



HAL
open science

Aligning sentence simplification with ESL learner's proficiency for language acquisition

Guanlin Li, Yuki Arase, Noel Crespi

► **To cite this version:**

Guanlin Li, Yuki Arase, Noel Crespi. Aligning sentence simplification with ESL learner's proficiency for language acquisition. Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL), Apr 2025, Albuquerque, New Mexico, United States. 10.48550/arXiv.2502.11457 . hal-04926354

HAL Id: hal-04926354

<https://hal.science/hal-04926354v1>

Submitted on 3 Feb 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Aligning Sentence Simplification with ESL Learner’s Proficiency for Language Acquisition

Guanlin Li^{1,2}, Yuki Arase², Noël Crespi¹

¹Samovar, Telecom SudParis, Institut Polytechnique de Paris, France

²School of Computing, Institute of Science Tokyo, Japan

{guanlin_li, noel.crespi}@telecom-sudparis.eu, arase@c.titech.ac.jp

Abstract

Text simplification is crucial for improving accessibility and comprehension for English as a Second Language (ESL) learners. This study goes a step further and aims to facilitate ESL learners’ language acquisition by simplification. Specifically, we propose simplifying complex sentences to appropriate levels for learners while also increasing vocabulary coverage of the target level in the simplifications. We achieve this without a parallel corpus by conducting reinforcement learning on a large language model. Our method employs token-level and sentence-level rewards, and iteratively trains the model on its self-generated outputs to guide the model to search for simplification hypotheses that satisfy the target attributes. Experiment results on CEFR-SP and TurkCorpus datasets show that the proposed method can effectively increase the frequency and diversity of vocabulary of the target level by more than 20% compared to baseline models, while maintaining high simplification quality. ¹

1 Introduction

Controlled text simplification considers audience-targeted attributes when generating simplified texts, so that the generated texts do not only meet the criteria of simplicity, but also preserve desired attributes for the targeted audiences. Recent studies on controlled text simplification aimed to help reading comprehension for language learners and employed school grade levels annotated in the training corpus as the simplification target (Scarton and Specia, 2018; Sheang and Saggion, 2021; Agrawal and Carpuat, 2023) or text features (number of words, character-level Levenshtein similarity etc.) between source and target sentences (Nishihara et al., 2019; Martin et al., 2020).

Different from these studies, we aim to aid language learning and education for English as a Sec-

ond Language (ESL) learners by simplifying sentences while preserving desirable attributes for language acquisition. We use the Common European Framework of Reference for Languages (CEFR), the world standard definition of language proficiency. Our method is motivated by two classic L2 learning theories: the input hypothesis (Krashen, 1981) and frequency effect (Ellis, 2002). The input hypothesis stated that in order for the language acquisition to happen, the textual input which is either too simple or too complex for learner comprehension will not be useful for acquisition. If a learner’s current competence is i , then comprehensible input should contain both i and $(i + 1)$ content (Mitchell et al., 2019). Frequency theory holds that the frequency of the words and phrases in the input is a key determinant of acquisition, and words with higher frequency in the usage tend to be easier to acquire (Ellis and Ferreira-Junior, 2009). The key challenge here is the lack of a parallel corpus for training that should provide complex-simple sentence pairs labelled their levels. Parallel sentences of this kind are naturally scarce, and worse, annotation of difficulty levels, in particular, CEFR, is non-trivial and requires language education experts (Arase et al., 2022).

To achieve sentence simplification for aiding language learning without a parallel corpus, we propose reinforcement learning on a pre-trained large language model (LLM). Based on the aforementioned L2 learning theories, the proposed method simplifies the complex sentences to the one corresponding to the learner’s proficiency level or one level higher (i and $i + 1$ levels) and increases the coverage (frequency and diversity) of the corresponding level’s vocabulary in the generated simplifications. Specifically, we reformulate the controlled simplification task as a lookahead search problem: in the decoding step t , the model searches for the token that satisfies the target vocabulary constraint while also ensuring that future tokens

¹Codes available at <https://github.com/JumpyPizza/align-sentence-simplification-with-ESL-learner>

increase the target vocabulary coverage as much as possible, and the final hypothesis falls into the desired CEFR level. We combine a simple word-match-based heuristic with the supervised sentence-level signal to guide decoding and train the model iteratively using gradient policy optimization to memorize the search strategy that maximizes the overall reward. Remarkably, we eliminate the need for a parallel corpus by utilizing LLMs’ language generation capacity for simplification via reinforcement learning. Experimental results show that the method significantly increases the coverage and diversity of the target vocabulary in the outputs by up to 20% compared to the baselines, while maintaining high simplification quality.

Our primary contributions are twofold. First, we propose the sentence simplification method that aligns generated simplifications with ESL learners’ proficiency level on word, phrase and sentence levels and preserves attributes effective for facilitating language learning. Second, our method is easy to deploy and does not require a parallel corpus that is often expensive to create.

2 Related Work

We briefly summarize two lines of simplification methods, controlled simplification and reinforcement learning based simplification.

Controlled Simplification attaches tokens or prompts to the input to control the simplification-related attributes during generation (Yang et al., 2023; Agrawal and Carpuat, 2023; Sheang and Saggion, 2021; Martin et al., 2020, 2022; Scarton and Specia, 2018; Chi et al., 2023). While these methods learn to control levels of simplified sentences using a parallel corpus with annotated difficulties, our method controls attributes useful for language learning without a parallel corpus. As opposed to the training-time controlling, Kew and Ebling (2022) adopted FUDGE (Yang and Klein, 2021), which adjusts the logits of the text generation model during decoding using a classifier, to directly control the attribute of the simplification in the decoding time.

Reinforcement Learning based Simplification has explored controllability by defining rewards based on simplicity-related criteria (Zhang and Lapata, 2017; Guo et al., 2018; Nakamachi et al., 2020; Laban et al., 2021). The rewards for the objectives are constructed using supervised or unsupervised

evaluation metrics for simplicity, adequacy and fluency. In contrast, we aim to control attributes useful for language learning and education. Furthermore, while RL tends to suffer from unstable training and sensitivity to the choice of hyperparameters, our method achieves training stability by adopting entropy regularization in the model optimization process and introducing a dynamic reward that adjusts based on the data distribution.

3 Problem Definition

We aim to facilitate language learning by simplification targeted at ESL learners. In this study, we use CEFR levels as a representative measure for the learners’ proficiency and model the target level based on the vocabulary² (words, phrases, idioms) and sentence CEFR levels³.

Our problem is thus defined as follows. We assume that learners know their own CEFR level i . Given a sentence above the learner’s level i , we generate its simplified version that (a) contains as much vocabulary of the level i and $i + 1$ as possible, and (b) corresponds to the target (learner’s) level i at the sentence level.

3.1 Constraint Formalization

Generating simplified texts subject to vocabulary constraints can be approached as a lexical-constrained text generation task (Zetsu et al., 2022). Traditionally, lexical constraints in text generation involve a *short* list of required words, which Lu et al. (2021) expressed as a Conjunctive Normal Form (CNF), such as $(D_1 \vee D_2 \vee \dots) \wedge \dots \wedge (D_{m-1} \vee D_m)$ in which D_m stands for a single constraint, and all clauses must be satisfied, imposing hard constraints on the generation process.

In our setting, however, this formulation is no longer applicable because the vocabulary constraint is *as large as the size of the vocabulary of a specific level*. In addition, we aim to satisfy *as many clauses as possible*. Therefore, we formalize constraints as Disjunctive Normal Form (DNF), indicating words and phrases suitable for the target proficiency level: $D = (D_1) \vee (D_2 \wedge D_3 \wedge \dots) \vee \dots \vee (D_m)$, where the form stands for the word list of the target

²<https://www.englishprofile.org/wordlists/evp>

³<https://www.coe.int/en/web/common-european-framework-reference-languages/table-1-cefr-3.3-common-reference-levels-global-scale>

language level, a single D_m represents word and the conjunctive clauses represent several words, namely phrases. Notably, this form of constraints allows for the control of discontinuous phrases, which is difficult in previous methods.

3.2 Optimization Function

Based on the DNF constraints, our task imposes soft constraints that aim to include as many clauses as possible. Given the simplification hypotheses $\{\text{seq}_1, \text{seq}_2, \dots, \text{seq}_n\}$, the goal is to maximize:

$$\sum_{j=1}^m \sum_{k=1}^n \text{count}(C_j, \text{seq}_k), \quad (1)$$

where $\text{count}(C_j, \text{seq}_k)$ indicates the number of clauses C_j satisfied by seq_k . Consequently, the target during the generation process is to search for the next token that:

- simplifies the original text;
- is contained in $\exists C_i \in D$;
- leads to future tokens that satisfy C_i ;
- leads to complete phrases or phrases with slots (discontinuity) that satisfy C_i .

4 Proposed Method

To search for a hypothesis that better satisfies pre-determined constraints, some previous methods use rollout in decoding that generates partial future sequences (Chaffin et al., 2022; Lu et al., 2022). These methods become infeasible for large models due to the inefficiency of sampling in decoding time and handling the large vocabulary constraints in our task. To effectively and efficiently search for the tokens that satisfy our constraints, we instead consider sampling in the training time and formulate the lookahead search problem using RL search (Fickinger et al., 2021) (see Fig. 1).

4.1 RL Search

Consider the text generation process as a Markov Decision Process (MDP), at each timestep t of the generation, the language model observes a state s_t , which is the generated partial sequence seq_{t-1} , and takes an action a_t to choose the token from its vocabulary. When the EOS token is reached, a reward R for the generated sequence is calculated and used to update the model. In this setting, the language model is the policy function that searches for a token $v_i \in \mathcal{V}$ where \mathcal{V} is the vocabulary, and we can use any policy gradient algorithm to guide the language model to search for the generations that

Algorithm 1 Training Procedure

Input: Complex sentences;

- 1: Generate simplified texts from the complex sentences using current policy (rollout);
 - 2: Evaluate the current policy and produce rewards to guide the search;
 - 3: Optimize the policy model using the rewards
 - 4: Iteratively perform steps 1-3 till converge.
-

maximize the constraint satisfaction. Algorithm 1 indicates our training procedure.

4.2 Policy Model

The policy model generates a simplified sentence seq given a complex counterpart as a prompt pmt . The policy model is initialized from an instruction tuned language model, which unsupervisedly provides robust text simplifications (Kew et al., 2023).

By design, the rewards for the policy model across different proficiency levels are varied. For instance, given the same model response, a positive reward for C level could correspond to a negative reward for A level. Therefore, using the original language model as the backbone, we train separate copies of the policy model for A, B and C levels by adding and updating distinct LoRA parameters to the backbone parameters (Hu et al., 2022), while keeping the backbone frozen.

4.3 Reward Models

Inspired by the L2 learning theories, we design two types of rewards at lexical and sentence levels.

4.3.1 Lexical Constraint Reward

We use a simple heuristic to guide the search for generations that satisfy the lexical constraints:

$$H(\text{seq}) = \sum_{C_j \in D} r(\text{count}(C_j, \text{seq})), \quad (2)$$

where C is a clause from D , r denotes the reward according to the number of satisfactions of C in seq , and H denotes the reward score for the generated sequence seq in the current decoding step. To calculate the match counts, we remove basic stop words from the sentence after lemmatization.

As a simple baseline, we define r as a constant value 1 for word and 1.5 for phrase to encourage the model to generate more phrases and idioms. However, we found that this simple baseline is easily hacked by the model after a few steps of training,

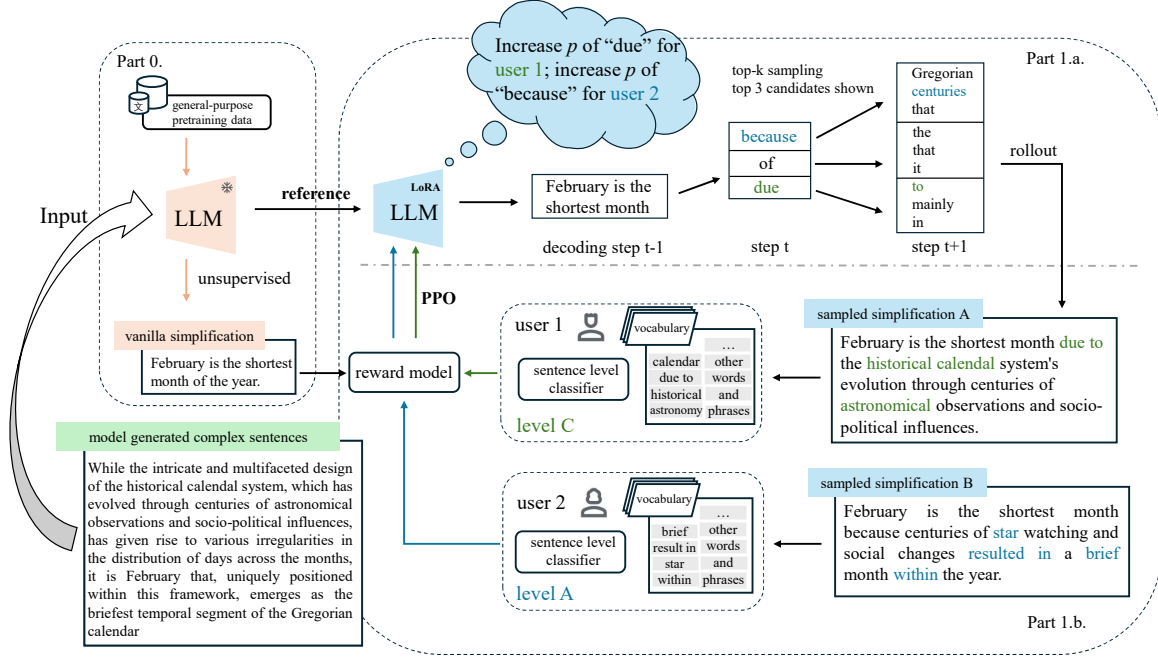


Figure 1: (better viewed in color) The overall framework of the proposed method: the simplification model is initialized from a pretrained large language model which is also used as a frozen (*) reference model to provide entropy regularization (part 0.); top-k sampling is adopted in the decoding process to sample varied simplifications for the complex sentence (part 1.a.); the generated simplifications are evaluated based on the language proficiency level (vocabulary level and sentence level) of the target audience, which is used as rewards to update the simplification model (part 1.b.) to adopt better decoding strategy.

i.e., the model only generates a limited set of frequent words that were learnt to produce rewards. To encourage the model to explore more diverse words and improve the overall coverage of target-level words in the generations, the reward should intuitively encourage maximizing the entropy for the clauses in D , so that all the clauses are evenly distributed. Accordingly, we adjust the reward r for the count of C_j as a **dynamic reward**:

$$r = \begin{cases} 1 & \text{if } 0 \leq p_j < \frac{1}{m}, \\ e^{-\alpha p_j} & \text{if } \frac{1}{m} \leq p_j \leq 1, \end{cases} \quad (3)$$

where p_j is:

$$p_j = \frac{\sum_{k=1}^n \text{count}(C_j, \text{seq}_k)}{\sum_{k=1}^n \sum_{j=1}^m \text{count}(C_j, \text{seq}_k)}, \quad (4)$$

to discourage the model from exploiting the same clause. Here, m denotes the total number of clauses in D and α is a constant to adjust the penalty degree for too frequent clauses. Eq. 4 is calculated after each epoch and the reward is adjusted accordingly. If matched clauses are above the target level, we give a constant negative score of -1 .

4.3.2 Sentence Level Reward

To go beyond words and guide the simplification model’s search for a sentence of the target level, we incorporate a sentence-level reward model by simulating human experts’ judgment for the sentence’s CEFR level. We use pairwise ranking loss to train the reward model, since the class distribution for the CEFR-SP data is imbalanced (Arase et al., 2022). The ranking loss has been shown to be able to encourage the model to only pay attention to the focused class (Henning et al., 2023), thus may mitigate the class imbalance problem.

Consequently, we construct sentence pairs prioritizing the level we focus on generating: for a collection of sentences $\mathcal{S} = \{s_1, s_2, \dots, s_n\}$, each sentence s_i is evaluated by human experts and annotated with a language level l . Given the level we want to generate, we select the sentences with the target level $\mathcal{S}_{\text{tgt}} = \{s_i \in \mathcal{S} \mid l_i = \text{level}_{\text{tgt}}\}$, and randomly sample sentences from other levels to construct a negative set $\mathcal{S}_{\text{non-tgt}} = \{s_j \in \mathcal{S} \mid l_j \neq \text{level}_{\text{tgt}}\}$. Then, we construct sentence pairs $\mathcal{P} = \{(s_i, s_j) \mid s_i \in \mathcal{S}_{\text{tgt}}, s_j \in \mathcal{S}_{\text{non-tgt}}\}$ by randomly selecting from \mathcal{S}_{tgt} and $\mathcal{S}_{\text{non-tgt}}$.

Notably, we do not require the pair to be parallel;

they just need to be at different levels. By this design, we disentangle the adequacy requirement for the simplification from the target-level search process. The former is handled by the underlying LLM, and the latter is dealt with by the reward model by level judgment.

With the constructed sentence pairs, we train a sentence-level reward model r_θ . The training objective is to minimize loss:

$$\mathcal{L}(\theta) = - \sum_{(s_i, s_j) \in \mathcal{P}} \log \sigma(r_\theta(s_i) - r_\theta(s_j)) \quad (5)$$

where σ is the sigmoid function. After training the reward model, for a generated sentence seq, we take $r_l = \sigma(r_\theta(s))$ as the reward, and use a linear combination of the lexical reward and the sentence-level reward as the overall reward:

$$R = \lambda r + \gamma r_l \quad (6)$$

4.4 Stabilized RL Training

The original instruct-tuned model is used as a frozen reference model, providing an entropy regularization for the updated policy model to ensure training stability during the search process. Specifically, the simplification seq' produced by the frozen backbone model f' is added as an entropy regularization to the overall reward:

$$R' = R - \log(p_{(f(\text{seq}|pmt)}/p_{(f'(\text{seq}'|pmt))}). \quad (7)$$

By doing so, we may keep the LLM’s strong paraphrasing ability while letting it acquire controllability in CEFR levels.

The policy model f , namely the simplification model is then updated to search for the generations that maximize the reward. In this study, we adopt Proximal Policy Optimization (Schulman et al., 2017) to update the policy model, which achieves stable training and faster convergence.

5 Experiment Settings

We aim to evaluate the effectiveness of the proposed model in generating high-quality simplifications that align with the target vocabulary and sentence-based CEFR level. This section provides details of the experiment settings.

5.1 Resource and Implementation

Sentence CEFR Level To train the sentence-level reward model, we used CEFR-SP (Arase et al., 2022), which provides labels of six CEFR levels

| | A1 | A2 | B1 | B2 | C1 | C2 |
|-------|-----|------|------|------|-----|----|
| Train | 248 | 1284 | 2479 | 2226 | 889 | 52 |
| Val | 79 | 276 | 485 | 336 | 149 | 40 |
| Test | 71 | 289 | 540 | 369 | 150 | 39 |

Table 1: Statistics on CEFR-SP w/o Newsela

for a total of $17k$ sentences annotated by experts. We only used the publicly available subset from the dataset (excluding data based on Newsela (Xu et al., 2015)), which resulted in $10k$ sentences with labels. The statistics of the dataset are described in Table 1. We fine-tuned the GPT-2 (Radford et al.) using the annotated CEFR levels.

Vocabulary List For the lexical constraint reward model, we need vocabulary lists per CEFR level. We downloaded the English Vocabulary Profile (EVP) data⁴ and used it as a dictionary of words and phrases annotated with their corresponding CEFR levels⁵. Since our goal is to generate the simplifications in i and $i + 1$ levels, we always aggregate the vocabulary lists in two levels. For clarity, we consider A1+A2, B1+B2, and C1+C2 levels. In total, we got 1076 words for A level, 3823 words for B level, 3612 words for C level.

Complex Sentence Collection We trained the policy model to iteratively learn to search for a hypothesis that maximizes rewards based on its own generations. The only requirement for our training corpus is a supply of complex sentences that warrant simplification, because sufficiently simple sentences without the need for simplification may disturb the learning. Cegin et al. (2023) showed that large language models are highly capable of paraphrasing. Following this study, we used GPT-4⁶ to synthesize complex sentences from the CEFR-SP training set to create our training corpus. We manually prepared prompts to ensure that the outputs are always at least as complex as the highest C2 level. More details are in Appendix B.

We trained separate models for A, B and C levels since different levels require different rewards (see Section 4.2). For computational efficiency, we adopted a relatively small Phi-3-mini-3b model (Abdin et al., 2024). More implementation details

⁴<https://www.englishprofile.org/wordlists/evp>

⁵EVP assigns CEFR levels to each word sense, so the same word can appear at different levels depending on its meaning. For simplicity, we did not conduct word sense disambiguation and assigned the lowest level.

⁶<https://openai.com/index/gpt-4/>

| CEFR-SP | A-Frequency | A-Diversity | B-Frequency | B-Diversity | C-Frequency | C-Diversity |
|-----------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Reference | 0.292 | 0.527 | 0.283 | 0.465 | 0.080 | 0.102 |
| phi3-3b-vanilla | 0.252 | 0.665 | 0.215 | 0.435 | 0.041 | 0.172 |
| T5+grade-A | 0.194 | 0.438 | 0.269 | 0.271 | 0.072 | 0.114 |
| FUDGE-A | 0.257 | 0.215 | 0.207 | 0.069 | 0.043 | 0.018 |
| phi3-A | 0.299 | 0.684 | 0.196 | 0.403 | 0.038 | 0.141 |
| T5+grade-B | 0.204 | 0.447 | 0.275 | 0.266 | 0.069 | 0.110 |
| FUDGE-B | 0.223 | 0.226 | 0.231 | 0.084 | 0.049 | 0.027 |
| phi3-B | 0.151 | 0.677 | 0.262 | 0.538 | 0.064 | 0.251 |
| T5+grade-C | 0.203 | 0.441 | 0.276 | 0.271 | 0.074 | 0.114 |
| FUDGE-C | 0.239 | 0.217 | 0.220 | 0.077 | 0.052 | 0.025 |
| phi3-C | 0.171 | 0.658 | 0.263 | 0.275 | 0.189 | 0.365 |

Table 2: Results on target attribute controllability on CEFR-SP-Test. For “Reference”, the frequency and diversity metrics were calculated using a subset of each grade level to show distributions in sentences of specific levels.

can be found in Appendix C.

5.2 Evaluation Datasets

To evaluate the simplification outputs, we need parallel corpora of complex and reference simple paraphrases. Below describes the resources we used for the evaluation.

CEFR-SP-Test As the formal evaluation dataset, we used CEFR-SP. We expanded its test set to be parallel because CEFR-SP is a non-parallel corpus. Specifically, we generated complex sentences for each sentence in the CEFR-SP test set using the same method described in Section 5.1. These complex sentences were input to models, and outputs were evaluated by comparing them to the original CEFR-SP sentences as references.

TurkCorpus To assess the applicability of the proposed method for a general simplification task, we also evaluated models on another widely-used dataset, TurkCorpus (Xu et al., 2016). We used the test set of the corpus, including 359 complex sentences, each has 8 human-written simplified sentences as references. It should be noted that TurkCorpus does not provide any level annotations.

5.3 Evaluation Metrics

We evaluated simplification outputs from two perspectives: **simplification quality** and **target audience attributes** by both automatic and human assessments. Simplification quality was assessed across three dimensions: **Simplicity; Fluency; and Adequacy**. As automatic metrics for simplicity, we employed LENS (Maddela et al., 2023) and SALSA (Heineman et al., 2023), which are two recently proposed model-based evaluation methods. For fluency and adequacy, we employed an instruction-tuned language model as an off-the-shelf evaluation model, which was shown to be

effective in automatic translation quality evaluations (Kocmi and Federmann, 2023). Target audience attributes were measured in terms of **target vocabulary coverage** and **sentence CEFR level**, in which vocabulary coverage includes both frequency and diversity of target vocabulary. For the evaluation of sentence CEFR-level, we used human evaluation. For more details on evaluation metrics, please refer to Appendix A.

5.4 Baselines

Overall, we choose two lines of work as the baselines for comparison.

Controlled Simplification There are limited variants in controlled simplification methods which mostly employ control tokens with supervised learning. Based on previous literature, we implemented two baselines for controlling the target level of the simplified texts: a supervised baseline of T5+grade (Scarton and Specia, 2018) that attaches CEFR levels as control tokens and an unsupervised baseline of FUDGE that uses a discriminator at decoding time (Yang and Klein, 2021).

Non-controlled Simplification The Turk corpus was used to evaluate the effectiveness of the proposed method in general simplification. As opposed to controlled simplification, this task does not consider controlling attributes, such as grade levels, during the simplification. For this line of models, we choose the following methods: DRESS (Zhang and Lapata, 2017), DMass (Zhao et al., 2018), EditNTS (Dong et al., 2019), ACCESS (Martin et al., 2020), IterativEdit (Kumar et al., 2020). We used outputs of these models shared in the EASSE package (Alva-Manchego et al., 2019). In addition, we also compare the vanilla phi3-3b instruction-tuned model as a baseline, under zero-shot setting without fine-tuning on simplification.

| TURK | A-Frequency | A-Diversity | B-Frequency | B-Diversity | C-Frequency | C-Diversity |
|-----------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Reference | 0.176 | 0.229 | 0.227 | 0.132 | 0.056 | 0.046 |
| phi3-3b-vanilla | 0.166 | 0.180 | 0.177 | 0.083 | 0.034 | 0.023 |
| T5+grade-A | 0.187 | 0.180 | 0.217 | 0.088 | 0.051 | 0.028 |
| FUDGE-A | 0.175 | 0.177 | 0.175 | 0.069 | 0.034 | 0.018 |
| phi3-A | 0.216 | 0.208 | 0.153 | 0.063 | 0.031 | 0.018 |
| T5+grade-B | 0.201 | 0.190 | 0.217 | 0.085 | 0.052 | 0.028 |
| FUDGE-B | 0.163 | 0.177 | 0.178 | 0.077 | 0.039 | 0.022 |
| phi3-B | 0.126 | 0.201 | 0.330 | 0.112 | 0.066 | 0.035 |
| T5+grade-C | 0.187 | 0.194 | 0.225 | 0.090 | 0.050 | 0.026 |
| FUDGE-C | 0.171 | 0.174 | 0.181 | 0.076 | 0.037 | 0.019 |
| phi3-C | 0.151 | 0.178 | 0.193 | 0.092 | 0.091 | 0.041 |

Table 3: Results on target attribute controllability on TurkCorpus

Complex Sentence The considerable distance, compounded by Jamie’s current condition of pregnancy, which inexorably engenders a state of increased fatigue, renders the prospect of ambulation to said location prohibitively challenging for her.

Ref. (level B) It is too far for Jamie to walk to, especially because she is pregnant and easily exhausted.

Simplifications

Level A: Jamie is *too tired to* walk far because she is pregnant.

Level B: Jamie’s pregnancy makes it very hard for her to walk to the location *due to* the long distance.

Level C: Jamie’s pregnancy leads to *fatigue*, making it hard for her to walk to the distant place.

Table 4: A randomly selected example from the simplification result of the proposed method. The target vocabulary of the corresponding level is marked in *italic* font.

6 Experiment Results

This section analyses experiment results of automatic and human evaluations, and ablation study.

6.1 Automatic Evaluation Results

Target Attributes Tables 2 and 3 show the evaluation results for the target vocabulary coverage. These results demonstrate that compared to the baseline models, the proposed model significantly increases the frequency of target vocabulary in simplified sentences while also improving vocabulary diversity. Notably, the proposed method successfully increases the frequency and diversity of A and C-level vocabulary, which should be harder than B-level due to the scarcity of level A and C samples (Arase et al., 2022).

Simplification Quality Tables 5 and 6 show the evaluation results for the simplification quality. Overall, these results indicate that our models can produce high-quality simplifications, greatly outperforming the baseline models. Remind that our model does not have reward to encourage the model to follow the adequacy requirement. We attribute these improvements to the benefits of using entropy regularization imposed by the reference model that allows the preservation of the high paraphrasing capability of LLMs. Table 4 shows a randomly

picked example of simplification by our method; Appendix E provides more.

6.2 Human Evaluation Results

We perform a human evaluation to assess the simplification quality from human perspectives. We recruited three graduate-level students majoring in linguistics to perform the evaluation. The evaluators were first trained with the background knowledge and then given a guideline to evaluate the following aspects of the samples: fluency, simplicity, adequacy, and CEFR sentence level.

We asked annotators to make binary judgements for fluency, simplicity, and adequacy. For sentence level, because CEFR-level judgements require expertise in language education, we simplified the task to collect reliable decisions. We asked the evaluators to judge if a simplified sentence matches the desired sentence level (denoted as “Level”). We showed a reference with its CEFR level and requested the evaluators to judge if the model output matches the reference’s simplicity. In addition, we asked them if a simplification output is preferable in terms of its CEFR level compared to the one generated by a model targeting a different level (denoted as “Prefer”). For example, an evaluator judges if an output of the A-level model is preferable to that of the C-level compared to the

| CEFR-SP | LENS | SALSA | Fluency | Adequacy |
|-----------------|--------------|--------------|--------------|--------------|
| Reference | 43.57 | 59.54 | 0.829 | 0.624 |
| phi3-3b-vanilla | 63.37 | 74.18 | 0.897 | 0.538 |
| T5+grade-A | 41.37 | 58.98 | 0.547 | 0.291 |
| FUDGE-A | 60.84 | 70.16 | 0.780 | 0.447 |
| phi3-A | 67.29 | 76.23 | 0.827 | 0.604 |
| T5+grade-B | 40.15 | 58.43 | 0.535 | 0.290 |
| FUDGE-B | 53.33 | 68.69 | 0.823 | 0.540 |
| phi3-B | 64.61 | 72.21 | 0.871 | 0.768 |
| T5+grade-C | 41.67 | 59.12 | 0.538 | 0.277 |
| FUDGE-C | 60.50 | 70.48 | 0.830 | 0.473 |
| phi3-C | 57.06 | 70.93 | 0.913 | 0.615 |

Table 5: Simplification quality on CEFR-SP-Test per levels; T5-grade, FUDGE and proposed method were evaluated using subsets of specific levels (A, B and C level references, respectively).

| TURK | LENS | SALSA | Fluency | Adequacy |
|-----------------|--------------|--------------|--------------|--------------|
| Reference | 35.20 | 64.96 | 0.732 | 0.901 |
| ACCESS | 49.90 | 62.68 | 0.576 | 0.780 |
| DMASS | 46.52 | 58.97 | 0.515 | 0.665 |
| DRESS | 59.76 | 62.63 | 0.807 | 0.615 |
| DRESS-LS | 60.56 | 62.92 | 0.838 | 0.657 |
| EditNTS | 57.71 | 64.86 | 0.752 | 0.710 |
| IterativEdit | 37.35 | 49.74 | 0.409 | 0.607 |
| phi3-3b-vanilla | 65.08 | 71.93 | 0.830 | 0.807 |
| phi3-A | 64.92 | 73.68 | 0.720 | 0.708 |
| phi3-B | 70.25 | 69.05 | 0.855 | 0.952 |
| phi3-C | 62.24 | 70.43 | 0.869 | 0.872 |

Table 6: Simplification quality on TurkCorpus; all models evaluated on the entire sentences as TurkCorpus does not annotate levels.

A-level reference.⁷ For each CEFR level, 30 simplifications of the CEFR-SP-Test were randomly sampled and annotated “Level” and “prefer” judgements. We report the ratios of positive judgements as evaluation scores. The details of the annotation guideline and interface are presented in Appendix D.

Table 7 shows the results; the simplicity score is generally high, close to 1, across models. This is expected as the source sentences were generated to be highly complex. The adequacy measurement results are consistent with automatic evaluation; identifying our proposed models as the most adequate. Furthermore, the proposed method achieves the best controllability on sentence levels compared to the baselines as indicated by significantly higher “Level” and “Prefer” scores.

6.3 Ablation Study

In this section, we show how each part of the proposed rewards contributes to the final performance. We compare the following models: vanilla phi3 model, reward using only target vocabulary counts,

⁷Preference score for the reference was judged by comparison with a sentence randomly chosen from another level

| Model | Simplicity | Adequacy | Fluency | Prefer | Level |
|-------------|------------|----------|---------|--------|-------|
| Reference | 1.00 | 0.89 | 0.99 | 0.87 | – |
| T5+grade-A | 0.83 | 0.16 | 0.47 | 0.40 | 0.10 |
| T5+grade-B | 0.90 | 0.13 | 0.50 | 0.43 | 0.17 |
| T5+grade-C | 0.80 | 0.16 | 0.60 | 0.40 | 0.17 |
| FUDGE-A | 1.00 | 0.50 | 0.80 | 0.50 | 0.43 |
| FUDGE-B | 0.96 | 0.43 | 0.83 | 0.57 | 0.47 |
| FUDGE-C | 1.00 | 0.47 | 0.83 | 0.57 | 0.33 |
| Ours phi3-A | 1.00 | 0.76 | 0.90 | 0.67 | 0.83 |
| Ours phi3-B | 1.00 | 0.83 | 0.90 | 0.70 | 0.63 |
| Ours phi3-C | 0.96 | 0.80 | 1.00 | 0.80 | 0.60 |

Table 7: Human evaluation results

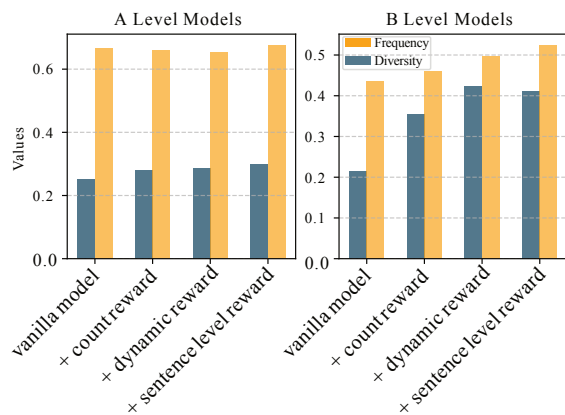


Figure 2: Reward effects on target vocabulary coverage

reward using dynamically adjusted vocabulary coverage rates, and reward using both dynamic vocabulary coverage rate and sentence levels (proposed method). The frequency and diversity evaluation results for A and B level models are presented in Fig. 2. Complete results can be found in Appendix E. It can be seen that changing the simple match count reward to a dynamically adjusted reward indeed encourages the model to increase the entropy inside the target vocabulary and largely improve the vocabulary diversity.

7 Conclusion

In this paper, we target ESL learners as audiences for text simplification aiming to facilitate the learning process of the foreign language. Referring to the input hypothesis and frequency effect theory in L2 learning, we propose a reinforcement learning method on LLM to control the simplification model to generate outputs that satisfy the vocabulary and sentence level constraints. Experiment results show that the proposed method can increase the target vocabulary coverage in the generated simplifications, and human evaluation results confirmed that the simplified texts generally preserve the targeted CEFR levels.

In practice, different individuals have varied lev-

els of knowledge for the language. We plan to extend the method to generate individual learner-targeted personalized simplifications in the future.

Limitations

This work assumes the target vocabulary for the learner is accessible, which in reality may not be the case as the target vocabulary varies with learner individuals and has to be first estimated. Although it is out of the scope of this paper, this direction constitutes our future work. Besides, currently, we do not control the frequency to a specific number, such as 95% i level and 5% $i + 1$ level, which is an important aspect to consider according to the L2 learning theory.

The control for target vocabulary and sentence is implemented individually for different levels rather than using one model altogether, causing heavier computational loads. In future, we seek to improve the design of reward model to integrate rewards for different proficiency levels into one model, and explore for a finer control over the frequency of the generated vocabulary.

References

- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. 2024. [Phi-3 technical report: A highly capable language model locally on your phone](#). *arXiv preprint arXiv:2404.14219*.
- 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.
- 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 (EMNLP-IJCNLP): System Demonstrations*, pages 49–54, Hong Kong, China. Association for Computational Linguistics.
- Fernando Alva-Manchego, Carolina Scarton, and Lucia Specia. 2021. [The \(un\)suitability of automatic evaluation metrics for text simplification](#). *Computational Linguistics*, 47(4):861–889.
- 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.
- Jan Cegin, Jakub Simko, and Peter Brusilovsky. 2023. [ChatGPT to replace crowdsourcing of paraphrases for intent classification: Higher diversity and comparable model robustness](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1889–1905, Singapore. Association for Computational Linguistics.
- Antoine Chaffin, Vincent Claveau, and Ewa Kijak. 2022. [PPL-MCTS: Constrained textual generation through discriminator-guided MCTS decoding](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2953–2967, Seattle, United States. 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.
- Liam Cripwell, Joël Legrand, and Claire Gardent. 2023. [Simplicity level estimate \(SLE\): A learned referenceless metric for sentence simplification](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12053–12059, Singapore. Association for Computational Linguistics.
- 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, Florence, Italy. Association for Computational Linguistics.
- Nick C Ellis. 2002. Frequency effects in language processing: A review with implications for theories of implicit and explicit language acquisition. *Studies in second language acquisition*, 24(2):143–188.
- Nick C Ellis and Fernando Ferreira-Junior. 2009. Construction learning as a function of frequency, frequency distribution, and function. *The Modern language journal*, 93(3):370–385.
- Arnaud Fickinger, Hengyuan Hu, Brandon Amos, Stuart Russell, and Noam Brown. 2021. [Scalable online planning via reinforcement learning fine-tuning](#). *Advances in Neural Information Processing Systems*, 34:16951–16963.
- Han Guo, Ramakanth Pasunuru, and Mohit Bansal. 2018. [Dynamic multi-level multi-task learning for sentence simplification](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 462–476, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

- David Heineman, Yao Dou, Mounica Maddela, and Wei Xu. 2023. [Dancing between success and failure: Edit-level simplification evaluation using SALSA](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3466–3495, Singapore. Association for Computational Linguistics.
- Sophie Henning, William Beluch, Alexander Fraser, and Annemarie Friedrich. 2023. [A survey of methods for addressing class imbalance in deep-learning based natural language processing](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 523–540, Dubrovnik, Croatia. Association for Computational Linguistics.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- 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.
- Tannon Kew and Sarah Ebling. 2022. [Target-level sentence simplification as controlled paraphrasing](#). In *Proceedings of the Workshop on Text Simplification, Accessibility, and Readability (TSAR-2022)*, pages 28–42, Abu Dhabi, United Arab Emirates (Virtual). Association for Computational Linguistics.
- J Peter Kincaid, Robert P Fishburne Jr, Richard L Rogers, and Brad S Chissom. 1975. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel.
- Tom Kocmi and Christian Federmann. 2023. [Large language models are state-of-the-art evaluators of translation quality](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203, Tampere, Finland. European Association for Machine Translation.
- Stephen Krashen. 1981. Second language acquisition. *Second Language Learning*, 3(7):19–39.
- 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, Online. Association for Computational Linguistics.
- Philippe Laban, Tobias Schnabel, Paul Bennett, and Marti A. Hearst. 2021. [Keep it simple: Unsupervised simplification of multi-paragraph text](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6365–6378, Online. Association for Computational Linguistics.
- Ximing Lu, Sean Welleck, Peter West, Liwei Jiang, Jungo Kasai, Daniel Khashabi, Ronan Le Bras, Lianhui Qin, Youngjae Yu, Rowan Zellers, Noah A. Smith, and Yejin Choi. 2022. [NeuroLogic a*esque decoding: Constrained text generation with lookahead heuristics](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 780–799, Seattle, United States. Association for Computational Linguistics.
- Ximing Lu, Peter West, Rowan Zellers, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. [NeuroLogic decoding: \(un\)supervised neural text generation with predicate logic constraints](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4288–4299, Online. Association for Computational Linguistics.
- Mounica Maddela, Yao Dou, David Heineman, and Wei Xu. 2023. [LENS: A learnable evaluation metric for text simplification](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16383–16408, Toronto, Canada. Association for Computational Linguistics.
- Louis Martin, Éric de la Clergerie, Benoît Sagot, and Antoine Bordes. 2020. [Controllable sentence simplification](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4689–4698, Marseille, France. European Language Resources Association.
- Louis Martin, Angela Fan, Éric 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, Marseille, France. European Language Resources Association.
- Rosamond Mitchell, Florence Myles, and Emma Marsden. 2019. *Second language learning theories*. Routledge.
- Akifumi Nakamachi, Tomoyuki Kajiwara, and Yuki Arase. 2020. [Text simplification with reinforcement learning using supervised rewards on grammaticality, meaning preservation, and simplicity](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing: Student Research Workshop*, pages 153–159, Suzhou, China. Association for Computational Linguistics.
- 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.
- 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, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners.
- 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.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *arXiv preprint arXiv:1707.06347*.
- Kim Cheng Sheang and Horacio Saggion. 2021. [Controllable sentence simplification with a unified text-to-text transfer transformer](#). In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 341–352, Aberdeen, Scotland, UK. Association for Computational Linguistics.
- Regina Stodden, Omar Momen, and Laura Kallmeyer. 2023. [DEplain: A German parallel corpus with intralingual translations into plain language for sentence and document simplification](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16441–16463, Toronto, Canada. Association for Computational Linguistics.
- Elior Sulem, Omri Abend, and Ari Rappoport. 2018. [Semantic structural evaluation for text simplification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 685–696, New Orleans, Louisiana. Association for Computational Linguistics.
- Renliang Sun, Hanqi Jin, and Xiaojun Wan. 2021. [Document-level text simplification: Dataset, criteria and baseline](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7997–8013, Online and Punta Cana, Dominican Republic. 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.
- Kevin Yang and Dan Klein. 2021. [FUDGE: Controlled text generation with future discriminators](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3511–3535, Online. Association for Computational Linguistics.
- Kexin Yang, Dayiheng Liu, Wenqiang Lei, Baosong Yang, Mingfeng Xue, Boxing Chen, and Jun Xie. 2023. [Tailor: A soft-prompt-based approach to attribute-based controlled text generation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 410–427, Toronto, Canada. 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.
- 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*.
- Xingxing Zhang and Mirella Lapata. 2017. [Sentence simplification with deep reinforcement learning](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 584–594, Copenhagen, Denmark. Association for Computational Linguistics.
- Sanqiang Zhao, Rui Meng, Daqing He, Andi Saptono, and Bambang Parmanto. 2018. [Integrating transformer and paraphrase rules for sentence simplification](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3164–3173, Brussels, Belgium. Association for Computational Linguistics.

A Evaluation Metrics

In this section, the evaluation metrics are explained in details: fluency, adequacy, target vocabulary frequency, diversity and target sentence level.

Previous studies have introduced various metrics for evaluating simplicity, which we summarize in Table 8. Among these metrics, SARI (Xu et al., 2016) is the most commonly employed in the literature. However, recent studies show that SARI may not be an optimal measure for assessing the quality of simplicity (Alva-Manchego et al., 2021;

Maddela et al., 2023; Stodden et al., 2023). Thus, we chose to use LENS (Maddela et al., 2023) and SALSA (Heineman et al., 2023), two recently proposed metrics, to measure simplicity.

| Metric | Scope | Reference |
|---------------------------------|--------------------------------|-----------|
| BLEU (Papineni et al., 2002) | semantic similarity | Y |
| FKGL (Kincaid et al., 1975) | readability | N |
| FKBLEU (Xu et al., 2016) | readability, similarity | Y |
| SARI (Xu et al., 2016) | keep, add, delete | Y |
| D-SARI (Sun et al., 2021) | keep, add, delete | Y |
| SAMSA (Sulem et al., 2018) | semantic structural similarity | N |
| BERTScore (Zhang* et al., 2020) | semantic similarity | Y |
| SLE (Cripwell et al., 2023) | human rating + FKGL | N |
| LENS (Maddela et al., 2023) | human rating | Y |
| SALSA (Heineman et al., 2023) | human rating | N |

Table 8: Metrics used in recent literature. Scope denotes the aspect that the metric aims to evaluate, and reference indicates whether the metric is computed based on references or not.

For adequacy and fluency, the ideal approach is human evaluation; however, this is impractical due to the large dataset size. Instead, we employed large language models to assess these two aspects. With a capable language model f , the generated simplification sentence s is evaluated as:

$$score(s) = \sum_{v \in V_y} f(v | (pmt, s)) \quad (8)$$

where pmt is a prompt designed for the model to output "yes" if the model evaluate s to be adequate or fluent, and V_y is a vocabulary subset for "yes" with $V_y = \{\text{YES, Yes, yes}\}$. We use Llama-3-8b-instruct⁸ model as the evaluation model in our experiment.

To measure target vocabulary frequency, we took the ratio between the total count of matched target words and the total generated words.

$$\sum_{j=1}^m \sum_{k=1}^n \text{count}(C_j, \text{seq}_k) / \sum_{k=1}^n \text{count}(\text{token}, \text{seq}_k) \quad (9)$$

To measure vocabulary diversity, we took the ratio between the number of matched words and number of words in the word list.

$$\sum_{j \in D} \mathbf{1} \left(\bigvee_{k=1}^n \mathbf{1}_{C_j}(\text{seq}_k) \right) / |D| \quad (10)$$

B Complex Sentence Generations

To generate complex sentences, we prompted the GPT-4 model⁹ to rephrase sentences of varied levels into highly complex sentences. To ensure the diversity of the generated complex data, we initially

⁸<https://llama.meta.com/llama3/>

⁹<https://openai.com/index/gpt-4/>

created a variety of seed prompts manually and instructed GPT-4 to generate additional prompts based on these seed prompts. GPT-4 was then prompted to generate complex sentences based on these diversified prompts. The 5 manually written seed prompts and 10 model generated prompts are presented in Table 9. In total 15 prompts were used to generate complex sentences, for each generation, one of the prompts was randomly selected. We present samples of the generated complex sentences together with simplifications in Table 10 and Table 11.

C Training Details

Implementation Details of Baselines We implemented the baseline models using the transformers library¹⁰. T5-s2s models require parallel corpus of complex-simple sentences, for which we used the pseudo-parallel sentences of generated complex sentences and their original sentences, and prepended level tokens for level controlling during training and evaluation. We implemented the FUDGE simplification control model with a Llama-3-8b-instruct model as the generation model, and its logits during the inference were adjusted using the CEFR level classification model released by (Arase et al., 2022).

Implementation Details of Proposed Method

The PPO training algorithm was implemented using the TRL library¹¹ with a learning rate of $3e-5$. For the dynamic reward model used in the training, we set the α to be 1.2, as we found a value slightly bigger than 1 was shown to have better performance empirically; the reward for phrases is always set to be 1.5 times more than words to reward more on the phrase generation. For the overall reward, λ was set to be 1.5 to compensate for the vocabulary reward penalty, and γ was set to 1. During training, we used the following prompt for the model to generate simplifications: "Given a complex sentence { }, generate a simplified version for it:".

Training Performance The performance of the sentence level reward model is shown in Fig. 3. Fig.4 shows mean reward and KL change over the training steps with and without the dynamic reward. Objective/KL indicates the deviation of the simplification model from the reference model, and

¹⁰<https://huggingface.co/docs/transformers/en/index>

¹¹<https://huggingface.co/docs/trl/main/en/index>

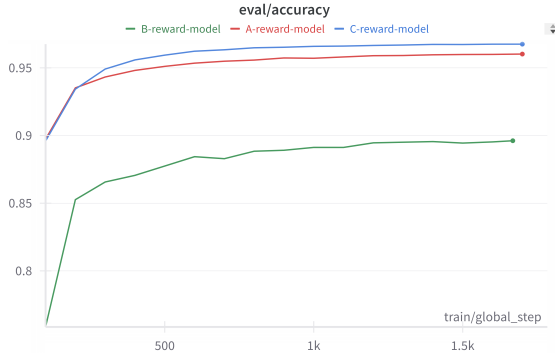


Figure 3: Sentence level reward model evaluation accuracy

an absurdly high KL indicates model collapse; a burst in mean reward indicates model collapses and only produces a limited set of vocabulary. It can be observed that using the dynamic reward helps stabilize the training, while using the match count alone causes the model to be over-optimized and collapse to a limited vocabulary subset.


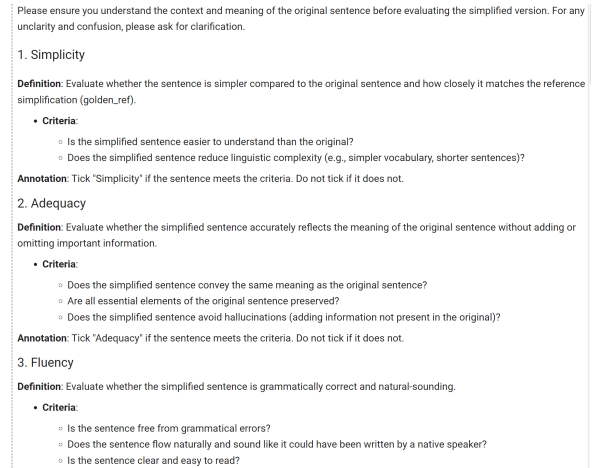


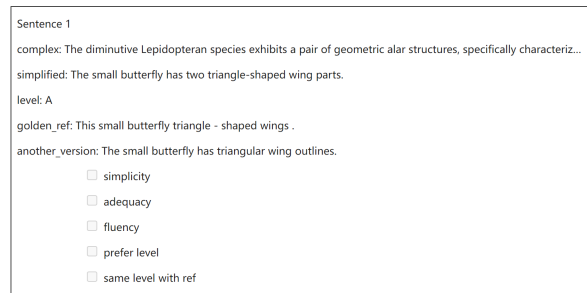
Figure 4: Training stability w/wo dynamic reward

D Human Evaluation Details

In this section, the annotation guidelines that evaluators used to evaluate the generated simplifications as well as the annotation interface are presented. The annotation guidelines contain the definition of the aspect to be evaluated, the criteria for the evaluation and indications for the annotation interface. The annotation interface is designed to be a binary-choice form, for each aspect to be evaluated, the evaluators chose to tick to indicate that the simplification contains the aspect to be evaluated, and does not satisfy the aspect otherwise. The evaluation results are then used to calculate the binary accuracy of the aspects to be evaluated. The evaluation guidelines are shown in Fig. 5(a) and the evaluation interface is shown in Fig. 5(b).



(a) Screenshot of annotation guidelines shown to the evaluators



(b) Screenshot of annotation interface shown to the evaluators

Figure 5: Screenshot of annotation guidelines

E More Evaluation Results

Fig. 6 shows the ablation study results on all three levels of models.

Tables 10 and 11 present example outputs: the complex sentences, reference sentences, and simplified sentences. For each complex sentences, there are three versions of the simplified sentences, corresponding to A, B and C levels generated by different models targeting the corresponding level, respectively.




Figure 6: Vocabulary targeting w/ different rewards, all levels

| |
|---|
| Manual Composed Prompts |
| You are an expert in academic writing, renowned for your ability to compose intricate and sophisticated sentences. Please rephrase the following sentence,so that it's a complex, hard to follow sentence that would usually appear in a journal article, without loss of original meaning: |
| You are an experienced English teacher. Please rephrase the following sentence,to make it a complicated, very hard sentence to read that a English learner may encounter in daily reading, without loss of original meaning: |
| You are a successful postmodernism theater and book critic. You used varied writing styles in your articles. Please rephrase the following sentence,to make it a complex and very difficult to understand sentence,without loss of original meaning: |
| You are a philosopher and literature professor. You usually make intricate perception and sharp insight in your writing. Please rephrase the following sentence,to make it a short but complex and very hard to follow,without loss of original meaning: |
| You are an editor of social and financial news and journals. Please rephrase the following sentence,so that the sentence has complex compositions and advanced words, that normal readers cannot understand, without loss of original meaning: |
| Model Generated Prompts |
| You are a legal scholar with extensive experience in drafting complex legal documents. Please rephrase the following sentence,to make it a complex and legally intricate sentence,without loss of original meaning: |
| You are a renowned scientist known for writing dense and comprehensive research papers. Please rephrase the following sentence,to make it a complex and highly technical sentence,without loss of original meaning: |
| You are a seasoned journalist known for crafting elaborate and detailed investigative reports. Please rephrase the following sentence,to make it a complex and deeply investigative sentence,without loss of original meaning: |
| You are a literary critic who writes for a prestigious literary journal, known for your sophisticated language. Please rephrase the following sentence,to make it a complex and highly sophisticated sentence,without loss of original meaning: |
| You are a historian known for your detailed and intricate historical analyses. Please rephrase the following sentence,to make it a complex and historically detailed sentence,without loss of original meaning: |
| You are an expert in technical writing, specializing in creating elaborate and detailed user manuals. Please rephrase the following sentence,to make it a complex and technically detailed sentence,without loss of original meaning: |
| You are a linguist with expertise in creating intricate and multifaceted linguistic analyses. Please rephrase the following sentence,to make it a complex and linguistically intricate sentence,without loss of original meaning: |
| You are a political theorist known for your dense and intricate political analyses. Please rephrase the following sentence,to make it a complex and politically intricate sentence,without loss of original meaning: |
| You are an economist renowned for your detailed and complex economic analyses. Please rephrase the following sentence,to make it a complex and economically detailed sentence,without loss of original meaning: |
| You are a theologian known for your intricate and deeply philosophical theological writings. Please rephrase the following sentence,to make it a complex and theologically intricate sentence,without loss of original meaning: |

Table 9: Prompts used to generate complex sentences

| |
|--|
| <p>Complex Sentence Let us endeavor to delve into the intricacies and nuances of the text, striving to comprehend the underlying themes and implications inherent within, as we embark on this journey of intellectual exploration.</p> |
| <p>Reference, level: A Let 's try to read .</p> |
| <p>Simplifications Let's try to understand the text well. Let's try to understand the text's main ideas and meanings as we read. Let's understand the text's themes and meanings.</p> |
| <p>Complex Sentence Are you experiencing any concerns or apprehensions regarding the composition, quality, or safety of the food you consume?</p> |
| <p>Reference, level: A Do you worry about what you eat ?</p> |
| <p>Simplifications Do you have worries about the food you eat? Do you have worries about the food you eat's quality or safety? Do you have concerns about your food's composition, quality, or safety?</p> |
| <p>Complex Sentence The considerable distance, compounded by Jamie's current condition of pregnancy, which inexorably engenders a state of increased fatigue, renders the prospect of ambulation to said location prohibitively challenging for her.</p> |
| <p>Reference, level: B It is too far for Jamie to walk to , especially because she is pregnant and easily exhausted .</p> |
| <p>Simplifications Jamie is too tired to walk far because she is pregnant. Jamie's pregnancy makes it very hard for her to walk to the location due to the long distance. Jamie's pregnancy leads to fatigue, making it hard for her to walk to the distant place.</p> |
| <p>Complex Sentence With meticulous care, the young girl positioned the doll on her bed, her actions imbued with a quiet reverence, before softly pressing her lips to its cheek, an act infused with such tenderness that one might easily imagine she believed the doll to be a living boy.</p> |
| <p>Reference, level: B The little girl placed the doll carefully on her bed and kissed its cheek as gently as if it were a real boy .</p> |
| <p>Simplifications The girl put the doll on her bed and kissed it gently. The girl carefully placed the doll on her bed, kissed its cheek gently, and seemed to love it like a boy. The girl gently kissed the doll out of affection.</p> |

Table 10: Complex sentences, references and model simplified sentences in A, B, C levels, respectively

| |
|---|
| <p>Complex Sentence</p> <p>Numerous physical characteristics of alkenes and alkanes exhibit remarkable congruity; manifesting as chromatically imperceptible, possessing nonpolarity, and exhibiting combustibility.</p> |
| <p>Reference, level: C</p> <p>Many of the physical properties of alkenes and alkanes are similar : they are colorless , nonpolar , and combustable .</p> |
| <p>Simplifications</p> <p>Alkenes and alkanes have similar physical traits. They are hard to see color differences, not attracted to water, and burn easily.</p> <p>Alkenes and alkanes have similar physical traits. They are colorless, nonpolar, and burn easily.</p> <p>Alkenes and alkanes share traits like colorlessness, nonpolarity, and flammability.</p> |
| <p>By the month of May in the year 2008, the phenomenon known as Rickrolling had ascended to the status of an internet sensation, infiltrating the digital landscape with such pervasive enthusiasm that it swiftly garnered the attention of mainstream media outlets, which began to extensively cover this viral trend.</p> |
| <p>Reference, level: C</p> <p>By May 2008 , Rickrolling had become an internet hit or phenomenon and very soon rickrolling was being covered on mainstream media .</p> |
| <p>Simplifications</p> <p>In 2008, Rickrolling became very popular on the internet and in the news.</p> <p>In 2008, Rickrolling became a big internet trend.</p> <p>In 2008, Rickrolling became an internet sensation, leading to widespread media coverage.</p> |

Table 11: Complex sentences, references and model simplified sentences in A, B, C levels, respectively