Demo of RAVEN: Realtime Accessibility in Virtual ENvironments for Blind and Low-Vision People

Xinyun Cao
xinyunc@umich.edu
University of Michigan
Ann Arbor, Michigan, USA

Kexin Phyllis Ju
kexinju@umich.edu
University of Michigan
Ann Arbor, Michigan, USA

Chenglin Li
lchengl@umich.edu
University of Michigan
Ann Arbor, Michigan, USA

Venkatesh Potluri
potluriv@umich.edu
University of Michigan
Ann Arbor, Michigan, USA

Dhruv Jain
profdj@umich.edu
University of Michigan
Ann Arbor, Michigan, USA

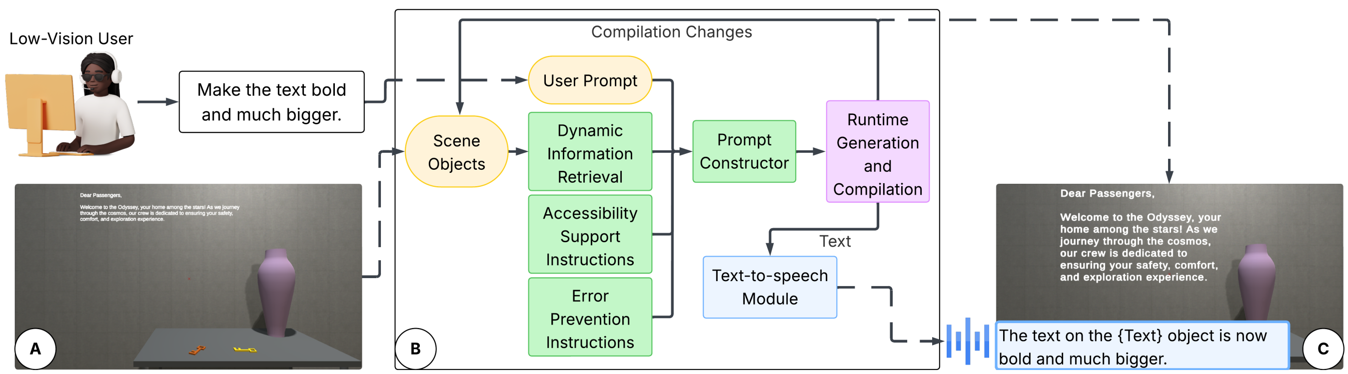

Figure 1: RAVEN is an interactive system that empowers BLV users to query and modify 3D scenes via natural language. The above image illustrates an example of an accessibility modification: A) A low-vision user types in a modification text. B) The system integrates runtime code generation LLM agent with dynamic scene information and instructions to apply accessibility-enhancing changes at runtime. C) The system compiles LLM-produced code to achieve modification while providing spoken response to the user.

ABSTRACT

As virtual 3D environments become prevalent, enabling presence and spatial exploration, equitable access is crucial for blind and low-vision (BLV) users who face challenges with spatial awareness, navigation, and interaction. To address these, previous work explored enhancing visual information or supplementing it with auditory and haptic modalities. However, these methods are static and might risk steep learning curves. In this work, we present RAVEN, a system that responds to query or modification prompts from BLV users to improve the accessibility of a 3D virtual scene at runtime. The system integrates LLMs with semantic scene data and runtime code generation to support iterative, dialogue-based user interactions. We evaluated the system with eight BLV people, uncovering key insights into the strengths and shortcomings of generative AI-driven accessibility in virtual 3D environments.

CCS CONCEPTS

• **Human-centered computing → Accessibility systems and tools**.

Conference'17, July 2017, Washington, DC, USA
2025. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
https://doi.org/10.1145/nnnnnnn.nnnnnnn

KEYWORDS

Accessibility, blind and low-vision, virtual 3D environment, generative AI.

1 INTRODUCTION

Virtual 3D environments have become pervasive, enabling rich interactions, offering users a sense of presence and spatial exploration [5, 15]. However, with their increased adoption arises the imperative of ensuring inclusive and equitable access, particularly for blind and low-vision (BLV) users who encounter substantial challenges in spatial understanding, navigation, and object interaction [29].

To address these challenges, prior work has proposed tools that modify or supplement visual information using alternative modalities, such as audio descriptions [11], haptic feedback [13, 32], and enhanced visual effects tailored for low vision [33]. Systems like SceneWeaver [2] offer users greater agency over when and how to consume scene descriptions. Commercial platforms have also adopted accessibility features - such as high-contrast display modes and spatial audio - to accommodate BLV players [4, 20].

However, current accessibility approaches share a fundamental limitation: they are largely developer-driven and static - such as fixed mappings for color changes or auditory overlays - and often fail to meet the nuanced, evolving needs of individual users [9]. They also typically require users to learn specific control mappings or adjust settings in non-intuitive ways, resulting in steep learning curves and limited support for dynamic, context-specific adaptation.

Recent advances in generative AI, particularly large language models (LLMs), present possibilities for new, flexible, and conversational interaction paradigms. Prior work has leveraged LLMs for accessibility tasks ranging from querying visual information [1] to runtime scene editing [14]. These developments suggest the potential for LLM-powered tools to equip BLV users to directly query and adapt 3D scenes using natural language, bypassing rigid developer-defined workflows and expanding user agency.

To leverage this emerging capability, we present RAVEN, an interactive system that empowers BLV users to engage with 3D scenes via conversational natural language interaction. RAVEN supports both scene querying (e.g., "What's around me?") and accessibility-related modifications (e.g., "Make the table brighter" or "Move the bench closer"). The system integrates LLMs with semantic scene data and runtime code generation tools to apply accessibility-enhancing changes at runtime, while also providing spoken responses to user queries. Interaction is iterative, allowing users to refine modifications through follow-up prompts in a dialogue-like flow.

2 SYSTEM DESIGN

RAVEN enables blind and low-vision (BLV) users to interact with 3D virtual environments using natural language. It supports two core capabilities: (1) querying information about the scene, and (2) modifying the scene *in real time* to improve accessibility. While the system is designed to handle a broad range of natural language prompts, this work focuses on six core prompt categories that reflect common accessibility needs and provide a structured basis for evaluation. We first introduce these categories (section 2.1), then describe the system architecture and technical components that support them (section 2.2).

2.1 Accessibility Improvement Categories

RAVEN supports prompts belonging to a set of improvement categories informed by prior research on 3D accessibility for BLV individuals [21, 22, 25, 26, 30, 33]. Four categories target visual accessibility, and two focus on auditory interaction. Our system is optimized to support both queries and modifications in each of these categories.

Visual Categories. These allow users to adapt the visual presentation of the scene to better match their preferences:

(1) *Color:* Retrieve and modify color schemes or colors for specific objects to aid recognition, especially for users with color blindness [21, 30].
(2) *Object Location:* Retrieve information about locations of objects such as furniture, characters or players, and reposition them to simplify navigation and object interaction, inspired by assistive game toolkits [27].
(3) *Size:* Resize objects or in-scene text to enhance visibility and readability [22, 26].
(4) *Scene Brightness:* Provide information about scene brightness, and adjust lighting across the scene, or for specific light sources, to accommodate individual comfort and visual sensitivity [33].

Auditory Categories. These features enhance awareness of environmental elements using sound:

(1) *Volume:* Increase or decrease the volume of specific sound sources to isolate or emphasize key elements [7, 23].
(2) *Pitch:* Alter pitch to distinguish sounds more effectively, highlight key content, or emulate common screen reader cues [7, 12, 19, 28].

2.2 RAVEN Architecture

RAVEN operates on minimally annotated 3D scenes. At development time, designers label key objects by attaching metadata that includes each object's name, visual and auditory descriptions, and whether it is physical (e.g., a bench) or non-physical (e.g., background music). Users interact with RAVEN by typing natural language prompts to query (e.g., "what is the color of the bench") or modify the scene (e.g., "make the text bigger"). To account for limited screen reader support in Unity environments, the system self-voices text input, prompts, and responses using built-in speech synthesis [16]. When a user submits a prompt, RAVEN follows a three-stage processing pipeline (fig. 1):

(1) **Dynamic Information Retrieval** extracts the semantic scene graph (SSG) with updated object properties.
(2) **Prompt Constructor** builds an LLM-compatible prompt using the user query, SSG, and task-specific instructional prompts.
(3) **Runtime Generation and Compilation**, powered by Gromit [14], generates and applies code for scene modifications or provides a verbal response.

See fig. 1 for the system workflow and an example use case.

2.2.1 *Dynamic Information Retrieval.*

RAVEN uses a semantic scene graph (SSG) similar to that described in [14], capturing a structured view of scene elements. For each object, the SSG stores its name, developer-provided descriptions (visual, auditory, functional), scripts, center point, scale, and child object relationships.

To support accessibility-related queries and modifications, the system augments the SSG with dynamic properties, including: (1) object color (HEX code), (2) text content and font size, (3) relative position or distance from the player, (4) light source density, and (5) audio source properties (mute/unmute, volume, pitch, range).

2.2.2 *Prompt Constructor.*

The Prompt Constructor synthesizes the user's input, scene context, and prompt-engineered instructions into a structured prompt for the LLM. To ensure the LLM produces contextually grounded and trustworthy outputs, especially in a critical domain like accessibility, we employ two key strategies:

Accessibility Support Instructions. To compensate for LLM limitations in applying accessibility design principles contextually [10], we embed rules that guide the model to reason about: (1) object position and direction, (2) size of objects and text, and (3) lighting and sound features.

Error Prevention Instructions. To reduce hallucinations and gracefully handle vague or unsupported requests, the system: (1) prompts the user for clarification if input is under-specified, (2) provides corrective suggestions when users report issues (e.g., "it's not working"), and (3) communicates feature limitations (e.g., no magnifiers or captioning) using an explicit out-of-scope list.

2.2.3 Runtime Generation and Compilation. RAVEN uses Gromit [14, 18], an open-source framework for real-time behavior generation in Unity. When given a constructed prompt, the LLM generates a textual response and, if applicable, Unity script code targeting a specific object. Gromit compiles and attaches the script at runtime, applying the modification in the scene. It then returns the LLM's textual response, which RAVEN reads aloud to the user. For our implementation, we upgraded Gromit to use GPT-4o for enhanced reasoning and generation quality.

3 USER STUDY

To assess the utility of on-demand accessibility modifications and the perceived usability of our system, we conducted a user evaluation of the RAVEN system with BLV participants across three scenes with increasingly open-ended tasks.

3.1 Scenes

We describe below the scenes that shaped participants' experiences with our system. Each scene involved tasks that offered progressively more open-ended interaction. The demo experience of this work consists of these scenes.

Scene 1, Guided Tutorial: In Scene 1 (see 2a), we provided a clear demonstration of the system and the categories in a simple scene resembling a room in a game. The researcher demonstrated the six categories (section 2.1), then invited participants to try similar prompts themselves to learn and explore the system.

Scene 2, Task-Driven Exploration: In Scene 2 (see 2b), we observed how participants applied the accessibility categories to achieve goals within a task-oriented context. This scene, set in a virtual park with nature items and sound sources, included six predefined tasks aligned with the categories from the guided tutorial. (See suppl. material.)

Scene 3, Open-Ended Exploration: In Scene 3 (see 2c), we sought to understand how participants used the system within an open-ended context. The scene featured a relatively complex environment - a spaceship-themed room with 16 objects of various shapes and colors, including three sound sources. Participants were given ten minutes to freely explore the scene and use RAVEN to support their accessibility needs in this process.

3.2 Method

We recruited eight BLV participants (5 men, 3 women) with an average age of 36.6 years and a diverse range of visual abilities. Two participants were blind with no vision, two had light perception, and four identified as low-vision or visually impaired. The study was conducted in person and lasted 1.5 hours per participant. The scenes ran as Unity projects on the researcher's device, and participants interacted with the system using a keyboard and stereo headphones.

Each session began with the three scene experiences, followed by a survey collecting Likert-scale ratings on the perceived usefulness of the six accessibility categories, confidence in achieving accessibility improvements, system intuitiveness, and responses

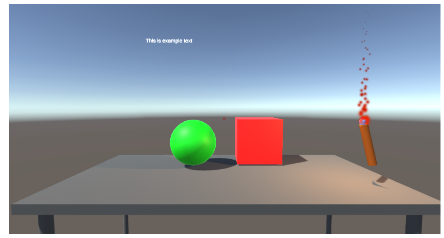

(a) Scene 1

(b) Scene 2

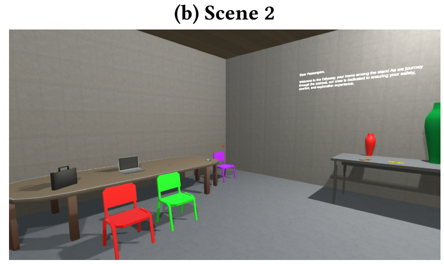

(c) Scene 3

Figure 2: Screenshots of the three scenes used in our evaluation. Scene 1 is a simple demo with basic elements and sound sources. Scene 2 is a park scene with cats and ambient nature sounds. Scene 3 is a spaceship-themed room with furniture, objects, and themed audio.

to the SUS questionnaire [6]. At the end of each session, we conducted a semi-structured interview to gather qualitative feedback on participants' experiences.

For analysis, we used descriptive statistics to summarize the questionnaire data and conducted two sets of reflexive thematic analysis [8] done by two coders on both interview transcripts and the prompts used. We also calculated the prompt success rate and categorized the reasons for errors.

3.3 Findings

On average, participants agreed that they felt confident using the tool to make scenes accessible to them (*mean*=4.1, *SD*=0.8, on a scale of 1: highly disagree to 5: highly agree). They also found the system intuitive to use (*mean*=4.3, *SD*=0.9, on a scale of 1: very unintuitive to 5: very intuitive). The average SUS usability score was 79.7, indicating good usability [3] with room for improvement. Participants' subjective feedback contextualize these ratings. Some (N=2) highlighted the system's robustness in handling typos (P2), responding to ambiguous queries (P5), and processing compound requests (P5). Others (N=3) credited the intuitiveness of the natural language interface, which provides improved learnability compared to existing methods with keyboard shortcuts. Several participants (N=2) also appreciated the system's open-endedness. As P8 noted, *"the sky was the limit in some of the things that I could ask [the system] to do"*.

Despite positive feedback, the system's error rate emerged as a key limitation. Out of 336 valid prompts, 253 (75.3%) resulted in correct query responses or intended modifications. Nine prompts were correctly flagged by the LLM as "out of scope," while 74 (22.0%) failed. Among these, 14 were intent errors, where the system misinterpreted the user's request. For instance, when P4 requested to *"extinguish the torch"*, the torch light was turned off but the burning sound persisted. The other 60 were technical errors, where the response did not match the scene state or the intended modification was not executed. These errors stemmed from LLM hallucinations, the model's lack of awareness regarding its access limitations (*e.g.,* actual volume or pitch of audio files), or failures during code compilation.

Most participants (N=7) expressed concerns about these errors. As P4 observed: *"Right now, it's not so foolproof as to say I completely trust it."* Participants pointed to a lack of transparency in failure cases and emphasized the need for more trustworthy verification methods. They also suggested several usability improvements, including: adaptive detail levels in scene descriptions (N=5), support for functional object interactions (N=3), undo capabilities (P4), more support for creative authoring (P2), and the ability to add localized sound cues for object identification (P1).

4 LIMITATIONS AND FUTURE WORK

While our findings highlight the promise of RAVEN in supporting BLV users through conversational accessibility modifications, they also reveal key limitations that inform directions for future work. First, the system is not yet ready for deployment in real-world applications due to a non-trivial error rate. Improving reliability will require advances in error prevention, safeguarding, and verification. One promising direction is the use of multiple AI agents to cross-validate outputs - such as ensemble methods [17, 31] or multi-agent debate strategies [24] - to enhance both correctness and user trust. Second, RAVEN currently lacks an understanding of object affordances. While it can retrieve an object's location and appearance, it does not reason about how that object can be used (e.g., recognizing a bench as a seating surface or localizing where to sit). Future work could incorporate richer semantic models of object functionality to support deeper interaction. Finally, our study was conducted in a controlled lab setting with predefined exploratory tasks. It did not evaluate performance in real-world 3D applications or time-sensitive environments (e.g., action or multiplayer games). Long-term, in-the-wild studies will be critical for assessing practical viability and sustained accessibility impact.

5 CONCLUSION

RAVEN explores a new frontier in accessible interaction - enabling blind and low-vision users to query and modify 3D virtual environments through natural language. By combining semantic scene understanding with real-time code generation, RAVEN shifts accessibility from a static, developer-defined feature to an interactive, user-driven experience. Our evaluation with eight BLV participants demonstrated the promise of this approach: users found the system intuitive, flexible, and empowering. However, the study also surfaced key limitations, particularly around reliability, error transparency, and the lack of affordance reasoning. As generative AI becomes more embedded in interactive systems, ensuring it supports accessibility with both accuracy and trust is essential. RAVEN marks an early step toward that goal, pointing to a future where immersive environments are not only richer, but fundamentally more inclusive.

REFERENCES

[1] Rudaiba Adnin and Maitraye Das. [n. d.]. "I look at it as the king of knowledge": How Blind People Use and Understand Generative AI Tools. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility* (New York, NY, USA, 2024-10-27) *(ASSETS '24)*. Association for Computing Machinery, 1–14. https://doi.org/10.1145/3663548.3675631

[2] Harshadha Balasubramanian, Cecily Morrison, Martin Grayson, Zhanat Makhataeva, Rita Faia Marques, Thomas Gable, Dalya Perez, and Edward Cutrell. [n. d.]. Enable Blind Users' Experience in 3D Virtual Environments: The Scene Weaver Prototype. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2023-04-19) *(CHI EA '23)*. Association for Computing Machinery, 1–4. https://doi.org/10.1145/3544549.3583909

[3] Aaron Bangor, Philip Kortum, and James Miller. 2009. Determining what individual SUS scores mean: adding an adjective rating scale. *J. Usability Studies* 4, 3 (May 2009), 114–123.

[4] Ben Bayliss. 2022. *Fortnite Accessibility — Menu Deep Dive*. https://caniplaythat.com/2022/04/26/fortnite-accessibility-menu-deep-dive/

[5] Doug A. Bowman, Chris North, Jian Chen, Nicholas F. Polys, Pardha S. Pyla, and Umur Yilmaz. [n. d.]. Information-rich virtual environments: theory, tools, and research agenda. In *Proceedings of the ACM symposium on Virtual reality software and technology* (New York, NY, USA, 2003-10-01) *(VRST '03)*. Association for Computing Machinery, 81–90. https://doi.org/10.1145/1008653.1008669

[6] john Brooke. [n. d.]. SUS: A 'Quick and Dirty' Usability Scale. In *Usability Evaluation In Industry*. CRC Press. Num Pages: 6.

[7] Ruei-Che Chang, Chia-Sheng Hung, Bing-Yu Chen, Dhruv Jain, and Anhong Guo. 2024. SoundShift: Exploring Sound Manipulations for Accessible Mixed-Reality Awareness. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (Copenhagen, Denmark) *(DIS '24)*. Association for Computing Machinery, New York, NY, USA, 116–132. https://doi.org/10.1145/3643834.3661556

[8] Victoria Clarke and Virginia Braun. [n. d.]. Thematic Analysis. In *Encyclopedia of Critical Psychology*, Thomas Teo (Ed.). Springer New York, 1947–1952. https://doi.org/10.1007/978-1-4614-5583-7_311

[9] Chris Creed, Maadh Al-Kalbani, Arthur Theil, Sayan Sarcar, and Ian Williams. [n. d.]. Inclusive AR/VR: Accessibility Barriers for Immersive Technologies. 23, 1 ([n. d.]), 59–73. https://doi.org/10.1007/s10209-023-00969-0 arXiv:2304.13465 [cs]

[10] Kate S Glazko, Momona Yamagami, Aashaka Desai, Kelly Avery Mack, Venkatesh Potluri, Xuhai Xu, and Jennifer Mankoff. 2023. An Autoethnographic Case Study of Generative Artificial Intelligence's Utility for Accessibility. In *Proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility* (New York, NY, USA) *(ASSETS '23)*. Association for Computing Machinery, New York, NY, USA, Article 99, 8 pages. https://doi.org/10.1145/3597638.3614548

[11] J.L. Gonzalez-Mora, A. Rodriguez-Hernandez, E. Burunat, F. Martin, and M.A. Castellano. [n. d.]. Seeing the world by hearing: Virtual Acoustic Space (VAS) a new space perception system for blind people.. In *2006 2nd International Conference on Information & Communication Technologies* (2006-04), Vol. 1. 837–842. https://doi.org/10.1109/ICTTA.2006.1684482

[12] William Grussenmeyer and Eelke Folmer. 2017. Accessible Touchscreen Technology for People with Visual Impairments: A Survey. *ACM Trans. Access. Comput.* 9, 2, Article 6 (Jan. 2017), 31 pages. https://doi.org/10.1145/3022701

[13] G. Jansson, H. Petrie, C. Colwell, D. Kornbrot, J. F"anger, H. K"onig, K. Billberger, A. Hardwick, and S. Furner. [n. d.]. Haptic Virtual Environments for Blind People: Exploratory Experiments with Two Devices. 4, 1 ([n. d.]), 8–17. https://doi.org/10.20870/IJVR.1999.4.1.2663 Number: 1.

[14] Nicholas Jennings, Han Wang, Isabel Li, James Smith, and Bjoern Hartmann. [n. d.]. What's the Game, then? Opportunities and Challenges for Runtime Behavior Generation. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2024-10-11) *(UIST '24)*. Association for Computing Machinery, 1–13. https://doi.org/10.1145/3654777.3676358

[15] Cher P. Lim, Darren Nonis, and John Hedberg. [n. d.]. Gaming in a 3D multiuser virtual environment: engaging students in Science lessons. 37, 2 ([n. d.]), 211–231. https://doi.org/10.1111/j.1467-8535.2006.00531.x _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8535.2006.00531.x.

[16] Microsoft 2025. Microsoft Speech API (SAPI) 5.3. Retrieved May 27, 2025 from https://learn.microsoft.com/en-us/previous-versions/windows/desktop/ms723627(v=vs.85)

[17] Ninad Naik. [n. d.]. Probabilistic Consensus through Ensemble Validation: A Framework for LLM Reliability. https://doi.org/10.48550/arXiv.2411.06535 arXiv:2411.06535 [cs]

[18] NicholasJJ. [n. d.]. *NicholasJJ/GROMIT*. https://github.com/NicholasJJ/GROMIT original-date: 2024-10-19T00:26:46Z.

[19] NVDA 2025. *NVDA 2024.4.2 User Guide*. Retrieved May 23, 2025 from https://download.nvaccess.org/documentation/userGuide.html

[20] Playstation 2024. The Last of Us Part II - Accessibility. Retrieved Jan 28, 2024 from https://www.playstation.com/en-us/games/the-last-of-us-part-ii/accessibility/

[21] Christine Rigden. [n. d.]. 'The Eye of the Beholder'— Designing for Colour-Blind Users. 17 ([n. d.]).

[22] Gary S. Rubin, Mary Feely, Sylvie Perera, Katherin Ekstrom, and Elizabeth Williamson. 2006. The effect of font and line width on reading speed in people with mild to moderate vision loss. *Ophthalmic and Physiological Optics* 26, 6 (2006), 545–554. https://doi.org/10.1111/j.1475-1313.2006.00409.x arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1475-1313.2006.00409.x

[23] Monika Rychtarikova. 2015. How do blind people perceive sound and soundscape. *Akustika* 23, 1 (2015), 6–9.

[24] Andries Smit, Paul Duckworth, Nathan Grinsztajn, Thomas D. Barrett, and Arnu Pretorius. [n. d.]. Should we be going MAD? A Look at Multi-Agent Debate Strategies for LLMs. https://doi.org/10.48550/arXiv.2311.17371 arXiv:2311.17371 [cs]

[25] Sarit Felicia Anais Szpiro, Shafeka Hashash, Yuhang Zhao, and Shiri Azenkot. 2016. How People with Low Vision Access Computing Devices: Understanding Challenges and Opportunities. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility* (Reno, Nevada, USA) *(ASSETS '16)*. Association for Computing Machinery, New York, NY, USA, 171–180. https://doi.org/10.1145/2982142.2982168

[26] Songül Atasavun Uysa and Tülin Düger. 2012. Writing and Reading Training Effects on Font Type and Size Preferences by Students with Low Vision. *Perceptual and Motor Skills* 114, 3 (2012), 837–846. https://doi.org/10.2466/15.10.11.24.PMS.114.3.837-846 arXiv:https://doi.org/10.2466/15.10.11.24.PMS.114.3.837-846 PMID: 22913024.

[27] VoiceAttack 2025. *VoiceAttack - Voice Recognition for your Games and Apps*. Retrieved Jan 22, 2025 from https://voiceattack.com/

[28] VoiceOver 2025. *Change your VoiceOver settings on iPhone*. Retrieved May 23, 2025 from https://support.apple.com/en-au/guide/iphone/iphfa3d32c50/ios

[29] Gareth R. White, Geraldine Fitzpatrick, and Graham McAllister. [n. d.]. Toward accessible 3D virtual environments for the blind and visually impaired. In *Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts* (New York, NY, USA, 2008-09-10) *(DIMEA '08)*. Association for Computing Machinery, 134–141. https://doi.org/10.1145/1413634.1413663

[30] Lee H. Wurm, Gordon E. Legge, Lisa M. Isenberg, and Andrew Luebker. [n. d.]. Color improves object recognition in normal and low vision. 19, 4 ([n. d.]), 899–911. https://doi.org/10.1037/0096-1523.19.4.899 Place: US Publisher: American Psychological Association.

[31] Han Yang, Mingchen Li, Huixue Zhou, Yongkang Xiao, Qian Fang, and Rui Zhang. [n. d.]. One LLM is not Enough: Harnessing the Power of Ensemble Learning for Medical Question Answering. https://doi.org/10.1101/2023.12.21.23300380 Pages: 2023.12.21.23300380.

[32] Yuhang Zhao, Cynthia L. Bennett, Hrvoje Benko, Edward Cutrell, Christian Holz, Meredith Ringel Morris, and Mike Sinclair. [n. d.]. Enabling People with Visual Impairments to Navigate Virtual Reality with a Haptic and Auditory Cane Simulation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2018-04-19) *(CHI '18)*. Association for Computing Machinery, 1–14. https://doi.org/10.1145/3173574.3173690

[33] Yuhang Zhao, Ed Cutrell, Christian Holz, Meredith Ringel Morris, Eyal Ofek, and Andrew D. Wilson. [n. d.]. SeeingVR: A Set of Tools to Make Virtual Reality More Accessible to People with Low Vision. https://www.microsoft.com/en-us/research/publication/seeingvr-a-set-of-tools-to-make-virtual-reality-more-accessible-to-people-with-low-vision-2/