Exposing real-world challenges using HRI in the wild

Aaron G. Cass, Kristina Striegnitz, Nick Webb, and Venus Yu
Union College
Schenectady, NY, USA 12308
{cassa,striegnk,webbn}@union.edu, venus0yu@gmail.com

ABSTRACT
We describe an experiment designed to determine actions a robot can perform to reliably enlist help from passers-by in a public space. Based on our experiences with this experiment, we argue in favor of carrying out Human Robot Interaction (HRI) studies in the wild versus in the lab. Our position is that field studies expose challenges that a robot will face that would be difficult to anticipate and simulate in a laboratory setting, and thus will produce results that better generalize to real-world situations.

ACM Classification Keywords
H.5.2 User Interfaces: Natural Language; I.2.9 Robotics: Autonomous Vehicles

Author Keywords
WoZ; Human-Robot Interaction; Field Studies

INTRODUCTION
We are engaged in an ongoing research program to design robot behaviors appropriate for interaction with people where the robot needs assistance to perform a sub-task of a larger goal. As part of this program we are interested in determining what robot actions allow the robot to reliably enlist help from passers-by in public spaces. To empirically evaluate these robot actions, we have recently designed and carried out a series of experimental trials in a busy public space on our campus.

In this paper, we use our experiences from those recent trials to argue in favor of doing such experiments "in the wild" versus in laboratory settings. We have found that our field studies expose challenges that a robot will face in real-world situations and which would be difficult to anticipate and simulate in a laboratory setting.

RELATED WORK
A few studies, such as [5, 1], evaluate robots that ask for help from people in public spaces. However, they have primarily focused on path planning and different approach trajectories. In other scenarios robots approach people for help who are seated and working or in a meeting [6, 4, 3]. In those scenarios, the main challenge is not attracting attention, but limiting annoyance at the interruption. However, the findings are still relevant and inform our experiment design. For example, they highlight the importance of analyzing the behavior of passers-by since people are less likely to help when they are busy [3], and they confirm that it is possible to design robot behaviors that express subtle information about the robot's goals [6]. We have used some of these ideas in designing our own robot experiments in the wild [2].

EXPERIMENTAL TRIALS
We have recently carried out a series of experimental trials in a popular thoroughfare on campus, a hallway where a significant number of people pass through en route to classrooms or other public spaces. Our robot SARAH[1] was controlled using a Wizard of Oz protocol with two experimental conditions: one with no movement from SARAH, and the other with a rotational movement that followed a subject we wished to engage. In both, SARAH would greet people with a verbal expression ("Hello! Can you please help me?"). If people registered interest by approaching the robot, they were asked to press a specific number on a keypad.

We recorded, for offline analysis, a stream from a webcam attached to SARAH. In total we collected approximately 14 hours of data over 14 sessions on 13 different days over the course of 5 weeks. During our experiment, 1658 people passed by SARAH. Of those, only 714 engaged with her in any way, including just looking at her. Of the 714, 108 completed our task. We found that movement of the robot increased the number of engagements (313 vs. 401) and the number of completed tasks (51 vs. 57), but only the difference in number of engagements was statistically significant.

CHALLENGES
These experiments exposed challenges for a robot in a populated public space, which would be difficult to anticipate and simulate in a laboratory setting. Therefore, our position is that many robot behaviors that attract attention from primed subjects in relatively tame laboratory settings would not work in public spaces. Thus results from a laboratory experiment would likely lead to proposed robot actions that work only in the lab but not in a real-world deployment.

[1] Socially Appropriate Robot that Approaches for Help, a ROS-enabled, human-scaled robot with LIDAR, Microsoft Kinect, and bump sensors, built on a Pioneer P3-DX base with an added "torso" supporting a monitor on which a face is displayed.

PubRob'18, September 03, 2018, Barcelona, Spain

ISBN .
DOI:

For example, our assumption had been that getting a subject's attention would lead to task completion. Consequently, we could have created tests in a controlled environment to demonstrate increased attention getting behaviors. However, when a robot in the wild employs the behaviors successful in the lab, it would not necessarily be able to attract more attention. In the following we discuss both the challenges encountered in field studies and the experiment validity benefits they provide.

In the wild, SARAH is competing for attention from people who don't expect to interact with a robot, who are busy going places, surrounded by other people traveling in many different directions, inundated with the din of activity around them, and engaged in conversation, using their mobile phones, or preoccupied with other activities. In a lab setting, subjects do not have the same demands on their attention and, in most cases, expect to interact with a robot.

In our experiments, SARAH interacted with a broad population. Some were drawn to her simply because she was a robot, while others avoided any contact with her. They seemed to be afraid of the robot, annoyed by it, or reluctant to draw attention to themselves in a public space (groups seemed more willing to approach SARAH). In a laboratory study on human-robot interaction, we suspect we would have a much narrower cross-section of the population because participant recruitment is necessarily self-selecting. We would likely get mostly subjects who want to engage with a robot. Robot behaviors that work with such a sample of people might not work for the wider populace. Furthermore, people may be more comfortable interacting with a robot in the privacy of the lab than in a public space.

In our experiment, we asked subjects to help SARAH by pressing a number on a keypad. However, it seems that some people were less inclined to help because they realized that this task did not really address a need of the robot or another person. We propose for future experiments a more meaningful task, such as calling the nearby elevator for SARAH. In a lab study, however, subjects would perform whatever task the experimenters specified. Also, in the wild the robot must make its needs clear, while in a lab the experimenters can motivate the task.

We also found that sound volume played a critical role. After the fact, we determined that the level of SARAH's voice was the same as the ambient noise in the busy hallway, thus making it difficult for some subjects to hear SARAH's requests for assistance. In a laboratory setting, with fewer distractions and subjects primed to pay attention to the robot, we suspect we would not have noticed any issues with sound volume.

For our experiments the biggest limitation we faced was the lack of experimenter awareness of activity around SARAH. We used only one camera mounted to SARAH with a limited field of view, and in addition there was a lag between the wizard interface and the robot. The wizard was relying on the webcam data to initiate a conversation, and as a result of the lag and limited field of view, people were often already past SARAH before she uttered her first greeting, and the opportunity to interact was missed. Thus, the timeliness and richness of the information available to the wizard (or autonomous robot) as well as the timing of the robot behavior is crucial. Clearly, additional onboard sensing and improved communication between the wizard interface and the robot can be used to improve on this situation. However, in a laboratory study, we may not have ever discovered this problem because the interactions would be more orchestrated.

In summary, we find that executing these kinds of experiments in the wild has illuminated genuine issues robots will face in public spaces. Our strongest position is that experiments have to be valid, in the sense that they should have a meaningful task, and engage with unprimed subjects in a real-world setting wherever possible. Further, we have to ensure that the robot encounters a wide range of participants. We are planning studies in spaces such as grocery stores and malls, where the robot will encounter a different audience than on a college campus. Finally, working in these environments exposes the robot to seemingly minor but critical challenges of timing, interacting at exactly the right moment, and volume, being heard without being overbearing.

REFERENCES

1. Andrea Bauer, Klaas Klasing, Georgios Lidoris, Quirin Mühlbauer, Florian Rohrmüller, Stefan Sosnowski, Tingting Xu, Kolja Kühnlenz, Dirk Wollherr, and Martin Buss. 2009. The Autonomous City Explorer: Towards Natural Human-Robot Interaction in Urban Environments. *International Journal of Social Robotics* 1, 2 (2009), 127–140.
2. Aaron G Cass, Eric Rose, Kristina Striegnitz, and Nick Webb. 2015. Determining appropriate first contact distance: trade-offs in human-robot interaction experiment design. In *Workshop on Designing and Evaluating Social Robots for Public Settings at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2015)*. IEEE, Hamburg, Germany, 6.
3. Helge Hüttenrauch and Kerstin Severinson Eklundh. 2003. To Help or Not To Help a Service Robot. In *The 12th IEEE International Workshop on Robot and Human Interactive Communication (ROMAN)*. IEEE, Millbrae, CA, USA, 379–384.
4. Stephanie Rosenthal, Manuela Veloso, and Anind K. Dey. 2011. Is Someone in this Office Available to Help Me? Proactively Seeking Help from Spatially-Situated Humans. *Journal of Intelligent Robotic Systems* 66 (2011), 205–221.
5. S. Satake, T. Kanda, D. F. Glas, M. Imai, H. Ishiguro, and N. Hagita. 2009. How to approach humans?-strategies for social robots to initiate interaction. In *Human-Robot Interaction (HRI), 2009 4th ACM/IEEE International Conference on*. IEEE, La Jolla, CA, USA, 109–116.
6. Paul Saulnier, Ehud Sharlin, and Saul Greenberg. 2011. Exploring minimal nonverbal interruption in HRI. In *Proceedings of the IEEE International Symposium on Robots and Human Interactive Communications(RO-MAN)*. IEEE, Atlanta, GA, USA, 79–86.