Considering Human Interaction and Variability in Automatic Text Simplification

Jenia Kim[1], Stefan Leijnen[1], Lisa Beinborn[2]

[1]HU University of Applied Sciences Utrecht, Netherlands
[2]University of Göttingen, Germany
Correspondence: jenia.kim@hu.nl

Abstract

Research into automatic text simplification aims to promote access to information for all members of society. To facilitate generalizability, simplification research often abstracts away from specific use cases, and targets a prototypical reader and an underspecified content creator. In this paper, we consider a real-world use case – simplification technology for use in Dutch municipalities – and identify the needs of the content creators and the target audiences in this scenario. The stakeholders envision a system that (a) assists the human writer without taking over the task; (b) provides diverse outputs, tailored for specific target audiences; and (c) explains the suggestions that it outputs. These requirements call for technology that is characterized by *modularity*, *explainability*, and *variability*. We argue that these are important research directions that require further exploration.

1 Introduction

Full participation in modern society requires reading and understanding a wide variety of written information. For example, drawing the right conclusion from a letter from the tax authority, or from instructions about how to apply for unemployment benefits, is crucial for active citizenship. Unfortunately, not everybody is equally skilled at reading. In the Netherlands, for example, about 2.5 million adults (one in six adults) have limited literacy, i.e., difficulty with reading and/or writing (Netherlands Court of Audit, 2016). To ensure fair access to crucial information for everyone, text simplification research aims to develop technology that can automatically identify sources of complexity in text and generate simplifications.

Simplification research often abstracts from specific use cases to facilitate the generalizability of the developed methods. Curated datasets and evaluation setups tend to target a prototypical reader and an underspecified content creator. In practice, however, technology does not exist in a vacuum; it is always interconnected with people. As users engage with technology, they gradually develop a mental model of its functioning, which subsequently shapes their further interaction and engagement (e.g., Baxter and Sommerville, 2011; Lee et al., 2024). Therefore, technology that does not meet the needs of its intended users and their preferences regarding the outputs and the interaction might result in unsuccessful deployment.

Text simplification is at its core a human-centered problem; it operates on a text generated by human writers and reduces its complexity for the sake of human readers. In this paper, we discuss how the preferences of the intended writers and the characteristics of the intended readers shape the properties of the required simplification technology. We do so by exploring a real-world use case: a simplification system that is meant to assist content creators in Dutch municipalities with writing accessible text.

2 Use Case Description

The public sector in the Netherlands is committed to promoting inclusive and accessible communication. For example, the City of Amsterdam published writing guidelines that instruct the employees to use "clear language" in all their written communication, and "simple language" in communication that targets audiences with limited literacy. Unfortunately, these efforts have proven to be insufficient. A recent study (Corsius et al., 2022) that evaluated 240 texts from 70 Dutch government organizations found that the texts – which discussed crucial information about payments and healthcare – were not understandable enough, due to lexical complexity, vague or indirect style, and the length of the text, among other factors.

The civil servants themselves indicate that implementing the guidelines in practice is difficult.

Proceedings of the Third Workshop on Text Simplification, Accessibility and Readability (TSAR 2024), pages 52–60

The target levels in the guidelines are described in terms of the Common European Framework of Reference for Languages (CEFR); B1-level is considered "clear language", and A2-level is considered "simple language". However, in a workshop conducted by the City of Amsterdam in 2022 (Pinhão and Gornishka, 2022), the participants indicated that it is not straightforward to understand what the A2/B1-level requirements mean in practice, and that they lack this expertise in the organization.

Furthermore, they indicated that it is challenging to write to a broad audience (the residents of Amsterdam) that consists of diverse groups with different linguistic needs. Among the discussed possible solutions, the participants mentioned that they would benefit from a tool that could review what they wrote, highlight potential difficulties, and provide suggestions on how to solve them.

In interviews conducted with representatives of other municipalities in the Netherlands, the interest in automated solutions surfaced as well.[1] When asked about their needs and concerns regarding the introduction of such technologies, the interviewees emphasized the need of the writers to remain in control of the text, mainly because of the concern that automated simplification might result in changes in meaning and loss of nuance.

To summarize, there are three main points raised by the stakeholders. First, the target audience consists of diverse groups with different linguistic needs, so there is no "one-fits-all" solution. Second, the technology is viewed as a source of knowledge about these diverse linguistic needs and how to accommodate them; by using the technology, writers expect to improve their own expertise on the subject. Third, the writers wish to remain in control of the task and to take responsibility for the final output. In other words, they envision a system that (a) assists the human writer without taking over the task; (b) provides diverse outputs, tailored for specific target audiences; and (c) explains the suggestions that it outputs.

3 Automated Text Generation vs. Human-AI Co-Creation

Automated generation of simplified text is not the type of technology that is envisioned by the stakeholders in our use case. The content creators do not want the technology to take over the simplification task; rather, they want to collaborate with the AI, while remaining in control of the writing task and its outputs. This type of assemblage, where a human and an AI algorithm work together on a creative task, is called a *co-creation system* (Allen et al., 1999; Lubart, 2005; Zhu et al., 2018; Guzdial and Riedl, 2019).

Co-creation systems aim to support the endeavor of writing accessible text, while maintaining human agency, control, and ownership. Writing is a creative activity, which people find meaningful and satisfying; this holds not only for creative, but also for professional writing (Brand and Leckie, 1988). Moreover, writers often have a strong personal connection and a feeling of ownership towards the work they produce (e.g., Nicholes, 2017). A successful human-AI collaboration should, therefore, aim to preserve the meaningfulness of the task for the human writers, and their sense of ownership, agency, and control (e.g., Zhou and Sterman, 2023; Biermann et al., 2022).

In addition to preserving the sense of meaningfulness and satisfaction for individual writers, co-creation systems may have benefits on the organizational and societal levels as well. Within organizations, over-reliance on fully automatically generated simplifications might result in the loss of knowledge and expertise among human employees (Gibbs et al., 2021); co-creation systems, on the other hand, have the potential to increase human expertise, as writers gradually learn from the system's feedback.

On a societal level, use of co-creation systems ensures a clearer allocation of responsibility between the authorities and the citizen. For example, an existing application endorsed by the Dutch government (Rijksoverheid, 2023) allows people to scan formal letters (e.g., from the tax authority) with their phone camera and instantly receive a simplified version of them. This type of technology places the responsibility for understanding the letter on the citizen, rather than on the tax authority. It has been shown that this expectation for self-reliance is detrimental for the ability of many citizens, especially from marginalized groups, to realize their basic human rights (Netherlands Institute for Human Rights, 2020). Use of co-creation systems, on the other hand, leaves the responsibility for accessibility and social inclusion with the institutions who create the content, instead of passing it on to the citizen.

[1]The interviews were conducted as part of the preliminary phase of the project *Duidelijke TAAL* (Clear Language), funded by the Dutch National Organisation for Practice-Oriented Research SIA; file number RAAK.PUB13.022.

4 One-Fits-All vs. Heterogeneous Audiences

The simplification technology in our use case targets all (adult) residents of the city, which corresponds to an extremely diverse audience, including people with various education backgrounds, people with cognitive disabilities or learning disorders, non-native speakers, etc.

One possible approach to the heterogeneity of the audience is to write the simplest possible version, which can be understood by (almost) everybody. This has been the dominant approach in the public sector in the Netherlands in the last decade; the guideline was to write on CEFR B1-level, with the assumption that this level is understandable to 95% of the Dutch population (Jansen, 2013). However, this approach has two main limitations.

First, simplifying to the lowest level possible necessarily involves some loss of meaning, or at least loss of nuance. This is not suitable for all contexts, since some communication requires a high degree of semantic precision. For example, Garimella et al. (2022) study simplification of legal text and find significant disagreement between legal experts on the required level of detail.

Second, using CEFR levels as a target is controversial (e.g., Jansen, 2013). CEFR is not a readability metric that is meant to evaluate text complexity; rather, it is meant to evaluate the skills of the learner.[2] These skills are not directly transferable to specific linguistic features of the text; in fact, to assign a CEFR level to a given sentence requires high level of expertise and experience in foreign language teaching (Arase et al., 2022). Furthermore, as this framework was created specifically for foreign language learners, it is unclear whether it is appropriate for other target groups, like low literate native speakers. It has been shown that readability needs to be measured differently for L1 and L2 readers (Beinborn et al., 2014).

As an alternative to the "one-fits-all" approach, one could aim at accommodating the diverse needs of the audience by customizing the outputs to different groups. This was the preferred direction mentioned by the participants of the workshop in our use case, who envisioned creating multiple versions of the same document or webpage, from which the readers can choose (Pinhão and Gornishka, 2022). To accomplish this, the particular linguistic requirements of different groups need to be identified. Furthermore, novel technical solutions need to be developed that would allow a more personalized government communication, in which the right type of content reaches each citizen. These challenges are discussed further in Section 6.

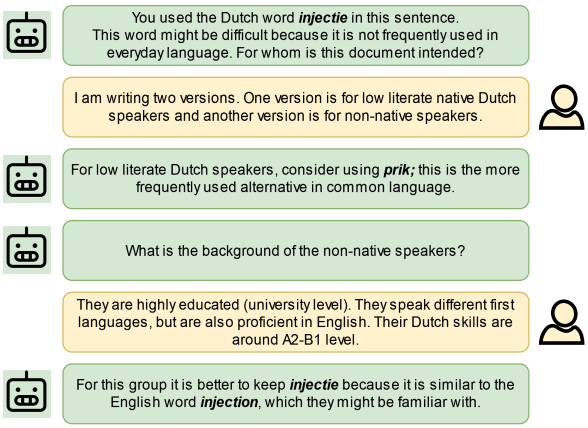

Figure 1: An imaginary co-creation assistant for accessible text, based on the requirements of our use case. It provides modular suggestions, accompanied by explanations, and tailored for various audiences.

5 Explainability, Modularity, Variability

The stakeholders in our use case envision a co-creation system that (a) assists the human writer without taking over the task; (b) provides diverse outputs, tailored for specific target audiences; and (c) explains the suggestions that it outputs. Figure 1 shows an imaginary example of an interaction that fulfills these requirements.

This entails certain characteristics of the underlying technology. First, simplification in this use case cannot be formulated as an end-to-end operation, i.e., rephrasing of a sentence to a simpler version. Rather, the output of the system has to be **modular**; it should suggest specific simplification operations (e.g., lexical substitution in Figure 1), leaving the decision which ones to accept and how to combine them in the hands of the human writer.

Second, the model needs to be able to generate **variable** simplified outputs for the same complexity. This is necessary for two purposes. First, it allows adaptability to different target audiences; e.g., in Figure 1, the system outputs two different synonyms, each of which is tailored to a different target group. Second, it allows adaptability to different writers as well. Simplification is not a well-

[2]For example, B1-level reading skills mean that the learner *"can understand texts that consist mainly of high frequency everyday or job-related language [...and..] can understand the description of events, feelings and wishes in personal letters"* (Council of Europe).

defined, closed-ended operation; it can be achieved through various strategies. Simplification strategies differ both on the inter- and intra-expert level: the proposed editing operations might vary across experts but also for an individual expert over different points in time, while the result may be equally acceptable (Xu et al., 2015; Alva-Manchego et al., 2021). A co-creation assistant should therefore be able to suggest various possible operations for the same complexity (e.g., splitting a complex sentence or reordering its parts); the human writer can choose the most suitable operation, according to their own style and preferences.

Third, the model's outputs and suggestions need to be **explainable**; i.e., they need to be motivated by expert knowledge (e.g., in Figure 1, the importance of cognates for foreign language readability). Our definition of explainability goes beyond visualizing which elements contributed to a model' complexity prediction based on post-hoc attribution methods (Garbacea et al., 2021; Hobo et al., 2023). While this can be an important first step, we envision explanations that one would expect to receive from a human expert. They should provide insights on why certain phenomena cause comprehension difficulties (for certain audiences), and how the suggestion reduces the complexity.

6 Discussion

In this section, we discuss how the requirements described above fit into existing research, and sketch promising directions for future work.

Human-Computer Interaction

Co-creation systems, like the one envisioned in our use case, are extensively researched by the Human-Computer Interaction (HCI) community. This field focuses on interfaces between people and technologies, putting in the center the experience of the users during the interaction, which is explored through user studies.

In the context of simplification, human-computer interaction was explored in reading-assistance systems, i.e., with readers as the intended users. For example, Rello et al. (2013) conducted user studies with people with dyslexia, and Alonzo et al. (2022) studied the preferences of deaf and hard-of-hearing individuals, who use simplification technology in the context of their work.

In our use case, on the other hand, the intended users of the technology are not the readers, but the writers. We therefore build on human-centered research on *interactive writing assistants*: a co-creation tool that assists people with improving the quality and effectiveness of their writing (e.g., Du et al., 2022). In a recent study, Lee et al. (2024) systematically review 115 articles about interactive writing assistants, and create a comprehensive taxonomy of the aspects that play a role in their design. This taxonomy can be a good starting point for a structured exploration of co-creation systems for text simplification. It is important that different use cases are described in a methodical way, and that design decisions for specific scenarios are grounded in user studies.

Personalized Simplification

The civil servants in our use case perceive their heterogeneous audience as a collection of different groups, with diverse linguistic needs. Indeed, research has shown that different target groups have different readability and simplification requirements; for example, people with dyslexia benefit more from seeing a number of synonyms for a complex word, rather than one simple synonym (Rello et al., 2013).

However, recent studies indicate that the perception of complexity varies on individual level, rather than group level. For example, Gooding and Tragut (2022) show that the judgments of non-native English speakers regarding lexical complexity depend not only on their proficiency in English, but also on (a combination of) idiosyncratic characteristics, like the reader's first language and reading experience. To address this, adaptive and personalized models can be created, which obtain individual data from users and learn user-specific simplifications (e.g., Bingel et al., 2018; Gooding and Tragut, 2022).

In the context of our use case, there are a few potential issues with individual-level personalization that need to be considered. First, from an ethical viewpoint, collecting individual data and training personalized models in the municipality context can be viewed as an infringement on the citizen's privacy. Second, such solutions would necessarily involve digital interfaces (e.g., websites, apps), which are not easily accessible for some groups in the Dutch population; in fact, the same vulnerable populations (e.g., people with lower education levels) often have difficulties both with complex texts and with digital literacy (Netherlands Institute for Human Rights, 2020). Lastly, this approach involves less human oversight and control over the

simplified content, compared to the co-creation setup. Therefore, we focus on personalization on a group level, which can be performed by the writers and does not require interactive input from the readers. On digital interfaces, readers could choose the desired version of the text themselves (similarly to the choice of language on websites). However, ensuring that the right content reaches every citizen in offline communication remains an open challenge.

Towards Explainability, Modularity, and Variability in Dutch Text Simplification

To the best of our knowledge, no current simplification technology for Dutch fully incorporates explainability, modularity, and variability, as we described them. Commercial writing assistants for text simplification offer a certain degree of modularity and explainability, but remain limited in terms of variability (see Appendix A).

To further advance the research towards explainability, modularity, and variability, a few promising directions can be explored. First, it is crucial to better understand the underlying (psycho-)linguistic and cognitive phenomena that affect text complexity for different target groups. Based on this knowledge, modular and audience-specific simplification operations that address these phenomena can be defined. Work on Dutch readability for specific target groups is limited; for example, Kleijn (2018) explores readability in adolescents (high-school students), Maat and Gravekamp (2022) analyze differences between people with different education levels, and Reichrath and Moonen (2022) study an heterogeneous sample of the residents of Amsterdam, which they divide into two categories of literacy based on the CEFR. We believe that further work in this line is needed, which specifically focuses on identifying differences in the linguistic needs of different groups.

For explainability and variability of outputs, the identified modular and audience-specific operations need to be incorporated into simplification models. It remains an open question how to implement such fine-grained control over the process. For example, for transformer language models, control tokens have been introduced that can modulate specific attributes of the model outputs (Martin et al., 2020). The approach can be applied to achieve audience-specific simplifications; e.g., simplification of English text for people with cognitive disabilities (Chamovitz and Abend, 2022) or simplification of Russian text for foreign language learners with diverse proficiency levels (Dmitrieva, 2023). Seidl and Vandeghinste (2024) apply control tokens to manipulate various lexical and syntactic attributes of Dutch output but do not explicitly connect their approach to specific target audiences.

For the application of control tokens, a parallel corpus of complex-simple sentence pairs in the target language is required to train the model. As no large manually annotated parallel corpora exist for Dutch, Seidl and Vandeghinste (2024) train their model on automatically translated data, and Vlantis et al. (2024) use an English simplification model as an intermediary and automatically translate the input and output sequences. Both approaches suffer from the limited quality of the translation engine.[3]

Large generative language models, like GPT-3.5 and Llama2, can generate simplifications in a few-shot or zero-shot setting. However, it is unclear whether the output can be controlled towards specific operations or target audiences; for example, Farajidizaji et al. (2024) try to steer model outputs towards different readability levels by using different prompts, but with limited success. To achieve explainability in this setup, methods such as chain-of-thought prompting can be explored (Wei et al., 2022; Cohen and Cohen, 2024).

Another promising research direction explores how to align language models outputs with human production variability. Giulianelli et al. (2023) show that text generation models produce lower variability than humans on a text simplification task. Further research into this problem can contribute to more human-like variability in simplification outputs.

7 Conclusion

We argued that considering human-centered aspects is a crucial step in technology development. We discussed a real-world use case, and showed how the type of human-machine interaction envisioned by the content creators, and the variability in the needs of the target readers shape the requirements from the simplification technology. Specifically, they call for technology that is characterized by *modularity*, *explainability*, and *variability*. How this can be achieved, for Dutch as well as other languages, remains an open question which we intend to explore in our future work.

[3]As an alternative to the use of machine translation, a large parallel corpus of synthetic simplification data has been recently generated by prompting ChatGPT (see Appendix B); however, no evaluation of this dataset is publicly available.

8 Lay Summary

In this article, we discuss technology that makes difficult text simpler; it is called "text simplification technology" or "text simplification tools". The goal of this technology is to make written information, like letters and websites, easy to understand for everyone. For example, if somebody is writing an email and uses a difficult word that not many people know, a simplification tool can recognize the difficult word and replace it with an easier word that has the same meaning.

When researchers design simplification tools, they usually try to create a general solution that can be used in many different cases. The problem is that general solutions are not always the best ones. For example, different groups of people need simple text: children, people with dyslexia, immigrants who learn a new language, and others. For each group, different things can be difficult, so there is no general solution that fits everybody.

In our research, we look at a specific case: a simplification tool for municipalities in the Netherlands. The people who work in the municipalities write a lot of important information, like letters about payments and healthcare, or websites about municipal services. These people want to have a tool that can help them write in a way that everybody can understand. It is important to them to stay in control of the writing; this means that the tool should make suggestions about how to improve the text but the final decision is done by the writer. They also want the tool to explain the suggestions that it gives, for example why something is difficult. In addition, they want the tool to provide different suggestions for different groups of readers, according to what each group needs.

The wishes of the writers in the municipalities require a certain type of simplification technology, which does not exist yet. We plan to work on solving this problem in our future research.

References

James E Allen, Curry I Guinn, and Eric Horvtz. 1999. Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications*, 14(5):14–23.

Oliver Alonzo, Lisa Elliot, Becca Dingman, Sooyeon Lee, Akhter Al Amin, and Matt Huenerfauth. 2022. Reading-assistance tools among deaf and hard-of-hearing computing professionals in the us: Their reading experiences, interests and perceptions of social accessibility. *ACM Transactions on Accessible Computing (TACCESS)*, 15(2):1–31.

Fernando Alva-Manchego, Carolina Scarton, and Lucia Specia. 2021. The (un)suitability of automatic evaluation metrics for text simplification. *Computational Linguistics*, 47(4):861–889.

Yuki Arase, Satoru Uchida, and Tomoyuki Kajiwara. 2022. CEFR-based sentence difficulty annotation and assessment. *arXiv preprint arXiv:2210.11766*.

Gordon Baxter and Ian Sommerville. 2011. Socio-technical systems: From design methods to systems engineering. *Interacting with computers*, 23(1):4–17.

Lisa Beinborn, Torsten Zesch, and Iryna Gurevych. 2014. Readability for foreign language learning: The importance of cognates. *ITL-International Journal of Applied Linguistics*, 165(2):136–162.

Oloff C Biermann, Ning F Ma, and Dongwook Yoon. 2022. From tool to companion: Storywriters want AI writers to respect their personal values and writing strategies. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference*, pages 1209–1227.

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.

Alice G Brand and Phoebe A Leckie. 1988. The emotions of professional writers. *The Journal of psychology*, 122(5):421–439.

Mark Breuker. 2022. Cefr labelling and assessment services. In *European Language Grid: A Language Technology Platform for Multilingual Europe*, pages 277–282. Springer International Publishing Cham.

Eytan Chamovitz and Omri Abend. 2022. Cognitive simplification operations improve text simplification. In *Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL)*, pages 241–265, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Cassandra A Cohen and William W Cohen. 2024. Watch your steps: Observable and modular chains of thought. *arXiv preprint arXiv:2409.15359*.

Mischa Corsius, Vera Lange, Yvette Linders, Henk Pander Maat, Els van der Pool, Nina Sangers, Keun Sliedrecht, Wouter Sluis-Thiescheffer, and Charlotte Swart. 2022. Monitor Begrijpelijkheid Overheidsteksten (Monitor Understandability of Government Texts). Accessed: 09-July-2024.

Council of Europe. Self-assessment Grids (CEFR). Accessed: 05-September-2024.

Anna Dmitrieva. 2023. Automatic text simplification of Russian texts using control tokens. In *Proceedings of the 9th Workshop on Slavic Natural Language Processing 2023 (SlavicNLP 2023)*, pages 70–77, Dubrovnik, Croatia. Association for Computational Linguistics.

Wanyu Du, Zae Myung Kim, Vipul Raheja, Dhruv Kumar, and Dongyeop Kang. 2022. Read, revise, repeat: A system demonstration for human-in-the-loop iterative text revision. *arXiv preprint arXiv:2204.03685*.

Asma Farajidizaji, Vatsal Raina, and Mark Gales. 2024. Is it possible to modify text to a target readability level? an initial investigation using zero-shot large language models. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 9325–9339, Torino, Italia. ELRA and ICCL.

Cristina Garbacea, Mengtian Guo, Samuel Carton, and Qiaozhu Mei. 2021. Explainable prediction of text complexity: The missing preliminaries for text simplification. 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)*, pages 1086–1097, Online. Association for Computational Linguistics.

Aparna Garimella, Abhilasha Sancheti, Vinay Aggarwal, Ananya Ganesh, Niyati Chhaya, and Nandakishore Kambhatla. 2022. Text simplification for legal domain: Insights and challenges. In *Proceedings of the Natural Legal Language Processing Workshop 2022*, pages 296–304, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Jennifer L Gibbs, Gavin L Kirkwood, Chengyu Fang, and J Nan Wilkenfeld. 2021. Negotiating agency and control: Theorizing human-machine communication from a structurational perspective. *Human-Machine Communication*, 2:153–171.

Mario Giulianelli, Joris Baan, Wilker Aziz, Raquel Fernández, and Barbara Plank. 2023. What comes next? evaluating uncertainty in neural text generators against human production variability. *arXiv preprint arXiv:2305.11707*.

Sian Gooding and Manuel Tragut. 2022. One size does not fit all: The case for personalised word complexity models. *arXiv preprint arXiv:2205.02564*.

Matthew Guzdial and Mark Riedl. 2019. An interaction framework for studying co-creative AI. *arXiv preprint arXiv:1903.09709*.

Eliza Hobo, Charlotte Pouw, and Lisa Beinborn. 2023. "geen makkie": Interpretable classification and simplification of Dutch text complexity. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 503–517, Toronto, Canada. Association for Computational Linguistics.

Carel Jansen. 2013. Taalniveau B1: de nieuwste kleren van de keizer (Language level B1: the emperor's newest clothes). *Onze Taal*, 82(2):56–57.

Suzanne Kleijn. 2018. *Clozing in on readability: How linguistic features affect and predict text comprehension and on-line processing*. Ph.D. thesis, Utrecht University.

Mina Lee, Katy Ilonka Gero, John Joon Young Chung, Simon Buckingham Shum, Vipul Raheja, Hua Shen, Subhashini Venugopalan, Thiemo Wambsganss, David Zhou, Emad A Alghamdi, et al. 2024. A design space for intelligent and interactive writing assistants. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–35.

Todd Lubart. 2005. How can computers be partners in the creative process: Classification and commentary on the special issue. *International journal of human-computer studies*, 63(4-5):365–369.

Henk Pander Maat and Jet Gravekamp. 2022. Kan een tekst te simpel zijn? hoe lager en hoger opgeleiden oordelen over eenvoudige taal. *Tijdschrift voor Taalbeheersing*, 44(2):62–90.

Louis Martin, Éric de la Clergerie, Benoît Sagot, and Antoine Bordes. 2020. Controllable sentence simplification. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4689–4698, Marseille, France. European Language Resources Association.

Netherlands Court of Audit. 2016. Aanpak van laaggeletterdheid (Tackling low literacy). Accessed: 09-July-2024.

Netherlands Institute for Human Rights. 2020. Iedereen op eigen kracht? Nederlanders over zelfredzaamheid en mensenrechten (Every person for themselves? Dutch people on self-reliance and human rights). Accessed: 09-July-2024.

Justin Nicholes. 2017. Measuring ownership of creative versus academic writing: Implications for interdisciplinary praxis. *Writing in Practice*, 3(1).

Cláudia Pinhão and Iva Gornishka. 2022. Workshop Results: Improving readability of city communications. Accessed: 09-July-2024.

Eva D Poort and Jennifer M Rodd. 2019. A database of dutch–english cognates, interlingual homographs and translation equivalents. *Journal of cognition*, 2(1).

Enid Reichrath and Xavier Moonen. 2022. Assessing the effects of language for all. *Nordic Journal of Linguistics*, 45(2):232–248.

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*, pages 1–10.

Rijksoverheid. 2023. Terugblik Demo Donderdag: Lees Simpel app versimpelt overheidsinformatie. Accessed: 05-September-2024.

Theresa Seidl and Vincent Vandeghinste. 2024. Controllable sentence simplification in dutch. *Computational Linguistics in the Netherlands Journal*, 13:31–61.

Anaïs Tack, Thomas François, Piet Desmet, and Cédrick Fairon. 2018. NT2Lex: A CEFR-graded lexical resource for Dutch as a foreign language linked to open Dutch WordNet. In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 137–146, New Orleans, Louisiana. Association for Computational Linguistics.

Charlotte van de Velde. 2023. Automatic sentence-level simplification for dutch. Master's thesis, KU Leuven.

Daniel Vlantis, Iva Gornishka, and Shuai Wang. 2024. Benchmarking the simplification of Dutch municipal text. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2217–2226, Torino, Italia. ELRA and ICCL.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

Wei Xu, Chris Callison-Burch, and Courtney Napoles. 2015. Problems in current text simplification research: New data can help. *Transactions of the Association for Computational Linguistics*, 3:283–297.

David Zhou and Sarah Sterman. 2023. Creative struggle: Arguing for the value of difficulty in supporting ownership and self-expression in creative writing. In *The Second Workshop on Intelligent and Interactive Writing Assistants (In2Writing)*.

Jichen Zhu, Antonios Liapis, Sebastian Risi, Rafael Bidarra, and G Michael Youngblood. 2018. Explainable ai for designers: A human-centered perspective on mixed-initiative co-creation. In *2018 IEEE conference on computational intelligence and games (CIG)*, pages 1–8. IEEE.

A Commercial Writing Assistants for Dutch Text Simplification

There are a couple of existing commercial writing assistants for text simplification in Dutch, such as *Klinkende Taal* and *Tolkie Schrijfhulp*. These tools offer a convenient co-creation interface by integrating into widely used software like Microsoft Word and Outlook; the control over the writing process remains mainly with the human writer, who gets feedback from the system and decides whether to incorporate it. The tools provide some degree of modularity and explainability by identifying specific problems in the text, including difficult words and complicated structures (e.g., passive sentences, long sentences); they provide general explanations about the complexity (e.g., *"passive sentences make it unclear who does what, and give the text a distant tone"*), and in some cases suggests alternatives (e.g., synonyms). The incorporation of variability is limited in these tools. Both perform their evaluation based on CEFR levels; while *Tolkie Schrijfhulp* focuses exclusively on the B1-level target, *Klinkende Taal* offers some variability by letting the user choose the target CEFR level herself. In addition, *Tolkie Schrijfhulp* offers one group-specific check: words that are difficult for people with dyslexia.

B Resources for Dutch Text Simplification

Description	Size	Source simplifications / annotations	Domain	Link	Reference
Parallel corpus	1,311 sentence pairs	Manual	Government	Link to GitHub	Vlantis et al. (2024)
Parallel corpus	1,267 sentence pairs	Automatic (LLM)	*unknown*	Link to HuggingFace	van de Velde (2023)
Parallel corpus	2.87M paragraph pairs	Automatic (LLM)	Wikipedia	Link to HuggingFace	*n/a*
Contextualized lexical simplifications	96 sentences	Manual	Government	Link to GitHub	Hobo et al. (2023)
Complex words and simpler alternatives	∼800 words / expressions	Manual	Government	link to City of Amsterdam	*n/a*
Complex words and simpler alternatives	∼130 words / expressions	Manual	Legal	Link to City of Amsterdam	*n/a*
Words and frequency distributions graded on CEFR levels	17,743 words / expressions	Automatic	*n/a*	Link to NT2Lex	Tack et al. (2018)
Texts graded on CEFR levels	1,200 texts	Manual	Various	Link to Edia	Breuker (2022)
Dutch-English cognates and homographs	∼200 words	Manual	*n/a*	Link to OSF	Poort and Rodd (2019)

Table 1: Overview of resources for Dutch text simplification