

MULTI-NEWS⁺: Cost-efficient Dataset Cleansing via LLM-based Data Annotation

Juhwan Choi¹, Jungmin Yun¹, Kyohoon Jin² and YoungBin Kim^{1,2}

¹Department of Artificial Intelligence, Chung-Ang University

²Graduate School of Advanced Imaging Sciences, Multimedia and Film, Chung-Ang University
{gold5230, cocoro357, fhzh123, ybkim85}@cau.ac.kr

Abstract

The quality of the dataset is crucial for ensuring optimal performance and reliability of downstream task models. However, datasets often contain noisy data inadvertently included during the construction process. Numerous attempts have been made to correct this issue through human annotators. However, hiring and managing human annotators is expensive and time-consuming. As an alternative, recent studies are exploring the use of large language models (LLMs) for data annotation.

In this study, we present a case study that extends the application of LLM-based data annotation to enhance the quality of existing datasets through a *cleansing* strategy. Specifically, we leverage approaches such as chain-of-thought and majority voting to imitate human annotation and classify unrelated documents from the Multi-News dataset, which is widely used for the multi-document summarization task. Through our proposed cleansing method, we introduce an enhanced MULTI-NEWS⁺. By employing LLMs for data cleansing, we demonstrate an efficient and effective approach to improving dataset quality without relying on expensive human annotation efforts.

1 Introduction

The significance of dataset quality in deep learning applications cannot be overstated as mislabeled or noisy data can severely degrade performance (Song et al., 2023). Datasets with incorrect labels, noise, or inconsistencies undermine the consistency and stability of model training. Cleansing these datasets contributes to enhancing model performance and generalization capabilities. Hence, ensuring the quality of the dataset by identifying and eliminating noisy data is imperative. In the realm of natural language processing, several researchers have attempted to improve the quality of noisy datasets (Jiang et al., 2020, 2022). For example, ReDoCRED (Tan et al., 2022) addressed issues such as

Source 1
Starting in 1996, alexa internet has been donating their crawl data to the internet archive. Flowing in every day, these data are added to the wayback machine after an embargo period.
Source 2
... For the first time in decades, researchers trying to develop a vaccine for malaria have discovered a new target they can use to attack this deadly and common parasite...
Source 3
Focused crawls are collections of frequently updated webcrawl data from narrow (as opposed to broad or wide) web crawls, often focused on a single domain or subdomain.
Summary
Researchers think they've found a promising new potential weapon in the fight against malaria in a fairly unlikely place: the blood of toddlers. In a paper published in science today, ...

Table 1: Examples of noisy documents in Multi-News dataset. Sources 1 and 3 do not contribute to the summary. We aim to identify such noisy documents without a human annotator.

false negatives in DocRED (Yao et al., 2019), a widely used dataset for relation extraction. Similarly, annotation inconsistencies were found in the MultiWOZ dataset (Budzianowski et al., 2018) for dialogue state tracking (Qian et al., 2021), leading to efforts to rectify these issues (Eric et al., 2020; Zang et al., 2020; Han et al., 2021; Ye et al., 2022a).

Despite these efforts, relying on human annotators to enhance datasets poses challenges such as high costs and time constraints. The quality of the annotation might also be affected by potential variations, such as subjective bias and the proficiency of the annotator (Rashtchian et al., 2010). Furthermore, cleansing a noisy dataset typically requires a larger budget, often involving majority voting by multiple annotators or validation by experts (Tan et al., 2022). Given the significance and necessity of enhancing the quality of existing datasets, these obstacles hinder practical efforts to cleanse datasets efficiently. Therefore, it is crucial to explore cost-effective methods that can cleanse the

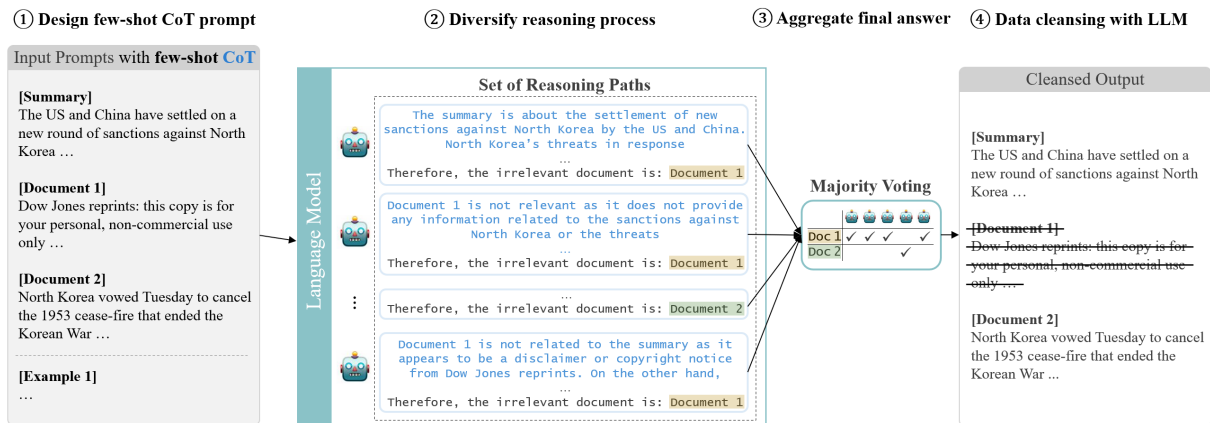


Figure 1: Overall framework for cleansing data and composing MULTI-NEWS⁺.

existing dataset, minimizing human involvement.

In this study, we propose leveraging large language model (LLM)-based annotation for dataset cleansing. Researchers have explored cost-efficient alternatives to human annotators by employing LLMs across various tasks (Wang et al., 2021; Ding et al., 2023; He et al., 2024; Bansal and Sharma, 2023; Zhang et al., 2023; Choi et al., 2024). However, the real-world applicability of LLM-based annotation on existing datasets is still less explored. Building on these insights, we extend the application of LLM-based annotations to denoise the existing dataset and improve its quality. Specifically, we conduct a case study to cleanse the Multi-News (Fabbri et al., 2019), a dataset for multi-document summarization tasks. This dataset consists of news articles crawled from the internet and is widely used in multi-document summarization research. However, as shown in Table 1, we identify several issues related to the noise in the dataset. For instance, the set of documents contained system messages from platforms such as Twitter, Wayback Machine, or Dow Jones that are unrelated to the summary and degrade the dataset quality.

To accomplish our purpose, we utilize LLMs to analyze the summary and associated documents, identifying and excluding any documents that are not relevant to the summary. Specifically, we employ approaches such as chain-of-thought (CoT), providing the rationale for decision-making with enhanced transparency and facilitating human investigation. We further enhance our cleansing process by incorporating self-consistency considerations, which mimic the majority voting process used by human annotators (Wang et al., 2023b). Based on our carefully designed framework, we introduce MULTI-NEWS⁺, an enhanced version of

the existing Multi-News dataset, achieved through our LLM-based cleansing strategy. To the best of our knowledge, this is the first attempt to exploit LLMs to enhance the quality of real-world datasets. Our experiments demonstrate the effectiveness of MULTI-NEWS⁺, providing a valuable resource for future research. We make MULTI-NEWS⁺ and our source code publicly available for further study.

2 Related Work

Dataset quality has been an interest to researchers because of its importance in ensuring the quality of the model trained with the dataset (Budach et al., 2022). Previous studies found that large amounts of data automatically crawled from the web may contain noisy documents, and proper filtering procedures can be an efficient solution against them (Xu and Koehn, 2017; Khayrallah and Koehn, 2018; Kryściński et al., 2019; Luccioni and Viviano, 2021; Kreutzer et al., 2022). Accordingly, several studies in text summarization investigated various strategies to filter out noisy data (Matsumaru et al., 2020; Nan et al., 2021; Guo et al., 2022) and released new datasets with better quality (Grusky et al., 2018; Urlana et al., 2022). However, their strategies are primarily composed of coarse rule-based methods and less interpretable model output, or costly human investigation has been applied for constructing new datasets. Furthermore, such strategies have not been applied to multi-document summarization datasets.

In the meantime, with the advancement of LLMs (Zhao et al., 2023), researchers have explored the usage of LLMs for data annotation, a task that traditionally relied on human annotators. Initial attempts have revealed the potential capabilities of models like GPT-3 for data annotation (Wang




Figure 2: Histogram comparing the amount of input articles in each dataset.

et al., 2021). These studies indicate that GPT-3 can annotate datasets more efficiently and cost-effectively than human annotators. This results in enhanced downstream task performance, with the model trained on the GPT-3 annotated dataset outperforming the one trained on the human-annotated dataset. Subsequent studies have further demonstrated the capabilities of GPT-3, showing its ability to generate labeled data using external knowledge or instructions about desired labels and domains (Ding et al., 2023). Additionally, researchers have examined the usefulness of newer models like GPT-3.5 and evaluated the effectiveness of CoT in improving annotation quality (He et al., 2024). LLM-based annotation has also been extended to low-resource languages where hiring human annotators is challenging (Choi et al., 2024).

In this work, we introduce a novel approach to filtering noisy documents from multi-document summarization dataset by extending cost-efficient LLM-based annotation beyond traditional data annotation tasks. By leveraging the capabilities of LLMs, our study facilitates real-world dataset cleansing, enhancing the quality of existing datasets. This attempt is noteworthy as it broadens the scope of LLM applications, offering effective solutions for improving dataset quality and streamlining its cleansing process, minimizing reliance on human annotations.

3 MULTI-NEWS⁺

The previous Multi-News dataset plays an important role in multi-document summarization research. It consists of sets of documents and their corresponding summaries. However, as shown in Table 1 and detailed in Appendix G and H, the

Multi-News dataset contains several noisy and irrelevant articles that are unrelated to the summary or other documents. This issue arises from their construction process, which relies on automated crawling from the Internet Archive.

To solve this issue and cleanse the dataset, we defined our problem as a classification task determining whether each document is relevant to the summary. To this end, we designed the prompt for the model as shown in Appendix J. We integrated CoT to enhance the model’s performance by evaluating the relevance of each document to the summary. Thus, a *rationale* for the decision can be made available, which marks the difference between LLM-based and human annotations. While traditional human annotation through crowdsourcing platforms like Amazon Mechanical Turk usually produces annotation results without underlying reasons due to additional costs, LLM-based annotators can easily offer explanations through CoT. These rationales can assist human managers in reviewing results and rectifying erroneous decisions.

Furthermore, we imitated the conventional dataset cleansing procedure which typically involves multiple human annotators and their collective judgments, primarily through majority voting. Similarly to the majority voting process used by human annotators, we applied this approach to the LLM-based annotators. In particular, we generated five individual LLM agents to read the summary and documents and determine if the document is relevant to the summary. This strategy based on self-consistency can boost the quality of annotations, by rectifying potential errors made by individual agents (Wang et al., 2023b). Figure 1 presents the summary of the overall process.

Based on the proposed method, we utilized five LLM agents to individually annotate 56,216 sets of summaries and documents from the Multi-News dataset. Specifically, we employed the GPT-3.5-turbo-0125 model¹, the most recent model at the time of this study. With a prompt designed for a 3-shot CoT, approximately 3,500 tokens were required to annotate the input summaries and articles, along with around 100 tokens for generating reasoning processes and annotation results. The cost per annotation sample amounted to approximately 0.01\$ (0.002\$ per agent), resulting in a total cost of approximately 550\$ to annotate the

¹GPT-3.5-turbo-0125 charges 0.0005\$ for the input of 1,000 tokens, and 0.0015\$ for the generation of 1,000 tokens.

Model Metric	<i>BART-large-cnn</i>				
	ROUGE-1	ROUGE-2	ROUGE-L	BERTScore	BARTScore
Multi-News	48.64	18.86	24.11	0.6401	-2.763
MULTI-NEWS ⁺	49.17	19.04	24.36	0.6418	-2.698
Ablation (Urlana et al., 2022)	47.48	18.27	23.81	0.6362	-2.767
Model Metric	<i>T5-base</i>				
	ROUGE-1	ROUGE-2	ROUGE-L	BERTScore	BARTScore
Multi-News	40.11	13.90	21.58	0.6003	-2.407
MULTI-NEWS ⁺	40.45	14.17	21.84	0.6027	-2.362
Ablation (Urlana et al., 2022)	39.30	13.65	21.42	0.5967	-2.457

Table 2: Performance comparison of the Multi-News and MULTI-NEWS⁺ datasets on two models. The ‘‘Ablation’’ row represents a version of the Multi-News dataset that has been cleansed using methods from previous study (Urlana et al., 2022).

entire Multi-News dataset.

After annotation, we found that 27,052 of the 153,091 articles can be considered noisy documents and do not contribute to the summarization. Subsequently, we constructed MULTI-NEWS⁺ by removing these noisy documents from Multi-News while preserving the train/valid/test split. Figure 2 presents the comparison of the Multi-News and MULTI-NEWS⁺ datasets in terms of the number of documents per set. More than 15% of the documents in Multi-News are irrelevant, diminishing the dataset’s quality and degrading the model’s performance. Furthermore, 379 sets have no relevant source articles, as shown in Appendix H. In contrast, by deleting noisy documents, MULTI-NEWS⁺ demonstrates enhanced quality.

4 Experiment

4.1 Experimental Design

To validate the efficacy of data cleansing and the development of MULTI-NEWS⁺ in filtering out noisy documents and improving the performance of downstream task models, we measured the multi-document summarization performance of models trained on each dataset, similar to previous study (Guo et al., 2022). Enhanced model performance indicates superior dataset quality (Ye et al., 2022b; Choi et al., 2024). We fine-tuned two different models, BART (Lewis et al., 2020) and T5 (Raffel et al., 2020) on Multi-News and MULTI-NEWS⁺. Performance evaluation metrics included the following metrics: ROUGE (Lin, 2004), BERTScore (Zhang et al., 2020), and BARTScore (Yuan et al., 2021). For a fair comparison, we used the test set of MULTI-NEWS⁺ for each model and reported the average performance across three random seeds.

4.2 Result

The results in Table 2 demonstrate the superiority of the MULTI-NEWS⁺ dataset in enhancing the performance of summarization models compared to the original Multi-News dataset. Across various metrics, models trained on MULTI-NEWS⁺ consistently outperform those trained on Multi-News, indicating better summarization quality with the refined dataset. This highlights the effectiveness of dataset cleansing in removing noisy and irrelevant documents, thereby enhancing the overall performance of summarization models. Additionally, we performed a human evaluation on the output of 379 sets that are classified as having no relevant source articles and found that 356 sets are correctly classified, which represents 93.9% of the human-machine agreement rate. We provide an example of error analysis in Appendix I.

Additionally, we conducted an ablation study using the cleansing method proposed by a previous study (Urlana et al., 2022), detailed in Appendix F. Our findings indicate that this method is ineffective in improving downstream task performance on the Multi-News dataset, which focuses on multi-document summarization and differs from the configuration used in the prior study. This underscores the effectiveness of our proposed method and the value of MULTI-NEWS⁺.

5 Discussion and Future Works

In this section, we discuss recent advancements in the field since the submission of the manuscript and propose strategies for incorporating them in future research.

Cutting-edge models. Although we employed five GPT-3.5-turbo-0125 models for our experiments, the field has seen the release of more

advanced models, such as GPT-4o (OpenAI, 2024b), GPT-4o-mini (OpenAI, 2024a), and OpenAI O1 (OpenAI, 2024c), along with the continued development of open-source models like LLaMA-3 (Dubey et al., 2024), Gemma-2 (Team et al., 2024), and Mistral Nemo (Mistral, 2024). Models such as GPT-4o-mini and other open-source alternatives offer reduced costs compared to GPT-3.5-turbo-0125, making their adoption promising for both lowering the expense of dataset cleansing and improving the accuracy of detecting noisy documents.

Weighted majority voting. The availability of high-performance yet cost-effective models like GPT-4o presents the opportunity to use them as *expert* annotators, given their superior capabilities compared to models like GPT-3.5-turbo-0125 or GPT-4o-mini. For example, rather than using five GPT-3.5-turbo-0125 models, we could employ three GPT-3.5-turbo-0125 models alongside one GPT-4o, with GPT-4o carrying double the weight of a GPT-3.5-turbo-0125 annotator. This approach positions GPT-4o as an expert, where agreement between at least one GPT-3.5-turbo-0125 model and GPT-4o would trigger document deletion.

Supervision from superior models. Another potential approach involves using more capable models to verify annotation results. In this scenario, GPT-4o would not participate in the initial annotation process but would instead verify the outcomes produced by GPT-3.5-turbo-0125 models. By taking the documents, summaries, and annotation results as input, GPT-4o acts as an expert reviewer overseeing the outputs of standard annotators.

Cost-efficient cleansing via pre-screening. In this paper, we applied the data cleansing strategy to every document in the dataset. However, a more cost-efficient approach could involve performing the annotation procedure only on documents likely to contain noise. Techniques such as dataset cartography (Swayamdipta et al., 2020) could serve as a pre-screening method to identify cleansing candidates, thereby reducing the overall cost of dataset cleansing.

6 Conclusion

In this study, we suggest deploying cost-efficient LLM-based data annotation to cleanse real-world

datasets by identifying and excluding irrelevant and noisy data. We conducted a case study using this strategy to cleanse the Multi-News dataset and proposed the improved MULTI-NEWS⁺ dataset. Our case study revealed that MULTI-NEWS⁺ provides superior data quality compared to the original Multi-News dataset. Additionally, we have made MULTI-NEWS⁺ publicly available, thereby supporting further research in the field of multi-document summarization.

Our work paves the road to extending our data cleansing strategy to other datasets, broadening the scope of utilizing LLMs. This extension would enhance the quality of existing datasets across various domains without the need to construct new datasets from scratch. As such, our approach not only contributes to the advancement of multi-document summarization research but also offers a cost-efficient solution for enhancing dataset quality. We are committed to extending our LLM-based method to other datasets, further solidifying its applicability to other tasks.

Limitations

We acknowledge several limitations regarding our proposed method. First, our method is primarily limited by the possibility of wrong classification even with majority voting and CoT. In the future, we may adopt various LLMs as agents and apply weighted majority voting according to their performance to alleviate this issue, as discussed in Section 5.

Secondly, the nature of the Multi-News dataset might exhibit a real-world case of automatic collection of documents from the web that are not always relevant to the summary. In other words, the inclusion of noisy documents might demonstrate the characteristics of real-world automatic crawling. For instance, the model trained on the Multi-News dataset may be more suitable for a real-time system that automatically crawls data from the web and summarizes them. However, we believe such a possibility can be dealt with through the reciprocal usage of our MULTI-NEWS⁺ and previous Multi-News dataset. For instance, one could utilize a previous Multi-News dataset when the trained model is expected to consistently deal with noisy documents for inference and there are no pre-defined strategies for filtering out these noisy documents at inference time. Otherwise, for cases where the model is expected to only handle clean documents,

it will be more beneficial to utilize our proposed MULTI-NEWS⁺ dataset for training the model.

Ethics Statement

As we are exploiting LLMs for classifying irrelevant documents rather than text generation, the ethical concern with our method is smaller than that of studies that utilize LLMs to generate texts. Nonetheless, recent studies suggest that the CoT technique may induce ethical bias in LLM (Shaikh et al., 2023). In future work, we plan to investigate this phenomenon’s appearance in our method.

Acknowledgements

This research was supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education(NRF-2022R1C1C1008534), and Institute for Information & communications Technology Planning & Evaluation (IITP) through the Korea government (MSIT) under Grant No. 2021-0-01341 (Artificial Intelligence Graduate School Program, Chung-Ang University).

References

- Parikshit Bansal and Amit Sharma. 2023. [Large language models as annotators: Enhancing generalization of nlp models at minimal cost](#). *arXiv preprint arXiv:2306.15766*.
- Lukas Budach, Moritz Feuerpfeil, Nina Ihde, Andrea Nathansen, Nele Noack, Hendrik Patzlaff, Felix Naumann, and Hazar Harmouch. 2022. The effects of data quality on machine learning performance. *arXiv preprint arXiv:2207.14529*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. [Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling](#). In *Proceedings of EMNLP*, pages 5016–5026.
- Juhwan Choi, Eunju Lee, Kyohoon Jin, and YoungBin Kim. 2024. [GPTs are multilingual annotators for sequence generation tasks](#). In *Findings of EACL*, pages 17–40.
- Bosheng Ding, Chengwei Qin, Linlin Liu, Yew Ken Chia, Boyang Li, Shafiq Joty, and Lidong Bing. 2023. [Is GPT-3 a good data annotator?](#) In *Proceedings of ACL*, pages 11173–11195.
- 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*.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, Adarsh Kumar, Anuj Goyal, Peter Ku, and Dilek Hakkani-Tur. 2020. [Multiwoz 2.1: A consolidated multi-domain dialogue dataset with state corrections and state tracking baselines](#). In *Proceedings of LREC*, pages 422–428.
- Alexander Richard Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. [Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model](#). In *Proceedings of ACL*, pages 1074–1084.
- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. [Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies](#). In *Proceedings of NAACL*, pages 708–719.
- Yanzhu Guo, Chloé Clavel, Moussa Kamal Eddine, and Michalis Vazirgiannis. 2022. [Questioning the validity of summarization datasets and improving their factual consistency](#). In *Proceedings of EMNLP*, pages 5716–5727.
- Ting Han, Ximing Liu, Ryuichi Takanabu, Yixin Lian, Chongxuan Huang, Dazhen Wan, Wei Peng, and Minlie Huang. 2021. [Multiwoz 2.3: A multi-domain task-oriented dialogue dataset enhanced with annotation corrections and co-reference annotation](#). In *Proceedings of NLPCC*, pages 206–218.
- Xingwei He, Zhenghao Lin, Yeyun Gong, A-Long Jin, Hang Zhang, Chen Lin, Jian Jiao, Siu Ming Yiu, Nan Duan, and Weizhu Chen. 2024. [Annollm: Making large language models to be better crowdsourced annotators](#). In *Proceedings of NAACL (Industry Track)*, pages 165–190.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Chao Jiang, Mounica Maddela, Wuwei Lan, Yang Zhong, and Wei Xu. 2020. [Neural crf model for sentence alignment in text simplification](#). In *Proceedings of ACL*, pages 7943–7960.
- Chao Jiang, Wei Xu, and Samuel Stevens. 2022. [arxivedit: Understanding the human revision process in scientific writing](#). In *Proceedings of EMNLP*, pages 9420–9435.
- Huda Khayrallah and Philipp Koehn. 2018. [On the impact of various types of noise on neural machine translation](#). In *Proceedings of ACL 2018 Workshop on Neural Machine Translation and Generation*, pages 74–83.
- Diederik P Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *Proceedings of ICLR*.

- Julia Kreutzer, Isaac Caswell, Lisa Wang, Ahsan Wahab, Daan van Esch, Nasanbayar Ulzii-Orshikh, Allahsera Tapo, Nishant Subramani, Artem Sokolov, Claytone Sikasote, et al. 2022. [Quality at a glance: An audit of web-crawled multilingual datasets](#). *Transactions of the Association for Computational Linguistics*, 10:50–72.
- Wojciech Kryściński, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Neural text summarization: A critical evaluation](#). In *Proceedings of EMNLP*, pages 540–551.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of ACL*, pages 7871–7880.
- Chin-Yew Lin. 2004. [Rouge: A package for automatic evaluation of summaries](#). In *Proceedings of ACL 2004 Workshop Text Summarization Branches Out*, pages 74–81.
- Alexandra Luccioni and Joseph Viviano. 2021. [What’s in the box? an analysis of undesirable content in the common crawl corpus](#). In *Proceedings of ACL*, pages 182–189.
- Kazuki Matsumaru, Sho Takase, and Naoaki Okazaki. 2020. [Improving truthfulness of headline generation](#). In *Proceedings of ACL*, pages 1335–1346.
- Mistral. 2024. [Mistral nemo](#). Accessed: Sep 21, 2024.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gülçehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence rnns and beyond](#). In *Proceedings of CoNLL*, pages 280–290.
- Feng Nan, Ramesh Nallapati, Zhiguo Wang, Cicero dos Santos, Henghui Zhu, Dejiao Zhang, Kathleen Mckeown, and Bing Xiang. 2021. [Entity-level factual consistency of abstractive text summarization](#). In *Proceedings of EACL*, pages 2727–2733.
- OpenAI. 2024a. [Gpt-4o mini: advancing cost-efficient intelligence](#). Accessed: Sep 21, 2024.
- OpenAI. 2024b. [Hello gpt-4o](#). Accessed: Sep 21, 2024.
- OpenAI. 2024c. [Introducing openai o1-preview](#). Accessed: Sep 21, 2024.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Proceedings of NeurIPS*.
- Kun Qian, Ahmad Beirami, Zhouhan Lin, Ankita De, Alborz Geramifard, Zhou Yu, and Chinnadhurai Sankar. 2021. [Annotation inconsistency and entity bias in multiwoz](#). In *Proceedings of SIGDIAL*, pages 326–337.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Cyrus Rashtchian, Peter Young, Micah Hodosh, and Julia Hockenmaier. 2010. [Collecting image annotations using amazon’s mechanical turk](#). In *Proceedings of NAACL 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*, pages 139–147.
- Omar Shaikh, Hongxin Zhang, William Held, Michael Bernstein, and Diyi Yang. 2023. [On second thought, let’s not think step by step! bias and toxicity in zero-shot reasoning](#). In *Proceedings of ACL*, pages 4454–4470.
- Hwanjun Song, Minseok Kim, Dongmin Park, Yooju Shin, and Jae-Gil Lee. 2023. [Learning from noisy labels with deep neural networks: A survey](#). *IEEE Transactions on Neural Networks and Learning Systems*, 34(11):8135–8153.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A Smith, and Yejin Choi. 2020. [Dataset cartography: Mapping and diagnosing datasets with training dynamics](#). In *Proceedings of EMNLP*, pages 9275–9293.
- Qingyu Tan, Lu Xu, Lidong Bing, Hwee Tou Ng, and Sharifah Mahani Aljunied. 2022. [Revisiting docred-addressing the false negative problem in relation extraction](#). In *Proceedings of EMNLP*, pages 8472–8487.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. [Gemma 2: Improving open language models at a practical size](#). *arXiv preprint arXiv:2408.00118*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Ashok Uralana, Nirmal Surange, Pavan Baswani, Priyanka Ravva, and Manish Shrivastava. 2022. [Tesum: Human-generated abstractive summarization corpus for telugu](#). In *Proceedings of LREC*, pages 5712–5722.
- Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. 2023a. [A survey on large language model based autonomous agents](#). *arXiv preprint arXiv:2308.11432*.

Shuohang Wang, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2021. [Want to reduce labeling cost? gpt-3 can help](#). In *Findings of EMNLP*, pages 4195–4205.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023b. [Self-consistency improves chain of thought reasoning in language models](#). In *Proceedings of ICLR*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of EMNLP (Demo Track)*, pages 38–45.

Hainan Xu and Philipp Koehn. 2017. [Zipporah: a fast and scalable data cleaning system for noisy web-crawled parallel corpora](#). In *Proceedings of EMNLP*, pages 2945–2950.

Yuan Yao, Deming Ye, Peng Li, Xu Han, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, Lixin Huang, Jie Zhou, and Maosong Sun. 2019. [Docred: A large-scale document-level relation extraction dataset](#). In *Proceedings of ACL*, pages 764–777.

Fanghua Ye, Jarana Manotumruksa, and Emine Yilmaz. 2022a. [Multiwoz 2.4: A multi-domain task-oriented dialogue dataset with essential annotation corrections to improve state tracking evaluation](#). In *Proceedings of SIGDIAL*, pages 351–360.

Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. 2022b. [Zerogen: Efficient zero-shot learning via dataset generation](#). In *Proceedings of EMNLP*, pages 11653–11669.

Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [BartScore: Evaluating generated text as text generation](#). In *Proceedings of NeurIPS*, pages 27263–27277.

Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020. [Multiwoz 2.2: A dialogue dataset with additional annotation corrections and state tracking baselines](#). In *Proceedings of ACL 2020 Workshop on NLP for Conversational AI*, pages 109–117.

Ruoyu Zhang, Yanzeng Li, Yongliang Ma, Ming Zhou, and Lei Zou. 2023. [LLMaAA: Making large language models as active annotators](#). In *Findings of EMNLP*, pages 13088–13103.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2020. [BertScore: Evaluating text generation with bert](#). In *Proceedings of ICLR*.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. [A survey of large language models](#). *arXiv preprint arXiv:2303.18223*.




Figure 3: A screenshot of a webpage that is relevant to the article in Appendix H. Multi-News includes the text in the red box instead of the desired content in the blue box.

A Dataset Statistics

MULTI-NEWS⁺ keeps the train/valid/test split of Multi-News, which is 80%, 10%, and 10%. Table 3 displays the number of articles per each split.

	Multi-News		MULTI-NEWS ⁺		% of modification	
	Sets	Articles	Sets	Articles	Sets	Articles
Train	44,972	125,417	44,668	102,057	0.7%	18.6%
Validation	5,622	15,367	5,585	12,509	0.7%	18.6%
Test	5,622	15,505	5,584	12,703	0.7%	18.1%

Table 3: Number of sets and articles per each split.

B Construction Process of Multi-News

In this section, we briefly explain the construction process of the Multi-News dataset. Multi-News is based on data from newser.com² that offers human-written summaries of news articles. Each summary is written by professional human editors and involves several outlinks to the original articles and relevant websites. Multi-News collected this human-written summary and documents from its outlinks, which behave as source documents for summarization. Notably, the authors of Multi-News archived every article leveraging Wayback Machine³, a system that supports archiving of the circumstances of a given website, to ensure the reproducibility and support future investigation. Contents of each document have been accessed and crawled from these Wayback-archived links.

²<https://newser.com>

³<https://web.archive.org>

However, this affected problems regarding the quality of the dataset. As shown in examples of noisy documents in Appendix G, several noisy documents consist of a message from Wayback Machine. Moreover, the failure to crawl the content of the webpage caused other problems. We investigated the case shown in Appendix H and found that it is a result of the crawling of the wrong part of the website. Figure 3 clearly showcases this phenomenon where the content in the red box is crawled instead of the content in the blue box, which is desired. Even though the content in the blue box is different for each article, the system wrongly crawled the shared red box, which resulted in five noisy documents that share the same content and do not contribute to the summary.

From the example above, we revealed the presence of the wrongly crawled documents, that affect the quality of the dataset. We believe such phenomena would be alleviated with the advancement of LLM-based autonomous agents (Wang et al., 2023a), as they could visit the website and only crawl the text relevant to the summary. Even though we leave this as future work, this research direction should be prompted.

C Implementation Details

We utilized PyTorch (Paszke et al., 2019) and Huggingface Transformers (Wolf et al., 2020) to implement and evaluate the model. Specifically, we employed *facebook/bart-large-cnn*⁴ and *google-t5/t5-base*, with 406M and 220M parameters, respectively, for BART and T5. Each model was trained using Adam (Kingma and Ba, 2015) with a learning rate of $2e-5$ over 3 epochs. We used a batch size of 4 and implemented a gradient accumulation step of 4, resulting in a practical batch size of 16. For evaluation, we utilized *bert-base-uncased* and *facebook/bart-large-cnn* for BERTScore and BARTScore, respectively. We reported BERTScore-F1 in Table 2. ROUGE scores were measured using the *rouge-score*⁵ library, with the F1 score of each metric. The training was conducted on a single NVIDIA A100 40GB GPU. We provide the source code and dataset to the public.⁶

For the human evaluation, we recruited three vol-

⁴Note that this model is already fine-tuned with the CNN/DM dataset (Nallapati et al., 2016), a single-document summarization dataset.

⁵<https://pypi.org/project/rouge-score/>

⁶https://github.com/c-juhwan/multi_news_plus

Model Metric	<i>Mistral-7B-Instruct-v0.2</i>	
	BERTScore	BARTScore
No Noisy Example	0.6004	-2.704
One Noisy Example	0.5976	-2.721
Two Noisy Examples	0.5954	-2.738
Model Metric	<i>Llama-2-7b-chat-hf</i>	
	BERTScore	BARTScore
No Noisy Example	0.6038	-2.507
One Noisy Example	0.6022	-2.521
Two Noisy Examples	0.6016	-2.539

Table 4: Performance of LLM-based summarization of Multi-News with different amounts of noisy examples. We only report two model-based metrics as the human-generated reference summary has a different form compared to the LLM-generated summary.

unteers and individually asked them to determine whether the decision of the model was correct or not given the summary, original articles, and rationale of the model. We defined the model made an incorrect decision when at least one human evaluator flagged the output as an incorrect classification.

D Manual Analysis

To perform a more detailed analysis of the accuracy of the proposed method, we randomly selected 60 instances from the validation set, which comprises 153 documents. A confusion matrix was defined to evaluate the classification for each document as follows:

- True Positive (TP): Relevant documents that were correctly classified as relevant.
- False Positive (FP): Documents classified as relevant but are not actually relevant.
- True Negative (TN): Irrelevant documents correctly classified as not relevant.
- False Negative (FN): Relevant documents incorrectly classified as not relevant.

Upon review, we found that 127 documents were classified as TP, 24 as TN, and 2 as FN. The annotation framework identified 26 documents as irrelevant and noisy, which accounts for approximately 17% of the total 153 documents. This aligns closely with the statistics in Table 3 of Appendix A, which indicates that 18.6% of documents in the validation set were classified as noisy.

From these results, the precision is 1.0, as there were no FP documents, while the recall is approximately 0.984. Additionally, we observed that 17 of the 24 TN documents could be classified as noisy system messages, such as “This will appear next to all of your comments; this will not appear anywhere on Newser,” as illustrated in Appendix G. The remaining 7 documents were irrelevant to the summary.

Furthermore, we investigated the two FN cases. In one instance, the summary included a portion related to the misclassified document at the very end. In the other, the misclassified document provided context for the summary but was not directly connected to it. These cases are consistent with the error patterns discussed in Appendix I.

It is important to note that while individual annotators occasionally made incorrect classifications, the majority voting process effectively corrected these errors. This highlights the efficacy of our proposed method in improving data annotation quality and ensuring thorough dataset cleansing.

E Additional Experiment with Large Language Models

This section introduces our additional experiment that investigates the influence of noisy examples for LLMs in a few-shot learning scheme. For this purpose, we used 7B-sized, instruction-tuned Llama2 (Touvron et al., 2023) and Mistral (Jiang et al., 2023). Specifically, we used *meta-llama/Llama-2-7b-chat-hf* and *mistralai/Mistral-7B-Instruct-v0.2* from Transformers (Wolf et al., 2020). In this experiment, we prompted the model to summarize the documents in the test set of Multi-News with two-shot examples selected from the training set of Multi-News. Additionally, we differentiated the number of noisy documents in the examples given as the prompt. Table 4 presents the experimental result. The result demonstrates that the inclusion of the noise in the example degrades the quality of the summary generated by the LLM. This suggests the significance of the exclusion and filtering of the noise for LLMs, which underscores the necessity of dataset cleansing presented in this paper.

F Analysis of Multi-News

Following the previous study of TeSum (Urlana et al., 2022), we apply filtering strategies and analyze the characteristics of Multi-News with these strategies. Table 5 exhibits the result of the analy-

	Multi-News
Dataset Size	56,216
Source Article Size	156,289
Avg Words in Source	433.62
Avg Sentences in Source	23.42
Avg Words in Summary	228.69
Avg Sentences in Summary	11.52
Empty Summary	0
Duplicated Summary	0
Prefixes Summary	0
Empty Source	570
Duplicated Source	544
Source < 4 Sentences	45
Source < 40 Words	7
Summary < 10 Words	0
Compression < 50%	31,994
Compression > 80%	390
Abtractivity < 10	496
Abtractivity > 80	126
Avg Abtractivity	41.42
Avg Compression	46.19%

Table 5: The result of analysis of Multi-News dataset with rule-based filtering methods (Urlana et al., 2022). We concatenated every source document to measure their average word and sentence length.

sis. First, we found that 0.7% of total source documents can be considered noisy documents as it is empty or duplicated from other source documents within the same set. Second, we found previous rule-based filtering methods are not very effective standards for the Multi-News dataset. For instance, there were no sets that had empty summaries, summaries that were duplicated with other summaries, or summaries that repeated the first few sentences of source documents. The only exception is *Compression < 50%*, which identified more than half of the dataset. However, it should be noted that Multi-News is a multi-document summarization dataset, which is different from datasets for previous studies. For instance, average compression is significantly lower than other single-document summarization datasets reported in the previous study (Urlana et al., 2022), as multiple source documents in Multi-News involve more information compared to the source document of single-document summarization datasets. In conclusion, this analysis demonstrates that previous filtering strategies are less practical for multi-document summarization datasets such as Multi-News and enlightens the necessity of novel approaches for these datasets.

G Examples of Noisy Documents

This section demonstrates several examples of noisy documents observed in the Multi-News dataset not related to the summary. Please refer to the released dataset file for details.

- Tweet with a location you can add location information to your tweets, such as your city or precise location, from the web and via third-party applications. You always have the option to delete your tweet location history. [Learn more](#)
 - Focused crawls are collections of frequently-updated webcrawl data from narrow (as opposed to broad or wide) web crawls, often focused on a single domain or subdomain.
 - Dow jones reprints: this copy is for your personal, non-commercial use only. To order presentation-ready copies for distribution to your colleagues, clients or customers, use the order reprints tool at the bottom of any article or visit www.djreprints.com
 - This crawl of online resources of the 115th us congress was performed on behalf of the united states national archives & records
 - The seed for this crawl was a list of every host in the wayback machine this crawl was run at a level 1 (urls including their embeds, plus the urls of all outbound links including their embeds) the warc files associated with this crawl are not currently available to the general public.
 - These crawls are part of an effort to archive pages as they are created and archive the pages that they refer to. That way, as the pages that are referenced are changed or taken from the web, a link to the version that was live when the page was written will be preserved.then the internet archive hopes that references to these archived pages will be put in place of a link that would be otherwise be broken, or
 - Please enable cookies on your web browser in order to continue. The new european data protection law requires us to inform you of the following before you use our website: we use cookies and other technologies to customize your experience, perform analytics and deliver personalized advertising on our sites, apps and newsletters and across the internet based on your interests. By clicking “i agree” below, you consent to the use by us and our third-party partners of cookies and data gathered from your use of our platforms. See our [privacy policy](#) and [third party partners](#) to learn more about the use of data and your rights. You also agree to our [terms of service](#).
 - Thank you for reading. Please purchase a subscription to continue reading. A subscription is required to continue reading. Thank you for reading 5 free articles. You can come back at the end of your 30-day period for another 5 free articles, or you can purchase a subscription and continue to enjoy valuable local news and information. If you are a current 7-day subscriber you are granted an all-access pass to the website and digital newspaper replica. Please click [sign up to subscribe](#), or [login](#) if you are already a member. Thank you for reading 5 free articles. You can come back at the end of your 30-day period for another 5 free articles, or you can purchase a subscription and continue to enjoy valuable local news and information. If you are a current 7-day subscriber you are granted an all-access pass to the website and digital newspaper replica. Please click [below to get started](#).
 - Add a location to your tweets when you tweet with a location, twitter stores that location. You can switch location on/off before each tweet and always have the option to delete your location history. [Learn more](#)
-

H Extreme Cases of Noisy Documents

In addition to examples of noisy documents, we discovered the following extreme case of noisy data in the Multi-News dataset. In this example, five documents have the same content but offer no information on the summary. Thus, it cannot generate a reasonable summary based on the given documents. We witnessed 379 similar cases during the dataset cleansing process, as reported in Figure 2. While they were excluded from training and testing, we included them in the dataset file for future investigation.

Summary

Note to tweeting politicians: watch what you post, because politwoops will remember it forever. The transparency-minded website is safeguarding politicians' deleted tweets, enabling the rest of us to giggle or ponder over them at our leisure, the atlantic reports. The site's current 6-month stash includes a few doozey deletions, including john mccain mocking vladimir putin's tears and rep. Jeff miller posting a link to a poll that asked, " was obama born in the united states? " a few deletions are more odd than obvious, begging us to ask what politicians were thinking. Why, for example, did rep. Tom graves remove a tweet about going out one night with his wife? or rep. Kathy hochul delete one about her visit to a cancer institute? perhaps rep. Stephen fincher's tweet comparing the bachelor to the hunger games is a more obvious case, but the online avenues of a politician's mind can be dimly lit indeed.

Document 1

An archive of the public statements deleted by u.s. Politicians. Explore the tweets they would prefer you couldn't see. If you aren't an elected official or running for office and feel your account is being tracked by mistake then please contact us.

Document 2

An archive of the public statements deleted by u.s. Politicians. Explore the tweets they would prefer you couldn't see. If you aren't an elected official or running for office and feel your account is being tracked by mistake then please contact us.

Document 3

An archive of the public statements deleted by u.s. Politicians. Explore the tweets they would prefer you couldn't see. If you aren't an elected official or running for office and feel your account is being tracked by mistake then please contact us.

Document 4

An archive of the public statements deleted by u.s. Politicians. Explore the tweets they would prefer you couldn't see. If you aren't an elected official or running for office and feel your account is being tracked by mistake then please contact us.

Document 5

An archive of the public statements deleted by u.s. Politicians. Explore the tweets they would prefer you couldn't see. If you aren't an elected official or running for office and feel your account is being tracked by mistake then please contact us.

I Error Analysis

Following the form of the previous study (Choi et al., 2024), we provide an error analysis to provide a more balanced view of the behavior and limitations of our proposed method. In the first example, we can observe that while Document 1 can be regarded as irrelevant to the summary except that there is a mention of fusion tv, Document 2 contains information about Mike Tyson and his new TV documentary series. However, the model predicted both documents are irrelevant to the summary, primarily because the model concentrated on the mention of the “world team tennis exhibition” from Document 2. From this insight, we hypothesize GPT-3.5 suffers from a mixture of irrelevant and relevant information in one document.

Summary

Over his career, former heavyweight champion mike tyson recorded 50 wins and six losses. But he recently notched another big loss in latin america — this time as a coach of a bird, reports the ap. Tyson traveled to suriname as part of the new fusion tv documentary series outpost, and was soundly beaten when he entered a bird in a songbird contest, a cherished local tradition. Cameras captured iron mike as he learned about the contest, located a bird to enter — he dubbed the tiny guy " little mike " — but then suffered a tko when a competing champion cheeped and peeped more than his bird did in the same 15-minute period. " little mike let us down, man. I was in his corner, though, " said tyson. " it was just amazing meeting the people, meeting the culture — i had a great time. " the series, kicking off on sunday with tyson’s episode, mixes travel adventure, history, and journalism to shine a light on global stories. The first season focuses on latin america and includes as hosts the late show with stephen colbert bandleader jon batiste, brain games star jason silva, and transgender model carmen carrera. Spanish versions air on unimas. Tyson was lured onto the show by the chance to visit a country he’d never heard of and his love of birds. The former boxer has loved pigeons and kept them since he was a kid in brooklyn. (sunday’s show recorded the moment tyson lovingly released his bird in suriname.) " my wife always says the reason i keep my pigeons is they connect me to my childhood, " tyson said. " once it’s in your blood, it never leaves. It’s just who you are. "

Document 1

Starting in 1996, alexa internet has been donating their crawl data to the internet archive. Flowing in every day, these data are added to the wayback machine after an embargo period. *[Abbreviated duplicated text]* Outpost shows you the world like you’ve never seen it. The series lives at the intersection of investigative journalism and adventure travel, bringing you a local perspective on faraway places and inviting you to explore. The series premieres march 26 @ 8 and 11 pm on fusion tv. In the first episode, transgender model carmen carrera travels to brazil, a place where rates of violence against lgbt people are some of the highest in the world, to find out what’s happening, what life is like for young transgendered people in brazil, and what the future might hold. Gabriel leigh takes us to el alto, bolivia, where some of the craziest architecture on earth is taking shape as part of a surge in indigenous purchasing power.

Document 2

[Abbreviated duplicated text] file - in this monday, oct. 10, 2016, file photo, mike tyson attends a world team tennis exhibition to benefit the elton john aids foundation in las vegas. Tyson traveled to suriname as part of the new fusion tv documentary series "outpost " and was soundly beaten when he entered a bird in a songbird... (associated press) *[Abbreviated duplicated text]* new york (ap) — over his career, former heavyweight champion mike tyson recorded 50 wins and six losses. But he recently notched another big loss in latin america — this time as a coach of a bird. Tyson traveled to suriname as part of the new fusion tv documentary series " outpost " and was soundly beaten when he

This second example also showcases the characteristics of GPT-3.5 model we used. In this example, it is obvious that Document 2 is less relevant to the summary, which is mainly about the relationship between Gwyneth Paltrow and Chris Martin. However, while it is not the main content of the document as well as Document 2, Document 1 contains a sentence that mentions the relationship between the two (“her amicable split from husband chris martin of coldplay”). Nonetheless, the model predicted Document 1 is also irrelevant to the summary, implying the model is stringent to the *partial* contribution of the document to the summary. However, it is important to note that we categorized these instances as errors based on rigorous human evaluation, and such errors constituted fewer than 10% of the total classifications, where a single flag by multiple human evaluators was sufficient to deem it an error. We are planning to manually revise these errors in the released version of MULTI-NEWS⁺.

Summary

Gwyneth paltrow continues to paint the sunniest of pictures of her post-conscious-uncoupling life with chris martin, but the description she gives glamour in a new interview may be the most interesting one so far. " we're still very much a family, even though we don't have a romantic relationship. He's like my brother, " she says, explaining that the two of them and their two kids still spend quite a bit of time together, even staying in one another's houses and spending holidays together (not to mention collaborating on songs together). " the ideal is to stay married. But if you can't stay married, wouldn't the ideal be that you could still be a family and you could put aside your own stuff long enough to explore — what is this new family and who am i in it? " paltrow muses. " and chris is a great ex-husband ' cause he's a very, very willing partner in how to do that. " she adds that, though she's " very independent, " she does see the value in having a husband, and though she's not quite divorced yet, she could perhaps see herself getting married again someday. (click to see what she has to say about her other famous exes.)

Document 1

Gwyneth paltrow is in a state of deep focus. The new goop office is under construction — "it's like a dust bowl, " she says with a laugh — so today she's helming her company from the kitchen island of her los angeles home. Fitting, considering it was at her kitchen table (then in london) that paltrow, 43, started goop as a newsletter to friends nearly eight years ago. Since then, she has built goop into a global brand: it has produced sought-after collaborations with valentino and stella mccartney; opened pop-up shops; and brought terms like conscious uncoupling and vaginal steaming to the masses (the first a description of her amicable split from husband chris martin of coldplay; the second, a way to cleanse one's uterus — don't try it at home). Her presence has also unwittingly exposed a dirty little secret: as fans, we provide actresses with wealth and fame, only to scoff when they actually lead that rich and famous lifestyle publicly. We want these stars to be "just like us. " but paltrow's life simply isn't. She won't pretend that she shops at the dollar store for beauty products or feeds her kids, apple, 11, and moses, 9, a steady diet of fast food;

Document 2

Gwyneth paltrow was definitely in the mood to share during her appearance on howard stern's siriusxm radio show on wednesday.... Especially when it came to her a-list exes. In the hour-long chat, stern of course wanted to know all about paltrow's ex-fiance brad pitt, who the shakespeare in love star was engaged to when she was 24 years old. The beautiful blondes eventually called it quits in 1997 after three years together. Getty images " i didn't think about it at the time, but i ' m sure it did help with my career, " the now 42-year-old actress admits about the start of all the paparazzi attention when the two got together on the set of seven. " i definitely fell in love with him. He was so gorgeous, and sweet – i mean, he was brad pitt, you know? " video: a history of gwyneth's former flames her parents, the late bruce paltrow and actress blythe danner, also 100 percent approved. " my father was so devastated when we broke up, " she dishes. " my father loved him like he was his son. " in hindsight, she blames the demise of their relationship on her youth. " i was such a kid, i was 22 when we met, " she explains.

J Prompt

This section describes the prompt used for dataset cleansing. We truncated several parts to save space and only displayed a 1-shot CoT example. Please refer to the released source code for the exact prompt.

System

You are a helpful assistant. Your job is to identify the document that is not relevant to the summary. The relevance is determined by the mention of the key points in the summary. For example, the alert message from the twitter, the location information, and the unrelated information are considered irrelevant. Please response with the resaoning process and the result as document number or 'None'. Do not end the response with a period. If multiple documents are irrelevant to the summary, please resonse in the form of 'Document 1|Document 3'

User

Given the summary and documents, please identify the document that is not relevant to the summary such as advertisements. If every document is relevant to the summary, output 'None'. Input document will be given in next dialogue.

User

[Summary]

The content of summary

[Document 1]

The content of document 1

[Document 2]

The content of document 2

[Document 3]

The content of document 3

[Document 4]

The content of document 4

Assistant

The summary discusses Tucker Carlson's controversial tweet regarding Sarah Palin and the subsequent backlash he faced from conservative women. Document 1 is related to Tucker Carlson's controversial tweet. Document 2 is unrelated as it contains information about adding location information to tweets and does not contribute to the discussion about Tucker Carlson's tweet and the reaction it received. Document 3 is a tweet message that calls Tucker Carlson a misogynist, related to the summary. Document 4 is related to the summary as it deals with Tucker Carlson's movement. Therefore, the irrelevant document is: Document 2
