



Science out of its Ivory Tower: improving accessibility with reinforcement learning

Haining Wang, Jason Clark, Hannah McKelvey, Leila Sterman, Zheng Gao, Zuoyu Tian, Sandra Kübler, Xiaozhong Liu

Wang, Haining, Jason Clark, Hannah McKelvey, Leila Sterman, Zheng Gao, Zuoyu Tian, Sandra Kübler, and Xiaozhong Liu. "Science out of its Ivory Tower: improving accessibility with reinforcement learning." *Scientometrics* 130, no. 8 (2025): 4519-4543.

Accessibility Disclaimer:

For a more accessible version of this document, please submit an accessibility request form through the Montana State University Library website.



Science out of its Ivory Tower: improving accessibility with reinforcement learning

Haining Wang¹ · Jason Clark² · Hannah McKelvey² · Leila Sterman² · Zheng Gao³ · Zuoyu Tian⁴ · Sandra Kübler¹ · Xiaozhong Liu⁵

Received: 7 November 2024 / Accepted: 8 July 2025 / Published online: 19 July 2025
© The Author(s) 2025

Abstract

A vast amount of scholarly work is published daily, yet much of it remains inaccessible to the general public due to dense jargon and complex language. To address this challenge in science communication, we introduce a reinforcement learning framework that fine-tunes a language model to rewrite scholarly abstracts into more comprehensible versions. Guided by a carefully balanced combination of word- and sentence-level accessibility rewards, our language model effectively substitutes technical terms with more accessible alternatives, a task which models supervised fine-tuned or guided by conventional readability measures struggle to accomplish. Our best model adjusts the readability level of scholarly abstracts by approximately six U.S. grade levels—in other words, from a postgraduate to a high school level. This translates to roughly a 90% relative boost over the supervised fine-tuning baseline, all while maintaining factual accuracy and high-quality language. An in-depth analysis of our approach shows that balanced rewards lead to systematic modifications in the base model, likely contributing to smoother optimization and superior performance. We envision this work as a step toward bridging the gap between scholarly research and the general public, particularly younger readers and those without a college degree.

Keywords Accessible language · Science communication · Language model · Text simplification · Reinforcement learning · Open science

Introduction

Accessible language in science communication

At first glance, the daily publication of tens of thousands of scientific papers—many freely accessible through open science and open access initiatives—suggests few barriers to knowledge dissemination. However, two key facts challenge this perception and show that significant barriers remain. A recent survey by the U.S. Department of Education found that more than half of U.S. adults aged 16 to 74 (54%, or 130 million people) read at or below a sixth-grade level (Rothwell, 2020). Meanwhile, an analysis of the readability of biomedical research abstracts published from 1881 to 2015 found that scientific writing

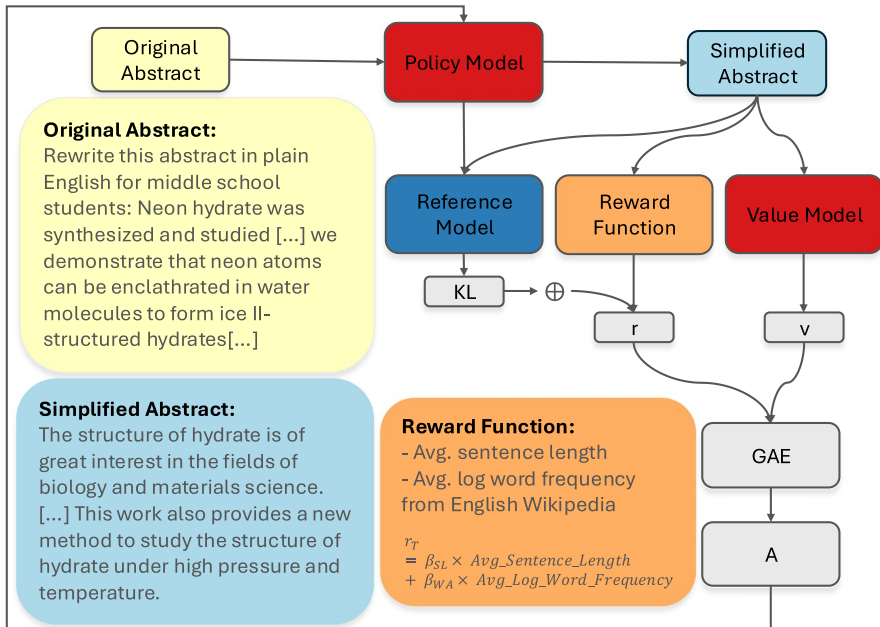


Fig. 1 RLAM training workflow: RLAM rewrites scholarly abstracts using PPO (Sect. “**Reinforcement learning from accessibility measures**”), guided by a reward function that balances average sentence length and word accessibility. The policy model is optimized iteratively through an actor-critic framework: it generates simplified abstracts, whose quality is assessed by the reward function and regularized through KL divergence from the frozen reference model (the supervised fine-tuning model). The reward signal is further contrasted with expected returns estimated by the value model, implemented as a linear head atop the policy model. The resulting advantage is distributed across tokens via Generalized Advantage Estimation (GAE). RLAM enables the lightweight Gemma-2B model to reduce abstracts from postgraduate to high school readability, achieving a 90% improvement over the supervised baseline

has become increasingly hard to read over time (Plavén-Sigraý et al., 2017). Even when intended to be accessible, scientific abstracts typically require a postgraduate level of reading comprehension due to jargon use and sentence structure (Wang & Clark, 2025). This discrepancy leaves a significant portion of the population—including young readers and adults without advanced degrees—unable to fully engage with scientific works, even if these are made freely available online. The “infodemic” surrounding COVID-19 highlighted this issue: the urgent need for understandable information about the virus clashed with the complex presentation of scientific findings—leading many to turn to more digestible but less reliable narratives on social media (Wang et al., 2019; Islam et al., 2020; Calleja et al., 2021) (Fig. 1).

While the legal and medical fields have long been encouraged to use accessible language as a clear conduit for public engagement (Mazur, 2000; Petelin, 2010), momentum for the adoption of accessible language within scientific communities has been building, roughly since the start of the open science movement (Schriver, 2017). For instance, the National Institutes of Health (NIH) advocates for “clear and simple” principles when communicating with audiences with limited health literacy, and the *Proceedings of the National Academy of Sciences of the United States of America* (PNAS) requires authors to submit a significance statement accessible to non-experts (Berenbaum, 2021; Pool et al.,

2021). However, there are inherent conflicts between the specialized nature of communication among disciplinary peer scholars and the public-oriented dissemination of scientific findings. Even assuming that communicating scholarly works in plain language is possible, it will inevitably increase the communication cost among domain experts and create confusion at the more advanced levels, compared to the use of jargon and technical terms. In an era of increasingly specialized scientific research, this conundrum is not easily addressed by scientists or disseminators. Moreover, expecting individual researchers to translate their findings for the public imposes an unrealistic burden, given their primary focus on advancing specialized knowledge. At the same time, public demand for clarity remains high, especially during fast-moving scientific crises where accessible information is urgently needed.

In response, we propose addressing the need for communicating scientific findings to a broader audience by *rewriting scholarly abstracts with simpler words and grammar using language models*. Since readability is key to comprehending scholarship (Flesch, 1946; DuBay, 2004; Kerwer et al., 2021), we envision the resulting accessible narratives as paving the way for the “last mile” of science, broadening access to scientific understanding and engagement, especially for younger readers and those without a college degree.

Challenges to effective simplification

Fine-tuning a language model using pairs of abstracts and their accessible versions is the de facto method for automating the rewriting of scholarly abstracts into more accessible versions (Xu et al., 2015; Goldsack et al., 2022; Joseph et al., 2023). Accordingly, we introduced the Scientific Abstract-Significance Statement (SASS) corpus (Wang & Clark, 2025), a dataset composed of paired abstracts and significance statements from diverse disciplines, with the latter targeting “an undergraduate-educated scientist outside their field of specialty” (Berenbaum, 2021; Pool et al., 2021). Although the simplified abstracts generated from language models fine-tuned on the SASS corpus are approximately three grade levels more readable than the original abstracts, as measured by U.S. grade-based readability scores (Wang & Clark, 2025, Sect. 6), the documents are still not sufficiently accessible; even the best models produce college-level texts. Additionally, because the vocabulary used in significance statements is often just as complex as that found in the abstracts themselves (Wang & Clark, 2025, Sect. 3), the readability improvements are primarily due to shorter sentences, and technical terms remain inadequately addressed.

Alternatively, the optimization of a language model can be guided by a chosen objective in an actor-critic manner (Ramamurthy et al., 2023). It is intuitive to choose an established document readability measure, such as the Automated Readability Index (ARI; see Sect. “[Simplification & accessibility](#)”), to assess the overall readability of the outputs generated by the language model.¹ However, we found that the optimization of language models guided by ARI is highly unstable, often resulting in the production of seemingly more accessible versions that still contain many technical terms. Inspired by Riddell and Igarashi (2021), we decomposed the measurement of document readability into two distinct

¹ The objective can also be reading difficulties ranked by human raters, similar to reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022) where the feedback focuses specifically on document readability rather than general preferences. While this type of feedback is of high quality, it is significantly more expensive than our reward functions. Additionally, document readability is more tangible and therefore easier to quantify than the more ambiguous concept of “human preference.” For these reasons, we did not consider using human raters in this study.

measures: one at the sentence level and one at the word level. We then prioritized word-level accessibility in the optimization to encourage the model to use more accessible words instead of simply shortening sentences.

Contribution

Our work aims to serve as a bridge between scholarly works and the general public, particularly benefiting younger readers and those without a college degree.

1. We address the common challenges in science communication by rewriting scholarly abstracts at a high school reading level using a language model.
2. We identify the challenges language models face in properly addressing jargon and propose Reinforcement Learning from Accessibility Measures (RLAM) as a means to improve the models' use of accessible terms in their rewrites. RLAM-trained language models can significantly reduce the reading level of a scholarly abstract from a post-graduate level to a high school level, achieving a 3-grade-level reduction—or about a 90% performance boost—compared to models fine-tuned using the same corpus.
3. We observe systematic differences between reinforcement learning models guided by different rewards and conclude that disproportionate weights for sentence-level rewards contribute to unstable training and lower simplification quality.

Our code, training logs, and model generations are available at <https://github.com/Wang-Haining/RLAM> under a permissive license.

Literature review

The common challenges of science communication

Science communication research uses several models to describe strategies for public understanding of and engagement with science, including the deficit model, the dialogue model, and the participatory model (Trench, 2008; Hetland, 2014; Sokolovska et al., 2019). The deficit model assumes that the public lacks scientific knowledge and that science communication should focus on simplifying complex scholarship through one-way information transmission. The dialogue model encourages two-way communication, where scientists and the public engage in conversations that foster mutual understanding (Trench, 2008; Joly & Kaufmann, 2008). The participatory model involves the public as active participants in the scientific process, recognizing their contributions as equally valuable to those of scientists (Brossard & Lewenstein, 2009). In practice, the dialogue model suggests that researchers should engage with the public through mediums such as social media (Davies, 2008; Hara et al., 2019; Knox & Hara, 2021), by crafting more digestible manuscripts in research (Maurer et al., 2021), and by developing plain language guidelines for practitioners (Greene et al., 2017). Communicators are encouraged to scaffold content by providing the audience with relevant, easily understandable material that either gradually increases in complexity or includes simplified explanations, examples, and references to better support the learning process (Wolfe, 2008; Landrum et al., 2014). These models often coexist and overlap in practice (Brossard & Lewenstein, 2009; Metcalfe, 2022), with practitioners employing multiple strategies to increase the readability of complex information.

In regular practice, two notable challenges persist for researchers when attempting to effectively communicate their findings to lay audiences. First, placing the burden of effective communication on scholars, who are trained in conducting research and not in effective communication, can be overwhelming. Second, translating scientific findings that heavily utilize specialized jargon often involves sacrificing some specificity in favor of accessibility. Balancing the precision required by the scientific community with the clarity needed for public understanding remains a significant challenge. However, in order to uphold the fundamental principles of open science and the goal of science communication, the public's understanding of scholarship should be considered a requisite part of the scholarly process (Burns et al., 2003; Fecher & Friesike, 2014).

Text simplification

The natural language processing community addresses document accessibility through automated *text simplification*, a process that rewrites complex documents using simpler grammar and vocabulary while retaining the original meaning (Chandrasekar et al., 1996; Alva-Manchego et al., 2020; Al-Thanyyan & Azmi, 2021). In the early days of text simplification, knowledge-based approaches were widely used (De Belder & Moens, 2010; Zhao et al., 2018; Hijazi et al., 2022). These methods relied on predefined sets of linguistic rules to transform complex sentences into simpler forms. For example, such systems would apply syntactic transformations, such as breaking down compound sentences into simpler ones, and lexical substitutions, where difficult words were replaced with more frequent, simpler synonyms. While effective in specific cases, rule-based systems were limited by the rigidity of their rules, often struggling with the variability of natural language.

With the rise of machine learning, approaches that learn simplification patterns from large corpora of text have become prominent. The most common method currently is fine-tuning a pre-trained language model with parallel corpora containing paired documents of different readability levels (Xu et al., 2015; Goldsack et al., 2022; Joseph et al., 2023). These parallel corpora allow the model to “translate” complex texts into simpler ones, typically using cross-entropy loss to align the predicted simplified version with the target text. This shift to data-driven methods allows for greater flexibility and more nuanced simplifications, as the models learn from a broader range of linguistic patterns.

The effectiveness that supervised fine-tuning can achieve depends on the quality of the parallel corpus, specifically whether the simplified documents are sufficiently accessible. As an example, Devaraj et al. (2021) achieved approximately a one-grade reading level improvement, comparable to the readability of the plain-language samples, by fine-tuning a language model [i.e., BART (Lewis et al., 2019)] with a corpus of pairs of technical and plain-language medical texts. This is expected: for sequence-to-sequence models, such as BART and T5, the objective of supervised fine-tuning is to minimize the Kullback–Leibler (KL) divergence between the predicted token distribution and the target distribution. Similarly, for causal language models, such as GPT and Gemma, minimizing KL divergence is equivalent to maximizing the joint probability of generating accessible sequences, conditioned on the abstracts. Consequently, the best outcome that models can achieve through supervised fine-tuning is to yield lay versions that as simple as—but *no simpler than*—the simplified examples in the corpus.

The problems is that, in practice, high-quality corpora for simplification are either unavailable or prohibitively expensive to obtain, and widely adopted text simplification targets are often only marginally more readable than their sources. For example, the Simple

English Wikipedia is, on average, less than two grade levels more readable than standard English Wikipedia articles (Isaksson, 2018). This issue also extends to abstract simplification tasks, where the simplified documents are typically less than one grade level more readable than the abstracts themselves (Wang & Clark, 2025). Carlson et al. (2018) proposed evaluating the readability of simplified texts using different versions of the Bible, which cover a wide spectrum of readability levels. This approach is specifically designed for evaluation purposes rather than for training a general-purpose simplification model.

Efforts to overcome this data bottleneck have focused on two main approaches. First, scholars have modified the cross-entropy objective to achieve additional improvements, often by incorporating knowledge on word-level difficulty (Nishihara et al., 2019; Devaraj et al., 2021; Yanamoto et al., 2022). For instance, by adding pre-determined terms to the cross-entropy objective to reduce the model's probability assigned to a set of technical words—identified using a classifier trained on essays of different readability levels—a BART model can achieve an additional one-grade reading level improvement over one trained with the standard cross-entropy objective (Devaraj et al., 2021). However, while including extra terms in the objective function can be useful, this approach tends to be rather rigid and may compromise language quality. Moreover, the performance gains from these improvements only improve readability by around one grade level. Additional performance gains have also been observed when fine-tuning larger language models (ranging from 2 to 7 billion parameters) (Wang & Clark, 2025). Having been exposed to a significantly larger number of documents, the model may land on a parameter landscape slightly better than the one represented by the target documents. Although using a larger language model results in about a three-grade-level boost in readability, it still falls approximately three grade levels short of making the abstracts generally accessible to our target audience—individuals without a college degree.

Our means of addressing the data bottleneck is to optimize language models by incorporating dynamically adjusted terms into the cross-entropy objective in an actor-critic manner. This method aligns with reinforcement learning techniques that integrate multiple rewards from different types of feedback; in long-form question-answering tasks, Wu et al. (2023) have demonstrated that incorporating various reward types—such as coherence, relevance, and completeness—outperforms approaches relying on a single, holistic reward. We adapted this approach by assessing document readability through a careful balance of two heuristics proposed by Riddell and Igarashi (2021); these estimate syntactic and lexical difficulty, respectively.

Scientific abstract-significance statement (SASS) corpus

We used the Scientific Abstract-Significance Statement (SASS) corpus in our experiments. This corpus is composed of 3430 abstract-significance statement pairs derived from *PNAS* and divided into training (3030 samples), validation (200 samples), and test sets (200 samples) (Wang & Clark, 2025). It covers a wide range of disciplines, ensuring diverse representation across various fields, as shown in Fig. 2. Corpus statistics are shown in Table 1; refer to Sect. “[Evaluation](#)” for a detailed description of the measures.

We observed that significance statements are semantically coherent with their corresponding abstracts. The corpus statistics indicate that significance statements are more readable than abstracts, as shown by lower mean values in the Automated Readability

Table 1 Corpus statistics for the Scientific Abstract-Significance Statement (SASS) corpus

Section	ARI↓	F-K↓	VOA	SL↓	WA	WL↓
Abstract	18.9 (2.8)	19.2 (2.4)	−0.43 (0.25)	25.4 (4.9)	12.0 (0.4)	5.3 (0.4)
Significance	18.1* (3.1)	18.6* (2.7)	−0.31* (0.26)	23.9* (5.3)	11.9* (0.4)	5.4* (0.4)

Metrics include: ARI (Automated Readability Index), F-K (Flesch-Kincaid readability test), VOA (log ratio of proportion of words found in the VOA1500 vocabulary), SL (average sentence length and number of sentences), WA (word accessibility; log frequency per 1 billion tokens in English Wikipedia), and WL (average word length). Measures whose names are followed by a down arrow symbol (↓) indicate that lower values correspond to a more readable document. Numeric values in parentheses are the corresponding standard deviations. Paired t-tests were conducted for each metric comparing the abstracts and significance statements, with *p*-values adjusted using the Bonferroni correction for multiple comparisons. The observed differences in each of the measurements are statistically significant (**p* < 0.05) after adjusting for the grouped *p*-values.

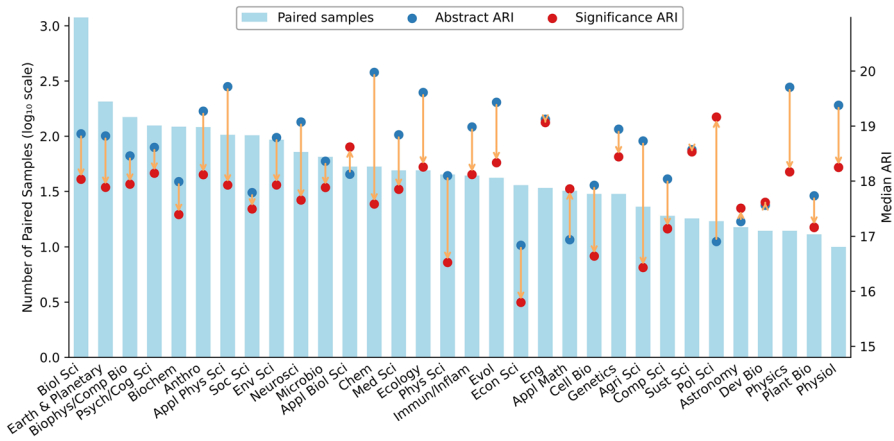


Fig. 2 Discipline and readability distributions of abstracts and significance statements found in the training set of the Scientific Abstract-Significance Statement corpus. The count of paired samples in different disciplines is shown in blue bars on a log10 scale (disciplines with fewer than ten samples are not shown). Readability is measured using the Automated Readability Index (ARI), which estimates the number of years of schooling required to understand a text. On average, abstracts have a readability slightly below 20 ARI, indicating a post-graduate level. Significance statements are generally more readable than their corresponding abstracts. Orange arrows indicate the change in readability from abstracts to significance statements

Index (ARI) and Flesch-Kincaid readability test (F-K). This suggests that the SASS corpus can be useful in simplifying scholarly abstracts across diverse disciplines.

We also observed that word accessibility (i.e., log frequency per 1 billion tokens found in English Wikipedia) and average word length suggest that significance statements can be less accessible at the word level than are their corresponding abstracts. Although the log ratio of words found in the VOA1500 vocabulary is slightly lower than in the corresponding abstracts, these 1500 words are very basic and include a high proportion of function words. Considering that significance statements use approximately 1.5 fewer words on average, the increased use of VOA words may be a consequence of the higher use of function words to maintain grammaticality.

Reinforcement learning from accessibility measures

Language modeling via proximal policy optimization

At the core of our approach is language modeling with Proximal Policy Optimization (PPO) (Schulman et al., 2017) guided by two accessibility measures. A causal language model trained on large corpora can generate the next token based on the current sequence, which is useful in the context of reinforcement learning for developing a policy model that determines the most appropriate next token to maximize the expected return in terms of document readability. We initiated the policy model using the weights from language models pretrained on large corpora (see Sect. “Training”) and then fine-tuned on the SASS corpus $\pi_\theta = \pi_{\theta_{\text{SFT}}}$.

The process begins with an input sequence $s_0 = (a_0, a_1, \dots, a_t)$, where each a_t is from a set of tokens W , and s_0 represents an abstract formatted in a simple template.² The language model π_θ then generates $a_0, a_1, \dots, a_{T-1} \sim \pi_\theta(\cdot | s_t)$, creating its accessible version until the maximum number of tokens T is reached, either due to the context length or an end-of-sentence token:

$$\pi_\theta(a_0, a_1, \dots, a_{T-1}) = \prod_{t=0}^{T-1} \pi_\theta(a_t | s_t) \quad (1)$$

Our objective is to learn a policy model that, given an abstract, models the joint probability of tokens leading to a high reward in terms of accessibility while maintaining semantic coherence. Formally, this is expressed as:

$$J(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^{T-1} (r(s_t, a_t) - \beta_{\text{KL}} \text{KL}(\pi_\theta(a_t | s_t) \parallel \pi_{\theta_{\text{SFT}}}(a_t | s_t))) \right] \quad (2)$$

here $J(\pi_\theta)$ represents the expected return when following policy π_θ . The reward $r(s_t, a_t)$ is estimated for each time step t in the trajectory $\tau = (s_0, a_0, s_1, a_1, \dots, s_{T-1}, a_{T-1})$, where s_t is the sequence of tokens at time t , formed as the concatenation of s_0 and the tokens a_0, a_1, \dots, a_{t-1} . This formula iteratively computes the rewards given the current sequence s_t and the token a_t chosen by the policy. The β_{KL} -weighted KL divergence term $\text{KL}(\pi_\theta(a_t | s_t) \parallel \pi_{\theta_{\text{SFT}}}(a_t | s_t))$ is applied at every step of sequence generation to ensure the policy does not deviate significantly from the supervised fine-tuned model. This is crucial because, without such a constraint, the policy model might quickly learn to output whatever the reward model favors to maximize its return, which can lead to undesirable behaviors. For instance, the model might repeatedly output a frequent word (“is is is ...”), achieving a high reward based solely on accessibility measures but lacking meaningful content. Following Stiennon et al. (2020), β_{KL} is dynamically adjusted by targeting a specific KL divergence between π_θ and $\pi_{\theta_{\text{SFT}}}$ using a capped proportional controller in logarithmic space. The benefit of using a dynamic KL control, as opposed to a fixed one, is that it allows the model to adapt more flexibly to different stages of training, accommodating varying levels of KL divergence between the policy and SFT model. For further details, see

² The input reads: “Rewrite this abstract in plain English for middle school students: {abstract}\n Lay summary: {significance}”.

“[Appendix A](#)” for the specific PPO implementation we are using and “[Appendix B](#)” for the dynamic proportional controller.

Reward function

The reward function evaluates the overall quality of the output (s_T). After the initial failures of testing a traditional readability measure (i.e., ARI) as the criterion, we decided to use a balance of two accessibility measures: average sentence length in words and word accessibility, adopted from Riddell and Igarashi (2021), to guide the optimization.

Word accessibility reward

A word’s accessibility is approximated by how frequently it appears in a large reference corpus. We chose the English Wikipedia corpus due to its domain similarity and applied a Moses tokenizer, yielding a vocabulary of 14.6 million types from a total of 3.6 billion tokens. If a token is among the most common 100,000 types, we report its frequency per billion tokens as its accessibility measure. Otherwise, we estimate its frequency using ridge regression with an ℓ_2 -norm coefficient equal to 1.0. This model allows us to make serviceable estimates of the frequency of arbitrary tokens, including tokens that do not appear in the reference corpus. This model takes as input the token’s length in Unicode code points, its byte unigrams, byte bigrams, and byte trigrams. The model estimates the token’s log frequency per 1 billion tokens. We used the natural logarithm of frequencies per billion tokens as the measure of word accessibility.³ For example, the accessibility score for “big” is 11.8, while “colossal” scores 7.3. Despite being comparable in meaning, the model’s production of the latter will receive fewer rewards. Coefficient β_{wa} is to control the scale of the credit given for word accessibility.

Sentence length reward

Sentence length is also determined by a Moses tokenizer, which preserves hyphenation and splits contractions.⁴ We negate the value of sentence length for intuitive calculation of the rewards for optimization.

³ We have faithfully followed the experiment of Riddell and Igarashi (2021) with three differences. First, our reference corpus is the English Wikipedia, whereas the original study used the Common Crawl News corpus. Second, we did not discard duplicated sentences as Riddell and Igarashi (2021) did, because we found that sentence duplication is not common in Wikipedia. Third, the original study reported word *inaccessibility* scores by negating the logarithm of frequency per billion. We report *accessibility*, without negation, because it is more naturally suited to serve as a reward. Refer to (Riddell & Igarashi, 2021, pp. 1186–1187) for the training of the ridge regression.

⁴ For the Moses rule-based tokenizer, we use the `sacremoses` Python package.

Experiment setup

Training

We initialized the policy models π_θ by adopting the Gemma-2B checkpoint reported by Wang and Clark (2025) ($\pi_{\theta_{\text{SFT}}}$). The original Gemma-2B was trained on three trillion tokens, consisting of publicly available data as well as proprietary datasets comprising “primarily English data from web documents” (Mesnard et al., 2024). The specific checkpoint we adopted was fine-tuned using the SASS corpus in a straightforward manner. It was chosen for its strong performance in simplification quality, its faithfulness, and its relatively compact size.

The two accessibility rewards were weighted as follows: the word accessibility reward was set to $\beta_{\text{WA}} = 4.0$, and we report four models with gradually increasing β_{SL} values of 0.05, 0.08, 0.1, and 0.2.⁵ For adaptive control of the per-token semantic reward, we started with an initial $\beta_{\text{KL}} = 0.2$ and targeted a KL divergence of 8.0 nats during the training course, capping it in the range between 0.15 and 0.25. We used a micro batch size of 4, with each sequence used to run the PPO algorithm for 4 epochs using importance sampling, with gradient accumulation steps set to 4. We used a clip range of 0.2 for the policy gradient and value function estimation to ensure stability. The value function coefficient was set to 0.1. The optimization used standard AdamW optimizer parameters ($\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1 \times 10^{-8}$) (Kingma & Ba, 2015; Loshchilov & Hutter, 2019). The learning rate was fixed at 1×10^{-6} . The training was conducted on 2 H100 (80GB) GPUs using mixed precision training in bfloat16. The sampling temperature was set to 0.7 in the rollout phase. Following Huang et al. (2024), we assigned unfinished roll-outs (indicated by the missing end of sequence token) with a fixed low score to encourage the model to generate complete simplified narratives. We also report a model trained as dictated by ARI,⁶ (i.e., RLARI; described in Sect. “Challenges to effective simplification”), with all other parameters kept the same as those guided by the accessibility measures.

We performed multiple runs for each reinforcement learning process and selected the checkpoint to report based on a balance of semantic retention and ARI score obtained on

⁵ Selecting hyperparameters for PPO training is a computationally intensive task, particularly due to the sensitivity of reward model weights. We grounded our search in intuitive defaults commonly adopted in PPO literature—learning rate, batch size, and value function coefficient—and verified training stability on a surrogate task (IMDb sentiment alteration). After confirming reliable PPO dynamics, we conducted a targeted random sweep over β_{WA} and found that values near $\beta_{\text{WA}} = 4.0$ consistently produced stable learning and favorable trade-offs between readability and semantic retention. We then refined the joint configuration of β_{KL} and β_{WA} to further stabilize optimization and prevent collapse and semantic drift. While we acknowledge that this hyperparameter configuration may not be globally optimal, it reflects a deliberate balance between simplicity and effectiveness. Rather than relying on a separate language model trained with extensive human annotations to approximate human evaluation (Ouyang et al., 2022) our reward function—implemented in just a few lines of Python—offers an elegant and computationally efficient solution for guiding the simplification process. In addition, because word accessibility is in logarithmic space, we subtracted 10 from the estimated word accessibility and reset any values lower than 10 to 0 to keep them within a reasonable range.

⁶ We began with pilot experiments using ARI as the sole reward, and observed that it encouraged the model to take shortcuts, often collapsing outputs into overly compressed summaries—sometimes as short as a single sentence. While such ultra-brief abstracts resemble the desired behavior in extreme summarization tasks (Narayan et al., 2018) they proved suboptimal for enhancing accessibility. These results underscored the need to explicitly guard against excessive sentence shortening; without such constraints, reinforcement learning risks compromising the semantic fidelity of the output.

the validation set. We observed that the readability of the generated text on the validation set often began to decline rapidly after plateauing for a while, typically accompanied by a surge in the standard deviation of average word accessibility among sentences. This signals that the language model was sacrificing language quality and semantic relevance for extra improvements in readability. Therefore, we reported the checkpoint immediately before such instability occurred.

Evaluation

We evaluated the simplified texts generated by the language models trained with the reinforcement learning framework using 200 abstracts from the SASS corpus test set for simplification. Though advanced decoding methods might further refine the quality of the outputs, we used multinomial sampling with the temperature set to zero to intentionally produce the most deterministic outputs. This approach helped us better understand the modeling of accessible language and made it easier for us to identify potential quirks. We assessed the quality of the generated simplified abstracts both quantitatively and qualitatively. Quantitatively, we measured the generated texts based on their semantic retention and accessibility using established relevance measures as well as readability and accessibility metrics.

Semantic retention

BERTScore calculates the cosine similarity between each token in the candidate sentence and each token in the reference sentence using contextual embeddings from a pre-trained language model; its results align well with human judgment on semantic similarity evaluation (Zhang et al., 2020). It is not directly influenced by lexical overlap, making it more suitable for evaluating simplification systems than are metrics that rely on matching words, such as BLEU (Papineni et al., 2002). For our evaluation, we used embeddings from the 18th layer of a BERT-large-uncased model and reported the F1 score. This choice is based on prior findings indicating that the 18th layer yields a strong Pearson correlation (0.72) on the WMT16 To-English benchmark (Zhang et al., 2020).

Simplification & accessibility

Accessibility can be measured with respect to the overall simplification quality (SARI); readability (ARI and Flesch-Kincaid); and other straightforward document complexity measures, including average sentence length, word accessibility (i.e., log frequency per 1 billion tokens found in English Wikipedia), the log ratio of its proportion of VOA Special English words (1517 types in total), and average word length.

SARI (System output Against References and against the Input sentence) is specifically designed to evaluate text simplification (Xu et al., 2016). It aims to measure how well a simplified text retains the original meaning while improving readability. SARI provides a balanced measure of how well a text simplification system performs by focusing on the necessary operations of adding, deleting, and retaining words.

ARI and Flesch-Kincaid readability tests assign a numerical score to text that reflects the U.S. grade level required for comprehension. Lower scores (1–13) indicate content suitable for kindergarten through twelfth grade, with each score corresponding to a subsequent grade level. Scores in the range of 14–18 suggest college-level readability, ranging from

first- to senior-year content appropriateness. Higher scores (19 and above) are associated with advanced college education. Both measures use average sentence length. Flesch-Kincaid uses syllables per word, while ARI uses characters per word for its linear combination with sentence length.

We harvested VOA Special English vocabulary comprising 1517 unique words (VOA1500).⁷ We calculate the ratio of words that appear in the VOA1500 to those that do not, then report the natural logarithm of this ratio for each generated sample. Values above 0 indicate that the text contains more Special English words than non-Special English words, and a higher value indicates a greater presence of “easy” words. We chose the VOA word list because it is a well-known, publicly available vocabulary list specifically designed for learners of English as a second language. Its use aligns with established heuristics in readability research, which suggest that texts using learner-friendly words tend to be more accessible.

Language quality, faithfulness, & completeness

We manually examined 5% of all generated samples, corresponding to a randomly chosen subset of the test set from the SASS corpus. Each generation is annotated with respect to language quality, faithfulness, and completeness using a rubric of Good, Acceptable, and Poor. We focused on fluency and grammaticality and hand-picked both good and problematic examples when evaluating language quality. For the evaluation of faithfulness, we conduct close readings to assess the extent to which a simplified abstract remains factually faithful to the original narrative. If uncertainty arises, we consult the corresponding manuscript, as our abstract simplification system must avoid producing misinformation. Completeness is also a key consideration, as it is essential to include the main findings and implications of the research for the general public, since this is the primary goal of scientific dissemination. We have documented the detailed evaluation rubrics, along with the backgrounds of the annotators, in “[Appendix C](#)”.

Token distribution shift analysis

To understand the impact of optimization guided by different rewards, we analyzed the shifts in token distribution in generations from reinforcement learning models compared to those generated by the supervised fine-tuning model from which they were initiated. Following the approach of Lin et al. (2023), we first generated a response using a model trained with reinforcement learning for a given query. We then used the supervised fine-tuning model to predict the most probable token at each position based on the context up to that point.

Token positions were categorized into three groups based on how their ranks shifted between the supervised fine-tuning model and the reinforcement learning-optimized models. The first group, unshifted positions, includes tokens that maintained their top rank in both models. The second group, marginal positions, consists of tokens that dropped slightly in rank, appearing as the second or third choice in the supervised fine-tuning model. Finally, shifted positions are those where a token’s rank fell outside the top three choices

⁷ We included VOA Special English Word Book Sections A–Z, Science Programs, and Organs of the Body hosted on Wikipedia (https://simple.wikipedia.org/wiki/Wikipedia:VOA_Special_English_Word_Book).

in the supervised fine-tuning model. This categorization allows us to identify which tokens are affected, thereby understanding the impact of reinforcement learning guided by different rewards on the model's decision-making process.

Findings & discussion

To evaluate the effectiveness of RLAM in simplifying scholarly abstracts, we present both quantitative and qualitative analyses. First, we assess whether RLAM can significantly improve readability without compromising content quality. Second, we compare RLAM outputs with those generated by supervised fine-tuning (SFT) and reinforcement learning guided solely by readability (RLARI). We also examine how reinforcement learning optimization shapes model behavior, as reflected in token distribution shifts, to better understand why RLAM outperforms RLARI.

Quantitative assessment

Table 2 summarizes the performance of Gemma-2B, tuned in different ways, when evaluated on the test set of the SASS corpus. The first scenario is the supervised fine-tuned baseline, which performs next-token prediction on the SASS corpus training set. The second and third scenarios are reinforcement learning through PPO guided by ARI or accessibility measures. We assessed the generation quality by considering both semantic retention and simplification, specifically using BERT score (BS), SARI, ARI, Flesch-Kincaid readability test (F-K), the log ratio of words in the VOA1500 vocabulary (VOA), sentence length (SL), word accessibility (WA), and word length (WL). A one-tailed paired t-test was conducted for each metric to compare observations between the reinforcement learning and supervised fine-tuning baselines, assuming improvement in document readability. Bonferroni correction was applied to each set of tests to maintain family-wise significance at 0.05 (* $p < 0.05$).

We observe that reinforcement learning models trained with different rewards exhibit a notable reduction in reading level, bringing abstracts down to high school levels. The model directly guided by ARI (RLARI) achieves an ARI of 12.6, while the most performant model guided by accessibility measures (RLAM, $\beta_{SL} = 4.0$ and $\beta_{WA} = 0.2$) reaches 12.5, both aligning with the readability level expected for individuals who have completed K-12 education (approximately ARI 13). However, RLARI and RLAM models achieve these readability improvements in different ways. For RLAM models, better readability is achieved mostly through improved token-level accessibility. RLAM models show an increase in word accessibility from 0.5 to 0.8 compared to the supervised baseline. This increase in the natural logarithm suggests that words generated by RLAM models are, on average, 1.6 to 2.2 times more frequent in the English Wikipedia corpus than those generated by the SFT model. In comparison, RLARI's 0.1-unit increase in word accessibility, although observed, does not result in a statistically significant change in word frequency compared to the SFT model. Similarly, the log ratio of the VOA1500 vocabulary in the RLAM models shows a significant improvement, with log ratios ranging from -0.02 to 0.08 . This implies that for every 100 non-VOA1500 (or more complex) words generated, RLAM models can produce approximately 98 to 106 VOA1500 basic words. In contrast, the SFT and RLARI models exhibit VOA log ratios of -0.26 and -0.17 , respectively, indicating that for every 100 non-VOA1500 words generated, these models produce only

Table 2 Comparison of Gemma-2B's performance across three approaches: the supervised fine-tuned baseline (SFT), reinforcement learning guided by ARI (using an intermediate checkpoint before significant policy gradient instability was observed, RLARD), and reinforcement learning guided by two accessibility measures (RLAM)

Model	β_{SL}	β_{WA}	ARI \downarrow	F-K \downarrow	SARI	VOA	SL \downarrow	WA	WL \downarrow	BS
SFT	–	–	15.5 (3.0)	16.5 (2.6)	39.1 (5.0)	–0.26 (0.30)	20.6 (4.1)	11.9 (0.5)	5.2 (0.4)	0.64 (0.06)
RLARI	–	–	12.6* (2.9)	14.3* (2.5)	40.1* (4.8)	–0.17* (0.31)	16.4* (3.7)	12.0 (0.5)	5.0* (0.4)	0.64 (0.05)
RLAM	0.05	4.0	13.5* (2.8)	14.8* (2.4)	39.8 (5.1)	0.08* (0.29)	21.0 (4.2)	12.7* (0.5)	4.8* (0.4)	0.62 (0.06)
RLAM	0.08	4.0	13.7* (3.0)	14.9* (2.6)	40.4* (5.2)	–0.02* (0.28)	20.2 (3.7)	12.5* (0.5)	4.9* (0.4)	0.64 (0.05)
RLAM	0.1	4.0	13.3* (2.6)	14.7* (2.3)	40.3 (5.0)	–0.02* (0.31)	19.3* (3.6)	12.4* (0.5)	4.9* (0.4)	0.64 (0.05)
RLAM	0.2	4.0	12.5* (2.9)	14.0* (2.5)	39.8 (5.0)	–0.01* (0.32)	17.7* (3.4)	12.4* (0.5)	4.9* (0.4)	0.63 (0.05)

SFT was fine-tuned using the Scientific Abstract-Significance Statement (SASS) corpus reported in Wang and Clark (2025). The columns labeled β_{WA} and β_{SL} pertain specifically to RLAM, where the rewards for average word accessibility and sentence length are balanced. The inference on the test split from SASS uses multinomial sampling. Metrics ARI, F-K, SARI, VOA, SL, WA, WL, and BS stand for Automated Readability Index, Flesch-Kincaid readability test, log ratio of VOA1500 vocabulary, sentence length, word accessibility, word length, and BERTScore (F1), respectively. Measures followed by a down arrow symbol (\downarrow) indicate that lower values are better. Numeric values in parentheses are the corresponding standard deviations. A paired two-tailed t-test was performed on observations of each measure between each model and the original abstracts. At a model-wise p -value of 0.05, measures that differ significantly from the SFT baseline are marked with an asterisk

around 77 (84) VOA1500 words. The average word length in characters for RLAM models ranges from 4.8 to 4.9, slightly shorter than RLARI’s and outperforming SFT. Overall, the above evidence suggests that RLAM models achieve better readability by using more common, simpler, and shorter words.

On the other hand, the RLARI model achieves better readability by producing much shorter sentences, with only a marginal boost in word-level accessibility. The RLARI model has the shortest average sentence length of 16.4 words, significantly outperforming the SFT model. In comparison, a significantly shorter average sentence length is only observed when RLAM’s sentence length reward coefficient (β_{SL}) exceeds 0.08. We also observed that, in pilot studies, if β_{SL} is set to a higher value, such as 0.5, the model’s optimization will collapse after only a few hundred steps, similar to what we consistently observed in the optimization of RLARI models. The improvement in the RLARI model’s word-level accessibility is inconsistent: while we observed significant gains in the VOA log ratio and word length, the word accessibility measure did not show statistically significant improvement after a Bonferroni correction. That said, although RLARI uses more VOA basic words and shorter words, their frequency in the English Wikipedia corpus is not high enough to result in a significant increase in word accessibility compared to the SFT baseline.

For semantic coherence with the human-written significance statements, the BERT scores from the reinforcement learning models are not significantly worse than those of the SFT model. This finding suggests that the generations from the reinforcement learning models are likely to remain semantically faithful and that the reinforcement learning process does not significantly degrade language quality, at least up to the point before unstable optimization signals were observed. SARI scores from the reinforcement learning models are significantly, yet only marginally, better than those of the SFT baseline. Since SARI is a measurement of a combination of deletions, additions, and retention operations, a straight-forward explanation is not available. However, the higher SARI scores confirm the greater simplification quality of the reinforcement learning models, which is likely due to a combination of simpler words and shorter sentences.

Qualitative analysis

We annotated 5% of the generated simplified abstracts from reinforcement learning models guided by ARI (RLARI) and accessibility measures (RLAM, with $\beta_{WA} = 4.0$ and varying β_{SL} values) with respect to language quality, faithfulness, and completeness, as shown in Fig. 3. The test abstracts included three from biological sciences and one each from




Fig. 3 Results from the annotation of 5% of the generated outputs from reinforcement learning models guided by ARI (RLARI) and accessibility measures (RLAM, with $\beta_{WA} = 4.0$ and varying β_{SL} values), assessing language quality, faithfulness, and completeness

chemistry; mathematics; evolutionary biology; environmental sciences; ecology; economic sciences; and earth, atmospheric, and planetary sciences.

Regarding overall quality, we found that reinforcement learning-trained models generally produced high-quality language. Compared to the SFT generations, RLAM outputs are often shorter and more semantically complete, due to the imposed token budget for newly generated content (241 tokens, the length of the longest significance statement in the training set). The main shortcoming is the presence of small trailing phrases that often deflate readability scores, such as “(PsycINFO Database Record),” “(show more),” and “All rights reserved.” The first artifact is also found in the SFT model generations and is hypothesized to be carried over from the corpora that Gemma-2B was previously exposed to, as this pattern does not appear in the SASS corpus. An informal review of the remaining generations suggests that this phenomenon is amplified, appearing even in samples without such trailers in the SFT generations. The latter two artifacts were newly observed and are caused by the reinforcement learning processes. Where any of these issues appear, we annotate them as “Acceptable,” even though the generations are otherwise fluent and grammatical.

We also found that the proportion of these artifacts steadily increases as training continues, and they are usually found in only a subset of samples across different experiments. Further, when β_{SL} is set to a larger value, such as 0.5, the generations tend to cheat by generating readability-deflating trailers within the first few hundred iterations of optimization. To address this problem, we experimented with an additional heuristic—the standard deviation of sentence-level word accessibility. When this measure exceeds approximately 0.6, it indicates that readability-deflating trailers are beginning to accumulate. We had limited success using this extra reward and hypothesized that it might be too indirect for the language model to effectively learn from amid other signals that may covary with this measure. During manual examinations across experiments, no repetition was observed, which we had previously identified as a major concern in the SFT models (Wang & Clark, 2025).

In assessments of faithfulness, we did not find any models hallucinating in the generations we examined. However, in informal examinations, we did find that reinforcement learning checkpoints may hallucinate by generating simple but overly hedging or short expressions when over-optimized. Generations from RLARI-trained models often remain unfinished, but they retain the main gist of the abstract. Although we do not observe trailer phrases in the RLARI-generated texts like those found in the RLAM models, this issue frequently arises in other RLARI runs. Subsequent checkpoints often exhibit similar problems and tend to deteriorate rapidly once the model starts cutting corners. This typically occurs shortly before the training process completely fails. However, the reported checkpoint happens to miss this characteristic.

Token distribution shift analysis for reinforcement learning models

We annotate each token in the generations according to the three categories defined in our token distribution shift analysis. These categories are: unshifted positions (where the supervised fine-tuning model’s top prediction agrees with that of the reinforcement learning model), marginal positions (where the reinforcement learning model’s top choice is the second or third choice of the supervised fine-tuning model), and shifted positions (where the reinforcement learning model’s top prediction falls outside the top three choices of the supervised fine-tuning model). This categorization allows us to closely examine how reinforcement learning with different rewards influences the behavior of the supervised fine-tuning model from which they were initiated.

Table 3 Token distribution analysis across different RLAM Models and the RLARI Model

Model	β_{SL}	β_{WA}	Marginal tokens	Shifted tokens	Unshifted tokens
RLARI	–	–	2303 (5.96%)	183 (0.47%)	36,128 (93.56%)
RLAM	0.05	4.0	2564 (10.12%)	354 (1.40%)	22,415 (88.48%)
RLAM	0.08	4.0	2297 (8.36%)	274 (1.00%)	24,893 (90.64%)
RLAM	0.1	4.0	2468 (8.00%)	311 (1.01%)	28,054 (90.90%)
RLAM	0.2	4.0	2406 (8.93%)	346 (1.28%)	24,176 (89.78%)

The table shows the counts and proportions of tokens classified as marginal, shifted, and unshifted for each model. RLAM models are evaluated with different β_{SL} values while keeping β_{WA} constant at 4.0. Values in parentheses represent their proportion relative to all generated tokens of the respective model




Fig. 4 Token distribution shift analysis for reinforcement learning models. The figure illustrates the average distribution of marginal and shifted tokens across RLAM models (top row) and RLARI models (bottom row). The left column represents the proportion of marginal tokens, while the right column shows the proportion of shifted tokens, both relative to the total token count at each position in the generated sequences

As shown in Table 3, we quantify the distribution of token types across models and their respective proportions relative to the total token count for each model. The table suggests that the RLARI model modifies the SFT baseline less frequently: 0.5% of tokens are shifted, ca. 6.0% are marginally shifted, while approximately 94.0% remain unaltered. In contrast, the RLAM models exhibit a higher proportion of both shifted (ranging from 1.0 to 1.4%) and marginal (8.0 to 10.1%) tokens: approximately twice as many shifted tokens as the RLARI model, and a noticeably increased number of marginal tokens. We hypothesize that the differences primarily originate from how the sentence length reward is incorporated into the optimization process. Compared to the coefficient of sentence length in the ARI formula, we assign much smaller weights to sentence length in the rewards guiding RLAM models. Since the supervised model has already demonstrated its ability to shorten sentences, it is more reasonable to encourage the model to focus on finding accessible word alternatives for a higher reward, rather than excessively incentivizing shorter sentences. Over-prioritizing sentence length can also exploit the reward system by nudging the model to generate benign but meaningless trailing phrases, further derailing it from producing more accessible words and causing the optimization to collapse prematurely. Another reason for the difference in token type distribution could be the different heuristics for word choice: RLAM models are incentivized to favor more common words, while RLARI prioritizes shorter words.

We also analyzed the distribution of token types with respect to their positions in the generated sequences. In Fig. 4, the top two subplots show the average distributions from the four RLAM sequences for clarity. The left column focuses on marginal tokens, while the right column visualizes shifted tokens. Each subplot illustrates the proportion of each token category relative to the total token count at each position. Our analysis reveals that the

RLAM models exhibit more marginal modifications at the initial positions, whereas shifted tokens tend to accumulate toward the end of the sequences. This pattern suggests that, compared to the supervised fine-tuning model, which tends to overemphasize implications (Wang & Clark, 2025), the RLAM models may truncate redundant content by employing more shifted tokens toward the end. Marginal shifts are used at the beginning to introduce main subjects or background information with simpler words.

In contrast, the RLARI model does not display such a systematic pattern; different types of tokens appear to be more evenly distributed across positions but lack the level of consistency observed in the RLAM models. We hypothesize that this difference may be caused by disproportionately large incentives for shorter sentences, which could hinder effective exploration of alternative, simpler word substitutions and ultimately contribute to the RLARI model's performance variability after the initial phases of optimization. Additionally, approximately 90 percent of the tokens generated by the reinforcement learning models remain unchanged, which roughly corroborates previous findings in RLHF studies, where around 8 percent are changed (Lin et al., 2023). The difference may arise from the nature of the task; we focus on rephrasing in an accessible manner, which often does not require altering many tokens, unlike Lin et al. (2023), who align language models to human preferences, where the expected answer can be quite different from what would otherwise be derived from the base model.

Conclusion

To improve the accessibility of scientific literature to the general public, we implemented reinforcement learning techniques to guide language models, extending beyond the traditional cross-entropy objective. Our study demonstrates that carefully balancing accessibility measures at the word and sentence levels can effectively guide Gemma-2B in simplifying scholarly abstracts, outperforming the supervised fine-tuning baseline by a large margin. This approach achieves these improvements without compromising language quality or faithfulness and mitigates the supervised fine-tuning model's tendency to overemphasize research implications. The best model trained using our method successfully adjusts the readability level of scholarly abstracts by approximately six U.S. grade levels—in other words, from a postgraduate to a high school level. Compared to the supervised fine-tuning model, the words generated by the model trained via our approach are proven to be more common (1.6 to 2.2 times more frequent), easier (with more VOA basic words), and shorter (by 0.3–0.4 characters). This improvement addresses a key limitation of existing corpora, in which the target distribution (i.e., significance statements) often does not adequately prioritize the accessibility of word choice. We also investigated the token distribution shift to better understand how reinforcement learning with different rewards influences the behavior of the SFT model from which it was initiated. Our analysis revealed systematic modifications in model outputs when using carefully balanced word- and sentence-level rewards, compared to traditional readability measures, which inevitably involve more aggressive rewards for reducing sentence length.

We considered the limitations of using average sentence length as a reward, as it can be a rather coarse measurement of syntactic complexity and perhaps also of cognitive comprehensibility. We observed that it is not uncommon for restrictive attributive clauses and other syntactic structures to appear in the generations from RLAM-trained models that could theoretically be split into separate sentences. However, increasing the coefficient

for the sentence length reward exacerbated optimization instability, as was observed in the RLARI models. We also recognize that average sentence length is a blunt proxy for syntactic complexity. Future work may benefit from integrating more principled syntactic metrics like dependency length (Futrell et al., 2015), which could offer more nuanced incentives for producing cognitively accessible sentence structures.

Another potential direction is to generate lay summaries from the full text, not just the abstract. With current large language models capable of handling contexts with several thousand tokens, potentially enhanced by rotary positional encoding (Su et al., 2024), this approach is feasible. This method also has the added benefit of reducing hallucinations by providing more comprehensive information from the entire manuscript. Also, reinforcement learning from accessibility measures is not limited to scientific literature; it can also be an effective solution for other simplification scenarios where a gold-standard corpus is unavailable. We hope this work contributes to bridging the gap between scholarship and a broader audience, advancing the understanding and development of better simplification systems, and ultimately fostering a more informed and engaged society.

Appendix A: Reinforcement learning using PPO

We fine-tune the policy model in an actor-critic manner: while the policy model (the actor) generates sequences of tokens based on the current sequence s_t , the critic is an additional linear layer that takes the output of the language model’s last layer and produces a scalar for time step t , estimating the expected cumulative reward of producing the token a_t , noted as $V_\rho(a_t)$. The problem is reduced to maximizing at every step the expected cumulative reward of taking action a_t in state s_t and following the policy π_θ thereafter ($V_\rho(s_{t+1}) = V_\rho(s_t, a_t)$), compared to the expected cumulative reward of being in state s ($V_\rho(s_t)$), termed as the advantage $d_t = V_\rho(s_{t+1}) - V_\rho(s_t)$.

The final reward of the entire generation (see Sect. “[Reward function](#)”) is back-propagated through the sequence using Temporal Difference (TD) and Generalized Advantage Estimation (GAE) to estimate the advantage of each token:

$$d_t = \sum_{t'=t}^T (\gamma \lambda)^{t'-t} (r_{t'} + \gamma V_\rho(s_{t'+1}) - V_\rho(s_{t'})) \tag{3}$$

where γ is the discount factor for future rewards, λ controls the bias-variance trade-off, and r_T is the final reward, which, in our case, is a linear combination of two accessibility measures. V_ρ is trained via minimizing a square-error loss (see Eq. 5).

We used the PPO (Schulman et al., 2017) clipped surrogate objective with importance sampling to more efficiently use offline samples to update the online policy. Importance sampling corrects for the discrepancy between the behavior policy that generated the samples and the current policy by weighting the samples using the ratio of their probabilities under both policies. The PPO algorithm introduces a clipping mechanism to balance exploration and exploitation while preventing large, potentially harmful updates to the policy (see Eq. 4).

The entire RLAM algorithm is illustrated in Algorithm 1.

Algorithm 1 Training with reinforcement learning from uncombined accessibility measures. The policy model is updated using the PPO clipped surrogate objective (Eq. 4), and

the value model is updated by minimizing a square-error objective (Eq. 5).

```

1: Input: initial policy model  $\pi_{\theta_{\text{SFT}}}$ , randomly initiated value head  $V_{\rho_{\text{init}}}$ , reward function  $R$  for the last token, weighted
   by  $\beta_{\text{KL}}$  for KL divergence term; task prompts  $X$ ; hyperparameters  $\gamma, \lambda, \epsilon$ 
2:  $\pi_{\theta} \leftarrow \pi_{\theta_{\text{SFT}}}, V_{\rho} \leftarrow V_{\rho_{\text{init}}}$ 
3: for step = 1, ...,  $M$  do
4:   Sample a batch  $\{s_0\}^n$  from  $X$ 
5:   Sample output sequences  $\{a_0, a_1, \dots, a_{T-1}\}^n \sim \pi_{\theta}(\cdot | s_0)$  for each prompt  $s_0$  in the batch ▷ Eq. 1
6:   Compute reward  $r_T^n$  for each sampled sequence  $\{a_0, a_1, \dots, a_{T-1}\}^n$  using reward function  $R$  ▷ Sec. 4.2
7:   Distribute the final reward  $r_T^n$  to each token in the sequence through GAE ▷ Eq. 3
8:   Compute advantages  $\{d_t | s_t\}_{t=0}^{T-1}$ , value targets  $\{V_{\text{targ}}(s_t)\}_{t=0}^{T-1}$  for each sequence with  $V_{\rho}$  and compute KL
   divergence penalty  $\text{KL}_t = \text{KL}(\pi_{\theta}(a_t | s_t) || \pi_{\theta_{\text{SFT}}}(a_t | s_t))$ 
9:   for PPO iteration = 1, ...,  $\mu$  do
10:    Update the policy model by maximizing the PPO clipped surrogate objective with KL penalty:
        
$$\theta \leftarrow \arg \max_{\theta} \frac{1}{n} \sum_{n=1}^n \frac{1}{T} \sum_{t=1}^T \min \left( \frac{\pi_{\theta_{\text{online}}}(a_t | s_t)}{\pi_{\theta_{\text{offline}}}(a_t | s_t)} (A_t - \beta_{\text{KL}} \text{KL}_t), \right.$$

(4)
        
$$\left. \text{clip} \left( \frac{\pi_{\theta_{\text{online}}}(a_t | s_t)}{\pi_{\theta_{\text{offline}}}(a_t | s_t)}, 1 - \epsilon, 1 + \epsilon \right) (A_t - \beta_{\text{KL}} \text{KL}_t) \right)$$

11:   end for
12:   Update the value model by minimizing a square-error objective:
        
$$\rho \leftarrow \arg \min_{\rho} \frac{1}{n} \sum_{n=1}^n \frac{1}{T} \sum_{t=1}^T (V_{\rho}(s_t) - V_{\text{targ}}(s_t))^2$$

(5)
13: end for
14: Output:  $\pi_{\theta}$ 

```

Appendix B: Adaptive KL controller

Following (Ziegler et al., 2019, Sect. 2.2), we dynamically adjust β_{KL} to target a specific KL divergence value, $\text{KL}_{\text{target}}$, using a log-space proportional controller.

The update rule is:

$$\beta_{\text{KL}_{t+1}} = \beta_{\text{KL}_t} \left(1 + K_{\beta} \cdot \text{clip} \left(\frac{\text{KL}(\pi_{\theta}, \pi_{\theta_{\text{SFT}}}) - \text{KL}_{\text{target}}}{\text{KL}_{\text{target}}}, -0.2, 0.2 \right) \right)$$

where K_{β} is the proportional gain, set to 0.01.

Appendix C: Manual evaluation

The evaluation involved three annotators: Jason Clark (academic librarian), Zuoyu Tian (computational linguist), and Haining Wang (information scientist). The following rubrics were iteratively developed and refined over two months of model monitoring and interim generation reviews.

- Language Quality: Evaluates the fluency and grammatical accuracy of the text.
 - Good: Natural and free from grammatical errors; expressions flow properly.
 - Acceptable: Natural and error-free, although it may contain formulaic endings (e.g., “All rights reserved.”) that have minimal impact on readability.

- Poor: Contains grammatical errors.
- Completeness: Assesses the coverage of the abstract’s main components, including context, purpose, main findings, and conclusions.
 - Good: Fully covers all four aspects without truncation.
 - Acceptable: Covers at least three aspects, though the generation may be truncated.
 - Poor: Covers fewer than three of the main aspects.
- Faithfulness: Assesses the scientific fidelity of the simplified text to the original abstract.
 - Good: Accurately reflects the main claims of the original abstract, with no factual errors.
 - Acceptable: Maintains semantic alignment, though the claims remain unverifiable after 15 min of manual searching without AI assistance. No factual errors are present.
 - Poor: Misrepresents findings, contains false claims, or includes factual errors.

Each annotator independently reviewed the selected samples using the finalized rubrics before beginning formal annotation. In the formal rating phase, discrepancies were resolved through discussion, referencing the original manuscripts and external sources as needed. While our evaluations of language quality and completeness were generally consistent, assessing faithfulness proved challenging given the broad disciplinary diversity of the source abstracts. For example, the original sentence states: “At high concentrations, Zn^{2+} , in addition, binds to a second site and inhibits the outward movement of the voltage sensor of Hv1.” A simplified version reads: At high levels, it also prevents the channel’s parts from moving.” While the simplified version is notably more accessible and preserves the core causal relationship, it remains unclear—without specialized domain expertise—whether the omission of “outward” specifies a subtle factual distortion (e.g., whether inward movement is also implicated). In these instances, we prioritized semantic alignment for a general audience. In retrospect, we acknowledge that incorporating inter-rater agreement metrics, such as Cohen’s Kappa, for language quality and completeness would have enhanced our methodological rigor.

We have made all generations, including intermediate outputs from the RL training process, publicly available in our project repository. We encourage readers and domain experts to evaluate the faithfulness of our outputs in light of their disciplinary expertise.

Acknowledgements We gratefully acknowledge the support of the Institute of Museum and Library Services (No. RE-246450-OLS-20) and the National Social Science Fund of China (No. 23 & ZD221). We also thank the organizers of the LIS Education and Data Science Integrated Network Group (LEADING), including Jane Greenberg, Erik Mitchell, Kenning Arlitsch, Jonathan Wheeler, and Samantha Grabus. Additionally, we are thankful to Coltran Hophan-Nichols and Alexander Salois from the University Information Technology Research Cyberinfrastructure at Montana State University for providing computational resources on the Tempest High Performance Computing System, Doralyn Rossmann for research support, Michael Hartwell for proofreading, Jacob Striebel for discussions on improving sentence-level rewards, and Deanna Zarrillo for early involvement in the project.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory

regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Al-Thanyyan, S. S., & Azmi, A. M. (2021). Automated text simplification: A survey. *ACM Computing Surveys*, *54*(2), 1–36.
- Alva-Manchego, F., Scarton, C., & Specia, L. (2020). Data-driven sentence simplification: Survey and benchmark. *Computational Linguistics*, *46*(1), 135–187.
- Berenbaum, M. R. (2021). On COVID-19, cognitive bias, and open access. *Proceedings of the National Academy of Sciences*, *118*(2), e2026319118.
- Brossard, D., & Lewenstein, B. V. (2009). A critical appraisal of models of public understanding of science: Using practice to inform theory. In *Communicating Science* (pp. 25–53). Routledge.
- Burns, T. W., O'Connor, D. J., & Stocklmayer, S. M. (2003). Science communication: A contemporary definition. *Public Understanding of Science*, *12*(2), 183–202.
- Calleja, N., AbdAllah, A., Abad, N., Ahmed, N., Albarracin, D., Altieri, E., Anoko, J. N., Arcos, R., Azlan, A. A., Bayer, J., Bechmann, A., Bezbaruah, S., Briand, S. C., Brooks, I., Bucci, L. M., Burzo, S., Czerniak, C., De Domenico, M., Dunn, A. G., ... Purnat, T. D. (2021). A public health research agenda for managing infodemics: Methods and results of the first WHO infodemiology conference. *JMIR Infodemiology*, *1*(1), e30979.
- Carlson, K., Riddell, A., & Rockmore, D. (2018). Evaluating prose style transfer with the Bible. *Royal Society Open Science*, *5*(10), 171920.
- Chandrasekar, R., Doran, C., & Srinivas, B. (1996). Motivations and methods for text simplification. In *COLING 1996 volume 2: The 16th international conference on computational linguistics*.
- Davies, S. R. (2008). Constructing communication: Talking to scientists about talking to the public. *Science Communication*, *29*(4), 413–434.
- De Belder, J., & Moens, M.-F. (2010). Text simplification for children. In *Proceedings of the SIGIR workshop on accessible search systems* (pp. 19–26).
- Devaraj, A., Marshall, I., Wallace, B., & Li, J. J. (2021). Paragraph-level simplification of medical texts. In K. Toutanova, A. Rumshisky, L. Zettlemoyer, D. Hakkani-Tur, I. Beltagy, S. Bethard, R. Cotterell, T. Chakraborty, & Y. Zhou (Eds.), *Proceedings of the 2021 conference of the North American chapter of the association for computational linguistics: Human language technologies* (pp. 4972–4984). Association for Computational Linguistics.
- DuBay, W. H. (2004). The principles of readability. Technical report, Impact Information, Costa Mesa, CA.
- Fecher, B., & Friesike, S. (2014). Open science: One term, five schools of thought. *Opening Science* (pp. 17–47). Springer.
- Flesch, R. (1946). *The art of plain talk* (1st ed.). Harper & Row.
- Futrell, R., Mahowald, K., & Gibson, E. (2015). Large-scale evidence of dependency length minimization in 37 languages. *Proceedings of the National Academy of Sciences*, *112*(33), 10336–10341.
- Goldsack, T., Zhang, Z., Lin, C., & Scarton, C. (2022). Making science simple: Corpora for the lay summarisation of scientific literature. arXiv preprint [arXiv:2210.09932](https://arxiv.org/abs/2210.09932)
- Grene, M., Cleary, Y., & Marcus-Quinn, A. (2017). Use of plain-language guidelines to promote health literacy. *IEEE Transactions on Professional Communication*, *60*(4), 384–400.
- Hara, N., Abbazio, J., & Perkins, K. (2019). An emerging form of public engagement with science: Ask me anything (AMA) sessions on Reddit r/science. *PLoS ONE*, *14*(5), e0216789.
- Hetland, P. (2014). Models in science communication: Formatting public engagement and expertise. *Nordic Journal of Science and Technology Studies*, *2*(2), 5–17.
- Hijazi, R., Espinasse, B., & Gala, N. (2022). GRASS: A syntactic text simplification system based on semantic representations. In *11th international conference on natural language processing*.
- Huang, S., Noukhovitch, M., Hosseini, A., Rasul, K., Wang, W., & Tunstall, L. (2024). The N+ implementation details of RLHF with PPO: A case study on TL; DR summarization. arXiv preprint [arXiv:2403.17031](https://arxiv.org/abs/2403.17031)
- Isaksson, F. (2018). Is simple wikipedia simple? A study of readability and guidelines. Bachelor's thesis, Linköping University. 18 ECTS, Cognitive Science, LIU-IDA/KOGVET-G–18/029–SE, Supervisor: Arne Jönsson, Examiner: Henrik Danielsson.
- Islam, A. N., Laato, S., Talukder, S., & Sutinen, E. (2020). Misinformation sharing and social media fatigue during COVID-19: An affordance and cognitive load perspective. *Technological Forecasting and Social Change*, *159*, 120201.

- Joly, P.-B., & Kaufmann, A. (2008). Lost in translation? The need for ‘upstream engagement’ with nano-technology on trial. *Science as Culture*, 17(3), 225–247.
- Joseph, S., Kazanas, K., Reina, K., Ramanathan, V., Xu, W., Wallace, B., & Li, J. J. (2023). Multilingual simplification of medical texts. In H. Bouamor, J. Pino, & K. Bali (Eds.), *Proceedings of the 2023 conference on empirical methods in natural language processing* (pp. 16662–16692). Association for Computational Linguistics.
- Kerwer, M., Chasiotis, A., Stricker, J., Günther, A., & Rosman, T. (2021). Straight from the scientist’s mouth—plain language summaries promote laypeople’s comprehension and knowledge acquisition when reading about individual research findings in psychology. *Collabra: Psychology*, 7(1), 18898.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In Y. Bengio & Y. LeCun (Eds.), *3rd international conference on learning representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, conference track proceedings*.
- Knox, E., & Hara, N. (2021). Public engagement with science via social media: A case of communicating the pandemic on Twitter. *Proceedings of the Association for Information Science and Technology*, 58(1), 759–761.
- Landrum, J., Gao, L., Jiang, Z., Hara, N., & Liu, X. (2014). Scaffolding of scientific publications with open educational resources (OER). *Proceedings of the American Society for Information Science and Technology*, 51(1), 1–5.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., & Zettlemoyer, L. (2019). BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. arXiv preprint [arXiv:1910.13461](https://arxiv.org/abs/1910.13461).
- Lin, B. Y., Ravichander, A., Lu, X., Dziri, N., Sclar, M., Chandu, K., Bhagavatula, C., & Choi, Y. (2023). The unlocking spell on base LLMs: Rethinking alignment via in-context learning. In *The twelfth international conference on learning representations*.
- Loshchilov, I. & Hutter, F. (2019). Decoupled weight decay regularization. In *International conference on learning representations*.
- Maurer, M., Siegel, J. E., Firminger, K. B., Lowers, J., Dutta, T., & Chang, J. S. (2021). Lessons learned from developing plain language summaries of research studies. *HLRP: Health Literacy Research and Practice*, 5(2), e155–e161.
- Mazur, B. (2000). Revisiting plain language. *Technical Communication*, 47(2), 205.
- Mesnard, T., Hardin, C., Dadashi, R., Bhupatiraju, S., Pathak, S., Sifre, L., Rivière, M., Kale, M. S., Love, J., et al. (2024). Gemma: Open models based on Gemini research and technology. arXiv preprint [arXiv:2403.08295](https://arxiv.org/abs/2403.08295)
- Metcalfe, J. (2022). Science communication: A messy conundrum of practice, research and theory. *Journal of Science Communication*, 21(7), C07.
- Narayan, S., Cohen, S. B., & Lapata, M. (2018). Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. In E. Riloff, D. Chiang, J. Hockenmaier, & J. Tsujii (Eds.), *Proceedings of the 2018 conference on empirical methods in natural language processing* (pp. 1797–1807). Association for Computational Linguistics.
- Nishihara, D., Kajiwara, T., & Arase, Y. (2019). Controllable text simplification with lexical constraint loss. In F. Alva-Manchego, E. Choi, & D. Khashabi (Eds.), *Proceedings of the 57th annual meeting of the association for computational linguistics: Student research workshop* (pp. 260–266). Association for Computational Linguistics.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Gray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, & A. Oh (Eds.), *Advances in neural information processing systems 35* (pp. 27730–27744). Curran Associates.
- Papineni, K., Roukos, S., Ward, T., & Zhu, W.-J. (2002). BLEU: A method for automatic evaluation of machine translation. In P. Isabelle, E. Charniak, & D. Lin (Eds.), *Proceedings of the 40th annual meeting of the association for computational linguistics* (pp. 311–318). Association for Computational Linguistics.
- Petelin, R. (2010). Considering plain language: Issues and initiatives. *Corporate Communications: An International Journal*, 15(2), 205–216.
- Plavén-Sigray, P., Matheson, G. J., Schiffler, B. C., & Thompson, W. H. (2017). Research: The readability of scientific texts is decreasing over time. *eLife*, 6, e27725.
- Pool, J., Fatehi, F., & Akhlaghpour, S. (2021). Infodemic, misinformation and disinformation in pandemics: Scientific landscape and the road ahead for public health informatics research. In *Public health and informatics* (pp. 764–768). IOS Press.

- Ramamurthy, R., Ammanabrolu, P., Brantley, K., Hessel, J., Sifa, R., Bauckhage, C., Hajishirzi, H., & Choi, Y. (2023). Is reinforcement learning (not) for natural language processing: Benchmarks, baselines, and building blocks for natural language policy optimization. In *The eleventh international conference on learning representations*.
- Riddell, A., & Igarashi, Y. (2021). Varieties of plain language. In R. Mitkov & G. Angelova (Eds.), *Proceedings of the international conference on recent advances in natural language processing (RANLP 2021)* (pp. 1180–1187).
- Rothwell, J. (2020, September 8). *Assessing the economic gains of eradicating illiteracy nationally and regionally in the United States*. Barbara Bush Foundation for Family Literacy.
- Schriver, K. A. (2017). Plain language in the US gains momentum: 1940–2015. *IEEE Transactions on Professional Communication*, 60(4), 343–383.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. arXiv preprint [arXiv:1707.06347](https://arxiv.org/abs/1707.06347)
- Sokolowska, N., Fecher, B., & Wagner, G. G. (2019). Communication on the science-policy interface: An overview of conceptual models. *Publications*, 7(4), 64.
- Stiennon, N., Ouyang, L., Wu, J., Ziegler, D., Lowe, R., Voss, C., Radford, A., Amodei, D., & Christiano, P. F. (2020). Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33, 3008–3021.
- Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., & Liu, Y. (2024). RoFormer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568, 127063.
- Trench, B. (2008). Towards an analytical framework of science communication models. In *Communicating science in social contexts: New models, new practices* (pp. 119–135). Springer.
- Wang, H. & Clark, J. A. (2025). Simplifying scholarly abstracts for accessible digital libraries using language models. In *Proceedings of the 24th ACM/IEEE joint conference on digital libraries*. Association for Computing Machinery.
- Wang, Y., McKee, M., Torbica, A., & Stuckler, D. (2019). Systematic literature review on the spread of health-related misinformation on social media. *Social Science & Medicine*, 240, 112552.
- Wolfe, J. (2008). Annotations and the collaborative digital library: Effects of an aligned annotation interface on student argumentation and reading strategies. *International Journal of Computer-Supported Collaborative Learning*, 3, 141–164.
- Wu, Z., Hu, Y., Shi, W., Dziri, N., Suhr, A., Ammanabrolu, P., Smith, N. A., Ostendorf, M., & Hajishirzi, H. (2023). Fine-grained human feedback gives better rewards for language model training. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, & S. Levine (Eds.), *Advances in neural information processing systems 36*. Curran Associates.
- Xu, W., Callison-Burch, C., & Napoles, C. (2015). Problems in current text simplification research: New data can help. *Transactions of the Association for Computational Linguistics*, 3, 283–297.
- Xu, W., Napoles, C., Pavlick, E., Chen, Q., & Callison-Burch, C. (2016). Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics*, 4, 401–415.
- Yanamoto, D., Ikawa, T., Kajiwara, T., Ninomiya, T., Uchida, S., & Arase, Y. (2022). Controllable text simplification with deep reinforcement learning. In Y. He, H. Ji, S. Li, Y. Liu, & C.-H. Chang (Eds.), *Proceedings of the 2nd conference of the asia-pacific chapter of the association for computational linguistics and the 12th international joint conference on natural language processing (volume 2: Short papers)* (pp. 398–404). Association for Computational Linguistics.
- Zhang, T., Kishore, V., Wu, F., Weinberger, K. Q., & Artzi, Y. (2020). BERTScore: Evaluating text generation with bert. In *International conference on learning representations*.
- Zhao, S., Meng, R., He, D., Saptono, A., & Parmanto, B. (2018). Integrating transformer and paraphrase rules for sentence simplification. In E. Riloff, D. Chiang, J. Hockenmaier, & J. Tsujii (Eds.), *Proceedings of the 2018 conference on empirical methods in natural language processing* (pp. 3164–3173). Association for Computational Linguistics.
- Ziegler, D. M., Stiennon, N., Wu, J., Brown, T. B., Radford, A., Amodei, D., Christiano, P., & Irving, G. (2019). Fine-tuning language models from human preferences. arXiv preprint [arXiv:1909.08593](https://arxiv.org/abs/1909.08593)

Authors and Affiliations

Haining Wang¹  · Jason Clark² · Hannah McKelvey² · Leila Sterman² · Zheng Gao³ · Zuoyu Tian⁴ · Sandra Kübler¹ · Xiaozhong Liu⁵

✉ Haining Wang
hw56@iu.edu

Jason Clark
jaclark@montana.edu

Hannah McKelvey
hannah.mckelvey@montana.edu

Leila Sterman
leila.sterman@montana.edu

Zheng Gao
zhgao3@coupang.com

Zuoyu Tian
ztian@macalester.edu

Sandra Kübler
skuebler@iu.edu

Xiaozhong Liu
xliu14@wpi.edu

¹ Indiana University, Indianapolis, IN, USA

² Montana State University, Bozeman, MT, USA

³ Coupang, Mountain View, CA, USA

⁴ Macalester College, Saint Paul, MN, USA

⁵ Worcester Polytechnic Institute, Worcester, MA, USA