

German Text Simplification: Finetuning Large Language Models with Semi-Synthetic Data

Lars Klöser and Mika Beele and Jan-Niklas Schagen and Bodo Kraft

University of Applied Sciences Aachen
{kloeser, beele, kraft}@fh-aachen.de
niklas.schagen@alumni.fh-aachen.de

Abstract

This study pioneers the use of synthetically generated data for training generative models in document-level text simplification of German texts. We demonstrate the effectiveness of our approach with real-world online texts. Addressing the challenge of data scarcity in language simplification, we crawled professionally simplified German texts and synthesized a corpus using GPT-4. We finetune *Large Language Models* with up to 13 billion parameters on this data and evaluate their performance. This paper employs various methodologies for evaluation and demonstrates the limitations of currently used rule-based metrics. Both automatic and manual evaluations reveal that our models can significantly simplify real-world online texts, indicating the potential of synthetic data in improving text simplification.

1 Introduction

In our modern and digitalized societies, access to information is essential for active participation. However, certain groups, such as individuals with intellectual disabilities or non-native speakers, often struggle to understand the local language, which can impede their social and civic engagement. Each group faces unique challenges in text comprehension. Integrating automatic text simplification tools can significantly benefit these groups by providing accessible information, thereby providing a pivotal means for greater inclusion.

Various linguistic initiatives, like (Netzwerk-Leichte-Sprache, 2022), have been established in German-speaking regions to address this, specifically designed for different target groups. Such efforts align with international legal frameworks like *Article 9 of the UN Convention on the Rights of Persons with Disabilities*¹, which advocates for the right to accessible communication.

¹<https://www.ohchr.org/en/instruments-mechanisms/instruments/convention-rights-persons-disabilities>

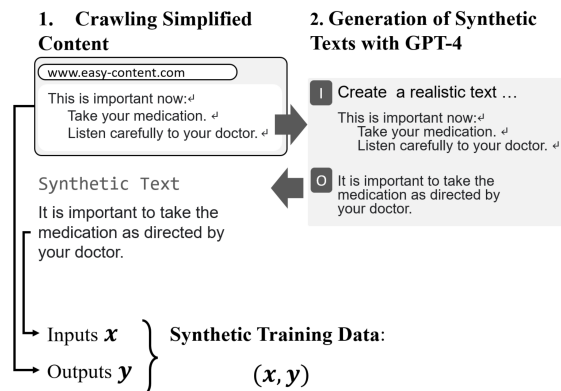


Figure 1: Illustration of synthetic data generation. Data is crawled from websites specializing in language simplification. GPT-4 generates texts in everyday language, ensuring the original content remains unaltered. We construct a simplification dataset where these texts serve as input while the crawled simplifications act as reference simplifications.

Creating simplified content manually is a labor-intensive and time-consuming process, significantly hindering its broad availability and accessibility. In contrast, *Large Language Models* (LLMs), especially smaller ones fine-tuned for text simplification, offer a viable and efficient alternative (Anschütz et al., 2023). These smaller models require fewer resources and are simpler to operate than larger LLMs, making them ideal for scaling up the process of automatic language simplification. A key challenge in finetuning LLMs for text simplification lies in the limited availability of parallel data (Anschütz et al., 2023; Toborek et al., 2022).

Our approach, as illustrated in Figure 1, tackles the challenge of data scarcity in language simplification by creating semi-synthetic data. This involves crawling various sources for already simplified web content and then utilizing GPT-4 to generate hypothetical original texts corresponding to these simplifications. Our dataset thus comprises GPT-4’s outputs as the inputs and the crawled con-

tent as the simplified outputs, forming a text simplification dataset. We apply this dataset as the basis for finetuning Large Language Models (LLMs) for automatic text simplification.

We publish all necessary resources to reproduce this paper’s results on a public GitHub repository². Our scientific contributions can be summarized as follows:

1. Creating a corpus of parallel text simplification data in German based on novel methodology.
2. Training, evaluating, and releasing LLM-based language simplification models for German texts.

2 Related Work

Various approaches and methodologies have been developed for automatic German text simplification. (Anschütz et al., 2023) proposed a two-step approach utilizing pretrained language models finetuned on German simplifications to diminish the requirement for parallel data. In contrast to our approach, the parallel data contains a mixture of summarization and simplification and targets only newspaper articles. (Spring and Rios, 2021) train German text simplification models by using labels to target specific language levels, ensuring model adaptations are level-appropriate and control copying behavior. Their dataset focuses on newspaper articles and sentence-level simplification.

Diverging from neural network-based methods, (Garain et al., 2019) introduced methodology based on parse trees. Similarly, (Praveen Kumar et al., 2022) offered a pattern-based syntactic simplification framework. (Kajiwara and Komachi, 2018) presented methods for text simplification in languages with limited simplified corpora, including lexical substitution and monolingual translation, focusing on resource-scarce languages like Japanese.

Regarding evaluation metrics, (Sulem et al., 2018) investigate the limitations of BLEU as a widespread evaluation metric for text generation tasks. (Alva-Manchego et al., 2021) explored the correlation between existing metrics and human judgments in multi-operation text simplifications, providing insights into the appropriateness of automatic metrics for assessing text simplification. (Maddela et al., 2022) introduced LENS, a

learnable evaluation metric for text simplification trained on modern language models, showing a better correlation with human judgment.

Regarding parallel corpora and resources, (Ebling et al., 2022) aggregated corpora for the automatic processing of simplified German, providing resources for training and evaluation. (Holmer and Rennes, 2023) create so-called *pseudo parallel* sentence pairs of simple and complex sentences from given sentence collections. (Hauser et al., 2022) introduced SNIML, a multilingual corpus of news articles in simplified language, and (Rios et al., 2021) showcased a dataset for document-level text simplification in German, including articles paired with simplified summaries. (Aumiller and Gertz, 2022) addressed the challenge of concurrently summarizing and simplifying longer texts, introducing a new dataset for joint text simplification and summarization. (Hewett, 2023) introduces a dataset with parallel sentence-level simplifications and additional information about the document’s rhetorical structure.

Text simplification corpora exist for various languages. For example, (Coster and Kauchak, 2011) introduced a dataset pairing English Wikipedia with Simple English Wikipedia, enabling the analysis of various simplification operations, including rewording, reordering, insertion, and deletion.

In Easy Language generation, (Deilen et al., 2023) investigated the feasibility of using ChatGPT to translate administrative texts into German *Leichte Sprache* (easy language), a highly regulated language variety with a focus on text simplification.

3 Task Definition

In this section, we give a task definition for language simplification. This definition outlines the requirements for the trained models and motivates our methodology for dataset creation.

The **inputs space** consists of editorially created German texts. Based on the selection of web sources, we assume predominantly grammatically complete sentences and quality-assured content. In the context of this work, we exclude user-created input and social media texts. These web documents contain multiple paragraphs and sentences.

The **output space** consists of simplifications of the input. The style and level of simplification correspond to the contents currently available in the German language, precisely as they are presently accessible. We aim to avoid modifications of the

²<https://github.com/MSLars/German-Text-Simplification>

Table 1: Word and document frequencies in the dataset across different sources, segregated into test and train sets

	Test		Train	
	Docs	Words	Docs	Words
<i>einfachstars</i>	317	45,444	2,213	307,307
<i>mdr</i>	10	1,696	85	13,285
<i>nachrichtenleicht</i>	298	44,138	2,147	318,069
<i>hurraki</i>	181	16,152	1,234	109,386
<i>ndr</i>	94	17,211	709	135,340
<i>kurier</i>	72	13,425	481	76,744
<i>leicht-kicken</i>	8	537	67	2,672
<i>einfach-teilhabe</i>	8	749	79	8,193
<i>stadt-koeln</i>	3	2,588	16	11,830
<i>inclusion_europe</i>	1	35	18	920
<i>bundesregierung</i>	5	1,337	22	9,212
<i>hamburg-de</i>	3	1,199	59	16,280
Σ	1000	144,511	7,130	1,009,238

content, like summarization. However, understandability may require additional explanations of certain concepts in the simplified texts.

Two concepts for language simplification have been established for the German language: *Einfache Sprache* (simple language) and *Leichte Sprache* (easy language). We seek to explain how our methodologies intersect and align with these well-established frameworks.

Simple language covers text simplification in general. Possible target groups include readers unfamiliar with the domain or used language, for example, in legal or medical texts or language learners. The target group can fundamentally understand the concepts. Linguistic complexity, however, makes understanding more difficult. Moreover, these simplifications can aid in language acquisition.

In contrast, easy language is aimed at people with severely limited text comprehension, such as those with intellectual disabilities. Fixed sets of simplification rules have been established (Netzwerk-Leichte-Sprache, 2022). These rules cover areas like syntactical, lexical, or typographical simplifications.

In our approach, we scrape texts from various sources, each characterized by its distinct language style. Our research hints that these target texts generally do not conform to the rules of easy language. Rather, many crawled texts may align more with the domain of simple language.

3.1 Dataset

We introduce a parallel corpus consisting of texts in everyday language and their corresponding simplifications as an instantiation of the task defined

in section 3.

We create semi-synthetic text pairs to overcome the challenge of training data scarcity. Based on the results of various benchmarks, we assume that GPT-4 can produce human-like texts in various domains (OpenAI, 2023). We crawl simplified texts by expanded versions of the crawlers used by (Anschütz et al., 2023). Our additional preprocessing standardizes the typography using rule-based methods. Subsequently, we use GPT-4 to create realistic synthetic source texts from the simplifications. To ensure the generated texts were sufficiently diverse, 15 distinct prompts were used. In the following, we will investigate these data in detail.

3.1.1 Synthetic Texts

Table 1 presents the scope and size of the semi-synthetic dataset created for German text simplification. A random sample of 1,000 examples has been reserved as test data. To our knowledge, this dataset is the first semi-synthetic approach to the German language simplification task. It is also noteworthy for being the most comprehensive dataset available for document-wide simplification and the only dataset focusing on document-level language simplification across various domains.

The performance of machine learning models is highly contingent on the quality of the training data, as indicated in various studies (Jain et al., 2020). We believe that the complexity and characteristics of the synthetic data used for training should closely mirror the real data from similar contexts in the specific field.

Table 2: Analysis of textual complexity across various domains. Synthetic web content is slightly more complex compared to real web content. The metrics support that the crawled simplifications are less complex than real and synthetic everyday web content.

	Sports	Celebrities	News
<i>Metric: avg. sentence length</i>			
Easy	10.38 ± 2	11.56 ± 2.45	10.79 ± 1.49
Synth.	24 ± 9.3	22.94 ± 6.59	21 ± 5.4
Com.	19.59 ± 3.65	21.17 ± 4.21	18.66 ± 2.7
<i>Metric: avg. commas per sentence</i>			
Easy	.09 ± .18	.17 ± .24	.00 ± .03
Synth.	1.69 ± 1.18	1.67 ± .81	1.52 ± .72
Com.	.48 ± .41	1.3 ± .52	.81 ± .35
<i>Metric: avg. distance verb compounds</i>			
Easy	.09 ± .15	.14 ± .14	.14 ± .11
Synth.	.34 ± .33	.36 ± .19	.34 ± .19
Com.	.27 ± .19	.3 ± .14	.3 ± .13

Table 2 offers an overview of various metrics

used to estimate the linguistic complexity of the crawled simplified texts, the synthetic data, and real German web content, examining three distinct domains as examples. The selected metrics rate the reconstructed texts slightly more complex than the crawled texts. While these metrics do not definitively determine whether the data is realistic and overcomplication in reconstruction cannot be ruled out, our preliminary conclusion is that the examples could be suitable for the task.

Given that the primary focus of this work is not on the realistic generation of web content, we do not delve deeper into these aspects. Instead, our research examines whether the trained models effectively reduce the complexity of real web content as evaluated in [subsection 5.3](#). This approach aligns with our goal to enhance the practical applicability of language simplification tools in real-world scenarios.

3.2 Crawled Simplifications

This section delves into the specifics of the crawled simplifications. [Table 2](#) categorizes various sources into domains. This offers a structured view of the different simplifications obtained from these domains.

Table 3: Comparative analysis between different styles of simplified news content.

Metrik	MDR	NDR	NL
Sentence length	12.39 ± 1.54	10.47 ± 1.43	12 ± 1.5
Commas per sentence	.04 ± .09	.00 ± .00	.22 ± .17
Distance verb compounds	.2 ± .12	.14 ± .11	.21 ± .14
Words per line	8.78 ± 22.12	8.55 ± 11.77	14.14 ± 24.15

[Table 3](#) investigates the variety inside a single domain. It comprehensively analyzes metrics related to simplified texts from various news providers. These metrics reveal notable differences in the style of simplifications among the providers. For instance, NDR’s texts stand out for their absence of commas, suggesting a preference for simpler sentence structures without subordinate clauses. In contrast, NL (nachrichtenleicht.de) frequently employs commas, indicating a higher likelihood of compound and complex sentences, often incorporating subordinate clauses. Additionally, NL’s texts, on average, contain longer sentences than other sources, highlighting a distinct approach to

text simplification. These findings underscore the stylistic diversity within the dataset, demonstrating that simplifications are not uniform but vary significantly across news providers.

4 Methodology

This study aims to perform task-specific fine-tuning of LLMs. The extensive volume of publications in this domain makes it impractical to evaluate all available models and training configurations against one another. Instead, we aim to justify our core design choices in this chapter and provide further justification through targeted evaluation in the subsequent chapter.

4.1 Language Modelling

We apply LLMs based on so-called *decoder-only* transformer models as introduced in ([Radford et al., 2018](#)). Decoder-only models are designed to model the probability of the subsequent token $P(x_{i+1}|x_1, \dots, x_i)$ in a given sequence $x = x_1, \dots, x_n$ of tokens that represent a text. We represent each text simplification sample as a sequence $x = (x_{source}, SEP, x_{target})$ where *SEP* is a special token that separates the source from its simplification target. As input, we provide x_{source} followed by the *SEP* token. The model, during training, attempts to maximize the probabilities

$$P(x_{i+1}^t | x_{source}, x_1^t, \dots, x_i^t), i = 1, \dots, m$$

of the tokens in the simplifications $x_{target} = x_1^t, \dots, x_m^t$ using the cross entropy loss.

We finetuned two distinct versions of two different pretrained LLMs. Specifically, we finetuned two German versions of GPT-2 ([Minixhofer et al., 2022](#)) and two versions of the Leo LM model ([Plüster, 2023](#)). We selected these models because they are decoder-only, available in different sizes and pretrained on German texts. Each of these models underwent finetuning on our training dataset for three epochs. For this process, we employed the HuggingFace³ library, a popular choice for machine learning and natural language processing tasks. The detailed configurations used for training, including parameters and environmental settings, are meticulously documented in [Appendix A](#).

³<https://huggingface.co/>

4.2 Decoding Algorithm

This section describes the methods of deriving concrete sequences from the probability distributions for individual follow-up tokens provided by LLMs, commonly called *decoding algorithms*. We compare four distinct approaches:

Greedy Approach: This method sequentially selects the token with the highest probability. It is straightforward and efficient but may not yield the most contextually appropriate sequence.

Beam Search Algorithm: This technique chooses the best alternative from a fixed number of possibilities, each with the currently highest probability. It balances between exploring various possibilities and focusing on the most probable options.

Sampling-Based Algorithm: Here, follow-up tokens are randomly selected based on the probability distribution of the LLMs. This approach introduces variability and can generate more diverse outputs (Holtzman et al., 2020).

Contrastive Search Approach: This novel method contrasts traditional search techniques. It considers the likelihood of individual tokens and evaluates the probability distribution over a set of potential sequences, aiming to balance between the most probable and contextually appropriate choices. This approach is useful in ensuring that the generated text maintains coherence and relevance (Su and Collier, 2022).

We utilized a fixed configuration for each approach as provided in [Appendix A](#). This comparative analysis offers insights into the efficacy and suitability of different decoding strategies.

We frequently observed prediction repetitions in our investigation, particularly with smaller models. In text generation, in general, a repetition penalty is frequently used. However, in this context, some repetition may be beneficial. Hence, we’ve devised an alternative approach that allows for a certain degree of repetition, recognizing its potential value in making texts clearer and more comprehensible.

To address this, we implemented a strategy to halt the generation of further tokens if the frequency of a token within a certain window exceeded a predefined threshold. This intervention was designed to enhance the quality of the generated text by preventing excessive repetition, which can detract from the readability and coherence of the output. Such a method is crucial in maintaining language’s natural flow and diversity, especially in scenarios where smaller models may struggle.

5 Evaluation

In this chapter, we provide a comprehensive evaluation of the finetuned models. Our analysis is twofold: firstly, we assess the performance of various model configurations on the semi-synthetic dataset. This evaluation will delve into how different configurations perform in terms of efficacy, which will be measured using a range of metrics.

Secondly, we extend our evaluation to include an analysis of crawled web content. This is a vital step towards demonstrating the real-world applicability of our models.

5.1 Evaluation Metrics

For automatic evaluation, we apply three rule-based metrics commonly used to evaluate simplification models. Each metric compares reference simplifications with model predictions mostly based on n-gram overlaps. N-grams are contiguous sequences of words in a text.

BLEU (Papineni et al., 2002) computes precision scores that measure the frequency of distinct n-grams in the reference simplification over the frequency of distinct n-grams in the model prediction. Typically, we use precision scores for uni, bi, tri, and tetra-grams. These are aggregated with a geometric mean and combined with a brevity penalty for too short predictions.

METEOR (Banerjee and Lavie, 2005) is based on matching unigrams of the model’s prediction with unigrams of the reference. It calculates precision and a heavily weighted recall on these matches. Additionally, it includes a fragmentation penalty that penalizes predictions with limited sequential overlap with the reference.

SARI (Xu et al., 2016) is designed to evaluate sentence-level text simplification systems. The metric compares n-gram operations between input on the one side and reference and predicted output on the other. It computes F-scores over added and kept n-grams. For deleted n-grams, the precision score is considered. The final score is the arithmetic mean.

5.2 Automatic Evaluation Results

The comprehensive results from our automatic evaluation, as detailed in [Table 4](#), provide insights into the performance of two variants of pretrained language models across different generation algorithms. In our analysis, all configurations exhibited the highest score for SARI, followed by METEOR

Table 4: Rule-based evaluation metrics computed on the test set. Scores are grouped by pretrained language model and generation algorithm. Metrics increase with model size. For the largest model, beam search is the best decoding algorithm.

	BLEU	METEOR	SARI
<i>Model: gpt2-wechsel-german</i>			
greedy	0.72	9.14	36.61
beam	1.49	13.03	36.80
sampling	0.96	11.31	37.62
contrastive	0.83	10.07	37.02
<i>Model: gpt2-xl-wechsel-german</i>			
greedy	6.77	23.58	46.49
beam	8.21	23.80	45.41
sampling	8.35	26.86	47.48
contrastive	6.99	23.87	46.74
<i>Model: leo-hessianai-7b</i>			
greedy	24.46	45.31	60.51
beam	25.97	46.17	61.35
sampling	23.79	44.97	60.23
contrastive	24.39	45.20	60.45
<i>Model: leo-hessianai-13b</i>			
greedy	24.53	45.32	60.52
beam	25.78	45.64	62.24
sampling	23.93	45.06	60.41
contrastive	24.64	45.57	60.66

and BLEU. This phenomenon is explored in greater detail in subsection 5.2.1.

Metrics increased with an increasing number of model parameters. Notably, the improvement in metrics was evident up to the transition from models with 7 billion to 13 billion parameters, beyond which we observed no significant differences in metrics. We investigate the behavior of these metrics in more detail in subsection 5.2.1. However, the superior performance of the 13 billion parameter model in other tasks suggests that the combination of automated metrics and our dataset may not be capable of discerning performance differences (Plüster, 2023). This could potentially lead to erroneous model selection in practical applications.

In many language generation applications, maximization-based methods like beam search are often deemed less suitable due to their propensity for monotonous and repetitive predictions, as opposed to sampling or contrastive search (Su and Collier, 2022). However, our results do not confirm this for our instantiation of text simplification. The findings suggest that the efficacy of these methods may vary depending on the specific nature of the language generation task.

5.2.1 Limitations of Rule-Based Metrics Employed within Simplification

In Figure 2, different simplification styles stand out. Since our dataset only includes a single refer-

Reference Simplification:

Those who have aphasia,
have difficulties to find words.
Or difficulties to speak words.
Or difficulties to understand words.

Model Prediction:

People with aphasia have problems:

- to find words
- to speak words
- to understand words

The • "breaks" all tetra-grams

Figure 2: In this example, a reference simplification and a model prediction, translated into English, are contextually similar but lack any shared tetra-grams, yielding a BLEU score of zero.

ence translation and the metrics focus on sequential overlaps, the stylistic variety of the dataset is not adequately considered in this automated and rule-based evaluation. This leads to lower scores for all metrics we used since all measure the sequential overlap.




Figure 3: N-gram precisions for predictions of the leo-hessianai-7b model on the complete test set. We observed significantly sloping precision scores for increasing n-gram sizes

The pair of reference simplification and model prediction in Figure 2 highlights the BLEU metric's key limitations in evaluating text simplification on our test set. The example showcases varying styles between the reference and the model prediction. The reference employs grammatically complete sentences linked with the conjunction "Or", while the model prediction opts for a listing format. This stylistic divergence, especially with short sentences in the model's output, leads to a lack of common 4-grams. BLEU combines n-gram precision with a geometric mean. The geometric mean is calculated by multiplying all the precision values and then taking the n th root (where n is the number of values). If any value in the dataset is 0, the product of all the

values becomes 0 as well. Figure 3 illustrates the n-gram precision scores. This leads to a BLEU score 0, which doesn't accurately reflect the simplified text's quality. By looking at the graph of n-gram precision scores, we can apply this understanding to most of our dataset.

Reference Simplification

The climate crisis is a crisis.
 In this, the climate is changing rapidly.
 And therefore, there are many problems.
 And there are dangers for people and animals.

Model Prediction

The climate around the world is changing.
 This is happening more and more quickly.
 This is bad for nature.
 This is called the climate crisis.

Figure 4: Example to illustrate a high fragmentation penalty due to varied placement of 'climate crisis', negatively impacting the METEOR Score.

In the given example in Figure 4, the reference simplification introduces the term *climate crisis* in the first sentence, whereas the model's simplification describes aspects of the climate crisis before introducing the term. Such rearrangements lead to cross-alignments and higher fragmentation penalties, which affect the METEOR score. This illustrates how n-gram intersections influence the metrics, particularly in the context of differing simplification styles within our corpus. METEOR caps the fragmentation penalty at 50 percent, which limits its influence on the final score.

As indicated in our analysis and shown in Table 4, the SARI score tends to rate the model solutions more favorably, aligning more closely with the positive manual evaluations of the models in Table 5. This suggests that SARI might be a more reliable indicator of text simplification quality in our context.

SARI was originally meant to be applied within sentence-level simplification. We apply SARI to our multi-sentence documents and we conducted further investigations on the composition of the SARI scores to conclude the metric's plausibility in its three categories, *add*, *keep*, and *delete* within our task.

In sentence-level text simplification, those operations are considered equally difficult (Xu et al., 2016), and therefore, they are weighted equally in the final arithmetic mean. However, due to stylistic transformations within simplification, n-grams of

the input are deleted at a much higher probability. For example, structural unigrams like bullet points and line breaks discontinuing original n-grams.




Figure 5: N-gram deletion precisions for predictions of the leo-hessianai-7b model on the complete test set. The median values of our observed SARI delete precision scores reach high values, especially for tri- and tetra-grams

Most n-grams from the input, especially tri-grams and tetra-grams, are considered as *correctly deleted* for reference simplification and model prediction. This results in a precision score for deleted tri-grams and tetra-grams close to one for most samples, as illustrated in Figure 5. Due to the arithmetic mean, this has a nearly constant and strong influence on the final value. SARI tends to be biased optimistically within our task.

Concerning the automated metrics, the 13 billion parameter model does not outperform the 7 billion one. One reason might be the stylistic diversity in the dataset. Given varied styles within even single domains, see Table 3, the model simplification and the reference simplification might be in different styles. This random factor may affect all metrics that measure sequential overlaps, limiting the overall scores. We might not be able to measure strong models abilities on the target task.

5.3 Evaluation on Real-World Data

We aimed to prove the practical relevance of our models by expanding our evaluation to include real-world data. As evaluation data, we use texts from a German news website⁴, a sports news website⁵, and a website for tabloid news⁶. These texts are not simplified.

We do not have reference simplifications for these texts. To evaluate the models' simplifications, we consider two types of metrics. Firstly, linguistic

⁴ tagesschau.de

⁵ transfermarkt.de

⁶ vip.de

simplification using the metrics already introduced. Secondly, the content similarity is done through a manual evaluation of 135 pairs of crawled and non-simplified texts with the simplifications of the models. In this process, pairs could be rated with 0 (no agreement), 1 (partial agreement), 2 (substantial agreement), and 3 (complete agreement). The results in Table 5 measure the language simplification capabilities of gpt2-xl-wechsel-german (gpt2-xl) and leo-hessianai-7b (leo-7b).

Table 5: Language complexity and content similarity metrics for model simplifications of real-world online data. *Human Evaluation* summarizes a manual evaluation of content similarity with scores from 0 (no similarity) to 3 (complete equality).

<i>Metrik</i>	gpt2-xl	leo-7b
<i>Sentence length</i>	16.35 ± 6.05	14.03 ± 3.47
<i>Commas per sentence</i>	.48 ± .68	.24 ± .28
<i>Words per line</i>	12.35 ± 6.08	10.35 ± 3.48
<i>Human Evaluation</i>	1.34 ± 1.11	2.68 ± 0.55

The complexity metrics sentence length, commas per sentence, and words per line indicate that gpt2-xl simplifies texts less than leo-7b. Regarding content accuracy, leo-7b outperformed gpt2-xl, demonstrating a more consistent replication of original content, as shown by the human evaluation scores.

Compared to the gpt2-xl model, the leo-7b model reproduces content much more accurately. On average, the content agreement of this model’s simplifications and the original text was rated at least as "substantial agreement".

These results on real data suggest that our models, trained on semi-synthetic data, significantly reduce text complexity while reliably retaining content. Semi-synthetic data is a promising way to train text simplification models and circumvent data scarcity problems.

6 Limitations

Our examination reveals that rule-based metrics have limited suitability for evaluating state-of-the-art models in document-level simplification. While our chosen evaluation methodology yields promising results, it lacks a targeted analysis of the end-users for whom the simplification is intended, a scope beyond the ambit of this study. Further-

more, alternative methods to simplify language using LLMs, such as in few-shot learning, merit a comparative analysis against our approach.

7 Conclusion

This study represents a significant stride in tackling the challenge of data scarcity in automatic text simplification. We have crafted a semi-synthetic dataset that has proven effective for training models, which are capable of simplifying complex texts. Notably, our models trained on this synthetic data have demonstrated proficiency in simplifying real web content, validating the practicality of our approach. Semi-synthetic data offers the opportunity to efficiently integrate large amounts of existing simplifications into supervised training without manual effort. This is an efficient and promising alternative to alignments or the manual creation of parallel data.

A vital contribution of this work is the open availability of both the dataset and the models, which serve as a foundational resource for researchers in the text simplification field. Integrating state-of-the-art LLMs with supervised learning has shown to be an efficient method for German text simplification. The limitations of current automated and rule-based metrics, such as BLEU, METEOR, and SARI, are increasingly apparent, particularly for our document-level simplification dataset. This suggests that state-of-the-art LLMs may advance to a point where more nuanced evaluation methodologies are required to accurately measure performance differences and select superior models.

Looking ahead, there are promising directions for future research. One crucial area is adapting generation to adhere more closely to the specific simplification styles, for example, following the rules of easy language and thereby tailoring simplifications more effectively to target audiences. This could involve exploring ways to influence the style of simplification, which would enhance the applicability of these models in real-world applications. By fine-tuning the models to align with the nuanced requirements of different user groups, we can make significant strides toward more inclusive and accessible digital content.

Acknowledgments

This project is funded by the German Federal Ministry for Family Affairs, Senior Citizens, Women and Youth under the grant number 3923406K05.

References

- Fernando Alva-Manchego, Carolina Scarton, and Lucia Specia. 2021. [The \(Un\)Suitability of Automatic Evaluation Metrics for Text Simplification](#). *Computational Linguistics*, 47(4):861–889. Place: Cambridge, MA Publisher: MIT Press.
- Miriam Anshütz, Joshua Oehms, Thomas Wimmer, Bartłomiej Jezierski, and Georg Groh. 2023. [Language Models for German Text Simplification: Overcoming Parallel Data Scarcity through Style-specific Pre-training](#).
- Dennis Aumiller and Michael Gertz. 2022. [Klexikon: A German Dataset for Joint Summarization and Simplification](#).
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- W. Coster and David Kauchak. 2011. [Simple English Wikipedia: A New Text Simplification Task](#).
- Silvana Deilen, Sergio Hernández Garrido, Ekaterina Lapshinova-Koltunski, and Christiane Maaß. 2023. [Using ChatGPT as a CAT tool in Easy Language translation](#).
- Sarah Ebling, Alessia Battisti, Marek Kostrzewa, Dominik Pfützte, Annette Rios, Andreas Säuberli, and Nicolas Spring. 2022. [Automatic Text Simplification for German](#). In *Frontiers in Communication*, volume 7, page 706718. ISSN: 2297-900X Journal Abbreviation: Front. Commun.
- Avishek Garain, Arpan Basu, Rudrajit Dawn, and Sudip Kumar Naskar. 2019. [Sentence Simplification using Syntactic Parse trees](#). *2019 4th International Conference on Information Systems and Computer Networks (ISCON)*, pages 672–676.
- Renate Hauser, Jannis Vamvas, Sarah Ebling, and M. Volk. 2022. [A Multilingual Simplified Language News Corpus](#).
- Freya Hewett. 2023. [APA-RST: A Text Simplification Corpus with RST Annotations](#). *Proceedings of the 4th Workshop on Computational Approaches to Discourse (CODI 2023)*, pages 173–179.
- Daniel Holmer and Evelina Rennes. 2023. [Constructing Pseudo-parallel Swedish Sentence Corpora for Automatic Text Simplification](#).
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The Curious Case of Neural Text Degeneration](#). ArXiv:1904.09751 [cs].
- Abhinav Jain, Hima Patel, Lokesh Nagalapatti, Nitin Gupta, Sameep Mehta, Shanmukha Guttula, Shashank Mujumdar, Shazia Afzal, Ruhi Sharma Mittal, and Vitobha Munigala. 2020. [Overview and Importance of Data Quality for Machine Learning Tasks](#). In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '20*, pages 3561–3562, New York, NY, USA. Association for Computing Machinery.
- Tomoyuki Kajiwara and Mamoru Komachi. 2018. [Text Simplification without Simplified Corpora](#). In *Journal of Natural Language Processing*, volume 25, pages 223–249. ISSN: 1340-7619, 2185-8314 Issue: 2 Journal Abbreviation: Journal of Natural Language Processing.
- Mounica Maddela, Yao Dou, David Heineman, and Wei Xu. 2022. [LENS: A Learnable Evaluation Metric for Text Simplification](#).
- Benjamin Minixhofer, Fabian Paischer, and Navid Reksabsaz. 2022. [WECHSEL: Effective initialization of subword embeddings for cross-lingual transfer of monolingual language models](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3992–4006, Seattle, United States. Association for Computational Linguistics.
- Netzwerk-Leichte-Sprache. 2022. [Die Regeln - Netzwerk Leichte Sprache](#).
- OpenAI. 2023. [GPT-4 Technical Report](#). ArXiv:2303.08774 [cs].
- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing Zhu. 2002. [Bleu: a Method for Automatic Evaluation of Machine Translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Björn Plüster. 2023. [LeoLM: Igniting German-Language LLM Research | LAION](#).
- Archana Praveen Kumar, Ashalatha Nayak, Manjula Shenoy K., Roshan Jacob Manoj, and Akansha Priyadarshi. 2022. [Pattern-Based Syntactic Simplification of Compound and Complex Sentences](#). *IEEE Access*, 10:53290–53306.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving Language Understanding by Generative Pre-Training](#).
- Annette Rios, Nicolas Spring, Tannon Kew, Marek Kostrzewa, Andreas Säuberli, Mathias Müller, and Sarah Ebling. 2021. [A New Dataset and Efficient Baselines for Document-level Text Simplification in German](#). *Proceedings of the Third Workshop on New Frontiers in Summarization*, pages 152–161.

Nicolas Spring and Annette Rios. 2021. [Exploring German Multi-Level Text Simplification](#). In *Proceedings of the Conference Recent Advances in Natural Language Processing - Deep Learning for Natural Language Processing Methods and Applications*, pages 1339–1349. INCOMA Ltd. Shoumen, BULGARIA.

Yixuan Su and Nigel Collier. 2022. [Contrastive Search Is What You Need For Neural Text Generation](#). Publisher: arXiv Version Number: 3.

Elior Sulem, Omri Abend, and Ari Rappoport. 2018. [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, Brussels, Belgium. Association for Computational Linguistics.

Vanessa Toborek, Moritz Busch, Malte Boßert, Christian Bauckhage, and Pascal Welke. 2022. [A New Aligned Simple German Corpus](#). arXiv.

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.

A Example Appendix

Table 6: Parameter settings for various algorithms

Parameter	Wert
<i>Finetuning</i>	
learning_rate	2e-5
weight_decay	0.05
batch_size	2
n_epochs	3
<i>Greedy</i>	
no_ngram_repeat_size	5
max_length	1024
<i>Beam Search</i>	
no_ngram_repeat_size	5
max_length	1024
num_beams	5
early_stopping	True
<i>Sampling</i>	
no_ngram_repeat_size	5
max_length	1024
do_sample	True
top_p	0.95
top_k	5
temperature	0.5
<i>Contrastive</i>	
no_ngram_repeat_size	5
max_length	1024
penalty_alpha	0.05
top_k	5