

MINWIKISPLIT: A Sentence Splitting Corpus with Minimal Propositions

Christina Niklaus¹³, André Freitas², and Siegfried Handschuh¹³

¹ University of St.Gallen

{christina.niklaus, siegfried.handschuh}@unisg.ch

² University of Manchester

andre.freitas@manchester.ac.uk

³ University of Passau

{christina.niklaus, siegfried.handschuh}@uni-passau.de

Abstract

We compiled a new sentence splitting corpus that is composed of 203K pairs of aligned complex source and simplified target sentences. Contrary to previously proposed text simplification corpora, which contain only a small number of split examples, we present a dataset where each input sentence is broken down into a set of minimal propositions, i.e. a sequence of sound, self-contained utterances with each of them presenting a minimal semantic unit that cannot be further decomposed into meaningful propositions. This corpus is useful for developing sentence splitting approaches that learn how to transform sentences with a complex linguistic structure into a fine-grained representation of short sentences that present a simple and more regular structure which is easier to process for downstream applications and thus facilitates and improves their performance.

1 Introduction

Sentences that present a complex linguistic structure can be *hard to comprehend by human readers*, as well as *difficult to analyze by semantic applications* (Saggion, 2017). Identifying such grammatical complexities in a sentence and transforming them into simpler structures, using a set of text-to-text rewriting operations, is the goal of syntactic text simplification (TS). One of the major types of operations that are used to perform this rewriting step is *sentence splitting*: it divides a sentence into several shorter components, with each of them presenting a simpler and more regular structure that is easier to process by both humans and machines (see Table 1).

Syntactic TS with a focus on the task of sentence splitting has been attracting growing interest in the natural language processing (NLP) community within the past few years. One line of work

targets reader populations with reading difficulties, such as people suffering from dyslexia, aphasia or deafness (Siddharthan and Mandya, 2014; Saggion et al., 2015; Ferrés et al., 2016), while the second line of work aims at generating an intermediate representation that is easier to process for downstream semantic tasks whose predictive quality deteriorates with sentence length and complexity. Prior work has established that applying syntactic TS as a preprocessing step can improve the performance of a variety of applications, including Machine Translation (Štajner and Popovic, 2016, 2018), Open Information Extraction (Cetto et al., 2018), or Text Summarization (Siddharthan et al., 2004; Bouayad-Agha et al., 2009).

2 Limitations of Existing Sentence Splitting Corpora

All of the TS approaches mentioned above make use of a set of *hand-crafted transformation rules* to decompose complex sentences into a sequence of structurally simplified components, requiring a complex rule engineering process. To overcome this expensive manual effort, Narayan et al. (2017) presented a first attempt at modelling a data-driven sentence splitting approach where simplification rewrites are learned automatically from examples of aligned complex source and simplified target sentences. Since previously compiled TS corpora (PWKP (Zhu et al., 2010), EW-SEW (Coster and Kauchak, 2011), and Newsela (Xu et al., 2015)) contain only a small number of split examples, they are ill-suited for learning to decompose sentences into shorter, syntactically simplified components. Therefore, Narayan et al. (2017) gathered a new dataset, WEBSPLIT, which is the first TS corpus that explicitly addresses the task of sentence splitting, while abstracting away from deletion-based and lexical simplification opera-

Complex source	The house was once part of a plantation and it was the home of Josiah Henson, a slave who escaped to Canada in 1830 and wrote the story of his life.
MINWIKISPLIT	The house was once part of a plantation. It was the home of Josiah Henson. Josiah Henson was a slave. This slave escaped to Canada. This was in 1830. This slave wrote the story of his life.
Complex source	Gary Goddard, founder of Gary Goddard Entertainment, a company that designs theme parks, attractions and upscale resorts, estimated that about half his work in the last few years has been in Asia and the Middle East.
MINWIKISPLIT	About half his work in the last few years has been in Asia. This was what Goddard estimated. About half his work in the last few years has been in the Middle East. This was what Gary Goddard estimated. Gary Goddard was founder of Gary Goddard Entertainment. Gary Goddard Entertainment was a company. This company designs theme parks. This company designs attractions. This company designs upscale resorts.
Complex source	The film is a partly fictionalized presentation of the tragedy that occurred in Kasaragod District of Kerala in India, as a result of endosulfan, a pesticide used on cashew plantations owned by the government.
MINWIKISPLIT	The film is a partly fictionalized presentation of the tragedy. This tragedy occurred in Kasaragod District of Kerala in India. This was as a result of endosulfan. Endosulfan is a pesticide. This pesticide is used on cashew plantations. These cashew plantations are owned by the government.

Table 1: Example instances from MINWIKISPLIT. A complex source sentence is broken down into a set of syntactically simplified sentences by decomposing clausal and phrasal elements and transforming them into self-contained propositions that present a simple structure in the form of a sequence of subject, predicate, and optionally an object, adverbial or complement (or simple combination thereof).

tions. It is composed of over one million tuples that map a single complex sentence to a sequence of structurally simplified sentences.

Aharoni and Goldberg (2018) criticized the data split proposed by Narayan et al. (2017). They observed that 99% of the simple sentences (which make up for more than 89% of the unique ones) contained in the validation and test sets also appear in the training set. Consequently, instead of learning to split and rephrase complex sentences, models that are trained on this dataset will be prone to learn to memorize entity-fact pairs. Hence, this split is not well suited for measuring a model’s ability to generalize to unseen input sentences. To fix this issue, Aharoni and Goldberg (2018) present a new train-development-test data split where nearly no simple sentence that is contained in the development or test set occurs verbatim in the training set.

Lately, Botha et al. (2018) discovered that the sentences from the WEBSPLIT corpus contain fairly *unnatural linguistic expressions* over only a *small vocabulary* and a rather *uniform sentence structure*, which is predominantly composed of a

sequence of coordinate clauses, occasionally augmented with a relative or adverbial clause (see Table 2). To overcome these limitations, they present WIKISPLIT, a dataset of one million sentences that were mined from Wikipedia edit histories. This corpus provides a rich and varied vocabulary over naturally expressed sentences showing a diverse linguistic structure, and their extracted splits. However, there is only a single split per source sentence in the training set (see Table 3). Accordingly, when a model is trained on this dataset, it is susceptible to exhibiting a *strong conservatism*, splitting each input sentence into exactly two output sentences only. Consequently, the resulting simplified sentences are still comparatively long and complex, mixing multiple, potentially semantically unrelated propositions that are difficult to analyze for downstream tasks.

3 MINWIKISPLIT Corpus

We improve on previously compiled sentence splitting corpora and present MINWIKISPLIT,¹ a

¹The MINWIKISPLIT dataset is publicly released under <https://github.com/Lambda-3/MinWikiSplit>.

- (1) A Loyal Character Dancer was published by Soho Press, in the United States, where some Native Americans live.
- (2) Dead Man’s Plack is in England and one of the ethnic groups found in England is the British Arabs.

Table 2: Characteristic example source sentences from the WEBSPLIT corpus.

Complex source	Starring Meryl Streep, Bruce Willis, Goldie Hawn and Isabella Rossellini, the film focuses on a childish pair of rivals who drink a magic potion that promises eternal youth.
Simplified output	Starring Meryl Streep, Bruce Willis, Goldie Hawn and Isabella Rossellini. The film focuses on a childish pair of rivals who drink a magic potion that promises eternal youth.
Complex source	The Assistant Attorney in Orlando investigated the modeling company, and decided that they were not doing anything wrong, and after Pearlman’s bankruptcy, the company emerged unscathed and was sold to a Canadian company.
Simplified output	The Assistant Attorney in Orlando investigated the modeling company, and decided that they were not doing anything wrong. After Pearlman’s bankruptcy, the modeling company emerged unscathed and was sold to a Canadian company.

Table 3: Pairs of aligned complex source and simplified target sentences from WIKISPLIT.

new dataset that can be used to train models for the task of decomposing sentences with a complex linguistic structure into a simplified representation that presents a more regular structure which is easier to process for downstream semantic applications and may support a faster generalization in machine learning tasks. This output may serve as an intermediate representation to facilitate and improve the performance of a wide range of artificial intelligence (AI) tasks.

Since shorter sentences are generally better processed by NLP systems (Narayan et al., 2017), we aimed at gathering a corpus where **each complex source sentence is broken down into a set of minimal propositions**, i.e. a sequence of sound, self-contained utterances, with each of them presenting a minimal semantic unit that cannot be further decomposed into meaningful propositions (Bast and Haussmann, 2013). Thus, we augment the Split-and-Rephrase task that was originally defined in Narayan et al. (2017) by the notion of *minimality*. In that way, we intend to overcome the conservatism exhibited by state-of-the-art structural TS approaches, which tend to retain the input rather than transforming it, and expect to improve the performance of a wide range of AI tasks.

4 Corpus Construction

MINWIKISPLIT is a large-scale sentence splitting corpus consisting of 203K complex source sentences and their simplified counterparts in the form of a sequence of minimal propositions. It was created by running DISSIM (Niklaus et al., 2019), a syntactic TS framework, over the one million complex input sentences from the WIKISPLIT corpus. DISSIM applies a small set of 35 hand-written transformation rules to decompose a wide range of linguistic constructs, including both clausal components (coordinations, adverbial clauses, relative clauses and reported speech) and phrasal elements (appositions, prepositional phrases, adverbial/adjectival phrases and coordinate noun phrases). In that way, a fine-grained output in the form of a sequence of minimal, self-contained propositions is produced. Some example instances are shown in Table 1.

To ensure that the resulting dataset is of high quality, we defined a set of dependency parse and part of speech based heuristics to filter out sequences that contain *grammatically incorrect* sentences, as well as sentences that *mix multiple semantic units* and, thus, are violating the specified minimality requirement. For instance, in order to verify that the simplified sentences are grammatically sound, we check whether the root node of the output sentence is a verb and whether one of its child nodes is assigned a dependency label that denotes a subject component. To test if the simplified sentences represent minimal propositions, we check whether the output does not contain a clausal modifier, such as a relative clause modifier,

	#T/S	#S/C	%SAME	LD _{SC}	SAMSA	SAMSA _{abl}
Complex	30.75	1.18	100	0.00	0.36	0.94
MINWIKISPLIT	12.12	3.84	0.00	17.73	0.40	0.48

Table 4: Results of the automatic evaluation procedure on a random sample of 1000 sentences.

adverbial clause modifier or a clausal modifier of a noun. Moreover, we ensure that no conjunction is included in the simplified output sentences. In the future, we will implement some further heuristics to *avoid uniformity* in the structure of the source sentences. In that way, we aim at guaranteeing a great structural variability in the input in order to enable systems that are trained on the MINWIKISPLIT corpus to learn splitting rewrites for a wide range of linguistic constructs.

After running the sentence simplification framework DISSIM over the sentences from the WIKISPLIT corpus and applying the set of heuristics that we defined to ensure grammaticality and minimality of the output, 203K pairs of input and corresponding output sequences were left.

5 Experiments

We performed both a manual analysis and an automatic evaluation to assess the quality of the produced corpus.

5.1 Automatic Metrics

To estimate the quality of the simplified target sentences of the MINWIKISPLIT corpus, we computed some basic statistics, including (i) the average sentence length of the simplified sentences in terms of the average number of tokens per output sentence (#T/S); (ii) the average number of simplified output sentences per complex input (#S/C); (iii) the percentage of sentences that are copied from the source without performing any simplification operation (%SAME), serving as an indicator for conservatism, i.e. the tendency to retain the input rather than transforming it; and (iv) the averaged word-based Levenshtein distance from the input (LD_{SC}), which provides further evidence for how reluctant the underlying system is in splitting the input into minimal semantic units.

Moreover, to measure the structural simplicity of the instances contained in MINWIKISPLIT, we calculated the SAMSA and SAMSA_{abl} scores of both the complex source and the simplified output sentences (Sulem et al., 2018b). They are the first metrics that explicitly target syntactic aspects

of TS. The SAMSA metric is based on the idea that an optimal split of the input is one where each predicate-argument structure is assigned its own sentence in the simplified output and measures to what extent this assertion holds for the input-output pair under consideration. Accordingly, the SAMSA score is maximized when each split sentence represents exactly one semantic unit in the input. SAMSA_{abl} does not penalize cases where the number of sentences in the simplified output is lower than the number of events contained in the input, indicating separate semantic units that should be split into individual target sentences for obtaining minimal propositions.²

These computations were carried out on a random sample of 1000 sentences from MINWIKISPLIT. The results are provided in Table 4. The scores demonstrate that on average our proposed sentence splitting corpus contains four simplified target sentences per complex source sentence, with every target proposition consisting of 12 tokens. Moreover, no input is simply copied to the output, but split into smaller components. Both the high averaged Levenshtein distance of almost 18 and the SAMSA score (0.40) confirm previous findings. The latter is highly correlated with structural simplicity and grammaticality, indicating that the output sentences contained in our corpus are grammatically sound and present a simpler syntax than the input. With 0.48, we reach a decent score for the simplified target sentences with regard to SAMSA_{abl}, too, which has a high correlation with meaning preservation.

5.2 Manual Analysis

In a second step, we randomly selected a subset of 300 sentences from MINWIKISPLIT, on which we conducted a manual analysis in order to get some deeper insights into the quality of the simplified sentences. Each input-output pair was rated by

²Prior work on syntactic TS commonly also reports average BLEU (Papineni et al., 2002) scores. However, Sulem et al. (2018a) recently demonstrated that this score is inappropriate for the evaluation of TS approaches when sentence splitting is involved. Therefore, we refrain from calculating BLEU scores.

2 non-native, but fluent English speakers according to three parameters: grammaticality, meaning preservation and structural simplicity (see Table 5).

G	Is the output fluent and grammatical?
M	Does the output preserve the meaning of the input?
S	Is the output simpler than the input, ignoring the complexity of the words?

Table 5: Questions for the human evaluation.

The inter-annotator agreement was computed using Cohens quadratic weighted κ (Cohen, 1968). The obtained rates were 0.24, 0.25 and 0.75 for grammaticality, meaning preservation and structural simplicity, respectively. System scores were calculated by averaging over the annotators’ scores and the 300 sentences.

G	M	S
4.36	4.10	3.43

Table 6: Averaged human evaluation ratings on a random sample of 300 sentences from MINWIKISPLIT. Grammaticality (G), meaning preservation (M) and structural simplicity (S) are measured using a 1 (very bad) to 5 (very good) scale.

The results of the human evaluation are displayed in Table 6. These scores show that we succeed in producing output sequences that reach a high level of grammatical soundness and almost always perfectly preserve the original meaning of the input. The third dimension under consideration, structural simplicity, which captures the degree of minimality in the simplified sentences, scores high values, too. However, our manual analysis revealed some room for improvement. Consequently, in the future, we plan to implement stricter heuristics for sorting out output sequences that still mix multiple semantically unrelated propositions.

6 Conclusion

We compiled MINWIKISPLIT, a sentence splitting corpus consisting of 203K complex source sentences and their split counterparts. This dataset can be used to train natural language generation applications that perform a syntactic TS, simplifying sentences with a complex linguistic structure into a fine-grained representation of short

sentences that present a simple and more regular structure. The thus generated output may serve as an intermediate representation that is easier to process for downstream semantic applications and may thus lead to a better performance of those tools. We intend to train a sentence simplification model on MINWIKISPLIT and compare it to previously proposed systems trained on the WEBSPLIT and WIKISPLIT corpora.

Moreover, we plan to improve the quality of the simplified target sentences in our corpus in accordance with the insights we gained through the analyses described above. First of all, we will perform a detailed error analysis of the output to determine the most common types of mistakes and get some starting points for further improving our heuristics for filtering out malformed simplifications. To enhance the syntactic correctness of the output, we will train a classifier on the recently proposed CoLA dataset (Warstadt et al., 2018) to eliminate instances with ungrammatical target sentences from our corpus. In addition, special attention will be given to improving the heuristics that ensure that each simplified target sentence represents a single semantic unit.

References

- Roei Aharoni and Yoav Goldberg. 2018. [Split and rephrase: Better evaluation and stronger baselines](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 719–724. Association for Computational Linguistics.
- Hannah Bast and Elmar Haussmann. 2013. Open information extraction via contextual sentence decomposition. In *2013 IEEE Seventh International Conference on Semantic Computing*, pages 154–159. IEEE.
- Jan A. Botha, Manaal Faruqui, John Alex, Jason Baldridge, and Dipanjan Das. 2018. [Learning to split and rephrase from wikipedia edit history](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 732–737. Association for Computational Linguistics.
- Nadjet Bouayad-Agha, Gerard Casamayor, Gabriela Ferraro, Simon Mille, Vanesa Vidal, and Leo Wanner. 2009. Improving the comprehension of legal documentation: the case of patent claims. In *Proceedings of the 12th International Conference on Artificial Intelligence and Law*, pages 78–87. ACM.
- Matthias Cetto, Christina Niklaus, André Freitas, and Siegfried Handschuh. 2018. [Graphene:](#)

- Semantically-linked propositions in open information extraction. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2300–2311. Association for Computational Linguistics.
- Jacob Cohen. 1968. Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological bulletin*, 70(4):213.
- William Coster and David Kauchak. 2011. Simple english wikipedia: A new text simplification task. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Short Papers - Volume 2*, HLT '11, pages 665–669, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Daniel Ferrés, Montserrat Marimon, Horacio Saggion, and Ahmed AbuRa'ed. 2016. Yats: Yet another text simplifier. In *Natural Language Processing and Information Systems*, pages 335–342, Cham. Springer International Publishing.
- Shashi Narayan, Claire Gardent, Shay B. Cohen, and Anastasia Shimorina. 2017. Split and rephrase. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 606–616. Association for Computational Linguistics.
- Christina Niklaus, Matthias Cetto, André Freitas, and Siegfried Handschuh. 2019. [Transforming complex sentences into a semantic hierarchy](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3415–3427, Florence, Italy. 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*, pages 311–318. Association for Computational Linguistics.
- Horacio Saggion. 2017. Automatic text simplification. In *Synthesis Lectures on Human Language Technologies*.
- Horacio Saggion, Sanja Štajner, Stefan Bott, Simon Mille, Luz Rello, and Biljana Drndarevic. 2015. [Making it simplext: Implementation and evaluation of a text simplification system for spanish](#). *ACM Trans. Access. Comput.*, 6(4):14:1–14:36.
- Advait Siddharthan and Angrosh Mandya. 2014. [Hybrid text simplification using synchronous dependency grammars with hand-written and automatically harvested rules](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 722–731. Association for Computational Linguistics.
- Advait Siddharthan, Ani Nenkova, and Kathleen McKeown. 2004. Syntactic simplification for improving content selection in multi-document summarization. In *Proceedings of the 20th international conference on Computational Linguistics*, page 896. Association for Computational Linguistics.
- 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.
- Sanja Štajner and Maja Popovic. 2018. Improving machine translation of english relative clauses with automatic text simplification. In *Proceedings of the First Workshop on Automatic Text Adaptation (ATA)*.
- Elior Sulem, Omri Abend, and Ari Rappoport. 2018a. [Bleu is not suitable for the evaluation of text simplification](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 738–744. Association for Computational Linguistics.
- Elior Sulem, Omri Abend, and Ari Rappoport. 2018b. [Semantic structural evaluation for text simplification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 685–696. Association for Computational Linguistics.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2018. Neural network acceptability judgments. *arXiv preprint arXiv:1805.12471*.
- 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.
- Zheming 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*, pages 1353–1361. Association for Computational Linguistics.