

How Does Delegation in Social Interaction Evolve Over Time? Navigation with a Robot for Blind People

Rayna Hata
Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
rhata@andrew.cmu.edu

Masaki Kuribayashi
Waseda University
Tokyo, Japan
Miraikan - The National Museum of
Emerging Science and Innovation
Tokyo, Japan
rugbykuribayashi@toki.waseda.jp

Allan Wang
Miraikan - The National Museum of
Emerging Science and Innovation
Tokyo, Japan
allan.wang@jst.go.jp

Hironobu Takagi
IBM Research
Tokyo, Japan
takagih@jp.ibm.com

Chieko Asakawa
Miraikan - The National Museum of
Emerging Science and Innovation
Tokyo, Japan
chiekoa@us.ibm.com



Figure 1: User Delegation vs. User Taking Control. Examples of how users managed social interactions during navigation. In (A), the user delegates to the robot, which asks people with a stroller to move. In (B), the user takes control by speaking and asking the two people to move. In (C), the user again takes control, deciding to move the blocking suitcases without asking for help.

Abstract

Autonomy and independent navigation are vital to daily life but remain challenging for individuals with blindness. Robotic systems can enhance mobility and confidence by providing intelligent navigation assistance. However, fully autonomous systems may reduce users' sense of control, even when they wish to remain actively involved. Although collaboration between user and robot has been recognized as important, little is known about how perceptions of this relationship change with repeated use. We present a repeated exposure study with six blind participants who interacted with a navigation-assistive robot in a real-world museum. Participants completed tasks such as navigating crowds, approaching lines, and

encountering obstacles. Findings show that participants refined their strategies over time, developing clearer preferences about when to rely on the robot versus act independently. This work provides insights into how strategies and preferences evolve with repeated interaction and offers design implications for robots that adapt to user needs over time.

CCS Concepts

• Human-centered computing → User studies; Accessibility systems and tools.

Keywords

Autonomous robots, Delegation, Social navigation, Blind Low-Vision, Assistive technology

ACM Reference Format:

Rayna Hata, Masaki Kuribayashi, Allan Wang, Hironobu Takagi, and Chieko Asakawa. 2026. How Does Delegation in Social Interaction Evolve Over Time? Navigation with a Robot for Blind People. In *Proceedings of the*



This work is licensed under a Creative Commons Attribution 4.0 International License. CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2278-3/26/04
<https://doi.org/10.1145/3772318.3791439>

2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3772318.3791439>

1 Introduction

Autonomy and independent navigation are vital components of daily life, yet they remain challenging for individuals with blindness. Robotic systems have been explored to enhance mobility and confidence by providing autonomous navigation assistance in complex environments [2, 6, 18, 49].

Nonetheless, there are situations where the robot cannot navigate successfully on its own, such as when encountering dense crowds or physical obstacles [50]. Collaboration between the user and the robot is essential for overcoming such challenging scenarios. Shared control, an emerging interaction paradigm in which both user input and robot autonomy jointly shape navigation or task completion decisions, provides a way to address situations that autonomous systems alone cannot resolve [1, 34, 50].

In human-computer interaction (HCI) and human-robot interaction (HRI), shared control has been widely studied as a way to balance efficiency with user agency. For blind users, robot autonomy is often prioritized for convenience and safety; however, prior findings suggest that many also value social autonomy and prefer to remain actively engaged in decision-making and interaction with the surroundings, choosing when and how to rely on robotic support [15].

On the spectrum between user-led and robot-led interaction in navigation-assistive contexts, *delegation* leans toward the robot-led side. For example, when blocked by a crowd, delegation involves the user pressing a button for the robot to announce “Excuse me” (Fig. 1–A), rather than directly asking people to move themselves (Fig. 1–B). Similarly, when an obstacle is blocking the path, delegation means asking the robot to request help from nearby staff (Fig. 1–C), as opposed to the user attempting to move the obstacle independently.

Because deciding what action to take requires an understanding of both the environment and the surrounding social context, prior work has identified several categories of queries made by BLV people [15]. However, many studies in assistive robotics for BLV users, such as [15], have relied on Wizard-of-Oz (WOZ) style descriptions rather than machine-generated ones. This approach is limited, as it is unrealistic to assume that a human operator will always be available in real-world deployments. Moreover, participants may recognize when a human is speaking through the robot, which could influence their responses and reduce the authenticity of the interaction.

Additionally, to evaluate how delegation is reflected in practice, longitudinal studies are necessary. Unlike many assistive robots that are designed to operate autonomously, shared control offers a high level of interactivity, and users may find it difficult to delegate actions to a robot. However, as the novelty effect fades, they may adjust their trust in the robot, refine their willingness to delegate tasks, and develop strategies for interpreting its behavior. Nonetheless, most prior research has examined only short-term or first-time interactions, leaving little understanding of how user perceptions, preferences, and interaction strategies evolve over time. Capturing these longitudinal changes is crucial for designing systems that not

only perform well in controlled trials but also support long-term adoption in everyday contexts.

To examine these dynamics, we developed a robot with a shared control interface and conducted a three-week longitudinal study with six blind participants in a real-world museum setting. The system incorporated interactive features to support joint decision-making, including GPT-based environmental descriptions for querying surroundings, obstacle-specific messages explaining why the robot had stopped, and robot-mediated social communication for negotiating crowds or requesting assistance.

Participants completed navigation tasks such as moving between exhibits, negotiating crowds, approaching lines, and handling physical obstacles. Across the weeks, they engaged with the robot both verbally and physically, making use of its interactive features while also exercising their own judgment about when to act independently. Our study focused in particular on *obstacles at choke points*, *crowds*, and *lines*, as these represent scenarios where autonomous navigation alone is insufficient and some form of social reasoning or interaction is required. Crowds demand negotiation of shared space, obstacles often require assistance from others to be moved, and lines involve subtle reasoning about when and how to engage.

Our findings show that participants shifted from exploratory to purposeful use of the robot’s functions, increasingly delegating social interactions to the robot in noisy or crowded environments while maintaining independence in quieter settings. This pattern contrasts with prior work done in this area, which has reported that participants preferred to take over interactions when faced with shared control, delegation, or tasks that the robot could not resolve on its own. Our results suggest that under certain real-world conditions, users may instead view robot-led interaction as the more effective option, highlighting the importance of context in shaping delegation preferences.

Participants also refined their use of environmental descriptions, moving from curiosity-driven activation to targeted requests for context-specific information. These behaviors demonstrate that delegation and collaboration are dynamic, shaped by both situational demands and evolving user trust. We also observed how participants used the robot to engage with dynamic crowds and real museum visitors, with responses ranging from people stepping aside immediately to cases where requests were misunderstood.

The main contributions of this paper are:

- A three-week longitudinal study with six blind participants using a navigation-assistive robot in a real-world museum setting, providing empirical insights into how perceptions of collaboration and social autonomy evolve with repeated interaction.
- Detailed analysis of participants’ evolving strategies for delegation, including when they chose to rely on the robot to initiate social interactions (asking people to move or requesting help) and when they acted independently, as well as how they used environmental descriptions to interpret surroundings and guide decisions.
- Concrete design implications for navigation-assistive robots, including the need for adaptive levels of support, interactive description systems that enable follow-up questions, and flexible communication styles that accommodate both individual preferences and dynamic social environments.

2 Related Works

2.1 Assistive Navigation Systems and Robots for Blind People

Various navigation systems have been developed to help blind people reach their destinations, typically relying on maps, infrastructure, and localization methods [20, 24]. They have been realized through smartphones [5, 16, 21, 43], wearable devices [29, 30], cane-extensions [35], and more recently, robotic platforms [10, 19, 26, 32, 58]. These robots have gained attention due to their automatic mobility and wayfinding capabilities [54].

Multiple robotic form factors have been explored, including companion-like [7], cane-like [41], quadruped [6, 22, 52, 58], and wheeled designs [10, 32, 58], each offering distinct benefits. Quadrupeds, for example, can navigate stairs and uneven outdoor terrain [3], while many blind users prefer wheeled robots for their stability and lower noise levels [51]. One widely deployed example is the suitcase-style *AI Suitcase* [46].

Even with these advances, navigation robots still face challenges in dense, dynamic environments such as busy public spaces [46]. This difficulty is often described as the “*freezing robot problem*” [48, 50]. Researchers have explored social navigation planners [50], as well as alerting methods to signal pedestrians [17], to mitigate this issue. More recently, shared control has emerged as a promising way to handle situations a robot cannot resolve autonomously, allowing users to directly shape navigation decisions [15, 25, 53].

Building on this line of work, we integrated an enhanced shared-control system that supports both navigation and on-demand situational awareness in challenging social environments. Prior suitcase-style robots provided notifications when being blocked by people and approaching exhibitions [47], but they did not allow users to request environmental information when needed. To address this gap, we incorporated a description feature that generates broad environmental summaries rather than only safety-related notifications. We designed a general-purpose prompt that produces descriptions of objects, people, spatial layout, and overall scene context. This approach was inspired by open-ended exploration tools such as Wanderguide [26], which offer contextual descriptions when user intent is not yet known.

We then used this enhanced system to investigate how blind participants practically engage with shared control in situations where the robot alone cannot resolve the environment, focusing on how users choose between acting independently or relying on robot-mediated social communication.

2.2 Delegation in Social Interactions

Although fully automatic robotic navigation may seem ideal for blind individuals, failures can leave users in situations where they cannot obtain assistance [15]. To address these risks, the shared control paradigm has emerged as a promising interaction method for navigation systems for the blind, as it is known to improve confidence, comfort, and trust [41, 58]. In shared control, user input and robot autonomy are dynamically combined to collaboratively complete navigation tasks.

Within this paradigm, delegation can be understood as a specific form of shared control, where authority is explicitly transferred

between the human and the robot. Rather than blending inputs continuously, the user may delegate subtasks, such as obstacle avoidance or route following, to the robot, while retaining high-level decision-making power. This framing highlights the importance of preserving human agency: the robot should enhance, rather than override, the user’s autonomy [15]. Compared to [15] our work contributes new insights by examining delegation in a real-world setting that involves longitudinal exposure, longer routes, more scenarios, and a fully autonomous system rather than a Wizard of Oz approach. This combination provides a richer understanding of how delegation preferences develop over time and how users negotiate control during everyday navigation tasks.

Prior efforts in the accessibility domain illustrate how such approaches can be applied. For example, shared-control systems have been used to support wheelchair navigation [40, 44, 45], and guide robots have enabled blind users to specify proceeding directions at intersections, while the robot provides automated guidance to the next landmark [13, 14, 25, 27]. Other work has focused on exploration support in unfamiliar environments, enabling blind users to flexibly switch between direct control and automated guidance [25, 26].

These examples reinforce the idea that delegation within shared control is not merely about dividing labor; it is about creating mutual support systems where both the human and the robot actively contribute to successful navigation. Such systems must be designed to preserve decision-making authority while offering reliable assistance, thereby promoting independence without diminishing user control.

2.3 Longitudinal Studies

Longitudinal studies allow participants to gain extended exposure to new technologies, providing insights into how systems are adopted, adapted, and sustained over time. Such studies are particularly valuable for moving beyond short-term novelty effects and capturing how interaction patterns evolve in real-world contexts.

Research on large display systems, VR learning environments, and other interactive technologies has shown that novelty effects are strongest immediately after introduction or when modifications are made, often producing an initial surge in engagement followed by stabilization [9, 12, 23, 28]. These findings highlight that short-term trials may misrepresent long-term adoption, since early behaviors often differ substantially from stable usage patterns. Multi-week deployments are therefore essential for understanding how strategies evolve once novelty diminishes and real-world routines emerge. These prior works show the importance of longitudinal approaches in revealing not only how novelty shapes early adoption but also how user behaviors and perceptions stabilize over time.

Although there is limited work on longitudinal navigation studies with blind participants, prior longitudinal studies in related domains, such as [28], where participants repeated similar tasks across sessions, informed our design choices. However, as noted in [38], the appropriate duration for longitudinal studies remains context dependent and is still considered an open methodological question within HRI.

Despite this, very few studies have applied longitudinal methods to assistive robots [54], particularly in navigation research. Existing efforts have mainly focused on non-navigation contexts, such as classrooms [36], gaming [33], or even non-robotic assistive technologies [42]. In navigation assistance, most evaluations still rely on short-term trials lasting only a few hours [54]. Although some projects develop navigation systems iteratively through hardware updates and deployments in large facilities [46], individual longitudinal familiarization and adaptation remain underexplored.

In this work, we address this gap by conducting a three session repeated-exposure study across three weeks with blind participants using a navigation-assistive robot. Our focus is on how delegation in social interactions evolves over time, going beyond single-session familiarization to investigate how participants balance independence with robot assistance in repeated, real-world use.

3 Longitudinal Study Methods

In this work, we conducted a three-week longitudinal study with six blind individuals, each of whom interacted with the navigation robot across three separate weeks. We first describe our robotic system, which adopts a shared control interface, followed by details of the study.

3.1 Robot Function and Interaction

We used an open-source robotic system [47], modified to serve as the guide robot for this study (Fig. 2a). The robot is equipped with LiDAR sensors and three RGB-D cameras (front, left, and right) to support navigation and obstacle detection. Its handle interface integrates a touch sensor and four buttons (top, bottom, left, and right), allowing participants to issue commands and receive haptic feedback. A smartphone mounted on the back of the robot communicated with a neck speaker worn by participants, which delivered GPT-generated obstacle detections and environmental descriptions. An external speaker on the robot was used for preset social phrases when participants delegated social interactions to the robot.

3.1.1 Surrounding GPT. The **top button** allowed participants to request a description of their surroundings, which we refer to as “*Surrounding GPT*”. Pressing this button captured a frame from the front RGBD camera and submitted it to GPT using the prompt “*Please describe this scene.*”. The prior version of the robot [47] did not provide users with a way to request information about their surroundings. It only delivered notifications about nearby obstacles and approaching exhibits, without accepting any user-initiated queries. To address this limitation, we introduced Surrounding GPT as a method for offering general situational awareness rather than only safety related alerts. We designed a generic prompt that encouraged the model to describe the scene broadly, informed by the principles of Wanderguide [26], which supports exploration by providing open-ended contextual descriptions when user intent is not yet known. This resulted in outputs that included visible objects, people, spatial layout, and higher level semantic impressions of the environment.

The **right button** enabled participants to delegate social interaction by having the robot say “*Excuse me, please move.*”, while the **left button** prompted the robot to vocalize “*Excuse me, please help me.*”. Together, these functions gave participants multiple ways to

manage social interactions, supporting both independent action and delegation to the robot depending on the situation. This feature was developed on top of the open sourced robot platform [47]. While the robot’s compact, seamless design allows users to blend naturally into their environment, this advantage also introduces limitations not typically encountered with more visibly identifiable navigation tools such as a white cane or guide dog, since bystanders may not immediately recognize the device or know to make space or offer assistance.

3.1.2 Obstacle Detection. The robot includes a built-in obstacle detection system that classifies obstacles into general categories such as *person*, *object in front*, or *unknown obstacle*. To enrich this capability, we developed an additional layer that provides more detailed contextual information. When one of the three categories is detected, the system captures a frame from the front RGB camera and sends it to GPT-4o along with the following prompt:

You are an assistant guiding a person with visual impairment. The robot is currently stopped for one of the following reasons: (1) there is a person in front, (2) it is trying to avoid a person, or (3) it has detected an obstacle. Look at the image and, following the rules below, state a concise one-sentence explanation corresponding to the reason for the stop.

(1) If there is a person in front, or (2) if the robot is trying to avoid a person: state only the approximate number of people, whether they are in a line or a crowd, and whether they are facing toward you or facing away. Example: “There is a crowd ahead with about three people, facing away.”

(3) If the robot has detected an obstacle: briefly state what the obstacle is. Example: “The robot has stopped due to an obstacle. There is a cone ahead.”

Please strictly follow the rules below:

- Do not provide any other information or descriptions.
- Do not use vague expressions or apologies such as “I do not know,” “I cannot judge,” or “I am sorry.”
- Do not include opinions or subjectivity; state only objective facts.
- Keep the explanation as short and concise as possible.

The prompt was provided in the local language to ensure that the generated output was also in that language.

3.1.3 Navigation and Interaction. Participants navigated by holding onto the robot’s handle. Once a destination was set through the Bluetooth-connected smartphone, participants indicated readiness to begin moving by pressing the right button. The robot then autonomously navigated to the destination, following waypoints and avoiding obstacles and other visitors.

In situations where the robot could not move past people, such as static crowds or physical obstacles, either the robot or the participant needed to initiate a social interaction with those nearby to resolve the problem. Participants were informed that they could either use the robot’s buttons to have the system ask people to

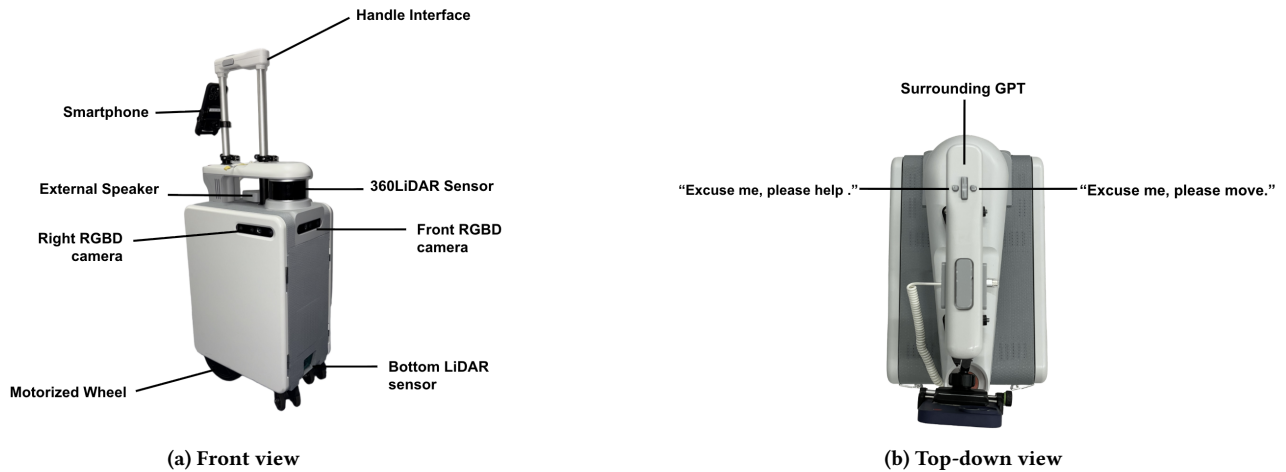


Figure 2: Guide Robot. The suitcase-shaped robot used in the study. It integrates a handle interface, smartphone, and external speaker for user interaction. It includes multimodal sensing through front and right RGBD cameras, a 360 LiDAR, a bottom LiDAR, and motorized wheels for navigation.

move or request assistance from nearby individuals, or they could choose to manage these interactions independently.

3.2 User Study Detail

Using our robot system, we conducted a three-week study with six blind participants, which was approved by our institutional review board (IRB).

3.2.1 Recruitment and Participants. The study was advertised to the local blind and low-vision community through a mailing list affiliated with our institution. Six participants (three women and three men) between the ages of 18 and 69 were recruited ($M = 42.0$, $SD = 23.6$). Table 1 presents a detailed breakdown of participant demographics, including age, gender, vision level, duration of impairment, navigation tools used, and frequency of independent outings. Given the difficulty of recruiting blind participants for long-term studies, our system was evaluated with six users. This sample size aligns with usability research indicating that a single study with approximately five users can reveal about 75% of usability issues [25, 26, 37].

3.2.2 Scenarios & Routes. We designed each route intentionally to expose participants to a diverse set of navigation challenges. The routes included narrow hallways, long corridors, chokepoints, queues, and crowded areas, allowing us to emulate the types of situations blind travelers commonly encounter in public indoor environments.

To guide this design process, we developed a set of principles and guidelines informed by a prior work on social navigation [8]. These guidelines emphasized creating situations where users would need to negotiate space, manage congestion, and determine when to rely on the robot for assistance versus when to take direct control.

We expanded and adapted the routes explored in [15]. Based on these guidelines, our routes were significantly longer and more complex than those used in prior work. The repeated-exposure

structure of our study allowed participants to experience multiple instances of social navigation patterns across different spaces. This enabled us to build on established navigation patterns while incorporating more complex and extended paths appropriate for longitudinal evaluation. By integrating both structured challenges and open-ended environments, the routes offered a comprehensive testbed for observing delegation behaviors and shared control in realistic conditions.

We designed three distinct scenarios in which participants were required to take an action to address the situation, following the framework established in a prior work [15].

- (1) **Crowd scenario:** A static crowd blocked the path. Participants could either speak directly or press the right button for the robot to say “*Excuse me, please move.*”
- (2) **Obstacle scenario:** An object blocked the path. Participants could press the left button for the robot to request help or do so themselves.
- (3) **Line scenario:** A line of people formed at the destination. Participants decided whether to initiate the interaction themselves or delegate it to the robot.

Fig. 4 illustrates the routes and locations where the scenarios were conducted over the three weeks. The *crowd scenario* was created by two experimenters who stood close together to form a choke point. They remained stationary unless the participant acted, either by commanding the robot to speak or verbally asking them to step aside.

For the *obstacle scenarios*, three different obstacles were introduced, one per week (Fig. 3). In the first week, a simple grass-like foam object was used. This object was taken directly from a museum exhibit to simulate a pop-up exhibition. In the second week, two suitcases were placed.

In the final week, a baby stroller accompanied by a man looking at his phone was positioned. These suitcase and stroller obstacles



(a) “The robot has stopped due to an obstacle. There is a green obstacle ahead.”



(b) “The robot has stopped due to an obstacle. There are suitcases ahead.”



(c) “The robot has stopped due to an obstacle. There is a baby stroller ahead.”

Figure 3: Route Object Obstacles. Examples of object obstacles encountered during navigation with GPT-generated descriptions, as captured by the robot’s front RGBD camera. Week 1 included a green foam block from a children’s play exhibit, Week 2 featured two suitcases, and Week 3 presented a baby cart.

Table 1: Participant demographics.

ID	Age	Gender	Vision Level	Impairment Duration	Navigation Tools	Independent Outings
P1	21	Man	Blind	>10 years	White cane	Every day
P2	66	Woman	Blind	>10 years	Guide dog, apps	Every day
P3	69	Man	Blind	>10 years	White cane, apps, guide	Every day (guide 3–4 times/month)
P4	18	Woman	Blind	<5 years	White cane, GPS apps	~5 times/week
P5	54	Man	Blind	>10 years	White cane, BeMyAI	Every day
P6	24	Woman	Low vision	>10 years	White cane, apps	About once/week

were selected because they are common items brought into the museum by visitors and therefore reflect realistic, everyday conditions. Unless participants attempted to move the obstacle, asked for assistance, or requested the stroller to be moved, the objects remained in place. An experimenter acted as a nearby bystander, ready to assist only when asked. Although the robot was deployed during regular museum hours and participants encountered natural crowds, some obstacles needed to be introduced because museum staff routinely cleared unattended objects. Adding controlled obstacles ensured that all participants experienced comparable navigation challenges in consistent locations, allowing fair comparison while still situating the study within a naturalistic environment.

The *line scenarios* varied across the three weeks. In the first and third weeks, a line blocked the area near the destination, preventing the robot from reaching it. Because robots often misclassified lines as crowds, participants were required to seek help from surrounding people to interpret the situation and wait patiently. The line moved forward one by one only after the participant recognized it as a line by asking the nearby bystander. In the second week, the line directly blocked the participant’s path, requiring them to ask people in line to move aside so they could pass. Each line consisted of two to three researchers. Importantly, in all weeks, participants also encountered real dynamic crowds of museum visitors, as the study took place during regular museum hours.


3.2.3 Procedure. Participants attended three sessions (approximately two hours each), scheduled once per week, and spent about 60 minutes per session walking routes with the robot. The overall

procedure remained consistent across the three weeks, with only minor details varying. Although there is limited work on longitudinal navigation studies with blind participants, our approach was informed by longitudinal studies in related domains where participants repeated similar tasks across sessions [28, 39], and the appropriate duration for such studies remains context dependent and an open question within HRI [38]. Prior works show that stable conditions help reduce novelty effects and allow participants to focus on using the system rather than adapting to new elements each week, leading to more meaningful feedback about the system itself.

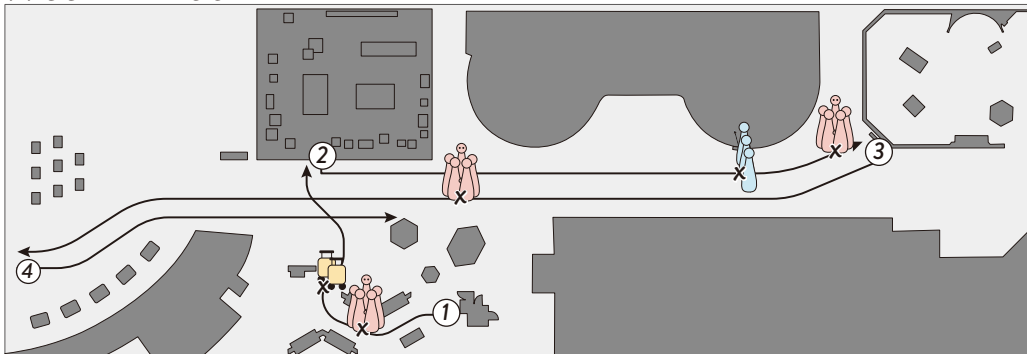
In the first session, participants completed consent forms and provided demographic information. Each week began with a training session. During Week 1, training lasted about 30 minutes to ensure participants became familiar with all robot functions. This included practicing walking with the handle, using the buttons, and experiencing each of the three scenarios while trying all possible actions. In Weeks 2 and 3, participants completed a shorter refresher training session to help them recall the robot’s functions.

After training, participants walked three predetermined routes during regular museum hours, encountering both staged and natural obstacles, as shown in Fig. 4. The labels refer to the planned obstacles and crowd scenarios introduced by the researchers. Each session concluded with a brief post-session interview about their experiences and reflections, which will be described in the following section.

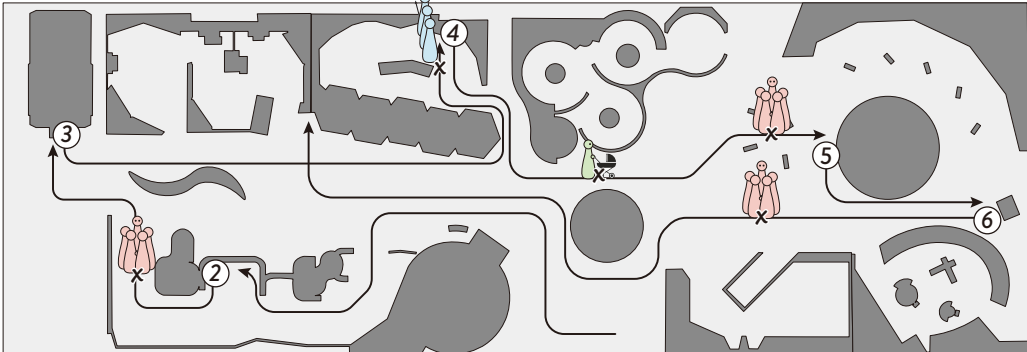
Week 1 Floor A



Week 2 Floor B



Week 3 Floor A








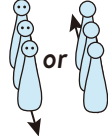
-  Route N Start
-  Object
-  Luggage
-  Babycart
-  Crowd
-  Line

Figure 4: Navigation Routes. Routes used in the study as mapped on museum floors. Week 1 and Week 3 were conducted on Floor A, while Week 2 was on Floor B. Across the routes, participants encountered a variety of obstacles, including crowds, lines, luggage, and baby carts.

Participants were not informed that certain scenarios would be performed by actors, and were asked to respond as they naturally would in a real-world setting. Actors were trained to follow consistent scripts for obstacle placement and crowd behavior, ensuring comparability across participants. The robot’s spoken output from the neck speaker was generally only audible to individuals standing very close to the participant, meaning that surrounding visitors were unlikely to hear or react to the robot’s requests unless they

were directly in the interaction space. This helped preserve the naturalistic feel of the environment while still allowing us to study how participants navigated semi-controlled situations.

3.2.4 Post-Session Interviews. At the end of each session, participants completed a short interview combining standardized survey measures and study-specific questions. The survey included the Robotic Social Attributes Scale (RoSAS) [4] and five Likert-scale

items developed for this study, each followed by an open-ended prompt to capture participants' reasoning.

The RoSAS was used to assess participants' perceptions of the robot's social attributes, with emphasis on Factor 2 (Competence) and Factor 3 (Discomfort). Participants rated each attribute on a 7-point Likert scale (1 = strongly disagree, 7 = strongly agree):

- **Competence-related:** capable, responsive, interactive, reliable, competent, knowledgeable
- **Discomfort-related:** scary, strange, awkward, aggressive, dangerous

Study-specific Likert items addressed delegation and environmental descriptions:

- (1) When the robot was stopped, I thought it was better for the robot (rather than myself) to ask people to move.
- (2) When the robot was stopped, I thought it was better for the robot (rather than myself) to ask for help.
- (3) I was able to understand my surroundings using the robot's descriptions.
- (4) I thought the way the robot explained its surroundings was appropriate.
- (5) Even in difficult situations, I was able to confidently decide how to act.

In weeks 2 and 3, participants were reminded of their earlier responses and asked if their views had changed. In the final week, they reflected on their overall experience and discussed how their perceptions of collaboration with the robot evolved over the three weeks.

4 Results

4.1 Delegating Social Interaction to the Robot

Overall, participants increasingly delegated social interactions to the robot over time, though individual trajectories varied, with some showing steady growth and others fluctuating across weeks.

Table 2: Delegation of Social Communication. Percentage of interactions in which participants delegated communication tasks (such as asking for help or requesting others to move) to the robot across three weeks. P2–P5 showed consistently high delegation, while P1 and P6 displayed more variability, with P6 gradually increasing reliance on the robot over time. These patterns highlight both stable and evolving strategies of delegation across participants.

Participant	Week 1	Week 2	Week 3
P1	0.0%	57.1%	12.5%
P2	83.3%	87.5%	92.6%
P3	85.7%	100%	100%
P4	100%	100%	100%
P5	20.0%	100%	100%
P6	0.0%	20.0%	54.5%

Table 2 shows the percentage of interactions in which participants chose to delegate social communication to the robot rather

than speaking themselves. This percentage was calculated by dividing the number of times a participant delegated the social interaction by the total number of times they were alerted to an obstacle.

Participants who began with low levels of delegation (P1, P5, P6) generally increased their reliance on the robot over time, while those who started high (P2, P3, P4) maintained consistently high levels across all three weeks. Participants showed similar delegation behaviors in both actor-created scenarios and naturally occurring situations (shown in Table 5).

For P1, delegation rose from none in Week 1 to 57.1% in Week 2, before falling back to 12.5% in Week 3. This fluctuation reflects his preference to act independently when possible, while still relying on the robot in particularly noisy or crowded contexts. By contrast, P2 maintained high delegation rates throughout (83.3%–92.6%), consistent with her reasoning that the robot's voice carried better in noisy environments:

“At first, I thought I should say it myself and I did, but it felt like people weren't really hearing me. With the robot, I felt like people were actually paying attention.” (P2, Week 3)

Similarly, P3 consistently delegated (85.7%–100%), noting that the robot's synthetic voice, while not natural, was often more effective at drawing attention than his own:

“I felt that the robot's voice carried better than my own. Especially in situations with lots of people and a noisy atmosphere, I thought the robot's voice was more likely to be heard than if I spoke myself.” (P3, Week 2)

P4 delegated 100% of the time across all weeks, which aligned with her self-described reluctance to directly ask for help in public situations:

“It's just my personality, but I'm not very good at calling out to people myself. For example, when I'm outside and get lost, I'm not the type who can easily say, 'Excuse me, could you help me?' In those situations, if the robot had that function, it would really help.” (P4, Week 1)

P5 showed one of the most dramatic increases, rising from 20% in Week 1 to 100% in Weeks 2 and 3. This trajectory suggests growing trust and comfort in letting the robot handle social interactions on his behalf.

Finally, P6 demonstrated a gradual increase, moving from 0% in Week 1 to 54.5% in Week 3. She explained that when she was uncertain whether people were nearby, the robot's voice provided a more reliable and less risky way to initiate interaction.

These trajectories show how quantitative patterns in delegation were closely tied to participants' lived experiences. Some participants, like P1 and P6, delegated selectively based on situational demands such as noise or uncertainty, while others, like P2, P3, and P4, consistently relied on the robot to overcome personal or environmental barriers to social interaction.

4.2 Use of Environmental Descriptions

Table 3 shows how often participants activated the “Surrounding GPT” function across the three weeks of the study. The numbers shown reflect the number of times each participant requested information about their surroundings by pressing the top button. Use varied across participants due to individual differences, but overall patterns followed a trajectory from exploration to purposeful engagement. Participants who began with little or no use (P1) gradually incorporated the function into their routines, while others (P6) relied on it heavily and consistently. Some, such as P2 and P4, experimented extensively in Week 2 before becoming more selective in Week 3, suggesting a shift from curiosity-driven exploration to targeted, experience-informed use. P2 and P4 also had increased exposure to obstacle-related GPT triggers during later sessions, which may have influenced their patterns of use.

Table 3: Use of Surrounding GPT. Number of times participants triggered the “Surrounding GPT” function across three weeks. Usage patterns varied widely: P2 and P4 peaked in Week 2, P6 consistently used the function most, while P5 discontinued after Week 1. These patterns reflect a broader shift from exploratory use to more purposeful engagement over time.

Participant	Week 1	Week 2	Week 3
P1	0	2	2
P2	1	8	4
P3	2	0	3
P4	2	8	1
P5	2	0	0
P6	7	5	6

For instance, P6 was the most consistent user, activating the function 7, 5, and 6 times across the three weeks. Her comments reflected this steady reliance:

“While walking, I often wanted a simple overview of what kinds of exhibits were around me, so I pressed the top button to hear the situation explained.” (P6, Week 3)

P1, by contrast, never used the feature in Week 1 but incorporated it twice per session in Weeks 2 and 3, showing how initially hesitant users grew to see value in the descriptions. Meanwhile, P2’s activations rose sharply from one in Week 1 to eight in Week 2 before dropping to four in Week 3. She explained that her interest shifted from general crowd awareness to exhibit information:

“When I’m walking through an area, I want to know the situation with the people around me. But when I want to know what kind of place it is, then it’s the exhibit information that I’m interested in.” (P2, Week 3)

Several participants also used the descriptions strategically when the robot repeatedly announced that a crowd was blocking its path. In these situations, they sought additional context via Surrounding GPT before deciding whether to act themselves or delegate to the

robot. Participants explained that when they received multiple obstacle alerts in quick succession, often due to a continuous flow of dynamic crowds, they activated Surrounding GPT to understand what was happening around them before choosing how to respond. For example, P6 in Week 3 frequently requested a description just before instructing the robot to ask the crowd to move.

Overall, participants were satisfied with the GPT-based descriptions, appreciating that they were concise and actionable. At the same time, preferences diverged: some felt the outputs were occasionally too broad and expressed interest in a more interactive system that would allow follow-up questions, while others valued the broader overviews as a way of situating themselves in the environment [26].

These results show that GPT-generated descriptions played a complementary role in collaboration: they not only provided situational awareness but also influenced participants’ decision-making about when and how to intervene, thereby supporting calibration of shared responsibility between user and robot.

4.3 Responses to Obstacle-Related GPT Triggers

Table 4 shows the number of obstacle-related GPT triggers and the percentage of those instances in which participants took action, either by using the robot or by acting themselves. To complement the percentage-based analysis in Table 4, we additionally report the raw counts of naturally occurring obstacles and instances where participants became truly stuck in Table 5. This table shows, for each participant and week, how many of the total Obstacle GPT messages were triggered by real environmental conditions such as dynamic crowds, stationary groups near exhibits, or bottleneck areas within the museum. It also shows how many times participants were unable to proceed for an extended period and either waited or initiated a social interaction. During trials with high crowd densities, we observe more frequent obstacle GPT calls such as in P2 Week 3, and P6 Week 3.

Overall, participants did not act on every alert, but instead developed strategies to decide when intervention was warranted. For example, P1 encountered 29 obstacle triggers in Week 2 but acted in only 24% of cases, often waiting to see if the robot would resolve the situation itself. By Week 3, although he encountered fewer triggers (19), his action rate rose to 42%, reflecting greater willingness to step in when the robot stalled. Similarly, P2’s triggers rose sharply from 10 in Week 1 to 47 in Week 3, yet her action percentage remained stable (60% to 57%). She explained that she often interpreted the robot’s small corrective motions—such as backing up or shifting slightly—as signs it was still searching for a path, and therefore chose to wait:

“When the robot stopped and moved a bit this way and that, I thought it was searching—looking for an open path—so I watched it that way.” (P2, Week 3)

P6 showed the opposite trajectory: she acted on 71% of seven triggers in Week 1, but her intervention rate dropped to 39% when she faced 28 triggers in Week 3. This suggests that as exposure increased, she became more selective about when to intervene, relying on the robot unless it appeared fully stuck. P4 described a similar decision-making process, noting that uncertainty about

Table 4: Obstacle-Related GPT Usage and Actions. Number of obstacle-related GPT triggers and corresponding action percentages across three weeks. Usage varied widely: P2 and P6 showed the highest number of triggers by Week 3, while action percentages were generally moderate, with P6 showing the highest early adoption but a decline over time, and P2, P4, and P5 reaching higher action rates by Week 3. These findings suggest that participants interpreted the robot’s signals and intervened only when it appeared stuck.

Participant	Week 1		Week 2		Week 3	
	Obstacle GPT	Action %	Obstacle GPT	Action %	Obstacle GPT	Action %
P1	15	27%	29	24%	19	42%
P2	10	60%	18	44%	47	57%
P3	12	58%	11	55%	18	44%
P4	12	33%	16	44%	14	57%
P5	11	45%	15	40%	14	57%
P6	7	71%	12	42%	28	39%

Table 5: Natural occurrences of obstacle alerts and instances where participants became stuck or chose to trigger an interaction. For each week, GPT counts the number of times Obstacle GPT was triggered by naturally occurring obstacles or crowds, and Action counts the number of times participants either remained stuck for a long time or decided they could not wait and triggered a social interaction.

Participant	Week 1		Week 2		Week 3	
	GPT	Action	GPT	Action	GPT	Action
P1	9	1	13	1	13	2
P2	6	0	15	2	42	15
P3	7	2	3	1	15	1
P4	6	0	6	1	8	2
P5	3	1	6	1	7	2
P6	3	0	7	0	21	7

nearby people often prompted her to act after waiting through multiple alerts:

“When it was moving in a way that seemed a bit uncertain—like stopping and then moving again—or when it seemed like there might be people nearby but I wasn’t sure, those were the times I used it.” (P4, Week 3)

Because crowd density varied by day and time, the number of obstacle triggers was not consistent across the sessions, and in some cases repeated alerts were generated for a single situation. This inflated the counts without necessarily prompting more actions.

These findings show that participants learned to interpret the robot’s signals and movements in order to calibrate their own involvement. Rather than treating every obstacle alert as a call to action, they assessed whether the robot was still attempting to resolve the situation or truly needed assistance, highlighting the importance of making the robot’s progress and state more transparent.

4.4 Learning and Adaptation Over Time

Over the course of the three weeks, participants noted changes in their understanding of the robot and its functions. They described how they began to anticipate and interpret the robot’s movements as well as adapt their use of its interactive features.

P2 mentioned that when the robot stopped, she learned to wait before expecting it to move again. She explained that when the robot stopped and made small corrective motions, she gradually recognized that it was searching for a way around an obstacle:

“It was the way the robot moved when there were people around. The small movements, like backing up a little, moving forward, backing up again, turning left, and trying different directions— I gradually came to understand that it was searching for a path.” (P2, Week 2)

P1 similarly reported that he was able to anticipate the robot’s movements and speed, particularly when it backed up. In addition to learning how to interpret the robot’s behavior, participants also changed how they approached delegating tasks to it. P2 reflected that in the first week she pressed the buttons simply to try them

out, but by the final week she was making decisions based on accumulated experience.

P1 initially resisted delegating to the robot, emphasizing his desire to act independently. In the first week, he reported relying only on himself:

“I think part of it comes from my personality, but for me it was like... from the start, my way of thinking has always been that—even though I can’t see—I’ve always wanted to do the same things as sighted people. Ever since I was little, I’ve felt that way. So since I can call out to people myself, I don’t really feel the need to rely on the robot if it’s something I can just do on my own. Because of that, I think I ended up doing all the calling out myself.” (P1, Week 1)

By the second week, however, P1 began using the buttons to delegate to the robot in noisy and crowded settings:

“The reason was that it was crowded today, with a lot of noise around, so my voice often didn’t reach people.” (P1, Week 2)

P2 and P3 expressed similar progressions, shifting from initial reliance on themselves to greater use of the robot in situations where their own voices were less effective.

P6 explained that in the first week she was so accustomed to asking people to move on her own that she did not think to use the robot. In later weeks, however, she began to delegate more to the robot:

“When it was noisy, I felt it was easier for people to notice and help because the robot spoke louder than my own voice.” (P6, Week 2)

P5 also described a similar shift. He noted that he preferred using his own voice when only a few people (two or three) were present, but found it difficult in larger crowds or outdoor environments. By the second and third weeks, he reported completely delegating to the robot, explaining that he had grown more accustomed to it, trusted it more, and preferred to let the robot manage these interactions.

As participants became more familiar with the robot, their use of environmental descriptions also changed. For example, P2 reported using the description button at first because she, in her words, “just wanted to try it out.”

4.5 Study-Specific Likert Questions and RoSAS

Participants’ responses to the four study-specific Likert-scale questions revealed clear trends over the course of the three weeks.

For statements (1) and (2), which asked whether participants preferred the robot to initiate social interactions (i.e., asking people to move or asking for help) instead of doing so themselves, we observed a shift toward stronger agreement over time (Fig. 5 and Fig. 6). Linear mixed-effects analyses for Statements 1 and 2 showed clear positive effects of Week on participants’ ratings. Week was a highly significant predictor for both items (Statement 1: $F(2, 10) = 13.57$, $p = .0016$; Statement 2: $F(2, 10) = 13.21$, $p = .0014$), indicating strong increases in participants’ comfort with the robot initiating social interactions over time.

For statement (3), which asked whether participants were able to understand their surroundings using the robot’s descriptions, ratings also increased across the three weeks (Fig. 7). Similarly, for statement (4), which asked whether the robot’s explanations were appropriate, participants expressed growing agreement (Fig. 8). Linear mixed-effects analyses showed that Week was also a significant predictor for Statement 3 ($F(2, 10) = 8.45$, $p = .0071$), indicating that participants increasingly felt able to understand their surroundings using the robot’s descriptions. For Statement 4, the effect of Week was positive but not significant ($F(2, 10) = 2.50$, $p = .13$), showing a modest but less consistent increase in perceived appropriateness of the robot’s explanations.

We observed a trend in responses to statement (5), which asked whether participants felt able to act confidently while using the robot. Overall, most participants reported increased confidence over time, though there was some variability across individuals. P6 remained consistent across all three weeks, while P1 showed a decrease in confidence in Week 3. Notably, P5 shifted from being neutral in Week 1 to strongly agreeing by Week 3, indicating a substantial gain in confidence. Similarly, P4 moved from a neutral position in Week 1 to moderate agreement in Weeks 2 and 3. P2 and P3 both maintained stable confidence levels in Weeks 1 and 2, but trended upward in Week 3 (Fig. 9). However, it is important to note that a mixed-effects model found that Week was not a significant predictor of confidence ($F(2, 10) = 1.25$, $p = .33$). This indicates that although several participants showed descriptive increases in confidence over time, these changes were not consistent enough across the group to produce a reliable overall trend.

In the answers to the RoSAS scale questions, participants were generally positive about the robot. On average, they showed high agreement with positive attributes such as being interactive, reliable, responsive, competent, and knowledgeable, and high disagreement with negative attributes such as scary, strange, awful, awkward, dangerous, and aggressive.

For the answers to the RoSAS scale, there was no clear trend of consistent increase or decrease across weeks. Instead, some participants’ ratings fluctuated, going up in one week and down in another. Several of these changes were tied to specific circumstances participants encountered during that session, such as larger crowds, the presence of more foreign visitors, or the perception of younger children running around in the environment.

4.6 Dynamic Crowds and Real-World Occurrences

In addition to the obstacles and crowds staged by actors, we observed several natural interactions between participants and other museum visitors. In some cases, people noticed the robot approaching and moved aside on their own. When visitors had their backs turned and did not see the robot, participants often relied on the robot’s voice to request that people move, which typically prompted a response.

There were also instances where visitors did not understand the robot’s request. Several participants noted that the museum had a large number of foreign visitors who were not fluent in the local language, raising concerns that the robot’s verbal prompts might go unnoticed or unrecognized. Participants suggested that the system




Figure 5: Ratings for Statement (1): Preference for the Robot to Ask for Help. Participant responses over three weeks showed a general upward trend, with P1, P4, and P5 exhibiting the strongest increases, indicating growing acceptance of the robot actively requesting assistance.




Figure 7: Ratings for Statement (3): Ability to Understand Surroundings Using the Robot's Descriptions. Participant responses across three weeks showed consistently high or increasing agreement. P2 and P6 reported maximum scores early on, while P5 demonstrated the largest improvement over time.

could be improved by offering the ability to switch languages or by using additional non-verbal signals, such as auditory cues (e.g., beeps), to attract attention and ensure that surrounding people understood the robot's presence and intention.

4.7 Unanticipated Obstacles and User Responses

Several participants experienced unexpected behavior from the robot due to limitations in its path planning. In multiple instances during Week 3 on Route 3, the robot struggled to navigate around a slanted wall that angled toward the walkway. The robot's sensors




Figure 6: Ratings for Statement (2): Preference for the Robot to Ask People to Move. Participant responses across three weeks showed steadily increasing agreement, with P1, P2, P4, and P5 exhibiting the largest gains, reflecting growing acceptance of robots directly requesting people to clear the way.




Figure 8: Ratings for Statement (4): Appropriateness of the Robot's Environmental Explanations. Participant responses across three weeks were generally high, with P1 consistently rating the maximum, P4 showing the largest upward trend, and only P6 decreasing slightly from 7 to 6.

did not detect the wall until it was too close, causing it to become completely stuck. In these cases, the GPT-based obstacle detection repeatedly announced: "There is an obstacle. There is a wall."

After calling for help, either via the robot or by using their own voice, participants then had to verbally indicate that they were stuck and request assistance. Actors approached and manually moved the robot only after participants explained the situation. During the post-study interviews, several participants reflected that if this occurred in a real-world setting, they would feel uncomfortable asking strangers for help, especially given the robot's weight. While




Figure 9: Ratings for Statement (5): Participants’ Confidence in Acting While Using the Robot. Responses over three weeks showed variation across individuals. P5 reported the strongest increase in confidence, P6 remained consistently high, while P1 declined slightly and others showed gradual improvements.

they acknowledged the difficulty of such situations, participants were uncertain how they might manage them outside the study environment.

4.8 Continued Usage

An important aspect of our longitudinal study was understanding how participants’ perceptions and comfort with the robot changed over multiple sessions. Experiencing the system across several days gave participants a clearer sense of how it might fit into their everyday routines. During the post-study interviews in the final week, we asked whether they would want to continue using the robot and in what contexts.

Participants offered varied perspectives. Several noted that they would not use the robot constantly but preferred it for situations that are more challenging or cognitively demanding. For example, P1 reported that the robot would be helpful in “crowded or complex places, like a mall,” where locating items and moving safely can be difficult. Others mentioned supermarkets, where dynamic obstacles and tight layouts often present challenges. One participant (P6) suggested the robot would be useful in environments that require careful navigation, such as “somewhere like a museum where you need to be careful about the objects around you.” Others described interest in using the robot in outdoor markets or small shop areas, where paths are irregular and human assistance may be inconsistent.

While participants did not view the robot as an all-purpose daily companion, they recognized clear situational value. Because the study exposed them to a variety of social navigation scenarios, participants were able to consider how the robot could support different navigation demands in their daily lives. Overall, participants expressed interest in using the robot selectively in more demanding

contexts, highlighting its potential role as a *contextual support tool* rather than a constant assistant.

5 Discussion

5.1 Delegating Social Interaction to the Robot

The results demonstrate that participants’ delegation to the robot was shaped both by individual differences and situational demands, as evidenced by many other similar systems [54]. Participants who began with high levels of delegation (P2, P3, P4) maintained this reliance across all three weeks, suggesting that for some, the robot immediately served as a valued partner in managing social interactions. Others, such as P1, P5, and P6, initially resisted delegating to the robot, emphasizing independence or relying on long-standing habits of asking for help themselves. However, over time these participants exhibited tendencies to delegate to the robot, particularly in noisy or crowded contexts where their own voices were less effective. While users could be given ways to amplify their own voices, several participants explained that they preferred not to initiate social interactions themselves. For example, P4 described the robot as a way to manage the situation without having to speak directly. Providing the option to delegate therefore supported both users who wanted to handle interactions independently and those who preferred the robot to take the lead. Prior work has noted that blind individuals often negotiate boundaries to avoid drawing unwanted attention from those around them [55]. In contrast, our findings show that the robot provided a way for participants to call attention to themselves in order to accomplish a task, thereby allowing them to sidestep the discomfort of initiating a direct social interaction. This progression illustrates a process of trust-building and pragmatic adaptation, in which participants came to view the robot as a reliable partner for managing social interactions.

Importantly, participants emphasized that the robot’s synthetic voice, though not natural, was effective because it projected loudly and captured attention. Prior work similarly demonstrates that synthetic or machine-generated voices, when well-designed, can match or even outperform human voices in noisy environments by improving clarity, intelligibility, and attention capture [11].

However, some participants raised concerns about phrasing and comprehension in multilingual or international settings, noting that visitors who did not understand the language sometimes failed to respond. These observations highlight the need for adaptable voice styles and language flexibility in future systems.

These findings show why delegation must be examined in longitudinal settings rather than single-session studies. Initial behaviors did not always reflect stable preferences; some participants who began with strong independence later adopted selective delegation as they gained experience, while delegation-first participants remained consistent across weeks. Longitudinal study designs reveal how preferences stabilize or shift as users build confidence in the system and understand its capabilities in real-world conditions.

Across participants, we observed three broad personality types shaping user delegation behavior:

- (1) **Independent-first participants.** Individuals such as P1 and P6 described themselves as wanting to maintain independence and “do everything that sighted people can do.” They were initially reluctant to delegate to the robot, preferring to

take initiative themselves. However, with longer exposure they recognized scenarios where their own voices were less effective, such as in noisy or crowded spaces, and began to selectively use the robot.

- (2) **Balanced participants.** A second group showed a more flexible approach from the outset, combining independent action with willingness to delegate to the robot when appropriate. These participants expressed comfort switching between modes from the start; however, they tended to use the robot more at the beginning.
- (3) **Delegation-first participants.** Participants like P4, who described themselves as shy or uncomfortable with asking strangers for help, relied on the robot consistently. For them, the robot lowered the social barrier to requesting assistance, enabling interactions they might otherwise have avoided.

These personas show that user delegation to the robot is both situational and personality-dependent: while some participants required time and repeated exposure before embracing the robot, others valued its role immediately as a social intermediary. The diversity of strategies highlights the importance of designing navigation-assistive robots that can flexibly support different interaction styles and evolve with users over time.

A prior study that examined user preferences for receiving environmental information highlighted the need to investigate how shared control between humans and robots might function when information is provided by an automated system rather than by a human in a Wizard-of-Oz (WoZ) setting [15]. In our study, we found that when participants queried general information about their surroundings using GPT-generated descriptions, they typically viewed this feature as most useful for gaining a broad sense of the environment or the overall “vibe,” rather than for highly localized details about what was directly in front of them. At the same time, we observed that some participants, particularly P6, began to repurpose the “Surrounding GPT” feature over time, increasingly activating it after being alerted to an obstacle. This shift suggests that the function not only supported orientation but also became a strategic tool for interpreting the robot’s behavior and deciding when to intervene.

Across the study, early behaviors did not always match later preferences. P1, who initially rejected delegation, later found the robot more effective at asking others to move, while P4 consistently preferred the robot to take the lead. Balanced participants remained flexible throughout. Despite these differences, all participants eventually expressed confidence in the robot’s ability to handle social requests, showing that repeated exposure strengthens trust and reliance beyond initial personality-based tendencies.

5.2 Learning and Adaptation Over Time

Prior work has highlighted the importance of examining long-term use when investigating how blind users interact with navigation-assistive robots, noting that preferences and strategies may evolve as familiarity increases [15]. Our findings build on this by showing that user adaptation was not simply a byproduct of repeated exposure, but instead reflected an evolving process of learning when and how to collaborate effectively with the robot. Our results extend prior work by showing how longitudinal exposure shapes

users’ understanding of the robot and alters their viewpoints over time. Repeated sessions revealed that participants often revised their initial opinions as they gained experience, demonstrating that preferences observed in one-time encounters may not reflect stable judgments.

Participants shifted from exploratory use of the robot’s features to more deliberate, context-sensitive strategies. For example, P1 initially rejected delegation to the robot altogether, framing independence as a matter of personal identity, yet later employed the robot selectively in noisy settings (Fig. 5 and Fig. 6). Similarly, P6 moved from no delegation to the robot in Week 1 to steadily greater reliance by Week 3 (Table 2), while P5 transitioned from primarily using his own voice to relying almost entirely on the robot. These trajectories illustrate how experiential learning, combined with situational pressures such as crowd density and background noise, fostered gradual trust in the robot’s social capabilities and shifted the boundaries of when users considered it advantageous to delegate social interactions to the robot.

Participants also reported developing interpretive frameworks for the robot’s behavior, such as recognizing corrective movements (e.g., backing up, pivoting, repeating small motions) as signs that the robot was “searching” for a path rather than fully stuck. This interpretability directly shaped collaboration: rather than intervening immediately, some participants chose to wait until the robot appeared unable to progress, suggesting that trust was dynamically calibrated in response to perceived robot effort (Table 4). In parallel, use of the “Surrounding GPT” function evolved from curiosity-driven exploration in early sessions to targeted requests for context-specific information in later sessions (Table 3). This reflects a broader process of experiential learning in which participants refined both their delegation strategies and their sense of when environmental descriptions provided actionable support.

Confidence ratings add further nuance to these findings (Fig. 9). While participants such as P5 and P4 reported increased confidence over time, attributing this to growing familiarity with the robot’s movements and descriptions, others (P1, P6) remained stable or showed declines in Week 3. These fluctuations underscore the influence of situational variability: differences in crowd density, background noise, and the unpredictability of museum environments all shaped how confident participants felt in acting alongside the robot. Importantly, confidence appeared to emerge from a combination of two processes: developing interpretive understanding of the robot’s signals and receiving timely, interpretable information about the surrounding environment.

Taken together, these findings suggest that collaboration with navigation-assistive robots is a dynamic, context-sensitive process shaped by both user learning and environmental contingencies. Adaptation occurred not only through repeated interaction but also through participants’ active construction of mental models about the robot’s behavior. Designing systems that explicitly support this interpretive process—by making corrective movements more transparent, providing adjustable levels of descriptive detail, and scaffolding confidence across longer-term use—may strengthen trust, sustain user agency, and ultimately promote real-world adoption.

5.3 Limitations and Implications

Our findings show that collaboration between user and robot is a dynamic process shaped by trust, learning, and situational demands. This underscores the importance of long-term, real-world evaluation: participants needed time not only to learn how the system functioned but also to refine their own preferences. Unplanned events in the museum, such as language mismatches, unexpected obstacles, and mapping errors, created challenges unlikely to appear in controlled settings.

In our system, the robot attempted to reroute when possible, but many scenarios intentionally focused on moments where rerouting was not feasible, reflecting the freezing robot problem and the need for user decision-making during social navigation. However, even if future systems can autonomously plan alternative paths around obstacles or crowded areas [31, 56], and not face the freezing robot problem, such rerouting may result in longer travel times or paths that do not align with a user's preferences. Users may still want control over where they go and how they navigate. Future work could examine whether, in longitudinal settings, participants would choose to rely on the robot to move a crowd rather than accept a longer detour. Our findings identified several scenarios in which blind participants preferred having the robot ask people to move, and showed how this preference varied across personality types and situational demands. Building on this, future systems could offer both options in real time, allowing users to compare the cost of a detour with the effort of engaging socially through the robot.

Several limitations emerged. First, the robot often struggled to distinguish between lines and crowds, leaving participants unsure whether they should wait or initiate an interaction. Prior work similarly notes that identifying socially meaningful group formations from a single image using off-the-shelf LLMs or vision-language models remains a challenging open problem in accessibility and computer vision, particularly when subtle motion cues are needed to interpret intent [57]. These difficulties help explain why the system could not consistently communicate such distinctions to participants.

Second, although participants appreciated GPT-generated summaries, the system did not support interactive question-and-answer functionality. We intentionally did not implement Q&A because LLM inference latency, combined with the hardware constraints of an on-device system and inconsistent network speeds needed for online processing, would have caused noticeable delays during navigation. However, interactive querying remains a promising direction for future work as models become faster, lighter, and better suited to real-time deployment on embedded platforms.

Third, unplanned obstacles occasionally caused the robot to become physically stuck (for example, against slanted walls or tight corner geometries). Participants noted that manually repositioning a heavy robot would be impractical, highlighting the need for more robust recovery mechanisms such as autonomous rerouting or more reliable assistance prompts.

Finally, language barriers sometimes reduced effectiveness in multilingual environments. Even so, participants reported that the robot's louder and more noticeable voice often helped draw attention better than their own speech, even when bystander comprehension was not guaranteed [17].

These limitations point to clear design directions. Navigation-assistive robots should (1) make obstacle states and corrective actions more transparent, (2) support user-driven queries for adaptive levels of detail, (3) provide flexible, multilingual communication, and (4) incorporate recovery strategies for unresolvable obstacles.

6 Conclusion

This work examined how blind participants engaged with a navigation-assistive robot over three weeks, focusing on how delegation and collaboration evolved through repeated interaction. Participants shifted from exploratory use to more deliberate strategies, selectively delegating social interactions to the robot and using GPT-generated descriptions to guide decision-making. These behaviors were shaped by both individual preferences and situational demands such as noise, crowd density, and the robot's corrective movements.

Our findings show that delegation is dynamic and adaptive: users develop trust, form mental models of the robot's behavior, and refine strategies over time. **While some users may prefer to handle social interactions independently, others may consistently rely on the robot, or adapt their approach depending on the circumstances.** Allowing for this flexibility is critical, as preferences vary widely across individuals and contexts. Therefore, our study demonstrates that assistive navigation systems for blind individuals should be evaluated in longitudinal settings, since preferences and usage patterns evolve over time and may differ substantially from those observed in first-time encounters.

Future work should build adaptive systems that personalize support, for example by enabling Q&A-style environmental descriptions, context-sensitive speech, or models that learn from prior user behavior. Extending this research to longer deployments and broader contexts such as outdoor navigation or public transit will further test how these dynamics scale. Additionally, equipping models that are more accurate at detecting obstacles may influence delegation behavior and should be further investigated.

Taken together, these findings show that the design of navigation-assistive robots must go beyond technical autonomy alone. Instead, robots must be built as adaptive partners that negotiate control fluidly with their users, supporting independence while remaining responsive to context. By prioritizing flexibility, transparency, and personalization, future systems can more effectively empower blind individuals and promote confident, long-term use.

References

- [1] Negin Amirshirzad, Asiye Kumru, and Erhan Oztop. 2019. Human adaptation to human-robot shared control. *IEEE Transactions on Human-Machine Systems* 49, 2 (2019), 126–136.
- [2] Mayara Bonani, Raquel Oliveira, Filipa Correia, André Rodrigues, Tiago Guerreiro, and Ana Paiva. 2018. What My Eyes Can't See, A Robot Can Show Me: Exploring the Collaboration Between Blind People and Robots. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility* (Galway, Ireland) (ASSETS '18). Association for Computing Machinery, New York, NY, USA, 15–27. doi:10.1145/3234695.3239330
- [3] Shaojun Cai, Ashwin Ram, Zhengtai Gou, Mohd Alqama Wasim Shaikh, Yu-An Chen, Yingjia Wan, Kotaro Hara, Shengdong Zhao, and David Hsu. 2024. Navigating Real-World Challenges: A Quadruped Robot Guiding System for Visually Impaired People in Diverse Environments. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 44, 18 pages. doi:10.1145/3613904.3642227

- [4] Colleen M. Carpinella, Alisa B. Wyman, Michael A. Perez, and Steven J. Stroessner. 2017. The Robotic Social Attributes Scale (RoSAS): Development and Validation. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction* (Vienna, Austria) (HRI '17). Association for Computing Machinery, New York, NY, USA, 254–262. doi:10.1145/2909824.3020208
- [5] Hsuan-Eng Chen, Yi-Ying Lin, Chien-Hsing Chen, and I-Fang Wang. 2015. Blind-Navi: A Navigation App for the Visually Impaired Smartphone User. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI EA '15). Association for Computing Machinery, New York, NY, USA, 19–24. doi:10.1145/2702613.2726953
- [6] Stacy A. Doore, Narit Trikasemsak, Alexandra Gillespie, Andrea Giudice, and Rayna Hata. 2024. Co-designing an Accessible Quadruped Navigation Assistant. In *2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*, 659–666. doi:10.1109/RO-MAN60168.2024.10731241
- [7] Catherine Feng, Shiru Azenkot, and Maya Cakmak. 2015. Designing a Robot Guide for Blind People in Indoor Environments. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts* (Portland, Oregon, USA) (HRI'15 Extended Abstracts). Association for Computing Machinery, New York, NY, USA, 107–108. doi:10.1145/2701973.2702060
- [8] Anthony Francis, Claudia Pérez-d'Arpino, Chengshu Li, Fei Xia, Alexandre Alahi, Rachid Alami, Aniket Bera, Abhijit Biswas, Joydeep Biswas, Rohan Chandra, et al. 2025. Principles and guidelines for evaluating social robot navigation algorithms. *ACM Transactions on Human-Robot Interaction* 14, 2 (2025), 1–65.
- [9] Sarah Gallacher, Jenny O'Connor, Jon Bird, Yvonne Rogers, Licia Capra, Daniel Harrison, and Paul Marshall. 2015. Mood Squeezer: Lightening up the Workplace through Playful and Lightweight Interactions. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (Vancouver, BC, Canada) (CSCW '15). Association for Computing Machinery, New York, NY, USA, 891–902. doi:10.1145/2675133.2675170
- [10] João Guerreiro, Daisuke Sato, Saki Asakawa, Huixu Dong, Kris M. Kitani, and Chieko Asakawa. 2019. CaBot: Designing and Evaluating an Autonomous Navigation Robot for Blind People. In *Proceedings of the 21st International ACM SIGACCESS Conference on Computers and Accessibility* (Pittsburgh, PA, USA) (ASSETS '19). Association for Computing Machinery, New York, NY, USA, 68–82. doi:10.1145/3308561.3353771
- [11] S. Hennig and R. Chellali. 2012. Expressive synthetic voices: Considerations for human robot interaction. In *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*. 589–595. doi:10.1109/ROMAN.2012.6343815
- [12] Elaine M Huang, Elizabeth D Mynatt, Daniel M Russell, and Alison E Sue. 2006. Secrets to success and fatal flaws: The design of large-display groupware. *IEEE Computer Graphics and Applications* 26, 1 (2006), 37–45.
- [13] Hochul Hwang, Tim Xia, Ibrahim Keita, Ken Suzuki, Joydeep Biswas, Sunghoon I. Lee, and Donghyun Kim. 2022. System Configuration and Navigation of a Guide Dog Robot: Toward Animal Guide Dog-Level Guiding Work.
- [14] Gaurav Jain, Yuanyang Teng, Dong Heon Cho, Yunhao Xing, Maryam Aziz, and Brian A. Smith. 2023. "I Want to Figure Things Out": Supporting Exploration in Navigation for People with Visual Impairments. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1, Article 63 (April 2023), 28 pages. doi:10.1145/3579496
- [15] Rie Kamikubo, Seita Kayukawa, Yuka Kaniwa, Allan Wang, Hermisa Kacorri, Hironobu Takagi, and Chieko Asakawa. 2025. Beyond Omakase: Designing Shared Control for Navigation Robots with Blind People. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 671, 17 pages. doi:10.1145/3706598.3714112
- [16] Yuka Kaniwa, Masaki Kuribayashi, Seita Kayukawa, Daisuke Sato, Hironobu Takagi, Chieko Asakawa, and Shigeo Morishima. 2024. ChitChatGuide: Conversational Interaction Using Large Language Models for Assisting People with Visual Impairments to Explore a Shopping Mall. *Proc. ACM Hum.-Comput. Interact.* 8, MHCI, Article 247 (Sept. 2024), 25 pages. doi:10.1145/3676492
- [17] Seita Kayukawa, Keita Higuchi, João Guerreiro, Shigeo Morishima, Yoichi Sato, Kris Kitani, and Chieko Asakawa. 2019. BBeep: A Sonic Collision Avoidance System for Blind Travellers and Nearby Pedestrians. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300282
- [18] Seita Kayukawa, Tatsuya Ishihara, Hironobu Takagi, Shigeo Morishima, and Chieko Asakawa. 2020. BlindPilot: A Robotic Local Navigation System that Leads Blind People to a Landmark Object. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–9. doi:10.1145/3334480.3382925
- [19] Seita Kayukawa, Daisuke Sato, Masayuki Murata, Tatsuya Ishihara, Akihiro Kosugi, Hironobu Takagi, Shigeo Morishima, and Chieko Asakawa. 2022. How Users, Facility Managers, and Bystanders Perceive and Accept a Navigation Robot for Visually Impaired People in Public Buildings. In *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. 546–553. doi:10.1109/RO-MAN53752.2022.9900717
- [20] Sulaiman Khan, Shah Nazir, and Habib Ullah Khan. 2021. Analysis of Navigation Assistants for Blind and Visually Impaired People: A Systematic Review. *IEEE Access* 9 (2021), 26712–26734. doi:10.1109/ACCESS.2021.3052415
- [21] Jee-Eun Kim, Masahiro Bessho, Shinsuke Kobayashi, Noboru Koshizuka, and Ken Sakamura. 2016. Navigating visually impaired travelers in a large train station using smartphone and bluetooth low energy. In *Proceedings of the 31st Annual ACM Symposium on Applied Computing* (Pisa, Italy) (SAC '16). Association for Computing Machinery, New York, NY, USA, 604–611. doi:10.1145/2851613.2851716
- [22] J. Taery Kim, Wenhao Yu, Yash Kothari, Jie Tan, Greg Turk, and Sehoon Ha. 2023. Transforming a Quadruped Into a Guide Robot for the Visually Impaired: Formalizing Wayfinding, Interaction Modeling, and Safety Mechanism.
- [23] Michael Koch, Kai von Luck, Jan Schwarzer, and Susanne Draheim. 2018. The Novelty Effect in Large Display Deployments – Experiences and Lessons-Learned for Evaluating Prototypes. In *Proceedings of 16th European Conference on Computer-Supported Cooperative Work - Exploratory Papers*. European Society for Socially Embedded Technologies (EUSSET). doi:10.18420/escsw2018_3
- [24] Bineeth Kuriakose, Raju Shrestha, and Frode Eika Sandnes. 2022. Tools and Technologies for Blind and Visually Impaired Navigation Support: A Review. *IETE Technical Review* 39, 1 (2022), 3–18. doi:10.1080/02564602.2020.1819893
- [25] Masaki Kuribayashi, Tatsuya Ishihara, Daisuke Sato, Jayakorn Vongkulbhisal, Karnik Ram, Seita Kayukawa, Hironobu Takagi, Shigeo Morishima, and Chieko Asakawa. 2023. PathFinder: Designing a Map-less Navigation System for Blind People in Unfamiliar Buildings. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 41, 16 pages. doi:10.1145/3544548.3580687
- [26] Masaki Kuribayashi, Kohei Uehara, Allan Wang, Shigeo Morishima, and Chieko Asakawa. 2025. WanderGuide: Indoor Map-less Robotic Guide for Exploration by Blind People. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 676, 21 pages. doi:10.1145/3706598.3713788
- [27] Gerard Lacey and Shane MacNamara. 2000. Context-Aware Shared Control of a Robot Mobility Aid for the Elderly Blind. *The International Journal of Robotics Research* 19, 11 (2000), 1054–1065. doi:10.1177/02783640022067968
- [28] Joomi Lee, Chen Crystal Chen, and Aryabrata Basu. 2025. From Novelty to Knowledge: A Longitudinal Investigation of the Novelty Effect on Learning Outcomes in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 31, 5 (2025), 3204–3212. doi:10.1109/TVCG.2025.3549897
- [29] Young Hoon Lee and Gérard Medioni. 2015. Wearable RGBD Indoor Navigation System for the Blind. In *Computer Vision - ECCV 2014 Workshops*, Lourdes Agapito, Michael M. Bronstein, and Carsten Rother (Eds.). Springer International Publishing, Cham, 493–508.
- [30] Bing Li, J. Pablo Muñoz, Xuejian Rong, Jizhong Xiao, Yingli Tian, and Aries Arditi. 2016. ISANA: Wearable Context-Aware Indoor Assistive Navigation with Obstacle Avoidance for the Blind. In *Computer Vision - ECCV 2016 Workshops*, Gang Hua and Hervé Jégou (Eds.). Springer International Publishing, Cham, 448–462.
- [31] Guan hong Liu, Tianyu Yu, Chun Yu, Haiqing Xu, Shuchang Xu, Ciyuan Yang, Feng Wang, Haipeng Mi, and Yuanchun Shi. 2021. Tactile Compass: Enabling Visually Impaired People to Follow a Path with Continuous Directional Feedback. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 28, 13 pages. doi:10.1145/3411764.3445644
- [32] Chen-Lung Lu, Zi-Yan Liu, Jui-Te Huang, Ching-I Huang, Bo-Hui Wang, Yi Chen, Nien-Hsin Wu, Hsueh-Cheng Wang, Laura Giarrè, and Pei-Yi Kuo. 2021. Assistive Navigation Using Deep Reinforcement Learning Guiding Robot With UWB/Voice Beacons and Semantic Feedback for Blind and Visually Impaired People. *Frontiers in Robotics and AI* 8 (2021). doi:10.3389/frobt.2021.654132
- [33] Oussama Metatla, Sandra Bardot, Clare Cullen, Marcos Serrano, and Christophe Jouffrais. 2020. Robots for Inclusive Play: Co-designing an Educational Game With Visually Impaired and sighted Children. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3313831.3376270
- [34] Selma Musić and Sandra Hirche. 2017. Control sharing in human-robot team interaction. *Annual Reviews in Control* 44 (2017), 342–354.
- [35] Arshad Nasser, Kai-Ning Keng, and Kening Zhu. 2020. ThermalCane: Exploring Thermotactile Directional Cues on Cane-Grip for Non-Visual Navigation. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility* (Virtual Event, Greece) (ASSETS '20). Association for Computing Machinery, New York, NY, USA, Article 20, 12 pages. doi:10.1145/3373625.3417004
- [36] Isabel Neto, Hugo Nicolau, and Ana Paiva. 2021. Community Based Robot Design for Classrooms with Mixed Visual Abilities Children. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 31, 12 pages. doi:10.1145/3411764.3445135

- [37] Jakob Nielsen and Thomas K. Landauer. 1993. A mathematical model of the finding of usability problems. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems* (Amsterdam, The Netherlands) (CHI '93). Association for Computing Machinery, New York, NY, USA, 206–213. doi:10.1145/169059.169166
- [38] Anastasia K. Ostrowski, Cynthia Breazeal, and Hae Won Park. 2022. Mixed-Method Long-Term Robot Usage: Older Adults' Lived Experience of Social Robots. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 33–42. doi:10.1109/HRI53351.2022.9889488
- [39] Maike Paetzel, Giulia Perugia, and Ginevra Castellano. 2020. The Persistence of First Impressions: The Effect of Repeated Interactions on the Perception of a Social Robot. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (Cambridge, United Kingdom) (HRI '20). Association for Computing Machinery, New York, NY, USA, 73–82. doi:10.1145/3319502.3374786
- [40] Xavier Perrin, Ricardo Chavarriaga, Francis Colas, Roland Siegwart, and José del R. Millán. 2010. Brain-coupled interaction for semi-autonomous navigation of an assistive robot. *Robotics and Autonomous Systems* 58, 12 (2010), 1246–1255. doi:10.1016/j.robot.2010.05.010
- [41] Vinitha Ranganeni, Mike Sinclair, Eyal Ofek, Amos Miller, Jonathan Campbell, Andrey Kolobov, and Edward Cutrell. 2023. Exploring Levels of Control for a Navigation Assistant for Blind Travelers. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (Stockholm, Sweden) (HRI '23). Association for Computing Machinery, New York, NY, USA, 4–12. doi:10.1145/3568162.3578630
- [42] Filipa Rocha, David Gonçalves, Ana Cristina Pires, Hugo Nicolau, and Tiago Guerreiro. 2025. Exploring Collaboration in Programming Activities with Children with Visual Impairments: a 10-Session Study in a School Setting. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–22.
- [43] Daisuke Sato, Uran Oh, João Guerreiro, Dragan Ahmetovic, Kakuya Naito, Hironobu Takagi, Kris M. Kitani, and Chieko Asakawa. 2019. NavCog3 in the Wild: Large-scale Blind Indoor Navigation Assistant with Semantic Features. *ACM Trans. Access. Comput.* 12, 3, Article 14 (Aug. 2019), 30 pages. doi:10.1145/3340319
- [44] H. Seki, S. Kobayashi, Y. Kamiya, M. Hikizu, and H. Nomura. 2000. Autonomous/semi-autonomous navigation system of a wheelchair by active ultrasonic beacons. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, Vol. 2. 1366–1371 vol.2. doi:10.1109/ROBOT.2000.844788
- [45] Mahendran Subramanian, Noyan Songur, Darrell Adjei, Pavel Orlov, and A Aldo Faisal. 2019. A.Eye Drive: Gaze-based semi-autonomous wheelchair interface. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 5967–5970. doi:10.1109/EMBC.2019.8856608
- [46] Hironobu Takagi, Kakuya Naito, Daisuke Sato, Masayuki Murata, Seita Kayukawa, and Chieko Asakawa. 2025. Field Trials of Autonomous Navigation Robot for Visually Impaired People. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 678, 8 pages. doi:10.1145/3706599.3706663
- [47] Hironobu Takagi, Kakuya Naito, Daisuke Sato, Masayuki Murata, Seita Kayukawa, and Chieko Asakawa. 2025. Field Trials of Autonomous Navigation Robot for Visually Impaired People. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 678, 8 pages. doi:10.1145/3706599.3706663
- [48] Peter Trautman and Andreas Krause. 2010. Unfreezing the robot: Navigation in dense, interacting crowds. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 797–803. doi:10.1109/IROS.2010.5654369
- [49] Andreas Wachaja, Pratik Agarwal, Mathias Zink, Miguel Reyes Adame, Knut Möller, and Wolfram Burgard. 2017. Navigating blind people with walking impairments using a smart walker. *Autonomous Robots* 41, 3 (2017), 555–573.
- [50] Allan Wang. 2023. *Social navigation with pedestrian groups*. Ph. D. Dissertation. PhD thesis, Carnegie Mellon University.
- [51] Luyao Wang, Qihe Chen, Yan Zhang, Ziang Li, Tingmin Yan, Fan Wang, Guyue Zhou, and Jiangtao Gong. 2022. Can Quadruped Navigation Robots Be Used as Guide Dogs?
- [52] Liyang Wang, Jinxin Zhao, and Liangjun Zhang. 2021. NavDog: robotic navigation guide dog via model predictive control and human-robot modeling. In *Proceedings of the 36th Annual ACM Symposium on Applied Computing* (Virtual Event, Republic of Korea) (SAC '21). Association for Computing Machinery, New York, NY, USA, 815–818. doi:10.1145/3412841.3442098
- [53] Yue Wang and Fumin Zhang. 2017. *Trends in Control and Decision-making for Human-robot Collaboration Systems*. Springer.
- [54] Yize Wei, Nathan Rocher, Chitralekha Gupta, Mia Huong Nguyen, Roger Zimmermann, Wei Tsang Ooi, Christophe Jouffrais, and Suranga Nanayakkara. 2025. Human Robot Interaction for Blind and Low Vision People: A Systematic Literature Review. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 276, 19 pages. doi:10.1145/3706598.3713438
- [55] Nancy Worth. 2013. Visual impairment in the city: Young people's social strategies for independent mobility. *Urban Studies* 50, 3 (2013), 574–586.
- [56] Shuchang Xu, Ciyuan Yang, Wenhao Ge, Chun Yu, and Yuanchun Shi. 2020. Virtual Paving: Rendering a smooth path for people with visual impairment through vibrotactile and audio feedback. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (2020), 1–25.
- [57] Xiantao Zhang. 2025. The Escalator Problem: Identifying Implicit Motion Blindness in AI for Accessibility. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*. 6635–6643.
- [58] Yan Zhang, Ziang Li, Haole Guo, Luyao Wang, Qihe Chen, Wenjie Jiang, Mingming Fan, Guyue Zhou, and Jiangtao Gong. 2023. "I am the follower, also the boss": Exploring Different Levels of Autonomy and Machine Forms of Guiding Robots for the Visually Impaired. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 542, 22 pages. doi:10.1145/3544548.3580884