



The Network Robot System: Enabling Social Human-Robot Interaction in Public Spaces

Dylan F. Glas, Satoru Satake, Florent Ferreri, Takayuki Kanda,
Norihiro Hagita

ATR Intelligent Robotics and Communication Laboratories
and

Hiroshi Ishiguro

Intelligent Robotics Laboratory, Osaka University

What kind of framework is needed to realize social robot applications in real public environments? Based on several years of experience in conducting field studies with social robots in public spaces, we identify key considerations that we have found to be indispensable in the development of robot services. We present a "network robot system" framework, which supports social robot services in five ways: observation of human behavior using environmental sensor networks, structured knowledge sharing, centralized resource and service allocation, global path planning for coordination between robots, and support for selected recognition and decision tasks by a human operator. We describe our network robot system implementation and present a demonstration in a shopping mall, illustrating how such a network robot system framework can be used to support heterogeneous teams of robots providing services in a real public environment.

Keywords: Human-robot interaction, network robot systems, field robotics, social robotics, service robots.

1 Introduction

Over the last decade, our laboratory has conducted many field trials in real-world environments, including a science museum (Shiomi, Kanda, Ishiguro, & Hagita, 2007), a train station (Shiomi et al., 2008), and shopping malls (Iwamura, Shiomi, Kanda, Ishiguro, & Hagita, 2011; Shiomi et al., 2009). Each deployment posed challenges in recognition, decision making, robot coordination, and information sharing. Through this experience, we have developed and refined a framework that addresses these challenges.

This framework is based on a "network robot system" (NRS) design approach, in which the robots themselves are merely the visible component of a network that integrates environmental sensor systems, central planning servers, cloud-based knowledge resources, and human users and supervisors. This framework has been successfully used by our research group (Figure 1(a)) and in collaboration with others (Figure 1(b)) in several field deployments. While similar systems have been developed for multi-robot task management, this is the first time this approach is being demonstrated in the domain of social robots.

In this paper, we will present our framework in the context of tasks such as guiding customers in a shopping mall. Our intention, however, is to share a general approach that can be useful in service robot deployment scenarios like those explored by other groups; for example, trash collection (Ferri et al., 2011); pedestrian guidance (Garrell & Sanfeliu, 2010); and assisting

Authors retain copyright and grant the Journal of Human-Robot Interaction right of first publication with the work simultaneously licensed under a Creative Commons Attribution License that allows others to share the work with an acknowledgement of the work's authorship and initial publication in this journal.



(a) A robot helping a customer with shopping.

(b) Collaboration between heterogeneous robots.

Figure 1. Example scenes using our Network Robot System framework.

people in hospitals (Mutlu & Forlizzi, 2008), supermarkets (Tomizawa, Ohya, & Yuta, 2006), and offices (Takayama, Marder-Eppstein, Harris, & Beer, 2011).

As this framework involves many technical components, we will often refer to publications that present those subsystems in greater depth.

2 Related Work

2.1 Social Robots in the Field

Research on robots developed for human-robot interaction (HRI) in public spaces, such as robots who act as museum guides (Burgard et al., 1998; Thrun et al., 1999) and shopping assistants (Gross et al., 2008), has addressed navigational and perceptual problems such as people-tracking and localization in public spaces. Other research, like Snackbot, a mobile robot deployed to deliver snacks in a university environment, has focused on elements such as dialogue content and social appropriateness (Lee et al., 2009). Robots have also been placed in busy public spaces to investigate social acceptance of robots (Kirby, Forlizzi, & Simmons, 2010; Weiss et al., 2008). These are examples of the types of robot services that our proposed framework aims to support in a modular and scalable way.

2.2 Ambient Intelligence

To augment on-board sensing, we often use “ambient intelligence” (AI) systems embedded in the environment. Laser range finders (LRFs) located in an environment can help a robot to find and approach pedestrians from a distance (Kanda, Glas, Shiomi, & Hagita, 2009). Environmental sensors have also been used with mobile robots to cover large areas for surveillance (Sanfeliu et al., 2010). The PEIS (“Physically Embedded Intelligent Systems”) Ecology project explored an architecture for integrating robots with various sensors and actuators into a smart-home environment (Saffiotti et al., 2008). As such systems can greatly aid robot perception and recognition, our proposed framework emphasizes integration with AI systems.

2.3 Multiple-Robot Coordination

Other research has addressed networked robots in traditional, non-social robot scenarios. Some architectures address decentralized cooperative behavior of swarm robots (Parker, 1998). Other studies have focused on coordinating multiple robots, addressing issues such as control of coverage (Hussein & Stipanovic, 2007) and formation (Monteiro & Bicho, 2010; Vig & Adams, 2006) and efficient exploration (Haumann, Listmann, & Willert, 2010). Centralized control approaches have also been taken, particularly when delegating heavy computation from mobile robots to servers (Inaba, Kagami, Kanehiro, Hoshino, & Inoue, 2000).

While these approaches can produce simple behaviors such as shape formation, social robot applications require more complex behaviors than these frameworks can provide.

2.4 Networked Social Robot Systems

Multiple-robot networks have also been developed for HRI. The Expo.02 work of Siegwart et al. (2002) was one example, where 11 robots shared their locations and some actions were coordinated to control the flow of visitors interacting with the robots at a robotics exhibit at the Swiss National Exposition. The DustBot project, which deployed a network of robots to perform tasks for urban hygiene management, was also designed to support multiple robot types, and it included communication with beacons (Ferri et al., 2011). However, these projects involved little or no coordination among robots and did not propose a general framework that could address various methods of coordination. By contrast, we are proposing a general framework to handle communication, robot coordination, and ways to define and allocate services between a central server and individual robots.

Many studies have addressed task allocation and coordination, summarized in Gerkey and Mataric (2004), where the goal is often task allocation for a large number of robots, and a primary concern is efficiency in allocation given utility of tasks. Our NRS's task allocation is relatively simple from an algorithmic point of view, but it addresses many practical considerations required for allocating social robots to provide services to human users. Hence, our work provides a successful example of how we can instantiate such theoretical ideas in a very concrete way.

3 Scenarios and Design Requirements

To provide a context for understanding the functionality of our framework, we will present three simplified robot service scenarios that we have actually implemented and demonstrated. These scenarios were chosen to showcase typical functionalities used in many of our field trials. Throughout the paper, we will refer back to these scenarios to place the system elements in context.

3.1 Scenario I: Multi-Robot Coordination

In this scenario, a customer at the mall entrance uses his mobile phone to request a robot to provide information about the shopping mall, while a nearby customer requests a robot to carry his bags (Figure 2). A central server assigns robots to these tasks based on their locations and capabilities, and plans coordinated paths for them to approach the customers. A humanoid robot approaches the first customer to offer information, and a robotic shopping cart approaches the other customer, offering to carry his shopping bags.

This scenario demonstrates the necessity of service allocation based on robot capabilities and coordinated path planning between robots.



(a) Different services are requested.



(b) Each robot engages in conversation with a customer.



Figure 2. Coordinating multiple robots.

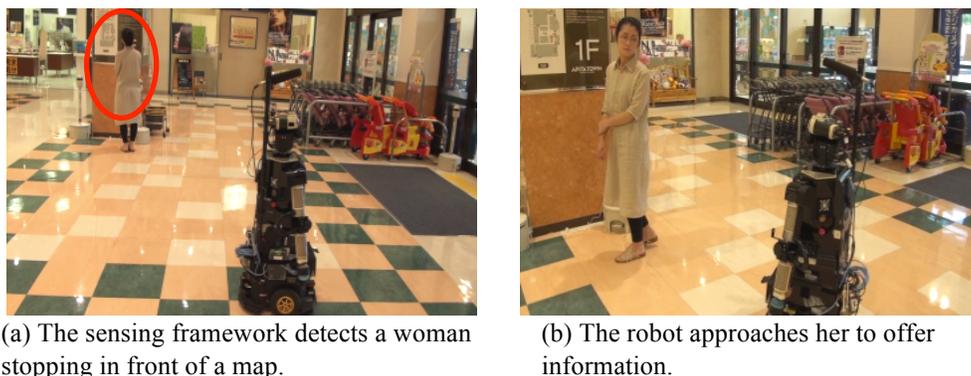


Figure 3. An example of context-aware service.

3.2 Scenario II: Context-Aware Service

The next scenario illustrates an example wherein robots proactively provide services to people based on their situational context, rather than responding to explicit requests. Near a large intersection in a shopping center, a robot is waiting to offer route guidance to customers. A woman stops in front of a map of the mall (Figure 3(a)). While looking at the map, she is approached and offered help by the robot, who says, “Are you looking for a particular shop?” (Figure 3(b)). The robot then answers any of her questions by giving directions or accompanying her to a destination. This scenario illustrates the need to recognize and anticipate people’s needs and the ability to allocate robot services accordingly.

3.3 Scenario III: Personalized Service

The third scenario shows how the NRS enables personalized services to be delivered to customers. From home, a registered customer uses her mobile phone to request a robot to help her with shopping and enters her shopping list (Figure 4(a)). Upon her arrival at the mall (Figure 4(b)), a robot comes to greet her and accompanies her through the supermarket carrying her basket (Figure 4(c)). Based on their location in the supermarket, the robot can remind her of items on her shopping list. This scenario illustrates how personal information can be used in services. It also shows the need for a means of personal identification of the service recipient.



Figure 4. An example of personalized service.

Table 1. System requirements.

<i>Navigation</i>	N1. Localization N2. Safety N3. Path planning and spatial resource allocation N4. Socially appropriate motion near people
<i>Conversation</i>	C1. Recognition support C2. Expert knowledge
<i>Anonymous Services</i>	A1. Recognize and anticipate people’s behavior A2. Assign services based on anticipated need
<i>Personalized Services</i>	P1. Identify individuals P2. Enable users to request services P3. Store customer data
<i>Modularity</i>	M1. Coordination of services M2. Support for multiple types of robots
<i>Safety and Supervision</i>	S1. Monitor for and correct recognition errors S2. Identify and intervene in problem situations S3. Enable one operator to supervise many robots

3.4 Design Requirements

The scenarios outlined here have illustrated several important requirements for an NRS framework. Other requirements, such as navigational safety and supervision by a remote operator, apply to all of the scenarios. Table 1 summarizes the key requirements that we considered in the design of our framework.

When we describe the elements of our system, we will refer back to Table 1 to show how those elements help satisfy this set of requirements.

4 System Implementation

In the process of preparing several field deployments of robots in commercial and public spaces, we developed a network robot system framework that addresses these requirements.

The high-level elements of our system are shown in Figure. 5. These include a *sensing framework*, several *information modules*, a system for *service allocation*, a *coordination module* for navigational coordination and path planning, and support from a *human supervisor*. Table 2 summarizes these elements, where the numbers in parentheses correspond to the requirements presented in Table 1.

Table 2. Key system elements and functions. Numbers in parentheses refer to requirements from Table 1.

<i>Sensing framework</i>	Robust tracking of people (N4) Recognize and anticipate people’s behavior (A1) Assist in robot localization (N1) Identify individuals (P1)
<i>Information modules</i>	Store information about robots (M2) Store information about environments (N2, N3) Store information about customers (P3)
<i>Path planning and spatial allocation</i>	Coordinate paths between robots to avoid conflicts and deadlock (N3) Ensure smooth locomotion near people (N4) Provide paths based on robot type (M2)
<i>Service allocation</i>	Coordinate robot services (M1) Enable users to request services (P2) Assign services based on anticipated need (A2)
<i>Support from a human operator</i>	Support for recognition (C1, S1) Direct control of robot (C2, S2) Ability to control multiple robots (S3)

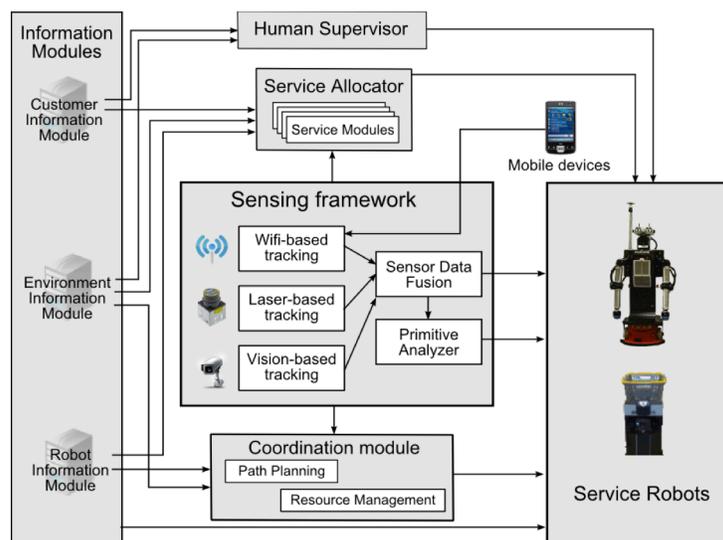


Figure 5. Overall architecture of our proposed Network Robot System framework.

In this section we will present the general design of each module of our system as well as specific instances of these modules from our implementations. Requirements from Table 1 addressed by a module are designated in parentheses; for example, (Table 1 – N1).

4.1 Sensing Framework

For robots to interact with humans, they need to be able to perceive and react to who people are and what they are doing. A sensing framework embedded in the environment can be used to track people in real time, to assist robot localization, and to anticipate people’s future behavior.

4.1.1 Robust tracking of people

Tracking people’s motion enables robots to react to a person’s identity and behavior, and motion data can help infer socially meaningful knowledge, such as direction of attention (Glas, Miyashita, Ishiguro, & Hagita, 2007), that can be useful for social interactions.

Sensors embedded in the environment can provide wide-area, high-precision tracking of people (Table 1 – N4), which can outperform people-tracking using on-board sensors in many situations: In crowded or cluttered areas, occlusions and noise can prevent stable tracking using on-board sensors, and resolution of laser or visual tracking drops off with distance from the robot. In our experience tracking pedestrians in a shopping mall with an on-board laser range finder, leg detections become indistinguishable from clutter and sensor noise beyond about 6 meters from the robot.

Sensors such as laser range finders (Glas et al., 2007) and video cameras (Matsumoto, Wada, Nishio, Miyashita, & Hagita, 2010) can be deployed to expand sensing coverage, minimize occlusions, and increase accuracy to track customers’ locations precisely. We have also used Wi-Fi fingerprinting (Bahl & Padmanabhan, 2000) for coarse localization of specific individuals.

4.1.2 Identifying individuals

Identification of individuals (Table 1 – P1) enables personalization of service, as illustrated in Scenario III, and continuity of service over time or between robots. In our work, we have used techniques such as radio-frequency identification (RFID) (Kanda, Shiomi, Miyashita, Ishiguro, & Hagita, 2009; Shiomi et al., 2007), visual face recognition (OMRON Corporation, n.d.), and Wi-Fi-based identification using smartphones for this task. By combining these results with the high-

precision position information available from the position-tracking systems, we can locate known individuals with high precision (Nishio et al., 2008).

In our demonstration of Scenario III, tracking was performed using a laser-based people tracker (high precision, but anonymous), together with localization using Wi-Fi fingerprinting (low precision, but able to identify registered individuals). On arrival at the mall, the customer's smartphone automatically connected to the Wi-Fi localization system, which estimated her rough location. By associating this location with the laser-based tracking data, the sensing framework could determine her precise location, enabling the robot to approach her and offer its service.

4.1.3 *Recognizing and anticipating human behavior*

Recognition of people's behavior is important in implementing smooth HRI. In particular, we are concerned with position and information derived from position; for example, "walking fast" (a motion primitive) and "in front of map" (a spatial primitive). In one field trial, we deployed robots to offer shop recommendations to visitors. Based on recognition of people's behavior, the robots avoided people who were hurrying and appeared busy, while approaching people who seemed to be walking slowly and might be open to talking with a robot (Shiomi et al., 2009).

Anticipation of human behavior is also important—if slow-moving robots can anticipate a person's future behavior, they can start moving early to approach potential customers (Satake et al., 2009). The work of Trautman and Krause (2010) has also used anticipation of human behavior to facilitate robot locomotion in crowds.

The presence of an environmental sensing framework enables us to use human motion data observed over long periods of time to generate detailed statistical models describing human behavior to make anticipation possible (Table 1 – A1). The primitive analyzer performs anticipation using location-based behavior likelihood models and an individual's behavior history over time (Kanda, Glas, et al., 2009).

In Scenario II, this technique was used to identify that the customer was performing a "stopping" behavior in front of the map.

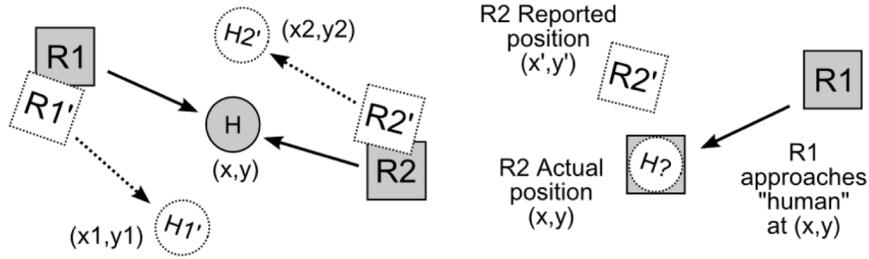
4.1.4 *Assisting robot localization*

Inconsistent localization between robots can cause a number of coordination problems. In the example shown in Figure 6(a), robots R1 and R2 have slight localization errors, perceiving their own poses to be R1' and R2'. Both robots believe they have identified separate people in need of help, but they have actually detected the same person, resulting in multiple robots offering services to the same person, as shown in the photo.

A similar problem, illustrated in Figure 6(b), is that one robot could mistake another robot for a pedestrian and try to initiate a social interaction with it. Since our robots are humanoid in form, they can be mistakenly detected as people by some sensor systems. These two problems were common in our early field trials.

Absolute localization in relation to known locations is important as well—knowing that the noodle shop is over *there* and the restrooms are over *there* is important for a robot giving directions.

In public spaces such as shopping malls, features used for map-based localization can change frequently as products and temporary displays are moved, as was the case in the supermarket entrance where we conducted field experiments during 2009–2010, shown in Figure 7 (Iwamura et al., 2011). In such environments, fixed references in the environment can often provide better estimates of a robot's position than the robot can obtain from map-matching. We have often used our human-tracking sensor systems to assist robots with localization by directly tracking the robots in the fixed reference frame of the sensors (Glas, Miyashita, Ishiguro, & Hagita, 2010) (Table 1 – N1).



(a) Example of multiple robots approaching the same person.



(b) Example of a robot mistaking another robot for a human and talking to that robot instead of an actual customer.

Figure 6. Social robot failures that can occur due to localization problems.



Figure 7. Example of an area in a shopping mall where features change from day to day. Top: photos on two different days. Bottom: laser scan maps of the area on two different days.

4.2 Information Modules

Roughly speaking, three categories of information are needed to support a network robot system.

The *Robot Information Module (RIM)* includes information about the capabilities of each robot, which can be used by a central planner for path planning within each robot's mobility

Table 3. Information module data summary.

Information Module	Data Provided
Robot Information Module (RIM)	Services offered Navigable terrain types Maximum clearance Maximum speed Ownership
Environment Information Module (EIM)	Localization map Dangerous areas Traversability map
Customer Information Module (CIM)	Customer name Mobile device ID Application content Personalization data

constraints and for appropriate allocation of robots to perform services (Table 1 – M2). In Scenario I, the planner used RIM information to allocate robots to services appropriate to their capabilities, such as assigning the cart robot to the baggage-carrying task.

The *Environment Information Module (EIM)* includes navigation and safety maps to be used for localization and path planning. In our implementation, navigation maps are generated through offline SLAM (simultaneous localization and mapping) using laser scan and odometry data recorded from robots, and safety zone maps are generated by hand (Table 1 – N2, N3).

The *Customer Information Module (CIM)* holds personal information about customers (or other service recipients) and is necessary for applications where personalized services are to be provided (Table 1 – P3). In Scenario III, the customer provides her shopping list information using her cell phone. This information is stored in the CIM alongside her name and known device ID.

A summary of the data stored in these information modules is shown in Table 3.

4.3 Path Planning and Spatial Allocation

The coordination module addresses the needs of path planning and of spatial coordination between robots, taking into consideration traversability constraints, which may vary by robot type. Some discussion of this module can be found in Shiomi et al. (2007).

4.3.1 Static space management

Multiple robots trying to go through a physically narrow space such as a supermarket aisle could encounter a deadlock situation (Figure 8). Robots might also compete for spaces based on application needs; for example, several cart robots waiting near the checkout counter to offer to carry shoppers' groceries, guide robots trying to occupy the space near a map, or advertising robots crowding near an entrance.

The coordination module helps to avoid deadlock and negotiate conflicts between robots for spatial resources (Table 1 – N3). When a robot requests use of a limited spatial resource, the coordination module will grant permission only if the space is not already in use.

4.3.2 Dynamic path planning

Robots operating in public spaces should move smoothly among people, avoiding collisions and strange behaviors like abrupt changes in direction. When information from a wide-area sensing framework is available, the coordination module can plan paths for the robot that are socially appropriate in the presence of pedestrians (Table 1 – N4).

When robots request paths to their destinations, the coordination module computes an efficient path appropriate for the service. The server periodically updates these paths, giving priority to robots with higher-priority services.



Figure 8. Deadlocked robots unable to pass each other in a corridor.

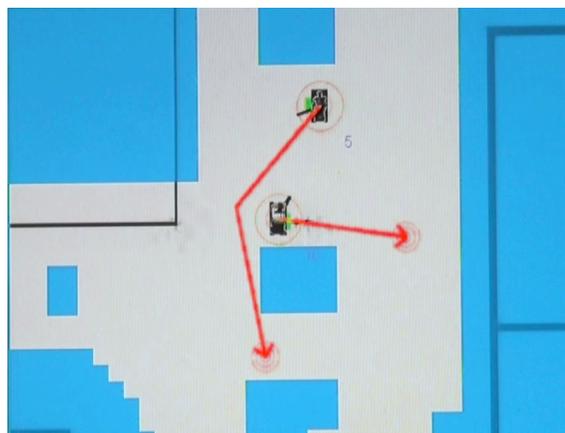


Figure 9. Paths planned for two robots in this example.

Social factors must also be considered in planning robot motion. For example, approaching someone from a frontal direction has been shown to be more psychologically acceptable than approaching from the side or back (Satake et al., 2009; Sisbot et al., 2005). The social information conveyed by interpersonal distance has also been considered in planning socially-appropriate behavior (Feil-Seifer & Mataric, 2012). A robot can also communicate intention through its locomotion, as illustrated by the example of “friendly patrolling” (Hayashi, Shiomi, Kanda, & Hagita, 2011). The coordination module can generate such special paths based on the robot’s current service task.

In Scenario I, each robot received the location of its assigned customer from the sensing framework and requested a path to that destination from the coordination module. The coordination module arranged the paths to achieve best efficiency of the robots (Figure 9). It computed paths for the robots in order of the priority of their services, while avoiding any potential conflict or collision among the robots, enabling each robot to safely reach its customer.

4.3.3 Traversability

The path planner uses information from the *RIM* and the *EIM* to plan paths for each robot based on traversability constraints (Table 1 – M2). For example, many public spaces have “detectable warning surface” blocks on the ground with raised bumps to help guide vision-impaired people, as shown in Figure 10. Some robots can traverse these uneven surfaces, while others cannot. To



Figure 10. Detectable warning surfaces for visually-impaired pedestrians are common in public spaces in Japan. Some robots cannot safely traverse these surfaces.

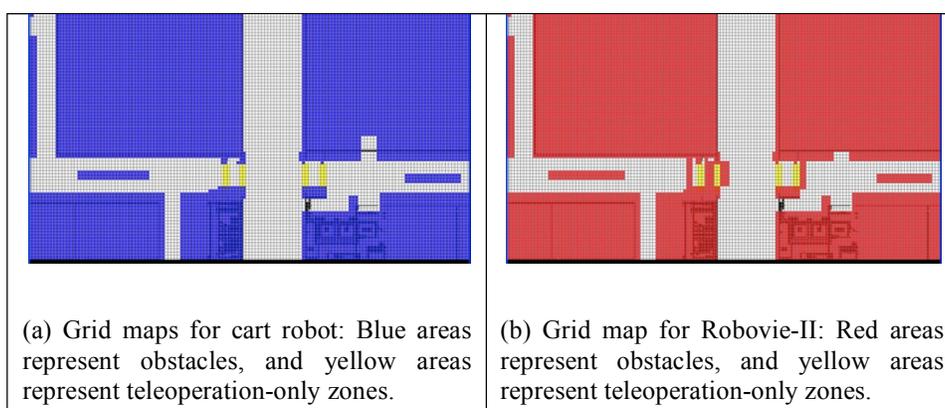


Figure 11. Traversability grid maps for each robot type.

address this problem, we create a traversability map for each robot type to be stored in the EIM. When a robot cannot safely traverse a region, it is designated as an “obstacle” zone.

Other areas may be physically traversable by robots but particularly dangerous—automatic glass doors, for example, as it is difficult for a robot to sense whether the doors are open and determine when it is safe to pass. We designate these areas as “teleoperation-only zones” through which robots may only pass if supervised by a remote operator (see Section 4.6).

Figure 11 shows an example of these zones taken from a shopping mall. The map on the left is for a cart robot, and the map on the right is for a humanoid robot, which is taller, less stable, and unable to traverse uneven surfaces. The corridor going from left to right goes through automatic doors in the center of the map. The map on the right contains “obstacle” areas around the automatic doors, where the floor is uneven. The coordination module will not plan any path through these areas for the humanoid robot. In the map on the left, however, such obstacle areas do not exist. The planner may thus plan a path for the cart robot that goes through the center of the map, but the robot will not be allowed to autonomously navigate through the “teleoperation-only areas”—a human operator must be called to assist the robot through these areas.

4.3.4 Safety

For robots operating in proximity to people in public spaces, safety is a very serious concern. While it is most important to avoid harming people, it is also important to avoid damaging the environment or the robot.

To avoid collisions with people, we use standard collision avoidance algorithms locally on the robots. We are currently using the “dynamic window approach” algorithm (Fox, Burgard &



Figure 12. Dangerous areas in a shopping mall.

Thrun, 1997). In crowded spaces, global path planning is used to direct the robots along routes that prevent them from coming too close to people detected by the sensing framework.

Additionally, public environments often contain dangerous areas that a robot cannot detect with its sensors. Transparent and reflective objects such as glass doors and mirrors, or drop-offs such as downward steps, can be difficult or impossible for robots to detect with laser range finders or cameras. Figure 12 shows some examples of obstacles that are difficult for robots to detect with ground-level laser range finders (a typical way for robots to detect obstacles). For safety and planning, maps of these invisible obstacles are provided by the EIM.

4.4 Service Allocation

The service allocator is the central planning mechanism that assigns services to robots and monitors the execution of those services. It handles service requests, identifies service opportunities, handles reservations for future services, and coordinates service allocation across multiple robots by considering service priorities and the capabilities and physical locations of the robots in its allocation algorithm (Table 1 – M1).

4.4.1 Services

In our framework, each service to be provided by robots is comprised of several service tasks, which are execution units managed by the server. The server contains logic determining which service tasks should be executed, under which conditions, by which types of robots. Once the server assigns a service task to a robot, the robot itself handles the details of service task execution, reporting back to the server upon success or failure.

On-demand services

For on-demand services, customers need a means to request services directly from the system; for example, via a smartphone or from the robots directly (Table 1 – P2). In either case, the requests are sent to the service allocator. Service requests can be for immediate service, as in Scenario I, or reservations for future services, as in Scenario III, where the customer schedules a shopping trip and reserves a robot to assist her at that time.

Proactive services

For proactive services, such as giving directions, recommending shops, or advertising services, the robots approach unknown people to offer their services. In such cases, the service

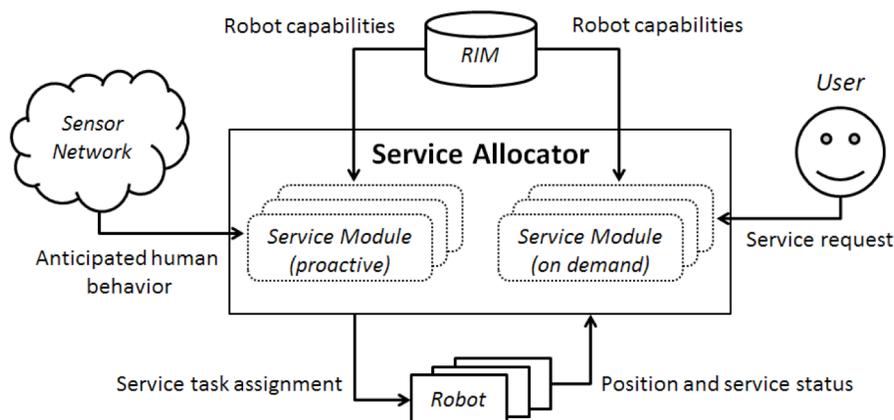


Figure 13. Basic flow of service task allocation

allocator must identify opportunities for providing services, rather than responding to requests, and allocation logic must be developed to assign robots to services based on anticipation of who will need or want the service (Table 1 – A2). To do this, it uses the statistical model of customer behavior provided by the primitive analyzer, described in Section 4.1.3.

For example, in the case of robots advertising for a shop, the system could be configured to target customers who are exhibiting “not busy” behavior in the spatial region in front of the shop. In the Scenario II example, the target behavior was “stopping” in the spatial region in front of the map.

The model of global behaviors can be used to predict customers who are likely to perform a target behavior in a certain area several seconds before they arrive, which gives the system time to allocate a robot and gives that robot time to move into the appropriate area. The robot is then sent to approach the person and give information about the shop.

4.4.2 Service allocator

Figure 13 shows the basic flow of service task allocation in the service allocator. The allocation logic depends on the specific needs of the service (Shiomi et al., 2007). For instance, a time-critical application would require a robot that is immediately available, but a proactive service would need a robot that could effectively initiate interaction with a customer; for example, by being able to approach from a frontal direction even if it required more time (Satake et al., 2009).

Thus, we designed the framework to enable developers to create their own algorithms for determining when a service should be performed. These are stored in the form of service modules within the service allocator. The service allocator uses the rules stored in its service modules to allocate available robots.

Concretely, the service allocator holds information about all available services in a list. For each service, it stores 1) the name of the service, 2) its priority, and 3) a link to a service module that handles the initiation of the service. For proactive services, the service allocator periodically checks the rules of each service module in order of priority, assigning service tasks to robots when necessary. For on-demand service, it calls the associated service module when it receives a request from a user, using the rules stored in that service module to determine which robot to assign.

To allocate a service to a robot, the service allocator performs a service matching algorithm, shown in Table 4. It queries the RIM for a list of all robots capable of providing that service. Next, the service allocator checks the location and service status (busy or available) reported by each robot. It provides this list of robots to the service module, which contains logic to decide whether to ask the service allocator to assign the service to one of the robots on the list.

Table 4. Service matching algorithm.

INPUT : The list of services *service_list* ordered by decreasing priority and the list of robots *robot_list*, the service allocator *service_allocator*

FOR each *service* in *service_list* **DO**
 IF *service* is proactive **THEN**
 initiation_rule_is_satisfied = *InitiationRuleFromServiceModule(service)*;
 ELSE
 IF *service* is on demand **THEN**
 initiation_rule_is_satisfied = *InitiationFromCustomerRequest(service)*;
 END IF
 END IF

 IF *initiation_rule_is_satisfied* **THEN**
 robot_capable_list = *GetServiceCapableRobotsFromRMI(robot_list, service)*;
 final_robot_list = *FilterRobotListByLocationAndStatus(robot_capable_list)*;
 selected_robot = *RequestRobotForAllocation(final_robot_list, service)*;

 IF *selected_robot* exists **THEN**
 AssignRobotToService(service_allocator, selected_robot, service);
 END IF
 END IF
END FOR

4.4.3 Service modules

When called, a service module (contained within the service allocator) performs its own computation, defined by developers. A service module for a proactive service typically uses information from the sensing framework. For each robot on the given list, it computes whether initiation of the service using the robot would be beneficial. If so, it requests the service allocator to assign the service to that robot. For example, a service module for an advertisement application might search for a person walking slowly, who may be a window-shopper, and try to find a robot that can approach that person quickly.

A service module for on-demand service typically performs a simple computation checking whether robots are available nearby so the service can be provided in timely manner.

4.4.4 Service execution

Service tasks are then executed on the individual robots. Once a service is allocated, the server requests the target robot to perform the first service task, and the robot updates its status to “busy” so that it will be not considered for other services.

The detailed procedure for each service task is pre-installed in each robot, and the service task is executed on the robot side, often using information from the server to provide the service, such as the location of a person to approach or the name of a registered user. The robot notifies the service allocator upon completion of each service task, and new service tasks required for the service are assigned until the entire service is complete, at which point the robot updates its status to “idle.”

4.4.5 Usage examples

In this section, we describe the service allocation process for the three scenarios presented in Section 3.

In Scenario I, illustrated in Figure 14, two customers requested different services, and so

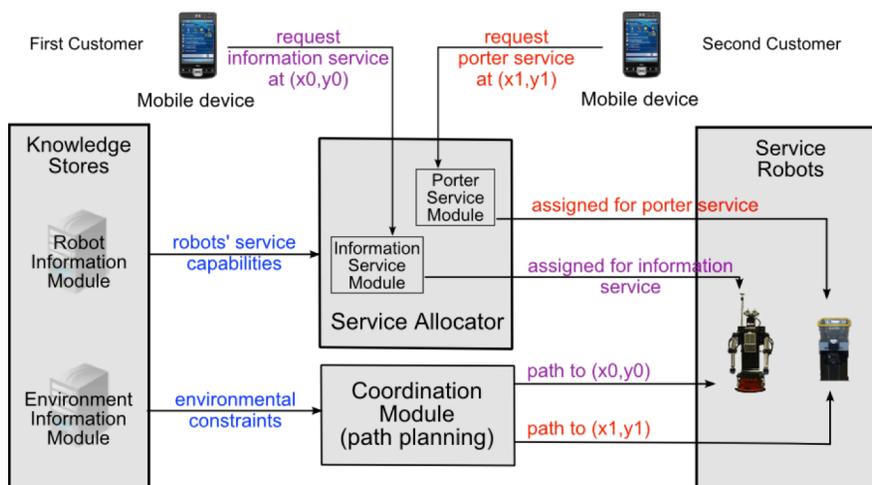


Figure 14. Data flow for coordinating multiple robots.

each request was handled by the corresponding service module. The request from the first customer was handled by the “information” service module. From the list of robots able to provide the “information” service (provided by the RIM), the service allocator selected the robot closest to that customer to perform the service.

The request from the second customer was handled by the “baggage-carrying” service module, which likewise selected the nearest robot with appropriate capabilities. The service allocator then assigned the cart robot to provide service for the second customer. The coordination module planned an appropriate path for each robot.

In Scenario II, the “route guidance” service module contained logic designating the space in front of the map as an area where “stopped” people would be likely to need the route guidance service. Thus, when the primitive analyzer reported a customer “stopping” at that location, the service module selected a robot for the “route guidance” service based on its capabilities and position to approach the customer and offer its service.

In Scenario III, the customer reserved a robot for a future shopping trip. Slightly before the appointed time, the service allocator assigned a robot to the “shopping support” service, beginning with a “wait for customer” service task. When the sensor network detected the customer’s arrival, the service allocator assigned the robot its next task, “approach customer,” to begin its service.

4.5 Support From a Human Operator

While robots today have greater capabilities for sensor recognition, dynamic planning, error detection, and error recovery than ever before, they are still far from ready to be deployed autonomously alongside humans in unstructured, public environments. For the foreseeable future, we expect that there will always be a human supervisor present behind the scenes to monitor and assist robots at some level.

In our system, supervisors typically use an interface such as that shown in Figure 15 to assist the robot’s recognition of sensor inputs, such as speech recognition or person identification (Table 1 – C1), and to monitor for and correct sensing errors, such as identifying dangerous situations or correcting localization (Table 1 – S1). The robot performs its own speech, gesture, and motion planning autonomously, and the role of the human is only to provide occasional sensor inputs.

In rare cases, an operator will need to control the robot directly to handle “uncovered” situations, such as an unexpected question from a customer (Table 1 – C2) or re-planning the robot’s path to avoid unmodeled obstacles (Table 1 – S2). In these cases, the robot cannot respond autonomously, so the human controls the robot directly.

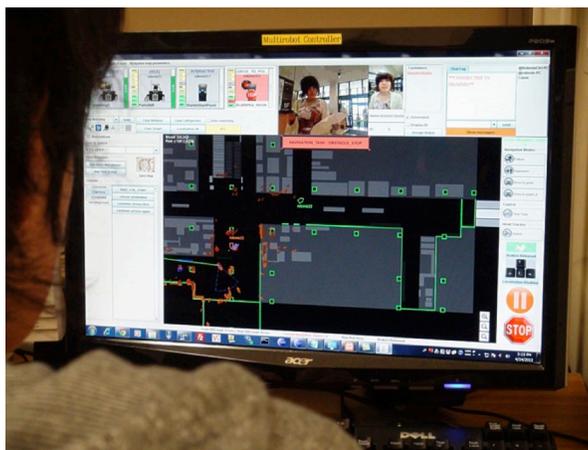


Figure 15. Operator using a teleoperation console to supervise four robots.

Finally, some mechanism is needed to enable one operator to manage multiple robots (Table 1 – S3). Techniques such as proactive timing control (Glas, Kanda, Ishiguro, & Hagita, 2012) and conversation fillers (Shiwa, Kanda, Imai, Ishiguro, & Hagita, 2009) can help improve performance of semi-autonomous robot teams in social interactions. These and other issues regarding supervisory teleoperation of multiple social robots are discussed in more depth in Glas et al. (2012).

5 Field Evaluation: Deployment in a Shopping Mall

Finally, to demonstrate the effectiveness of our NRS approach in enabling social robot applications in the field, we deployed four robots to provide services in a shopping mall.

5.1 Field Environment

The evaluation was conducted in the section of a shopping mall shown in Figure 16. We set up a laser-based human-tracking system in a high-traffic area near the entrance of the shopping mall to identify pedestrians for the robots to approach and to assist in localization of the robots.

The experiment was performed over a span of 3.5 hours on a Saturday afternoon in April. The day was not crowded, but the flow of pedestrians was steady. Figure 17 shows scenes from the field trial.

5.2 Robot Services

For this evaluation, we implemented both proactive and on-demand robot services to demonstrate the flexibility of our framework.

Two kinds of robots were deployed: two Robovie 2 humanoid robots and two cart robots, shown in Figure 17. The humanoid robots provided conversation-based services: For adults, the robots gave directions to locations in the mall, and for children, the robots performed entertaining play behaviors such as “rock, paper, scissors” and a guessing game. The cart robots provided an on-demand baggage-carrying service, carrying a customer’s bags to various destinations in the mall.

The state transition diagram in Figure 18 shows how these services were implemented in terms of “service tasks” assigned by the service allocator.

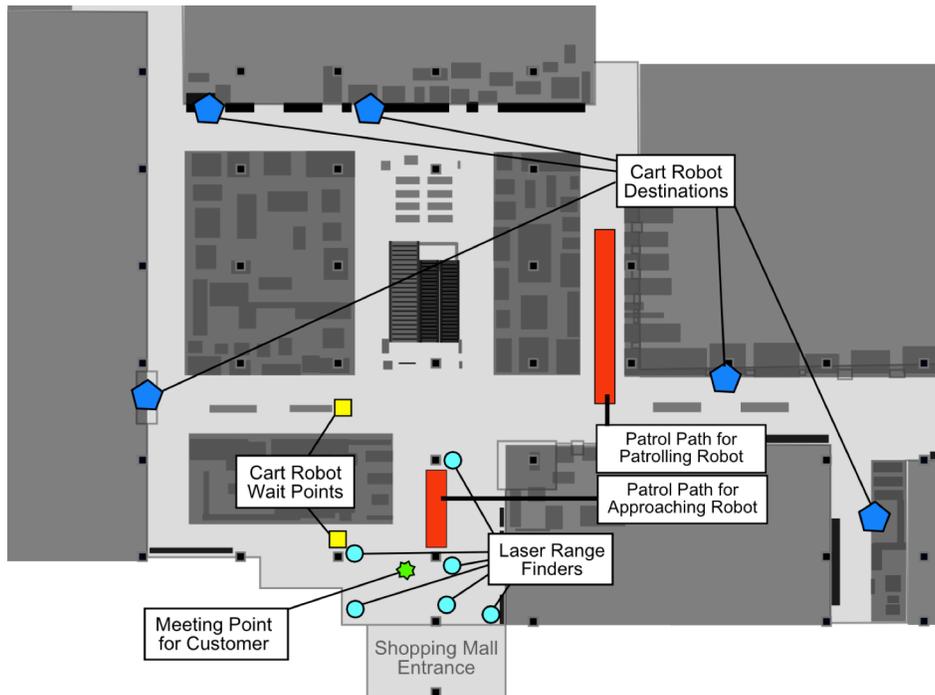


Figure 16. Map of the shopping mall where we conducted our field trial.



Figure 17. Scenes from our field trial. Top: a participant interacting with a cart robot. Bottom: shopping mall customers interacting with Robovie.

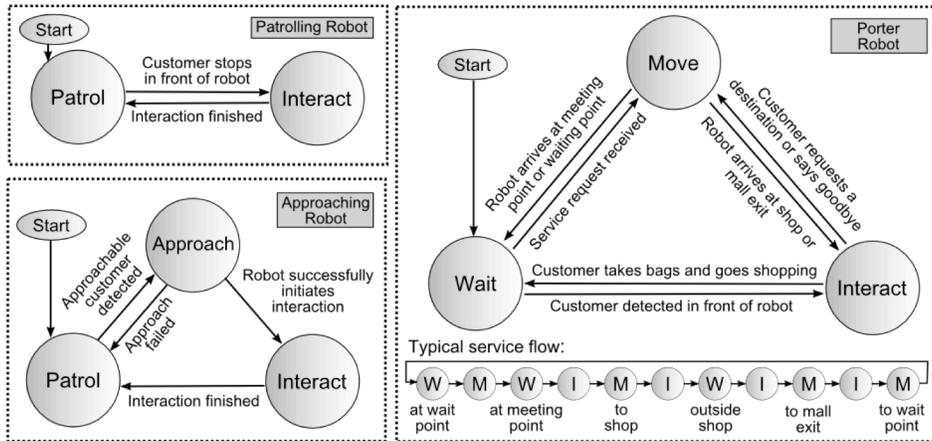


Figure 18. State transitions for service tasks in our field trial.

5.2.1 Anonymous service without sensor network support

One of the two humanoid robots (the “patrolling robot”) was placed in the shopping area and set to patrol along a fixed route, relying only on its on-board sensors for localization and human detection. This condition was chosen to illustrate the flexibility of the NRS framework—not all service areas can be covered by sensor systems, but the other functions of the framework can still be used in these areas.

For this service, the central server provided path planning, and the human operator assisted the robot with localization and speech recognition, but service allocation was handled locally—when customers stopped in front of the robot, it detected them and offered its services, reporting updates in its service state to the service allocator.

5.2.2 Context-aware service using the sensor network

The other humanoid robot (the “approaching robot”) was placed in the entrance area covered by the human-tracking system. Based on the information from this sensor network, the robot was assigned to actively approach customers to offer its services.

The primitive analyzer was used to predict which customers would be good approach candidates. The coordination module then calculated an approach path for the robot. The human operator assisted with speech recognition and localization support, and the robot was also able to use the sensor network to assist in localization and tracking people.

5.2.3 Personalized service: Baggage-carrying robots

Two non-humanoid cart robots, shown in Figure 17 (top), were also deployed. These robots were equipped with a camera, speaker, and microphone for simple conversational interactions. They provided a baggage-carrying service to registered customers who requested the service from their mobile devices.

The coordination module computed paths for the robots to requested destinations around the shopping mall, using safety and traversability maps from the EIM, and it dynamically allocated standby locations for the carts for idle times between customer requests.

The CIM provided photos for the operator to use in confirming the identity of the customer, and it provided the person’s name so that the robot could greet the customer personally. The sensor network provided identification and coarse localization for the customer carrying the smartphone, and forwarded the customer’s requests to the service allocator. The human supervisor assisted the

robots with speech recognition, localization, occasional obstacle avoidance, and confirming the identity of the registered customers.

5.3 Procedure

5.3.1 *Proactive services*

The humanoid robots approached and interacted with real customers in the shopping mall. After each interaction finished, a member of our laboratory briefly interviewed the customer to learn his or her impression of the robot and the quality of its service.

5.3.2 *Personalized services*

Since use of the personalized cart service required registration in the system and installation of a mobile application, it was impractical for us to test the service with general customers for a short-term deployment. Instead, we hired 15 participants to interact with the robots and give us their opinions of the service. Participants were native Japanese speakers recruited through an online service, with the following demographic characteristics: 8 male and 7 female, average age 25.3 years (SD 14.2 years), and paid 1000 yen per hour. Six had previously interacted with robots.

Each participant was instructed to use the cart's services twice, requesting a different destination each time, placing their baggage on the cart, and walking with it to the destination as if they were really using the service while shopping. After the participants had completed this task, they were interviewed to learn their impressions of the robot and the quality of its service.

5.4 Allocation of Functionality

To provide these services, functionality was distributed between the robot system, the NRS servers, and the human operator.

5.4.1 *Server Tasks*

The service allocator assigned services to robots. For the cart robots, this was done based on information from the participants' mobile devices, and for the humanoids it was done based on detection of people, either by the environmental sensor system in the case of the approaching robot, or by the on-board detection system in the case of the patrolling robot.

The coordination module computed paths for all robots, based on destination requests from the baggage-carrying robots and on human positions for the approaching robot. It sent a fixed path to the patrolling robot.

5.4.2 *Robot Tasks*

The robots executed on-board logic for carrying out the services assigned to them. The contents of these services were designed as conditional flowcharts, using the Interaction Composer software presented in Glas, Satake, Kanda, and Hagita (2011). When a path was needed or a service phase was completed, the robot would notify the central servers.

The robots also performed on-board sensing to detect nearby people and obstacles. For the patrolling robot, when a person was detected, the robot sent a request to the service allocator to begin providing a service to that person. For the other robots, the detection of a person was used in on-board processing for service execution.

5.4.3 *Operator Tasks*

The tasks of the human operator were as follows:

Confirm the identities of customers. For the cart robots, the operator matched a photo of the customer to the video feed from the robot's camera to guarantee proper user identification for personalized services.

Perform manual speech recognition. At each point where the robot’s internal behavior flow expected a speech input, the operator was presented with a list of valid candidates (e.g., “yes” and “no”), and the operator would listen to the audio stream from the robot’s microphone and click the button corresponding to what was said.

Select guide destinations from a map. When the humanoid robots were asked to give directions to locations, the operator would listen to the requested location and click the appropriate spot on a map. We have found this interface to be more effective than selection from a list, which is used for other speech recognition tasks.

Correct the robot’s localization. The operator periodically corrected the robot’s position on a map if the robot was improperly localized.

Assist the robot in obstacle avoidance. The operator sometimes manually controlled the robot’s locomotion for short periods of time if the robot became stuck due to objects in the environment, such as baggage carts or baby strollers stopped in its path.

5.5 NRS Framework Achievements

Overall, the NRS framework successfully enabled four robots to simultaneously perform services in a commercial space, taking advantage of centralized planning, sensor networks, and the availability of a human operator. Various metrics of the system’s performance are described here and summarized in Table 5.

The human-tracking system and primitive analyzer identified 431 trajectories of people walking through the entrance area over the 3.5-hour span of our experiment—an average of about two people per minute. The Wi-Fi device tracker successfully identified the participants’ mobile devices and successfully received 100% of service requests, enabling the service allocator to send a robot whenever a customer requested one.

In response to 30 service requests, two from each of the 15 participants, the service allocator successfully allocated a cart robot to the customer every time. It also successfully directed the humanoid robots to initiate 27 interactions with anonymous customers. One interaction is defined as a complete conversation with a unique customer.

The coordination module successfully guided the cart robots to and from 30 destinations, traveling a combined total distance of 3564 m. Figure 19 shows an example of the paths generated for the four robots during our field trial. For the approaching robot, the path planner provided successful approach paths to allow it to talk with 18 people.

Table 5. Summary of results from the field trial.

Baggage-Carrying Robots	
Number of interactions	30
Total distance traveled	3564 m
Participants reporting positive impressions	73%
Conversational Robots	
Number of interactions	27
Participants reporting positive impressions	68%
Sensor network	
Anonymous people tracked	431
Registered IDs detected	15

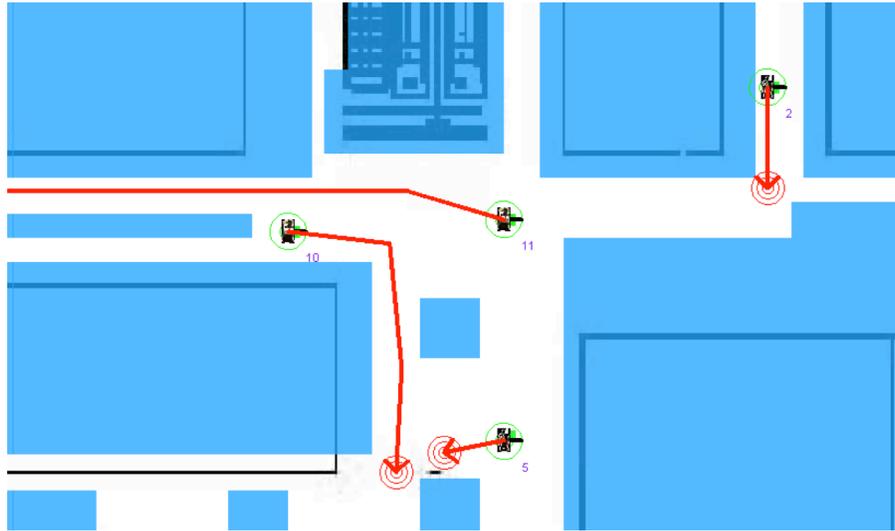


Figure 19. Example of paths generated for four robots during our field evaluation.

5.6 Customer Impressions

Of the 42 people who interacted with the robots (27 anonymous shopping mall customers who freely interacted with the humanoids, and 15 hired participants who interacted with the carts), we were able to interview 40 about their impressions of the robots. These interviews were recorded and transcribed, and two independent coders categorized the responses from the customers. To measure coding consistency, Cohen's kappa was evaluated for each set of responses coded, and then the coders were asked to discuss and reach agreement on any responses they had disagreed upon, to produce a final, consistent data set.

5.6.1 Humanoid robots

Of the 27 anonymous customers who interacted with the humanoid robots, we were able to interview 25. We asked these customers two questions: (1) what were their impressions of the robots, and (2) whether they would like to use such robot services if they were available. Results are shown in Figure 20.

Impressions from customers' interactions with the robots:

A majority of 76.0% (19/25) of the customers reported positive feelings, including interest, happiness, and amazement. However, 40.0% (10/25) also reported negative feelings, which included frustration with the conversational interaction and a comment from one customer that the robot's appearance was scary. Cohen's kappa for this question was 0.83, indicating very good agreement between the coders.

Aside from general positive comments like saying the robot was "fun" or "interesting," four customers commented that they liked that the robot made eye contact and followed them with its gaze. By contrast, two customers gave negative comments that the robot was not making eye contact with them but was looking somewhere else. These remarks underscore the importance of gaze control in conversational interactions. Also, one customer specifically commented on the robot's approach behavior, saying that the fact that the robot drove up to him to speak created a friendly feeling.

Five customers reported that the robot's response speed was too slow, most likely due to the operator taking a long time for some operation, such as finding a shop on the map.

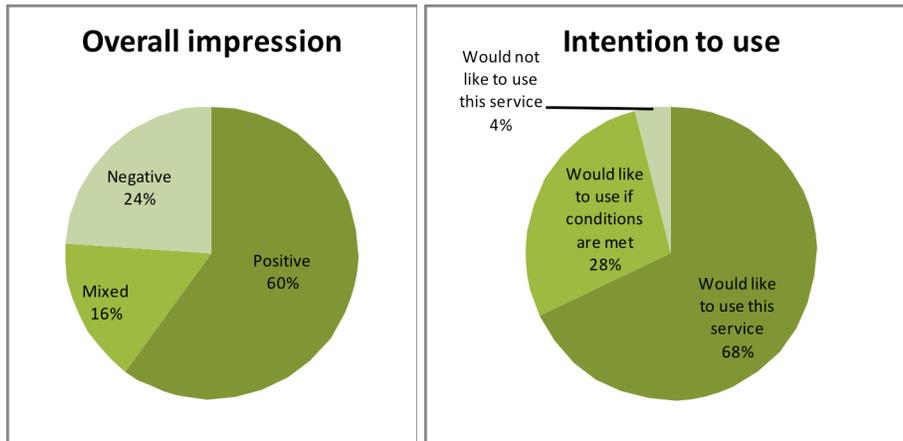


Figure 20. Interview results from 25 shopping mall customers who freely interacted with the humanoid robots.

Other negative comments included, “the robot seemed condescending,” “the robot’s appearance is too angular and scary,” and, “I felt self-conscious while talking to the robot in front of other people, as if I was talking to myself.” These comments suggest areas where the robot’s behavior and appearance could be improved for better interactions.

Intention to use this robot service if it is available in the future:

An overwhelming 96.0% (24/25) of the customers who interacted with the robot stated that they would like to use this type of robot service in the future, although some (7/25) said this was contingent upon improved conversational capability of the robot. Cohen’s kappa was 0.65, indicating good agreement.

Three customers suggested that the robot would be particularly useful for first-time visitors to the shopping mall, and seven thought that the robot would be particularly fun for children to talk to. Two suggested that it would be useful for elderly people, although one of them raised the concern that speech recognition would be particularly difficult for elderly customers, since they tend to speak in heavy local dialects.

5.6.2 *Baggage-carrying robots*

For the 15 paid participants who used the cart robots, we asked three questions: (1) what were their impressions of the robots, (2) whether they would like to use such robot services if they were available, and (3) how they felt about using the smartphone interface to interact with the robots. These results are shown in Figure 21.

Impressions from their interactions with the robots:

Most participants reported both positive and negative feelings about the baggage-carrying robot service. In total, 66.7% (10/15) reported positive feelings (happiness, friendliness, amazement, and approval), and 86.7% (13/15) reported negative feelings (anxiety, frustration, and difficulty in communicating with the robot). Cohen’s kappa was 0.67, indicating good agreement between the coders.

Of the negative comments, the most common complaint (6/15) was that the robot’s speed (0.75 m/s) was too slow. Other complaints were about the robot’s speech itself—three people said that its pronunciation was strange, and three said that its response was too slow. Two said that they enjoyed talking with the robot and wished they could have talked with it more (the cart robots spoke very little, and spent most of their time driving to the destination in silence).

We had expected that a personalized service that called the customer by name would get a positive response, and indeed, three participants did say they were happy that the robot called

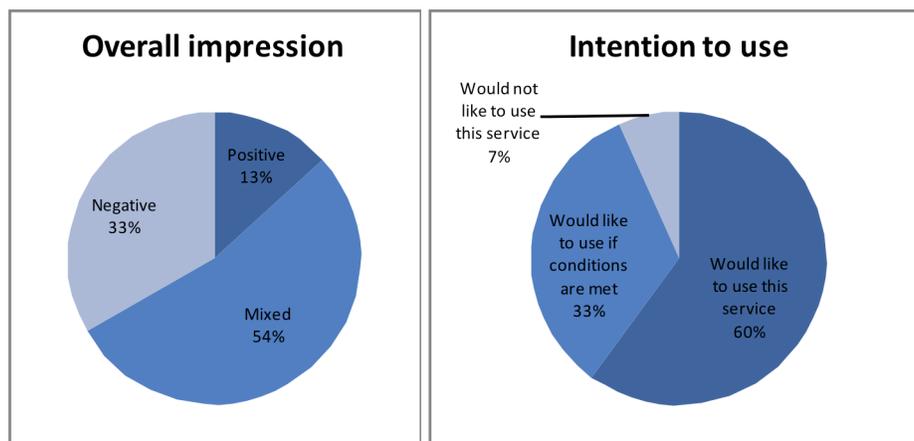


Figure 21. Interview results from 15 hired participants who interacted with the baggage-carrying cart robots.

them by name; however, another said, “I was a bit embarrassed when the robot said my name very loudly,” indicating that robot behaviors should be designed carefully not to make people feel self-conscious if other people might be listening.

Finally, three participants said the conversation seemed mechanical and lacked feeling, and that the interactions seemed artificial or awkward; for example, “the cart has no feeling when it talks, so when it says ‘thank you’ I don’t feel anything.” It is possible that these comments reflect an effect similar to the findings in Iwamura et al. (2011), where customers were more comfortable with a humanoid robot than with a cart robot as a shopping companion.

Intention to use this robot service if it is available in the future:

The response of the participants to this question was overwhelmingly positive: 93.3% (14/15) of participants indicated that they would like to use such a robot in the future, although 33.3% (5/15) said they would do so only under certain conditions, such as “if the system is improved,” “only when I have heavy baggage,” and “when I am elderly, but not while I am young.” Many of the participants (5/15) said in particular that they thought this service would be useful for elderly people or people with disabilities. Cohen’s kappa for this question was 0.76, indicating very good agreement.

Feelings about interacting with the robot using a mobile phone:

For the cart robots, we asked an additional question to get feedback on people’s impressions of the mobile phone interface. In this case, 73.3% (11/15) of participants reported positive feelings, including ease of use, convenience, and approval. The other 26.7% (4/15) reported negative feelings, such as frustration with waiting for the robot to arrive after requesting the service and the necessity of owning a smartphone to use the service. Cohen’s kappa for this question was 1.0, indicating perfect agreement.

5.7 Discussion

5.7.1 Cross-platform collaboration

One strength of our NRS framework is the ability to easily integrate new robots, new sensors, or new services into the system. Using simple protocols to connect to the planning, coordination, and localization servers, third-party robots can be integrated easily into the NRS without deep or complex software integration, enabling collaboration between very different kinds of robots.



Figure 22. Collaboration between Robovie and DustCart, supported by the NRS.

We have conducted collaborative work with the DustCart robot from the EU DustBot project (Figure 22) and with Honda’s ASIMO robot (Figure 1). In these demonstrations, Robovie talked with visitors and initiated a collaborative task, wherein the other robot performed some physical task, such as serving a drink or carrying baggage, while Robovie continued talking to the visitor, chatting or offering verbal instruction.

In each case, only 1–2 weeks of implementation and testing were necessary to integrate the new robots with the NRS platform and prepare a collaborative robot demonstration. We have also conducted other NRS demonstrations with robots such as Mitsubishi’s Wakamaru and Toshiba’s ApriPoco. An interface component for ROS (Robot Operating System) has also been developed for the latest version of the platform.

5.7.2 Practical considerations

Aside from the target functionalities presented in this paper, our experiences have shown a number of practical benefits provided by the modular design of the NRS framework.

When robots experience hardware problems—hard drive crashes, electrical failures, etc., the modularity of the NRS framework makes it easy to swap a backup robot into the system, enabling the experiment or demonstration to continue in a nearly seamless way. Rather than statically specifying services, paths, etc. within individual robots, the NRS dynamically allocates paths and services, which minimizes the code or settings that need to be modified when the composition of the robot team changes.

We have even replaced robots with different robot models—in one field trial we had some hardware problems with a Robovie-R3 robot, and we were able to seamlessly replace it with a Robovie-R2 (a robot with a very different design) for an important demonstration. This was possible because differences in hardware components and internal implementations of gestures and poses are hidden beneath the abstraction layer of “service tasks,” enabling the different robots to operate interchangeably within the network robot system.

The addition of new sensor types is also facilitated by our modular design. We have developed several versions of our human-tracking system; for example, using RGBD sensors such as the Microsoft Kinect. Although developed by an independent team, these new sensor systems can be used with our robots if they support the data protocols in our NRS framework design. This flexibility has been extremely helpful in managing the complexity of heterogeneous robot deployments in multiple environments with different sensor systems.

Finally, the EIM makes it possible to easily switch environments. We often move robots between our lab and various field trial environments, and the local NRS at each environment enables the robots to automatically make use of the latest navigation maps and receive path planning and service allocation for that environment.

6 Conclusions

We have presented a framework for a “network robot system” in which mobile robots, planning servers, and sensors embedded in an environment are integrated to provide robot services to people in social contexts. The requirements for this framework, motivated by our experiences in several years of field trials, primarily include the need for recognition and anticipation of people’s behavior, identification of individuals, coordination of services and navigation paths between robots, and supervision by a human operator.

We presented a field experiment showcasing the capabilities of this framework by providing services with four robots in a shopping mall, and our results showed that not only was the technical framework successful in supporting the robot services, but that people who used the robots responded in a positive way, with a great majority indicating that they would like to use services like these in the future. This underscores the worth of conducting research and field experiments to investigate and develop social robot services in real-world environments, and we submit that the NRS approach is an effective and practical way to make such robot deployments a reality.

Acknowledgments

We would like to thank Dr. Masahiro Shiomi for his contribution in designing and building this framework. We would also like to thank Dr. Satoshi Koizumi and Benoit Toulmé for their help in the organization and realization of the field trial at the APITA shopping mall, and Yasuhiko Hato for assisting with the data analysis. This research was supported by the Ministry of Internal Affairs and Communications of Japan.

References

- Bahl, P., & Padmanabhan, V. N. (2000). *Enhancements to the RADAR user location and tracking system* (Technical Report MSR-TR-2000-12). Microsoft Research.
- Burgard, W., Cremers, A. B., Fox, D., Hahnel, D., Lakemeyer, G., Schulz, D., . . . Thrun, S. (1998). The interactive museum tour-guide robot. In *National Conference on Artificial Intelligence (AAAI1998)* (pp. 11–18).
- Feil-Seifer, D., & Matarić, M. (2012). Distance-based computational models for facilitating robot interaction with children. *Journal of Human-Robot Interaction*, *1*(1), 55–77. doi:10.5898/JHRI.1.1.Feil-Seifer
- Ferri, G., Manzi, A., Salvini, P., Mazzolai, B., Laschi, C., & Dario, P. (2011). Dustcart, an autonomous robot for door-to-door garbage collection: From Dustbot project to the experimentation in the small town of Peccioli. In *IEEE International Conference on Robotics and Automation (ICRA2011)* (pp. 655–660). doi:10.1109/ICRA.2011.5980254
- Fox, D., Burgard, W., & Thrun, S. (1997). The dynamic window approach to collision avoidance. *Robotics & Automation Magazine, IEEE*, *4*(1), 23–33. doi:10.1109/100.580977
- Garrell, A., & Sanfeliu, A. (2010). Model validation: Robot behavior in people guidance mission using DTM model and estimation of human motion behavior. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2010)* (pp. 5836–5841). doi:10.1109/IROS.2010.5651685
- Gerkey, B. P., & Matarić, M. (2004). A formal analysis and taxonomy of task allocation in multi-robot systems. *International Journal of Robotics Research*, *23*(9), 939–954. doi:10.1177/0278364904045564
- Glas, D. F., Miyashita, T., Ishiguro, H., & Hagita, N. (2007). Laser tracking of human body motion using adaptive shape modeling. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 602–608). doi:10.1109/IROS.2007.4399383

- Glas, D. F., Miyashita, T., Ishiguro, H., & Hagita, N. (2010). Automatic position calibration and sensor displacement detection for networks of laser range finders for human tracking. *Proceedings of 2010 IEEE International Conference on Intelligent Robots and Systems, Taipei, Taiwan*, 2938–2945. doi:10.1109/IROS.2010.5652272
- Glas, D. F., Satake, S., Kanda, T., & Hagita, N. (2011). *An interaction design framework for social robots*. Paper presented at the meeting of Robotics: Science and Systems, Los Angeles, CA, USA. Retrieved from <http://www.roboticsproceedings.org/rss07/p14.html>
- Glas, D. F., Kanda, T., Ishiguro, H., & Hagita, N. (2012). Teleoperation of multiple social robots. *IEEE Transactions on Systems, Man, and Cybernetics -- Part A: Systems and Humans*, 4(3), 530–544. doi:10.1109/TSMCA.2011.2164243
- Gross, H.-M., Boehme, H.-J., Schroeter, C., Mueller, S., Koenig, A., Martin, C., . . . Bley, A. (2008). Shopbot: progress in developing an interactive mobile shopping assistant for everyday use. In *IEEE International Conference on Systems, Man, and Cybernetics (SMC2008)* (pp. 3471–3478). doi:10.1109/ICSMC.2008.4811835
- Hayashi, K., Shiomi, M., Kanda, T., & Hagita, N. (2011). Friendly patrolling: A model of natural encounters. In *Proceedings of Robotics: Science and Systems (RSS 2011)*. Retrieved from <http://www.roboticsproceedings.org/rss07/p18.html>
- Haumann, A. D., Listmann, K. D., & Willert, V. (2010). DisCoverage: A new paradigm for multi-robot exploration. In *IEEE International Conference on Robotics and Automation (ICRA2010)* (pp. 929–934). doi:10.1109/ROBOT.2010.5509993
- Hussein, I. I., & Stipanovic, D. M. (2007). Effective coverage control for mobile sensor networks with guaranteed collision avoidance. *IEEE Transactions on Control Systems Technology*, 15, 642–657. doi:10.1109/TCST.2007.899155
- Inaba, M., Kagami, S., Kanehiro, F., Hoshino, Y., & Inoue, H. (2000). A platform for robotics research based on the remote-brained robot approach. *International Journal of Robotics Research*, 19, 933–954. doi:10.1177/02783640022067878
- Iwamura, Y., Shiomi, M., Kanda, T., Ishiguro, H., & Hagita, N. (2011). Do elderly people prefer a conversational humanoid as a shopping assistant partner in supermarkets? *6th ACM/IEEE International Conference on Human-Robot Interaction (HRI2011)* (pp. 449–456). doi:10.1145/1957656.1957816
- Kanda, T., Shiomi, M., Miyashita, Z., Ishiguro, H., & Hagita, N. (2009). An affective guide robot in a shopping mall. In *4th ACM/IEEE International Conference on Human-Robot Interaction (HRI2009)* (pp. 173–180). doi:10.1145/1514095.1514127
- Kanda, T., Glas, D. F., Shiomi, M., & Hagita, N. (2009). Abstracting people's trajectories for social robots to proactively approach customers. *IEEE Transactions on Robotics*, 25, 1382–1396. doi:10.1109/TRO.2009.2032969
- Kirby, R., Forlizzi, J., & Simmons, R. (2010). Affective social robots. *Robotics and Autonomous Systems*, 58, 322–332. doi:10.1016/j.robot.2009.09.015
- Lee, M. K., Forlizzi, J., Rybski, P. E., Crabbe, F., Chung, W., Finkle, J., . . . Kiesler, S. (2009). The snackbot: Documenting the design of a robot for long-term human-robot interaction. In *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction* (pp. 7–14). La Jolla, California, USA: ACM. doi:10.1145/1514095.1514100
- Matsumoto, Y., Wada, T., Nishio, S., Miyashita, T., & Hagita, N. (2010). Scalable and robust multi-people head tracking by combining distributed multiple sensors. *Intelligent Service Robotics*, 3(1), 29–36. doi:10.1007/s11370-009-0056-5

- Monteiro, S., & Bicho, E. (2010). Attractor dynamics approach to formation control: Theory and application. *Autonomous Robots*, *29*, 331–355. doi:10.1007/s10514-010-9198-8
- Mutlu, B., & Forlizzi, J. (2008). Robots in organizations: The role of workflow, social, and environmental factors in human-robot interaction. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI2008)* (pp. 287–294). doi:10.1145/1349822.1349860
- Nishio, S., Hagita, N., Miyashita, T., Kanda, T., Mitsunaga, N., Shiomi, M., & Yamazaki, T. (2008). Structuring information on people and environment for supporting robotic services. In *Proceedings of 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2008)* (pp. 2637–2642). doi:10.1109/IROS.2008.4651198
- OMRON Corporation (n.d.). OKAO vision. In *OMRON Global*. Retrieved August 19, 2012, from http://www.omron.com/r_d/coretech/vision/okao.html
- Parker, L. E. (1998). ALLIANCE: An architecture for fault tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, *14*, 220–240. doi:10.1109/70.681242
- Saffiotti, A., Broxvall, M., Gritti, M., LeBlanc, K., Lundh, R., Rashid, J., . . . Cho, Y. J. (2008). The PEIS-Ecology project: Vision and results. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 2329–2335). doi:10.1109/IROS.2008.4650962
- Sanfeliu, A., Andrade-Cetto, J., Barbosa, M., Bowden, R., Capitán, J., Corominas, A., . . . Spaan, M. T. J. (2010). Decentralized sensor fusion for ubiquitous networking robotics in urban areas. *Sensors*, *10*, 2274–2314. doi:10.3390/s100302274
- Satake, S., Kanda, T., Glas, D. F., Imai, M., Ishiguro, H., & Hagita, N. (2009). How to approach humans? Strategies for social robots to initiate interaction. In *4th ACM/IEEE International Conference on Human-Robot Interaction (HRI2009)* (pp. 109–116). doi:10.1145/1514095.1514117
- Shiomi, M., Kanda, T., Ishiguro, H., & Hagita, N. (2007). Interactive humanoid robots for a science museum. *IEEE Intelligent Systems*, *22*(2), 25–32. doi:10.1109/MIS.2007.37
- Shiomi, M., Sakamoto, D., Kanda, T., Ishi, C. T., Ishiguro, H., & Hagita, N. (2008). A semi-autonomous communication robot: A field trial at a train station. In *ACM/IEEE 3rd Annual Conference on Human-Robot Interaction (HRI2008)* (pp. 303–310). doi:10.1145/1349822.1349862
- Shiomi, M., Kanda, T., Glas, D. F., Satake, S., Ishiguro, H., & Hagita, N. (2009). Field trial of networked social robots in a shopping mall. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2009)* (pp. 2846–2853). doi:10.1109/IROS.2009.5354242
- Shiwa, T., Kanda, T., Imai, M., Ishiguro, H., & Hagita, N. (2009). How quickly should a communication robot respond? Delaying strategies and habituation effects. *International Journal of Social Robotics*, *1*(2), 141–155. doi:10.1007/s12369-009-0012-8
- Siegwart, R., Arras, K. O., Bouabdallah, S., Burnier, D., Froidevaux, G., Greppin, X., . . . Tomatis, N. (2003). Robox at Expo.02: A large scale installation of personal robots. *Robotics and Autonomous Systems*, *42*, 203–222. doi:10.1016/S0921-8890(02)00376-7
- Sisbot, E. A., Alami, R., Simeon, T., Dautenhahn, K., Walters, M., Woods, S., . . . Nehaniv, C. (2005). Navigation in the presence of humans. In *Proceedings of IEEE International Conference on Humanoid Robots* (pp. 181–188). doi:10.1109/ICHR.2005.1573565
- Takayama, L., Marder-Eppstein, E., Harris, H., & Beer, J. M. (2011). Assisted driving of a mobile remote presence system: System design and controlled user evaluation. In *IEEE International Conference on Robotics and Automation (ICRA2011)* (pp. 1883–1889). doi:10.1109/ICRA.2011.5979637

- Thrun, S., Bennewitz, M., Burgard, W., Cremers, A. B., Dellaert, F., Fox, D., . . . Schulz, D. (1999). Minerva: A second-generation museum tour-guide robot. In *IEEE International Conference on Robotics and Automation (ICRA1999)* (pp. 1999–2005). doi:10.1109/ROBOT.1999.770401
- Tomizawa, T., Ohya, A., & Yuta, S. (2006). Remote shopping robot system, -Development of a hand mechanism for grasping fresh foods in a supermarket. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2006)* (pp. 4953–4958). doi:10.1109/IROS.2006.282457
- Trautman, P., & Krause, A. (2010). Unfreezing the robot: Navigation in dense, interacting crowds. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2010)* (pp.797–803). doi:10.1109/IROS.2010.5654369
- Vig, L., & Adams, J. A. (2006). Multi-robot coalition formation. *IEEE Transactions on Robotics*, 22(4), 637–649. doi:10.1109/tro.2006.878948
- Weiss, A., Bernhaupt, R., Tscheligi, M., Wollherr, D., Kühnlenz, K., & Buss, M. (2008). A methodological variation for acceptance evaluation of human-robot interaction in public places. In *IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN2008)* (pp. 713–718). doi:10.1109/ROMAN.2008.4600751

Authors' names and contact information: D. F. Glas, ATR Intelligent Robotics and Communication Laboratories, Kyoto, Japan. Email: dylan@atr.jp; S. Satake, ATR Intelligent Robotics and Communication Laboratories, Kyoto, Japan. Email: satoru@atr.jp; F. Ferreri, ATR Intelligent Robotics and Communication Laboratories, Kyoto, Japan. Email: florent@atr.jp; T. Kanda, ATR Intelligent Robotics and Communication Laboratories, Kyoto, Japan. Email: kanda@atr.jp; H. Ishiguro, Intelligent Robotics Laboratory, Osaka University, Toyonaka, Japan. Email: ishiguro@sys.es.osaka-u.ac.jp; N. Hagita, ATR Intelligent Robotics and Communication Laboratories, Kyoto, Japan. Email: hagita@atr.jp