

Leveraging Digital Accessibility Using Generative AI

RENÉ SCHÄFER, RWTH Aachen University, Germany

SARAH SAHABI, RWTH Aachen University, Germany

PAUL PREUSCHOFF, RWTH Aachen University, Germany

JAN BORCHERS, RWTH Aachen University, Germany

Generative AI (GenAI), and especially large language models (LLMs), offer intriguing opportunities to improve accessibility in the digital world. In this position paper, we present our research findings and ideas for next steps in two areas of ongoing work in this field. First, we show that LLMs can adjust HTML to mitigate deceptive patterns and manipulative designs in online user interfaces, without prior training. For this, we prompted GPT-4o with HTML code, asked it iteratively to make it less manipulative, and manually analyzed 2,600 redesigns that demonstrate the potential of this approach as a technical countermeasure against deceptive patterns. Second, we outline how GenAI approaches can augment screen readers used by blind and low-vision users by enabling natural, conversational interactions and generating missing metadata. A particular advantage of GenAI for this use case lies in its capability to adapt to individual user preferences and needs.

Additional Key Words and Phrases: deceptive patterns, dark patterns, countermeasures, large language models

1 Introduction

Interactions in the digital world are rarely friction-free. Some interfaces are designed to manipulate users. Others lack qualities of universal usability, making interactions unnecessarily difficult for user groups with particular capabilities. In both cases, digital accessibility technologies can help remove these frictions and improve the user experience.

In our first example, digital accessibility can support users when they encounter content with *deceptive patterns*. These patterns manipulate users in their decision-making in favor of a service owner [16]. A common example is a cookie banner with a colorful and large button to accept all cookies, while the decline option is barely visible at all. Deceptive patterns are common in various contexts such as shopping websites [10, 16], games [27], social media [17], and apps [6, 12]. Because of this prevalence, researchers call for effective countermeasures to eventually help users [2, 7]. One such approach is technical countermeasures that, e.g., automatically detect deceptive patterns in an interface [5, 16] and then highlight them with an explanation [16, 21] or remove them completely [22]. Generative AI (GenAI) has also started to shift into focus for this purpose [18, 20]. However, large language models (LLMs) have not been used to counter deceptive patterns directly on the web. Therefore, we developed an approach that prompts an LLM with the task of making a given website less manipulative by adjusting the underlying HTML code and CSS classes.

Digital accessibility also plays a crucial role in our second example. There are over 270 million blind and low-vision (BLV) individuals worldwide [3], many of whom are relying on assistive technology, such as screen readers and voice assistants [23]. However, these tools have limitations, such as missing metadata [23]. Moreover, they may come with a range of usability issues: For example, the sequential output of screen readers results in inefficient interactions, and it is difficult to remember the numerous keyboard shortcuts or voice commands, which increases cognitive load [13, 24].

Authors' Contact Information: René Schäfer, rschaefer@cs.rwth-aachen.de, RWTH Aachen University, Aachen, Germany; Sarah Sahabi, sahabi@cs.rwth-aachen.de, RWTH Aachen University, Aachen, Germany; Paul Preuschoff, preuschoff@cs.rwth-aachen.de, RWTH Aachen University, Aachen, Germany; Jan Borchers, borchers@cs.rwth-aachen.de, RWTH Aachen University, Aachen, Germany.



This work is licensed under a Creative Commons Attribution 4.0 International License.




Fig. 1. We prompted the LLM GPT-4o with the HTML source of UI designs like the one on the left and prompted it to “Make that less manipulative”. For example, a fictional ticket shop (left side) uses several deceptive patterns to manipulate users into purchasing VIP tickets for a festival. After a few iterations with our prompt, the LLM provided a fair design (right side).

Researchers have addressed these issues by introducing conversational screen readers that shift away from the linear nature of traditional screen readers and, hence, promise faster learning progress and usage [1, 24]. However, they have been limited in their ability to understand and communicate natural language and adapt to user preferences over time. Recent progress in GenAI, particularly LLMs, holds the potential to address some of the current limitations [19].

For example, research has identified BLV users’ desires for natural conversational interactions, as well as systems capable of learning and adapting to user behavior and individual preferences over time (e.g., speech rate, information order, or interaction granularity) [19]. Accordingly, recent studies augmented traditional screen readers with LLMs, enabling more efficient natural language interactions [13, 25]. Instead of performing a series of keyboard shortcuts through voice commands, they enable the use of natural language dialogs to achieve the same in a single step. Compared to conventional practices like the JAWS screen reader¹, user studies with BLV participants revealed significant improvements in interaction efficiency and usability. To further explore the potential of GenAI to improve digital accessibility for BLV users, we interviewed 13 BLV participants and collected qualitative data from 19 BLV participants in an online survey. Participants expressed their wish for AI-based conversational screen readers that provide additional functionalities beyond their traditional counterparts. Examples include engaging in natural conversations, as shown by previous research, offering descriptions of visual elements lacking alt text, and completing actions on the user’s behalf.

In the GenAICHI workshop, we want to spark discussions about using GenAI to leverage digital accessibility. To support this, we bring the above two concrete examples from our ongoing work to the table. In the remainder of this position paper, we therefore provide results from our current work and ideas for future research in these two contexts:

- (1) Mitigating the effects of deceptive patterns
- (2) Enhancing digital accessibility tools for BLV users

2 Using Generative AI to mitigate deceptive patterns

Deceptive patterns have become widespread [6, 10, 14]. Because of this, researchers have classified and unified existing patterns [8] and called for effective countermeasures [7]. Proposed countermeasures range from educating users or designers to regulatory approaches and technical countermeasures [2]. Such technical countermeasures often require reliable automatic detection of deceptive patterns. However, automatic detection and classification are hard to solve, and some patterns might not be automatically detectable at all [5]. To bridge this gap, we introduce an approach for a technical countermeasure that does not rely on explicit a priori detection and classification. Instead, it feeds the HTML source of web elements into a large language model (LLM) and iteratively prompts it to make these elements less

¹<https://www.freedomscientific.com/products/software/jaws/>, Accessed: February 20, 2025




Fig. 2. The average value of each iteration of the redesigns compared to the initial design for all deceptive patterns using the *minimal* and *improved prompt*. Both prompts achieved the least manipulative designs after three iterations.

manipulative. We evaluated our approach with OpenAI’s GPT-4o² LLM and both self-created deceptive and fair web elements, and real websites.

We started with a *minimal prompt* “Make that less manipulative”. This zero-shot prompt already successfully removed several patterns and worked surprisingly well, given that we did not provide any further information to the LLM. Still, several problems occurred repeatedly in the LLM’s redesigns. Sometimes, the LLM introduced additional deception, e.g., by making one of several buttons look grayed out. In other cases, the LLM hallucinated additional facts or options. Based on the problems we observed using our *minimal prompt*, we constructed an *improved prompt* that contained several rules that the LLM had to adhere to, such as not changing any numbers. This *improved prompt* drastically reduced hallucinations across trials. Both prompt variants yielded the best results after three iterations (Figure 2). With our improved prompt, 91% of deceptive elements were less manipulative than our initial designs at that point. To evaluate this approach in the field, we performed preliminary tests with our *improved prompt* on real websites. We extracted the source of the websites using the browser’s developer tools and reintegrated the output of the LLM back into the browser. These first tests appear promising, and the generated HTML visually fits the remaining website. Nevertheless, automatically altering content could raise issues of user consent and data handling that should be investigated further.

Our findings suggest that LLMs can defuse certain deceptive patterns without prior model training, promising a major advance in fighting these manipulations. We openly share the labeled dataset of our 26 initial designs and 2,600 generated redesigns, including all responses and justifications provided by GPT-4o, on an Open Science Framework repository³. Overall, we observed surprisingly positive results from our tests with GPT-4o, and we see this approach as promising to assist people facing manipulative designs in their daily online interactions.

3 Augmenting Screen Readers with Generative AI

Through a short online survey, we asked 19 BLV participants about their currently used AI assistive technologies, potential use cases of AI-enhanced assistance, anticipated challenges when using such technologies daily, and desired features for enhancing their usefulness. Additionally, we interviewed 13 regular screen reader users online for more in-depth insights. Since we recruited participants of varying technology proficiency levels, we created a simple demonstrator of a GenAI-based voice assistant using GPT-4o to present a potential use case of GenAI. Participants could interact with the demonstrator by speaking or typing before receiving a read-out answer from the underlying model. The UI was kept simple, containing only one button for starting and stopping voice recording, and provided

²<https://platform.openai.com/docs/models/gpt-4o>, Accessed: February 20, 2025

³<http://doi.org/10.17605/osf.io/tgrw9>

natural voice and auditory feedback. After familiarization with the demonstrator, participants performed four example tasks we had prepared to foster a similar level of exposure across participants. We defined tasks based on Marchionini's classification of search actions [15]. In a semi-structured interview, we discussed participants' thoughts, experiences, and ideas on GenAI-enhanced assistive technology. We analyzed the transcripts of our audio recordings and the written survey data using reflexive thematic analysis [4]: We inductively coded the data and iteratively refined and grouped the results into broader themes encapsulating current challenges, visions, and opportunities expressed by participants.

A recurring topic was the potential of augmenting screen readers with GenAI-based functionality. As identified in related research, our participants stated that, even after a long learning phase to learn how to control a screen reader, the interaction feels cumbersome because elements are mostly read out sequentially. Instead, natural language processing (NLP) capabilities of LLMs integrated into a screen reader might enable a conversational, natural interaction [9, 13]. For example, instead of using keyboard shortcuts to navigate to and open a specific file, the user could just request to open the presentation file attached to yesterday's email from the financial department.

Moreover, some participants mentioned the possibility of an AI-supported screen reader analyzing the screen content and generating short summaries to provide an initial overview of what is displayed. Users could ask more detailed follow-up questions about specific content in a natural dialog instead of manually locating information with a screen reader. In particular, one participant envisioned metaphorical auditive feedback based on the screen content. For example, a clock's ticking sound could be played to represent loading times.

P05: "Maybe it can make a scrolling noise, and then when it looks like it's done or it gets to know that website, play me a little song or something."

Additionally, participants reported that required metadata, such as alt text or labels of UI elements, is often missing or inaccurate on websites, limiting screen readers' functionality. These elements can be identified by object detection and image recognition models [26, 28], and corresponding metadata can be generated after the fact. Again, the conversational nature of the GenAI-based screen reader would allow users to ask more questions, e.g., about the image for which the alt text was generated/adapted. Besides this, the capabilities of AI to learn over time could be used to personalize the metadata for individual users depending on their personal preferences and respective contexts.

Furthermore, participants emphasized the importance of the screen reader output:

P02: "So for everybody, it should be possible to make some changes on these things [voice and speed]. Because everybody has different preferences."

While basic customizations are already available, participants envisioned more dynamic adaptations to transform the user experience of assistive systems into something that feels more like a personified assistant. This includes matching the voice to a specific character, role, or mood, which could be realized via a GenAI-backed voice synthesis [11].

Our user study revealed several opportunities to augment traditional screen readers with GenAI-based functionality from BLV users' perspectives. However, participants emphasized that AI should not replace existing assistive technology but extend its functionalities instead.

4 Conclusion

We see potential for GenAI, and especially large-language models (LLMs), to forge a new path of online browsing accessibility. As examples, we showed how current LLMs can be used to automatically remove certain deceptive patterns that are embedded in HTML code and how GenAI may improve and expand the user experience of current screen readers. With these impulses, we hope to spark fruitful and exciting discussions at GenAICHI 2025.

Acknowledgments

This work was funded in part by the German B-IT Foundation.

References

- [1] Vikas Ashok, Yevgen Borodin, Yury Puzis, and I. V. Ramakrishnan. 2015. Capti-Speak: A Speech-Enabled Web Screen Reader. In *Proceedings of the 12th International Web for All Conference*. ACM, Florence Italy, 1–10. <https://doi.org/10.1145/2745555.2746660>
- [2] Kerstin Bongard-Blanchy, Arianna Rossi, Salvador Rivas, Sophie Doublet, Vincent Koenig, and Gabriele Lenzi. 2021. "I Am Definitely Manipulated, Even When I Am Aware of It. It's Ridiculous!" - Dark Patterns from the End-User Perspective. In *Designing Interactive Systems Conference (Virtual Event, USA)*. ACM, New York, 763–776. <https://doi.org/10.1145/3461778.3462086>
- [3] Rupert RA Bourne, Jaimie Adelson, Seth Flaxman, Paul Briant, Michele Bottone, Theo Vos, Kovin Naidoo, Tasanee Braithwaite, Maria Cicinelli, Just Jonas, et al. 2020. Global Prevalence of Blindness and Distance and Near Vision Impairment in 2020: Progress Towards the Vision 2020 Targets and What the Future Holds. *Investigative Ophthalmology & Visual Science* 61, 7 (2020), 2317–2317.
- [4] Virginia Braun and Victoria Clarke. 2006. Using Thematic Analysis in Psychology. *Qualitative Research in Psychology* 3 (01 2006), 77–101. <https://doi.org/10.1191/1478088706qp0630a>
- [5] Andrea Curley, Dymrna O’Sullivan, Damian Gordon, Brendan Tierney, and Ioannis Stavrakakis. 2021. The Design of a Framework for the Detection of Web-Based Dark Patterns. In *The Fifteenth International Conference on Digital Society (ICDS 2021)*. IARIA, Nice, France, 24–30. <https://doi.org/10.21427/20g8-d176>
- [6] Linda Di Geronimo, Larissa Braz, Enrico Fregnan, Fabio Palomba, and Alberto Bacchelli. 2020. UI Dark Patterns and Where to Find Them: A Study on Mobile Applications and User Perception. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376600>
- [7] Colin M. Gray, Johanna T. Gunawan, René Schäfer, Nataliia Bielova, Lorena Sanchez Chamorro, Katie Seaborn, Thomas Mildner, and Hauke Sandhaus. 2024. Mobilizing Research and Regulatory Action on Dark Patterns and Deceptive Design Practices. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 482, 6 pages. <https://doi.org/10.1145/3613905.3636310>
- [8] Colin M. Gray, Cristiana Teixeira Santos, Nataliia Bielova, and Thomas Mildner. 2024. An Ontology of Dark Patterns Knowledge: Foundations, Definitions, and a Pathway for Shared Knowledge-Building. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 289, 22 pages. <https://doi.org/10.1145/3613904.3642436>
- [9] Yanchu Guan, Dong Wang, Zhixuan Chu, Shiyu Wang, Feiyue Ni, Ruihua Song, and Chenyi Zhuang. 2024. Intelligent Agents with LLM-based Process Automation. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (Barcelona, Spain) (KDD '24)*. Association for Computing Machinery, New York, NY, USA, 5018–5027. <https://doi.org/10.1145/3637528.3671646>
- [10] Johanna Gunawan, Amogh Pradeep, David Choffnes, Woodrow Hartzog, and Christo Wilson. 2021. A Comparative Study of Dark Patterns Across Web and Mobile Modalities. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 377 (Oct 2021), 29 pages. <https://doi.org/10.1145/3479521>
- [11] Juhye Ha, Hyeon Jeon, Daeun Han, Jinwook Seo, and Changhoon Oh. 2024. CloChat: Understanding How People Customize, Interact, and Experience Personas in Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 305, 24 pages. <https://doi.org/10.1145/3613904.3642472>
- [12] Shun Hidaka, Sota Kobuki, Mizuki Watanabe, and Katie Seaborn. 2023. Linguistic Dead-Ends and Alphabet Soup: Finding Dark Patterns in Japanese Apps. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (Hamburg, Germany) (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 3, 13 pages. <https://doi.org/10.1145/3544548.3580942>
- [13] Satwik Ram Kodandaram, Utku Uckun, Xiaojun Bi, Iv Ramakrishnan, and Vikas Ashok. 2024. Enabling Uniform Computer Interaction Experience for Blind Users through Large Language Models. In *The 26th International ACM SIGACCESS Conference on Computers and Accessibility*. ACM, St. John's NL Canada, 1–14. <https://doi.org/10.1145/3663548.3675605>
- [14] Alexander Löbel, René Schäfer, Hanna Püschel, Esra Güney, and Ulrike Meyer. 2024. Access Your Data... if You Can: An Analysis of Dark Patterns Against the Right of Access on Popular Websites. In *Privacy Technologies and Policy (Karlstad, Sweden)*. Springer-Verlag, Berlin, Heidelberg, 23–47. https://doi.org/10.1007/978-3-031-68024-3_2
- [15] Gary Marchionini. 2006. Exploratory Search: From Finding to Understanding. *Commun. ACM* 49, 4 (April 2006), 41–46. <https://doi.org/10.1145/1121949.1121979>
- [16] Arunesh Mathur, Gunes Acar, Michael J. Friedman, Eli Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 81 (Nov 2019), 32 pages. <https://doi.org/10.1145/3359183>
- [17] Thomas Mildner, Merle Freye, Gian-Luca Savino, Philip R. Doyle, Benjamin R. Cowan, and Rainer Malaka. 2023. Defending Against the Dark Arts: Recognising Dark Patterns in Social Media. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (Pittsburgh, PA, USA) (DIS '23)*. Association for Computing Machinery, New York, NY, USA, 2362–2374. <https://doi.org/10.1145/3563657.3595964>
- [18] Stuart Mills and Richard Whittle. 2023. Detecting Dark Patterns Using Generative AI: Some Preliminary Results. , 33 pages. <https://doi.org/10.2139/ssrn.4614907>

- [19] Mahika Phutane, Crescentia Jung, Niu Chen, and Shiri Azenkot. 2023. Speaking with My Screen Reader: Using Audio Fictions to Explore Conversational Access to Interfaces. In *The 25th International ACM SIGACCESS Conference on Computers and Accessibility*. ACM, New York NY USA, 1–18. <https://doi.org/10.1145/3597638.3608404>
- [20] Yasin Sazid, Mridha Md Nafis Fuad, and Kazi Sakib. 2023. Automated Detection of Dark Patterns Using In-Context Learning Capabilities of GPT-3. In *2023 30th Asia-Pacific Software Engineering Conference (APSEC)*. IEEE, IEEE, Seoul, Republic of Korea, 569–573. <https://doi.org/10.1109/APSEC60848.2023.00072>
- [21] René Schäfer, Paul Miles Preuschoff, and Jan Borchers. 2023. Investigating Visual Countermeasures Against Dark Patterns in User Interfaces. In *Proceedings of Mensch Und Computer 2023 (Rapperswil, Switzerland) (MuC '23)*. Association for Computing Machinery, New York, NY, USA, 161–172. <https://doi.org/10.1145/3603555.3603563>
- [22] René Schäfer, Paul Miles Preuschoff, René Röpke, Sarah Sahabi, and Jan Borchers. 2024. Fighting Malicious Designs: Towards Visual Countermeasures Against Dark Patterns. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 296, 13 pages. <https://doi.org/10.1145/3613904.3642661>
- [23] Abigale Stangl, Meredith Ringel Morris, and Danna Gurari. 2020. "Person, Shoes, Tree. Is the Person Naked?" What People with Vision Impairments Want in Image Descriptions. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376404>
- [24] Alexandra Vtyurina, Adam Fourney, Meredith Ringel Morris, Leah Findlater, and Ryen W. White. 2019. VERSE: Bridging Screen Readers and Voice Assistants for Enhanced Eyes-Free Web Search. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*. ACM, Pittsburgh PA USA, 414–426. <https://doi.org/10.1145/3308561.3353773>
- [25] Bryan Wang, Gang Li, and Yang Li. 2023. Enabling Conversational Interaction with Mobile UI Using Large Language Models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–17. <https://doi.org/10.1145/3544548.3580895>
- [26] Jason Wu, Rebecca Krosnick, Eldon Schoop, Amanda Swearngin, Jeffrey P Bigham, and Jeffrey Nichols. 2023. Never-ending Learning of User Interfaces. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (San Francisco, CA, USA) (UIST '23)*. Association for Computing Machinery, New York, NY, USA, Article 113, 13 pages. <https://doi.org/10.1145/3586183.3606824>
- [27] José P Zagal, Staffan Björk, and Chris Lewis. 2013. Dark Patterns in the Design of Games. In *Foundations of Digital Games 2013*.
- [28] Xiaoyi Zhang, Lilian De Greef, Amanda Swearngin, Samuel White, Kyle Murray, Lisa Yu, Qi Shan, Jeffrey Nichols, Jason Wu, Chris Fleizach, Aaron Everitt, and Jeffrey P Bigham. 2021. Screen Recognition: Creating Accessibility Metadata for Mobile Applications from Pixels. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–15. <https://doi.org/10.1145/3411764.3445186>