

Controlling Reading Ease with Gaze-Guided Text Generation

Andreas Säuberli^{1,2} Darja Jepifanova¹ Diego Frassinelli¹ Barbara Plank^{1,2}

¹MaiNLP, Center for Information and Language Processing, LMU Munich, Germany

²Munich Center for Machine Learning (MCML), Munich, Germany

andreas.saeuberli@lmu.de

Abstract

The way our eyes move while reading can tell us about the cognitive effort required to process the text. In the present study, we use this fact to generate texts with controllable reading ease. Our method employs a model that predicts human gaze patterns to steer language model outputs towards eliciting certain reading behaviors. We evaluate the approach in an eye-tracking experiment with native and non-native speakers of English. The results demonstrate that the method is effective at making the generated texts easier or harder to read, measured both in terms of reading times and perceived difficulty of the texts. A statistical analysis reveals that the changes in reading behavior are mostly due to features that affect lexical processing. Possible applications of our approach include text simplification for information accessibility and generation of personalized educational material for language learning.

1 Introduction

Controlling the readability of texts is essential for many practical applications. In second language acquisition, teachers need to provide students with material that is simple enough to be understood, but difficult enough to achieve a learning effect (Krashen, 1985; Nation and Wang Ming-Tzu, 1999). Making texts easier to read can also facilitate information access, for example, for people with communication impairments (Maaß, 2020; Karreman et al., 2007; Rello et al., 2013; Alonzo et al., 2020; Säuberli et al., 2024).

Recently, natural language processing methods and large language models (LLMs) in particular have been applied to adapt the readability of texts automatically, mostly using prompting techniques (Feng et al., 2023; Kew et al., 2023; Fang et al., 2025). While these approaches have shown some success, the generated texts are often misaligned with the desired readability levels (Imperial and

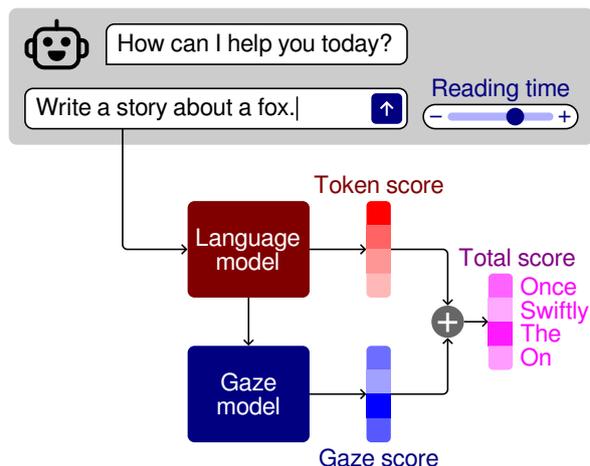


Figure 1: **Example application of gaze-guided text generation in a chatbot, and overview of the method.** Our approach allows steering language model outputs to elicit certain reading behaviors – for example, increased or decreased reading time. It integrates a gaze model predicting eye-tracking measures into the decoding process to modify the token probabilities generated by the language model.

Tayyar Madabushi, 2023; Uchida, 2025), and there is a lack of control and transparency in the generation process (Sun et al., 2023; Zhao et al., 2024). To counteract some of these limitations, we propose eye-tracking data as a more direct signal for controlling the reading ease of LLM outputs.

The way a person’s eyes move while reading a text contains rich information about both the text and the reader (Kuperman et al., 2018). For example, longer reading times on a word reflect greater effort in processing that word (Just and Carpenter, 1980; Rayner, 1998). Consequently, eye tracking has been used as a measure to assess the readability of texts (Rello et al., 2013; Singh et al., 2016; Vajjala et al., 2016; Gruteke Klein et al., 2025).

In this paper, we exploit the relationship between eye movements and reading ease from a different perspective. Instead of using gaze data to *measure* the readability of existing texts, we use it to *gen-*

erate texts with controllable reading ease. Specifically, we propose a method for guiding language models to produce outputs that will elicit certain reading behaviors, allowing us to steer the generated texts to be more or less easily readable. We achieve this by integrating a *gaze model* into the decoding process (see Figure 1). The gaze model predicts eye-tracking measures for candidate tokens generated by the language model. These predictions are then used to re-rank the token candidates. Our human evaluation study demonstrates the success of this method both in terms of the effect on readers’ eye movements and their subjective perception of overall difficulty. The contributions of this paper are:

1. We introduce a method for using eye-tracking data to guide text generation with language models.
2. We present a new publicly available eye-tracking-while-reading dataset involving texts generated using this method.
3. We analyze generated texts, eye-tracking and human rating data to demonstrate the effectiveness of the method.

2 Related work

Some controlled text generation approaches rely on fine-tuning a language model to condition the output on a class or attribute given at inference time (e.g., Keskar et al., 2019). This has several disadvantages: (1) fine-tuning large models is resource-intensive, (2) the method is inflexible because adapting the control mechanism (e.g., adding classes/attributes) requires re-training, and (3) this is often not possible with closed LLMs without access to model parameters. Therefore, several approaches have been introduced to enable *modular* control at inference-time. These approaches (including ours) involve training additional, usually smaller models to steer an off-the-shelf, unmodified language model (Dathathri et al., 2020; Krause et al., 2021; Liu et al., 2021; Arora et al., 2022; Li et al., 2023; Liu et al., 2024). Among these, the most similar to our approach is FUDGE (Yang and Klein, 2021), which involves training a binary classifier to predict the desired attribute (e.g., topic or formality) for completion candidates and combining the predictions with the probabilities from the language model. Kew and Ebling (2022) successfully applied this approach to text simplification

using a classifier trained to discriminate between simple and complex sentences. Our approach differs in that we use a regression model to predict continuous variables such as reading times.

While we are the first – to the best of our knowledge – to adapt inference-time controlled text generation to eye-tracking data, recent research has proposed alternative methods for integrating eye-movement information into text generation models, for example, by fine-tuning language models directly on gaze data to improve human alignment (Lopez-Cardona et al., 2025a). Specifically, Kiege-land et al. (2024b) fine-tuned language models to optimize their alignment to human psychometric data. Kiege-land et al. (2024a) used eye-movements as a source of feedback in direct preference optimization for controlled sentiment generation. Finally, Lopez-Cardona et al. (2025b) used synthetic eye-tracking data to model human preferences as a reward for reinforcement learning.

Previous work on the prediction of eye movements while reading include models from computational psycholinguistics such as E-Z Reader (Reichle et al., 2003) and SWIFT (Engbert et al., 2005). In natural language processing, various neural network approaches have been applied to predict scan-paths or reading measures, including fine-tuned transformer-based language models (Hollenstein et al., 2021), recurrent neural networks (Deng et al., 2023), and diffusion models (Bolliger et al., 2023, 2025a).

3 Gaze-guided text generation

The high-level idea of our method is to manipulate the probability distribution generated by a language model to favor tokens that are likely to elicit a specific reading behavior (e.g., longer/shorter fixations, or more/fewer regressions). It comprises an ensemble of two models:

1. a **language model** that predicts probabilities for the next token, and
2. a **gaze model** that predicts an aggregated eye-tracking measure (e.g., reading time or number of regressions) for the sequence generated so far.

Crucially, both models are unidirectional, i.e., predictions are made based on preceding context only, to allow for autoregressive text generation.

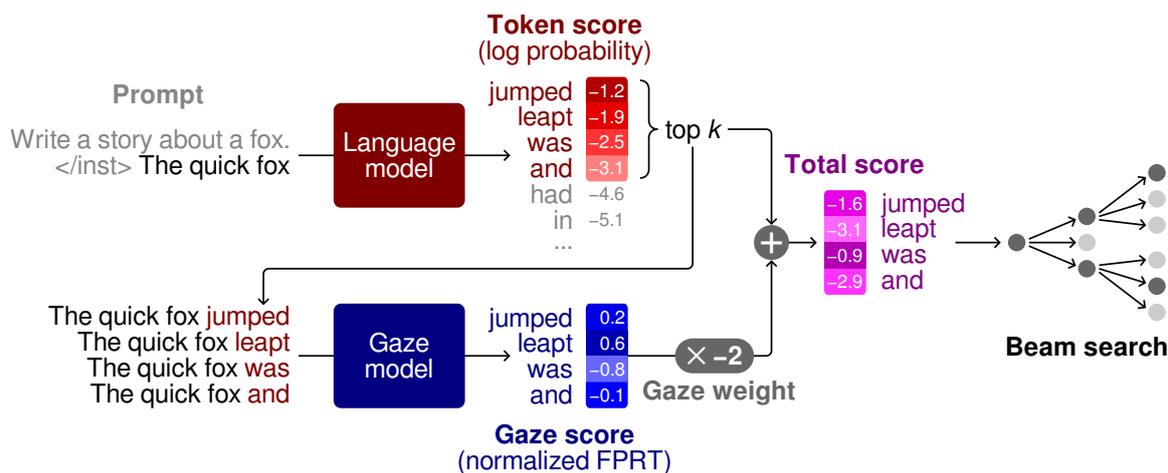


Figure 2: Gaze-guided text generation approach. (1) An off-the-shelf language model predicts a probability distribution (*token scores*) for the next token. (2) A fine-tuned gaze model predicts an eye-tracking measure (*gaze score*; in this case: FPRT, first-pass reading time) for each of the top k candidates. (3) The gaze scores are multiplied with a user-defined *gaze weight* (in this case: -2) and added to the token scores. (4) The resulting *total scores* are used for decoding with beam search.

The text generation procedure is visualized in Figure 2. First, we prompt the language model with any user-defined prompt to predict the probability distribution for the next token. We then calculate the log probability for each candidate sequence by summing up the log probabilities for each generated token – these are the **token scores**. We find the top k candidates sequences and use the gaze model to predict the eye-tracking measure for each sequence – the **gaze scores**. The final **total scores** are obtained by multiplying the gaze scores with a user-defined **gaze weight** and adding them to the corresponding token scores. We use the total scores to select the next token through beam search (or any other decoding strategy).

To exemplify this, consider a gaze model that predicts reading time as our gaze score. In this case, a positive gaze weight means that tokens that yield long reading times will be favored over tokens that yield short reading times. On the contrary, a negative gaze weight (as in Figure 2) will favor tokens with shorter reading times, thereby potentially enhancing reading ease. A gaze weight of 0 corresponds to disabling the gaze model and only using the token scores from the language model.

The modular nature of this setup means that both the language model and the gaze model can be independently replaced without retraining the other model. However, since the two models may require different tokenizers, there is a chance that partial words generated by the language model do not match the gaze model’s tokenization, leading

to unreliable gaze scores for those subwords. In our case, we alleviated this problem by applying beam search (with beam size k). This prevents the model from relying too much on gaze score predictions of subwords that might not be part of the gaze model’s vocabulary, and allows it to ‘revise’ previously generated tokens as the gaze model’s tokenization of the final word stabilizes.¹

For our experiments, we chose word-level **first-pass reading time** (FPRT; sometimes referred to as *gaze duration*) as the eye-tracking measure, which is the sum of fixation durations on a word before the gaze moves away from the word for the first time (see Section 4.1.2). Since the gaze model is required to predict a single value for each candidate sequence, we use the sum of predicted word-level FPRTs across the entire candidate sequence as our gaze score.

4 Experimental setup

The goal of our study is to measure the effect of gaze-guided text generation, reading behavior, and subjective perception by human readers. To this end, we generated texts with different gaze weights and presented these texts to L1 and L2 readers, tracking their eye movements while reading and collecting ratings from them.

¹A similar issue arises in token-level language model ensembling. See Chen et al. (2025) for a discussion of more advanced solutions in this context.

4.1 Text generation

We opted to generate short fictional stories to minimize individual differences in participants' topic familiarity and to increase motivation. All the texts were in English due to the availability of eye-tracking data for this language, which was required to train the gaze model.

4.1.1 Language model

We used the off-the-shelf instruction-tuned Llama 3.2 model with 3 billion parameters (Grattafiori et al., 2024) due to its strong performance relative to its size and speed. The user messages for the language model included a story title and prompt, and the instruction to generate a short story of around 500 words. To come up with the story prompts, we used GPT-4o (OpenAI, April 2025)² and manually selected six prompts to represent a variety of topics. All story and instruction prompts can be found in Appendix A.

4.1.2 Gaze model

As mentioned in Section 3, we used word-level FPRT as the eye-tracking measure for the gaze model. The reason for this choice is that (1) FPRT is commonly used in psycholinguistic studies to measure the effort involved in early cognitive processes (Godfroid, 2019, pp. 221–222), and (2) it can be plausibly predicted based on preceding context only, since it only includes fixations made during the first pass over the word.³

We fine-tuned a GPT-2 model with 124 million parameters (Radford et al., 2019) to predict z -score-normalized FPRT for each word. The training data is taken from the EMTeC dataset (Bolliger et al., 2025a), which contains eye movements of 107 L1 English speakers recorded while reading texts generated by language models with various decoding strategies. This makes it an ideal dataset for our use case.

Mapping word-level reading times to subword-level tokens is a problem inherent to the processing of eye-tracking data with modern language models (Beinborn and Hollenstein, 2024, p. 43). To solve this issue in our training data, we distributed each word's FPRT uniformly across all its subwords. At inference time, the sum of predicted FPRTs over all tokens represents our gaze score.

²<https://chatgpt.com/>

³In preliminary experiments, we also explored other eye-tracking measures such as go-past time, which includes regression path durations, but found that the gaze model performed substantially worse for these measures (see Appendix B).

The resulting model achieved 35.5% explained variance (R^2) on a held-out set of 20% of texts, outperforming a linear regression baseline with word length and frequency as features ($R^2 = 31.4\%$). For more details on the model and the training process, refer to Appendix B.

4.1.3 Decoding and gaze weights

Based on preliminary exploration, we used the top $k = 8$ candidates for gaze prediction and applied beam search with a beam size of 8 for decoding.

For each story prompt, we generated texts for several gaze weights between -5 (decrease FPRT) and $+5$ (increase FPRT). For our chosen language model, gaze weights beyond ± 3 tended to cause repetitive failure modes in the output (see Section 5.1 and Figure 3). Therefore, we selected the gaze weights -2 , 0 , and $+2$ for measuring the effect of gaze guidance in the eye-tracking study. Table 1 contains sample outputs for different gaze weights.

4.2 Eye-tracking study

In a new eye-tracking study, we recorded participants' eye movements while reading the texts generated with gaze weights -2 , 0 , and $+2$ (a total of 18 texts). We recruited both L1 and L2 speakers of English, since the gaze model was trained on L1 data, while one of the main applications of the approach is L2 language learning.

4.2.1 Participants

Twelve participants in the L1 group and twelve participants in the L2 group were recruited primarily through university mailing lists and courses. Informed consent was obtained before conducting the study. Participants received 10 euros or participated as part of a university course.

Participants' ages ranged between 19 and 53 years (mean: 27.2 ± 7.0). Participants were categorized as L1 when they self-reported that English was the first language they learned, or that they grew up multilingually with English and had native-level proficiency. In the L1 group, reported language varieties were American (5), British (5), Indian (1), and Singaporean English (1). In the L2 group, native languages included German (8), Chinese (2), Russian (1) and Turkish (1), while self-reported English proficiency levels ranged from C1 (9) to C2 (3) in the Common European Framework of Reference (CEFR).

Gaze weight	Output
−5	Emily had always been the girl who came in second. It was as if she had a magnet on her back that drew all the best things in life to the people in front of her. In school, she was always the one who came in second in the class, the one who got the second prize in the science fair, and the one who lost to her best friend in every game of tag. [...]
−2	For as long as she could remember, Emma had come in second place. It wasn't just in school or on the sports field – it was in every aspect of her life. Whether it was a contest, a competition, or even a relationship, she always seemed to fall just short of the top spot. [...]
0	Emily had always been the bridesmaid, never the bride. In school, she was the runner-up in every science fair, the second-place winner in every local art competition, and the silver medalist in every regional sports tournament. [...]
2	Emma had always been the perpetual bridesmaid, the perpetual runner-up, the perpetual almost-but-not-quite. It wasn't just limited to contests, either. Relationships, friendships, even everyday accomplishments – she consistently found herself finishing second. [...]
5	Emma's perpetual existence revolved around mediocrity. Every endeavor, every competition, every social interaction, she consistently ranked second-best-in-the-back-of-the-pack-to-last-but-not-quite-second-place-but-closer-than-the-last-place-[...]

Table 1: Sample outputs generated with different gaze weights. Note the differences in word length, lexical complexity, and the repetitive failure mode at *gaze weight 5*.

4.2.2 Apparatus

We recorded participants' eye movements with an *EyeLink 1000 Plus* eye tracker (SR Research, Kanata, Ontario, Canada) with a sampling rate of 1000 Hz. Only the dominant eye was tracked, which was determined in a test at the start of the session. We used a chin and forehead rest to minimize head movements and to maintain a constant eye-screen distance of 66 cm. The stimuli were presented on a 24.5-inch monitor with a resolution of 1920×1080 pixels and a refresh rate of 240 Hz. All content was presented within 28×19 cm or ~25×17 degrees of visual angle. We performed 9-point calibration at the beginning and, if necessary, between texts. Texts were typeset in *Noto Sans*⁴ in black on a light gray background, with an x-height of 4 mm or ~0.35 degrees of visual angle and double line spacing. We used *PsychoPy* (Peirce et al., 2019) to implement the stimulus presentation software.

4.2.3 Procedure

Data was collected in a single session of 30–45 minutes per participant. After reading instructions on the screen and a practice trial, participants read six texts (two for each gaze weight). Each participant saw one text for each of the six story prompts. Depending on the length, the texts were presented on 4 or 5 pages. Participants could advance to the next page by pressing a key on a keyboard, but could not return to previous pages. Each text was followed by a multiple-choice comprehension question for attention checking and three 4-point Likert-style ratings on the difficulty, naturalness/fluency, and

interestingness of the text. The texts were assigned such that each participant saw exactly one story for each prompt, and the order was randomized for each participant.

4.2.4 Data cleaning and processing

We follow common practice in performing post-hoc corrections of eye-tracking data (Godfried, 2019, pp. 257–259). We manually defined thin-plate spline transformations on the raw data to correct vertical misalignment with the text lines where necessary and possible. Trials or parts of trials with very poor quality were removed.

The participants' answers to the comprehension questions indicated that the texts were read attentively overall. All participants achieved at least 5 out of 6 correct answers (average: 94.4% accuracy overall). Therefore, we did not remove any additional trials on this basis.

We then used *pymovements* (Krakowczyk et al., 2023) to process the cleaned data. We applied the I-DT algorithm for fixation detection (Salvucci and Goldberg, 2000) with a dispersion threshold of 1.0 degrees of visual angle and a minimum fixation duration of 100 ms.

In accordance with Jakobi et al. (2024), we report more detailed data quality metrics in Appendix C and publish both the raw and the cleaned and pre-processed versions.

5 Results

We examine the effects of gaze-guided text generation from three perspectives:

⁴<https://notofonts.github.io>

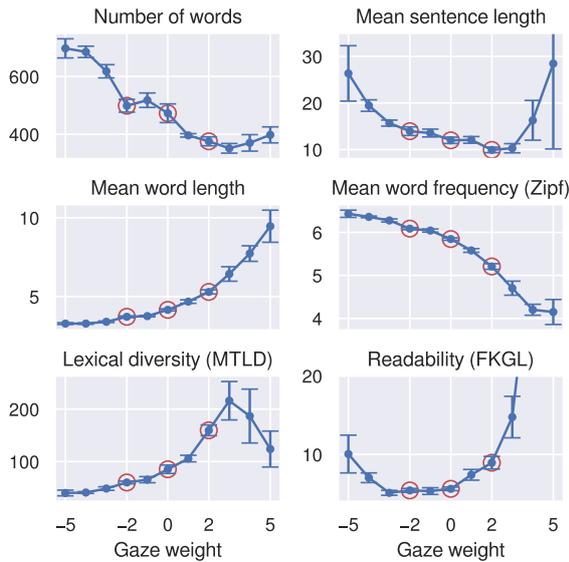


Figure 3: Text characteristics across different gaze weights. Error bars show standard error of the mean from six generated texts. Points circled in red represent the texts shown to participants in the eye-tracking study.

1. **Effects on text characteristics:** How does changing the gaze weight affect linguistic features in the generated texts?
2. **Effects on reading behavior:** How does changing the gaze weight affect first-pass reading times in L1 and L2 readers?
3. **Effects on subjective perception:** How does changing the gaze weight affect perceived difficulty, naturalness, and interestingness in L1 and L2 readers?

5.1 Effects on text characteristics

We calculated word-level, sentence-level, and document-level statistics across different gaze weights, including word frequency (Zipf score, based on the Python package *wordfreq*⁵, Speer, 2022), lexical diversity (measure of textual lexical diversity [MTLD], McCarthy, 2005; Python package *lexical-diversity*⁶), and readability (Flesch-Kincaid grade level [FKGL], Kincaid et al., 1975; Python package *textstat*⁷). Figure 3 shows several trends. Increasing the gaze weight tends to lead to texts with fewer but longer and less frequent words, shorter sentences, more lexical diversity, and a higher reading level. Gaze weights beyond

⁵<https://pypi.org/project/wordfreq/>

⁶<https://pypi.org/project/lexical-diversity/>

⁷<https://pypi.org/project/textstat/>



Figure 4: Length-normalized gaze and token scores predicted by the models, compared to first-pass reading times measured in the eye-tracking study. Token scores are token log probabilities predicted by the language model. Gaze scores are normalized first-pass reading times predicted by the gaze model. Error bars show standard error of the mean.

± 3 sometimes produced highly repetitive texts, resulting in extreme values in the case of sentence length, lexical diversity, and readability. This degeneration in output quality is somewhat expected – with larger gaze weights, the gaze model overpowers the language model’s output probabilities, potentially assigning high probability to contextually implausible words. In preliminary experiments, smaller language models were more prone to this behavior.

5.2 Effects on reading behavior

In Figure 4, we compare the gaze and token scores predicted by the models with the observed reading times in the eye-tracking experiment (recall that readers saw texts with gaze weights ± 2). As expected, gaze scores increase with positive gaze weights and decrease with negative gaze weights, while the token score is maximal when no gaze model is used. FPRT measurements roughly follow the predicted gaze scores, although the effect is stronger for L2 readers compared to L1 readers. A linear mixed-effects model with fixed effects for positive and negative gaze weights and random intercepts for readers confirms significant effects for the gaze weights in both reader groups (all $p < 0.001$).

Psycholinguistic studies have established that simple features like word length or prevalence (i.e., the proportion of native speakers who know the word) are highly predictive of processing effort (Barton et al., 2014; Brysbaert et al., 2016). To

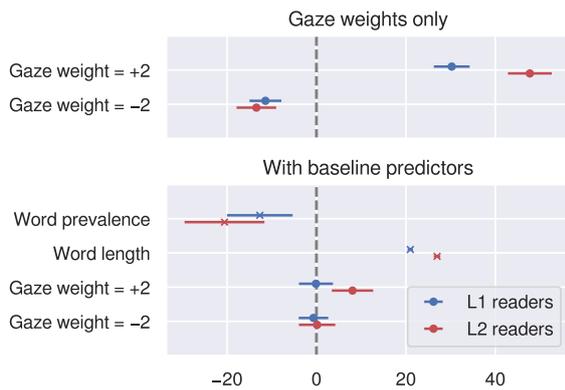


Figure 5: Coefficients and 95% confidence intervals of linear mixed-effects models fitted to observed first-pass reading times. **Upper panel:** Model with gaze weight as predictor (compared to the reference with gaze weight 0). **Lower panel:** Model with word length and prevalence as additional baseline predictors. Both models include random intercepts for readers.

determine how much these features contributed to the differences in reading time, we fitted additional models with word length (number of characters) and prevalence (according to Brysbaert et al., 2018⁸) as additional fixed effects. As shown by Figure 5, we find negative effects for word prevalence (i.e., well-known words are read more quickly) and positive effects for word length (i.e., longer words are read more slowly), as expected. However, adding these baseline predictors also mostly renders the gaze weight effects insignificant, with the exception of positive gaze weight in L2 readers. In other words, in the L1 case, differences in word length and prevalence fully explain the changes in reading time induced by the gaze guidance, whereas in the L2 case, increasing the gaze weight led to increases in reading time *beyond* length and prevalence effects. Overall, this shows that the gaze model mostly captures features related to lexical processing. More details on the mixed-effects models can be found in Appendix D.

5.3 Effects on subjective perception

Mean difficulty, naturalness, and interestingness ratings are visualized in Figure 6. In both groups, we observe a consistent trend of rising difficulty ratings with increasing gaze weights, and a substantial increase in difficulty for L2 readers. L2 readers rated texts with a positive gaze weight as less natu-

⁸Words that are not part of Brysbaert et al.’s (2018) list were excluded from the analysis. This amounted to an exclusion of 3.6% of all data points.

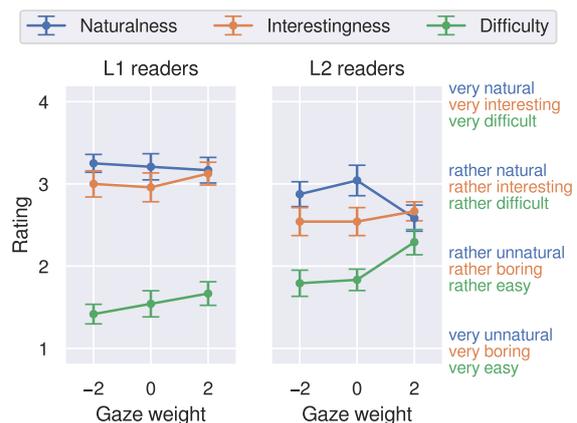


Figure 6: Mean difficulty, naturalness, and interestingness ratings across gaze weights. Error bars show standard error of the mean.

ral, likely reflecting the increased perceived difficulty as opposed to poor fluency. Interestingness remained largely unaffected in both groups.

These results show that readers perceived an increase in difficulty with larger gaze weights, and that the gaze guidance did not substantially affect text quality.

5.4 Gaze model performance

To analyze the gaze model’s performance in more depth, we consider how well its word-level predictions on the six generated texts correlate with the observed reading times in our eye-tracking study. The Pearson correlation coefficient for the L2 group is slightly higher than for the L1 group ($r = 0.585$ vs. $r = 0.552$). For comparison, the correlation on the held-out data split during training is $r = 0.596$. Figure 7 shows that the correlation the model also performs better with low-prevalence words.

6 Discussion

6.1 Effectiveness of gaze guidance

Based on the results above, we conclude that our method is successful in adapting language model outputs to different levels of reading ease without impacting text quality. The approach was less effective for reducing FPRT than for increasing it. This is easily explained by the fact that there is a lower bound, but not an upper bound on processing effort in humans – in other words, making a text more difficult to read is always possible and easier to do, while simplification is only possible to a limited degree and more challenging to achieve.

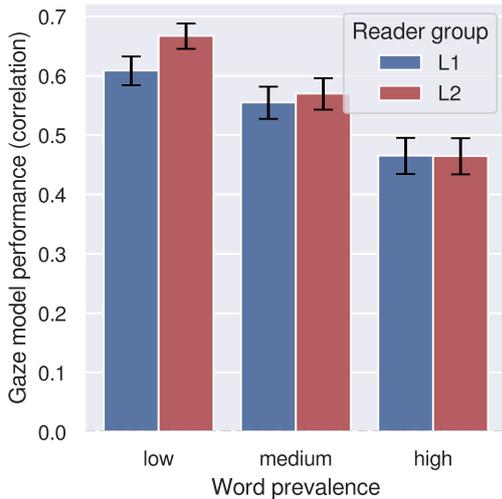


Figure 7: Pearson correlation between word-level gaze model predictions and observed first-pass reading time in the eye-tracking study. Error bars represent 95% confidence intervals. The low, medium, and high prevalence buckets are equally sized quantiles.

An important limitation revealed by our experiments is that the effects introduced by the gaze model mostly reflect low-level lexical information like word length and prevalence. This can be attributed to our choice of eye-tracking measure. FPRT and other measures based on fixation durations are known to be strongly influenced by these local, word-level features (Kliegl et al., 2006), while measures like rereading time are commonly associated with later cognitive processes at the level of syntax (Godfroid, 2019, p. 225) and may therefore be better suited as a proxy for sentence-level reading ease (Vajjala et al., 2016). However, reliably predicting these measures with a gaze model is more difficult and implausible with a unidirectional language model. Therefore, future research could improve our approach by integrating more complex generative models of eye movements (e.g. Deng et al., 2023; Bolliger et al., 2023, 2025b) and using a decoding strategy that supports more backtracking to revise previously generated tokens.

6.2 Differences between L1 and L2 readers

Overall, we see similar effects for the two reader groups, apart from expected differences like faster overall reading time and lower perceived difficulty for L1 readers.

The most notable difference is the observation that word length and prevalence effects cannot entirely explain the increase in L2 readers’ FPRT on

texts with gaze weight 2. One possible reason for this could be that the prevalence values are based on native speakers’ vocabulary knowledge, while the L2 readers in our study may have encountered a large number of unknown words in those texts, which could affect reading times beyond the linear effect of prevalence. This would also explain the sharp increase in difficulty ratings and the lower naturalness ratings for texts with gaze weight 2. Another possibility could be higher complexity at the syntactic level, although we could not find direct evidence for this.

7 Conclusion and future work

We presented a method for generating texts with an inference-time mechanism to control reading ease. The approach is modular: the language model does not require fine-tuning, and the gaze model can be easily replaced, making it both efficient and flexible. We conducted an eye-tracking experiment to assess the effects on readers’ cognitive processing effort and perceived difficulty, and release the collected dataset for further research. The results demonstrated that the approach is successful in steering the language model towards more or less easily readable texts, and that this effect is primarily due to changes in lexical complexity.

Manipulating reading ease is of use in educational scenarios for language learning, and in text simplification to enhance communication accessibility. Both of these applications involve users with diverse needs, and can thus benefit substantially from personalization (Bingel et al., 2018; Tetzlaff et al., 2020). Eye movements while reading can capture – at least to some degree – differences in how well readers understand a text, how they approach a reading task, as well as information on language knowledge and reading difficulties (Koornneef and Mulders, 2016; Ahn et al., 2020; Raatikainen et al., 2021; Skerath et al., 2023; Shubi et al., 2024, 2025). While the gaze model in the present work was trained on a general group of L1 speakers, the training data can be easily replaced to reflect a different user group, provided that sufficient eye-tracking data exists for that group. Thus, through further research, gaze-guided text generation could enable personalized generation of educational or informational material, adapted to the reading behavior of the specific reader (or reader group) at hand. Our work represents a first step in this direction.

Data availability and reproducibility

The collected dataset, including raw and processed eye-tracking data, is publicly available for research purposes through this link: <https://osf.io/rhgbk>. The dataset is also integrated into the *pymovements* Python library⁹ (Krakowczyk et al., 2023, 2025) for easy access.

All datasets and code libraries used in this project are open-source and received due citations. The code and data for reproducing the results and figures in this paper is available from this repository: <https://github.com/mainlp/gaze-guided-text-generation/>.

Limitations

English only. We only performed experiments with English texts, limiting conclusions about other languages. The main reason for this is that the dataset we used for training the gaze model only contains English, and that eye-tracking-while-reading data is generally sparse in other languages.

Model selection. We only tested one combination of language and gaze models. We restricted our experiments to this basic setup due to the resources required for collecting eye-tracking data. Testing more models would require multiple sessions for each participant, or recruiting a larger number of participants.

Downstream applicability. Our experiments only evaluated the approach with the FPRT reading measure, which mostly captures low-level lexical features. This limits the use for text simplification applications, where simplification of syntactic, semantic and discourse-level aspects is usually required. While we believe that our method can also be applied to reading measures that reflect these aspects, this has not been shown empirically in the present work.

User groups. Although we list text simplification in the context of accessibility as a potential use case, we did not include primary user groups for this application. We were careful not to make any statements about our method’s usefulness for communication accessibility, but this remains an important limitation of our work.

⁹<https://pymovements.readthedocs.io/>

Ethical considerations

Eye-tracking data. Experiments involving human participants require careful consideration. Our study participants took part on a voluntary basis, either as part of a university course or as paid participants (10 euros for a 30–45 min session). We obtained informed consent from participants and allowed them to withdraw their participation at any point. We provided detailed information to participants beforehand. We only publish fully pseudonymized data. While eye-tracking data is considered biometric and identifying individuals from it is technically possible (Jäger et al., 2020), doing so would require disproportionate effort.

Usage of generative models. We used GitHub Copilot for coding assistance, and ChatGPT to support writing the story prompts and comprehension questions. All generated content was thoroughly checked and tested. We did not use generative models for interpreting the results or writing this paper.

Acknowledgements

We would like to thank the three anonymous reviewers, whose constructive comments contributed significantly to the final paper. We are also grateful to members of MaiNLP lab, particularly Silvia Casola and Shijia Zhou, as well as Yang Janet Liu for helpful feedback on early drafts of the manuscript.

This research is in parts supported by the ERC Consolidator Grant DIALECT 101043235.

References

- Seoyoung Ahn, Conor Kelton, Aruna Balasubramanian, and Greg Zelinsky. 2020. [Towards predicting reading comprehension from gaze behavior](#). In *ACM Symposium on Eye Tracking Research and Applications*, ETRA ’20. ACM.
- Oliver Alonzo, Matthew Seita, Abraham Glasser, and Matt Huenerfauth. 2020. [Automatic text simplification tools for deaf and hard of hearing adults: Benefits of lexical simplification and providing users with autonomy](#). In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI ’20, pages 1–13. ACM.
- Kushal Arora, Kurt Shuster, Sainbayar Sukhbaatar, and Jason Weston. 2022. [DIRECTOR: Generator-classifiers for supervised language modeling](#). In *Proceedings of the 2nd Conference of the Association for Computational Linguistics and the 12th International Joint Conference*

- on *Natural Language Processing (Volume 1: Long Papers)*, pages 512–526. Association for Computational Linguistics.
- Jason J. S. Barton, Hashim M. Hanif, Laura Eklinder Björnström, and Charlotte Hills. 2014. [The word-length effect in reading: A review](#). *Cognitive Neuropsychology*, 31(5–6):378–412.
- Lisa Beinborn and Nora Hollenstein. 2024. *Cognitive Plausibility in Natural Language Processing*. Springer International Publishing.
- Joachim Bingel, Gustavo Paetzold, and Anders Søgaard. 2018. [Lexi: A tool for adaptive, personalized text simplification](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 245–258, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Lena Bolliger, David Reich, Patrick Haller, Deborah Jakobi, Paul Prasse, and Lena Jäger. 2023. [ScanDL: A diffusion model for generating synthetic scanpaths on texts](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15513–15538. Association for Computational Linguistics.
- Lena S. Bolliger, Patrick Haller, Isabelle C. R. Cretton, David R. Reich, Tannon Kew, and Lena A. Jäger. 2025a. [EMTeC: A corpus of eye movements on machine-generated texts](#). *Behavior Research Methods*, 57(7).
- Lena S. Bolliger, David R. Reich, and Lena A. Jäger. 2025b. [ScanDL 2.0: A generative model of eye movements in reading synthesizing scanpaths and fixation durations](#). *Proceedings of the ACM on Human-Computer Interaction*, 9(3):1–29.
- Marc Brysbaert, Paweł Mandera, Samantha F. McCormick, and Emmanuel Keuleers. 2018. [Word prevalence norms for 62,000 English lemmas](#). *Behavior Research Methods*, 51(2):467–479.
- Marc Brysbaert, Michaël Stevens, Paweł Mandera, and Emmanuel Keuleers. 2016. [The impact of word prevalence on lexical decision times: Evidence from the Dutch Lexicon Project 2](#). *Journal of Experimental Psychology: Human Perception and Performance*, 42(3):441–458.
- Zhijun Chen, Jingzheng Li, Pengpeng Chen, Zhuoran Li, Kai Sun, Yuankai Luo, Qianren Mao, Dingqi Yang, Hailong Sun, and Philip S. Yu. 2025. [Harnessing multiple large language models: A survey on LLM ensemble](#). *arXiv*.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. [Plug and play language models: A simple approach to controlled text generation](#). In *International Conference on Learning Representations*.
- Shuwen Deng, David R. Reich, Paul Prasse, Patrick Haller, Tobias Scheffer, and Lena A. Jäger. 2023. [Eyettention: An attention-based dual-sequence model for predicting human scanpaths during reading](#). *Proceedings of the ACM on Human-Computer Interaction*, 7(ETRA):1–24.
- Ralf Engbert, Antje Nuthmann, Eike M. Richter, and Reinhold Kliegl. 2005. [SWIFT: A dynamical model of saccade generation during reading](#). *Psychological Review*, 112(4):777–813.
- Dengzhao Fang, Jipeng Qiang, Yi Zhu, Yunhao Yuan, Wei Li, and Yan Liu. 2025. [Progressive document-level text simplification via large language models](#). *arXiv*.
- Yutao Feng, Jipeng Qiang, Yun Li, Yunhao Yuan, and Yi Zhu. 2023. [Sentence simplification via large language models](#). *arXiv*.
- Aline Godfroid. 2019. *Eye Tracking in Second Language Acquisition and Bilingualism: A Research Synthesis and Methodological Guide*. Routledge.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The Llama 3 herd of models](#). *arXiv*.
- Keren Gruteke Klein, Shachar Frenkel, Omer Shubi, and Yevgeni Berzak. 2025. [Surprisal takes it all: Eye tracking based cognitive evaluation of text readability measures](#). *arXiv*.
- Nora Hollenstein, Emmanuele Chersoni, Cassandra L. Jacobs, Yohei Oseki, Laurent Prévot, and Enrico Santus. 2021. [Cmcl 2021 shared task on eye-tracking prediction](#). In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*, pages 72–78. Association for Computational Linguistics.
- Joseph Marvin Imperial and Harish Tayyar Madabushi. 2023. [Flesch or fumble? Evaluating readability standard alignment of instruction-tuned language models](#). In *Proceedings of the Third Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 205–223, Singapore. Association for Computational Linguistics.
- Deborah N. Jakobi, Daniel G. Krakowczyk, and Lena A. Jäger. 2024. [Reporting eye-tracking data quality: Towards a new standard](#). In *Proceedings of the 2024 Symposium on Eye Tracking Research and Applications*, ETRA '24, pages 1–3. ACM.
- Marcel A. Just and Patricia A. Carpenter. 1980. [A theory of reading: From eye fixations to comprehension](#). *Psychological Review*, 87(4):329–354.

- Lena A. Jäger, Silvia Makowski, Paul Prasse, Sascha Liehr, Maximilian Seidler, and Tobias Scheffer. 2020. *Deep Eyedetection: Biometric Identification Using Micro-movements of the Eye*, pages 299–314. Springer International Publishing.
- Joyce Karreman, Thea Van Der Geest, and Esmee Bursink. 2007. Accessible website content guidelines for users with intellectual disabilities. *Journal of Applied Research in Intellectual Disabilities*, 20(6):510–518.
- Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. **CTRL: A conditional transformer language model for controllable generation**. *arXiv*.
- Tannon Kew, Alison Chi, Laura Vásquez-Rodríguez, Sweta Agrawal, Dennis Aumiller, Fernando Alva-Manchego, and Matthew Shardlow. 2023. **BLESS: Benchmarking large language models on sentence simplification**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13291–13309. Association for Computational Linguistics.
- Tannon Kew and Sarah Ebling. 2022. **Target-level sentence simplification as controlled paraphrasing**. In *Proceedings of the Workshop on Text Simplification, Accessibility, and Readability (TSAR-2022)*, pages 28–42. Association for Computational Linguistics.
- Samuel Kiegl, David Robert Reich, Ryan Cotterell, Lena Ann Jäger, and Ethan Wilcox. 2024a. **The pupil becomes the master: Eye-tracking feedback for tuning LLMs**. In *ICML 2024 Workshop on LLMs and Cognition*.
- Samuel Kiegl, Ethan Wilcox, Afra Amini, David Robert Reich, and Ryan Cotterell. 2024b. **Reverse-engineering the reader**. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9367–9389. Association for Computational Linguistics.
- J. Peter Kincaid, Robert P. Fishburne Jr, Richard L Rogers, and Brad S. Chissom. 1975. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. Technical report.
- Reinhold Kliegl, Antje Nuthmann, and Ralf Engbert. 2006. **Tracking the mind during reading: The influence of past, present, and future words on fixation durations**. *Journal of Experimental Psychology: General*, 135(1):12–35.
- Arnout Koornneef and Iris Mulders. 2016. **Can we ‘read’ the eye-movement patterns of readers? Unraveling the relationship between reading profiles and processing strategies**. *Journal of Psycholinguistic Research*, 46(1):39–56.
- Daniel G. Krakowczyk, David R. Reich, Jakob Chwastek, Deborah N. Jakobi, Paul Prasse, Assunta Süß, Oleksii Turuta, Paweł Kasprowski, and Lena A. Jäger. 2023. **pymovements: A Python package for eye movement data processing**. In *2023 Symposium on Eye Tracking Research and Applications*, ETRA ’23, pages 1–2. ACM.
- Daniel G. Krakowczyk, David R. Reich, Andreas Säuberli, Iza Škrjanec, Isabelle C. R. Cretton, Deborah N. Jakobi, Sergiu Nisioi, Paul Prasse, and Lena A. Jäger. 2025. **The more the merrier: Boost your dataset visibility and discover eye-tracking datasets with pymovements**. In *Proceedings of the 2025 Symposium on Eye Tracking Research and Applications*, ETRA ’25, pages 1–3. ACM.
- Stephen D. Krashen. 1985. *The Input Hypothesis: issues and implications*. Laredo.
- Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq Joty, Richard Socher, and Nazneen Fatema Rajani. 2021. **GeDi: Generative discriminator guided sequence generation**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*. Association for Computational Linguistics.
- Victor Kuperman, Kazunaga Matsuki, and Julie A. Van Dyke. 2018. **Contributions of reader- and text-level characteristics to eye-movement patterns during passage reading**. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(11):1687–1713.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. 2023. **Contrastive decoding: Open-ended text generation as optimization**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.
- Alisa Liu, Xiaochuang Han, Yizhong Wang, Yulia Tsvetkov, Yejin Choi, and Noah A. Smith. 2024. **Tuning language models by proxy**. In *First Conference on Language Modeling*.
- Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A. Smith, and Yejin Choi. 2021. **DExperts: Decoding-time controlled text generation with experts and anti-experts**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics.
- Angela Lopez-Cardona, Sebastian Idesis, and Ioannis Arapakis. 2025a. **Integrating cognitive processing signals into language models: A review of advances, applications and future directions**. In *2025 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Angela Lopez-Cardona, Carlos Segura, Alexandros Karatzoglou, Sergi Abadal, and Ioannis Arapakis. 2025b. **Seeing eye to AI: Human alignment via gaze-based response rewards for large language models**. In

- The Thirteenth International Conference on Learning Representations.*
- Christiane Maaß. 2020. *Easy Language – Plain Language – Easy Language Plus: Balancing Comprehensibility and Acceptability.* Frank & Timme.
- Philip M McCarthy. 2005. *An assessment of the range and usefulness of lexical diversity measures and the potential of the measure of textual, lexical diversity (MTLD).* Ph.D. thesis, The University of Memphis.
- Paul Nation and Karen Wang Ming-Tzu. 1999. Graded readers and vocabulary. *Reading in a Foreign Language*, 12(2):255–380.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. *Scikit-learn: Machine learning in Python.* *Journal of Machine Learning Research*, 12(85):2825–2830.
- Jonathan Peirce, Jeremy R. Gray, Sol Simpson, Michael MacAskill, Richard Höchenberger, Hiroyuki Sogo, Erik Kastman, and Jonas Kristoffer Lindeløv. 2019. *PsychoPy2: Experiments in behavior made easy.* *Behavior Research Methods*, 51(1):195–203.
- Peter Raatikainen, Jarkko Hautala, Otto Loberg, Tommi Kärkkäinen, Paavo Leppänen, and Paavo Nieminen. 2021. *Detection of developmental dyslexia with machine learning using eye movement data.* *Array*, 12:100087.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI*.
- Keith Rayner. 1998. Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, 124(3):372–422.
- Erik D. Reichle, Keith Rayner, and Alexander Pollatsek. 2003. *The E-Z Reader model of eye-movement control in reading: Comparisons to other models.* *Behavioral and Brain Sciences*, 26(4):445–476.
- 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, W4A '13*, pages 1–10. ACM.
- Dario D. Salvucci and Joseph H. Goldberg. 2000. *Identifying fixations and saccades in eye-tracking protocols.* In *Proceedings of the symposium on Eye tracking research & applications - ETRA '00*, ETRA '00, pages 71–78. ACM Press.
- Skipper Seabold and Josef Perktold. 2010. *Statsmodels: Econometric and statistical modeling with Python.* In *Proceedings of the 9th Python in Science Conference, SciPy*, pages 92–96. SciPy.
- Omer Shubi, Cfir Avraham Hadar, and Yevgeni Berzak. 2025. *Decoding reading goals from eye movements.* In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5616–5637. Association for Computational Linguistics.
- Omer Shubi, Yoav Meiri, Cfir Avraham Hadar, and Yevgeni Berzak. 2024. *Fine-grained prediction of reading comprehension from eye movements.* In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3372–3391, Miami, Florida, USA. Association for Computational Linguistics.
- Abhinav Deep Singh, Poojan Mehta, Samar Husain, and Rajkumar Rajakrishnan. 2016. *Quantifying sentence complexity based on eye-tracking measures.* In *Proceedings of the Workshop on Computational Linguistics for Linguistic Complexity (CLALC)*, pages 202–212, Osaka, Japan. The COLING 2016 Organizing Committee.
- Lina Skerath, Paulina Toborek, Anita Zielińska, Maria Barrett, and Rob Van Der Goot. 2023. *Native language prediction from gaze: a reproducibility study.* In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 152–159. Association for Computational Linguistics.
- Robyn Speer. 2022. *wordfreq.* Zenodo.
- Jiao Sun, Yufei Tian, Wangchunshu Zhou, Nan Xu, Qian Hu, Rahul Gupta, John Wieting, Nanyun Peng, and Xuezhe Ma. 2023. *Evaluating large language models on controlled generation tasks.* In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing.* Association for Computational Linguistics.
- Andreas Säuberli, Franz Holzknecht, Patrick Haller, Silvana Deilen, Laura Schiffel, Silvia Hansen-Schirra, and Sarah Ebling. 2024. *Digital comprehensibility assessment of simplified texts among persons with intellectual disabilities.* In *Proceedings of the CHI Conference on Human Factors in Computing Systems, CHI '24*, pages 1–11. ACM.
- Leonard Tetzlaff, Florian Schmiedek, and Garvin Brod. 2020. *Developing personalized education: A dynamic framework.* *Educational Psychology Review*, 33(3):863–882.
- Satoru Uchida. 2025. *Generative AI and CEFR levels: Evaluating the accuracy of text generation with ChatGPT-4.* *Vocabulary Learning and Instruction*, 14(1):2078.
- Sowmya Vajjala, Detmar Meurers, Alexander Eitel, and Katharina Scheiter. 2016. *Towards grounding computational linguistic approaches to readability: Modeling reader-text interaction for easy and difficult texts.* In *Proceedings of the Workshop on Computational Linguistics for Linguistic Complexity (CLALC)*, pages 38–48, Osaka, Japan. The COLING 2016 Organizing Committee.

Kevin Yang and Dan Klein. 2021. [FUDGE: Controlled text generation with future discriminators](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.

Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. 2024. [Explainability for large language models: A survey](#). *ACM Transactions on Intelligent Systems and Technology*, 15(2):1–38.

A Language model prompts

We used the user message below to prompt the language model. The system prompt is the default defined in the Hugging Face implementation for meta-llama/Llama-3.2-3B.

```
Write a short story based on the following
title and prompt.
Title: {title}
Prompt: {prompt}
```

```
The story should not be longer than 500 words
Keep in mind that the reader will not see the
prompt, only the story itself.
Do not include the title.
```

The specific titles and prompts substituted into this template are listed below.

1. **The Forgotten Message:** A character finds an old voicemail on their phone from a number they don't recognize. The message is short, cryptic, and chilling.
2. **Breakfast with a Stranger:** While eating at a local diner, a character is mistakenly given someone else's breakfast order—and a note tucked under the plate that simply reads, "It's time."
3. **One Last Wish:** A child claims their elderly pet goldfish granted them a wish before it died. The next day, something unexpected happens.
4. **Delayed Flight:** Two strangers, both stuck in an airport during a storm, strike up a conversation that changes the course of their lives.
5. **Second Place:** A lifelong second-place finisher in everything—from contests to relationships—finally wins something. But the prize comes with a catch.
6. **The Last Light On:** In a neighborhood black-out, only one house still has power — but no one seems to live there.

B Gaze model training details

GPT-2 fine-tuning:

- Base model checkpoint: openai-community/gpt2
- Regression head: 1 linear layer
- Loss function: mean squared error
- Batch size: 10
- Learning rate: 1×10^{-4}

- Early stopping patience: 3 epochs
- Optimizer: AdamW

Linear regression baseline:

- Features:
 - Word length (in characters) of current word and two previous words
 - Word frequency (according to *wordfreq*; Speer, 2022) of current word and two previous words
- Implementation: *scikit-learn* (Pedregosa et al., 2011)

Evaluation metrics on the held-out 20% test set are listed in Table 2.

Model	First-pass reading time		
	MSE ↓	MAE ↓	R^2 ↑
Transformer	0.631	0.597	0.355
Linear regression	0.671	0.615	0.314

Model	Go-past time		
	MSE ↓	MAE ↓	R^2 ↑
Transformer	0.715	0.576	0.292
Linear regression	0.843	0.600	0.165

Table 2: Evaluation metrics for the two gaze models.

C Eye-tracking data quality

- Average calibration accuracy: 0.50° (min: 0.29°, max: 0.82° across participants)
- 7/144 trials removed
12/144 trials partially removed
- 2.47% data loss in raw data (e.g., blinks) (min: 0.47%, max: 16.30% across participants)
- 4.42% data loss in cleaned data (including manually removed areas) (min: 0.07%, max: 38.90% across participants)
- Average comprehension question accuracy: 94.4% (min: 5/6, max: 6/6 questions across participants)

D Linear mixed-effects model details

We used *statsmodels* (Seabold and Perktold, 2010) to fit the linear mixed-effects models. Table 3 shows the model specifications and coefficients for both reader groups. In the model formulas, *fprt* is the first-pass reading time on the area of interest (word), *gaze_weight* is a categorical variable with three levels (reference/0, positive/+2, and negative/-2), *word_length* is the number of characters in the area of interest (including punctuation), *word_prevalence* is the word's prevalence value according to Brysbaert et al. (2018), and *reader* is the reader's ID.

Model	Intercept	<i>gaze_weight</i> [+2]	<i>gaze_weight</i> [-2]	<i>word_length</i>	<i>word_prevalence</i>	R^2
L1 readers:						
<i>fprt</i> ~ <i>gaze_weight</i> + (1 <i>reader</i>)	90.1 ^{101.3} _{78.9}	30.3 ^{34.3} _{26.2}	-11.4 ^{-7.8} _{-15.0}	—	—	0.036
<i>fprt</i> ~ <i>word_length</i> + <i>word_prevalence</i> + (1 <i>reader</i>)	26.8 ^{46.7} _{6.8}	—	—	21.0 ^{21.6} _{20.4}	-12.7 ^{-5.4} _{-20.0}	0.179
<i>fprt</i> ~ <i>gaze_weight</i> + <i>word_length</i> + <i>word_prevalence</i> + (1 <i>reader</i>)	27.1 ^{47.2} _{6.9}	-0.1 ^{3.7} _{-3.9}	-0.7 ^{2.6} _{-4.0}	21.0 ^{21.6} _{20.4}	-12.7 ^{-5.3} _{-20.0}	0.179
L2 readers:						
<i>fprt</i> ~ <i>gaze_weight</i> + (1 <i>reader</i>)	129.9 ^{146.9} _{112.9}	47.7 ^{52.6} _{42.8}	-13.4 ^{-9.0} _{-17.9}	—	—	0.054
<i>fprt</i> ~ <i>word_length</i> + <i>word_prevalence</i> + (1 <i>reader</i>)	63.3 ^{88.9} _{37.7}	—	—	27.4 ^{28.1} _{26.7}	-22.3 ^{-13.4} _{-31.2}	0.206
<i>fprt</i> ~ <i>gaze_weight</i> + <i>word_length</i> + <i>word_prevalence</i> + (1 <i>reader</i>)	59.1 ^{85.0} _{33.3}	8.1 ^{12.7} _{3.4}	0.1 ^{4.2} _{-3.9}	27.0 ^{27.7} _{26.2}	-20.6 ^{-11.6} _{-29.5}	0.206

Table 3: Linear mixed-effects model formulas, effects with 95% confidence intervals, and explained variance (conditional R^2).