

MCTS: A Multi-Reference Chinese Text Simplification Dataset

Ruining Chong¹, Luming Lu¹, Liner Yang^{1†}, Jinran Nie^{1*},
Zhenghao Liu², Shuo Wang³, Shuhan Zhou^{1*}, Yaoxin Li^{1*} and Erhong Yang¹

¹School of Information Science, Beijing Language and Culture University

²Department of Computer Science and Technology, Northeastern University

³Department of Computer Science and Technology, Institute for AI, Tsinghua University
lineryang@gmail.com

Abstract

Text simplification aims to make the text easier to understand by applying rewriting transformations. There has been very little research on Chinese text simplification for a long time. The lack of generic evaluation data is an essential reason for this phenomenon. In this paper, we introduce MCTS, a multi-reference Chinese text simplification dataset. We describe the annotation process of the dataset and provide a detailed analysis. Furthermore, we evaluate the performance of several unsupervised methods and advanced large language models. We additionally provide Chinese text simplification parallel data that can be used for training, acquired by utilizing machine translation and English text simplification. We hope to build a basic understanding of Chinese text simplification through the foundational work and provide references for future research. All of the code and data are released at <https://github.com/blcuicall/mcts/>.

Keywords: Chinese Text Simplification, Multi-Reference, Language Resource

1. Introduction

The task of text simplification aims to make the text easier to understand by performing multiple rewriting transformations. It can provide reading assistance for children (Kajiwara et al., 2013), non-native speakers (Paetzold, 2016), and people with language disorders (Carroll et al., 1998; Paetzold, 2016; Evans et al., 2014). Moreover, text simplification can also be used as a data augmentation method to benefit downstream natural language processing (NLP) tasks (Van et al., 2021).

For a long time, the research of text simplification systems mainly depends on large-scale parallel corpora for training, such as WikiLarge (Zhang and Lapata, 2017) and Newsela (Xu et al., 2015). But due to the limitation of existing data in language and domain, recent work on text simplification systems has started to focus on unsupervised methods and achieves good results (Surya et al., 2019; Kumar et al., 2020; Martin et al., 2022), which makes it possible to build Chinese text simplification systems independent of large-scale parallel corpora. In this case, evaluating the Chinese text simplification systems becomes a problem that needs to be solved. On the other hand, large language models have the ability to solve various NLP tasks (Thoppilan et al., 2022; Chowdhery et al., 2022). Recently a series of large language models represented by ChatGPT¹ perform well on many tasks (Qin et al.,

2023; Jiao et al., 2023; Bang et al., 2023). In English text simplification, Feng et al. (2023) find that large language models outperform state-of-the-art methods and are judged to be on par with human annotators. Nevertheless, whether these models can achieve the same excellent results in Chinese text simplification remains unclear.

To solve these problems, we introduce MCTS, a multi-reference dataset for the evaluation of Chinese text simplification models. MCTS comprises 3,615 human simplifications associated with 723 original sentences selected from the Penn Chinese Treebank (Xue et al., 2005) (5 simplifications per original sentence), and it encompasses a diverse array of rewriting transformations. To our knowledge, it is the first published multi-reference Chinese text simplification evaluation dataset. We aim for this dataset to serve as a gauge for assessing the current state of Chinese text simplification and to furnish benchmarks that will inform and guide future scholarly inquiries.

We have also designed several simple, unsupervised Chinese text simplification methods and evaluated them using our newly proposed dataset. These approaches are intended to establish a foundation for subsequent research endeavors. Moreover, we assess the capabilities of advanced large language models, such as ChatGPT, in performing Chinese text simplification. Our findings indicate that these large language models surpass our established unsupervised baselines. Nonetheless, there remains a discernible disparity in quality compared to human-crafted simplifications. In summary, our contributions are listed below:

[†]Corresponding author: Liner Yang.

^{*}The work was completed during their studies at Beijing Language and Culture University.

¹<https://chat.openai.com/chat>

- We manually annotate a dataset that can be used for the evaluation of Chinese text simplification. It is a multi-reference dataset that encompasses multiple types of rewriting transformations. To our knowledge, it is the first published multi-reference Chinese text simplification evaluation dataset.
- We provide several text features and conduct a detailed analysis of the dataset, which could help to understand the characteristics of human Chinese text simplification.
- On the proposed dataset, we evaluate the performance of several unsupervised methods and advanced large language models, which could serve as the baselines for future research.
- As a by-product, we additionally release the Chinese text simplification parallel dataset suitable for training purposes, obtained through machine translation and the English text simplification model. The dataset comprises 691,474 sentence pairs. To our knowledge, it is the largest Chinese text simplification training dataset.

2. Related Work

2.1. Evaluation Data for English Text Simplification

In the field of English text simplification, the research on system evaluation has accumulated rich experience and results, providing valuable references and insights for the construction of Chinese text simplification evaluation data.

Early evaluation data for English text simplification mainly consist of sentence pairs obtained from English Wikipedia and Simple English Wikipedia through automatic sentence alignment. However, the Simple English Wikipedia was argued to contain many inadequate or inaccurate simplifications (Yasseri et al., 2012; Xu et al., 2015). Furthermore, evaluation data with only a single reference is insufficient to ensure reliability and comprehensiveness in highly subjective text simplification tasks.

For the above reasons, Xu et al. (2016) introduced TurkCorpus, a multi-reference dataset for evaluating English text simplification. They first collected 2,359 original sentences from English Wikipedia and then obtained 8 manual reference references for every original sentence via crowdsourcing. The dataset can be used for evaluation metrics requiring multiple references, such as BLEU (Papineni et al., 2002) and SARI (Xu et al., 2016). However, the rewriting transformations involved in TurkCorpus are simplistic. Annotators

were asked to simplify a sentence mainly by lexical paraphrasing but without deleting content or splitting the sentences. Another multi-reference dataset for English text simplification, HSplit (Sulem et al., 2018), uses the same original sentences from the TurkCorpus test set but only focuses on sentence splitting.

The simplicity of the rewriting transformations hinders research into the model’s ability to perform a wide range of rewriting operations. To solve this problem, Alva-Manchego et al. (2020) created the ASSET dataset. Using the same origin sentences, They extended TurkCorpus through crowdsourcing. They defined a wider range of rewriting transformations, with lexical paraphrasing (lexical simplification and reordering), sentence splitting, and compression (deleting unimportant information). Through experimentation, it was discovered that the ASSET dataset features more abstractive simplifications, which are deemed simpler compared to those found in other evaluation corpora. Now ASSET has been adopted as a standard dataset for evaluating English text simplification systems.

Similar to ASSET, MCTS is a dataset with multiple references and multiple rewriting transformations. To our knowledge, this dataset is the largest and most comprehensive evaluation dataset for the Chinese text simplification task to date.

2.2. Unsupervised Text Simplification

Unsupervised text simplification methods do not require aligned complex-simple sentence pairs. After Surya et al. (2019) introduced the first unsupervised neural text simplification system by importing adversarial and denoising auxiliary losses, the field of unsupervised text simplification has rapidly developed.

Currently, the research on unsupervised text simplification mainly adopts two strategies. One approach is to train models using text mining or automatically generated pseudo-data pairs, and the other is to adopt cross-language transfer learning methods for text simplification.

MUSS (Martin et al., 2022) is an outstanding representative of the first methods class. It could effectively generate simplified text by using text paraphrasing data obtained through data mining and realize the difficulty control with ACCESS (Martin et al., 2020) during inference. Brunato et al. (2016) mined pseudo-parallel simplification data from a wide amount of texts extracted from the web using metrics such as semantic similarity and readability. Lu et al. (2021) constructed pseudo-parallel data for text simplification from large-scale machine translation data pairs.

The second category of methods, cross-language transfer learning, explores how to transfer

Original	为适应大西南进出口物资快速增长的需要，北部湾沿海四市开始了新一轮建港热潮。 In order to adapt to the rapid growth of import and export materials in the southwest, four coastal cities in the Beibu Gulf have started a new wave of port construction.
Reference	大西南进口和出口物资的需要快速增长，北部湾沿海四座城市兴起了新一轮港口建设。 The demand for imported and exported materials in the southwest has grown rapidly. Four coastal cities in the Beibu Gulf have begun a new round of port construction.
Original	中国又一条煤炭运输大通道——连接天津蓟县与天津港之间的蓟港铁路日前破土动工。 Another major coal transportation corridor in China - the Jigang Railway connecting Tianjin Jixian County and Tianjin Port - has recently broken ground.
Reference	蓟港铁路连接天津蓟县和天津港，用于煤炭运输，前几天开始建造。 Jigang Railway connects Tianjin Jixian County and Tianjin Port for coal transportation, and construction began a few days ago.
Original	按设计，“进步M-24”号载重飞船可同轨道站在无人操纵的情况下进行自动对接。 According to the design, the "Progress M-24" heavy-duty spacecraft can automatically dock with the orbital station under unmanned control.
Reference	按照设计，无人操控的“进步M-24”号飞船可以自动对接轨道站。 According to the design, the unmanned "Progress M-24" spacecraft can automatically dock with the orbital station.

Table 1: Examples of simplifications collected for MCTS

the knowledge of English text simplification to other languages. Mallinson et al.(2020) used a unique model architecture and multi-task learning strategy to realize cross-language zero-shot learning for German text simplification.

In this study, we construct several unsupervised models for Chinese text simplification using pseudo-data construction and cross-language transfer learning as the main ideas, which can provide important baseline references for future research work.

3. Creating MCTS

In this section, we describe MCTS in more detail. In section 3.1, we introduce the preparation of original sentences. In section 3.2, we introduce the annotation process of MCTS.

3.1. Data Preparation

We choose Penn Chinese Treebank (CTB) as the source of the original sentence in the dataset. CTB is a phrase structure tree bank built by the University of Pennsylvania. It includes Xinhua news agency reports, government documents, news magazines, broadcasts, interviews, online news, and logs. We concentrate on news content because of its wide accessibility and significant influence. The rich corpus in the news area also makes it conducive to further work advancement.

We initially employ a filter based on the HSK² average vocabulary difficulty, setting a threshold for the average vocabulary level to filter out simple sentences, ensuring the selected original

sentences possess sufficient complexity. Subsequently, through manual screening of the remaining sentences, we selected 723 sentences to serve as the original sentences for our study.

3.2. Annotation Process

MCTS stands as a fully manually annotated evaluation dataset. Here is a detailed description of the annotation process.

Annotator Recruitment All the annotators we recruited are native Chinese speakers and are undergraduate or graduate students in school. Most of them have a background in linguistics or computer science. All annotators need to attend a training course and take the corresponding Qualification Test (see more details below) designed for our task. Only those who have passed the Qualification Test could enter the Annotation Round.

Simplification Instructions We provide the exact instructions for annotators for the Qualification Test and the Annotation Round. In the instructions, we define three types of rewriting transformations.

- *Paraphrasing*: Replacing complex words or phrases with simple formulations.
- *Compression*: Deleting repetitive or unimportant information from the sentence.
- *Structure Changing*: Modifying complex sentence structures into simple forms.

Compared to rewriting transformations involved in ASSET, we have made the following two

²<https://www.chinesetest.cn/>

changes : (1) We employ *Paraphrasing*, which extends beyond the narrow scope of mere *Lexical Paraphrasing*. The paraphrasing transformation in Chinese is much more flexible than in English. It not only includes synonym substitution but also covers the interpretation of complex phrases or idioms. (2) We replaced *Sentence Splitting* with *Structural Changing*. In Chinese text simplification, it's more common to convert complex sentences into compound sentences with several clauses separated by commas rather than splitting them into separate sentences with periods. The narrowly defined Sentence Splitting is not well-suited to this situation. The category of *Structure Changing* encapsulates a variety of shifts, including but not limited to the conversion from passive to active voice, the simplification of double negatives into affirmative constructions, the substitution of synonyms, the interpretation of complex phrases or idioms, and the rephrasing of rhetorical questions into declarative statements.

These two changes ensure that the defined rewriting transformations are both non-redundant and exhaustive, and they align with the practical context of Chinese text simplification.

For every rewriting transformation, we provide several examples. Annotators could decide for themselves which types of rewriting to execute in any given original sentence.

Qualification Test At this stage, we provided 20 sentences to be simplified. Annotators needed to simplify these sentences according to the instructions given. We checked all submissions to filter out annotators who could not perform the task correctly. Of the 73 people who initially registered, only 35 passed the Qualification Test (48%) and worked on the task.

Annotation Round Annotators who passed the Qualification Test had access to this round. To facilitate annotating work, we provided a platform that can display the difficulty level of words in a text. We collected 5 simplifications for each of the 723 original sentences. Table 1 presents a few examples of simplifications in MCTS, together with their English translation.

4. Dataset Analysis

Following ASSET (Alva-Manchego et al., 2020), we analyze various text features within the MCTS dataset to examine the nature of its simplifications comprehensively. We provide a detailed introduction in sections 4.1 and 4.2. In section 4.3, we report a detailed comparison between MCTS and another recently proposed dataset.

4.1. Text Features

We calculate several low-level features for all simplification examples to measure the rewriting transformations included in MCTS. These features are listed below.

- Number of sentence splits: The difference between the number of sentences in the simplification and the number of sentences in the original sentence.
- Compression level: The number of characters in the simplification divided by the number of characters in the original sentence.
- Replace-only Levenshtein distance: The character-level Levenshtein distance (Levenshtein et al., 1966) for replace operations only divided by the length of the shorter string in the original sentence and simplification. As described in ASSET, ignoring insertions and deletions can make this feature independent of compression level and serve as a proxy for measuring the lexical paraphrases of the simplification.
- Proportion of words deleted, words added and words reordered: The number of words deleted/reordered from the original sentence divided by the number of words in the original sentence; and the number of words that were added to the original sentence divided by the number of words in the simplification.
- Lexical complexity score ratio: We compute the score as the mean squared lexical difficulty level in HSK. The ratio is then the value of this score on the simplification divided by that of the original sentence, which can be considered as an indicator of lexical simplification.
- Dependency tree depth ratio: The ratio of the depth of the dependency parse tree of the simplification relative to that of the original sentence. Following ASSET (Alva-Manchego et al., 2020), we perform parsing using spaCy³. This feature can reflect structural simplicity to a certain extent.

4.2. Results and Analysis

The density of all these features is shown in Figure 1. We can see that sentence splitting operation appears not frequently on MCTS. By observing the data, we believe this is due to the characteristics of the Chinese language. Compound sentences are commonly used in Chinese, and one sentence consists of two or more independent clauses. As we

³<https://github.com/explosion/spaCy>

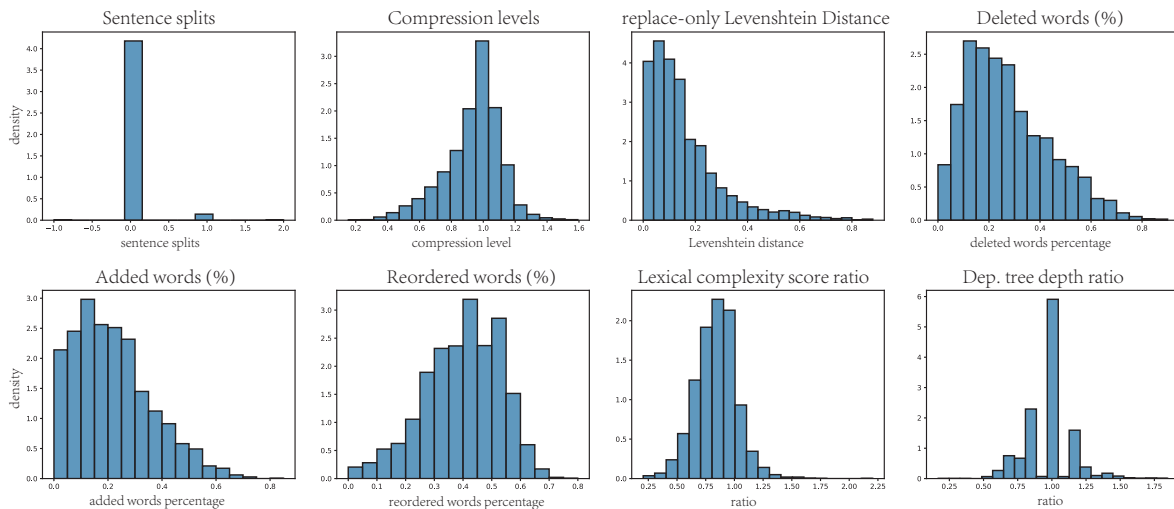


Figure 1: Density of text features in simplifications from MCTS

mentioned in section 3.2, during the simplification, annotators tend to rewrite a complex sentence with nested clauses into compound sentences rather than multiple simple sentences. So this is not to say that Chinese text simplification rarely involves sentence structure change, but that the way of structural change is not limited to sentence splitting.

Although we have introduced compression as a rewriting transformation in the simplification instructions, the compression ratio is not too concentrated on the side less than 1.0. The reason is that, on the one hand, the annotators tend to retain as much semantic information as possible, and on the other hand, more characters may be added when paraphrasing.

By analyzing replace-only Levenshtein distance, we can see that the simplifications in MCTS have a considerable degree of paraphrasing the input as simplifications are distributed at all levels. Regarding the distribution of deleted, added, and reordered words, we can find that the peaks all occur at positions greater than 0.0. This further reveals the plentiful rewriting operations contained in MCTS.

In terms of lexical complexity, we can clearly see the high density of ratios less than 1.0, indicating that simplification has significantly lower lexical complexity compared to the original sentence. Some instances have a lexical complexity ratio greater than 1.0, which may be due to deleted simple words in the process of sentence compression.

Finally, the dataset shows a high density of a 1.0 ratio in dependency tree depth. This may indicate that significant structural changes were not made.

4.3. Comparison with CSS

We have also noticed another recent work by Yang et al. (2023a), which also attempts to construct a

multi-reference Chinese text simplification evaluation dataset independently. However, their work is not sufficiently comprehensive. The created CSS dataset lacks a validation set, and each original sentence has only two references, constraining its universality. We report a detailed comparison of two datasets in Tables 2 and 3. We use the sentence-transformers toolkit (Reimers and Gurevych, 2020) to calculate semantic similarity and use the python-Levenshtein package to calculate edit distance.

In table 2, we observe that MCTS significantly surpasses CSS in terms of data volume, both in the number of references (5 vs. 2) and the quantity of source sentences (723 vs. 383). By analyzing the sentence length, we can find that MCTS is slightly longer than CSS in terms of the original sentence length. From the length changes before and after simplification, it can be seen that MCTS is more inclined to perform vocabulary level replacement or deletion operations compared to CSS. Finally, by analyzing semantic similarity and editing distance, we find that MCTS has made greater simplification operations than CSS.

In table 3, we measure the difference between two datasets in lexical complexity using the HSK Level (Kong et al., 2022) metric. We provide a detailed introduction to this metric in section 5.2 In simple terms, the higher the proportion of words in levels 1-3 (7+), the easier (more challenging) the text is understood. We can find that the original text of MCTS uses more difficult vocabulary compared to CSS. While observing the simplified version, MCTS is simpler than CSS. This is enough to demonstrate that MCTS has undergone a more in-depth simplification.

Datasets	NRef.	Sent.	Length (Char)		Length (Token)		SS	LD
			Ori.	Ref.	Ori.	Ref.		
CSS	2	383	49.49	47.3	31.37	28.06	93.03	17.15
MCTS	5	723	50.13	50.05	32.6	30.04	90.13	27.33

Table 2: Comparison of the MCTS and CSS datasets. NRef. refers to the number of simplified sentences corresponding to each original sentence. Sent. refers to the number of original sentences. Length (Char) and Length (Token) refer to the average character length and the average token length of each sentence, respectively. SS refers to Semantic Similarity. LD refers to Levenshtein Distance. Ori. represents the score of the original sentences, and Ref. represents the score of the reference sentences.

Datasets	Ori.		Refs.	
	L1-3(%)	L7+(%)	L1-3(%)	L7+(%)
CSS	40.93	44.74	44.3	41.79
MCTS	39.64	45.05	45.85	39.65

Table 3: Comparison of MCTS and CSS on HSK Level Scores. L1-3 and L7+, respectively, represent the proportion of vocabulary at levels 1-3 and levels 7+.

5. Experiment

In order to measure the development status of Chinese text simplification and provide references for future research, we conduct a series of experiments on the proposed MCTS.

5.1. Methods

We attempt several unsupervised Chinese text simplification methods and large language models and provide their results on MCTS. The first three are unsupervised methods that utilize automatic machine translation technology. We use Google Translator⁴ to translate. These unsupervised methods can be used as the baselines for future work.

We divide all the 723 sentences in MCTS into two subsets: 366 for validation and 357 for testing the Chinese text simplification models. For those model training methods, the valid set is used to select the best-performing checkpoints, and the final score is calculated on the test set.

Direct Back Translation As high-frequency words tend to be generated in the process of neural machine translation (Lu et al., 2021), back translation is a potential unsupervised text simplification method. We translate the original Chinese sentences into English and then translate them back to obtain simplified results. We choose English as the bridge language because of the rich bilingual translation resources between Chinese and English.

Translated Wiki-Large Translating existing text simplification data into Chinese is a simple way to construct pseudo data. We translated English sentence pairs in Wiki-Large into Chinese sentence pairs and used them to train a Chinese version BART-based (Shao et al., 2021) model as one of our baselines.

We have followed the concept of previous work (Martin et al., 2022) and uniformly described similar methods for constructing pseudo data as unsupervised methods.

Cross-Lingual Pseudo Data In addition to the above two methods, we also design a way to construct pseudo data for Chinese text simplification, which can leverage the knowledge from English text simplification models. As shown in Figure 2, we first collect a large amount of Chinese sentence data, for example, the People’s Daily Corpus. Then, we translate these sentences into English and simplify them using existing English text simplification models. Finally, we translate the simplified English sentences back into Chinese and align them with the original Chinese sentences to obtain parallel data. Due to the tendency of machine translation to use simple vocabulary and sentence structures, there may be a decrease in text difficulty during both the translation and English simplification stages.

To ensure data quality, we filter the obtained parallel data from three aspects: simplicity, fluency, and semantic retention. For simplicity, we calculate the average lexical difficulty level for both the original sentence and the simplified sentence using HSK. Only when the difficulty level of the simplified sentence is significantly reduced compared to the original sentence, this parallel sentence pair will be retained. For fluency, we calculate the perplexity

⁴<https://translate.google.com/>




Figure 2: Pseudo data acquisition process

for the simplified sentences and filter out sentences above the preset threshold. For semantic retention, we use sentence-transformers toolkit (Reimers and Gurevych, 2020) to calculate the semantic similarity between the original sentence and simplified sentence, and also filter out sentences that exceed the preset threshold.

By filtering more than 4 million sentence pairs, we finally get 691,474 sentence pairs that could be used for training. By randomly selecting and manually observing 50 entries from these data, we find that the generated sentences had significant simplification operations compared to the original sentences. Using the filtered data, we train a Chinese BART-base model.

Large Language Models We choose two advanced large language models to conduct experiments: *gpt-3.5-turbo* and *text-davinci-003*. Both of them are based on GPT-3.5. The former is the most capable GPT-3.5 model and is optimized for chatting. The latter is the previous model, which can execute any language task according to instructions. We translate the simplification prompt used by Feng et al. (2023) as our prompt. More details about the prompt can be found in Table 4. Recent studies have shown that the in-context few-shot method can enhance the performance of LLM (Brown et al., 2020; Wei et al., 2022). To further investigate the ability of LLM on Chinese text simplification tasks, we tested the 3-shot performance of the *gpt-3.5-turbo* method by adding 3 instances at the end of our prompt, which is named *gpt-3.5-turbo-few-shot*.

Our Prompt
我想让你把我的复杂句子替换成简单的句子。你要保持句意不变，但使句子更简单。
复杂句：{Complex Sentence}
简单句：{Simplified Sentence(s)}

Table 4: Prompt for Chinese text simplification

5.2. Automatic Metrics

Following previous work, we choose three metrics for evaluation: SARI (Xu et al., 2016), BLEU (Pa-

pineni et al., 2002) and HSK Level (Kong et al., 2022).

SARI SARI (Xu et al., 2016) is a commonly used evaluation metric for text simplification. Comparing system outputs to multiple simplification references and the original sentences, SARI calculates the mean of the n-gram F1 scores of *add*, *keep*, and *delete*. In our experiment, we tokenize sentences using Stanford CoreNLP⁵ and use the EASSE toolkit⁶ (Alva-Manchego et al., 2019) to calculate SARI.

BLEU BLEU (Bilingual Evaluation Understudy) (Papineni et al., 2002) was initially used to evaluate the quality of machine translation. By calculating the N-gram and counting the times that can be matched, BLEU can reflect the closeness between system outputs and references. Just like calculating SARI, we use the EASSE toolkit to calculate the BLEU score.

HSK Level In order to measure the complexity of Chinese sentences, we import HSK Level. HSK is the Chinese proficiency test designed for non-native speakers. It provides a vocabulary of nine levels from easy to difficult, which includes over 6,000 entries, covering the most common Chinese words. In our proposed dataset (source sentences), 85.14% of the words could be found in the HSK list, with many of the rest being proper nouns. Following previous work (Kong et al., 2022), we count the proportion of words at levels 1-3 and 7+ in system outputs. The higher the proportion of words in levels 1-3 (7+), the easier (more challenging) the outputs are understood. Our specific implementation of this metric is the same as Kong et al. (2022).

5.3. Human Evaluation

In order to obtain more comprehensive evaluation results, we further conduct human evaluation. Following the previous work (Dong et al., 2019; Kumar

⁵<https://github.com/stanfordnlp/CoreNLP>

⁶<https://github.com/feralvam/easse>

Method	SARI \uparrow	BLEU \uparrow	L1-3 (%) \uparrow	L7+ (%) \downarrow
Source	22.37	84.75	40.24	44.90
Gold Reference	48.11	61.62	46.25	39.50
Direct Back Translation	40.37	48.72	39.19	45.44
Translated Wiki-Large	28.30	82.20	40.32	44.92
Cross-Lingual Pseudo Data	38.49	<u>63.06</u>	41.57	44.24
gpt-3.5-turbo-few-shot	43.95	56.46	44.44	40.67
gpt-3.5-turbo	<u>42.39</u>	49.22	<u>43.68</u>	<u>41.29</u>
text-davinci-003	37.97	36.18	38.80	45.32

Table 5: The automatic evaluation results on the test set of MCTS. Source refers to the score calculated by taking the original sentence directly as the output of the system. \uparrow The higher, the better. \downarrow The lower, the better. **Bold** means the best result, and underline means the second-best result.

et al., 2020), we evaluate the Chinese text simplification systems on three dimensions:

- Fluency: Is the output grammatical?
- Adequacy: How much meaning from the original sentence is preserved?
- Simplicity: Is the output simpler than the original sentence?

We provide simplifications generated by different systems for the recruited volunteers. We ask the volunteers to fill out a five-point Likert scale (1 is the worst, 5 is the best) about these simplifications for each dimension. Additionally, following Feng et al.’s work (2023), we measure the volunteers’ subjective choices by ranking the simplifications to focus on actual usage rather than evaluation criteria.

6. Results

In this section, we report the evaluation results of the test set of MCTS.

6.1. Results of Automatic Evaluations

The results of automatic evaluations are shown in Table 5. In addition to the model results, we also report the score of the source and gold reference. The source scores are calculated on the unedited original sentence. We calculate the gold reference scores by evaluating each reference against all others in a leave-one-out scenario and then averaging the scores.

To our surprise, direct back translation gets the best SARI score among the unsupervised methods. But regarding HSK level, the performance of direct back translation is not good, even worse than the source. We find that many rewrite operations were generated during the back translation

process, which is highly correlated with the SARI score. However, due to the lack of control over simplicity, direct back translation is more like a sentence paraphrase method than text simplification. This may be why it performs poorly on the HSK level.

The translated Wiki-Large method gets the best BLEU score but the lowest SARI score among all methods. In fact, the system output has hardly changed compared to the original sentence. As the unedited source gets the highest BLEU score of 84.75, we believe the single BLEU value cannot be used as an excellent indicator of text simplification. Because there is a significant overlap between the original sentence and the references. As for the poor performance of the translated Wiki-Large method, we believe it is due to the large amount of noise contained in the translated training data.

The SARI score of the cross-lingual pseudo data method is 38.49, which is between the other two unsupervised methods. But it performs better on the HSK level than the other two. This may be because the model learned simplification knowledge from pseudo data that was transferred from the English text simplification model.

In terms of the large language models, the gpt-3.5-turbo significantly performs better than text-davinci-003. However, although the gpt-3.5-turbo-few-shot under the few-shot setting can further improve performance, it is still insufficient compared to the gold reference.

6.2. Results of Human Evaluations

We conduct human evaluations on three representative methods, namely direct back translation, cross-lingual pseudo data, and gpt-3.5-turbo. We recruit three volunteers to conduct the evaluation. All of them have a background in linguistics. We select 30 sentences from the test set of MCTS for each volunteer and provide them with the original

Method	Simplicity \uparrow	Fluency \uparrow	Adequacy \uparrow	Avg. \uparrow	Rank \downarrow
Direct Back Translation	3.42 \pm 0.87	4.36 \pm 0.78	4.72 \pm 0.56	4.17	2.88
Cross-Lingual Pseudo Data	4.11 \pm 0.81	<u>4.46</u> \pm 0.65	3.88 \pm 0.96	4.15	2.86
gpt-3.5-turbo	<u>4.17</u> \pm 0.89	<u>4.46</u> \pm 0.70	<u>4.43</u> \pm 0.78	<u>4.35</u>	<u>2.29</u>
Gold Reference	4.20 \pm 1.08	4.68 \pm 0.55	4.31 \pm 0.93	4.40	1.97

Table 6: The human evaluation results. Avg. means the average score of fluency, adequacy, and simplicity. Rank means the subjective ranking for the simplifications. Each annotator annotates all 30 sentences. The values in the table after the \pm symbol are the standard deviation.

sentences and the outputs of these methods. For the convenience of comparison, a randomly selected reference for each sentence is additionally provided. Volunteers are asked to rate the simplification of these four groups. The results of the human evaluation are shown in Table 6.

We can see that the gold reference gets the best average score and rank. It is significantly superior to the output results of other simplification systems. For detail, it gets the best simplicity score of 4.20 and the best fluency score of 4.68. Due to some degree of sentence compression, it does not achieve the best adequacy score but only 4.31. As for the direct back translation method, despite its excellent performance in adequacy, it achieves the lowest simplicity score due to the lack of corresponding control measures. On the contrary, the cross-lingual pseudo-data method performs well in terms of simplicity but does not perform well in terms of adequacy. Because it tends to perform more sentence compression, which removes lots of semantic information. These two unsupervised methods get a similar average score and rank.

The gpt-3.5-turbo gets the second-best results among all metrics. By analyzing the average score and the rank score, we can find that it is significantly better than the two unsupervised simplification methods. But compared to the gold reference, there is still a certain gap. Our experiment has shown that there is still room for further improvement in the large language model’s Chinese text simplification ability.

By comparing the results of automatic evaluation and manual evaluation, we can find that the SARI and HSK Level results fit well with the manual assessment rankings. Methods that performed better on these two metrics also received higher rankings in the human evaluation, while the BLEU scores were almost independent of simplification quality. Thus we consider SARI and the HSK Level to be valid in assessing the quality of Chinese text simplification.

Besides, several features can be observed in the statistical results of Simplicity, Fluency, and Adequacy:

- (1) Gold reference has the highest standard deviation

in the simplicity statistics. This suggests that manual simplification may not always tend to be as simple as possible, and fluctuates more in simplicity compared to automatic simplification methods.

- (2) On fluency, gold reference receives the highest score and lowest standard deviation, suggesting that gold reference not only strives to keep sentences fluent but also maintains a high level of consistency.

- (3) Gold reference and cross-lingual pseudo data obtained large standard deviations and low scores in adequacy. This suggests that these two methods may tend to perform more deletion operations.

7. Conclusion

In this paper, we introduce the MCTS, a human-annotated dataset for the validation and evaluation of Chinese text simplification systems. It is a multi-reference dataset that contains multiple rewriting transformations. By calculating the low-level features for simplifications, we have shown the rich simplifications in MCTS, which may be of great significance for understanding the simplification and readability of Chinese text from a linguistic perspective. Furthermore, we test the Chinese text simplification ability of some unsupervised methods and advanced large language models using the proposed dataset. We find that even advanced large language models are still inferior to human simplification under the zero-shot and few-shot settings. Finally, we hope our work can motivate the development of Chinese text simplification systems and provide references for future research.

8. Acknowledgments

This work was supported by the MOE (Ministry of Education in China) Project of Humanities and Social Sciences (No. 23YJCZH264), the funds of Research Project of the National Language Commission (No. ZDA145-17) and Research Funds of Beijing Language and Culture University (No. 23YCX131).

9. Bibliographical References

- Fernando Alva-Manchego, Louis Martin, Carolina Scarton, and Lucia Specia. 2019. *EASSE: Easier automatic sentence simplification evaluation*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 49–54, Hong Kong, China. Association for Computational Linguistics.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Love-nia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023. A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Dominique Brunato, Andrea Cimino, Felice Dell’Orletta, and Giulia Venturi. 2016. Paccss-it: A parallel corpus of complex-simple sentences for automatic text simplification. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 351–361.
- John Carroll, Guido Minnen, Yvonne Canning, Siobhan Devlin, and John Tait. 1998. Practical simplification of english newspaper text to assist aphasic readers. In *Proceedings of the AAI-98 Workshop on Integrating Artificial Intelligence and Assistive Technology*, pages 7–10. Association for the Advancement of Artificial Intelligence.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Yue Dong, Zichao Li, Mehdi Rezagholizadeh, and Jackie Chi Kit Cheung. 2019. Editnts: An neural programmer-interpreter model for sentence simplification through explicit editing. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3393–3402.
- Richard Evans, Constantin Orasan, and Iustin Dor-nescu. 2014. An evaluation of syntactic simplification rules for people with autism. Association for Computational Linguistics.
- Yutao Feng, Jipeng Qiang, Yun Li, Yunhao Yuan, and Yi Zhu. 2023. Sentence simplification via large language models. *arXiv preprint arXiv:2302.11957*.
- Biyang Guo, Xin Zhang, Ziyuan Wang, Minqi Jiang, Jinran Nie, Yuxuan Ding, Jianwei Yue, and Yupeng Wu. 2023. How close is chatgpt to human experts? comparison corpus, evaluation, and detection. *arXiv preprint arXiv:2301.07597*.
- Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. 2023. Is chatgpt a good translator? a preliminary study. *arXiv preprint arXiv:2301.08745*.
- Tomoyuki Kajiwara, Hiroshi Matsumoto, and Kazuhide Yamamoto. 2013. Selecting proper lexical paraphrase for children. In *Proceedings of the 25th Conference on Computational Linguistics and Speech Processing (ROCLING 2013)*, pages 59–73.
- Cunliang Kong, Yun Chen, Hengyuan Zhang, Liner Yang, and Erhong Yang. 2022. Multitasking framework for unsupervised simple definition generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5934–5943.
- Dhruv Kumar, Lili Mou, Lukasz Golab, and Olga Vechtomova. 2020. Iterative edit-based unsupervised sentence simplification. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7918–7928.
- Vladimir I Levenshtein et al. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710. Soviet Union.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- Xinyu Lu, Jipeng Qiang, Yun Li, Yunhao Yuan, and Yi Zhu. 2021. An unsupervised method for building sentence simplification corpora in multiple languages. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 227–237.

- Jonathan Mallinson, Rico Sennrich, and Mirella Lapata. 2020. Zero-shot crosslingual sentence simplification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5109–5126.
- Louis Martin, Éric Villemonte De La Clergerie, Benoît Sagot, and Antoine Bordes. 2020. Controllable sentence simplification. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4689–4698.
- Louis Martin, Angela Fan, Éric Villemonte De La Clergerie, Antoine Bordes, and Benoît Sagot. 2022. Muss: Multilingual unsupervised sentence simplification by mining paraphrases. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 1651–1664.
- Gustavo Henrique Paetzold. 2016. *Lexical simplification for non-native english speakers*. Ph.D. thesis, University of Sheffield.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiaao Chen, Michihiro Yasunaga, and Diyi Yang. 2023. Is chatgpt a general-purpose natural language processing task solver? *arXiv preprint arXiv:2302.06476*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners.
- Nils Reimers and Iryna Gurevych. 2020. Making monolingual sentence embeddings multilingual using knowledge distillation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 4512–4525.
- Luz Rello, Clara Bayarri, Azuki Górriz, Ricardo Baeza-Yates, Saurabh Gupta, Gaurang Kanvinde, Horacio Saggion, Stefan Bott, Roberto Carlini, and Vasile Topac. 2013. Dyswebxia 2.0! more accessible text for people with dyslexia. In *Proceedings of the 10th International Cross-Disciplinary Conference on Web Accessibility*, pages 1–2.
- Yunfan Shao, Zhichao Geng, Yitao Liu, Junqi Dai, Fei Yang, Li Zhe, Hujun Bao, and Xipeng Qiu. 2021. Cpt: A pre-trained unbalanced transformer for both chinese language understanding and generation. *arXiv preprint arXiv:2109.05729*.
- Sai Surya, Abhijit Mishra, Anirban Laha, Parag Jain, and Karthik Sankaranarayanan. 2019. Unsupervised neural text simplification. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2058–2068.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*.
- Hoang Van, Zheng Tang, and Mihai Surdeanu. 2021. How may i help you? using neural text simplification to improve downstream nlp tasks. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4074–4080.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Shiping Yang, Renliang Sun, and Xiaojun Wan. 2023a. A new dataset and empirical study for sentence simplification in chinese. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8306–8321.
- Xianjun Yang, Yan Li, Xinlu Zhang, Haifeng Chen, and Wei Cheng. 2023b. Exploring the limits of chatgpt for query or aspect-based text summarization. *arXiv preprint arXiv:2302.08081*.
- Taha Yasseri, András Kornai, and János Kertész. 2012. A practical approach to language complexity: a wikipedia case study. *PloS one*, 7(11):e48386.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

10. Language Resource References

- Alva-Manchego, Fernando and Martin, Louis and Bordes, Antoine and Scarton, Carolina and

- Sagot, Benoît and Specia, Lucia. 2020. *ASSET: A Dataset for Tuning and Evaluation of Sentence Simplification Models with Multiple Rewriting Transformations*.
- Sulem, Elinor and Abend, Omri and Rappoport, Ari. 2018. *BLEU is Not Suitable for the Evaluation of Text Simplification*.
- Wei Xu, Chris Callison-Burch, and Courtney Napoles. 2015. Problems in current text simplification research: New data can help.
- Xu, Wei and Napoles, Courtney and Pavlick, Ellie and Chen, Quanze and Callison-Burch, Chris. 2016. *Optimizing statistical machine translation for text simplification*. MIT Press.
- Xue, Naiwen and Xia, Fei and Chiou, Fu-Dong and Palmer, Marta. 2005. *The penn chinese treebank: Phrase structure annotation of a large corpus*. Cambridge University Press.
- Zhang, Xingxing and Lapata, Mirella. 2017. *Sentence Simplification with Deep Reinforcement Learning*. Association for Computational Linguistics.