

Analysing Zero-Shot Readability-Controlled Sentence Simplification

Abdullah Barayan^{1,2}, Jose Camacho-Collados¹ and Fernando Alva-Manchego¹

¹ School of Computer Science and Informatics, Cardiff University, UK

² Faculty of Computing and Information Technology, King Abdulaziz University, Jeddah, Saudi Arabia

{barayanas, camachocolladosj, alvamanchegof}@cardiff.ac.uk

Abstract

Readability-controlled text simplification (RCTS) rewrites texts to lower readability levels while preserving their meaning. RCTS models often depend on parallel corpora with readability annotations on both source and target sides. Such datasets are scarce and difficult to curate, especially at the sentence level. To reduce reliance on parallel data, we explore using instruction-tuned large language models for zero-shot RCTS. Through automatic and manual evaluations, we examine: (1) how different types of contextual information affect a model’s ability to generate sentences with the desired readability, and (2) the trade-off between achieving target readability and preserving meaning. Results show that all tested models struggle to simplify sentences (especially to the lowest levels) due to models’ limitations and characteristics of the source sentences that impede adequate rewriting. Our experiments also highlight the need for better automatic evaluation metrics tailored to RCTS, as standard ones often misinterpret common simplification operations, and inaccurately assess readability and meaning preservation.

1 Introduction

Text simplification (TS) consists of rewriting a text into an easier-to-understand version while preserving its meaning (Alva-Manchego et al., 2020; Al-Thanyyan and Azmi, 2021). Readability-controlled TS (RCTS), in particular, aims to generate simplifications in specified lower readability levels (Scarton and Specia, 2018). This is useful when creating language learning materials for children (Gala et al., 2020; Javourey-Drevet et al., 2022) and non-native speakers (Zou and Long, 2019) to assist with reading comprehension. Figure 1 shows how an RCTS model can rewrite a text from a high readability level to different lower ones.

Most methods for sentence-level RCTS rely on fully supervised training approaches (Nishi-

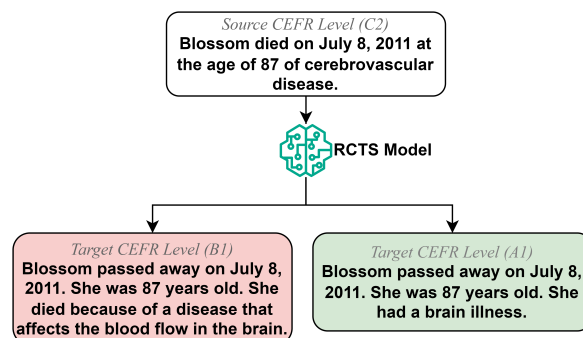


Figure 1: Example of an RCTS model rewriting a text with CEFR level C2 to either A1 or B1 levels.

hara et al., 2019; Spring et al., 2021; Rios et al., 2021; Yanamoto et al., 2022; Zetsu et al., 2022; Agrawal and Carpuat, 2023), and are hindered by the scarcity of high-quality annotated datasets that include original-simplified sentence pairs alongside readability levels. Parallel datasets such as Newsela (Xu et al., 2015) and OneStopEnglish (Vajjala and Lučić, 2018) contain readability annotations only at the document level, which cannot be directly transferred to individual paragraphs or sentences in each text. To address this, we investigate using instruction-tuned large language models (LLMs) for RCTS on English sentences without parallel annotated data (i.e. in a zero-shot setting). In particular, we use the CEFR-SP dataset (Arase et al., 2022), since the CEFR standard offers multiple readability targets, which challenge controlled generation models. We exploit LLMs’ ability to follow natural language instructions by providing different types of contextual information about the desired readability, including descriptions of readers’ skills at the target CEFR level and examples of sentences with the required readability.

We add to existing work on LLMs for RCTS (Imperial and Tayyar Madabushi, 2023; Farajidizaji et al., 2023; Agrawal and Carpuat, 2023) by: (1) analysing models’ performance at the sentence

level rather than the passage level; (2) examining a higher number of LLMs (with a total of 8 models), and (3) measuring not just a model’s ability to meet the target readability, but also how this affects preserving the original meaning using both automatic metrics and human assessment. We find that:

- Generating simplifications at the specified readability levels was difficult for all models, especially when the gap between the source and target readability levels was considerable;
- Providing specific information about the target readability does not have a consistent impact across all models;
- Improving evaluation metrics for RCTS is crucial, as current methods often overlook key simplification operations that are well regarded by human judges. Consequently, automatic readability predictors tend to overestimate text complexity, and semantic similarity metrics penalize effective paraphrasing.

2 Sentence-Level RCTS Approach

We leverage the instruction-tuned LLMs’ ability to follow natural language instructions, which we use to specify the desired readability levels and provide contextual details about them.

2.1 Target Readability Levels

As our target readers, we chose learners of English as a second language. One of the most recognised metrics for measuring language learning is the Common European Framework of Reference for Languages (CEFR),¹ which categorizes language proficiency on a six-point reference scale: A1, A2, B1, B2, C1 and C2, denoting increasing levels of difficulty. For our study, sentences in CEFR levels C2, C1 and B2 are considered as “difficult” for our readers, and we analyse the LLMs’ ability to rewrite them to “simpler” levels B1, A2, and A1.

2.2 Contextual Information

We explore how contextual details related to target readability level affect LLMs’ ability to control readability in generated simplifications, particularly examining two types of information:

- **CEFR Descriptor:** specifies what readers at the target CEFR level can understand, according to the overall reading comprehension

¹<https://www.coe.int/en/web/common-european-framework-reference-languages>

scale in the communicative language activities CEFR Descriptor Scheme.²

- **CEFR Examples:** examples of sentences that readers at the desired CEFR level can understand. For each level, we choose three random sentences from the training split.³

We designed four prompts based on the same general instruction template that details the task and target CEFR level, but with variations in the additional information provided: (P1) target, (P2) target+descriptor, (P3) target+examples, and (P4) target+descriptor+examples. Following previous work in TS (Kew et al., 2023), we adopt the detailed prompting style that describes the task. The specific prompts used can be found in Appendix C.

2.3 Instruction-Tuned LLMs

We investigate 8 instruction-tuned LLMs with diverse sizes and two types of architecture: an encoder-decoder and a standalone decoder, including open-weight and closed-weight models. For the open-weight models, we use LLMs in the 80M-8B range: **Flan-T5** (Chung et al., 2022) in four size variants Small (80M), Base (250M), Large (780M), and XL (3B); **Llama-3-instruct-8B** (AI@Meta, 2024); **OpenChat_3.5-7B** (Wang et al., 2023); and **Mistral-7B-Instruct** (Jiang et al., 2023). For closed models, we use **GPT-3.5-Turbo** and **GPT-4-Turbo** from OpenAI.⁴ More information on the models can be found in Appendix A.

3 Evaluation Setup

3.1 Test Set

For our study, we relied on the CEFR-based Sentence Profile corpus (CEFR-SP, Arase et al., 2022). This is a sentence-level readability assessment dataset, containing 17k English sentences sourced from Newsela-Auto and Wiki-Auto (as collected by Jiang et al. (2020)), and SCoRE (Chujo et al., 2015). As such, its sentences belong to three different domains: news, encyclopedic (Wikipedia) and English language learning, respectively. Each sentence in the dataset was annotated by two English-education experts with a corresponding CEFR standard scale. To identify “complex” instances for

²Each CEFR level descriptor can be found in Appendix B.

³Studies on the effect of the number of examples or smarter strategies for selecting them are left for future work.

⁴<https://platform.openai.com/docs/models>

Level	Num. Sents.	Avg. Len.	Avg. TD	Avg. AoA
B2	142	18.43	6.38	4.94
C1	187	17.74	6.11	4.95
C2	72	18.87	6.65	5.17

Table 1: Test set statistics: the total number of sentences per level (Num. Sents), average sentence length (Avg Len), average tree depth (Avg TD) and average Kuperman age of acquisition score (Avg AoA).

our simplification experiments, we selected sentences from the test split of CEFR-SP whose readability annotations are publicly-available (i.e. from WikiAuto and SCoRE), and where both annotators agreed on their readability being above B1. This resulted in a total of 401 sentences. As shown in Table 1, around 65% of sentences in our test set belong to levels C1 and C2, indicating that models will need to deal with mostly “difficult” input sentences. Surprisingly, the readability levels show minimal variation in syntactic and lexical complexities, as indicated by measures of syntactic tree depth and age of acquisition, respectively. However, sentences at the C2 level (the highest readability) present the highest values in both indices.

3.2 Automatic Evaluation Metrics

We evaluated the generated simplifications using automatic metrics to check if the target readability was met and the original meaning preserved:

Adjacency CEFR Accuracy: measures the percentage of sentences for which the system’s output aligns closely with the target CEFR level. Specifically, it considers outputs successful if their CEFR level is within one level of the specified target. To predict CEFR levels, we use the state-of-the-art sentence-level estimator of (Arase et al., 2022).⁵ In preliminary experiments, LLMs and other models had lower performances when predicting CEFR levels (Appendix E).

Root Mean Squared Error (RMSE): measures the average error between the estimated and target CEFR levels of the system output, by calculating the square root of the average squared differences.

Spearman Correlation (ρ): measures rank-order correlation between target and estimated CEFR levels of the system outputs.

⁵<https://zenodo.org/records/7234096>

Semantic Textual Similarity (STS): measures the cosine similarity between the original and generated simplification embeddings. Higher scores reflect better meaning preservation in the output. We use the sentence transformer library (Reimers and Gurevych, 2019) to compute the sentence embeddings using the **all-MiniLM-L6-v2** model.

Exact Copies: measures the percentage of generated simplifications that are exact matches to their corresponding original texts.

We also assessed fluency, but scores were high across all models and evaluation settings, so we only describe them in Appendix F.

3.3 Baselines

COPY: System that simply copies the original text for each target readability level, setting the lowest standard for metric evaluation.

SUPERVISED: Following Scarton and Specia (2018), we trained a level-based supervised baseline that incorporates a special token to indicate the target CEFR level (e.g., <A1>, <A2>, <B1>) at the beginning of each original sentence. Due to the unavailability of original-simplified sentence pairs annotated with CEFR levels, we fine-tuned a pre-trained T5-Base (Raffel et al., 2020) model on a specially curated unsupervised corpus.⁶

Table 2 presents the scores obtained by all selected models in our test set, with the different designed prompts, and across all automatic metrics. The following sections analyze these results in terms of two core aspects: (1) the achievement of target readability levels (Sec. 4), which ensures the simplified text meets intended accessibility goals, and (2) the preservation of the original meaning (Sec. 5), essential for maintaining fidelity to the source content. This dual focus is crucial for evaluating the effectiveness of the models in producing accessible, yet accurate, simplified texts.

4 Are Target Readability Levels Met?

According to the scores in Table 2 (left side), all models have difficulties in reaching the target CEFR levels, independently of the prompt being used. GPT-4 with P4 emerges as the top-performing model, but it only achieves a CEFR

⁶Details about model training, training dataset creation, and inference settings are provided in Appendix D.

Models	ρ (\uparrow)	AdjAcc (%) (\uparrow)	RMSE (\downarrow)	STS (%) (\uparrow)	Copies (%) (\downarrow)
<i>Baselines</i>					
COPY SUPERVISED	0	24.15	2.45	100	100
	0.01 \pm 0.02	39.43 \pm 1.44	2.08 \pm 0.04	84.51 \pm 0.75	12.69 \pm 0.91
<i>P1: target level</i>					
Flan-T5-small	-0.01 \pm 0.02	34.69 \pm 0.76	2.21 \pm 0.01	86.70 \pm 0.58	12.97 \pm 1.12
Flan-T5-base	-0.01 \pm 0.01	28.01 \pm 0.32	2.38 \pm 0.02	95.41 \pm 0.41	30.15 \pm 2.24
Flan-T5-large	-0.00 \pm 0.01	28.04 \pm 0.49	2.36 \pm 0.01	96.40 \pm 0.25	25.68 \pm 1.58
Flan-T5-xl	-0.00 \pm 0.01	27.96 \pm 0.89	2.38 \pm 0.01	97.17 \pm 0.20	30.09 \pm 1.39
GPT-3.5-Turbo-Instruct	0.05 \pm 0.01**	39.87 \pm 1.03	2.01 \pm 0.03	85.12 \pm 0.36	0.00 \pm 0.00
GPT-4-Turbo	<u>0.11\pm0.01**</u>	<u>51.68\pm0.79</u>	<u>1.67\pm0.02</u>	76.96 \pm 0.33	0.00 \pm 0.00
OpenChat_3.5	-0.01 \pm 0.02	42.87 \pm 0.29	1.96 \pm 0.01	85.35 \pm 0.28	0.11 \pm 0.14
Llama-3-8B-Instruct	0.07 \pm 0.00**	44.81 \pm 1.57	1.89 \pm 0.03	84.32 \pm 0.23	0.58 \pm 0.37
Mistral-7B-Instruct-v0.2	0.06 \pm 0.01**	41.40 \pm 0.83	1.98 \pm 0.01	83.63 \pm 0.20	0.00 \pm 0.00
<i>P2: target level + description</i>					
Flan-T5-small	0.02 \pm 0.03	33.25 \pm 0.92	2.23 \pm 0.03	86.99 \pm 0.40	13.13 \pm 1.29
Flan-T5-base	-0.01 \pm 0.01	27.60 \pm 0.70	2.39 \pm 0.02	95.54 \pm 0.45	29.54 \pm 2.27
Flan-T5-large	0.01 \pm 0.02	28.29 \pm 0.88	2.36 \pm 0.01	96.19 \pm 0.35	25.96 \pm 3.25
Flan-T5-xl	0.01 \pm 0.01	28.49 \pm 1.03	2.37 \pm 0.02	96.95 \pm 0.33	28.43 \pm 1.05
GPT-3.5-Turbo-Instruct	0.08 \pm 0.02**	41.51 \pm 0.80	1.98 \pm 0.02	84.95 \pm 0.23	0.00 \pm 0.00
GPT-4-Turbo	<u>0.18\pm0.02**</u>	<u>52.04\pm0.69</u>	1.66\pm0.01	77.24 \pm 0.25	0.00 \pm 0.00
OpenChat_3.5	0.03 \pm 0.02	45.08 \pm 1.08	1.90 \pm 0.02	83.04 \pm 0.41	0.17 \pm 0.11
Llama-3-8B-Instruct	0.17 \pm 0.01**	46.33 \pm 1.01	1.83 \pm 0.01	83.29 \pm 0.35	0.33 \pm 0.19
Mistral-7B-Instruct-v0.2	0.09 \pm 0.01**	41.65 \pm 0.88	1.96 \pm 0.01	83.45 \pm 0.31	0.00 \pm 0.00
<i>P3: target level + examples</i>					
Flan-T5-small	-0.00 \pm 0.01	38.79 \pm 0.89	2.12 \pm 0.02	78.54 \pm 0.49	6.48 \pm 1.28
Flan-T5-base	-0.01 \pm 0.01	29.01 \pm 0.90	2.35 \pm 0.01	93.62 \pm 0.32	22.94 \pm 1.42
Flan-T5-large	-0.01 \pm 0.00	28.43 \pm 0.40	2.36 \pm 0.01	95.75 \pm 0.17	23.30 \pm 3.17
Flan-T5-xl	-0.01 \pm 0.01	28.10 \pm 0.55	2.37 \pm 0.01	96.86 \pm 0.26	31.78 \pm 1.35
GPT-3.5-Turbo-Instruct	0.08 \pm 0.00**	42.12 \pm 0.67	1.98 \pm 0.01	84.97 \pm 0.55	0.00 \pm 0.00
GPT-4-Turbo	<u>0.15\pm0.00**</u>	<u>50.73\pm1.17</u>	<u>1.69\pm0.01</u>	78.13 \pm 0.19	0.00 \pm 0.00
OpenChat_3.5	0.06 \pm 0.01*	43.11 \pm 0.68	1.94 \pm 0.01	84.63 \pm 0.27	0.42 \pm 0.27
Llama-3-8B-Instruct	0.11 \pm 0.02**	46.35 \pm 0.73	1.84 \pm 0.01	83.46 \pm 0.38	0.61 \pm 0.31
Mistral-7B-Instruct-v0.2	0.08 \pm 0.01**	40.23 \pm 0.64	2.03 \pm 0.01	85.66 \pm 0.28	0.00 \pm 0.00
<i>P4: target level + description + examples</i>					
Flan-T5-small	-0.02 \pm 0.00	38.51 \pm 1.02	2.13 \pm 0.01	77.83 \pm 0.70	7.70 \pm 1.08
Flan-T5-base	-0.01 \pm 0.02	30.34 \pm 1.31	2.33 \pm 0.03	91.63 \pm 0.28	19.59 \pm 1.21
Flan-T5-large	-0.01 \pm 0.01	29.01 \pm 0.53	2.35 \pm 0.01	95.51 \pm 0.27	22.58 \pm 1.71
Flan-T5-xl	-0.01 \pm 0.00	28.70 \pm 0.69	2.36 \pm 0.02	96.57 \pm 0.24	28.90 \pm 1.84
GPT-3.5-Turbo-Instruct	0.08 \pm 0.01**	43.34 \pm 1.15	1.93 \pm 0.02	84.40 \pm 0.23	0.00 \pm 0.00
GPT-4-Turbo	0.18 \pm 0.01**	52.45\pm0.96	1.66\pm0.01	77.81 \pm 0.27	0.00 \pm 0.00
OpenChat_3.5	0.06 \pm 0.01*	44.00 \pm 0.80	1.91 \pm 0.02	83.74 \pm 0.38	0.25 \pm 0.18
Llama-3-8B-Instruct	0.19\pm0.01**	48.21 \pm 1.23	1.80 \pm 0.02	82.86 \pm 0.16	0.50 \pm 0.19
Mistral-7B-Instruct-v0.2	0.09 \pm 0.00**	42.01 \pm 0.76	1.98 \pm 0.01	84.20 \pm 0.21	0.00 \pm 0.00

Table 2: Results on the test set, across the four prompts. Metrics: Spearman’s ρ , AdjAcc, RMSE, STS, and Copy Percentage. Boldface indicates the best overall performance across all prompts, while underlined results represent the best performance for each individual prompt. Significance: * ($p < 0.01$), ** ($p < 0.001$).

accuracy of 52.45% and a low Spearman’s rank correlation of 0.18. Its RMSE of 1.66 suggests that the model tends to generate simplifications with more than one CEFR level difference from the target. Llama-3-8B-Instruct also shows a similar performance, with accuracies ranging from 44.81% to 48.21%, ρ values between 0.07 and 0.19, and RMSE scores are close to 2.

Compared to the SUPERVISED baseline (accuracy of 39.43%, near-zero correlation, and RMSE

of 2.08), the previous two models show better readability control. Other models, such as Openchat, GPT-3.5-turbo-instruct and Mistral-7B-Instruct-v0.2, show moderate improvements over SUPERVISED, with accuracies between 40.23% and 45.08%, low positive correlations, and RMSE values ranging from 1.90 to 2.03. In contrast, Flan-T5 models demonstrate poor performance in the task, under-performing even against the SUPERVISED baseline with accuracies ranging from




Figure 2: Heatmaps depicting RMSE scores of the best models based on source and target readability levels.

27.60% to 38.79%, mostly negative or negligible correlations, and high RMSE values (between 2.12 and 2.39). Actually, these models have performances close to the COPY baseline, indicating a conservative behavior. The following sections provide further insights into these results.

4.1 The Effect of Target CEFR Information

Providing information about the target readability, in the form of level descriptions and example sentences, resulted in a slight enhancement of performance. This is evidenced by the increase in accuracy and decrease in RMSE when comparing prompts P1 to P2, P3 and P4. Most models achieved their best performance when combining both CEFR descriptions and example sentences (P4). However, Openchat_3.5 exhibited greater improvements from descriptions alone rather than from examples, attaining its highest accuracy and lowest RMSE with prompt P2. Similarly, the Mistral-7B-Instruct-v0.2 model showed a positive response to descriptions, with better performance in P2 compared to P1 and P3, and reached its peak accuracy with the full information provided in P4. In contrast, the Flan-T5 small and base models experienced a decline in performance when provided exclusively with CEFR descriptions (P2), performing slightly worse than with the base prompt (P1). Overall, responses to additional contextual information about target readability varies across models.

4.2 The Effect of Source vs Target Readability

Figure 2 presents the RMSE scores for the top five performing models across all combinations of source and target CEFR levels. The heatmaps indicate that simplification becomes increasingly difficult (evidenced by higher RMSE) as target levels decrease. Models achieve their best performance when source and target CEFR levels are close, especially in simplifications from B2 to B1, as shown by the blue sections. This suggests that models handle moderate simplifications with fewer complexity adjustments more effectively. Conversely, models struggle most as the gap between source and target levels widens, especially when simplifying from C2 to A1, highlighting their limited ability to handle large reductions in complexity.

4.3 Are Generated Outputs “Simple”?

We analyse if the generated outputs are still “simple”, even if they did not meet their target readability levels. In particular, we look at the average syntactic tree depth and the average Kuperman Age of Acquisition (AoA) score,⁷ as proxies for syntactic and lexical complexity, respectively. We compute these scores considering all original sentences, and all generated simplifications in each target CEFR level. As a reference, we also compute them for all sentences in the CEFR-SP dataset manually annotated with the target readability levels (B1, A2 and A1). Figure 3 compares these scores for the simplifications of the top five models, the SUPERVISED

⁷Calculated with LFTK (Lee and Lee, 2023).




Figure 3: Line graphs comparing syntactic and lexical complexities of the generated simplifications from the top-performing models to the source and target level syntactic and lexical complexity.

Model	Levels	AccuracyH(%)	AccuracyA(%)
GPT-4-Turbo	B1	100.00	60.00
	A2	86.67	46.67
	A1	53.33	20.00
	Overall	80	42.2
Llama-3-8B-Instruct	B1	93.33	60.00
	A2	93.33	66.67
	A1	66.67	6.67
	Overall	84.4	44.4

Table 3: Human readability assessment results for GPT-4-Turbo and Llama-3-8B-Instruct. AccuracyH: Human-rated accuracy; AccuracyA: Classifier-rated accuracy.

baseline, and both source and target references.

Regarding syntactic complexity, simplifications from the SUPERVISED baseline exhibit minimal variation across all target CEFR levels, suggesting limited adaptability. In contrast, GPT-4-Turbo and Llama-3-8B-Instruct adjust syntactic complexity more effectively as the target CEFR level decreases. Regarding lexical complexity, most LLMs reduced complexity at lower CEFR levels. The SUPERVISED baseline’s lexical complexity remained relatively flat, again indicating a lack of adaptability to varying levels of proficiency. Notably, GPT-4 and GPT-3.5-Turbo-Instruct were particularly effective at adapting lexical complexity. However, comparing these scores to the ones from the target levels, reveals significant room for improvement.

4.4 Do Humans Agree with the Metrics?

We randomly sampled 45 simplifications (15 per target level) from GPT-4-Turbo and Llama-3-8B-Instruct, totaling 90 instances. We recruited three evaluators with experience in teaching English as a second language, and asked them to: (1) assess the CEFR level that a student would need to understand each generated simplification, and (2) judge whether the corresponding source sentences could actually be adjusted to match the target CEFR lev-

els. This aimed to establish if the fact that an output did not reach the target level was due to issues with the model itself or the nature of the source text.⁸

Once again, we compute Adjacency CEFR Accuracy, but this time using the manual judgements instead of predicted levels. The majority label was set as the CEFR level for each instance. If no clear majority existed, we checked if the annotated levels differed by more than one from the pre-specified target. If so, the simplification was considered incorrect; otherwise, it was considered correct.

Models mostly reach the CEFR targets. Table 3 shows that both GPT-4-Turbo and Llama-3-8B-Instruct produced sentences that aligned with the target CEFR levels in 80% and 84.4% of cases, respectively, based on human evaluations (AccuracyH). This indicates that in the majority of cases, human raters judged the generated sentences to match closely the expected readability levels for the intended audience.

Models struggle with lower CEFR levels. Both models struggled when simplifying sentences to lower CEFR levels, particularly A1, confirming the finding in 4.2. GPT-4 failed in 7 A1 cases and 2 A2 cases out of 9 total failures, while Llama 3 failed in 4 A1 cases and 1 each for A2 and B1 out of 6 total failures. However, for failed cases, especially at A1 level, evaluators judged most source sentences unsuitable for adjustment: GPT-4 had 6 out of 7 A1 cases, and Llama 3 had 3 out of 4 A1 cases judged as nonadjustable. Overall, while both models excelled at higher CEFR levels, simplifying to lower levels remains challenging due to both model limitations and the inherent constraints of limited language proficiency at lower readability levels.

⁸Details on evaluator selection, inter-rater agreement, and the exact instructions provided can be found in Appendix H.1.

Sentence	RatingH	RatingA
Primary hypersomnias, which are conditions that make people very sleepy, are not common.	B1	C1
When a woman is giving birth, the water bag (amniotic sac) breaks and the water (amniotic fluid) comes out from her body.	A2	C1
The students were happy because the Japanese Women’s Team won. They stayed awake late. They cleaned up before they went home.	A2	B1

Table 4: Comparison of Human and Classifier Ratings

Automatic CEFR prediction does not account for all types of simplification operations. To understand the gap between AccuracyH and AccuracyA, we reviewed some instances that show differences between classifier and human evaluations. For the first example in Table 4, human evaluators rated it as B1, while the classifier rated it as C1. However, if we removed the technical term “Primary hypersomnias”, the classifier re-rates the sentence as B1. This shows that the classifier tends to assign a higher level when technical terms are present, even if they are explained in the sentence itself. Similarly, for the second example of Table 4, human evaluators rated it as A2, while the classifier rated it as C1. When the medical terms “amniotic sac” and “amniotic fluid” are removed, the classifier revised the rating to B1, again indicating that it overestimates complexity when presented with technical terms, even if explanations are provided. For the third example, human evaluators rated it as A2, while the classifier rated it as B1. However, when each sentence is evaluated individually, the classifier assigns ratings of A2, A1, and A2, respectively, aligning more closely with human judgment.

This highlights key areas where the classifier’s approach differs from human judgment, and suggests the need for further improvement. Elaboration and sentence splitting are vital simplification operations, particularly at lower CEFR levels, and failing to account for them leads to inaccurate difficulty assessments. Training the classifier with parallel sentences—where the same text is written at different CEFR levels (e.g., A1, A2, B1) could help the model better understand sentence complexity and align more closely with human evaluations.

5 Is the Original Meaning Preserved?

5.1 Vector-based Semantic Similarity

According to the scores in Table 2 (right side), models show varied performances. Flan-T5 models

Original	Simplified	STS
(C2) Pharmacoepidemiology bridges pharmacology and epidemiology and borrows their methodology from both of them.	(B1) (GPT-4) Pharmacoepidemiology combines the study of drugs (pharmacology) and the study of health in populations (epidemiology). It uses methods from both these fields.	0.82
(B2) Nothing you do or say will ever stop Sally from chasing her dream of becoming a veterinarian.	(A1) (Llama3) Sally will always try to become a vet.	0.78
(C1) There are two main kinds of charter: bareboat and skippered.	(B1) (GPT-4) There are two main types of boat rental: one where you rent the boat without a crew (bareboat), and one where you rent the boat with a captain (skippered).	0.60
(C1) However, some people with schizophrenia have the ability to tickle themselves.	(B1) (GPT-4) Some people who have schizophrenia can tickle themselves.	0.97

Table 5: Examples of the impact of simplification on STS scores.

(particularly Flan-T5-large and Flan-T5-x1) achieve the highest STS scores, ranging from 95.51% to 96.95%. This superior performance is likely attributed to their conservative behavior (high percentage of Copies). Other models such as Llama-3-8B-Instruct, Openchat, Mistral-7B-Instruct-v0.2, and GPT-3.5-turbo-instruct achieve scores ranging from 83.04% to 85.66%. However, GPT-4, while excelling in readability control, often shows lower STS scores, ranging from 76.96% to 78.13%.

Through manual inspection, we found that STS tends to over-penalize text changes that could be considered appropriate in RCTS. In order to make a text more accessible, simplification can sometimes involve removing or inserting new information, such when executing compression or elaboration operations, respectively (Alva-Manchego et al., 2020). Such forms of rewriting might result in STS scores that underestimate the true quality of the simplifications. For example, the first three simplifications in Table 5 could be considered of high quality since they accurately capture the meaning of their original counterparts. However, they did not receive the nearly-perfect STS scores that the last conservative simplification did. Experiments with BERTscore (Zhang et al., 2020) and AlignScore (Zha et al., 2023), typically used for meaning preservation and factuality evaluation, yielded similar results. As shown in Table 6, BERTscore favored outputs that minimally modify or directly copy the input text. On the other hand, AlignScore penalized the addition of supporting details (such as elaboration and examples) deeming them incorrect (as illustrated in the first and second examples

Original	Simplified	BERTscore
(C2) Pharmacoepidemiology bridges pharmacology and epidemiology and borrows their methodology from both of them.	(B1) (GPT-4) Pharmacoepidemiology combines the study of drugs (pharmacology) and the study of health in populations (epidemiology). It uses methods from both these fields.	0.52
(C1) Rivlin had pulmonary fibrosis, a lung disease.	(B1) (GPT-4) Rivlin had a lung disease called pulmonary fibrosis.	0.86
(C2) Dimethylmercury is extremely toxic and dangerous to handle.	(B1) (Mistral-7B) Dimethylmercury is very harmful and should be handled with care.	0.75
(C2) Dimethylmercury is extremely toxic and dangerous to handle.	(B1) (Flan-T5-L) Dimethylmercury is extremely dangerous and toxic to handle.	0.95

Table 6: Examples of the impact of simplification on BERTscore scores.

Original	Simplified	AlignScore
(C1) Primary hypersomnias are very rare.	(A1) (GPT-4) Hypersomnias are sleep conditions where someone sleeps a lot or feels very sleepy during the day. Primary hypersomnias are a type of this condition and they do not happen often. They are rare.	0.003
(C1) Bisques are heavy cream soups traditionally prepared with shellfish, but can be made with any type of seafood or other base ingredients.	(A1) (GPT-4) Bisques are thick soups with cream. They are usually made with shellfish like crab or lobster. But you can use other seafood or different ingredients too.	0.44
(B2) Unfortunately, she is a parent whose rules must be enforced through guilt and manipulation instead of reason and trust.	(A1) (Flan-T5-small) This is a parent.	0.97

Table 7: Examples illustrating the impact of simplification on AlignScore scores.

in Table 7), while rewarding factual accuracy even when simplifications are inappropriate (as seen in the third example in Table 7).⁹ This highlights the necessity for meaning preservation metrics that account for the nuances of the editing operations executed while simplifying texts.

5.2 Human Judgements

We randomly sampled 45 simplifications from GPT-4-Turbo and Llama-3-8B-Instruct, totaling 90 instances. Three fluent English speakers evaluated whether the original meaning was retained in the simplified sentences. We adapted the error annotation scheme for evaluating factuality in TS proposed by Devaraj et al. (2022): instead of scoring each operation independently, we asked annotators to provide an overall score for each simplification. Annotators rated each simplification using the fol-

⁹More details and results for BERTscore and AlignScore are available in Appendix G

Model	STS(%) Avg	0 (%)	0.5 (%)	1 (%)	1.5 (%)	2 (%)	
GPT-4-Turbo	67.41	0.94	8.88	13.33	62.22	11.85	3.70
Llama-3-8B-Instruct	79.92	0.95	17.03	22.96	25.92	20.0	14.07

Table 8: Human evaluation of meaning preservation.

lowing scale: 0 indicates *no or trivial changes*; 1 signifies that *non-trivial changes have occurred, but the main idea is maintained*; and 2 indicates that *the main idea has not been preserved in the simplified sentence*. We also allowed intermediate scores (0.5 and 1.5).¹⁰

As shown in Table 8, Llama-3 obtained a higher percentage of minimal (score 0.5) or no changes (score 0) compared to GPT-4-Turbo, which explains its higher STS score. However, Llama-3 also had more instances where meaning was significantly altered or lost (14.07% scored 2) compared to GPT-4-Turbo (3.70%). This contrasts with the STS results, which showed GPT-4-Turbo having the lowest meaning preservation. In contrast, GPT-4-Turbo made more non-trivial changes while still preserving the original meaning, with 62.22% of its simplifications receiving a score of 1, which may explain its lower STS rating. This discrepancy between STS scores and human evaluation underscores the challenge of assessing meaning preservation in simplification tasks, highlighting the need for better metrics that balance meaning retention and substantial rewriting.

6 Related Work

Readability-Controlled Simplification. To control the readability of simplified text, two main approaches have been proposed. The first modifies the models’ training process by (a) adding special tokens (Scarton and Specia, 2018; Nishihara et al., 2019; Spring et al., 2021; Rios et al., 2021; Yanamoto et al., 2022; Zetsu et al., 2022; Agrawal and Carpuat, 2023) or (b) instruction prompts (Chi et al., 2023) to the source sentences during training to guide the model’s behavior. The second approach employs a post-processing control module to control readability during inference (Zetsu et al., 2022). All these methods rely on fully supervised training approaches, requiring a large amount of sentence pairs annotated with readability levels. However, a major challenge lies in the difficulty of obtaining such annotated data. Our work sought to

¹⁰Details on evaluator selection, inter-rater agreement, and the instructions provided can be found in Appendix H.2.

reduce dependence on parallel data by leveraging instruction-tuned LLMs for sentence-level RCTS in zero-shot settings, only relying on contextual information for each target readability level.

Text Simplification with LLMs. Few studies exist on LLMs’ ability to perform TS (Feng et al., 2023; Ryan et al., 2023; Kew et al., 2023), with limited exploration of their ability for controlling the readability of simplifications. Imperial and Tayyar Madabushi (2023) examined various instruction-tuned LLMs in interpreting and applying readability specifications from prompts in their output for story completion and narrative simplification. However, their evaluation was limited to one readability level and focused on readability achievement. Also, Farajidizaji et al. (2023) investigated the potential of ChatGPT and Llama-2 for modifying passages for different readability levels. Different from these works, our study included a higher number and range of LLMs, and focused on sentence-level RCTS rather than passage-level. Furthermore, we not only assessed models on their ability to meet the target readability, but analysed how this affects meaning preservation using both automatic metrics and human assessment.

7 Conclusion

We explored sentence-level zero-shot RCTS to analyse various instruction-tuned LLMs in simplifying sentences to specified CEFR levels. Our experiments showed that LLMs have difficulties with this task, especially when there is a large gap between source and target readability levels. Providing contextual information about the target readability (i.e. CEFR descriptions and example sentences) improves performance but not uniformly.

We also demonstrated that automatically assessing models for RCTS is challenging due to the nature of the task. Standard similarity metrics penalise removing or adding information even if it improves readability, while CEFR predictors often misjudge the true difficulty of simplified texts, failing to account for rewritings via elaboration and sentence splitting. This gap between human and automated assessments underscores the need for more nuanced evaluation mechanisms in RCTS. Refined metrics would not only enhance the evaluation process but also drive significant improvements in RCTS development by better supporting output filtering to ensure both readability and meaning preservation, or by serving as reward signals in

reinforcement learning approaches.

Limitations

Due to our focus on sentence-level RCTS, our study was constrained to the data in the CEFR-SP dataset, the only available resource with sentence-level readability annotations from experts. As such, our experiments were exclusively conducted using English data and limited to the CEFR standard for establishing readability levels.

Additionally, our study focused on zero-shot scenarios. While obtaining a large parallel dataset is challenging, collecting a few (synthetic) samples might be feasible, allowing for studies on few-shot or small-scale fine-tuning setting. This is considered as future work.

Ethics Statement

For our experiments, we only used subsets of the CEFR-SP dataset with non-restrictive licenses. In particular, the Wiki-Auto dataset is publicly available under the CC BY-SA 3.0 license, and the SCoRE dataset is available under the CC BY-NC-SA 4.0 license.

For readability assessments, we paid the evaluators \$65.99 for completing the task, which took an average of 1 hour to complete. Similarly, for meaning preservation assessments, evaluators received a flat rate of \$33 to complete the task, with an average task completion time of 40 minutes.

References

- Sweta Agrawal and Marine Carpuat. 2023. [Controlling pre-trained language models for grade-specific text simplification](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12807–12819, Singapore. Association for Computational Linguistics.
- AI@Meta. 2024. [Llama 3 model card](#).
- Suha S. Al-Thanyyan and Aqil M. Azmi. 2021. [Automated text simplification: A survey](#). *ACM Comput. Surv.*, 54(2).
- Fernando Alva-Manchego, Carolina Scarton, and Lucia Specia. 2020. [Data-driven sentence simplification: Survey and benchmark](#). *Computational Linguistics*, 46(1):135–187.
- Yuki Arase, Satoru Uchida, and Tomoyuki Kajiwara. 2022. [CEFR-based sentence difficulty annotation and assessment](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6206–6219, Abu Dhabi, United

- Arab Emirates. Association for Computational Linguistics.
- Alison Chi, Li-Kuang Chen, Yi-Chen Chang, Shu-Hui Lee, and Jason S. Chang. 2023. [Learning to paraphrase sentences to different complexity levels](#). *Transactions of the Association for Computational Linguistics*, 11:1332–1354.
- Kiyomi Chujo, Kathryn Oghigian, and Shiro Akasegawa. 2015. [A corpus and grammatical browsing system for remedial EFL learners](#). In Agnieszka Lenko-Szymańska and Alex Boulton, editors, *Multiple Affordances of Language Corpora for Data-driven Learning*, volume 69 of *Studies in Corpus Linguistics*, pages 109–128. John Benjamins.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *Preprint*, arXiv:2210.11416.
- Ashwin Devaraj, William Sheffield, Byron Wallace, and Junyi Jessy Li. 2022. [Evaluating factuality in text simplification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7331–7345, Dublin, Ireland. Association for Computational Linguistics.
- Asma Farajidizaji, Vatsal Raina, and Mark Gales. 2023. [Is it possible to modify text to a target readability level? an initial investigation using zero-shot large language models](#). *Preprint*, arXiv:2309.12551.
- Yutao Feng, Jipeng Qiang, Yun Li, Yunhao Yuan, and Yi Zhu. 2023. [Sentence simplification via large language models](#). *Preprint*, arXiv:2302.11957.
- Núria Gala, Anaïs Tack, Ludivine Javourey-Drevet, Thomas François, and Johannes C. Ziegler. 2020. [Alector: A parallel corpus of simplified French texts with alignments of misreadings by poor and dyslexic readers](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 1353–1361, Marseille, France. European Language Resources Association.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Joseph Marvin Imperial and Harish Tayyar Madabushi. 2023. [Flesch or fumble? evaluating readability standard alignment of instruction-tuned language models](#). In *Proceedings of the Third Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 205–223, Singapore. Association for Computational Linguistics.
- Ludivine Javourey-Drevet, Stéphane Dufau, Thomas François, Núria Gala, Jacques Ginestíe, and Johannes C. Ziegler. 2022. [Simplification of literary and scientific texts to improve reading fluency and comprehension in beginning readers of french](#). *Applied Psycholinguistics*, 43(2):485–512.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Chao Jiang, Mounica Maddela, Wuwei Lan, Yang Zhong, and Wei Xu. 2020. [Neural CRF model for sentence alignment in text simplification](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7943–7960, Online. Association for Computational Linguistics.
- Tannon Kew, Alison Chi, Laura Vásquez-Rodríguez, Sweta Agrawal, Dennis Aumiller, Fernando Alva-Manchego, and Matthew Shardlow. 2023. [BLESS: Benchmarking large language models on sentence simplification](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13291–13309, Singapore. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2017. [Adam: A method for stochastic optimization](#). *Preprint*, arXiv:1412.6980.
- Kalpesh Krishna, John Wieting, and Mohit Iyyer. 2020. [Reformulating unsupervised style transfer as paraphrase generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 737–762, Online. Association for Computational Linguistics.
- Sachin Kumar, Eric Malmi, Aliaksei Severyn, and Yulia Tsvetkov. 2021. [Controlled text generation as continuous optimization with multiple constraints](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 14542–14554. Curran Associates, Inc.
- Bruce W. Lee and Jason Lee. 2023. [LFTK: Handcrafted features in computational linguistics](#). In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 1–19, Toronto, Canada. Association for Computational Linguistics.
- Tarek Naous, Michael J. Ryan, Anton Lavrouk, Mohit Chandra, and Wei Xu. 2023. [Readme++: Benchmarking multilingual language models for multi-domain readability assessment](#). *Preprint*, arXiv:2305.14463.

- Daiki Nishihara, Tomoyuki Kajiwara, and Yuki Arase. 2019. [Controllable text simplification with lexical constraint loss](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 260–266, Florence, Italy. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- 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](#). In *Proceedings of the Third Workshop on New Frontiers in Summarization*, pages 152–161, Online and in Dominican Republic. Association for Computational Linguistics.
- Michael Ryan, Tarek Naous, and Wei Xu. 2023. [Revisiting non-English text simplification: A unified multilingual benchmark](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4898–4927, Toronto, Canada. Association for Computational Linguistics.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *International Conference on Learning Representations*.
- Carolina Scarton and Lucia Specia. 2018. [Learning simplifications for specific target audiences](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 712–718, Melbourne, Australia. Association for Computational Linguistics.
- Nicolas Spring, Annette Rios, and Sarah Ebling. 2021. [Exploring German multi-level text simplification](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 1339–1349, Held Online. IN-COMA Ltd.
- Sowmya Vajjala and Ivana Lučić. 2018. [On-eStopEnglish corpus: A new corpus for automatic readability assessment and text simplification](#). In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 297–304, New Orleans, Louisiana. Association for Computational Linguistics.
- Guan Wang, Sijie Cheng, Xianyuan Zhan, Xiangang Li, Sen Song, and Yang Liu. 2023. [Openchat: Advancing open-source language models with mixed-quality data](#). *arXiv preprint arXiv:2309.11235*.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krma Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sampat, Siddhartha Mishra, Sujan Reddy A, Sumanta Patro, Tanay Dixit, and Xudong Shen. 2022. [Super-NaturalInstructions: Generalization via declarative instructions on 1600+ NLP tasks](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5085–5109, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2018. [Neural network acceptability judgments](#). *arXiv preprint arXiv:1805.12471*.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. [Finetuned language models are zero-shot learners](#). In *International Conference on Learning Representations*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Wei Xu, Chris Callison-Burch, and Courtney Napoles. 2015. [Problems in current text simplification research: New data can help](#). *Transactions of the Association for Computational Linguistics*, 3:283–297.

Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. [Optimizing statistical machine translation for text simplification](#). *Transactions of the Association for Computational Linguistics*, 4:401–415.

Daiki Yanamoto, Tomoki Ikawa, Tomoyuki Kajiwara, Takashi Ninomiya, Satoru Uchida, and Yuki Arase. 2022. [Controllable text simplification with deep reinforcement learning](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 398–404, Online only. Association for Computational Linguistics.

Tatsuya Zetsu, Tomoyuki Kajiwara, and Yuki Arase. 2022. [Lexically constrained decoding with edit operation prediction for controllable text simplification](#). In *Proceedings of the Workshop on Text Simplification, Accessibility, and Readability (TSAR-2022)*, pages 147–153, Abu Dhabi, United Arab Emirates (Virtual). Association for Computational Linguistics.

Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. [AlignScore: Evaluating factual consistency with a unified alignment function](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11328–11348, Toronto, Canada. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

Yiyang Zou and Shaoyun Long. 2019. [Effects of graded reading on middle school students' reading comprehension](#). *English Language and Literature Studies*, 9:39.

A Selected Instruction-Tuned LLMs

We investigate 8 instruction-tuned LLMs with diverse architectures and sizes, including open-weight and closed-weight models. Taking into account hardware resource limitations, we select open-source models within the maximum limit of 8 billion parameters.

Flan-T5 (Chung et al., 2022) is an advanced open-source instruction-fine-tuned language model built on top of T5 (Raffel et al., 2020). In our study, we used four Flan-T5 variants: Small¹¹, Base¹², Large¹³, and XL¹⁴ with 80M, 250M, 780M, and 3B

¹¹<https://huggingface.co/google/flan-t5-small>

¹²<https://huggingface.co/google/flan-t5-base>

¹³<https://huggingface.co/google/flan-t5-large>

¹⁴<https://huggingface.co/google/flan-t5-xl>

parameters, respectively. These models were fine-tuned across 1,836 fine-tuning tasks from instruction datasets including T0-SF (Sanh et al., 2022), Muffin (Wei et al., 2022), and Natural Instructions V2 datasets (Wang et al., 2022).

Mistral-7B-Instruct is another open-source instruction-fine-tuned LLM created by fine-tuning the Mistral 7 billion parameters model (Jiang et al., 2023) on various publicly shared instruction datasets in HuggingFace. We use the improved Instruct version 0.2 model.¹⁵

OpenChat-7B (Wang et al., 2023) is a collection of open-source language models, fine-tuned with the Conditioned-RLFT (C-RLFT) approach using mixed-quality data. For this study, we used the OpenChat-3.5 7B parameter model built on top of Mistral-7B.¹⁶

Llama-3-instruct-8B (AI@Meta, 2024) represents the latest advancement in Meta's instruction-tuned language models series. For this task, we use the 8B version finetuned using supervised fine-tuning (SFT) and reinforcement learning with human feedback (RLHF) on publicly available instruction datasets, along with over 10 million human-annotated training data.¹⁷

OpenAI-GPT in this study we use GPT-3.5-turbo and GPT-4-turbo from OpenAI. For GPT-3.5-turbo we used the version fine-tuned with proprietary training data up to Sep 2021 with a 4K context window.¹⁸ For GPT-4-turbo we used the version fine-tuned with proprietary training data up to April 2023 with a 128K context window.¹⁹

B CEFR Levels Descriptors

Table 9 outlines the descriptors for each CEFR level, based on the overall reading comprehension scale used in the communicative language activities of the CEFR Descriptor Scheme.²⁰

¹⁵<https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.2>

¹⁶https://huggingface.co/openchat/openchat_3.5

¹⁷<https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>

¹⁸[gpt-3.5-turbo-instruct: https://platform.openai.com/docs/models/gpt-3-5-turbo](https://platform.openai.com/docs/models/gpt-3-5-turbo)

¹⁹[gpt-4-1106-preview: https://platform.openai.com/docs/models/gpt-4-turbo-and-gpt-4](https://platform.openai.com/docs/models/gpt-4-turbo-and-gpt-4)

²⁰<https://www.coe.int/en/web/common-european-framework-reference-languages/cefr-descriptors>

CEFR levels	Descriptors
A1	Can understand very short, simple texts a single phrase at a time, picking up familiar names, words and basic phrases and rereading as required.
A2	Can understand short, simple texts containing the highest frequency vocabulary, including a proportion of shared international vocabulary items.
B1	Can read straightforward factual texts on subjects related to their field of interest with a satisfactory level of comprehension.
B2	Can read with a large degree of independence, adapting style and speed of reading to different texts and purposes, and using appropriate reference sources selectively. Has a broad active reading vocabulary, but may experience some difficulty with low-frequency idioms.
C1	Can understand a wide variety of texts including literary writings, newspaper or magazine articles, and specialised academic or professional publications, provided there are opportunities for rereading and they have access to reference tools. Can understand in detail lengthy, complex texts, whether or not these relate to their own area of speciality, provided they can reread difficult sections.
C2	Can understand a wide range of long and complex texts, appreciating subtle distinctions of style and implicit as well as explicit meaning. Can understand virtually all types of texts including abstract, structurally complex, or highly colloquial literary and non-literary writings.

Table 9: CEFR descriptors, describing readers’ overall comprehension abilities for each CEFR level.

C Prompts

Figure 4 illustrates the four prompts provided for LLMs RCTS.

D Supervised Baseline Training and Inference Settings

D.1 Supervised Baseline Training Dataset

This corpus was generated by applying the CEFR classifier (Arase et al., 2022) to the sentence alignments from the OneStopEnglish (Vajjala and Lučić, 2018) dataset and Wiki-Manual (aligned, partialAligned) splits (Jiang et al., 2020). Then we select only sentence pairs with different CEFR levels and swap their source-target roles when the target sentence has a higher CEFR level. Additionally, we ensured there was no overlap between sentences used in testing and those extracted from the Wiki-Manual. This process resulted in the selection of 2557 sentence pairs — 1,255 from OneStopEnglish and 1,302 from the Wiki-Manual. 20.94% of sentence Alignments from OneStopEnglish and 44.75% of sentence alignments from Wiki-Manual were removed because the source-simple pairs had been classified at the same CEFR level. Moreover, one sentence alignment was removed from Wiki-Manual as it overlapped with the test set. Then we split the dataset into 80% training and 20% validation sets. Table 10 provides the distribution of sentence target levels in both data splits.

	A1	A2	B1	B2	C1	Total
Train	96	473	1264	202	10	2045
Validation	24	118	316	51	3	512

Table 10: Distribution of sentence target levels in training and validation sets.

D.2 Supervised Baseline Training

For the supervised baseline model training, the batch size was 16, and the optimizer used was Adam (Kingma and Ba, 2017) with a learning rate of $3e - 5$, weight decay of 0.1 and 383 warm up steps. The fine-tuning continued for 10 epochs, and a checkpoint with the highest SARI (Xu et al., 2016) score on the validation set was chosen as the final model.

D.3 Inference Settings

We perform inference for open-weight LLMs and the supervised baseline model on NVIDIA GeForce RTX 4090 GPU, with 24 GB of memory, through the Transformers library (Wolf et al., 2020). For closed-weight models, we use OpenAI APIs for inference. In terms of the generation hyper-parameters, we specify the model to generate maximum output length of 100 tokens, and implement Nucleus Sampling (Holtzman et al., 2020) with a probability of 0.9, a temperature of 1.0. The remaining inference hyper-parameters are set to the default values in Hugging Face and OpenAI APIs.

To ensure robustness and account for variability

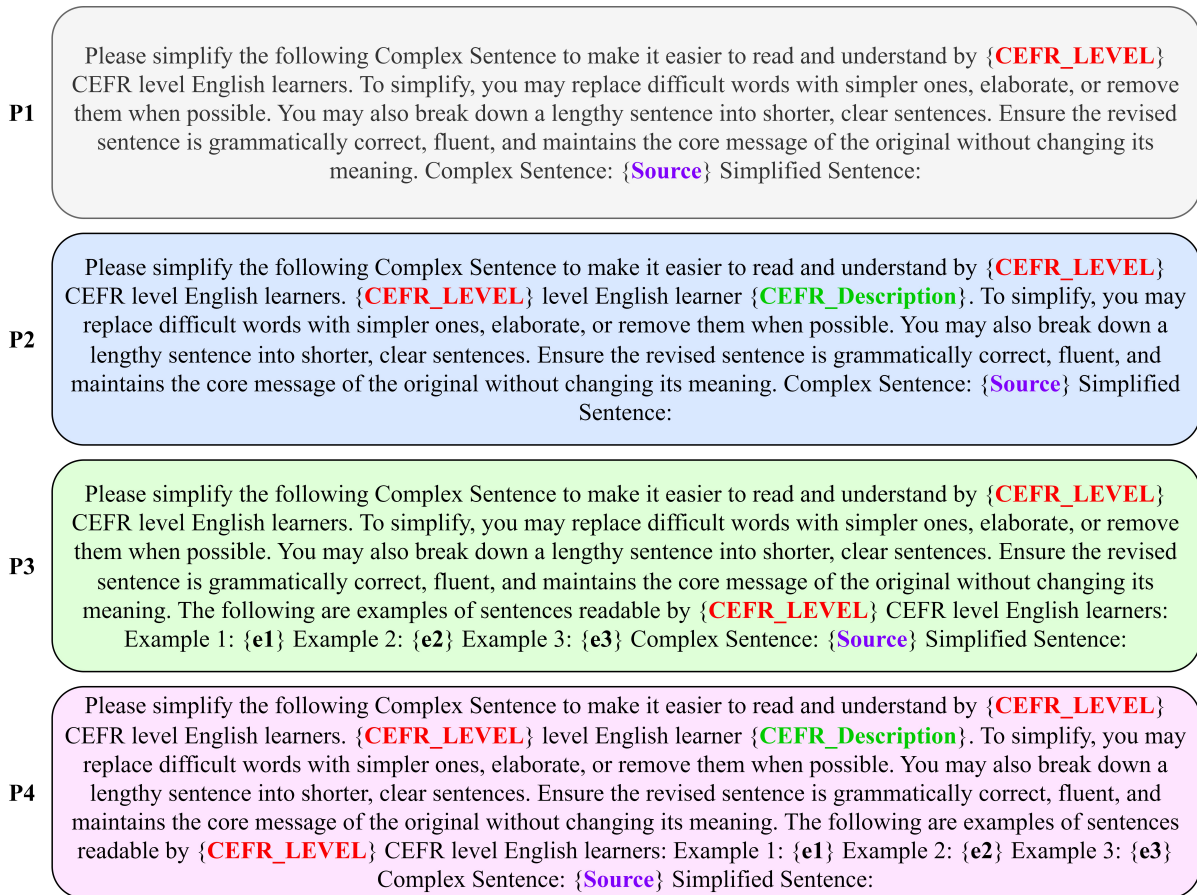


Figure 4: Prompts provided to instruction-tuned LLMs.

in generation, we perform each inference run with three different random seeds. After that, we calculate the average and standard deviation results for all metrics.

E CEFR Level Assessment Performance

Table 11 provides a comparative analysis of CEFR level assessment performance on the publicly available CEFR-SP test set from SCoRE and Wiki-Auto. This comparison compares various LLMs including the GPT-4 model, OpenChat-3.5, and Llama-3-8B-Instruct, as well as the README sentence-level CEFR level estimator (Naous et al., 2023), and the CEFR sentence-based estimator used in this study Arase et al. (2022). The results indicate that the performance of these models was lower than that achieved by the model developed by Arase et al. (2022). Moreover, Tables 12, 13 and 14 show the prompt provided to the GPT-4, Llama-3-8B-Instruct and OpenChat-3.5 models for the assessment, respectively.

Model	A1	A2	B1	B2	C1	C2	Average
GPT4	0.60	0.73	0.63	0.77	0.67	0.18	0.60
OpenChat-3.5	0.46	0.69	0.78	0.53	0.07	0.00	0.42
Llama-3-8B-Instruct	0.05	0.07	0.76	0.58	0.00	0.00	0.24
README	0.63	0.69	0.59	0.53	0.36	0.50	0.55
Arase et al. (2022)	0.82	0.83	0.87	0.87	0.89	0.93	0.87

Table 11: Comparison of Macro-F1 scores (%) across CEFR levels for various LLMs, the README estimator, and the CEFR sentence-based estimator used in this study, evaluated on the publicly available CEFR-SP test set from SCoRE and Wiki-Auto.

F Fluency Evaluation

Following (Krishna et al., 2020; Kumar et al., 2021), we evaluate the fluency of generated text using a RoBERTa-large model that has been fine-tuned on the CoLA corpus (Warstadt et al., 2018). This model evaluates each generated output for grammatical correctness. The final fluency score is then determined by averaging these evaluations across all outputs, providing an overall measure of the text’s grammatical quality. Table 15 shows that all models achieve high fluency scores, in-

dependently of the prompt used. Most of them (i.e. GPT-4, GPT-3.5-turbo-instruct, Llama-3-8B-Instruct, and Openchat) achieve near-perfect fluency scores. Flan-T5 models also obtain high fluency, with scores between 0.89 and 0.99.

G Meaning Preservation Supplemental Results

Table 16 presents the BERTScore and AlignScore results obtained by all selected models in our test set, with the different designed prompts.

BERTScore (Zhang et al., 2020): We compute the BERTScore precision and F1 scores for the predictions against the source sentences. The BERTScore is calculated using Huggingface’s implementation with `lang="en"` and `rescale_with_baseline=True`.

AlignScore (Zha et al., 2023): We compute the AlignScore for the predictions against the source sentences. The score is calculated using the original implementation <https://github.com/yuh-zha/AlignScore> with the following settings: `model='roberta-base'` and `evaluation_mode='nli'`.

H Human Evaluation Details

H.1 Readability Assessment

The human readability assessment was conducted by three evaluators with ESL experience, each having varying levels of experience: one native speaker with 1 year of experience, one non-native speaker with 3 years of experience, and one native speaker with 25 years of experience. To ensure that only qualified evaluators who are familiar with, and understand the CEFR descriptor scheme are included, we implemented a qualification task. This task comprised 30 sentences annotated with CEFR levels from the CEFR-SP dataset (Arase et al., 2022). Evaluators had to achieve at least 80% accuracy on this set to qualify. We measured inter-annotator agreement using two metrics: Percentage of Majority Agreement (i.e., the proportion of instances that had a clear majority vote) and Krippendorff’s alpha (ordinal level). The evaluators achieved 96.66% and 88.88% Majority Agreement in the qualification and rating tasks, respectively. Krippendorff’s alpha indicated a high level of agreement ($\alpha = 0.85$) in the qualification task and moderate agreement ($\alpha = 0.43$) in the rating task. Additionally, each evaluator’s agreement with the annotated examples in the qualification task was

measured using Krippendorff’s alpha ($\alpha = 0.86$, $\alpha = 0.86$, $\alpha = 0.87$), indicating strong agreement. Figure 5 shows screenshots of the human readability assessment.

H.2 Meaning Evaluation

The human meaning evaluation was conducted by three fluent English speakers: two native speakers and one non-native but fluent speaker. To ensure that only evaluators who are qualified and understand the scoring scheme are included, we implemented a qualification task. This task consisted of 20 sentence pairs with substantial agreement ($\alpha = 0.79$) among three of the authors. Evaluators were required to achieve at least ($\alpha = 0.75$) average agreement with the authors on this set. The selected participants achieved Krippendorff’s alpha scores of .79, .82, and .82, based on an ordinal level of measurement. We also determined inter-annotator agreement using Percentage of Majority Agreement and Krippendorff’s alpha (ordinal level). The evaluators achieved 90.0% and 82.2% Majority Agreement in the qualification and rating tasks, respectively. Krippendorff’s alpha indicated a high level of agreement ($\alpha = 0.86$) in the qualification task and moderate agreement ($\alpha = 0.31$) in the rating task. Figure 6 and Figure 7 display screenshots of the human meaning evaluation and the evaluation criteria provided to the evaluators.

What is the lowest CEFR level that a student should have to understand "The students celebrated the Japanese Women's Team win and cleaned up before they went home.?"

A A1
 B A2
 C B1
 D B2
 E C1
 F C2

Could "The students celebrating the Japanese Women's Team win stayed up late into the night but cleaned up their mess before they left ." be rewritten to the target level A1 ?*

A Yes
 B No

Figure 5: Screenshot of the Human Readability Assessment

2 → Source Sentence: "Opiuroids in general are mostly scavengers or detritivores ." *

Instructions for Evaluation:

1. Read the source sentence carefully.
2. You can use additional resources (e.g., dictionaries, encyclopedias) if any concepts or terms are unclear.
3. Evaluate the two simplified sentences produced by different systems independently.
4. Rate the two simplified sentences below on a scale from 0 to 2, following the guidelines provided in the table.

	0 (No or Trivial Change)	0.5	1 (Nontrivial, Main Idea Preserved)	1.5	2 (Main Idea Not Preserved)
Starfish-like creatures usually eat dead plants and animals.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Opiuroids are mostly scavengers or eat dead things.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Score	Change Type	Criteria
0	No or Trivial Change, Maintain the Original Idea	<ul style="list-style-type: none"> There is no new information included in the simplified text, or the new information is trivial (e.g., a single insignificant word) No information is removed from the complex sentence, or what is removed is very minor and insignificant to the main idea of the complex sentence. There is no altered information in the simplified text at all.
1	Nontrivial, Main Idea Preserved	<ul style="list-style-type: none"> New information is introduced in the simplified text, but this information does not introduce a new main idea. Instead, this information merely supports the same main idea expressed in the complex sentence (e.g., elaboration, example). Information is removed from the complex text, but its removal does not obfuscate the main idea of the complex text. There exists a piece of altered information in the simplified sentence, but the main idea is still intact.
2	Main Idea Not Preserved	<ul style="list-style-type: none"> A new main idea is introduced in the simplified text. Information critical to the main idea of the complex text is removed. The main idea of the complex text is altered in the simplified text.

Figure 6: Screenshot of the Human Meaning Evaluation

Role	Content
System	<p>Assess the CEFR level (A1, A2, B1, B2, C1, C2) required for an English learner to read and comprehend the provided sentence. Then, return the CEFR level in JSON format {'cefr_level': '<level>'}</p> <p>A1: Can understand very short, simple texts a single phrase at a time, picking up familiar names, words, and basic phrases and rereading as required.</p> <p>A2: Can understand short, simple texts containing the highest frequency vocabulary, including a proportion of shared international vocabulary items.</p> <p>B1: Can read straightforward factual texts on subjects related to their field of interest with a satisfactory level of comprehension.</p> <p>B2: Can read with a large degree of independence, adapting style and speed of reading to different texts and purposes, and using appropriate reference sources selectively. Has a broad active reading vocabulary, but may experience some difficulty with low-frequency idioms.</p> <p>C1: Can understand in detail lengthy, complex texts, whether or not these relate to their own area of speciality, provided they can reread difficult sections.</p> <p>C2: Can understand virtually all types of texts including abstract, structurally complex, or highly colloquial literary and non-literary writings.</p>
User	Sentence: [Sentence]

Table 12: Prompt provided to the GPT-4 model to assess readability assessment performance.

Role	Content
System	<p>Assess the CEFR level (A1, A2, B1, B2, C1, C2) required for an English learner to read and comprehend the provided sentence. Then, ONLY return the determined CEFR level.</p> <p>A1: Can understand very short, simple texts a single phrase at a time, picking up familiar names, words, and basic phrases and rereading as required.</p> <p>A2: Can understand short, simple texts containing the highest frequency vocabulary, including a proportion of shared international vocabulary items.</p> <p>B1: Can read straightforward factual texts on subjects related to their field of interest with a satisfactory level of comprehension.</p> <p>B2: Can read with a large degree of independence, adapting style and speed of reading to different texts and purposes, and using appropriate reference sources selectively. Has a broad active reading vocabulary, but may experience some difficulty with low-frequency idioms.</p> <p>C1: Can understand in detail lengthy, complex texts, whether or not these relate to their own area of speciality, provided they can reread difficult sections.</p> <p>C2: Can understand virtually all types of texts including abstract, structurally complex, or highly colloquial literary and non-literary writings.</p>
User	Sentence: [Sentence]. CEFR level:

Table 13: Prompt provided to the Llama-3-8B-Instruct model to assess readability assessment performance.

Role	Content
User	<p>Assess the CEFR level (A1, A2, B1, B2, C1, C2) required for an English learner to read and comprehend the provided sentence. Then, return the CEFR level.</p> <p>A1: Can understand very short, simple texts a single phrase at a time, picking up familiar names, words, and basic phrases and rereading as required.</p> <p>A2: Can understand short, simple texts containing the highest frequency vocabulary, including a proportion of shared international vocabulary items.</p> <p>B1: Can read straightforward factual texts on subjects related to their field of interest with a satisfactory level of comprehension.</p> <p>B2: Can read with a large degree of independence, adapting style and speed of reading to different texts and purposes, and using appropriate reference sources selectively. Has a broad active reading vocabulary, but may experience some difficulty with low-frequency idioms.</p> <p>C1: Can understand in detail lengthy, complex texts, whether or not these relate to their own area of speciality, provided they can reread difficult sections.</p> <p>C2: Can understand virtually all types of texts including abstract, structurally complex, or highly colloquial literary and non-literary writings.</p>
User	Sentence: [Sentence]

Table 14: Prompt provided to the OpenChat-3.5 model to assess readability assessment performance.

Models	ρ (\uparrow)	Accuracy (%) (\uparrow)	RMSE (\downarrow)	Fluency (\uparrow)	STS (%) (\uparrow)	Copies (%) (\downarrow)
<i>Baselines</i>						
COPY SUPERVISED	0	24.15	2.45	0.99	100	100
	0.01 \pm 0.02	39.43 \pm 1.44	2.08 \pm 0.04	0.98 \pm 0.01	84.51 \pm 0.75	12.69 \pm 0.91
<i>P1: target level</i>						
Flan-T5-small	-0.01 \pm 0.02	34.69 \pm 0.76	2.21 \pm 0.01	0.91 \pm 0.02	86.70 \pm 0.58	12.97 \pm 1.12
Flan-T5-base	-0.01 \pm 0.01	28.01 \pm 0.32	2.38 \pm 0.02	0.97 \pm 0.01	95.41 \pm 0.41	30.15 \pm 2.24
Flan-T5-large	-0.00 \pm 0.01	28.04 \pm 0.49	2.36 \pm 0.01	0.98 \pm 0.01	96.40 \pm 0.25	25.68 \pm 1.58
Flan-T5-xl	-0.00 \pm 0.01	27.96 \pm 0.89	2.38 \pm 0.01	0.99 \pm 0.01	97.17 \pm 0.20	30.09 \pm 1.39
GPT-3.5-Turbo-Instruct	0.05 \pm 0.01**	39.87 \pm 1.03	2.01 \pm 0.03	0.99 \pm 0.00	85.12 \pm 0.36	0.00 \pm 0.00
GPT-4-Turbo	<u>0.11\pm0.01**</u>	<u>51.68\pm0.79</u>	<u>1.67\pm0.02</u>	0.99 \pm 0.01	76.96 \pm 0.33	0.00 \pm 0.00
OpenChat_3.5	-0.01 \pm 0.02	42.87 \pm 0.29	1.96 \pm 0.01	0.98 \pm 0.01	85.35 \pm 0.28	0.11 \pm 0.14
Llama-3-8B-Instruct	0.07 \pm 0.00**	44.81 \pm 1.57	1.89 \pm 0.03	0.99 \pm 0.01	84.32 \pm 0.23	0.58 \pm 0.37
Mistral-7B-Instruct-v0.2	0.06 \pm 0.01**	41.40 \pm 0.83	1.98 \pm 0.01	0.98 \pm 0.01	83.63 \pm 0.20	0.00 \pm 0.00
<i>P2: target level + description</i>						
Flan-T5-small	0.02 \pm 0.03	33.25 \pm 0.92	2.23 \pm 0.03	0.91 \pm 0.01	86.99 \pm 0.40	13.13 \pm 1.29
Flan-T5-base	-0.01 \pm 0.01	27.60 \pm 0.70	2.39 \pm 0.02	0.97 \pm 0.01	95.54 \pm 0.45	29.54 \pm 2.27
Flan-T5-large	0.01 \pm 0.02	28.29 \pm 0.88	2.36 \pm 0.01	0.98 \pm 0.01	96.19 \pm 0.35	25.96 \pm 3.25
Flan-T5-xl	0.01 \pm 0.01	28.49 \pm 1.03	2.37 \pm 0.02	0.99 \pm 0.01	96.95 \pm 0.33	28.43 \pm 1.05
GPT-3.5-Turbo-Instruct	0.08 \pm 0.02**	41.51 \pm 0.80	1.98 \pm 0.02	0.99 \pm 0.01	84.95 \pm 0.23	0.00 \pm 0.00
GPT-4-Turbo	<u>0.18\pm0.02**</u>	<u>52.04\pm0.69</u>	1.66\pm0.01	0.99 \pm 0.01	77.24 \pm 0.25	0.00 \pm 0.00
OpenChat_3.5	0.03 \pm 0.02	45.08 \pm 1.08	1.90 \pm 0.02	0.98 \pm 0.01	83.04 \pm 0.41	0.17 \pm 0.11
Llama-3-8B-Instruct	0.17 \pm 0.01**	46.33 \pm 1.01	1.83 \pm 0.01	0.99 \pm 0.01	83.29 \pm 0.35	0.33 \pm 0.19
Mistral-7B-Instruct-v0.2	0.09 \pm 0.01**	41.65 \pm 0.88	1.96 \pm 0.01	0.98 \pm 0.01	83.45 \pm 0.31	0.00 \pm 0.00
<i>P3: target level + examples</i>						
Flan-T5-small	-0.00 \pm 0.01	38.79 \pm 0.89	2.12 \pm 0.02	0.89 \pm 0.01	78.54 \pm 0.49	6.48 \pm 1.28
Flan-T5-base	-0.01 \pm 0.01	29.01 \pm 0.90	2.35 \pm 0.01	0.95 \pm 0.01	93.62 \pm 0.32	22.94 \pm 1.42
Flan-T5-large	-0.01 \pm 0.00	28.43 \pm 0.40	2.36 \pm 0.01	0.98 \pm 0.01	95.75 \pm 0.17	23.30 \pm 3.17
Flan-T5-xl	-0.01 \pm 0.01	28.10 \pm 0.55	2.37 \pm 0.01	0.99 \pm 0.00	96.86 \pm 0.26	31.78 \pm 1.35
GPT-3.5-Turbo-Instruct	0.08 \pm 0.00**	42.12 \pm 0.67	1.98 \pm 0.01	0.99 \pm 0.01	84.97 \pm 0.55	0.00 \pm 0.00
GPT-4-Turbo	<u>0.15 \pm 0.00**</u>	<u>50.73\pm1.17</u>	<u>1.69\pm0.01</u>	0.99 \pm 0.00	78.13 \pm 0.19	0.00 \pm 0.00
OpenChat_3.5	0.06 \pm 0.01*	43.11 \pm 0.68	1.94 \pm 0.01	0.98 \pm 0.01	84.63 \pm 0.27	0.42 \pm 0.27
Llama-3-8B-Instruct	0.11 \pm 0.02**	46.35 \pm 0.73	1.84 \pm 0.01	0.99 \pm 0.00	83.46 \pm 0.38	0.61 \pm 0.31
Mistral-7B-Instruct-v0.2	0.08 \pm 0.01**	40.23 \pm 0.64	2.03 \pm 0.01	0.98 \pm 0.00	85.66 \pm 0.28	0.00 \pm 0.00
<i>P4: target level + description + examples</i>						
Flan-T5-small	-0.02 \pm 0.00	38.51 \pm 1.02	2.13 \pm 0.01	0.89 \pm 0.01	77.83 \pm 0.70	7.70 \pm 1.08
Flan-T5-base	-0.01 \pm 0.02	30.34 \pm 1.31	2.33 \pm 0.03	0.96 \pm 0.01	91.63 \pm 0.28	19.59 \pm 1.21
Flan-T5-large	-0.01 \pm 0.01	29.01 \pm 0.53	2.35 \pm 0.01	0.98 \pm 0.01	95.51 \pm 0.27	22.58 \pm 1.71
Flan-T5-xl	-0.01 \pm 0.00	28.70 \pm 0.69	2.36 \pm 0.02	0.98 \pm 0.01	96.57 \pm 0.24	28.90 \pm 1.84
GPT-3.5-Turbo-Instruct	0.08 \pm 0.01**	43.34 \pm 1.15	1.93 \pm 0.02	0.99 \pm 0.01	84.40 \pm 0.23	0.00 \pm 0.00
GPT-4-Turbo	0.18 \pm 0.01**	52.45\pm0.96	1.66\pm0.01	0.99 \pm 0.01	77.81 \pm 0.27	0.00 \pm 0.00
OpenChat_3.5	0.06 \pm 0.01*	44.00 \pm 0.80	1.91 \pm 0.02	0.99 \pm 0.01	83.74 \pm 0.38	0.25 \pm 0.18
Llama-3-8B-Instruct	0.19 \pm 0.01**	48.21 \pm 1.23	1.80 \pm 0.02	0.99 \pm 0.01	82.86 \pm 0.16	0.50 \pm 0.19
Mistral-7B-Instruct-v0.2	0.09 \pm 0.00**	42.01 \pm 0.76	1.98 \pm 0.01	0.98 \pm 0.01	84.20 \pm 0.21	0.00 \pm 0.00

Table 15: Results on the test set, across the four prompts. Metrics: Spearman’s ρ , Accuracy, RMSE, Fluency, STS, and Copy Percentage. Boldface indicates the best overall performance across all prompts, while underlined results represent the best performance for each individual prompt. Significance: * ($p < 0.01$), ** ($p < 0.001$).

Models	BERTScore <i>precision</i> (%) (↑)	BERTScore <i>f1</i> (%) (↑)	AlignScore (%) (↑)
<i>Baselines</i>			
COPY	100	100	99.88
SUPERVISED	74.99±0.89	65.15±0.98	80.56±1.09
<i>P1: target level</i>			
Flan-T5-small	73.76±0.83	66.13±0.75	73.57±1.69
Flan-T5-base	87.50±0.77	84.25±0.84	87.09±1.44
Flan-T5-large	88.56±0.42	84.91±0.57	93.10±0.42
Flan-T5-xl	90.29±0.29	87.29±0.44	94.06±0.31
GPT-3.5-Turbo-Instruct	62.78±0.55	58.38±0.48	83.62±0.68
GPT-4-Turbo	49.23±0.40	46.99±0.43	71.87±0.60
OpenChat_3.5	65.05±0.23	62.04±0.23	83.00±0.55
Llama-3-8B-Instruct	65.65±0.47	57.93±0.43	80.46±0.74
Mistral-7B-Instruct-v0.2	55.73±0.41	53.03±0.39	79.91±0.39
<i>P2: target level + description</i>			
Flan-T5-small	73.62±0.57	66.41±0.68	72.98±1.45
Flan-T5-base	87.54±0.91	84.49±1.10	86.77±1.88
Flan-T5-large	88.48±0.63	84.42±0.87	92.92±0.64
Flan-T5-xl	89.90±0.44	86.62±0.62	94.00±0.22
GPT-3.5-Turbo-Instruct	62.68±0.33	58.20±0.39	83.84±0.66
GPT-4-Turbo	49.76±0.48	47.01±0.37	71.98±0.84
OpenChat_3.5	61.55±0.53	58.01±0.52	81.56±0.37
Llama-3-8B-Instruct	64.11±0.52	56.32±0.50	80.46±0.36
Mistral-7B-Instruct-v0.2	56.29±0.95	52.97±0.66	79.55±0.84
<i>P3: target level + examples</i>			
Flan-T5-small	62.99±0.45	54.53±0.39	59.03±1.53
Flan-T5-base	84.72±0.40	80.85±0.48	82.53±0.93
Flan-T5-large	87.59±0.30	83.46±0.32	92.51±0.58
Flan-T5-xl	90.25±0.30	86.90±0.28	94.14±0.21
GPT-3.5-Turbo-Instruct	63.32±0.52	58.33±0.52	83.56±0.58
GPT-4-Turbo	51.37±0.54	48.44±0.40	73.77±0.75
OpenChat_3.5	64.14±0.43	60.64±0.41	82.32±0.51
Llama-3-8B-Instruct	63.87±0.46	55.95±0.46	80.37±0.82
Mistral-7B-Instruct-v0.2	59.61±0.51	56.36±0.55	82.29±0.66
<i>P4: target level + description + examples</i>			
Flan-T5-small	61.84±1.14	53.75±1.15	58.23±2.36
Flan-T5-base	81.41±0.63	76.86±0.61	78.20±0.78
Flan-T5-large	87.08±0.58	82.73±0.67	92.09±0.92
Flan-T5-xl	89.51±0.45	85.82±0.45	93.85±0.50
GPT-3.5-Turbo-Instruct	63.13±0.33	57.32±0.25	83.76±0.62
GPT-4-Turbo	51.06±0.21	47.59±0.25	73.03±0.59
OpenChat_3.5	63.33±0.72	59.17±0.61	82.05±0.31
Llama-3-8B-Instruct	63.80±0.60	55.28±0.47	80.25±1.04
Mistral-7B-Instruct-v0.2	58.66±0.25	54.59±0.28	81.29±0.38

Table 16: BERTScore and AlignScore results on the test set, across the four prompts.

Score	Change Type	Criteria
0	No or Trivial Change, Maintain the Original Idea	<ul style="list-style-type: none"> • There is no new information included in the simplified text, or the new information is trivial (e.g., a single insignificant word) • No information is removed from the complex sentence, or what is removed is very minor and insignificant to the main idea of the complex sentence. • There is no altered information in the simplified text at all.
1	Nontrivial, Main Idea Preserved	<ul style="list-style-type: none"> • New information is introduced in the simplified text, but this information does not introduce a new main idea. Instead, this information merely supports the same main idea expressed in the complex sentence (e.g., elaboration, example). • Information is removed from the complex text, but its removal does not obfuscate the main idea of the complex text. • There exists a piece of altered information in the simplified sentence, but the main idea is still intact.
2	Main Idea Not Preserved	<ul style="list-style-type: none"> • A new main idea is introduced in the simplified text. • Information critical to the main idea of the complex text is removed. • The main idea of the complex text is altered in the simplified text.

Figure 7: Criteria Provided to the Evaluators for the Meaning Evaluation