

On the Role of Summary Content Units in Text Summarization Evaluation

Marcel Nawrath^{1*}, Agnieszka Nowak^{1*}, Tristan Ratz^{1*}, Danilo C. Walenta^{1*}, Juri Opitz^{2*}, Leonardo F. R. Ribeiro^{3†}, João Sedoc⁴, Daniel Deutsch⁵, Simon Mille⁶, Yixin Liu⁷, Lining Zhang⁴, Sebastian Gehrmann⁸, Saad Mahamood⁹, Miruna Clinciu¹⁰, Khyathi Chandu¹¹, Yufang Hou^{12*‡}

¹TU Darmstadt, ²University of Zurich, ³Amazon AGI, ⁴New York University, ⁵Google Research, ⁶ADAPT Centre, DCU, ⁷Yale University, ⁸Bloomberg, ⁹trivago N.V., ¹⁰University of Edinburgh, ¹¹Allen Institute for AI, ¹²IBM Research Europe, Ireland

Abstract

At the heart of the Pyramid evaluation method for text summarization lie human written summary content units (SCUs). These SCUs are concise sentences that decompose a summary into small facts. Such SCUs can be used to judge the quality of a candidate summary, possibly partially automated via natural language inference (NLI) systems. Interestingly, with the aim to fully automate the Pyramid evaluation, Zhang and Bansal (2021) show that SCUs can be approximated by automatically generated semantic role triplets (STUs). However, several questions currently lack answers, in particular: i) Are there other ways of approximating SCUs that can offer advantages? ii) Under which conditions are SCUs (or their approximations) offering the most value? In this work, we examine two novel strategies to approximate SCUs: generating SCU approximations from AMR meaning representations (SMUs) and from large language models (SGUs), respectively. We find that while STUs and SMUs are competitive, the best approximation quality is achieved by SGUs. We also show through a simple sentence-decomposition baseline (SSUs) that SCUs (and their approximations) offer the most value when ranking short summaries, but may not help as much when ranking systems or longer summaries.

1 Introduction

Judging the quality of a summary is a challenging task. Besides being short and faithful to its source document, a summary should particularly excel in *relevance*, that is, a summary should select only the most relevant or salient facts from a source document. An attractive method for assessing such notion of relevance is the *Pyramid*-method (Nenkova and Passonneau, 2004) that is based on so-called

Summary Content Units (SCUs) which decompose a reference summary into concise human-written English sentences. With SCUs available from one or different reference summaries, we can then more objectively assess the degree to which a generated summary contains the relevant information. With the aim to fully automate the Pyramid method, Zhang and Bansal (2021) suggest that the required human effort can be partially and even fully alleviated, by i) automatically generating SCUs and ii) validating the relevance of a generated summary with a natural language inference (NLI) system that checks how many SCUs are entailed by the generated summary.

Since strong NLI systems are available off-the-shelf and are known to be useful in natural language generation (NLG) evaluation (Chen and Eger, 2023; Steen et al., 2023), the generation of SCUs appears as the most challenging and least-understood part of an automated Pyramid method. Indeed, while Zhang and Bansal (2021) show that SCUs can be approximated by phrasing semantic role triplets using a semantic role labeler and coreference resolver, we lack availability and understanding of possible alternatives as well as their potential impact on downstream-task summary evaluation in different scenarios.

In this work, we proposed two novel approaches to approximate SCUs: *semantic meaning units* (SMUs) that are based on abstract meaning representation (AMR) and *semantic GPT units* (SGUs) that leverage SoTA large language models (LLMs). We carry out experiments to systematically evaluate the intrinsic quality of SCUs and their approximations. On the downstream task evaluation, we find that although SCUs remain the most effective metric to rank different systems or generated summaries across three meta-evaluation datasets, surprisingly, an efficient sentence-splitting baseline already yields competitive results when compared to SCUs. In fact, the sentence-splitting baseline

* Equal contributions.

† Work done prior to joining Amazon.

‡ Correspondence to yhou@ie.ibm.com.

outperforms the best SCU approximation method on a few datasets when ranking systems or long generated summaries.

In summary, our work provides important insights into the application of automation of the Pyramid method in different scenarios for evaluating generated summaries. We make the code publicly available at <https://github.com/tristanratz/SCU-text-evaluation/>.

2 Related work

Over the past two decades, researchers have proposed a wide range of human-in-the-loop or automatic metrics to assess the quality of generated summaries in different dimensions, including linguistic quality, coherence, faithfulness, and content quality. For more in-depth surveys on this topic, please refer to Howcroft et al. (2020), El-Kassas et al. (2021), and Gehrmann et al. (2023).

In this work, we focus on evaluating the content quality of a generated summary that assesses whether the summary effectively captures the salient information of interest from the input document(s). In the reference-based metrics, content quality is assessed by comparing system-generated summaries to human-written reference summaries. The Pyramid method (Nenkova and Passonneau, 2004) is regarded as a reliable and objective approach for assessing the content quality of a generated system summary. Below we provide a brief overview of the Pyramid method and highlight several previous efforts to automate this process.

Pyramid Method. The original Pyramid method (Nenkova and Passonneau, 2004) comprises two steps: SCUs creation and system evaluation. In the first step, human annotators exhaustively identify Summary Content Units (SCUs) from the reference summaries. Each SCU is a concise sentence or phrase that describes a single fact. The weight of an SCU is determined by the number of references in which it occurs. In the second step, the presence of each SCU in a system summary is manually checked. The system summary’s Pyramid score is calculated as the normalized sum of the weights of the SCUs that are present. Later, Shapira et al. (2019) introduce a revised version of the original Pyramid method where they eliminate the merging and weighting of SCUs, thereby enabling SCUs with the same meaning to coexist.

Automation of the Pyramid Method. Given the high cost and expertise required for implementing the Pyramid method, there have been attempts to automate this approach in recent years. Peyrard et al. (2017) propose an automatically learned metric to directly predict human Pyramid scores based on a set of features. Zhang and Bansal (2021) propose a system called *Lite³ Pyramid* that uses a semantic role labeler to extract semantic triplet units (STUs) to approximate SCUs. They further use a trained natural language inference (NLI) model to replace the manual work of assessing SCUs’ presence in system summaries. In our work, we explore two new methods to approximate SCUs. We further investigate the effectiveness of the automated Pyramid method in different scenarios.

3 SCU approximation I: SMU from AMR

Abstract Meaning Representation (AMR) (Banasescu et al., 2013) is a widely-used semantic formalism employed to encode the meaning of natural language text in the form of rooted, directed, edge-labeled, and leaf-labeled graphs. The AMR graph structure facilitates machine-readable explicit representations of textual meaning.

Motivated by Zhang and Bansal (2021)’s observation that STUs based on semantic roles cannot well present single facts in long reference summary sentences that contain a lot of modifiers, adverbial phrases, or complements, we hypothesize that AMR has the potential to capture such factual information more effectively. This is because, in addition to capturing semantic roles, AMR models finer nuances of semantics, including negations, inverse semantic relations, and coreference.

To generate semantic meaning units (SMUs) from a reference summary, we employ a **parser** that projects each sentence of the reference summary onto an AMR graph, then split the AMR graph into meaningful, event-oriented subgraphs. Finally, we use a **generator** to generate a text (an SCU approximation) from each subgraph.¹

While there may be various conceivable ways to *extract subgraphs*, for our experiment we use simple and intuitive **splitting** heuristics. Given an AMR graph, we first extract all predicates to

¹For **parser** and **generator**, we use strong off-the-shelf models from <https://github.com/bjascob/amr-lib-models>: `parse_xfm_bart_large` and `generate_t5wtense`. `parse_xfm_bart_large` is fine-tuned on `BART_large`. The model achieves a high Smatch score on the standard AMR benchmark (83.7 SMATCH on the AMR-3 test dataset).

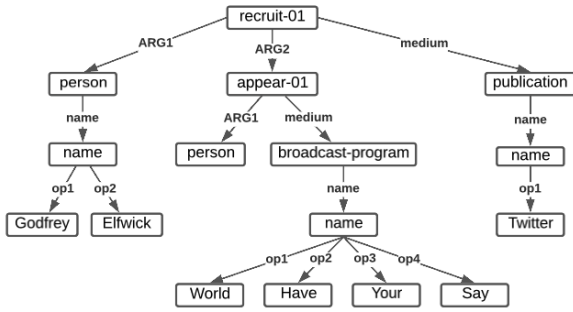


Figure 1: The AMR graph for the sentence “*Godfrey Elfwick recruited via Twitter to appear on World Have Your Say*”

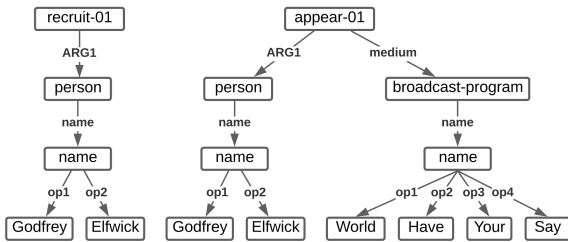


Figure 2: Two AMR sub-graphs for the sentence “*Godfrey Elfwick recruited via Twitter to appear on World Have Your Say*”

discern their semantic meaning as we view them to form the core of a sentence’s meaning. Subsequently, the argument connections within the predicates were examined. If a predicate is connected to at least one core role (CR), indicated by ARG_n edge labels, we extract a sub-graph for every CR of this predicate containing the CR and the underlying connections. Figure 2 shows an example of two extracted subgraphs from the AMR graph in Figure 1 for the input sentence “*Godfrey Elfwick recruited via Twitter to appear on World Have Your Say*” based on the **parser** and **splitting** steps. Finally, we generate two SMUs by applying the **generator** to the two subgraphs in Figure 2:

- SMU 1: *Godfrey Elfwick was recruited.*
- SMU 2: *Godfrey Elfwick will appear on World Have Your Say.*

4 SCU approximation II: SGU from LLM

Recently, it became widely known that pre-trained large language models (LLMs) are able to generate high-quality output according to prompts given by humans, optionally exploiting a few shown examples through in-context learning (Brown et al., 2020). Therefore, we try to approximate SCUs

using GPT models from OpenAI, calling the resulting units Semantic GPT Units (SGUs). Specifically, we use GPT-3.5-Turbo which is built on Instruct-GPT (Ouyang et al., 2022) and GPT-4 (OpenAI, 2023) to generate SGUs (SGUs_3.5 and SGUs_4) for each reference summary using the same prompt and a one-shot example. Please refer to Appendix A.1 for more details.

5 Experiments

5.1 Dataset and NLI models

Datasets. We run the experiments on four existing English meta-evaluation datasets: (1) TAC08 (NIST, 2008), (2) TAC09 (NIST, 2009), (3) REAL-Summ (Bhandari et al., 2020) and (4) PyrXSum (Zhang and Bansal, 2021). We use TAC08 for development purposes and evaluate the results on the last three datasets. Each dataset contains one or multiple reference summaries, the corresponding human-written SCUs, the generated summaries from different systems, and the human evaluation result for each summary/system based on the Pyramid method. Table 1 shows some statistics of the reference summaries across different datasets. In general, PyrXSum contains short and abstractive summaries, while RealSumm and TAC09 contain long and extractive summaries. More details on the datasets can be found in appendix A.2.

NLI Models. We use the NLI model from Zhang and Bansal (2021)² that was fine-tuned on TAC08’s SCU presence gold annotations based on a NLI model from Nie et al. (2020). Following Zhang and Bansal (2021), we use the fine-tuned NLI model with the probability of presence label to calculate the Pyramid score of a generated summary. Please refer to Appendix A.4 for more details.

5.2 Baselines

STUs are short sentences that are based on semantic role (SR) triples from an SR-labeling and coreference system (Zhang and Bansal, 2021).

Sentence splitting is a baseline that may shed light on the overall usefulness of SCUs in summary evaluation. We split every reference summary into sentences and treat them as SCU approximations.

N-grams consist of phrases randomly extracted from a reference summary. For each sentence from

²The model can be downloaded from the HF model hub: <https://huggingface.co/shiyue/roberta-large-tac08>.

	RealSumm	PyrXSum	TAC09
Avg. # sent.	4.73	2.02	27.22
Avg. # words	63.71	20.56	386.82
Avg. # words/sent	13.47	10.18	14.21
# ref summary	1	1	4
Avg. # SCUs	10.56	4.78	31.63

Table 1: Statistics of the reference summaries from different datasets.

Metrics	RealSumm		PyrXSum		TAC09	
	R	P	R	P	R	P
sentence split	.54	.67	.41	.54	.50	.54
ngrams	.41	.52	.38	.52	.46	.39
STUs	.66	.68	.54	.65	.61	.53
SMUs	.56	.58	.53	.58	.52	.48
SGUs_3.5	.58	.67	.58	.63	.36	.48
SGUs_4	.61	.69	.61	.66	.52	.61

Table 2: Intrinsic evaluation results. R is the recall-oriented simulation easiness score from Zhang and Bansal (2021), while P is our precision-oriented score that is computed in the reverse direction.

the summary, we naively generate all possible combinations of $\{3, 4, 5\}$ consecutive words. We randomly select a subset from these combinations, which accounts for 5% of all combinations.

5.3 Intrinsic Evaluation

As proposed by Zhang and Bansal (2021), we evaluate approximation quality with an *easiness score*. The score is built by iterating over each SCU-SxU pair and averaging over the maximum ROUGE-1-F1 score found for each SCU. Naturally this score is recall-biased, and therefore, we also present the score calculated in the reverse direction, to evaluate the precision of our approximated SCUs (c.f. Appendix A.3 for more details). The results are shown in Table 2. We find that best approximation quality for RealSumm is achieved by STUs, while for PyrXSum, SGU_4 performs best. Considering the longer texts of TAC09, STUs excel in recall, while SGU_4 excels in precision.

5.4 Extrinsic Evaluation

Our downstream evaluation consists of two parts: summary quality evaluation at the system and summary levels, respectively. System-level correlation assessment evaluates the ability of the metric to compare different summary systems individually. In contrast, summary-level evaluation determines the metric’s ability to compare summaries created

by different systems for a common set of documents. Following Zhang and Bansal (2021), we use Pearson r and Spearman ρ to evaluate the correlations between metrics with gold human labeling scores. Pearson measures linear correlation and Spearman measures ranking correlation. Please refer to Appendix A.4 for more details about using the NLI model to score a generated summary and how to calculate these two types of correlations.

The results are shown in Table 3.³ In general, SGUs offer the most useful SCU approximation, with the exception of TAC09 (summary-level), where STUs remain the best approximation method, slightly outperforming our simple sentence-splitting baseline. However, SGUs still lack the usefulness of true SCUs, which seem to remain the most useful way to evaluate system summary quality (if resources permit). Interestingly, however, to discriminate the quality of systems, it is enough to use any approximation, even the sentence split baseline is sufficient to accurately discriminate between systems.

5.5 Human Evaluation

For a representative sample of human results of our experiment, three authors evaluated the quality of SCUs, STUs, SMUs and SGUs_4 for 10 reference summaries randomly sampled from REALSumm and PyrXSum, annotating each of 40 examples according to 3 dimensions: Well-formedness, Descriptiveness and Absence of hallucination, amounting to a total of 240 annotation hits. Please refer to Appendix A.5 for more details of the annotation scheme.

Overall, Cohen’s κ scores among three annotators range from 0.37 to 0.87. After a thorough check, we found that all annotators agree on the general trend (i.e., SCUs and SGUs are generally better than SMUs and STUs). One annotator appeared to diverge from the other two by slightly favoring SMUs over STUs. To increase the power of the experiment, two annotators then annotated another 20 summaries each, resulting in an additional 480 annotation hits.

The findings shown in Figure 3 illustrate that the quality of SMUs is comparable to the STUs.

³Note that we do not include comparisons with the recent automatic evaluation metrics based on LLMs such as GPTScore (Fu et al., 2023). Recent studies pointed out that these automatic metrics are not as effective as the traditional automatic evaluation metrics, such as ROUGE-1, to compare the summaries of different systems in terms of content coverage (Liu et al., 2023).

Metrics	System-Level						Summary-Level					
	RealSumm		PyrXSum		TAC09		RealSumm		PyrXSum		TAC09	
	r	ρ	r	ρ	r	ρ	r	ρ	r	ρ	r	ρ
SCUs	.95	.95	.98	.98	.99	.97	.59	.58	.70	.69	.76	.70
SCU Approximations												
- sentence split	.93	.95	.97	.97	.97	.94	.48	.46	.37	.36	.73	.66
- ngrams	.90	.92	.94	.82	.96	.92	.36	.35	.38	.38	.65	.61
- STUs	.92	.94	.95	.95	.98	.95	.51	.50	.46	.44	.73	.67
- SMUs	.94	.94	.96	.94	.98	.96	.50	.48	.46	.44	.70	.64
- SGUs_3.5	.93	.95	.97	.93	.96	.88	.49	.46	.56	.55	.54	.49
- SGUs_4	.92	.94	.97	.95	.98	.96	.54	.52	.58	.56	.71	.66

Table 3: Results of different metrics on three datasets. Best numbers among all SCU approximations are bolded.

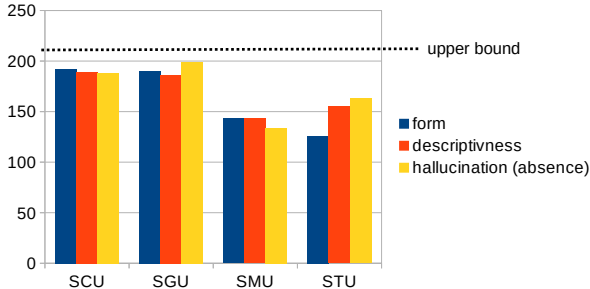


Figure 3: Human evaluation results. Each bar represents the sum of scores aggregated over all annotators. Upper-bound indicates the best possible result (each annotator always assigns the maximum quality score).

But the revealed overall trend is clearly that SCUs (human written units) and SGUs (LLM generated units) achieve *similar* and *very high* quality, while STUs (triplet units based on SRL) and SMUs (units from AMR semantic graph) are similarly of lower quality. To see if these differences are significant, we calculate the Wilcoxon signed-rank test. For all categories (descriptiveness, well-formedness, and absence of hallucination), the human SCUs and SGUs are not of significantly different quality ($p < 0.005$). However, both SCUs and SGUs are of significantly better quality than STUs and SMUs ($p < 0.005$). Within SMUs and STUs, the categories of descriptiveness and well-formedness are not of significantly different quality ($p < 0.005$). However, STUs are significantly better in reduction of hallucination compared with SMUs ($p < 0.005$), an outcome that could be explained either by destruction of coherent information when splitting the AMR graph, or hallucination of the AMR models.⁴

The result of the human annotation, however,

⁴While AMR parsers nowadays achieve impressive scores on benchmarks (Lam et al., 2021; Lee et al., 2022), recent research shows that they still make crucial errors (Opitz and Frank, 2022; Groschwitz et al., 2023).

must not be taken as proof that there is quality parity of SCUs and SGUs. Indeed, when contrasting the finding of the human evaluation, where SCUs and SGUs appear of similar high quality, against the empirical finding that SCUs provide substantially better downstream performance *for shorter texts in summary-level evaluation*, we have reason to believe that there is a quality aspect of SCUs that both LLM/ChatGPT and our annotation setup failed to measure.

6 Discussion and conclusions

This work focuses on automating the Pyramid method by proposing and evaluating two new methods to approximate SCUs. We found out that there are more effective ways of approximating SCUs than with STUs only, and our extrinsic evaluation suggests that costly SCUs and approximations may even be unnecessary for system comparisons.

There are several aspects worth discussing. Firstly, as shown by comparing, e.g., the results of STUs and SMUs in Table 2 and Figure 3, it appears that ROUGE-1-F1 exhibits a weak correlation with human evaluation. This raises concerns about the effectiveness of using this metric in previous studies to evaluate the quality of SCU approximations. Secondly, it seems that we may not need the costly SCUs and their approximations to compare summarization systems or rank long generated summaries (TAC09). Surprisingly, a simple sentence splitting baseline already achieves competitive results compared to SCUs on these tasks, while automatically obtained SGUs generally score high both on system- and summary-level evaluations. Finally, SCUs and their approximations offer the most value for summary-level evaluation, especially when summaries are rather short (PyrXSum and RealSumm).

Limitations

First, we would have liked to achieve better performance with SMUs generated from an AMR. In theory, using AMR graph splitting would ideally decompose a textual meaning into parts, and the AMR generation systems promise to phrase any such sub-graph into natural language. Inspecting all three parts of our SMU generation pipeline (parsing, splitting, and generating), we find that some issues may be due to our manually designed splitting strategy being too naive. While the rules are simple and their creation has profited from communication with AMR-knowledgeable researchers, a main problem is that there are countless possibilities of how to split an AMR graph, and the importance of rules can depend on the graph context. Therefore, we believe that future work can strongly improve the AMR-based approach by learning how to better split meaning representation graphs.

Second, our NLI system was fine-tuned on gold SCUs extracted from the development data (TAC08), since this was found to work best by [Zhang and Bansal \(2021\)](#). While in principle this does not affect the evaluation of SxUs, which was the focus of this paper, it is not unlikely that by training the NLI system on each SxU type separately, the results of SxUs may further improve. Therefore the results for human written SCUs can be considered slightly optimistic. In general, the interaction of NLI and SCUs in an automated Pyramid method needs to be better understood. Other recent findings ([Chen and Eger, 2023](#); [Steen et al., 2023](#)) suggest that NLI models may play an underestimated role in NLG evaluation. As a check, we repeated the evaluation with an NLI system without SCU fine-tuning, and observe significant performance drops across the board, indicating that (i) SCU results are likely not too over-optimistic in comparison to SxUs, and (ii) the effective adaptation strategy of the NLI system may be the second cornerstone of an accurate automatic Pyramid method and thus should be explored in future work.

Finally, although our results offer insights into the design choices when applying the automatic Pyramid method for text summarization evaluation in different scenarios (short vs. long summaries), we do not explore its applications beyond summary and summarizer evaluation and ranking, such as using basic meaning units to evaluate the factual precision in long-form text generation ([Min et al., 2023](#)). We leave this for future studies.

References

- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking. In *Proceedings of the 7th linguistic annotation workshop and interoperability with discourse*, pages 178–186.
- Manik Bhandari, Pranav Narayan Gour, Atabak Ashfaq, Pengfei Liu, and Graham Neubig. 2020. [Re-evaluating evaluation in text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9347–9359, Online. 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*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Yanran Chen and Steffen Eger. 2023. [MENLI: Robust evaluation metrics from natural language inference](#). *Transactions of the Association for Computational Linguistics*, 11:804–825.
- Wafaa S. El-Kassas, Cherif R. Salama, Ahmed A. Rafea, and Hoda K. Mohamed. 2021. [Automatic text summarization: A comprehensive survey](#). *Expert Systems with Applications*, 165:113679.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. [GpScore: Evaluate as you desire](#). *arXiv preprint arXiv:2302.04166*.
- Sebastian Gehrmann, Elizabeth Clark, and Thibault Selam. 2023. [Repairing the cracked foundation: A survey of obstacles in evaluation practices for generated text](#). *Journal of Artificial Intelligence Research*, 77:103–166.
- Jonas Groschwitz, Shay Cohen, Lucia Donatelli, and Meaghan Fowlie. 2023. [AMR parsing is far from solved: GrAPES, the granular AMR parsing evaluation suite](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10728–10752, Singapore. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *NIPS*.

- David M. Howcroft, Anya Belz, Miruna-Adriana Clinciu, Dimitra Gkatzia, Sadid A. Hasan, Saad Mahamood, Simon Mille, Emiel van Miltenburg, Sashank Santhanam, and Verena Rieser. 2020. [Twenty years of confusion in human evaluation: NLG needs evaluation sheets and standardised definitions](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 169–182, Dublin, Ireland. Association for Computational Linguistics.
- Hoang Thanh Lam, Gabriele Picco, Yufang Hou, Young-Suk Lee, Lam M. Nguyen, Dzung T. Phan, Vanessa López, and Ramon Fernandez Astudillo. 2021. [Ensembling graph predictions for amr parsing](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*.
- Young-Suk Lee, Ramón Astudillo, Hoang Thanh Lam, Tahira Naseem, Radu Florian, and Salim Roukos. 2022. [Maximum Bayes Smatch ensemble distillation for AMR parsing](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5379–5392, Seattle, United States. Association for Computational Linguistics.
- Yixin Liu, Alex Fabbri, Pengfei Liu, Yilun Zhao, Linyong Nan, Ruilin Han, Simeng Han, Shafiq Joty, Chien-Sheng Wu, Caiming Xiong, and Dragomir Radev. 2023. [Revisiting the gold standard: Grounding summarization evaluation with robust human evaluation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4140–4170, Toronto, Canada. Association for Computational Linguistics.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FACTScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. 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 Rebecca Passonneau. 2004. [Evaluating content selection in summarization: The pyramid method](#). In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 145–152, Boston, Massachusetts, USA. Association for Computational Linguistics.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- NIST. 2008. <https://tac.nist.gov/publications/2008/papers.html>.
- NIST. 2009. <https://tac.nist.gov/publications/2009/papers.html>.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Juri Opitz and Anette Frank. 2022. [Better Smatch = better parser? AMR evaluation is not so simple anymore](#). In *Proceedings of the 3rd Workshop on Evaluation and Comparison of NLP Systems*, pages 32–43, Online. Association for Computational Linguistics.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. [Training language models to follow instructions with human feedback](#). *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Maxime Peyrard, Teresa Botschen, and Iryna Gurevych. 2017. [Learning to score system summaries for better content selection evaluation](#). In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 74–84, Copenhagen, Denmark. Association for Computational Linguistics.
- Ori Shapira, David Gabay, Yang Gao, Hadar Ronen, Ramakanth Pasunuru, Mohit Bansal, Yael Amsterdamer, and Ido Dagan. 2019. [Crowdsourcing lightweight pyramids for manual summary evaluation](#). 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 682–687, Minneapolis, Minnesota. Association for Computational Linguistics.
- Julius Steen, Juri Opitz, Anette Frank, and Katja Markert. 2023. [With a little push, NLI models can robustly and efficiently predict faithfulness](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 914–924, Toronto, Canada. Association for Computational Linguistics.
- Shiyue Zhang and Mohit Bansal. 2021. [Finding a balanced degree of automation for summary evaluation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6617–6632, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

A Appendix

A.1 Obtaining SGUs from the GPT Models

Below we show an example of the prompt we use to obtain SGUs from the GPT-3.5-Turbo and GPT-4 models through OpenAI APIs⁵. We randomly choose a summary and the corresponding human-written SCUs from the TAC08 dataset as the one-shot example. We did not alter the standard parameters except for temperature which we set to 0 to ensure reproducibility.

System Instruction You split the provided input in small sentences separated by an #. The split sentences represent subsentences of the original sentences.

Example Input Irish PM Ahern said the main goal of the US-brokered Good Friday pact of 1998, a joint Catholic-Protestant administration in Northern Ireland, could be revived only with a complete end of IRA weapons use. The landmark peace deal led to a virtual end of violence in that area. Sinn Fein leader Gerry Adams has appealed to IRA members to end their armed struggle in favor of democratic politics. Hopes are rising in Northern Ireland that the IRA will disarm. British PM Blair and Ahern will chair a review of the Northern Ireland situation in London.

Example Output Good Friday pact was agreed in 1998 # Good Friday pact was a peace pact # Good Friday pact set up a joint Catholic-Protestant administration in Northern Ireland # Good Friday pact was mediated by the US # Irish Republican Army increased activity # Irish PM Ahern called to end violence # Sinn Fein Adams called to end violence # Hope in Northern Ireland that the IRA will disarm # British PM Blair and Ahern will chair a review of the Northern Ireland situation in London

Testing Reference Summary Netherlands midfielder Wesley Sneijder has joined French Ligue 1 side Nice on a free transfer.

Output From GPT-3.5-Turbo *Netherlands midfielder Wesley Sneijder has joined Nice # Sneijder was a free transfer # Nice is a French Ligue 1 side*

Output From GPT-4 *Netherlands midfielder Wesley Sneijder # Sneijder joined French Ligue 1 side Nice # Joined on a free transfer*

⁵<https://openai.com/blog/openai-api>

A.2 Dataset Details

In general, all datasets (TAC08, TAC09, RealSumm, PyrXSum) contain: a) human written reference summaries; b) human expert written SCUs that are derived from the human written reference summaries; c) automatic summaries generated from different systems; d) SCU-presence labels for all system summaries that are labeled using either in-house annotators or Amazon Mechanical Turk (AMT).

The TAC08 dataset includes 96 examples and outputs from 58 systems, while TAC09 contains 88 examples and outputs from 55 systems. Both datasets contain multiple reference summaries for each example, as well as the corresponding SCU annotations.

The REALSumm dataset contains 100 test examples from the CNN/DM dataset (Hermann et al., 2015) and 25 system outputs. The SCUs are labeled by the authors and SCU-presence labels are collected using Amazon Mechanical Turk (AMT).

PyrXSum (Zhang and Bansal, 2021) includes 100 test examples from the XSum dataset (Narayan et al., 2018), which contains short and abstractive summaries. Similar to REALSumm, the SCUs are manually labeled by the authors and SCU-presence labels are collected for summaries generated by 10 systems through AMT.

A.3 Intrinsic Evaluation Details

We calculate two intrinsic evaluation metrics: A recall-based easiness score and a precision-based easiness score, denoted by $EasinessR$ and $EasinessP$. They evaluate how accurately the generated SxU units resemble human written SCUs. For a sentence with N human-written SCUs,

$$EasinessR = \frac{\sum Acc_j}{N},$$

where

$$Acc_j = \max_m Rouge1_{F1}(SCU_j, SxU_m).$$

In the above formula, Acc_j finds the SxU unit that has the max $Rouge1_{F1}$ score with SCU_j . $EasinessR$ corresponds to the easiness score defined in Zhang and Bansal (2021). To complement the recall-based easiness score, we introduce a precision-based $EasinessP$ that is calculated as:

$$EasinessP = \frac{\sum Acc_j}{N},$$

where

$$Acc_j = \max_m Rouge1_{F1}(SxU_j, SCU_m).$$

This time, Acc_j finds the SCU unit that has the max $Rouge1_{F1}$ score with SxU_j .

A.4 Extrinsic Evaluation Details

Details about NLI models For extrinsic evaluation, we follow the previous method proposed in [Zhang and Bansal \(2021\)](#) and use the NLI model they fine-tuned on the TAC 2008 dataset. More specifically, based on the NLI model, a system summary s will be scored as:

$$Score_s = \sum P_{NLI}(e|SxU_j, s)/N,$$

where N is the total number of SxUs extracted from the gold reference summary or summaries, and $P_{NLI}(e|SxU_j, s)$ is the probability of the entailment class from the underlying NLI model that tells us how likely the unit SxU_j is entailed by the system summary s . [Zhang and Bansal \(2021\)](#) explored different ways of using the NLI model, including a standard 3-class setting of the NLI model (entail/neutral/contradict) and a fine-tuned version of a 2-class setting (present/not-present), as well as using either the output probability of entailment/present class or the predicted 1 or 0 entailment/present label. They reported that using the fine-tuned model with the probability of the presence label works the best. We use this setup in our work.

Details about calculating correlations *System-level* correlation assesses the metric’s ability to compare different summarization systems. This is denoted as K and measures the correlation between human scores (h), the metric (m), and the generated summaries (s) for N examples across S systems in the meta-evaluation dataset. The system-level correlation is then defined as:

$$K_{m,h}^{sys} = K([\frac{1}{N} \sum_{i=1}^N m(s_{i1}), \dots, \frac{1}{N} \sum_{i=1}^N m(s_{iS})], [\frac{1}{N} \sum_{i=1}^N h(s_{i1}), \dots, \frac{1}{N} \sum_{i=1}^N h(s_{iS})])$$

Summary-level correlation assesses the metric’s ability to compare summaries produced by different systems for a common document(s). The summary-level correlation is then defined as:

$$K_{m,h}^{sum} = \frac{1}{N} \sum_{i=1}^N K([m(s_{i1}), \dots, m(s_{iS})], [h(s_{i1}), \dots, h(s_{iS})])$$

A.5 Human annotated evaluation

The text units of each example were analyzed regarding Well-formedness, Descriptiveness and Absence of hallucination. For each dimension, we classified it into one of three categories based on the evaluator’s satisfaction with the system’s output. These categories ranged from “1 - Unhappy with system output”, “2 - implying dissatisfaction or a less than satisfactory result”, to “3 - Okay with system output (3)”. Below we denote ASCU for approximated summary content units (e.g., STUs, SMUs, and SGUs_4) and SCUs. We provide a detailed definition for each evaluation category:

- Well-formedness (surface quality)
 - 1: Many ASCUs are not concise English sentences
 - 2: Some ASCUs are not concise English sentences
 - 3: Almost all or all ASCUs are concise English sentences
- Descriptiveness (meaning quality I)
 - 1: Many meaning facts of the summary have not been captured well by the ASCUs
 - 2: Some meaning facts of the summary have not been captured by the ASCUs
 - 3: Almost every or every meaning fact of the summary has been captured by an ASCU
- Absence of hallucination (meaning quality II)
 - 1: Many ASCUs describe meaning that is not grounded in the summary
 - 2: There is some amount of ASCUs that describe meaning that is not grounded in the summary
 - 3: Almost no or no ASCU describes meaning that is not grounded in the summary

In the following, we show an example of the reference summary from PyrXSum and the corresponding SCUs and their approximations:

- **Reference summary:** West Ham say they are “disappointed” with a ruling that the terms of their rental of the Olympic Stadium from next season should be made public.
- **SCUs:** West Ham are “disappointed” with a ruling # The ruling is that their rental terms should be made public # West Ham will rent the Olympic Stadium from next season
- **STUs:** West Ham say they are “disappointed” with a ruling that the terms of their rental of the Olympic Stadium from next season should be made public # They are “disappointed” with a ruling that the terms of their rental of the Olympic Stadium from next season should be made public # should made public
- **SMUs:** West Ham say they are disappointed by the ruling that their terms of rental for the Olympic Stadium next season should be public # The ruling that the terms of West Ham’s Olympic Stadium rental next season should be public was disappointing # West Ham rent the Olympic Stadium # West Ham will rent the Olympic Stadium next season
- **SGUs_4:** West Ham is disappointed with a ruling # Terms of their Olympic Stadium rental should be made public # Olympic Stadium rental starts next season