

# GUMSum: Multi-Genre Data and Evaluation for English Abstractive Summarization

Yang Janet Liu and Amir Zeldes

Department of Linguistics

Georgetown University

{yl1879, amir.zeldes}@georgetown.edu

## Abstract

Automatic summarization with pre-trained language models has led to impressively fluent results, but is prone to ‘hallucinations’, low performance on non-news genres, and outputs which are not exactly summaries. Targeting ACL 2023’s ‘Reality Check’ theme, we present GUMSum, a small but carefully crafted dataset of English summaries in 12 written and spoken genres for evaluation of abstractive summarization. Summaries are highly constrained, focusing on substitutive potential, factuality, and faithfulness. We present guidelines and evaluate human agreement as well as subjective judgments on recent system outputs, comparing general-domain untuned approaches, a fine-tuned one, and a prompt-based approach, to human performance. Results show that while GPT3 achieves impressive scores, it still underperforms humans, with varying quality across genres. Human judgments reveal different types of errors in supervised, prompted, and human-generated summaries, shedding light on the challenges of producing a good summary.

## 1 Introduction

Recent advances in supervised summarization models as well as prompt-based approaches using large pre-trained language models have led to substantial improvements in summary fluency, with prompt-based outputs now surpassing supervised approaches in human evaluation (Goyal et al., 2022). At the same time, researchers in the field repeatedly note that the most commonly used datasets, such as CNN/DailyMail (CNN/DM, Hermann et al. 2015) and Extreme Summarization (XSum, Narayan et al. 2018), which are large-scale ‘found’ datasets not designed to facilitate high quality summarization, are problematic, and in many cases contain texts which are not summaries, are incomplete or unfaithful to the texts they relate to, add information not present in texts, or any combination of the above (Reiter, 2022; Liu et al.,

2022a). Existing datasets are also limited to mainly newswire text (cf. Zopf et al. 2016), which is a fraction of extant genres in general and on the Web.

The main contributions of this paper are in providing and evaluating a very genre-diverse dataset and guidelines for carefully crafted, rather than ‘found’ summaries, which follow the same design across text types. Building on the UD English GUM treebank (Zeldes, 2017), which contains 213 spoken and written texts balanced across 12 different genres, our summaries target three goals: 1) to be **substitutive** (i.e. informative, functioning as a substitute for reading a text, cf. Edmundson 1969; Nenkova and McKeown 2011) rather than indicative (e.g. ‘clickbait’ designed to attract readership); 2) to be **faithful** to the text, adhering to original formulations wherever possible; 3) to be **hallucination-free**, meaning summaries make a strong effort not to add any information (even if it is likely to be true), mentioning only entities and events actually contained in the text, thereby preventing typical errors associated with datasets such as XSum (Narayan et al., 2018). Instructions on obtaining the dataset and responses from the human evaluation study as well as evaluation code can be found at <https://github.com/janetlaueung/GUMSum4EVAL>.<sup>1</sup>

## 2 Related Work

The problem of mitigating factuality and faithfulness issues in Natural Language Generation (NLG) has recently received considerable attention, with studies proposing auxiliary tasks using the Multi-Task Learning approach to constrain models, such as overlapping entities (Nan et al., 2021), encoding of predicate triples from source documents (Zhu et al., 2021) or encouraging systems to incorporate or copy entities from source documents (Xiao and

<sup>1</sup>Data is also available from the corpus website at <https://gucorpling.org/gum/> and guidelines at <https://wiki.gucorpling.org/en/gum/summarization>.

Carenini, 2022; Maddela et al., 2022). In addition, Tang et al. (2022) present a thorough investigation of factual errors in summarization and propose a taxonomy of error types with a focus on entity and predication errors, while Thomson et al. (2023) examine types of accuracy errors made by neural systems and contrast them with human errors.

These papers share concerns about the nature of widely used datasets for English, such as XSum and CNN/DM, but are limited by the lack of evaluation data specifically targeting genre-diverse texts with high-quality summaries: ones which ideally maximize faithfulness, rule out hallucinations, and follow consistent guidelines for what constitutes a summary. Although there are some non-news single-document summarization datasets covering Reddit (Kim et al., 2019) and Podcast data (Reza-pour et al., 2022), text types are still quite limited and data is often not publicly available (Tang et al., 2022). This motivates our work to create open-access, multi-genre data with consistent guidelines across text types.

### 3 Dataset

**Contents** GUMSum covers the 213 documents (amounting to  $\sim 200K$  tokens) from the 12-genre UD English GUM corpus (Zeldes 2017; specifically GUM V9), which provides gold syntax trees, entity types, coreference resolution, and discourse parses for the data. For this paper, we added summaries to each document in the corpus, by the authors and students in a Computational Linguistics course as part of a class-based project,<sup>2</sup> guided by general and genre-specific instructions. Although the range of  $\sim 20$  human summarizers is broad as a result, we defined guidelines to constrain summaries and ensure they are maximally ‘reality-checked’, i.e. faithful and factual, as evaluated below. Documents vary in length, ranging between 167 and 1,878 tokens (mean=957, sd=249.6), and cover the genres in Table 1. Because of the classroom context in which summaries are collected and the natural variation in student styles and adherence to guidelines, all summaries are thoroughly checked by a teaching assistant and the course instructor. For the 24 documents in the UD treebank’s official test set of GUM V9, we provide two summaries to support inter-annotator agreement and multiple-reference evaluation.

<sup>2</sup>Consent to release data was given by all students.

Genres	Source	Docs	Toks	ØSum.Len (sd)
Interviews	Wikinews	19	18,190	49 (6.3)
News stories	Wikinews	23	16,145	51 (9.0)
Travel guides	Wikivoyage	18	16,514	59 (8.9)
How-to guides	WikiHow	19	17,081	67 (6.5)
Academic	various	18	17,169	35 (11.2)
Biographies	Wikipedia	20	18,213	44 (9.8)
Fiction	various	19	17,510	47 (10.3)
Web forums	Reddit	18	16,364	50 (8.7)
Conversations	SBC	14	16,416	41 (13.7)
Speeches	various	15	16,720	46 (9.2)
Vlogs	YouTube	15	16,864	50 (11.8)
Textbooks	OpenStax	15	16,693	51 (8.9)
total / average		213	203,879	50 (12.2)

Table 1: Overview and Statistics of GUMSum.

**Guidelines** Previous literature has characterized ‘good’ summaries primarily as ones that are concise, accurate, fluent, and coherent (Fabbri et al., 2021). What these qualities mean varies depending on the summary’s objective: whether it is domain-specific or general, indicative (enticing readers to read the text) or informative (aiming to substitute reading it, Nenkova and McKeown 2011) etc. GUMSum’s summaries explicitly target a **domain-general, substitutive, maximally concise** format, which is therefore constrained to:

- [1] have at most one sentence / 380 characters<sup>3</sup>
- [2] have the goal of replacing reading the text
- [3] give participants/time/place/manner of events
- [4] form a sentence rather than a fragment
- [5] omit distracting information
- [6] avoid entities or information not present in the text, even if we are fairly sure it is true
- [7] reject synonyms for words in the text

For instance, the summary in (1) for a story involving ‘robbers plundering a vault’ follows guidelines by providing a declarative-sentence (criteria [1], [4]), synopsis of events, participants (exactly *five robbers*), time (a date) and place (*Poughkeepsie*) ([3]), as well as additional details (exact name of the bank, mode of escape). (2) is underspecified (we do not know when or where the event occurred, criterion [3]). (3) paraphrases the robbers’ escape by introducing an entity *not* in the original text (uncaught by *police*, violating [6]), and substitutes ‘robbed’ for ‘plundered’, a near synonym but a deviation from the original text’s style ([7]).

<sup>3</sup>We follow XSum in targeting 1-sentence summaries, and we aimed for a maximum of 5 lines in a PEP8-compliant IDE, but in practice no summary exceeded 380 characters.

- (1) 📍 On March 23, 1999, five bank robbers plundered the vault of First National Bank in Poughkeepsie, NY and escaped in a bus they had stolen.
- (2) 📍 Bank robbers plundered a vault and escaped.
- (3) 📍 Bank robbers who robbed a bank in Poughkeepsie were never caught by police.

Although these examples illustrate newswire language, GUMSum covers very different spoken and written text types as well:

- (4) *Some people debate whether the original 3 hour cut of Snyder’s movie about Batman and Superman should have been released instead of the shorter version, which prioritized getting to the action faster in order to appeal to a general audience. (Reddit)*
- (5) *Ash tells about her day, which includes a yoga class, marketing brand management class, doing some work while having coffee at Saxby’s, and finally cooking pasta with peppers for dinner together with her boyfriend Harry. (YouTube CC-BY vlog)*

The summary in (4) follows the guidelines by not mentioning that the discussion is on Reddit ([6], the interlocutors are simply ‘people’), since Reddit is not mentioned. Similarly, while Zack Snyder’s film *Batman v Superman: Dawn of Justice* is most likely being discussed, it is not named explicitly, leading to the formulation ‘Snyder’s movie about Batman and Superman’. In (5), the summary focuses on conveying events which happen over the course of a vlog, but again, the unmentioned word ‘vlog’ is avoided, while specific details about the participants and circumstances (people, but also the type of class) are prioritized. Summaries are thus highly constrained to remain faithful and avoid even minor potential hallucinations, such as completing the title of a film. For more on genre-specific guidelines and examples, see Appendix A.

## 4 Evaluation

**Automatic Evaluation** To evaluate how well current neural approaches produce ‘reality-checked’ summaries approaching the ones in GUMSum, we obtain system outputs from two recent supervised systems, BRIO (Liu et al., 2022b) and SimCLS

(Liu and Liu, 2021), as well as prompt-based outputs using a GPT3 model (Brown et al., 2020), GPT3-text-davinci-002 (GPT3-DV2), with the prompt ‘Summarize the text above in one sentence.’. We chose system models trained on the XSum dataset, since it has one-sentence summaries more in line with the GUMSum data. However, because systems have never seen data in many of GUMSum’s genres, we also add an additional experiment in which we fine-tune the higher-scoring supervised system, i.e. BRIO’s trained-model on XSum for *generation*, by continuing training it on the 165 documents in the UD treebank’s train set of the underlying GUM V9 corpus (BRIO-FT in Table 2; details/splits and system output selection can be found in Appendix B). Scores are compared to a second human-written summary obtained from a human evaluation study, using the same guidelines.

	R-1	R-2	R-L	BS	MS	METEOR	BLEU	BLEURT
SimCLS	23.1	6.2	17.2	86.0	12.1	13.4	2.1	31.9
BRIO	27.8	10.2	21.2	87.2	15.9	18.0	3.7	36.3
GPT3-DV2	31.1	<b>12.1</b>	25.1	88.5	21.1	20.8	3.8	42.2
BRIO-FT*	<b>37.3</b>	12.0	<b>27.1</b>	<b>88.7</b>	<b>27.4</b>	<b>27.6</b>	<b>6.1</b>	<b>44.3</b>
Human 2	<b>38.9</b>	<b>12.7</b>	<b>28.4</b>	<b>88.8</b>	<b>28.5</b>	<b>33.0</b>	<b>7.5</b>	<b>50.2</b>

Table 2: Automatic Evaluation Metrics of System Outputs and Human Agreement (\* = 3 run average).

Table 2 shows that while systems have impressive scores for ROUGE (Lin, 2004), BERTScore (BS, Zhang et al. 2020), MoverScore (MS, Zhao et al. 2019), METEOR (Banerjee and Lavie, 2005), BLEURT (Sellam et al., 2020), and BLEU (Papineni et al., 2002), they still lag behind the human summaries across the board. Reproducing findings by Goyal et al. (2022), GPT3-DV2 outperforms supervised systems trained on XSum, though our data contains much more diverse genres than those in that paper. However, fine-tuning on even a small amount of GUMSum data (165 documents) in this paper already outperforms GPT3-DV2. This strongly suggests that a major problem with supervised systems in domain-general settings is simply the training data itself. Qualitative inspection of outputs suggests fine-tuning was particularly helpful for summarizing conversations, Reddit, and how-to guides, on which all systems struggled. For humans, genre differences were much less pronounced, with lowest scores surprisingly for news. Figure 2 gives a detailed breakdown of BLEURT scores (Sellam et al., 2020) by genre for each scenario. Human scores lead in every genre except academic, news, and interview, and generally vary less by genre than systems. BRIO-FT is improved




Figure 1: Barplots of Human Evaluations on Preferences, Substitutive Potential, Hallucination, and Faithfulness.

especially on genres that diverge from XSum, such as conversations, travel guides from Wikivoyage, and how-to guides from Wikihow.




Figure 2: BLEURT Scores by Genre for Each Scenario (whow=how-to guides, voyage=travel guides).

Finally, the human scores provide some numbers for ceiling performance as reflected by automatic metrics. Comparing human numbers to the best-system numbers suggests that there is a substantial gap for systems which have never been trained on in-domain data. However, for the the fine-tuning (FT) scenario, we notice that ROUGE scores are neck-and-neck with the second human summary, likely because the system is trained with an objective averaging R1, R2, and R-L, on which it excels. By contrast, metrics more focused on verbatim overlap, such as BLEU, or semantic similarity, such as BLEURT, retain a more substantial gap, with FT results on BLEURT being close to GPT3-DV2 and still nearly 6 points below human performance.

It is an established finding however that metrics do not tell the whole story (Novikova et al., 2017; Reiter, 2018; Marasović, 2018; Gehrmann et al., 2022). In fact, we regularly observe hallucinations, especially in XSum-trained systems, such as prefixing generic leads (e.g. ‘In our series of letters from British journalists ...’, when there are no journalists

involved) or inserting entities and events not mentioned in the text. We thus conduct a human evaluation of system outputs below, focusing on substitutivity, hallucinations, and faithfulness, and more importantly, apply the same evaluation criteria to the human-written summaries for a more targeted evaluation, as advocated by Liu et al. (2022a).

**Human Evaluation** We asked 12 Linguistics students to evaluate the full texts and the summaries of the 24 documents in the test set of the source GUM V9 corpus and to produce an additional summary for their assigned texts (see detailed instructions in Appendix C).<sup>4</sup> Figure 1 shows humans overwhelmingly preferred the human-written summary (1(a), 83%, with exceptions citing gold summaries as less pleasant to read), and also found it best at substituting reading the text (1(b), 79%). Pretrained supervised systems were judged to be highly non-substitutive (88% for SimCLS, 79% for BRIO), while 71% of GPT3-DV2 outputs were judged moderately so.

While all systems exhibited some hallucinations and unfaithfulness, GPT3-DV2 performed best, in part because its outputs tended to be short (mean 138 characters vs. human 272 characters) and general, giving fewer chances for issues. At the same time, hallucination types varied substantially. Human violations in both categories were rare and subtle, resulting from evaluators adhering to guidelines very literally: for example, one evaluator proposed that a human summary’s use of the pronoun ‘she’ in reference to a vlogger whose pronouns had not been stated is a form of hallucination, while another pointed out that a mention of ‘Washington’ in a news article was a faithfulness issue, since with-

<sup>4</sup>The hourly pay is \$20.29/hour based on the pay rate of the 2022 / 2023 academic year for graduate students at Georgetown University. It took about 1.5 hours in total for each annotator to complete all the tasks for the two documents.



out specifying ‘DC’, the place is ambiguous. Hallucinations from GPT3-DV2 were more pronounced (e.g. designating a speaker mentioning retirement as an attendee of a seminar about retirement, which was not mentioned), while XSum-trained systems had more extreme cases, such as incorrectly attributing a speech about New Zealand to its former Prime Minister John Key (BRIO), claiming a fictional short story is a BBC documentary (SimCLS), or adding to a textbook excerpt on the Civil War by calling it the longest, most expensive conflict in US history (BRIO and SimCLS). Below we provide a comparison of outputs for two documents and a qualitative analysis.

We also asked evaluators whether they could tell if summaries were NLG outputs, and learned that while ‘NLG’ guesses were correct, and most human summaries were also recognized, humans could not tell for certain in 56% of the outputs they evaluated (incl. 8% of human-written cases).

**Qualitative Analysis** Figure 3 shows two human-written and several system-generated summaries, for a conversation in (a) and for a news text in (b).<sup>5</sup> Note the typical hallucinated lead about journalists in the first BRIO output, which disappears after fine-tuning, and a similar insertion about a Nigerian writer in the output for SimCLS. GPT3-DV2 does not show similar issues, but misses important context information, e.g. the purpose of the conversation revolving around whether speakers should go to a specific dance class, and why or why not.

The news output is substantially better for all systems. BRIO disagrees with SimCLS and GPT3 on the number of ‘remaining’ space shuttles: three remained to be retired, but there were four total in the article, including the already retired shuttle Discovery. All pre-trained system outputs are substantially less detailed than the human summaries, which add information about time and place of the announcement, or the list of space shuttles. Human 2 commits a similar hallucination error to BRIO in identifying the already retired Discovery as being retired at document creation time. However, both human summaries agree that a prominent part of the story involves the disappointment or criticism from sites that were not selected to house retired shuttles, a topic to which most of the latter half of the original story is dedicated. The fine-tuned model successfully adds more details in line with

<sup>5</sup>The PDFs of the full-text of these two documents are provided in the repository of the paper for reference.

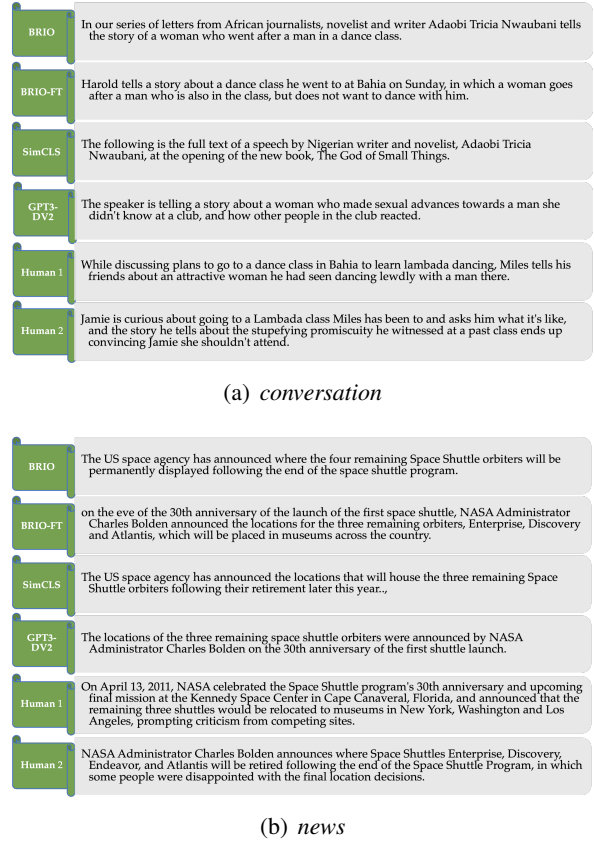


Figure 3: Sample Summary Outputs of a *conversation* Text and a *news* Text from Each Evaluation Scenario.

the human summaries, but also fails to capture the site controversy in the second half of the document.

## 5 Conclusion

The dataset and guidelines introduced in this paper make a step towards consistent and constrained multi-genre evaluation of factual summarization. Our results show that domain-general summarization is still hampered by serious reliability and factuality problems, which may only become apparent when confronted with a dataset with strict ‘reality check’ constraints and diverse text types. Even small amounts of such data can be used to fine-tune pre-trained systems, with measurable improvements for system outputs.

The human evaluation study also revealed that pre-trained systems are bad at delivering substitutive summaries, perhaps because, as pointed out in Reiter (2022), “summarisation datasets should contain summaries,” but often they do not. Meanwhile, human identification of possibly more minor hallucinations in human-written summaries also suggests that more work is needed in delimiting what a ‘reality check’ for summaries should include.

## Limitations

GUMSum is designed to constrain summaries to one sentence for all 12 genres, which raises the question of whether one-sentence summaries are useful for all possible genres or long-document summarization. This is a complex topic that needs in-depth investigation. For GUMSum, as mentioned in Section 3, document length is limited to 167–1,878 tokens. Moreover, in analyzing human evaluators’ responses to two open-ended questions ([1] and [2] in Appendix C), we noticed that virtually all evaluators mentioned that limiting the summary to one-sentence is very difficult and that some genres were easier than others. For example, one evaluator who was given a vlog and a travel guide commented that,

*“The travel guide was much more difficult than the vlog, likely because it was longer and denser. [...] the travel guide packed a lot more information into its pages and within each sentence.”*

This indicates that genre differences at the summary-level is not trivial due to the style of the original text.

Additionally, this paper examined a specific subset of pre-trained systems and one version of GPT3’s pretrained language model (i.e. GPT3-text-davinci-002), producing findings which may not generalize to other settings. The dataset used for the evaluation is also substantially smaller than those used in most work on summarization, due to the fact that it was carefully crafted based on both general and genre-specific guidelines to be substitutive and to avoid hallucinations and faithfulness issues, rather than originating in a found dataset, in order to conduct a more targeted evaluation, as recommended by Liu et al. (2022a). While it is inevitable that more data would lead to different results, we do not believe that system rankings or overall findings would be substantially different, so long as the guidelines and genres examined here remain stable.

Finally, we must raise a further limitation involving text type and language: our study encompasses 12 specific written and spoken genres available in the UD English GUM corpus, but does not capture findings for other genres, or indeed other languages, which deserve more attention in future studies.

## Ethics Statement

The data produced in this paper is made openly available in accordance with the original licenses of the underlying resources and academic fair use. We are keenly aware that NLP, and particularly NLG technology can be misused adversely, for example to generate fake news, we believe the risks posed by models which are not ‘reality-checked’ outweigh those associated with improving models to prevent factuality and generalization issues across domains. The latter issue is particularly relevant, since technologies limited to particular domains and styles will primarily benefit actors in sectors engaged with that data (e.g. news, for example, financial reporting), while underserving the public in other areas (e.g. computer-mediated communication). We therefore concur with this year’s ACL theme that work towards ‘reality checking’ our outputs is a net positive.

## Acknowledgements

The human evaluation study was funded by a GSAS-GradGov Research Project Award (GRPA) towards graduate students’ research and professional development endeavors at Georgetown University. We thank the following participants for their valuable participation and insightful feedback in our human evaluation study (alphabetically ordered by last names): Kris Cook, Jessica Cusi, Helen Dominic, Luke Gessler, Caroline Gish, Lauren Levine, Cynthia Li, Kristina Lignell, Davide Locatelli, Emma Manning, and others who prefer to stay anonymous. We thank Nathan Schneider and the anonymous reviewers for their feedback.

## References

- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- 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 (NIPS)*, volume 33 of *NIPS'20*, pages 1877–1901, Red Hook, NY, USA. Curran Associates, Inc.
- H. P. Edmundson. 1969. [New Methods in Automatic Extracting](#). *J. ACM*, 16(2):264–285.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. [SummEval: Re-evaluating summarization evaluation](#). *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Sebastian Gehrmann, Elizabeth Clark, and Thibault Selam. 2022. [Repairing the cracked foundation: A survey of obstacles in evaluation practices for generated text](#). *arXiv*.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. [News Summarization and Evaluation in the Era of GPT-3](#). *arXiv*.
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, page 1693–1701, Cambridge, MA, USA. MIT Press.
- Byeongchang Kim, Hyunwoo Kim, and Gunhee Kim. 2019. [Abstractive summarization of Reddit posts with multi-level memory networks](#). 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 2519–2531, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- Yixin Liu, Alexander R. Fabbri, Pengfei Liu, Yilun Zhao, Linyong Nan, Ruilin Han, Simeng Han, Shafiq Joty, Chien-Sheng Wu, Caiming Xiong, and Dragomir Radev. 2022a. [Revisiting the gold standard: Grounding summarization evaluation with robust human evaluation](#). *arXiv*.
- Yixin Liu and Pengfei Liu. 2021. [SimCLS: A simple framework for contrastive learning of abstractive summarization](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 1065–1072, Online. Association for Computational Linguistics.
- Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. 2022b. [BRIO: Bringing order to abstractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2890–2903, Dublin, Ireland. Association for Computational Linguistics.
- Mounica Maddela, Mayank Kulkarni, and Daniel Preotiuc-Pietro. 2022. [EntSUM: A data set for entity-centric extractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3355–3366, Dublin, Ireland. Association for Computational Linguistics.
- Ana Marasović. 2018. [NLP's generalization problem, and how researchers are tackling it](#). *The Gradient*.
- Feng Nan, Ramesh Nallapati, Zhiguo Wang, Cicero Nogueira dos Santos, Henghui Zhu, Dejiao Zhang, Kathleen McKeown, and Bing Xiang. 2021. [Entity-level factual consistency of abstractive text summarization](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2727–2733, Online. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Ani Nenkova and Kathleen R. McKeown. 2011. [Automatic Summarization](#). *Foundations and Trends in Information Retrieval*, 5(2-3):103–233.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. [Why we need new evaluation metrics for NLG](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2241–2252, Copenhagen, Denmark. 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 on Association for Computational Linguistics*, ACL '02, page 311–318, USA. Association for Computational Linguistics.
- Ehud Reiter. 2018. [A Structured Review of the Validity of BLEU](#). *Computational Linguistics*, 44(3):393–401.
- Ehud Reiter. 2022. [Summarisation datasets should contain summaries!](#)
- Rezvaneh Rezapour, Sravana Reddy, Rosie Jones, and Ian Soboroff. 2022. [What makes a good podcast summary?](#) In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '22, page 2039–2046, New York, NY, USA. Association for Computing Machinery.



Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.

Liyan Tang, Tanya Goyal, Alexander R. Fabbri, Philippe Laban, Jiacheng Xu, Semih Yahvuz, Wojciech Kryściński, Justin F. Rousseau, and Greg Durrett. 2022. [Understanding Factual Errors in Summarization: Errors, Summarizers, Datasets, Error Detectors](#). *arXiv*.

Craig Thomson, Ehud Reiter, and Barkavi Sundararajan. 2023. [Evaluating factual accuracy in complex data-to-text](#). *Computer Speech & Language*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2019. [HuggingFace’s Transformers: State-of-the-art Natural Language Processing](#). *arXiv*.

Wen Xiao and Giuseppe Carenini. 2022. [Entity-based SpanCopy for Abstractive Summarization to Improve the Factual Consistency](#). *Arxiv Preprint*.

Amir Zeldes. 2017. [The GUM Corpus: Creating Multilayer Resources in the Classroom](#). *Language Resources and Evaluation*, 51(3):581–612.

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*.

Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. 2019. [MoverScore: Text generation evaluating with contextualized embeddings and earth mover distance](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 563–578, Hong Kong, China. Association for Computational Linguistics.

Chenguang Zhu, William Hinthorn, Ruochen Xu, Qingkai Zeng, Michael Zeng, Xuedong Huang, and Meng Jiang. 2021. [Enhancing factual consistency of abstractive summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 718–733, Online. Association for Computational Linguistics.

Markus Zopf, Maxime Peyrard, and Judith Eckle-Kohler. 2016. [The next step for multi-document summarization: A heterogeneous multi-genre corpus built with a novel construction approach](#). In *Proceedings of COLING 2016, the 26th International*

*Conference on Computational Linguistics: Technical Papers*, pages 1535–1545, Osaka, Japan. The COLING 2016 Organizing Committee.

## A Genre-specific Guidelines

The following excerpts from genre-specific guidelines exemplify instructions which were given to annotators working on documents in those specific genres. The full guidelines can be viewed at <https://wiki.gucorpling.org/gum/summarization>.

### A.1 Biographies

Summaries for biographies and other texts centered around an individual:

- typically take the form “Kim is/was a French X who ...”
- typically include information about what this person is/was known for (“... best known for ...”)
- information about the time period and place is typically included (“a Japanese X”, “a German X living in France”, “a 19th century Kenyan X”)

Examples:

- Jared Padalecki is an award winning American actor who gained prominence in the series *Gilmore Girls*, best known for playing the role of Sam Winchester in the TV series *Supernatural*, and for his active role in campaigns to support people struggling with depression, addiction, suicide and self-harm.
- Jenna Nicole Mourey, better known as Jenna Marbles, is a very successful American YouTube personality, vlogger, comedian and actress, known for her videos “How To Trick People Into Thinking You’re Good Looking” and “How To Avoid Talking To People You Don’t Want To Talk To”.

### A.2 Fiction

- In non-metalinguistic texts (i.e. fiction itself, not texts about fiction), summarize the text as if it is a literal, true story; for example, “Huckleberry Finn is fishing”, not “In this extract from the novel *Huckleberry Finn*, fictional character Huck is...”



- Even if described events are factually incorrect, or involve science fiction or imaginary contexts, we summarize without commenting on this (e.g. “Three unicorns chat and decide to go fishing”)
- Unnamed active protagonists should be referred to as “a/the protagonist”
- An unnamed narrator who is not an agent in the story can be referred to as “a/the narrator”

Examples:

- Jacques Chalmers, a starfighter pilot for the Empire, is terrified of overwhelming enemy forces as he leaves his deployment carrier together with his comrades, and later narrowly escapes the Enemy after witnessing the destruction of the Kethlan system.
- Santa Claus’s second wife, Betty Moroz, plays online video games with her friends Williams and Gomez while making dinner on Christmas Eve, and is then disappointed when Santa gets a call from his secretary Ginny and goes out to take care of the children of the world, missing dinner.

### A.3 Vlogs

- Typically a present tense third person style is used, and events are ordered in sequence, for example: “Ash tells about her day, which includes a yoga class, marketing brand management class, doing some work while having coffee at Saxby’s, and finally cooking pasta with peppers for dinner together with her boyfriend Harry.”
- As in conversations, people other than the vlogger who play a significant role in the vlog should be mentioned, but if their name is not mentioned within the excerpt being annotated, then they can only be referred to using generic terms (“a friend/relative/...”)
- If the vlogger does not mention that they are a vlogger in the video, or that this is a vlog, do not refer to them as such (e.g. “Jasmine tells about ...”, not “YouTube vlogger Jasmine tells ...”)

Examples:

- Jasmine tells about how she tested positive for Covid on December 16th after she spent time without a mask with her sister, who also tested positive, and recounts her symptoms over several days, starting from a sore throat, then fever and congestion, and finally a partial loss of smell and taste and shortness of breath.

## B Experiment Details

### B.1 Fine-tuning on BRIO

All three fine-tuning sessions were conducted using 1 NVIDIA A100 40GB GPU on Google Cloud Platform, which cost \$2.8 per hour.<sup>6</sup> The configurations of BRIO for XSum<sup>7</sup> were used except that the default number of epochs was increased to 1000 from 100 in order to achieve better validation performance on GUMSum’s dev set. Specifically, we take BRIO’s *generation* model checkpoint on XSum from Huggingface’s Transformers (Wolf et al., 2019).<sup>8</sup> The average training time for a single run was about 7 hours. Table 3 shows the validation performance of each run on the documents from the dev set of GUM V9. Both dev and test partitions contain 24 documents, 2 for each genre, leaving 165 documents for training.<sup>9</sup>

	VAL_LOSS	VAL_R-1	VAL_R-2	VAL_R-L	BEST_epoch
RUN 1	72.3	39.3	14.5	29.3	899
RUN 2	71.9	39.9	15.3	29.2	799
RUN 3	73.0	38.3	14.1	28.6	849
AVG.	72.4	39.1	14.6	29.0	—

Table 3: FT Validation Performance on 24 dev docs.

### B.2 GPT3 Output Selection

We use OpenAI’s GPT3-text-davinci-002<sup>10</sup> with the prompt *Summarize the text above in one sentence.* and keep the default settings. Due to the nondeterministic nature and in order to ensure a fair comparison, we generated 3 summaries for each text and computed average ROUGE scores (the mean of R-1/2/L) against the human-written summaries and selected the summary with the middle average ROUGE score. At the time, the Davinci

<sup>6</sup><https://cloud.google.com/compute/docs/gpus#a100-40gb>

<sup>7</sup><https://github.com/yixinL7/BRIO/blob/main/config.py#L37-L71>

<sup>8</sup><https://huggingface.co/Yale-LILY/brio-xsum-cased>

<sup>9</sup>The complete list of train/dev/test document names is provided in the repository.

<sup>10</sup>GPT3-text-davinci-003 was not available at the time.

model costs \$0.0200 / 1K tokens. To avoid repetitive computation and to facilitate further research, we release all the GPT3-generated summaries for GUMSum. No post-editing was made on the GPT3-generated summaries.

### B.3 BRIO-/SimCLS- Generated Summaries

We use BRIO’s *generation* model checkpoint on XSum available on Huggingface (i.e. Yale-LILY/brio-xsum-cased) to obtain BRIO-generated summaries for GUMSum’s texts. For SimCLS (Liu and Liu, 2021), we use the checkpoint on XSum provided by the authors in their GitHub repository.<sup>11</sup> Although some BRIO-/SimCLS-generated summaries contain trailing punctuation, no post-editing was made on these system outputs.

## C Human Evaluation Details

We recruited 12 students who are native speakers of English to participate in this human evaluation study. Each student was assigned two documents from two different genres. They were given 4 weeks to work on a series of tasks for each document, as shown in Figure 4 below. Every student received a Google Form for each assigned text.

The screenshot shows a Google Form titled "Section 2 of 5". The first section is "OVERVIEW" which states: "Overall you will be assigned two texts (from different genres) and asked to accomplish the following 4 main tasks for each text, which in total should take 1-1.5 hours. You will receive two different Google Forms for the two texts." It then lists four tasks: 1. read the document, 2. write a one sentence summary for the document, 3. evaluate other existing summaries for the same document, and 4. answer general and aspect-specific questions about these summaries regarding their quality. Below this, it says "The rest of the form is organized as follows:" and lists: Page 3 contains tasks 1-2, Page 4 contains tasks 3-4, and Page 5 wraps up the evaluation study. The second section is "ACKNOWLEDGEMENT" with a red asterisk. It says: "I have read the description above and will be able to accomplish all the tasks for the two documents assigned to me by [redacted]". There is a radio button for "Yes".

Figure 4: Overview of Task Description.

**Tasks 1 and 2** Students were asked to review both general and genre-specific guidelines before writing their own one-sentence summary for the assigned document. We also asked for their consent

to release their written-summaries to GUMSum to facilitate multiple-reference evaluation and inter-annotator agreement, as shown in Figure 5.

The screenshot shows a Google Form with three sections. The first section is "Task 1: Read the text" with a red asterisk. It has two radio buttons: "Yes: I have read the text carefully" and "No". The second section is "Task 2: Write a one-sentence summary for the text you just read". It says: "Please follow both the general and genre-specific guidelines described at [redacted] carefully where sample summaries are also provided." The third section asks: "Have you reviewed the guidelines (both the general and your assigned genre-specific) carefully before writing your summary?" with a red asterisk. It has two radio buttons: "Yes" and "No". The fourth section says: "Please write your summary for the text you just read." with a red asterisk. It includes a note: "It must be one sentence long and may not exceed 380 characters." Below this is a text input field labeled "Long answer text". The fifth section is "CONSENT" with a red asterisk. It says: "We would like to request your consent to publish the summary you wrote above under a CC0 license, meaning that the written summary above will be included in a dataset that will be made publicly available later. Your work will be acknowledged, but your name and information will NOT be linked to the individual summary you wrote." It has two radio buttons: "Yes: I consent to make my written summary available in this dataset for research purposes" and "No: I do not consent to make my written summary available in this dataset for research purposes".

Figure 5: A Screenshot of Tasks 1 and 2.

**Tasks 3 and 4** Students were presented both system-generated and human-written summaries in order to evaluate various aspects of each summary candidate. The order of outputs shown to the evaluators was randomized for each source text, and we also ask them to not modify their written summary after viewing the presented ones. In addition, we ask the evaluators to justify their decisions in a few sentences for certain questions:

[1] **Please choose your most and least preferred summaries respectively.** You can select more than one for each category below if multiple summaries are equally most or least preferred by you.

- Please justify your decisions above in a few sentences below. *For instance, you could say, "I prefer summary X over summary Y because X doesn't contain the main point (while a minor one is included) or Y contains incorrect information" etc. The more detailed the justifications, the better!*

<sup>11</sup><https://github.com/yixinL7/SimCLS>

- [2] **How substitutive is each summary candidate?** According to the guidelines, substitutive summaries replace reading the text as best as possible in one sentence - they are not just meant to attract readers to reading the text; they are meant to save you the trouble of reading it)
- [3] Does the summary include information **NOT PRESENT** in the text **even if you happen to know that it is factually correct**?
- Please justify your decisions (esp. the ones you chose YES for) above in a few sentences below. For instance, you can list the relevant information below.
- [4] Does the summary include **INCORRECT information**? (i.e. information **PRESENT** in the original text but used or interpreted *in a different, misleading, or incorrect way in the summary*; in other words, this summary is not faithful to the original text)
- Please justify your decisions (esp. the ones you chose YES for) above in a few sentences below. For instance, you can list the relevant information below.
- [5] **Is the summary written in good English?** (e.g. no grammar errors or incomplete sentences etc.)
- [6] **Can you tell which summary is human-written and which one is computer-generated?** If you are very unsure about this (confidence level at or below 50%), then choose the "can't tell" category.

- Please justify your decisions above in a few sentences below. *In particular, if you have a very strong opinion about a specific summary or certain summaries, we'd highly appreciate it if you could share your valuable thoughts with us.*

**Wrapping-up** The last part of the evaluation study is to ask evaluators to first rate the level of difficulty of the entire evaluation task on a scale of 1 to 5 where 1 means 'Not difficult at all' and 5 means 'Extremely difficult'. We also collect their responses to the following open-ended questions in order to help us get a better idea of the challenges of producing a good summary for various text types, which are very valuable insights to guide future

research on designing more specifically defined guidelines and targeted evaluation.

- [1] Based on your experience here, what's the most difficult or challenging thing you found when writing a one-sentence summary for the genre you are assigned?
- [2] Is there anything else you would like to share regarding your experience of writing a summary and/or evaluating other existing summaries?

### C.1 Additional Plots of Responses from the Human Evaluation Study

Figure 6 shows additional responses on English fluency quality for selected systems vs. human performance, as well as a breakdown of annotators' guesses as to whether they were looking at human or system summaries.





Figure 6: Barplots of Human Evaluations on English Quality and Source of Summaries.