Towards Inclusive Reading: A Neural Text Generation Framework for Dyslexia Accessibility

ELHAM MADJIDI and CHRISTOPHER CRICK, Oklahoma State University, USA

This paper presents a novel approach to generating dyslexia-friendly text using neural text generation techniques. We propose a framework that leverages transformer-based language models, specifically GPT and T5, and incorporates syllable and morphological analysis to enhance the readability and comprehension of text for dyslexic readers. Our approach involves fine-tuning the language models on a curated dataset of dyslexia-friendly text, validated through human assessments and feedback from individuals with dyslexia.

We conduct a two-phase experiment with 14 undergraduate students with dyslexia to evaluate the effectiveness of our generated text. The results demonstrate improvements in reading time for participants presented with the refined dyslexia-friendly passages, while also highlighting the importance of individual preferences and text engagement. Furthermore, we provide insights into the specific challenges faced by dyslexic readers and propose targeted approaches to address these issues.

This research contributes to the advancement of text accessibility by automating the process of converting standard text into dyslexia-friendly formats. The insights gained from this study inform the design of dyslexia-friendly materials and emphasize the importance of a holistic approach to text accessibility. Our framework has the potential to increase the availability of dyslexia-friendly content and support individuals with dyslexia in accessing written information.

Additional Key Words and Phrases: dyslexia, text accessibility, neural text generation, language models, text simplification

ACM Reference Format:
Elham Madjidi and Christopher Crick. 2024. Towards Inclusive Reading: A Neural Text Generation Framework for Dyslexia Accessibility. In *11th International Conference on Software Development and Technologies for Enhancing Accessibility and Fighting Info-exclusion (DSAI 2024), November 13–15, 2024, Abu Dhabi, UAE.* ACM, New York, NY, USA, 8 pages. https://doi.org/.....

1 INTRODUCTION

Dyslexia is a neurodevelopmental disorder characterized by difficulties in reading, writing, and spelling, despite having normal intelligence and adequate educational opportunities [11]. It is estimated that around 10% of the population worldwide is affected by dyslexia, making it one of the most common learning disabilities [26]. Individuals with dyslexia often struggle with accessing written information, leading to challenges in academic, professional, and personal spheres [16].

The advent of digital technologies has opened up new possibilities for supporting individuals with dyslexia, with various assistive tools and technologies being developed to enhance reading experiences [21]. However, the vast majority of written content available online and in print is not designed with dyslexia in mind, creating significant barriers for dyslexic readers [13]. This highlights the need for effective and efficient methods to convert standard text into dyslexia-friendly formats.

Existing approaches to creating dyslexia-friendly text have primarily focused on manual simplification and formatting techniques, such as using specific fonts, colors, and layouts [22]. While these methods have shown promise in improving readability for dyslexic readers, they are time-consuming, labor-intensive, and often rely on human expertise. Automating the process of generating dyslexia-friendly text has the potential to greatly increase the availability and accessibility of written content for individuals with dyslexia.

Recent advancements in natural language processing (NLP) and machine learning have paved the way for the development of sophisticated text generation models. Transformer-based language models, such as GPT [18] and T5 [19], have demonstrated remarkable capabilities in generating coherent and contextually relevant text. These models learn from vast amounts of diverse text data and can be fine-tuned for specific tasks, such as text simplification and style transfer [30].

In this paper, we propose a novel approach to generating dyslexia-friendly text using neural text generation techniques. Our approach leverages the power of transformer-based language models, specifically GPT and T5, and incorporates syllable and morphological analysis to create more readable and comprehensible text for dyslexic readers. By fine-tuning these models on a carefully curated dataset of dyslexia-friendly text, we aim to automate the process of converting standard text into dyslexia-friendly versions. The main contributions of our work are as follows:

- We present a novel framework for generating dyslexia-friendly text using state-of-the-art neural text generation techniques, combining transformer-based language models with linguistic analysis.
- We conduct a two-phase experiment with undergraduate students with dyslexia to assess the effectiveness of our generated dyslexia-friendly text in terms of reading time, readability, and comprehension.
- We provide valuable insights into the specific challenges and preferences of dyslexic readers, informing the design of dyslexia-friendly materials and highlighting the importance of a holistic approach to text accessibility.

The remainder of this paper is organized as follows: Section 2 provides an overview of related work on dyslexia, text accessibility, and neural text generation. Section 3 describes our proposed methodology, including the dataset, language models, and linguistic analysis techniques used. Section 4 presents the experimental setup, results, and analysis of our study. Finally, Section 5 concludes the paper and discusses future research directions.

DSAI 2024 , November 13–15, 2024, Abu Dhabi, UAE

ACM ISBN .....
https://doi.org/.....

2 RELATED WORK

Dyslexia is a neurodevelopmental disorder characterized by difficulties in reading, spelling, and language processing, despite normal intelligence and adequate educational opportunities [25]. Individuals with dyslexia face significant challenges in accessing written information, which can hinder their academic success and social inclusion [7]. This section reviews current research and practices aimed at enhancing text accessibility for individuals with dyslexia, focusing on text formatting, writing style, content simplification, and the potential of neural text generation techniques.

Text Formatting

Text formatting plays a crucial role in improving readability for individuals with dyslexia. Research has shown that certain font styles, such as OpenDyslexic, Dyslexie, Lexia Readable, and Sylexiad e Read Regular, can enhance readability [2]. However, mainstream fonts like Helvetica, Courier, Arial, Verdana, and Computer Modern Unicode have also proven effective in facilitating reading performance [20]. Other formatting considerations include using a minimum 12pt or 14pt font size, dark-colored text, lowercase letters, and avoiding excessive capitalization [2]. Highlighting important text through bold or colored boxes, maintaining a left-aligned, non-hyphenated layout, and providing sufficient line spacing further contribute to creating dyslexia-friendly materials [2].

Writing Style and Content Simplification

The way information is written and the simplification of text are critical factors in facilitating comprehension for individuals with dyslexia [6, 14, 24, 27]. Best practices include keeping sentences and paragraphs short, simple, and concise, with clear instructions and minimal lengthy explanations [6, 14, 24, 27]. Breaking text into shorter, easily digestible units, employing bullet points and numbers, and avoiding lengthy prose passages enhance accessibility [6, 14, 24, 27]. Addressing word complexity by replacing complex words with simpler synonyms [22] and condensing lengthy paragraphs into concise sentences are effective text simplification strategies. Tools like SeeWord [32] can help differentiate words with similar spellings, which can be challenging for individuals with dyslexia.

Neural Text Generation

Recent advancements in deep learning and the availability of large textual corpora have led to the development of neural text generation techniques. These techniques learn representations of human-written texts using deep neural network models, such as Recurrent Neural Networks (RNNs) [23], Long Short-Term Memory (LSTM) [8], and Transformers [29]. Encoder-decoder architectures, like sequence-to-sequence learning [28] and attention mechanisms [1], have shown promise in various text generation tasks. Models like BERT [9], XLNet [31], RoBERTa [10], and GPT [5, 17, 18] have achieved remarkable performance in language understanding and generation. ChatGPT, a dialogue-based model built on GPT-3.5 [15], demonstrates the potential for human-like interaction.

While these models excel in generating fluent and contextually relevant text, their application in enhancing text accessibility for individuals with dyslexia remains an open research question. Exploring the potential of neural text generation techniques to simplify and adapt text to the specific needs of dyslexic readers could lead to novel solutions in this domain.

3 MODEL

3.1 Dyslexia-Friendly Dataset and Transformer-based Language Models

The foundation of our approach to generating dyslexia-friendly text lies in the combination of a specialized dataset and state-of-the-art transformer-based language models.

We developed a unique dataset of dyslexia-friendly text [12] which forms the basis for fine-tuning our models. This dataset was created through a rigorous process involving Amazon Mechanical Turk workers who rewrote standard texts following dyslexia-friendly guidelines, followed by evaluation and validation by individuals with dyslexia.

Our methodology leverages this dataset to train transformer-based models such as such as BERT [9], RoBERTa [10], GPT [17], and T5 [19]. These models have revolutionized natural language processing with their ability to capture and generate human-like text.

By fine-tuning these pre-trained models on our specialized dataset, we enable them to learn the specific modifications and adaptations required to generate text that is more accessible to individuals with dyslexia.

Transformer-based models employ an encoder-decoder architecture, where the encoder processes the input text and the decoder generates the output text. This architecture enables them to perform a wide range of NLP tasks, including text classification, question answering, summarization, and machine translation. One of the key strengths of these models is their unsupervised pre-training on large corpora of unlabeled text data, allowing them to acquire a deep understanding of language structure, semantics, and context. Among the transformer-based models, we specifically utilize GPT and T5 for our task. The Text-to-Text Transfer Transformer (T5), developed by Google AI Language, stands out for its versatility and performance. T5 has been pre-trained on a massive corpus of diverse text-to-text tasks, enabling it to generate coherent and contextually relevant text. GPT, on the other hand, excels in generating human-like text across a wide range of topics and styles.

By fine-tuning GPT and T5 on our dyslexia-friendly text dataset, we harness their generative capabilities to automatically produce text that adheres to dyslexia-friendly guidelines. This approach mirrors the manual process undertaken by the AMT workers who created our dataset, leveraging the models' pre-existing language knowledge and adapting it to the specific needs of dyslexic readers.

The fine-tuning process involves training the models on paired examples of original and dyslexia-friendly text from our dataset. This allows the models to learn the patterns, structures, and characteristics that make text more readable for individuals with dyslexia. The result is a system capable of automatically transforming standard text into a more accessible format, potentially improving reading experiences for those with dyslexia.

3.2 Syllable and Morphological Analysis

In addition to utilizing pre-trained language models like GPT and T5, we incorporated syllable and morphological analysis into our model to identify and simplify complex word structures that may be challenging for individuals with dyslexia. Syllable and morphological analysis involves examining the internal structure of words to determine their complexity and applying simplification strategies to make them more readable and manageable [22].

Syllable analysis focuses on breaking down words into their constituent syllables, which are units of pronunciation that typically consist of a vowel sound with or without surrounding consonants [3]. By identifying the syllable structure of words, we can assess their complexity and potential difficulty for individuals with dyslexia. Words with a higher number of syllables or complex syllable structures may be more challenging to read and process [22].

Morphological analysis, on the other hand, examines the internal structure of words and identifies their constituent morphemes, which are the smallest meaningful units of language, such as roots, prefixes, and suffixes [4]. By analyzing the morphological structure of words, we can identify complex word formations that may pose difficulties for individuals with dyslexia, such as words with multiple affixes or compound words formed by combining multiple roots [4]

One possible approach to implement syllable and morphological analysis is to utilize the NLTK library for tokenizing and analyzing the words in the text. By leveraging NLTK's functionalities, we can identify complex words based on predefined criteria. For syllable analysis, we can set a threshold, such as considering words with more than three syllables as complex. Similarly, for morphological analysis, we can establish criteria to determine the complexity of words based on their morphological structure, such as the presence of multiple affixes or compound formations.

Here's an example of how we applied syllable and morphological analysis to a sample text:

> Original text: "The musician's virtuosic performance left the audience in awe."

Syllable analysis:

- mu-si-cian's
- vir-tu-o-sic
- per-for-mance
- au-di-ence

Morphological analysis:

- musician: root word
- virtuosic: root "virtuo-" + suffix "-sic"
- performance: root "perform" + suffix "-ance"
- audience: root word

Based on the analysis, we identified "virtuosic" as a complex word due to its multiple syllables and the presence of the suffix "-sic". To simplify this word, we referred to a predefined synonym dictionary that maps complex words to simpler alternatives. In this case, "virtuosic" was replaced with "skilled", resulting in the following simplified version:

> Simplified text: "The musician's skilled performance left the audience in awe."

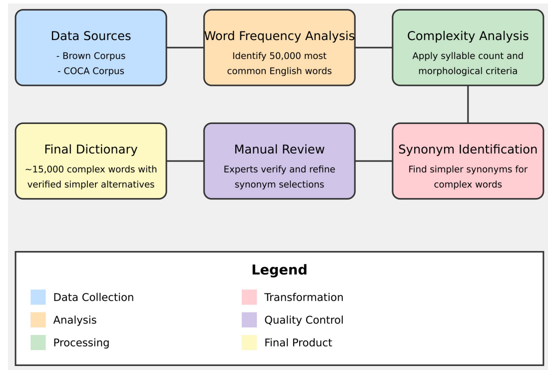

Fig. 1. Syllable and Morphological Analysis–dictionary creation process

The dictionary we used to swap out complex words with simpler counterparts was developed using the Brown Corpus and the Corpus of Contemporary American English (COCA) to get word frequency information. We identified the 50,000 most common English words and manually found synonyms for these words, focusing particularly on those identified as complex by our criteria. This manual process involved careful review to ensure the quality and appropriateness of the synonyms, resulting in a final dictionary of approximately 15,000 complex words with verified simpler alternatives Figure 1.

While there are other ways to tackle complex word substitution, such as using machine learning and deep learning techniques, we haven't explored these in the current study. These approaches, particularly those utilizing contextual word embeddings or transformer-based models, could potentially enhance this work in the future by providing more context-sensitive simplifications.

In our study, we applied this simplification as a post-processing step after the GPT model generation. This decision was made to ensure that the carefully curated simple words from our dictionary were not inadvertently replaced or modified by the GPT model. However, we acknowledge that this approach may limit the model's ability to maintain coherence in the simplified text.

3.3 Addressing Specific Challenges for Individuals with Dyslexia

Our interviews with undergraduate students with dyslexia revealed specific word-related challenges beyond general complexity. These challenges can be categorized into two main types: **Visual Confusion**, Words with confusing letter combinations, such as 'b' and 'd' appearing together (e.g., "abdominal") or contains 'p' and 's' sounds that might be difficult to distinguish (e.g., "Hypothesis"), and **Phonological Complexity**, Words with multiple occurrences of similar-sounding letters (e.g., "accessibility" with multiple 'c' and 's' sounds)

To address these challenges, we developed a targeted word replacement approach. We chose this strategy because it directly addresses the source of difficulty while maintaining the overall structure and meaning of the text. Our process involved several interconnected steps, each building on the insights gained from our interviews with dyslexic students.

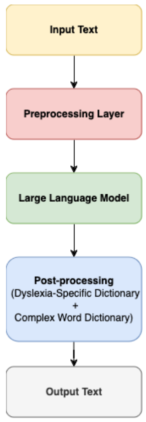

Fig. 2. Architecture of the Dyslexia-Friendly Text Simplification System. The system processes input text through multiple stages: preprocessing, a large language model, and post-processing using specialized dictionaries, to produce simplified output text for individuals with dyslexia

We began by compiling a comprehensive list of letter combinations and patterns known to be challenging for individuals with dyslexia, based on the feedback and insights gathered from our interviews. This list served as the foundation for creating a set of rules that define the criteria for identifying words containing these problematic patterns. To implement these rules effectively, we utilized regular expressions and developed custom algorithms capable of analyzing word structure and letter combinations in detail.

With these rules in place, we then turned our attention to our existing database of 50,000 common English words, which we had previously compiled. We systematically applied our rules to this database, identifying words that exhibited the problematic features we had defined. For each of these identified words, we undertook the task of finding suitable synonyms. This process was not simply about finding any synonym, but rather about carefully selecting alternatives that avoided the same problematic features while maintaining the original meaning as closely as possible.

The outcome of this process was a specialized dictionary designed to address the specific challenges faced by individuals with dyslexia. Similar to our complex word dictionary, we incorporated this specialized dictionary into our text simplification system as a post-processing step.

It's important to note that this method has limitations. The context-insensitive nature of dictionary-based replacement can sometimes lead to inappropriate simplifications, and in some cases, nuances of meaning may be lost. Furthermore, the effectiveness of this approach can vary depending on the subject matter and intended audience of the text.

To summarize our methodology, Figure 2 presents the overall architecture of our dyslexia-friendly text simplification system. This diagram illustrates how the various components we've discussed - from the initial input text through the preprocessing, large language model, and post-processing stages - work together to produce the final simplified output.

4 EXPERIMENT

The experiment was conducted in two phases to assess the effectiveness of dyslexia-friendly text modifications on reading comprehension and performance among undergraduate students with dyslexia. The first phase involved seven participants, and based on their responses, the materials were refined for the second phase, which included an additional seven participants. A total of 14 native English-speaking undergraduate students with dyslexia were interviewed for this study.

Both groups were presented with three pairs of passages sourced from the GRE exam and the textbook "Understanding Psychology." Each pair consisted of an original passage and its dyslexia-friendly counterpart, with the order of presentation randomized to minimize bias. Participants were asked to read the passages aloud, and their voice was recorded to analyze the number of mistakes made and the time taken to read each passage.

After reading each passage, participants were asked three comprehension questions to gauge their understanding of the material. The Informal Reading Inventory (IRI) was used to evaluate each participant's performance, as illustrated in Figure 3. This inventory captured reading mistakes, reading level, pace, time taken to complete the passage, and responses to comprehension questions. Following the completion of both passages in a pair, participants

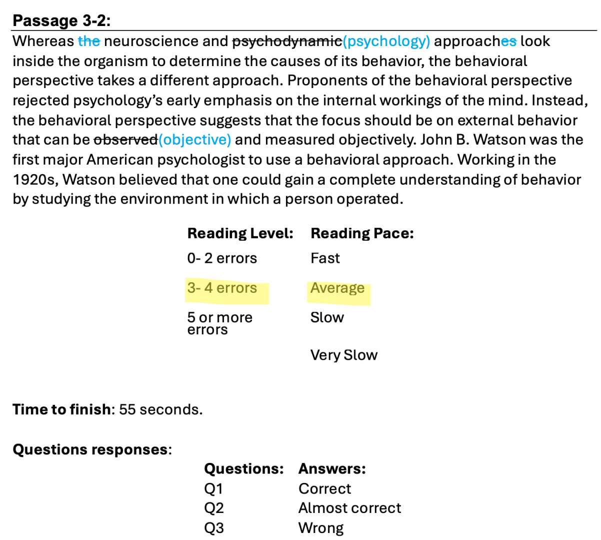

Fig. 3. Example of the Informal Reading Inventory (IRI) for one passage after the interview with a participant. Blue parts indicate the participant's mistakes while reading, and strikeouts represent the correct text. Reading level and pace are highlighted according to the participant's performance. The time taken to complete the passage is recorded in seconds, and responses to each question are marked as correct, wrong, or partially correct.

were asked to rate the readability and comprehension difficulty of each passage on a scale from 1 (very easy) to 5 (very difficult), allowing for non-integer values. Additionally, participants provided their overall impressions of the two passages.

The results from the first phase indicated room for improvement in the dyslexia-friendly modifications, as the impact on comprehension and reading speed was not as substantial as anticipated. For the second phase, the dyslexia-friendly passages were further refined by incorporating shorter sentences, simpler vocabulary, and targeted replacements for words that proved challenging for most participants in the first phase. These enhancements were achieved through fine-tuning the language model on summarized passages and employing syllable and morphological analysis to identify and replace complex words.

Figures 4, 5, and 6 present the data collected during the interviews, comparing the original passages (1-2, 2-2, 3-2) with their dyslexia-friendly counterparts (1-1, 2-1, 3-1). Figure 4 displays the reading time in seconds for each participant, while Figure 5 shows the participants' readability ratings for each passage. Figure 6 illustrates the comprehension scores based on the number of correct answers to the three questions asked about each passage, with a score of 3 indicating all questions answered correctly and a score of 0 indicating no correct answers or skipped questions due to lack of understanding.

The primary difference between the two groups was observed in the reading time, with participants in the second group generally requiring less time to read the modified passages. However, the readability ratings and comprehension scores did not exhibit a straightforward improvement. Some participants found the original text more engaging and appreciated the craftsmanship in explaining the subject matter, while others preferred the concise nature of the dyslexia-friendly versions. Overall, comprehension showed a slight improvement based on the number of correct answers, but individual preferences varied regarding the engagement level of the modified text.

Finally, all participants were asked about the visual differences between the interviewer's version and the interviewee's version of the passages. The interviewer's version was presented on white paper, justified, with single spacing and a 14-point sans-serif font. The interviewee's version featured a light orange/warm yellow background, left-aligned text, 1.5 line spacing, and a similar font size and style. Twelve participants reported that the yellow background enhanced readability, with some mentioning a preference for warmer color backgrounds when selecting books and stationery. Regarding text justification and spacing, all participants unanimously agreed that the 1.5 line spacing improved readability, and nine participants found left-aligned text easier on their eyes, reducing the likelihood of mixing up lines while reading.

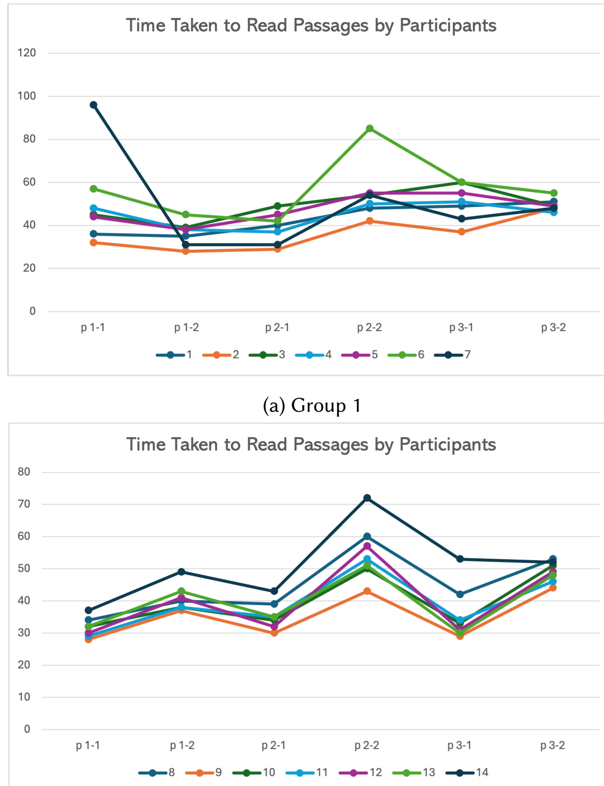

(a) Group 1

(b) Group 2

Fig. 4. Reading time for original and dyslexia-friendly passages. The line charts display the reading time in seconds for each participant across the six passages. The dyslexia-friendly passages (1-1, 2-1, 3-1) are compared to their original counterparts (1-2, 2-2, 3-2), showing the differences in reading time between the two versions for each participant.

5 CONCLUSION

In this paper, we presented a novel approach to generating dyslexia-friendly text using neural text generation techniques. By leveraging the power of transformer-based language models, specifically GPT and T5, and incorporating syllable and morphological analysis, we developed a model capable of automatically converting standard text into dyslexia-friendly versions. Our approach aimed to enhance the readability and comprehension of written material for individuals with dyslexia.

The study utilized a carefully curated dataset of dyslexia-friendly texts, validated through human assessments and feedback from college students with dyslexia. This dataset served as a valuable resource for fine-tuning the language models and ensuring the quality and relevance of the generated text.

Through a two-phase experiment involving 14 undergraduate students with dyslexia, we evaluated the effectiveness of our dyslexia-friendly text modifications. The results demonstrated improvements in reading time, with participants in the second group, who were presented with refined dyslexia-friendly passages, generally requiring less time to read the modified texts. While readability ratings and comprehension scores did not exhibit a consistent improvement, the study highlighted the importance of individual preferences and the role of text engagement in the reading experience.

The insights gained from the interviews with participants also shed light on the specific challenges faced by dyslexic readers, such as confusing letter combinations and similar-sounding letters. By incorporating targeted approaches to address these challenges, our model demonstrated a commitment to providing comprehensive solutions that cater to the unique needs of individuals with dyslexia.

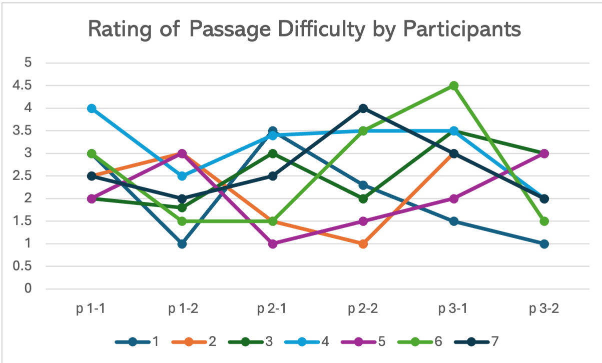

(a) Group 1

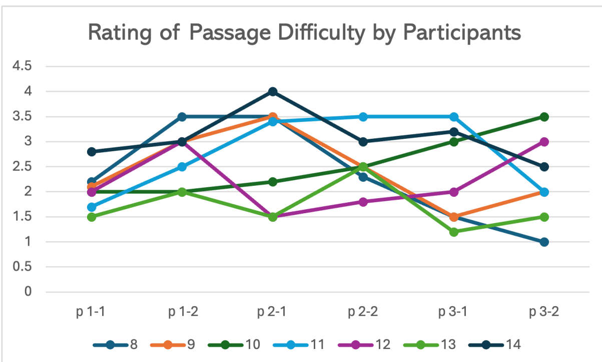

(b) Group 2

Fig. 5. Readability ratings for original and dyslexia-friendly passages. The line charts present the readability ratings (on a scale of 1-5) provided by each participant for the six passages. A rating of 1 indicates that the passage was perceived as very easy to read, while a rating of 5 suggests that the passage was very difficult to read. The charts allow for a comparison of the readability ratings between the dyslexia-friendly passages (1-1, 2-1, 3-1) and their original counterparts (1-2, 2-2, 3-2) for each participant.

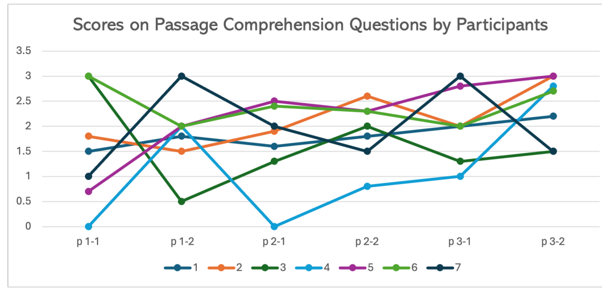

(a) Group 1

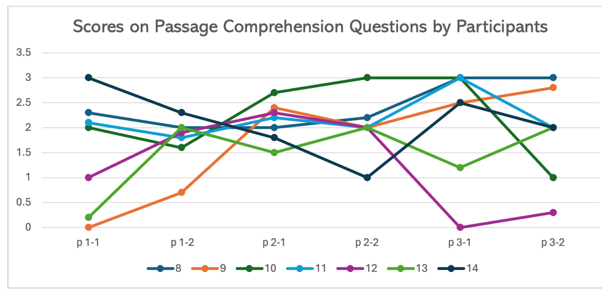

(b) Group 2

Fig. 6. Comprehension scores for original and dyslexia-friendly passages. The line charts illustrate the comprehension scores achieved by each participant in both groups for the six passages. The comprehension scores range from 0 to 3, with 0 indicating no correct answers or skipped questions due to lack of understanding, and 3 indicating all three questions were answered correctly. The charts enable a comparison of the comprehension scores between the dyslexia-friendly passages (1-1, 2-1, 3-1) and their original counterparts (1-2, 2-2, 3-2) for each participant.

Moreover, the study revealed the significance of visual factors in enhancing readability for dyslexic readers. The preferences for warmer background colors, left-aligned text, and increased line spacing underscore the importance of considering visual presentation alongside text simplification techniques.

The contributions of this work are twofold. Firstly, we introduced a novel approach to generating dyslexia-friendly text using state-of-the-art neural text generation techniques. Our model showcased the potential of leveraging advanced language models and incorporating linguistic analysis to create more accessible written content. Secondly, we provided valuable insights into the specific needs and preferences of individuals with dyslexia, informing the design of dyslexia-friendly materials and highlighting the importance of a holistic approach to text accessibility.

However, it is important to acknowledge the limitations of our study. The sample size of 14 participants, while providing valuable qualitative insights, may not be representative of the entire population of individuals with dyslexia. Future research could benefit from larger-scale studies with more diverse participant groups to further validate and refine the findings.

Additionally, while our model demonstrated promising results in generating dyslexia-friendly text, there is room for further improvement and exploration. Future work could investigate the integration of more advanced linguistic features, such as semantic analysis and discourse structure, to enhance the coherence and readability of the generated text. Incorporating user feedback and personalization options could also help tailor the text simplification process to individual needs and preferences.

In conclusion, our research presents a significant step forward in the field of generating dyslexia-friendly text using neural text generation techniques. By combining state-of-the-art language models with linguistic analysis and considering the specific needs of dyslexic readers, we have developed a model that has the potential to greatly improve text accessibility. The insights gained from this study not only contribute to the advancement of natural language processing techniques but also have practical implications for the design of educational materials, digital content, and assistive technologies for individuals with dyslexia.

As we move forward, it is crucial to continue exploring innovative approaches to enhance text accessibility and empower individuals with dyslexia. By bridging the gap between cutting-edge technology and the unique needs of dyslexic readers, we can create a more inclusive and accessible world of written information. Our work serves as a foundation for future research and development in this area, paving the way for more effective and personalized solutions that promote equal access to knowledge and opportunities for all individuals, regardless of their reading abilities.

REFERENCES

[1] Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*.

[2] M. Ball and W. McCormack. 2013. *Dyslexia at Second Level: Factsheets for Teachers*. http://www.dyslexia.ie/information/computers-and-technology/making-information-accessible-dyslexia-friendly-style-guide/ [Accessed 24th March 2015].

[3] Susan Bartlett, Grzegorz Kondrak, and Colin Cherry. 2009. On the syllabification of phonemes. In *Proceedings of human language technologies: The 2009 annual conference of the north american chapter of the association for computational linguistics*. 308–316.

[4] Laurie Bauer. 2003. *Introducing linguistic morphology*. Edinburgh university press.

[5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.

[6] Vagner Figueredo de Santana, Rosimeire de Oliveira, Leonelo Dell Anhol Almeida, and Marcia Ito. 2013. Firefixia: An Accessibility Web Browser Customization Toolbar for People with Dyslexia. In *Proceedings of the 10th International Cross-Disciplinary Conference on Web Accessibility* (Rio de Janeiro, Brazil) *(W4A '13)*. Association for Computing Machinery, New York, NY, USA, Article 16, 4 pages. https://doi.org/10.1145/2461121.2461137

[7] Jack M Fletcher, G Reid Lyon, Lynn S Fuchs, and Marcia A Barnes. 2018. *Learning disabilities: From identification to intervention*. Guilford Publications.

[8] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.

[9] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT*. 4171–4186.

[10] 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* 32 (2019).

[11] G Reid Lyon, Sally E Shaywitz, and Bennett A Shaywitz. 2003. A definition of dyslexia. *Annals of dyslexia* 53 (2003), 1–14.

[12] Elham Madjidi and Christopher Crick. 2023. Enhancing Textual Accessibility for Readers with Dyslexia through Transfer Learning. In *Proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility*. 1–5.

[13] Jacob E McCarthy and Sarah J Swierenga. 2010. What we know about dyslexia and web accessibility: a research review. *Universal Access in the Information Society* 9 (2010), 147–152.

[14] Beth A O'Brien, J Stephen Mansfield, and Gordon E Legge. 2005. The effect of print size on reading speed in dyslexia. *Journal of Research in Reading* 28, 3 (2005), 332–349.

[15] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155* 32 (2022).

[16] Marco Pino and Luigina Mortari. 2014. The inclusion of students with dyslexia in higher education: A systematic review using narrative synthesis. *Dyslexia* 20, 4 (2014), 346–369.

[17] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. 33 (2018).

[18] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners.

[19] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research* 21, 140 (2020), 1–67. http://jmlr.org/papers/v21/20-074.html

[20] Luz Rello and Ricardo Baeza-Yates. 2013. Good fonts for dyslexia. In *Proceedings of the 15th international ACM SIGACCESS conference on computers and accessibility*. 1–8.

[21] Luz Rello and Ricardo Baeza-Yates. 2016. The Effect of Font Type on Screen Readability by People with Dyslexia. *ACM Trans. Access. Comput.* 8, 4, Article 15 (may 2016), 33 pages. https://doi.org/10.1145/2897736

[22] Luz Rello, Ricardo Baeza-Yates, Stefan Bott, and Horacio Saggion. 2013. Simplify or help? Text simplification strategies for people with dyslexia. In *Proceedings of the 10th International Cross-Disciplinary Conference on Web Accessibility*. 1–10.

[23] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. 1986. Learning representations by back-propagating errors. *nature* 323, 6088 (1986), 533–536.

[24] Francisca Serrano and Sylvia Defior. 2008. Dyslexia speed problems in a transparent orthography. *Annals of Dyslexia* 58, 1 (2008), 81 – 95. https://search.ebscohost.com/login.aspx?direct=true&db=aph&AN=35395316&site=ehost-live

[25] Sally E Shaywitz, Bennett A Shaywitz, Robert K Fulbright, Pawel Skudlarski, W Einar Mencl, R Todd Constable, Kenneth R Pugh, John M Holahan, Karen E Marchione, Jack M Fletcher, et al. 2003. Neural systems for compensation and persistence: young adult outcome of childhood reading disability. *Biological psychiatry* 54, 1 (2003), 25–33.

[26] Linda S Siegel. 2006. Perspectives on dyslexia. *Paediatrics & child health* 11, 9 (2006), 581–587.

[27] British Dyslexia Association. Dyslexia style guide. January 2012. *http://www.bdadyslexia.org.uk/* (January 2012).

[28] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. *Advances in neural information processing systems* 27 (2014).

[29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).

[30] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *7th International Conference on Learning Representations, ICLR 2019*.

[31] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems* 32 (2019).

[32] Yeliz Yesilada and Simon Harper. 2019. *Web Accessibility: A Foundation for Research*. Springer.

6 APPENDIX

6.1 Group 1 detailed tables

Table 1. Reading time (in seconds) for original and dyslexia-friendly passages

	Passage Numbers					
Participants_ids	1-1	1-2	2-1	2-2	3-1	3-2
1	36	35	40	55	49	51
2	32	28	29	42	37	48
3	45	39	49	54	60	49
4	48	38	37	50	51	46
5	44	38	45	55	55	49
6	57	45	42	85	60	55
7	96	31	31	54	43	48

Table 2. Readability ratings (1-5) for original and dyslexia-friendly passages

	Passage Numbers					
Participants_ids	1-1	1-2	2-1	2-2	3-1	3-2
1	3	1	3.5	2.3	1.5	1
2	2.5	3	1.5	1	3	2
3	2	1.8	3	2	3.5	3
4	4	2.5	3.4	3.5	3.5	2
5	2	3	1	1.5	2	3
6	3	1.5	1.5	3.5	4.5	1.5
7	2.5	2	2.5	4	3	2

Table 3. Comprehension scores (0-3) for original and dyslexia-friendly passages

	Passage Numbers					
Participants_ids	1-1	1-2	2-1	2-2	3-1	3-2
1	1.5	1.8	1.6	1.8	2	2.2
2	1.8	1.5	1.9	2.6	2	3
3	3	0.5	1.3	2	1.3	1.5
4	0	2	0	0.8	1	2.8
5	0.7	2	2.5	2.3	2.8	3
6	3	2	2.4	2.3	2	2.7
7	1	3	2	1.5	3	1.5