

Text Simplification of Scientific Texts for Non-Expert Readers

Notebook for the SimpleText Lab at CLEF 2023

Björn Engemann¹, Fabian Haak¹, Christin Katharina Kreutz¹, Narjes Nikzad Khasmakhi¹ and Philipp Schaer¹

¹TH Köln – University of Applied Sciences, Cologne, Germany

Abstract

Reading levels are highly individual and can depend on a text's language, a person's cognitive abilities, or knowledge on a topic. Text simplification is the task of rephrasing a text to better cater to the abilities of a specific target reader group. Simplification of scientific abstracts helps non-experts to access the core information by bypassing formulations that require domain or expert knowledge. This is especially relevant for, e.g., cancer patients reading about novel treatment options.

The SimpleText lab hosts the simplification of scientific abstracts for non-experts (Task 3) to advance this field. We contribute three runs employing out-of-the-box summarization models (two based on T5, one based on PEGASUS) and one run using ChatGPT with complex phrase identification.

Keywords

Text Simplification, Natural Language Processing, Transformer Language Models, ChatGPT, Prompt Engineering, Non-Expert Readers

1. Introduction

Even with the ever-increasing number of scientific publications, accessing information from scientific texts for non-expert readers remains a problem, e.g., cancer patients investigating possible treatment options [1]. Complex scientific texts pose challenges for non-expert readers for various reasons: jargon, abbreviations, and technical details hinder accessing such content. Text simplification attempts to mitigate these challenges, making the content more comprehensible to a broader range of readers [2] by reducing language complexity in terms of vocabulary and sentence structure without distorting its original intent and meaning [3].

Estimating a text's readability is subjective and reader-dependent. For example, substituting complex terms could be considered a good solution for non-native speakers, while shortening and rephrasing the text may be more suitable for children who have limited literacy skills [4].

SimpleText¹ [5], now in its third iteration as a part of CLEF [6, 7], intends to advance

CLEF 2023: Conference and Labs of the Evaluation Forum, September 18–21, 2023, Thessaloniki, Greece

✉ bjoern.engelmann@th-koeln.de (B. Engemann); fabian.haak@th-koeln.de (F. Haak); christin.kreutz@th-koeln.de (C. K. Kreutz); narjes.nikzad_khasmakhi@th-koeln.de (N. Nikzad Khasmakhi); philipp.schaer@th-koeln.de (P. Schaer)

🆔 0009-0000-7074-9066 (B. Engemann); 0000-0002-3392-7860 (F. Haak); 0000-0002-5075-7699 (C. K. Kreutz); 0000-0003-3536-1343 (N. Nikzad Khasmakhi); 0000-0002-8817-4632 (P. Schaer)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

¹<http://simpletext-project.com/2023/clef/>

simplification of scientific text for non-experts. The lab consists of three tasks: selecting passages to include in a simplified summary (task 1), identifying and explaining difficult concepts for general audiences (task 2), and scientific text simplification (task 3) [5].

This work exclusively focuses on task 3, which has already been tackled in previous iterations of SimpleText and beyond [8, 9, 10]. Recent advancements in language modeling, particularly transformer-based architectures like T5 [11] and different GPT variants [12, 13, 14], have enabled text simplification methods to incorporate semantic features in addition to lexical and syntactic ones. Notable works [15, 16, 8] primarily focus on using T5 models, while Rubio and Martínez [9] leverage BART [17]. Additionally, a GPT-2-based text simplification model was proposed in a zero-shot manner [18]. Recently, ChatGPT [19] has emerged as a powerful tool that can be used without training.

In one of our approaches for text simplification for non-experts, we use prompt engineering and complex phrase identification preprocessing with ChatGPT. Further, we contribute approaches based on the T5 and PEGASUS summarization models.

2. Dataset

The dataset for task 3 of the SimpleText consists of short texts taken from scientific publications, primarily single sentences, e.g., *"The malicious alteration of machine time is a big challenge in computer forensics."* (sentence ID G09.1_567125047_1). The SimpleText lab offers a train and a test dataset split. The test split is categorized into small, medium, and large datasets, with the texts from the small being included in the medium and the large split containing all texts from the other two. In our work, we utilized the training set to engineer prompts and implemented our approaches on the large dataset comprising 152,072 source texts. We encountered several issues with the large dataset. The dataset includes duplicated samples. Removing them reduces the number of unique texts in the dataset to 135,540. Some texts (e.g., all sentences of the document with the doc_ID 2787296320) appear up to five times in the dataset. These duplicate entries are identical, with a few exceptions where the query texts differ. Additionally, the dataset frequently contains utf-8 hexcodes within the texts, as shown in Table 4. The next issue concerns samples where some texts are incomplete, as highlighted in the example mentioned in subsection 3.1. These issues pose challenges for the summarization models described in subsection 3.2. Further, a small number of texts are either empty or contain punctuation marks. This can result in AI hallucination or generating empty simplified texts, as illustrated in Table 1. The phenomenon of hallucination is an already recognized problem of large language models [20].

3. Methodological Approaches to Simplification

This section outlines the methodology employed in each of our submitted runs. We generated four runs, including one using ChatGPT prompting and three based on out-of-the-box text summarization models.

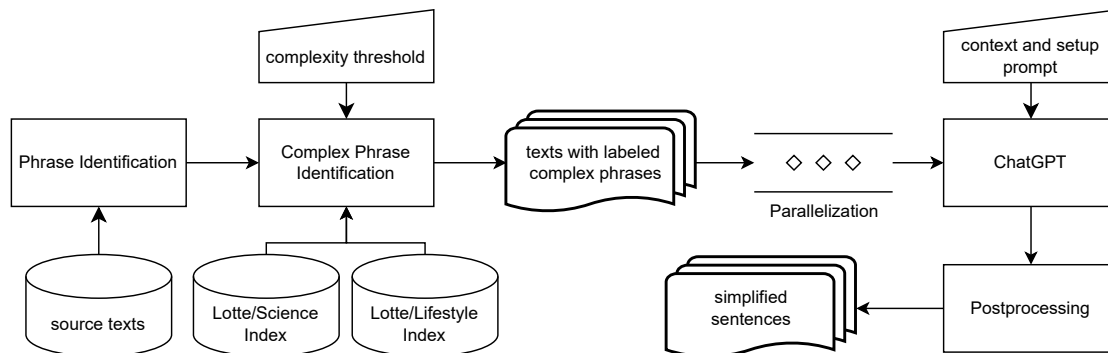


Figure 1: Pipeline of our ChatGPT run including the complex phrase identification component.

3.1. ChatGPT for Simplification

In our `irgc_task_3_ChatGPT_2stepTurbo` run (called ChatGPT in the remainder of this paper), we employed a combination of complex phrase identification and simplification with ChatGPT.

The architecture of our two-step simplification approach using ChatGPT conjoined with complex phrase identification is depicted in Figure 1. In brief, we extract keyphrases from the source text before identifying complex phrases. The complex phrases are then labeled in the source text, which is given to ChatGPT with a prompt asking it to simplify the sentences with special focus on the labeled parts.

In our approach, we avoided performing any preprocessing. However, we conducted some initial testing to gain better insights into the samples and the functionality of ChatGPT. We witnessed that our model sometimes generated information not present in the source texts. We encountered this situation primarily when the sentences in the source text were sparse and a query text was provided in the prompt. For example, by including the query text *"digital assistant"* along with the text *"This approach is ma..."* (sentence ID `G01.1_3006661050_2`), we noted that the text was rephrased as *"This way of doing things is commonly used in digital assistants."* To address this issue, we decided not to include the provided query texts in our approach. When we tried to simplify multiple source texts at once using ChatGPT, the model inferred information from consecutive sentences despite the prompt telling the model to treat these texts separately. Consequently, we randomized the order of the source texts.

For our submitted run, the first step of the approach consists of phrase identification and complex phrase identification. These labeled texts are then provided to the ChatGPT module as a second step for simplification. Finally, the simplified texts are post-processed. Since the complex phrase identification, simplification, and post-processing steps are fundamental components of our approach, we will elaborate on them in detail.

Complex Phrase Identification One aspect of our simplification approach is to avoid highly complex phrases, i.e., phrases that are difficult to understand for non-experts. We first identify keyphrases from the source text to make this possible. Keyphrases are those phrases of a text

which contain core information [21]. We assume that especially keyphrases are relevant for understanding texts, so we try to express this core information more easily. Since identifying keyphrases is an already established task, we use *KBIR-inspec*, a pre-trained model which is available on Huggingface².

However, not all recognized phrases are complex, so we introduce a mechanism to distinguish complex from non-complex phrases. To do this, we evaluate the complexity of all terms in a phrase (Equation 1) and whether the terms' aggregated complexity value exceeds a manually specified threshold. Since the domain of the task is scientific texts, we evaluate the complexity of terms using the term statistics of two text corpora: a dataset that is comprised of texts from lifestyle forums and one that is comprised of texts from science-focused forums [22]. We assume that terms whose relative frequency in scientific texts is greater than those from lifestyle texts also tend to be more complex. We obtain our complexity estimate by the difference of the inverse document frequency (*idf*) of term t from the lifestyle dataset with the *idf* value of the scientific dataset. The complexity of a phrase $\langle t_1, \dots, t_n \rangle$ is defined by the function ϕ :

$$\phi(\langle t_1, \dots, t_n \rangle) = \max_{i \in \{1, \dots, n\}} \left\{ \log_{N_{lf}} \left(\frac{N_{lf}}{df_{lf}(t_i)} \right) - \log_{N_{sc}} \left(\frac{N_{sc}}{df_{sc}(t_i)} \right) \right\}. \quad (1)$$

Here $df_{lf}(t)$ is the number of documents from the lifestyle dataset in which the term t appears, $df_{sc}(t)$ for the science dataset respectively. The total number of documents is N . We set the complexity threshold to 0.01. So every phrase above this threshold gets tagged as complex. In our case, tagging means annotating phrases using square brackets (e.g., second row in Table 3). Thus, the more frequently the terms of a phrase occur in scientific texts, and the less frequently they occur in lifestyle texts, the more complex we estimate the phrase to be. To index and analyze the datasets and term statistics, we use the PyTerrier [23] framework.

ChatGPT module The ChatGPT simplification module involves simplifying the sentences, with a focus on the labeled complex phrases. In this work, we employ the *gpt-3.5-turbo-0301* model accessed via API³.

The model is given a first instruction prompt⁴. The prompt was iteratively engineered by testing, with the goal of circumventing common problems encountered with ChatGPT. In addition to a general description of the task to simplify a list of input texts, the prompt contains further specifications. The prompt includes an instruction to treat a list of input terms separately to prevent information from diffusing between sentences. We further provide information on the number of expected input texts. This asserts that the output will not split input texts with more than one sentence. The essential part of the initial prompt is explaining how complex phrases are marked by square brackets, along with an example. The system is requested to either explain or replace these complex phrases. The list of ten input sentences with the complex phrases tagged by the complex phrase identification module is provided to the system as a

²<https://huggingface.co/ml6team/keyphrase-extraction-kbir-inspec>

³<https://platform.openai.com/docs/models/gpt-3-5>

⁴The used prompt is: "You are given 10 texts, TREAT THEM SEPARATELY. Complex technical or scientific terms are indicated by square brackets (e.g. [convolutional neural network]). For each text, replace or explain the terms in square brackets to make it understandable to non-experts and to simplify for non-expert readers. Make sure that you return exactly one enumerated list of 10 simplified texts and that all terms in square brackets are simplified or explained!"

separate user-tagged message as a Markdown-formatted list. The marking of complex phrases by square brackets seems to work well with ChatGPT.

Postprocessing As a first postprocessing step, we remove list tags and split the text generated by ChatGPT into separate simplified texts. This is necessary because ChatGPT does not reliably return a cohesively structured list of results despite the same prompt structure being used. The results were in all cases formatted as a list of texts, but sometimes the lists were enumerated, while some contained dashed or other itemization marks. Furthermore, ChatGPT sometimes separated the list elements by empty lines. Despite the detailed description in the prompt, ChatGPT sometimes did not remove terms in square brackets, and in some of these instances, the square brackets were not removed either. We apply a postprocessing step to cope with these cases, removing all square brackets from all simplified texts.

3.2. Summarization Models for Simplification

In addition to the ChatGPT run, we also conducted three additional runs using out-of-the-box sequence-to-sequence models popular for summarization tasks. The main reason for employing text summarization techniques is to attempt to simplify the text snippets. In this scenario, we have two runs based on the T5 [11] model and one run that uses PEGASUS [24]. For both models we have the exact same processing steps: We instruct the model to *summarize* all source texts. We do not consider the queries but only the source sentences, so the *source_snt*. We use batch processing of texts which we want to simplify, set the maximum input length to 512 tokens per example, pad the inputs, and determine outputs to have a maximum length of 100 tokens. This leads to a model occasionally creating multiple sentences in one summarization, each fully summarizing the original text. Therefore, out of these merged simplifications (summarizations), we pick the shortest sentence as the simplified version of the original text.

In our `irgc_task_3_t5` run (called `t5` for simplicity) we use the T5 [11] base model with 220 million parameters available on Huggingface⁵.

For our `irgc_task_3_t5_noaron` run (called `t5_noaron` for simplicity), we use the output from the `t5` run but exclude a reoccurring hallucination. In the `t5` runs the simplified sentences contain the name of a specific healthcare researcher several times, indicating that person stated the following simplification of the source text, e.g., text ID `G04.3_2094757405_4` "*aaron carroll: modelling of an infection network between viral and cellular proteins will provide a conceptual framework*". In the `t5_noaron` run we remove this hallucination by removing the expression "*aaron carroll:* " from all texts produced by the `t5` run.

In our `irgc_task_3_pegasusTuner007plus_plus` (called `pegasus` in the remainder of this paper for simplicity) run we use the PEGASUS [24] summarization model, which is provided by Tuner007⁶ as a fine-tuned version of the original PEGASUS.

⁵<https://huggingface.co/t5-base>

⁶https://huggingface.co/tuner007/pegasus_summarizer/discussions

4. Implementation and Generated Runs

4.1. Implementation Specifications

We implemented all approaches in Google Colab⁷. The pegasus and t5 runs were generated on an A100 GPU. Rephrasing the sentences using T5 took approximately 90 minutes, using PEGASUS took approximately 4.5 hours.

In order to reduce both the cost of ChatGPT calls and the duration of execution, we took the following approaches. A call to ChatGPT was designed to encapsulate a prompt containing ten sentences as one request to reduce overhead. In addition, we split the dataset into 15 chunks to better handle unexpected problems. To reduce execution time, we processed each chunk using 75 processes. This allowed us to reduce the expected total computation time from 60 to 20 hours. We delayed the start of each process because the ChatGPT API has a request limit of 3,500 requests/minute and 90,000 tokens/minute. Therefore, exceeding the rate limit yields an unusable result for the request. According to our estimation, processing ten texts takes about 14 seconds (ten seconds for pre- and post-processing and four seconds for the ChatGPT API call). The ChatGPT API incurred a cost of about \$30 for the entirety of our runs.

4.2. Results

In order to thoroughly analyze the outputs of our runs, we conducted an *automatic* and an example-based *manual* evaluation of generated simplifications in comparison to the source texts. These evaluations lead to our final estimation of *priorities* for the produced runs. The priorities indicate our estimation of the quality of the produced results and which runs' results should be considered in SimpleText's manual evaluation.⁸

4.2.1. Automatic Evaluation

In our automatic evaluation, we primarily focused on assessing the readability of the generated text rather than other metrics such as semantic accuracy. We believe that in the context of text simplification, readability measures play a crucial role in evaluating the quality of the output.

We observed readability measures that only require text snippets and no gold-labeled simplifications. Table 1 indicates the following averaged readability measures for the source texts as well as our four runs on the large dataset. The number of empty texts indicates how many of the texts a run produced did not contain any content. The following measures have all been averaged by the number of texts which are not empty. The compression ratio indicates by how much the original source text's size has been reduced in the simplified variant. The following readability measures have all been implemented using the Python library textstat⁹: The Flesch readability index [25] quantifies the reading ease of a text based on word and sentence length.

⁷ChatGPT: <https://colab.research.google.com/drive/10LyoZPzxUlqFxHkXyfjxezO469c1ou9z?usp=sharing>
Summarization Models: <https://colab.research.google.com/drive/1dI0rGH2mPMJ8OdsGnrqAz0HWQzfQCxK0?usp=sharing>

⁸We assumed, that a manual evaluation would be part of the evaluation of the submissions, which is not the case. However, we believe manual evaluations as described in our manual evaluation in section 4.2.2 better reflect the quality of the produced texts than the metrics used in the SimpleText evaluation.

⁹<https://pypi.org/project/textstat/>

Table 1

Averaged readability scores and their standard deviations for source texts as well as our four run variants. ↓ indicates the measure is better the lower it is, ↑ indicates the measure is better the higher it is.

measure	source texts	ChatGPT	t5	t5_noaron	pegasus
empty texts ↓	-	3	148	148	1
compression ↑	-	1.16 ± 0.72	1.53 ± 2.3	1.55 ± 2.34	1.0 ± 1.6
Flesch ↑	36.7 ± 23.33	42.41 ± 21.61	46.4 ± 38.65	46.3 ± 38.85	38.79 ± 23.28
Dale-Chall ↓	12.68 ± 2.18	12.35 ± 1.96	12.41 ± 3.05	12.37 ± 3.05	11.44 ± 2.62
difficult words ↓	8.51 ± 4.32	7.55 ± 3.33	6.23 ± 2.86	6.2 ± 2.88	7.99 ± 3.83
reading time ↓	1.94 ± 0.99	1.71 ± 0.68	1.45 ± 0.59	1.44 ± 0.59	2.24 ± 1.03
syllable count ↓	39.12 ± 20.07	34.4 ± 13.81	29.22 ± 12.22	29.11 ± 12.25	45.23 ± 20.86
lexicon count ↓	22.51 ± 11.05	20.26 ± 8.04	16.88 ± 7.04	16.8 ± 7.07	26.18 ± 11.73
sentence count ↓	1.05 ± 0.58	1.07 ± 0.29	1.44 ± 0.56	1.44 ± 0.56	1.32 ± 0.62

Table 2

Artificially constructed example which has not been taken from the dataset showcasing the non-trivial relation between the length of a simplified text and its source text.

type	text	syllables	lexicons	sentences
source	The WHO reprimands the potentate for using UVAs.	13	8	1
manually simplified	The world health organization spoke a warning. The ruler should not be using unmanned areal vehicles (drones).	27	17	2

Scores between 30 and 49 indicate difficult text. The new Dale-Chall [26] score indicates the reading level of a text as a grade corresponding to the familiarity of persons from that grade with a list of the 3000 most common English words. Scores are defined up to a value between 9 and 9.9 which corresponds to the reading level of an average college student. The higher the score, the more difficult a text. The number of difficult words indicates the number of words with ≥ 3 syllables and which are not in a predefined list of easy words. The reading time indicates the seconds required to read a text; each character taking 14.69ms [27]. The syllable count gives the number of syllables in a text. The lexicon count gives the number of different words in a text. The sentence count indicates how many sentences a text consists of. For these last three measures, one could argue that a simplified text should be shorter in terms of containing fewer syllables, words and sentences compared to the original text which also is in line with previous work [28]. The artificially constructed example in Table 2 shows that this oversimplified assumption does not always hold.

From our strictly numerical analysis of the runs, we found that all our runs improved the source texts' readability in all measures except for the number of sentences. The t5_noaron run yielded among the best results for many of the observed measures. The ChatGPT run did not produce the numerically best results. It seemed to produce texts with more syllables and lexicals than the T5 models which negatively impacted the reading time as well.

As a result of our automatic evaluation of readability measures which do only depend on text snippets and not on some form of gold labels, we would expect the t5_noaron run to provide the best results.

4.2.2. Manual Evaluation

To support our decision-making process on the estimation of the quality of runs we conducted an example-based manual evaluation. We randomly selected texts to assess perceived readability, grammar, and vocabulary and how the simplified texts cope with complex scientific terms. This evaluation process was similar to the human evaluation described by Ermakova et al. [6].

Table 3 shows an example for a source text along with the tagged text produced by the complex phrase identification module of the ChatGPT approach as well as the texts of the runs.

Our manual evaluation yielded several interesting observations. The results did not align with the findings of our automatic evaluation. Furthermore, the PEGASUS model produces texts, that are almost identical to the original text, which has been described as a common problem in last year’s iteration of SimpleText [29]. Another finding is that the text produced by the T5 model is significantly shorter and grammatically simpler but omits important information from the source sentence, e.g., the impact on drug discovery. Further, the text still contains scientific formulations ("*conceptual framework*") and a case of AI hallucination in the form of reporting the text as a quote, which has been partially filtered out in the t5_noaron approach.

From our perspective, the ChatGPT approach yielded the best results by far. The identification of phrases that need to be replaced or explained for a non-expert reader was successful. The ChatGPT model replaced or explained most of the phrases with some exceptions such as "*infection network*", which sub-optimally was replaced with "*infection pathways*". Overall, however, the text produced by the model conveys all essential information from the source sentence in a simpler, concise way.

Another factor influencing our decision to rank the ChatGPT approach as best performing in the manual evaluation is the ability of the system to cope with the problems in the dataset discussed in section 2 as well as correctly interpreting abbreviations and scientific terms without further context. The example shown in Table 4 illustrates these advantages, showing how the model correctly interprets the utf-8 fragments ("*NHPu0027s*") as well as interpreting "*NHP*" as "*non-human primates*".

4.2.3. Priorities of Runs

Due to the task description asking to simplify texts for non-experts, we weigh the overall soundness, completeness, and textual quality evaluated in the manual evaluation more than the results of the automatic evaluation. This decision is consolidated by the fact that the priority determines which of the runs will be manually evaluated. Further, the metrics for the automatic evaluation mostly reflect the readability for those with reduced reading capability, such as children, the most common application of text simplification. However, non-experts should be no less capable of reading texts with, e.g., longer words or generally higher complexity, rather the overall scientific specificity needs to be reduced. Accordingly, we give the ChatGPT run the highest priority, followed by the t5_noaron, the pegasus, and the t5 runs in that order.

Table 3

Examples for the output of our runs, including the complex term identification module of the ChatGPT approach, for an exemplary source sentence (ID G04.3_2094757405_4). In the manual evaluation, we found the ChatGPT run to produce the best simplifications for a non-expert audience.

source text	Indeed, modelling of an infection network between viral and cellular proteins will provide a conceptual and analytic framework to efficiently formulate new biological hypothesis at the proteome scale and to rationalize drug discovery.
tagged text	Indeed, modelling of an [infection network] between viral and cellular proteins will provide a conceptual and analytic framework to efficiently formulate new [biological hypothesis] at the [proteome] scale and to rationalize [drug discovery].
ChatGPT	Building a model of infection pathways between viral and cellular components can help scientists generate new ideas for drugs at a large scale.
t5	aaron carroll: modelling of an infection network between viral and cellular proteins will provide a conceptual framework.
t5_noaron	modelling of an infection network between viral and cellular proteins will provide a conceptual framework.
pegasus	Modelling of an infection network between viral and cellular proteins will provide a conceptual and analytic framework to efficiently formulate new biological hypothesis at the proteome scale and to rationalize drug discovery.

Table 4

Example for a source text with problematic formatting issues (ID G03.2_2395758311_4). The ChatGPT approach did not only effectively deal with this issue but correctly interpreted the meaning of "NHP" as "non-human primates" from the given context of the source text.

source text	ChatGPT
There are two challenges in analyzing the NHPu0027s surveillance video: the NHPu0027s behaviors are lack of regularity and intention, and serious occlusions are brought by the fences of the cages.	The irregular behavior of non-human primates and occlusions caused by cages' fences make it challenging to analyze their surveillance video footage.

5. Conclusion

Our research focuses on text simplification, particularly in the context of scientific texts for non-experts. We present one run employing ChatGPT with complex phrase identification and three runs using out-of-the-box summarization models (two based on T5 and one based on PEGASUS).

While our automatic evaluation did not rank ChatGPT as the best run, a manual analysis evaluated the texts produced through ChatGPT as the best. Although we did not explicitly evaluate the inclusion of complex phrase identification in the ChatGPT run, we found it to improve the system's effectiveness. The identified complex terms indicate that the datasets for constructing the complex phrase identification system were a reasonable choice.

During our implementation, we encountered several issues. Labeling complex phrases using square brackets in our approach may pose challenges when the input text already contains square brackets. This problem could be circumvented by implementing an additional preprocessing

step. Another challenge is hallucinated content, possibly attributed to the absence of context in data. A possible improvement for the summarization model approaches would be flagging difficult words as ones we want to exclude in the simplified (summarized) variant.

References

- [1] M. M. Afsar, T. Crump, B. H. Far, An exploration on-demand article recommender system for cancer patients information provisioning, in: *FLAIRS 2021*, 2021. URL: <https://doi.org/10.32473/flairs.v34i1.128339>.
- [2] S. S. Al-Thanyyan, A. M. Azmi, Automated text simplification: A survey, *ACM Comput. Surv.* 54 (2021). URL: <https://doi.org/10.1145/3442695>.
- [3] K. C. Sheang, H. Saggion, Controllable sentence simplification with a unified text-to-text transfer transformer, in: *INLG 2021, ACL, Aberdeen, Scotland, UK, 2021*, pp. 341–352. URL: <https://aclanthology.org/2021.inlg-1.38>.
- [4] S. Agrawal, M. Carpuat, How to control text simplification? an empirical study of control tokens for meaning preserving controlled simplification, *CoRR abs/2305.14993* (2023). URL: <https://doi.org/10.48550/arXiv.2305.14993>. arXiv: 2305.14993.
- [5] L. Ermakova, E. SanJuan, S. Huet, O. Augereau, H. Azarbonyad, J. Kamps, Overview of simpletext - clef-2023 track on automatic simplification of scientific texts, in: *CLEF 2023, CEUR Workshop Proceedings, CEUR-WS.org, 2023*.
- [6] L. Ermakova, P. Bellot, J. Kamps, D. Nurbakova, I. Ovchinnikova, E. SanJuan, É. Mathurin, S. Araújo, R. Hannachi, S. Huet, N. Poinso, Automatic simplification of scientific texts: Simpletext lab at CLEF-2022, in: *ECIR 2022, volume 13186 of LNCS, Springer, 2022*, pp. 364–373. URL: https://doi.org/10.1007/978-3-030-99739-7_46.
- [7] L. Ermakova, P. Bellot, P. Braslavski, J. Kamps, J. Mothe, D. Nurbakova, I. Ovchinnikova, E. SanJuan, Overview of simpletext 2021 - CLEF workshop on text simplification for scientific information access, in: *CLEF 2021, volume 12880 of LNCS, Springer, 2021*, pp. 432–449. URL: https://doi.org/10.1007/978-3-030-85251-1_27.
- [8] J. Monteiro, M. Aguiar, S. Araújo, Using a pre-trained simplet5 model for text simplification in a limited corpus, in: *CLEF2022, volume 3180 of CEUR Workshop Proceedings, CEUR-WS.org, 2022*, pp. 2826–2831. URL: <https://ceur-ws.org/Vol-3180/paper-241.pdf>.
- [9] A. Rubio, P. Martínez, HULAT-UC3M at simpletext@clef-2022: Scientific text simplification using BART, in: *CLEF 2022, volume 3180 of CEUR Workshop Proceedings, CEUR-WS.org, 2022*, pp. 2845–2851. URL: <https://ceur-ws.org/Vol-3180/paper-243.pdf>.
- [10] E. Hasler, A. de Gispert, F. Stahlberg, A. Waite, B. Byrne, Source sentence simplification for statistical machine translation, *Comput. Speech Lang.* 45 (2017) 221–235. URL: <https://doi.org/10.1016/j.csl.2016.12.001>.
- [11] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, *J. Mach. Learn. Res.* 21 (2020) 140:1–140:67. URL: <http://jmlr.org/papers/v21/20-074.html>.
- [12] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin,

- S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, CoRR abs/2005.14165 (2020). URL: <https://arxiv.org/abs/2005.14165>. arXiv:2005.14165.
- [13] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, Improving language understanding with unsupervised learning, 2018. URL: <https://openai.com/research/language-unsupervised>.
- [14] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, Language models are unsupervised multitask learners, 2019. URL: <https://d4mucfpksywv.cloudfront.net/better-language-models/language-models.pdf>.
- [15] S. Wu, H. Huang, CYUT team2 simpletext shared task report in CLEF-2022, in: CLEF 2022, volume 3180 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 2862–2866. URL: <https://ceur-ws.org/Vol-3180/paper-246.pdf>.
- [16] L. Talec-Bernard, Is using an AI to simplify a scientific text really worth it?, in: CLEF 2022, volume 3180 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 2858–2861. URL: <https://ceur-ws.org/Vol-3180/paper-245.pdf>.
- [17] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, L. Zettlemoyer, BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension, in: *ACL 2020*, 2020, pp. 7871–7880. URL: <https://doi.org/10.18653/v1/2020.acl-main.703>.
- [18] F. Mostert, A. Sampatsing, M. Spronk, D. Rau, J. Kamps, University of amsterdam at the CLEF 2022 simpletext track, in: CLEF 2022, volume 3180 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 2832–2844. URL: <https://ceur-ws.org/Vol-3180/paper-242.pdf>.
- [19] Openai. ChatGPT, <https://chat.openai.com/>, 2022.
- [20] M. Zhang, O. Press, W. Merrill, A. Liu, N. A. Smith, How language model hallucinations can snowball, CoRR abs/2305.13534 (2023). URL: <https://doi.org/10.48550/arXiv.2305.13534>.
- [21] M. Kulkarni, D. Mahata, R. Arora, R. Bhowmik, Learning rich representation of keyphrases from text, in: *NAACL-HLT (Findings) 2022*, *ACL*, 2022, pp. 891–906. URL: <https://doi.org/10.18653/v1/2022.findings-naacl.67>.
- [22] K. Santhanam, O. Khattab, J. Saad-Falcon, C. Potts, M. Zaharia, Colbertv2: Effective and efficient retrieval via lightweight late interaction, arXiv preprint arXiv:2112.01488 (2021). URL: <https://arxiv.org/abs/2112.01488>.
- [23] C. Macdonald, N. Tonello, Declarative experimentation in information retrieval using pyterrier, in: *ICTIR 2020*, *ACM*, 2020, pp. 161–168. URL: <https://doi.org/10.1145/3409256.3409829>.
- [24] J. Zhang, Y. Zhao, M. Saleh, P. J. Liu, PEGASUS: pre-training with extracted gap-sentences for abstractive summarization, in: *ICML 2020*, volume 119 of *Proceedings of Machine Learning Research*, *PMLR*, 2020, pp. 11328–11339. URL: <http://proceedings.mlr.press/v119/zhang20ae.html>.
- [25] R. F. Flesch, A new readability yardstick., *The Journal of applied psychology* 32 3 (1948) 221–33. URL: <https://doi.org/10.1037/h0057532>.
- [26] J. Chall, E. Dale, *Readability Revisited: The New Dale-Chall Readability Formula*, Brookline Books, 1995. URL: <https://books.google.de/books?id=2nbuAAAAMAAJ>.
- [27] V. Demberg, F. Keller, Data from eye-tracking corpora as evidence for theories of syntactic processing complexity, *Cognition* 109 (2008) 193–210. URL: <https://doi.org/10.1016/j.cognition.2008.05.003>.

cognition.2008.07.008.

- [28] L. Martin, É. de la Clergerie, B. Sagot, A. Bordes, Controllable sentence simplification, in: LREC 2020, European Language Resources Association, 2020, pp. 4689–4698. URL: <https://aclanthology.org/2020.lrec-1.577/>.
- [29] L. Ermakova, E. SanJuan, S. Huet, O. Augereau, H. Azarbonyad, J. Kamps, CLEF 2023 simpletext track - what happens if general users search scientific texts?, in: ECIR 2023, volume 13982 of LNCS, Springer, 2023, pp. 536–545. URL: https://doi.org/10.1007/978-3-031-28241-6_62.