Need Finding: A Tool for Directing Robotics Research and Development

Caroline Pantofaru and Leila Takayama
Willow Garage, Inc.
Menlo Park, California, USA
Email: {pantofaru,takayama}@willowgarage.com

***Abstract*—Directing research in any field is difficult, but personal robotics presents a unique challenge. Many standard tools such as user studies, market research and incremental technology research rely on having a clear system design or a well-defined task. Personal robotics, however, is relatively new and poorly defined; personal robots can do anything that helps people. In this paper, we present *need finding* as a tool for exploring the personal robotics design and application space, and share our experiences with using this method for understanding robotics hobbyists as a user group. We expect that need finding will become widely used alongside other techniques to help answer the question: What should robots do?**

I. Introduction

One of the pressing challenges for personal robotics, and robotics more broadly, is finding answers to the burning questions: What should robots do? What are robots good for? Who will actually buy and use them? Finding answers to these questions not only gives direction to robotic products, but increases the relevance of robotics research and moves the field closer to making a difference in society. Indeed, large research funding agencies like the National Science Foundation judge research proposals upon the criteria of broader impact and intellectual merit.

For mature product markets such as factory automation, there exist clear methods and metrics for developing robotic products that involve building a list of technical requirements and specifications [9]. Unfortunately, these techniques are not applicable to the emerging market of personal robotics which is poorly defined in the areas of market segmentation, user needs (which the users themselves might not even know [7]), and technological capabilities to fulfill those needs.

For personal robotics to become a widespread reality, it is critical for us to not only look for new markets, but also to create them. To do this, we need a different set of research and design tools. Fortunately, there are some great lessons to be learned from product designers, who work at the forefront of technology development. In this paper, we share our perspectives on working with one such product design tool called *need finding* [1] to explore possible designs and applications for personal robotics. *Need finding* is a methodological tool for both grounding and inspiring research that has broader impact. The power of need finding is that it allows us to quickly enter a user's world, casting the user as an expert who will teach us about the meaning of tasks, objects and technology to them. Stepping back from the space of robots and into the space of tasks and meaning generates intuition, context, and more creative ideas for the future direction of personal robotics.

II. Current methods for directing research

Current methods for determining research and development directions in the robotics community can be roughly divided into two categories: user-driven and researcher-driven. User-driven methods involve direct interactions with potential customers of a technology, including market research, surveys and user studies. Researcher-driven methods focus on the researcher's (or developer's) expertise and intuition, including determining the state of the art in technology and its possible direct applications, determining the next incremental step in developing a technology, or intuitively deciding what is interesting. All of these methods have distinct strengths and weaknesses, which we discuss here. Need finding is by no means a replacement for these other tools, but it can serve as an additional source of information.

A. User-driven methods

Market research includes analyzing customers, markets and the risks of bringing a product to market. The encompassed techniques are effective for products near completion, for which a market already exists. Market research is, however, difficult to do for a disruptive technology like personal robotics which does not yet have a market.

Surveys are an effective way of reaching a wide audience with low cost, obtaining a large and varied sample. The recent growth of crowd-sourcing venues such as Amazon's Mechanical Turk has made surveys even easier to distribute. The difficultly in assessing survey responses, however, lies in separating what people say versus what they actually do. Among the many drawbacks of relying too heavily upon self-reported data is that people are more prone to exaggerate, oversimplify, or simply forget aspects of their story. A potentially even more problematic issue with purely self-reported data is the social desirability bias, which is an inclination for respondents to say things that they think will be viewed upon favorably by others. For example, if you ask parents what kinds of foods they feed their children, they are more likely to tell you about healthy meals rather than unhealthy snacks, but actually seeing their fridge and pantry might show you a more complete version of their story. There is even a standardized scale for measuring how much a person is inclined toward the

social desirability bias [2]. These problems can be so serious that some usability experts will suggest completely ignoring what users say, and instead relying much more heavily upon watching what they do [7].

Many human-robot interaction studies have taken the form of user studies, allowing users to interact with an existing system or a mock-up of a future system (Wizard-of-Oz) [3]. User studies are fundamentally time consuming, limiting the number of participants. To generate statistically significant results with a small number of participants necessitates limiting the number of variables being examined. Practically speaking, this implies that the system to be studied must be well-specified, with only a very small number of components to be tested and only a small number of options for each component. This is an excellent method for almost-complete systems, but a poor method for radical new directions.

B. Researcher-driven methods

The majority of robotics research is driven by considering the next logical step. Given the state of the field, which new tasks seem plausible to address, and what are the incremental pieces of technology needed to address those tasks? This method ensures steady progress, but rarely results in dramatic changes in direction. More problematically, this method is only informed by the state of technology and not by the needs of real users, so the adoption of resulting technologies is questionable.

A similar approach is the proverbial hammer searching for a nail. A researcher or company may possess a technique which has been successful in the past, so it is applied to an ever-increasing array of problems. This has many of the same problems as incremental development, with the added downside that the technology may be inappropriate for the chosen application.

Finally, pure intuition can drive design and in fact an entire research field or industry. This is an extremely powerful design motivation, especially since domain experts spend years building intuition around their technology and market. Disruptive technologies are by definition unforeseen by users; they are created by the researchers and developers and pushed onto users. This is a very risky and potentially very rewarding approach. Inspiration, however, does not happen everyday, and so more structured exploration of a design space can be useful.

III. Need Finding

A. What is need finding?

Need finding is a method that comes from the product design research community [1]. Its goal is to identify fundamental user needs of a community of people that a product aims to satisfy. In the case of robotics, need finding is a great complement to advanced technology development (technology push) in that need finding presents us with opportunities to generate technology pull.

On the surface, need finding resembles ethnographic methods in cultural anthropology, making observations and interviewing users, but these two methods are quite different in that need finding involves more than "thick description" of a culture [5]. Need finding also involves generating design imperatives and product solutions, which are developed along with empathy and frameworks for understanding the culture [1]. Design imperatives are the high-level set of must-have features or capabilities that provide real value to customers by addressing their needs. Design solutions are concepts, products, and prototypes of products that meet real user needs and directly follow the design imperatives.

Need finding is somewhat similar to needs assessment [6] in its goals. Needs assessment emerged from the organizational level of analysis, particularly in the domains of educational institutions, governments, and corporations. While the goals of need finding (from design research) are quite similar to the goals of needs assessment, they have different histories and slightly different protocols. Both approaches clearly separate needs apart from solutions; needs are used to drive toward solutions, but are not solutions themselves. This distinction is critical in that it explicitly identifies needs as existing without presupposed solutions. Here is one canonical example: Astronauts need to record notes while in outer space. They do not necessarily need high tech pens that operate in 0G to write on paper (which is a presupposed solution); they might also simply use pencils on paper or voice recorders.

There are also several fundamental differences between need finding and needs assessment. Needs assessment aims to identify gaps between results and consequences, whereas need finding aims to identify gaps between use, usability, and meaning. The scope of needs assessment work is typically very broad, working at many levels of an organization (mega, macro, and micro). In contrast, the need finding work presented here is oriented toward products and driving research and development work that will ultimately have broad impact for end-users.

B. How is need finding done?

Typically, need finding begins with identifying a set community of people who might be potential users (or extreme users) of products that a company produces. A very small team goes out to visit a sample of members from that community in the places where they do the activities that the products might support. Immersion in the interviewee's environment is one of the keys to success, inspiring the participant to discuss details they might have otherwise forgotten, and allowing the interviewer to quickly reconcile the interviewee's words with reality.

As compared to other types of user engagements, the sample size is relatively small; a sample size of eight to sixteen interviews is typical. Cultural Consensus Theory [8] states that for a sufficiently culturally coherent group, a small number of interviews will yield responses which have a high probability of being representative of the whole group. So determining how much the chosen user group knows about and agrees on their culture, in turn determines how many interviews are required to obtain representative information.

Each interview and visit to a user's space might involve doing observational work, as well as interviews (e.g. talking with people about their experiences), and walk-throughs of spaces or tasks. The interview takes place in the participant's environment so that the participant's stories can be compared against reality immediately. An interesting difference versus traditional ethnography is that the interviewers may direct the conversation to keep on topic and make the process time-efficient. The goal is to extract useful information, not simply compile observations. A visit with a need finding participant might last anywhere from one to many hours.

Along with getting a sense of the physical, cultural, and social environments, artifacts, and activities, the need finding team is focusing upon hearing good stories. The types of stories and observations that can be particularly interesting and fruitful for insight are contradictions (e.g., parents saying they want to feed healthy food to their children, but in practice actually feeding them unhealthy foods), spoken and unspoken social norms (e.g., not ending a personal relationship via text message), and success or failure stories [4]. Stories that interviewees volunteer may reflect situations in which reality did not match their mental model of a situation, thereby expressing an unfulfilled need.

Data gathered from the interviews forms the basis of the four components of the need finding process: generating empathy, creating frameworks, generating design imperatives and finally, creating solutions and prototypes. Generating empathy for the interviewees and potential users is best done by actually speaking to them and visiting their environment, so it is valuable to have researchers and engineers actually go out into the field. However, as it is not possible for each person to talk to each interviewee, each interview should be presented to the entire team through videos, images, stories, or other appropriate media.

Frameworks are structured higher-level concepts extracted from the field data. They can be formulated as timelines, hierarchies of needs, 2-dimensional axes, or any other format that expresses an idea seen across multiple interviews.

Design imperatives are generated from the frameworks and represent attributes which a system must have (or omit) in order to satisfy the user group. Even at this stage, it is preferable not to presuppose solutions. Returning to the example of astronauts recording notes in space, a design imperative might be that the astronauts must be able to record information while floating; another imperative may be that they need to recall the information within two seconds even if they have moved to another location; finally an astronaut must always feel in control. None of these imperatives presupposes that the astronaut must actually write down the information, leaving open other solutions such as voice recording or relating the information to another astronaut.

With empathy for the user group, frameworks representing their needs and viewpoints, and design imperatives making these needs concrete, the process of generating solutions and prototypes is informed and responsive to the pull of user needs. The prototypes generated are hypotheses of what users want. Promising prototypes can be subjected to the traditional exercise of user testing to obtain more concrete feedback.

Although this method seems linear, it is usually not. Articulating stories, framework, insights, and solutions is not a cleanly sequential exercise; most teams bounce back and forth between these sets.

C. Why do need finding?

Although need finding takes some time (e.g., several weeks for a user group), money, and effort, we have found that its benefits have been well worth these expenses.

First, the research and development team gains a sense of empathy with the user group at hand. We faced the original stereotypes and assumptions we had made about this user group and replaced as much as we could with grounded observations, real people, and real relationships with them.

Second, researchers gain inspiration and insight into potential research directions which can make a difference to users by satisfying important (and often unspoken) user needs. Not only do they generate specific solutions and prototypes, but they also gain generative frameworks from which to produce more solutions and design imperatives with which to drive research decisions. This inspiration may also result not in pursuing an application itself, but rather an enabling basic technology that would make an application possible in the future. Even such basic research profits from user insight. For example, object recognition research would benefit from knowing which objects it needs to recognize for a specific task, instead of simply resorting to readily available, randomly chosen objects.

Third, the research team forms relationships with the user community that it is aiming to support and can bounce ideas and prototypes with those users early and often through the iterative design process. As we have seen in the human-computer interaction community, rapid iterative designs coupled with user testing, feedback, and evaluations can be a very fruitful way to make easy-to-use products that people actually want to use. The robotics industry can leverage those lessons learned from the computer industry to its advantage.

IV. The Reality of Need Finding for Robotics

To gain insight into the practice of need finding, we sent 14 researchers and developers out into the field to perform 7 interviews with a common user group of robotics hobbyists. These interviews were performed in user-chosen spaces that reflected where most of their hobby work was done, such as living rooms, garages and communal labs. Each session consisted of an interview portion and a tour of the participant's workspace. Some of the sessions also included demonstrations of the participant's current work. After concluding the interviews, the entire need finding team reconvened for several hours to share and reconcile observations. The user community needs were converted into frameworks and a set of design imperatives for robotics hobbyists. Finally, prototypes of possible research projects and products were generated that reflected the design imperatives. Below are some of the lessons learned from this initial exploration.

A. Challenges in executing need finding

Need finding for robotics as conducted by roboticists presented a set of interesting challenges. Non-industrial robotics is a nascent field at best, so interviews on the topic required careful wording. Participants who were well informed about the state of the art in robotics were extremely pessimistic about robot abilities. On the other hand, participants who were less informed tended to have completely unrealistic expectations. A quick Internet search shows that people have very strong opinions about what robots are and what they should do. To mitigate these effects, it was important to understand the participants knowledge of the field. Attempts to work around the problem by asking participants to share their desires for robotics regardless of their belief in the state of the art were unsuccessful. Note that were we discussing a mature or existing technology, such as a dish washer, we could expect reasonable answers to questions such as, "What would you improve about this product? or, "What should this product be used for? Asking these questions about robots, however, is futile.

We were much more successful when discussing ideas at a *task* level and discussing the desired end-result, without asking which technology should be used to get there. A particularly successful approach consisted of asking participants how they typically execute a task, to give an example of a successful execution and an example of an execution that went awry. For example, we asked the robotics hobbyists about a typical day of working on their robot, an example of good progress, and an example of a frustrating work session. Requesting specific examples and stories eliminated some of the self-reporting bias.

Allowing roboticists to execute need finding interviews provides challenges in itself. By their nature and to their advantage, roboticists want to solve problems and have opinions about how to do so. In some of the interviews, these characteristics manifested themselves through the interviewer (roboticist) asking leading questions. In the most extreme cases, the interviewer would attempt to engineer solutions to the participants problems during the interview. Unfortunately, this negates the goal of the interview which is to obtain previously unknown information about people and tasks. If the roboticists themselves are to perform the interviews, we strongly recommend training in listening techniques, especially how to orient toward the interviewee as the expert, who is teaching the interviewer about his or her activities. If interviewers take on a dominant position in relation to the interviewee (e.g., asking questions that feel to the interviewee like, "Why aren't you doing things in my far-superior way?"), that discourages interviewees from opening up and sharing stories. Keeping interview questions at a task level instead of about technology can ameliorate roboticist bias as well.

B. Positive experiences with need finding

Despite the challenges, need finding was a very positive experience. In fact, 12 of the 14 researchers and developers have volunteered to participate in a second need finding initiative (the other two were subject to time constraints).

The data gathering process was not very time consuming. Each interview took 2-3 hours, with a different team assigned to each participant. The most laborious task was summarizing the interviews for the other teams, including creating short video clips and a very short overview presentation. This communication had to be done well to generate empathy for the interviewee. The teams then spent a day comparing notes and coming to shared conclusions, frameworks and design imperatives. This exercise was excellent for team building, as well.

By the end of data amalgamation, the results were surprisingly coherent. The majority of basic needs discovered in the interviews were the same, despite coming from different individuals interviewed by different teams. Generating these shared results was also an excellent team building exercise for researchers and developers. Instead of disagreeing on technologies and methods (as researchers are trained to do), the team could agree on user needs.

The process also generated empathy for the users and a shared vocabulary among the researchers. During brainstorming and prototyping sessions after need finding, researchers could relate specific components of their ideas to a user need. This elevated the level of the conversation and reduced the time spent arguing over semantics. The roboticists have also kept in contact with the participants, with opinions being requested and given in both directions.

The results of this first need finding expedition generated buy-in from researchers and developers to undertake another, larger effort.

V. Conclusions

Deciding what we should do next as a research community is often difficult and controversial. In this paper, we have presented need finding as an approach to exploring the research and design space. Need finding is an additional tool to add to the toolbox that sits alongside other user- and researcher-driven exploration methods. It will not answer all of our questions, but need finding can help to improve intuition, generate ideas and direction, and most importantly, create a shared context that provides a basis for future conversations.

References

[1] S. Beckman and M. Barry. Innovation as a Learning Process: Embedded Design Thinking. *California Management Review*, 50(1):25–56, 2007. URL http://insideinnovation.haas.berkeley.edu/Presentations/2007Conference/Beckman.pdf.

[2] D. P. Crowne and D. Marlowe. A new scale of social desirability independent of psychopathology. *Journal of Consulting Psychology*, 24:349–354, 1960.

[3] N. Dahlback, A. Jonsson, and L. Ahrenberg. Wizard of Oz Studies: Why and How. In *Proceedings of the*

1st International Conference on Intelligent User Interfaces, pages 193–200, 1993. URL http://portal.acm.org/ft_gateway.cfm?id=169968&type=pdf.

[4] M. Fellman. Breaking Tradition. *Marketing Research*, 1: 20–25, 1999.

[5] C. Geertz. From the Native's Point of View: On the Nature of Anthropological Understanding. *Culture Theory: Essays on Mind, Self, and Emotion*, pages 123–136, 1984.

[6] R. Kaufman and F. W. English. *Needs Assessment: Concept and Application*. Educational Technology Publications, 1979.

[7] J. Nielsen. First rule of usability: Dont listen to users. http://www.useit.com/alertbox/20010805.html, 2001.

[8] AK Romney, SC Weller, and WH Batchelder. Culture as consensus: A theory of culture and informant accuracy. *American Anthropologist*, 88(2):313–38, 1986.

[9] K. Ulrich. *Product Design and Development*. McGraw-Hill/Irwin, 1999.