

Navig-AI-tion: Navigation by Contextual AI and Spatial Audio

Mathias N. Lystbæk

mathiasl@cs.au.dk

Google

Mountain View, USA

Aarhus University

Aarhus, Denmark

Eric J Gonzalez

ejgonz@google.com

Google

Seattle, USA

Andrea Colacco

andreacolaco@google.com

Google

Mountain View, USA

Haley Adams

haleyadams@google.com

Google

Mountain View, USA

Luca Ballan

ballanlu@google.com

Google

Mountain View, USA

Peter Tan

ptan@google.com

Google

Mountain View, USA

Ranjith Kagathi Ananda

ranjithkagathi@google.com

Google

Mountain View, USA

Qiuxuan Wu

qiuxuanwu@google.com

Google

Mountain View, USA

Mar Gonzalez-Franco

margon@google.com

Google

Seattle, USA

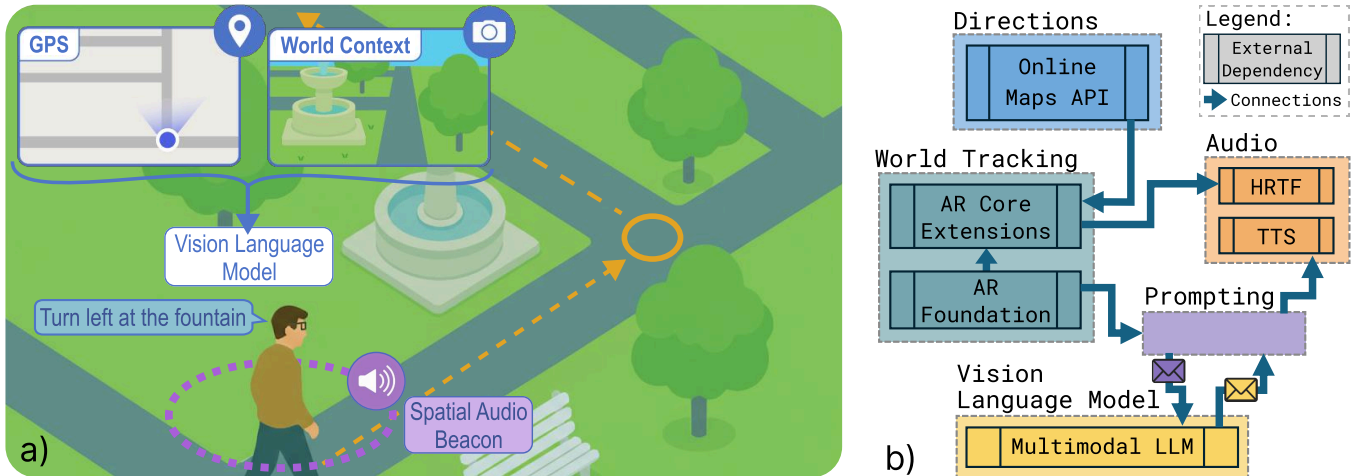


Figure 1: Augmented navigation using AI and spatial audio for display-less smart glasses (a). A simplified overview of the system architecture (b).

Abstract

Audio-only walking navigation can leave users disoriented, relying on vague cardinal directions and lacking real-time environmental context, leading to frequent errors. To address this, we present a novel system that integrates a Vision Language Model (VLM) with a spatial audio cue. Our system extracts environmental landmarks to anchor navigation instructions and, crucially, provides a directional spatial audio signal when the user faces the wrong direction, indicating the precise turn direction. In a user study ($n=12$), the spatial audio cue with VLM reduced route deviations compared to both

VLM-only and Google Maps (audio-only) baseline systems. Users reported that the spatial audio cue effectively supported orientation and that landmark-anchored instructions provided a better navigation experience over audio-only Google Maps. This work serves as an initial look at the utility of future audio-only navigation systems for incorporating directional cues, especially real-time corrective spatial audio.

CCS Concepts

• **Computing methodologies** → **Information extraction**; **Natural language generation**; • **Human-centered computing** → **Auditory feedback**.

Keywords

Artificial Intelligence, Navigation, Spatial Audio, Contextual Awareness

CHI EA '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*, April 13–17, 2026, Barcelona, Spain, <https://doi.org/10.1145/3772363.3799295>.

ACM Reference Format:

Mathias N. Lystbæk, Haley Adams, Ranjith Kagathi Ananda, Eric J Gonzalez, Luca Ballan, Qiuxuan Wu, Andrea Colacço, Peter Tan, and Mar Gonzalez-Franco. 2026. Navig-AI-tion: Navigation by Contextual AI and Spatial Audio. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3772363.3799295>

1 Introduction

Traditional GPS navigation has enabled a mobility revolution, yet there are legitimate questions about how much engagement with the environment current systems provide [12]. Modern turn-by-turn navigation excels at getting us from A to B, but growing evidence suggests that this could come at the cost of understanding and remembering the spaces we navigate [6, 7, 11]. Without coupling with the environment, these systems risk reducing navigation to blindly following an endless list of instructions [8].

Head-worn devices with cameras, like AR glasses, open new possibilities for enhancing navigation, particularly in complex walking scenarios where minimal screen interactions are desirable. Prior work has shown that spatial audio cues can provide more seamless orientation guidance, subtly assisting users in re-orienting themselves [2, 5, 10, 13]. Meanwhile, head-attached cameras combined with multimodal Large Language Models (LLMs), or Vision Language Models (VLMs), can generate contextual, landmark-based navigation instructions, e.g., replacing “turn left in 50 feet” with “turn left at the fountain” (cf. Figure 1a).

The combination of VLM-powered visual context and spatial audio has the potential to create a less disruptive navigation experience grounded in both sensory input and the environment. This is precisely the approach we prototype and evaluate in this paper.

While our goal is to support orientation in audio-only contexts, our preliminary study did not find improvements in spatial awareness (pointing accuracy). However, we observe reductions in route deviations and distance walked.

2 Related Work

The pervasive use of turn-by-turn navigation apps (e.g., Google Maps, Apple Maps, Waze) is a testament to their effectiveness for reaching destinations [16]. However, this convenience can come at a cognitive cost: the more people rely on GPS navigation, the worse their perceived navigation skills and environmental knowledge tends to become [15, 22].

2.1 Audio cues for navigation

Audio cues have been shown to effectively support walking navigation as an alternative to screen-based instructions [5, 10, 14, 20, 21, 23], with spatial audio in particular helping users find the correct direction [5, 19, 20]. Notably, Clemenson et al. introduced Soundscape [5], which guided users via continuous spatial audio and improved spatial recollection of points of interest (POIs). Our approach is similar, however, we only present the spatial audio cue when the user is not facing the suggested direction, lowering auditory load.

2.2 VLMs for navigation

Recent advances in LLMs enable richer, more context-aware descriptions of surroundings and more natural interactions [1], especially multimodal variants (VLMs) that integrate visual and textual data. To the best of our knowledge, few systems integrate VLMs for wayfinding. One example is NaviGPT [25], which combines Apple Maps (route guidance), LiDAR (obstacle detection), and GPT-4 (information processing and instructions) for real-time navigation on iPhone. In contrast to NaviGPT, our system is designed exclusively for audio-based delivery and supports display-less devices, anticipating future hardware form factors (e.g., Meta Ray-Bans). Additionally, we present a user evaluation yielding key design insights into VLM integration for wayfinding that can guide future development in this emerging field.

3 System

The system augments navigation instructions with context about the user and their environment, including orientation and a point-of-view photo. Furthermore, the system provides spatial auditory feedback when the user is not facing the next turn, giving the user an indication of the direction that they should walk. The spatial audio cue is otherwise turned off, allowing the user to listen to their own audio while walking (e.g., music, videos, or podcasts).

The system consists of five components working in sequence to (1) collect navigation instructions, (2) track the user’s position and orientation, (3) present audio feedback, (4) capture the user’s perspective and prompt the VLM, and (5) extract landmarks and augment the original instruction. See Figure 1b for a diagram overview.

3.1 Directions

Walking directions are queried based on a list of POI GPS coordinates (in this work, we used the Google Maps Platform API), returning JSON-formatted steps with walking instructions for the desired origin and destination.

3.2 World Tracking

The World Tracking component tracks the user and keeps track of where turn waypoints and POIs are relative to the user, also checking if the user has reached a turn. The system only captures an image from the user’s point-of-view when it is about to prompt the VLM. AR Foundation provides the core functionality for phone tracking, camera capture, and spatial audio positioning based on user orientation. AR Core Extensions adds accurate GPS coordinate capture and virtual anchor placement. User orientation is determined via Google’s Visual Positioning System (VPS).

3.3 Audio

The audio component controls spatial audio cue playback and Text-To-Speech (TTS) feedback (Android TTS API running on-device). To accurately simulate real-world audio characteristics, we utilize Head-Related Transfer Functions (HRTFs), simulating how sound passes through human ears based on the relative position of the audio cue and human ear topology (in this work, we use the Steam Audio SDK). Similar to [5], the spatial audio source is placed 1 meter from the user (cf. Figure 1a), conveying turn direction, but not distance. The spatial audio cue is presented as a repeating “waterfall”

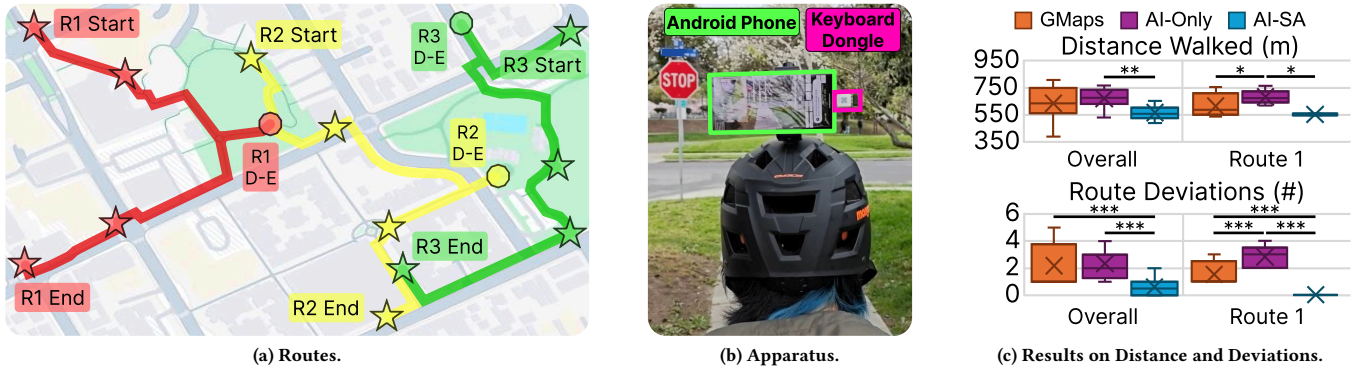


Figure 2: (a) Three color-coded routes, marked with route start and end for each route R_X , where “ X ” is 1 to 3, along with color-matched stars and circles indicating landmarks along the routes. “ R_X D-E” indicates the dead end for each route. (b) The apparatus setup for the prototype and evaluation. (c) Box plots showing results on Distance Walked (m) and Route Deviations (count). Whiskers indicate minimum and maximum. Statistical significance is shown as * for $p < 0.05$, ** for $p < 0.01$, and *** for $p < 0.001$.

sound, which we found suitable for individuals to accurately locate (validated through a training phase). As mentioned, the spatial audio cue is not presented to the user at all times, activating when the user’s facing direction deviates by more than 25° off from the current turn waypoint (either to the left or the right) and stops playing once the user is facing the current turn waypoint within 25° .

3.4 Prompting

The prompting component serves as an intermediary between the World Tracking, Audio, and VLM components. It captures the context (including camera image, position, orientation, direction and distance to the next turn, and information about the current navigation step), which is sent to the VLM component. Note that the model is initialized with system instructions to steer the model. Prompting occurs in two scenarios: (1) if the user is stationary while looking around and at some point looks in the direction of the current turn waypoint (within 25°), or (2) when walking close to the turn waypoint (within 30 meters). When a response is received from the VLM, the contextual navigation instruction is provided to the user via TTS through the Audio component.

3.5 Vision Language Model

The VLM component describes the structure of the system instructions provided when creating an instance of the VLM (in this work, we use Gemini 2.0 Flash Exp). The system instructions tell the model what role it should “play”, provides brief information about the system concept and task description, and the expected input format of an image and metadata. As a brief overview, the VLM is instructed to: Provide egocentric instructions instead of cardinal instructions. Detect potential landmarks and include the landmarks’ position relative to the user in the instruction, along with detecting a bounding box to mitigate left/right reversal hallucinations, i.e., landmarks on the left mentioned as being on the right.. If the user is not facing the correct direction, the VLM is instructed to tell the

user to turn and position the landmark on their side or behind them. If the distance to the next turn is too far (> 60 meters), then the VLM is instructed to tell the user to continue forward. To make the user aware of their next turn more quickly, the VLM is instructed to present the turn guidance first, followed by the contextual guidance. Furthermore, the VLM is explicitly instructed to not hallucinate. On each prompt, the model parses based on the system instruction, image, and metadata, and outputs the most likely response.

4 Preliminary Evaluation

This study serves as a preliminary evaluation, aiming to gather an initial understanding of user behavior and experience with our contextual AI and spatial audio system. Participants walked predefined routes (cf. Figure 2a) while following navigation instructions and audio cues. We compare our system, with VLM instruction augmentation and spatial audio (AI-SA), against a version with just VLM instruction augmentation (AI-Only), and a state-of-the-art baseline in the form of Google Maps (audio-only) with detailed walking navigation enabled, (GMaps).

The study used a within-subject design with one independent variable *condition* (AI-SA, AI-Only, and GMaps) in counterbalanced order to limit order effects. To eliminate route learning, we prepared three different routes, although in the same area, which were also counterbalanced.

The three routes were similar in distance (550-650 meters), estimated time (8-9 minutes), included the same number of POIs (5), and a single “dead end” each (see Table 1 for an overview). The dead end required participants to return along the path they came from, exposing a core challenge when routes are not “linear” and include backtracking.

To capture the participants’ point-of-view as best as possible, an Android phone was mounted to a helmet (cf. Figure 2b), simulating the head-worn device form factor, which was especially important for the spatial audio, as the orientation of the phone determines how the spatial audio updates.

Participants were tasked to walk to each of the POIs along the routes and remember the POIs along the route for a pointing task in the end. Notably, participants were not informed about the inclusion of dead ends. Once ready and the study had started, they were told to walk in the direction they believed they should. The experiment conductor walked closely behind the participant to monitor the application, ensuring that the study went smoothly. Whenever the participants walked too far off the route (> 10 meters), the conductor intervened, noted the deviation, and told the participant that they walked in the wrong direction. The participant would then try another direction. Upon reaching the end of the route, the participants pointed with their head towards the beginning POI, then subsequent POIs along the route, with each angular offset being logged to gauge differences in spatial awareness.

4.1 Results

We recruited 12 participants (6 male, 5 female, 1 non-binary), consisting mainly of software engineers. Participants’ ages ranged from 18 to 55. All participants had normal, or corrected-to-normal, vision.

ANOVA was used to analyze normally distributed quantitative data (with Greenhouse-Geisser correction), while Friedman tests (with post-hoc Conover tests) were used for non-normally distributed data (all Holm-Bonferroni corrected). [Figure 2c](#) shows only significant results for brevity, see [Appendix A](#) in supplementary materials for additional graphs and user feedback. Route-level analyses use only the subset of participants assigned to that route under the counterbalanced design, resulting in smaller degrees of freedom. See [Appendix B](#) for all recorded paths annotated with deviations. **Distance Walked**, measured by GPS ($F(2, 22) = 5.017, p = 0.016, \eta^2 = 0.313$), shows that our participants overall walked shorter paths with AI-SA than AI-Only ($p = 0.004$), indicating that they followed the paths more accurately. Particularly on Route 1 ($F(2, 4) = 40.959, p = 0.002, \eta^2 = 0.953$), we found that participants walked longer paths with AI-Only compared to both GMaps ($p = 0.02$) and AI-SA ($p = 0.048$). Participants overall had fewer **Route Deviations** ($> 10m$ off-route) ($\chi^2 = 14.55, p < 0.001, W = 0.606$) with AI-SA than the other conditions (both $p < 0.001$). On Route 1 ($\chi^2 = 6, p = 0.05, W = 1$), AI-Only exhibited more route deviations than both GMaps ($p < 0.001$) and AI-SA ($p < 0.001$), while GMaps exhibited more route deviations than AI-SA ($p < 0.001$). **Pointing Accuracy** was measured to gauge spatial awareness, as the angular error between the participant’s pointing direction and the actual direction of the four POIs from the last POI. We found no differences in Pointing Accuracy overall or any of the routes. Prompt-to-response latency was measured at 3.31s on average ($\pm 0.81s$, from 1.48s to 11.29s).

Participants generally rated AI-SA highest (only once rating the condition 2nd), while AI-Only (rated 3rd three times) was generally rated higher than GMaps (rated 3rd 9 times). Several participants described the spatial audio cue as helpful for verifying if they were on the right track and for orienting themselves. The continuous nature of the spatial audio cue was appreciated by some, making it easier to follow. However, some found the spatial audio cue could be drowned out by other sounds or become easy to ignore if not focused. Participants mentioned appreciating specific and notable

Table 1: An overview of route characteristics. \pm indicates standard deviation. Note that the 180° turn at the dead end is excluded from the mean turn angle.

Measure	Route 1	Route 2	Route 3
Distance	650	550	600
Estimated time	9 min.	8 min.	9 min.
# of intersections	9	10	10
Mean alternative paths	2.78 (± 0.83)	2.5 (± 0.53)	2.4 (± 0.7)
Mean intersection turn angle	67.28° ($\pm 28.95^\circ$)	64.97° ($\pm 23.05^\circ$)	80.65° ($\pm 12.76^\circ$)
Turns	8	9	10

landmarks as they were easy to see and provided clear reference points. However, more common landmarks were often found confusing and unhelpful, making it difficult to identify the correct one. The use of cardinal directions in GMaps was widely perceived as unhelpful, especially without a map or compass at hand. Participants often did not know which direction was north, south, east, or west, leading them to guess.

5 Discussion & Future Work

While the addition of landmarks to navigation instructions has been shown to improve user performance and experience [25], our results are initially not better than the state-of-the-art GMaps, other than route deviations; we saw no differences in pointing accuracy. However, users notably preferred the VLM-augmented instructions (AI-Only and AI-SA) over GMaps, mostly due to the inclusion of landmarks and conversion of cardinal directions to egocentric directions, while especially the spatial audio cue of AI-SA rated highly with participants.

5.1 Route dissimilarities

Despite our efforts to keep the three routes comparable, Route 1 appears more complex, particularly in the AI-Only condition. Post-hoc comparisons, summarized in [Table 1](#) (inspired by [18]), reveal that while Route 1 had fewer intersections and turns, it had a higher average number of alternative paths per intersection. The spatial audio cue inherently disambiguates such choices, whereas turn-by-turn conditions may require participants to think and make assumptions, potentially leading to more deviations. As shown in [Appendix B](#), the dead-end on Route 1 caused notably more deviations in AI-Only than in the other conditions (dead-end deviations: AI-Only = 8, GMaps = 3, AI-SA = 0), likely due to its placement at an intersection (see [Figure 2a](#)), which left participants uncertain about which way to turn.

5.2 Spatial Sound

Our results highlight that the spatial audio cue was critical to keeping users on their path, while the VLM-augmented instructions served more as preemptive information to prepare participants for an upcoming turn. Rather than relying on cardinal or egocentric directions, the spatial audio gives a clear, continuously updating signal for where the next turn is. We therefore believe spatial audio

would play a critical role in display-less scenarios, such as when the phone is pocketed or when wearing smart glasses.

5.3 Choice of landmarks

Participants noted a preference for more unique landmarks, as simply being told to “turn left at the tree on your right” is unhelpful, especially in a park with many trees. We believe that these cases could be mitigated through prompt engineering and task-specific fine-tuning. One participant mentioned wanting more consistency in landmark choice when looking in the same area, which could potentially be achieved by sending a history of recent prompts and responses to the model.

5.4 Hallucinations

Although we instructed the VLM to avoid hallucinations, others have found them inherent to LLMs [24]. Despite attempts to correct left/right reversals via bounding box detection, such errors still occurred, causing participant confusion, as did occasional object misidentification.

5.5 Prompt latency

The average latency of ≈ 3.3 s is largely attributable to Gemini 2.0 Flash prioritizing speed over quality. At an average walking speed of ≈ 1.38 m/s [3], this equates to ≈ 4.5 m of travel, meaning landmarks closer than ≈ 5 m risk causing confusion rather than helping. While manageable on foot, the issue scales significantly at higher speeds: ≈ 18.3 m at 20km/h (biking) and ≈ 45.8 m at 50km/h (city driving), making latency reduction critical for broader use cases.

5.6 Street View Augmentation

The quality of images captured from the user’s perspective greatly affects what the VLM can extract. Existing data from tools like Google Street View could partially address this [4, 9, 17], enabling landmark extraction at route initialization and reducing latency, while also reducing battery consumption for camera-less devices. Prompts could be adapted to favor static landmarks, such as large buildings. However, since Street View may be outdated, landmarks could become unavailable or occluded from the user’s view. Perhaps a hybrid approach could find landmarks online and verify validity in real time.

5.7 Variants of hardware degradation

Although our system requires stereo audio, it could be adapted for reduced functionality to broaden accessibility, including mono audio devices. The core parts needed are a navigation directions source (e.g., Google Maps or OpenStreetMap), a way of capturing the environment (be it in real time with a camera or online through Street View or similar), a low-latency VLM, and text-to-speech. Correct or incorrect facing directions could also be conveyed via auditory icons.

5.8 Limitations

As this is a small exploratory study ($n = 12$), per-route comparisons rely on small subsets and should be interpreted with caution. The

routes could have been more similar, as significant differences between conditions were only indicated on Route 1, though we believe varied layouts are important for capturing real-world variability.

The model used, Gemini 2.0 Flash Exp, has since been deprecated, which limits interpretability. However, we expect newer model versions to perform at least as well given architectural improvements.

Finally, while we report prompt-to-response latency, component-level breakdowns could have offered additional insight, though the prompt-to-response measure captures the major portion of overall latency.

6 Conclusion

In this work, we introduced a novel audio-only navigation system that leverages spatial audio and contextual awareness to improve the user experience. Our findings indicate that the integration of particularly spatial audio cues (with VLM-generated landmark-anchored instructions) leads to improved navigation efficiency compared to VLM-only systems and traditional navigation apps like Google Maps (audio-only), although no improvement to spatial awareness was found. Our users preferred egocentric directions, from the VLM, over cardinal directions, from Google Maps. Similarly, users noted that the landmark-anchored instructions were easier to follow, although mainly when the landmark is unique. This work indicates benefits to be had from utilizing directional and environmental context to better guide users without access to a screen.

References

- [1] Steven Abreu, Tiffany D. Do, Karan Ahuja, Eric J. Gonzalez, Lee Payne, Daniel McDuff, and Mar Gonzalez-Franco. 2024. PARSE-Ego4D: Personal Action Recommendation Suggestions for Egocentric Videos. [doi:10.48550/arXiv.2407.09503](https://arxiv.org/abs/2407.09503)
- [2] Christopher C Berger, Mar Gonzalez-Franco, Ana Tajadura-Jiménez, Dinei Florencio, and Zhengyou Zhang. 2018. Generic HRTFs may be good enough in virtual reality. Improving source localization through cross-modal plasticity. *Frontiers in Neuroscience* 12 (2018), 21. [doi:10.3389/fnins.2018.00021](https://doi.org/10.3389/fnins.2018.00021)
- [3] Richard W Bohannon and A Williams Andrews. 2011. Normal walking speed: a descriptive meta-analysis. *Physiotherapy* 97, 3 (2011), 182–189. [doi:10.1016/j.physio.2010.12.004](https://doi.org/10.1016/j.physio.2010.12.004)
- [4] Howard Chen, Alane Suhr, Dipendra Misra, Noah Snavely, and Yoav Artzi. 2019. TOUCHDOWN: Natural Language Navigation and Spatial Reasoning in Visual Street Environments. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, New York, NY, USA, 12530–12539. [doi:10.1109/CVPR.2019.01282](https://doi.org/10.1109/CVPR.2019.01282)
- [5] Gregory D Clemenson, Antonella Maselli, Alexander J Fiannaca, Amos Miller, and Mar Gonzalez-Franco. 2021. Rethinking GPS navigation: creating cognitive maps through auditory clues. *Scientific reports* 11, 1 (Apr 2021), 10 pages. [doi:10.1038/s41598-021-87148-4](https://doi.org/10.1038/s41598-021-87148-4)
- [6] Louisa Dahmani and Véronique D Bohbot. 2020. Habitual use of GPS negatively impacts spatial memory during self-guided navigation. *Scientific reports* 10, 1 (Apr 2020), 14 pages. [doi:10.1038/s41598-020-62877-0](https://doi.org/10.1038/s41598-020-62877-0)
- [7] Aaron L Gardony, Tad T Brunyé, Caroline R Mahoney, and Holly A Taylor. 2013. How navigational aids impair spatial memory: Evidence for divided attention. *Spatial Cognition & Computation* 13, 4 (2013), 319–350. [doi:10.1080/13875868.2013.792821](https://doi.org/10.1080/13875868.2013.792821)
- [8] Mar Gonzalez-Franco, Gregory D Clemenson, and Amos Miller. 2021. How GPS weakens memory—and what we can do about it. <https://www.scientificamerican.com/article/how-gps-weakens-memory-mdash-and-what-we-can-do-about-it/>
- [9] Karl Moritz Hermann, Mateusz Malinowski, Piotr Mirowski, Andras Banki-Horvath, Keith Anderson, and Raia Hadsell. 2020. Learning to Follow Directions in Street View. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 07 (Apr. 2020), 11773–11781. [doi:10.1609/aaai.v34i07.6849](https://doi.org/10.1609/aaai.v34i07.6849)
- [10] Simon Holland, David R Morse, and Henrik Gedenryd. 2002. AudioGPS: Spatial audio navigation with a minimal attention interface. *Personal and Ubiquitous computing* 6, 4 (Sep 2002), 253–259. [doi:10.1007/s007790200025](https://doi.org/10.1007/s007790200025)
- [11] Toru Ishikawa, Hiromichi Fujiwara, Osamu Imai, and Atsuyuki Okabe. 2008. Wayfinding with a GPS-based mobile navigation system: A comparison with

- maps and direct experience. *Journal of environmental psychology* 28, 1 (2008), 74–82. doi:10.1016/j.jenvp.2007.09.002
- [12] Gilly Leshed, Theresa Velden, Oya Rieger, Blazej Kot, and Phoebe Sengers. 2008. In-car gps navigation: engagement with and disengagement from the environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) (CHI '08). Association for Computing Machinery, New York, NY, USA, 1675–1684. doi:10.1145/1357054.1357316
- [13] Tiffany Liu, Javier Hernandez, Mar Gonzalez-Franco, Antonella Maselli, Melanie Kneisel, Adam Glass, Jarnail Chudge, and Amos Miller. 2022. Characterizing and Predicting Engagement of Blind and Low-Vision People with an Audio-Based Navigation App. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI EA '22). Association for Computing Machinery, New York, NY, USA, Article 411, 7 pages. doi:10.1145/3491101.3519862
- [14] David McGookin, Stephen Brewster, and Pablo Priego. 2009. Audio Bubbles: Employing Non-speech Audio to Support Tourist Wayfinding. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 5763 LNCS (2009), 41–50. doi:10.1007/978-3-642-04076-4_5
- [15] Laura Miola, Veronica Muffato, A Boldrini, Francesca Pazzaglia, and Chiara Meneghetti. 2024. Development of a self-report measure of GPS uses and its relationship with environmental knowledge and self-efficacy and pleasure in exploring. *Cognitive Research: Principles and Implications* 9, 1 (Nov 2024), 78. doi:10.1186/s41235-024-00605-2
- [16] Laura Miola, Veronica Muffato, Enrico Sella, Chiara Meneghetti, and Francesca Pazzaglia. 2024. GPS use and navigation ability: A systematic review and meta-analysis. *Journal of Environmental Psychology* 99 (2024), 102417. doi:10.1016/j.jenvp.2024.102417
- [17] Piotr Mirowski, Andras Banki-Horvath, Keith Anderson, Denis Teplyashin, Karl Moritz Hermann, Mateusz Malinowski, Matthew Koichi Grimes, Karen Simonyan, Koray Kavukcuoglu, Andrew Zisserman, and Raia Hadsell. 2019. The StreetLearn Environment and Dataset. *CoRR* abs/1903.01292 (2019), 13 pages. arXiv:1903.01292 <http://arxiv.org/abs/1903.01292>
- [18] Martin Raubal and Max J Egenhofer. 1998. Comparing the complexity of wayfinding tasks in built environments. *Environment and Planning B: Planning and Design* 25, 6 (1998), 895–913.
- [19] Brian D. Simpson, Douglas S. Brungart, Ronald C. Dallman, Jacque Joffrion, Michael D. Presnar, and Robert H. Gilkey. 2005. Spatial Audio as a Navigation Aid and Attitude Indicator. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 49, 17 (9 2005), 1602–1606. doi:10.1177/154193120504901722
- [20] Simone Spagnol, György Wersényi, Michał Bujacz, Oana Bălan, Marcelo Herrera Martínez, Alin Moldoveanu, and Runar Unnthorsson. 2018. Current Use and Future Perspectives of Spatial Audio Technologies in Electronic Travel Aids. *Wireless Communications and Mobile Computing* 2018, 1 (2018), 3918284. doi:10.1155/2018/3918284
- [21] Steven Strachan, Parisa Eslambolchilar, Roderick Murray-Smith, Stephen Hughes, and Sile O'Modhrain. 2005. GpsTunes: controlling navigation via audio feedback. In *Proceedings of the 7th International Conference on Human Computer Interaction with Mobile Devices & Services* (Salzburg, Austria) (MobileHCI '05). Association for Computing Machinery, New York, NY, USA, 275–278. doi:10.1145/1085777.1085831
- [22] Alexis Topete, Chuanxiuyue He, John Protzko, Jonathan Schooler, and Mary Hegarty. 2024. How is GPS used? Understanding navigation system use and its relation to spatial ability. *Cognitive research: principles and implications* 9, 1 (Mar 2024), 16. doi:10.1186/s41235-024-00545-x
- [23] Nigel Warren, Matt Jones, Steve Jones, and David Bainbridge. 2005. Navigation via continuously adapted music. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems* (Portland, OR, USA) (CHI EA '05). Association for Computing Machinery, New York, NY, USA, 1849–1852. doi:10.1145/1056808.1057038
- [24] Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. 2025. Hallucination is Inevitable: An Innate Limitation of Large Language Models. arXiv:2401.11817 [cs.CL] doi:10.48550/arXiv.2401.11817
- [25] He Zhang, Nicholas J. Falletta, Jingyi Xie, Rui Yu, Sooyeon Lee, Syed Masum Billah, and John M. Carroll. 2025. Enhancing the Travel Experience for People with Visual Impairments through Multimodal Interaction: NaviGPT, A Real-Time AI-Driven Mobile Navigation System. In *Companion Proceedings of the 2025 ACM International Conference on Supporting Group Work*. ACM, New York, NY, USA, 29–35. doi:10.1145/3688828.3699636

A Additional Results

This appendix contains additional results that could not make it into the main paper, including additional figures and user feedback.

A.1 Additional quantitative figures

Figure 3 shows the results for Walked Distance across all routes.

Figure 4 shows the results for Route Deviations across all routes.

Figure 5 shows the results on Pointing Accuracy across all routes.

Figure 6 shows the results for Ranking.

A.2 Additional participant feedback

The following contains additional, and more detailed, participant feedback. Several participants described the spatial audio cue as helpful for verifying if they were on the right track and for orienting themselves. P2 noted that the “[audio cue] is clear about direction. Whenever lost track then can listen for the audio”, and P4 stated it “helped verify that I was pointing in the right direction”. P11 mentioned that it provided “almost constant feedback” on whether they were headed the right way. The continuous nature of the spatial audio cue was appreciated by some, making it easier to follow. P5 noted they rely more on the spatial audio cue because it is faster, and P6 described it as “more accurate, more continuous”. However, some found the spatial audio cue could be drowned out by other sounds or become easy to ignore if not focused. P8 preferred a “louder part of the clip” so it was clearer, and P12 likened it to “elevator music” that could be “zoned out”. P7 mentioned needing to get used to the sound initially. P9 noted that the spatial audio cue helped with direction, while spoken instructions allowed for more fluid movement as it helped them prepare.

Participants mentioned appreciating specific and notable landmarks like a “sign” (P12), “arches” (P1), “Wells Fargo” (P12), “stairs” (P8), or a “silver car” (P12) as they were easy to see and provided clear reference points. Instructions referencing nearby items like a “trash can” were also clear (P5 and P7). However, more common objects/landmarks such as “tree” (P2, P5, P6, P9, P10), “fire hydrant” (P1), “white building” (P1), or “lamppost” (P4, P6) were often found confusing and unhelpful because there could be multiple such objects nearby, making it difficult to identify the correct one. Some participants walked past landmarks without noticing them (P12). Participants preferred prioritizing unique landmarks over common ones (P6, P10). Instructions telling participants to turn towards a specific street were considered helpful when street signs were visible (P7, P12). However, participants mentioned problems occurring when one “does not know the neighborhood” (P11) or “could not see street signs” (P9), making these instructions less helpful.

The use of cardinal directions in GMaps was widely perceived as unhelpful (P1, P2, P5, P6, P8, P10, P11, P12), especially without a map or compass at hand (P10, P11). Participants often did not know which direction north, south, east, or west was, leading them to guess.

Participants had mixed opinions about distance estimates. Some found instructions like “next turn is in how many meters” to be helpful (P5, GMaps). Distance affirmation was also appreciated (P2, P7). However, many participants stated they were not good at judging distances and sometimes felt they had to turn sooner than the indicated distance (P4, P8). P10 mentioned that “If I can be told

that you’ll pass two intersections along the way turn right after the third intersection or third block or whatever that that is much more useful than the distance.”

B GPS Paths

This appendix provides a visual overview of the paths that our participants walked along with annotations for deviations.




Figure 3: Results on the Distance Walked overall and for each route separately.

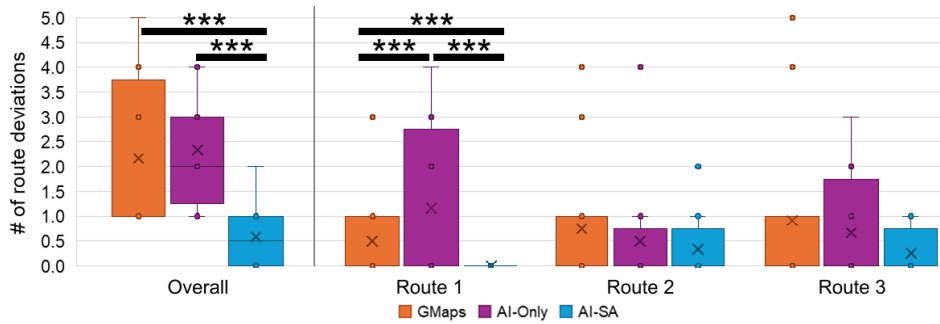


Figure 4: Results on the Number of Deviations overall and for each route separately.




Figure 5: Results on Pointing Accuracy overall and for each route separately.






Figure 6: Results on user preference/rankings of the three conditions.

P1, Route 1, GMaps - Distance: 583 m



P1, Route 2, AI-SA - Distance: 581 m



P1, Route 3, AI-Only - Distance: 746 m






Figure 7: P1's walked paths with 1 deviation in Route 1, 2 in Route 2, and 2 in Route 3.

P2, Route 1, AI-SA - Distance: 554 m



P2, Route 2, AI-Only - Distance: 552 m



P2, Route 3, GMaps - Distance: 803 m






Figure 8: P2's walked paths with 0 deviations in Route 1, 1 in Route 2, and 5 in Route 3.

P3, Route 1, GMaps - Distance: 584 m



P3, Route 2, AI-SA - Distance: 486 m



P3, Route 3, AI-Only - Distance: 702 m





Figure 9: P3's walked paths with 1 deviation in Route 1, 0 in Route 2, and 1 in Route 3.

P4, Route 1, AI-SA - Distance: 559 m



P4, Route 2, AI-Only - Distance: 527 m




P4, Route 3, GMaps - Distance: 388 m




Figure 10: P4's walked paths with 0 deviations in Route 1, 1 in Route 2, and 1 in Route 3. Note that the last part of P4's GPS data is missing.

P5, Route 1, AI-Only - Distance: 684 m



P5, Route 2, GMaps - Distance: 610 m



P5, Route 3, AI-SA - Distance: 642 m






Figure 11: P5's walked paths with 2 deviations in Route 1, 1 in Route 2, and 1 in Route 3.

P6, Route 1, GMaps - Distance: 536 m



P6, Route 2, AI-SA - Distance: 530 m



P6, Route 3, AI-Only - Distance: 675 m






Figure 12: P6's walked paths with 1 deviation in Route 1, 1 in Route 2, and 2 in Route 3.

P7, Route 1, AI-Only - Distance: 662 m



P7, Route 2, AI-SA - Distance: 520 m



P7, Route 3, GMaps - Distance: 752 m






Figure 13: P7's walked paths with 3 deviations in Route 1, 1 in Route 2, and 4 in Route 3.

P8, Route 1, AI-Only - Distance: 618 m



P8, Route 2, GMaps - Distance: 551 m



P8, Route 3, AI-SA - Distance: 603 m






Figure 14: P8's walked paths with 2 deviations in Route 1, 1 in Route 2, and 0 in Route 3.

P9, Route 1, GMaps - Distance: 752 m



P9, Route 2, AI-SA - Distance: 510 m



P9, Route 3, AI-Only - Distance: 740 m




Figure 15: P9's walked paths with 3 deviations in Route 1, 0 in Route 2, and 3 in Route 3.

