Implicature-Based Inference for Socially-Fluent Robotic Teammates

Claire Liang, Julia Proft, and Ross A. Knepper
Department of Computer Science
Cornell University
Ithaca, NY, USA
cyl48@cornell.edu, {jproft,rak}@cs.cornell.edu

***Abstract*—Actions performed in a collaborative scenario often have a functional goal paired with an implicit communicative goal that can only be understood in context. Robots must be able to reason about implicit communication in order to fully comprehend the meaning of actions performed by human teammates and the implicit signals conveyed by their own actions. To study implicit communication and its impact on team effectiveness in a simple domain, we implemented two AIs for playing *Hanabi*, a collaborative card game in which players are limited to a small set of actions but can choose their actions strategically in order to implicitly convey information that cannot be stated outright. Preliminary results from simulated games show that the AI that leverages implicit communication coordinates with teammates 46% more effectively than the one that does not, as measured by game score. We discuss the implications of these results and our plans to extend this work to the domain of robotics.**

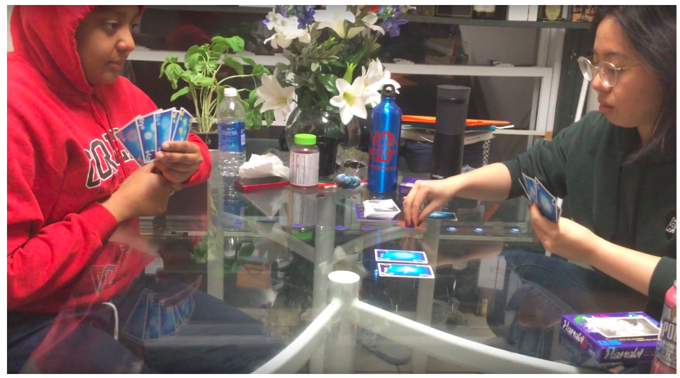

Fig. 1. Players of the collaborative card game *Hanabi* often use implicit communication to convey hidden information to their teammates.

I. Introduction

An important area of modern robotics research involves enabling collaborative behavior for robots working in close partnership with humans. When humans work together, they frequently employ implicit communication for reasons such as efficiency and group cohesion. For example, in a scenario where Alan and Beth are moving furniture together, Alan might take hold of one end of a couch, which achieves the functional objective of readying him to lift the couch. Since Alan cannot lift the couch by himself, the action also implies that Beth should move to the other side and prepare to lift the couch with him; no explicit request for this cooperative action is needed. We believe that robots need to be able to understand and generate this kind of implicit communication that results from actions being situated in context in order to be effective partners and integrate well with human teams.

Implicit communication manifests itself in different ways and is identified by different names depending on the context. In this paper, we focus on *conversational implicature*, first described by Grice [1]. We provide a primer for conversational implicature and describe how implicature can be used in a simple collaborative domain. Finally, we summarize our preliminary results and describe the next steps we will take with this work to examine other methods of implicit communication of interest in robotics.

II. Conversational Implicature

Implicature comes from pragmatics, the linguistics subfield that studies language in context. Grice [1] was the first to illuminate the mechanics behind speech that provides information beyond what is explicitly stated, or entailed. His *cooperative principle* asserts that speakers should contribute what is required by the accepted purpose of the conversation. He gives four maxims that describe how to conduct cooperative speech:

- Quality: only contribute information that is true.
- Quantity: provide all necessary information, but not more.
- Relation: make your contribution relevant.
- Manner: avoid ambiguity; be clear, orderly, and brief.

Instances of implicature arise when the maxims are flouted. Consider the following example from Recanati [4]:

Annie: "Can you cook?"
Bernard: "I am French."

In order to make sense of Bernard's response, Annie must apply the following inference steps:

(a) *Contextual premise:* Bernard is able to answer the question of whether he can cook.
(b) *Contextual premise:* It is common knowledge that the French are known for their ability to cook.
(c) Assume Bernard adheres to the cooperative principle and the four maxims.
(d) By (a), Bernard can completely resolve Annie's question, and by (c), he will.
(e) Only the propositions that Bernard can or cannot cook can fully resolve the question.
(f) By stating a fact seemingly irrelevant to the propositions of (e), Bernard flouts the maxim of Relation.

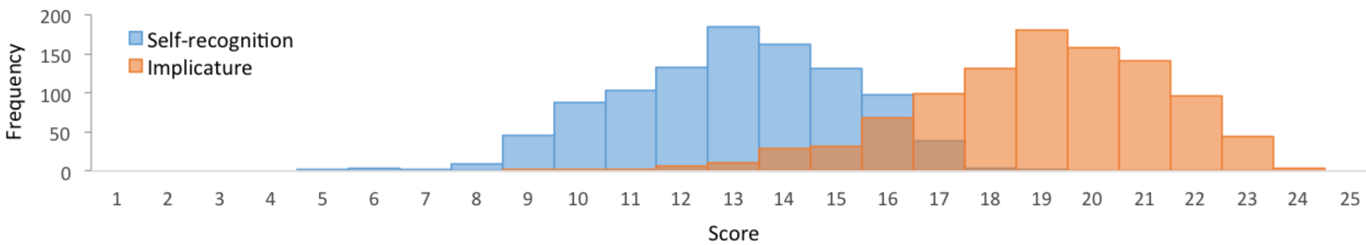

Fig. 2. Histogram of final scores of games played by AIs employing Osawa's self-recognition strategy [3] (blue) and our implicature-based strategy (orange).

(g) Thus Annie must search over a plausible set of facts in the relevant common ground to find fact (b) and conclude that Bernard is implicating that he is able to cook.

The search for meaning given in Lines (d)-(g) is described in greater detail in our previous paper [2].

III. Implicit Communication in Hanabi

Hanabi (see Figure 1) is a cooperative card game in which players can view all teammates' hands but their own. Each card has a color and a number value. Players have a restricted set of possible moves for each turn: playing a card, discarding a card, or providing a hint to a teammate. Hints are the only allowed method for conveying information about the cards a teammate possesses, and they must indicate all cards of one color or all cards of one numeric value in the teammate's hand. Cards must be played in sequence. With each successful play, the score increments. We consider two strategies for a two-player game of *Hanabi* with different hint interpretation methods:

(1) The *self-recognition strategy* presented by Osawa [3] focuses on using hints to minimize the entropy of one's own hand. Upon receiving a hint, the player estimates the contents of his or her own hand by simulating the teammate's worldview. Importantly, the player will not attempt to play a card until receiving enough information to be certain that the card is playable.

(2) Our *implicature-based strategy* prioritizes implied card playability over entropy reduction, thus preferring to interpret hints as actionable information whenever possible. In other words, the player uses the Gricean maxims and common knowledge about the game state to assume unstated information about the playability of a card from a hint, which often allows cards to be played sooner.

Before giving a hint, the player simulates how the teammate will interpret it. We note that both strategies involve the nested belief structures described by Vogel et al. [5]. The self-recognition strategy is a *level-one listener*, meaning that only the literal meaning of a hint is considered. Our implicature-based strategy, in contrast, is a *level-two listener* because it incorporates inference about the outcome-oriented motivation of the speaker in choosing which hint to give.

An example of an implicature-based hint would be Adrian informing Becky that she has one red card. The fact that Adrian gave a hint about a single card rather than a hint about multiple cards violates the maxim of Quantity and implicates that Adrian would like Becky to play the red card now.

IV. Preliminary Results and Future Work

We evaluated each strategy described above by teaming two instances of an AI based on the strategy and simulating 1,000 two-player games of *Hanabi*. Using final score as our metric for collaborative effectiveness, our preliminary results (see Figure 2) show that our implicature-based AI performs on average about 46% better than the AI employing self-recognition. In addition, the implicature-based AI is approximately 13.6% faster: it takes turns in about 3.71 seconds whereas the self-recognition AI takes 4.30 seconds. These results indicate that the use of implicature benefits a cooperative task.

Our observations of human *Hanabi* games suggest that most new players initially adopt a naïve entropy-minimization strategy in giving and interpreting hints but shift to an implicature-based strategy by the end of their first game. The rapidity with which people shift strategy suggests that people have an innate ability to perform implicature-based inference. We believe that the skills inherent in implicit inference for hint-giving in a simple domain like *Hanabi* will extrapolate to a variety of more complex human-robot teaming scenarios. Our first step in exploring these scenarios will be to test whether our preliminary AI-AI findings transfer to human-AI teams.

Whereas the official rules of *Hanabi* state that information can be conveyed only through restricted hints about card color or numeric value, in practice players often implicate additional information through out-of-band communication, sometimes without realizing it. These out-of-band implicit hints are frequently nonverbal and include saccadic eye motion and facial expressions. Arguably this kind of out-of-band communication is key to making the game fun and rewarding, and it likely contributes strongly to team cohesiveness.

In light of this, an important direction of our future work is to equip a *Hanabi*-playing robot with the ability to generate nonverbal implicit communicative actions alongside verbal implicatures and then to compare it to a *Hanabi*-playing robot that communicates implicitly only when giving verbal hints. We expect that the former robot will be perceived as more helpful and intuitively understandable than the latter. Such a result would indicate that it is essential for robots to leverage both nonverbal and verbal implicit communication channels in order to collaborate most effectively and naturally with people.

Acknowledgment

This material is based upon research supported by the Office of Naval Research under Award Number N00014-16-1-2080.

References

[1] Herbert P Grice. Logic and conversation. *Syntax and Semantics*, pages 41–58, 1975.

[2] Ross A Knepper, Christoforos I Mavrogiannis, Julia Proft, and Claire Liang. Implicit communication in a joint action. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, pages 283–292. ACM, 2017.

[3] Hirotaka Osawa. Solving Hanabi: Estimating hands by opponent's actions in cooperative game with incomplete information. In *Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

[4] François Recanati. What is said. *Synthese*, 128(1):75–91, 2001.

[5] Adam Vogel, Christopher Potts, and Dan Jurafsky. Implicatures and nested beliefs in approximate decentralized-pomdps. In *ACL (2)*, pages 74–80, 2013.