

Med-EASi: Finely Annotated Dataset and Models for Controllable Simplification of Medical Texts

Chandrayee Basu¹, Rosni Vasu², Michihiro Yasunaga¹, Qian Yang³

¹ Stanford University

² University of Zurich

³ Cornell University

cbasu@stanford.edu, rosni@ifi.uzh.ch, myasu@stanford.edu, qy242@cornell.edu

Abstract

Automatic medical text simplification can assist providers with patient-friendly communication and make medical texts more accessible, thereby improving health literacy. But curating a quality corpus for this task requires the supervision of medical experts. In this work, we present **Med-EASi** (**M**edical dataset for **E**laborative and **A**bstractive **S**implification), a uniquely crowdsourced and finely annotated dataset for supervised simplification of short medical texts. Its *expert-layman-AI collaborative* annotations facilitate *controllability* over text simplification by marking four kinds of textual transformations: elaboration, replacement, deletion, and insertion. To learn medical text simplification, we fine-tune T5-large with four different styles of input-output combinations, leading to two control-free and two controllable versions of the model. We add two types of *controllability* into text simplification, by using a multi-angle training approach: *position-aware*, which uses in-place annotated inputs and outputs, and *position-agnostic*, where the model only knows the contents to be edited, but not their positions. Our results show that our fine-grained annotations improve learning compared to the unannotated baseline. Furthermore, *position-aware* control generates better simplification than the *position-agnostic* one. The data and code are available at <https://github.com/Chandrayee/CTRL-SIMP>.

Introduction

Health literacy refers to our knowledge, and our ability to obtain, process, and understand health information and services to make appropriate health decisions (Kindig et al. 2004). Low health literacy has several adverse effects including poor patient self-care, lack of timely communication of health issues, and even increased risk of hospitalization and mortality (King 2010; Berkman et al. 2011; Tajdar et al. 2021). Low digital health literacy makes it harder for consumers to disambiguate between reliable medical information (NIA 2018; Savery et al. 2020) and unreliable ones, accelerating the spread of medical misinformation, as in the case of COVID-19 (Bin Naeem and Kamel Boulos 2021). Despite the promise of automated medical text simplification in mitigating this problem, there are limited datasets and open-source libraries for this task (Siddharthan 2014; Phan et al. 2021; Cao et al. 2020; Van den Bercken, Sips,

and Lofi 2019). In this work, we present the very first finely annotated dataset *Med-EASi* for medical text simplification and describe two different T5-based models *ctrlSIM* and *ctrlSIM_{ip}* that enable controllable simplification. We define *controllability* as the ability of a user to selectively simplify the contents of a short medical text. Existing unsupervised models (Prabhumoye et al. 2018; Subramanian et al. 2018; Shen et al. 2017; Alva-Manchego et al. 2020) and hierarchical tagging and text editing models (Malmi et al. 2019; Mallinson et al. 2022, 2020) can be trained directly on unlabeled datasets like (Cao et al. 2020) for medical text simplification. However, we observe the need for fine-grained annotations to enable word or phrase level *controllability* over the model outputs and for elaborative simplification (Srikanth and Li 2020).

Crowdsourcing annotations for medical texts is challenging, which explains the dearth of datasets and research in this domain, compared to general text simplification. To obtain high-quality annotations, we must recruit a specific subpopulation of workers with domain expertise (Nye et al. 2018). Furthermore, when it comes to evaluating the quality of simplification, only experts can judge the correctness and relevance of the added content. On the other hand, only the laymen audience can validate the readability and comprehensibility of the model outputs. Therefore, we deploy a novel data annotation format for this research, involving both medical experts and layman crowd-workers.

We make the following contributions:

- **Dataset:** We finely annotate two existing parallel medical text simplification corpora with four kinds of textual transformations, viz. elaboration, replacement, deletion, and insertion of new content.
- **Expert-layman-AI annotation:** We divide the annotation tasks between medical experts and layman crowd-workers depending on the syntactic and domain-specific complexity of the example texts. We assist layman crowd-workers by providing AI-generated annotations.
- **Controllable models:** We train two multi-angle (Tafjord and Clark 2021) text simplification models with T5-large backbone. Our models perform controllable simplifications with satisfactory outputs.

Related Work

Text simplification in the medical domain is mostly restricted to paraphrasing (Abrahamsson et al. 2014). More recent transformer-based approaches follow developments in general text simplification (Zhang and Lapata 2017; Nisioi et al. 2017; Jiang et al. 2020). These models treat text simplification as monolingual translation without controllability and deploy tricks like auto-completion (Van, Kauchak, and Leroy 2020) and unlikelihood training (Welleck et al. 2019; Devaraj et al. 2021) to mitigate for low resources, and common pitfalls of LLMs, e.g. hallucination, and neural text degeneration.

Meanwhile, research in automatic non-medical text simplification has been burgeoning, with the introduction of large parallel corpora (Zhu, Bernhard, and Gurevych 2010; Woodsend and Lapata 2011; Coster and Kauchak 2011; Xu, Callison-Burch, and Napoles 2015; Paetzold and Specia 2017). The state-of-the-art models can be classified into edit-based (Malmi et al. 2019; Mallinson et al. 2020; Agrawal, Xu, and Carpuat 2021; Omelianchuk, Raheja, and Skurzhanyskiy 2021; Cumbicus-Pineda, Gonzalez-Dios, and Soroa 2021; Mallinson et al. 2022) and text-to-text models. The creation of multi-references has enabled the models to explicitly learn different kinds of textual transformations, viz. *lexical changes* (e.g. paraphrasing), *syntactic modifications* (e.g. reordering of concepts, splitting texts, reducing sentence length, etc.) and *compression* (e.g. deleting peripheral information irrelevant to the target domain) (Alva-Manchego et al. 2020). Our approach is similar to the text-to-text models that exert controllability using task-specific prompts (Keskar et al. 2019; Dathathri et al. 2019; Kariuk and Karamshuk 2020; Brown et al. 2020; Reif et al. 2021; Lyu et al. 2021; Xu, Peng, and Liu 2022).

However, unlike other models that control for specific attributes of the generated simplification like compression ratio, word rank (Martin et al. 2019b), level of paraphrasing (Maddela, Alva-Manchego, and Xu 2020), grade-level (Nishihara, Kajiwara, and Arase 2019) etc., we develop a single end-to-end model that can perform all the edits that simplification entails, similar to edit-based models, while also allowing users to select the content to be edited and the desired form of edit.

Features of Text Simplification

We define text simplification as the process of reducing the linguistic complexity of a text, while still retaining the original information content and meaning (Siddharthan 2014). A domain-specific text undergoes various kinds of edits to reach the final simple form. We explore four kinds of textual transformations, defined as follows:

Deletion: removal of any word, phrase, sub-statement, or full sentence from the expert text

Insertion: addition of words, phrases, and sentences that just change the style of the text or fix errors. They do not provide any extra information about any term in the text.

Replacement: replacing any complex word or phrase in the expert text with simpler words or phrases. Unlike elaboration, the original term is missing in the simple text.

Elaboration: extra information or definitions of original content in the expert text, added as a word or a phrase to the simple text. Typically, the term being elaborated is also present in the simple, optionally replaced by its synonym. We consider these as two distinct types: elaborations where part of the original phrase is preserved (*type 1*) and elaborations where original content is fully replaced by its synonyms (*type 2*).

While elaboration can be broken down into keep, delete, and insert content and replacement can be broken down into delete and insert content (Mallinson et al. 2020), we treat these transformations differently for more human-interpretable controllability.

Creating Med-EASi

Existing Parallel Corpora

To create Med-EASi, we leverage existing parallel corpora ([expert *E*, simple *S*] text pairs) for medical text simplification. Based on the evaluation by Basu et al. (2021), we select two publicly available datasets for annotation: SIMPWIKI (Van den Bercken, Sips, and Lofi 2019) and MSD (Cao et al. 2020). We sorted the text pairs by Levenshtein Similarity and Compression ratio (Martin et al. 2019a) and selected samples for annotation, to cover a wide range of these metrics.

Expert-layman-AI Crowdsourcing

We recruited experts and layman crowd-workers to annotate Med-EASi with four textual edits (see section Features of Text Simplification), depending on the complexity of the texts. Many SIMPWIKI text pairs have high Levenshtein similarity (> 0.7), while MSD examples are more dissimilar and complex. We, therefore, allocated most of SIMPWIKI to the layman crowd-workers and left all of MSD and more difficult SIMPWIKI examples for the medical experts.

Annotations

Crowdsourcing annotations from layman workers is advantageous due to the sheer number of available annotators. This, however, acts contrary to the whole purpose of medical text simplification: the texts are already inaccessible to layman workers. To make the annotation task easier, we, therefore ask the layman workers to choose between two possible AI-generated annotations. Specifically, we use difflib (Python 2022) to identify how text A was transformed into text B using four types of edits: replace, keep (equal), delete and insert. We post-processed difflib’s output to generate annotation as *Uterine cancer, also known as womb cancer, is any type of cancer that <rep>emerges<by>starts</rep> from the tissue of the uterus.*, where <rep> stands for replace.

We used a SpanBERT-based coreference resolution model (Joshi et al. 2019) to identify spans in the simple text that referred to some entity in the expert text. We passed concatenated expert and simple texts to this model. We marked a span in the simple text as elaboration if it was longer than the reference span in the expert text. Resultant annotation looked like this: *In his*

Expert Text	Simple Text
People with positive fecal occult blood tests require colonoscopy, as do those with lesions seen during a sigmoidoscopy or an imaging study.	People with positive fecal occult blood tests require a colonoscopy, as do those with sores seen during a sigmoidoscopy or an imaging study.
Monozygotic twins have a concordance of about 45 %.	When comparing monozygotic twins, we found that they had a concordance of about 45 %.
During childhood, she suffered from partially collapsed lungs twice, had pneumonia four to five times a year, as well as a ruptured appendix and a tonsillar cyst.	During childhood, suffered from collapsed lungs twice, had pneumonia four or five, (ruptured) appendix and a tonsillar cyst.
CT also is required to accurately assess skull base bony changes, which are less visible on MRI.	Computed tomography (CT) also helps to assess these changes.

Figure 1: Med-EASi is the first finely annotated dataset for controllable medical text simplification. Our T5-based models with multi-angle training: *ctrlSIM* and *ctrlSIM_{ip}* allow users to selectively edit contents of the complex medical text by *elaboration*, *replacement*, *deletion* or *insertion*. The above examples are generated by our best model *ctrlSIM_{ip}*. Replacements are shown in magenta and elaborations in blue. The incorrect conversions are shown in red.

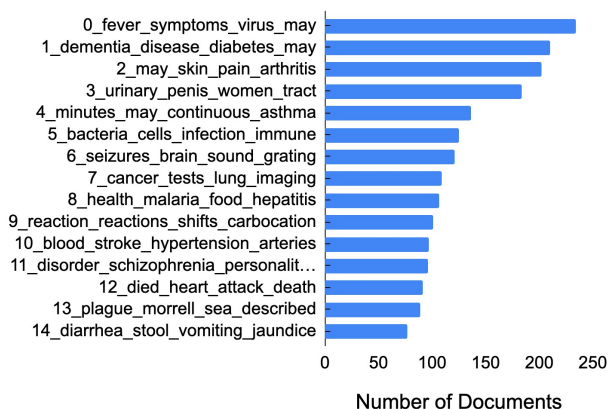


Figure 2: Topic distribution of Med-EASi with top-3 tokens on the x-axis.

<elab>Nobel lecture</by>Nobel Prize lecture</elab>, <elab>Lewis</by>E.B. Lewis</elab> said "Ultimately, comparisons of the (control complexes) throughout the animal kingdom should provide a picture of how the organisms, as well as the (control genes) have evolved.", where <elab> stands for elaboration.

Layman Crowdsourcing

We recruited layman crowd-workers who were fluent in English from Toloka (Toloka 2022) and followed a three-step training protocol. In step 1, we asked crowd-workers to watch a video describing the purpose of the project, our dataset, and the annotation tasks. The video was followed by 2 experience-related questions and 1 exam question. We asked workers to indicate the frequency of their interaction with medical data in English (messages with doctors, medical bills, Wikipedia, etc.), and their level of experience in NLP tasks. In the exam question, we asked them to select the correct annotation format for replacement. Workers who had watched 85 % of the video, had experience interpreting medical data in English, and were able to answer the

exam question correctly were selected for the second round of training (83 % acceptance rate).

In step 2, we assigned 12 tasks. Each task contained an expert-simple text pair and two possible aforementioned annotations. We asked the workers to choose the correct annotation from the two options or select None of the above. Workers with > 80 % correct answers were automatically selected for the next stage of training (80 % acceptance rate).

In step 3, all the tasks selected required corrections from the workers. These were examples where our automatic annotation schemes had failed. We manually checked the correct annotations and accepted any annotation that had at most one missing annotation (e.g. missed one replacement) and high overlap with ground truth spans. Workers with > 70 % accepted answers were selected for the actual annotation task (15 % acceptance rate).

The final annotation task design matched that of training step 2. We offered a bonus for writing the correct annotation. We got 3 annotations for each data pair from 3 workers. We aggregated the annotations using Dawid-Swene aggregation method (Dawid and Skene 1979) and automatically accepted the ones with at least 90% confidence. We passed the ones with lower confidence to an expert.

Expert Crowdsourcing

We recruited 5 experts from Upwork: two medical doctors, two medical students, and one biomedical research scientist, all with experience in data annotation for NLP. We asked the experts to watch the full instruction video and gave them 20 annotation tasks. The annotations were checked by two of our in-house team members. We relied on their medical knowledge significantly and clarified any disagreement with further discussion. All 5 experts preferred to provide correct annotations directly without referring to AI-generated ones.

Dataset Statistics

We annotated all of MSD, because of its clinical nature and diversity of transformations and approximately 1500 text pairs from SIMPWIKI. The resulting Med-EASi dataset contains a total of 1979 expert-simple text pairs.

Topic distribution of the dataset: To understand the domain-complexity of the data, we identified the medical concepts in each text pair and their Unified Medical Language Systems (UMLS) (Bodenreider 2004) representations using QuickUMLS (Soldaini and Goharian 2016). Med-EASi covers a total of 3909, 3304 unique medical concepts in the expert and simple texts respectively, and a total of 4478 concepts across all text pairs. The topic distribution computed using BERTopic (Grootendorst 2022) shows wide coverage of medical subdomains like infectious diseases, cardiology, neurology, etc. (Figure 2).

Quality: Following Basu et al. (2021) we measure data diversity using reference-less quality metrics from EASSE library (Martin et al. 2018). The Levenstein similarity (0.689 ± 0.22), the fraction of words added (9.347 ± 11.57), deleted (10.574 ± 12.38), kept (12.792 ± 10.01), and the compression ratio (1.025 ± 0.59) were computed. Med-EASi has overall acceptability 98.989% for expert text and 98.737% for layman texts computed using COLA-trained DistilBERT classifier (Morris et al. 2020; Warstadt, Singh, and Bowman 2019). The readability is estimated using the Flesch Kincaid readability grade (FKGL), measured as the minimum education level required to read and understand a text, expressed as an empirical function of total words, total sentences, and total syllables (Kincaid et al. 1975). The expert texts (12.47 ± 5.28) have statistically significant differences in the readability grade (paired t-test with $p < 0.001$) compared to the layman text (10.491 ± 4.98).

Multi-Angle Controllable Simplification Model

Any model performing simplification must have three capabilities: 1. predicting the span of the expert text that must be altered, 2. predicting the alteration or operation on each span, and 3. predicting the additional (in case of elaboration) or alternative (for replacement) contents. In *Controllable text simplification*, a user should be able to indicate the complex contents that must be deleted, replaced, or elaborated. In such cases, instead of asking the model to predict the span that must be transformed, the user can provide the span as input and the model should be able to incorporate the requested transformation into the generated simplification. Following our definition of *controllability*, we aim to develop a model that can simplify short medical texts, both with and without controllability instructions.

Med-EASi contains diverse training pairs, each with a different set of textual transformations. Therefore, we utilize the flexible nature of seq-2-seq models (Raffel et al. 2020) and finetune T5-large with a combination of left-to-right generation and infilling. We use variable task descriptions to accommodate heterogeneous forms of inputs and outputs, similar to MACAW (Tafjord and Clark 2021). This approach is called multi-angle training, where each component of the input and the output is a **slot** and the input-output combination is called an **angle**. Some examples of slots include expert text, phrases to be elaborated, or content that must be replaced.

Slots and Angles

We consider each example as a set of slots S_i and corresponding values V_i . Angle $A_i = S_{S_i} \rightarrow S_{T_i}$ is a combination of source or input slots S_{S_i} and target or output slots S_{T_i} . The slots are a way to decompose the task descriptions and prompts into human-interpretable instructions. It also allows easy recovery of the final simplification from the generated output with standard post-processing. The requested output slots are concatenated as task descriptions at the beginning with no associated values. We keep the `simple` slot as the last requested slot. We hoped that the generated simplification will be more accurate if conditioned on the predicted edit spans. While identifying angles, we use uppercase abbreviations of the slot names: We identify the slots with uppercase letters, D: deletion, I: insertion, R: replacement, X: elaboration, E: expert text, S: simple text, Ea: annotated expert text, Sa: annotated simple text. See Figure. 3 for color-coded examples of slots and angles.

Model Versions

We train one baseline model without annotated data. We hypothesize that controllability is not only desirable but can also improve text-to-text model performance by providing additional supervision signals. To test this, we develop four versions of T5-large: two control-free and, two controllable versions.

The **control-free** versions are referred to as *SIM* and *SIM_{ip}* respectively (*ip* stands for in-place annotation). The former predicts all the spans that must be transformed and the corresponding transformations. This version of the model also learns to predict what transformations are invalid for a given example. The latter directly generates the annotated simplification, from which the simple text is extracted with post-processing.

Like the control-free models, in case of the **controllable models** too, we experiment with two different input-output formats: *position-agnostic* and *position-aware*. We refer to the controllable models as *ctrlSIM* and *ctrlSIM_{ip}* respectively. *ctrlSIM* allows users to input the words or phrases they expect to be edited by replacement or elaboration, while *ctrlSIM_{ip}* requires users to highlight the same contents in place.

Training

Data, Angles, and Slots

We allocated 75 % or 1398 data points for training and set aside 10 % or 197 data points as dev set for hyperparameter tuning. The dev set matched the data distribution of the test set with some unseen UMLS terms. In our training data, we have a total of 445 elaborations, 2044 replacements, 512 insertions, and 905 deletions. So approximately 50% are replacement operations.

SIM and *SIM_{ip}* are trained on fixed angles: E→RXDIS and E→Sa respectively. In the case of examples with missing slots, all the empty slots contain the same token `<extra_id.0>` as value. While the input-output format for

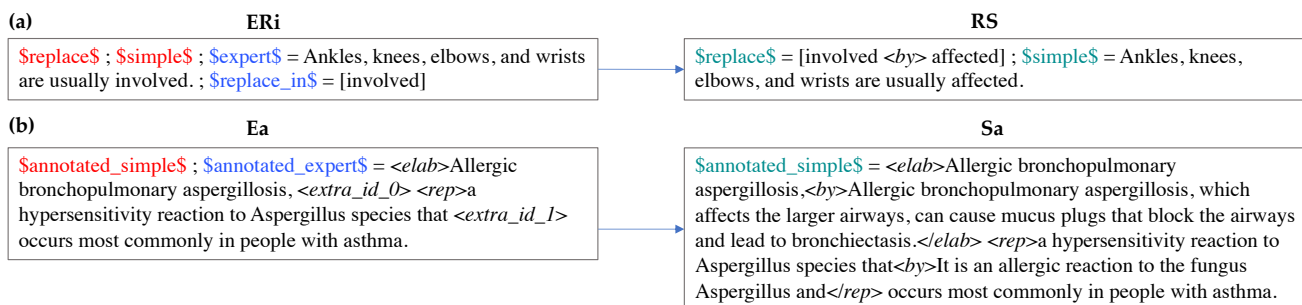


Figure 3: The table shows two forms of instructions for controllability, (a) $ctrlSIM$ for angle $ERi \rightarrow RS$: contents of the expert text that must be edited without reference to their position in the expert text and (b) $ctrlSIM_{ip}$ for angle $Ea \rightarrow Sa$: expert text with in-place annotation of spans to be edited and the desired edit types. The task descriptions are color-coded with red, the input slot names with blue, and the output slot names with teal. The values corresponding to the slots are in regular text color.

$ctrlSIM$ is similar to SIM , the model here only outputs slots present in the example and is not required to know which slots are empty. Table 1 shows all the training angles of $ctrlSIM$. $ctrlSIM_{ip}$ is trained on two different angles $E \rightarrow Sa$ and $Ea \rightarrow Sa$. Annotated expert text marks the beginning of the content to be transformed with the corresponding edit type tag ($\langle rep \rangle$ or $\langle elab \rangle$) and adds a special token at the end of the content. The token is different for elaboration and replacement. We do not expect the user to know what content must be inserted or deleted and leave that to the model to figure out.

Pretrained Models and Hardware

We experimented with two different pretrained models as backbones for our text simplification models: T5-large (Raffel et al. 2020) and SciFive-large(+pubmed+pmc) (Phan et al. 2021). While T5 is trained on 750 GB of web-crawled data, i.e. Colossal Clean Crawled Corpus (C4) (Raffel et al. 2020), it is not specifically fine-tuned for medical text generation. SciFive, on the other hand, is T5 retrained on combinations of C4, PubMed database of 32 million citations and abstracts (NIH 2022), and PubMed Central (PMC), a corpus of free full-text articles in the domain of biomedical and life sciences (NCBI 2022). We trained our models on two different servers, one with two GeForce RTX 3080 and another with three GeForce RTX 2080 Ti.

Decoding

Greedy decoding worked well for our output generation. We tested several beam sizes for beam search. A beam size of 10 handled the open-ended nature of elaboration and replacement tasks adequately, where the model is expected to leverage the knowledge of the pretrained backbone, beyond our training corpus. Nucleus sampling did not improve SARI score (next section) much and brought down the other scores.

Metrics

We evaluate the quality of simplification using SARI: a well-known metric that measures the goodness of words added,

deleted, and kept by the simplification model, by comparing the generated text against the input and multiple references (Xu et al. 2016).

We report the results of our two controllable models in terms of SARI and its sub-scores ADD, DEL, and KEEP, wherever applicable, and individually for each type of output slot. We considered several possible errors in our models including missing slots or predicting undesired slots. Although rare, we had a few examples with such errors. While computing SARI-based metrics we only considered examples with slots that have values in the true label as well as generated output. We ignored the following: examples that have slots with no values in the true label or in the generated output, examples that have slots with values in the true label but no values in the output, or vice versa.

Slot-wise we computed SARI for generated simplification, ADD score for insertion, KEEP and ADD for elaboration type 1 and, ADD and DEL for replacement. In order to predict whether $ctrlSIM$ is able to identify the content that must be elaborated or replaced and the corresponding valid additions, we split the output for these two slots into two parts: 1. textual span that must be edited (S_{pre}), and 2. the resulting span after the transformation (S_{post}), separated by $\langle by \rangle$ in model inputs and outputs. SARI DEL can be applied to S_{post} to evaluate whether certain content is present in the expert text, but not in the simple text. It, however, cannot evaluate if a model is able to predict the content to be deleted correctly. The same problem exists in the case of S_{pre} for replacement. To account for the deleted span of the expert text we modified the DEL score into ALT-DEL score. Following notations of Xu, Callison-Burch, and Napoles (2015) (ignoring the n-grams g for brevity), ALT-DEL score can be written as:

$$p_{altdel} = \frac{(I \cap O) \cap \bar{R}}{O} \quad (1)$$

$$r_{altdel} = \frac{(I \cap O) \cap \bar{R}}{(I \cap \bar{R})} \quad (2)$$

where p and r refer to precision and recall respectively. The predicted span (O) to be deleted must be present in the

expert text (I) and absent from the simple text (R). Please note that for the rest of the paper, we revert back to using I and R for insertion and replacement respectively.

We also evaluate our controllable models’ performance using ROUGE-L (Lin 2004) for recall. Since the ultimate goal is to generate simplifications that are accessible to laymen with different levels of medical knowledge, we report their readability using FKGL.

Hyperparameter Tuning

We fine-tuned the baseline, as well as all 4 versions of our models for 30 epochs, each with batch sizes 4, 8, 16, 32, and 64, and a constant learning rate of 8e-06.

Since the primary goal of our model is to output correct simplification, we use SARI instead of dev set error, for model selection. When computing SARI for each dev example, we ignored the examples where the model was unable to output all the requested slots. We observed that while dev set error goes up with more training, the model is able to understand the meaning of the slots better as well as able to add simpler words with further training. While KEEP, DEL, and SARI did not change much beyond epoch 15, ADD score and FKGL improved with further training.

Results

We allocated 15 % of Med-EASi or 300 text pairs for evaluation. We sampled the test set to cover wide ranges of Levenshtein Similarity and representative UMLS terms. To understand whether our model is able to retrieve its pretraining knowledge in the text simplification task, we kept 34 % of the 300 text pairs with 672 previously unseen UMLS terms. We sampled the rest 200 pairs at random from the original data distribution. Based on dev set performance we picked the top model checkpoints and evaluated their performances on the above test set.

We test the following features of both *ctrlSIM* and *ctrlSIM_{ip}* with our test results.

Overall Quality

We trained the best model *ctrlSIM_{ip}* once again, for 60 epochs, with 10 % warm-up steps and a cosine learning rate decay from epoch 27 onwards. The new dev set performance showed 41.36 SARI score and 9.05 ADD score, an improvement over the prior version. The model achieves this result with a batch size of 4 at epoch 13.

Overall, we obtained SARI scores comparable with the performances of the top text simplification models on other datasets (Mallinson et al. 2020).

Controllability

We test if the controllable models are able to perform the task requested by the users. To do so, we split the test results of the two models by angles.

Figure 4 shows the average SARI score of the simplifications generated by *ctrlSIM*. We observe that the angle $ERi \rightarrow RS$ (when the user provides the content that must be replaced, Ri in addition to the expert text E) achieves the

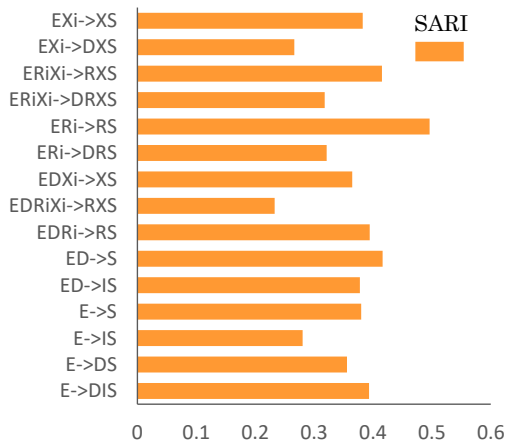


Figure 4: SARI scores of *ctrlSIM*’s outputs arranged by angles

highest score of 0.49. Such a high score could be attributed to the large number of instances of this angle in the training data. Furthermore, when the model predicts other slots prior to the simplification, like replacement or elaborations, it improves the quality of the generated simplification, as seen for angles $EXi \rightarrow XS$, $ERiXi \rightarrow RXS$ and $ERi \rightarrow RS$. Likewise, instructing the model to remove certain content results in a simplification with the highest DEL score.

ctrlSIM_{ip} demonstrates the above phenomenon more clearly. The average SARI scores of the simplifications vary significantly by the angles, 0.22 for $E \rightarrow Sa$ and 0.46 for $Ea \rightarrow Sa$.

Elaboration and Replacement

We ask: How are the models performing in open-ended and potentially more complicated generation tasks like elaboration and longer replacements?

We split the test results of the two models by slots. For replacement, we report the average ADD score of the generated replacements. We measure the elaboration quality by the mean of ADD score and KEEP score and call it the elaboration score. *ctrlSIM* on test data produces an average ADD score of 13.24 for replacement, an ADD score of 9.85, a KEEP score of 23.63, and an overall elaboration score of 16.74 across all the data points. Overall, we observed that *ctrlSIM* is able to detect the contents to be deleted much better than addition with an average ALTDEL score of 39.27. Our low ADD score of the simplified texts, on an average and slot-wise, suggest that the model is unable to add content. One potential solution would be to concatenate additional context to the examples as another input slot. We have such a provision in our existing models that researchers can use in future work.

ctrlSIM_{ip} performed better deletion and replacement than *ctrlSIM* with ALTDEL score of 51.96 for deletion and average ADD score of 18.08 for replacement. The model displayed similar elaboration skill as *ctrlSIM* with an overall elaboration score of 16.53. Note that these results are

models	base	angle	SARI	ADD	DEL	KEEP	FKGL
Baseline	T5	E→S	39.14	9.34	67.8	40.29	11.1
<i>SIM</i>	T5	E→RXDIS	36.47	6.9	62.25	40.25	10.05
<i>SIM_{ip}</i>	T5	E→Sa	37.61	7.15	64.65	7.15	11.87
<i>ctrlSIM</i>	T5	<i>multi</i>	39.28	7.13	66.94	43.78	11.04
<i>ctrlSIM</i>	Sci5	<i>multi</i>	38.07	7.53	65.39	41.3	11.28
<i>ctrlSIM_{ip}</i>	T5	<i>multi_{ip}</i>	40.89	6.58	75.16	40.94	11.41
<i>ctrlSIM_{ip}</i>	Sci5	<i>multi_{ip}</i>	39.42	4.89	73.6	39.78	10.87

Table 1: Dev set performances for various model versions, *multi*: [E→S, E→DIS, ERi→DRS, ED→IS, EDXi→XS, ERi→RS, ERiXi→DRXS, E→DS, EXi→XS, ERiXi→RXS, EDRi→RS, EDRiXi→RXS, E→IS, ED→S, EXi→DXS], *multi_{ip}*: [E→Sa, Ea→Sa]

models	base	angle	SARI	ADD	DEL	KEEP	FKGL	ROUGE-I
<i>ctrlSIM</i>	T5	<i>multi</i>	39.63	8.99	70.67	39.2	10.55	0.41
<i>ctrlSIM_{ip}</i>	T5	<i>multi_{ip}</i>	40.2	8.51	70.07	42.04	11.09	0.43

Table 2: Test set performances for various model versions, *multi*: [E→S, E→DIS, ERi→DRS, ED→IS, EDXi→XS, ERi→RS, ERiXi→DRXS, E→DS, EXi→XS, ERiXi→RXS, EDRi→RS, EDRiXi→RXS, E→IS, ED→S, EXi→DXS], *multi_{ip}*: [E→Sa, Ea→Sa]

only from the test examples with controllable angle $Ea \rightarrow Sa$.

Qualitatively, we observe that many times *ctrlSIM_{ip}* replaces a text with the correct alternative, and other times fails to perform any replacement. Likewise, when asked to elaborate, sometimes *ctrlSIM_{ip}* is able to retrieve full forms of medical abbreviations from its prior knowledge, but at other times elaboration is just treated as a change of style (see Figure 1). Since some of these errors can only be obvious to humans, we resort to human evaluation, in particular expert evaluation.

Human Evaluation

We recruited three medical/biomedical experts to evaluate the quality of *ctrlSIM_{ip}*'s outputs. We also hired two layman native speakers to rate fluency and grammatical correctness. We sampled 50 outputs randomly from our test data for this purpose.

For fluency, we used a 4-point scale and a yes/no for grammaticality. We specifically asked our experts to judge the output annotation quality by answering 4 questions: one yes/no question and the remaining 3 to be rated on a 4-point scale (0-3). We asked,

- Did the model perform what the user asked for?
- Do the replacements in the output annotation match the replaced contents?
- Are the elaborations in the output annotation relevant to the content elaborated?
- Are the elaborations satisfactory?

Fluency: There is a moderate agreement between the 2 experts ($\tau = 0.397$, $p = 0.0018$). The mean scores of 2.32 and 2.22 indicate that the generated text is mostly fluent. **Grammaticality:** Both experts agreed that 40% of the model output is grammatically correct. **Output annotation Quality:** According to the experts, 22% of the time,

the model performed as expected. The mean scores from each expert (2.36, 1.86, and 2.04) indicate that model-generated content *mostly matches* the replaced content. Furthermore, the added content for elaboration is *mostly relevant* to the content elaborated (mean scores 2.33, 2.0, and 2.47), however, the elaborations were only *somewhat satisfactory* (mean scores 2.27, 1.2, and 1.73).

Summary

We are constantly fed with medical information online through news articles, popular science magazines, and social media posts. Despite volumes of medical articles being generated every day, low health literacy remains a challenge in healthcare. To make the medical texts more accessible, we create a finely annotated dataset Med-EASi for medical text simplification. Med-EASi consists of several pairs of expert medical texts and their annotated simplifications. It covers a wide range of medical topics and textual complexities and is annotated with four kinds of textual transformations: deletion, insertion, elaboration, and replacement.

We leverage the power and flexibility of large LMs like T5 to enable controllable text simplification, where the user can instruct the model to selectively simplify contents of a short medical text. We test two different kinds of controllability, one where the user inputs the content they want to be altered by a specific type of edit, and another, where the user can highlight in-place the same content and mark it with the desired transformation. Both of our controllable models perform at par with the top text simplification models. Our in-place controllable model displays promising results, generating mostly fluent and correct texts. The model is able to replace complex medical content appropriately.

Overall, we hope that Med-EASi can foster open research in AI-assisted medical text simplification. One of the potential future directions would be to improve the model's elaboration skills by supplying additional contexts like definitions and descriptions of complex medical terms.

Acknowledgements

We are grateful to Bhavana Dalvi Mishra (Senior Research Scientist at Allen Institute for Artificial Intelligence) for her invaluable guidance throughout the duration of this project. We would also like to thank the anonymous reviewers for their insightful feedback.

We acknowledge Toloka AI for partially funding our crowdsourcing tasks. This work has been partially supported by the Swiss National Science Foundation (SNSF) under contract number 200020.184994 (CrowdAlytics Research Project).

References

- Abrahamsson, E.; Forni, T.; Skeppstedt, M.; and Kvist, M. 2014. Medical text simplification using synonym replacement: Adapting assessment of word difficulty to a compounding language. In *Proceedings of the 3rd Workshop on Predicting and Improving Text Readability for Target Reader Populations (PITR)*, 57–65.
- Agrawal, S.; Xu, W.; and Carpuat, M. 2021. A non-autoregressive edit-based approach to controllable text simplification. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 3757–3769.
- Alva-Manchego, F.; Martin, L.; Bordes, A.; Scarton, C.; Sagot, B.; and Specia, L. 2020. ASSET: A Dataset for Tuning and Evaluation of Sentence Simplification Models with Multiple Rewriting Transformations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4668–4679. Online: Association for Computational Linguistics.
- Basu, C.; Vasu, R.; Yasunaga, M.; Kim, S.; and Yang, Q. 2021. Automatic Medical Text Simplification: Challenges of Data Quality and Curation. In *HUMAN@ AAAI Fall Symposium*.
- Berkman, N. D.; Sheridan, S. L.; Donahue, K. E.; Halpern, D. J.; and Crotty, K. 2011. Low health literacy and health outcomes: an updated systematic review. *Annals of internal medicine*, 155(2): 97–107.
- Bin Naeem, S.; and Kamel Boulos, M. N. 2021. COVID-19 misinformation online and health literacy: a brief overview. *International journal of environmental research and public health*, 18(15): 8091.
- Bodenreider, O. 2004. The unified medical language system (UMLS): integrating biomedical terminology. *Nucleic acids research*, 32(suppl_1): D267–D270.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Cao, Y.; Shui, R.; Pan, L.; Kan, M.-Y.; Liu, Z.; and Chua, T.-S. 2020. Expertise style transfer: A new task towards better communication between experts and laymen. *arXiv preprint arXiv:2005.00701*.
- Coster, W.; and Kauchak, D. 2011. Simple English Wikipedia: A New Text Simplification Task. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 665–669. Portland, Oregon, USA: Association for Computational Linguistics.
- Cumbicus-Pineda, O. M.; Gonzalez-Dios, I.; and Soroa, A. 2021. A Syntax-Aware Edit-based System for Text Simplification. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, 324–334.
- Dathathri, S.; Madotto, A.; Lan, J.; Hung, J.; Frank, E.; Molino, P.; Yosinski, J.; and Liu, R. 2019. Plug and play language models: A simple approach to controlled text generation. *arXiv preprint arXiv:1912.02164*.
- Dawid, A. P.; and Skene, A. M. 1979. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 28(1): 20–28.
- Devaraj, A.; Marshall, I.; Wallace, B.; and Li, J. J. 2021. Paragraph-level Simplification of Medical Texts. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4972–4984. Online: Association for Computational Linguistics.
- Grootendorst, M. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794*.
- Jiang, C.; Maddela, M.; Lan, W.; Zhong, Y.; and Xu, W. 2020. Neural CRF model for sentence alignment in text simplification. *arXiv preprint arXiv:2005.02324*.
- Joshi, M.; Levy, O.; Weld, D. S.; and Zettlemoyer, L. 2019. BERT for coreference resolution: Baselines and analysis. *arXiv preprint arXiv:1908.09091*.
- Kariuk, O.; and Karamshuk, D. 2020. Cut: Controllable unsupervised text simplification. *arXiv preprint arXiv:2012.01936*.
- Keskar, N. S.; McCann, B.; Varshney, L. R.; Xiong, C.; and Socher, R. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Kincaid, J. P.; Fishburne Jr, R. P.; Rogers, R. L.; and Chissom, B. S. 1975. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. Technical report, Naval Technical Training Command Millington TN Research Branch.
- Kindig, D. A.; Panzer, A. M.; Nielsen-Bohlman, L.; et al. 2004. *Health literacy: a prescription to end confusion*. Washington, DC: National Academies Press.
- King, A. 2010. Poor health literacy: a ‘hidden’ risk factor. *Nature Reviews Cardiology*, 7(9): 473–474.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Lyu, Y.; Liang, P. P.; Pham, H.; Hovy, E.; Póczos, B.; Salakhutdinov, R.; and Morency, L.-P. 2021. StylePTB: A Compositional Benchmark for Fine-grained Controllable Text Style Transfer. *arXiv preprint arXiv:2104.05196*.
- Maddela, M.; Alva-Manchego, F.; and Xu, W. 2020. Controllable text simplification with explicit paraphrasing. *arXiv preprint arXiv:2010.11004*.
- Mallinson, J.; Adamek, J.; Malmi, E.; and Severyn, A. 2022. EdiT5: Semi-Autoregressive Text-Editing with T5 Warm-Start. *arXiv preprint arXiv:2205.12209*.
- Mallinson, J.; Severyn, A.; Malmi, E.; and Garrido, G. 2020. Felix: Flexible text editing through tagging and insertion. *arXiv preprint arXiv:2003.10687*.
- Malmi, E.; Krause, S.; Rothe, S.; Mirylenka, D.; and Severyn, A. 2019. Encode, tag, realize: High-precision text editing. *arXiv preprint arXiv:1909.01187*.
- Martin, L.; Humeau, S.; Mazaré, P.; Bordes, A.; de la Clergerie, É. V.; and Sagot, B. 2019a. Reference-less Quality Estimation of Text Simplification Systems. *CoRR*, abs/1901.10746.
- Martin, L.; Humeau, S.; Mazare, P.-E.; De La Clergerie, É. V.; Bordes, A.; and Sagot, B. 2018. Reference-less Quality Estimation of Text Simplification Systems. In *Proceedings of the 1st Workshop on Automatic Text Adaptation (ATA)*, 29–38.

- Martin, L.; Sagot, B.; de la Clergerie, E.; and Bordes, A. 2019b. Controllable sentence simplification. *arXiv preprint arXiv:1910.02677*.
- Morris, J. X.; Lifland, E.; Yoo, J. Y.; Grigsby, J.; Jin, D.; and Qi, Y. 2020. Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in nlp. *arXiv preprint arXiv:2005.05909*.
- NCBI. 2022. PMC PubMed Central. <https://www.ncbi.nlm.nih.gov/pmc>. Accessed: 2023-04-01.
- NIA, N. 2018. Online Health Information: Is It Reliable? <https://www.nia.nih.gov/health/online-health-information-it-reliable>. Accessed: 2023-04-01.
- NIH. 2022. PubMed.gov. <https://pubmed.ncbi.nlm.nih.gov/>. Accessed: 2023-04-01.
- Nishihara, D.; Kajiwar, T.; and Arase, Y. 2019. Controllable text simplification with lexical constraint loss. In *Proceedings of the 57th annual meeting of the association for computational linguistics: Student research workshop*, 260–266.
- Nisioi, S.; Štajner, S.; Ponzetto, S. P.; and Dinu, L. P. 2017. Exploring neural text simplification models. In *Proceedings of the 55th annual meeting of the association for computational linguistics (volume 2: Short papers)*, 85–91.
- Nye, B.; Li, J. J.; Patel, R.; Yang, Y.; Marshall, I. J.; Nenkova, A.; and Wallace, B. C. 2018. A corpus with multi-level annotations of patients, interventions and outcomes to support language processing for medical literature. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2018, 197. NIH Public Access.
- Omelianchuk, K.; Raheja, V.; and Skurzshanskiy, O. 2021. Text simplification by tagging. *arXiv preprint arXiv:2103.05070*.
- Paetzold, G. H.; and Specia, L. 2017. A survey on lexical simplification. *Journal of Artificial Intelligence Research*, 60: 549–593.
- Phan, L. N.; Anibal, J. T.; Tran, H.; Chanana, S.; Bahadroglu, E.; Peltekian, A.; and Altan-Bonnet, G. 2021. SciFive: a text-to-text transformer model for biomedical literature. *arXiv:2106.03598*.
- Prabhumoye, S.; Tsvetkov, Y.; Salakhutdinov, R.; and Black, A. W. 2018. Style transfer through back-translation. *arXiv preprint arXiv:1804.09000*.
- Python. 2022. difflib — Helpers for computing deltas. <https://docs.python.org/3/library/difflib.html>. Accessed: 2023-04-01.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P. J.; et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140): 1–67.
- Reif, E.; Ippolito, D.; Yuan, A.; Coenen, A.; Callison-Burch, C.; and Wei, J. 2021. A recipe for arbitrary text style transfer with large language models. *arXiv preprint arXiv:2109.03910*.
- Savery, M.; Abacha, A. B.; Gayen, S.; and Demner-Fushman, D. 2020. Question-driven summarization of answers to consumer health questions. *Scientific Data*, 7(1): 1–9.
- Shen, T.; Lei, T.; Barzilay, R.; and Jaakkola, T. 2017. Style transfer from non-parallel text by cross-alignment. *arXiv preprint arXiv:1705.09655*.
- Siddharthan, A. 2014. A survey of research on text simplification. *ITL-International Journal of Applied Linguistics*, 165(2): 259–298.
- Soldaini, L.; and Goharian, N. 2016. Quickumls: a fast, unsupervised approach for medical concept extraction. In *MedDR workshop, sigir*, 1–4.
- Srikanth, N.; and Li, J. J. 2020. Elaborative simplification: Content addition and explanation generation in text simplification. *arXiv preprint arXiv:2010.10035*.
- Subramanian, S.; Lample, G.; Smith, E. M.; Denoyer, L.; Ranzato, M.; and Boureau, Y.-L. 2018. Multiple-attribute text style transfer. *arXiv preprint arXiv:1811.00552*.
- Tafjord, O.; and Clark, P. 2021. General-Purpose Question-Answering with Macaw. *ArXiv*, abs/2109.02593.
- Tajdar, D.; Lüthmann, D.; Fertmann, R.; Steinberg, T.; van den Bussche, H.; Scherer, M.; and Schäfer, I. 2021. Low health literacy is associated with higher risk of type 2 diabetes: a cross-sectional study in Germany. *BMC public health*, 21(1): 1–12.
- Toloka. 2022. Powering data-centric AI development. <https://toloka.ai/>. Accessed: 2023-04-01.
- Van, H.; Kauchak, D.; and Leroy, G. 2020. AutoMeTS: the autocomplete for medical text simplification. *arXiv preprint arXiv:2010.10573*.
- Van den Bercken, L.; Sips, R.-J.; and Lofi, C. 2019. Evaluating neural text simplification in the medical domain. In *The World Wide Web Conference*, 3286–3292.
- Warstadt, A.; Singh, A.; and Bowman, S. R. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7: 625–641.
- Welleck, S.; Kulikov, I.; Roller, S.; Dinan, E.; Cho, K.; and Weston, J. 2019. Neural text generation with unlikelihood training. *arXiv preprint arXiv:1908.04319*.
- Woodsend, K.; and Lapata, M. 2011. Learning to Simplify Sentences with Quasi-Synchronous Grammar and Integer Programming. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, 409–420. Edinburgh, Scotland, UK.: Association for Computational Linguistics.
- Xu, M.; Peng, M.; and Liu, F. 2022. Text style transfer between classical and modern chinese through prompt-based reinforcement learning. *World Wide Web*, 1–18.
- Xu, W.; Callison-Burch, C.; and Napoles, C. 2015. Problems in current text simplification research: New data can help. *Transactions of the Association for Computational Linguistics*, 3: 283–297.
- Xu, W.; Napoles, C.; Pavlick, E.; Chen, Q.; and Callison-Burch, C. 2016. Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics*, 4: 401–415.
- Zhang, X.; and Lapata, M. 2017. Sentence simplification with deep reinforcement learning. *arXiv preprint arXiv:1703.10931*.
- Zhu, Z.; Bernhard, D.; and Gurevych, I. 2010. A Monolingual Tree-based Translation Model for Sentence Simplification. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, 1353–1361. Beijing, China: Coling 2010 Organizing Committee.