

Prompt-based Learning for Text Readability Assessment

Bruce W. Lee^{1,2,‡}, Jason Hyung-Jong Lee²

¹University of Pennsylvania - PA, USA

²LXPER AI Research - Seoul, South Korea

brucelws@seas.upenn.edu

jasonlee@lxper.com

Abstract

We propose the novel adaptation of a pre-trained seq2seq model for readability assessment. We prove that a seq2seq model – T5 or BART – can be adapted to discern which text is more difficult from two given texts (pairwise). As an exploratory study to prompt-learn a neural network for text readability in a *text-to-text* manner, we report useful tips for future work in seq2seq training and ranking-based approach to readability assessment. Specifically, we test nine input-output formats/prefixes and show that they can significantly influence the final model performance.

Also, we argue that the combination of text-to-text training and pairwise ranking setup 1) enables leveraging *multiple* parallel text simplification data for teaching readability and 2) trains a neural model for the general concept of readability (therefore, better cross-domain generalization). At last, we report a 99.6% pairwise classification accuracy on Newsela and a 98.7% for OneStopEnglish, through a joint training approach. Our code is available at github.com/brucewlee/prompt-learning-readability.

1 Introduction

Readability assessment evaluates the reading difficulty of a given piece of text (Vajjala, 2021). The early traditional readability assessment methods like Flesch-Kincaid Grade Level (Kincaid et al., 1975) utilized a linear regression formula fitted to data from large-scale reading experiments on human subjects. More recently, readability assessment has often been viewed as a classification task (Feng et al., 2010). Under this classification-based task formulation, models using various handcrafted features (Xia et al., 2016; Vajjala and Meurers, 2012), computer-generated features (Martinc et al., 2021; Imperial, 2021), or both (Lee et al., 2021)

have been reported. Showing the potential that neural modeling is more suitable than handcrafted features in holistically capturing the inherent linguistic properties that affect readability.

Among the varying approaches to readability assessment, fine-tuning deep transformer models (Vaswani et al., 2017), that are pre-trained with language modeling objectives (e.g. *Masked Language Modelling*, *Next Sentence Prediction*), has proven highly effective in multiple reports (Lee and Vajjala, 2022; Lee et al., 2021). So far, encoder-only transformer architectures like BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) have been the go-to approach, and few reports discuss the applicability of other architecture types. Further, there is no report on how readability assessment can be cast in a text-to-text task formulation (§2). But recent reports (Raffel et al., 2020) show that a text-to-text is promising for multiple downstream tasks.

The main contribution of this paper is that we fine-tune full encoder-decoder transformer architectures (also referred to as *sequence-to-sequence*) and check if they can learn about text readability. A sequence-to-sequence model has been previously fine-tuned for downstream tasks like document ranking (Nogueira et al., 2020), but few reports discuss whether the architecture can learn about readability.

We fine-tune BART (Lewis et al., 2020) and T5 (Raffel et al., 2020) on the popular OneStopEnglish (Vajjala and Lučić, 2018) and Newsela (Xu et al., 2015) data. Then, we measure their performance on two other datasets with readability annotations for US and CEFR curricula, respectively.

We also conduct methodological explorations on how a sequence-to-sequence model can be well-trained to learn text readability. This includes how input and output format should be structured, considering that the fine-tuning for a sequence-to-sequence model is naturally cast in a text-to-text

[‡]Woong Sung (Bruce) Lee was on leave from the University of Pennsylvania during the research period.

format (Nogueira et al., 2020). We summarize our research questions in two:

- 1) Can a sequence-to-sequence model be fine-tuned for text readability – with a parallel text simplification dataset?
- 2) If so, can the performance generalize across domains? In other words, is the model learning the *dataset* or the concept of *readability*?

These research questions and our task approach are intentionally formulated in reference to what previous literature (Vajjala, 2021; Lee et al., 2021) have proposed as future directions. As we elaborate further in the following sections, our approach simplifies some inherent problems that we had in the readability assessment. But our study has limitations and requires further explorations (§5).

2 Background Knowledge

2.1 Sequence-to-Sequence Transformers

Pre-trained sequence-to-sequence transformers essentially incorporate both encoder and decoder parts from the original Transformer (Vaswani et al., 2017) architecture. The most notable examples are T5 and BART, pre-trained using document denoising strategies (i.e. during pre-training, input is an intentionally corrupted document, and output is a recovered document or corrupted portions).

Though BART allows some flexibility in altering the model architecture for downstream tasks (Lewis et al., 2020), T5 is built with the intention of unifying all NLP tasks into a text-to-text-format (Raffel et al., 2020). Here, text-to-text means that both input and output are always text strings, unlike encoder-only, BERT-style models that only output either a class label ([CLS] token) or an input span.

2.2 Existing Downstream Tasks

Pre-trained sequence-to-sequence models can be fine-tuned to most NLP downstream tasks, including neural machine translation (Wang et al., 2022) and abstractive summarization (Saito et al., 2020). Then, a training instance is formatted in a way that is much like *telling* the model what to do by adding task-specific prefix (Raffel et al., 2020).

If a model were to be fine-tuned for translation, the input format could be “*translate English to German: That is good.*” and the target output is “*Das ist gut.*” When applied to document ranking, Nogueira et al. (2020) proposed a slightly different format. The input format was written “*Query: ...*

Document: ... Relevant:” so that the target output tokens – “*true*” or “*false*” – can naturally come after the input format. Our work is the most influenced by this formatting approach.

3 Experimental Setup

3.1 Methods

All our experiments are based on T5 and BART, both obtained from the respective online repositories through Huggingface (Wolf et al., 2019). Since the text-to-text formulation has never been used to fine-tune text readability¹, here we limit to the simple task of comparing the readability of two given texts. Following our input format (Table 1), we fed two text snippets of varying difficulties to the model every instance. Then, the model was trained to give the corresponding target output. We tested nine input-output formats, as shown in Table 1.

3.2 Datasets and Preparation

3.2.1 Data Type and Permutation Methods

The datasets that we use in this study are of two types. *Parallel* type contains multiple reading-level versions of a text (mostly through human expert paraphrasing). We call a grouping of text in multiple reading levels a *slug*. On the other hand, there is *distinct* type of datasets. *Distinct* type is a more common format where each text is given a readability level with no multiple versions of the same content. Our naming and permutation strategies are inspired by existing work on pairwise ranking for readability (Lee and Vajjala, 2022).

A parallel dataset $D_{parallel}$ can be expressed as a row-wise collection of i slugs $D_{parallel} = [S_1, \dots, S_i]$, where a slug is a column-wise collection of j pairs of a text and a reading level $S_i = [(x_{i,1}, y_{i,1}), \dots, (x_{i,j}, y_{i,j})]$. For parallel dataset, we perform permutation jP_2 on the slug level, creating $S'_i = [((x_{i,1}, x_{i,2}), (y_{i,1}, y_{i,2})), \dots, ((x_{i,j-1}, x_{i,j}), (y_{i,j-1}, y_{i,j}))]$. A pair like $((x_{i,1}, x_{i,2}), (y_{i,1}, y_{i,2}))$ is considered an instance for train/dev/test. Then, $D'_{parallel} = [S'_1, \dots, S'_i]$, where all S' are flattened to make $D'_{parallel}$ an iterable collection of tuples. This setup is intended to be robust to future implementations of paraphrase-based text simplification datasets where the standards for readability annotation/level are only consistent within a slug.

A distinct dataset $D_{distinct}$ can be expressed as a collection of i pairs of a text and a reading

¹Lee et al. (2021) trains BART for readability but uses class labels (BERT-style) as target output, instead of text tokens

Type	Input Format	Target Output
Question	"Which Text is more difficult? Text 1: ... Text 2: ..."	"Text 1" or "Text 2"
Statement	"Text 1 is more difficult than Text 2. Text 1: ... Text 2: ..."	"True" or "False"
Follow-up	"Text 1: ... Text2: ... More difficult:"	"Text 1" or "Text 2"
Reverse-Question	"Which Text is easier? Text 1: ... Text 2: ..."	"Text 2" or "Text 1"
Reverse-Statement	"Text 1 is easier than Text 2. Text 1: ... Text 2: ..."	"False" or "True"
Reverse-Follow-up	"Text 1: ... Text2: ... Easier:"	"Text 2" or "Text 1"
Alternate-Question	"Which Text is harder? Text 1: ... Text 2: ..."	"Text 1" or "Text 2"
Alternate-Statement	"Text 1 is harder than Text 2. Text 1: ... Text 2: ..."	"True" or "False"
Alternate-Follow-up	"Text 1: ... Text2: ... Harder:"	"Text 1" or "Text 2"

Table 1: Input-Output format candidates we tested. The text-to-text formulation is intuitive internally (model) and externally (human) because the model input and output are both representations of some semantic concept.

level $D_{distinct} = [(x_1, y_1), \dots, (x_i, y_i)]$. Then, we can perform permutation iP_2 to create $D'_{distinct} = [((x_1, x_2), (y_1, y_2)), \dots, ((x_{i-1}, x_i), (y_{i-1}, y_i))]$. This setup is under the postulation that readability annotation is consistent throughout the dataset. Pairwise instances of two same levels are removed.

3.2.2 Datasets

Two *parallel* and two *distinct* datasets are used. Also, we agree with Vajjala (2021)'s concern that "(in some datasets) articles tagged with different reading levels don't share the same topical content (...) question what the readability assessment models learn - is it a notion of text complexity, or topical differences among texts?". Hence, we only use parallel data – NEWS and OSEN – for training.

Our dataset processing strategy (§3.2.1) and pairwise comparison approach force the model to learn label-agnostic, the global concept of relative difficulties of texts. That is, the model learns that a text annotated level 5 should be harder than a level 3 (or a level 2 or a level 1) within a slug or a dataset. Such a setup is inherently robust against cross-domain usage (Table 3). Further, it enables combining multiple datasets of various slug sizes or readability annotation standards for joint training (§4).

Newsela(NEWS) is a *parallel* text simplification dataset introduced by Xu et al. (2015). It consists of 2,154 slugs, each item re-written 4 or 5 times for children at different grade levels. Hence, a total of 10,786 texts are contained, and 43,316 pairwise instances are created after data permutation (§3.2.1). Random shuffling split these instances into 6:2:2 for train/test splits. We provide reproducible scripts for all datasets through code.

OneStopEnglish(OSEN) is a *parallel* dataset intended for both text simplification and readability assessment research (Vajjala and Lučić, 2018). It

Format Type	OSEN		NEWS	
	T5	BART	T5	BART
Question	0.815(30)	0.965(18)	0.981(3)	-
Statement	0.639(29)	0.978(30)	-	-
Follow-up	0.784(25)	0.960(27)	-	-
Reverse-Q	0.793(30)	0.991(30)	-	0.993(3)
Reverse-S	0.524(28)	0.978(26)	-	-
Reverse-F	0.828(30)	0.991(30)	-	0.993(5)
Alternate-Q	0.811(30)	0.960(25)	-	-
Alternate-S	0.617(29)	0.978(30)	-	-
Alternate-F	0.789(28)	0.938(29)	-	-

Table 2: Validation set accuracy reports on NEWS and OSEN. The best epoch is reported in brackets. NEWS is only reported for the best format type due to data size.

consists of 189 slugs, each item in 3 paraphrases at different reading levels. A total of 567 texts are contained, and 1,134 pairwise instances are created. We use a 6:2:2 split ratio through random shuffling.

Common Core Standards(CCSB) is a *distinct* collection of exemplary official texts with readability annotations in U.S. Common Core Standards. We scraped data from the source ourselves. We used 69 story-type texts in 6 reading levels. After permutation, 3,846 pairwise instances are created.

Cambridge English Readability(CAMB) is a *distinct* dataset of reading passages from main suite Cambridge English Exams (Xia et al., 2016). All 331 texts are labeled A2, B1, B2, C1, or C2 reading levels, following the CEFR standards. After permutation, 87,574 pairwise instances are created.

3.3 Training

The batch size is fixed at 8, both for training and inference. The learning rate is fixed at 1e-5 for T5 and BART. We fine-tune OSEN for 30 epochs and NEWS for 3 epochs. We report the best epoch performance based on the validation set in Table 2. For joint training (Table 3), we take an OSEN-trained model and then fine-tune further using NEWS for 3 more epochs.

Model / Fine-Tune Data	Test Data			
	OSEN	NEWS	CCSB	CAMB
Flesch-Kincaid / None	0.978	0.986	0.798	0.808
T5 / OSEN	<u>0.784</u>	<u>0.518</u>	<u>0.509</u>	<u>0.492</u>
BART / OSEN	<u>0.978</u>	<u>0.871</u>	<u>0.639</u>	<u>0.629</u>
T5 / NEWS	<u>0.907</u>	<u>0.967</u>	<u>0.747</u>	<u>0.764</u>
BART / NEWS	0.987	<u>0.993</u>	<u>0.793</u>	0.883
T5 / OSEN + NEWS	<u>0.992</u>	<u>0.987</u>	<u>0.771</u>	<u>0.778</u>
BART / OSEN + NEWS	<u>0.983</u>	0.996	<u>0.790</u>	<u>0.865</u>

Table 3: In-domain and *cross-domain* accuracies across datasets. For OSEN and NEWS, test sets (§3.2.2) are used. The best result per dataset is in **bold**. T5 is trained with `Question` format, whereas BART is trained with `Reverse-F` format. Flesch-Kincaid refers to the popular Flesch-Kincaid Grade Level formula published in Kincaid et al. (1975). We use the implementation in `github.com/textstat/textstat`.

4 Results

1. **A pretrained sequence-to-sequence model could be fine-tuned for text readability, in a text-to-text style.** Table 2 shows that the concept of readability could be fine-tuned in a text-to-text task formulation, some setups with decent accuracies of > 0.9 . For a smaller dataset (OSEN), BART significantly outperformed T5, but their performance deviation was little on a larger dataset (NEWS). We believe this is caused due to difference in pre-training methodologies that caused T5 to require more training steps to learn about our downstream task. Also, BART always generalized better than T5 across unseen datasets in Table 3.

2. **Input-output format significantly affected the final performance, especially when fine-tuning T5 with lesser training steps (OSEN).** Among the nine input-output formats we tested, T5 and BART performed best under `Question` and `Reverse-Q/F` types, respectively. Performance deviations caused by input-output format changes were larger than we expected. Further, no certain format generally ensured good results across models. This raises the need for additional "format-tuning" processes when exploring new models. However, it must be noted that several observations point to T5 being under-trained for the general concept of readability at the data size of OSEN (see Table 2 and Table 3). The input-output format's influence is lesser for setups where models learned better about readability.

3. **Joint training has the potential to help both in-domain and cross-domain performances.** Joint training of multiple datasets for a single model is an under-explored concept in readabil-

ity assessment. This is because human experts annotate existing datasets with varying standards. Dataset construction can also differ (e.g. different number of classes or too difficult to map classes). Hence, it was unknown if combining datasets of varying labeling standards improves performance.

This work solves the problem by re-casting the task into a simple, universal question of comparing two texts' difficulties (§3.2.2). Table 3 shows that in- and cross-domain performances can improve through joint training. For example, in-domain accuracies for OSEN increased to 0.208 when the model was further fine-tuned with a larger extra data, NEWS. However, a NEWS-only model generally performed better than the OSEN+NEWS model in Table 3. We expect that OSEN, which is almost 40 times smaller, only confused the model.

4. **Exposing the model to more texts with a wider range of readability helped fine-grained readability comparison.** Importantly, we showed that readability assessment models fine-tuned with *parallel* datasets could be generalized across *distinct* datasets (e.g. OSEN \rightarrow CCSB). But model performances varied depending on label distance. Models performed better when the two compared texts' readability labels were larger apart (i.e. the model is more likely to guess level 1 vs level 4 correctly than level 1 vs level 2). This problem worsened when the model was trained using OSEN. Using NEWS as training data or extra data helped. We want to point out that a slug size in NEWS is 4 or 5, exposing the model to more permutations.

5. **Text-to-text style fine-tuning required more training steps than expected.** The majority of our OSEN fine-tuning experiments showed that the model's validation set performance continues to increase up to epoch 30. This is contrastive to how usual classification approaches, using encoder-only models, only fine-tune up to epoch 3~5 even on smaller datasets like OSEN or CAMB (Lee et al., 2021). Intuitively speaking, there is potential that better performance can be achieved if fine-tuned further. We will explore this concept in the future.

6. **Though often overlooked, traditional readability formulas provide challenging baselines.** The traditional readability formulas are criticized for their low performances in multi-class ranking or regression-based readability task formulation (Lee and Lee, 2023). However, they provide surprisingly strong baselines for pairwise difficulty comparisons, as seen in Table 3.

5 Conclusion

So far, we have reported our exploratory work on casting readability assessment tasks in a text-to-text formulation for BART and T5. We summarized our observations into five categories in §4, which can serve as base guidelines for future work. Our experimental setup and data permutation methods allow the joint training of more than one dataset, regardless of whether the dataset construction is *parallel* or *distinct* (§3). Using NEWS as extra training data further to fine-tune an OSEN-trained model greatly improved model performance. However, we did not train the other way around (NEWS → OSEN), which should be proved in the future.

6 Limitations

Our limitations are in input text length and output labels. Though our novel task formulation allows the application of essential concepts like *joint training* or *cross-domain evaluation* in the field of readability assessment, it is based on a pairwise classification method. Since the pairwise approach only allows the readability ranking of two texts (e.g. which is easier?), it lacks practicality compared to regression or multi-label classification-based models. Though we achieve an almost perfect accuracy of 99.6% in Newsela data, knowing which is easier out of texts has little use as a real-world system. Hence, further research must be conducted to generate **more useful output labels** and process longer sequences. Like Nogueira et al. (2020), we are looking into using a sliding window to generate output labels for longer input sequences. We are also researching neural models pre-trained specifically for readability assessment using the prompt-based learning method introduced in the paper. Such a model can be leveraged for multi-class classification.

References

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Lijun Feng, Martin Jansche, Matt Huenerfauth, and Noémie Elhadad. 2010. A comparison of features for automatic readability assessment. In *Coling 2010: Posters*, pages 276–284.
- Joseph Marvin Imperial. 2021. Bert embeddings for automatic readability assessment. *arXiv preprint arXiv:2106.07935*.
- 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. Technical report, Naval Technical Training Command Millington TN Research Branch.
- Bruce W. Lee, Yoo Sung Jang, and Jason Lee. 2021. [Pushing on text readability assessment: A transformer meets handcrafted linguistic features](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10669–10686, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Bruce W Lee and Jason Hyung-Jong Lee. 2023. Traditional readability formulas compared for english. *arXiv preprint arXiv:2301.02975*.
- Justin Lee and Sowmya Vajjala. 2022. A neural pairwise ranking model for readability assessment. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3802–3813.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Matej Martinc, Senja Pollak, and Marko Robnik-Šikonja. 2021. [Supervised and Unsupervised Neural Approaches to Text Readability](#). *Computational Linguistics*, pages 1–39.
- Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. 2020. Document ranking with a pre-trained sequence-to-sequence model. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 708–718.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.

- Itsumi Saito, Kyosuke Nishida, Kosuke Nishida, and Junji Tomita. 2020. Abstractive summarization with combination of pre-trained sequence-to-sequence and saliency models. *arXiv preprint arXiv:2003.13028*.
- Sowmya Vajjala. 2021. Trends, limitations and open challenges in automatic readability assessment research. *arXiv preprint arXiv:2105.00973*.
- Sowmya Vajjala and Ivana Lučić. 2018. Onestopenglish corpus: A new corpus for automatic readability assessment and text simplification. In *Proceedings of the thirteenth workshop on innovative use of NLP for building educational applications*, pages 297–304.
- Sowmya Vajjala and Detmar Meurers. 2012. [On improving the accuracy of readability classification using insights from second language acquisition](#). In *Proceedings of the Seventh Workshop on Building Educational Applications Using NLP*, pages 163–173, Montréal, Canada. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Wenxuan Wang, Wenxiang Jiao, Yongchang Hao, Xing Wang, Shuming Shi, Zhaopeng Tu, and Michael Lyu. 2022. Understanding and improving sequence-to-sequence pretraining for neural machine translation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2591–2600.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Menglin Xia, Ekaterina Kochmar, and Ted Briscoe. 2016. [Text readability assessment for second language learners](#). In *Proceedings of the 11th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 12–22, San Diego, CA. 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.

A Obtaining Dataset

We obtained Newsela by requesting academic access at newsela.com/data.

OneStopEnglish dataset is freely available at github.com/nishkalavallabhi/OneStopEnglishCorpus.

Cambridge English Readability dataset is freely available at ilexir.co.uk/datasets/index.html.

We crawled Common Core Appendix B from www.corestandards.org/assets/Appendix_B.pdf.

B Training Details

B.1 Models, GPU, Train Time

All pre-trained models are retrieved from Huggingface. Fine-tuning code is written in PyTorch. We used a single NVIDIA RTX 2080 GPU for all our training.

T5

- huggingface.co/t5-base
- 1,020 seconds for 30 epochs on permuted OneStopEnglish, in `Question` format
- 6,994 seconds for 3 epochs on permuted Newsela, in `Question` format

BART

- huggingface.co/facebook/bart-base
- 532 seconds for 30 epochs on permuted OneStopEnglish, in `Reverse-F` format
- 4,084 seconds for 3 epochs on permuted Newsela, in `Reverse-F` format

B.2 More on Input Sequence Length

The current experimental setup is inherently weak against long texts. This is because we input a set of two texts with format prefixes (Table 1). Hence, for a model with 512 max token length our actual token length limitation per text is ≤ 256 , with the exact number depending on chosen format type. Though both BART and T5 support longer input sequences, in comparison to other popular models like BERT or RoBERTa, we must empirically confirm if good performance can be achieved when BART and T5 are used with longer max sequence length. In this paper, the max sequence length was set to 512 (which means that ≤ 256 token limit existed per text), and most texts from OSEN had to be truncated before training.