

Can summarization approximate simplification? A gold standard comparison

Giacomo Magnifico Eduard Barbu

Institute of Computer Science

University of Tartu

Estonia

{giacomo.magnifico, eduard.barbu}@ut.ee

Abstract

This study explores the overlap between text summarization and simplification outputs. While summarization evaluation methods are streamlined, simplification lacks cohesion, prompting the question: how closely can abstractive summarization resemble gold-standard simplification? We address this by applying two BART-based BRIO summarization methods to the Newsela corpus, comparing outputs with manually annotated simplifications and achieving a top ROUGE-L score of 0.654. This provides insight into where summarization and simplification outputs converge and differ.

1 Introduction

Text simplification can operate at various linguistic levels—semantic, syntactic, or lexical—using diverse strategies to achieve specific goals (Pellow and Eskenazi, 2014; Paetzold and Specia, 2016; Chen et al., 2017; Van et al., 2021). In practice, Automatic Text Simplification (ATS) transforms complex text into simpler versions by splitting sentences, shortening length, and simplifying vocabulary and grammar. The best English-language ATS models rely on parallel corpora like WikiSmall (Zhu et al., 2010; Zhang and Lapata, 2017), aligning complex and simple sentences from standard and Simple English Wikipedias (originally 108,000 instances from 65,133 articles, currently 89,042). The most valuable resource for text simplification is the Newsela corpus Xu et al. (2015), which includes 9,565 news articles professionally rewritten at multiple reading levels, with 1,913 original articles and four levels of simplification. However, it lacks the volume needed to train advanced deep-learning models effectively.

Simplification lacks standardized procedures and a common algorithm, partly due to the absence

of a "native speaker of simplified language" (Siddharthan, 2014). The subjective nature of simplification also makes consistent methodology difficult (Grabar and Saggion, 2022). The evaluation metrics for simplification are similarly inconsistent. Some, like BLEU or Levenshtein distance (Levenshtein, 1965; Papineni et al., 2002), focus on intrinsic grammatical features and struggle with semantic changes, while others, such as cosine distance, emphasize semantic similarity. By contrast, summarization metrics are well-established, even when imperfectly applied (Grusky, 2023). Furthermore, while the two tasks present some divergences in their focus (e.g. the relevance of information ordering, the choice of domain-agnostic lexicon, and the preference for short active forms instead of long passive forms), they remain convergent in producing shorter and poignant text. Given the state of things, we believe that comparing simplification with summarization could provide insights into their convergence.

This study investigates whether a state-of-the-art (SotA) summarization system can approximate manual simplification by comparing annotated simplifications with automated summarization. Starting with Newsela’s English documents, we process original articles with BRIO (Liu et al., 2022), a SotA abstractive summarizer, applying document-wide and paragraph-by-paragraph summarization methods. We then evaluate each output set against the four simplification levels using ROUGE-L scores to measure similarity. Results indicate an average performance difference of 0.444, with paragraph-by-paragraph summarization achieving the highest score (0.654) at level 1, gradually decreasing through levels 2 to 4. While paragraph-by-paragraph summarization does not equate to manual simplification, it may serve as an effective preparatory step for manual annotators.

Background and related research are discussed

in Section 2, with the experimental setup and findings detailed in Sections 3 and 4. A summary of the presented work, followed by the limits of the scope and suggestions for future research, are provided in Section 5.

2 Related work

The multifaceted nature of implementing text simplification has led to multiple works that share the goal of rewriting complex documents with simpler, more straightforward language. This is ultimately achieved by modifying the original text both lexically and syntactically as defined in Truică et al. (2023), either in an automated or manual way. Multiple works in the field have tackled different applications, from aiding people with disabilities (Rello et al., 2013; Chen et al., 2017), low-literacy adults (Watanabe et al., 2009; Paetzold and Specia, 2016), non-native learners (Allen, 2009; Pellow and Eskenazi, 2014) to auxiliary systems to improve the effectiveness of other NLP tasks (Stymne et al., 2013; Wei et al., 2014; Štajner and Popovic, 2016).

Due to the wide range of applications, a major subjectivity issue emerges when evaluating the different methods for simplification (Grabar and Saggion, 2022). Different scoring methods that have been utilized for simplification include: *BLEU* (Papineni et al., 2002); *TERp*, Translation Edit Rate plus, which computes the number of the three edit operations plus the inverse (Snover et al., 2009); *OOV*, Out Of Vocabulary, which measures the rate of oov words from a chosen simple vocabulary (e.g. Basic English list) (Vu et al., 2014); *changed*, measuring the percentage of the test examples where the system suggested some change (Horn et al., 2014); *potential*, computing the proportion of instances in which at least one of the candidates generated is in the gold-standard (Paetzold and Specia, 2016); *SARI*, the most recent, which performs a similar comparison to BLEU but is considered more reliable (Xu et al., 2016).

The general approach to text summarization is more streamlined, aiming to produce a shorter text than the input one while keeping all relevant information, defined as abstract or summary (Moiyadi et al., 2016). The most common approaches are naïve Bayes (Kupiec et al., 1995; Gambhir and Gupta, 2017), swarm algorithms (Jarraya and Bouri, 2012; Izakian and Mesgari, 2015), and sequence-to-sequence models (Sutskever et al.,

2014; Zhang et al., 2020).

A further distinction can be made between abstractive and extractive summarization methods (Nazari and Mahdavi, 2019). Where extractive methods produce text by concatenating selected parts of the original document, abstractive methods apply language generation techniques to produce a shorter document (Jeæek and Steinberger, 2008; Gupta and Lehal, 2010). Standard scoring methods for text summarisation are precision/recall measures and various instances of *ROUGE* (Lin and Och, 2004a; Grusky, 2023), some examples being *ROUGE-n*, *ROUGE-L*, and the most recent *ROUGE-SEM* (Zhang et al., 2024).

The Newsela corpus is a collection of 1,130 articles rewritten and simplified by professional editors, aimed at children of different grade levels (Xu et al., 2015). From each individual article, four different versions have been derived through manual simplification process and labelled with a number from 1 to 4, representative of the level of simplification. Label 4 represents the most simplified output, suitable for a 3rd grader; label 3 represents an output suitable for a 4th grader; labels 2 and 1 identify outputs suitable for 6th and 7th graders. The original articles are suitable for 12th graders.

Considering possible modifications to the dataset past the authors' presentation of their work, the corpus currently consists of 9,565 documents, of which 1,913 original articles.

3 Experimental setup

For the purpose of this work, the architecture chosen to perform the summarization procedure was BRIO, a system presented in Liu et al. (2022) and based both on the BART architecture (Lewis et al., 2020) and the PEGASUS architecture (Zhang et al., 2020). The choice was motivated by its state-of-the-art performance in summarization tasks, its ease of availability and implementation, and the double-model-based system that it employs. The dual nature of BRIO is the result of fine-tuning two different architectures on two different datasets with a specific training paradigm. Since the two datasets were characterized by longer texts (Hermann et al., 2015) and shorter texts (Narayan et al., 2018), the two backbones for the architecture keep these properties. Therefore, the BART-based BRIO was chosen as a

summarizer for its performance with longer texts, as suggested by the original authors.

The original articles from the Newsela corpus were then processed through the summarization model. For each article, two procedures were followed to produce different output documents: document-wide summarization and paragraph-by-paragraph summarization, as explained below. A graphic representation of the general procedure is provided in Figure 1

Document-wide. The more intuitive application of text summarization, this method involved the generation of a single string containing the whole text by joining the various paragraphs and subsequently processing it with the summarizer model. Once the architecture produced an output string, it was written in a separate `*.txt` file.

Paragraph-by-paragraph. This summarization approach stems from the visual structure of academic texts, which usually separate topics and changes in content by dividing the document into paragraphs. Thus, the intuition was to make the architecture follow a similar pattern to preserve the content and produce a more effective summarization. This method implemented splitting the original text into paragraphs and processing each paragraph separately with the summarizer model. The resulting outputs were subsequently rejoined and written as a single document in a separate `*.txt` file.

Both procedures were applied to each of the original 1,913 English articles in the Newsela corpus, and the resulting two sets of summarized documents were compared to the simplified version produced by the editors. This was done by iterating through the different levels of simplification (1, 2, 3, and 4) and calculating the *precision*, *recall* and *ROUGE F1* score between each simplified version of the document and the summarized version of it. The resulting evaluation was stored, and the average was calculated level-wise for each metric with the scores from the whole set. Then, the scoring procedure was repeated for the remaining summarized set. The chosen evaluation score was *ROUGE-L* as it was both a part of the original BRIO publication (Liu et al., 2022) and a statistic based on Long Common Sequence (LCS) (Lin and Och, 2004b), which made it well suited to measure the grammatical integrity, keyword conservation and coherence in the summarized texts.

4 Results

The average scores for the three evaluation metrics used in comparing the human-produced simplification and the automated summarization are available in Table 1. To provide an easier analysis, the scores have been divided by the level of simplification taken under scrutiny and the type of summarization procedure performed on the original articles. The upper section of Table 1 provides the average evaluation score between all the documents summarized with the first method mentioned in Section 3 and their simplified equivalent for each level. The second summarization method, paragraph-by-paragraph, is evaluated in the lower part of the Table.

Level	Precision	Recall	ROUGE-L
DOCUMENT			
label 1	0.058	0.918	0.109
label 2	0.061	0.884	0.113
label 3	0.066	0.811	0.122
label 4	0.078	0.731	0.141
PARAGRAPH			
label 1	0.731	0.615	0.654
label 2	0.721	0.561	0.616
label 3	0.703	0.461	0.541
label 4	0.699	0.354	0.451

Table 1: Average precision, recall and ROUGE-L scores when comparing the summarization output against the different levels of manually simplified articles. The table is divided according to the two types of summarization techniques presented, *document-level* and *paragraph-level*.

To make the gap in scores and the variability in summarization performance through the different processes more apparent, two graphic representations of the average scores are provided in Figure 2. The data corresponds to the document-wide summarization method on the left side and the paragraph-by-paragraph method on the right.

When comparing the results from the two processes, the overall difference in balance between precision/recall for the document-wide summarization method is immediately noticeable. Even considering the progressive improvement of the precision rate and the lowering of the recall score, the minimum gap between the two is 0.653. The first hypothesis was that it was due to the summarizer generating lengthy and repetitive summaries;

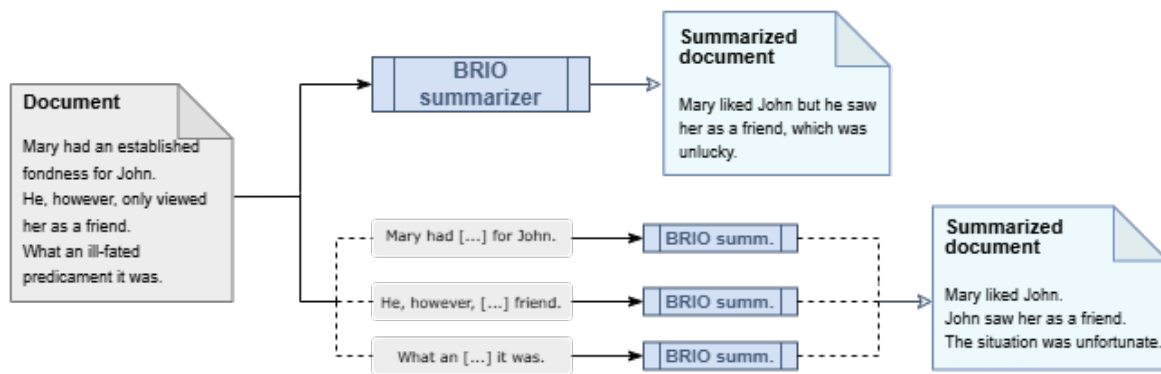


Figure 1: Representation of the processing pipeline for each article, showing the document-wide method (upper side) and paragraph-by-paragraph (lower side).




Figure 2: Comparison between the different levels of simplified text (1 to 4, left to right) and the two automated types of summarization. On the left is the performance of the document-wide summarization, on the right the performance of the paragraph-by-paragraph method.

however, a quick analysis of the outputs confirmed the variety in length and the production of shorter documents than their input. Therefore, the more plausible hypothesis is that while the longest common sentences between manual simplification and automated summarization are recalled in the text (most likely the keywords), the structural lexicon and syntactical choices of the simplified version would not appear through document-wide summarization. Consequently, this can lead to the poor similarity between the two document types and the convolution of information through summarization, a hypothesis corroborated by the low *ROUGE-L* score.

On the other side of Figure 2, the scores provide a better-looking picture of the paragraph-by-paragraph performance. With a *ROUGE-L* score of 0.566 averaged between all levels of simplification, the similarity between the simplified and summarized versions is noticeable. Although they perform better when compared to lower levels

of simplification than to more simplified documents, the summarized outputs obtained through paragraph-by-paragraph processing perform well enough to justify further investigation and analysis. Our hypothesis for the better performance of the paragraph-by-paragraph, when compared to the document-wide processing, lies in the nature of the process: a block-by-block iteration might be more similar to the manually performed annotation than a text-wide transformation is.

Worth of notice for the production of these results was the difference in time requirements between the first summarization method and the second when operating on an average machine (16 GB RAM, 8 cores, 2,90 GHz CPU). The time elapsed for the paragraph-by-paragraph processing method was greatly increased, ranging between 10x and 50x more for each iteration and thus requiring several minutes instead of seconds. While the reason behind this issue requires more investigation, with the current implementation,

performing such a method on a large-scale dataset without some optimization or access to a powerful machine is not recommended.

5 Conclusions, limitations and future work

In this work, the similarities between simplified and summarized text have been analysed through the automated summarization of articles from the Newsela corpus, performed with two different methods and compared to four levels of professional manual simplification representative of diverse school grade levels. By examining the results obtained by a ROUGE-L scoring comparison between our output and the manual standard, it is shown that the proposed paragraph-by-paragraph method is superior to a document-wide approach, with the highest score being 0.654. Hence, it is possible to claim that while automated summarization does not produce text similar enough to simplified documents to justify its substitution, it still produces text similar enough to be used as a baseline to perform simplification on - instead of starting from the original text.

However, there are important limitations to the currently chosen metric. As ROUGE-L cannot measure semantic similarity between instances, all sequences that are semantically correct but lexically different would not compute as "similar". Since abstractive summarization could generate text that is lexically different from the simplification golden standard but still effectively simplified, further analysis with semantically relevant metrics should be conducted. In addition, future work in this direction should implement ulterior thorough analyses with more refined metrics, such as ROUGE-SEM or SARI, along with a comparison between manual simplification, automated summarization and automated simplification algorithms. In particular, the latter could shed some light on the intrinsic similarities between simplification and summarization and help further investigate the potential interdisciplinary approaches to the text simplification field of research.

Further investigation into optimization procedures to make the most-performing methods available for lower-end machines should also be conducted to allow for wider access to the tools and improved effectiveness of summarizers as a simplification helping tool.

6 Acknowledgments

This research has been supported by the EKT55 project "Teksti lihtsustamine eesti keeles".

References

- David Allen. 2009. A study of the role of relative clauses in the simplification of news texts for learners of english. *System*, 37:585–599.
- Ping Chen, John Rochford, David Kennedy, Soussan Djamasbi, Peter Fay, and Will Scott. 2017. Automatic text simplification for people with intellectual disabilities. In *Artificial Intelligence Science and Technology*, pages 725–731.
- Mahak Gambhir and Vishal Gupta. 2017. Recent automatic text summarization techniques: a survey. *Artificial Intelligence Review*, 47.
- Natalia Grabar and Horacio Saggion. 2022. Evaluation of Automatic Text Simplification: Where are we now, where should we go from here. In *Traitement Automatique des Langues Naturelles*, pages 453–463, Avignon, France. ATALA.
- Max Grusky. 2023. Rogue scores. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1914–1934, Toronto, Canada. Association for Computational Linguistics.
- Vishal Gupta and Gurpreet Lehal. 2010. A survey of text summarization extractive techniques. *Journal of Emerging Technologies in Web Intelligence*, 2.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *ArXiv*, abs/1506.03340.
- Colby Horn, Cathryn Manduca, and David Kauchak. 2014. Learning a lexical simplifier using Wikipedia. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 458–463, Baltimore, Maryland. Association for Computational Linguistics.
- Z. Izakian and M. Mesgari. 2015. Fuzzy clustering of time series data: A particle swarm optimization approach. *Journal of AI and Data Mining*, 3(1):39–46.
- Bilel Jarraya and Abdelfatteh Bouri. 2012. Meta-heuristic optimization backgrounds: A literature review. *International Journal of Contemporary Business Studies*, 3:31–44.
- Karel Jeek and Josef Steinberger. 2008. Automatic text summarization (the state of the art 2007 and new challenges).

- Julian Kupiec, Jan Pedersen, and Francine Chen. 1995. A trainable document summarizer. In *Proceedings of the 18th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '95, page 68–73, New York, NY, USA. Association for Computing Machinery.
- Vladimir I. Levenshtein. 1965. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet physics. Doklady*, 10:707–710.
- 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 the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Chin-Yew Lin and FJ Och. 2004a. Looking for a few good metrics: Rouge and its evaluation. In *Ntcir workshop*.
- Chin-Yew Lin and Franz Josef Och. 2004b. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, pages 605–612, Barcelona, Spain.
- Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. 2022. Brio: Bringing order to abstractive summarization.
- Hamza Shabbir Moiyadi, Harsh Desai, Dhairya Pawar, Geet Agrawal, and Nilesh M. Patil. 2016. Nlp based text summarization using semantic analysis. *International Journal of Advanced Engineering, Management and Science*, 2(10).
- 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.
- N. Nazari and M. A. Mahdavi. 2019. A survey on automatic text summarization. *Journal of AI and Data Mining*, 7(1):121–135.
- Gustavo Paetzold and Lucia Specia. 2016. Benchmarking lexical simplification systems. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3074–3080, Portorož, Slovenia. European Language Resources Association (ELRA).
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Annual Meeting of the Association for Computational Linguistics*.
- David Pellow and Maxine Eskenazi. 2014. An open corpus of everyday documents for simplification tasks. In *Proceedings of the 3rd Workshop on Predicting and Improving Text Readability for Target Reader Populations (PITR)*, pages 84–93, Gothenburg, Sweden. Association for Computational Linguistics.
- Luz Rello, Ricardo Baeza-Yates, and Horacio Saggion. 2013. The impact of lexical simplification by verbal paraphrases for people with and without dyslexia. In *Proceedings of the 14th International Conference on Computational Linguistics and Intelligent Text Processing - Volume 2, CICLing'13*, page 501–512, Berlin, Heidelberg. Springer-Verlag.
- Advait Siddharthan. 2014. A survey of research on text simplification. *ITL - International Journal of Applied Linguistics*, 165:259–298.
- Matthew Snover, Nitin Madnani, Bonnie Dorr, and Richard Schwartz. 2009. Ter-plus: paraphrase, semantic, and alignment enhancements to translation edit rate. *Machine Translation*, 23:117–127.
- Sanja Štajner and Maja Popovic. 2016. Can text simplification help machine translation? In *Proceedings of the 19th Annual Conference of the European Association for Machine Translation*, pages 230–242.
- Sara Stymne, Jörg Tiedemann, Christian Hardmeier, and Joakim Nivre. 2013. Statistical machine translation with readability constraints. In *Proceedings of the 19th Nordic Conference of Computational Linguistics (NODALIDA 2013)*, pages 375–386. Linköping University Electronic Press, Sweden. 19th Nordic Conference of Computational Linguistics, NODALIDA 2013 ; Conference date: 22-05-2013 Through 24-05-2013.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks.
- Ciprian-Octavian Truică, Andrei-Ionuț Stan, and Elena-Simona Apostol. 2023. Simplex: a lexical text simplification architecture. *Neural Computing and Applications*, 35(8):6265–6280.
- Hoang Van, Zheng Tang, and Mihai Surdeanu. 2021. How may I help you? using neural text simplification to improve downstream NLP tasks. *CoRR*, abs/2109.04604.
- Tu Thanh Vu, Giang Binh Tran, and Son Bao Pham. 2014. Learning to simplify children stories with limited data. In *Intelligent Information and Database Systems*, pages 31–41, Cham. Springer International Publishing.
- William Watanabe, Arnaldo Junior, Vinícius Uzêda, Renata Fortes, Thiago Pardo, and Sandra Aluisio. 2009. Facilita: Reading assistance for low-literacy readers. pages 29–36.

- Chih-Hsuan Wei, Robert Leaman, and Zhiyong lu. 2014. Simconcept: A hybrid approach for simplifying composite named entities in biomedical text. volume 19.
- Wei Xu, Chris Callison-Burch, and Courtney Napoles. 2015. Problems in current text simplification research: New data can help. *Transactions of the Association for Computational Linguistics*, 3:283–297.
- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics*, 4:401–415.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. PEGASUS: Pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11328–11339. PMLR.
- Ming Zhang, Chengzhang Li, Meilin Wan, Xuejun Zhang, and Qingwei Zhao. 2024. Rouge-sem: Better evaluation of summarization using rouge combined with semantics. *Expert Systems with Applications*, 237:121364.
- Xingxing Zhang and Mirella Lapata. 2017. Sentence simplification with deep reinforcement learning. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 584–594, Copenhagen, Denmark. Association for Computational Linguistics.
- Zhemin Zhu, Delphine Bernhard, and Iryna Gurevych. 2010. A monolingual tree-based translation model for sentence simplification. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 1353–1361, Beijing, China. Coling 2010 Organizing Committee.