Computational Modeling of Non-Visual Vibrotactile Touchscreen Exploration

Saehui Hwang
Mechanical Engineering
Stanford University
Stanford, California, USA
saehui@stanford.edu

Robert J. Moss
Computer Science
Stanford University
Stanford, California, USA
mossr@stanford.edu

Danyang Fan
Mechanical Engineering
Stanford University
Stanford, California, USA
danfan17@stanford.edu

Sean Follmer
Mechanical Engineering
Stanford University
Stanford, California, USA
sfollmer@stanford.edu

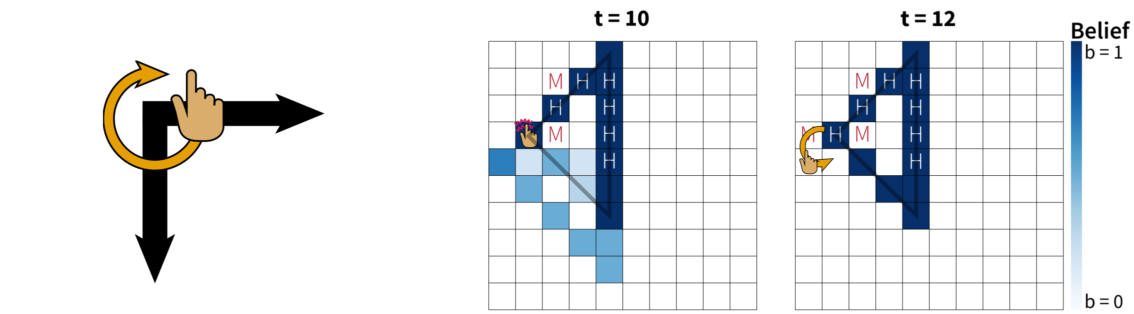

(a) Human behavior: Circling a vertex point. **(b) Learned POMDP behavior: Circling a vertex point.**

Abstract

Using vibration feedback on a touchscreen is a promising method to provide blind and low-vision (BLV) users access to graphical content. While prior studies have explored the design space of vibrotactile rendering of graphics, findings do not generalize to complex shapes, and comprehensive standards for vibrotactile graphics comparable to those for tactile graphics are yet to be defined. To address this gap, we present a computational model for non-visual vibrotactile touchscreen exploration using a partially observable Markov decision process (POMDP) framework. Preliminary simulations of a triangle-tracing task demonstrate that empirically observed exploration strategies, such as circling or crossing around a point, emerge as adaptive behaviors under this framework. The model can further incorporate factors such as memory limitations and observation uncertainty, providing a new approach for analyzing exploration behaviors influenced by environmental and user-specific variables. This framework introduces a tool to understand non-visual exploration strategies and inform vibrotactile graphics design.

CHI EA '25, Yokohama, Japan

ACM ISBN 979-8-4007-1395-8/25/04
https://doi.org/10.1145/3706599.3719851

CCS Concepts

• **Human-centered computing** → **Accessibility design and evaluation methods**.

Keywords

Computational Rationality, Haptic, Accessibility, Touchscreen, Non-Visual Computing

ACM Reference Format:
Saehui Hwang, Robert J. Moss, Danyang Fan, and Sean Follmer. 2025. Computational Modeling of Non-Visual Vibrotactile Touchscreen Exploration. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25), April 26–May 01, 2025, Yokohama, Japan.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3706599.3719851

1 Introduction

Blind and Low-Vision (BLV) users often rely on tactile and touchscreen graphics to consume images and diagrams. Unlike conventional tactile graphics that use physical deformations like raised lines and dots, touchscreen-based graphics are refreshable and use vibrations to convey information on a flat surface. Whereas tactile graphics benefit from well-established design principles codified by organizations like Braille Authority of North America [2], vibrotactile graphics on touchscreens do not have comparable guidelines [10, 11, 17]. Insights from tactile perception studies do not seamlessly transfer to touchscreen-based vibrotactile graphics due to fundamental differences, such as the inability to perceive direction and common reliance on a single point of haptic feedback

[4, 36]. Furthermore, there is limited understanding of which types of graphics are best suited for this medium and how they should be effectively presented.

Modeling can be instrumental in offering rigorous, systematic, and generalizable insights into the complexities of non-visual interaction [3]. In Human-Computer Interaction (HCI), modeling has demonstrated its value in simulating user behavior, assessing interface effectiveness, and refining design concepts. Some work extends these methods to users with impairments. For instance, Touch-WLM examines speed-accuracy trade-offs in text entry for individuals with dyslexia, tremor, or memory dysfunction [31]. Li et al. [21] model how blind users select menu items using auditory feedback alone.

Our work is inspired by the theoretical framework of computational rationality [25], which posits that human behavior emerges as the result of optimizing expected utility under internal constraints (e.g., cognitive and physical limitations) and external constraints (e.g., the physical environment). We propose a computational model of non-visual vibrotactile touchscreen exploration, framing graphic exploration as a partially observable Markov decision process (POMDP) [15]. In this framework, an agent (which simulates a non-visual user) performs sequential actions to trace an underlying shape while relying on haptic feedback (vibrotactile sensations) rather than visual information. The agent maintains probabilistic beliefs about the shape's location, updates these beliefs based on observations, and selects actions to maximize cumulative rewards.

Our preliminary simulations of a triangle tracing task demonstrates that empirically observed exploration strategies, such as crossing or circling around a point [11, 29], emerge naturally as a result of the rationality assumption, even with a simplified framework. This indicates that such strategies may arise as adaptive behaviors to constraints like the absence of vision. Furthermore, the POMDP framework is highly flexible, allowing for the incorporation of additional cognitive factors, such as working memory limitations, observation uncertainty, and prior knowledge through customizing the belief updater, observation model, and prior distribution, respectively. This adaptability allows for the simulation of diverse user experiences, including personal differences in cognitive and sensory abilities.

As a next step, collecting empirical data will be crucial for validating and refining the model's assumptions and parameters. Ultimately, our approach can help researchers and practitioners in HCI deepen their understanding of non-visual interaction, guide the design of accessible interfaces, and evaluate vibrotactile displays for diverse user populations.

2 Related Work

2.1 Vibrotactile Touchscreen Interaction for BLV Users

Vibrotactile rendering is an effective method for providing non-visual feedback to BLV users on touchscreen interfaces [10, 11, 36]. Modern touchscreen devices are inherently multipurpose, multi-sensory, and equipped with accessibility features such as screen readers and gesture interactions, making them a viable platform for BLV-friendly interfaces [4]. Vibrotactile feedback, which uses the built-in vibration motor to provide haptic output, can deliver tactile cues without requiring additional hardware [5] and offer more privacy compared to audio feedback [36].

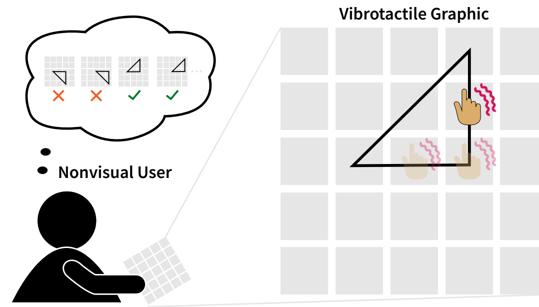

Figure 2: An overview of a non-visual user interacting with a vibrotactile touchscreen. We propose that as the user receives partial (non-visual) information at each time step, her belief of the composition of the underlying shape is updated.

Early systems demonstrated the feasibility of vibrotactile rendering for conveying graphical information, allowing BLV users to interpret road networks [28], graphs [7], shapes [5], and maps [13] through vibration and auditory cues. Vibrotactile touchscreens have been particularly effective for rendering simple graphics [8], enabling users to navigate lines and shapes with higher accuracy and preference compared to audio-only feedback [36].

Unlike embossed maps or pin-array systems, which allow users to directly perceive attributes like a line's orientation, thickness, or elevation through touch, a touchscreen's smooth surface offers no tactile information about these attributes [36]. Furthermore, vibration based touchscreens only support single finger interaction as commercial hardware has only one on-board vibration motor [4]. Therefore, determining a line's orientation requires active finger movement [11, 17].

Due to these differences, graphic materials on touchscreen-based interfaces should be schematized and rendered differently from traditional tactile graphics [27]. However, compared to extensive design guidelines for tactile graphics [2], few comparable resources exist to assist designers in effectively applying vibration to shapes and paths [11]. Addressing these challenges is essential for advancing vibrotactile interface design and improving usability.

Vibrotactile exploration strategies have been identified for simple graphics, such as circling at junctions and making a cross-like motion to avoid straying as shown in fig. 3 [11, 29]. Strategies for raised-line graphics have been shown to vary by graphic type, task focus, and individual user preferences [1]. However, no equivalent insight exists for vibrotactile graphics, highlighting a gap in our understanding of the interface. Our work aims to extend upon these heuristics through a modeling approach with potential to predict and understand strategies for more complex patterns. By building a modeling framework that predicts and explains how these strategies may emerge as rational responses to environmental constraints, our goal is to understand existing heuristics and predict behavior for more complex scenarios.

2.2 Haptic Search Strategies

Haptic search strategies have been empirically studied in non-visual contexts, focusing on tasks such as line tracing, distance comparison, and local and global searches among distractions [26, 39]. Research shows that search performance improves significantly with the use of multiple hands and fingers, particularly for individuals who are blind [23]. Blind individuals often demonstrate superior haptic performance compared to sighted participants, likely due to their ability to effectively leverage multi-hand and multi-finger strategies honed through experience [6]. These strategies include systematic approaches, such as parallel sweeps and spirals, as well as random methods like ballistic or Brownian motion, which are modulated by factors such as task complexity and detection radius [22]. These empirical studies provide valuable insights into simple search tasks but fall short of addressing advanced scenarios involving diverse shapes and user-specific constraints. Our computational model simulates adaptive exploration strategies, offering both a computational explanation for these behaviors and a foundation for addressing further cognitive or physical constraints.

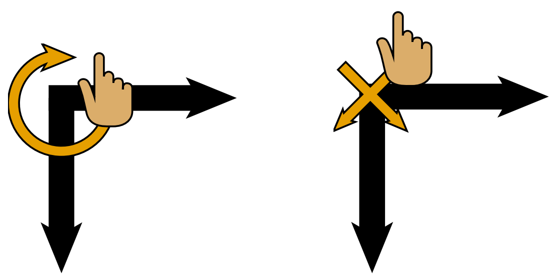

(a) Circling a vertex point **(b) Crossing a vertex point**

Figure 3: Common strategies for non-visual exploration of vibrotactile graphics on a touchscreen. Adapted from Gorlewicz et al. [11]. These strategies are thought to be used to determine a line's orientation.

3 Modeling Non-visual Vibrotactile Exploration as Bounded Optimal Control

The partially observable Markov decision process (POMDP) provides a flexible and powerful framework for modeling sequential decision-making under uncertainty [15, 18]. This approach is particularly well-suited for capturing how users make decisions when interacting with interfaces that offer limited feedback (i.e., partial observability). Like a user interacting with a new shape, the agent in the POMDP model cannot directly observe the true state of the world. Instead, it builds a *belief*—a probabilistic representation of the world—based on partial observations and interactions with small parts of the environment [18].

The following approach was inspired by the words of Lederman and Klatzky [20], "Hand movements can serve as "windows," through which it is possible to learn about the underlying representation ..." We model non-visual vibrotactile touchscreen exploration as a haptic tracing problem in a grid world, where the user sequentially interacts with the touchscreen to uncover and understand the location and shape of the underlying structure. The

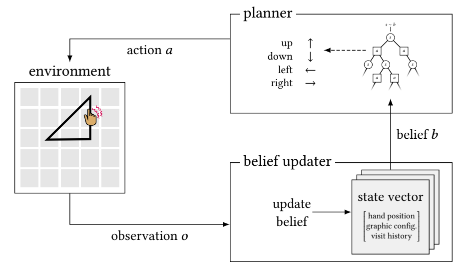

Figure 4: The proposed computational model. An online planner simulates how non-visual users perform the task of following the underlying shape.

shape is a collection of grid cells that would provide vibrational feedback when the position of the agent (non-visual user) coincides with the shape. While these shapes are limited to the resolution of the grid, this is an appropriate simplification given that so far, only basic shapes (lines, triangles, and polygons) have been effectively identified using vibrational feedback [5]. Furthermore, basic shapes form the basis from which more complicated graphics are created [36].

We formally describe the problem as follows:

> Given an n-by-n grid that contains a shape made up of K cells, denoted by $S_{\text{shape}} = \{(x_1, y_1), \ldots, (x_K, y_K)\}$, visit every cell in the shape in the fewest number of steps as possible using movement actions in the cardinal directions.

State Space. At each time step t, the environment is in a state $s \in \mathcal{S}$, which encodes the agent's current position on the grid (x, y), the cell positions comprising the shape $S_{\text{shape}} = (x_1, y_1), \ldots, (x_K, y_K)$, and the agent's history H of visited positions with their corresponding observations. As users often do not know shape and location a priori, the shape configuration is hidden from the agent. We assume that the agent's location and history is deterministically known to the agent, which is an assumption that can be relaxed to reflect a more realistic user, as discussed in section 6.1.

Observation Space and Observation Model. As the agent cannot directly observe where the shape is located, the agent relies on observations limited to a single cell corresponding to whether their finger is vibrating to estimate the location of the shape. The observation space O is defined as: $O = \{\text{vibration}, \text{no vibration}\}$. The observation model $O(s', a, o) = P(o \mid a, s')$ is computed based on the agent's current position and the S_{shape} configuration. Observations are deterministic, signifying that the user is certain of a vibration. Prior research indicates that vibrations are effective at indicating whether a user is "on" or "off" a graphical element [37]. This deterministic assumption can also be relaxed to instead convey uncertainty in the vibration as discussed in section 6.1.

Action Space. The action space $\mathcal{A}$ consists of hand movements in the four cardinal directions: $\mathcal{A} = \{\text{up}, \text{down}, \text{left}, \text{right}\}$. After taking an action, the environment enters a new state according to the state-transition model.

State-Transition Model. The state-transition model, $T(s, a, s') = P(s' \mid s, a)$, is deterministic. Given a cardinal movement action, the new position (x', y') is computed, and history H is updated accordingly. The shape position does not change over time, representing a static graphic on the touchscreen. If all cells composing the underlying shape have been visited, ($S_{\text{shape}} \subseteq H$), the process transitions to a terminal state.

Reward Function. The reward function encodes the agents' preferences and goals. The reward function $R(s, a, s')$ assigns a single large positive reward for visiting all cells that make up the underlying shape (set to 1000 in our experiments). In POMDP planning and reinforcement learning (RL), it is well understood that poorly defined reward functions can lead to unintended behaviors, such as "reward gaming" [34]. To address this, we define the reward to directly reflect the true objective—revealing the full shape—rather than a proxy goal, such as revealing parts of the shape. This ensures that the agent's behavior remains aligned with the actual objective of the task.

Discount Factor. A discount factor $\gamma = 0.9$ is used to prefer immediate rewards over future ones, reflecting the tendency of human decision-making to prioritize short-term outcomes, especially in uncertain or sensory-constrained environments.

3.1 Belief Update

In a POMDP, the belief of the agent can be represented as a probability distribution over the underlying (hidden) state. Observations and actions taken by the agent can be used to update this belief. That is, the user updates where they think the underlying state is based on their observations. For instance, when they encounter a cell with no vibration, they rule out the possibility that the underlying shape involves that cell. We used a particle filter for state estimation (the shape configuration), which can represent a broad range of distributions [38].

While state estimators such as the Kalman Filter [16] are widely used, Kalman filtering requires the assumption that the belief is Gaussian. Instead, particle filtering[9] represents the belief as a collection of state particles that are a plausible approximation of the underlying state. Since the observations are discrete, we use a particle filter with rejection. To update the belief, each particle s_i from the current belief is transitioned to a new state s'_i using the transition model and the selected cardinal action a from the planner (see section 3.2). Observations are then sampled from the observation model using each new state and the selected action, $o_i \sim P(\cdot \mid a, s'_i)$. Each of these observations are compared to the true observation o received from the environment, and the states whose associated observations that do not match the true observation are rejected from the filter.

In our model, the agent has a uniform prior belief over the possible shape configurations. Initial beliefs can be customized to reflect the agent's prior knowledge about the underlying graphic.

3.2 Approximately Optimal Control

To solve the POMDP and get a policy $a = \pi(b)$ that maps beliefs to actions, we used the POMCP algorithm (partially observable Monte Carlo planning), an online planning algorithm based on Monte Carlo tree search [33]. This choice was motivated by the need to handle large state spaces efficiently and the algorithm's suitability for cognitive modeling. Prior research suggests that memory-based and goal-directed simulations play a role in human cognitive processes, particularly in predicting future experiences and guiding actions [32].

4 Experiment

To simulate an agent exploring various shapes, we specify the grid size, the agent's starting position, the ground truth shape configuration, and the agent's initial belief. Triangles were chosen as the ground truth shape for our simulated display, building on findings from recent empirical studies that successfully rendered basic polygons for participants with visual impairments [11] and blindfolded participants [35]. We simulated an agent with prior knowledge of the shape type (right triangle) and size (outline composed of 12 cells), exploring a 10 by 10 grid world through observations. The grid size was chosen to be small for computational efficiency while being large enough to render the shape. A particle filter with 10000 particles was sufficient to prevent particle depletion. The agent's initial belief was represented as a uniform probability distribution across all possible 12-point triangles that could be drawn on the grid. The agent's initial position was defaulted to $(5, 5)$.

5 Preliminary Findings

Empirical studies by Raja [29], later confirmed by Gorlewicz et al. [11], indicate that during non-visual exploration of geometric patterns, users employ an exploration strategy known as circling and crossing, as illustrated in fig. 3. Results from the experiment show that circling and cross strategies are frequently observed as a means to optimally trace the shape (fig. 5). In our model, these patterns naturally emerge as a consequence of optimizing for a single reward criterion—successfully tracing the shape, which provides a reward of +1000—showing that these strategies effectively help reduce uncertainty and narrow beliefs.

Gorlewicz et al. [11] and Tennison and Gorlewicz [35] observed that participants spent the majority of their exploration time at key points of interest, such as vertices. We qualitatively verify such exploratory behavior in fig. 6, through a heatmap over the agent's visit frequency on the grid. The initial agent position was set to $(1, 10)$ to demonstrate the planner's ability to navigate to points of interest.

6 Discussion and Future Work

6.1 Model Enhancement

Conducting an empirical study to compare simulation results with real-world behavior is an essential next step. There are several opportunities for enhancing the model's ability to generalize findings to real-world touchscreen interactions. While the current model assumes small, discrete state/action spaces and deterministic transition/observation models, these can be readily adjusted to capture the complexities of human behavior and cognition. To better reflect touchscreen interactions, we plan to expand the action space for greater movement flexibility and/or increase grid resolution in the state space. At higher grid resolutions, the four-direction movement may be a sufficient approximation to capture curved trajectories.

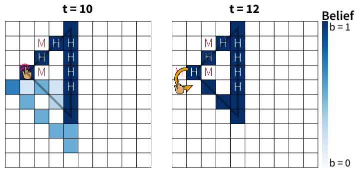

(a) POMDP behavior of circling a vertex point.

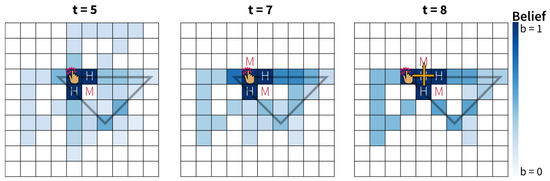

(b) POMDP behavior of crossing a vertex point.

Figure 5: Circling and cross-like strategies are frequently employed to optimally trace the shape. Blue areas indicate the agent's belief about the shape's location, with darker regions representing greater certainty. The agent's movement effectively reduces uncertainty in the shape's location

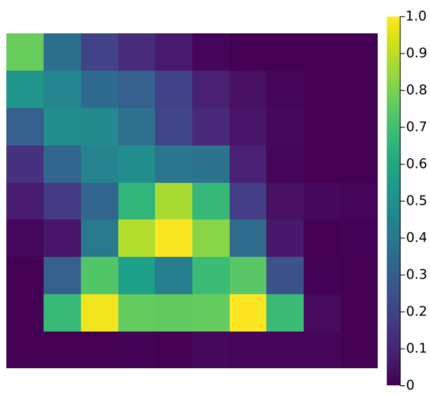

Figure 6: A normalized heatmap of the agent's visit count over the grid world. The agent successfully navigates toward the shape, and higher visit counts are assigned to the vertices of the triangle. This bears similarity with finger position data experimentally collected by Gorlewicz et al. [11].

Higher-resolution grids could enable more complex graphics beyond abstract shapes such as bar charts [40] and street maps [4]. Additionally, we will incorporate sensory fatigue [17, 36] and multi-hand exploration in the observation model. Finally, we propose introducing stochastic transitions to account for the variability in proprioceptive tracking and motor control [17]. To further account for cognitive load, we will integrate memory constraints using custom belief updaters or existing approaches from the literature [14].

6.2 Applications of the Model

This model provides a versatile tool for understanding how users adapt their exploration strategies under various environmental constraints (e.g., grid density, shape complexity) and cognitive limitations (e.g., memory constraints, proprioceptive uncertainty). By simulating these adaptations, the model can help predict the impact of design decisions, such as optimal feedback frequency or layout structure, on user performance and experience. Models like ours can complement human-subject experiments by exploring broader design scenarios quickly.

Modeling in HCI is a powerful approach for explaining variations in user behavior based on individual attributes and preferences [24]. For example, prior studies have observed unexpected strategies in vibrotactile exploration, such as broad scanning versus short scanning patterns [35]. These behaviors could be analyzed as potential adaptations to factors like memory limits, motor constraints, or environmental conditions. Understanding these strategies could inform the development of personalized user interfaces that adapt to individual needs, enhancing usability and accessibility.

6.3 Limitations

The strategies identified in this work are specific to the goal of haptic tracing; however, different contexts may require entirely different strategies. For instance, tasks such as discriminating between shapes, comparing sizes, or identifying key features might not necessitate visiting every part of a shape, leading to alternative exploration behaviors. Nonetheless, with minor modifications to our reward function (which dictates the goal of the agent), such tasks can still be effectively modeled as a POMDP, allowing us to use off-the-shelf solvers to solve them within the same formal sequential decision-making framework.

Additionally, exploring all parts of a graphic does not necessarily equate to understanding it. Advances in cognitive science, such as causal inference [12] and Probabilistic Program Induction [19, 30], provide promising frameworks for understanding processes such as abstraction of concepts in the graphic.

7 Conclusion

In an era where touchscreen devices are ubiquitous, understanding vibrotactile interaction is essential for advancing accessibility and usability [10]. We have proposed a flexible computational model of non-visual vibrotactile touchscreen exploration using a partially observable Markov decision process (POMDP) framework. The proposed framework provides a flexible and generalizable approach for investigating how various environmental factors, such as graphic shape and grid size, and user specific factors, such as memory, and proprioceptive acuity, influence exploration behavior and strategies. Preliminary results reveal that characteristic strategies, such as circling or crossing, naturally emerge from rational decision-making under uncertainty. Without empirical data, we do not claim that our model perfectly represents human cognitive processes. Rather, we argue that even a simplified rational model can capture important aspects of how blind and low-vision (BLV) users might explore non-visual interfaces.

Beyond reproducing well-documented user behaviors, this modeling approach offers a way to systematically probe the impact of internal and environmental constraints of the interface. Ultimately, our approach offers a structured way to understand exploration strategies under various constraints and guide vibrotactile interface design.

Code Availability. The code can be found in the github repository https://github.com/saehuihwang/GraphExploration

Acknowledgments

The authors thank Jean-Peïc Chou, Yujie Tao and Olivia Tomasetti for their thoughtful feedback on this work. We also thank the Stanford Intelligent Systems Laboratory for their development of the POMDPs.jl ecosystem. Saehui Hwang was supported by graduate fellowship awards from Knight-Hennessy Scholars at Stanford University.

References

[1] Sandra Bardot, Marcos Serrano, Bernard Oriola, and Christophe Jouffrais. 2017. Identifying how Visually Impaired People Explore Raised-line Diagrams to Improve the Design of Touch Interfaces. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 550–555. doi:10.1145/3025453.3025582

[2] Braille Authority of North America. 2022. Guidelines and Standards for Tactile Graphics. https://www.brailleauthority.org/guidelines-and-standards-tactile-graphics. Accessed January 24, 2025.

[3] Xiuli Chen, Gilles Bailly, Duncan P Brumby, Antti Oulasvirta, and Andrew Howes. 2015. The emergence of interactive behavior: A model of rational menu search. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 4217–4226.

[4] Nicholas A Giudice, Benjamin A Guenther, Nicholas A Jensen, and Kaitlyn N Haase. 2020. Cognitive mapping without vision: Comparing wayfinding performance after learning from digital touchscreen-based multimodal maps vs. embossed tactile overlays. *Frontiers in Human Neuroscience* 14 (2020), 87.

[5] Nicholas A. Giudice, Hari Prasath Palani, Eric Brenner, and Kevin M. Kramer. 2012. Learning non-visual graphical information using a touch-based vibro-audio interface. In *Proceedings of the 14th international ACM SIGACCESS conference on Computers and accessibility*. ACM, Boulder Colorado USA, 103–110. doi:10.1145/2384916.2384935

[6] Daniel Goldreich and Ingrid M Kanics. 2003. Tactile acuity is enhanced in blindness. *Journal of Neuroscience* 23, 8 (2003), 3439–3445.

[7] Cagatay Goncu and Kim Marriott. 2011. GraVVITAS: generic multi-touch presentation of accessible graphics. In *Human-Computer Interaction–INTERACT 2011: 13th IFIP TC 13 International Conference, Lisbon, Portugal, September 5-9, 2011, Proceedings, Part I 13*. Springer, 30–48.

[8] Cagatay Goncu and Kim Marriott. 2012. Comparison of Haptic and Non-Speech Audio Feedback. In *7th Int. Conference on the Theory and Application of Diagrams (Diagrams' 2012), Canterbury, UK, LNCS*, Vol. 7352. 26–32.

[9] Neil J. Gordon, David J. Salmond, and Adrian F.M. Smith. 1993. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEEE for Radar and Signal Processing*, Vol. 140.

[10] Jenna L Gorlewicz, Jennifer L Tennison, Hari P Palani, and Nicholas A Giudice. 2018. The graphical access challenge for people with visual impairments: Positions and pathways forward. In *Interactive multimedia-multimedia production and digital storytelling*. IntechOpen.

[11] Jenna L Gorlewicz, Jennifer L Tennison, P Merlin Uesbeck, Margaret E Richard, Hari P Palani, Andreas Stefik, Derrick W Smith, and Nicholas A Giudice. 2020. Design guidelines and recommendations for multimodal, touchscreen-based graphics. *ACM Transactions on Accessible Computing (TACCESS)* 13, 3 (2020), 1–30.

[12] Thomas L Griffiths and Joshua B Tenenbaum. 2009. Theory-based causal induction. *Psychological review* 116, 4 (2009), 661.

[13] William Grussenmeyer, Jesel Garcia, and Fang Jiang. 2016. Feasibility of using haptic directions through maps with a tablet and smart watch for people who are blind and visually impaired. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services*. 83–89.

[14] Jussi PP Jokinen, Sayan Sarcar, Antti Oulasvirta, Chaklam Silpasuwanchai, Zhenxin Wang, and Xiangshi Ren. 2017. Modelling learning of new keyboard layouts. In *Proceedings of the 2017 CHI conference on human factors in computing systems*. 4203–4215.

[15] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence* 101, 1-2 (1998), 99–134.

[16] R. E. Kalman. 1960. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering* 82, 1 (March 1960), 35–45. doi:10.1115/1.3662552

[17] Roberta L Klatzky, Nicholas A Giudice, Christopher R Bennett, and Jack M Loomis. 2014. Touch-screen technology for the dynamic display of 2D spatial information without vision: Promise and progress. *Multisensory research* 27, 5-6 (2014), 359–378.

[18] Mykel J. Kochenderfer, Tim A. Wheeler, and Kyle H. Wray. 2022. *Algorithms for Decision Making*. MIT Press.

[19] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. 2015. Human-level concept learning through probabilistic program induction. *Science* 350, 6266 (2015), 1332–1338.

[20] Susan J Lederman and Roberta L Klatzky. 1987. Hand movements: A window into haptic object recognition. *Cognitive psychology* 19, 3 (1987), 342–368.

[21] Zhi Li, Yu-Jung Ko, Aini Putkonen, Shirin Feiz, Vikas Ashok, IV Ramakrishnan, Antti Oulasvirta, and Xiaojun Bi. 2023. Modeling touch-based menu selection performance of blind users via reinforcement learning. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–18.

[22] Valerie S Morash. 2015. Detection radius modulates systematic strategies in unstructured haptic search. In *2015 IEEE World Haptics Conference (WHC)*. IEEE, 1–6.

[23] Valerie S Morash, Allison E Connell Pensky, Steven TW Tseng, and Joshua A Miele. 2014. Effects of using multiple hands and fingers on haptic performance in individuals who are blind. *Perception* 43, 6 (2014), 569–588.

[24] Antti Oulasvirta, Jussi PP Jokinen, and Andrew Howes. 2022. Computational rationality as a theory of interaction. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–14.

[25] Antti Oulasvirta, Jussi P. P. Jokinen, and Andrew Howes. 2022. Computational Rationality as a Theory of Interaction. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 359, 14 pages. doi:10.1145/3491102.3517739

[26] KE Overvliet, JBJ Smeets, and E Brenner. 2007. Haptic search with finger movements: using more fingers does not necessarily reduce search times. *Experimental Brain Research* 182 (2007), 427–434.

[27] Hari Prasath Palani, Jennifer L Tennison, G Bernard Giudice, and Nicholas A Giudice. 2019. Touchscreen-based haptic information access for assisting blind and visually-impaired users: Perceptual parameters and design guidelines. In *Advances in Usability, User Experience and Assistive Technology: Proceedings of the AHFE 2018 International Conferences on Usability & User Experience and Human Factors and Assistive Technology, Held on July 21–25, 2018, in Loews Sapphire Falls Resort at Universal Studios, Orlando, Florida, USA 9*. Springer, 837–847.

[28] Benjamin Poppinga, Charlotte Magnusson, Martin Pielot, and Kirsten Rassmus-Gröhn. 2011. TouchOver map: audio-tactile exploration of interactive maps. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*. 545–550.

[29] Monoj Kumar Raja. 2011. *The development and validation of a new smartphone based non-visual spatial interface for learning indoor layouts*. Ph. D. Dissertation. Citeseer.

[30] Joshua S Rule, Joshua B Tenenbaum, and Steven T Piantadosi. 2020. The child as hacker. *Trends in cognitive sciences* 24, 11 (2020), 900–915.

[31] Sayan Sarcar, Jussi PP Jokinen, Antti Oulasvirta, Zhenxin Wang, Chaklam Silpasuwanchai, and Xiangshi Ren. 2018. Ability-based optimization of touchscreen interactions. *IEEE Pervasive Computing* 17, 1 (2018), 15–26.

[32] Daniel L Schacter, Donna Rose Addis, Demis Hassabis, Victoria C Martin, R Nathan Spreng, and Karl K Szpunar. 2012. The future of memory: remembering, imagining, and the brain. *Neuron* 76, 4 (2012), 677–694.

[33] David Silver and Joel Veness. 2010. Monte-Carlo planning in large POMDPs. *Advances in neural information processing systems* 23 (2010).

[34] Joar Skalse, Nikolaus Howe, Dmitrii Krasheninnikov, and David Krueger. 2022. Defining and Characterizing Reward Gaming. *Advances in Neural Information Processing Systems (NeurIPS)* 35 (2022), 9460–9471.

[35] Jennifer L Tennison and Jenna L Gorlewicz. 2016. Toward non-visual graphics representations on vibratory touchscreens: Shape exploration and identification. In *Haptics: Perception, Devices, Control, and Applications: 10th International Conference, EuroHaptics 2016, London, UK, July 4-7, 2016, Proceedings, Part II 10*. Springer, 384–395.

[36] Jennifer L Tennison and Jenna L Gorlewicz. 2019. Non-visual perception of lines on a multimodal touchscreen tablet. *ACM Transactions on Applied Perception (TAP)* 16, 1 (2019), 1–19.

[37] Jennifer L Tennison, P Merlin Uesbeck, Nicholas A Giudice, Andreas Stefik, Derrick W Smith, and Jenna L Gorlewicz. 2020. Establishing vibration-based tactile line profiles for use in multimodal graphics. *ACM Transactions on Applied Perception (TAP)* 17, 2 (2020), 1–14.

[38] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. 2005. *Probabilistic Robotics*. MIT Press.

[39] Vonne van Polanen, Wouter M Bergmann Tiest, and Astrid ML Kappers. 2011. Movement strategies in a haptic search task. In *2011 IEEE World Haptics Conference*. IEEE, 275–280.

[40] Zhuohao Zhang, John R Thompson, Aditi Shah, Manish Agrawal, Alper Sarikaya, Jacob O Wobbrock, Edward Cutrell, and Bongshin Lee. 2024. ChartA11y: Designing accessible touch experiences of visualizations with blind smartphone users. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*. 1–15.