, abs/2402.01832.
- Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip Torr, Song Bai, and Xiaojuan Qi. 2023. [Is synthetic data from generative models ready for image recognition?](#) In *The Eleventh International Conference on Learning Representations*.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022. [TRUE: Re-evaluating factual consistency evaluation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3905–3920, Seattle, United States. Association for Computational Linguistics.
- Mohammad Javad Hosseini, Andrey Petrov, Alex Fabrikant, and Annie Louis. 2024. [A synthetic data approach for domain generalization of NLI models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2212–2226, Bangkok, Thailand. Association for Computational Linguistics.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Sosuke Kobayashi. 2018. [Contextual augmentation: Data augmentation by words with paradigmatic relations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 452–457, New Orleans, Louisiana. Association for Computational Linguistics.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- Xian Li, Ping Yu, Chunting Zhou, Timo Schick, Omer Levy, Luke Zettlemoyer, Jason E Weston, and Mike Lewis. 2024. [Self-alignment with instruction back-translation](#). In *The Twelfth International Conference on Learning Representations*.
- Zhuoyan Li, Hangxiao Zhu, Zhuoran Lu, and Ming Yin. 2023. [Synthetic data generation with large language](#)

- models for text classification: Potential and limitations. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- 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.
- Kevin Lin, Oyvind Tafjord, Peter Clark, and Matt Gardner. 2019. [Reasoning over paragraph effects in situations](#). In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 58–62, Hong Kong, China. Association for Computational Linguistics.
- Alisa Liu, Swabha Swayamdipta, Noah A. Smith, and Yejin Choi. 2022. [Wanli: Worker and ai collaboration for natural language inference dataset creation](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel Weld. 2020. [S2ORC: The semantic scholar open research corpus](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4969–4983, Online. Association for Computational Linguistics.
- Yu Meng, Jiaxin Huang, Yu Zhang, and Jiawei Han. 2022. Generating training data with language models: Towards zero-shot language understanding. *Advances in Neural Information Processing Systems*, 35:462–477.
- Shreyash Mishra, S Suryavardan, Amrit Bhaskar, Parul Chopra, Aishwarya N. Reganti, Parth Patwa, Amitava Das, Tanmoy Chakraborty, A. Sheth, and Asif Ekbal. 2022. [Factify: A multi-modal fact verification dataset](#). In *DE-FACTIFY@AAAI*.
- Jingwei Ni, Minjing Shi, Dominik Stammach, Mrinmaya Sachan, Elliott Ash, and Markus Leippold. 2024. [Afacta: Assisting the annotation of factual claim detection with reliable llm annotators](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. [CoQA: A conversational question answering challenge](#). *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Tobias Schimanski, Jingwei Ni, Mathias Kraus, Elliott Ash, and Markus Leippold. 2024. [Towards faithful and robust LLM specialists for evidence-based question-answering](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1913–1931, Bangkok, Thailand. Association for Computational Linguistics.
- Mohamed El Amine Seddik, Swei-Wen Chen, Soufiane Hayou, Pierre Youssef, and Mérouane Debbah. 2024. [How bad is training on synthetic data? a statistical analysis of language model collapse](#). *ArXiv*, abs/2404.05090.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Umang Sharma, Sidarth Saran, and Dr Shankar M. Patil. 2023. [Fake news detection using machine learning algorithms](#). *2023 International Conference on New Frontiers in Communication, Automation, Management and Security (ICCAMS)*, 1:1–7.
- Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson. 2023. [The curse of recursion: Training on generated data makes models forget](#). *ArXiv*, abs/2305.17493.
- Patrice Y Simard, Yann A LeCun, John S Denker, and Bernard Victorri. 2002. Transformation invariance in pattern recognition—tangent distance and tangent propagation. In *Neural networks: tricks of the trade*, pages 239–274. Springer.
- Krishnakant Singh, Thanush Navaratnam, Jannik Holmer, Simone Schaub-Meyer, and Stefan Roth. 2024. Is synthetic data all we need? benchmarking the robustness of models trained with synthetic images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2505–2515.
- Zhen Tan, Dawei Li, Song Wang, Alimohammad Beigi, Bohan Jiang, Amrita Bhattacharjee, Mansoor Karami, Jundong Li, Lu Cheng, and Huan Liu. 2024. [Large language models for data annotation and synthesis: A survey](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 930–957, Miami, Florida, USA. Association for Computational Linguistics.
- Liyan Tang, Philippe Laban, and Greg Durrett. 2024. [Minicheck: Efficient fact-checking of llms on grounding documents](#). *ArXiv*, abs/2404.10774.
- Ruixiang Tang, Xiaotian Han, Xiaoqian Jiang, and Xia Hu. 2023. [Does synthetic data generation of llms help clinical text mining?](#) *ArXiv*, abs/2303.04360.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. Fever: a large-scale dataset for fact extraction and verification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language*

- Technologies, Volume 1 (Long Papers)*, pages 809–819.
- Pablo Villalobos, Anson Ho, Jaime Sevilla, Tamay Be-siroglu, Lennart Heim, and Marius Hobbhahn. 2024. [Position: Will we run out of data? limits of LLM scaling based on human-generated data](#). In *Forty-first International Conference on Machine Learning*.
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. [Fact or fiction: Verifying scientific claims](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7534–7550, Online. Association for Computational Linguistics.
- Yuwei Wan, Yixuan Liu, Aswathy Ajith, Clara Grazian, Bram Hoex, Wenjie Zhang, Chunyu Kit, Tong Xie, and Ian Foster. 2024. [Sciqag: A framework for auto-generated science question answering dataset with fine-grained evaluation](#).
- Shuohang Wang, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2021. [Want to reduce labeling cost? GPT-3 can help](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4195–4205, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- William Yang Wang and Diyi Yang. 2015. [That’s so annoying!!!: A lexical and frame-semantic embedding based data augmentation approach to automatic categorization of annoying behaviors using #petpeeve tweets](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2557–2563, Lisbon, Portugal. Association for Computational Linguistics.
- Xinyi Wang, Hieu Pham, Zihang Dai, and Graham Neubig. 2018. [SwitchOut: an efficient data augmentation algorithm for neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 856–861, Brussels, Belgium. Association for Computational Linguistics.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, et al. 2022. Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 214–229.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. 2015. [Towards ai-complete question answering: A set of prerequisite toy tasks](#). *arXiv: Artificial Intelligence*.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- Chien-Sheng Wu, Andrea Madotto, Wenhao Liu, Pascale Fung, and Caiming Xiong. 2022. [QAConv: Question answering on informative conversations](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5389–5411, Dublin, Ireland. Association for Computational Linguistics.
- Ziang Xie, Sida I. Wang, Jiwei Li, Daniel Lévy, Aiming Nie, Dan Jurafsky, and Andrew Y. Ng. 2017. [Data noising as smoothing in neural network language models](#). In *International Conference on Learning Representations*.
- Ying Xu, Dakuo Wang, Mo Yu, Daniel Ritchie, Bingsheng Yao, Tongshuang Wu, Zheng Zhang, Toby Jia-Jun Li, Nora Bradford, Branda Sun, Tran Bao Hoang, Yisi Sang, Yufang Hou, Xiaojuan Ma, Diyi Yang, Nanyun Peng, Zhou Yu, and Mark Warschauer. 2022. [Fantastic questions and where to find them: FairytaleQA – an authentic dataset for narrative comprehension](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 447–460, Dublin, Ireland. Association for Computational Linguistics.
- Yifan Yao, Jinhao Duan, Kaidi Xu, Yuanfang Cai, Zhibo Sun, and Yue Zhang. 2024. A survey on large language model (llm) security and privacy: The good, the bad, and the ugly. *High-Confidence Computing*, page 100211.
- Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. 2022. [ZeroGen: Efficient zero-shot learning via dataset generation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11653–11669, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Adams Wei Yu, David Dohan, Quoc Le, Thang Luong, Rui Zhao, and Kai Chen. 2018. [Fast and accurate reading comprehension by combining self-attention and convolution](#). In *International Conference on Learning Representations*.
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. [AlignScore: Evaluating factual consistency with a unified alignment function](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11328–11348, Toronto, Canada. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

A Supplemental Figures

We present a detailed set of figures and tables to supplement the results presented in the main text.

Main Experiments: For figures in the main text where only one task is shown (Figure 1 and Figure 2), we provide the complete figures with both tasks (Figure 6 and Figure 7). We also provide the individual performance curves for these experiments (Figure 8 and Figure 10).

Robustness to choice of QA metric: To verify the robustness of the results, we show that the QA results are not an artifact of the choice of metric (Figure 3 and Figure 4) by using Exact Match, String Inclusion, ROUGE-1 (Lin, 2004) and BERTScore (Zhang et al., 2020). There is overwhelming agreement between all metrics on the rankings of models.

Addressing spurious correlations: We show that the performance gains afforded by human generated data cannot be explained by a spurious correlation between the human generated train and test splits. This would occur when there are significant annotation artifacts that are not relevant to the task, but are correlated with the correct output. We conduct an out-of-domain experiment (Table 5), using different datasets to source the training data and testing on a single hold out dataset. Using more synthetic data leads to performance declines even in the OOD setting, showing that human data is of higher quality and the results from the main text cannot be explained by a spurious correlation between the human test and human training samples. Interestingly, in the OOD setting the decline is steady, and we do not observe the phenomenon of a small amount of human data having a disproportionate impact on performance. This suggests that the disproportionate impact of human data occurs when the human data is in-domain. We leave a further exploration of the OOD generalization abilities of synthetic vs. human data to future work.

Multilingual Experiments: We replicate our experiment using the Arabic, Georgian, and Indonesian splits of the XFact dataset. We observe (Figure 9) the same trend as those from earlier experiments, confirming that our results are not limited to the English language. While the phenomenon is reproduced, the threshold of replacement at which we observe a precipitous decline is not the same across languages. We hypothesize that the language-specific threshold at which a little human data leads to significant performance increases is

dependent on how low resource the language is. The study of synthetic data in the multilingual setting has unique considerations that we have not addressed in this work; we leave a focus on these problems to future work.

Ablations: We show that the same trends can be seen (Figures 11 to 15) when using a different fine-tuning model (Mistral-7B), models of varying scales (from 1B parameter models to 30B parameter models. different prompting models (GPT-4 and Claude-3.5) and a more sophisticated prompting strategy (Chain-of-Thought Prompting). Across all configurations, we see a consistent decrease in performance when moving from 95% to 100% synthetic data, confirming that models trained on purely synthetic data can be improved by including just 125 real data points. For Chain-of-Thought Prompting, the authors manually annotated 3 examples with rationales per dataset to serve as the prompts. The complete examples and pipeline are provided with the code: github.com/dhananjayashok/littlehumandata

We additionally show that these trends hold across data scales (Figures 16 and 17), replicating the experiment with $n=3000$ and $n=1000$. While the trend is clearly visible in both cases, the results for $n=1000$ have more variance and hence have a minority of cases where the relationship does not hold.

Tradeoff Experiment: The main text shows results for the experiment detailed in Section 4 on the WANLI dataset (Figure 4), here we show results on the remaining three datasets (Figure 18) and provide (Figure 1) the number of additional synthetic points needed to match the performance gains of 200 additional real points (average, median and standard deviation for each dataset). ROPES shows similar results to WANLI, however FairyTaleQA and FEVER present different trends. On FEVER, we are able to reach the saturation point, after which additional data (whether synthetic or real) does not increase performance. Even in this case, we are able to reach this point of diminishing marginal return more rapidly when using a small amount of synthetic data. On a base synthetic training set of size 3000, adding 200 real data points drives the test accuracy to 89.25%, a score that is only matched once we add at least 2000 synthetic data points (an order of magnitude larger). On FairyTaleQA, we get enormous estimates for the number of additional synthetic points needed (a mean of $2.8e5$). We do not interpret these numbers

literally, rather seeing this as a sign that human generated data may occasionally boost performance to an extent that could be fundamentally unachievable by purely synthetic data.

B Synthetic Data Generation

In our implementation (Figure 19), we use few-shot learning with $k = 3$, i.e., three examples per query, with each example drawn randomly (with replacement) from the training set of the specific dataset.

We generate one synthetic point for every real point in the dataset, using the evidence text it is associated with. This gives us a total of n synthetic data points for every n real data points in a dataset.

We observed that if we did not correct for label shift, the prompt model would be heavily biased towards True claims, i.e., it would generate a dataset containing 90% True claims, while original datasets have proportions between 33%–60% True.

For the synthetic datasets used in our experiments, we correct for this label shift by specifying the label of the claim we wish to generate and providing only examples of claims with that specific label in the prompt.

For all datasets, we verify that the diversity of the generated claims/questions/answers are comparable to that of the human generated texts (see Appendix D)

This setting is generous towards synthetic data generation. In practice, we might only have three fixed examples to use in the prompt, potentially reducing the diversity of synthetic data generated. We verify that this does not affect the results of our experiment in the Chain-of-Thought ablation (Figure 15), where we use a fixed set of examples to generate all synthetic data points. We may also not know the correct label proportion to ask for and suffer a significant label shift when using synthetic data generation.

C Datasets Used

All datasets used below are released under open use licenses, authorizing their use in this research. For each dataset, we discuss the potential of dataset leakage (i.e., whether the data has been exposed to GPT-3.5-Turbo during its training) as well as the extent of automation involved in the generation of each dataset. However, across all experiments and ablations, these factors do not seem to have any

discernable effect on the trends discovered in this work.

C.1 Fact Verification Datasets

FEVER (Thorne et al., 2018) is a dataset of claims about specific entities, generated by altering sentences extracted from Wikipedia. The evidence passages are sentences from Wikipedia articles relevant to the entity in question. This dataset has been well established for a long time before the release of the prompting models used in this work, increasing the chance that it has been exposed to the prompt model ahead of time.

SciFact (Wadden et al., 2020) is a fact verification dataset for the scientific domain, which uses the abstracts of scientific articles as evidence texts. The corpus is collected from S2ORC (Lo et al., 2020), a publicly-available corpus of millions of scientific articles. Annotators are shown a source citation in the context of an article, and are asked to write up to three claims based on the content of the citation.

The above datasets are popular NLP challenge sets that were well known even before the release of GPT-3.5-Turbo (Brown et al., 2020), the prompting model used in this work. The following two datasets were released after the official training date cut-off, guaranteeing that the data has not been seen ahead of time.

WANLI (Liu et al., 2022) is an NLI dataset of 108K examples created through a hybrid worker and AI collaboration approach. The creators first study MultiNLI (Williams et al., 2018) and use dataset cartography to automatically identify examples that demonstrate challenging reasoning patterns. and then instruct GPT-3 to compose new examples with similar patterns. Machine generated examples are then automatically filtered, and finally revised and labeled by human crowd workers. While GPT-3.5-Turbo has not been trained on this data, it is worth noting that the data is partially synthetically generated.

FACTIFY (Mishra et al., 2022) is a dataset on multi-modal fact verification. It contains images, textual claims, reference textual documents and reference images. The dataset marks some examples that can be verified using text only; we use this sample in our experiments. This dataset was released after the training cut-off date for GPT-3.5 and takes its evidence texts/claims from human-written news or editorial articles. This ensures that the prompt models studied have not seen the data before training.

Label mapping for NLI and FV: While all of the above datasets contain labels for Supports, Refutes, and Not Enough Information (or Entails, Contradicts, Neutral), we consider the stricter formulation of Fact Verification used by [Honovich et al. \(2022\)](#) and [Zha et al. \(2023\)](#), considering a claim to be factual if the label is Supports (Entails), and non-factual otherwise.

C.2 Question Answering Datasets

ROPES ([Lin et al., 2019](#)) is a QA dataset which tests a system’s ability to apply knowledge from a passage of text to a new situation. The evidence context contains causal or qualitative relation(s) (e.g., “animal pollinators increase efficiency of fertilization in flowers”), and a novel situation that uses this background. The question requires reasoning about effects of the relationships in the background passage in the context of the situation.

CoQA ([Reddy et al., 2019](#)) is a dataset for building Conversational Question Answering systems. CoQA measures the ability of machines to understand a text passage and answer a series of interconnected questions that appear in a conversation. In our experiments, we extract only the first question in the series and use this to obtain our (context, question, answer) data points.

QAConv ([Wu et al., 2022](#)) focuses on informative conversations, including business emails, panel discussions, and work channels. The creators collect QA pairs with both human-written and machine-generated questions. They use a question generator and a dialogue summarizer as auxiliary tools to collect and recommend questions. While the arXiv version of the paper appears before the GPT-3 cut-off data (April 2021 to the cut-off date of Sept 2021), the paper itself appeared only at ACL 2022. It is still possible that the training data was compromised, and owing to the lack of clarity on the training data used for GPT-3 we have no way to confirm or deny this speculation.

FairyTaleQA ([Xu et al., 2022](#)) is a dataset focusing on narrative comprehension of kindergarten to eighth-grade students. The evidence texts are derived from children-friendly stories which serve as evidence texts. The questions are both explicit and implicit, covering seven types of narrative elements or relations. This dataset was released after the GPT-3 training cut-off date, ensuring that it has not been seen by our prompt model before.

D Detailed Discussion on Differences Between Synthetic and Human Data

To compute the extent to which the evidence sentences ‘contain’ the questions, answers, and claims, we measure the BLEU of the generation with **each individual sentence** of the evidence texts, plotting the maximum of these BLEU scores in Figure 21. We find that synthetic generations have a far higher n-gram overlap with the evidence sentences than human generations. This suggests that synthetic data generation produces data points that are more extractive, while humans are more likely to abstract from the evidence. We also use the position of the evidence sentence that achieves the highest BLEU score as a proxy for the source location of the synthetic generation, and find that synthetic data generation chooses more diverse sources for the question and answer content, with human annotation overwhelmingly more likely to create questions whose answers lie in the start of the evidence texts (Figure 22). Finally, the main text shows the size length comparison for a single dataset. Here we provide a larger sample (Figure 20). We explore the errors created by the models trained on 0% and 100% data, searching for trends or divergences between the input instances that achieve a low prediction accuracy or score. Our investigation finds no major distinguishing factors between them, leaving a more fine-grained study of the effect of purely synthetic data on model decision-making to future work.

E Implementation Details

While our full code implementation can be seen in the GitHub repository (to be released after review), we list the key implementation details below.

Hardware and Systems Used: The experiments were run on a cluster that included nodes with: five A40 GPUs (48GB), three RTX 2080Tis, and a separate machine using a single A100 GPU.

Prompt Models used: We used GPT-3.5-Turbo and GPT-4-Turbo Batch APIs from OpenAI. Generations were obtained at various points from August 2024 to September 2024.

Fine-Tuning Models Used: We used two fine-tuning models in our experiments. Llama3 used the Llama3.1-8B HuggingFace Checkpoint, and Mistral used the Mistral7B-Instruct-v0.2 HuggingFace Checkpoint. We did not conduct an extensive hyperparameter search, however we tried various epochs on smaller samples of the FEVER and

ROPES datasets, selecting that number for every dataset on all experiments. Fact verification models used Adam Optimization with a learning rate of $1e-5$ for two epochs, while QA datasets used a learning rate of $1e-2$ for five epochs.

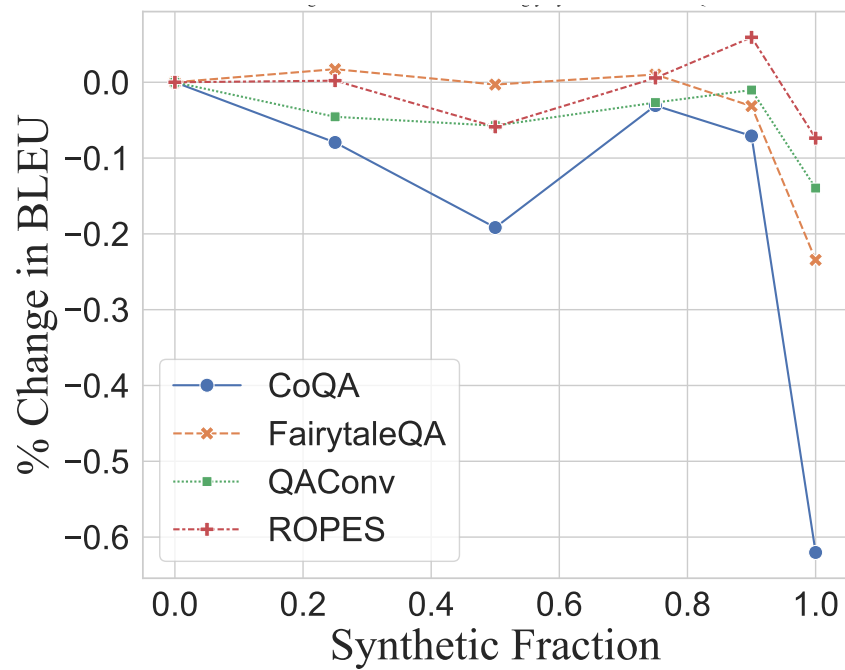



Figure 6: Change in model performance as the proportion of synthetic points in the training data is increased. Across datasets, the performance decrease when moving from 0% to 90% synthetic data is often less than that of moving from 90% to purely synthetic data.

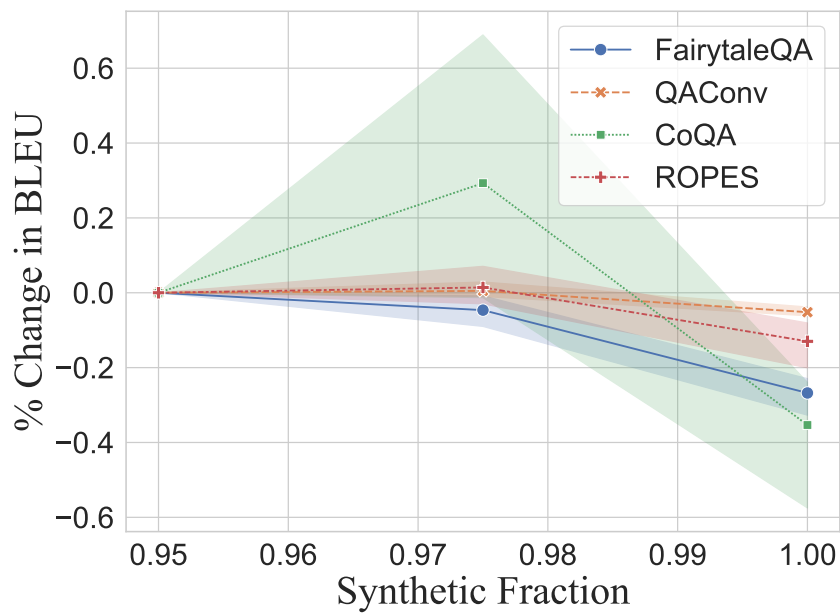
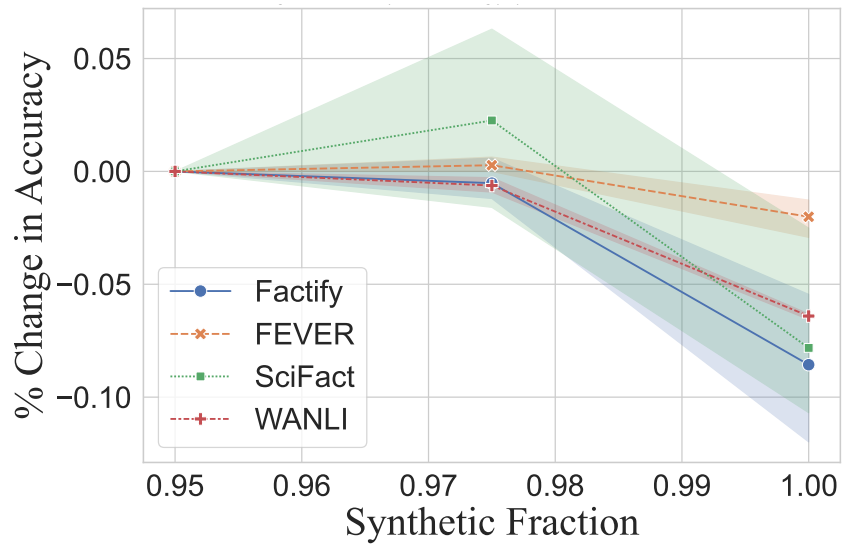


Figure 7: Model performance as the synthetic proportion of the training data varies from 95% to 100%. Across all datasets and random seeds, having **just 2.5%** of the training dataset being human generated boosts performance.

Dataset	Synthetic %	EM	Inc	R Inc	BLEU	ROUGE	BERTScore
CoQA	0	40.6	52.2	60.8	47.9	64.0	87.8
	25	35.0	51.8	54.8	44.1	62.0	85.9
	50	31.6	42.8	60.6	38.7	54.7	79.9
	75	39.2	50.4	69.0	46.4	62.1	79.9
	90	36.2	50.2	58.2	44.5	60.8	85.4
	100	13.6	26.0	58.2	18.2	26.6	52.8
FairytaleQA	0	0.0	0.0	0.0	39.3	55.3	90.8
	25	0.0	0.0	0.0	40.0	56.1	90.5
	50	0.0	0.0	0.0	39.2	55.3	88.9
	75	0.0	0.0	0.0	39.7	55.4	90.6
	90	0.0	0.0	0.0	38.1	54.2	89.9
	100	0.0	0.0	0.0	30.1	49.5	88.4
QAConv	0	29.4	36.3	49.4	35.1	51.6	89.9
	25	28.6	34.5	48.0	33.5	48.7	89.6
	50	28.0	34.0	47.1	33.1	48.7	89.5
	75	28.6	35.9	48.5	34.2	50.5	90.0
	90	29.0	36.0	49.9	34.8	50.6	89.0
	100	23.5	34.3	41.0	30.2	45.2	87.2
ROPES	0	66.8	67.4	72.2	67.0	72.7	96.0
	25	66.8	67.5	69.8	67.2	71.1	96.2
	50	62.8	63.4	65.5	63.1	66.6	95.2
	75	66.8	68.0	68.8	67.4	70.8	96.2
	90	70.6	71.5	71.8	71.0	73.0	96.8
	100	60.8	63.9	61.2	62.1	64.9	95.3

Table 3: Full Results for the QA datasets. There is overwhelming agreement between all metrics on the ranking between models trained on different synthetic fractions. EM: Exact Match, Inc: String Inclusion, R Inc: Reverse String Inclusion

Run	Dataset	Synthetic %	BLEU	ROUGE	BERTScore
0	FairytaleQA	95	38.78	54.80	90.34
		97.5	37.19	52.09	86.17
		100	26.10	43.82	77.03
	QAConv	95	34.45	51.23	90.35
		97.5	34.31	51.89	89.80
		100	32.33	49.39	89.77
	CoQA	95	25.33	35.64	59.39
		97.5	42.78	57.88	78.84
		100	19.11	27.57	51.44
	ROPES	95	72.89	77.26	97.22
		97.5	70.74	73.83	96.58
		100	58.28	60.82	94.37
1	FairytaleQA	95	39.50	54.67	90.57
		97.5	35.95	53.10	89.73
		100	29.75	47.05	81.86
	QAConv	95	38.94	56.41	90.79
		97.5	40.06	57.64	90.45
		100	37.46	54.11	88.57
	CoQA	95	34.84	47.16	63.35
		97.5	41.91	56.53	85.30
		100	14.81	24.58	50.02
	ROPES	95	69.61	71.97	96.23
		97.5	69.72	72.80	96.58
		100	62.08	65.44	95.03
2	FairytaleQA	95	37.97	53.98	90.34
		97.5	37.65	52.44	87.83
		100	29.26	49.70	88.56
	QAConv	95	38.07	54.49	90.13
		97.5	37.70	54.64	88.61
		100	35.94	51.84	89.49
	CoQA	95	30.80	42.03	62.72
		97.5	30.40	41.02	56.83
		100	23.42	37.20	63.31
	ROPES	95	63.46	65.80	95.28
		97.5	67.94	71.47	96.16
		100	58.34	61.62	94.16

Table 4: Results on $n = 5000$ from 95% to 100% for the QA datasets. There is overwhelming agreement between all metrics on the ranking between models trained on different synthetic fractions.

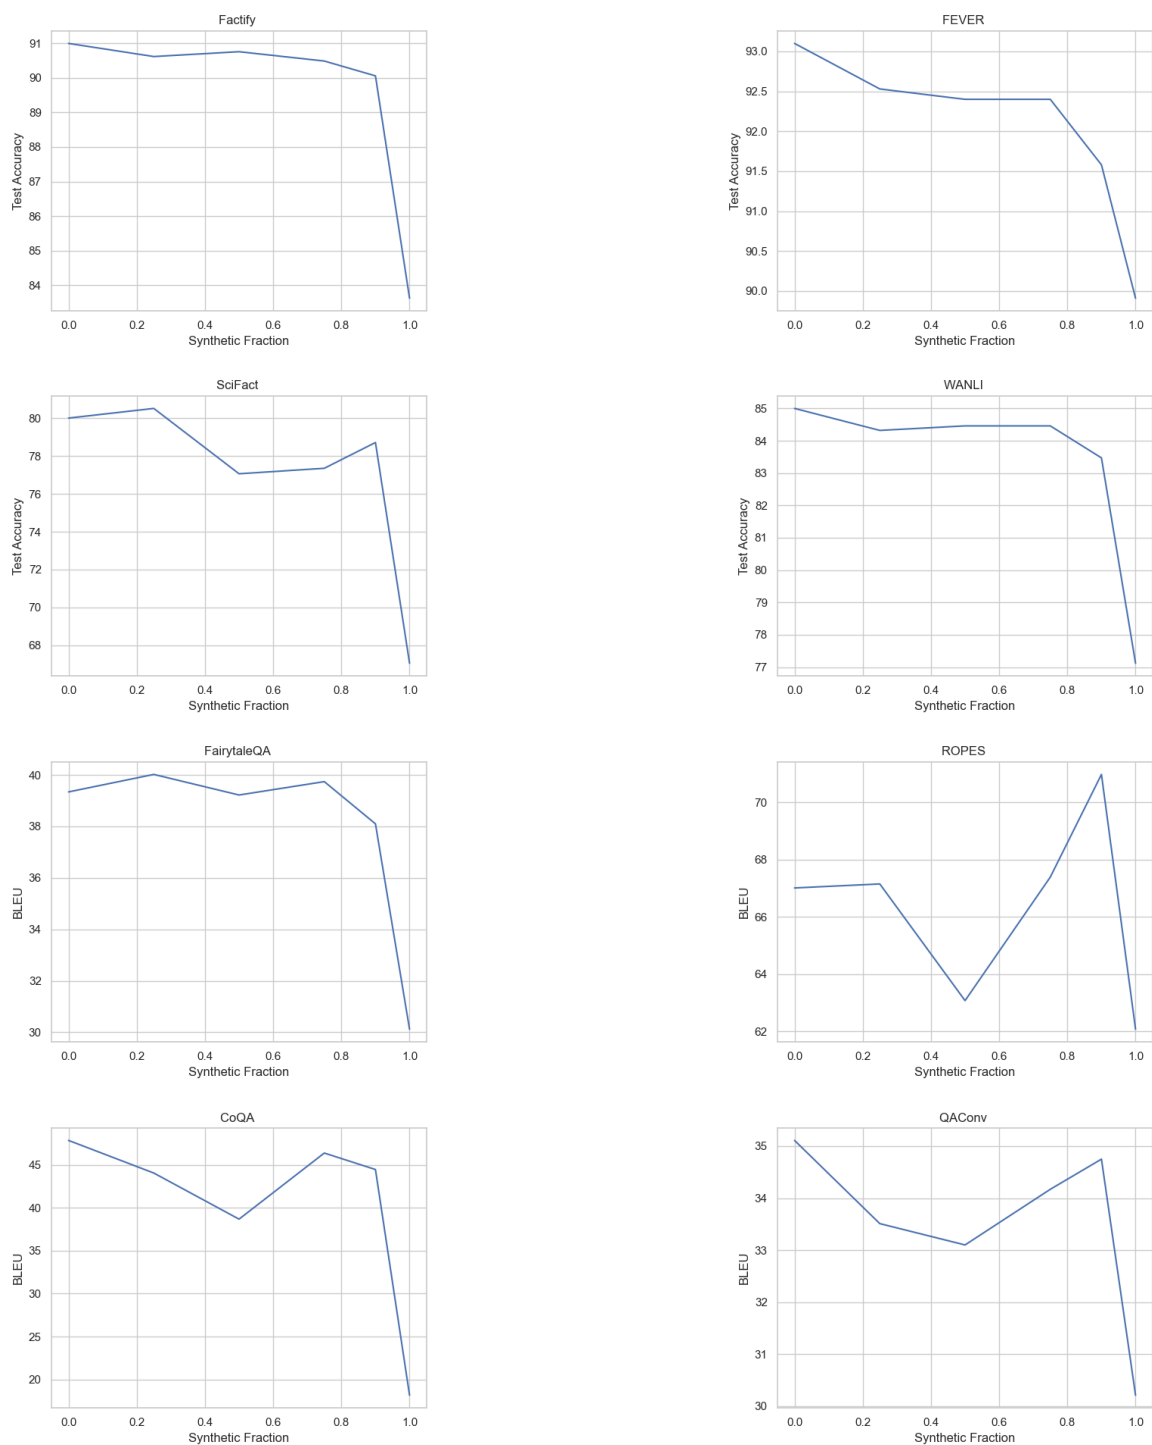


Figure 8: Change in model performance as the proportion of synthetic points in the training data is varied. Across datasets, the performance decrease when moving from 0% to 90% synthetic data is often less than that of moving from 90% to purely synthetic data.

Train Sets	Test Set	Synthetic %	Test Accuracy
FEVER, SciFact	WANLI	0	69.98
		25	67.56
		50	65.86
		75	64.82
		100	64.36
WANLI, SciFact	FEVER	0	83.01
		25	80.64
		50	79.22
		75	78.94
		100	76.01
FEVER, WANLI	SciFact	0	71.76
		25	69.75
		50	69.82
		75	66.42
		100	64.41

Table 5: Test accuracy when replacing human data with synthetic data in the out-of-distribution setting. Using more synthetic data leads to performance declines even in the OOD setting, showing that human data is of higher quality and the results from the main text cannot be explained by a spurious correlation between the human test and human training samples.

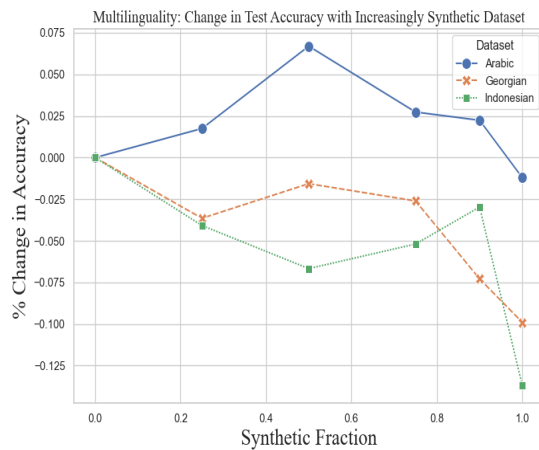


Figure 9: Change in model performance as the proportion of synthetic points in the training data is increased on multilingual fact verification datasets (splits of X-Fact). We observe the same trend as those from earlier experiments, confirming that our results are not limited to the English language. While the phenomenon is reproduced, the threshold of replacement at which we observe a precipitous decline is not the same across languages.




Figure 10: Model performance as the synthetic proportion of the training data varies from 95% to 100%. Across all datasets and random seeds, having **just 2.5%** of the training dataset being human generated boosts performance.

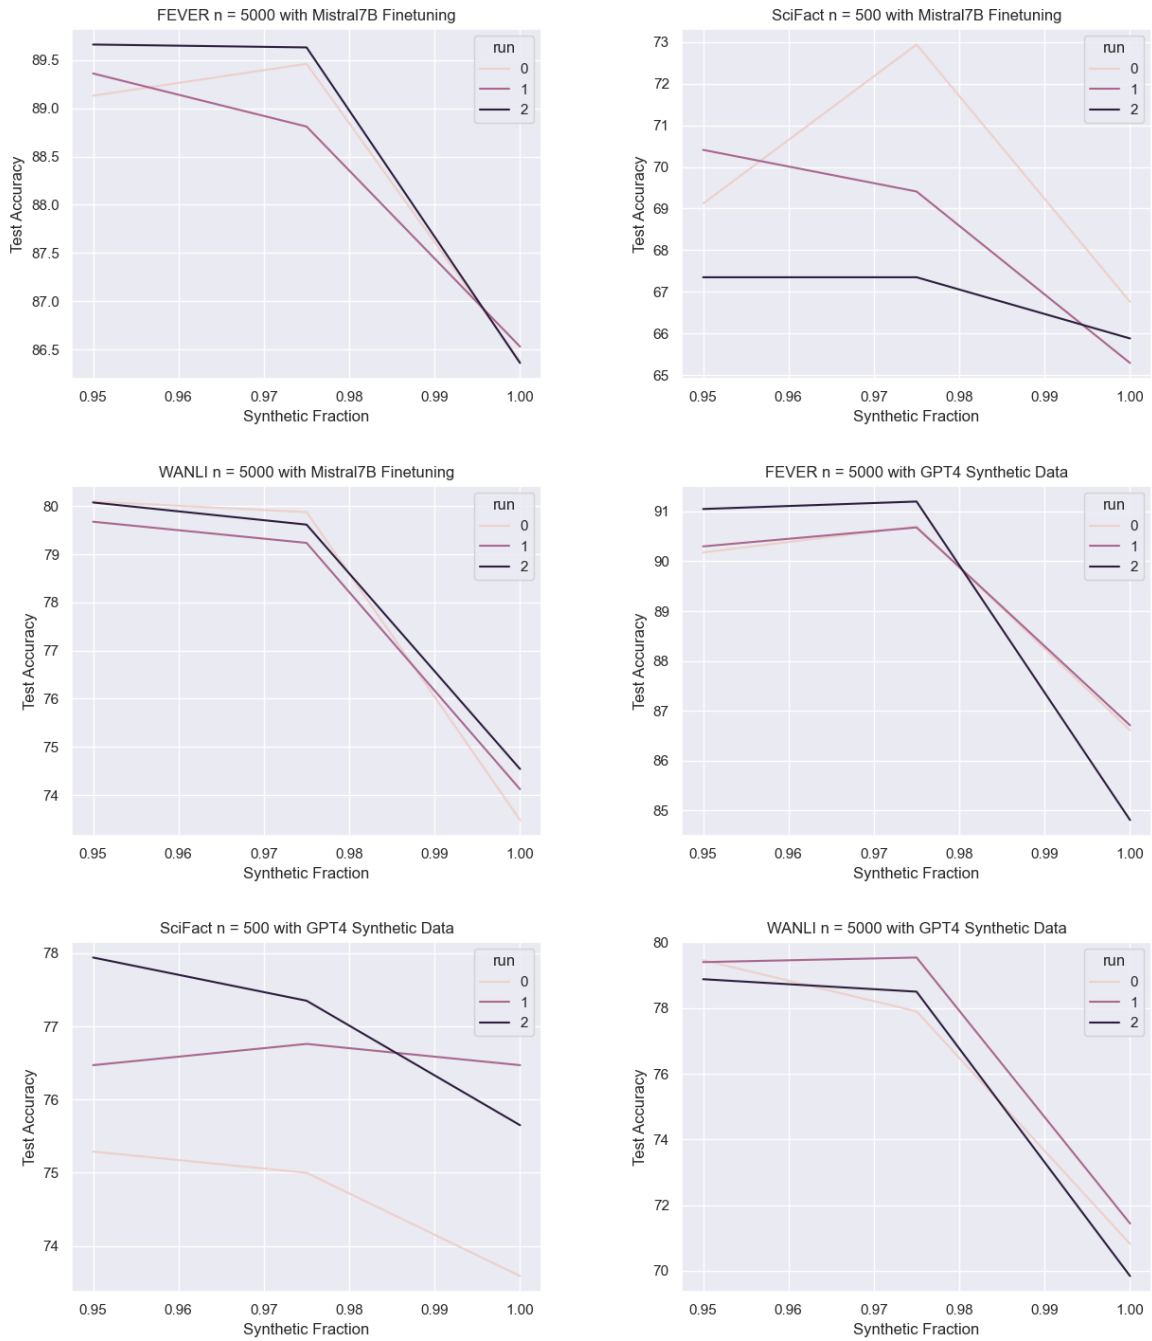


Figure 11: Results hold consistently on Fact Verification datasets when using Mistral7B as the fine-tuning model and GPT-4 as the prompting model.

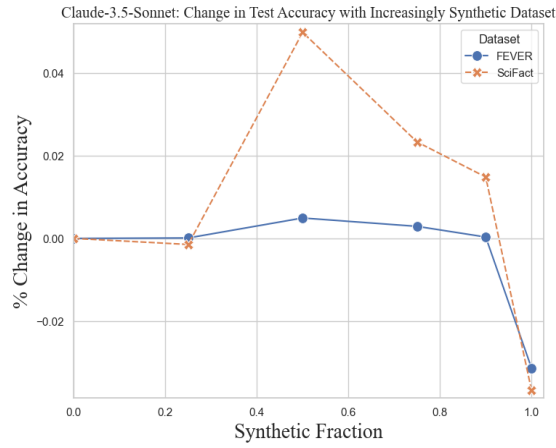


Figure 12: Results hold when using Claude-3.5-Sonnet as the prompting model, showing that the phenomenon is not particular to Synthetic Data from GPT based models.

Dataset	Claim / Question	
	Synthetic	Human
FEVER	35.78	42.76
WANLI	15.15	20.10
SCIFACT	7.12	20.92
FACTIFY	14.50	23.93
NarrativeQA	30.92	8.25
CoQA	7.08	8.39
FairyTaleQA	22.59	16.85
ROPES	28.73	41.42

Table 6: 4-Gram overlap % between all synthetic and human generated claims / questions for each dataset. On several datasets, synthetic claims have a lower overlap

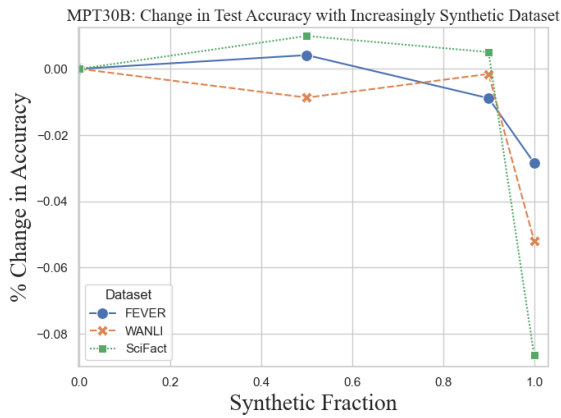
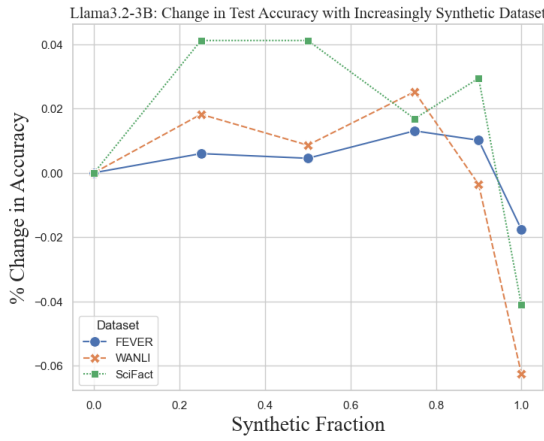
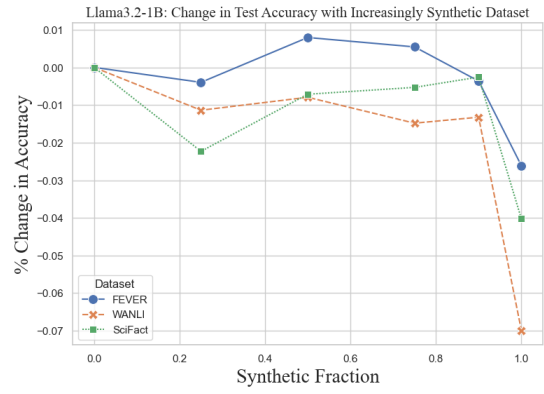


Figure 13: Results hold consistently on Fact Verification datasets when using models of different scales.




Figure 14: Results hold consistently on Question Answering datasets when using Mistral7B as the fine-tuning model and GPT-4 as the prompting model




Figure 15: Results hold when using Chain-Of-Thought Prompting on GPT-3.5

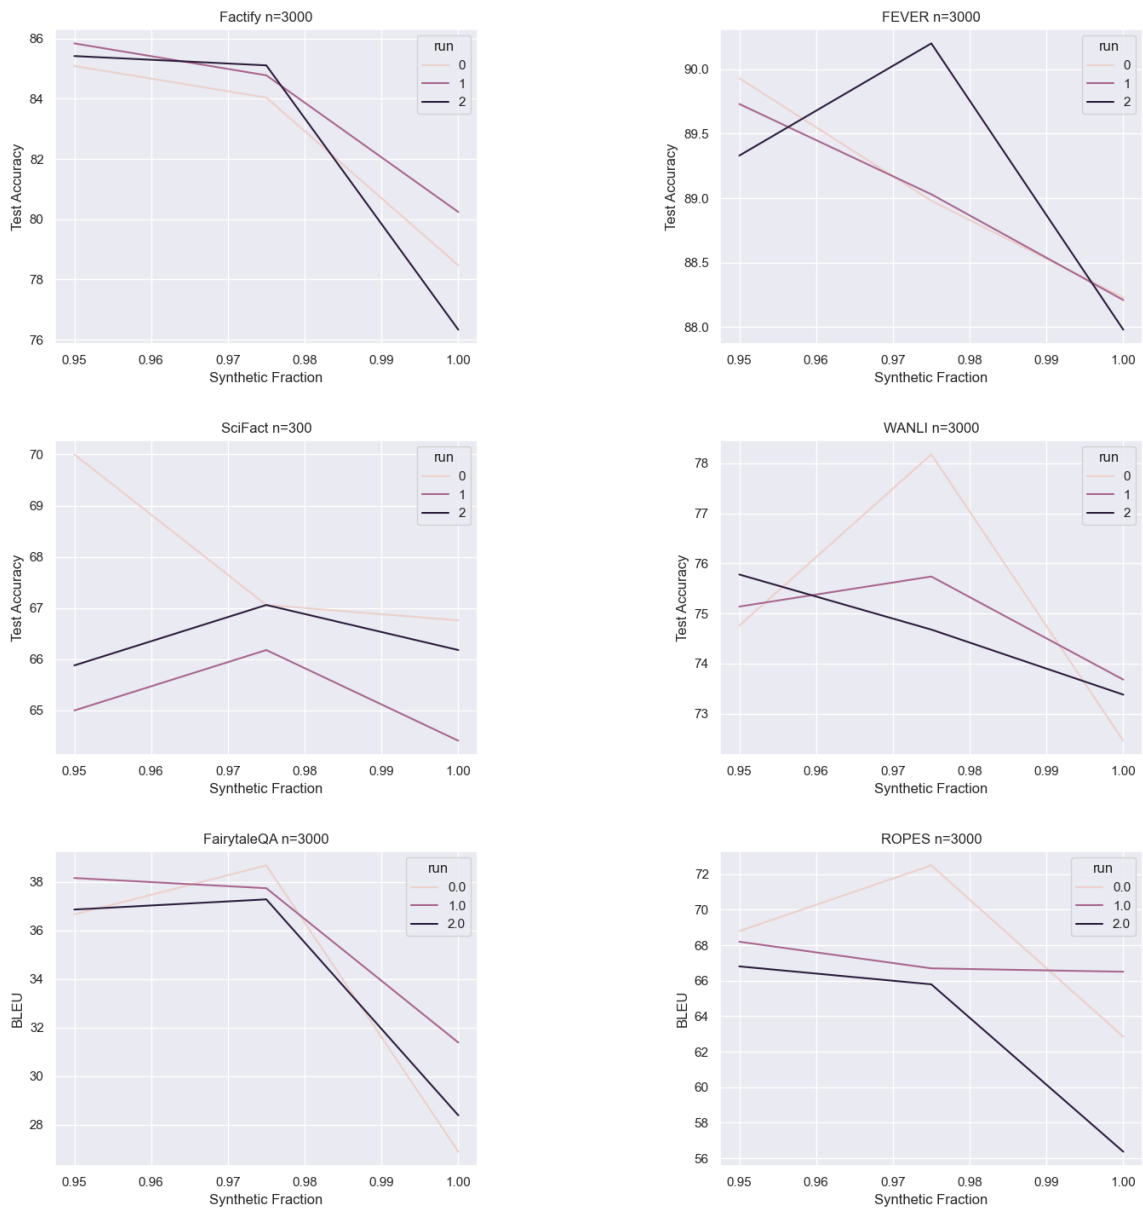


Figure 16: Model performance as the synthetic proportion of the training data varies from 95% to 100% with total number of points $n = 3000$. Across all runs on all datasets including **just 75 real datapoints** can boost performance.

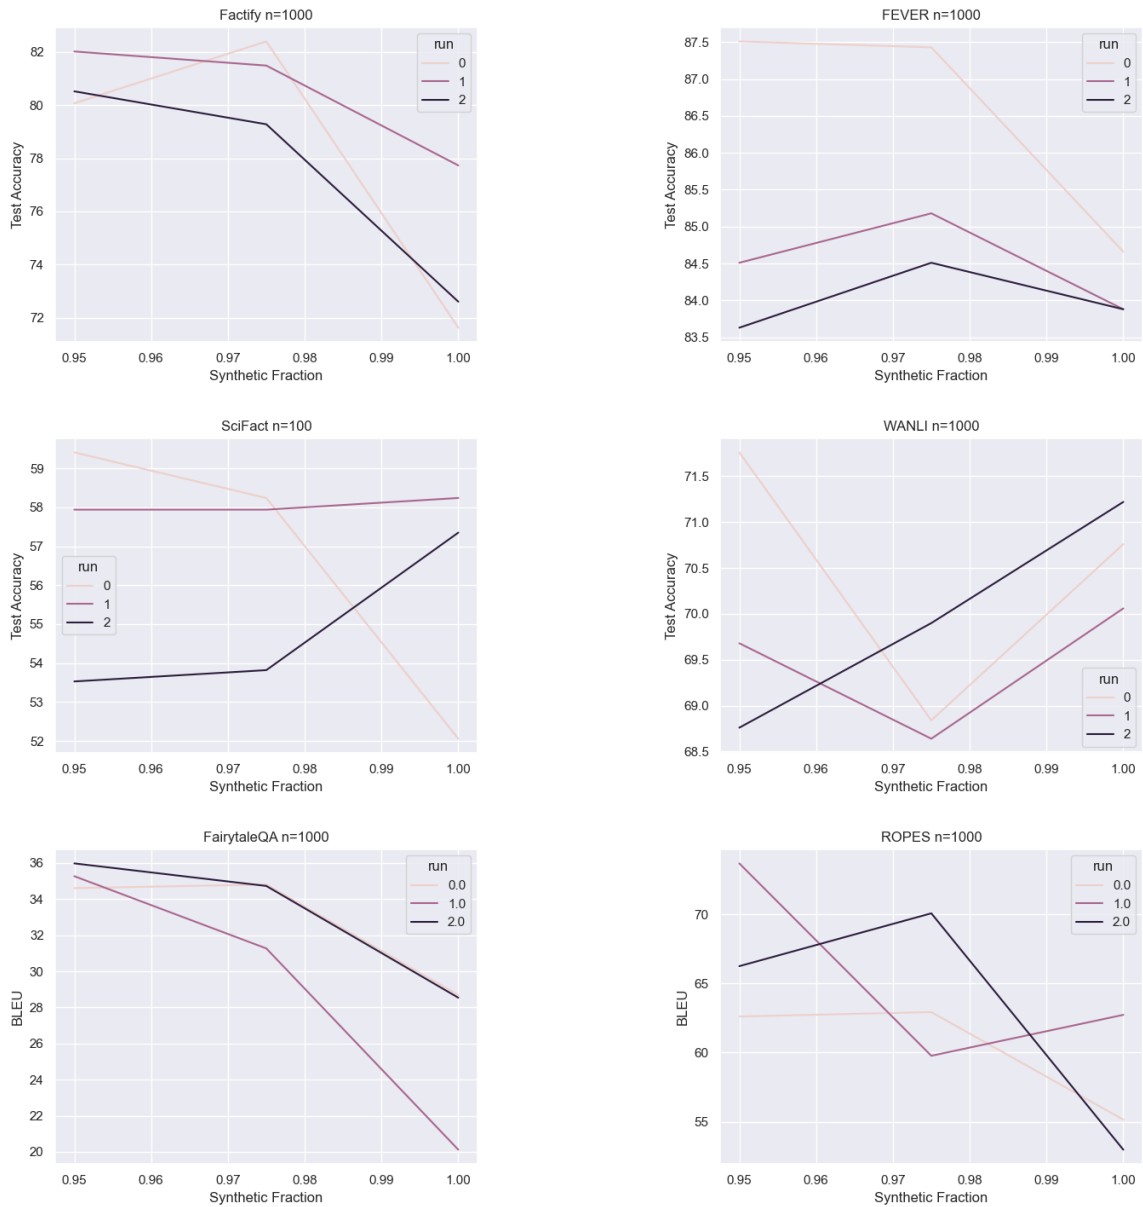


Figure 17: Model performance as the synthetic proportion of the training data varies from 95% to 100% with total number of points $n = 1000$. While the most common trend is that including real data improves performance, the results are much more unstable.




Figure 18: Adding 200 real data points is as effective as adding an order of magnitude more synthetic data points.

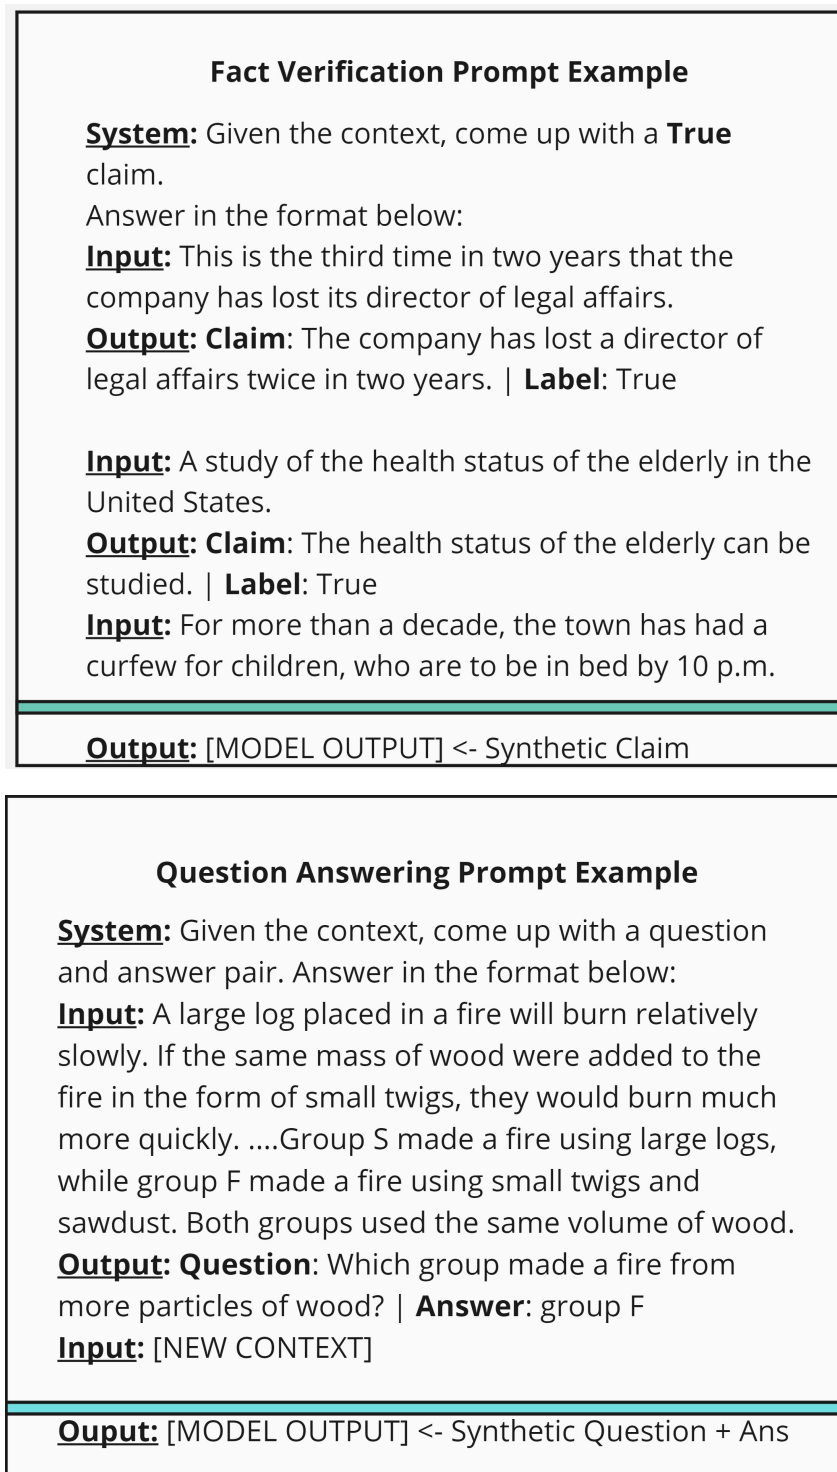


Figure 19: Example prompts used to synthetically generate (claim, label) or (question, answer) pairs using a new context / evidence text.

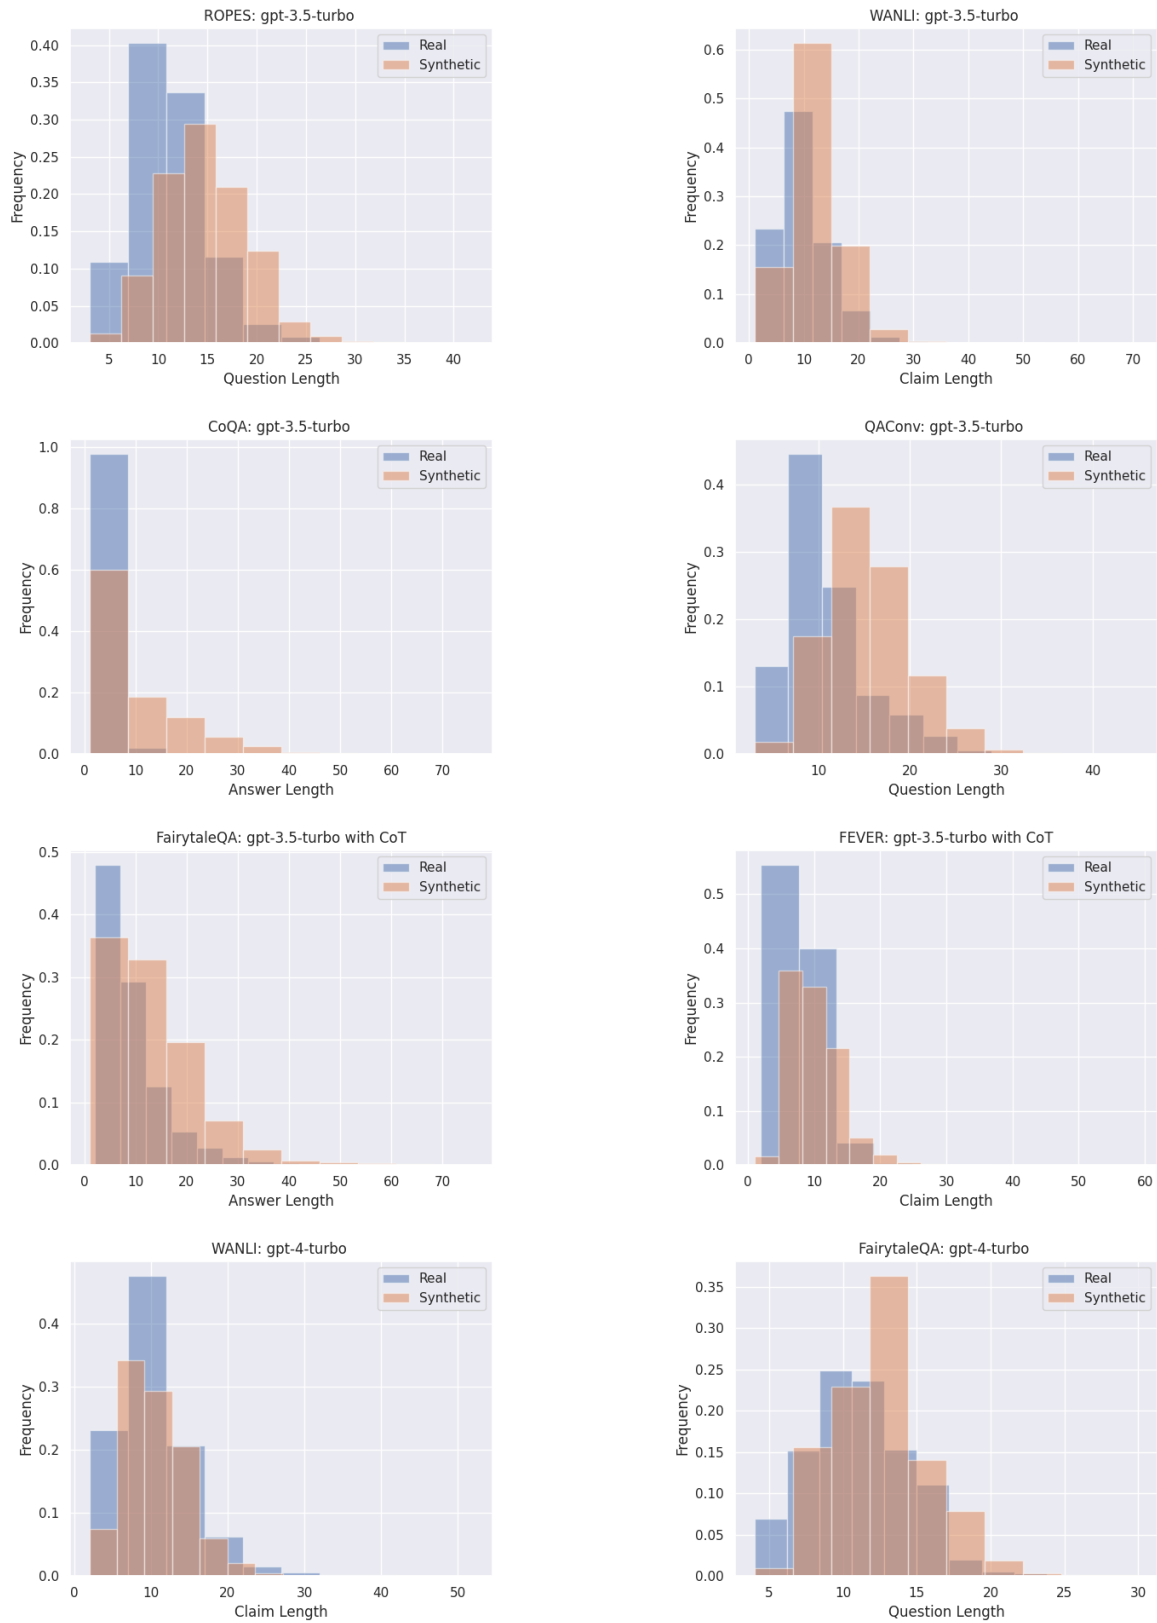


Figure 20: Synthetic data is, on average, longer than its human generated counterpart. This trend can be seen on FV (claims) and QA (claims and questions), and holds across prompt models and strategies.




Figure 21: Synthetic data generally exhibits a higher maximum BLEU score measured against sentences from the context. This suggests that synthetic questions, answers, and claims are more extractive than their human generated counterparts




Figure 22: Synthetic data typically chooses more diverse sources (in terms of answer location or claim location in the evidence text), while humans tend to favor the start of the evidence text.