Not All Who Wander Are Lost: A Localization-Free System for In-the-Wild Mobile Robot Deployments

Amal Nanavati*
University of Washington
Seattle, USA
amaln@cs.washington.edu

Nick Walker*
University of Washington
Seattle, USA
nswalker@cs.washington.edu

Lee Taber
University of California
Santa Cruz, USA
ltaber@ucsc.edu

Christoforos Mavrogiannis
University of Washington
Seattle, USA
cmavro@cs.washington.edu

Leila Takayama
University of California
Santa Cruz, USA
takayama@ucsc.edu

Maya Cakmak
University of Washington
Seattle, USA
mcakmak@cs.washington.edu

Siddhartha Srinivasa
University of Washington
Seattle, USA
siddh@cs.washington.edu

***Abstract*—It is difficult to run long-term in-the-wild studies with mobile robots. This is partly because the robots we, as human-robot interaction (HRI) researchers, are interested in deploying prioritize expressivity over navigational capabilities, and making those robots autonomous is often not the focus of our research. One way to address these difficulties is with the Wizard of Oz (WoZ) methodology, where a researcher teleoperates the robot during its deployment. However, the constant attention required for teleoperation limits the duration of WoZ deployments, which in-turn reduces the amount of in-the-wild data we are able to collect. Our key insight is that several types of in-the-wild mobile robot studies can be run without autonomous navigation, using *wandering* instead. In this paper we present and share code for our wandering robot system, which enabled Kuri, an expressive robot with limited sensor and computational capabilities, to traverse the hallways of a 28,000 ft^2 floor for four days. Our system relies on informed direction selection to avoid obstacles and traverse the space, and periodic human help to charge. After presenting the outcomes from the four-day deployment, we then discuss the benefits of deploying a wandering robot, explore the types of in-the-wild studies that can be run with wandering robots, and share pointers for enabling other robots to wander. Our goal is to add wandering to the toolbox of navigation approaches HRI researchers use, particularly to run in-the-wild deployments with mobile robots.**

***Index Terms*—robot navigation, in-the-wild deployment, wandering, Wizard of Oz, robots asking for help**

Fig. 1. Kuri in the halls of our academic building. Kuri is an expressive and engaging robot with low sensing and computation capabilities. Its lidar can only see up to 7.5 ft, insufficient to localize in long, wide hallways. Further, a significant portion of its compute gets used by localization algorithms. Motivated by these limited capabilities, in this paper we describe a system for "wandering" that enabled Kuri to traverse our large hallways for four days.

I. Introduction

Over the last several years, the human-robot interaction (HRI) community has been moving towards in-the-wild studies. In-the-wild studies take place directly in "settings where people are and will increasingly engage with robots" and reveal insights about "how people will respond to robots in complex social settings and how robots will affect social dynamics in situ" [1]. In fact, in-the-wild studies have become so important to our community that the entire theme of HRI20 was "Real World Human-Robot Interaction."

However, in-the-wild studies with mobile robots are difficult to run. This is partly because our research community's foci are often on interaction, so we are interested in deploying robots that prioritize expressiveness and interactivity over navigational capabilities. For example, consider Mayfield Robotics' Kuri. Kuri has an engaging physical design and uses expressive animations composed of head movements, lights and vocalizations to portray emotions (e.g., happiness, sadness, sleepiness). However, Kuri's lidar has a max range 7.5 ft (Fig. 1), which prevents it from localizing in wide hallways using range-based techniques. Further, its on-board computer has 1.5GB RAM, which makes it challenging to localize using vision-based techniques [2]. Kuri is one example of several interactive mobile robots (Sec. II-A) that have difficulty running off-the-shelf autonomous navigation algorithms.

To account for the difficulty of running in-the-wild stud-

This work was (partially) funded by the NSF GRFP (DGE-1762114), the National Science Foundation IIS (#2007011), National Science Foundation DMS (#1839371), the Office of Naval Research, US Army Research Laboratory CCDC, Amazon, and Honda Research Institute USA.

*Authors contributed equally.

ies with interactive mobile robots, the HRI community has relied on teleoperation—also referred to as Wizard of Oz (WoZ) [3]—to set robots in motion. This can be seen in the in-the-wild papers we publish; over the last 5 years, less than half of the mobile robot in-the-wild papers used autonomous robots, compared to over 90% for stationary robots[1]. However, teleoperating a robot requires constant attention, which limits the duration of in-the-wild deployments. These shorter deployments, in-turn, limit the types of insights and amount of data we can collect from them.

In this paper, we present an alternative method for in-the-wild deployments of mobile robots, that enables longer-term deployments without the constant attention required by WoZ. Specifically, our key insight is that *several types of in-the-wild studies of mobile robots can be run without autonomous navigation, using wandering instead*. To demonstrate this, we begin with the aforementioned Kuri, which is highly expressive but cannot localize in the wide hallways of our 28,000 ft^2 floor. We then present and share the code[3] for our wandering robot system, which enables Kuri to traverse the hallways over several days. This system relies on informed direction selection to avoid obstacles and navigate long hallways, and periodic human help via a chatbot to charge. This system uses known techniques—wandering robot behavior and human help—to achieve novel functionality—a multi-day deployment of a low-spec mobile robot in a large indoor environment.

Our four-day deployment demonstrated that the robot was able to traverse the 1,200 ft of hallways in the building with little human help (around 0.5 hours over the 32 hour deployment). We conclude this paper with a discussion of the benefits of wandering robot deployments, the types of user studies that can use wandering robots, and how to develop wandering systems for other robot platforms.

II. Related Works

A. Interactive Robots

The interactive robots in our community (e.g., Kismet [4], Keepon [5], Flobi [6], Geminoid [7], Leonardo [8], Cozmo, Furhat, Jibo, Kuri, Pepper) typically prioritize communicative expressivity through their industrial [9], [10] and interaction designs [11], [12] as well as in their engineering design (e.g., using more actuators in the face than body, using more sensors for recognizing human behaviors). While this prioritiziation of expressivity makes sense when one is trying to study and produce robots that engage with people, these designers and engineers are also limited by financial budgets (e.g., research grant limits, hardware cost targets) and technical budgets (e.g., power, compute, weight). As such, they make trade-offs with regard to where to spend precious time, money, and effort. While it's not impossible to develop robots that are capable in expressive behaviors, human interaction, and dynamic navigation (e.g., the design goals set by Nexi [13]), that is an exception, not the norm. Most interactive robots prioritize certain design goals (e.g., human interactivity) over other system design goals (e.g., robust navigation).

In many cases, a robot's expressiveness trades off against its navigation capabilities. Many interactive robots are not mobile at all (e.g., ElliQ, Furhat, Jibo, Keepon, Mabu), likely because their use cases do not require mobility. However, of the interactive *mobile* robots, some cannot navigate autonomously in spaces any larger than a table top (e.g., Cozmo avoid cliffs, but has no localization or navigation capabilities). Others struggle with localizing and navigating in large spaces, mostly because of short 3D sensor ranges and/or narrow sensor fields of view (e.g., Kuri localizes and navigates through home residences, but cannot localize in large offices or warehouses). Indeed, many "social robots" require additional sensors and compute in order to get them to localize and navigate. The NAO has an accessory lidar head unit for improving its navigation capabilities. Pepper also needs third-party lidar to navigate through large spaces [14]. It is these mobile, interactive robots that are the focus of the current work.

B. In-The-Wild Studies

In-the-wild studies have been used for many aspects of human-robot interaction, including: investigating human reactions to deployed robots [15]–[22]; designing and testing robots' interaction, engagement, and learning techniques [17], [23]–[25]; and collecting datasets that can be used to develop algorithms that work in-the-wild [20], [26], [27]. However, as was mentioned above, several of the in-the-wild studies that involve mobile robots use teleoperation to move the robot around (the Wizard of Oz methodology [3])[1]. For example, Taylor et al. [27] uses a teleoperated mobile robot to gather ego-centric data on human group movement around a robot, Fallatah et al. [15] used a teleoperated mobile robot to investigate human responses to a help-seeking robot, and Palinko et al. [28] uses a teleoperated mobile robot to investigate the impact of robot gaze on who interacted with it. However, relying on a wizard to teleoperate the robot throughout its deployment takes up valuable researcher time and attention, thereby limiting the duration of the deployment.

A few notable examples of in-the-wild deployments that involve *autonomous* mobile robots include: CoBot, an office service robot [29]; SPENCER, an airport guide robot [18]; Hobbit, an in-home assistive robot for older adults [19]; and Robovie, a shopping mall service robot [17]. These robots required considerable researcher effort to build, design, and maintain, which may not be feasible or practical, particularly for labs whose speciality is not full-stack robot development. We believe that the wandering robot system we present will enable researchers to more easily run longer in-the-wild deployments of mobile robots, deepening our collective understanding of in-the-wild mobile robot interactions.

[1]Based on papers from the *HRI* conference and *Transactions on Human-Robot Interaction* from 2017-2021 that included the keyword "in the wild" and had a robot deployment (source: ACM Digital Library). Of the 33 papers found, 9 involved autonomous mobile robots, 11 involved teleoperated mobile robots, 13 involved autonomous stationary robots, and 1 involved a teleoperated stationary robot (1 involved both mobile and stationary robots).

C. Indoor Robot Navigation

Most frequently, indoor robot navigation is divided into a localization problem, that estimates a position based on kinematic and sensor models as well as a map of the environment, and a planning and controls problem, that computes a path for the robot and velocity commands to follow that path [30]. Because some robots lack the sensors to precisely localize, or the computational power to run localization algorithms [31], [32], research has developed techniques for localization-free coverage, or patrol. These approaches use only contact-sensors, and have the robot move straight until hitting an obstacle, and then rotate [33]–[35]. Due to these techniques' efficacy, low setup time, and minimal hardware cost, similar methods have been used in commercially successful robot vacuum cleaners for over a decade [36]. However, these methods are undesirable if the robot must avoid collisions.

Recently, there have been several research efforts to move away from requiring range sensors to localize. One such direction involves using visual and inertial information to track the robot [37]. A visual tracking estimate of sufficient quality, coupled with an exploration method, can enable a robot to complete a coverage task [38]. Another such direction involves using learned visual representations of the environment to enable point-to-point navigation [39]. Recent work has shown that it is possible to create topological maps of environments, enabling robots to reach even far-away image goals by navigating through a sequence of sub-goal images [40]. These methods hold promise for expanding the space of mobile robots that can effectively localize and navigate, particularly because they use relatively cheap sensors. However, they currently either require too much compute to run on-board some mobile robots, require a significant amount of preparatory work to train the models, and/or are not robust to in-the-wild challenges such as moving people or lighting changes.

a) Wandering: The notion of a wandering robot has existed in the robotics community since at least the 1980s [41]. A well-known commercial instantiation of this concept is the early versions of the Roomba, a robot vacuum cleaner that covered a room by moving straight until it hit a wall, turning away from the wall, and continuing [36]. In these wandering robots, however, not needing human intervention was a design requirement [41]. However, particularly with the rise of notions of human help, or symbiotic autonomy [29], in the HRI community, we believe that integrating wandering behavior with periodic human help can increase the types of mobile robots we can use for in-the-wild deployments, and lengthen the duration of those deployments.

D. Human Help

Multiple works have suggested that robots use human help to overcome failures and to handle unexpected circumstances, a concept also known as symbiotic autonomy [29]. For example, the CoBot project demonstrated that a robot could effectively navigate a multi-floor office building over several years by relying on co-located humans to provide help like operating the elevator, moving furniture, and completing tasks the robot could not (e.g., leaving notes on doors, retrieving food from fridges) [29], [42]. Other projects have demonstrated that robots with rudimentary navigation skills are able to successfully navigate outdoor spaces by either actively asking co-located humans for help [43] or by passively waiting for co-located humans to help it [44]. Another work studied remote (not co-located) human helpers, and found that helpers had large individual variation in their responses to help requests and got annoyed if the robot asked too much [45]. Our work extends the concept of human help to wandering robots, and shares open-source code[3] for enabling a robot to ask for help using commonly-used messaging applications.

III. Wandering Robot System

A. Robot: Kuri

We developed our system atop the Mayfield Kuri robot, a small, differential-drive social robot. Designed as a consumer product, Kuri is an expressive robot that can embody emotions like happiness, sadness and tiredness through eye movements, head movements, chest-light patterns, and beeps. However, as a product aimed at a \$700 retail pricepoint, Kuri also has limited sensing and compute capabilities. It is equipped with a custom-designed low-power lidar sensor with a horizontal field-of-view of 100°, a max range of 7.5 feet for walls, and a max range of 4.5 ft for human legs (Fig. 1). It can struggle to perceive dark surfaces until they are inches away. Kuri also has a monocular RGB camera, with a horizontal field-of-view of 87.5° and a vertical field-of-view of 47°. Its computer is a low-power Intel single-board computer (Intel(R) Atom(TM) x5-Z8350 CPU @ 1.44GHz). Conventional workloads like running a localization particle filter or a vision-based localization technique can consume the majority of available compute.

Although no longer commercially available, at least 48 university labs possess Kuris[2]. Because Kuri was designed to be a relatively affordable consumer product that uses technologies that are 5+ years old, its capabilities likely represent a lower bound on those that we can expect in deployed mobile robots. In other words, researchers will likely use mobile robots at least as powerful as Kuri for deployments. Hence, we chose Kuri as a platform for our wandering robot system.

B. System Requirements

Our goal was to develop a mobile robot system that could

- *be deployed for several days, with only periodic human intervention*. Unlike the Wizard of Oz paradigm [3], which typically involves a researcher teleoperating or providing instructions to the robot at all times, we wanted a human helper to be able to continue doing a full day of work while only periodically helping the robot.
- *effectively traverse the space it was deployed in*. We wanted the robot's position over its deployment to be distributed over the navigable spaces in the building, as opposed to being concentrated in a few places.

[2] According to personal communication with a former Mayfield Robotics employee

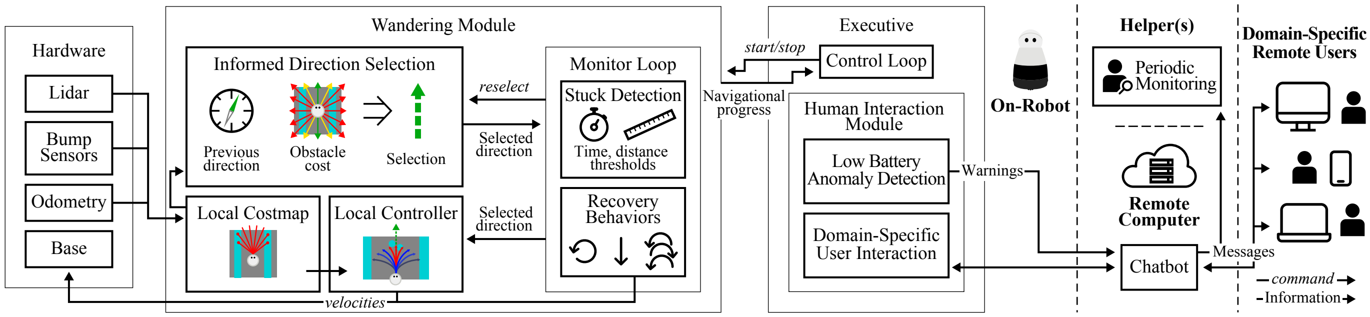

Fig. 2. Our system consists of a "wandering module" and a "human interaction module." The wandering module's consists of: an informed direction selection component that takes in the local costmap and the robot's previous direction and selects the robot's next direction; a local controller that takes in the selected direction and converts it to velocity commands; and a monitor loop that detects if the robot is stuck and attempts recovery behaviors. The human interaction module consists of a low battery anomaly detection loop that run on the robot and notifies the chatbot (running on a remote computer) when it detects a need for human help. The chatbot then notifies the helper(s). Together, these components enabled Kuri to wander the halls of our academic building for four days.

- *move at a reasonable speed*. We did not want the robot to move so slowly that it would hinder walking traffic, nor so fast that it may be dangerous or scary to passersby.

C. Domain

The floor in the academic building that we targeted for deployment is around 28,000 ft^2, with 1,200+ ft of hallways for the robot to navigate in. Hallways range from 6-10 ft wide and in many instances continue for 130 ft with minimal distinguishing geometry (see Fig. 4a). Several walls are made of glass (Fig. 4b), a material that cannot be detected by most lidar systems. Further, there is a black, chain-link bannister around many walkways that is also difficult to detect using lidar (Fig. 4b). The perimeter of the banister is marked on the floor with a cliff, which can lead to robots getting stuck. These notable challenges aside, the space typifies common office interiors into which one might deploy a mobile robot.

D. Early Attempts

We attempted to implement multiple conventional navigation approaches before developing our wandering system.

1) Lidar-Based Localization: We attempted to use Kuri's default navigation stack, developed by Mayfield Robotics to enable Kuri to navigate in home interiors. However, due to the large hallways and difficult-to-perceive materials, our deployment setting differed from domestic environments enough that this solution was unusable. Kuri's lidar could detect few surfaces—and none in some locations—which prevented it from building a map. We tried supplying a map created using a powerful lidar (manually edited to remove materials Kuri couldn't perceive), but found that the sparse sensor readings were also insufficient to perform Monte Carlo Localization against this map. Localization estimates diverged within 30ft, leading to unpredictable behavior. In practice, when using these lidar-based localization approaches, the robot often drove up to walls or difficult-to-see banisters, and sometimes collided as it futilely attempted to "go around."

2) Vision-Based Localization: We then sought out and evaluated vision-based localization techniques [46], [47]. Based on Kuri's monocular camera, and the need to close loops as it moved through the hallways, we first considered ORB-SLAM2 [2]. However, perhaps due to a combination of the camera's limited field-of-view, the robot's weak on-board computer, and our environment's lack of distinguishing features in certain hallways, we found that Kuri had to move extremely slowly (< 0.1 m/s) to stay localized. We also found that passing humans could cause Kuri to get delocalized, and it would sometimes relocalize far from its previous estimate (despite not moving). This led us to techniques that perform fusion with other sensors that can track the robot's motion [48]. We first tried VINS-MONO, a technique that merges IMU data with a monocular camera to maintain a localization estimate [49]. However, this technique assumes a static link between the IMU and camera, which isn't the case due to Kuri's pan/tilt head. Further, because Kuri moves along the ground, its IMU is not as informative of a sensor. This led us to a technique that uses wheel encoders instead, and allows for a dynamic link between the camera and wheel encoders [50]. However, building the packages in the robot's software distribution would have required large modifications to the source for library compatibility, which led us to pursue other approaches.

3) Adding Additional Sensors: Although we considered adding additional sensors to Kuri, we rejected it for three reasons. First, many candidate sensors would have strained Kuri's battery, which lasts for three hours under minimally demanding workloads. Second, Kuri lacked space within its chassis to mount additional sensors internally, and external sensors would impact its expressivity. Third, we felt that mounting sensors would make our system less reproducible.

4) Wandering Robot: Inspired by Brooks [41], we finally settled on using a localization-free approach and developing a wandering robot. Based on earlier approaches, we characterized a wandering robot as a system that iteratively selects a straight-line direction and follows it until an event triggers direction reselection. We therefore characterized the space of wandering systems in terms of two questions: what event triggers direction reselection, and how a new direction is selected. Brook's [41] robot, for example, triggers direction reselection after 10 secs have elapsed, and selects directions

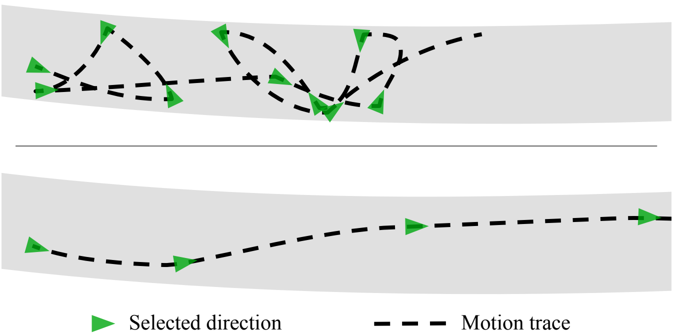

Fig. 3. How a robot with the same initial direction would move under different wandering behaviors. *Top*: Our first attempt at wandering had the robot move up to 3m in a direction and then pick a different direction, which prevented the robot from traversing long halls. *Bottom*: Using informed direction selection, the robot is able to successfully traverse long halls. In practice, it is also able to stay further from obstacles than the first attempt, because the robot rotates less and can therefore maintain a more accurate costmap.

uniformly at random. The Roomba, on the other hand, triggers direction reselection when it collides with a wall, and selects a direction by approximately reflecting off of the colliding surface [36]. With TweenBot [44], direction reselection is triggered when a human decides to rotate the robot, and what direction is selected is also up to the human.

In our case, we did not want the robot to collide with obstacles, due to the aforementioned potential of getting stuck near bannisters in our environment. Therefore, we initially tried triggering direction reselection after the robot had moved 3m, and selecting a direction by uniformly randomly sampling in $[0, 2\pi)$. However, we found that this technique resulted in the robot moving back-and-forth across a few-meter distance (Fig. 3 Top). Even after lowering the sampling range to be maximally 90° from its current direction, after a few iterations of resampling the robot would turn almost 180° from its original heading. This repetitive motion lead the robot to get stuck in hallways and prevented it from traversing the building.

E. System Implementation

Our final system[3] (Fig. 2) consists of a wandering module that is in charge of robot motion and a human interaction module that is in charge of communicating with the user(s).

1) Wandering Module: The wandering module operates in two layers: *informed direction selection* and *local control*. The informed direction selection layer (Alg. 1) uses local context in the form of a costmap, as well as the robot's previously selected direction, to pick the robot's next direction. The costmap, M, is defined in the base frame with the robot at the center at a 0-rad angle. Both lidar and bump-sensor readings are used to populate the costmap to account for the fact that Kuri's lidar alone may not pick up some obstacles. Given a costmap, the layer selects evenly spaced directions around its edge ($\texttt{generateDirections}(\eta)$). It then selects the direction that would encounter the least costly obstacles

Algorithm 1 Informed Direction Selection

Input: Current local costmap M,
number of directions to consider η,
previously selected direction θ_{prev}
Output: Next direction θ_{next}
$dirs \leftarrow \texttt{generateDirections}(\eta)$
$costs \leftarrow [(\texttt{obsCost}(\theta, M)\ ,\ |\theta - \theta_{\text{prev}}|)\ \mid\ \theta \in dirs]$
$i \leftarrow \texttt{argmin}(costs)$ {Using lexicographic ordering}
return $dirs[i]$

($\texttt{obsCost}(\theta, M)$), breaking ties by favoring the direction closest to the previously selected one as assessed in the robot's odometry frame. This ensures that Kuri avoids obstacles—lowering its chances of getting stuck—while navigating long hallways by moving in a similar direction (Fig. 3 Bottom). In practice, we found that a 3.5m^2 costmap updating at 2Hz (buffering the 10Hz lidar readings) paired with η=24 worked well with the robot's speed capped to .15m/s.

The selected direction is then passed to a local control layer, which generates velocity commands for the robot to follow that direction. Any local controller can be used; we use ROS's `dwa_local_planner`[4]. When the local controller can no longer generate velocity commands (e.g., due to the complete obstruction of that direction), the informed direction selection layer is engaged to reselect a direction.

The robot monitors its progress in the odometry frame and engages recovery behaviors if it detects that it hasn't moved a meter in thirty seconds. These behaviors are to: (a) clear the costmap and rotate in-place for eight seconds; (b) move backwards for ten seconds; and (c) alternately rotate left and right for ten seconds. In contrast with the other components of our system, these recovery behaviors were tailored to our deployment environment. Although simple, they effectively address failure modes we observed during testing, including getting stuck on furniture or trapped with a tread off of a cliff.

2) Human Interaction Module: The human interaction module enables Kuri to use a chatbot to contact humans, either to help it or for domain-specific interaction purposes. Specifically, the robot requests help when it is low on battery, since it cannot return to its charger autonomously due to its lack of localization. As Kuri's battery dips below set thresholds, it sends a message to designated "low battery helper(s)" (in our case, a researcher). This message tells recipients Kuri's battery level, optionally includes a picture of Kuri's surroundings, and asks them to put Kuri on its charger.

Although our system only uses low battery help messages, chatbots in general can enable rich forms of communication through buttons, open-ended text responses, emoji reactions, and much more. To illustrate some of these interaction modalities, our code includes a sample "where am I" help message where the robot shows users a picture of its current camera view, asks them to click a button corresponding to where it is, and/or asks them to type in its location using open-ended

[3] https://github.com/hcrlab/kuri_wandering_robot

[4] http://wiki.ros.org/dwa_local_planner

text. Although the system does not autonomously decide when to request this type of help, and does not use it to localize, this sample message illustrates the potential for rich remote human-robot interactions via chatbots. We use Slack as the platform for this chatbot.

A final component of our system includes researcher(s) periodically monitoring Kuri's camera feed to determine whether it is moving. Researchers do this infrequently—once every few hours—but it is important to catch the few situations where Kuri is unable to get unstuck using recovery behaviors. We did this by visualizing the robot's camera stream in RVIZ[5], although in principle this could also be done using a webstream or an extension to the interaction module where the user requests the robot's current view.

IV. Findings from a Multi-Day Deployment

To understand and illustrate the potential value of a wandering robot system, we ran a multi-day deployment in our large academic building (Sec. III-C). In addition to the aforementioned system requirements (Sec. III-B), another goal we had was for our system to be extensible to domain-specific scenarios. Therefore, we needed a domain-specific scenario for this deployment. This was around the time that our department was trying to boost morale and create a sense of shared community, spurred by Covid-19 work-from-home restrictions. Conversations around this departmental goal resulted in the idea of a robot photographer, designed to enable users to feel a sense of connection for a place by sharing images with them. In this section, we provide an overview of extensions we made to the system to adapt it to this domain-specific scenario, and an evaluation of the robot's wandering behavior. Additional details and findings from the deployment can be found in Appendix A and in our video[6].

A. Deployment Scenario: A Robot Photographer

In our deployment scenario, Kuri wandered the hallways of our academic building and took pictures to share with remote users. These users were distinct from the designated helper who periodically helped the robot charge its battery, although both interacted with the robot using the same Slack workspace. Kuri's goal was to take images of the building that it thought the remote users would like, share them with users, and get their feedback so it could improve its photo sharing. This deployment was approved by our university's Institutional Review Board (IRB) and the building manager, and participants were recruited from the population of people who used this building and the associated Slack workspace.

B. Extensions of the System

For this deployment, we extended the wandering robot system with a photo-taking module and additional chatbot interactions.

1) Photo-Taking Module: We extended the wandering module (Sec. III-E1) by enabling the robot to stop wandering when it wants to take a picture, and to execute precise orienting motions to fine-tune its image view. Specifically, as the robot wanders, it analyzes its image stream to detect objects and determine which images users might like (more details on the robot's model for human preferences can be found in Appendix A). When the robot detects images users might like, it stops its wandering motion. It then segments the image into a region of interest, using the approach in Vázquez and Steinfeld [51], rotates its head to center the region of interest, and then captures the picture. Note that although the extension to wandering behavior in this deployment is limited to starting and stopping wandering behavior for domain-specific purposes, in principle it can also be used to integrate wandering motions with other forms of motion (e.g., stop wandering when you see a person and move towards them).

2) Chatbot Interaction Design: We extended the human interaction module (Sec. III-E2) by enabling the robot to to interact with more users beyond the designated helper(s). Specifically, Kuri tells users it took an image for them, shares the image, and asks them to click a checkmark or x-mark button based on whether they like the image (Fig. 4c). This user feedback is then sent back to the robot, to better learn the user's preferences. To further engage the user, Kuri periodically asks an open-ended followup question, such as why they liked the image. More details about the interaction design can be found in Appendix A.

C. Wandering Findings

We deployed Kuri in our academic building for a period of four days in Summer 2021. In total, the robot ran for 32 hours, where it traversed all of the 1,200+ ft of hallways of the floor (Fig. 4a). The robot never ran out of battery. Its system of notifying a helper (a member of the research team) when it was low on battery enabled it to get charged in a timely fashion the 12 times it needed to over the course of the study. The robot's recovery behaviors enabled it to get unstuck most of the times it encountered environmental hazards; it needing manual rescue 4 times over the 4 days due to getting stuck on the cliff near the banisters (Sec. III-C). The helper noticed that the robot was stuck by periodically checking the robot's camera feed and realizing that it was not moving.

Overall, the system required around half an hour of the helper's time over the course of its 32 hour deployment (16 total instances of help, where most of those 30 minutes went towards the putting the robot on its charger). This is a tiny fraction of the researcher time that would have been required to teleoperate it under a WoZ design.

V. Discussion

In this paper, we presented wandering robots as a way of enabling multi-day, in-the-wild deployments of mobile robots that might otherwise face challenges navigating autonomously. We shared the design and code[3] for the system, which relies on informed direction selection and human help, and presented

[5]http://wiki.ros.org/rviz
[6]https://youtu.be/EsL9108-QYM

Fig. 4. (a) Kuri's movement around the floor plan of the building. Each dot represents a location Kuri was in when it took a picture, as annotated by researchers for a random 20% of the pictures it took. (b) Photos of Kuri as it wandered the environment. The environment contains materials that are difficult to perceive with depth sensors, like glass, chrome-finish metal, and black chain-link banisters. Large windows also varied the lighting, which impacts vision-based navigation approaches. (c) An example chatbot prompt from the deployment where the robot asked a user whether they like an image it took.

	Autonomous	WoZ	Wandering
Computation	Substantial	Minimal	Some
Connectivity	Not required	Required	Not required
Setup	Substantial	Minimal	Some
Supervision	Minimal	Substantial	Some
Robustness	Variable	Substantial	Variable
Goals	Yes	Yes	No

TABLE I
SUMMARY OF THE TRADE-OFFS BETWEEN NAVIGATION APPROACHES FOR MOBILE ROBOT DEPLOYMENTS.

outcomes from a four-day deployment in a large academic building.

A. The Benefits of Wandering Robots

Wandering robots have several benefits. First, *wandering robots expand the space of robots that can be deployed in-the-wild*. Even robots with poor sensors and computational capacity can wander. Despite only being able to perceive obstacles up to 7.5 ft away, Kuri was able to successfully traverse the large floor by wandering (Sec. IV).

Second, *wandering mobile robots require less development and setup time*. Probabilistic localization approaches typically require building a map of the environment, which takes time and familiarity with the foibles of map-building SLAM approaches. Learning-based navigation systems, which maintain implicit representations of the robot's location, require the collection of in-situ training data to be effective, and are challenging to deploy on low-compute platforms. In contrast, our open-source implementation[3] enables quick deployment of a robot like Kuri.

Third, *wandering robots can be sufficiently autonomous for long-term deployments*. Low-autonomy approaches like WoZ are generally valuable in HRI, although WoZ in-the-wild deployments require one or more researchers to constantly monitor or teleoperate the robot. In contrast, a wandering robot frees researchers' time: through careful system design, teleoperation can be avoided and human help may only be needed a few times throughout the deployment. This makes it more feasible to deploy the robot for a longer period of time.

Finally, when compared to stationary in-the-wild robots, *wandering robots enable researchers to more deeply explore a domain* because the robot will naturally interact with more people in varied contexts. Several in-the-wild studies have deployed stationary robots in particular indoor settings—a mall kiosk [21], in front of an elevator [20], at the entrance to a building [52]—and studied users' reactions to them. Yet, we know that the way users interact with a robot depends heavily on the context [1]. Thus, being able to easily vary the context (locations, times, direction of motion, etc.) of interaction by wandering could deepen our perspectives on how users interact with a mobile robot in the building.

Table I presents some of the tradeoffs between autonomous, Wizard of Oz, and wandering navigation approaches for mobile robot deployments. Each approach has its strengths and shortcomings. Our goal is to add wandering to the space of navigational approaches considered by HRI researchers when running in-the-wild mobile robot deployments.

B. User Studies Where Wandering Robots Can Be Used

The notion of a wandering robot may seem counter to the goals of mobile robotics—a field that has focused on domains like item pickup and delivery [29], [53], [54], guiding users [55]–[59], and taking inventory [60]–[62]. However, we contend that numerous in-the-wild human-robot interaction user studies can be run with a wandering robot. These include:

- *studies that investigate human reactions to an in-the-wild mobile robot*. For example, these can be exploratory studies or studies that investigate the impact of robot design on humans' in-the-wild reactions.

- *studies that investigate a robot's interactions with bystanders*. For example, these can involve investigating communication modes (e.g., natural language, expressive beeps, screens, etc.) or how to engage bystanders.
- *studies that investigate aspects of remote human-robot interaction*. For example, these can investigate how robots share information with remote operators, how they can elicit feedback from remote humans, or how they can engage users through disembodied communication.

Wandering robots also expand the types of robots that can be used for such studies to include robots with low sensor or computational capabilities. This lowers the cost and development barriers for deploying a system, enabling more researchers to run in-the-wild mobile robot studies.

C. Generalizing to Other Robots

One reason we developed this system on Kuri is because there are 72 Kuris at 48 different universities[2], so this system can be used off-the-shelf by dozens of labs to run in-the-wild deployments. However, because the HRI community uses a variety of robots, in this section we share pointers on developing a wandering system for other mobile robots.

A crucial part of our wandering module is the costmap. Maintaining the costmap requires range sensors that can estimate the distance to nearby obstacles, and odometry accurate enough that observations cohere as the robot moves. Importantly, neither capability need be excellent; Kuri's lidar is limited compared to today's alternatives, and the robot is not capable of dead reckoning for more than a meter before there is noticeable error.

Some auxiliary sensors are useful, but are not be required to implement wandering. Bump sensors were valuable for Kuri, whose low-range lidar gave the robot a proclivity for close encounters with hard-to-see surfaces. Most robots which use modern, commercially available depth sensors or even low-cost sonar arrays should have sufficient range to not need a bump sensor. We found that cliff sensors were unnecessary in our environment as fatal ledges, like stairwells, were behind doors. While not universal, this is true of many office buildings due to fire protection measures for egress routes.

Like extra sensors, recovery behaviors suited to the deployment environment aren't necessary but increase the amount of time the robot can be expected to go without requiring assistance. While piloting our deployment, we quickly discovered environmental hazards that would predictably ensnare the robot. Our procedure was to then attempt to recover with manual teleoperation, and if that was successful, to implement a matching scripted motion as a recovery behavior.

The chatbot, by virtue of running on a remote machine and exposing a general-purpose HTTP interface, can work as-is for other robots. It can be readily extended to account for other types of help and forms of user interaction, as demonstrated by the sample "where am I" help message provided by our code. Our code is particular to Slack, but other major messaging platforms (e.g., Microsoft Teams, Discord, Matrix, Telegram, etc.) provide their own bot APIs, so our system can be adapted to whatever platform is used in a particular deployment.

D. Limitations and Future Work

There are multiple interesting directions for extending our system's capabilities. Although our system only leverages human help when low on battery, it would be valuable to elicit human help in additional situations, such as when the robot is stuck. This could also enable investigations into how humans' willingness to help a robot is impacted by the *type* of help requested. Further, our system is unable to perform goal-directed motion, which can be required for some deployments or user studies. An interesting way to engineer goal-directed motion into this system would be to extend the human-interaction module to enable helpers to localize the robot.

Furthermore, we would be interested in extending the evaluation of our system. In this work, we tested our system on a Kuri. Several dozen labs also have a Kuri and could directly use our code. However, to further demonstrate the value of wandering robots in facilitating in-the-wild deployments, it would be valuable to test our system on additional platforms. This could involve collaborating directly with labs that use the system to understand their use case, gain insights into the challenges of extending the system to other robots and environments, and modify the system to make it easier to generalize. Finally, due to Covid-19 restrictions, our study did not involve in-person interactions (beyond the designated "helper"). We would be excited to run a long-term deployment study with a wandering robot while people are in the building. This would allow us to extract insights about co-located humans' reactions to an expressive in-the-wild mobile robot.

ACKNOWLEDGMENT

We thank Honda Research Institute for their generous support, Nikita Filippov for his contributions to the chatbot, the former employees of Mayfield Robotics for their help with this project, the deployment participants for their insights, and the anonymous reviewers for their helpful feedback.

REFERENCES

[1] M. Jung and P. Hinds, "Robots in the wild: A time for more robust theories of human-robot interaction," *J. Hum.-Robot Interact.*, vol. 7, no. 1, may 2018.

[2] R. Mur-Artal and J. D. Tardós, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *IEEE Trans. Robot.*, vol. 33, no. 5, pp. 1255–1262, 2017.

[3] L. D. Riek, "Wizard of oz studies in hri: a systematic review and new reporting guidelines," *Journal of Human-Robot Interaction*, vol. 1, no. 1, pp. 119–136, 2012.

[4] C. L. Breazeal, *Designing Sociable Robots*. MIT press, 2002.

[5] H. Kozima, M. P. Michalowski, and C. Nakagawa, "Keepon," *International Journal of Social Robotics*, vol. 1, no. 1, pp. 3–18, 2009.

[6] I. Lütkebohle, F. Hegel, S. Schulz, M. Hackel, B. Wrede, S. Wachsmuth, and G. Sagerer, "The bielefeld anthropomorphic robot head "flobi"," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 3384–3391.

[7] S. Nishio, H. Ishiguro, and N. Hagita, "Geminoid: Teleoperated android of an existing person," *Humanoid Robots: New Developments*, vol. 14, pp. 343–352, may 2007.

[8] C. Breazeal, C. D. Kidd, A. L. Thomaz, G. Hoffman, and M. Berlin, "Effects of nonverbal communication on efficiency and robustness in human-robot teamwork," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2005, pp. 708–713.

[9] C. F. DiSalvo, F. Gemperle, J. Forlizzi, and S. Kiesler, "All robots are not created equal: the design and perception of humanoid robot heads," in *Proceedings of the 4th conference on Designing interactive systems: processes, practices, methods, and techniques*, 2002, pp. 321–326.

[10] F. Hegel, F. Eyssel, and B. Wrede, "The social robot 'flobi': Key concepts of industrial design," in *19th International Symposium in Robot and Human Interactive Communication*. IEEE, 2010, pp. 107–112.

[11] G. Hoffman and W. Ju, "Designing robots with movement in mind," *Journal of Human-Robot Interaction*, vol. 3, no. 1, pp. 91–122, 2014.

[12] D. Glas, S. Satake, T. Kanda, and N. Hagita, "An interaction design framework for social robots," in *Robotics: Science and Systems*, vol. 7, 2012, p. 89.

[13] J. K. Lee and C. Breazeal, "Human social response toward humanoid robot's head and facial features," in *CHI '10 Extended Abstracts on Human Factors in Computing Systems*, ser. CHI EA '10. New York, NY, USA: Association for Computing Machinery, 2010, p. 4237–4242.

[14] R. Groot, "Autonomous exploration and navigation with the pepper robot," Master's thesis, Utrecht University, 2018.

[15] A. Fallatah, B. Chun, S. Balali, and H. Knight, "Semi-ethnographic study on human responses to a help-seeker robot," in *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 2020, pp. 640–640.

[16] W.-Y. Lee, Y. T.-Y. Hou, C. Zaga, and M. Jung, "Design for serendipitous interaction: Bubblebot-bringing people together with bubbles," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 759–760.

[17] D. Brščić, H. Kidokoro, Y. Suehiro, and T. Kanda, "Escaping from children's abuse of social robots," in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, 2015, pp. 59–66.

[18] M. Joosse and V. Evers, "A guide robot at the airport: First impressions," in *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, 2017, pp. 149–150.

[19] M. Bajones, D. Fischinger, A. Weiss, P. D. L. Puente, D. Wolf, M. Vincze, T. Körtner, M. Weninger, K. Papoutsakis, D. Michel *et al.*, "Results of field trials with a mobile service robot for older adults in 16 private households," *ACM Trans. Hum.-Rob. Interact.*, vol. 9, no. 2, pp. 1–27, 2019.

[20] D. Bohus, C. W. Saw, and E. Horvitz, "Directions robot: in-the-wild experiences and lessons learned," in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*. Citeseer, 2014, pp. 637–644.

[21] T. Kanda, M. Shiomi, Z. Miyashita, H. Ishiguro, and N. Hagita, "An affective guide robot in a shopping mall," in *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction*, 2009, pp. 173–180.

[22] S. Šabanović, S. M. Reeder, and B. Kechavarzi, "Designing robots in the wild: In situ prototype evaluation for a break management robot," *Journal of Human-Robot Interaction*, vol. 3, no. 1, pp. 70–88, feb 2014.

[23] E. S. Short, M. L. Chang, and A. Thomaz, "Detecting contingency for hri in open-world environments," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018, pp. 425–433.

[24] E. S. Short, A. Allevato, and A. L. Thomaz, "Sail: simulation-informed active in-the-wild learning," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 468–477.

[25] K. Sasabuchi, K. Ikeuchi, and M. Inaba, "Agreeing to interact: Understanding interaction as human-robot goal conflicts," in *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018, pp. 21–28.

[26] A. Ben-Youssef, C. Clavel, S. Essid, M. Bilac, M. Chamoux, and A. Lim, "Ue-hri: a new dataset for the study of user engagement in spontaneous human-robot interactions," in *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, 2017, pp. 464–472.

[27] A. Taylor, D. M. Chan, and L. D. Riek, "Robot-centric perception of human groups," *ACM Trans. Hum.-Rob. Interact.*, vol. 9, no. 3, pp. 1–21, 2020.

[28] O. Palinko, K. Fischer, E. Ruiz Ramirez, L. Damsgaard Nissen, and R. M. Langedijk, "A drink-serving mobile social robot selects who to interact with using gaze," in *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 2020, pp. 384–385.

[29] M. Veloso, J. Biswas, B. Coltin, and S. Rosenthal, "Cobots: Robust symbiotic autonomous mobile service robots," in *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

[30] B. Siciliano, O. Khatib, and T. Kröger, "Mobile and distributed robotics," in *Springer Handbook of Robotics*. Springer, 2008, pp. 795–958.

[31] D. L. Rizzini, F. Monica, S. Caselli, and M. Reggiani, "Addressing complexity issues in a real-time particle filter for robot localization," in *Proceedings of the Fourth International Conference on Informatics in Control, Automation and Robotics, Robotics and Automation*, may 2007, pp. 355–362.

[32] S. B. Samsuri, H. Zamzuri, M. A. A. Rahman, S. A. Mazlan, and A. Rahman, "Computational cost analysis of extended kalman filter in simultaneous localization and mapping (ekf-slam) problem for autonomous vehicle," *ARPN Journal of Engineering and Applied Sciences*, vol. 10, no. 17, pp. 153–158, 2015.

[33] T. Alam, L. Bobadilla, and D. A. Shell, "Minimalist robot navigation and coverage using a dynamical system approach," in *2017 First IEEE International Conference on Robotic Computing (IRC)*, 2017, pp. 249–256.

[34] J. S. Lewis, D. A. Feshbach, and J. M. O'Kane, "Guaranteed coverage with a blind unreliable robot," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 7383–7390.

[35] A. Q. Nilles, Y. Ren, I. Becerra, and S. M. LaValle, "A visibility-based approach to computing non-deterministic bouncing strategies," *The International Journal of Robotics Research*, vol. 0, no. 0, p. 0278364921992788, 0.

[36] B. Bennett, "Why your roomba takes a weird path to keep your floors clean," Aug 2021. [Online]. Available: https://www.cnet.com/home/kitchen-and-household/this-is-why-your-roombas-random-patterns-actually-make-perfect-sense/

[37] M. R. U. Saputra, A. Markham, and N. Trigoni, "Visual slam and structure from motion in dynamic environments: A survey," *ACM Comput. Surv.*, vol. 51, no. 2, Feb. 2018.

[38] A. Kim and R. M. Eustice, "Active visual slam for robotic area coverage: Theory and experiment," *The International Journal of Robotics Research*, vol. 34, no. 4-5, pp. 457–475, 2015.

[39] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, "Target-driven visual navigation in indoor scenes using deep reinforcement learning," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 3357–3364.

[40] X. Meng, N. Ratliff, Y. Xiang, and D. Fox, "Scaling local control to large-scale topological navigation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 672–678.

[41] R. Brooks, "A robust layered control system for a mobile robot," *IEEE J. Robot. Autom.*, vol. 2, no. 1, pp. 14–23, 1986.

[42] S. Rosenthal, M. Veloso, and A. K. Dey, "Is Someone in this Office Available to Help Me?: Proactively Seeking Help from Spatially-Situated Humans," *Journal of Intelligent & Robotic Systems*, vol. 66, no. 1-2, pp. 205–221, Apr. 2012.

[43] A. Weiss, J. Igelsböck, M. Tscheligi, A. Bauer, K. Kühnlenz, D. Wollherr, and M. Buss, "Robots asking for directions—the willingness of passers-by to support robots," in *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2010, pp. 23–30.

[44] K. Kinzer, "Tweenbots." [Online]. Available: http://www.tweenbots.com/

[45] T. Fong, C. Thorpe, and C. Baur, "Multi-robot remote driving with collaborative control," *IEEE Trans. Ind. Electron.*, vol. 50, no. 4, pp. 699–704, Aug. 2003.

[46] A. Li, X. Ruan, J. Huang, X. Zhu, and F. Wang, "Review of vision-based simultaneous localization and mapping," in *2019 IEEE 3rd Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)*. IEEE, 2019, pp. 117–123.

[47] A. Buyval, I. Afanasyev, and E. Magid, "Comparative analysis of ros-based monocular slam methods for indoor navigation," in *Ninth International Conference on Machine Vision (ICMV 2016)*, vol. 10341. International Society for Optics and Photonics, 2017, p. 103411K.

[48] C. Chen, H. Zhu, M. Li, and S. You, "A review of visual-inertial simultaneous localization and mapping from filtering-based and optimization-based perspectives," *Robotics*, vol. 7, no. 3, p. 45, 2018.

[49] T. Qin, P. Li, and S. Shen, "Vins-mono: A robust and versatile monocular visual-inertial state estimator," *IEEE Trans. Robot.*, vol. 34, no. 4, pp. 1004–1020, 2018.

[50] F. Zheng and Y.-H. Liu, "Visual-odometric localization and mapping for ground vehicles using se (2)-xyz constraints," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3556–3562.

[51] M. Vázquez and A. Steinfeld, "An assisted photography method for street scenes," in *2011 IEEE Workshop on Applications of Computer Vision (WACV)*. IEEE, 2011, pp. 89–94.

[52] R. Gockley, A. Bruce, J. Forlizzi, M. Michalowski, A. Mundell, S. Rosenthal, B. Sellner, R. Simmons, K. Snipes, A. C. Schultz *et al.*, "Designing robots for long-term social interaction," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2005, pp. 1338–1343.

[53] J. López, D. Pérez, E. Zalama, and J. Gómez-García-Bermejo, "Bellbot-a hotel assistant system using mobile robots," *International Journal of Advanced Robotic Systems*, vol. 10, no. 1, p. 40, 2013.

[54] G. Hollinger, D. Ferguson, S. Srinivasa, and S. Singh, "Combining search and action for mobile robots," in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 952–957.

[55] J. Guerreiro, D. Sato, S. Asakawa, H. Dong, K. M. Kitani, and C. Asakawa, "Cabot: Designing and evaluating an autonomous navigation robot for blind people," in *The 21st International ACM SIGACCESS conference on computers and accessibility*, 2019, pp. 68–82.

[56] W. Burgard, A. B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun, "Experiences with an interactive museum tour-guide robot," *Artificial Intelligence*, vol. 114, no. 1-2, pp. 3–55, 1999.

[57] M. Joosse, M. Lohse, and V. Evers, "How a guide robot should behave at an airport insights based on observing passengers," Centre for Telematics and Information Technology (CTIT), Tech. Rep. TR-CTIT-15–01, 2015.

[58] I. R. Nourbakhsh, C. Kunz, and T. Willeke, "The mobot museum robot installations: A five year experiment," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, vol. 4. IEEE, 2003, pp. 3636–3641.

[59] A. Tamai, T. Ikeda, and S. Iwaki, "A method for guiding a person combining robot movement and projection," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 1265–1270.

[60] I. Ehrenberg, C. Floerkemeier, and S. Sarma, "Inventory management with an rfid-equipped mobile robot," in *2007 IEEE International Conference on Automation Science and Engineering*. IEEE, 2007, pp. 1020–1026.

[61] J. Zhang, Y. Lyu, T. Roppel, J. Patton, and C. Senthilkumar, "Mobile robot for retail inventory using rfid," in *2016 IEEE International Conference on Industrial Technology (ICIT)*. IEEE, 2016, pp. 101–106.

[62] M. Paolanti, M. Sturari, A. Mancini, P. Zingaretti, and E. Frontoni, "Mobile robot for retail surveying and inventory using visual and textual analysis of monocular pictures based on deep learning," in *2017 European Conference on Mobile Robots (ECMR)*. IEEE, 2017, pp. 1–6.

[63] B. Dumitrascu, K. Feng, and B. E. Engelhardt, "Pg-ts: Improved thompson sampling for logistic contextual bandits," *arXiv preprint arXiv:1805.07458*, 2018.

[64] D. Russo, B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen, "A tutorial on thompson sampling," *arXiv preprint arXiv:1707.02038*, 2017.

Appendix A
Robot Photographer Deployment Details

As was mentioned above, the primary goal of this deployment was to test the robot's wandering system and to provide a case-study in extending that system for domain-specific scenarios. However, the interactions that ensued between the robot and users in our domain-specific scenario, a robot photographer, revealed initial insights that might be interesting to others in the HRI community. Hence, in this appendix we describe more details and initial insights from the deployment, particularly focused on the specific robot photographer scenario.

A. Recruitment

Participants were recruited through email and Slack messages posted in communication channels that were likely to reach people affiliated with our department. To be eligible, participants had to either have an office in or have attended meetings/classes in the academic building the Kuri was deployed in.

B. Modeling Users' Photo Preferences

The robot modeled human photo preferences by estimating the humans' preferences over objects they like to see in pictures. Specifically, the robot converted each image to a vector I, where I_i is the probability that object i is in the image (as outputted by Amazon Rekognition). The robot then assumes that user k has a corresponding preference vector θ_k, where $\theta_{k,i}$ indicates how much user k likes seeing object i in an image. The robot further assumes that the probability that user k reacts to the image with a checkmark is $(1 + e(-\theta_k \cdot I))^{-1}$. The robot begins with a Gaussian prior over every θ_k, seeded by pilot studies with building users, and updates its belief over a particular user's θ_k using Laplace Approximation every time it gets a response from them. This focus on objects is because of the aforementioned goal of using the robot photographer to strengthen people's feelings of connection to the building; therefore, the robot sought to understand which objects or views in the building users most liked.

We integrate this model of photo preferences into a larger 2-Arm Logistic Contextual Bandits formulation, where for each image it sees Kuri must decide whether or not to capture it for a particular user (i.e., the robot is solving a separate Contextual Bandits problem per user, where the arms are to "capture" or "not capture" the image). Kuri uses Laplace-TS, a Thompson Sampling based approach, to solve this problem. See Dumitrascu et al. [63] and Russo et al. [64] for details on Logistic Contextual Bandits and Laplace-TS.

C. Interaction Design

When sharing images with users, Kuri wrote that it took an image for them, and asked them to react based on whether or not they liked at least one object in the image. To further draw users' attention to the objects in the image, the message also listed a few objects that Kuri detected in the image. See Fig. 4c for a sample interaction message.

This interaction was developed over several pilot tests. One insight from the pilots was that care must be taken when selecting the emojis users use to respond to robot messages, due to multiple and possibly conflicting prior connotations of emojis (for example, does "thumbs up" mean the person liked the image, or that they are acknowledging having seen the image?). Another insight was that periodic followup questions could help engage users and make the task feel less like a CAPTCHA-style labeling task. Hence, the robot sometimes asked open-ended followup questions, like "Can you explain more about why you (dis)liked this photo? Any objects that you (dis)liked?" These followup questions also helped us gain initial insights into the willingness for everyday users, who are not designated helpers, to help the robot improve.

D. Deployment Procedure

Although the deployment lasted 4 days (32 hours), users interacted with the robot in two 3-day batches. This was to avoid user fatigue and accommodate user's schedules, while enabling the robot to reach a mature level of performance in terms of learning each users' preferences. Each day, the robot engaged with the user 4 times (at 2 hour intervals), each time sending a batch of the 5 captured images that it felt the user would most like. Participants could respond to the pictures at anytime, by clicking either the checkmark or x-mark buttons. Users were told that they were participating in the "Seeing the World Through the Eyes of the Robot" project, where the robot's goal was to learn the types of objects in the building that they likes and share pictures of those objects from the robot's perspective. At the end of each day, user's completed a survey that asked both quantitative questions (e.g., indicate (dis)agreement with statements like "I would interact with Kuri again in the future" or "Kuri learnt what types of objects I like to see in images") and qualitative questions (e.g., "Did interacting with Kuri help you feel a sense of connection with the building Kuri was photographing?," "In your interactions with Kuri, did you feel like you were interacting with something closer to a robot or a chatbot?," etc.). Participants read and signed an informed consent form before participating, and were compensated with a $25 Amazon gift card after participating.

E. Initial Insights

During the deployment, the robot interacted with n=31 remote users (10 female, 14 male, 1 prefer not to state, and 6 who didn't respond to that question). It sent them a total of 1,860 images (219 unique ones), out of which users liked 1,002 (53.9%), disliked 736 (39.6%), and did not respond to 122 (6.6%). All participants were sent all surveys and survey questions, but some chose not to respond. For the below analysis, we removed users who did not complete all surveys, for a total of n=22 users.

User perception of the system was overall positive: 84% of users said they "would interact with Kuri more," and (a different) 84% agreed with the statement "as I interacted with Kuri more, it shared images I liked more." Below, we delve more into selected insights from users' qualitative perceptions of the data. All qualitative responses were coded by a researcher, and responses where users did not answer the question were ignored.

1) Connection to the Building: Users were asked two open-ended questions about their connection to the building: “Did interacting with Kuri help you feel a sense of nostalgia for working in the building Kuri was photographing? Please explain why/why not” and “Did interacting with Kuri help you feel less like you were working from home? Please explain why/why not” We coded these responses as “positive,” “mixed”, and “negative.” In response to the nostalgia question, 66% of users responded positively. For example, one user wrote “Yes, it reminded me of the labs where my friends worked from, places I used to walk around when in the building and a sense of nostalgia taking me back to the time I spent there.” 28% responded negatively, with one user writing “No; at some point the pictures felt repetitive and the images were not of areas that held particular significance to me.” However, the results were flipped on the working from home question, with 5% saying Kuri helped them feel like they were working from home less and 77% indicating that it did not. The negative responses often discussed how Slack was not sufficient to help them overcome the feeling of working from home, with one user saying ”No, it [Kuri] was pretty disconnected from my reality”. This indicates that remote robots can help users feel connected to spaces that they no longer work in, although that is not enough to overcome the feeling of working from home, which would be an exciting space to explore in future work.

2) Chatbot vs Robot: To gain more insight into remote human-robot interaction, we also asked users if they felt Kuri was more like a robot or chatbot. The results were mixed: 10 users felt that Kuri was more like a robot, 9 like a chatbot, and 2 mentioned aspects of both. Two themes arose in users' responses—the presence of a physical body and the interaction style. In terms of physical body, one participant mentioned that Kuri had a physical body, which made it more like a robot. However, another participant mentioned that they had never experienced Kuri's physical body, which made it more like a chatbot. In terms of interaction style, one participant wrote “It felt like interacting with a robot. The interaction was sparse and it did not feel spontaneous and real-time like a conversation with a chatbot.” However, another participant wrote “[Kuri was more like a] Chatbot, it had that strangely polite form of speech and all interactions were done over text.” This indicates that whether users focus more on the embodied or remote parts of the interaction depends heavily on their prior associations of chatbots and robots. User perceptions of systems that have embodied and remote components is an exciting area for future work.

3) Emerging Patterns in User Responses to Followup Questions: Although the robot's learning system did not use users' responses to followup questions, users' did not know that. Despite that, we noticed that some users volunteered information that would help the robot better learn their object preferences. For example, in response to the open-ended question of what objects they (dis)liked, one user wrote “I like the stairs and the geometric aspect,” while another wrote “I do not care about the flooring and the corridor is pretty ugly.” Further, multiple users began writing their response in the same comma-separated format that the robot used to describe objects in the picture. For example, one user wrote that they liked the “Skylight, architecture, window.” This indicates that users might be searching for ways to make their open-ended responses more useful to the robot, and given our interaction design they settled on a comma-separated format. This tendency to search for ways to help the robot was elaborated on by one user in the survey, who wrote “I appreciated being able to give more detail about what I liked/disliked in the images but I wasn't certain how to phrase my feedback in a way that would be useful.” This indicates that open-ended responses might hold promise for enabling robots to get help and improve learning, although future work is required to systematically study the implications of and learning potential for open-ended followup questions.