Strategies and Mechanisms to Enable Dialogue Agents to Respond Appropriately to Indirect Speech Acts

Gordon Briggs† and Matthias Scheutz‡

***Abstract*— Humans often use indirect speech acts (ISAs) when issuing directives. Much of the work in handling ISAs in computational dialogue architectures has focused on correctly identifying and handling the underlying non-literal meaning. There has been less attention devoted to how linguistic responses to ISAs might differ from those given to literal directives and how to enable different response forms in these computational dialogue systems. In this paper, we present ongoing work toward developing dialogue mechanisms within a cognitive, robotic architecture that enables a richer set of response strategies to non-literal directives.**

I. Introduction

Language-enabled autonomous agents are becoming increasingly prevalent in society, enabling human users to interact with and direct these agents using speech and other forms of natural language (NL), instead of more traditional human-machine interfaces (e.g. touch UIs). This trend is closely related to the growing interest in the development of "social robots" and other forms of artificial agents marketed for their "social" human-interaction capabilities.[1] However, even in narrow, task-based interactions, there is still a large gap between the current state of language-interaction capabilities of these new "social" robots and the richness and complexity of human communicative capabilities.

One key facet of human language is the ability to communicate a specific proposition or proposal many different ways. The simplest, and most relevant, example in task-based human-agent dialogues is, "Do X," where the expectation is for the agent to understand and successful carry out the commands over an increasingly large set of tasks. One major challenge, however, in this linguistic domain is that when people wish to communicate the desire for the addressee to "Do X," they often do not phrase it literally as "Do X." Rather, humans tend to use what are known as *indirect speech acts* (ISAs) [1]. For instance, an example of an ISA to request a cup of coffee could be, "Can you get me a coffee?" Instead of a literal question as to the addressee's ability to obtain a coffee for the speaker, this is usually interpreted as a directive for the addressee to serve the speaker coffee. Both understanding and generating these ISAs falls under the domain of *pragmatics*.

Humans balance a variety of pragmatic goals when formulating a directive, such as those expressed in the Grice's conversational maxims [2] and (sometimes competing) social normative considerations such as *politeness* [3]. People also balance these pragmatic considerations when responding to directives. There are a variety of ways one could signal acceptance of an ISA such as "Can you get me a coffee?", including, but not limited to: "Okay", "Yes, I can.", "Yes, I can. I'll do that right away." Different responses such as these represent different response strategies that consider different pragmatic goals.

In this paper, we present ongoing work toward developing dialogue mechanisms within a cognitive, robotic architecture that enables a richer set of response strategies to non-literal directives. We begin by presenting a set of different response strategies to non-literal directives and discuss how they differ with regard to what pragmatic goals they fulfill. We then present the representations and mechanisms that we have developed to enable this richer set of response strategies. Next, we present a proof-of-concept demonstration of these mechanisms in a fully-autonomous, integrated, human-robot interaction (HRI) scenario and also discuss the generality of the proposed mechanisms showing how they would scale to larger interaction domains. Finally, we compare our proposed system to the closest existing systems and discuss important differences, followed by a discussion of future work.

II. Motivation

Given the prevalence in the media of fictional robots or AIs with quite sophisticated NL abilities, it does not seem far-fetched to suppose that people would be inclined to communicate with robots using a lot of non-literal speech. Indeed, this intuition is supported by ideas such as the CASA hypothesis [4], which presents the notion that people fall back on human-human social interaction patterns with technology, regardless of their deeper philosophical views regarding the actual social agency of such artifacts. The usage of non-literal directives by humans instructing robotic agents has also been confirmed in recent HRI studies [5].

Over the years, several projects have sought to develop mechanisms for understanding and correctly acting on ISAs [6], [7], [8], [9], [5]. These projects focused primarily on providing mechanisms to correctly identify and comply with the underlying directive found in each processed ISA, rather than on the precise linguistic formulation of the addressee's response to the ISA speaker. Differing response formulations can be produced depending on which pragmatic goals are favored. We discuss some relevant pragmatic considerations below.

†NRC Postdoctoral Fellow, Naval Research Laboratory, Washington, DC USA, `gordon.briggs.ctr@nrl.navy.mil`

‡Human-Robot Interaction Laboratory Tufts University, Medford, MA USA, `matthias.scheutz@tufts.edu`

[1]Ex. https://www.jibo.com/

A. Pragmatic Goals

Below we list a set of pragmatic goals. We do not intend this list to be exhaustive, but rather a basis to compare and contrast various response strategies to ISAs. These pragmatic goals include:

1) *Correctness* - An utterance that conveys correct and supportable information is preferable to one that does not. This is representative of Grice's Maxim of Quality [2].
2) *Informativeness* - An utterance that conveys more information is preferable to one that conveys less. This is representative of Grice's Maxim of Quantity [2].
3) *Brevity* - An utterance that conveys its primary meaning in less time or fewer words is preferable to one that conveys its primary meaning in more time or words. This is representative of Grice's Maxim of Manner [2].
4) *Politeness* - An utterance that conform with social norms pertaining to respect for the addressee's feelings, wants, and autonomy is more preferable to one that does not [3], [10].

B. Response Strategies

Given a literal directive from speaker S, it is necessary for the addressee H to respond with either an acceptance or rejection of this directive. However, in the case of a non-literal directive, there exists both the underlying intended directive as well as the surface speech act. Below we present different strategies that are possible when responding to ISAs:

1) Strategy 1 - Respond to Non-literal Content Only: The first strategy is simply to respond to a non-literal directive in the same manner as one would with a literal directive. Many current ISA handling dialogue agents tend to follow this strategy [8], [9]. While this strategy may be sufficient to enable task-based interactions, and tends to optimize over the pragmatic objective of *brevity*, it is not ideal along a couple of dimensions.

The first is politeness. Prior psycholinguistic studies on politeness have demonstrated that people prefer (and consider it more polite) when both the literal and non-literal meanings of a non-literal directive are attended to in responses [11], [12]. That is to say, when given a request such as "Could you give me an X?", people would prefer and perceive as more polite and helpful responses such as "Yes. Here you go," over a simple "here you go," or "okay." Indeed, Leech (2014) also affirms that a positive answer via restatement of the query model is also considered a polite response strategy (e.g. "I could. Here you go.").

This form of response also is not particularly informative, at least in an explicit sense. Sometimes, speakers may use ISAs that take the form of questions because they also are genuinely uncertain about whether certain necessary capabilities or permissions (or other "felicity conditions") are satisfied [13]. A positive response to such an ISA might implicitly answer such a question, but would not be explicitly informative. Also, a negative response and explanation may not address the original query. For instance, responding, "Sorry, I am not permitted to serve coffee," to the ISA "Can you get me a coffee?" does not provide any information as to the ability of the robot to serve coffee, which is independent of the permissibility of this task.

2) Strategy 2A - Serial Processing and Response: The second strategy is to first process and respond the literal aspect for the ISA, then to process and respond to the non-literal aspect.[2]

This strategy, however, runs into trouble in certain cases. Consider the indirect request, "Would you mind getting off my foot?" Given that the intent of this ISA is to get the addressee's weight removed from the speaker's foot, responding to the literal query is either superfluous in the negative case (e.g. "No (I don't mind)."), or a violation of politeness in the positive case (e.g. "Yes (I mind)"). In such a case, Strategy 1 would be preferable.

Also, imagine a scenario in which a robot is currently busy with a high priority task, but is given the following ISA, "Could you grab me a glass of water?" Assuming the robot is capable and permitted to achieve this task, but it is of lower priority than its current one, the robot would potentially respond in the following manner, "Yes. Sorry, I am busy with a higher priority task." This is an awkward sounding response. The human addressee would interpret the initial positive response as an implicit acceptance of the directive, which is subsequently contradicted (a violation of correctness). This effect can be mitigated by use of linguistic cues that subtly acknowledge this polarity contradiction.

3) Strategy 2B - Serial Processing and Response with Polarity Acknowledgments: This modified strategy involves still processing and responding to both the literal and non-literal aspects of the ISA in serial order, but tracks the polarity of the first response and augments the subsequent response to the non-literal directive to acknowledge any contradictory polarity information. For instance, the response in the previous scenario with the server robot would be, "Yes. *But* now I am busy with a higher priority task." The "but" acknowledges the contradictory polarity information, and makes it clearer for the addressee that the initial positive response pertained to the literal query, while the subsequent negative response pertains to the non-literal directive.

4) Strategy 3 - Process Both, Then Respond: The final potential strategy is to process and determine the responses to both literal and non-literal aspects of the ISA, then determine the optimal response based on the particular social context and which subset of the pragmatic goals are currently most important. For instance, if brevity is most important, then responses similar to those generated by Strategy 1 could be favored. If politeness is most important, then responses similar to those generated by Strategies 2A and 2B could be

[2]We do not consider the reverse order, as it often results in a unnecessarily redundant (violation of brevity) response (e.g. "Sure. Yes, I could.")

favored. Such a strategy would require a mechanism to both ascertain a preference ordering of pragmatic goals based on the current situational context, as well as a means of selecting an appropriate response form based on this current ordering. To the authors' knowledge, there is no existing computational dialogue system that implements this ISA response strategy. It is also an area of current and future research for our dialogue architecture.

Nevertheless, one feature that is necessary in Strategies 2A, 2B, and 3 is the explicit maintenance and handling of both the literal and non-literal aspects of ISAs. In the following section, we introduce the pragmatic representation and mechanisms we utilize in order to implement strategies 1, 2A, and 2B (and will use to implement Strategy 3).

III. Pragmatic Interpretation Mechanisms

We start by providing a brief overview of the employed pragmatic framework implemented within a dialogue component in our cognitive robotic architecture and then provide details on the pragmatic inference algorithms.

A. Utterance Representation

We assume that the natural language processing subsystem of the agent architecture (consisting of the speech recognizer, the part-of-speech tagger, as well as the syntactic and semantic parsers) can generate utterance type representations with surface semantics in the following form:

$$U = UtteranceType(\alpha, \beta, X, M)$$

where $UtteranceType$ denotes the speech act classification, α denotes the speaker, β denotes the addressee, X denotes an initial semantic analysis, while M denotes a set of sentential modifiers (e.g., "now", "still", "really", "please"). These utterance type representation then form the basis for representations of dialogue exchange patterns and pragmatic inference rules.

B. Pragmatic Reasoning

The aim of the pragmatics reasoning mechanism is twofold. It determines a set of beliefs ($\mathbf{B}$) to infer upon receiving a natural language utterance U in a given context C and then determines what the best and most appropriate utterance form U is to generate given the goal of communicating the formula ϕ in context C. In order to facilitate both the understanding ($U + C \rightarrow \mathbf{B}$) and generation ($\phi + C \rightarrow U$) processes, a representation for pragmatic rules is needed.

We employ a representational form of pragmatic rules described below:

$$[[U]]_C := \langle \mathbf{B_{lit}}, \mathbf{B_{int}}, \theta \rangle$$

The rule associates a particular utterance form U in context C with a tuple containing the set of beliefs $\mathbf{B_{int}}$ to be inferred based on the intended meaning of the utterance (which is potentially non-literal), the set of beliefs to be inferred based on the literal meaning of the utterance $\mathbf{B_{lit}}$, as well as the degree θ to which the utterance can be considered a face-threatening act (FTA) in context C [3]. The idea of modeling the degree an utterance can be a FTA has been used in previous systems [14], [15], [9]. Both belief sets are represented in order to determine whether or not the pragmatic rule corresponds to a literal form, which affects the modulation of natural language generation based on politeness considerations. Additionally, as has been previously discussed, it is often the case that interlocutors expect that both the literal and non-literal aspects of utterances to be addressed and acknowledged. As such, both need to be explicitly represented so that both interpretations can be appropriately handled.

In order to facilitate rule specification, we use some notational shorthand in order to reduce the amount of information that needs to be explicitly specified for each rule. For instance, in cases where the intended meaning and the literal meaning are equal, it would be redundant to have to specify two equal sets of conjoined belief predicates. Therefore, in the case were two sets of semantics are not specified, it is assumed that $\mathbf{B_{int}} = \mathbf{B_{lit}}$.

C. Response Mechanisms

We have discussed the representation used for rules in the pragmatics reasoner. Here, we describe how these rules are applied and used to both update the belief state of the robot as well as begin the process of appropriate response generation. Specifically, given a parsed and semantically analyzed natural language utterance u, the main responsibility of pragmatic reasoning is to determine the set of updates to the agent's beliefs based on u. Subsequently, an appropriate response must be generated as well, possibly including reasoning about what information to include in the response and hot to formulate it. Based on u and the generated response, the belief state of the robot must then also be updated.

Figure 1 depicts the pragmatic interpretation process. The response strategy currently implemented in the dialogue component of our robotic architecture is Strategy 2B. This is achieved by first generating a response to the literal meaning of the query $\mathbf{B_{lit}}$, tracking the polarity of this response (either positive, negative, or neutral), then (if there exists a difference between $\mathbf{B_{int}}$ and $\mathbf{B_{lit}}$) generating a response to the non-literal meaning of the utterance $\mathbf{B_{int}}$. In certain cases, Strategy 1 behavior can be achieved by having rules in which both the literal and the non-literal belief update sets reflect the content of a non-literal directive. This is desirable for handling cases such as "Would you mind..." questions.

IV. Example Interaction

In this section, we present a simple HRI scenario in which we give a Nao robot simple instructions (both non-literal and literal).[3] Note that this walkthrough is primarily intended to demonstrate how the dialogue architecture processes both the literal and non-literal aspects of utterances in accordance with Strategy 2A/2B. The details that involve how the goal adoption process and rejection explanation mechanisms

[3] Video of this interaction can be found at: https://youtu.be/geOQ5oinZlw

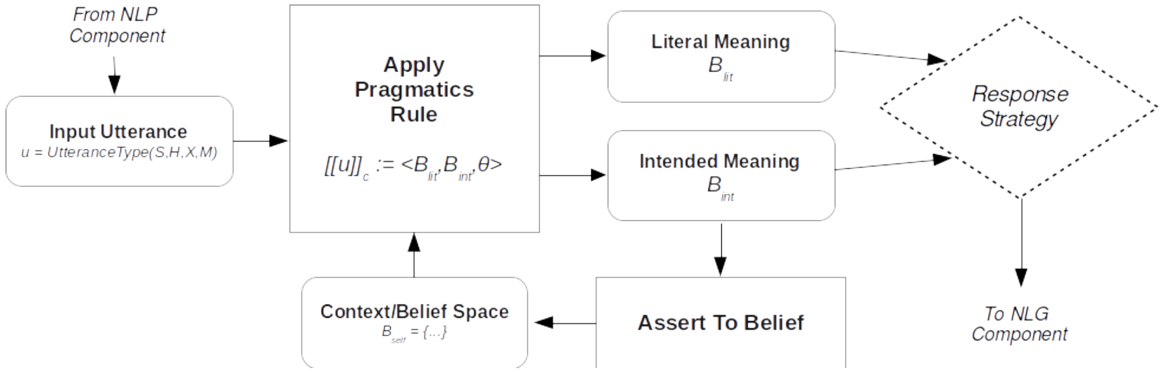

Fig. 1. Diagram depicting pragmatic rule application mechanism within the dialogue component of the robotic architecture.

function are described in other works. The transcript of this interaction is found below:

(D1) Person (CommX): Do you have a name?
Robot: Yes. My name is Shafer.
(D2) Person (CommX): I need you to turn right.
Robot: Okay. <turns right>.
(D3) Person (CommX): Can you stop?
Robot: Yes. <stops turning>.
(D4) Person (CommX): Go straight please.
Robot: Okay...Sorry, I cannot do that as there is no support ahead.
Person (CommX): The obstacle is not solid.
Robot: Okay.
Person (CommX): Walk forward.
Robot: Okay <walks forward>.

A. Initial Configuration

We begin by specifying the initial belief state and relevant pragmatic rules that loaded into the belief component and dialogue component in [our cognitive, robotic architecture], respectively.

1) Example Pragmatic Rules: For the sake of brevity, we omit the θ values associated with each rule. When a speaker α commands an addressee β to go straight, this is interpreted as indicative of a desire of the speaker α for β to achieve the goal of having moved forward one meter (where there is no intended non-literal meaning).

$$Instruct(\alpha,\beta,do(\beta,go(\beta,straight)),\{please\}) := \langle\{bel(\beta,want(\alpha,movedOneMeter(\beta)))\}\rangle \quad (1)$$

The interpretation of this utterance is the same for the version that does not have a "please" modifier. The modifier only modulates the associated θ-value of the rule.

When a speaker α asks an addressee β whether or not β has a name, this has the literal interpretation that α wants to know whether or not β possess a name. The non-literal interpretation is that α wants to know β's name.

$$AskYN(\alpha,\beta,have(\beta,name),\{\}) := \langle\{bel(\beta,itk(\alpha,have(\beta,name)))\}, \{bel(\beta,itkRef(\alpha,nameOf(\beta)))\}\rangle \quad (2)$$

When a speaker α tells an addressee β that α needs β to turn in a direction γ, this has the literal interpretation that α wants β to believe that α needs this particular state of affairs to be true. It has the additional non-literal interpretation that α wants β to turn in direction γ:

$$Stmt(\alpha,\beta,need(\alpha,turn(\beta,\gamma)),\{\}) := \langle\{bel(\beta,want(\alpha,bel(\beta,need(\alpha,turn(\beta,\gamma)))))\}, \{bel(\beta,want(\alpha,turning(\beta,\gamma)))\}\rangle \quad (3)$$

When a speaker α asks an addressee β whether or not β "can stop," this has the literal interpretation that α wants to know if β has the ability of stopping, whereas the non-literal interpretation is that α wants β to stop:

$$AskYN(\alpha,\beta,capableOf(\beta,stop(\beta)),\{\}) := \langle\{bel(\beta,itk(\alpha,capableOf(\beta,stopped(\beta))))\}, \{bel(\beta,want(\alpha,stopped(\beta)))\}\rangle \quad (4)$$

Finally, in the general case, when a speaker α tells an addressee β a proposition ϕ, then this has the literal interpretation that α wants β to belief that ϕ holds (with no additional non-literal interpretation):

$$Stmt(\alpha,\beta,\phi,\{\}) := \langle\{bel(\beta,want(\alpha,bel(\beta,\phi)))\}\rangle \quad (5)$$

2) Belief State: The robot begins with set of initial beliefs that describe the relevant capabilities and facts pertaining to this interaction.

$$\mathbf{B_{self}} = \{name(self,shafer), bel(self,capableOf(self,stopped(self))), bel(self,capableOf(self,turning(self,right))), bel(self,capableOf(self,turning(self,left))), bel(self,capableOf(self,movedOneMeter(self)))\}$$

B. Walkthrough

(D1) The interaction begins when the human interactant asks the robot, "Do you have a name?" The initial components in the NL input stream (i.e. the speech recognizer and NLP component) identify this utterance as being of of the form: $AskYN(commX,self,have(self,name),\{\})$. As such, this triggers pragmatic rule 2, which yields a distinct literal interpretation $\mathbf{B_{lit}}$ and non-literal interpretation $\mathbf{B_{int}}$. Based on the literal interpretation, the robot believes that the speaker has an intention-to-know if it possess a name ($itk(commX, have(self, name))$). This prompts a query to the belief component to see if the fact $have(self, name)$ is supported. Because $name(self, shafer) \Rightarrow have(self, name)$, $have(self, name)$ is return as being supported. As such, the robot responses in the affirmative.

The previous query was answered immediately by the dialogue component, as the literal semantics corresponded to an immediate dialogue obligation. In addition to this process, dialogue component periodically query the belief component to see if there are any intentions of other agents to know information it could potentially supply. Because the non-literal semantics $bel(self, itkRef(commX, nameOf(self)))$ have been submitted to the belief component, this prompts the dialogue component to supply this additional information. As such the robot answers the non-literal query, "My name is Shafer."

(D2) The human interactant then tells the robot: "I need you to turn right." The initial NL components identify this utterance as being of the form: $Stmt(commX, self, need(commX, turn(self, right)), \{\})$. This triggers the application of pragmatic rule 3. Because the utterance is of type $Stmt$, the dialogue obligation with regard to the literal interpretation is minimal. The robot simply acknowledges, "Okay."

The non-literal interpretation ($want(commX, turning(self, right))$), however, is submitted to the belief component. This desire of the human interactant triggers the robot to formulate and adopt a goal corresponding to this desire. This reasoning regarding obligations and permissibility that occurs at this stage is described in more detail in [16].

(D3) Next, as the robot is still turning right, the human asks the robot, "Can you stop?" Based on the identification of this NL utterance as: $AskYN(commX, self, capableOf(self, stop(self)), \{\})$, the pragmatic rule 4 is triggered. Once again, two distinct behaviors are caused by the distinct literal and non-literal meanings of this utterance. The dialogue component immediately seeks to answer the yes-no query, which it answers affirmatively, as $capableOf(self, stopping(self)$, is supported in the belief component.

The robot adopts the goal to stop based on the non-literal interpretation of the utterance ($want(commX, stopped(self))$), and subsequently stops.

(D4) Then, the human interlocutor instructs the robot to go straight, "Go straight please." Based on the identification of this NL utterance as: $Instruct(commX, self, do(self, go(self, straight)), \{\})$, the pragmatic rule 1 is triggered. The predicate $want(commX, movedOneMeter(self))$ is asserted in the robot's belief state and the inferred goal $goal(self, movedOneMeter(self))$ is submitted to the robot's goal manager. The goal status of the submitted goal is returned as INPROGRESS at this point, so the robot tentatively accepts the directive: "Okay..."

However, when the goal status is returned as FAILED, due to the presence of an obstacle ahead, the dialogue component follows up the initial acceptance with a rejection (including explanatory information), "Sorry, I can not do that as there is an obstacle ahead." Note that because this was a correction to a tentative acceptance of a literal command after a pause, the "but" modifier is not used.

The details of how this rejection process is grounded in reasoning regarding obligations and permissions is described in [16], as are the details of how these objections can be circumvented by providing new facts about the environment (that entail different conclusions regarding the permissibility or obligations surrounding certain goals). As the main intention of this example was to highlight the handling of both the literal and non-literal aspects of utterances, and the remaining utterances do not have significant non-literal content, we conclude our analysis of the scenario. Another example of ISAs handling can be found at: https://youtu.be/If4v4ohCByk.

V. Discussion

The contribution of this work is twofold. First, we highlight the issue of having multiple potential response patterns when handling ISAs. Second, we present a pragmatic interpretation mechanism that can be used to generate these different response patterns. Below, we highlight the differences between our approach and other NL architectures.

A. Other Approaches

In contrast to the system we present in this paper, other NL architectures used in human-robot interaction settings usually do not recognize dialogue obligations generated by the literal/surface forms of the received utterances in addition to those generated by non-literal semantics. Rather they primarily enable robots to detect non-literal directives, and then respond accordingly, either answering only the literal non-directive or the non-literal directive. For instance, Wilske and Kruijff (2006) implement this strategy to avoid having to have the robot say "I'm unavailable/busy," in the case when it receives a non-literal directive when the robot is busy or in a non-operational mode. Their system simply responds to the literal form of the utterance in order to satisfy discourse obligations [8]. As previously alluded to, this may lead to an awkward incongruity between the positive polarity of the response to the literal question or statement, and the implicit rejection of the directive. In the case of ISAs that take the form of questions (e.g. "Can you get me a coffee?"), this would indicate that the . However, in the case of statements about the speaker's desires and intentions (e.g. "I would like you to get me a coffee."), only acknowledging the statement without subsequent action might also be construed as dismissive and impolite.

Another prominent NL architecture that engages in some form of pragmatic reasoning during dialogue interactions is the TRIPS architecture. Allen and colleagues (2001) describe the TRIPS system handling an ISA in which both the literal surface form and non-literal interpretation are processed. However, the only details given regarding how the literal semantics are processed is that the system recognizes and processes an obligation to "RESPOND-TO" the surface utterance [17]. As such, it is difficult to make a direct comparison of TRIPS with our approach. It is unclear what response strategies TRIPS would use (and under what contextual conditions if there is variation in response strategy). TRIPS may have the representational sophistication to employ some of the response strategies we present in this paper, but it is not clear how they would be realized.

B. Future Work

As previously mentioned, the ultimate aim is to have a flexible, deliberative response mechanism (Strategy 3), rather than the relatively fixed ones represented by Strategies 1, 2A-B. This is a challenging problem. One potential approach, which has been demonstrated in the case of directive generation, is ranking candidate utterances according to individual pragmatic and sociolinguistic goals and merging these individual rankings to obtain the "best" utterance for a particular interaction [18]. These rankings can be differentially weighted depending on what pragmatic goals are more important for a specific social context. For instance, if a robot were deployed in a service role that involved interacting with members of the public (e.g. museum tour guide, reception waiter, etc.) then criteria such as *politeness* and *informativity* could be favored. In contrast, if a robot were deployed in a task-oriented setting with dedicated human interaction partners, then *directness* and *brevity* could be favored.

While this proposed approach is one potential way of realizing a general response strategy, there are still a variety of research problems that need to be addressed for such an approach to be truly generalizable. One challenge is that general operationalizations of various pragmatic criteria would need to be developed for not only single utterances (as in [18]), but also for utterance sequences such as those involved in responding to both the literal and non-literal aspects of directives. Another challenge is that mechanisms that recognize different social contexts and adjust the proposed pragmatic goal weightings would need to be developed.

VI. Conclusion

People will utilize both literal and non-literal utterances in order to communicate with artificial agents in task-based interactions. Previous work in handling non-literal utterances in computational dialogue systems has focused more on the correct understanding of these ISAs and subsequent correct task execution, rather than how to formulate the linguistic responses during such interactions. We have highlighted the fact that there are multiple possible response forms to ISAs, and made the case that explicit pragmatic interpretation representations and mechanisms are necessary in order to achieve a broader range of ISA responses. We describe the representations and architectural mechanisms in the DIARC/ADE robotic architecture that allow a range of strategies to be achieved. Finally, we have demonstrated these mechanisms in a couple proof-of-concept human-robot interactions, and compared our approach to other approaches for ISA handling.

VII. Acknowledgments

The work presented was completed while the first author was a member of the Human-Robot Interaction Laboratory at Tufts University. The views expressed in this paper are solely those of the authors and should not be taken to reflect any official policy or position of the United States Government or the Department of Defense.

References

[1] J. Searle, "Indirect speech acts," in *Syntax and Semnatics, 3: Speech Acts*, P. Cole and J. Morgan, Eds., 1975.

[2] P. Grice, "Logic and conversation," in *Syntax and Semnatics, 3: Speech Acts*, P. Cole and J. Morgan, Eds., 1975.

[3] P. Brown and S. C. Levinson, *Politeness: Some universals in language usage*. Cambridge University Press, 1987.

[4] C. Nass, J. Steuer, and E. Tauber, "Computers are social actors," in *Proceedings of the SIGCHI conference on Human factors in computing systems: celebrating interdependence*. ACM, 1994, pp. 72–78.

[5] G. Briggs, T. Williams, and M. Scheutz, "Enabling robots to understand indirect speech acts in task-based interactions," *Journal of Human-Robot Interaction*, vol. 6, pp. 64–94, 2017.

[6] C. R. Perrault and J. F. Allen, "A plan-based analysis of indirect speech acts," *Comput. Linguist.*, vol. 6, no. 3-4, pp. 167–182, July 1980. [Online]. Available: http://dl.acm.org/citation.cfm?id=972537.972540

[7] E. A. Hinkelman and J. F. Allen, "Two constraints on speech act ambiguity," in *Proceedings of the 27th annual meeting on Association for Computational Linguistics*, ser. ACL '89. Stroudsburg, PA, USA: Association for Computational Linguistics, 1989, pp. 212–219. [Online]. Available: http://dx.doi.org/10.3115/981623.981649

[8] S. Wilske and G.-J. M. Kruijff, "Service robots dealing with indirect speech acts," in *IROS'06*, 2006, pp. 4698–4703.

[9] G. Briggs and M. Scheutz, "A hybrid architectural approach to understanding and appropriately generating indirect speech acts," in *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, 2013.

[10] G. N. Leech, *The pragmatics of politeness*. Oxford University Press, USA, 2014.

[11] H. H. Clark and D. H. Schunk, "Polite responses to polite requests," *Cognition*, vol. 8, no. 2, pp. 111–143, 1980.

[12] R. W. Gibbs Jr and R. A. Mueller, "Conversational sequences and preference for indirect speech acts," *Discourse Processes*, vol. 11, no. 1, pp. 101–116, 1988.

[13] R. W. Gibbs Jr, "What makes some indirect speech acts conventional?" *Journal of memory and language*, vol. 25, no. 2, pp. 181–196, 1986.

[14] S. Gupta, M. A. Walker, and D. M. Romano, "How rude are you?: Evaluating politeness and affect in interaction," in *Affective Computing and Intelligent Interaction*. Springer, 2007, pp. 203–217.

[15] C. A. Miller, P. Wu, and H. B. Funk, "A computational approach to etiquette: Operationalizing brown and levinson's politeness model," *Intelligent Systems, IEEE*, vol. 23, no. 4, pp. 28–35, 2008.

[16] G. Briggs and M. Scheutz, ""Sorry, I Can't Do That": Developing Mechanisms to Appropriately Reject Directives in Human-Robot Interactions," in *AAAI Fall Symposium Series: Artificial Intelligence for Human-Robot Interaction*, Arlington, VA, 2015, pp. 32–36.

[17] J. F. Allen, D. K. Byron, M. Dzikovska, G. Ferguson, L. Galescu, and A. Stent, "Toward conversational human-computer interaction," *AI magazine*, vol. 22, no. 4, p. 27, 2001.

[18] F. Gervits, G. Briggs, and M. Scheutz, "The Pragmatic Parliament: A Framework for Socially-Appropriate Utterance Selection in Artificial Agents," in *39th Annual Meeting of the Cognitive Science Society*, London, UK, 2017.