

AutoMeTS: The Autocomplete for Medical Text Simplification

Hoang Van, David Kauchak*, Gondy Leroy

University of Arizona, *Pomona College

{vnhh, gondyleroy}@email.arizona.edu, david.kauchak@pomona.edu

Abstract

The goal of text simplification (TS) is to transform difficult text into a version that is easier to understand and more broadly accessible to a wide variety of readers. In some domains, such as healthcare, fully automated approaches cannot be used since information must be accurately preserved. Instead, semi-automated approaches can be used that assist a human writer in simplifying text faster and at a higher quality. In this paper, we examine the application of autocomplete to text simplification in the medical domain. We introduce a new parallel medical data set consisting of aligned English Wikipedia with Simple English Wikipedia sentences and examine the application of pretrained neural language models (PNLMs) on this dataset. We compare four PNLMs (BERT, RoBERTa, XLNet, and GPT-2), and show how the additional context of the sentence to be simplified can be incorporated to achieve better results (6.17% absolute improvement over the best individual model). We also introduce an ensemble model that combines the four PNLMs and outperforms the best individual model by 2.1%, resulting in an overall word prediction accuracy of 64.52%.

1 Introduction

The goal of text simplification is to transform text into a variant that is more broadly accessible to a wide variety of readers while preserving the content. While this has been accomplished using a range of approaches (Shardlow, 2014), most text simplification research has focused on fully-automated approaches (Xu et al., 2016; Zhang and Lapata, 2017; Nishihara et al., 2019). However, in some domains, such as healthcare, using fully-automated text simplification is not appropriate since it is critical that the important information is preserved fully during the simplification process. For example, Shardlow et al. (2019) found that fully automated approaches omitted 30% of critical information when used to simplify clinical texts. For these types of domains, instead of fully-automated approaches, interactive text simplification tools are better suited to generate more efficient and higher quality simplifications (Kloehn et al., 2018).

Autocomplete tools suggest one or more words during text composition that could follow what has been typed so far. They have been used in a range of applications including web queries (Cai et al., 2016), database queries (Khoussainova et al., 2010), texting (Dunlop and Crossan, 2000), e-mail composition (Chen et al., 2019), and interactive machine translation, where a user translating a foreign sentence is given guidance as they type (Green et al., 2014). Our work is most similar to interactive machine translation. Autocomplete tools can speed up the text simplification process and give full control over information preservation to users, which is required in some domains, such as health and medical.

In this paper, we explore the application of pretrained neural language models (PNLMs) to the autocomplete process for sentence-level medical text simplification. Specifically, given (a) a difficult sentence a user is trying to simplify and (b) the simplification typed so far, the goal is to correctly suggest the next simple word to follow what has been typed. Table 1 shows an example of a difficult sentence along with a simplification that the user has partially typed. An autocomplete model predicts the next word to assist in finishing the simplification, in this case a verb like “take”, which might be continued to a

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

Difficult	Lowered glucose levels result both in the reduced release of insulin from the beta cells and in the reverse conversion of glycogen to glucose when glucose levels fall.
Typed	This insulin tells the cells to _____

Table 1: An example of text simplification Autocomplete task. The user is simplifying the difficult sentence on top and has typed the words on the bottom so far. The example is taken from a medical parallel English Wikipedia sentence pair in Table 2.

Difficult	Lowered glucose levels result both in the reduced release of insulin from the beta cells and in the reverse conversion of glycogen to glucose when glucose levels fall.
Simple	This insulin tells the cells to take up glucose from the blood.

Table 2: An example of sentence pair in Medical Wikipedia parallel corpus.

partial simplification of “take up glucose”. By suggesting the next word, the autocomplete models provide appropriate guidance while giving full control to human experts in simplifying text. We explore this task in the health and medical domain where information preservation is a necessity.

We make three main contributions:

1. We introduce a new parallel medical data set consisting of aligned English Wikipedia and Simple Wikipedia sentences, which is extracted from the commonly used general Wikipedia parallel corpus (Kauchak, 2013). The resulting medical corpus has 3.3k sentence pairs. This corpus is larger than previously generated corpora (by over 1k sentence pairs) and has stricter quality control (Van den Bercken et al., 2019). Our corpus requires a medical sentence to contain 4 or more medical words and belong to medical titles as compared to the no title requirement and needing to contain only 1 medical word, as described in Van den Bercken et al. (2019).
2. We examine the use of PNLMs for the autocomplete task on sentence-level text simplification and provide an initial analysis based on four recent models on this new medical corpus. In traditional autocomplete tasks, only the text being typed is available. For text simplification, the additional context of the difficult sentence is also available. We show how this additional information can be integrated into the models to improve the quality of the suggestions made.
3. We introduce an ensemble model that combines the output of the four PNLMs and outperforms all of the individual models. The ensemble approach is not application specific and may be used in other domains where PNLMs have been employed.

2 Medical Parallel English Wikipedia Corpus Creation

We automatically extracted medical sentence pairs from the sentence-aligned English Wikipedia corpus from Kauchak et al. (2013). Table 2 shows an example medical sentence pair. To identify the medical sentence pairs, we first created a medical dictionary with 260k medical terms selected from the Unified Medical Language System (UMLS) (Bodenreider, 2004) by selecting all UMLS terms that were associated with the semantic types of: Disease or Syndrome, Clinical Drug, Diagnostic Procedure, and Therapeutic or Preventive Procedure.

We extracted sentences from the English Wikipedia corpus based on the occurrence of terms in the medical dictionary. Specifically, a sentence pair was identified as medical if both the title of the article and the English Wikipedia sentence had 4 or more terms that matched the medical keywords. A term was considered a match to the UMLS dictionary if it had a similarity score higher than 0.85 using QuickUMLS (Soldaini and Goharian, 2016). We then reviewed the sentence pairs and removed all non-medical sentence

pairs. The final medical parallel corpus has 3.3k aligned sentence pairs¹.

Van den Bercken et al. (2019) also created a parallel medical corpus by filtering sentence pairs from Wikipedia. Our corpus is significantly larger (45% larger; 2,267 pairs vs. 3,300 pairs) and uses a stricter criteria for identifying sentences: they only required a single word match in the text itself (not the title) and used a lower similarity threshold of 0.75 (vs. our approach of 0.85).

3 Approach

We had three goals for our analysis: understand the effect of incorporating the additional context of the difficult sentence into autocomplete text simplification models, explore a new application for PNLMs, and evaluate our new ensemble approach to the autocomplete text simplification.

3.1 Autocomplete Approach For Medical Text Simplification

We pose the autocomplete text simplification problem as a language modeling problem: given a difficult sentence that a user is trying to simplify, $d_1d_2\dots d_m$, and the simplification typed so far, $s_1s_2\dots s_i$, the autocomplete task is to suggest word s_{i+1} . Table 1 gives an example of the autocomplete task. To evaluate the models, we calculated how well the models predicted the next word in a test sentence, given the previous words. A simple test sentence of length n , $s_1s_2\dots s_n$, would result in $n - 1$ predictions, i.e., predict s_2 given s_1 , predict s_3 given s_1s_2 , etc. For example, Table 2 shows a difficult sentence from English Wikipedia and the corresponding simplification from the medical Simple English Wikipedia. Given this test example, we generate 12 prediction tasks, one for each word in the simple sentence after the first word. Table 3 shows these 12 test prediction tasks.

3.2 Increasing Information Through Incorporation of Additional Context

Unlike other autocomplete tasks, for text simplification, the difficult sentence provides very explicit guidance about what words and information should be included as the user types. As a baseline, we compare the models without the context of this additional information, i.e., we predict word s_{i+1} given only the simplification typed so far, $s_1s_2\dots s_i$. We take a straightforward approach to incorporate the context of the difficult sentence: we concatenate the difficult sentence in front of a simplification typed so far, i.e., predict s_{i+1} given $d_1d_2\dots d_m.s_1s_2\dots s_i$. This has the advantage of incorporating the difficult sentence and biasing the predictions towards those found in the encoded context from difficult sentences, but is still model-agnostic, allowing us to apply it to all the different PNLMs without modifying the underlying architecture.

3.3 Transformer-based Language Models

We examined four PNLMs based on the Transformer network (Vaswani et al., 2017): BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019), XLNet (Yang et al., 2019), and GPT-2 (Radford et al., 2019). For each of the models, we examine versions that only utilize the text typed so far, denoted “No Context”, as well as “context-aware” variants that utilize the difficult sentence. We fine-tuned all four models on the 160k sentence pair general parallel English Wikipedia (Kauchak, 2013) (excluding the development and testing data) and then further fine-tuned them on the separate medical training set described in section 2. Note that none of the test sentences were in a dataset used for fine-tuning.

BERT: Bidirectional Encoder Representations from Transformers (Devlin et al., 2018) is a method for learning language representations using bidirectional training. BERT has been shown to produce state-of-the-art results in a wide range of generation and classification applications. We use the base original BERT² model pre-trained on the BooksCorpus (Zhu et al., 2015) and English Wikipedia.

RoBERTa: A Robustly Optimized BERT Pretraining Approach (Liu et al., 2019). The RoBERTa uses the same model architecture as BERT. However, the differences between RoBERTa and BERT are that

¹https://github.com/vanh17/MedTextSimplifier/tree/master/data_processing/data

²<https://github.com/huggingface/bert/>

Difficult sentence	Lowered glucose levels result both in the reduced release of insulin from the betacells and in the reverse conversion of glycogen to glucose when glucose levels fall.	
Prediction Task	Simplification typed so far	Predict
1	This	insulin
2	This insulin	tells
3	This insulin tells	the
4	This insulin tells the	cells
5	This insulin tells the cells	to
6	This insulin tells the cells to	take
7	This insulin tells the cells to take	up
8	This insulin tells the cells to take up	glucose
9	This insulin tells the cells to take up glucose	from
10	This insulin tells the cells to take up glucose from	the
11	This insulin tells the cells to take up glucose from the	blood
12	This insulin tells the cells to take up glucose from the blood	.

Table 3: The resulting prediction tasks that are generated from the example in Table 2.

RoBERTa does not use Next Sentence Prediction during pre-training and RoBERTa uses larger mini-batch size. We used the publicly released base RoBERTa³ with 125M parameters model.

XLNet: Generalized Auto-regressive Pretraining Method (Yang et al., 2019). Like BERT, XLNet benefits from bidirectional contexts. However, XLNet does not suffer limitations of BERT because of its auto-regressive formulation. In this work, we used publicly available base English XLNet⁴ with 110M parameters model.

GPT-2: Generative Pretrained Transformer 2 (Radford et al., 2019). Like BERT, GPT-2 is also based on the Transformer network, however, GPT-2 uses unidirectional left-to-right pre-training process. We use the publicly released GPT-2⁵ model, which has 117M parameters and is trained on web text.

3.4 Ensemble Models

Each of the models above utilizes different network variations and was pretrained on different datasets, and therefore they do not always make the same predictions. Ensemble approaches combine the output of different systems to try and leverage these differences to create a single model that outperforms any of the individual models. We examined three ensemble approaches that combine the output of the four models.

Majority Vote: As a baseline ensemble approach, we examined a simple majority vote on what the next word should be. We take the top 5 suggestions from each of the models and do a majority count on the pool of combined suggestions. The output of the model is the suggestion with highest count. If there is a tie, we randomly select one of the top suggestions. We picked the top 5 suggestions since this was the cutoff where the models tended to peak on the development data (for example, see the accuracy@N as shown in Table 7. Having more than 5 suggestions did not improve performance much but slowed the computation.

4-Class Classification (4CC): The ensemble problem can be viewed as a classification problem where the goal is to predict which system output should be used given a difficult sentence and the words typed so far, i.e., the autocomplete example. We posed this as a supervised classification problem. Given an autocomplete text simplification example, we can generate training data for the classifier by comparing

³<https://github.com/huggingface/roberta>

⁴<https://github.com/huggingface/xlnet>

⁵<https://github.com/huggingface/gpt2>

Prediction Task	Class
(Difficult sentence). This (MASK)	RoBERTa
(Difficult sentence). This insulin (MASK)	BERT
(Difficult sentence). This insulin tells (MASK)	XLNet

Table 4: An example of training data for the 4CC model. Class can be one of the four option: BERT, RoBERTa, XLNet, GPT-2.

the output of each system to the correct answer. If a system does get the example correct, then we include an example with that system as the label. Table 4 shows three such examples, where RoBERTa correctly predicted the first example, BERT the second, and XLNet the third. If multiple systems get the example right, we then randomly assign the label to one of the systems.

We train a neural text classifier implemented by huggingface⁶ with this training set to the PNLMs given the next-word prediction task. To make use of the models’ confidence on top of the results from model selection, we designed a scoring system for output selection as follows:

$$Score(w, X) = \alpha * P(w|X) + \theta * I(X, S)$$

where $P(w|X)$ is model X ’s confidence on predicted word w ; $I(X, S)$ is an identity function, which returns 1 if $X = S$ and 0 otherwise; S is the predicted model from model selector; and α and θ are scoring parameters. We use 0.5 for both α and θ . At testing time, we pick the highest score and output the word w , given a prediction task.

Autocomplete for Medical Text Simplification (AutoMeTS): RoBERTa performed significantly better than the other three individual models at the simplification autocomplete task. As a result, there was a strong bias toward RoBERTa in the training data for the 4CC ensemble model. To mitigate this effect, we also developed an ensemble approach based on a multi-label classifier for model selector, which we denote the AutoMeTS ensemble model. This choice of model selector, to our knowledge, is novel to transformer-based ensemble models. For this choice of classifier, each prediction task is given a sequence of 4 binary labels. Each label represents the correctness of each of the individual PNLMs, with a 0 representing an incorrect prediction on the task and a 1 representing a correct prediction. Table 5 shows an example of this dataset with the labels in order “RoBERTa BERT XLNet GPT-2”. For the first example, RoBERTa, XNLET, and GPT-2 correctly predicted the next word, while BERT did not.

We trained a neural multi-label classifier implemented by huggingface on this training dataset. To make use of the models’ confidence on top of the results from model selection, we designed a scoring system for output selection as follows:

$$Score(w, X) = \beta * P(w|X) + \sigma * S(X, Ls)$$

where $P(w|X)$ is model X ’s confidence on predicted word w ; $S(X, Ls)$ is a function, which returns 0.25 if model X is in Ls and 0 otherwise; Ls is the predicted sequence of labels from the model selector; and β and σ are scoring parameters. We use 0.5 for both β and σ . At testing time, we output the word w with the highest score, given a prediction task.

4 Experimental Setup

We compare the performance of the models on the medical autocomplete text simplification task. We used our medical parallel English Wikipedia corpus with 70% of the sentence pairs for training, 15% for development, and 15% for testing. We fine-tuned individual PNLMs using huggingface⁷ with a batch-size of 8, 8 epochs, and a learning rate of $5e^{-5}$. Early stopping was used based on the second time a decrease in the accuracy was seen.

⁶<https://github.com/huggingface/transformers>

⁷<https://github.com/huggingface/>

Prediction Task	Sequence of Labels
(Difficult sentence). This (MASK)	1 0 1 1
(Difficult sentence). This insulin (MASK)	0 1 0 0
(Difficult sentence). This insulin tells (MASK)	1 1 1 1

Table 5: An example of training data for the AutoMeTS model. For a prediction task, a sequence of 4 labels is give in the order "RoBERTa BERT XLNet GPT-2". The value of 1 means the model correctly predicted the right word, and 0 otherwise.




Figure 1: Accuracy for the context-aware models based on the length of the difficult sentences: very short (≤ 5 tokens), short (6 – 15), medium (16 – 19), and long (≥ 20).

We used two metrics to evaluate the quality of the approaches. First, we used standard accuracy, where a prediction is counted correct if it matches the test prediction word. Accuracy is pessimistic in that the predicted word must match exactly the word seen in the simple sentence, and as such it does not account for other possible words, such as synonyms, that could be correctly used in the context. Since the parallel English Wikipedia corpus does not offer multiple simplified versions for a given difficult sentence, accuracy is the best metric that considers automated scoring, simplification quality, and information preservation. Accuracy-based metrics can help offset an expensive manual evaluation while providing the best approximation of how the autocomplete systems work. We do not use BLEU (Papineni et al., 2002) and SARI (Xu et al., 2016) scores, which are widely used in text simplification domain, because the two metrics are specifically designed for fully-automated models that predict an entire sentence at a time. For autocomplete, the models only predict a single word at a time and then, regardless of whether the answer is correct or not, use the additional context of the word that the user typed next to make the next prediction.

Autocomplete models can suggest just the next word, or they can be used to suggest a list of alternative words that the user could select from (since the models are probabilistic they can return a ranked list of suggestions). To evaluate this use case, and to better understand what words the models are predicting, we also evaluated the models using accuracy@N. Accuracy@N counts a model as correct for an example as long as it suggests the correct word within the first N suggestions.

Model	No Context	Context-Aware
Single PNLMs		
RoBERTa	56.23	62.40
BERT	50.43	53.28
XLNet	45.70	46.20
GPT-2	23.2	49.00
Ensemble Models		
Majority Vote	39.75	48.25
4CC	52.27	59.32
AutoMeTS	57.89	64.52
Upper bound	60.22	66.44

Table 6: Accuracy of pretrained neural language models (PNLMs), both with and without the context of the difficult sentence on the medical parallel English Wikipedia corpus.

	RoBERTa	BERT	XLNet	GPT-2
accuracy@1	62.40	53.28	46.20	49.00
accuracy@2	67.20	54.50	46.90	49.44
accuracy@3	70.00	56.20	49.20	52.57
accuracy@4	72.10	58.00	51.30	54.32
accuracy@5	73.20	59.40	53.50	56.12

Table 7: Accuracy@N of the RoBERTa, BERT, XLNet, and GPT-2 with context on next word prediction.

5 Results and Discussion

We first analyze the results of the individual PNLMs on the medical text simplification autocompletion task and then explore the ensemble approaches. We also include a number of post-hoc analyses to better understand what the different models are doing and limitations of the models.

5.1 Individual Language Models

Accuracy Table 6 shows the results for the four different variants (RoBERTa, BERT, XLNet, and GPT-2 with and without context). Even without any context, many of the models get every other word correct (accuracy of around 50%). With the additional context of the difficult sentence, all models improve. GPT-2 improves drastically (more than doubling the accuracy) and RoBERTa also achieves a reasonable improvement of 6% absolute. Both with and without context, RoBERTa is the best performing model achieving significantly higher results than the other models

Accuracy@N Table 7 shows the accuracy@N from PNLMs on next word prediction. By allowing the autocomplete system user the option to pick from a list of options, the correct word is much more likely to be available. Even just showing three options, results in large improvements, e.g., 7.5% absolute for RoBERTa. When five options are available, increases range from 6–10.8%.

Impact of the difficult sentence length To better understand the models, we compared the average performance of the models based on the length of the sentence that was being simplified. We divided the test sentence into four different groups based on length: very short (≤ 5 tokens), short (6 – 15 tokens), medium (16 – 19 tokens), and long (≥ 20 tokens). Figure 1 shows the test accuracy of the context-aware models broken down into these 4 different groups. RoBERTa, BERT, and XLNet are fairly consistent regardless of the difficult sentence lengths; only for long sentences does the performance drop. GPT-2 performs poorly on short sentences, but well for other lengths. We hypothesize that the training data for GPT-2 (web text) may require more context for this more technical task.




Figure 2: Accuracy for the four context-aware models based on the number of words typed so far.

Impact of the number of words typed We also analyzed the average performance of the models based on how many of the words of the simplified sentence had been typed so far, i.e., the length of $s_1s_2\dots s_i$. Figure 2 shows the performance accuracy of the models based on how many words of the simplification the model has access to. Early on, when the sentence is first being constructed, all models perform poorly. As more words are typed, the accuracy of all models increases. GPT-2 performs the best early on, but then levels off after about 7 words. Both BERT and RoBERTa continue to improve as more context is added, which may partially explain their better performance overall.

5.2 Ensemble Models

Although RoBERTa performs the best overall, as Figures 1 and 2 show, different models perform better in different scenarios. The ensemble models try and leverage these differences to generate a better overall model.

As shown in Table 6, the majority vote ensemble model does not perform better than the best individual PNLN. The 4CC does outperform the majority vote approach by 11.07% and does perform better than three of the four PNLNs, but it still fails to beat RoBERTa. AutoMeTS, by viewing the problem as a multi-label problem, is able to avoid some of the biases in the training data that 4CC has, resulting in an absolute improvement of 2.1% over the best individual model (RoBERTa).

To understand the performance differences of the ensemble models, Table 8 shows the percentage that each of the four PNLNs was used for each of the ensemble approaches. The problem with the majority vote is that it tends to utilize all of the systems, regardless of their quality. For example, it shows a high percentage of XLNet, even though its performance was the worst. Because RoBERTa is the best performing model, the 4CC approach had a very strong bias towards RoBERTa. The multi-label selector reduces the bias towards using RoBERTa (a 11.25% decrease in the appearance of RoBERTa) and is able to leverage predictions from the other model when appropriate.

To understand the limits of an ensemble approach, we also calculated the upper bound that the ensemble approach could achieve. Specifically, as long as at least one model among the four PNLNs correctly predicts the next word, we mark it as correct for the upper bound. This means that no other possible combination of the four PNLNs can perform better. Here this upper bound is 66.44% (Table 6), which is about a 2% improvement over our ensemble approach; there is a bit of room for improvement, but also better language modeling techniques also need to be explored.




Figure 3: Accuracy for the RoBERTa, AutoMeTS, and upper bound models based on the number of words typed so far.

	Majority Vote	4CC	AutoMeTS
RoBERTa	47.29%	71.00%	59.75%
BERT	20.25%	12.45%	18.09%
XLNet	15.41%	5.72%	7.06%
GPT-2	17.05%	10.83%	15.10%

Table 8: The appearance frequency of PNLMs in Majority Vote, 4CC, AutoMeTS ensemble models.

Figure 3 shows the accuracy for RoBERTa, AutoMeTS, and the upper bound based on size of the context, i.e., words typed so far. For small context, the ensemble approach performs much better than RoBERTa. This likely can be attributed to selecting one of the other models that performs better for small context, e.g., GPT-2. As the context size increases, however, RoBERTa and the ensemble model perform similarly. The upper bound is consistently above the ensemble approach across all context sizes.

6 Conclusions

In this paper, we introduced a new medical parallel English Wikipedia corpus for text simplification, which contains 3.3K medical sentence pairs. Further, we proposed a new autocomplete application for PNLMs for medical text simplification. Such autocomplete models can assist users in simplifying text with improved efficiency and higher quality results in domains where information preservation is especially critical, such as healthcare and medicine, and where fully-automated approaches are not appropriate. We examined four recent PNLMs: BERT, RoBERTa, XLNet, and GPT-2, and showed how the additional context of the sentences being simplified could be incorporated into the autocomplete simplification process. Further, we introduced AutoMeTS, an ensemble method that combines the advantages of each of the different PNLMs. The AutoMeTS model outperforms the best individual model, RoBERTa, by 2.1%. Longer term, we envision that this new application could lead to other interesting model adaptations and advance text simplification in medical domains.

Acknowledgements

Research reported in this paper was supported by the National Library of Medicine of the National Institutes of Health under Award Number R01LM011975. The content is solely the responsibility of the

authors and does not necessarily represent the official views of the National Institutes of Health.

References

- Olivier Bodenreider. 2004. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32(suppl_1):D267–D270.
- Fei Cai, Maarten De Rijke, et al. 2016. A survey of query auto completion in information retrieval. *Foundations and Trends® in Information Retrieval*.
- Mia Xu Chen, Benjamin N Lee, Gagan Bansal, Yuan Cao, Shuyuan Zhang, Justin Lu, Jackie Tsay, Yinan Wang, Andrew M Dai, Zhifeng Chen, et al. 2019. Gmail smart compose: Real-time assisted writing. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2287–2295.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Mark D Dunlop and Andrew Crossan. 2000. Predictive text entry methods for mobile phones. *Personal Technologies*, 4(2-3):134–143.
- Spence Green, Sida I. Wang, Jason Chuang, Jeffrey Heer, Sebastian Schuster, and Christopher D. Manning. 2014. Human effort and machine learnability in computer aided translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1225–1236, Doha, Qatar, October. Association for Computational Linguistics.
- David Kauchak. 2013. Improving text simplification language modeling using unsimplified text data. In *Proceedings of the 51st annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 1537–1546.
- Nodira Khoussainova, YongChul Kwon, Magdalena Balazinska, and Dan Suciu. 2010. Snipsuggest: Context-aware autocompletion for sql. *Proceedings of the VLDB Endowment*.
- Nicholas Kloehn, Gondy Leroy, David Kauchak, Yang Gu, Sonia Colina, Nicole P. Yuan, and Debra Revere. 2018. Improving consumer understanding of medical text: Development and validation of a new subsimplify algorithm to automatically generate term explanations in english and spanish. *Journal of Medical Internet Research (JMIR)*.
- 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*.
- 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.
- 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 on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Matthew Shardlow and Raheel Nawaz. 2019. Neural text simplification of clinical letters with a domain specific phrase table. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 380–389.
- Matthew Shardlow. 2014. A survey of automated text simplification. *International Journal of Advanced Computer Science and Applications*, 4(1):58–70.
- Luca Soldaini and Nazli Goharian. 2016. QuickUMLS: a fast, unsupervised approach for medical concept extraction. In *MedIR Workshop, SIGIR*, pages 1–4.
- Laurens Van den Bercken, Robert-Jan Sips, and Christoph Lofi. 2019. Evaluating neural text simplification in the medical domain. In *The World Wide Web Conference*, pages 3286–3292.

- 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*, pages 5998–6008.
- 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*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764.
- 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.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, pages 19–27.