

Do You Really Want to Know? Display Questions in Human-Robot Dialogues *A Position Paper*

Maxim Makatchev and Reid Simmons

{mmakatch, reids}@cs.cmu.edu

Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213

Abstract

Not all questions are asked with the same intention. Humans tend to address the implicit meaning of the question (that contributes to its *pragmatic force*), which requires knowledge of the context and a degree of common ground, more so than addressing the explicit propositional content of the question. Is recognizing the pragmatic force in today's human-robot dialogue systems worth the trouble? We focus on display questions (questions to which the asker already knows the answer) and argue that there are realistic human-robot interaction scenarios in existence today that would benefit from the deeper intention recognition. We also propose a method for obtaining display question annotations by embedding an elicitation question into the dialogue. The preliminary study of our robot receptionist shows that at least 16.7% of interactions with the embedded elicitation question include a display question.

Introduction

Humans produce utterances that convey intentions beyond those implied by their apparent propositional meaning. For example, some utterances formed as questions are not seeking information explicitly referred to by their semantic content, but instead are intended as assertions (*rhetorical questions* such as “Who do you think you are?”), social rituals (*phatic questions* such as “How are you?”), or are seeking to test the receiver's knowledge (*display questions* such as “Where is my office?”). As robots become more socially engaging, users may feel more inclined to produce utterances with the range of *pragmatic meaning* (briefly, a kind of meaning that depends on the context) that is closer to human-human dialogues. Robots, in turn, should be designed to respond to such utterances, either by solely interpreting the semantic content, or going deeper and addressing the pragmatic component, such as whether the question is an assertion, a greeting, or a test of the listener's knowledge.

Most present day human-robot interaction (HRI) studies are conducted in a controlled laboratory environment. It has been argued that studying HRI “in the wild” can uncover patterns that are not otherwise observed in the laboratory settings, e. g. (Sabanovic, Michalowski, and Simmons 2006).

Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

```

1 U: Jane Smith
2 R: Jane Smith is in 4001 Wean Hall.
   Would you like directions?
3 U: no she is not1
4 R: OK
5 U: Smith is not where you say she is

```

Figure 1: A verbatim fragment of a dialogue containing user's feedback. The person's name and room are modified for privacy, and the labels “U:” and “R:” are added to denote the user and robot turns.

For example Hala, the robot receptionist stationed in Qatar, routinely receives relationship status questions and marriage proposals (Makatchev et al. 2010). In the case of Roboceptionist, a robot receptionist located at the entrance of a university building (Gockley et al. 2005), users exhibit behaviors such as swearing (Lee and Makatchev 2009), joking, and asking display questions, i. e. questions for which they already know the answers.

Consider the fragment of a dialogue shown in Figure 1 between a human and the Roboceptionist. Clearly, the user has some knowledge of where Jane's office is (the robot's answer is actually incorrect). Why did the user ask this question to the robot in the first place? What was the user's intention? Is this the best possible response, to provide an answer referring to the semantic content of the query? In this paper, we argue that these are important questions to ask and address a problem essential for development of display question detection methods: obtaining ground-truth display question annotations.

Display questions are extensively studied in education research as an important part of a canonical Western-style classroom instruction (Boyd and Naclér 2001) (e.g. “What is the capital of Cuba?”). Outside of the “teacher talk” context, much like rhetorical and phatic questions (Malinowski 1923), display questions are not answered based solely on their semantic content (Koshik 2003). However, unlike rhetorical questions, they do usually require a recipient's response.

We hypothesize relations between display questions and user intentions, the process of assigning participant roles, and a degree of trust the user attributes to the robot. A major difficulty for evaluating these hypotheses is the lack of an

1 U: where is hamburg hall
 2 R: Sorry. I could not get the information. If you happen to know the answer to your question, can you tell it to me?
 3 U: yes
 4 R: Ok, what is the answer?
 5 U: go out of the double doors behind you and walk to your left, the path will lead you to the back of hamburg hall.

Figure 2: A verbatim dialogue fragment containing the embedded elicitation question in turns 2 and 4.

annotated corpus of display questions. As a possible solution to this problem, we propose an interactive approach to label display questions by incorporating an explicit elicitation request to the user within the dialogue itself. The advantage of our method, as compared to user interviews and questionnaires, is that it neither requires a robot to step out of the character nor breaks the impression that the environment is unmonitored. In the following section we describe the method in more detail and provide preliminary results of an application of this method to display question annotation.

Eliciting the Display Question Label

The Roboceptionist dialogues occasionally include fragments where users provide explicit feedback, or even answers to their own questions after the robot fails to answer them properly. For example, dialogues such as the one in Figure 1 suggest that (a) users do ask display questions, (b) some users are willing to provide detailed feedback after the robot fails to produce a satisfactory answer. We attempt to exploit the tendencies by making the robot at random purposely fail to answer a question. After each user question, we randomly choose to present an elicitation question (but not more than once per dialogue). Hence, longer dialogues are more likely to have an elicitation question than shorter ones. A fragment of such dialogue is shown in Figure 2. The subset of display questions that is labeled this way serves as an estimate of a lower bound on the fraction of interactions that contain a display question among all the interactions with an elicitation question. This is the estimate of a lower bound since (a) the user may choose not to answer the elicitation question even if she knows the answer, and (b) the elicitation intervention may be applied to a question that is not a display question within a dialogue that contains a display question.

Note that the lower bound holds even though the presence of the elicitation question affects the rest of the interaction. An experiment conducted over a period of 3 months shows that at least 16.7% ($SE = 4.6\%$) of interactions that included an elicitation question contained a display question.

Conclusion

We attempted to draw attention to the pragmatic forces behind questions asked to a robot receptionist. We hypothesized the relations between display questions and user intentions, role assignments and levels of trust. These hypoth-

esis are open for experimental verification, provided there is a method for reliable annotation and detection of display questions. We suggested one candidate method to label display questions that neither breaks the robot's character nor the impression that the environment is unmonitored. The approach, which involves eliciting user's answer to her own question shows that at least 16.7% of interactions that included the robot's elicitation question also contained a display question asked by the user. It remains to be seen how this data can be used to reliably determine user intentions and their effect on the user-robot interactions.

If the hypotheses are true, there are a number of ways in which the dialogue can be adapted to the user that produced a display question. To mention a few, (a) the user's true intention (e. g. exploring the robot's capabilities) can be addressed, (b) the robot can behave in accordance to the role or relative status that the user assigns to the robot, or try to affect these assignments, and (c) combining the two above, the robot may address the issue of low trust from the human by explaining its own capabilities, the robot's and the user's roles, and the shared task.

Acknowledgements

The authors would like to thank Sanako Mitsugi, Antonio Roque, and Suzanne Wertheim. This publication was made possible by the support of an NPRP grant from the Qatar National Research Fund. The statements made herein are solely the responsibility of the authors.

References

- Boyd, S., and Naclér, K. 2001. Sociocultural aspects of bilingual narrative development in Sweden. In Verhoeven, L. T., and Strömquist, S., eds., *Narrative development in a multilingual context*. Amsterdam: John Benjamins B.V.
- Gockley, R.; Bruce, A.; Forlizzi, J.; Michalowski, M.; Mundell, A.; Rosenthal, S.; Sellner, B.; Simmons, R.; Snipes, K.; Schultz, A. C.; ; and Wang, J. 2005. Designing robots for long-term social interaction. In *Proc. Int. Conf. on Intelligent Robots and Systems*, 2199–2204.
- Koshik, I. 2003. Wh-questions used as challenges. *Discourse Studies* 5:51–77.
- Lee, M. K., and Makatchev, M. 2009. How do people talk with a robot? An analysis of human-robot dialogues in the real world. In *Proc. of CHI*, 3769–3774.
- Makatchev, M.; Fanaswala, I. A.; Abdulsalam, A. A.; Browning, B.; Gazzawi, W. M.; Sakr, M.; and Simmons, R. 2010. Dialogue patterns of an arabic robot receptionist. In *Proc. of HRI*. Osaka, Japan: ACM.
- Malinowski, B. 1923. The problem of meaning in primitive languages. In Ogden, C. K., and Richards, I. A., eds., *The Meaning of Meaning: A Study of Influence of Language Upon Thought and of the Science of Symbolism*. New York: Harcourt, Brace and World. 296–336.
- Sabanovic, S.; Michalowski, M. P.; and Simmons, R. 2006. Robots in the wild: Observing human-robot social interaction outside the lab. In *Proc. of Int. Workshop on Advanced Motion Control*.