

# The Snackbot: Documenting the Design of a Robot for Long-term Human-Robot Interaction

Min Kyung Lee<sup>1</sup>, Jodi Forlizzi<sup>1,3</sup>, Paul E. Rybski<sup>2</sup>, Frederick Crabbe<sup>4</sup>, Wayne Chung<sup>3</sup>,  
Josh Finkle<sup>3</sup>, Eric Glaser<sup>3</sup>, Sara Kiesler<sup>1</sup>

<sup>1</sup>HCI Institute, <sup>2</sup>Robotics Institute  
Carnegie Mellon University  
5000 Forbes Ave.  
Pittsburgh, PA 15213  
{mklee, forlizzi, prybski, kiesler}  
@cs.cmu.edu

<sup>3</sup>School of Design  
Carnegie Mellon University  
5000 Forbes Ave.  
Pittsburgh, PA 15213  
{wcchung, jfinkle, eglaser}@cmu.edu

<sup>4</sup>United States Naval Academy  
Computer Science Department  
572C Holloway Rd. Stop 9F  
Annapolis, MD 21402  
crabbe@usna.edu

## ABSTRACT

We present the design of the Snackbot, a robot that will deliver snacks in our university buildings. The robot is intended to provide a useful, continuing service and to serve as a research platform for long-term Human-Robot Interaction. The design process, which occurred over 24 months, included a study of snacking and snack services, prototyping, empirical user testing, and form factor design studies. We describe the organization of the project, and the interdisciplinary decision making that led to a holistic design for a useful robot that will behave appropriately in a public setting, and will hopefully encourage healthy snacking.

## Categories and Subject Descriptors

A.m. [Miscellaneous]: Human Robot Interaction – *Social Robots*

## General Terms

Design, Human factors, Documentation

## Keywords

Social robot, design process, interaction design, holistic design

## 1. INTRODUCTION

Experimental systems, including receptionists, assistants, guides, tutors, and social companions, have been developed as platforms for research and technology development [3][4][5][6][8][10][11][13][14][17][21][25][29][31][33][34]. Much of this work has taken place in the research laboratory, but a few systems have made the successful transition to real world settings such as museums and educational institutions [15][17][22][27][28]. Real world settings raise the bar to fluid, natural interaction with robotic systems.



Figure 1. The Snackbot robot.

Robots in real settings also need to interact with people appropriately. Safe interactions are necessary to contribute to ethical research in the field, to improve people's trust in and comfort with robotic technology, and to ensure safety and reliability for all who come into contact with this technology. Socially appropriate interaction behaviors are needed so people like the robot and are interested in interacting with it over time.

Our research group seeks to develop robots that travel around and near people, and that support people in real-world environments. We are interested in developing robots that act as social assistants, with the ability to use speech and gesture, and engage people in a social manner. A major goal is to create robots that are not just mobile but also interact over a period of time, performing a service.

Many questions about long-term HRI are unanswered. How do people's perception and attitudes towards a robot evolve over time? What interaction design strategies will reinforce a positive long-term relationship between people and a robot? Will

employees use a robot in the way that designers intended or will they appropriate the robot in new ways, as has happened with other technologies [9]? Could robots deliver services that are beneficial to people over the long term?

To address these challenges, we designed and developed the Snackbot, a robust robot that will roam semi-autonomously in campus buildings, offering snacks to office residents and passersby (Figure 1). We designed the Snackbot not just as a service, but also as a research platform to investigate questions related to long-term interaction with social robots.

Our future work with the Snackbot will involve field trials with the robot in its actual context of use. Such research poses several technical, interaction, and design challenges. First, the robot must be robust and powerful enough to operate autonomously and interact with multiple users for extended periods of time. The technology should also be flexible enough to accommodate technical improvements and new applications. To test different approaches to human-robot interaction over time, researchers should be able to manipulate aspects of the robot's physical appearance and behavior. We are particularly interested in how a robot delivers a service after the initial novelty effect has worn off.

In this paper, we present our design process for the Snackbot, shaped by our initial design goals, constraints we discovered along the way, and design decisions guided by interim empirical studies. We describe design process recommendations for designing robotic products and services for long-term human-robot social interaction.

## 2. CONTEXT OF USE

Robotic advances are being directed towards special populations, including elders, those with physical and cognitive disabilities, and others. We want to design robots that can interact with almost everyone, regardless of any dispositions to using technology. To satisfy this goal, we are interested in how a robot can deliver a service within a work environment.

We chose to design a robot that would provide snack deliveries in the two connected buildings in which we work. By "snack" we mean light food eaten between meals. Snacks include "junk food" such as food offered in vending machines, and "healthy" snacks such as fruit and nuts. Snacking is practiced by a majority of people in the developed world [1][24][32]. In workplaces, people snack in their offices and labs as well as in halls, cafeterias, and food vending areas.

A robot delivering snacks must have a wide range of mobility. The buildings are large, ranging between 4 and 8 floors. About 1000 people work or visit these buildings each day. Because the buildings offer only prepackaged snacks in convenient locations, we felt a snack service that offered higher quality snacks would be a useful application for a long-term product and service in these buildings. Most snacks that do exist are highly caloric, and the robot could include healthier snacks in its offerings. We felt many technical and design research questions could be discovered in understanding how a robotic snack service might succeed within the social and environmental context of our buildings, how it would differ from traditional vendors and vending machines, and how it could support people's goals such as taking a break from work and delivering snacks as gifts to people. We have described some of the research supporting these decisions in a separate paper [18]. This research, combined with our overall research

goals in HRI, led to the three design goals that anchored our design process.

## 3. DESIGN GOALS

We had three design goals for development of the Snackbot robot:

The first was to *develop the robot holistically*. Rather than advancing technology *per se* or focusing on one aspect of design or interaction, such as a dialogue system, we took a design approach that considered the robot at a human-robot-context systems level [23]. Such an approach allowed us to think about the emergent qualities of the product and service, which might not be recognized if the system were analyzed in component parts rather than holistically.

The second goal was to *simultaneously develop a robotic product and service*. By this we mean that the robot as a product would have to be more than sociable and attractive; it would need to deliver something useful to people. We adopted this goal to increase the likelihood that people would continue to be interested in interacting with the robot over a period of time [20]. By developing a snack delivery service that worked with wireless service points in the building, we could collect and record knowledge about people's snack preferences, and use these to further enhance the service we provide to them.

The third goal was to *develop interaction designs that would help to evoke social behavior*. Because the robot was meant to serve as a research platform that would be used by people over time, decisions about functions and features were made supporting the interest of promoting sociability. For instance, we aimed to have the robot interact with people using natural language. Other research has shown that people interact with a robot longer when it exhibits social cues [5][12]. Other aspects of sociability that we plan to explore and extend include personalization of the service, and robot politeness and non-verbal behaviors [2].

## 4. SNACKBOT TEAM

The Snackbot team consisted of 5 faculty, 5 graduate students, and 7 undergraduate students drawn from several disciplines including design, behavioral sciences, computer science, and robotics. Because of the wide range of expertise, we frequently had members from one group attending the meetings of the other. For instance, the designers worked on the form studies but they often interacted with the engineers, and everyone helped out with the empirical studies. Organization of this group was assisted through the use of an on-line forum called the Kiva ([www.thekiva.org](http://www.thekiva.org)), hosted on a website accessible to team members from anywhere on the Internet. We found this web facility to be more useful than email because all of the information was organized and presented in a searchable, threaded format to the entire team.

The team was organized in a hierarchical fashion with the lead design faculty and robotics faculty taking on the roles of organizing the designers and engineers, respectively, and the lead graduate student organizing the behavioral studies. In each case, each group maintained its own set of requirements and goals. The group lead would manage the efforts of the individuals assigned to their specific group and would be in charge of making sure that their contribution was properly integrated in the whole plan. For instance, the technology lead was responsible for developing the overall system architecture and ensuring that the individuals who were contributing to that followed the engineering standards that were defined for that system.

A great deal of emphasis was placed on good documentation of process as well as code and interim prototypes. The Snackbot system will be worked on by future generations of students, and each new person on the project must be able to follow in the footsteps of those that worked on it before.


## 5. SYSTEM OVERVIEW

To give the reader a snapshot of what the design process has achieved and where we plan to go, we present an overview of the system depicted in Figure 1. At the time of this writing, the Snackbot robot is in its final assembly as the four-and-a-half-foot tall, semi-autonomous humanoid robot shown in the figure. It will traverse the hallways of our buildings, delivering snacks to residents in offices and labs. The Snackbot will have its own “office” where people can send email or IM for ordering snacks (or sending snacks as gifts to others). We also plan extensions of the service. For example, the Snackbot might visit a group’s lounge area and invite the group to socialize around a snack appropriate for that purpose (e.g., bagels and cream cheese in the morning).

### 5.1 Hardware

The Snackbot robot is based on the existing CMAssist platform [26], augmented with some commercial hardware and software and new elements and code. The Snackbot uses a MobileRobots Inc. Pioneer 3 DX base for mobility. Bumpers, sonars, and a SICK laser are used to detect and avoid collisions and to detect position within an environment. A Hokuyo URG laser is mounted in the robot’s chest to detect potential collisions with higher objects, and to detect people by torso.

The Snackbot currently has non-functional arms that hold a tray, used for carrying snacks (Figure 2). The tray is equipped with 12 load cells; each is capable of measuring a weight range of 13 to 763 grams. This functionality will be used so the robot knows when someone has removed or replaced a snack on the tray.



**Figure 2. Sketch of the Snackbot tray, showing the load cell configuration.**

An Acoustic Magic microphone array is mounted under the tray. It serves as the primary audio input source for the robot’s natural spoken language and dialog processing system. The robot’s head is mounted on a Directed Perception pan/tilt unit, affording a 360 degree pan range and a 111 degree tilt range. A Point Grey Bumblebee 2 stereo camera is mounted behind the robot’s eyes; a monocular Point Grey Dragonfly2 camera is mounted on the top of the head and is fitted with a 180-degree fisheye lens from

Omnitech Robotics. The Snackbot also has two 2.4GHz Intel laptops running Ubuntu Linux for data processing.

### 5.2 Software

The Snackbot uses MobileRobot’s ARIA API that works with the Pioneer base. ARNL provides functionality for map construction, and path planning. A distributed software architecture developed by the CMAssist project [26] interfaces with the behavior control modules and the speech processing interface. When the Snackbot moves through its environment, it will track its current position by comparing the current set of laser scans and an odometry estimate against a previously programmed map.

We use an Augmented Transition Network (ATN) manager for our dialogue system. This will allow for a flexible discourse structure, but will require more work by a dialogue designer. We also use an open source Sphinx4 speech recognizer system (<http://cmusphinx.sourceforge.net/sphinx4/>), written in Java, and the Cepstral speech synthesizer [19].

### 5.3 Form

The form of the robot is made of cast fiberglass and is custom designed to fit the Pioneer base and an internal structure that anchors the laptop and other components. It has a semi-humanoid form and uses simple geometric shapes. There are three exterior pieces: one for the head, one for the torso, and one for the base.

### 5.4 Interaction

There are several basic modes of interaction with the robot. In stationery mode, the robot is positioned in a social space and people can approach the robot to help themselves to a snack. In roaming mode, the robot uses the map to visit people’s offices and to deliver snacks. Snacks can be ordered in advance (using a web page, email, or IM) or selected during the visit.

To interact with the Snackbot, people eventually will engage in natural dialogue with the dialogue system. Visual feedback will occur through an LED mouth, which will indicate when the robot is “talking.” Sound will be used as an additional informational cue.

## 6. DESIGN PROCESS

To holistically conceive of the robot as a product and service, we had to consider many aspects of the design process concurrently: its form, how it would interact with people, and how it would fit to the physical constraints and social norms of an office environment. We also wanted to ensure our design would truly foster long-term interaction. For this reason, at the outset, we decided on a semi-humanoid form for the robot, since other work has shown that humanoid robots are interesting to people [30].


### 6.1 Context Research

We conducted research on snacking in our office buildings, described in more detail elsewhere [18]. Our environmental research took the form of a campus survey to document all of the places where people can get snacks. From candy dishes in administrative offices to vending machines in the basements of building, we mapped site lines and studied each site for accessibility. We also mapped distances to, and popularity of, nearby locations that are popular for snack breaks — for instance, a local coffee shop that is frequented by members of the campus community. One of the findings from this work was that people mainly choose convenience over snack quality, but they do not mind walking for a snack if social interaction is part of the activity

(and especially if the snack is free). Based on our observations, we created two basic modes of interaction for the robot: mobile and stationary. We decided that the robot in mobile mode should offer to deliver healthy snack choices such as fruit, and that in stationary mode should offer high quality snacks in communal locations that would attract groups. These decisions support our overall design goal of evoking social behavior, and ensure that we are not making a robot that will only bring fattening snacks to inactive people.

We also did product research on types of robot form. Our product research took the form of collecting and analyzing images of existing social robots, which ranged from animal to abstract to humanoid forms. We categorized these into four types: humanoid, abstract, hybrid, and other. Humanoid robots were of interest, because they mimic the anatomy and form of the human figure. However, humanoids are mechanically complex, and for our research, may not be robust enough for long-term use in the field. Abstract robots were less relevant because they have a mechanical aesthetic, showing tracks, wheels, and other parts that do not invite human interaction at an intimate level. Hybrid robots were of greatest interest, because they combine simple geometric forms with human cues. This was a good choice for further investigation, as the housing design would then allow for the holistic combination of hardware and aesthetic components.

We generated sketches based on the hybrid concept. Two types of sketches were initially explored: more industrial, mechanical forms with wide shoulders, aggressive stance, and masculine proportions, and more playful, cute forms with rounded proportions and childlike faces (Figure 3). To support our goal of social interaction by making the robot approachable by everyone, we merged these two styles to create a gender-neutral, friendly, yet professional-looking form to fit the context of our university.



**Figure 3. Sketches for the robot housing: a) machine-like, b) rounded and friendly, c) concepts combined.**

## 6.2 Iterative Design Studies

We conducted empirical studies to investigate and support our design process. Here, we describe four of them as examples: an early technology feasibility study, an early interaction study, a dialogue study, and a height and approach study. Each of these

tests was conducted in support of our overarching design goals for holistic development, product and service, and social behavior. Each generated design implications for our robot and tradeoffs with other aspects of the system. The process and results of these iterative studies are described in this section.

### 6.2.1 Early Technology Feasibility Study

We assembled some of the robot's key capabilities on an existing mobile robot platform, the CMAssist robot [26]. The goal of this study was to test and verify the basic functionality of the major components of the system, and to ensure that it would work smoothly with the wireless network in our buildings.

The robot, partly tele-operated, traversed hallways for two-hour periods, over a two week period, in the two campus buildings described above, and prompted passersby to take free snacks. The robot was controlled by an experimenter using a joystick about 20 yards away. The dialogue system was also run using streamed audio and five human-controlled utterances, allowing us to quickly understand the timing and robustness for this type of dialogue system. To help our technology prototype look like an aesthetic robot design, we created a housing with vacuum-formed materials and foam core components, so the robot had a clean, unified appearance (Figure 4).

This early trial helped us learn about many tradeoffs we would face in the future design of the hardware, software, and interaction design of the robot. We subsequently decided to use a commercially-available base for the Snackbot. A Pioneer base would be more reliable than a home-built base, and would provide mobile functions that would be easily replicable. It would also be quieter and less distracting to office residents. One drawback of using this kind of base is that it would create a set of constraints for the final industrial design of the robot housing. Such constraints included the dimensions of the robot, the availability (or lack thereof) of mounting points for the torso, and maximum load that could be carried. Our plans for the torso and other electronics exceeded the recommended weight limit of the Pioneer, and so later experiments were performed to learn the maximum reasonable weight the robot could carry while still having reasonable operational battery life.

In terms of software, we learned that it would be feasible to entirely automate the dialogue structure using a finite number of preset phases, because conversation with the robot quickly revealed stable patterns. We also learned that we would need to devise ways to deal with network lag or drop-off and still preserve the idea of a sociable, fluidly interacting robot. This led us to pursue the interaction study described in the next section.



**Figure 4. Prototype used for early technology feasibility study.**

### 6.2.2 *Early Interaction Study*

Our early interaction design study took the form of three semi-structured trials with the first robot prototype in two campus buildings. Here, our goal was to come up with archetypical dialogue structures for interacting with the Snackbot, to support our design goals of product and service and robust social interaction.

We used Wizard of Oz methods, where a remote dialogue operator used Skype and interactively “chatted” with snack customers. A separate operator performed motor control of the robot using a joystick and tether. We adopted the convention of American ice cream trucks, and developed a 30-second melody and a cheery “Hello!” for the robot to announce itself in the hallways. Interaction with customers was structured in that the Skype operator had a script to follow, but could deviate from it in real time if needed.

We learned that people found the melody and greeting to be too annoying for use in an office building. This was partly due to the fact that the sound was played from a low-quality speaker, and therefore distorted, but the social norms of an office environment also play a role. We also learned that a minimal, straightforward design of the dialogue would be all that is needed, because people readily filled in dialogue and other social cues, such as indicating which snack that they intended to take off the tray by showing it to the robot’s eye cameras, and by politely repeating phrases during their interactions. These findings suggested methods for collecting speech and environmental sound as input for the dialogue system, and gave us ideas for how to specifically design and study the dialogue system, which we describe in the next section.

### 6.2.3 *Dialogue Study*

We next conducted a study to verify our design of the dialogue structures and scenarios. Our overarching goal was to discover how to provide dialogue with the robot in a way that evokes social behavior and allows the service to proceed as intended.

We created general dialogue excerpts and ran them in a Wizard of Oz study with 12 participants. One experimenter ran the robot’s dialogue scripts in a remote location, and another noted what the participant said in response to the scripts. We used the stationery mode as a scenario for the study — passersby approached the robot and discussed what snacks were available that day.

We learned several things about our first iteration of the dialogue design. First, nearly half of the phrases we designed were unsuitable in that that people frequently deviated from the script as we designed it. We added phrases to control for unintelligible speech or users wandering off topic. We also learned that people liked to play with the dialogue structure to see where it might fail. For example, if the Snackbot asked, “Is this your first visit?” a participant might answer “I have been here lots of times but I have never seen you,” instead of giving a simple yes or no answer. Although we tried to structure the dialog to discourage such behavior, we were unsuccessful. We subsequently added phrases to try to smooth over these communication breakdowns.

We found that care needs to be taken in constructing the output phrases so that they are intelligible and imitate human intonation. Although our synthesizer is state-of-the art, certain words, phrases, and spellings can result in difficult to understand speech. The synthesizer has trouble particularly with the rise of voice expected when people ask questions. For example, “Would you

like an apple?” sounds strange with synthesized speech intonation. Thus we learned the Snackbot should instead say, “I want to know if you would like an apple,” to eliminate intonation issues.

We found that some participants used visual cues much more than others, thereby minimizing the use of dialogue. In particular, they tended to examine the tray rather than asking what snacks were offered, and to simply remove the item without verbally indicating what they would like, despite a direct question. We learned that we would need to tightly couple the dialogue system with the sensor system to adequately track all of the non-verbal communication in support of evoking social behavior.

Other interesting social interactions were observed, such as groups of people interacting with the robot. Group conversation presents a difficulty for the speech recognition system, which is unlikely to differentiate person-to-person conversations from those targeted towards the Snackbot. Some of this difficulty can be mitigated with careful integration with other sensors. To best understand where to place these sensors, we undertook a height and approach study described in the next section.

### 6.2.4 *Height and Approach Study*

Rather than arbitrarily deciding the height of the robot, we wanted to learn whether the height of the robot affects people’s approach interactions with the robot, and affects the data collected. To our knowledge, there have been no formal studies about the body size of a robot. Therefore, we conducted a study to discover what an appropriate height might be for the Snackbot.

We conducted a between-subjects experiment with 72 participants using the technology feasibility prototype described earlier. The robot had three height conditions, 44 inches (112 centimeters), 52.5 in (129 cm), and 56.5 in (144 cm). We chose these three heights as deviations from the average height of a short human being with an average reach of lower arm length, so it would be comfortable to approach and take a snack from the robot even in the shortest condition. We did not want to make the robot taller than people in order not to be threatening.

The study was conducted in a public area of our campus. We offered free snacks for participating in the survey. We used a 5-point Likert scale to understand how friendly and intelligent people felt the robot was, and how they responded to the height of the robot. An open-ended question asked participants to list the personality traits that they ascribe to the robot. We also asked participants their gender, age, and height.

Using a 5-point scale where 1 = much too small and 5 = much too tall, participants preferred the tallest robot most on a,  $F [1,71] = 4.10, p < .02$ . The smallest and mid-sized robots averaged 2.4, meaning the robot was between “too small” (score of 2) and “just right” (score of 3). The tallest robot was almost just right with a mean of 2.9. Participants liked the fact that they could make eye contact with the tallest robot, and disliked that they had to bend to interact with the smaller two robots. There were also interesting differences in the personality traits people attributed to each prototype. The smallest robot most frequently was described as servile, obedient, and submissive. We felt that to best support our goal of evoking social behavior, the robot should be seen more as a co-worker than as a servant. In our university culture, even the least skilled workers are given respect more as peers than as servants. This consideration also indicated that the tallest robot would be most appropriate.

Armed with the findings from the early technology feasibility study, the early interaction study, the dialogue study, and the height and approach study, we built the second prototype of the Snackbot.

## 7. SECOND PROTOTYPE

We embarked on designing a more robust, refined system, using a Pioneer base. This decision was made to support our design goals of offering a product and service in our office environment, by reducing the distracting noise, and ensuring operation over long periods of time and a variety of floor types.

From our interaction study, we learned that we would need to develop a set of sensors that would allow us to know when a snack was taken. Because we did not want to overload the vision system, which would eventually support person recognition, we added a mid-chest laser and pressure sensors to the robot's tray. These additions would also support natural social behavior between Snackbot and its customers.

To develop the second prototype, we focused on the development of the housing, the design of the tray, and the development of an internal structure to anchor the sensors, laptops, and housing to the Pioneer base. We also focused on the expressive qualities of the robot's face, and finalizing the interaction design.

### 7.1 Housing

Working from the early sketches described above, we built a number of quarter-scale and half-scale models of the robot. After ascertaining correct proportions for the dimensions of the Pioneer, laptops, and other internal components, we constructed a full-scale mock-up to proportions, radii, and design details (Figure 5). Using this prototype, our design team was able to address dimensions, hardware placement, configurations, assembly, and tray size and arm options.




Figure 5. Full-scale mock-up of the robot.

To check responses to the full-scale model, we placed it in a hallway in our building and conducted a survey with 59 participants to understand positive or undesirable associations to the design. Participants rated the robot as friendly and likable (mean 3.88 and 3.87), and neither intelligent or unintelligent (mean 3.42). The robot evoked descriptions of service jobs such as a waiter or waitress, or general Sci-Fi characters such those from the Jetsons. Based on these responses, we felt that the final form design supported our three design goals.

One of the issues with the housing was weight. The Pioneer has a recommended payload limit of 50 lbs. for carrying additional

weight. From our experiments, we determined that between 70 and 80 lbs. of weight was still reasonable for the robot to carry and still have an acceptable operational lifetime. This drove the selection of fiberglass as the material for the outer housing and aluminum 80/20 as material for the inner housing. We also segmented the base and made a variety of cuts in the torso to reduce weight. The resulting housing is lightweight, strong, and easy to add attachments for internal materials. Other components were constructed out of neoprene fabric to create a holistic aesthetic form.

After generating a number of color studies for the housing, we selected a color scheme of medium gray and orange. Both hues do not cause gender attributions or strong attributions of service type in the U.S. culture. For example, a blue robot might connote a medical service, due to the ubiquitous use of the color blue in the health sector. A green robot might connote a sustainable product. Orange is also often associated with food and restaurants. Together, the orange, gray, and dark gray of the neoprene created a distinctive, impressive form.

### 7.2 Tray

The tray design (Figure 2) was developed for providing food or snacks at the appropriate delivery height, but also as an input system for measuring the weight and presence or absence of items on the tray. The tray has movable slots that can be configured in a number of ways to hold different snacks. The tray is made of aluminum, styrene, 12 load cells, and a cloth covering on which snacks will be placed.

### 7.3 Internal Structure

The design for the internal structure continued to evolve as the external housing design was finalized. To minimize weight while providing maximum strength, three vertical struts of extruded aluminum were used as the base for the design. We augmented these with a custom aluminum plate at the top of the Pioneer base and one at the shoulder, to mount internal components. These additions allowed for retrofitting to a variety of components using off-the-shelf brackets and anchors. These design decisions afford modularity, which support our overall goal of holistic design.

### 7.4 Head and Face

Our overall goal was to create an expressive head that would serve as a locus of interaction, relying on appropriate features that convey the right level intelligence and functionality to the robot [10].

The final head design features a simple form that is wider than tall, suggesting a young, friendly robot. The width of the Bumblebee camera also determined the width of the head and the placement and size of the eye sockets. The necessity of the camera suggested a simple design for the face. We also felt that by minimizing complexity and detail in the eyes, Snackbot customers would not develop false perceptions about the intelligence of the robot.

A 3 x 12 LED display was developed for the mouth, serving as an expressive focus for interaction. The mouth is programmed with a series of animations that show verbal and emotional feedback in the form of lip shapes, colors and movement. While the robot does not have functional ears, we added ears to the head design, so that customers would understand that the robot can hear them.

## 7.5 Interaction Design

The final interaction design for the robot includes stationary and mobile delivery modes that provide a variety of services to our university community and support social behavior. We have designed a basic interaction infrastructure, so we can use and vary these modalities to conduct experiments once the robot is fully implemented. Snacks can also be ordered through a web site or IM service. The interaction design, of course, will evolve as we conduct field trials with the Snackbot.

## 8. LESSONS LEARNED

We have spent almost two years on the holistic development of a robotic system, and we have learned several lessons. We articulate them here, relative to our overall design goals.

Our first lesson was to understand how the robot will actually work in a context of people, other products, a physical environment, and social norms. Then, in the service of holistic design, the design of particular subsystems can be undertaken. For instance, our first technology, dialogue, and interaction designs did not support the demands of our office environment. Of course, the general point is not new. Many others have articulated the need to design for the context (e.g., Jones and Hinds [16]).

Our second lesson in terms of holistic design was to design for modularity. Functions should be developed individually, but with an eye to the constraints caused by other aspects of the system. Modularity also means that components can be upgraded or changed as new and better systems become available. For example, the selection of a Pioneer base created weight constraints, which became an issue in the design of the housing. Again, others in various fields have recommended designing in modularity (e.g., Cai and Sullivan [7]).

In terms of product and service, we learned the robot should offer capabilities that add value to people's lives, and allow them to add value themselves through interacting with the robot. This idea drove our choice to offer healthy snacks, and to provide a stationary mode that invites people to take a walk to the robot. Future experiments will be done to understand whether and how the robot's interaction design can be modified to best support people.

In terms of social behavior, we learned to work to make a social robot sociable within the limitations of current technology. For example, we needed to make iterative changes to the dialogue system to both support fluid and natural social interaction while working with the constraints that the wireless network provided.

None of these lessons, taken independently, are new, as the HRI community will recognize. What we think is a contribution is our showing how we tried to tackle all these lessons together. The larger lesson is that designing for all these goals is what is really hard. It requires a design team dedicated to an interdisciplinary holistic design process.

## 9. CONCLUSION AND FUTURE WORK

In this paper, we have presented the design and development process for the Snackbot robot, a robot that is designed for long-term delivery of snacks in our building. We had three overarching goals for the development of this system: to develop a robot holistically, to develop the robot as a product and a service within our building, and to design the robot in a way that evokes social behavior.

As of this writing, we have yet to entirely assemble the robot as designed. For this reason, we cannot currently evaluate whether the Snackbot has met our overall design goals. Our future research will assess our success relative to these goals through extensive field trials. We have also designed the robot in a flexible manner so that we can change the form and behavior for more controlled study.

## 10. ACKNOWLEDGEMENTS

We gratefully acknowledge the support of the National Science Foundation through grants IIS-0121426, IIS-0624275, and CRI 0709077, and the support of Microsoft Research. The Kwanjeong Educational Foundation provided a graduate fellowship to the first author. We also acknowledge the fine work of Outlaw Performance, Rubright's Frame, Lombardi Productions, and the assistance of Tarun Agarwal, Ryan Anderson, Alvin Barton, Levi Boyles, Andrea Bravi, Andrea Irwin, Jessica Jones, Jaldert Rombouts, Aubrey Shick, Scott Smith, Jeremy Stolarz, Kevin Yoon, Matthijs Zwinderman, Tijs Zwinkels, and Scott Hudson.

## 11. REFERENCES

- [1] Bellisle, F., Dalix, A.M., Mennen, L., Galan, P., Hercberg, S., de Castro, J.M., & Gausseres, N. (2003). Contribution of snacks and meals in the diet of French adults: A diet-diary study. *Physiological Behavior*, 79, 183-189.
- [2] Bickmore, T. W., and Picard, R. W. (2005). Establishing and maintaining long-term human-computer relationships. *ACM Transactions on Computer-Human Interaction*, 12, 293-327.
- [3] Blow, M., Dautenhaun, K., Appleby, A., Nehaniv, C.L., and Lee, D. (2006). The art of designing robot faces: Dimensions for human-robot interaction. *Proceedings of HRI'06*. New York, NY: ACM Press, 331-332.
- [4] Breazeal, C. (2003). Emotion and sociable humanoid robots. *International Journal of Human-Computer Studies*, 59, 119-155.
- [5] Bruce, A., Nourbakhsh, I., and Simmons, R. (2002). The role of expressiveness and attention in human-robot interaction. *Proceedings of ICRA'02*, 4138-4142.
- [6] Burgard, W., Cremers, A.B., Fox, D., Hähnel, D., Lakemeyer, G., Schulz, D., Steiner, W., and Thrun, S. (1999). Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114, 3-55.
- [7] Cai, Y. and Sullivan, K. J. (2005). A value-oriented theory of modularity in design. *ACM Sigsoft Software Engineering Notes*, 30, 4.
- [8] Dautenhaun, K., Walters, M., Woods, S., Koay, K.L., Nehaniv, C.L., Sisbot, A., Alami, A. and Simeon, T. (2007). How May I Serve You?: A robot companion approaching a seated person in a helping context. *Proceedings of HRI'07*. New York, NY: ACM Press, 172-179.
- [9] DeSanctis, G., and Poole, M. S. (1994). Capturing the complexity in advanced technology use: Adaptive structuration theory. *Organization Science*, 5, 121-147.
- [10] DiSalvo, C., Gemperle, F., Forlizzi, J., and Kiesler, S. (2002). All robots are not created equal: The design and perception of humanoid robot heads. *Proceedings of DIS'02*. New York, NY: ACM Press, 321-326.

- [11] Fong, T., Nourbakhsh, I., and Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42, 143-166.
- [12] Forlizzi, J. (2007). How robotic products become social products: An ethnographic study of cleaning in the home. *Proceedings of HRI'07*. New York, NY: ACM Press, 129-136.
- [13] Gockely, R., Simmons, R., and Forlizzi, J. (2007). Natural person-following behavior for social robots. *Proceedings of HRI'07*. New York, NY: ACM Press, 17-24.
- [14] Gockley, R., Simmons, R., and Forlizzi, J. (2006). Interactions with a moody robot. *Proceedings of HRI'06*. New York, NY: ACM Press, 186-193.
- [15] Hayashi, K., Sakamoto, D., Kanda, T., Shiomi, M., Koizumi, S., Ishiguro, H., Ogasawara, T., and Hagita, N. (2007). Humanoid robots as a passive-social medium: A field experiment at a train station. *Proceedings of HRI'07*. New York, NY: ACM Press, 137-144.
- [16] Jones, H., and Hinds, Pamela. (2002). Extreme work teams: Using SWAT teams as a model for coordinating distributed robots. *Proceedings of the Computer Supported Cooperative Work Conference '02*. NY: ACM Press.
- [17] Kanda, T., Takayuki, H., Eaton, D., and Ishiguro, H. (2004). Interactive robots as social partners and peer tutors for children: A field trial. *Human-Computer Interaction (Special Issue on Human-Robot Interaction)*, 19, 61-86.
- [18] Lee, M.K., Kiesler, S., and Forlizzi, J. (2008). How do people snack? Understanding the context of a mobile robot snack service. Unpublished ms. Carnegie Mellon Univ., Pittsburgh, PA. 15213.
- [19] Lenzo, K.A., and Black, A.W., Theta, Cepstral, <http://www.cepstral.com>.
- [20] Morelli, N. (2002). Designing product/service systems. A methodological exploration. *Design Issues*, 18, 3-17.
- [21] Mutlu, B., Osman, S., Forlizzi, J., Hodgins, J., and Kiesler, S. (2006). Perceptions of ASIMO: An exploration of co-operation and competition with humans and humanoid robots. *Proceedings of HRI'06*. New York, NY: ACM Press, 351-352.
- [22] Nabe, S., Kanda, T., Hiraki, K., Ishiguro, H., Kogure, K., and Hagita, N. (2007). Analysis of human behavior to a communication robot in an open field. *Proceedings of HRI'07*. New York, NY: ACM Press, 234-241.
- [23] Nelson, H.G., and Stolterman, E. (2003). *The Design Way*. Englewood Cliffs, NJ: Educational Technology Publications.
- [24] Ovaskainen, M.L., Reinivuo, H., Tapanainen, H., Hannila, M-L., Korhonen, T. and Pakkala, H. (2006). Snacks as an element of energy intake and food consumption, *European Journal of Clinical Nutrition*, 60, 494-501.
- [25] Pollack, M.E., Engberg, S., Matthews, J. T., Thrun, S., Brown, L., Colbry, D., Orosz, C., Peintner, B., Ramakrishnan, S., Dunbar-Jacob, J., McCarthy, C., Montemerlo, M., Pineau, J., and Roy, N. (2002). Pearl: A mobile robotic assistant for the elderly. *Proceedings of AAAI Workshop on Automation as Eldercare*, AAAI, 2002.
- [26] Rybski, P. E., Yoon, K., Stolarz, J., and Veloso, M. M. (2007). Interactive robot task training through dialog and demonstration. *Proceedings of HRI'07*. New York, NY: ACM Press, 49-56.
- [27] Severinson-Eklundh, K., Green, A., and Huttenrauch, H. (2003). Social and collaborative aspects of interaction with a service robot. *Robotics and Autonomous Systems*, 42, 223-234.
- [28] Shiomi, M., Kanda, T., Ishiguro, H., and Hagita, N. (2006). Interactive humanoid robots for a science museum. *Proceedings of HRI'06*. New York, NY: ACM Press, 305-312.
- [29] Sidner, C.L., Lee, C., Morency, L.-P., and Forlines, C. (2007). The effect of head-nod recognition in human-robot conversation. *Proceedings of HRI'07*. New York, NY: ACM Press, 290-296.
- [30] Siino, R. and Hinds, P. (2004). Making sense of new technology as a lead-in to structuring: The case of an autonomous mobile robot. *Academy of Management Best Paper Proceedings*, New Orleans.
- [31] Simmons, R., Goldberg, D., Goode, A., Montemerlo, M., Roy, N., Sellner, B., Urmson, C., Schultz, A., Abramson, M., Adams, W., Atrash, A., Bugajska, M., Coblenz, M., MacMahon, M., Perzanowski, D., Horswill, I., Zubek, R., Kortenkamp, D., Wolfe, B., Milam, T., and Maxwell, B. (2003). Grace: An autonomous robot for the AAAI robot challenge. *AAAI Magazine*, 24, 51-72.
- [32] Verplanken, B. (2006) Beyond frequency: Habit as mental construct, *British Journal of Social Psychology*, 45, 639-656.
- [33] Wada, K. Shibata, T. Saito, T., and Tanie, K. (2002). Robot assisted activity for elderly people and nurses at a day service center. *Proceedings of ICRA'02*, 1416 - 1421.
- [34] Yamaoka, F., Kanda, T., Ishiguro, H., and Hagita, N. (2006). How contingent should a communication robot be? *Proceedings of HRI'06*. New York, NY: ACM Press, 313-320.