The Social Co-Robotics Problem Space: Six Key Challenges

Laurel D. Riek
Department of Computer Science and Engineering
University of Notre Dame
Email: lriek@nd.edu

***Abstract*—In order to realize the long-term vision of intelligent co-robots capable of competent proxemic interaction with humans, it is important that our research community fully define the problem space and articulate its inherent challenges. In particular, we must recognize that many problems in the space may not be computable or definable, and must determine ways to address this challenge moving forward. This paper broadly sketches six key challenges in the social co-robotics problem space and suggests several paths toward solving them, such as Wizard-of-Oz constraints and problem satisfication.**

I. Introduction

Three recent U.S. Government reports and funding initiatives in robotics - the CCC Robotics Roadmap [26], the National Intelligence Council Global Trends 2030 [5], and the National Robotics Initiative [17] - all strongly emphasize the theme that in order to realize the vision of intelligent, capable co-robots, robots must be able to operate intelligently in close proximity to (and with) humans.

The co-robotics problem domain includes both proximate and remote interaction [12], and covers a wide range of human environments. In this paper, we focus specifically on problems relating to co-robots in *human social environments* (HSE). These are any environments in which a robot operates proximately with humans. We define these robots as *social co-robots*, in that they operate in an HSE, are physically embodied, and have at least some degree of autonomy.

It is worth noting that social co-robots are not necessarily *sociable* [2] - they do not necessarily need to interact with us interpersonally. For example, a service robot that empties the dishwasher may not be sociable, but because it operates in an HSE with the aforementioned characteristics it is a social co-robot.

Fig. 1 depicts the broad application space for co-robots in human environments, and emphasizes a few application areas where social co-robots in HSEs are warranted. Exemplar areas include personal assistive robots (physical and cognitive), educational robots, robots for leisure activities, service robots, and robots in clerical domains. These application areas are not meant to be mutually exclusive, but the majority of problems in, for example, biomolecular or field robotics, are not usually social in nature.

This paper will define six challenges unique to the social co-robotics problem domain (Section II), then suggest some possible avenues to explore for addressing them (Section III).

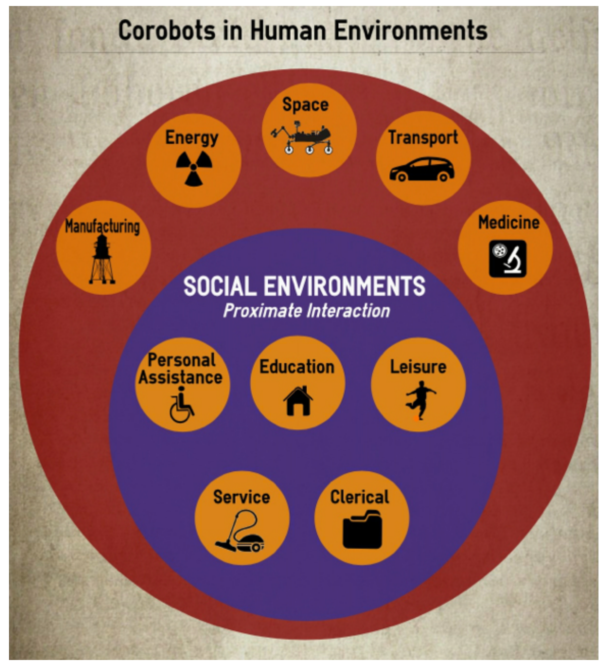

Fig. 1. The co-robotics problem domain can be divided into two domains: human social environments, where proximate interaction is required, and others, where remote interaction is more common. *None of these areas are intended to be mutually exclusive.*

II. Six Key Challenges in Social Co-Robotics

Fig. 2 depicts several of the unique set of challenges that social co-robotics faces. While many of these problem domains are tied to the three traditional robotics problems (perception, cognition, and action), they are much more complex in scope. These problems may not be computable, or even definable. Further, several of these problems resolve into fundamental AI-complete problems, such as natural language understanding [14], making them intractable.

It is critical as a community we articulate the inherent hardness of these problems, and recognize there is no silver bullet for solving them [24].

A. Problem 1: Dynamic Spaces

One thing unique to social co-robots is that they must operate in HSEs with humans present. Humans by their very existence create unforeseen challenges to robotics that are still relatively new to the field. A workshop parallel to this one

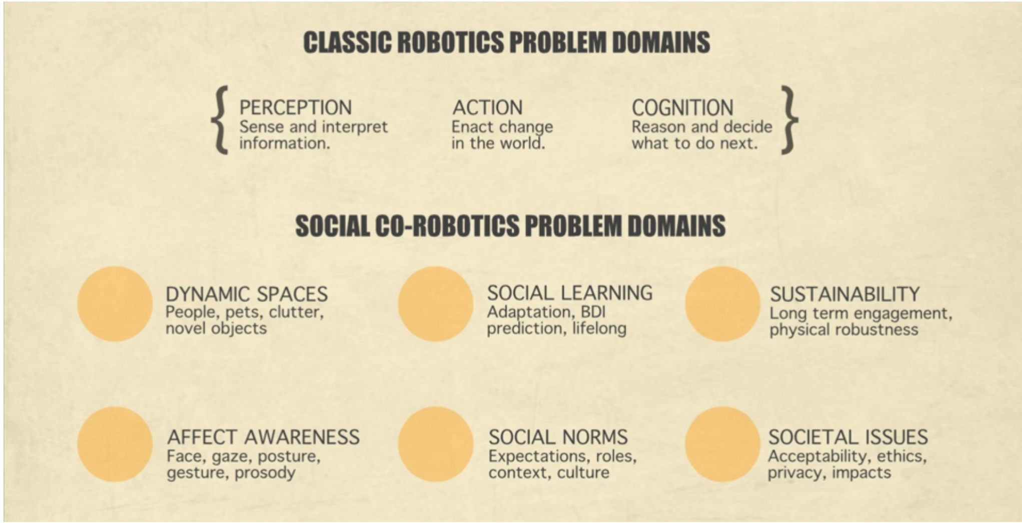

Fig. 2. Social co-robotics has a unique set of challenges, many of which are entirely unlike traditional robotics problems in terms of their level of complexity as well as the challenge in adequately defining them. Six problem domains are listed above.

at RSS is called "Robots in Clutter", and its CFP describes many of these challenges: vision under cases of clutter and occlusion, dynamic navigation, action planning on the fly, blind manipulation, and "sporadic user involvement" [35].

HSEs are often highly fluid and variable in nature, which makes it difficult to anticipate or plan for these environmental alterations as a robot. Traditional control paradigms, even those that are well suited to and robust within dynamic environments (e.g., [3, 11]) do not scale well to these complexities.

Part of the problem is that the way in which a robot contends with a dynamic HSE is closely tied to its task and embodiment, thus making it difficult to generalize the problem. As robotics researchers, we can clearly envision a general solution to the 2005 DARPA Grand Challenge that is vehicle-independent. We cannot, however, envision a general solution to a personal assistance robot that helps a person with severe physical disabilities compete their daily tasks[1]. Solving such a problem is multifaceted, contains potentially an infinite number of dynamically changing subtasks, and any solution attempt is intricately tied to the robot's embodiment and capabilities.

Thus, this problem area may not even be definable for robotics, let alone computable. However, perhaps one way to address it is through the use of social learning, described in the next section.

B. Problem 2: Social Learning

One of the most critical things a social co-robot needs to be able to do is learn and adapt to not only its environment, but the co-humans within it. However, because HSEs are highly fluid (as are people), it is important that social co-robots are capable of lifelong learning [31]. This learning may occur under direct human tutelage or independently.

In 2010, Carlson et al. [6] introduced NELL, the Never Ending Language Learner. This project is teaching a machine to "read the web", by extracting facts while crawling textual webpages 24 hours a day, 7 days a week. Another project, RobotsFor.Me [32], may also be suitable to help enable lifelong learning by letting people remotely log into a PR2 24/7 and teach it in an embodied, situated way.

However, even with these projects, learning about and adapting to humans, and anticipating and modeling their BDIs (beliefs, desires and intentions) is quite challenging, and suffers from the same problem as the *Dynamic Spaces* problem. Models for one paradigm with one robot with one set of capabilities (e.g., a PR2 in an office environment, with end-effectors and a Kinect) do not necessarily extend to others. BDI modeling may be computable within a virtual agent space, but may reach a point of being non-computable (or non-definable) when embodied on a co-robotic system.

C. Problem 3: Sustainability

Another problem in social co-robotics concerns sustainability, or long-term interaction [8]. What happens when the novelty of a robot wears off? How does a robot adapt to changes in the preferences of co-humans sharing the HSE?

Many of these problems inherently require a robot to exhibit some degree of creativity, be attuned to the moods of its proximate humans, and keep an inordinate sense of time and history. However, creativity is likely NP-Complete [29], mood awareness is at best AI-Complete (see Section II-D), and the requisite granulaity of storage is not easily definable - the representation strategy alone is likely AI-Complete [33].

D. Problem 4: Affect and Social Signal Awareness

One of the more commonly explored problems in social robotics concerns the recognition and synthesis of affect [21] and social signals [34]. While the precise definition of the terms "affect" and "social signal" are frequently debated in

[1]To appreciate the deep complexity of this problem, readers are encouraged to watch this time lapse video of a person with Muscular Dystrophy performing his morning routine: http://youtu.be/aSDxVG0fVg4 [10]

the affective science and social computing communities, most agree that these terms denote the visual and aural channels of human communication. Thus, a socially aware machine infers meaning from a human's face, gaze, posture, gestures, proxemics, and prosody, and can generate these signals itself.

There is certainly a clear need for robots operating in HSEs to be able to recognize and synthesize affect [26]. However, Picard, recognized as the founder of the field of affective computing, calls facial recognition alone one of the hardest, most complex problems in computer science [22]. Indeed, if one considers all of visual communication to be a component of natural language (as most do [30]), the problem of affect recognition is at best an AI-Complete problem: it inherently requires natural language understanding [14].

The problem of affect synthesis is unfortunately no better off from a complexity standpoint. Even in an asymmetrical dialogue, visual communication is intimately tied to dialogue content [1], situational context, *a priori* knowledge, expected norms, and, of course, *Social Learning*. Thus, this problem, too becomes at best an AI-complete problem, but at worst is not even definable let alone computable.

E. Problem 5: Social Norms

Social norms are "a standard, customary, or ideal form of behavior to which individuals in a social group try to conform." [4]. In the social co-robotics problem domain, this encompasses several things. Social norms place constraints on a robot's actions, in that they must conform to people's expectations and the situation they are in [15]. This is not to say co-humans will not forgive robot mistakes, our own research suggests people may be willing to overlook a robot's social missteps [25]. Nonetheless, from a technology acceptance perspective, there is motivation to program robots to be aware of social norms.

The *Social Norm* awareness problem is perhaps a superset of the *Affect and Social Signal Awareness* problem, because it requires additional knowledge to contextualize and classify observed human behavior. A person screaming alone in one's house is quite different than screaming while attending a sporting event. This problem space, too, is infinite, and thus not easily definable or computable.

F. Problem 6: Societal Issues

As social co-robots share HSEs with people, they inherently raise a plethora of societal issues, including privacy, security, acceptability, and so on. While these issues are not necessarily unique to social co-robotics, some of them may raise an alarm of intrusiveness other co-robotics domains do not need to face. To again draw on the example of biomolecular robotics; while a person may express concerns over the use of micro-scale drug delivery robots on an abstract level, they are unlikely to experience the same vitriolic response as they might to an embodied robot with agency in their home.

In his recent book *Robot Futures* [18], Nourbakhsh describes some of these vitriolic responses to robots, even in places as innocuous as a science museum. This use, misuse, and abuse of automated agents is not new, it dates back at least 60 years to Asimov's writings [16]. It has recently garnered attention in the HCI community by Parasuraman and Riley [20], who stress the importance of closely examining these attitudes and incorporating them into the design process.

This problem places a burden on social roboticists, because in addition to contending with the plethora of computational challenges that face our field, we further need to be concerned about public opinion of our robots if we ever want them to be purchased and used. Thus, iterative design and technology acceptance is critical, even at early stages of research.

III. Paths Forward

In social co-robotics, we tend to address these monumental challenges in three ways: we ignore them, we "wizard away" the problem by having a human compute the solution, or we severely constrain the problem space.

As a community, ignoring these problems does not help us advance our discipline. Instead, we suggest judicious use of human computation (wizards) and developing new techniques for problem satisficing as paths forward. We discuss these in further detail below.

A. Judicious Use of Wizard-of-Oz

Presently, many researchers in social co-robotics use Wizard of Oz (WoZ) control extensively, to solve most of the aforementioned problems, such as natural language processing, social understanding, dynamic space operation, etc. [23]. The original idea behind the WoZ paradigm was to be part of an iterative design process, a small aid in development as other components came to fruition [13]. The paradigm was not intended to be an end of and in itself; however, it has lately been re-tasked to enable robotics researchers to "project into the future" [8], enabling experiments which would be impossible with present technology.

While this at first is a compelling idea, at closer inspection the majority of problems a wizard is simulating in these robotic systems actually fall within the aforementioned six major problem domains of social co-robots. In other words, *wizards are simulating AI-Complete, non-definable, and non-computable problems.*

In no way do we argue for the total abdication of the WoZ paradigm from social co-robotics research; instead, we suggest roboticists be more careful in how they employ it. For a roboticist intentionally designing a semi-autonomous robot that will have human help when making decisions, it may make sense to employ WoZ as a kind of real-time human computation [27]. Other areas in the field of Artificial Intelligence facing NP-hard problems have embraced this paradigm, such as image labeling and machine translation - why not social co-robots?

But for researchers using WoZ purely as a method for testing complex psychological hypotheses involving co-robot acceptance, we urge judicious use of its employment. We may never have robots capable of some of these tasks, so it may be scientifically disingenuous and unrealistic to run experiments assuming their existence.

B. Problem Satisficing: "Good Enough" Social Co-robots

It may be the case that we cannot solve these monumental challenges facing our field; perhaps these problems will remain non-definable and non-computable forever. Nonetheless, it may be possible that we can build "good enough" social co-robots, where we solve problems well enough to enable adequate operation in HSEs.

Davis [9] (invoking Simon [28]) writes "nature is a satisficer, not an optimizer". Organisms solve problems in an acceptable or satisfactory (though not necessarily optimal) way. Why not robots? Certainly people have considered satisficing controllers for robotics problems in the past (c.f. [19], [7]), thus, we may be able to imagine ways to satisfice within the social co-robotics problem space without over-reduction.

It is not yet clear what the dimensions of such a satisfication might look like for social co-robots in HSEs; however, this seems to be a wide open area of research. What is the minimum functionality a social co-robot needs to complete its tasks in HSEs? What level of failure are co-humans willing to tolerate and excuse? There are many interesting questions in this domain, and we look forward to exploring them.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-1253935.

References

[1] J.B. Bavelas, L. Coates, and T Johnson. Listeners as co-narrators. *J Pers Soc Psychol*, 2000.

[2] C. Breazeal. *Designing Sociable Robots*. MIT Pr., 2004.

[3] R A Brooks. *Intelligence without representation*. MIT Press, 1991.

[4] M. Burke and H. P. Young. Norms, customs and conventions. In *Handbook of social economics*. 2010.

[5] M. Burrows. Global Trends 2030: Alt. Worlds. 2012.

[6] A Carlson, J Betteridge, B Kisiel, B Settles, E Hruschka, and T Mitchell. Toward an architecture for never-ending language learning. *AAAI*, 2010.

[7] J A Conlin. Getting around: making fast and frugal navigation decisions. *Prog Brain Res.*, 174, 2009.

[8] K. Dautenhahn. Methodology and themes of human-robot interaction: A growing research field. *International Journal of Advanced Robotic Systems*, 4(1), 2007.

[9] R. Davis. What Are Intelligence? And Why? *American Association for Artificial Intelligence*, 1998.

[10] ExtraAmpersand. My Disability: a real look at my life with Becker's MD. URL http://youtu.be/aSDxVG0fVg4.

[11] D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics and Automation*, 4(1):23–33, 1997.

[12] M. A Goodrich and A. C Schultz. Human-robot interaction: a survey. *Foundations and Trends in Human-Computer Interaction*, 1(3):203–275, 2007.

[13] T. D Kelley and L. N Long. Deep Blue Cannot Play Checkers: The Need for Generalized Intelligence for Mobile Robots. *Journal of Robotics*, (21), 2010.

[14] L. Lee. I'm sorry Dave, I'm afraid I can't do that": Linguistics, Statistics, and Natural Language Processing. *Computer Science: Reflections on the Field*, 2004.

[15] M. Lohse. *Investigating the influence of situations and expectations on user behavior: empirical analyses in human-robot interaction*. PhD thesis, Bielefeld University, 2010.

[16] L. McCauley. Countering the Frankenstein Complex. *American Association for Artificial Intelligence*, 2007.

[17] National Science Foundation. National Robotics Initiative. URL http://nsf.gov/nri.

[18] I. Nourbakhsh. *Robot Futures*. MIT Pr., 2013.

[19] T. J. Palmer and M. A. Goodrich. Satisficing anytime action search for behavior-based voting. 1, 2002.

[20] R. Parasuraman and V Riley. Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 1997.

[21] R. W. Picard. *Affective computing*. MIT Press, 1997.

[22] R W Picard. Emotion Technology: From Autism to Customer Experience and Decision Making. *Microsoft Research Cambridge*, 2009.

[23] L. D. Riek. Wizard of Oz Studies in HRI: A Systematic Review and New Reporting Guidelines. *Journal of Human Robot Interaction*, 1(1):119–136, 2012.

[24] L. D. Riek and P. Robinson. Challenges and Opportunities in Building Socially Intelligent Machines. *IEEE Signal Processing*, 28(3):146–149, 2011.

[25] L. D. Riek, T. Rabinowitch, P. Bremner, A.G. Pipe, M. Fraser, and P. Robinson. Cooperative gestures: Effective signaling for humanoid robots. *5th ACM/IEEE Int'l Conference on Human-Robot Interaction (HRI)*, 2010.

[26] Robotics-VO. A Roadmap for U.S. Robotics: From Internet to Robotics. Technical report, 2013.

[27] D Shahaf and E Amir. Towards a Theory of AI Completeness. *8th International Symposium on Logical Formalizations of Commonsense Reasoning*, 2007.

[28] H. A. Simon. Rational choice and the structure of the environment. *Psychological Review*, 63(2), 1956.

[29] William Squires. Creative computers: Premises and promises. *Art Education*, 36(3):21–23, 1983.

[30] J Streeck. Gesture as communication I: Its coordination with gaze and speech. *Communication Monographs*, 60 (4), 1993.

[31] S. Thrun and T. Mitchell. Lifelong robot learning. *Robotics and Autonomous Systems*, 15, 1995.

[32] R Toris and S Chernova. RobotsFor.Me and Robots For You. *IUI Interactive Machine Learning Workshop*, 2013.

[33] M C Torrance and L A Stein. Communicating with martians (and robots). Technical report, 1997.

[34] A. Vinciarelli, M Pantic, and H. Bourlard. Social signal processing: Survey of an emerging domain. *Image and Vision Computing*, 27(12):1743–1759, 2009.

[35] M Zillich, M. Bennewitz, M. Fox, J. Piater, and D. Pangercic. 2nd Workshop on Robots in Clutter: Preparing robots for the real world , June 2013. URL http://workshops.acin.tuwien.ac.at/clutter2013/.